






# Transparent Batchable Time-lock Puzzles and Applications to Byzantine Consensus

Shravan Srinivasan<sup>1</sup> , Julian Loss<sup>2</sup> , Giulio Malavolta<sup>3</sup>, Kartik Nayak<sup>4</sup> ,  
Charalampos Papamanthou<sup>5</sup> , and Sri AravindaKrishnan Thyagarajan<sup>6</sup> 

<sup>1</sup> University of Maryland

<sup>2</sup> CISA Helmholtz Center for Information Security

<sup>3</sup> Max Planck Institute for Security and Privacy

<sup>4</sup> Duke University

<sup>5</sup> Yale University

<sup>6</sup> NTT Research

**Abstract.** Time-lock puzzles (TLP) are a fascinating type of cryptographic problem that is easy to generate, but takes a certain time to solve, even when arbitrary parallel speedup is allowed. TLPs have wide-ranging applications including fairness, round efficient computation, and more. To reduce the effort needed to solve large numbers of TLPs, prior work has proposed batching techniques to reduce the cost of solving. However, these proposals either require: (1) a trusted setup or (2) the puzzle size linear in the maximum batch size, which implies setting an a priori bound on the maximum size of the batch. Any of these limitations restrict the utility of TLPs in decentralized and dynamic settings like permissionless blockchains.

In this work, we demonstrate the feasibility and usefulness of a TLP that overcomes all of the above limitations. Our construction is based on *indistinguishable obfuscation* and shows that there are no fundamental barriers in achieving such a TLP construction. As a main application of our TLP, we show how to improve the resilience of consensus protocols toward network-level adversaries in the following two settings: (1) We show a generic compiler that boosts the resilience of a Byzantine broadcast protocol  $\Pi$  as follows: if  $\Pi$  is secure against  $t < n$  weakly adaptive corruptions, then the compiled protocol is secure against  $t < n$  strongly adaptive corruptions. Here, ‘strong’ refers to adaptively corrupting a party and deleting messages that it sent while still honest. Our compiler is round and communication preserving, and gives the *first* expected constant-round Byzantine broadcast protocol against a strongly adaptive adversary for the dishonest majority setting. (2) We adapt the Nakamoto consensus protocol to a weak model of synchrony where the adversary can adaptively create minority partitions in the network. Unlike prior works, we do not assume that *all* honest messages are delivered within a known upper bound on the message delay. To the best of our knowledge, this is the first work to show that it is possible to achieve consensus in the permissionless setting even after relaxing the standard synchrony assumption.

## 1 Introduction

A *Time-Lock Puzzle* (TLP) is a cryptographic primitive that allows a sender to lock a message as a computational puzzle in a manner where the receiver will be able to solve the puzzle after a stipulated time  $\mathbf{T}$ . In terms of efficiency, a sender should be able to generate a puzzle substantially faster than the time required to solve, and in terms of security, an adversary should not be able to solve the puzzle faster than the stipulated time, even with parallel computation. Rivest, Shamir, and Wagner (RSW) [45] proposed the first TLP construction based on the sequentiality of repeated modular squaring in the RSA group. Many other TLP constructions [8,39,49] have followed suit in different settings but require the same flavor of sequential operations during solving. TLPs have found a wide variety of applications including sealed-bid auctions [45,39], timed-commitments [33,49], e-voting [14,39], fair contract signing [9,39], zero-knowledge proofs [18], cryptocurrency payments [48],

distributed consensus [50], blockchain front-running prevention [1], and more applications continue to emerge.

In many TLP applications involving multiple users, it is often the case that a user is required to solve the puzzles of all other users, and record all of the solutions. Say an auction house has to open all the time-locked bids and declare them publicly before announcing the winner. Batching and solving the puzzles is essential for scalability in such TLP applications that have large number of participating users. Intuitively, batch solving of TLPs allows a receiver to solve multiple puzzles simultaneously (at the price of solving of one puzzle) without needing to solve each puzzle separately. Specifically, the total running time of the batch-solve operation is bounded by some  $p(\lambda, \mathbf{T}) + \tilde{p}(\lambda, n)$  for some fixed polynomials  $p$  and  $\tilde{p}$ , where  $\lambda$  is the security parameter,  $n$  is the number of puzzles to be batched, and  $\mathbf{T}$  is the timing hardness of a single puzzle.

Modern TLP constructions achieve the seemingly impossible batch-solve property under various settings and assumptions [39,48,49]. However, *all* existing batch solving TLP schemes suffer from the following limitations: (1) requires a trusted setup to generate structured reference string, and/or (2) individual puzzle size scales linearly in the maximum number of puzzles that can be batched, which further implies that there is an a priori upper bound on the number of puzzles that can be batched (which is set upfront during puzzle generation).

Yet emerging blockchain applications like Miner Extractable Value (MEV) prevention [1], cryptocurrency payments [48], distributed consensus [50], etc., either for security or performance, require the TLP scheme to have a *transparent setup*, the puzzle size *independent* of the number of puzzles batched and thereby support batching an *unbounded* number of puzzles. These requirements arise in blockchain systems especially in the permissionless setting for the following reasons:

- First, it is often impossible or impractical to rely on a trusted party to generate the public parameters. Moreover, a compromised setup with trapdoors can violate the security of the system. Precisely for these reasons, such trusted parties are not assumed to exist and usage of cryptosystems requiring such trusted setups are actively discouraged in permissionless blockchain systems.
- Second, in permissionless systems (like Bitcoin), nodes can join and leave the network at will, and there are no mechanisms to authenticate any participant. So the exact number of participants  $n$  is unknown at any point in time. However, existing TLP schemes require the bound on the number of users at the time of puzzle generation. Thus, restricting the ability to accommodate more participants on demand after the puzzle generation phase.
- Third, large puzzle size increases the total communication costs. For instance, in a setting where all participants have to exchange puzzles with each other, a linear-sized puzzle of the prior constructions blows up the total communication overhead to  $O(n^3)$ .

Motivated by these open problems and applications, we ask the following question:

*Is it possible to build a TLP scheme with batch solving that has a transparent setup, puzzle size independent of the batch size, and therefore allows unbounded batching?*

In this work, we affirmatively answer this question and use our new TLP scheme to solve two elusive problems in distributed consensus. Namely,

1. The problem of expected *constant round* Byzantine broadcast under *corrupt majority* and *strongly adaptive model*.
2. The problem of permissionless consensus in *any* network model that does not assume that all messages will be delivered by a known upper bound on the network delay.

Below, we motivate how our TLP with the above properties can enhance security and reduce communication costs in the two applications we consider.

At a high level, in both applications, TLPs help to defend against a powerful network adversary that has the ability to delay or delete messages from the network. For instance, a powerful network adversary can simply learn the contents of the message sent by any honest party before deciding to corrupt or deliver the message. However, by time-lock encrypting the message, the adversary cannot learn the contents of the message before time  $\mathbf{T}$  without doing sequential work. In the meantime, all the honest messages would have been delivered. Thus, the adversary is forced to corrupt an honest node without inspecting the contents of the message. Prior works [50,15] showed the feasibility of the distributed consensus in the presence of a network level adversary using TLPs (without batching property) at the cost of polylogarithmic blowup in round complexity. However, batch solving and compact puzzles aid in improving the round complexity and communication costs, respectively, in these applications.

**Round-efficient Byzantine broadcast.** Byzantine broadcast (BB) is a well-studied problem in distributed consensus, and, in recent times, BB has emerged as a fundamental building block in blockchains [26]. Despite decades of study in improving the round efficiency, no prior BB protocol has expected constant round-complexity under strongly adaptive and dishonest majority setting. In the strongly adaptive model, an adversary can observe all the honest messages, corrupt honest nodes on the fly, and perform *after-the-fact* removal, that is, the adversary can delete any honest message *in-flight* before it reaches any other honest nodes.

Current success in constant round BB is in the weakly adaptive model where the adversary’s power is severely limited [51]. Wan et al. proposed the first sub-linear BB protocol under strongly adaptive and dishonest majority setting [50]. Their work used TLPs to prevent the adversary from inspecting the contents of the message before corrupting a node. Since their TLP construction did not support batching, they proposed a sub-protocol with polylogarithmic round complexity to distribute task of opening all puzzles to honest nodes, rather than solving the puzzles individually. We observe that by using a TLP with batching property and puzzle size independent of the batch size as a building block, it is possible to achieve expected constant round Byzantine broadcast under strongly adaptive and dishonest majority setting.

**Permissionless protocol in the mobile sluggish model.** Guo et al. introduced a relaxation of the synchronous model, which was subsequently called the *mobile sluggish model* [29,3]. In the mobile sluggish model, a subset of honest nodes, called *sluggish* nodes, can arbitrarily lose synchrony, but they faithfully follow the rest of the protocol. The remaining honest nodes, called *prompt* nodes, are synchronous and faithfully follow the protocol. Moreover, sluggishness can be *mobile*, that is, any honest can become sluggish over the protocol execution, and if a sluggish node becomes prompt by regaining synchrony, it will receive all the backlogged messages. This model is stronger than the partially synchronous and asynchronous model but weaker than the synchronous model. Pass and Shi showed that it is *impossible to achieve permissionless consensus* in a partially synchronous or asynchronous network [43]. Unfortunately, Nakamoto consensus is vulnerable to consistency violations even in the mobile sluggish model, as we show in this work. Specifically, even a *single* mobile sluggish fault can effectively reduce the collective mining rate of honest nodes by half!

One way to defend against a mobile sluggish adversary is to let an honest block winner simply time-lock encrypt the message before sending it, and other honest nodes time-lock encrypt a *decoy* to distract the adversary from spotting the block winner. Since the adversary cannot learn the contents of the puzzle without spending sufficient time, by setting the TLP duration slightly greater than

the round duration, the adversary is now forced to corrupt or deliver the message randomly. At the end of the round, honest nodes can batch solve the TLPs they received and update their chains. Unfortunately, *no* prior TLP with batch solving works in this application, due to the requirements and challenges in the permissionless setting: we cannot rely on a trusted setup [45,39,48], we do not know the number of users in the network a priori, and we do not want to blowup the round [45] and communication complexity [39,48,49].

## 1.1 Our Contributions

We give the first TLP construction (§4) that *simultaneously* achieves:

- *Transparent Setup*: Requires an one-time *transparent (public-coin)* setup.
- *Batch Solving*: Supports batch solving of *any* polynomial number of puzzles even after puzzle generation, and the size of individual puzzles is *independent* of the number of puzzles to be batched.

Our construction is based on *Indistinguishable Obfuscation (IO)* [24] where users’ puzzles are obfuscated programs. We employ new techniques to achieve compactness in the puzzle size while supporting unbounded batch sizes. Even though our construction is far from being practically efficient, our construction crucially shows that there are no fundamental barriers from achieving a TLP with the above properties and lays a blueprint for future work to instantiate our new techniques with more efficient tools. In §2, we briefly explain why existing techniques for TLP fail to achieve the desired properties, along with giving a brief overview of the key techniques used in our TLP scheme.

We use our TLP construction as a fundamental building block and overcome other challenges to solve two longstanding open problems in consensus:

1. **Round-efficient Byzantine broadcast:** In the years of distributed consensus research, we present (§5) the first *expected constant round* Byzantine broadcast under strongly adaptive and corrupt majority setting. To realize our result, we develop a generic compiler that uses *any* batch solving TLP construction to *convert any* broadcast protocol secure against a *weakly* adaptive adversary [51] into a broadcast protocol secure against a *strongly* adaptive adversary in a *round preserving* way, which could also be of independent interest. With our TLP, this compiler is additionally *communication* preserving. We formally prove the security (Thm. 3) of our compiler in the *programmable random oracle* (RO) model.
2. **Permissionless protocol in the mobile sluggish model:** We first show an attack to illustrate that Nakamoto consensus is not secure *even* in the mobile sluggish model (§6.1). We then present a proof-of-work based permissionless protocol (§6.3) which does not assume that the network is synchronous or all honest messages arrive on time. To the best of our knowledge, this is the *first* work to show that it is possible to achieve consensus in the permissionless setting even after *relaxing the standard synchrony assumption!* To do this, we develop a novel proof-of-work based *decoy* mechanism that uses TLPs to defend against a mobile sluggish adversary that can arbitrarily delay a subset of honest messages. We formally analyze our protocol to prove that it achieves *consistency* and *liveness*. Specifically, we show that our protocol realizes the standard properties namely, *chain growth* (Thm. 5), *chain quality* (Thm. 6), and *common prefix* (Thm. 7).

## 1.2 Related Work

**Time-lock puzzles.** Bitanski et al. [8] proposed a different approach to construct time-lock puzzles, assuming the existence of succinct randomized encodings [7] and non-parallelizable languages. Similar to RSW puzzles, during solving each puzzle has to be solved individually to obtain their solutions. Liu et al. [38] combine (extractable) witness encryption [25] and a public reference clock like the Bitcoin blockchain. In their construction, one can batch open many puzzles as the blockchain reaches a certain height as the computational effort is shared by the entire blockchain network in mining new blocks. Their construction relies on Succinct Non-Interactive Argument of Knowledge (SNARKs) [6] and thus non-falsifiable assumptions. Our construction on the other hand does not require such assumptions and does not rely on a global reference clock like a blockchain. In fact, one of the main applications of our TLP construction is to build such a blockchain. Malavolta and Thyagarajan introduced Homomorphic time-lock puzzles [39]. Their construction allowed homomorphic function evaluations to be performed on puzzles to obtain a single puzzle that embeds the function of all the original solutions. Subsequently, their solution was shown [48] to be useful also to batch a bounded number of solutions into a single puzzle, provided there is a trusted setup. In a recent work, Thyagarajan et al. [49] proposed a Class group based construction that gets rid of the trusted setup and only requires a transparent public-coin setup. In Table 1, we compare these TLP constructions with our construction in this work.

Recently, Burdges and Feo [11] proposed a related but a new notion called *delay encryption*. On a high level, users encrypt their messages to some common previously unpredictable identity ID using an Identity-Based Encryption (IBE) scheme. The decryption key for the identity ID can be derived anyone but the derivation is a delayed operation, meaning that it takes time  $\mathbf{T}$  to derive the key. We can batch decrypt several encryptions provided they are wrt. to the same ID. The drawback of their construction is the requirement of a trusted setup which is considered a strong assumption in the applications of our interest.

**Table 1:** Comparison with other batch TLP schemes.  $\lambda$  is the security parameter and  $\mathbf{T}$  is time hardness parameter. Compactness of puzzles here refers to the size of puzzles being independent of the batch size.

Scheme	Transparent setup	Unbounded batching	Compact puzzles	One-time Setup time	Practical efficiency
RSA-based [48,39]	✗	✗	✗	$O(\log(\mathbf{T}), \lambda)$	✓
Class-groups based [49]	✓	✗	✗	$O(\mathbf{T}, \lambda)$	✓
IO-based ( <b>This work</b> )	✓	✓	✓	$O(\mathbf{T}, \lambda)$	✗

**Strongly adaptive Byzantine broadcast.** Despite the rich literature in BB, Wan et al. proposed the first expected sub-linear round protocol in the strongly adaptive setting using Public-Key Infrastructure (PKI) and time-lock puzzles [50]. Subsequently, Cohen, Garay, and Zikas [15] explored the feasibility of *fair* broadcast in the strongly adaptive setting for both property-based definition and simulation-based definitions. Whereas, our focus is on achieving expected constant-round Byzantine broadcast under strongly adaptive and dishonest majority setting. We relate to other prior works in App. A.

**Protocols in the mobile sluggish model.** Guo et al. first introduced the mobile sluggish model as “weakly synchronous” model and showed that it is impossible for a Byzantine broadcast protocol to tolerate majority faults (Byzantine or sluggish) [29]. Subsequently, Abraham et al. presented Byzantine Fault Tolerant blockchain protocol that can tolerate minority corruptions in the mobile

sluggish model [3]. Kim et al. [35] observed that many proof-of-stake protocols, such as Dfinity [30], Streamlet [12], OptSync [47], can support mobile sluggish faults. These prior techniques heavily relied on using messages (votes) from a majority of the nodes (certificates) to establish communication with sluggish nodes and ensure safety of the protocol. Since Nakamoto consensus does not rely on using such certificates, their techniques do not apply in our setting.

**Nakamoto style protocols.** Prior works can be categorized based on the flavor of synchrony used to analyze Nakamoto consensus. In the lock-step model of synchrony, Garay et al. formally analyzed Nakamoto consensus [21]. Subsequently, Pass et al. and Kiffer et al. showed that the Nakamoto consensus is secure even in the non-lock-step synchrony model where the message delay is bounded and the time proceeds in discrete rounds [41,21,34,54]. Ren discarded the notion of discrete rounds and proved the security of Nakamoto consensus in the continuous model [44]. Even parallelly composed Nakamoto protocols such as Prism and OHIE are also in the lock-step model of synchrony [4,53]. Unfortunately, all these analyses assume that *any* honest message reaches other honest nodes in  $\Delta$  time regardless of the flavor of synchrony. Our analysis is in the mobile sluggish model, which assumes that a subset of honest nodes can violate the  $\Delta$ -assumption. However, the prompt nodes in our setting are assumed to be in lock-step model of synchrony.

**Network-adversary lower bounds and impossibilities.** Abraham et al. showed that a sub-quadratic protocol could not be resilient against a strongly adaptive adversary that can perform *after-the-fact* removal [2]. In Nakamoto consensus, delaying an honest block has the same effects as deleting the block. For example, if a newly mined block is delayed for a sufficiently long time, it could end up as an orphan block, which eventually gets pruned after the main chain stabilizes. Moreover, sluggishness can be mobile, thus making the sluggish adversary more powerful than the strongly adaptive adversary. Pass and Shi showed that it is impossible to achieve permissionless consensus in the partially synchronous or asynchronous network [43]. However, since the mobile sluggish model is stronger than the partially synchronous model, this impossibility does not apply to our setting.

## 2 Technical Overview

We give an overview of our TLP construction that supports batch solving for unbounded number of puzzles. Furthermore, we give an overview of our protocols that use our TLP construction as a fundamental building block to tolerate network-level adversaries in the Byzantine broadcast and Nakamoto consensus.

### 2.1 Time-Lock Puzzles with Batch Solving

**Bounded batching of TLPs.** Before delving into the specific of our construction, we show how standard techniques [39,49] readily give a construction of TLP with *bounded* batched solving, i.e., where the number  $n$  of batched solutions is fixed at puzzle generation time. Given a Linearly Homomorphic Time-Lock puzzle scheme LHP with homomorphism over  $\mathbb{Z}_q$  and a large enough  $q$ , we can homomorphically evaluate the packing algorithm. In more detail, we are given  $n$  puzzles  $Z_1, \dots, Z_n$  (of the LHP scheme) each encoding  $\lambda$ -bit values with timing hardness of each puzzle being  $\mathbf{T}$ . To batch solve these puzzles, we first homomorphically evaluate the following linear function:

$$f(x_1, \dots, x_n) = \sum_{i=1}^n 2^{(i-1) \cdot \lambda} \cdot x_i.$$



The resultant evaluated puzzle  $Z^*$  is then solved in time  $\mathbf{T}$  to obtain all the  $n$  values where each of these values were originally encoded as  $\lambda$ -bit values. Importantly, this means that the bit-representation of the plaintext space must be large enough to accommodate all  $n$ -secrets, i.e.,  $\log(q) \approx n \cdot \lambda$ . Since the domain has to be fixed at the time of puzzle generation, this means that each puzzle scales linearly with  $n$ . In settings with  $n$  parties, where each party generates a puzzle and broadcasts it to the other parties, the total communication is  $O(n^3)$ , assuming a total of  $O(n^2)$  communication for a broadcast of a single bit and ignoring factors that depend on  $\lambda$ .

**Unbounded batching?** The question that we set out to answer is whether it is possible to construct a TLP that supports *unbounded* batch-solving. One approach to do that is to “defer” the choice of the plaintext space at the solving time, so that the solver can select the appropriate domain, depending on how many puzzles need to be batched. A naive idea is to define a program  $\mathbf{P}$  that, on input the batch size  $n$ , outputs a LHP puzzle  $Z$  embedding the user’s message  $m$  and where the message space is sufficiently large to accommodate packing of  $n$  puzzles. This solution is clearly insecure as it reveals  $m$  in the plain, so to amend this we *obfuscate* the program  $\tilde{\mathbf{P}} := i\mathcal{O}(\mathbf{P})$ , using indistinguishability obfuscation (IO) [24,32]. Setting a super-polynomial upper bound on  $n \approx 2^{\omega(\log(\lambda))}$  allows one to batch any polynomial number of puzzles.

Unfortunately, this simple construction runs into issues when proving security. A natural strategy when proving security would be to hybrid over all possible  $n$ , hardwire the corresponding puzzle in the description of the circuit and the swap with for a puzzle encoding a fixed string (say 0) appealing to the security of the time-lock puzzle. However it is not hard to see that this would quickly run into issues: As  $n$  grows to super-polynomial, the size of the corresponding puzzle, and consequently of the obfuscated circuit, would also be super-polynomial. This is not only an issue of the security proof, since the actual obfuscated circuit must be padded to the maximum size of the circuits that is defined in the analysis. To get this strategy to work, our construction would yield a super-polynomial size puzzle!

