Blockchain Security when Messages are Lost

Taha Ameen, Suryanarayana Sankagiri, Bruce Hajek
University of Illinois at Urbana-Champaign

Abstract

Security analyses for consensus protocols in blockchain research have primarily focused on the synchronous model, where point-to-point communication delays are upper bounded by a known finite constant. These models are unrealistic in noisy settings, where messages may be lost (i.e. incur infinite delay). In this work, we study the impact of message losses on the security of the proof-of-work longest-chain protocol. We introduce a new communication model to capture the impact of message loss called the 0-∞ model, and derive a region of tolerable adversarial power under which the consensus protocol is secure. The guarantees are derived as a simple bound for the probability that a transaction violates desired security properties. Specifically, we show that this violation probability decays almost exponentially in the security parameter. Our approach involves constructing combinatorial objects from blocktrees, and identifying random variables associated with them that are amenable to analysis. This approach improves existing bounds and extends the known regime for tolerable adversarial threshold in settings where messages may be lost.

1 Introduction

Blockchain is the data structure used by peers (miners) in a peer-to-peer network to maintain a common ledger in a decentralized manner. The consistency of this ledger is ensured through consensus protocols such as the longest-chain protocol. Following this protocol, an honest miner groups transactions into a block and appends its block to the longest chain in its view, before broadcasting the new blockchain to all other peers. Further, the system may have adversarial users that deviate from the protocol arbitrarily. Despite adversarial users attempting to disrupt the system and peer-to-peer communication incurring message delays, the parties following the protocol must agree on a consistent ledger.

Blockchain security has been studied under various consensus protocols (see [BSAB$^+$19, GK20] for a survey). Of these, the longest-chain protocol is of great interest, due its heavy use in modern blockchain implementations. The longest-chain protocol has been modeled under various assumptions: for example, discrete time is used in [GKR20, BKM$^+$20], and continuous time dynamics is used in [LGR21, Ren19, DKT$^+$20]. Further, the protocol has also been studied for a variety of leader election mechanisms in the consensus protocol. For instance, [PSS17, Ren19, GKL20] assume the proof-of-work mechanism, whereas [PS17, KRDO17, FZ17] assume a proof-of-stake mechanism. All these works establish security of the longest-chain protocol for the synchronous communication model, where communication delays are upper bounded by a known finite constant. A common theme among these results is that in the synchronous delay model, the longest-chain protocol is 'secure' under sufficient honest representation, with high probability.

In this work, we analyze the impact of message losses on the security of the longest-chain protocol following proof-of-work leader election, by introducing and analyzing an appropriate communication network model. We motivate this by reviewing some existing communication models in the literature and the known security guarantees associated with them.

1.1 Related Work

The underlying communication network can delay the successful delivery of peer-to-peer message broadcasts. Popular blockchains such as Bitcoin use the Internet as their communication network. Since this communication is subject to delay, it is natural to model the delays incurred by each block, and study the impact of delay on the security of the longest-chain protocol.

Let $0 \leq i < j$. Let b_i represent the i-th mined honest block. Let $\mathsf{delay}\,(b_i \to b_j)$ denote the time taken for block b_i to reach the miner of block b_j, and let β represent the fraction of adversarial computational power in the system. Finally, let λ be the rate at which blocks are mined in the system. Various descriptions of $\mathsf{delay}\,(b_i \to b_j)$ lead to different communication network models:

Instantaneous Model The original white-paper by Satoshi Nakamoto [Nak08] assumes an ideal communication channel, i.e. $\mathsf{delay}\left(b_i \to b_j\right) = 0$. In this model, the longest-chain protocol is provably secure when the honest computational power in the system exceeds the adversarial computational power, i.e. when $\beta < 1 - \beta$, or equivalently, when $\beta < 1/2$.

Synchronous Model The model assumes a deterministic delay for each block that is upper bounded by a known constant Δ, i.e., $\mathsf{delay}\left(b_i \to b_j\right) \leq \Delta < \infty$. This delay effectively reduces the growth rate of the chain held by an honest user. Even so, it has been proved [DKT+20, GKR20] that the synchronous model is secure with high probability if and only if

$$\beta < \frac{1}{1 + (1-\beta)\,\lambda\Delta}\left(1 - \beta\right),$$

where λ is the total mining rate of the honest users.

Partially Synchronous Model The partially synchronous model assumes the existence of some unknown and adversarially chosen 'Global Stabilization Time (GST)' such that the delays are unbounded before GST, but bounded after it [DLS88]. Therefore, at any time t, the delay satisfies $\mathsf{delay}\left(b_i \to b_j\right) \leq \Delta + \max\left(0, \mathsf{GST} - t\right)$. If certain conditions are met, the partially synchronous model is known to be secure with high probability after the Global Stabilization Time [NTT21].

Sleepy Model The sleepy model considers the setting where miners may either be online or offline, and their participation status may change during the execution of the protocol [PS17]. Let h_i denote the miner of block b_i. The incurred delay is thus

$$\mathsf{delay}\left(b_i \to b_j\right) = \begin{cases} 0 & h_j \text{ is awake when } b_i \text{ is mined} \\ \infty & h_j \text{ is asleep when } b_i \text{ is mined} \end{cases}.$$

Pass and Shi [PS17] showed that consensus can be achieved in the sleepy model with high probability, if a majority of the awake miners at any point in time are honest.

Random Delay Model The random delay model assumes that the point-to-point delays are independent and identically distributed, i.e. $\mathsf{delay}\left(b_i \to b_j\right) \sim \mathsf{X}$, where X is some known distribution. The longest-chain protocol is shown to be secure with high probability in the random delay model, if the delay distribution satisfies certain conditions and the adversarial representation in the system is below a certain threshold [SGH21].

Except for the random delay model, none of the above models account for the possibility that point-to-point communication may incur infinite delay, i.e. messages may be lost at random. For instance, the sleepy model allows infinite delay for users that are offline, but does not account for noise in the communication process. In contrast, we introduce and analyze a new communication model to study the impact of lost messages on blockchain security.

1.2 Contributions

0-∞ **Model** We introduce the 0-∞ model, where the delays are independent and identically distributed over the set $\{0, \infty\}$. Specifically, for any i, $j \geq 0$ such that $i < j$:

$$\mathsf{delay}\left(b_i \to b_j\right) = \begin{cases} 0 & \text{with probability } 1-d \\ \infty & \text{with probability } d \end{cases}.$$

This simple model postulates that a message sent point-to-point is either immediately received or permanently lost. This delay is independent for each user, and for each block. The modeling choice aligns with our objective of studying the effect of message losses.

We remark that the 0-∞ model is a special case of the i.i.d. random delay model introduced in [SGH21], which identifies a region of tolerable adversarial power as a function of the delay distribution. Specifically, if d is the probability of message loss and β is the fraction of computational power in the system that is adversarial, it is shown that the 0-∞ model is secure with high probability when $\beta < \frac{1-2d}{2(1-d)}$. However, this characterization is not tight for the 0-∞ model, and the analysis in [SGH21]

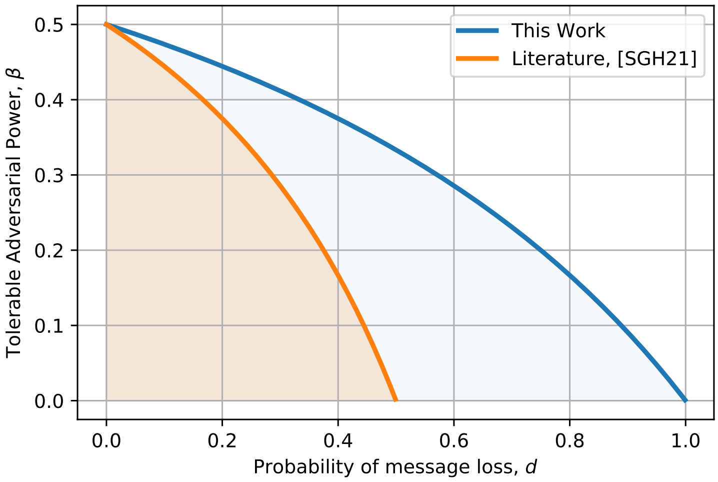

Figure 1: Improvement in the region of tolerable adversarial power

breaks down in the high-noise regime. For example, security of the model cannot be established when $d > 1/2$, i.e. more than half the messages are lost on average.

It is reasonable to wonder if adversarial computational power can at all be tolerated in the high-noise regime, for instance, when almost all messages are lost. Our work answers this question in the affirmative, by expanding the known security threshold for the 0-∞ model. In particular, our sufficient condition for security is $\frac{\beta}{1-\beta} < 1 - d$. Figure 1 shows this improvement.

Our method of analysis is significantly different from that in [SGH21]: we introduce a transmission-graph that captures the history of communication delays between blocks, and identify special paths in the graph that are linked to random variables which are amenable to analysis. Specifically, we identify special objects such as forward-special and backward-special blocks, and associate with them random variables such as forward-unheard and backward-unheard. Our technique also presents a new approach to infer the inclusion of special blocks in the chain held by an honest user through the concept of user-unheard-criterion. The method of analysis is inspired from [DKT$^+$20], where security of the synchronous model is established by considering races between honest and adversarial chains. However, our approach does not rely on message delays being finite, and we hope that the tools we introduce are of utility in the study of consensus mechanisms in more general settings, as well as of independent mathematical interest. Our contributions are summarized as:

- We introduce the 0-∞ model as a playground for studying the impact of message losses. This model provides a starting step for more complex models involving message losses.
- We introduce combinatorial objects of independent interest such as the transmission-graph. We also identify random variables (forward-unheard, backward-unheard) associated with this graph that are amenable to analysis, and introduce the user-unheard-criterion. These concepts may be utilized in security analysis of blockchain protocols in more general settings.
- We prove that the longest-chain protocol is secure in the 0-∞ model if certain conditions are met. These conditions are fairly general, and considerably extend the known threshold of tolerable adversarial power. In this regime, we show that the probability of security violation decays almost exponentially in the security parameter.

2 System Model

In this section, we describe our system model. We consider the setting where infinitely many miners participate in the longest-chain protocol for an infinite duration, and use proof-of-work as the leader election mechanism.

Ledgers, Transactions, Miners, and Blocks Blockchain is the data structure at the heart of the decentralized mechanism to maintain and update a ledger. The ledger is simply an ordered list of transactions. Transactions are assumed to be available to all the miners as soon as they are made.

Miners verify the validity of transactions, and update the ledger by grouping the transactions into blocks and linking blocks to form a blockchain. A block is an abstract data structure that contains a hash pointer to a parent block, a cryptographic signature of the block's miner, transactions, and other metadata. The first block in the system is called the genesis block.

Longest-Chain Protocol and Proof-of-Work Miners follow the proof-of-work longest-chain protocol for consensus. Following this protocol, a miner groups any and all transactions that are not included in this longest chain into a block, and attempts to append the block to the longest chain in its view. To do so, it must solve a hash puzzle and include the solution as proof-of-work. If the miner is successful, it broadcasts its chain as a message to other miners over a peer-to-peer network, subject to a communication delay. Upon receiving this message, an honest miner adopts the new chain if the received chain is longer than the chain in the miner's memory. Ties are broken using any deterministic rule, for example, by choosing the chain that terminates in the block that hashes to a lower value. The process continues indefinitely. We assume there are infinitely many miners, and at any finite time, a miner who successfully solves the hash puzzle is doing so for the first time almost surely.

Parties in the Protocol We refer to parties in the protocol as users. Users that contribute to modifying the ledger through appending blocks to the blockchain are called miners. A miner is either honest or corrupt. Honest miners follow the longest-chain protocol, whereas corrupt miners may deviate from the protocol. For simplicity, all corrupt users are grouped into a single entity called the adversary. This allows corrupt miners to communicate instantaneously, and captures the strong setting of perfect coordination between corrupt miners. The adversary can mine on any previously mined block, but its block must contain the proof-of-work to be valid. It can reveal its chain to any subset of honest users, and can delay its message by arbitrary amounts of time. It can also not include all the transactions it knows about that were not in ancestor blocks. We use β to represent the fraction of computational power in the system that is adversarial.

Mining Process The mining process is abstracted as follows. Let λ denote the total mining rate of the system. We consider a continuous time model where blocks are mined as a Poisson process with rate λ. Since β denotes the fraction of power that is adversarial and since successive mining instances are independent, adversarial block arrivals follow a Poisson process with rate $\beta\lambda$.

Blockchains and Blocktrees From any block, a unique sequence of blocks leading back to the genesis block can be identified via the hash pointers. We call this sequence a blockchain, or simply a chain. The convention is that the genesis block is the first block of a chain, and the terminating block is called the tip. At any given slot, honest users store a single chain in their memory.

Communication Delays We consider the setting where messages are either instantaneously delivered or permanently lost in an independent and identically distributed manner. Let $\mathsf{delay}\left(b_i \to b_j\right)$ denote the delay incurred by block b_i to reach the miner of block b_j. It is assumed that:

$$\mathsf{delay}\left(b_i \to b_j\right) = \begin{cases} 0, & \text{with probability } 1-d \\ \infty, & \text{with probability } d \end{cases}.$$

Here, d is the probability of message loss in an instance of point-to-point communication.

3 Main Result

In this section, we outline the desired security properties and present our main result. We define security on the level of transactions. It is desirable that a transaction eventually makes it to the ledger, and stays permanently at the same position in the ledger. This notion is formalized in Definition 3.1.

Definition 3.1 (Security)**.** *Let $\tau > 0$. Let $\mathcal{H}$ be any set of honest users. For any $h \in \mathcal{H}$, let $\mathcal{C}_h(t)$ denote the chain held by user h at time t. We say that a transaction* tx *made at some time s is $(\tau, \mathcal{H})$-secure if for any $h_1, h_2 \in \mathcal{H}$ and any $s_1, s_2 > s + \tau$, it holds that* tx *is included in a block b that is at the same position in $\mathcal{C}_{h_1}(s_1)$ and $\mathcal{C}_{h_2}(s_2)$.*

In the literature, security of a transaction is often defined as the confluence of *persistence* and *liveness*. A transaction satisfies liveness if it is eventually added to the ledger, and it satisfies persistence if it remains in the same position in the ledger for all future time, after a confirmation time. We remark that our definition of security implies these notions of persistence and liveness, and is consistent with existing definitions of security, such as in [GKL15, DKT+20]. Specifically, if a transaction tx satisfies $(\tau, \mathcal{H})$-security, then it is part of the chain held by all users in $\mathcal{H}$ before a confirmation time τ time has elapsed. Furthermore, once this confirmation time elapses, the transaction remains at the same position in the ledger for all future time. Our main result shows that if certain conditions are satisfied, then any transaction tx satisfies $(\tau, \mathcal{H})$ security except with a probability that decays almost exponentially in the confirmation time and scales linearly in the size of $\mathcal{H}$. It is stated as Theorem 3.2.

Theorem 3.2 (Main Result). *Let β be the fraction of computational power in the system that is adversarial, and d be the probability of message loss. If $\frac{\beta}{1-\beta} < (1-d)$, then for every $\varepsilon > 0$, there exist positive constants a and b such that for all $\tau \geq 0$ and for any honest transaction* tx *and any finite set of honest users $\mathcal{H}$:*

$$\mathbb{P}\left(\mathsf{tx} \textit{ violates } (\tau, \mathcal{H})\textit{-security}\right) \leq \exp\left(-a\tau^{1-\varepsilon}\right) + |\mathcal{H}| \exp\left(-b\tau\right).$$

The result states that under a certain threshold of tolerable adversarial power, the probability of security violation for any transaction and any finite set of users decays (almost) exponentially in the confirmation time. Hence, this violation probability can be made arbitrarily small by appropriately selecting the confirmation time. The sufficient condition $\frac{\beta}{1-\beta} < 1-d$ significantly improves the known threshold of tolerable adversarial power for the 0-∞ model (Figure 1). We also remark that our bound for the probability violation comprises of two terms, the latter of which scales linearly in $|\mathcal{H}|$. This linear scaling is expected, because no single message is successfully transmitted to all users in the model. Therefore, requiring a larger set of users to permanently adopt a transaction in their ledger requires a larger waiting time.