**Our solution.** To understand our solution, we first discuss a way to circumvent the above issue. We change the output of the obfuscated circuit to output the  $n$  dimensional vector  $(0, \dots, m, \dots, 0)$ , where  $m$  is inserted in the  $i$ -th slot, masked by the output of a puncturable pseudorandom function (PRF)  $F$  and a LHP puzzle  $Z$  encoding the PRF key  $k$ .

$$\begin{bmatrix} 0 \\ \vdots \\ m \\ \vdots \\ 0 \end{bmatrix} + \begin{bmatrix} F(k, 1) \\ \vdots \\ F(k, i) \\ \vdots \\ F(k, n) \end{bmatrix} = \begin{bmatrix} F(k, 1) \\ \vdots \\ F(k, i) + m \\ \vdots \\ F(k, n) \end{bmatrix}$$

This structured solution allows us to solve the proof problem by puncturing each position individually, thereby avoiding an exponential blow-up in the size of the obfuscated circuit. To solve this puzzle, the solver can arrange the various masked plaintexts (from other obfuscated circuits of other users) diagonally, and sum up all the results, to obtain

$$\begin{bmatrix} F(k_1, 1) + m_1 \\ F(k_1, 2) \\ \vdots \\ F(k_1, n) \end{bmatrix} + \begin{bmatrix} F(k_2, 1) \\ F(k_2, 2) + m_2 \\ \vdots \\ F(k_2, n) \end{bmatrix} + \dots + \begin{bmatrix} F(k_n, 1) \\ F(k_n, 2) \\ \vdots \\ F(k_n, n) + m_n \end{bmatrix}$$

Here  $k_i$  and  $m_i$  are the PRF key and message, respectively, of the  $i$ -th obfuscated program. However at this point it is not clear how to batch solve the resulting puzzles, since each party will use an independent PRF key  $k_i$ . This means that we shifted the problem from recovering the  $n$  messages to recovering  $n$  PRF keys, bringing us back to square one. Our last idea is to use instead a *key-homomorphic* PRF  $\bar{F}$ . Assuming a suitable instantiation [10], we have the above expression evaluate to

$$= \begin{bmatrix} \sum_i \bar{F}(k_i, 1) + m_1 \\ \sum_i \bar{F}(k_i, 2) + m_2 \\ \vdots \\ \sum_i \bar{F}(k_i, n) + m_n \end{bmatrix} \approx \begin{bmatrix} \bar{F}(\sum_i k_i, 1) + m_1 \\ \bar{F}(\sum_i k_i, 2) + m_2 \\ \vdots \\ \bar{F}(\sum_i k_i, n) + m_n \end{bmatrix}$$

We also have a LHP puzzle from each party encoding  $k_i$ . The solver can now add all keys homomorphically and solve the resulting LHP puzzle in time  $\mathbf{T}$  to obtain  $\sum_i k_i$ . Once the key is known, the solver can simply unmask all the values in the above vector by evaluating the PRF at points  $(1, \dots, n)$  using the key  $\sum_i k_i$ . Subtracting the output yields the vector of  $n$  plaintexts.

This gives an outline of our construction. In the actual scheme, extra care is needed to match the modulus of the TLPs with the key space of the PRF, to set the parameters to account for the imperfect homomorphism, and deal with the lack of imperfect correctness of punctured keys in the proof (see §4 for more details). Notice that the size of each puzzle is dominated by the obfuscated program  $\tilde{\mathbf{P}}$ , which can be implemented to be of size logarithmic in  $n$  (the batch size). Therefore, in a multi-party setting, we get a total communication of  $\tilde{O}(n^2)$  assuming  $O(n^2)$  communication for single bit broadcast and ignoring factors that depend on  $\lambda$ .

## 2.2 Application 1: Efficient Byzantine Broadcast

Byzantine broadcast (BB) is a classical problem in distributed consensus, where a *designated sender* holds a bit  $b$  and wants to transmit  $b$  to all  $n$  nodes in the presence of  $t$  faults. A BB protocol is secure if it can guarantee *consistency* (all honest nodes output the same bit) and *validity* (if the designated sender is honest, all honest nodes output the designated sender’s input  $b$ ).

With increasing applications for BB (cryptography, blockchains, etc.), we study the round-efficiency of BB under the dishonest majority setting. Prior BB protocols in the dishonest majority setting can broadly tolerate: (1) weakly adaptive or (2) strongly adaptive adversary. Both strongly and weakly adaptive adversary can corrupt honest nodes *on the fly*. But, a weakly adaptive adversary cannot perform *after-the-fact* removal.

Despite decades of study, the state-of-art round-efficient of BB in the dishonest majority is in the weakly adaptive setting [51]. Thus, it raises the question:

*Is it possible to achieve an expected constant-round Byzantine broadcast under strongly adaptive and corrupt majority?*

We affirmatively answer this using PKI, RO, and *any* batch solvable TLP construction. Our solution is a *generic round preserving compiler* that can convert any weakly adaptive BB protocol into a strongly adaptive one. With our batch solvable TLP, our compiler is additionally *communication preserving* as well!

The key ingredient in our compiler is that we use TLPs to hide the contents of the messages sent by the underlying protocol so that the strongly adaptive adversary cannot learn the contents



of any message before honest nodes receives it. Prior works use the RSW puzzles to defend against a strongly adaptive adversary [50,15]. But, due to the inability to batch solve RSW proofs, opening all puzzles collectively adds an overhead of polylogarithmic rounds to the any protocol [50]. An alternate approach is to use batch solvable TLPs defined in [49,48] to remove the polylogarithmic communication overhead incurred by RSW puzzles. But this increases communication complexity by a linear factor since prior batch solvable TLPs are not compact. Since our TLP is compact and batchable, we can solve all puzzles in one round without increasing the communication complexity.

Even though TLPs can prevent the adversary from inspecting the contents of the message, the primary challenge is in proving that the compiled protocol is secure against an adversary that can perform *after-the-fact* removal. This is because TLPs, apart from hiding the message contents for  $\mathbf{T}$ , also serve as a commitment to the message inside the puzzle, which prevents the simulator from simulating the honest nodes without knowing the actual contents of the puzzle! Cohen et al. encountered a similar problem in the context of simulation-based *fair* BB and they proposed a non-committing TLP to overcome this challenge [15, Theorem 5].

**Non-committing TLPs.** Informally, it allows the simulator to equivocate a time-lock puzzle. That is, the simulator first generates and sends a “fake” TLP to the network, which can be later “opened” be to *any* message. Thus, when the simulator is asked to explain the contents, it programs the RO to open the desired message. In §4.3, we show how to achieve this property with our TLP construction as well.

**Compiler overview.** Abstractly, in a weakly adaptive protocol  $\Pi_{\text{bb-wa}}$ , a node performs three *basic steps*: In every round, (1) receives the messages sent by other nodes, (2) performs the state transition based on the messages received and computes the messages to send, and (3) sends the messages to other nodes. Our compiler interleaves each step of  $\Pi_{\text{bb-wa}}$  with TLP operations to obtain a strongly adaptive protocol,  $\Pi_{\text{bb-sa}}$ .

In a bit more detail, before sending a message in the compiled  $\Pi_{\text{bb-sa}}$ , a node uses non-committing TLPs to encrypt the message it wants to send and computes the puzzle with proof of well-formedness of the puzzle. Thus, instead of sending the plaintext in  $\Pi_{\text{bb-wa}}$ , a node in  $\Pi_{\text{bb-sa}}$  sends the puzzle, ciphertext, and the proof of well-formedness to other nodes. When a node receives puzzles, ciphertexts, proofs of well-formedness from the network, instead of opening one puzzle at time,  $\Pi_{\text{bb-sa}}$  uses the batchable TLP proposed in this work to obtain all the solutions simultaneously without incurring additional round complexity or communication complexity to open all the puzzles. Thus, after opening all the puzzles, a node in  $\Pi_{\text{bb-sa}}$  invokes the state transition function just like a node in  $\Pi_{\text{bb-wa}}$ . This process is repeated for every round.

We remark that our compiler is round preserving for *any* batch solvable TLP. Thus, our compiler can be realized in practice using the constructions based on RSA or Class-groups [49,48]. We defer the details of our compiler to §5.2.

### 2.3 Application 2: Permissionless Consensus in the Mobile Sluggish Model

Nakamoto’s protocol, used in Bitcoin, achieves consensus over the Internet in a *permissionless* setting, where: any node can join and leave the system at any time, the exact number of participating nodes is unknown, and the nodes have to communicate over unauthenticated channels. However, to realize consistency and liveness, the protocol assumes that the network is synchronous – all honest messages get delivered to one another within a known upper bound  $\Delta$ .

Unfortunately, assuming that an Internet scale protocol is synchronous is excessively optimistic. Moreover, Pass and Shi [43] showed that it is *impossible* to achieve permissionless consensus in

an asynchronous or even in a partially synchronous network [5], which are relaxations of the synchronous model. Thus, to deploy the protocol in the real-world, the protocol designers are compelled to choose a loose upper bound  $\Delta$  as the network delay to accommodate nodes with slow network.

In this work, we relax the standard synchrony assumption and study Nakamoto consensus under the mobile sluggish model [3,29]. For Internet scale protocols, the sluggish model is a pragmatic trade-off between the synchronous model and partially synchronous/asynchronous model. Thus, we ask the following question:

*Is it possible to achieve consensus in a permissionless setting  
in the presence of mobile sluggish faults?*

We affirmatively answer this question by proposing a protocol that uses our TLP construction from §4 as a fundamental building block to show that it is possible to achieve *consistency* (any two prompt chains can differ only in the last few blocks) and *liveness* (every prompt node eventually commits all transactions) even in the presence of mobile sluggish faults.

In this subsection, we show how to adapt the Nakamoto consensus to defend against a mobile sluggish adversary using the our TLP. In our protocol, we use the following ideas: (1) All honest nodes time-lock encrypt any message they transmit, (2) all honest nodes send *decoys* to protect the block winner from getting caught by the adversary, (3) restrict the adversary from flooding with decoys, and (4) ignore malformed puzzles sent by the adversary.

Formally, we define a round, a super-round, and duration of a round in §6.2 and §6.3. However, as a warm-up, we present strawman solutions to illustrate the inadequacies of the well-known approaches.

**Strawman solutions.** The first straightforward solution is to use RSW puzzles to time-lock encrypt any message with a duration equal to the network delay before transmitting across the network [45]. Unfortunately, this approach does not work for the following reasons:

- Recall that in a protocol like the Nakamoto consensus, only the block winner sends a message to the network. Thus, the adversary can easily stop the one message transmitted, whether or not the message is encrypted.
- Say the other honest nodes send out time-lock encrypted dummy messages, which acts as a *decoy* to protect an honest block winner from getting caught. Unfortunately, the honest parties have to open all the puzzles to find the winning block. Thus, the honest parties either have to open all the puzzles individually or open them using the *distributed-solve* primitive proposed by Wan et al. [50, Section 4.2]. Both these approaches increase the round complexity of the protocol by linear and polylogarithmic rounds, respectively.

An alternate approach is to use TLPs with batch solving property defined in [49,48], but we would suffer from large communication costs and fixed batch size problem as explained before. Instead, we can now use our TLP that gets rid of the these issues. Below we give an overview of other challenges we encounter in designing our permissionless consensus protocol.

**Decoys, spam prevention, and malformed puzzles.** Since the Nakamoto consensus is in the permissionless setting, there are no identities to tackle Sybil attacks. This setting raises an important question: how to stop the adversary from spawning multiple identities to send decoys? We resort to proof-of-work to tackle the Sybil attack!

Say the difficulty threshold to mine a block is  $T$ , then we set the threshold to mine a decoy as  $T_c$ , such that  $T < T_c$ . Each RO query made by a node simultaneously tries to mine a block and a *decoy*. That is, say  $h$  is the output of the hash function. If  $h < T$ , then a *block* is mined, else

if  $T \leq h < T_c$ , then *a decoy* is mined. This is the “2-for-1 POW” trick introduced by Garay et al. [21,42,4]. The parameter  $T_c$  presents an interesting trade-off:  $T_c$  should be sufficiently high so that honest nodes mine enough decoys whereas the adversary should not be able to overwhelm the honest nodes with many decoy puzzles.

One of the challenges is that nodes do not know the exact number of decoys mined at a given time. However, since our TLP construction can batch a variable number of puzzles, nodes can flexibly batch puzzles on demand. Observe that  $T_c$  restricts the number of decoys that the adversary (and the honest nodes) can mine. But, this does not stop the adversary from flooding the honest nodes with malformed puzzles. Batching malformed puzzles along with honest puzzles prevents a node from obtaining the solutions to honest puzzles. To circumvent this problem we equip our TLP with a verifiability property that allows an honest node to reject a puzzle that is not *well-formed* according to the puzzle generation algorithm. Thus, a valid proof guarantees that the plaintext can be obtained by solving the puzzle.

**Mine phase and solve phase.** Since the mining process is stochastic, the arrival times of a decoy and a block are random. Say if an honest node sends the puzzle as soon as it finds the block, it is unlikely that the rest of the honest nodes will also be sending the decoy puzzles at the same time. If enough honest nodes do not provide “cover” to the block winner, then the probability of the adversary guessing the block winner is high. However, if all honest nodes wait until a pre-determined time to send the respective puzzles, then block winner will have the best chance of not being detected by the adversary.

In order to capture this intuition, we have two phases in our protocol:

- *Mine phase*: All nodes spend a sequence of  $m$  rounds mining a block or decoy without sending or receiving any messages.
- *Solve phase*: This phase begins as soon as the mine phase ends and consists of *two* rounds. In the first round, nodes send and receive the puzzles they have mined in the *mine phase*, and check the well-formedness of the received puzzles. In the second round, nodes will batch solve the TLPs to find the block, if any, and update the longest chain.

We generically denote the duration of the solve phase as  $D$  rounds. If one employs RSW puzzles and the *distributed-solve* primitive from Wan et al. [50] instead of our TLP construction, then  $D$  can be thought of as the number of rounds required to perform *distributed-solve* procedure.

However, when our protocol is instantiated with our TLP construction we have  $D = 2$ .

Thus, the duration of a super-round is  $(m + D)$  rounds.

**Putting it all together.** In summary, by using our TLP, the decoy mechanism, and super-rounds, our protocol works as follows: Every honest node performs the following steps in every super-round: (1) Receive transactions from the environment, (2) choose the longest chain it has seen so far and break ties arbitrarily, (3) mine for  $m$  rounds (mine phase), and (4) solve for  $D$  rounds (solve phase) and update the longest chain. We defer the details of the protocol to §6.3.

### 3 Cryptographic Background

We denote by  $\lambda \in \mathbb{N}$  the security parameter. We say that a function  $\mu$  is negligible if it vanishes faster than any polynomial. The notation  $[n]$  denotes a set  $\{1, \dots, n\}$ . Background and notations relevant to the two applications are deferred to §6.2 and §5.1, respectively.

### 3.1 Time-Lock Puzzles

In the following we give a definition for the main object of interest of this work, namely time-lock puzzles (TLPs) [45]. The syntax follows the standard notation for time-lock puzzles except that we consider an additional setup phase that depends on the hardness parameter but not on the secret.

**Definition 1 (Time-Lock Puzzles).** *Let  $\mathcal{S}$  be a finite domain. A time-lock puzzle (TLP) with solution space  $\mathcal{S}$  is tuple of four algorithms (PSetup, PGen, PSol) defined as follows.*

- $\text{pp} \leftarrow \text{PSetup}(1^\lambda, \mathbf{T})$  a probabilistic algorithm that takes as input a security parameter  $1^\lambda$  and a time hardness parameter  $\mathbf{T}$ , and outputs public parameters  $\text{pp}$ .
- $Z \leftarrow \text{PGen}(\text{pp}, s)$  a probabilistic algorithm that takes as input public parameters  $\text{pp}$ , and a solution  $s \in \mathcal{S}$ , and outputs a puzzle  $Z$ .
- $s \leftarrow \text{PSol}(\text{pp}, Z)$  a deterministic algorithm that takes as input public parameters  $\text{pp}$  and a puzzle  $Z$  and outputs a solution  $s$ .

**Definition 2 (Correctness).** *A TLP scheme (PSetup, PGen, PSol) is correct if for all  $\lambda \in \mathbb{N}$ , all polynomials  $\mathbf{T}$  in  $\lambda$ , all secrets  $s \in \mathcal{S}$ , and all  $\text{pp}$  in the support of  $\text{PSetup}(1^\lambda, \mathbf{T})$ , it holds that:  $\Pr[\text{PSol}(\text{pp}, \text{PGen}(\text{pp}, s)) = s] = 1$ .*

Security requires that the solution of the puzzles is hidden for all adversaries that run in (parallel) time less than  $\mathbf{T}$ .

**Definition 3 (Security).** *A TLP scheme (PSetup, PGen, PSol) is secure with gap  $\varepsilon < 1$  if there exists a polynomial  $\tilde{\mathbf{T}}(\cdot)$  such that for all polynomials  $\mathbf{T}(\cdot) \geq \tilde{\mathbf{T}}(\cdot)$  and every polynomial-size adversary  $(\mathcal{A}_1, \mathcal{A}_2) = \{(\mathcal{A}_1, \mathcal{A}_2)_\lambda\}_{\lambda \in \mathbb{N}}$  where the depth of  $\mathcal{A}_2$  is bounded from above by  $\mathbf{T}^\varepsilon(\lambda)$ , there exists a negligible function  $\mu(\cdot)$ , such that for all  $\lambda \in \mathbb{N}$  it holds that*

$$\Pr \left[ \begin{array}{l} b \leftarrow \mathcal{A}_2(\text{pp}, Z, \text{st}) \\ \wedge (s_0, s_1) \in \mathcal{S}^2 \end{array} : \begin{array}{l} \text{pp} \leftarrow \text{PSetup}(1^\lambda, \mathbf{T}(\lambda)) \\ (\text{st}, s_0, s_1) \leftarrow \mathcal{A}_1(1^\lambda, \text{pp}) \\ b \leftarrow \{0, 1\}, Z \leftarrow \text{PGen}(\text{pp}, s_b) \end{array} \right] \leq \frac{1}{2} + \mu(\lambda)$$

*Homomorphic Time-Lock Puzzles.* We also recall the definition of homomorphic time-lock puzzles [39], which allows one to compute functions on secrets homomorphically, without solving the puzzles first.

**Definition 4 (Homomorphic TLPs).** *Let  $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$  be a family of circuits (together with their respective representations). A TLP scheme (PSetup, PGen, PSol) is homomorphic if the syntax is augmented with the following interface:*

- $Z' \leftarrow \text{PEval}(C, \text{pp}, Z_1, \dots, Z_n)$  a probabilistic algorithm that takes as input a circuit  $C \in \mathcal{C}_\lambda$ , public parameters  $\text{pp}$  and a set of  $n$  puzzles  $(Z_1, \dots, Z_n)$  and outputs a puzzle  $Z'$ .

Homomorphic TLPs must satisfy the following notion of evaluation correctness.

**Definition 5 (Evaluation Correctness).** *Let  $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$  be a family of circuits (together with their respective representations). An homomorphic TLP scheme (PSetup, PGen, PSol, PEval) is correct (for the class  $\mathcal{C}$ ) if for all  $\lambda \in \mathbb{N}$ , all polynomials  $\mathbf{T}$  in  $\lambda$ , all circuits  $C \in \mathcal{C}_\lambda$  and respective inputs  $(s_1, \dots, s_n) \in \mathcal{S}^n$ , all  $\text{pp}$  in the support of  $\text{PSetup}(1^\lambda, \mathbf{T})$ , and all  $Z_i$  in the support of  $\text{PGen}(\text{pp}, s_i)$ , the following conditions are satisfied:*

– It holds that

$$\Pr [\text{PSol}(\text{pp}, \text{PEval}(C, \text{pp}, Z_1, \dots, Z_n)) = C(s_1, \dots, s_n)] = 1.$$

– There exists a fixed polynomial  $p(\cdot)$  such that the runtime of  $\text{PSol}$  is bounded by  $p(\lambda, \mathbf{T})$  and the runtime of  $\text{PEval}$  is bounded by  $p(\lambda)$ .

We require homomorphic TLPs specifically that support homomorphic evaluations of linear functions over the puzzles, that are secure against depth bounded but sub-exponential size adversaries. We have such constructions from RSA groups [39] and Class groups with imaginary quadratic order [49]. These constructions are proven secure against such adversaries by conjecturing the hardness of the *sequential squaring assumption* [39,37] against depth bounded but sub-exponential size adversaries.

### 3.2 Puncturable Pseudorandom Functions

A puncturable pseudorandom function (PRF) is an augmented PRF that has an additional puncturing algorithm. Such an algorithm produces a punctured version of the key that can evaluate the PRF at all points except for the punctured one. It is required that the PRF value at that specific point is pseudorandom even given the punctured key. A puncturable PRF can be constructed from any one-way function [27].

**Definition 6 (Puncturable PRFs).** *A puncturable family of PRFs is a tuple of polynomial-time algorithms  $(\text{Setup}, \text{KGen}, \text{Punc}, \text{F})$  defined as follows.*

- $\text{pp} \leftarrow \text{Setup}(1^\lambda)$  a probabilistic algorithm that takes as input the security parameter  $1^\lambda$  and outputs public parameters  $\text{pp}$ . Public parameters  $\text{pp}$  are taken as input in all other algorithms.
- $K \leftarrow \text{KGen}(\text{pp})$  a probabilistic algorithm that takes as input the public parameters  $\text{pp}$  and outputs a key  $K$ .
- $K_i \leftarrow \text{Punc}(K, i)$  a deterministic algorithm that takes as input a key  $K \in \mathcal{K}$  and a position  $i \in \mathcal{X}$  and returns a punctured key  $K_i$ .
- $y \leftarrow \text{F}(K, i)$  a deterministic algorithm that takes as input a key  $K$  and a string  $i \in \mathcal{X}$  and returns a string  $y \in \mathcal{Y}$ .

**Definition 7 (Correctness).** *For all  $\lambda \in \mathbb{N}$ , for all outputs  $K \leftarrow \text{KGen}(1^\lambda)$ , for all points  $i \in \mathcal{X}$  and  $x \in \mathcal{X} \setminus i$ , and for all  $K_{-i} \leftarrow \text{Punc}(K, i)$ , we have that  $\text{F}(K_{-i}, x) = \text{F}(K, x)$ .*

We require that punctured points are pseudorandom to the eyes of any efficient distinguisher.