4 Definitions and Preliminaries

This section introduces key quantities that are used extensively in the analysis. In Section 4.1, we introduce the combinatorial objects on which the analysis is performed, such as the main-blocktree and the transmission-graph. Properties of these objects are presented alongside to motivate their purpose. In Section 4.2, key random variables that are amenable to analysis, and associated with special paths in the transmission-graph are identified. These quantities are illustrated through an example in Section 4.3. In Section 4.4, these random variables and are used to define 'catch-up events', and the notion of 'η-Nakamoto blocks'. Finally, Section 4.5 introduces the 'user-unheard-criterion', which is used to infer useful information about the blockchain held by a user.

4.1 Graphs and Trees

Three combinatorial objects at the core of our analysis are the main-blocktree and the transmission-graph.

Main-blocktree Any block can be uniquely traced back to the first block in the system (called genesis block). The set of all blocks generated (mined) up till time t forms a directed tree, which we refer to as the main-blocktree and denote it by $\mathsf{MB}(t)$. Here, $\mathsf{MB}(t) = (V_t, E_t)$, where the vertex set V_t is the set of all blocks mined up till time t and the set of directed edges E_t comprises all parent-to-child block pairs. $\mathsf{MB}(t)$ represents the global information about the system, and both honest and adversarial blocks are included in it. Figure 2(a) shows an example of the vicinity of the j-th honest block, b_j in some $\mathsf{MB}(t)$.

Definition 4.1 (Heights). *Let b_j be the j-th honest block. The height of b_j in a blocktree is the length of the directed path (counting edges) from the genesis block to b_j. We denote the height of b_j in the main-blocktree by* $\mathsf{height}_{\mathsf{MB}}(b_j)$.

Transmission-graph At any time t, we associate with the main-blocktree $\mathsf{MB}(t)$, a graph consisting of only honest blocks that we call the transmission-graph. The transmission-graph at time t, denoted $\mathsf{TG}(t)$ is a directed acyclic graph that represents the history of network delays among the honest miners. Here, $\mathsf{TG}(t) = (V_t, E_t)$, where the vertex set V_t is the set of honestly mined blocks up till time t. An edge between b_i and b_j is present if $\mathsf{delay}(b_i \to b_j) = 0$. A useful observation about $\mathsf{TG}(T)$ is presented in Lemma 4.2.

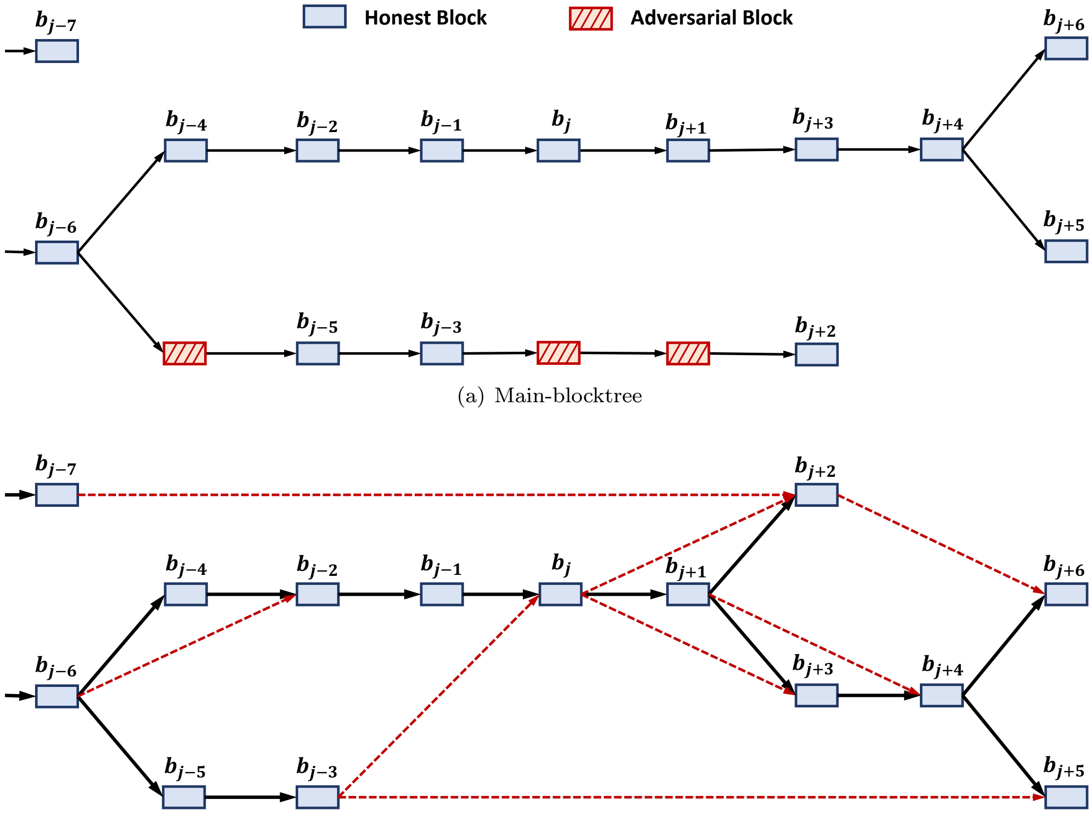

(a) Main-blocktree

(b) Transmission-graph

Figure 2: Graphs and trees used in analysis

Lemma 4.2. *Let b_i and b_k be the i-th and k-th honest blocks such that $i < k$. At any time t, suppose there exists a path $\mathcal{A}_{i,k}$ of length n from b_i to b_k in $\mathsf{TG}(t)$, i.e.*

$$\mathcal{A}_{i,k}\colon b_i = v_0 - v_1 - \cdots - v_n = b_k.$$

Then,

$$\mathsf{height}_{\mathsf{MB}}\left(b_k\right) - \mathsf{height}_{\mathsf{MB}}\left(b_i\right) \geq n$$

Proof. If v_{i-1} and v_i are two blocks in $\mathsf{TG}(t)$ such that there is an edge from v_{i-1} to v_i, then the miner of block v_i has heard of block v_{i-1}. Therefore, it adds its block at a height greater than that of v_{i-1}, so we have

$$\mathsf{height}_{\mathsf{MB}}\left(v_i\right) \geq \mathsf{height}_{\mathsf{MB}}\left(v_{i-1}\right) + 1$$

Repeatedly applying this inequality over the path from b_i to b_k yields the desired result. □

4.2 Special Sequences of Honest Blocks

Relative to the j-th honest block b_j, we define sequences of special blocks that correspond to forward and backward paths in $\mathsf{TG}(t)$. We also define notions of 'forward unheard' and 'backward unheard'.

4.2.1 Forward Special Blocks

Relative to the j-th honest block b_j, we define a sequence of 'forward special (FS)' blocks as follows.

Definition 4.3 (j-FS Sequence)**.** *Let $j \geq 0$ and let b_j be the j-th honest block. The j-FS sequence is a sequence of blocks $(b_j^0, b_j^1, b_j^2, \cdots)$ such that $b_j^0 = b_j$, and for all $k \geq 1$, b_j^k is the first block to hear of b_j^{k-1}.*

We refer to blocks in the j-FS sequence as j-FS blocks. For $0 \leq j < k$, denote by $\mathsf{FS}_j\left(b_k\right)$ the number of j-FS blocks mined between b_j and b_k (inclusive). Note that $\mathsf{FS}_j\left(b_{k-1}\right) \geq 1$, because b_j is always a j-FS block. The j-FS sequence associated with the transmission-graph in Figure 2(b) is shown in Figure 3. Notice that $\mathsf{FS}_j\left(b_{j+5}\right) = 3$ and $\mathsf{FS}_j\left(b_{j+6}\right) = 4$. This example is explored in more detail in Section 4.3.

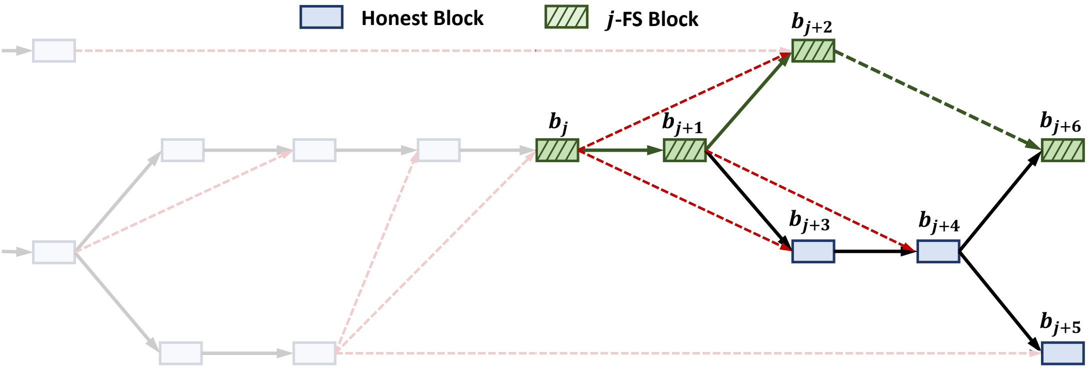

Figure 3: j-FS Sequence

Remark 4.4. *Let $0 \leq j \leq k$. Let d be the probability of message loss. The random variable $\mathsf{FS}_j\left(b_k\right)$ has the same distribution as $1+\sum_{i=1}^{k-j} \mathsf{Be}_i\left(1-d\right)$, where $\mathsf{Be}_i\left(1-d\right)$ are i.i.d. Bernoulli random variables with success probability equal to $1-d$. This is because b_j is a j-FS block, and every subsequent block is independently j-FS with probability $1-d$.*

Lemma 4.5. *Let $j \geq 0$. The j-FS sequence $(b_j^0, b_j^1, b_j^2, \cdots)$ is a forward directed path in the transmission-graph. Further, if $k > i \geq 0$, then the heights of the j-FS blocks b_j^i and b_j^k satisfy:*

$$\mathsf{height}_{\mathsf{MB}}\left(b_j^k\right) - \mathsf{height}_{\mathsf{MB}}\left(b_j^i\right) \geq k-i$$

Proof. Let $k \geq 0$. Since the miner of b_j^{k+1} has heard of b_j^k, there is an edge from b_j^k to b_j^{k+1}. The conclusion follows from Lemma 4.2. □

Definition 4.6 (Forward Unheard)**.** *Let $k \geq j \geq 0$. The Forward Unheard for block b_k with respect to block b_j is denoted $\mathsf{FU}_j\left(b_k\right)$ and defined as the number of consecutive j-FS blocks that the miner of b_k has not heard of, going backwards along the j-FS sequence from the last such block mined before b_k. If the miner of b_k has not heard of any j-FS block, then we toss independent biased coins (with failure probability equal to the probability of message loss) and continue to increment the count until a success is encountered.*

Remark 4.7. *Let $0 \leq j \leq k$. Let d be the probability of message loss. The random variable $\mathsf{FU}_j\left(b_k\right)$ has the same distribution as $\mathsf{Geom}\left(1-d\right)-1$, where $\mathsf{Geom}(1-d)$ is a geometric random variable, with minimum value 1. Further, if $k' \geq 0$ such that $k' \neq k$, then $\mathsf{FU}_j\left(b_k\right)$ and $\mathsf{FU}_j\left(b_{k'}\right)$ are independent.*

The intuition for defining $\mathsf{FU}_j\left(b_k\right)$ as above is illustrated through an example in Section 4.3.

4.2.2 Backward Relative Special Honest

Relative to the j-th honest block b_j, we define a sequence of 'backward special (BS) blocks as follows.

Definition 4.8 (j-BS Sequence)**.** *Let $j \geq 0$, and let b_j be the j-th honest block. The j-BS sequence is a sequence of blocks $(b_j^0, b_j^{-1}, b_j^{-2}, \cdots)$ such that $b_j^0 = b_j$ and b_j^{-i} is the most recently mined block heard by $b_j^{-(i-1)}$ for all $i \geq 1$.*

For $0 \leq i \leq j$, denote by $\mathsf{BS}_j\left(b_i\right)$ the number of j-BS blocks mined between b_i and b_j (inclusive). Note that $\mathsf{BS}_j\left(b_i\right) \geq 1$, since b_j is always a j-BS block. The j-BS sequence associated with the transmission-graph in Figure 2(b) is shown in Figure 4. Notice that $\mathsf{BS}_j\left(b_{j-5}\right) = 4$ and $\mathsf{BS}_j\left(b_{j-6}\right) = 5$. This example is explored in more detail in Section 4.3.

Remark 4.9. *Let $0 \leq i \leq j$. Let d be the probability of message loss. The random variable $\mathsf{BS}_i\left(b_j\right)$ has the same distribution as $1+\sum_{k=1}^{j-i} \mathsf{Be}_k\left(1-d\right)$, where $\mathsf{Be}_k\left(1-d\right)$ are i.i.d. Bernoulli random variables with success probability equal to $1-d$. This is because b_j is a j-BS block, and every previous block is independently j-BS with probability $1-d$.*

Lemma 4.10. *Let $j \geq 0$. The j-BS sequence $(b_j^0, b_j^{-1}, b_j^{-2}, \cdots)$ is a backward directed path in transmission-graph. Further, if $k > i \geq 0$, then the heights of the j-BS blocks b_j^{-i} and b_j^{-k} satisfy:*

$$\mathsf{height}_{\mathsf{MB}}\left(b_j^{-k}\right) - \mathsf{height}_{\mathsf{MB}}\left(b_j^{-i}\right) \geq k-i$$

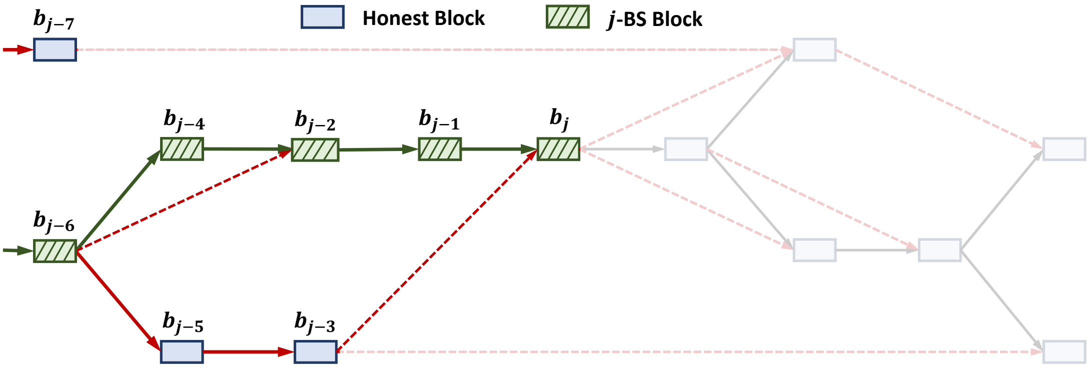

Figure 4: j-BS Sequence

Proof. Let $i \geq 0$. Since the miner of $b_j^{-(i-1)}$ has heard of b_j^{-i}, there is an edge from b_j^{-i} to $b_j^{-(i-1)}$. The conclusion follows from Lemma 4.2. □

Definition 4.11 (Backward Unheard)**.** *Let $j \geq i \geq 0$. The Backward Unheard for block b_i with respect to block b_j is denoted $\mathsf{BU}_j(b_i)$ and defined as the number of consecutive j-BS blocks whose miners have not heard of b_i, going backwards along the j-BS sequence from the first such block mined after b_i. If none of the miners of j-BS blocks have heard of b_i, then we toss independent biased cones (with failure probability equal to the probability of message loss) and continue to increment the count until a success is encountered.*

Remark 4.12. *Let $0 \leq i \leq j$. Let d be the probability of message loss. The random variable $\mathsf{BU}_j(b_i)$ has the same distribution as $\mathsf{Geom}(1-d) - 1$, where $\mathsf{Geom}(1-d)$ is a geometric random variable, with minimum value 1. Further, if $i' \geq 0$ such that $i' \neq i$, then $\mathsf{BU}_j(b_i)$ and $\mathsf{BU}_j(b_{i'})$ are independent.*

The intuition for defining $\mathsf{BU}_j(b_i)$ as above is illustrated through an example in Section 4.3. The usefulness of these quantities is evident from Lemma 4.13.