**Definition 8 (Pseudorandomness at Punctured Points).** *For all  $\lambda \in \mathbb{N}$  and for every PPT adversaries  $(\mathcal{A}_1, \mathcal{A}_2)$  there is a negligible function  $\mu(\cdot)$ , such that*

$$\Pr \left[ \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda), (i, \tau) \leftarrow \mathcal{A}_1(\text{pp}) \\ b \leftarrow \mathcal{A}_2(\tau, K_i, i, y) : K \leftarrow \text{KGen}(\text{pp}), K_i \leftarrow \text{Punc}(K, i), b \leftarrow \{0, 1\} \\ \text{if } b = 0 \text{ then } y \leftarrow \mathcal{Y}, \text{ else } y \leftarrow \text{F}(K, i) \end{array} \right] \leq \frac{1}{2} + \mu(\lambda).$$

*Key Homomorphism.* We also assume the existence of constructions of puncturable PRFs that satisfy key homomorphism [10].

**Definition 9 ( $\gamma$ -Almost Key-Homomorphic PRF).** *Let function  $\text{F} : \mathcal{K} \times \mathcal{X} \rightarrow \mathbb{Z}_p^m$  be such that  $(\mathcal{K}, +)$  is a group. Then the tuple  $(\text{F}, +)$  is said to be  $\gamma$ -almost key-homomorphic PRF if the following two conditions hold:*

- $F$  is a (puncturable) pseudorandom function.
- For all  $k_1, k_2 \in \mathcal{K}$  and all  $x \in \mathcal{X}$ , there exists a vector  $\mathbf{e} \in [0, \gamma]^m$  such that

$$F(k_1, x) + F(k_2, x) = F(k_1 + k_2, x) + \mathbf{e} \pmod{p}.$$

The scheme presented in [10] satisfies (additive) key-homomorphism over  $\mathbb{Z}_q^n$ , which we also use in this work. Their scheme satisfies a weaker notion of correctness, which we state below.

**Definition 10 (Computational Functionality Preservation).** For all  $\lambda \in \mathbb{N}$  and all PPT adversaries  $(\mathcal{A}_1, \mathcal{A}_2)$ , there exists a negligible function  $\mu(\cdot)$ , such that

$$\Pr \left[ \begin{array}{l} x^* \leftarrow \mathcal{A}_2^{F(K, \cdot)}(1^\lambda, K_{i^*}, \tau) \wedge \\ x^* \neq i^* \wedge \\ F(K, x^*) \neq F(K_{i^*}, x^*) \end{array} : \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda), K \leftarrow \text{KGen}(\text{pp}) \\ (i^*, \tau) \leftarrow \mathcal{A}_1(\text{pp}) \\ K_{i^*} \leftarrow \text{Punc}(K, i^*) \end{array} \right] \leq \mu(\lambda).$$

For our purposes, we require the above property of the key-homomorphic puncturable PRF from [10] to hold against super-polynomial adversaries, which is possible assuming the hardness of LWE against super-polynomial adversaries.

### 3.3 Indistinguishability Obfuscation

We recall the definition of indistinguishability obfuscation (iO) for circuits from [24].

**Definition 11 (iO for Circuits [24]).** A uniform PPT machine  $i\mathcal{O}$  is an indistinguishable obfuscator for circuit class  $\{\mathcal{C}_\lambda\}$ , if the following are satisfied:

- For all  $\lambda \in \mathbb{N}$ , or all  $C \in \mathcal{C}_\lambda$ , for all inputs  $x$ , we have

$$\Pr [C'(x) = C(x) : C' \leftarrow i\mathcal{O}(\lambda, C)] = 1$$

- For all  $\lambda \in \mathbb{N}$ , all pairs of circuit  $(C_0, C_1) \in \mathcal{C}_\lambda$  such that  $|C_0| = |C_1|$  and  $C_0(x) = C_1(x)$  on all inputs  $x$ , it holds that the distributions  $\{i\mathcal{O}(\lambda, C_0)\}$  and  $\{i\mathcal{O}(\lambda, C_1)\}$  are computationally indistinguishable.

## 4 Time-Lock Puzzles with Batch Solving

In this section we formally present the notion and constructions for time-lock puzzles with batched solving.

### 4.1 Definition

We define the notion of TLPs with batched solving. We borrow the standard interfaces of a TLP from §3.1 and append it with an interface to allow for batched solving of  $n$  puzzles.

**Definition 12 (Batch Solving).** A TLP scheme  $(\text{PSetup}, \text{PGen}, \text{PSol})$  supports batch solving with the aid of an additional interface defined below

- $(s_1, \dots, s_n) \leftarrow \text{BatchPSol}(\text{pp}, Z_1, \dots, Z_n)$  a deterministic algorithm that takes as input public parameters  $\text{pp}$  and puzzles  $Z_1, \dots, Z_n$ , and outputs solutions  $s_1, \dots, s_n$ .



**Definition 13 (Batch Solving Correctness).** An TLP scheme  $(\text{PSetup}, \text{PGen}, \text{PSol})$  with batch solving interface  $\text{BatchPSol}$  is correct if for all  $\lambda \in \mathbb{N}$ , all polynomials  $\mathbf{T}$  in  $\lambda$ , all polynomials  $n$  in  $\lambda$ , all solutions  $(s_1, \dots, s_n) \in \mathcal{S}^n$ , all  $\text{pp}$  in the support of  $\text{PSetup}(1^\lambda, \mathbf{T})$ , and all  $Z_i$  in the support of  $\text{PGen}(\text{pp}, s_i)$ , the following conditions are satisfied:

- There exists a negligible function  $\mu(\cdot)$  such that

$$\Pr [\text{BatchPSol}(\text{pp}, Z_1, \dots, Z_n) \neq (s_1, \dots, s_n)] \leq \mu(\lambda).$$

- There exist fixed polynomials  $p(\cdot), \tilde{p}(\cdot)$  such that the size complexity of the circuit evaluating  $\text{BatchPSol}(\text{pp}, Z_1, \dots, Z_n)$  is bounded by  $p(\lambda, \mathbf{T}) + \tilde{p}(\lambda, n)$ .

Notice that the above definition rules out trivial solutions, where you solve the  $n$  puzzles individually and output the solutions. This is because, in this solution the size scales with  $n \cdot \mathbf{T}$ , while the definition above only permits the scale to be  $n + \mathbf{T}$ . One can view  $\tilde{p}(\lambda, n)$  as capturing the time taken to read and process the  $n$  puzzles, and returning the  $n$  solutions. The factor  $p(\lambda, \mathbf{T})$  captures the solving of a single puzzle and itself is independent of  $n$ .

## 4.2 Bounded Batching of TLPs

As hinted to in §2.1, given a linearly homomorphic time-lock puzzle with homomorphism over  $\mathbb{Z}_q$ , that has a large enough message space, it was shown in [48,49] that we can homomorphically pack several puzzles into a single puzzle using standard techniques. Solving the single puzzle reveals the solutions to all the  $n$  puzzles that we started out with. A crucial requirement for the above batch solving to work is that the message space of the homomorphic time-lock puzzle must be large enough to accommodate all the  $n$   $\lambda$ -bit values. It was shown [48] that this is indeed possible for by instantiating the Paillier-based linearly homomorphic time-lock puzzle construction from [39] in the same way Damgård-Jurik [16] extended the Paillier cryptosystem [40]. That is, instantiate the Paillier-based linearly homomorphic time-lock puzzle from [39] with a modulus  $N^s$  instead of modulus  $N^2$ , for a large enough value  $s$ . A similar domain extension was also shown in the settings of class groups of imaginary quadratic orders [49] that only require a public-coin (transparent) setup.

More formally, the LHP scheme from [39] has the  $\text{LHP.PSetup}$  algorithm output  $\text{pp}_{\text{LHP}} = (\mathbf{T}, N, g, h)$ , where  $N = pq$  for some  $\lambda$ -bit primes  $p$  and  $q$ ,  $g$  is the generator of  $\mathbb{J}_N^*$ , and  $h := g^{2^{\mathbf{T}}} \bmod N$ . Here  $\mathbb{J}_N^*$  denotes the elements in  $\mathbb{Z}_N^*$  with Jacobi symbol  $+1$ . The puzzle  $Z$  embedding message  $m$  is of the form  $(u, v)$  where

$$u := g^r \bmod N \text{ and } v := h^{r \cdot N} (1 + N)^m \bmod N^2,$$

with randomness  $r \leftarrow [N^2]$ . The security of the scheme follows from the *sequential squaring assumption* [45,39]. The Damgård-Jurik extension from [39,48] lets the puzzle generation algorithm additionally choose  $s \in \mathbb{Z}$ , and set the puzzle  $Z := (u, v)$  where

$$u := g^r \bmod N \text{ and } v := h^{r \cdot N^{s-1}} (1 + N)^m \bmod N^s.$$

Here, the message space is  $\mathbb{Z}_{N^{s-1}}$  while the puzzle component  $v$  is in  $\mathbb{Z}_{N^s}$ .

Consider  $n$  puzzles  $Z_1, \dots, Z_n$  each encoding  $\lambda$ -bit values with timing hardness  $\mathbf{T}$ , and each of these puzzles are of the Damgård-Jurik extended form. The  $\text{LHP.BatchPSol}$  algorithm internally evaluates the following linear function

$$f(x_1, \dots, x_n) = \sum_{i=1}^n 2^{(i-1) \cdot \lambda} \cdot x_i$$

homomorphically over the puzzles using the LHP.PEval algorithm. The effect of this evaluation is that the resultant puzzle  $Z^*$  embeds the  $\lambda$ -bit values of  $(x_1, x_2, \dots, x_n)$ . The LHP.BatchPSol algorithm proceeds to solve the resultant puzzle  $Z^*$  in time  $\mathbf{T}$  to obtain the  $n$  values encoded as  $\lambda$ -bit values.

However, both of the above constructions only support bounded batching as they require the size of each puzzle  $Z_i$  (the  $v$  component) to scale linearly with the maximum batch size. Also, since the domain extension factor  $s$  has to be fixed at puzzle generation time, it determines an upper bound on the input size of the function  $f$  and therefore the number of puzzles we can batch solve later.

### 4.3 Unbounded Batching of TLPs

In this section we present a new TLP scheme with batched solving where overcome both drawbacks of the scheme above. Namely, our new construction allows for batching where the size of the time-lock puzzles output by puzzle generation algorithm are independent of the number of puzzles to be batched. As a consequence, we have unbounded batching (bounded above by a super polynomial  $2^{\omega(\log \lambda)}$ ) meaning that any polynomial number of puzzles can be batched with a fixed one-time setup.

Our construction uses the following ingredients:

- A linearly homomorphic TLP scheme  $\text{LHP} := (\text{LHP.PSetup}, \text{LHP.PGen}, \text{LHP.PSol}, \text{LHP.PEval})$  where the homomorphism is over  $\mathbb{Z}_q$ .
- A puncturable PRF  $(\text{Setup}, \text{KGen}, \text{Punc}, \text{F})$  denoted in short by  $\text{F}$ .
- An indistinguishable obfuscator  $i\mathcal{O}$  for circuits.
- A  $\gamma$ -almost key-homomorphic puncturable PRF  $(\overline{\text{Setup}}, \overline{\text{KGen}}, \overline{\text{Punc}}, \overline{\text{F}})$  (denoted in short by  $\overline{\text{F}}_i$ ) with key space  $\mathbb{Z}_q^n$  and where the noise bound is  $\gamma$  such that  $p = 2^{\omega(\log \lambda)} \cdot \gamma$ .

We set  $N = 2^{\omega(\log \lambda)}$  denote an upper bound on the number of participants.

Our construction  $(\text{PSetup}, \text{PGen}, \text{BatchPSol})$  is shown in Fig. 1. For simplicity we consider the messages encoded to be in  $\{0, 1\}$ , and argue that its straightforward to extend the construction for multiple bits.

A puzzle in our case is an obfuscation of the program  $\mathbf{P}$  which has the message  $m$ , a PRF  $\text{F}$  key  $k$ , and the public parameters  $\text{pp}$  hardwired in it. The program  $\mathbf{P}$  takes as input three values:  $n$  indicating number of puzzles to be batched,  $i$  “index” of the current puzzle and  $j$  “index” of other puzzles. It is important to note that the exact indices for each puzzle are only set later during batch solving. Let  $i$  be the symbolic index of the puzzle being generated now (whose concrete value will be set during batch solving). The program internally generates the PRF key  $\bar{k}$  for the key-homomorphic puncturable PRF which then is embedded inside the LHP puzzle  $Z$ . In case the indices  $i$  and  $j$  are the same, a ciphertext  $c$  is set to encrypt the message  $m$  using the value  $\overline{\text{F}}(\bar{k}, j)$  as the masking factor. In any other case,  $c$  encrypts 0 with  $\overline{\text{F}}(\bar{k}, j)$  as the masking factor. The program returns the puzzle  $Z$  and the ciphertext  $c$ .

The batch solving algorithm in the beginning, locally indexes the  $n$  puzzles in some order based on some rule (e.g., lexicographic ordering). We have them now ordered  $(Z_1, \dots, Z_n)$  where the  $i$ -th puzzle is an obfuscated program denoted by  $Z_i := \tilde{\mathbf{P}}_i$ . Then, for all  $i \in [n]$ , we execute the program

- PSetup( $1^\lambda, \mathbf{T}$ ):
  - Run  $\text{pp}_{\text{LHP}} \leftarrow \text{LHP.PSetup}(1^\lambda, \mathbf{T})$ .
  - Run  $\text{pp}_{\text{F}} \leftarrow \text{Setup}(1^\lambda)$  and  $\text{pp}_{\overline{\text{F}}} \leftarrow \overline{\text{Setup}}(1^\lambda)$ .
  - Return  $\text{pp} := (\text{pp}_{\text{LHP}}, \text{pp}_{\text{F}}, \text{pp}_{\overline{\text{F}}})$ .
- PGen( $\text{pp}, m$ ):
  - Generate  $k \leftarrow \text{KGen}(\text{pp}_{\text{F}})$ .
  - Define  $\mathbf{P}_{m,k,\text{pp}}(n, i, j)$  as the following circuit:<sup>a</sup>
    - \* Ensure  $i, j \in [n]$ .
    - \* Compute  $(\bar{r}, r) \leftarrow \text{F}(k, (n, i))$ .
    - \* Compute  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_{\overline{\text{F}}}; \bar{r})$ .
    - \* Compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .
    - \* If  $j = i$  set  $c = \bar{\text{F}}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
    - \* Else if  $j \leq n \leq N$  set  $c = \bar{\text{F}}(\bar{k}, j)$ .
    - \* Return  $(Z, c)$ .
  - Return  $\tilde{\mathbf{P}} := i\mathcal{O}(1^\lambda, \mathbf{P}_{m,k,\text{pp}})$ .
- BatchPSol( $\text{pp}, Z_1, \dots, Z_n$ ):
  - For each  $i \in [n]$ ,
    - \* Parse  $Z_i := \tilde{\mathbf{P}}_i$ .
    - \* For each  $j \in [n]$ , compute  $(Z_i^*, c_{i,j}) \leftarrow \tilde{\mathbf{P}}_i(n, i, j)$ .
    - \* Define  $\mathbf{c}_i = (c_{i,1}, \dots, c_{i,n}) \in \mathbb{Z}_p^n$ .
  - Set  $Z^* \leftarrow \text{LHP.PEval}(+, \text{pp}, Z_1^*, \dots, Z_n^*)$ .
  - Compute  $k^* \leftarrow \text{LHP.PSol}(\text{pp}, Z^*)$ .
  - Compute  $\mathbf{f}^* = (\bar{\text{F}}(k^*, 1), \dots, \bar{\text{F}}(k^*, n))$ .
  - Compute  $\mathbf{c}^* = \sum_{i=1}^n \mathbf{c}_i$ .
  - Return  $\mathbf{c}^* - \mathbf{f}^*$  rounded component-wise.

<sup>a</sup> The circuit is padded to the maximum size of the circuits among those defined in the security proof. We refer the reader to the end of this Section for a discussion on the size of this circuit.

**Fig. 1:** Our construction for TLP with unbounded batch solving.

$\tilde{\mathbf{P}}_i$  on values  $(n, i, j)$  for all  $j \in [n]$ . In the end we obtain a LHP time-lock puzzle  $Z_i^*$  and ciphertexts  $c_{i,j}$  for each  $j \in [n]$ . Recall that when  $i = j$  the program  $\tilde{\mathbf{P}}_i$  sets  $c_{i,j}$  to encrypt the message  $m_i$  (where  $m_i$  is the message inside puzzle  $Z_i$ ), and for all  $i \neq j$ , the ciphertext  $c_{i,j}$  encrypts 0. We then obtain a LHP puzzle  $Z^*$  by homomorphically adding the puzzles  $Z_i^*$  for all  $i \in [n]$  and solving  $Z^*$  returns a PRF key  $k^*$  of the key-homomorphic puncturable PRF. We retrieve the message  $m_j$  (for all  $j \in [n]$ ) by doing the following: (1) compute  $c_j^* = \sum_{i=1}^n c_{i,j}$ , (2) evaluate  $\bar{\text{F}}(k^*, j)$ , (3) set  $m_j$  as the rounding of  $(c_j^* - \bar{\text{F}}(k^*, j))$ . The correctness and security of our construction is formalized in the theorems below and the formal proofs are deferred to [App. B](#).

**Theorem 1.** *Let LHP be a linearly homomorphic TLP scheme where the homomorphism is over  $\mathbb{Z}_q$ , let F be a puncturable PRF, let  $i\mathcal{O}$  be an indistinguishable obfuscator for circuits and let  $\bar{\text{F}}$  be a  $\gamma$ -almost key-homomorphic puncturable PRF with key space  $\mathbb{Z}_q^n$  and where the noise bound is  $\gamma$*

such that  $p = 2^{\omega(\log \lambda)} \cdot \gamma$ . If all the above primitives are perfectly correct, then the TLP scheme with batch solving from Fig. 1 is perfectly correct.

**Theorem 2.** *Let LHP be secure against depth  $\mathbf{T}^\varepsilon(\lambda)$ -bounded adversaries with sub-exponential advantage,  $F$  be a sub-exponentially secure puncturable PRF,  $\bar{F}$  be a sub-exponentially secure  $\gamma$ -almost key-homomorphic puncturable PRF and  $i\mathcal{O}$  be a sub-exponentially secure indistinguishable obfuscator. Then, the construction from Fig. 1 is a secure time-lock puzzle with batch solving against all depth  $\mathbf{T}^\varepsilon(\lambda)$ -bounded adversaries.*

**Size of the obfuscated circuit.** Observe that at any point in the proof, we only hardwire information of size bounded by a fixed polynomial in the security parameter, and in particular independent of the number of parties. Since the size of the obfuscated circuit must be padded to the maximum size of the circuit at any point in the security proof, it follows that the overhead is also independent of the number of parties.

**Instantiations and Setup assumptions.** We can instantiate: the linearly homomorphic time-lock puzzle LHP with the Class group based scheme from [49]. We can instantiate the puncturable PRF  $F$  with the GGM based PRF [46,27], the  $i\mathcal{O}$  scheme with the scheme from [32], and the  $\gamma$ -almost key-homomorphic puncturable PRF with the scheme from [10]. Notice that the above instantiations, do not require trusted setups and therefore our TLP scheme does not require a trusted setup and instead only requires a one-time transparent public-coin setup (for the LHP from [49]).

**TLP runtime.** The runtime of PGen is dominated by the obfuscation of the circuit  $\mathbf{P}$  which is polynomial in  $\lambda$  and size of  $\mathbf{P}$ . Moreover, the size of  $\mathbf{P}$  is independent of the batch size  $n$  or the number of users. Thus, the total runtime is polynomial in  $\lambda$  and size of the time-locked message. The BatchPSol involves executing  $n$  obfuscated circuits, combining their outputs homomorphically using LHP.PEval, and solving the resulting TLP using LHP.PSol. The runtime of the first two operations is  $poly(\lambda, n)$ , whereas the last operation is  $poly(\lambda, T)$ .

**Verifiable TLPs.** We can support verifiability for our puzzles where the puzzle generator along with the puzzle also outputs a proof, that convinces a verifier that the puzzle is well-formed. In applications of our TLP scheme (including the ones in later sections), verifying whether a puzzle is in the support of the PGen is paramount for the correctness (Definition 13) of batched solving to hold. To provide such verifiability, we can add two new interfaces: PProve( $Z, m, r$ ) run by the puzzle generator that outputs a proof  $\pi$  to ascertain a puzzle  $Z$  is well-formed (with message  $m$  and randomness  $r$ ), and PVer( $Z, \pi$ ) run by a verifier that validates the proof w.r.t. the puzzle. In terms properties we want that a verifier shouldn't be convinced of a malformed puzzle and that the proof does not help in solving the puzzle any faster. For a formal definition and a discussion on concrete instantiations, we refer the reader to App. C.

**Non-committing TLPs.** A non-committing TLP lets a simulator generate a puzzle first and later “explain” the puzzle as committing to a message  $m$  by opening it to reveal  $m$ . Note that a TLP is committing to the message once the puzzle is generated. Cohen et al. [15] showed a generic approach to build such non-committing TLPs in the programmable random oracle model (PROM) and we can transform our TLP scheme into one that is non-committing in the same way. The idea is to run  $Z' \leftarrow \text{PGen}(\text{pp}, x)$  for some random  $r$ , and the final puzzle  $Z$  is set as  $Z := (Z', c)$  where  $c := H(r) \oplus m$ . The simulator when required to equivocate  $Z$  as a puzzle embedding the message  $m$ , sets the value  $H(r) := c \oplus m$  on the fly, as  $H(\cdot)$  is modeled as a PROM. We can modify the

construction from §4.3 by having letting PGen run as before, but output  $(\tilde{\mathbf{P}}, H(r) \oplus m)$  as the final puzzle, where  $\tilde{\mathbf{P}} := i\mathcal{O}(\lambda, \mathbf{P}_{r,k,pp})$ . Specifically this means that the PROM computation is outside the  $i\mathcal{O}$ .