Lemma 4.13. *Let $i < j < k$.*

(i) If $\mathsf{FS}_j(b_{k-1}) > \mathsf{FU}_j(b_k)$, then $\mathsf{height}_{\mathsf{MB}}(b_k) - \mathsf{height}_{\mathsf{MB}}(b_j) \geq \mathsf{FS}_j(b_{k-1}) - \mathsf{FU}_j(b_k)$.

(ii) If $\mathsf{BS}_j(b_{i+1}) > \mathsf{BU}_j(b_i)$, then $\mathsf{height}_{\mathsf{MB}}(b_j) - \mathsf{height}_{\mathsf{MB}}(b_i) \geq \mathsf{BS}_j(b_{i+1}) - \mathsf{BU}_j(b_i)$.

Proof. We prove statement (i). Let $m = \mathsf{FS}_j(b_{k-1}) - \mathsf{FU}_j(b_k) > 0$. Consider the j-FS sequence $(b_j^0, b_j^1, \cdots, b_j^m, \cdots)$. Note that $\mathsf{FS}_j(b_{k-1})$ is the number of j-FS blocks mined before b_k, and $\mathsf{FU}_j(b_k)$ is the number of consecutive blocks from this sequence that were not heard by b_k going backward. Therefore, b_k has heard of b_j^m, and will mine at a greater height. Since the j-FS blocks are all mined at different heights, we have

$$\mathsf{height}_{\mathsf{MB}}(b_k) - \mathsf{height}_{\mathsf{MB}}(b_j) > \mathsf{height}_{\mathsf{MB}}\left(b_j^m\right) - \mathsf{height}_{\mathsf{MB}}\left(b_j^0\right) \geq m = \mathsf{FS}_j(b_{k-1}) - \mathsf{FU}_j(b_k),$$

as desired. The proof of statement (ii) is essentially the same, because the forward sequences FS_j and FU_j map to the backward sequences BS_j and BU_j under reversing the directions of the edges in the transmission-graph. □

4.3 An Example

Some of the concepts introduced above are best understood through an example. Consider the main-blocktree in Figure 2(a) and its associated transmission-graph in Figure 2(b). The same transmission-graph is shown again in Figure 5, where the j-FS and j-BS blocks are highlighted. Recall that in the transmission-graph, a directed edge from b_j to b_k indicates that the miner of b_k has heard of the block b_j.

j-FS Sequence By definition, b_j is a j-FS block. The miner of the next block b_{j+1} has heard of the most recent j-FS block b_j, so $b_j^1 = b_{j+1}$ is a j-FS block. The miner of b_{j+2} has heard of b_j^1, so we have that $b_{j+2} = b_j^2$ is also a j-FS block. However, the miners of b_{j+3}, b_{j+4} and b_{j+5} have not heard of b_j^2, so these blocks are not j-FS. Finally, the miner of b_{j+6} has heard of b_j^2, so we have that $b_j^3 = b_{j+6}$. Notice that no two j-FS blocks can share the same height.

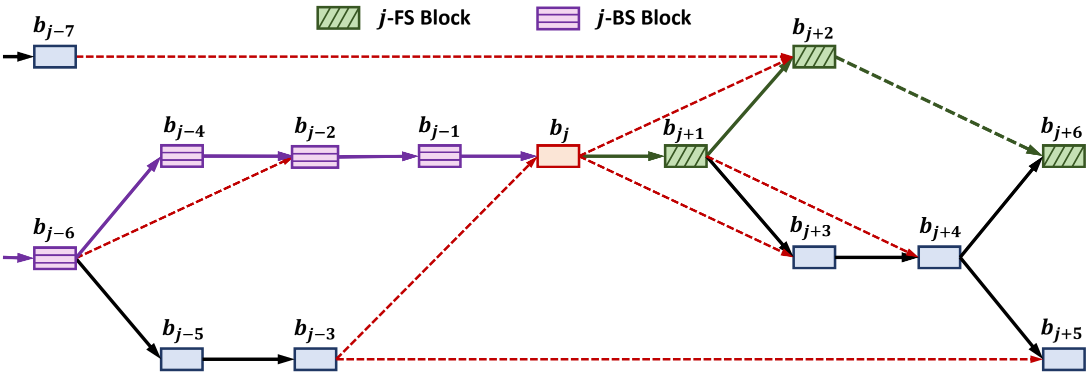

Figure 5: Example of j-FS and j-BS sequence

Forward Unheard Let us consider $\mathsf{FU}_j(b_k)$ with $k = j + 4$. Starting from b_k, we count the number of consecutive j-FS blocks that the miner of b_k has not heard of, going backwards in the j-FS sequence. The j-FS sequence at the mining time of b_k is (b_j, b_{j+1}, b_{j+2}). Going backwards in this sequence, we see that the miner of b_k has not heard of b_{j+2}, but has heard of b_{j+1}. Therefore, we stop the count and have $\mathsf{FU}_j(b_{j+4}) = 1$.

j-BS Sequence By definition, b_j is a j-BS block. Since the miner of the most recent j-BS block (b_j) has heard of the previous block b_{j+1}, we have that $b_{j-1} = b_j^{-1}$ is a j-BS block. Similarly, the miner of b_j^{-1} has heard of b_{j-2}, so we have that $b_{j-2} = b_j^{-2}$ is a j-BS block. However, the miner of b_j^{-2} has not heard of b_{j-3}, and so b_{j-3} is not a j-BS block. Notice that no two j-BS blocks can share the same height.

Backward Unheard Let us consider $\mathsf{BU}_j(b_i)$ with $i = j - 3$. Starting from the first j-BS block mined after b_i, we count the number of consecutive j-BS blocks that have not heard of b_i, going backwards in the j-BS sequence. The j-BS sequence at the mining time of b_i is b_j, b_{j-1}, b_{j-2}. Going backwards in this sequence, we see that the miners of b_{j-2} and b_{j-1} have not heard of b_i, but the miner of b_j has. Therefore, we stop the count and have $\mathsf{BU}_j(b_{j-3}) = 2$.

4.4 Catch-up Events and η-Nakamoto Blocks

In this section, we define catch-up events and η-Nakamoto blocks.

Adversarial arrivals Let $a(b_i, b_j)$ denote the number of adversarial blocks mined between the mining times of the i-th and j-th honest blocks. Similarly, let $a(b_j^i, b_j^k)$ denote the number of adversarial blocks mined between the i-th and k-th j-FS blocks.

Definition 4.14 (Catch-up Events)**.** *Let* $0 \le i < j < k$, *and let* $0 < \eta \le 1$. *The forward and backward catch-up events are respectively defined as:*

$$\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)} : a(b_j, b_k) \ge \eta \cdot \mathsf{FS}_j(b_{k-1}) - \mathsf{FU}_j(b_k) \tag{1}$$

$$\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)} : a(b_i, b_j) \ge \eta \cdot \mathsf{BS}_j(b_{i+1}) - \mathsf{BU}_j(b_i) \tag{2}$$

These events are catch-up events in the following sense. If the event $\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}$ occurs for some $j < k$ and $0 < \eta \le 1$, then more adversarial blocks have been mined in $[\tau_j, \tau_k]$ than *effective* j-FS blocks. For instance, when $\eta = 1$, more adversarial blocks have been mined in $[\tau_j, \tau_k]$ than the number of blocks in the j-FS sequence from b_j up to the last j-FS block heard by the miner of b_k. If these adversarial blocks were to form a side chain rooted at b_j, then there is a possibility that b_k mines on this side chain. However, such an attack would fail if the catch-up events did not occur, because there would necessarily be a chain longer than the adversarial side chain that the miner of b_k is aware of. Here, η is a robustness measure: If the catch-up event does not occur for a small value of η, then the honest blocks have a considerable lead over adversarial side chains.

Longest Chains of the main-blocktree A longest chain of the main-blocktree at time t is a path in $\mathsf{MB}(t)$ whose length is no shorter than any path in $\mathsf{MB}(t)$. Notice that $\mathsf{MB}(t)$ can have multiple longest chains.

η-Nakamoto Blocks A desirable property for honest block b_j is the existence of some $0 < \eta \leq 1$ such that none of the events $\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}$ occur for any $i < j$ and none of the events $\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}$ occur for any $k > j$. Intuitively, this ensures that adversarial chains rooted at any block b_i with $i < j$ are never long enough for any block b_k with $k > j$ to extend them. In turn, this suggests that block b_j would be in every longest chain of the main-blocktree right from being mined. We formalize this property using the notion of η-Nakamoto blocks below, and make this intuition rigorous in Section 5.

Definition 4.15 (η-Nakamoto Block)**.** *Let $j \geq 0$ and $0 < \eta \leq 1$. The j-th honest block b_j is said to be an η-Nakamoto block if the event*

$$\mathsf{N}_j^{(\eta)}: \left[\bigcap_{i:\, i<j} \left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^{\mathsf{c}}\right] \bigcap \left[\bigcap_{k:\, k>j} \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}\right] \tag{3}$$

occurs.

We remark that an η-Nakamoto block is also an η'-Nakamoto block, for any $\eta' \in [\eta, 1]$.

4.5 User-unheard-criterion

η-Nakamoto blocks are useful because they belong to every longest chain of the main-blocktree. However, this is not equivalent to η-Nakamoto blocks belonging to any *user's* chain for all future time. This is because at any given time, a given user has not necessarily heard of all the blocks in $\mathsf{MB}(t)$. In this section, we introduce the tools that are relevant to analyzing the state of a user's chain with respect to the main-blocktree. We begin by introducing 'Unheard' with respect to a user.

Definition 4.16 (User-unheard, $\mathsf{Unheard}_h\left(b_j^k\right)$)**.** *Let $j \geq 0$ and let b_j be the j-th honest block. Let $(b_j^0, b_j^1, \cdots)$ denote the j-FS sequence. Let h be an honest user. For any $k \geq 0$, we define $\mathsf{Unheard}_h\left(b_j^k\right)$ as the number of consecutive j-FS blocks that the user h has not heard of, going backwards along the j-FS sequence from b_j^k. If h has not heard of any j-FS block, then we toss independent biased coins (with failure probability equal to the probability of message loss) and continue to increment the count until a success is encountered.*

Remark 4.17. *Let $j, k \geq 0$. Let h be any honest user and d be the probability of message loss. Let $\mathsf{Geom}(1-d)$ be a geometric random variable, with minimum value 1. The distribution of the random variable $\mathsf{Unheard}_h\left(b_j^k\right)$ depends on h:*

- *If h has not mined any block after the mining time of b_j, then $\mathsf{Unheard}_h\left(b_j^k\right)$ has the same distribution as $\mathsf{Geom}\left(1-d\right) - 1$.*
- *If h is a miner of a j-FS block, then h has heard of its own block as well as the j-FS block that immediately preceded it. However, all other delays from the miners of j-FS blocks to h are still independent and identically distributed. In this case, $\mathsf{Unheard}_h\left(b_j^k\right)$ is stochastically dominated by $\mathsf{Geom}(1-d)$.*
- *If h is a miner of a block b_ℓ with $\ell > j$, such that b_ℓ is not a j-FS block, then the delay from the most recent j-FS block before b_ℓ to h is infinity. However, all other delays from the miners of j-FS blocks to h are still independent and identically distributed. In this case, $\mathsf{Unheard}_h\left(b_j^k\right)$ is stochastically dominated by $\mathsf{Geom}(1-d)$.*

In all cases, we have $\mathsf{Unheard}_h\left(b_j^k\right) \leq \mathsf{Geom}\left(1-d\right)$, where the $\leq$ sign indicates stochastic domination.

We compare 'forward-unheard' (denoted $\mathsf{FU}_j\left(b_k\right)$) and 'user-unheard' (denoted $\mathsf{Unheard}_h\left(b_j^k\right)$). Although similar in spirit, the quantity $\mathsf{FU}_j\left(b_k\right)$ counts the number of consecutive j-FS blocks not heard by the *miner* of block b_k, going backwards in the j-FS sequence from the last such block before b_k, whereas the quantity $\mathsf{Unheard}_h\left(b_j^k\right)$ fixes an honest user h and similarly counts the number of consecutive such blocks not heard by h. Since b_k and b_ℓ are mined by two different miners, then $\mathsf{FU}_j\left(b_k\right)$ and $\mathsf{FU}_j\left(b_\ell\right)$ are independent random variables. However, $\mathsf{Unheard}_h\left(b_j^k\right)$ and $\mathsf{Unheard}_h\left(b_j^\ell\right)$ need not be independent.

Next, we introduce the user-unheard-criterion, which will later allow us to infer useful information about the state of a user's chain from the main-blocktree.

Definition 4.18 (User-unheard-criterion)**.** *Let h be an honest user. Let $j \geq 1$, $k_0 \geq 1$, and $0 < \eta < 1$. We say that the (h, j, η, k_0)-user-unheard-criterion is satisfied if*

$$\mathsf{Unheard}_h\left(b_j^k\right) < \left(\frac{1-\eta}{2}\right) k, \quad \forall k \geq k_0, \tag{4}$$

5 Proof Outline

In this section, we present an outline of the proof of Theorem 3.2. First, deterministic and probabilistic results are stated. These results are used as building blocks in the proof sketch of our main result, Theorem 3.2 which is presented in Section 5.3. Rigorous proofs of all the results, including the main result are relegated to the appendix.

5.1 Deterministic Results

Theorem 5.1. *Let $j \geq 1$ and $0 < \eta \leq 1$. Let b_j be the j-th honest block, and τ_j be the mining time of b_j. Suppose b_j is an η-Nakamoto block. Then:*

(i) b_j is the unique honest block at its height in $\mathsf{MB}(t)$ for all $t \geq \tau_j$.

(ii) b_j is in every longest chain of $\mathsf{MB}(t)$ for all $t \geq \tau_j$.

(iii) b_j is the unique block at its height in $\mathsf{MB}(\tau_j)$.

(iv) For any k such that $k > j$, b_k is a descendant of b_j in the main-blocktree.