## 5 Application 1: Byzantine Broadcast

In this section, we present our generic efficiency preserving compiler to transform any Byzantine Broadcast (BB) protocol secure against weakly adaptive adversaries to one that is secure against strongly adaptive adversaries.

### 5.1 Model and Definitions

In our setting, there are  $n$  nodes, numbered 1 to  $n$ , running a distributed protocol where the identity of each node is known to one another through a PKI.

**Communication model.** We consider the standard synchronous model of communication where there is a known upper bound on the message delay ( $\Delta$ ). Moreover, we also assume that each node has access to a shared global clock and all parties are connected by a pairwise reliable channel. Also, each node has access to the functionalities: **RECEIVE** to receive messages and **SEND** to send messages to any recipient.

**Synchronous protocol.** We consider the protocols that are executed in a round-based fashion, where the duration of each round is  $\Delta$ . Any message sent by an honest node in a round reaches all other honest nodes by the beginning of the next round. Moreover, we also assume that each round is sufficiently long enough to perform standard cryptographic operations except **BatchPSol** and **PSol**.

Let  $\mathcal{P}$  be the set of possible internal states of a node and  $\mathcal{M}$  be the set of possible messages that can be sent and received by a node.

**Definition 14 ( $\Delta$ -secure Synchronous protocol).** Let  $\mathcal{F}_n$  denote the family of transition functions such that:

$$\mathcal{F}_n = \{f_{r,u} : \mathcal{P} \times \mathcal{M}^n \rightarrow \mathcal{P} \times \mathcal{M}^n : u \in [n], r \in \mathbb{Z}\}$$

A synchronous protocol  $\Pi_{\text{sync}}$  is executed by  $n$  nodes and proceeds in rounds. In every round  $r$ , every node  $u \in [n]$ , reads the messages addressed to it using the **RECEIVE** functionality, updates its state and computes the messages to be sent using  $f_{r,u}$ , and sends the messages to intended recipients using the **SEND** functionality.

#### Protocol $\Pi_{\text{sync}}(\lambda, \Delta)$

**Setup.**

- Let  $S_{0,u}$  be the initial state of node  $u \in [n]$
- Generate and publish public parameters

**Protocol.** A node  $u \in [n]$ , for each round  $r$ :

- Fetch messages from each sender:  $m := (m_1, \dots, m_n) \leftarrow \text{RECEIVE}()$
- Compute next state and messages:  $(S_{r+1,u}, m' := (m'_1, \dots, m'_n)) \leftarrow f_{r,u}(S_{r,u}, m)$
- Send messages: **SEND**( $m'$ , recipients)

**Adversary model.** The adversary can make at most  $t$  out of  $n$  nodes to arbitrarily deviate from the protocol execution, where  $t < n$ . Moreover, we assume that the adversary controls the delivery of all the messages in the network.

- We consider a *strongly adaptive adversary* that can corrupt nodes on the fly and perform *after-the-fact* removal.
- Whereas, a *weakly adaptive adversary* can *only* corrupt nodes on the fly, but cannot prevent the delivery of any message that was already sent.

Additionally, we consider a *rushing* adversary that can inspect the messages sent by any honest node before delivering it to the other nodes. Moreover, we assume that honest nodes can irrecoverably erase (part of) its state and memory at any time.

**Computational model.** All honest nodes are sequential, random access PPT, but the adversary is a non-uniform probabilistic parallel machine with polynomially bounded parallelism running in polynomially bounded parallel steps.

## 5.2 Protocol

For an  $n$  node protocol  $\Pi$ , we define a deterministic function called *output derivation function* for each node  $u \in [n]$ . This function allows a node to compute its output bit for  $\Pi$  based on the transcript of public messages exchanged by the participants and public parameters.

**Definition 15 (Output derivation function).** *Let  $\Pi$  be an  $n$  node protocol and  $\mathcal{Y}$  denote the public transcript space of the protocol  $\Pi$ , then  $\mathcal{G}_n$  denote the family **output derivation functions** such that:*

$$\mathcal{G}_n = \{g_u : \mathcal{Y} \rightarrow \{0, 1\} : u \in [n]\}$$

*Functions in  $\mathcal{G}_n$ , despite being deterministic, may not be efficiently computable without a party's keys.*

We recall the definition of a secure Byzantine broadcast protocol below.

**Definition 16 (( $\Delta, t$ )-secure Byzantine broadcast).** *Let  $\lambda$  be the security parameter,  $\Delta$  be the known upper-bound on the network delay, and node  $d \in [n]$  be the designated sender. A protocol  $\Pi$  executed by  $n$  nodes with specified family of functions  $\mathcal{G}_n$ , where the designated sender holds an input bit  $b \in \{0, 1\}$ , is a ( $\Delta, t$ )-secure broadcast protocol tolerating at most  $t$  corruptions if it satisfies the following properties with probability  $1 - \text{negl}(\lambda)$ :*

- *Consistency: If two honest nodes output bit  $b_i$  and  $b_j$  respectively, then  $b_i = b_j$ .*
- *Validity: If the designated sender is honest, then every honest node outputs the designated sender's input bit  $b$ .*
- *Termination: Every honest node  $u$  outputs a bit from  $g_u(\text{transcript})$ , where transcript is the transcript from running  $\Pi$ .*

*If the protocol can tolerate corruptions by a strongly adaptive and a weakly adaptive adversary, then it is **strongly adaptive** ( $\Delta, t$ )-secure and **weakly adaptive** ( $\Delta, t$ )-secure, respectively.*

Let  $\Pi_{\text{bb-wa}}$  be a weakly adaptive protocol, we formally describe  $\Pi_{\text{bb-sa}}$  below:

**Protocol  $\Pi_{\text{bb-sa}}(\lambda, \Delta, \Pi_{\text{bb-wa}}, \mathcal{G}_n)$**

Text in gray indicates the instructions from  $\Pi_{\text{bb-wa}}$ .

**Setup.**



- Let  $S_{0,u}$  be the initial state of node  $u \in [n]$
- For each round  $r$ ,  $\text{pp} \leftarrow \text{PSetup}(1^\lambda, \Delta)$
- Generate and publish public parameters

**Input.**

- Let  $b \in \{0, 1\}$
- If designated sender,  $d$ , then  $S_{0,d} := S_{0,d} \cup b$

**Protocol.** A node  $u \in [n]$ , for each round  $r$ :

- Fetch messages from each sender:  $m := (m_1, \dots, m_n) \leftarrow \text{RECEIVE}()$
- Parse message  $m_v$  as puzzle  $Z_v$ , ciphertext  $C_v$ , proof of well-formed  $\pi_v$  for all  $v \in [n]$
- Check  $\pi_v$ 's to verify if  $Z_v$ 's are well-formed by  $\text{PVer}(\text{pp}, Z_v, \pi_v)$
- Extract the individual solutions  $(s_1, \dots, s_n) \leftarrow \text{BatchPSol}(\text{pp}, Z_1, \dots, Z_n)$
- Decrypt  $C_v$ 's, set  $m_v := C_v \oplus H(s_v)$  for all  $v \in [n]$ , and  $m := (m_1, \dots, m_n)$
- Set internal state for round  $r$  as  $S_{r,u} := m_u$
- Compute next state and messages:  $(S_{r+1,u}, m' := (m'_1, \dots, m'_n)) \leftarrow f_{r,u}(S_{r,u}, m)$
- Pick  $s \in \mathcal{S}$ ,  $Z \leftarrow \text{PGen}(\text{pp}, s)$ , and compute  $\pi$  to prove that  $Z$  is well-formed.
- Reassign  $m'_u := (Z, S_{r+1,u} \oplus H(s), \pi)$  and  $m'_v := (Z, m'_v \oplus H(s), \pi)$  for all  $v \in [n] \setminus \{u\}$
- Set output messages as  $m' := (m'_1, \dots, m'_n)$  and erase  $S_{r+1,u}$ ,  $s$ ,  $\pi$
- Send messages:  $\text{SEND}(m', \text{recipients})$

**Output.**

- Let **transcript** be the public transcript of the protocol execution
- Return  $b \leftarrow g_u(\text{transcript})$

**Theorem 3.** Let  $\Pi_{\text{bb-wa}}$  be a weakly adaptive  $(\Delta, t)$ -secure Byzantine broadcast protocol with output derivation functions  $\mathcal{G}_n$  and  $\Pi_{\text{bb-sa}}$  be the compiled strongly adaptive  $(\delta, t)$ -secure protocol with output derivation functions  $\mathcal{G}_n$ , such that  $\Delta = 2\delta$ . If an  $\mathcal{A}$  violates  $\Pi_{\text{bb-sa}}$  with probability at least  $p$ , then there exists an adversary  $\mathcal{B}$  that violates  $\Pi_{\text{bb-wa}}$  with probability at least  $p$ .

**Analysis.** Suppose  $\exists$  an  $\mathcal{A}$  that can break  $\Pi_{\text{bb-sa}}$ , then we build another adversary  $\mathcal{B}$  that breaks  $\Pi_{\text{bb-wa}}$ . At a high level, we show that every attack by  $\mathcal{A}$  on  $\Pi_{\text{bb-sa}}$  can be translated to an attack on  $\Pi_{\text{bb-wa}}$ . Observe that  $\mathcal{B}$  is as powerful as  $\mathcal{A}$ , except  $\mathcal{B}$  cannot perform *after-the-fact* removal. Thus, to translate the *after-the-fact* removal,  $\mathcal{B}$  must know whether  $\mathcal{A}$  delivers or removes messages in  $\Pi_{\text{bb-sa}}$ .  $\mathcal{B}$  can know this only by waiting for  $\delta$  steps to see  $\mathcal{A}$ 's actions! Hence,  $\mathcal{B}$  starts the simulation  $\delta$  steps ahead of  $\Pi_{\text{bb-wa}}$ . But, when the simulation begins,  $\mathcal{B}$  doesn't yet have the real-world messages from  $\Pi_{\text{bb-wa}}$  that *can be copied* to  $\Pi_{\text{bb-sa}}$ . So  $\mathcal{B}$  sends non-committing TLPs to equivocate the contents of the puzzle (possible because of PROM). When  $\mathcal{A}$  solves the TLP and queries the RO, actual messages from  $\Pi_{\text{bb-wa}}$  will be available, and  $\mathcal{B}$  programs the RO to open the corresponding message from  $\Pi_{\text{bb-wa}}$ . Since the duration between when the messages are sent and the contents learned by the honest nodes should be the same in the simulation and the real-world, we set  $\Delta = 2\delta$ . Thus, asymptotically,  $\Pi_{\text{bb-sa}}$  is round preserving (as  $\Delta = 2\delta$ ) and communication preserving (due to compactness of our TLP). We present the detailed analysis in [App. D](#).


**Expected constant-round Byzantine broadcast.** Wan et al. proposed an expected constant round BB protocol under a weakly adaptive and dishonest majority setting [51]. Thus, using the compiler presented in the §5.2, we can obtain resilience in the strongly adaptive setting!

## 6 Application 2: Nakamoto Consensus Secure Against a Mobile Sluggish Adversary

In this section, we first show an attack against the Nakamoto consensus in the mobile sluggish model and later show how to secure the Nakamoto consensus using our TLP.

### 6.1 Attack on Nakamoto Consensus in the Mobile Sluggish Model.

In Nakamoto consensus, a chain forks when two distinct blocks extend the same parent block. Forks are inherently bad for security as it splits the honest mining efforts across the two branches of the tree. A benign example is when two blocks are mined less than  $\Delta$  time apart. Since the messages take  $\Delta$  to reach others, the winner of the second block would not have been aware of the previous block. Nakamoto consensus is parameterized in a way that the inter-arrival time between two blocks is much longer than the time to transmit between any two farthest nodes in the system. The security threat posed by forks is the exact reason the Nakamoto consensus is secure only in the synchronous model.



**Fig. 2:** Double spend attack: This plot depicts average block arrival times. Assuming 52 honest nodes (51 prompt + 1 sluggish) and 48 adversarial nodes, the average inter-arrival time of honest blocks and adversarial blocks in Bitcoin is 19.2 and 20.8 minutes, respectively. Observe that over  $19.2 \times 2$  minutes, even though the honest nodes have mined two blocks, due to sluggishness, the honest chain has grown only by one block.

The mobile sluggish model allows a subset of messages to be delayed for an arbitrary time. Since the sluggishness can be mobile, any node can be sluggish during the protocol execution. Thus, whenever an honest node mines a block, the adversary can simply delay the block propagation until another block extends the same header (see Fig. 2). At this point, the adversary can release both the blocks simultaneously to split the honest mining efforts. The adversary can sustain the forks as long as it has sufficient sluggish budget. Since the adversary is responsible for message delivery and sluggishness can be mobile, it could perform this attack repeatedly. In the meantime, adversarial nodes will continue to extend their chain in private. Using this strategy, even a single mobile sluggish fault has the ability to reduce the honest mining rate by *half*! Thus, an honest majority assumption may not be sufficient to get consistency and liveness in this model. In App. E, we detail the attack illustrated in Fig. 2.

### 6.2 Model

Let  $n$  be the total number of nodes,  $d$  be the maximum number of sluggish nodes, and  $t$  be the maximum number of adversarial nodes. Thus, there are at least  $n - t$  honest nodes and at least  $n - d - t$  prompt nodes.

**Sluggish network model.** We assume that the time proceeds in rounds. Moreover, we assume that the adaptivity of the adversary is static. That is before the protocol execution, the adversary

picks the set of nodes to corrupt. Moreover, we also assume that every node has access to a *shared global clock* and a pairwise reliable channel between any two parties.

The standard (lock-step) model of synchrony assumes that any message sent in round  $r$  reaches other nodes by  $r + 1$ . We consider a generalization of this model called the *mobile sluggish* model. In this model, if a node is prompt at round  $r$ , then any message sent by the node in round  $\leq r$  reaches all the nodes that are prompt in round  $r + 1$  by round  $\leq r + 1$ . Due to mobility of the sluggishness, set of prompt nodes in any two adjacent rounds need not be the same.

The adversary is responsible for message delivery. Thus, an adversary can reorder or delay messages (according to prompt and sluggish delay requirement), but *cannot* delete messages. Moreover, any message sent to a prompt node *by a prompt or an adversarial node* reaches all prompt nodes. We can relax this assumption by assuming that nodes gossip/echo any message they receive [41, See footnote 4]. The adversary inspects all messages (including puzzles and blocks) first before delivering to any node.

**Round duration.** We assume that the duration of a round is  $O(\Delta)$ . Specifically, we assume that a round is sufficiently long enough to send/receive messages and perform cryptographic operations (such as verifying a hash of a message, generating and verifying a zero-knowledge proof of well-formedness of a TLP, computing PGen/PEval, and signing and verifying a signature), *except* PSol, BatchPSol, and RO invocations to mine a block or a decoy.

**Computational model.** We adopt the *flat* model of computation introduced by Garay et al. [21]. In this model, all nodes are assumed to have the same computational power. Moreover, any node can make at most  $q$  proofs-of-work invocations to the RO in a round. Thus, the adversary can perform  $t \cdot q$  RO queries in each round. We remark that each node has an unlimited number of proof-of-work verification queries to the RO [21].

Additionally, we assume that all honest nodes are sequential, random access PPT, but the adversary is a non-uniform probabilistic parallel machine with polynomially bounded parallelism running in polynomially bounded parallel steps.

### 6.3 Protocol

**Super-round.** Since our protocol proceeds in two phases: (1) Mine phase ( $m$  rounds) and (2) Solve phase ( $D$  rounds), a super-round consists of a mine phase followed by a solve phase. Thus, the duration is  $(m + D)$  rounds.

#### Mobile Sluggish Nakamoto Protocol

**Input.**

- pp, TLP public parameters with  $\mathbf{T}$  as one round
- $m$ , duration of mine phase
- $D$ , duration of solve phase
- $q$ , maximum number of RO queries per round
- $T$ , difficulty threshold to mine a block
- $T_c$ , difficulty threshold to mine a decoy where  $T < T_c$

**Initialize.** Chain  $C$  containing agreed-upon genesis block  $C[0]$

**Protocol.** Every super-round  $R$  (which consists of  $(m + D)$  rounds)

- Get the payload from the environment

- Let  $h_{-1} := \text{H}(\text{C}[-1])$  be the hash of the last block on the longest chain  $\text{C}$
- Let  $B = \perp$  be an empty block
- For  $m$  rounds of **mine phase**:
  - For  $q$  RO queries:
    - \* Pick random  $\eta \in \{0, 1\}^\lambda$  and compute  $h := \text{H}(h_{-1}, \text{payload}, \eta)$
    - \* If  $h < T_c$  (mined a decoy)
      - Overwrite  $B := (h_{-1}, \text{payload}, \eta)$
    - \* If  $h < T$  (mined a block)
      - Overwrite  $B := (h_{-1}, \text{payload}, \eta)$
      - Set  $\text{C} := \text{C}||B$
      - Break out of the  $q$  and  $m$  loop
- If  $B \neq \perp$ 
  - Compute the TLP  $Z := \text{PGen}(\text{pp}, B)$
  - Compute proof of well-formed  $\pi := \text{PProve}(\text{pp}, Z, B)$
- **Solve phase** for  $D = 2$  rounds:
  - First round, multicast  $(Z, \pi)$  (if one exists), receive all the  $w$  puzzles from the network  $Z_1, \dots, Z_w$ , and check their well-formedness.
  - Second round, batch solve  $(s_1, \dots, s_w) := \text{BatchPSol}(Z_1, \dots, Z_w)$ .
- Update the chain  $\text{C}$  based on output from the solve phase

**Assumptions.** Let a block mined in a super-round  $R$  be a *prompt block*, if mined by an honest node and the node was prompt *at the beginning* of solving phase of *both*  $R - 1$  and  $R$ . Moreover, let  $f$  be the probability of one or more prompt blocks were mined in a super-round,  $c$  be the probability of every honest node mining at least one decoy in a super-round,  $\varepsilon, \delta \in (0, 1)$  be parameters, and  $p$  be the probability of a RO query mining a block. Our analysis assumes that:

$$\frac{(m + D)t + md}{cm(n - 2d - t)} \leq (1 - \delta) \quad (1) \qquad pqm(n - 2d - t) < 1/2 \quad (3)$$

$$\varepsilon + f < \delta/3 \quad (2) \qquad \frac{2\varepsilon}{1 - \varepsilon} < \delta^2 \quad (4)$$

**Analysis.** At a high level, our analysis extends the formal tools proposed by Garay et al. [21]. But there are several differences due to mobile sluggish faults and the use of TLPs:

1. The adversary can deviate from the protocol and invoke RO queries even during the solve phase. Intuitively, Eq. 1 quantifies the required advantage of the prompt nodes over sluggish and adversarial nodes for our protocol to be secure. Specifically, the numerator captures the computational advantage enjoyed by the adversarial nodes due to additional RO queries during the solve phase (the term  $(m + D)t$ ) and the loss in honest mining efforts due to sluggish nodes (the term  $md$ ). Large values of  $D$  decreases  $t$  (assuming other values can remain the same). But, due to the batch solving property of our TLP,  $D = 2$  in our protocol. Thus, the impact of  $D$  is minimal.
2. The mobility of the sluggishness provides the adversary timing based opportunities to reduce the contributions to the “prompt” chain. The adversary with  $d$  sluggish budget can toggle the sluggishness of  $2d$  nodes. If the adversary toggles the sluggishness when the honest nodes release TLPs at the end of the mining phase, it can reduce the number of nodes contributing to the prompt chain to  $(n - 2d - t)$ . This is because the  $d$  nodes that are sluggish through the mining

phase of a super-round may not be mining on the longest chain, and at the end of the mining phase, the adversary can use its mobility to make  $d$  prompt node sluggish (See Remark 1).

3. The sluggish nodes can inadvertently contribute to the adversarial chain. This is because the sluggish nodes may only have access to the view provided to them by the adversary.
4. Coordinated release of TLPs. Observe that from Eq. 3, large values of  $m$  decreases  $p$ , thus reducing the block arrival frequency. But, a bounded  $p$  ensures that the honest nodes do not fork one another and there are sufficient “convergence opportunities” to resolve forks [21,41]. Moreover, no prior permissionless protocol is secure under mobile-sluggish faults even under reduced performance.
5. Impact of decoys. In Eq. 1, the security impact of mining decoys by honest nodes is captured by  $c$ . We set the probability of mining a decoy such that honest nodes can mine sufficiently many decoys while simultaneously bounding the total number of decoys mined. Recall that our batch solvable TLP allows simultaneously opening a polynomial number of puzzles.

Notice that our analysis is a generalization of [21], thus by substituting  $m = 1, c = 1, d = 0$ , and  $D = 0$ , our analysis, in principle, collapses to [21]’s analysis. We prove liveness and consistency by assuming that the mining-hardness parameter is appropriately set in Eqs. 1 to 4. We present a complete analysis of the protocol in App. F.