Proof. See Appendix A.1 □

Theorem 5.2. *Let h be an honest user. Let $j \geq 1$ and $0 < \eta < 1$. Let $k_0 = \left\lceil \frac{2\eta}{1-\eta} \right\rceil$. Let $\mathcal{C}_h(t)$ denote the chain held by h at time t. If b_j is an η-Nakamoto block, and if the (h, j, η, k_0)-user-unheard-criterion is satisfied, then $b_j \in \mathcal{C}_h(t)$ for all $t \geq \tau_j^{k_0}$, where $\tau_j^{k_0}$ is the mining time of the k_0-th j-FS block.*

Proof. See Appendix A.2. □

5.2 Probabilistic Results

Lemma 5.3. *Let $0 \leq i < j \leq j' < k$, and $0 \leq \eta \leq 1$. The events:*

- *$\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}$ and $\overrightarrow{\mathsf{B}}^{(\eta)}_{j',k}$ are independent.*
- *$\overrightarrow{\mathsf{B}}^{(\eta)}_{i,j}$ and $\overleftarrow{\mathsf{B}}^{(\eta)}_{j',k}$ are independent.*

Proof. See Appendix B.1. □

Theorem 5.4. *Let $j \geq 1$ and $0 \leq \eta \leq 1$. Recall that b_j is the j-th honest block and $\mathsf{N}_j^{(\eta)}$ is the event that b_j is an η-Nakamoto block. If $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$, then there exists a positive constant $p_0 > 0$ such that*

$$\mathbb{P}\left(\mathsf{N}_j^{(\eta)}\right) \geq p_0 > 0$$

Proof. See Appendix B.2. □

Lemma 5.5. *Let $0 < \eta \leq 1$ and $i < j < k$. Let $\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}$ and $\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}$ be the catch-up events defined in (1), (2). If $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$, there exists a constant $c > 0$ such that*

$$\mathbb{P}\left(\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right) \leq e^{-c(k-j)}$$
$$\mathbb{P}\left(\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right) \leq e^{-c(j-i)}$$

Proof. See Appendix B.3. □

Theorem 5.6. *Let $0 < \eta \leq 1$. Let β be the fraction of computational power in the system that is adversarial and d be the probability of message loss. Let $B^{(\eta)}_{s,s+t}$ be the event that there are no η-Nakamoto blocks in $[s, s+t]$. If $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$, then there exists a constant $c_0 > 0$ such that for any $s, t \geq 0$,*

$$\mathbb{P}\left(B^{(\eta)}_{s,s+t}\right) \leq e^{-c_0\sqrt{t}}$$

Proof. See Appendix B.4. □

Theorem 5.7. *Let $0 < \eta \leq 1$. Let β be the fraction of computational power in the system that is adversarial and d be the probability of message loss. Let $B^{(\eta)}_{s,s+t}$ be the event that there are no η-Nakamoto blocks in $[s, s+t]$. If $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$, then for every $\varepsilon > 0$, there exist positive constants A, a such that for any $s, t > 0$,*

$$\mathbb{P}\left(B^{(\eta)}_{s,s+t}\right) \leq A \exp\left(-at^{1-\varepsilon}\right).$$

Proof. See Appendix B.5. □

Theorem 5.8. *Let $0 < \eta < 1$. Suppose the fraction β of computational power in the system that is adversarial, and the probability d of message loss satisfies $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$. Given $s \geq 0$, let b_J be the first η-Nakamoto block mined after time s. There exist constants C, $c > 0$ such that for any honest user h and for all $k' \geq 1$,*

$$\mathbb{P}\left((h, J, \eta, k')\text{-user-unheard-criterion fails}\right) \leq Ce^{-ck'}$$

Proof. See Appendix B.6. □

5.3 Proof Sketch

η-Nakamoto blocks are special blocks that are part of every longest chain in the main-blocktree, for all time after they are mined (Theorem 5.1). Therefore, if a transaction is included in an η-Nakamoto block or in any of its ancestors, then it will be included in every longest chain of the main-blocktree. Furthermore, if an honest user is up-to-date with the main-blocktree (specifically, by satisfying a relevant user-unheard-criterion, and therefore having heard of more forward special blocks with respect to the η-Nakamoto block, than adversarial blocks), then it is reasonable to expect that the honest user will also include the transaction in its chain. This idea is formalized in the user-unheard-criterion, and made rigorous in Theorem 5.2.

The core idea behind the proof of the main result (Theorem 3.2) is to show that η-Nakamoto blocks occur frequently with high probability (Theorem 5.7). In turn, this implies that a transaction is highly likely to be included in an η-Nakamoto block or its ancestor, soon after it is made. Theorem 5.8 then shows that it is also highly likely that a relevant user-unheard-criterion will hold for the honest user, once a waiting time has elapsed.

To prove Theorem 5.7, we bootstrap from a milder version of it, which is stated as Theorem 5.6. Inspired by the proof strategy in [DKT$^+$20], we prove Theorem 5.6 by separating catch-up events into long and short-term catch-up events. We show that long term catch-ups occur rarely as a consequence of Lemma 5.5, and short-term catch-up probabilities are bounded using a lower bound on the probability that the j-th block is an η-Nakamoto block (Theorem 5.4), and the fact that non-overlapping catch-up events are independent (Lemma 5.3). This establishes that a transaction is highly likely to be included in the main-blocktree for all future time, once a waiting period has elapsed. This proves thereom 5.7.

Next, we show that an honest user is likely to always be up-to-date with the main-blocktree, forever after a waiting time. This is done by explicitly bounding the probability that a relevant user-unheard-criterion is violated using tools from stochastic analysis. This proves Theorem 5.8.

Theorems 5.7 and 5.8 together imply the main result. The details of this proof are presented in Appendix B.7.

6 Conclusion

In this work, we introduced the 0-∞ model: a framework to study the impact of random message losses on the security of the proof-of-work longest-chain protocol. We investigated the security of this protocol by analyzing the transmission-graph, a dynamically evolving graph that captures the delays incurred by the blocks mined by honest miners. Specifically, we studied special sequences of blocks and identified random variables associated with them that are amenable to analysis. These random variables were used to define useful objects and desirable events, such as η-Nakamoto blocks and user-unheard-criterion respectively. These ideas allowed us to generalize analysis techniques from the synchronous delay model to a setting where delays are possibly infinite. We showed that the condition $\frac{\beta}{1-\beta} < 1-d$ is sufficient for a transaction to satisfy desired security properties except with a probability that decays almost exponentially in the security parameter. This greatly improved the known threshold of the fraction of adversarial power that is tolerable for a given probability of message loss in an instance of point-to-point communication.

7 Acknowledgement

This work was supported in part by NSF Grant Grant CCF 19-00636.

References

[BKM$^+$20] Erica Blum, Aggelos Kiayias, Cristopher Moore, Saad Quader, and Alexander Russell. The combinatorics of the longest-chain rule: Linear consistency for proof-of-stake blockchains. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1135–1154. SIAM, 2020.

[BSAB$^+$19] Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, Patrick McCorry, Sarah Meiklejohn, and George Danezis. Sok: Consensus in the age of blockchains. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pages 183–198, 2019.

[DKT$^+$20] Amir Dembo, Sreeram Kannan, Ertem Nusret Tas, David Tse, Pramod Viswanath, Xuechao Wang, and Ofer Zeitouni. Everything is a race and Nakamoto always wins. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 859–878, 2020.

[DLS88] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)*, 35(2):288–323, 1988.

[FZ17] Lei Fan and Hong-Sheng Zhou. A scalable proof-of-stake blockchain in the open setting (or, how to mimic nakamoto's design via proof-of-stake). Technical report, Cryptology ePrint Archive, Report 2017/656, 2017.

[GK20] Juan Garay and Aggelos Kiayias. Sok: A consensus taxonomy in the blockchain era. In *Cryptographers' Track at the RSA Conference*, pages 284–318. Springer, 2020.

[GKL15] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 281–310. Springer, 2015.

[GKL20] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. Full analysis of nakamoto consensus in bounded-delay networks. *Cryptology ePrint Archive*, 2020.

[GKR20] Peter Gaži, Aggelos Kiayias, and Alexander Russell. Tight consistency bounds for bitcoin. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 819–838, 2020.

[Kin64] John F Kingman. A martingale inequality in the theory of queues. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 60, pages 359–361. Cambridge University Press, 1964.

[KRDO17] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference*, pages 357–388. Springer, 2017.

[LGR21] Jing Li, Dongning Guo, and Ling Ren. Close latency-security trade-off for the nakamoto consensus. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 100–113, 2021.

[MU17] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomization and probabilistic techniques in algorithms and data analysis.* Cambridge University Press, 2017.

[Nak08] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.

[NTT21] Joachim Neu, Ertem Nusret Tas, and David Tse. Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 446–465. IEEE, 2021.

[PS17] Rafael Pass and Elaine Shi. The sleepy model of consensus. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 380–409. Springer, 2017.

[PSS17] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2017.

[Ren19] Ling Ren. Analysis of nakamoto consensus. *IACR Cryptol. ePrint Arch.*, 2019:943, 2019.

[SGH21] Suryanarayana Sankagiri, Shreyas Gandlur, and Bruce Hajek. The longest-chain protocol under random delays. *arXiv preprint arXiv:2102.00973*, 2021.

A Proofs of Deterministic Results

A.1 Proof of Theorem 5.1

Theorem 5.1. *Let $j \geq 1$ and $0 < \eta \leq 1$. Let b_j be the j-th honest block, and τ_j be the mining time of b_j. Suppose b_j is an η-Nakamoto block. Then:*

(i) b_j is the unique honest block at its height in $\mathsf{MB}(t)$ for all $t \geq \tau_j$.

(ii) b_j is in every longest chain of $\mathsf{MB}(t)$ for all $t \geq \tau_j$.

(iii) b_j is the unique block at its height in $\mathsf{MB}(\tau_j)$.

(iv) For any k such that $k > j$, b_k is a descendant of b_j in the main-blocktree.

Proof. Let $\mathsf{N}_j^{(\eta)}$ be the event that b_j is an η-Nakamoto block. Recall that

$$\mathsf{N}_j^{(\eta)} = \left[\bigcap_{i:\, i<j} \left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^{\mathsf{c}}\right] \bigcap \left[\bigcap_{k:\, k>j} \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}\right]$$

Suppose $\mathsf{N}_j^{(\eta)}$ occurs, so that b_j is an η-Nakamoto block.

(i) Let $t \geq \tau_j$. We show that b_j is the unique honest block at its height in $\mathsf{MB}(t)$, i.e. for any $i, k \geq 0$ such that $i < j < k$, we show that

$$\mathsf{height}_{\mathsf{MB}}(b_i) < \mathsf{height}_{\mathsf{MB}}(b_j) < \mathsf{height}_{\mathsf{MB}}(b_k). \tag{5}$$

We begin by proving the first inequality in (5). Since $\mathsf{N}_j^{(\eta)}$ occurs, it follows that the event $\left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^{\mathsf{c}}$ occurs. Therefore, we have

$$\begin{aligned}
\mathsf{height}_{\mathsf{MB}}(b_j) - \mathsf{height}_{\mathsf{MB}}(b_i) &\overset{(a)}{\geq} \mathsf{BS}_j(b_{i+1}) - \mathsf{BU}_j(b_i) \\
&\overset{(b)}{\geq} \eta \cdot \mathsf{BS}_j(b_{i+1}) - \mathsf{BU}_j(b_i) \\
&\overset{(c)}{>} a(b_i, b_j),
\end{aligned}$$

where (a) is from Lemma 4.13, (b) follows from the fact that $\eta \leq 1$, and (c) is the definition of $\left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^{\mathsf{c}}$. Since $a(b_i, b_j) \geq 0$, it follows that $\mathsf{height}_{\mathsf{MB}}(b_i) < \mathsf{height}_{\mathsf{MB}}(b_j)$. Similarly, the occurrence of $\left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}$ implies that $\mathsf{height}_{\mathsf{MB}}(b_j) < \mathsf{height}_{\mathsf{MB}}(b_k)$. Thus, b_j is the unique honest block at its height in $\mathsf{MB}(t)$.

(ii) We show that b_j is in every longest chain of $\mathsf{MB}(t)$. Let $\mathcal{V}$ be any longest chain in $\mathsf{MB}(t)$, i.e.

$$\mathcal{V} = v_0 - v_1 - \cdots - v_{\mathsf{height}_{\mathsf{MB}}(b_j)} - \cdots - v_{m-1} - v_m,$$

where v_0 is the genesis block, $v_{\mathsf{height}_{\mathsf{MB}}(b_j)} =: v$ is the block at the height of b_j, and m is the height of $\mathsf{MB}(t)$. It suffices to show that $v = b_j$. In fact, since b_j is the unique honest block at its height, it suffices to only show that v must be an honest block. Starting from v and traversing blocks in $\mathcal{V}$ in the backward direction, let b_i denote the first honest block encountered in $\mathcal{V}$, not including v itself. Since the genesis block is honest, such a b_i exists. Similarly, we also want to traverse blocks in $\mathcal{V}$ along the forward direction. Here, we have two cases:

Case 1: There is an honest block in $\mathcal{V}$ after v. In this case, let b_k be the first honest block in $\mathcal{V}$ after v. Let H denote $\mathsf{height}_{\mathsf{MB}}\left(b_k\right)$, and $\mathcal{V}_1$ be the sequence $b_i - \cdots - v - \cdots - b_k$.

Case 2: There is no honest block in $\mathcal{V}$ after v. In this case, let b_k be the first honest block mined after time t. Let $H = m$ denote the height of the main-blocktree, and $\mathcal{V}_1$ be the sequence $b_i - \cdots - v - \cdots - v_m$.

In both cases, let N denote the number of blocks in $\mathcal{V}_1$, not including b_i. Notice that the blocks are at consecutive heights. We have:

$$
\begin{aligned}
N &= H - \mathsf{height}_{\mathsf{MB}}\left(b_i\right) \\
&\overset{(a)}{\geq} \mathsf{height}_{\mathsf{MB}}\left(b_j\right) + \mathsf{FS}_j\left(b_{k-1}\right) - \mathsf{FU}_j\left(b_k\right) - \mathsf{height}_{\mathsf{MB}}\left(b_i\right) \\
&\overset{(b)}{\geq} \left(\mathsf{FS}_j\left(b_{k-1}\right) - \mathsf{FU}_j\left(b_k\right)\right) + \left(\mathsf{BS}_j\left(b_{i+1}\right) - \mathsf{BU}_j\left(b_i\right)\right) \\
&\overset{(c)}{\geq} \left(a(b_i, b_j) + 1\right) + \left(a(b_j, b_k) + 1\right) \\
&> a(b_i, b_k) + 1,
\end{aligned}
$$

where (a) follows from Lemma 4.13 in case 1, and the fact that j-FS blocks are mined at different heights in case 2. Further, (b) is due to Lemma 4.13 and the fact that $\eta \leq 1$, and (c) follows because the events $\left[\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right]^{\mathsf{c}}$ and $\left[\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right]^{\mathsf{c}}$ occur. However, by definition of b_i and b_k, all blocks in $\mathcal{V}_1$ after b_i except for b_k and possibly v are adversarial. Therefore, there must be at least one honest block in $\mathcal{V}_1$ strictly between b_i and b_k. Since the only possibility for this is v, it follows that v must be an honest block. Since b_j is the unique honest block at the height of v, we conclude that b_j in every longest chain of $\mathsf{MB}(t)$, as desired.

(iii) Let v be a block in $\mathsf{MB}(\tau_j)$ such that $\mathsf{height}_{\mathsf{MB}}\left(v\right) = \mathsf{height}_{\mathsf{MB}}\left(b_j\right)$. Let $\mathcal{T}$ be the tine of ancestors of v. Traversing blocks along the backward direction starting from v, let b_i denote the first honest block encountered. The portion of $\mathcal{T}$ between b_i and v consists of only adversarial blocks. However,

$$
\begin{aligned}
\mathsf{height}_{\mathsf{MB}}\left(v\right) - \mathsf{height}_{\mathsf{MB}}\left(b_i\right) &= \mathsf{height}_{\mathsf{MB}}\left(b_j\right) - \mathsf{height}_{\mathsf{MB}}\left(b_i\right) \\
&\overset{(a)}{\geq} \eta \cdot \mathsf{BS}_j\left(b_{i+1}\right) - \mathsf{BU}_j\left(b_i\right) \\
&\overset{(b)}{>} a(b_i, b_j),
\end{aligned}
$$

where (a) is from Lemma 4.13, and (b) is true because b_j is an η-Nakamoto block. Since v is the only possibility for an honest block in the portion of $\mathcal{T}$ after b_i, it follows that v must be an honest block. From (i), we conclude that $v = b_j$. This concludes the proof.