## References

1. Time-lock: Block producer extractable value - tezos, <https://tezos.gitlab.io/alpha/timelock.html>
2. Abraham, I., Chan, T.H.H., Dolev, D., Nayak, K., Pass, R., Ren, L., Shi, E.: Communication Complexity of Byzantine Agreement, Revisited. In: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing
3. Abraham, I., Malkhi, D., Nayak, K., Ren, L., Yin, M.: Sync HotStuff: Simple and Practical Synchronous State Machine Replication. In: 2020 IEEE Symposium on Security and Privacy (SP)
4. Bagaria, V., Kannan, S., Tse, D., Fanti, G., Viswanath, P.: Prism: Deconstructing the blockchain to approach physical limits. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security
5. Barak, B., Canetti, R., Lindell, Y., Pass, R., Rabin, T.: Secure Computation Without Authentication. In: Advances in Cryptology – CRYPTO 2005
6. Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., Virza, M.: Snarks for c: Verifying program executions succinctly and in zero knowledge. In: Annual cryptology conference. pp. 90–108
7. Bitansky, N., Garg, S., Lin, H., Pass, R., Telang, S.: Succinct randomized encodings and their applications. In: Proceedings of the forty-seventh annual ACM symposium on Theory of Computing. pp. 439–448
8. Bitansky, N., Goldwasser, S., Jain, A., Paneth, O., Vaikuntanathan, V., Waters, B.: Time-lock puzzles from randomized encodings. In: Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science. pp. 345–356
9. Boneh, D., Naor, M.: Timed Commitments. In: Advances in Cryptology — CRYPTO 2000
10. Brakerski, Z., Vaikuntanathan, V.: Constrained key-homomorphic prfs from standard lattice assumptions. In: Theory of Cryptography Conference. pp. 1–30
11. Burdges, J., Feo, L.D.: Delay encryption. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 302–326
12. Chan, B.Y., Shi, E.: Streamlet: Textbook Streamlined Blockchains. In: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies
13. Chan, T.H.H., Pass, R., Shi, E.: Sublinear-Round Byzantine Agreement Under Corrupt Majority. In: Public-Key Cryptography – PKC 2020
14. Chen, H.C., Deviani, R.: A Secure E-Voting System Based on RSA Time-Lock Puzzle Mechanism. In: 2012 Seventh International Conference on Broadband, Wireless Computing, Communication and Applications
15. Cohen, R., Garay, J., Zikas, V.: Adaptively Secure Broadcast in Resource-Restricted Cryptography. Cryptology ePrint Archive, Report 2021/775

16. Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of paillier’s probabilistic public-key system. In: International workshop on public key cryptography. pp. 119–136
17. Dolev, D., Strong, H.R.: Authenticated Algorithms for Byzantine Agreement. *SIAM Journal on Computing*
18. Dwork, C., Naor, M.: Zaps and their applications. In: Proceedings 41st Annual Symposium on Foundations of Computer Science
19. Feldman, P., Micali, S.: An Optimal Probabilistic Protocol for Synchronous Byzantine Agreement. *SIAM Journal on Computing*
20. Fitzi, M., Nielsen, J.B.: On the Number of Synchronous Rounds Sufficient for Authenticated Byzantine Agreement. In: Distributed Computing
21. Garay, J., Kiayias, A., Leonardos, N.: The Bitcoin Backbone Protocol: Analysis and Applications. In: Advances in Cryptology - EUROCRYPT 2015
22. Garay, J.A., Katz, J., Koo, C.Y., Ostrovsky, R.: Round Complexity of Authenticated Broadcast with a Dishonest Majority. In: 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)
23. Garay, J.A., Katz, J., Kumaresan, R., Zhou, H.S.: Adaptively secure broadcast, revisited. In: Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing - PODC '11
24. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate Indistinguishability Obfuscation and Functional Encryption for all Circuits. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science
25. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. In: Proceedings of the forty-fifth annual ACM symposium on Theory of computing. pp. 467–476
26. Gilad, Y., Hemo, R., Micali, S., Vlachos, G., Zeldovich, N.: Algorand: Scaling Byzantine Agreements for Cryptocurrencies. In: Proceedings of the 26th Symposium on Operating Systems Principles
27. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)
28. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for nizk. In: Annual International Cryptology Conference. pp. 97–111
29. Guo, Y., Pass, R., Shi, E.: Synchronous, with a Chance of Partition Tolerance. In: Advances in Cryptology – CRYPTO 2019
30. Hanke, T., Movahedi, M., Williams, D.: DFINITY Technology Overview Series, Consensus System
31. Hirt, M., Zikas, V.: Adaptively Secure Broadcast. In: Advances in Cryptology – EUROCRYPT 2010
32. Jain, A., Lin, H., Sahai, A.: Indistinguishability obfuscation from well-founded assumptions. In: Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing. pp. 60–73
33. Katz, J., Loss, J., Xu, J.: On the Security of Time-Lock Puzzles and Timed Commitments. In: Theory of Cryptography
34. Kiffer, L., Rajaraman, R., Shelat, A.: A Better Method to Analyze Blockchain Consistency. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security
35. Kim, J., Mehta, V., Nayak, K., Shrestha, N.: Making Synchronous BFT Protocols Secure in the Presence of Mobile Sluggish Faults. *Cryptology ePrint Archive, Report 2021/603*
36. Li, J., Guo, D.: On Analysis of the Bitcoin and Prism Backbone Protocols
37. Lin, H., Pass, R., Soni, P.: Two-Round and Non-Interactive Concurrent Non-Malleable Commitments from Time-Lock Puzzles. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)
38. Liu, J., Jager, T., Kakvi, S.A., Warinschi, B.: How to build time-lock encryption. *Designs, Codes and Cryptography* pp. 2549–2586
39. Malavolta, G., Thyagarajan, S.A.K.: Homomorphic Time-Lock Puzzles and Applications. In: Advances in Cryptology – CRYPTO 2019
40. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: International conference on the theory and applications of cryptographic techniques. pp. 223–238
41. Pass, R., Seaman, L., Shelat, A.: Analysis of the Blockchain Protocol in Asynchronous Networks. In: Advances in Cryptology – EUROCRYPT 2017
42. Pass, R., Shi, E.: FruitChains: A Fair Blockchain. In: Proceedings of the ACM Symposium on Principles of Distributed Computing
43. Pass, R., Shi, E.: Rethinking Large-Scale Consensus. In: 2017 IEEE 30th Computer Security Foundations Symposium (CSF)
44. Ren, L.: Analysis of Nakamoto Consensus. *Cryptology ePrint Archive, Report 2019/943*
45. Rivest, R.L., Shamir, A., Wagner, D.A.: Time-Lock Puzzles and Timed-Release Crypto. Tech. rep.
46. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. *SIAM Journal on Computing*



47. Shrestha, N., Abraham, I., Ren, L., Nayak, K.: On the Optimality of Optimistic Responsiveness. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security
48. Thyagarajan, S.A.K., Bhat, A., Malavolta, G., Döttling, N., Kate, A., Schröder, D.: Verifiable timed signatures made practical. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security
49. Thyagarajan, S.A.K., Castagnos, G., Laguillaumie, F., Malavolta, G.: Efficient cca timed commitments in class groups. In: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security
50. Wan, J., Xiao, H., Devadas, S., Shi, E.: Round-Efficient Byzantine Broadcast Under Strongly Adaptive and Majority Corruptions. In: Theory of Cryptography
51. Wan, J., Xiao, H., Shi, E., Devadas, S.: Expected Constant Round Byzantine Broadcast Under Dishonest Majority. In: Theory of Cryptography
52. Wikipedia contributors: Moment-generating function — Wikipedia, the free encyclopedia (2019), [Online; accessed 02-Feb-2022]
53. Yu, H., Nikolic, I., Hou, R., Saxena, P.: OHIE: Blockchain Scaling Made Simple. In: 2020 IEEE Symposium on Security and Privacy (SP)
54. Zhao, J., Tang, J., Li, Z., Wang, H., Lam, K.Y., Xue, K.: An Analysis of Blockchain Consistency in Asynchronous Networks: Deriving a Neat Bound. In: 2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS)

## A More Related Work

Dolev and Strong showed that it is possible to have a deterministic Byzantine Broadcast (BB) protocol that terminates in  $t + 1$  rounds even if  $t < n$  nodes are corrupt assuming PKI and a synchronous model of communication [17]. Since the seminal work by Dolev and Strong, numerous works have studied the round complexity of BB under various settings and assumptions. Notably, these approaches use randomized protocols [19] to overcome the linear round complexity barrier from [17].

Under the dishonest majority setting, prior works in BB can be categorized into: (1) weakly adaptive protocols or (2) strongly adaptive protocols. In the weakly adaptive setting, Garay et al. [23] showed the feasibility of BB in the atomic message delivery. Later, Chan et al. proposed the first expected sublinear-round complexity protocol [13]. Their protocol takes polylogarithmic (in security parameter) rounds. Subsequently, Wan et al. proposed the state-of-the-art protocol, which takes an expected constant number of rounds to achieve BB [51].

In the strongly adaptive setting, Hirt and Zikas showed that it is impossible to achieve BB assuming just PKI [31]. Garay et al. [22] and Fitzi et al. proposed an expected constant-round protocol when  $t = n/2 + O(1)$  [20]. Unfortunately, for any arbitrary  $t = n/2 + k$ , Garay et al. requires  $O(k^2)$  rounds [22] and Fitzi et al. requires  $O(k)$  rounds [20]. Thus, their protocols can provide expected constant-round complexity only in a narrow regime.

Subsequently, Wan et al. proposed the first expected sub-linear round protocol in the strongly adaptive setting using PKI and time-lock puzzles [50]. Their protocol takes polylogarithmic number of rounds even when the  $t = (1 - \varepsilon)n$ , where  $\varepsilon \in (0, 1)$ . Subsequently, Cohen, Garay, and Zikas [15, Theorem 5] explored the feasibility of *fair* broadcast in the strongly adaptive setting for both property-based definition and simulation-based definition under different assumptions (PKI, TLP, and RO). Moreover, their work uses a *non-committing time-lock encryption* primitive to overcome the BB impossibility for a simulation based definition.

## B Correctness and Security Analysis of Our TLP Construction

*Proof (Thm. 1).* We have for each  $i \in [n]$ , the tuple  $(Z_i^*, \mathbf{c}_i)$  where  $Z_i^* \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}_i; r_i)$  and  $\mathbf{c}_i = (c_{i,1}, \dots, c_{i,n})$  where for all  $j \neq i$ , we have  $c_{i,j} = \bar{F}(\bar{k}_i, j)$  and we have  $c_{i,i} = \bar{F}(\bar{k}_i, i) + m_i$ .

$\lceil p/2 \rceil \pmod{p}$ . By the correctness of LHP we have:

$$Z^* = \text{LHP.PGen}(\text{pp}, \sum_{i=1}^n \bar{k}_i) = \text{LHP.PEval}(+, \text{pp}, Z_1^*, \dots, Z_n^*)$$

and  $k^* = \sum_{i=1}^n \bar{k}_i = \text{LHP.PSol}(\text{pp}, Z^*)$ . Let  $\mathbf{c}^* = (c_1^*, \dots, c_n^*)$  and we have

$$\begin{aligned} c_j^* &= \sum_{i=1}^n c_{i,j} = \sum_{i=1, i \neq j}^n c_{i,j} + c_{j,j} = \sum_{i=1, i \neq j}^n \bar{F}(\bar{k}_i, j) + \bar{F}(\bar{k}_j, j) + m_j \cdot \lceil p/2 \rceil \pmod{p} \\ &= \sum_{i=1}^n \bar{F}(\bar{k}_i, j) + m_j \cdot \lceil p/2 \rceil \pmod{p} \\ &\leq \bar{F}\left(\sum_{i=1}^n \bar{k}_i, j\right) + \gamma \cdot n + m_j \cdot \lceil p/2 \rceil \pmod{p} \quad (\text{key-homomorphism of } \bar{F}) \\ &= \bar{F}(k^*, j) + \gamma \cdot n + m_j \cdot \lceil p/2 \rceil \pmod{p} \quad \left(\text{where } k^* = \sum_{i=1}^n \bar{k}_i\right) \end{aligned}$$

By definition  $\mathbf{f}^* = (\bar{F}(k^*, 1), \dots, \bar{F}(k^*, n))$  and therefore

$$c_j^* - \bar{F}(k^*, j) = \gamma \cdot n + m_j \cdot \lceil p/2 \rceil \pmod{p}$$

Since  $n$  is at most a polynomial and  $p = 2^{\omega(\log \lambda)} \cdot \gamma$ , we have that  $\gamma \cdot n$  is always smaller than  $p/4$ . Therefore rounding of the value to 0 or 1, if  $m_j = 0$  we get  $0 := \text{Round}(\gamma \cdot n \pmod{p})$  and if  $m_j = 1$  we get  $1 := \text{Round}(\gamma \cdot n + \lceil p/2 \rceil \pmod{p})$ . □

*Proof (Thm. 2).* The proof follows by a hybrid argument, where we define a (sub-exponentially long) series of hybrids and we argue about the indistinguishability of the neighbouring distributions. The changes in the execution in a hybrid compared to the previous hybrid is highlighted in **red**.

Hyb $_{\ell}$  ( $0 \leq \ell \leq N$ ): In Hyb $_{\ell}$  we change the program that we obfuscate to the following circuit.

$\mathbf{P}_{m,k,\text{pp}}^{\ell}(n, i, j)$ :

- \* Ensure  $i, j \in [n]$ .
- \* Compute  $(\bar{r}, r) \leftarrow \mathbf{F}(k, (n, i))$ .
- \* Compute  $\bar{k} \leftarrow \mathbf{KGen}(\text{pp}_{\bar{F}}; \bar{r})$ .
- \* Compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .
- \* **If  $j = i$  and  $n \leq \ell$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .**
- \* Else if  $j = i$  set  $c = \bar{F}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
- \* Else if  $j \leq n \leq N$  set  $c = \bar{F}(\bar{k}, j)$ .
- \* Return  $(Z, c)$ .

Clearly Hyb $_N$  corresponds to the distribution where  $b = 0$  and Hyb $_0$  is the execution where  $b = 1$ . To prove indistinguishability between Hyb $_{\ell-1}$  and Hyb $_{\ell}$  (for  $1 \leq \ell \leq N$ ) we design the following sub-hybrids.

Hyb $_{\ell-1,\rho}$ , ( $\rho \in [\ell]$ ): Is defined as the same execution as Hyb $_{\ell-1}$  except that we further modify the obfuscated circuit.

$\mathbf{P}_{m,k,\text{pp}}^{\ell-1,\rho}(n, i, j)$ :

- \* Ensure  $i, j \in [n]$ .
- \* Compute  $(\bar{r}, r) \leftarrow F(k, (n, i))$ .
- \* Compute  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_{\bar{F}}; \bar{r})$ .
- \* Compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .
- \* If  $j = i$  and  $n \leq \ell - 1$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* **Else If  $j = i$  and  $i \leq \rho$  and  $n = \ell$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .**
- \* Else if  $j = i$  set  $c = \bar{F}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
- \* Else if  $j \leq n \leq N$  set  $c = \bar{F}(\bar{k}, j)$ .
- \* Return  $(Z, c)$ .

It is easy to see that hybrid  $\text{Hyb}_{\ell-1,0}$  is identical to  $\text{Hyb}_{\ell-1}$  and hybrid  $\text{Hyb}_{\ell-1,\ell}$  is identical to  $\text{Hyb}_{\ell}$ . We defer the indistinguishability between  $\text{Hyb}_{\ell-1,\rho-1}$  and  $\text{Hyb}_{\ell-1,\rho}$  to [Lemma 1](#). By a standard hybrid argument we have that hybrid  $\text{Hyb}_{\ell-1,0}$  and hybrid  $\text{Hyb}_{\ell-1,\ell}$  are indistinguishable for all depth  $\mathbf{T}^\varepsilon(\lambda)$ -bounded adversaries. This proves that  $\text{Hyb}_{\ell-1}$  and  $\text{Hyb}_{\ell}$  are indistinguishable and by a standard hybrid argument we have that hybrid  $\text{Hyb}_0$  and  $\text{Hyb}_N$  are indistinguishable for all depth bounded adversaries. This concludes the proof.  $\square$

**Lemma 1.** *Hybrid executions  $\text{Hyb}_{\ell-1,\rho-1}$  and  $\text{Hyb}_{\ell-1,\rho}$  are indistinguishable for any depth  $\mathbf{T}^\varepsilon$ -bounded adversary.*

*Proof.* Consider the following set of sub-hybrids:

$\text{Hyb}_{\ell-1,\rho-1}^1$ : Is defined as the same execution as  $\text{Hyb}_{\ell-1,\rho-1}$  except that value  $(\bar{r}^*, r^*) \leftarrow F(k, (\ell, \rho))$  is hardwired in  $\mathbf{P}_{m,k,\text{pp}}^{\ell-1,\rho-1}(n, i, j)$  and the PRF  $F$  is punctured at position  $(\ell, \rho)$  and the corresponding punctured key  $k_{(\ell,\rho)} \leftarrow \text{Punc}(k, (\ell, \rho))$  is hardwired in the program. The resultant program is defined below.

$\mathbf{P}_{\bar{r}^*, r^*, m, k_{(\ell,\rho)}, \text{pp}}^{\ell-1,\rho-1,1}(n, i, j)$ :

- \* Ensure  $i, j \in [n]$ .
- \* **If  $(n, i) \neq (\ell, \rho)$  then compute  $(\bar{r}, r) \leftarrow F(k_{(\ell,\rho)}, (n, i))$ , else if  $(n, i) = (\ell, \rho)$  then compute  $(\bar{r}, r) = (\bar{r}^*, r^*)$ .**
- \* Compute  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_{\bar{F}}; \bar{r})$ .
- \* Compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .
- \* If  $j = i$  and  $n \leq \ell - 1$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* **Else If  $j = i$  and  $i \leq \rho - 1$  and  $n = \ell$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .**
- \* Else if  $j = i$  set  $c = \bar{F}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
- \* Else if  $j \leq n \leq N$  set  $c = \bar{F}(\bar{k}, j)$ .
- \* Return  $(Z, c)$ .

The indistinguishability between  $\text{Hyb}_{\ell-1,\rho-1}$  and  $\text{Hyb}_{\ell-1,\rho-1}^1$  is computational and follows from the security of  $i\mathcal{O}$ . Notice that the only difference between the executions is the obfuscated circuits. The two circuits are functionally equivalent given the PRF  $F$  is functionality preserving under puncturing. Therefore a distinguisher between the two hybrids can be used to violate the security of the  $i\mathcal{O}$  scheme.

Hybrid  $\text{Hyb}_{\ell-1,\rho-1}^2$ : Is defined as the same execution as  $\text{Hyb}_{\ell-1,\rho-1}^1$  except that the value  $(\bar{r}^*, r^*)$  is sampled uniformly.

The indistinguishability between  $\text{Hyb}_{\ell-1,\rho-1}^1$  and  $\text{Hyb}_{\ell-1,\rho-1}^2$  is computational and follows from the pseudorandomness of the PRF  $F$ . Notice that the only difference between the hybrids is in how

$(\bar{r}^*, r^*)$  is generated (either using PRF  $F$  or sampled uniformly at random). Therefore a distinguisher between the hybrids can be used to violate the pseudorandomness at the punctured point of the PRF  $F$ .

$\text{Hyb}_{\ell-1, \rho-1}^3$  Is defined as the same execution as  $\text{Hyb}_{\ell-1, \rho-1}^2$  except now the puzzle

$Z' \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$  is hardwired in the program. The new program is defined below.

$\mathbf{P}_{\bar{r}^*, r, m, k(\ell, \rho), Z', \text{pp}}^{\ell-1, \rho-1, 3}(n, i, j)$

- \* Ensure  $i, j \in [n]$ .
- \* If  $(n, i) \neq (\ell, \rho)$  then compute  $(\bar{r}, r) \leftarrow F(k_{(\ell, \rho)}, (n, i))$ , else if  $(n, i) = (\ell, \rho)$  then compute  $(\bar{r}, r) = (\bar{r}^*, r^*)$ .
- \* Compute  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_{\bar{F}}; \bar{r})$ .
- \* **If  $(n, i) = (\ell, \rho)$  compute  $Z := Z'$ .**
- \* **Else compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .**
- \* If  $j = i$  and  $n \leq \ell - 1$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* Else If  $j = i$  and  $i \leq \rho - 1$  and  $n = \ell$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* Else if  $j = i$  set  $c = \bar{F}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
- \* Else if  $j \leq n \leq N$  set  $c = \bar{F}(\bar{k}, j)$ .
- \* Return  $(Z, c)$ .

The indistinguishability between  $\text{Hyb}_{\ell-1, \rho-1}^2$  and  $\text{Hyb}_{\ell-1, \rho-1}^3$  is computational and follows from the security of  $i\mathcal{O}$ . Notice that the only difference between the executions is the obfuscated circuits. The two circuits are functionally equivalent as the only difference is in the hardwiring of  $Z'$ . Therefore a distinguisher between the two hybrids can be used to violate the security of the  $i\mathcal{O}$  scheme.

$\text{Hyb}_{\ell-1, \rho-1}^4$ : Is defined as the same execution as  $\text{Hyb}_{\ell-1, \rho-1}^3$  except now the puzzle  $Z'$  embeds the all 0 string rather than the key  $\bar{k}$ . The new program is defined below.

$\mathbf{P}_{\bar{r}^*, r, m, k(\ell, \rho), Z', \text{pp}}^{\ell-1, \rho-1, 4}(n, i, j)$

- \* Ensure  $i, j \in [n]$ .
- \* If  $(n, i) \neq (\ell, \rho)$  then compute  $(\bar{r}, r) \leftarrow F(k_{(\ell, \rho)}, (n, i))$ , else if  $(n, i) = (\ell, \rho)$  then compute  $(\bar{r}, r) = (\bar{r}^*, r^*)$ .
- \* Compute  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_{\bar{F}}; \bar{r})$ .
- \* **If  $(n, i) = (\ell, \rho)$  compute  $Z := Z'$ .**
- \* Else compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .
- \* If  $j = i$  and  $n \leq \ell - 1$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* Else If  $j = i$  and  $i \leq \rho - 1$  and  $n = \ell$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* Else if  $j = i$  set  $c = \bar{F}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
- \* Else if  $j \leq n \leq N$  set  $c = \bar{F}(\bar{k}, j)$ .
- \* Return  $(Z, c)$ .

The indistinguishability between  $\text{Hyb}_{\ell-1, \rho-1}^3$  and  $\text{Hyb}_{\ell-1, \rho-1}^4$  is against depth  $\mathbf{T}^\epsilon$ -bounded adversaries. Notice that the only difference between the hybrids is in the puzzle  $Z'$  that is output by the respective circuits. This immediately reduces to the security of the linearly homomorphic time-lock puzzle LHP against depth  $\mathbf{T}^\epsilon$  bounded adversaries.

$\text{Hyb}_{\ell-1, \rho-1}^5$ : Is defined as the same execution as  $\text{Hyb}_{\ell-1, \rho-1}^4$  except now the value  $s^* \leftarrow \bar{F}(\bar{k}, \rho)$  and  $s^* + m \cdot \lceil p/2 \rceil$  is hardwired in the program where  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_{\bar{F}}; \bar{r}^*)$  and the corresponding key  $\bar{k}$

is punctured at the position  $\rho$  and the corresponding punctured key  $\bar{k}_\rho \leftarrow \overline{\text{Punc}}(\bar{k}, \rho)$  is hardwired in the program. The new program is defined below.

$\mathbf{P}_{\bar{r}^*, r, m, k(\ell, \rho), \bar{k}_\rho, s^* + m \cdot \lceil p/2 \rceil, \text{pp}}^{\ell-1, \rho-1, 5}(n, i, j)$

- \* Ensure  $i, j \in [n]$ .
- \* If  $(n, i) \neq (\ell, \rho)$  then compute  $(\bar{r}, r) \leftarrow \text{F}(k(\ell, \rho), (n, i))$ , else if  $(n, i) = (\ell, \rho)$  then compute  $(\bar{r}, r) = (\bar{r}^*, r^*)$ .
- \* **If  $(n, i) \neq (\ell, \rho)$  then compute  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_F; \bar{r})$ , else if  $(n, i) = (\ell, \rho)$  then compute  $\bar{k} = \bar{k}_\rho$ .**
- \* If  $(n, i) = (\ell, \rho)$  compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, 0^\lambda; r)$ .
- \* Else compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .
- \* If  $j = i$  and  $n \leq \ell - 1$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* Else If  $j = i$  and  $i \leq \rho - 1$  and  $n = \ell$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* **Else if  $j = i$  and  $i = \rho$  and  $n = \ell$  set  $c = s^* + m \cdot \lceil p/2 \rceil \pmod{p}$ .**
- \* Else if  $j = i$  set  $c = \bar{F}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
- \* Else if  $j \leq n \leq N$  set  $c = \bar{F}(\bar{k}, j)$ .
- \* Return  $(Z, c)$ .

To establish the indistinguishability between the two hybrids, we first establish that the punctured programs are, with all but negligible probability, functionally equivalent.

**Lemma 2.** *Assuming PRF  $\bar{F}$  is functionality computationally preserving against sub-exponential adversaries [10], there exists no position  $i^* \in [n] \setminus \rho$ , such that  $\bar{F}(\bar{k}, i^*) \neq \bar{F}(\bar{k}_\rho, i^*)$ , where  $\bar{k}_\rho \leftarrow \overline{\text{Punc}}(\bar{k}, \rho)$ .*

*Proof.* The functionality preservation property of PRF  $\bar{F}$  states that there exists a negligible function  $\text{negl}$ , for all sub-exponential adversary  $\mathcal{A}$ , and it holds that

$$\Pr \left[ \begin{array}{l} i^* \neq \rho \wedge \\ \bar{F}(k, i^*) \neq \bar{F}(\bar{k}_\rho, i^*) \end{array} : i^* \leftarrow \mathcal{A}^{\bar{F}(\bar{k}, \cdot)}(1^\lambda, \bar{k}_\rho) \right] \leq \text{negl}(\lambda)$$

where  $\bar{k}_\rho \leftarrow \overline{\text{Punc}}(\bar{k}, \rho)$  and  $\mathcal{A}$  gets to choose position  $\rho$  before obtaining  $\bar{k}_\rho$ .

Suppose there exists a point  $i^* \in [n] \setminus \rho$ , then since  $n \in [N]$  where  $N = 2^{\omega(\log \lambda)}$ , an adversary can query the oracle  $\bar{F}$ , at most  $n$  times (with queries  $(1, \dots, n)$ ) to find  $i^*$  since the adversary can run in sub-exponential time. Clearly the adversary violates the assumption with non-negligible probability. □

We now argue that if hybrids  $\text{Hyb}_{\ell-1, \rho-1}^4$  and  $\text{Hyb}_{\ell-1, \rho-1}^5$  are distinguishable, we can construct a reduction against the security of  $i\mathcal{O}$ . Notice that the only difference between the two hybrids is in the obfuscated programs. More specifically, in  $\text{Hyb}_{\ell-1, \rho-1}^5$ , we puncture the key  $\bar{k}$  at position  $\rho$  (if  $(n, i) = (\ell, \rho)$ ). By Lemma 2, we see that it is computationally infeasible to find a point  $i^* \in [n] \setminus \rho$  such that the evaluation of the PRF  $\bar{F}$  is different with the normal key  $\bar{k}$  and the punctured key  $\bar{k}_\rho$ . Therefore the behavior of the two circuits is functionally equivalent and we can reduce the indistinguishability of the hybrids to the security of  $i\mathcal{O}$  against sub-exponential adversaries.

$\text{Hyb}_{\ell-1, \rho-1}^6$ : Is defined as the same execution as  $\text{Hyb}_{\ell-1, \rho-1}^5$  except the value  $s^*$  is sampled uniformly from the output space of the PRF  $\bar{F}$  and is hardwired in the program. The new program is denoted by  $\mathbf{P}_{\bar{r}^*, r, m, k(\ell, \rho), \bar{k}_\rho, s^* + m \cdot \lceil p/2 \rceil, \text{pp}}^{\ell-1, \rho-1, 6}(n, i, j)$  which is syntactically the same circuit as defined for  $\text{Hyb}_{\ell-1, \rho-1}^5$ .

The indistinguishability between  $\text{Hyb}_{\ell-1,\rho-1}^5$  and  $\text{Hyb}_{\ell-1,\rho-1}^6$  is computational and follows from the pseudorandomness of PRF  $\bar{F}$ .

$\text{Hyb}_{\ell-1,\rho-1}^7$ : Is defined as the same execution as  $\text{Hyb}_{\ell-1,\rho-1}^6$  except that when  $j = i$  and  $i \leq \rho$  and  $n = \ell$ , we set  $c = s^* + 0$  in the program. The new program is defined below.

$\mathbf{P}_{\bar{r}^*, r, m, k(\ell, \rho), \bar{k}_\rho, s^*, \text{pp}}^{\ell-1, \rho-1, 7}(n, i, j)$

- \* Ensure  $i, j \in [n]$ .
- \* If  $(n, i) \neq (\ell, \rho)$  then compute  $(\bar{r}, r) \leftarrow \mathbf{F}(k(\ell, \rho), (n, i))$ , else if  $(n, i) = (\ell, \rho)$  then compute  $(\bar{r}, r) = (\bar{r}^*, r^*)$ .
- \* If  $(n, i) \neq (\ell, \rho)$  then compute  $\bar{k} \leftarrow \overline{\text{KGen}}(\text{pp}_{\bar{F}}; \bar{r})$ , else if  $(n, i) = (\ell, \rho)$  then compute  $\bar{k} = \bar{k}_\rho$ .
- \* If  $(n, i) = (\ell, \rho)$  compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, 0^\lambda; r)$ .
- \* Else compute  $Z \leftarrow \text{LHP.PGen}(\text{pp}_{\text{LHP}}, \bar{k}; r)$ .
- \* If  $j = i$  and  $n \leq \ell - 1$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* Else If  $j = i$  and  $i \leq \rho - 1$  and  $n = \ell$  set  $c = \bar{F}(\bar{k}, j) + 0 \pmod{p}$ .
- \* Else if  $j = i$  and  $i = \rho$  and  $n = \ell$  set  $c = s^* + 0 \pmod{p}$ .
- \* Else if  $j = i$  set  $c = \bar{F}(\bar{k}, j) + m \cdot \lceil p/2 \rceil \pmod{p}$ .
- \* Else if  $j \leq n \leq N$  set  $c = \bar{F}(\bar{k}, j)$ .
- \* Return  $(Z, c)$ .

The executions in both hybrids  $\text{Hyb}_{\ell-1,\rho-1}^6$  and  $\text{Hyb}_{\ell-1,\rho-1}^7$  are identically distributed as the only change is when  $j = i$ ,  $i = \rho$  and  $n = \ell$  and given  $s^*$  is a string chosen uniformly at random.

Notice that for  $i \leq \rho$  and  $n \leq \ell$ , the ciphertext  $c$  encrypts  $b = 0$ . This is almost what we have in the execution of  $\text{Hyb}_{\ell-1,\rho}$ . If we keep what we have in hybrid  $\text{Hyb}_{\ell-1,\rho-1}^7$  and undo the changes starting from  $\text{Hyb}_{\ell-1,\rho-1}^6$  to  $\text{Hyb}_{\ell-1,\rho-1}^1$ , we end up with an execution where the obfuscated program is exactly  $\mathbf{P}_{m, k, \text{pp}}^{\ell-1, \rho}(n, i, j)$  from the hybrid  $\text{Hyb}_{\ell-1,\rho}$ . The same indistinguishability arguments hold while doing the undo process over these hybrids. This implies the indistinguishability between the hybrids  $\text{Hyb}_{\ell-1,\rho-1}$  and  $\text{Hyb}_{\ell-1,\rho}$ . This concludes the proof of the Lemma.  $\square$

## C Verifiable TLPs

**Definition 17 (Verifiable TLP with Batch Solving).** A TLP scheme with batch solving  $(\text{PSetup}, \text{PGen}, \text{PSol}, \text{BatchPSol})$  is verifiable with the aid of the following interfaces:

- $\pi \leftarrow \text{PProve}(\text{pp}, Z, m, r)$ : a PPT algorithm that takes as input the public parameters  $\text{pp}$ , a puzzle  $Z$ , a message  $m$ , and the randomness  $r$  used in puzzle generation, and outputs a proof  $\pi$ .
- $0/1 \leftarrow \text{PVer}(\text{pp}, Z, \pi)$ : a deterministic algorithm that takes as input the public parameters  $\text{pp}$ , a puzzle  $Z$  and a proof  $\pi$ , and outputs 1 if the puzzle is well-formed according to  $\pi$ , and 0 otherwise.

**Definition 18 (Verifiability).** A verifiable TLP scheme with batch solving  $(\text{PSetup}, \text{PGen}, \text{PSol}, \text{BatchPSol}, \text{PProve}, \text{PVer})$  satisfies verifiability if for all  $\lambda \in \mathbb{N}$ , all polynomials  $\mathbf{T}$  in  $\lambda$ , all  $\text{pp}$  in the support of  $\text{PSetup}$ , all  $(Z, \pi)$ , such that  $Z$  is not in the support of  $\text{PGen}(\text{pp}, \cdot)$ , then it holds that  $\Pr[\text{PVer}(\text{pp}, Z, \pi) = 1] = 0$ .

In terms of secrecy, we want that the secrecy property (Definition 3) except now the adversary additionally gets the proof  $\pi \leftarrow \text{PProve}(\text{pp}, Z, s_b, r)$  for the challenge puzzle  $Z$ . The adversary as before is required to guess  $b$  correctly with the additional information of  $\pi$ . We omit the formal definition since the extension is straightforward.



**Construction.** Verifiability can be achieved by augmenting the constructions from §4.2 and §4.3 with PProve and PVer algorithms. In both cases, we can instantiate PProve with a perfectly sound non-interactive zero-knowledge (NIZK) proof, for e.g., the one from [28], that certifies that the puzzle is in the support of PGen. That is, PProve internally runs the NIZK prover for the language  $\mathcal{L} := \{Z : \exists s \in \mathcal{S}, r \in \{0, 1\}^\lambda, s.t. Z \leftarrow \text{PGen}(\text{pp}, s; r)\}$ , where  $\mathcal{S}$  is the solution space of the TLP scheme. Correspondingly, PVer internally runs the NIZK verifier which verifies the NIZK proof  $\pi$  for the language  $\mathcal{L}$  and outputs 0 or 1. In protocols, before batch solving puzzles, parties may run PVer to verify each of the puzzles are well-formed for a correct execution of BatchPSol. We require a NIZK with perfect soundness for the notion of verifiability we defined in Definition 17. However we could also have used a computationally sound variant of the NIZK (even a zk-SNARK) by changing the definition of verifiability appropriately for computational adversaries.

## D Byzantine Broadcast: Analysis

*Proof (Thm. 3).* Assume that  $\Pi_{\text{bb-sa}}$  is *not*  $(\delta, t)$ -secure protocol secure against  $\mathcal{A}$ , then we build adversary  $\mathcal{B}$  that uses  $\mathcal{A}$  to show that  $\Pi_{\text{bb-wa}}$  is *not*  $(\Delta, t)$ -secure. At a high-level, the simulator  $\mathcal{B}$ :

1. Simulates honest nodes to  $\mathcal{A}$  in the execution of  $\Pi_{\text{bb-sa}}$
2. Sends the honest messages in the simulation identical to the messages sent by the honest nodes in  $\Pi_{\text{bb-wa}}$ .
3. Transforms every attack on  $\Pi_{\text{bb-sa}}$  by  $\mathcal{A}$  in the simulation to an attack on  $\Pi_{\text{bb-wa}}$ .
4. Observes the order, delay, contents of the messages sent by corrupted nodes in the simulation and replicates the same behavior in the execution of  $\Pi_{\text{bb-wa}}$ .

$\mathcal{B}$  accomplishes this by:

- Using non-committing time-lock encryption and programming the RO to explain the contents of message while decryption;
- Using a smaller network delay parameter  $\delta$  in the simulation and starting the simulation  $\delta$  time-steps ahead of  $\Pi_{\text{bb-wa}}$ .

Observe that the compiled protocol  $\Pi_{\text{bb-sa}}$  sends and receives same messages as  $\Pi_{\text{bb-wa}}$ , except that the messages in  $\Pi_{\text{bb-sa}}$  are encrypted. Hence, it is necessary for  $\mathcal{B}$  to send the honest messages in the simulation identical to the messages sent by the honest nodes in  $\Pi_{\text{bb-wa}}$ .

Let the upper bound on the network delay in  $\Pi_{\text{bb-wa}}$  be  $\Delta = 2\delta$  and the TLP hardness in  $\Pi_{\text{bb-sa}}$  be  $\mathbf{T} = \delta$ . Recall that  $\mathcal{B}$  is as powerful as  $\mathcal{A}$ , *except*  $\mathcal{A}$  can perform *after-the-fact* removal in the execution of  $\Pi_{\text{bb-sa}}$ , which *cannot* be performed by  $\mathcal{B}$  in the execution of  $\Pi_{\text{bb-wa}}$ . Thus, attacks by  $\mathcal{A}$  is translated by  $\mathcal{B}$  as follows:

- **General corruptions:** Whenever,  $\mathcal{A}$  corrupts a node in the simulation,  $\mathcal{B}$  corrupts the corresponding node in the execution of  $\Pi_{\text{bb-wa}}$ . Specifically,  $\mathcal{B}$  hands  $\mathcal{A}$ : all private coins (including the secret keys) of the corrupted node. Since honest nodes in  $\Pi_{\text{bb-sa}}$  securely erase the internal state,  $\mathcal{A}$  has to wait for  $T$  time-steps to learn the internal state of the corrupted node.
- **After-the-fact removal:**  $\mathcal{B}$  handles this attack by starting the simulation  $\delta$  time-steps ahead of the execution of  $\Pi_{\text{bb-wa}}$ . Whenever  $\mathcal{A}$  deletes any message in-flight sent by an honest node in  $\Pi_{\text{bb-sa}}$ ,  $\mathcal{B}$  corrupts the corresponding node in  $\Pi_{\text{bb-wa}}$ . Since the simulation is  $\delta$  time-steps ahead, the corresponding round in  $\Pi_{\text{bb-wa}}$  is *yet to begin*. Thus, the honest node in  $\Pi_{\text{bb-wa}}$  is corrupted even before it sends any messages.

Let *honest* denote the set of nodes currently honest in the execution of  $\Pi_{\text{bb-sa}}$ . The simulation works as follows (Fig. 3):

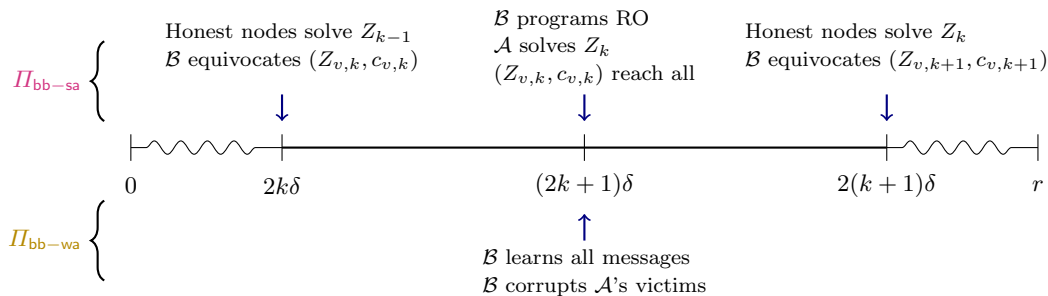
For every  $k \in \{0, 1, 2, \dots\}$ :

- At  $t = 2k\delta$ :
  - If  $t = 0$ ,  $\mathcal{B}$  starts the simulation of  $\Pi_{\text{bb-sa}}$
  - If  $t \neq 0$ , every honest node in the simulation opens the  $c_{v,k-1}$ 's to  $m_{v,k-1}$ 's for all  $v \in \text{honest}$
  - $\mathcal{B}$  sends non-committing TLPs as:
    - For every  $v \in \text{honest}$ :
      - \* Pick  $x_{v,k} \leftarrow \mathcal{S}$  and  $c_{v,k} \leftarrow \mathcal{C}$
      - \*  $Z_{v,k} \leftarrow \text{PGen}(\text{pp}, x_{v,k})$ , computes  $\pi_{v,k}$  as the proof of “well-formed-ness” of  $Z_{v,k}$ , and send  $(Z_{v,k}, c_{v,k}, \pi_{v,k})$
- Since  $\mathcal{A}$  is rushing, it learns all the messages instantaneously, checks  $\pi_{v,k}$ 's, combines all valid puzzles into a single puzzle, and starts solving the puzzle.
- At  $t = (2k + 1)\delta$ 
  - If  $t = \delta$ ,  $\mathcal{B}$  starts the execution of  $\Pi_{\text{bb-wa}}$
  - As  $\mathcal{B}$  is rushing, it can learn all the messages that are sent by the honest nodes in  $\Pi_{\text{bb-wa}}$ .
  - $\mathcal{A}$  delivers the messages to all honest nodes in the simulation. Also, every honest node checks  $\pi_{v,k}$ 's, combines all valid puzzles into a single puzzle, and starts solving the puzzle.
  - Since  $\delta$  time has elapsed since  $\mathcal{A}$  started solving the puzzle,  $\mathcal{A}$  solves the individual solutions  $(s_{v,k})_{v \in \text{honest}}$  from the puzzle, and performs  $(H(s_{v,k}))_{v \in \text{honest}}$ .
  - As  $\mathcal{B}$  learnt the honest messages  $(m_{v,k})_{v \in \text{honest}}$ , it programs  $H(s_{v,k})$  to return  $w_{v,k}$ , where  $w_{v,k} = c_{v,k} \oplus m_{v,k}$ .
- At every time step,
  - $\mathcal{B}$  identically maps the order, delay, contents of the messages sent by corrupted nodes in the simulation to the behavior of the corrupted nodes in execution of  $\Pi_{\text{bb-wa}}$ .
  - If  $\mathcal{A}$  corrupts an honest party or removes a message from an honest sender,  $\mathcal{B}$  marks the corresponding node as corrupted in  $\Pi_{\text{bb-wa}}$  and transfers all private coins of the corrupted node to  $\mathcal{A}$ .

Using the above simulation,  $\mathcal{B}$  can transform *any* attack on  $\Pi_{\text{bb-sa}}$  by  $\mathcal{A}$  into an attack on  $\Pi_{\text{bb-wa}}$ . From the strategy of  $\mathcal{A}$ , honest parties in  $\Pi_{\text{bb-sa}}$  determine their outputs based on the same values as what the honest parties see in the execution of  $\Pi_{\text{bb-wa}}$ . Moreover, by construction, the honest parties infer their outputs in the execution of  $\Pi_{\text{bb-sa}}$  deterministically using the same output derivation  $g_u$  from  $\Pi_{\text{bb-wa}}$ . Recall that the function  $g_u$ 's are fully defined by a parties public values at the beginning of the protocol. It follows that honest parties in the execution of  $\Pi_{\text{bb-wa}}$  output the same as in the execution of  $\Pi_{\text{bb-sa}}$ . Thus, if  $\Pi_{\text{bb-sa}}$  is *not*  $(\delta, t)$ -secure secure against  $\mathcal{A}$  with probability  $p$ , then  $\Pi_{\text{bb-wa}}$  is also *not*  $(\Delta, t)$ -secure against  $\mathcal{B}$  with probability  $p$ .  $\square$

## E Illustrating an Attack on Nakamoto Consensus under a Mobile Sluggish Adversary

**Nakamoto consensus in a nutshell.** Nakamoto consensus is a permissionless protocol that uses proof-of-work. Every node in the network continuously tries to add a new block to the chain. The protocol is surprisingly straightforward: at any point, every node picks the longest chain it has seen



**Fig. 3:** Timing of execution and messages in the simulation by  $\mathcal{B}$

so far (breaks ties arbitrarily), tries to extend the chain by solving a proof-of-work puzzle, and if a node succeeds, it sends out the latest block right away.