(iv) Let $k > j$. Since b_j is an η-Nakamoto block, we have from the proof of (i) that $\mathsf{height}_{\mathsf{MB}}\left(b_k\right) > \mathsf{height}_{\mathsf{MB}}\left(b_j\right)$. Let v denote the ancestor of b_k at the height of b_j in the main-blocktree. Starting from v and traversing blocks along the ancestors of v, let b_i denote the first honest block encountered. This is exactly case 1 in the proof of (ii), and it follows that $v = b_j$.

□

A.2 Proof of Theorem 5.2

Theorem 5.2. *Let h be an honest user. Let $j \geq 1$ and $0 < \eta < 1$. Let $k_0 = \left\lceil \frac{2\eta}{1-\eta} \right\rceil$. Let $\mathcal{C}_h(t)$ denote the chain held by h at time t. If b_j is an η-Nakamoto block, and if the (h, j, η, k_0)-user-unheard-criterion is satisfied, then $b_j \in \mathcal{C}_h(t)$ for all $t \geq \tau_j^{k_0}$, where $\tau_j^{k_0}$ is the mining time of the k_0-th j-FS block.*

Proof. Let $|\mathcal{C}_h(t)|$ denote the number of blocks in $\mathcal{C}_h(t)$. Let τ_j denote the mining time of b_j, and let τ_j^k denote the mining time of b_j^k. Since b_j is an η-Nakamoto block, we know that the event $\left[\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right]^{\mathsf{c}}$ occurs for all $k > j$. In turn, this implies

$$a(b_j^0, b_j^{k+1}) < \eta \cdot (k+1) \quad \forall k > 0.$$

However, $\eta\,(k+1) \le \left(\frac{1+\eta}{2}\right) k$ whenever $k \ge k_0$. Further, since (h, j, η, k_0)-user-unheard-criterion holds, we have that $\mathsf{Unheard}_h\left(b_j^k\right) < \left(\frac{1-\eta}{2}\right) k$ whenever $k \ge k_0$. These facts imply:

$$a(b_j^0, b_j^{k+1}) < k - \mathsf{Unheard}_h\left(b_j^k\right) \quad \forall k \ge k_0. \tag{6}$$

Fix $k \ge k_0$. Let t be such that $\tau_j^k \le t < \tau_j^{k+1}$. We show that h includes b_j in its chain at time t.

Since $a(b_j^0, b_j^{k+1}) \ge 0$, it follows from (6) that $\mathsf{Unheard}_h\left(b_j^k\right) < k$. Therefore, h has heard of at least one j-FS block. Therefore, $|\mathcal{C}_h(t)| \ge \mathsf{height}_{\mathsf{MB}}\,(b_j)$. Let $v \in \mathcal{C}_h(t)$ be the block at the height of b_j. To show that $b_j \in \mathcal{C}_h(t)$, it suffices by statement (iv) in Theorem 5.1 to prove that $\mathcal{C}_h(t)$ contains at least one honest block mined in $[\tau_j^k, \infty)$.

Let $\mathcal{C}_h(t)^{\lceil v}$ denote the sub-chain of $\mathcal{C}_h(t)$ starting from v, i.e. $\mathcal{C}_h(t)^{\lceil v}$ contains blocks of $\mathcal{C}_h(t)$ that are at height no less than that of v. We show that (6) implies that $\mathcal{C}_h(t)^{\lceil v}$ cannot contain all adversarial blocks. We know from statement (iii) of Theorem 5.1 that b_j is the unique block at its height in $\mathsf{MB}(\tau_j)$. Therefore, all blocks in $\mathcal{C}_h(t)^{\lceil v}$ are mined at or after time τ_j. Since $t < \tau_j^{k+1}$, we have

$$a(b_j^0, b_j^{k+1}) \overset{(a)}{<} k - \mathsf{Unheard}_h\left(b_j^k\right) \overset{(b)}{\le} |\mathcal{C}_h(t)^{\lceil v}|,$$

where (a) is the same as (6) and (b) follows from the fact that h has heard of b_j^m, where $m = k - \mathsf{Unheard}_h\left(b_j^k\right)$. Therefore, h adopts a chain that has length at least $\mathsf{height}_{\mathsf{MB}}\,(v) + m$. We conclude that there exists an honest block $b_k \in \mathcal{C}_h(t)^{\lceil v}$. From statement (iv) in Theorem 5.1, it follows that $\mathcal{C}_h(t)$ contains b_j.

The above argument is true for all $t \in [\tau_j^k, \tau_j^{k+1})$, so it follows that $h \in \mathcal{C}_h(t)$ for all $t \in [\tau_j^k, \tau_j^{k+1})$. Since this is true for all $k \ge k_0$, it follows that $h \in \mathcal{C}_h(t)$ for all $t \ge \tau_j^{k_0}$, as desired. □

B Proofs of Probabilistic Results

B.1 Proof of Lemma 5.3

Lemma 5.3. *Let $0 \le i < j \le j' < k$, and $0 \le \eta \le 1$. The events:*

- $\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}$ *and* $\overrightarrow{\mathsf{B}}_{j',k}^{(\eta)}$ *are independent.*
- $\overrightarrow{\mathsf{B}}_{i,j}^{(\eta)}$ *and* $\overleftarrow{\mathsf{B}}_{j',k}^{(\eta)}$ *are independent.*

Proof. We prove only the first statement, since the second uses a similar argument. For any $m \in \{i, j, j', k\}$, let τ_m denote the mining time of b_m. The LHS of the event

$$\overrightarrow{\mathsf{B}}_{j',k}^{(\eta)}\colon a(b_{j'}, b_k) \ge \eta \cdot \mathsf{FS}_{j'}\,(b_{k-1}) - \mathsf{FU}_{j'}\,(b_k)$$

depends on the number of adversarial arrivals in $[\tau_{j'}, \tau_k]$. Further, the RHS depends on the delays from $b_{j'}$ to all the honest blocks mined in $[\tau_{j'}, \tau_k)$.

In contrast, the LHS of the event

$$\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\colon a(b_i, b_j) \ge \eta \cdot \mathsf{BS}_j\,(b_{i+1}) - \mathsf{BU}_j\,(b_i)$$

depends on the number of adversarial arrivals in $[\tau_i, \tau_j]$. Further, the RHS depends on the delays from all the honest blocks mined in $[\tau_i, \tau_j)$ to b_j.

Since honest and adversarial arrivals are independent Poisson processes, and since the delay associated with any two blocks in $[\tau_i, \tau_j]$ is independent of the delay associated with any two blocks in $[\tau_{j'}, \tau_k]$, and since the two intervals do not overlap, it follows that $\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}$ and $\overrightarrow{\mathsf{B}}_{j',k}^{(\eta)}$ are independent. □

B.2 Proof of Theorem 5.4

Theorem 5.4. *Let $j \ge 1$ and $0 \le \eta \le 1$. Recall that b_j is the j-th honest block and $\mathsf{N}_j^{(\eta)}$ is the event that b_j is an η-Nakamoto block. If $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$, then there exists a positive constant $p_0 > 0$ such that*

$$\mathbb{P}\left(\mathsf{N}_j^{(\eta)}\right) \ge p_0 > 0$$

Proof. Fix j. Recall that the event $\mathsf{N}_j^{(\eta)} = \overleftarrow{\mathsf{E}}_j^{(\eta)} \bigcap \overrightarrow{\mathsf{E}}_j^{(\eta)}$, where $\overleftarrow{\mathsf{E}}_j^{(\eta)} = \bigcap_{i\colon i<j} \left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^{\mathsf{c}}$ and $\overrightarrow{\mathsf{E}}_j^{(\eta)} = \bigcap_{k\colon k>j} \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}$. From Lemma 5.3, it follows that $\overleftarrow{\mathsf{E}}_j^{(\eta)}$ and $\overrightarrow{\mathsf{E}}_j^{(\eta)}$ are independent events. Therefore,

$$\mathbb{P}\left(\mathsf{N}_j^{(\eta)}\right) = \mathbb{P}\left(\overleftarrow{\mathsf{E}}_j^{(\eta)} \bigcap \overrightarrow{\mathsf{E}}_j^{(\eta)}\right) = \mathbb{P}\left(\overleftarrow{\mathsf{E}}_j^{(\eta)}\right)\mathbb{P}\left(\overrightarrow{\mathsf{E}}_j^{(\eta)}\right).$$

Thus, it suffices to show the existence of $p > 0$ such that $\mathbb{P}\left(\overleftarrow{\mathsf{E}}_j^{(\eta)}\right) \geq p$ and $\mathbb{P}\left(\overrightarrow{\mathsf{E}}_j^{(\eta)}\right) \geq p$.

Notation: If an i.i.d. random process is a sequence of random variables with known distribution, say Geometric or Bernoulli with parameter q, we refer to the i-th random variable in the sequence as $\mathsf{Geom}_i(q)$ or $\mathsf{Be}_i(q)$ respectively.

Showing $\mathbb{P}\left(\overrightarrow{\mathsf{E}}_j^{(\eta)}\right) \geq p > 0$: By definition, the event $\overrightarrow{\mathsf{E}}_j^{(\eta)}$ occurs if

$$a(b_j, b_k) - \eta \cdot \mathsf{FS}_j\left(b_{k-1}\right) < -\mathsf{FU}_j\left(b_k\right) \ \forall k > j. \tag{7}$$

Notice that the LHS are RHS are independent random variables. The random variables inn the LHS are:

$$a(b_j, b_k) = \sum_{i=1}^{k-j}\left(\mathsf{Geom}_i\left(1-\beta\right) - 1\right), \tag{8}$$

$$\mathsf{FS}_j\left(b_{k-1}\right) = 1 + \sum_{i=1}^{k-j-1} \mathsf{Be}_i\left(1-d\right), \tag{9}$$

where equation (8) follows from the fact that there are $\mathsf{Geom}\left(1-\beta\right) - 1$ adversarial arrivals between two successive honest arrivals, and the fact that the number of adversarial arrivals in disjoint intervals is independent. Equation (9) follows from Remark 4.4.

Let $(\mathsf{X}_i\colon i \geq 1)$ be an i.i.d. random process, with $\mathsf{X}_i \sim \mathsf{Geom}_i\left(1-\beta\right) - 1 - \eta \cdot \mathsf{Be}_i\left(1-d\right)$. Let $\mathsf{W}_j = \sum_{i=1}^{j} \mathsf{X}_i$ denote the sum of the first j terms of the process $(\mathsf{X}_i\colon i \geq 1)$. Finally, it follows from Remark 4.7 that $(\mathsf{FU}_j\left(b_k\right) : k > j)$ is identical to the i.i.d. random process $(\mathsf{Y}_i\colon i > 0)$, with $\mathsf{Y}_i \sim \mathsf{Geom}_i\left(1-d\right) - 1$. Using these random variables, equation (7) can be equivalently stated as

$$\sum_{i=1}^{k-j} \mathsf{X}_i < -\mathsf{Y}_{k-j} \quad \forall k > j. \tag{10}$$

Let γ be a constant to be determined later, such that $0 < \gamma < \eta$. The inequality in (10) holds if $\mathsf{E_A} \cap \mathsf{E_B}$ occur, where

$$\mathsf{E_A}\colon \sum_{j=1}^{i} \mathsf{X}_j < -\gamma i \quad \forall i \geq 1$$

$$\mathsf{E_B}\colon -\gamma i < -\mathsf{Y}_i \quad \forall i \geq 1.$$

Here, $\mathsf{E_A}$ and $\mathsf{E_B}$ are independent events. Therefore, we have $\mathbb{P}\left(\overrightarrow{\mathsf{E}}_j^{(\eta)}\right) \geq \mathbb{P}\left(\mathsf{E_A} \cap \mathsf{E_B}\right) = \mathbb{P}\left(\mathsf{E_A}\right)\mathbb{P}\left(\mathsf{E_B}\right)$. It suffices to show the existence of $p_1, p_2 > 0$ such that $\mathbb{P}\left(\mathsf{E_A}\right) \geq p_1$ and $\mathbb{P}\left(\mathsf{E_B}\right) \geq p_2$. First, we bound $\mathbb{P}\left(\mathsf{E_B}\right)$.

$$\mathbb{P}\left(\mathsf{E_B}\right) = \mathbb{P}\left(\bigcap_{i \geq 1}\left\{\mathsf{Y}_i \leq \gamma i\right\}\right) = \prod_{i=1}^{\infty}\left(1 - d^{1+\gamma i}\right) =: p_2 > 0,$$

since $(\mathsf{Y}_i\colon i \geq 1)$ is an i.i.d process with $1 + \mathsf{Y}_i \sim \mathsf{Geom}_i\left(1-d\right)$. Notice that $p_2 > 0$ for all $0 \leq d < 1$ and for all $\gamma > 0$.

Next, we bound $\mathbb{P}\left(\mathsf{E_A}\right)$. Fix some $\ell \in \mathbb{N}$, and let c be a positive constant given by $c = \left(\eta - \gamma\right)\ell$. Consider the following two desirable events:

$${}_1\colon \mathsf{W}_\ell \leq -\gamma\ell - c$$

$${}_2\colon \max_{m \geq 0}\left(\mathsf{W}_{\ell+m} - \mathsf{W}_\ell + \gamma m\right) < c$$

It is clear that ${}_1\cap{}_2 \implies \mathsf{E}_\mathsf{A}$, and that ${}_1$ and ${}_2$ are independent events, since Poisson arrivals over disjoint intervals are independent. Therefore, it suffices to find constants $p_{11}, p_{12} > 0$ such that $\mathbb{P}\left({}_1\right) \geq p_{11}$ and $\mathbb{P}\left({}_2\right) \geq p_{12}$.

$$\begin{aligned}
\mathbb{P}\left({}_1\right) &= \mathbb{P}\left(\sum_{i=1}^{\ell} \mathsf{X}_i \leq -\gamma\ell - c\right) && (11)\\
&= \mathbb{P}\left(\sum_{i=1}^{\ell} \left(\mathsf{Geom}_i\left(1-\beta\right) - 1 - \eta \cdot \mathsf{Be}_i\left(1-d\right)\right) \leq -\gamma\ell - c\right) && (12)\\
&= \mathbb{P}\left(\sum_{i=1}^{\ell} \left(\mathsf{Geom}_i\left(1-\beta\right) - \eta \cdot \mathsf{Be}_i\left(1-d\right)\right) \leq \left(1-\gamma\right)\ell - c\right) && (13)\\
&\geq \mathbb{P}\left(\bigcap_{i=1}^{\ell} \left[\left\{\mathsf{Geom}_i\left(1-\beta\right) = 1\right\} \cap \left\{\mathsf{Be}_i\left(1-d\right) = 1\right\}\right]\right) && (14)\\
&= p_{11} > 0, && (15)
\end{aligned}$$

for fixed ℓ, since $\gamma < \eta$ and since c is chosen to be sufficiently small. Here, the inequality in (14) holds because one way to satisfy the inequality in the event in (13) is when both LHS and RHS of the inequality equal $\left(1-\eta\right)\ell$.

It remains to show $p_{12} > 0$. Consider the random process $\left(\mathsf{Z}_i \colon i \geq 1\right)$, where $\mathsf{Z}_i = \mathsf{X}_i + \gamma$. Then, $\left(\sum_{i=1}^{m} \mathsf{Z}_i \colon m \geq 1\right)$ and $\left(\mathsf{W}_{\ell+m} - \mathsf{W}_\ell + \gamma m \colon m \geq 1\right)$ follow the same distribution. The Kingman bound [Kin64] yields,

$$\mathbb{P}\left({}_2\right) = \mathbb{P}\left(\max_{m\geq 0} \sum_{i=1}^{m} \mathsf{Z}_i < c\right) \geq 1 - e^{-\theta^* c} =: p_{12},$$

where

$$\theta^* = \sup\left\{\theta > 0 \colon \mathbb{E}\left[e^{\theta \mathsf{Z}_1} \leq 1\right]\right\}.$$

Since $c > 0$, we see that $p_{12} > 0$ if $\theta^* > 0$. We show that $\theta^* > 0$ if $\frac{\beta}{1-\beta} < \eta\left(1-d\right)$. A simple computation yields

$$\mathbb{E}\left[\mathsf{Z}_1\right] = \frac{\beta}{1-\beta} - \eta \cdot \left(1-d\right) + \gamma,$$

If $\frac{\beta}{1-\beta} < \eta\left(1-d\right)$, then there exists $0 < \gamma < \eta$ such that $\mathbb{E}\left[\mathsf{Z}_1\right] < 0$. Since $\mathbb{E}\left[e^{\theta \mathsf{Z}_1}\right]_{\theta=0} = 1$, and $\frac{d}{d\theta}\mathbb{E}\left[e^{\theta \mathsf{Z}_1}\right]_{\theta=0} = \mathbb{E}\left[\mathsf{Z}_1\right] < 0$, we know that there exists $\theta_1 > 0$ such that $\mathbb{E}\left[e^{\theta \mathsf{Z}_1}\right]_{\theta=\theta_1} < 1$, which then implies $\theta^* > \theta_1 > 0$.

The above argument is summarized as

$$\mathbb{P}\left(\overrightarrow{\mathsf{E}}_j^{(\eta)}\right) \geq \mathbb{P}\left(\mathsf{E}_\mathsf{A}\right)\mathbb{P}\left(\mathsf{E}_\mathsf{B}\right) \geq \mathbb{P}\left({}_1\right) \cdot \mathbb{P}\left({}_2\right) \cdot \mathbb{P}\left(\mathsf{E}_\mathsf{B}\right) \geq p_{11} \cdot p_{12} \cdot p_2 =: p > 0. \tag{16}$$

Showing $\mathbb{P}\left(\overleftarrow{\mathsf{E}}_j^{(\eta)}\right) \geq p > 0$: For any $i, k \geq 0$ such that $j - i = k - j$, it follows from Remarks 4.4, 4.7, 4.9 and 4.12 that $\mathbb{P}\left(\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right) = \mathbb{P}\left(\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right)$. By symmetry, we have

$$\begin{aligned}
\mathbb{P}\left(\overleftarrow{\mathsf{E}}_j^{(\eta)}\right) &= \mathbb{P}\left(\bigcap_{i=0}^{j} \left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^\mathsf{c}\right)\\
&= \mathbb{P}\left(\bigcap_{k=j+1}^{2j+1} \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^\mathsf{c}\right)\\
&\geq \mathbb{P}\left(\bigcap_{k\colon k>j} \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^\mathsf{c}\right)\\
&= \mathbb{P}\left(\overrightarrow{\mathsf{E}}_j^{(\eta)}\right)\\
&\geq p,
\end{aligned}$$

where the constant $p > 0$ is the same as in (16). This concludes the proof. □

B.3 Proof of Lemma 5.5

Lemma 5.5. *Let $0 < \eta \leq 1$ and $i < j < k$. Let $\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}$ and $\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}$ be the catch-up events defined in* (1), (2). *If $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$, there exists a constant $c > 0$ such that*

$$\mathbb{P}\left(\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right) \leq e^{-c(k-j)}$$
$$\mathbb{P}\left(\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right) \leq e^{-c(j-i)}$$

Proof. First, we show the existence of $c > 0$ such that $\mathbb{P}\left(\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right) \leq e^{-c(k-j)}$. Recall that the event $\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}$ is defined as

$$\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}: \quad a(b_j, b_k) \geq \eta \cdot \mathsf{FS}_j\left(b_{k-1}\right) - \mathsf{FU}_j\left(b_k\right).$$

From Remarks 4.4, 4.7, 4.9 and 4.12, these random variables are characterized as:

$$\begin{aligned} a(b_j, b_k) &= \sum_{i=1}^{k-j} \left(\mathsf{Geom}_i\left(1-\beta\right) - 1\right), \\ \mathsf{FS}_j\left(b_{k-1}\right) &= 1 + \sum_{i=1}^{k-j-1} \mathsf{Be}_i\left(1-d\right), \\ \mathsf{FU}_j\left(b_k\right) &= \mathsf{Geom}\left(1-d\right) - 1. \end{aligned}$$

Further, the three random variables $a(b_j, b_k), \mathsf{FS}_j\left(b_{k-1}\right), \mathsf{FU}_j\left(b_k\right)$ are mutually independent. Consider the following desirable events associated with them.

$$\begin{aligned} {}_1 &: a(b_j, b_k) < \frac{\beta}{1-\beta}\left(k-j\right)\left(1+\varepsilon\right) \\ {}_2 &: \mathsf{FS}_j\left(b_{k-1}\right) > \left(1-d\right)\left(k-j\right)\left(1-\delta\right) \\ {}_3 &: \mathsf{FU}_j\left(b_k\right) \leq \left(k-j\right)\left[\left(1-\delta\right)\eta \cdot \left(1-d\right) - \frac{\beta}{1-\beta}\left(1+\varepsilon\right)\right] \end{aligned}$$

Clearly, ${}_1 \cap {}_2 \cap {}_3 \implies \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}$. Therefore, we have

$$\mathbb{P}\left(\left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}\right) \geq \mathbb{P}\left({}_1 \cap {}_2 \cap {}_3\right) = \mathbb{P}\left({}_1\right)\mathbb{P}\left({}_2\right)\mathbb{P}\left({}_3\right). \tag{17}$$

It suffices to find bounds for each term separately in the RHS of (17).

Bounding ${}_1$: Since $\mathbb{E}\left[a(b_j, b_k)\right] = \frac{\beta}{1-\beta}\left(k-j\right)$, it follows from the Chernoff bound that

$$\mathbb{P}\left({}_1^{\mathsf{c}}\right) = \mathbb{P}\left(a(b_j, b_k) \geq \frac{\beta}{1-\beta}\left(k-j\right)\left(1+\varepsilon\right)\right) \leq e^{-c_1(k-j)},$$
$$\text{where } c_1 = \frac{\beta}{1-\beta}\left(1+\varepsilon\right)\log\left(1+\varepsilon\right) + \frac{1}{1-\beta}\left(1+\beta\varepsilon\right)\log\left(1+\beta\varepsilon\right) > 0$$

Bounding ${}_2$: It follows from the Hoeffding bound that

$$\begin{aligned} \mathbb{P}\left({}_2^{\mathsf{c}}\right) = \mathbb{P}&\left(1 + \sum_{i=1}^{k-j-1} \mathsf{Be}_i\left(1-d\right) \leq \left(1-d\right)\left(k-j\right)\left(1-\delta\right)\right) \\ &\leq \mathbb{P}\left(\sum_{i=1}^{k-j} \mathsf{Be}_i\left(1-d\right) \leq \left(1-d\right)\left(k-j\right)\left(1-\delta\right)\right) \\ &\leq e^{-c_2(k-j)}, \end{aligned}$$

where $c_2 = \frac{\delta^2(1-d)}{2} > 0$.

Bounding $_3$: Consider the quantity $c_4 = \left[(1-\delta)\,\eta\cdot(1-d) - \frac{\beta}{1-\beta}\,(1+\varepsilon)\right]$. If $\frac{\beta}{1-\beta} < \eta\,(1-d)$, then there exist $\varepsilon, \delta > 0$ such that $c_4 > 0$. Since $1 + \mathsf{FU}_j\,(b_k)$ follows a geometric distribution, we have

$$\begin{aligned}\mathbb{P}\left({}^{\mathsf{c}}_{3}\right) &= \mathbb{P}\left(\mathsf{Geom}\,(1-d) - 1 > c_4\,(k-j)\right)\\ &= \mathbb{P}\left(\mathsf{Geom}\,(1-d) > 1 + c_4\,(k-j)\right)\\ &= 1 - d^{c_4(k-j)}\\ &\geq 1 - e^{c_3(k-j)},\end{aligned}$$

for some $c_3 > 0$.

Combining these facts together, we revisit (17). We have

$$\begin{aligned}\mathbb{P}\left(\left[\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right]^{\mathsf{c}}\right) &\geq \mathbb{P}\left(_1\right)\mathbb{P}\left(_2\right)\mathbb{P}\left(_3\right) && (18)\\ &= \left(1-\mathbb{P}\left({}^{\mathsf{c}}_{1}\right)\right)\left(1-\mathbb{P}\left({}^{\mathsf{c}}_{2}\right)\right)\left(1-\mathbb{P}\left({}^{\mathsf{c}}_{3}\right)\right) && (19)\\ &\geq \left(1-e^{-c_1(k-j)}\right)\left(1-e^{-c_2(k-j)}\right)\left(1-e^{-c_3(k-j)}\right) && (20)\\ &\geq \left(1-e^{-c_0(k-j)}\right)^3 && (21)\\ &\geq 1-e^{-c(k-j)}, && (22)\end{aligned}$$

where $c_0 = \max\{c_1, c_2, c_3\} > 0$, and subsequently $c > 0$. We conclude the existence of $c > 0$ for which $\mathbb{P}\left(\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right) \leq e^{-c(k-j)}$.

It remains to show that $\mathbb{P}\left(\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right) \leq e^{-c(j-i)}$. The proof is very similar, so the details are omitted. Recall that

$$\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}:\quad a(b_i, b_j) \geq \eta\cdot\mathsf{BS}_j\,(b_{i+1}) - \mathsf{BU}_j\,(b_i)\,,$$

where the random variables involved may be written as

$$\begin{aligned}a(b_i, b_j) &= \sum_{k=1}^{j-i}\left(\mathsf{Geom}_k\,(1-\beta) - 1\right),\\ \mathsf{BS}_j\,(b_{i+1}) &= 1 + \sum_{k=1}^{j-i-1}\mathsf{Be}_k\,(1-d)\,,\\ \mathsf{BU}_j\,(b_i) &= \mathsf{Geom}\,(1-d) - 1,\end{aligned}$$

Thus, if $k - j = j - i$, we see that

- $a(b_i, b_j)$ and $a(b_j, b_k)$ follow the same distribution.
- $\mathsf{BS}_j\,(b_{i+1})$ and $\mathsf{FS}_j\,(b_{k-1})$ follow the same distribution.
- $\mathsf{BU}_j\,(b_i)$ and $\mathsf{FU}_j\,(b_k)$ follow the same distribution.

Therefore, the same concentration inequalities apply, and we conclude that for the same constant c in (22), we have $\mathbb{P}\left(\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right) \leq e^{-c(j-i)}$. □

B.4 Proof of Theorem 5.6

Theorem 5.6. *Let $0 < \eta \leq 1$. Let β be the fraction of computational power in the system that is adversarial and d be the probability of message loss. Let $B^{(\eta)}_{s,s+t}$ be the event that there are no η-Nakamoto blocks in $[s, s+t]$. If $\frac{\beta}{1-\beta} < \eta\cdot(1-d)$, then there exists a constant $c_0 > 0$ such that for any $s, t \geq 0$,*

$$\mathbb{P}\left(B^{(\eta)}_{s,s+t}\right) \leq e^{-c_0\sqrt{t}}$$

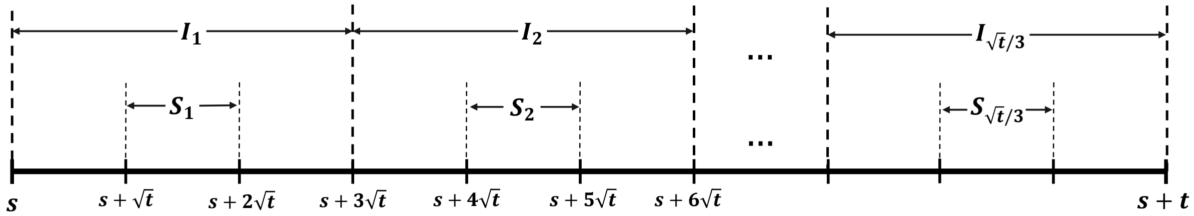

Figure 6: Partitioning $[s, s+t]$ into sub-intervals

Proof. For any $i \geq 0$, let τ_i denote the mining time of block b_i. Partition the interval $[s, s+t]$ into $\sqrt{t}$ intervals of length $\sqrt{t}$ each. Group these sub-intervals into threes, so that there are $\sqrt{t}/3$ groups of sub-intervals, namely I_1, I_2, $\cdots$, $I_{\sqrt{t}/3}$. Thus, $I_\ell = \left[s + 3\left(\ell - 1\right)\sqrt{t}, s + 3\ell\sqrt{t}\right]$. Further, let S_ℓ represent the middle sub-interval of I_ℓ, so that $S_\ell = [s + (3\ell - 2)\sqrt{t}, s + (3\ell - 1)\sqrt{t}]$, as shown in Figure 6.

Consider the following desirable events.

$$_1\colon \bigcap_{j\colon \tau_j \in [s+\sqrt{t}, s+t-\sqrt{t}]} \left[\left(\bigcap_{\substack{i<j \\ \tau_i < \tau_j - \sqrt{t}}} \left[\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right]^{\mathsf{c}}\right) \bigcap \left(\bigcap_{\substack{k>j \\ \tau_k > \tau_j + \sqrt{t}}} \left[\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right]^{\mathsf{c}}\right)\right]$$

$$_2\colon \bigcup_{\ell=1}^{\sqrt{t}/3} \mathsf{H}_\ell, \text{ where } \mathsf{H}_\ell\colon \bigcup_{j\colon \tau_j \in S_\ell} \left[\left(\bigcap_{\substack{i<j \\ \tau_i \geq \tau_j - \sqrt{t}}} \left[\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right]^{\mathsf{c}}\right) \bigcap \left(\bigcap_{\substack{k>j \\ \tau_k \leq \tau_j + \sqrt{t}}} \left[\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right]^{\mathsf{c}}\right)\right]$$

First, observe that $_1 \cap _2 \implies \left[B^{(\eta)}_{s,s+t}\right]^{\mathsf{c}}$. This is because $_1$ ensures that no catch-up events (neither $\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}$, nor $\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}$) occur when b_i and b_k are separated in time by more than $\sqrt{t}$ from b_j, for any b_j mined in $[s+\sqrt{t}, s+t-\sqrt{t}]$. This means that the existence of a block b_j in this interval for which no catch-up event occurs whenever i and k are within $\sqrt{t}$ of τ_j is sufficient to ensure that b_j is a Nakamoto block. This is exactly the event $_2$. Since $\cup_\ell S_\ell \subset \left[s+\sqrt{t}, s+t-\sqrt{t}\right]$, it follows that $_1 \cap _2 \implies \left[B^{(\eta)}_{s,s+t}\right]^{\mathsf{c}}$. Thus,

$$\mathbb{P}\left(B^{(\eta)}_{s,s+t}\right) \leq \mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right) + \mathbb{P}\left(\mathsf{G}_2^{\mathsf{c}}\right).$$

Next, we bound the probability of each term in the RHS separately.