**Double spend attack.** To illustrate this attack, consider the idealized example depicted in the Fig. 2. Say there are 100 nodes in total with identical computational power, 52 are honest (51 prompt and 1 sluggish), and 48 are adversarial nodes. Assume that the average inter-arrival time between two blocks in the system is 10 minutes. Thus, the average inter-arrival time of honest blocks and adversarial blocks is 19.2 and 20.8 minutes, respectively. Moreover, assume that prompt nodes do not have *any* network delay. Let  $k$  be the height of the block  $\mathbf{B}_k$ . The adversarial strategy is as follows:

- Say the adversary wants to double spend a transaction  $\text{txn}$  with  $\text{txn}'$ , it sends out the  $\text{txn}$  to the entire network at  $t = 0$  and in-private it mines chain that includes  $\text{txn}'$ .
- Since the honest computational power is more than the computation power of the adversarial nodes, honest nodes are likely to mine  $\mathbf{B}_k$  with  $\text{txn}$  before the adversarial nodes mine  $\mathbf{A}_k$  with  $\text{txn}'$ .
- When an honest node mines  $\mathbf{B}'_{k+1}$ , the adversary makes the block winner sluggish. Thus, the rest of the honest nodes will spend their efforts on extending  $\mathbf{B}_k$ .
- Eventually, when an honest node mines the block  $\mathbf{B}_{k+1}$ , the adversary decides to release this block to the network as it does not have the budget to make both  $\mathbf{B}_{k+1}$  and  $\mathbf{B}'_{k+1}$  sluggish. Thus, all nodes receive  $\mathbf{B}_{k+1}$ .
- In the meantime, the adversary continues to mine its private chain. Since the adversary's computational power is not wasted on orphan blocks, despite having less computational power, the adversary can grow its chain faster than the honest chain.
- The adversary can continue this attack repeatedly every time when honest nodes mine a block.

Eventually, the adversary's chain will be longer than the honest chain. At this point, the adversary can release its longer private chain to discard the honest mining efforts. Observe that 38.4 minutes after  $\mathbf{B}_k$ , honest nodes have mined  $\mathbf{B}'_{k+1}$  and  $\mathbf{B}_{k+1}$ , but the honest chain has grown only by one block. Thus, even a single mobile sluggish fault has the ability to reduce the honest mining rate by *half!*

**Other attack strategies.** Rather than releasing only  $\mathbf{B}_{k+1}$ , the adversary can additionally release  $\mathbf{B}'_{k+1}$  such that honest nodes are split across  $\mathbf{B}_{k+1}$  and  $\mathbf{B}'_{k+1}$ . Moreover, the adversary can tactically release its private chain to further divide the honest computational power. For instance, the adversary can release  $\mathbf{A}_{k+1}$  along with  $\mathbf{B}_{k+1}$  and  $\mathbf{B}'_{k+1}$  to create a three-way fork. The attack vector with sluggish faults also opens the possibility of a *history rewrite attack*. Say at an arbitrary time  $T$  the

---

$\lambda$	Security parameter
$\mathbf{T}$	Hiding duration of a time-lock puzzle. Set to one round.
$m$	Duration of the mining phase
$D$	Duration of the batch solving phase (= 2 in our protocol)
$R$	An index to denote a super-round
$ R $	Duration of a super-round is $(m + D)$
$p$	Probability of a RO query mining a block
$p_c$	Probability of a RO query mining a decoy
$q$	Number of RO queries that can be made a node in a <i>round</i>
$n$	Total number of nodes
$d, t$	At any time, maximum number of sluggish nodes and adversarial nodes, respectively
$c$	Probability of every honest node mining at least one decoy in a super-round
$f$	Probability of one or more prompt blocks were mined in a super-round
$\delta$	A constant in $(0, 1)$
$\varepsilon$	Concentration inequality parameter
$\mathbf{pp}$	All public parameters including parameters from the TLPs
$\pi$	NIZK proof of well-formedness of the TLP
$T$	Target hash value to mine a block
$T_c$	Target hash value to mine a decoy

---

**Table 2:** List of notations used in the analysis.

adversary decides to build an alternate chain from genesis. Since the sluggish fault can reduce the honest mining rate by half, the adversary can eventually create a longer chain from genesis, with sufficiently many adversarial nodes.

## F Mobile Sluggishness in Permissionless Setting: Analysis

We analyze the protocol presented in §6.3. We adopt the framework and formal tools proposed by Garay et al. [21] and later improved by Li et al. to formally analyze our protocol [36]. For the sake of clarity, we exhaustively define all the variables used in our analysis in Table 2.

**Definitions.** Let:

$$\begin{aligned}
 p &= T/2^\lambda & p_c &= T_c/2^\lambda \\
 c' &= 1 - (1 - p_c)^{qm} & c &= (c')^{(n-t)} = (1 - (1 - p_c)^{qm})^{(n-t)}
 \end{aligned}
 \tag{5}$$

where,

- $p$  and  $p_c$  are the probabilities of a RO query mining a block and decoy, respectively
- $c'$  is the probability of *an* honest node mining at least one decoy in a super-round
- $c$  is the probability of *every* honest node mining at least one decoy in a super-round

Recall that due to the mobility of sluggishness, an honest node can toggle between being prompt and sluggish throughout the execution. Thus, we define a block mined in a super-round  $R$  as:

- A *prompt block*, if mined by an honest node and the node was prompt *at the beginning* of solving phase of *both*  $R - 1$  and  $R$ . Recall that in our protocol (§6.3), in the first round of the solve phase, nodes send and receive puzzles. Thus, intuitively, being prompt at the beginning of the solve phase of both  $R - 1$  and  $R$  implies that a prompt block extends the longest chain in the view of the prompt nodes and all prompt nodes receive this block before the mining phase of the next super-round.
- A *sluggish block*, if mined by an honest node, but not a *prompt* block.

- An *honest block*, if it is a *prompt* or a *sluggish* block.
- An *adversarial block*, if mined by an adversarial node.

We define the random variables  $X, Y, Z, W, T$  as follows:

In super-round  $R$ ,

$$\begin{aligned}
 X[R] &= \begin{cases} 1, & \text{at least one prompt block was mined and} \\ & \text{all honest nodes mined a decoy} \\ 0, & \text{otherwise} \end{cases} \\
 Y[R] &= \begin{cases} 1, & \text{exactly one prompt block was mined and} \\ & \text{all honest nodes mined a decoy} \\ 0, & \text{otherwise} \end{cases} \\
 W[R] &= \begin{cases} 1, & \text{at least one sluggish block was mined} \\ 0, & \text{otherwise} \end{cases} \\
 Z[R] &= \text{Number of adversarial blocks mined} \\
 T[R] &= Z[R] + W[R]
 \end{aligned}$$

**Overview of the random variables.** Since the adversary can make  $d$  nodes sluggish, if an honest node does not mine a decoy, then the probability of the adversary guessing the block winner increases. Thus, in order to hide the prompt block winner, all honest nodes have to mine decoys. This requirement is captured by the term  $c$  in variables  $X[R]$  and  $Y[R]$ .

At a high-level, the variable  $X[R]$  captures the rate of growth of prompt chains after accounting for the probability of all honest nodes mining a decoy. Similarly,  $Y[R]$  captures the *uniquely successful* super-round for prompt nodes. The variable  $Y[R]$  is analogous to the notion of uniquely successful rounds/convergence opportunities proposed in the prior works [21,41]. A uniquely successful super-round guarantees that all prompt nodes have the opportunity to converge on a chain.

Despite being disconnected from the majority of the network sluggish nodes can build a chain amongst themselves. Since  $d$  sluggish nodes are *honest*, they do not collude amongst themselves when they invoke the RO. This, is captured by  $W[R]$ .

Random variable  $T[R]$  captures the maximum length by which any non-prompt node can increase its chain. Since, the adversary controls the messages received by sluggish nodes, potentially, the adversary can trick the sluggish nodes to extend the adversary's chain. Thus, the sluggish nodes can inadvertently extend the adversarial chain.

However, observe that in our mining procedure, honest nodes abruptly stop mining in a super-round if they have mined a block in the super-round. Thus, each honest node (sluggish or prompt) mines at most one block in a super-round. Since, the honest nodes invoke the RO without colluding, if two honest nodes mine a block in a super-round, then they cannot extend one another.

*Remark 1 (Impact of sluggishness on the prompt chain).* Observe that there are  $(n - d - t)$  prompt nodes, however because of the mobility of the sluggishness, effectively only  $(n - 2d - t)$  prompt nodes work on extending the longest prompt chain. This is because the adversary has a budget of  $d$  sluggish nodes at any given time. Thus, an adversary can toggle the sluggishness of at most  $2d$  honest nodes. Say  $d$  sluggish nodes have been trying to extend a shorter fork of the longest prompt

chain. These sluggish nodes are unlikely to add any useful work to the longest prompt chain as they are working on the staler chain. When the honest nodes send out the puzzles at the end of the mining process, the adversary can make a disjoint subset of  $d$  honest nodes sluggish. Thus, effectively removing the mining efforts of  $d$  nodes from  $(n - d - t)$  prompt nodes. Newly prompt nodes send and receive the backlog of all messages. However, since they were on the minority fork, the network will discard their prior efforts to extend the shorter chain.

**Expectation of random variables.** Let

$$f = 1 - (1 - p)^{qm(n-2d-t)} \leq pqm(n - 2d - t) \quad (6)$$

where  $f$  is the probability of one or more prompt blocks were mined in a super-round. Since the probability of mining one or more prompt blocks is independent of the probability of all honest nodes mining a decoy,  $\mathbb{E}[X[R]] = c(1 - (1 - p)^{qm(n-2d-t)})$ . Similarly, the probability of mining exactly one prompt block is independent of the probability of all honest nodes mining a decoy. Thus,  $\mathbb{E}[Y[R]] = cqm(n - 2d - t)p(1 - p)^{qm(n-2d-t)-1}$ .

$$\mathbb{E}[Z[R]] = pq(m + D)t$$

$$\mathbb{E}[W[R]] = 1 - (1 - p)^{qmd}$$

$$\mathbb{E}[T[R]] = \mathbb{E}[Z[R]] + \mathbb{E}[W[R]]$$

By linearity of expectations,

$$\mathbb{E}[T[R]] = \mathbb{E}[Z[R]] + \mathbb{E}[W[R]]$$

Since the adversary can eclipse the view for the sluggish nodes, the adversary can trick the sluggish nodes into extending the adversary's chain. Thus, the mining advantage of adversarial nodes is at most the sum of  $Z[R] + W[R]$ .

To summarize,

$$\begin{aligned} \mathbb{E}[X[R]] &= c(1 - (1 - p)^{qm(n-2d-t)}) & X[S, R] &= \sum_{i=S}^{R-1} X[i] \\ \mathbb{E}[Y[R]] &= cqm(n - 2d - t)p(1 - p)^{qm(n-2d-t)-1} & Y[S, R] &= \sum_{i=S}^{R-1} Y[i] \\ \mathbb{E}[T[R]] &= \mathbb{E}[Z[R]] + \mathbb{E}[W[R]] & T[S, R] &= \sum_{i=S}^{R-1} T[i] \end{aligned}$$

We compute bounds on the random variables  $X, Y, T$  as follows:

**Upper bound on  $\mathbb{E}[X[R]]$**

Using [Eq. 6](#)

$$\begin{aligned} cf &\leq cpqm(n - 2d - t) \\ \mathbb{E}[X[R]] &\leq cpqm(n - 2d - t) \end{aligned}$$



**Lower bound on  $\mathbb{E}[X[R]]$**

$$\begin{aligned}\frac{f}{1-f} &= \frac{1 - (1-p)^{qm(n-2d-t)}}{1 - 1 + (1-p)^{qm(n-2d-t)}} = \frac{1 - (1-p)^{qm(n-2d-t)}}{(1-p)^{qm(n-2d-t)}} \\ &= (1-p)^{-qm(n-2d-t)} - 1 = \frac{1}{(1-p)^{qm(n-2d-t)}} - 1\end{aligned}$$

Since,  $(1-p)(1+p) < 1$ , where  $p \neq 0$

$$> (1+p)^{qm(n-2d-t)} - 1$$

Using Bernoulli's inequality  $(1+x)^r \geq 1+rx$ , where  $x > -1$  and  $r \geq 1$

$$\frac{f}{1-f} > pqm(n-2d-t) \quad (7)$$

Multiply by  $c$

$$\begin{aligned}\frac{cf}{1-f} &> cpqm(n-2d-t) \\ \mathbb{E}[X[r]] = cf &> (1-f)cpqm(n-2d-t)\end{aligned} \quad (8)$$

$$(1-f)cpqm(n-2d-t) < \mathbb{E}[X[r]] = cf \leq cpqm(n-2d-t)$$

**Lower bound on  $\mathbb{E}[Y[R]]$**

$$\mathbb{E}[Y[R]] = cqm(n-2d-t)p(1-p)^{qm(n-2d-t)-1}$$

Using Bernoulli's inequality  $(1+x)^r \geq 1+rx$ , where  $x > -1$  and  $r \geq 1$

$$\begin{aligned}&\geq cpqm(n-2d-t)(1-pqm(n-2d-t) + p) \\ &> cpqm(n-2d-t)(1-pqm(n-2d-t)) \\ &> cf(1-f)\end{aligned} \quad (9)$$

Let  $g(x) = x(1-x)$ . Notice that  $g(x)$  is increasing when  $x \in [0, 1/2]$ . From [Eq. 3](#) and [Eq. 6](#),  $1/2 > pqm(n-2d-t) \geq f$ , thus [Eq. 9](#) is true.

**Rate of non-prompt vs prompt mining.**

$$\begin{aligned}\mathbb{E}[T[R]] &= pq(m+D)t + 1 - (1-p)^{qmd} \\ &\leq pq(m+D)t + pqmd\end{aligned}$$

Multiply and divide by  $(n-2d-t)$  and  $m$

$$\leq pq \frac{(m+D)t + md}{m(n-2d-t)} m(n-2d-t)$$

From [Eq. 1](#)

$$\leq (1-\delta)cpqm(n-2d-t)$$

From [Eq. 8](#)

$$< \frac{(1-\delta)}{(1-f)} \mathbb{E}[X[R]]$$

From [Eq. 2](#),  $\delta > f \Rightarrow 1-\delta < 1-f \Rightarrow \frac{1-\delta}{1-f} < 1$

$$\mathbb{E}[T[R]] < \mathbb{E}[X[R]] \quad (10)$$

Also, note that

$$\mathbb{E}[T[R]] \leq pq \frac{(m+D)t + md}{m(n-2d-t)} m(n-2d-t)$$

Using [Eq. 7](#)

$$< \frac{(m+D)t + md}{m(n-2d-t)} \frac{f}{(1-f)} \quad (11)$$

Below we define typical execution events for the blockchain. Informally, typical execution guarantees that random variables  $X, Y, T$  are close to expected values.

**Definition 19 (Typical execution).** For all integers  $1 \leq S < R$  and  $\varepsilon \in (0, 1)$ , let

$$E[S, R] = E_1[S, R] \cap E_2[S, R] \cap E_3[S, R]$$

where

$$E_1[S, R] := (1 - \varepsilon)\mathbb{E}[X[S, R]] < X[S, R] < (1 + \varepsilon)\mathbb{E}[X[S, R]] \quad (12)$$

$$E_2[S, R] := (1 - \varepsilon)\mathbb{E}[Y[S, R]] < Y[S, R] \quad (13)$$

$$E_3[S, R] := T[S, R] < \mathbb{E}[T[S, R]] + \varepsilon\mathbb{E}[X[S, R]] \quad (14)$$

**Typical execution properties.** Under typical execution, the following properties hold true.

**Lemma 3.**

$$(1 - \varepsilon)cf(R - S) < X[S, R] < (1 + \varepsilon)cf(R - S) \quad (15)$$

$$\left(1 - \frac{\delta}{3}\right)cf(R - S) < (1 - \varepsilon)cf(1 - f)(R - S) < Y[S, R] \quad (16)$$

$$T[S, R] < \mathbb{E}[T[S, R]] + \varepsilon\mathbb{E}[X[S, R]] < \left(1 - \frac{2\delta}{3}\right)cf(R - S) \quad (17)$$

$$\begin{aligned} T[S, R] &< \left(1 + \frac{\delta}{2}\right) \frac{(m+D)t + md}{cm(n-2d-t)} X[S, R] + \varepsilon\mathbb{E}[X[S, R]] \\ &< \left(1 - \frac{\delta}{2}\right) X[S, R] \end{aligned} \quad (18)$$

$$T[S, R] < Y[S, R] \quad (19)$$

*Proof.* We prove the lemma as follows:

**For [Eq. 15](#).** Substituting the expectation of  $X[S, R]$  in [Eq. 12](#) to get [Eq. 15](#).

**For [Eq. 16](#).** This can be proved as follows:

$$(1 - \varepsilon)(1 - f) > (1 - \varepsilon - f + \varepsilon f) > 1 - \varepsilon - f > \left(1 - \frac{\delta}{3}\right)$$

**For [Eq. 17](#).** First inequality follows the proof of [Eq. 14](#).

For second inequality, it is sufficient to verify the following:

Using upper bound of  $T[S, R]$  from [Eq. 11](#)

$$\frac{(m+D)t+md}{m(n-2d-t)} \frac{f}{(1-f)} (R-S) + \varepsilon cf(R-S) \stackrel{?}{<} \left(1 - \frac{2\delta}{3}\right) cf(R-S)$$

From [Eq. 1](#)

$$\begin{aligned} \frac{(1-\delta)cf}{(1-f)} + \varepsilon cf &\stackrel{?}{<} \left(1 - \frac{2\delta}{3}\right) cf \Rightarrow \frac{(1-\delta)}{(1-f)} + \varepsilon \stackrel{?}{<} \left(1 - \frac{2\delta}{3}\right) \\ (1-\delta) + \varepsilon(1-f) &\stackrel{?}{<} \left(1 - \frac{2\delta}{3}\right)(1-f) \end{aligned} \quad (20)$$

The upper bound of LHS of [Eq. 20](#) can be simplified as:

$$(1-\delta) + \varepsilon(1-f) < 1 - \delta + \varepsilon$$

The lower bound of RHS of [Eq. 20](#) can be simplified as:

$$\left(1 - \frac{2\delta}{3}\right)(1-f) > 1 - \frac{2\delta}{3} - f$$

To verify [Eq. 20](#), it is enough to check:

$$1 - \delta + \varepsilon \stackrel{?}{<} 1 - \frac{2\delta}{3} - f \Rightarrow -3\delta + 3\varepsilon + 3f \stackrel{?}{<} -2\delta \Rightarrow 3\varepsilon + 3f < 3\delta - 2\delta = \delta$$

**For [Eq. 18](#) (first inequality).** It is sufficient to verify the following:

$$\begin{aligned} &\frac{(m+D)t+md}{m(n-2d-t)} \frac{f}{(1-f)} (R-S) + \varepsilon \mathbb{E}[X[S, R]] \\ &\stackrel{?}{<} \left(1 + \frac{\delta}{2}\right) \frac{(m+D)t+md}{cm(n-2d-t)} X[S, R] + \varepsilon \mathbb{E}[X[S, R]] \end{aligned}$$

Using first inequality of [Eq. 15](#),  $f(R-S) < X[S, R]/(c(1-\varepsilon))$

$$\begin{aligned} \frac{(m+D)t+md}{cm(n-2d-t)} \frac{1}{(1-\varepsilon)(1-f)} &\stackrel{?}{<} \frac{(m+D)t+md}{cm(n-2d-t)} \left(1 - \frac{\delta}{2}\right) \\ \frac{1}{(1-\varepsilon)(1-f)} &\stackrel{?}{<} \left(1 + \frac{\delta}{2}\right) \end{aligned}$$

Since  $1 - \frac{\delta}{3} < (1-\varepsilon)(1-f)$ , to verify the above, it is sufficient to check the following:

$$1 \stackrel{?}{<} \left(1 - \frac{\delta}{3}\right) \left(1 + \frac{\delta}{2}\right) \Rightarrow 1 \stackrel{?}{<} 1 - \frac{\delta}{3} + \frac{\delta}{2} - \frac{\delta^2}{6} \Rightarrow 1 \stackrel{?}{<} 1 + \frac{\delta}{6} - \frac{\delta^2}{6}$$

Since  $\delta \in (0, 1)$ ,  $\frac{1}{6}\delta(1-\delta) > 0$ . Thus, first inequality in [Eq. 18](#) is true.

**For [Eq. 18](#) (second inequality).** It is sufficient to verify the following:

$$\left(1 + \frac{\delta}{2}\right) \frac{(m+D)t+md}{cm(n-2d-t)} X[S, R] + \varepsilon \mathbb{E}[X[S, R]] \stackrel{?}{<} \left(1 - \frac{\delta}{2}\right) X[S, R]$$

From Eq. 1:

$$\left(1 + \frac{\delta}{2}\right)(1 - \delta)X[S, R] + \varepsilon\mathbb{E}[X[S, R]] \stackrel{?}{<} \left(1 - \frac{\delta}{2}\right)X[S, R]$$

Using second inequality of Eq. 15,  $\mathbb{E}[X[S, R]] = cf(R - S) < X[S, R]/(1 - \varepsilon)$

$$\begin{aligned} \left(1 + \frac{\delta}{2}\right)(1 - \delta)X[S, R] + \frac{\varepsilon}{1 - \varepsilon}X[S, R] &\stackrel{?}{<} \left(1 - \frac{\delta}{2}\right)X[S, R] \\ 1 + \frac{\delta}{2} - \delta - \frac{\delta^2}{2} + \frac{\varepsilon}{1 - \varepsilon} &\stackrel{?}{<} 1 - \frac{\delta}{2} \Rightarrow 1 - \frac{\delta}{2} - \frac{\delta^2}{2} + \frac{\varepsilon}{1 - \varepsilon} \stackrel{?}{<} 1 - \frac{\delta}{2} \end{aligned}$$

From Eq. 4,  $\frac{2\varepsilon}{1 - \varepsilon} < \delta^2$ , thus the above equation is true.

**For Eq. 19.**

From Eq. 16 and Eq. 17

$$\begin{aligned} \left(1 - \frac{\delta}{3}\right)cf(R - S) < Y[S, R]T[S, R] < \left(1 - \frac{2\delta}{3}\right)cf(R - S) \Rightarrow \\ T[S, R] < Y[S, R] \end{aligned}$$

□

**Theorem 4.** *An execution is typical according to Definition 19 with probability  $1 - 4e^{-\frac{\varepsilon^2}{6}cf(R-S)}$ .*

*Proof.*

$$\begin{aligned} \Pr[E[S, R]] &= 1 - \Pr[E^c[S, R]] \\ &\geq 1 - \Pr[E_1^c[S, R]] - \Pr[E_2^c[S, R]] - \Pr[E_3^c[S, R]] \end{aligned}$$

$$\Pr[E_1^c[S, R]] = \Pr[|X[S, R] - \mathbb{E}[X[S, R]]| \geq \varepsilon\mathbb{E}[X[S, R]]]$$

Using Eq. 27

$$\leq 2e^{-\frac{\varepsilon^2}{3}cf(R-S)}$$

$$\Pr[E_2^c[S, R]] = \Pr[Y[S, R] \leq (1 - \varepsilon)\mathbb{E}[Y[S, R]]]$$

Using Eq. 26

$$\leq e^{-\frac{\varepsilon^2}{3}cf(1-f)(R-S)}$$

From Eq. 3,  $f < 1/2 \Rightarrow 1/2 < (1 - f)$

$$< e^{-\frac{\varepsilon^2}{6}cf(R-S)}$$

$$\Pr[E_3^c[S, R]] = \Pr[T[S, R] \geq \mathbb{E}[T[S, R]] + \varepsilon\mathbb{E}[X[S, R]]]$$

Using Eq. 10

$$\begin{aligned} &\leq \Pr [T[S, R] \geq \mathbb{E}[T[S, R]] + \frac{\varepsilon}{2}\mathbb{E}[T[S, R]] + \frac{\varepsilon}{2}\mathbb{E}[X[S, R]]] \\ &\leq \Pr [T[S, R] \geq (1 + \frac{\varepsilon}{2})\mathbb{E}[T[S, R]] + \frac{\varepsilon}{2}\mathbb{E}[X[S, R]]] \end{aligned}$$

Using Eq. 28

$$= \frac{\mathbb{E}[e^{uT[S, R]}]}{e^{u\left((1+\frac{\varepsilon}{2})\mathbb{E}[T[S, R]] + \frac{\varepsilon}{2}\mathbb{E}[X[S, R]]\right)}}$$

Let  $\beta = pq((m + D)t + md)(R - S)$

Since  $T[S, R] = X[S, R] + W[S, R]$ , using Eq. 31,  $\mathbb{E}[e^{uT[S, R]}] = \mathbb{E}[e^{uZ[S, R]}] \times \mathbb{E}[e^{uW[S, R]}]$ . Observe that  $\mathbb{E}[Z[R]]$  is a Binomial random variable and  $\mathbb{E}[W[R]]$  is a Bernoulli random variable, thus using Eq. 29 and Eq. 30:

$$\leq \frac{(1 + p(e^u - 1))^{q(m+D)t(R-S)} \cdot (1 + pqmd(e^u - 1))^{(R-S)}}{e^{u\left((1+\frac{\varepsilon}{2})\beta + \frac{\varepsilon}{2}cf(R-S)\right)}}$$

Note that  $(1 + x) \leq e^x$ . Thus,  $(1 + p(e^u - 1)) \leq e^{p(e^u - 1)}$  and  $(1 + pqmd(e^u - 1)) \leq e^{pqmd(e^u - 1)}$

$$\begin{aligned} &\leq \frac{e^{(e^u - 1)\beta}}{e^{u\left((1+\frac{\varepsilon}{2})\beta + \frac{\varepsilon}{2}cf(R-S)\right)}} \\ &\leq e^{(e^u - 1)\beta - u\left((1+\frac{\varepsilon}{2})\beta + \frac{\varepsilon}{2}cf(R-S)\right)} \\ &\leq e^{\left(e^u - 1 - u\left(1 + \frac{\varepsilon}{2}\right)\right)\beta - u\frac{\varepsilon}{2}cf(R-S)} \end{aligned}$$

Let  $u = \log(1 + \frac{\varepsilon}{2})$

$$\begin{aligned} &\leq e^{\left(1 + \frac{\varepsilon}{2} - 1 - \log(1 + \frac{\varepsilon}{2})\left(1 + \frac{\varepsilon}{2}\right)\right)\beta - \frac{\varepsilon}{2}\log(1 + \frac{\varepsilon}{2})cf(R-S)} \\ &\leq e^{\left(\frac{\varepsilon}{2} - \log(1 + \frac{\varepsilon}{2})\left(1 + \frac{\varepsilon}{2}\right)\right)\beta - \frac{\varepsilon}{2}\log(1 + \frac{\varepsilon}{2})cf(R-S)} \end{aligned}$$

$(1 + x) \log(1 + x) > x, \forall x > 0$

$$\begin{aligned} &\leq e^{\left(\frac{\varepsilon}{2} - \frac{\varepsilon}{2}\right)\beta - \frac{\varepsilon}{2}\log(1 + \frac{\varepsilon}{2})cf(R-S)} \\ &\leq e^{-\frac{\varepsilon}{2}\log(1 + \frac{\varepsilon}{2})cf(R-S)} \end{aligned}$$

For  $x > -1$ ,  $\frac{x}{1+x} \leq \log(1 + x)$

$$\leq e^{-\frac{\varepsilon}{2}\frac{\frac{\varepsilon}{2}}{1+\frac{\varepsilon}{2}}cf(R-S)} \leq e^{-\frac{\varepsilon}{2}\frac{\varepsilon}{2+\varepsilon}cf(R-S)}$$

Since  $0 < \varepsilon < 1 \Rightarrow \frac{\varepsilon}{3} < \frac{\varepsilon}{2+\varepsilon}$

$$< e^{-\frac{\varepsilon^2}{6}cf(R-S)}$$

$$\Pr[E_1^c[S, R]] \leq 2e^{-\frac{\varepsilon^2}{3}cf(R-S)} < 2e^{-\frac{\varepsilon^2}{6}cf(R-S)}$$

$$\Pr[E_2^c[S, R]] < e^{-\frac{\varepsilon^2}{6}cf(R-S)}$$

$$\Pr[E_3^c[S, R]] < e^{-\frac{\varepsilon^2}{6}cf(R-S)}$$

Thus,

$$\Pr[E[S, R]] > 1 - 4e^{-\frac{\varepsilon^2}{6}cf(R-S)}$$

□

**Upper bound on decoys.** Let,  $p_c = \frac{k}{q(m+D)}$ , where  $k < q(m+D)$ . Moreover, let  $k = \log n$  and  $z$  be the random variable to denote the number of decoys mined by the adversarial nodes in a super-round.

$$\begin{aligned}\mathbb{E}[z] &= \frac{k}{q(m+D)} \cdot t \cdot q(m+D) = kt \\ \Pr[z \geq (1+\varepsilon)\mathbb{E}[z]] &\leq e^{-\frac{\varepsilon^2 kt}{3}} \leq e^{-\Omega(\varepsilon^2 n \log n)}\end{aligned}$$

Since  $n$  is polynomial in  $\lambda$ , the probability of adversary overwhelming honest nodes with decoys is low.

**Probability of honest nodes mining a decoy.** Let,  $p_c = \frac{k}{q(m+D)}$ , where  $k < q(m+D)$ . From Eq. 5,  $c$  is the probability of every honest node mining at least one decoy in a super-round.

$$\begin{aligned}c &= (1 - (1 - p_c)^{qm})^{(n-t)} = (1 - (1 - \frac{k}{q(m+D)})^{qm})^{(n-t)} \\ &> (1 - (1 - \frac{k}{qm})^{qm})^{(n-t)}\end{aligned}$$

$$(1 - w/x)^x \leq e^{-w}$$

$$\begin{aligned}&> (1 - e^{-k})^{(n-t)} \\ &> (1 - e^{-k})^n\end{aligned}$$

Recall, that  $k = \log n$ .

$$> (1 - e^{-\log n})^n$$

Let  $a = \frac{1}{\ln 2} > 1$ , we compute the limit:

$$\begin{aligned}\lim_{n \rightarrow \infty} (1 - e^{-\log n})^n &= \lim_{n \rightarrow \infty} (1 - e^{-\frac{\ln n}{\ln 2}})^n = \lim_{n \rightarrow \infty} \left(1 - \left(e^{\ln n}\right)^{-a}\right)^n \\ &= \lim_{n \rightarrow \infty} \left(1 - \frac{1}{n^a}\right)^n = 1\end{aligned}$$

## F.1 Consistency and Liveness

A blockchain  $C$  is an ordered sequence of blocks. Let:

- $C[0]$  be the commonly agreed genesis block and  $|C|$  be the length of the chain
- $C^{\lceil k}$  denote the chain obtained by removing the last  $k$  blocks from  $C$
- $C \preceq C'$  to denote that  $C$  is a prefix of the chain  $C'$
- $C[-i : ]$  be the last  $i$  blocks of  $C$  inclusive of  $C[-i]$  and the last block of  $C$



**Blockchain properties.** We adapt the commonly used definitions of *chain growth*, *chain quality*, and *common prefix* for our setting [21,41,42]. Recall that a super-round is a consecutive sequence of rounds of block mining and batch-solving of puzzles.

**Definition 20 (Chain Growth).** *For any sequence of consecutive super-rounds in  $[S, R]$ , the length of a chain held by a node which is prompt at the beginning of the solve phase of  $R$  has grown by at least  $\tau \cdot (R - S)$  since the super-round  $S$ , where  $\tau \in [0, 1]$  and  $R \in \mathbb{N}$ .*

**Definition 21 (Chain Quality).** *For any sequence of at least  $k$  consecutive blocks in a chain held by a prompt node, the ratio of blocks mined by honest nodes is at least  $\mu \cdot k$ , where  $\mu \in [0, 1]$  and  $k \in \mathbb{N}$ .*

**Definition 22 (Common Prefix).** *Let  $C_1$  and  $C_2$  be the chain adopted by any two prompt nodes at the beginning of super-rounds  $R_1$  and  $R_2$ , respectively, such that  $R_1 \leq R_2$ . The chain obtained by removing last  $k$  blocks of  $C_1$  is a prefix of the chain  $C_2$ ,  $C_1^{\lceil k} \preceq C_2$ , where  $k \in \mathbb{N}$ .*

**Lemma 4.** *Suppose an honest node at the beginning of super-round  $S$  has a chain of length  $k$ , which was prompt at the beginning of the solve phase of  $S - 1$ . Then, by the beginning of super-round  $R > S$ , every honest node which was prompt at super-round  $R - 1$  will have a chain of at least  $k + X[S, R]$ .*

*Proof.* We prove this using induction.

- Base case ( $R = S + 1$ ): We assume that prompt nodes implicitly echo all the messages (including adversarial messages). This implies that a node prompt at round  $r$  will be up to date with prompt and adversarial messages sent before  $r$ . Thus, in every super-round, nodes prompt at the beginning of the solve phase, will have a chain of identical length by the end of solve phase. Since the boolean random variable  $X[S]$  indicates whether at least one *prompt block* was mined in super-round  $S$  (and all honest nodes mined a decoy), then by the start of super-round  $S + 1$  all prompt nodes (at the beginning of solve phase of  $S$ ) will have a chain of length at least  $k + X[S]$  after batch solving the puzzles at the end of  $S$ .
- Induction step: Assume that the hypothesis is true for super-round  $R - 1$ . Thus, at the beginning of  $R - 1$ , all prompt nodes will have a chain of length at least  $k + X[S, R - 1]$ . Observe that  $X[S, R] = X[S, R - 1] + X[R - 1]$ . When  $X[R - 1] = 1$ , it implies that at least one or more prompt block was mined in  $R - 1$ . Thus, all prompt nodes (during the solve phase of  $R - 1$ ) would have received this block in a TLP puzzle. By the end of  $R - 1$  all prompt nodes would have opened the puzzle to extend the chain. Thus, the new length of the prompt chain is at least  $k + X[S, R]$ .  $\square$

**Lemma 5.** *Under the event  $E[S, R]$  as described in Definition 19 and  $k \geq 2cf(R - S)$ , every node's last  $k$  consecutive blocks at the beginning of super-round  $R$  must have been mined after super-round  $S$ .*

*Proof.* The upper bound of chain growth is:

From Eq. 15 and Eq. 17

$$X[S, R] + T[S, R] < (1 + \varepsilon)cf(R - S) + (1 - \frac{2\delta}{3})cf(R - S) \quad (21)$$

$$< (2 + \varepsilon - \frac{2\delta}{3})cf(R - S) \quad (22)$$

$$\begin{aligned}
\text{From Eq. 2, } 6\varepsilon + 6f < 2\delta &\Rightarrow 3\varepsilon + 6f < 2\delta \Rightarrow \varepsilon - 2\delta/3 < -2f \\
&< (2 - 2f)cf(R - S) < (1 - f)2cf(R - S) & (23) \\
&< k & (24)
\end{aligned}$$

□

**Theorem 5 (Chain Growth).** *Under the event  $E[S, R]$  and  $(R - S) > \frac{2}{cf}$ , the length of the chain of every node prompt at the beginning of solve phase of  $R - 1$  grows by at least  $(1 - \varepsilon)cf(R - S)$  after the beginning of super-round  $S$ .*

*Proof.* From Eq. 15, we know that the length of a prompt chain grows by at least  $(1 - \varepsilon)cf(R - S)$ . Say, the length of the chain adopted by a prompt node at the beginning of super-round  $S$  is  $k$ . Then, by Lemma 4 all prompt nodes would have chain of length at least  $k + X[S, R]$  at beginning of super-round  $R$ . Since the duration  $(R - S) > 2/cf$  is sufficiently long, typical execution is applicable. Thus, every prompt chain grows by at least  $(1 - \varepsilon)cf(R - S)$ . □

**Theorem 6 (Chain Quality).** *Say a node which was prompt at the end of beginning of solve phase of  $R - 1$  has chain of more than  $k$  at the beginning of super-round  $R$ . Under the event  $E[S, R]$  and  $k \geq 2cf(R - S)$ , the ratio of prompt blocks in the last  $k$  consecutive blocks is at least  $\frac{\delta}{2}$ .*

*Proof.* Let,

- $C$  be the chain adopted by a prompt node at the beginning of super-round  $R$  such that  $|C| > k$ .
- $C[-k']$  be the last prompt block before  $C[-k]$  such that  $k' > k$ . If no such block exists, then  $C[-k'] = C[0]$ .
- $L = |C| - k' + 1$  be the length of the chain  $C[-k' + 1 : ]$
- $R'$  be the super-round in which  $C[-k']$  was mined by a prompt node
- $x$  be the number of prompt blocks in  $C[-k : ]$

Observe that the number of prompt blocks in  $C[-k' + 1 : ]$  is also  $x$ . Since  $k \geq 2cf(R - S)$  and  $k' > k$ , then by Lemma 5,  $R' < S$  and typical execution is applicable in super-rounds  $[R' + 1, R)$ . Thus:

$$L - x < T[S, R] < (1 - \delta/2)X[S, R] < (1 - \delta/2)L < L - \frac{L\delta}{2} < L - \frac{k\delta}{2} \Rightarrow x > \frac{\delta}{2}k$$

The second inequality is from Eq. 18, the third inequality is from Thm. 5, and the fourth inequality uses the fact that  $k < L$ . □

**Lemma 6.** *Suppose the  $k$ -th block  $B$  of  $C$  is a prompt block mined in a uniquely successful super-round. Then, the  $k$ -th block  $B'$  of  $C'$  is either  $B$  or a non-prompt block (sluggish or adversarial).*

*Proof.* Say  $B' (\neq B)$  is another prompt block at the  $k$ -th position. However, since  $B$  was mined in a uniquely successful super-round,  $B'$  could not have been mined in the same super-round, by definition. Say, the block  $B$  was mined in an earlier super-round  $R$  before  $B'$  was mined. Since,  $B$  is a prompt block, all prompt nodes would have adopted a chain of length at least  $k$  by the end of super-round  $R$ . This contradicts the assumption that  $B'$  was mined as the  $k$ -th block in a chain by a prompt node, as all prompt nodes would be extending a chain of length  $k$  by the end of  $R$ . □

**Theorem 7 (Common Prefix).** *Let  $C_1$  be the chain adopted by any honest node at the beginning of super-round  $R_1$  (which was prompt at the beginning of the solve phase of  $R_1 - 1$ ) and  $C_2$  be the chain adopted by any honest node at beginning of super-round  $R_2$  (which was prompt at the beginning of the solve phase of  $R_2 - 1$ ). Under event  $E[S, R_1]$ ,  $k \geq 2cf(R_1 - S)$ , and  $R_2 \geq R_1$ , the chain obtained by removing last  $k$  blocks of  $C_1$  is a prefix of the chain  $C_2$ , i.e.,  $C_1^{\lceil k} \preceq C_2$ .*

*Proof.* We prove this by contradiction. Assume that  $C_1^{\lceil k} \not\preceq C_2$ . Let  $R$  be the earliest super-round after  $R_1$  in which a prompt node  $P$  adopted a chain  $C'_2$  such that  $C_1^{\lceil k} \not\preceq C'_2$  and  $R \geq R_1$ . Thus, by definition  $R_2 \geq R \geq R_1$ . Moreover, let  $C'_1$  be the longest chain held by  $P$  at  $R - 1$  before adopting  $C'_2$ . Thus, by definition  $C_1 \preceq C'_1$ .

Let  $R^*$  be the super-round in which the last common prompt block between  $C'_1$  and  $C'_2$  was mined. Observe that all chains share a common genesis block, thus  $R^*$  is well-defined. Since the  $C_1^{\lceil k} \not\preceq C'_2$  and  $k \geq 2cf(R_1 - S)$ , thus by [Lemma 5](#),  $S > R^*$ .

After the protocol execution from  $R^* + 1$  to  $R$ , the chain obtained by removing the last  $k$  blocks from  $C'_1$  is not a prefix of the chain  $C'_2$ . Thus, we make the following claim:

$$T[R^* + 1, R] \geq Y[R^* + 1, R]$$

To see why this claim is valid when the prefix doesn't match, we count the number of non-prompt (adversarial + sluggish) blocks mined whenever there is a uniquely successful super-round in  $[R^*, R)$ . For a uniquely successful round  $U \in [R^*, R)$ , let  $j_U$  be the position of the corresponding block in the chain from the genesis.

$$J := \{j_U : U \text{ is a super-round where } Y[U] = 1 \wedge U \in [R^* + 1, R)\}$$

Note that both  $C'_1$  and  $C'_2$  has a block at every  $j \in J$ . Since the block winner of a  $j$  is the sole winner in that super-round and the block winner was prompt at the end of  $m$  rounds of mining, the time-lock encrypted block reaches all other prompt nodes by the beginning of the solve phase. Hence, by the end of solve phase and super-round, all other prompt nodes would have solved the time-lock puzzle to obtain the block and adopted the longer chain. This implies that  $|C'_1| \geq \max J$  and since  $P$  adopts  $C'_2$  in super-round  $R$ ,  $|C'_2| \geq |C'_1|$ .

Now we argue that for all  $j \in J$ , the  $j$ -th block in  $C'_1$  or in  $C'_2$  can either be identical or different. If the blocks are identical, then the block is adversarial or sluggish by the definition of  $R^*$ . On the other hand, if the blocks are different, then both blocks cannot be prompt. This follows from [Lemma 6](#). Thus, by counting the number of non-prompt blocks mined during uniquely successful super-rounds in each of the above possibilities, it is evident that  $T[R^* + 1, R] \geq Y[R^* + 1, R]$ .

However, since  $R_2 \geq R \geq R_1 > S > R^*$  and  $k \geq 2cf(R_1 - S) \Rightarrow 2cf(R - R^* - 1) > 2cf(R_1 - S)$ , the conditions required for typical execution  $E[R^* + 1, R]$  is applicable. Under  $E[R^* + 1, R]$ ,  $T[R^* + 1, R] < Y[R^* + 1, R]$  from the proof of [Eq. 19](#). Thus, this contradicts the assumption that  $C_1^{\lceil k} \not\preceq C_2$ .  $\square$

## F.2 Well-Known Facts used in the Analysis

**Proposition 1 (Chernoff Bounds).** *Let  $X = \sum_{i=1}^n X_i$ , where  $X_i = 1$  with probability  $p_i$  and  $X_i = 0$  with probability  $1 - p_i$ , and all  $X_i$  are independent. Let  $\mu = \mathbb{E}[X] = \sum_{i=1}^n p_i$ . Then,*

$\forall \delta \in (0, 1)$ ,

$$\Pr[X \geq (1 + \delta)\mu] \leq e^{-\delta^2\mu/2} \quad (25)$$

$$\Pr[X \leq (1 - \delta)\mu] \leq e^{-\delta^2\mu/3} \quad (26)$$

$$\Pr[|X - \mu| \geq \delta\mu] \leq 2e^{-\delta^2\mu/3} \quad (27)$$

$$\Pr[X \geq a] \leq \frac{\mathbb{E}[e^{uX}]}{e^{ua}}, \forall u > 0 \quad (28)$$

**Proposition 2 (Moment-Generating Function [52]).** *Moment Generating Function of random variable  $X$ , denoted as  $M_X(u)$ , is:*

$$M_X(u) = \mathbb{E}[e^{uX}] = (1 + p(e^u - 1))^m \quad \text{where } X \sim \text{Bin}(m, p) \quad (29)$$

$$M_X(u) = \mathbb{E}[e^{uX}] = (1 + p(e^u - 1)) \quad \text{where } X \sim \text{Bernoulli}(p) \quad (30)$$

as long as this expectation exists for  $v$  in some neighborhood of 0. That is, there is an  $v > 0$  such that for all  $u$  in  $-v < u < v$ ,  $\mathbb{E}[e^{uX}]$  exists [52].

Moreover, if  $X = \sum_{i=1}^n a_i X_i$ , where  $X_i$  are independent random variables and  $a_i$  are constants, then:

$$M_X(u) = M_{X_1}(a_1 u) M_{X_2}(a_2 u) \cdots M_{X_n}(a_n u) \quad (31)$$