Bounding $\mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right)$: Fix $\delta > 0$. Consider the following events:

$$\begin{aligned} &\mathsf{D}_1\colon \left\{\#\left\{i\colon \tau_i \in [s, s+t]\right\} > 2\lambda_h t\right\}, \\ &\mathsf{D}_2\colon \left\{\exists\ \tau_i, \tau_k \in [s, s+t]\colon\ (k-i) < (1-\delta)\lambda_h\sqrt{t}, \tau_k - \tau_i > \sqrt{t}\right\}. \end{aligned}$$

By the tail bound for Poisson random variables, we know that $\mathbb{P}(\mathsf{D}_1) \leq e^{-c_0 t}$ for some $c_0 > 0$. We now show that $\mathbb{P}(\mathsf{D}_2) \leq e^{-c_1\sqrt{t}}$ for some $c_1 > 0$. Let $T_{i,k} := \tau_k - \tau_i$ be the random variable denoting the time between the i-th and k-th mining. Let $M(t) = (1-\delta)\lambda_h\sqrt{t}$. Notice that

$$\bigcap_{i \in [s, s+t]} \left\{T_{i,i+M(t)} < \sqrt{t}\right\} \implies \mathsf{D}_2^{\mathsf{c}}.$$

For any honest arrival time τ_i, we have

$$\mathbb{E}\left[T_{i,i+M(t)}\right] = \mathbb{E}\left[\sum_{j=i}^{i+M(t)} (\tau_j - \tau_{j-1})\right] = \frac{M(t)}{\lambda_h} = (1-\delta)\sqrt{t}.$$

Applying the Chernoff bound, we see that there exists $c > 0$ such that

$$\mathbb{P}\left(T_{i,i+M(t)} > \sqrt{t}\right) \leq \mathbb{P}\left(T_{i,i+M(t)} > (1-\delta)\sqrt{t} + \delta\sqrt{t}\right) \leq e^{-c\sqrt{t}}.$$

Thus, we get

$$\begin{aligned}
\mathbb{P}\left(\mathsf{D}_2\right) &= \mathbb{P}\left(\bigcup_{i:\ \tau_i\in[s,s+t]} \left\{T_{i,i+M(t)} > \sqrt{t}\right\}\right) \\
&\le \mathbb{P}\left(\left(\bigcup_{i:\ \tau_i\in[s,s+t]} \left\{T_{i,i+M(t)} > \sqrt{t}\right\}\right) \cap \mathsf{D}_1^{\mathsf{c}}\right) + \mathbb{P}\left(\mathsf{D}_1\right) \\
&\le \left(\sum_{i=1}^{2\lambda_h t} \mathbb{P}\left(T_{i,i+M(t)} > \sqrt{t}\right)\right) + e^{-c_0 t} \\
&\le e^{-c_1\sqrt{t}},
\end{aligned}$$

for some $c_1 > 0$. Thus, $\mathbb{P}\left(\mathsf{D}_2\right) \le e^{-c_1\sqrt{t}}$. We may therefore bound $\mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right)$ as

$$\begin{aligned}
\mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right) &\le \mathbb{P}\left(\mathsf{D}_1 \cup \mathsf{D}_2\right) + \mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}} \cap \mathsf{D}_1^{\mathsf{c}} \cap \mathsf{D}_2^{\mathsf{c}}\right) \\
&\le \mathbb{P}\left(\mathsf{D}_1\right) + \mathbb{P}\left(\mathsf{D}_2\right) + \mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}} \cap \mathsf{D}_1^{\mathsf{c}} \cap \mathsf{D}_2^{\mathsf{c}}\right) \\
&\le e^{-c_0 t} + e^{-c_1\sqrt{t}} + \sum_{j=1}^{2\lambda_h t}\left[\left(\sum_{i=0}^{j-M(t)} \mathbb{P}\left(\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right)\right) + \left(\sum_{k=j+M(t)}^{\infty} \mathbb{P}\left(\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right)\right)\right] \\
&\le e^{-c_0 t} + e^{-c_1\sqrt{t}} + \sum_{j=1}^{2\lambda_h t}\left[\left(\sum_{i=0}^{j-M(t)} e^{-c(j-i)}\right) + \left(\sum_{k=j+M(t)}^{\infty} e^{-c(k-j)}\right)\right] \\
&\le e^{-c_0 t} + e^{-c_1\sqrt{t}} + \sum_{j=1}^{2\lambda_h t}\left(2\sum_{m=M(t)}^{\infty} e^{-cm}\right) \\
&= e^{-c_0 t} + e^{-c_1\sqrt{t}} + \frac{4\lambda_h t}{1-e^{-c}} e^{-cM(t)} \\
&= e^{-c_0 t} + e^{-c_1\sqrt{t}} + \frac{4\lambda_h t}{1-e^{-c}} e^{-c(1-\delta)\lambda_h\sqrt{t}} \\
&\le e^{-c_3\sqrt{t}},
\end{aligned}$$

for some $c_3 > 0$.

Bounding $\mathbb{P}\left(\mathsf{G}_2^{\mathsf{c}}\right)$: We have $\mathsf{G}_2^{\mathsf{c}} = \bigcap_{\ell=1}^{\sqrt{t}/3} \mathsf{H}_\ell^{\mathsf{c}}$. Notice that $\mathsf{H}_\ell^{\mathsf{c}}$ are mutually independent for distinct ℓ by Lemma 5.3. Recall that

$$\mathsf{H}_\ell = \bigcup_{j:\ \tau_j\in S_\ell} \mathsf{R}_j^\ell,\ \text{where}\ \mathsf{R}_j^\ell\colon \left(\bigcap_{\substack{i<j \\ \tau_i \ge \tau_j - \sqrt{t}}} \left[\overleftarrow{\mathsf{B}}^{(\eta)}_{i,j}\right]^{\mathsf{c}}\right) \bigcap \left(\bigcap_{\substack{k>j \\ \tau_k \le \tau_j + \sqrt{t}}} \left[\overrightarrow{\mathsf{B}}^{(\eta)}_{j,k}\right]^{\mathsf{c}}\right).$$

Let M_ℓ be the number of honest blocks mined in S_ℓ, and N_ℓ be the number of η-Nakamoto blocks mined in S_ℓ. Since H_ℓ is contained in the event $M_\ell \ge 1$, we have for each $\ell \in \left\{1, 2, \cdots, \sqrt{t}/3\right\}$:

$$\begin{aligned}
\mathbb{P}\left(\mathsf{H}_\ell\right) &= \mathbb{P}\left(\bigcup_{j:\ \tau_j\in S_\ell} \mathsf{R}_j^\ell\right) \\
&= \mathbb{P}\left(\left(\bigcup_{j:\ \tau_j\in S_\ell} \mathsf{R}_j^\ell\right) \bigcap \left\{M_\ell \ge 1\right\}\right) \\
&\ge \mathbb{P}\left(\bigcup_{j:\tau_j\in S_\ell} \mathsf{N}_j^{(\eta)}\right) \\
&\ge \frac{p_0^2}{2},
\end{aligned} \tag{23}$$

where p_0 is the lower bound on the probability that b_j is an η-Nakamoto block, obtained in Theorem 5.4. The inequality in (23) deserves some elaboration: Since M_ℓ is a non-negative integer valued random variable, we have from the second moment method that

$$\mathbb{P}\left(N_\ell \geq 1\right) = \mathbb{P}\left(N_\ell > 0\right) \geq \frac{\left(\mathbb{E}\left[N_\ell\right]\right)^2}{\mathbb{E}\left[N_\ell^2\right]} \geq \frac{\left(p_0 \lambda_h \sqrt{t}\right)^2 t}{\lambda_h \sqrt{t} + \left(\lambda_h \sqrt{t}\right)^2} \geq \frac{\lambda_h^2 p_0^2 t}{2\lambda_h^2 t} \geq \frac{p_0^2}{2},$$

for sufficiently large t. Here, we used the fact that $\mathbb{E}\left[N_\ell\right] = p_0 \lambda_h \sqrt{t}$, and $\mathbb{E}\left[N_\ell^2\right] \leq \mathbb{E}\left[M_\ell^2\right] = \lambda_h \sqrt{t} + \left(\lambda_h \sqrt{t}\right)^2$.

Thus, $\mathbb{P}\left(\mathsf{H}_\ell^{\mathsf{c}}\right) \leq 1 - \frac{p_0^2}{2} < 1$, which yields

$$\mathbb{P}\left(\mathsf{G}_2^{\mathsf{c}}\right) = \mathbb{P}\left(\bigcap_{\ell=1}^{\sqrt{t}/3} \mathsf{H}_\ell^{\mathsf{c}}\right) = \prod_{\ell=1}^{\sqrt{t}/3} \mathbb{P}\left(\mathsf{H}_\ell^{\mathsf{c}}\right) \leq \left(1 - \frac{p_0^2}{2}\right)^{\sqrt{t}/3} \leq e^{-c_4 \sqrt{t}},$$

for some $c_4 > 0$, since the H_ℓ's are mutually independent events. Therefore, we conclude that

$$\mathbb{P}\left(B_{s,s+t}^{(\eta)}\right) \leq \mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right) + \mathbb{P}\left(\mathsf{G}_2^{\mathsf{c}}\right) \leq e^{-c_3\sqrt{t}} + e^{-c_4\sqrt{t}} \leq e^{-c_0\sqrt{t}},$$

for some $c_0 > 0$, as desired. □

B.5 Proof of Theorem 5.7

Theorem 5.7. *Let $0 < \eta \leq 1$. Let β be the fraction of computational power in the system that is adversarial and d be the probability of message loss. Let $B_{s,s+t}^{(\eta)}$ be the event that there are no η-Nakamoto blocks in $[s, s+t]$. If $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$, then for every $\varepsilon > 0$, there exist positive constants A, a such that for any $s, t > 0$,*

$$\mathbb{P}\left(B_{s,s+t}^{(\eta)}\right) \leq A \exp\left(-a t^{1-\varepsilon}\right).$$

Proof. Fix $m > 1$. Consider the following statement for m:

$$\mathbb{S}[m] \; : \forall \; \theta \geq m, \; \exists \; a_\theta > 0, A_\theta > 0 \text{ such that } \mathbb{P}\left(B_{s,s+t}^{(\eta)}\right) \leq A_\theta \exp\left(-a_\theta t^{1/\theta}\right)$$

In Theorem 5.6, we proved that $\mathbb{S}[2]$ is true. Next, we show the following:

$$\mathbb{S}[m] \implies \mathbb{S}\left[\frac{2m-1}{m}\right].$$

Assume $\mathbb{S}[m]$ is true. For any $i \geq 0$, let τ_i denote the mining time of block b_i. Partition the interval $[s, s+t]$ into $t^{\frac{m-1}{2m-1}}$ intervals of length $t^{\frac{m}{2m-1}}$ each. Group these sub-intervals into threes, so that there are $\frac{t^{\frac{m-1}{2m-1}}}{3}$ groups of sub-intervals, namely $I_1, I_2, \cdots, I_{t^{\frac{m-1}{2m-1}}/3}$. Thus,

$$\begin{aligned} I_\ell \colon &= \left[s + 3\left(\ell - 1\right) t^{\frac{m}{2m-1}}, s + 3\ell t^{\frac{m}{2m-1}}\right], \\ S_\ell \colon &= \left[s + \left(3\ell - 2\right) t^{\frac{m}{2m-1}}, s + \left(3\ell - 1\right) t^{\frac{m}{2m-1}}\right]. \end{aligned}$$

Consider the following desirable events.

$${}_1\colon \quad \bigcap_{j\,:\,\tau_j \in \left[s+t^{\frac{m}{2m-1}}, s+t-t^{\frac{m}{2m-1}}\right]} \left[\left(\bigcap_{\substack{i<j \\ \tau_i < \tau_j - t^{\frac{m}{2m-1}}}} \left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^{\mathsf{c}}\right) \bigcap \left(\bigcap_{\substack{k>j \\ \tau_k > \tau_j + t^{\frac{m}{2m-1}}}} \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}\right)\right]$$

$${}_2\colon \quad \bigcup_{\ell=1}^{t^{\frac{m-1}{2m-1}}/3} \mathsf{H}_\ell, \text{ where } \mathsf{H}_\ell\colon \bigcup_{j\,:\,\tau_j \in S_\ell} \left[\left(\bigcap_{\substack{i<j \\ \tau_i \geq \tau_j - t^{\frac{m}{2m-1}}}} \left[\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right]^{\mathsf{c}}\right) \bigcap \left(\bigcap_{\substack{k>j \\ \tau_k \leq \tau_j + t^{\frac{m}{2m-1}}}} \left[\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right]^{\mathsf{c}}\right)\right]$$

Observe that ${}_1 \cap {}_2 \implies \left[B_{s,s+t}^{(\eta)}\right]^{\mathsf{c}}$.

Bounding $\mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right)$: We show that $\mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right) \leq e^{-c_3 t^{\frac{m}{2m-1}}}$ for some $c_3 > 0$. This is done following the steps in Theorem 5.6. Fix $\delta > 0$ and consider the following events:

$$\begin{aligned}
&\mathsf{D}_1 \colon \left\{\# \left\{i \colon \tau_i \in [s, s+t]\right\} > 2\lambda_h t\right\}, \\
&\mathsf{D}_2 \colon \left\{\exists\, \tau_i, \tau_k \in [s, s+t] \colon (k-i) < (1-\delta)\, \lambda_h t^{\frac{m}{2m-1}}, \tau_k - \tau_i > t^{\frac{m}{2m-1}}\right\}.
\end{aligned}$$

By the tail bound for Poisson random variables, we know that $\mathbb{P}\left(\mathsf{D}_1\right) \leq e^{-c_0 t}$ for some $c_0 > 0$. We now show that $\mathbb{P}\left(\mathsf{D}_2\right) \leq e^{-c_1 t^{\frac{m}{2m-1}}}$ for some $c_1 > 0$. Let $T_{i,k} := \tau_k - \tau_i$ be the random variable denoting the time between the i-th and k-th mining. Let $M(t) = (1-\delta)\, \lambda_h t^{\frac{m}{2m-1}}$. Notice that

$$\bigcap_{i \in [s, s+t]} \left\{T_{i, i+M(t)} < \sqrt{t}\right\} \implies \mathsf{D}_2^{\mathsf{c}}.$$

For any honest arrival time τ_i, we have

$$\mathbb{E}\left[T_{i,i+M(t)}\right] = \mathbb{E}\left[\sum_{j=i}^{i+M(t)} \left(\tau_j - \tau_{j-1}\right)\right] = \frac{M(t)}{\lambda_h} = (1-\delta)\, t^{\frac{m}{2m-1}}.$$

Applying the Chernoff bound, we see that there exists $c > 0$ such that

$$\mathbb{P}\left(\mathsf{T}_{i,i+M(t)} > t^{\frac{m}{2m-1}}\right) \leq \mathbb{P}\left(T_{i,i+M(t)} > (1-\delta)\, t^{\frac{m}{2m-1}} + \delta t^{\frac{m}{2m-1}}\right) \leq e^{-ct^{\frac{m}{2m-1}}}.$$

Thus, we get

$$\begin{aligned}
\mathbb{P}\left(\mathsf{D}_2\right) &= \mathbb{P}\left(\bigcup_{i \colon \tau_i \in [s, s+t]} \left\{T_{i,i+M(t)} > t^{\frac{m}{2m-1}}\right\}\right) \\
&\leq \mathbb{P}\left(\left(\bigcup_{i \colon \tau_i \in [s, s+t]} \left\{T_{i,i+M(t)} > t^{\frac{m}{2m-1}}\right\}\right) \cap \mathsf{D}_1^{\mathsf{c}}\right) + \mathbb{P}\left(\mathsf{D}_1\right) \\
&\leq \left(\sum_{i=1}^{2\lambda_h t} \mathbb{P}\left(T_{i,i+M(t)} > t^{\frac{m}{2m-1}}\right)\right) + e^{-c_0 t} \\
&\leq e^{-c_1 t^{\frac{m}{2m-1}}},
\end{aligned}$$

for some $c_1 > 0$. Thus, $\mathbb{P}\left(\mathsf{D}_2\right) \leq e^{-c_1 t^{\frac{m}{2m-1}}}$. We may therefore bound $\mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right)$ as

$$\begin{aligned}
\mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}}\right) &\leq \mathbb{P}\left(\mathsf{D}_1 \cup \mathsf{D}_2\right) + \mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}} \cap \mathsf{D}_1^{\mathsf{c}} \cap \mathsf{D}_2^{\mathsf{c}}\right) \\
&\leq \mathbb{P}\left(\mathsf{D}_1\right) + \mathbb{P}\left(\mathsf{D}_2\right) + \mathbb{P}\left(\mathsf{G}_1^{\mathsf{c}} \cap \mathsf{D}_1^{\mathsf{c}} \cap \mathsf{D}_2^{\mathsf{c}}\right) \\
&\leq e^{-c_0 t} + e^{-c_1 t^{\frac{m}{2m-1}}} + \sum_{j=1}^{2\lambda_h t} \left[\left(\sum_{i=0}^{j-M(t)} \mathbb{P}\left(\overleftarrow{\mathsf{B}}_{i,j}^{(\eta)}\right)\right) + \left(\sum_{k=j+M(t)}^{\infty} \mathbb{P}\left(\overrightarrow{\mathsf{B}}_{j,k}^{(\eta)}\right)\right)\right] \\
&\leq e^{-c_0 t} + e^{-c_1 t^{\frac{m}{2m-1}}} + \sum_{j=1}^{2\lambda_h t} \left[\left(\sum_{i=0}^{j-M(t)} e^{-c(j-i)}\right) + \left(\sum_{k=j+M(t)}^{\infty} e^{-c(k-j)}\right)\right] \\
&\leq e^{-c_0 t} + e^{-c_1 t^{\frac{m}{2m-1}}} + \sum_{j=1}^{2\lambda_h t} \left(2 \sum_{m=M(t)}^{\infty} e^{-cm}\right) \\
&= e^{-c_0 t} + e^{-c_1 t^{\frac{m}{2m-1}}} + \frac{4\lambda_h t}{1 - e^{-c}} e^{-cM(t)} \\
&= e^{-c_0 t} + e^{-c_1 t^{\frac{m}{2m-1}}} + \frac{4\lambda_h t}{1 - e^{-c}} e^{-c(1-\delta)\lambda_h t^{\frac{m}{2m-1}}} \\
&\leq e^{-c_3 t^{\frac{m}{2m-1}}},
\end{aligned}$$

for some $c_3 > 0$.

Bounding $\mathbb{P}(\mathsf{G}_2^{\mathsf{c}})$: We have $\mathsf{G}_2^{\mathsf{c}} = \bigcap_{\ell=1}^{t^{\frac{m-1}{2m-1}}/3} \mathsf{H}_\ell^{\mathsf{c}}$. Notice that $\mathsf{H}_\ell^{\mathsf{c}}$ are mutually independent for distinct ℓ by Lemma 5.3. Since $\mathbb{S}[m]$ is true, we have

$$\mathbb{P}(\mathsf{H}_\ell^{\mathsf{c}}) \leq \mathbb{P}\left(B^{(\eta)}_{s+(3\ell-2)t^{\frac{m}{2m-1}},s+(3\ell-1)t^{\frac{m}{2m-1}}}\right) \leq A_m \exp\left(-a_m\left(t^{\frac{m}{2m-1}}\right)^{\frac{1}{m}}\right) \leq A_m \exp\left(-a_m t^{1/(2m-1)}\right)$$

Therefore, it follows that

$$\mathbb{P}(\mathsf{G}_2^{\mathsf{c}}) = \mathbb{P}\left(\bigcap_{\ell=1}^{t^{\frac{m-1}{2m-1}}/3} \mathsf{H}_\ell^{\mathsf{c}}\right) = \prod_{\ell=1}^{t^{\frac{m-1}{2m-1}}/3} \mathbb{P}(\mathsf{H}_\ell^{\mathsf{c}}) \leq \left(A_m \exp\left(-a_m t^{1/(2m-1)}\right)\right)^{t^{\frac{m-1}{2m-1}}/3} \leq e^{-c_4 t^{\frac{m}{2m-1}}},$$

for some $c_4 > 0$, since the H_ℓ's are mutually independent events. Therefore, we have

$$\mathbb{P}\left(B^{(\eta)}_{s,s+t}\right) \leq \mathbb{P}(\mathsf{G}_1^{\mathsf{c}}) + \mathbb{P}(\mathsf{G}_2^{\mathsf{c}}) \leq e^{-c_3 t^{\frac{m}{2m-1}}} + e^{-c_4 t^{\frac{m}{2m-1}}} \leq A_{m'} e^{-a_{m'} t^{1/m'}},$$

where $m' = \frac{m}{2m-1}$, and $a_{m'}, A_{m'} > 0$. In turn, this implies that $\mathbb{S}\left[\frac{2m-1}{m}\right]$ is true.

Finally, consider the recursion given by $m_{k+1} = \frac{2m_k-1}{m_k}$, and the initial condition $m_1 = 2$. We have proved that $\mathbb{S}[m_1]$ is true and that $\mathbb{S}[m_{k+1}]$ is true whenever $\mathbb{S}[m_k]$ is true. By induction, it follows that $\mathbb{S}[m_k]$ is true for all $k \in \mathbb{N}$. Since $m_k = \frac{k+1}{k}$ and $\lim_{k\to\infty} m_k = 1$, we conclude that for every $\varepsilon > 0$, there exist positive constants A, a such that $\mathbb{P}\left(B^{(\eta)}_{s,s+t}\right) \leq A\exp\left(-at^{1-\varepsilon}\right)$. This concludes the proof. □

B.6 Proof of Theorem 5.8

Theorem 5.8. *Let $0 < \eta < 1$. Suppose the fraction β of computational power in the system that is adversarial, and the probability d of message loss satisfies $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$. Given $s \geq 0$, let b_J be the first η-Nakamoto block mined after time s. There exist constants C, $c > 0$ such that for any honest user h and for all $k' \geq 1$,*

$$\mathbb{P}\left((h, J, \eta, k')\text{-user-unheard-criterion fails}\right) \leq Ce^{-ck'}$$

Proof. By the union bound, we have

$$\mathbb{P}\left(\bigcup_{k=k'}^{\infty}\left\{\mathsf{Unheard}_h\left(b_J^k\right) > \left(\frac{1-\eta}{2}\right)k\right\}\right) \leq \sum_{k=k'}^{\infty} \mathbb{P}\left(\mathsf{Unheard}_h\left(b_J^k\right) > \left(\frac{1-\eta}{2}\right)k\right). \tag{24}$$

From remark 4.17, it follows that $\mathsf{Unheard}_h\left(b_J^k\right) \leq \mathsf{Geom}(1-d)$. Thus, for all $k \geq 1$, we have

$$\begin{aligned}\mathbb{P}\left(\mathsf{Unheard}_h\left(b_J^k\right) > \left(\frac{1-\eta}{2}\right)k\right) &\leq \mathbb{P}\left(\mathsf{Geom}(1-d) > \left(\frac{1-\eta}{2}\right)k\right)\\ &\leq C_0 \cdot e^{-ck},\end{aligned}$$

for some positive constants C_0 and c. Combining this with (24) yields

$$\mathbb{P}\left(\bigcup_{k=k'}^{\infty}\left\{\mathsf{Unheard}_h\left(b_J^k\right) > \left(\frac{1-\eta}{2}\right)k\right\}\right) \leq \sum_{k=k'}^{\infty} C_0 \cdot e^{-ck} = C \cdot e^{-ck'},$$

where $C = \frac{C_0}{1-e^{-c}}$ is a positive constant. This concludes the proof. □

B.7 Proof of Theorem 3.2

Before we prove our main result, we recall a useful lemma about Poisson random variables.

Lemma B.1. *Let X be a Poisson random variable with mean μ. Then*

(i) $\mathbb{P}(X \geq 2\mu) \leq e^{-\frac{4}{3}\mu}$.

(ii) $\mathbb{P}\left(X \leq \frac{1}{2}\mu\right) \leq e^{-\frac{1}{8}\mu}$.

Proof. The proof follows from Theorem 4.5 in [MU17]. □

We are now ready to state and prove our main result.

Theorem 3.2. *Let β be the fraction of computational power in the system that is adversarial, and d be the probability of message loss. If $\frac{\beta}{1-\beta} < (1-d)$, then for every $\varepsilon > 0$, there exist positive constants a and b such that for all $\tau \geq 0$ and for any honest transaction* tx *and any finite set of honest users $\mathcal{H}$:*

$$\mathbb{P}\left(\mathsf{tx} \text{ violates } (\tau, \mathcal{H})\text{-security}\right) \leq \exp\left(-a\tau^{1-\varepsilon}\right) + |\mathcal{H}| \exp\left(-b\tau\right).$$

Proof. Let η be such that $\frac{\beta}{1-\beta} < \eta \cdot (1-d)$. Let $k_0 = \left\lceil \frac{2\eta}{1-\eta} \right\rceil$, and fix an honest subset of users $\mathcal{H}$. The idea of the proof is as follows: if an η-Nakamoto block b_J is mined in the time interval $(s, s+t_1)$, and k_0 number of J-FS blocks are mined before time $s + t_1 + t_2$, and if all users $h \in \mathcal{H}$ satisfy the (h, J, η, k_0)-user-unheard-criterion, then Theorems 5.1 and 5.2 together imply that the η-Nakamoto block b_J is included in $\mathcal{C}_h(t)$ for all $t \geq s + t_1 + t_2$. Since tx must be included in either b_J or its ancestors, tx satisfies $(t_1 + t_2, \mathcal{H})$-security.

Let b_J be the first η-Nakamoto block mined after time s. Let $T_J = \tau_J - s$ denote the time between s and the mining time of the first η-Nakamoto block. For the r-th honest block b_r and any time t, let $N_r(t)$ denote the number of r-FS blocks mined until time t. Consider the following events:

$$\begin{aligned}
&\mathsf{E}_1 \colon T_J > t_1 \\
&\mathsf{E}_2^h \colon \text{The } (h, J, \eta, k_0)\text{-user-unheard-criterion is violated} \\
&\mathsf{E}_3 \colon \bigcup_{r : \tau_r \in [s, s+t_1]} \left\{ N_r\left(s + t_1 + t_2\right) < k_0 \right\}.
\end{aligned}$$

By the first paragraph of the proof, the union bound gives

$$\mathbb{P}\left(\mathsf{tx} \text{ violates } (t_1 + t_2, \mathcal{H})\text{-security}\right) \leq \mathbb{P}\left(\mathsf{E}_1\right) + \sum_{h \in \mathcal{H}} \mathbb{P}\left(\mathsf{E}_2^h\right) + \mathbb{P}\left(\mathsf{E}_3\right). \tag{25}$$

From Theorem 5.7, we have that for any $\varepsilon > 0$, there exist positive constants A', a' such that $\mathbb{P}\left(\mathsf{E}_1\right) \leq A'\left(\exp -a' t_1^{1-\varepsilon}\right)$.

From Theorem 5.8, we have that there exist positive constants C, c' such that $\mathbb{P}\left(\mathsf{E}_2^h\right) \leq C \exp\left(-c' k_0\right)$ for all $h \in \mathcal{H}$.

It remains to bound $\mathbb{P}\left(\mathsf{E}_3\right)$. Let λ_h denote the aggregate mining rate of the honest users. Let $t_2 > \frac{2k_0}{(1-d)\lambda_h}$. Let M be the number of honest miners in $[s, s+t_1]$, so that M has the Poisson distribution with mean $\lambda_h\left(1-d\right) t_1$. Therefore, we have from Lemma B.1 that

$$\mathbb{P}\left(M > 2\lambda_h t_1\right) \leq \exp\left(\frac{-4\lambda_h t_1}{3}\right).$$

For $r \geq 1$, consider the r-th honest miner b_r after time s and consider the r-FS sequence. Let U_r denote the number of r-FS blocks mined in $[\tau_r, \tau_r + t_2]$. Then, U_r has the Poisson probability distribution with mean $\lambda_h\left(1-d\right) t_2$. Applying Lemma B.1, we get

$$\mathbb{P}\left(U_r \leq k_0\right) \leq \mathbb{P}\left(U_r \leq \frac{1}{2}\lambda_h\left(1-d\right) t_2\right) \leq \exp\left(\frac{-\lambda_h\left(1-d\right) t_2}{8}\right).$$

Therefore, we have

$$\begin{aligned}
\mathbb{P}\left(\mathsf{E}_3\right) &= \mathbb{P}\left(\mathsf{E}_3 \cap \left\{M > 2\lambda_h t_1\right\}\right) + \mathbb{P}\left(\mathsf{E}_3 \cap \left\{M < 2\lambda_h t_1\right\}\right) \\
&\leq \mathbb{P}\left(N > 2\lambda_h t_1\right) + \sum_{r=1}^{2\lambda_h t_1} \mathbb{P}\left(U_r \leq k_0\right) \\
&\leq \exp\left(\frac{-4\lambda_h t_1}{3}\right) + 2\lambda_h t_1 \cdot \exp\left(\frac{-\lambda_h\left(1-d\right) t_2}{8}\right).
\end{aligned}$$

Therefore, combining (25) and the bounds on $\mathbb{P}(\mathsf{E}_1)$, $\mathbb{P}\left(\mathsf{E}_2^h\right)$, and $\mathbb{P}(\mathsf{E}_3)$ yields

$$\begin{aligned}
\mathbb{P}(\mathsf{tx}\text{ violates }(t_1+t_2,\mathcal{H})\text{-security}) &\le \mathbb{P}(\mathsf{E}_1) + \sum_{h\in\mathcal{H}} \mathbb{P}\left(\mathsf{E}_2^h\right) + \mathbb{P}(\mathsf{E}_3) \\
&\le A'\exp\left(-a't_1^{1-\varepsilon}\right) + \sum_{h\in\mathcal{H}} C\exp(-ct_2) + \exp\left(\frac{-4\lambda_h t_1}{3}\right) \\
&\qquad + 2\lambda_h t_1 \exp\left(\frac{-\lambda_h(1-d)t_2}{8}\right) \\
&\le A\exp\left(-at_1^{1-\varepsilon}\right) + B(|\mathcal{H}|+t_1)\exp(-bt_2)
\end{aligned}$$

for some positive constants A, a, B, and b. The following lemma therefore completes the proof of Theorem 3.2.

Lemma B.2. *Let $\varepsilon > 0$. Suppose there exist positive constants A, a, B, b such that for all $t_1 > 0$, $t_2 > 0$ and for any honest transaction tx and any finite set of honest users $\mathcal{H}$:*

$$\mathbb{P}(\mathsf{tx}\textit{ violates }(t_1+t_2,\mathcal{H})\textit{-security}) \le A\exp\left(-at_1^{1-\varepsilon}\right) + B(|\mathcal{H}|+t_1)\exp(-bt_2). \tag{26}$$

Then there exist positive constants a''' and b''' such that for all $\tau \ge 0$ and for any honest transaction tx and any finite set of honest users $\mathcal{H}$:

$$\mathbb{P}(\mathsf{tx}\textit{ violates }(\tau,\mathcal{H})\textit{-security}) \le \exp\left(-a'''\tau^{1-\varepsilon}\right) + |\mathcal{H}|\exp(-b'''\tau). \tag{27}$$

Proof. The lefthand side of (26) is zero if $\mathcal{H}=\emptyset$ so assume without loss of generality that $|\mathcal{H}| \ge 1$. Given $\tau \ge 0$, let $t_1 = t_2 = \tau/2$. Then (26) yields

$$\mathbb{P}(\mathsf{tx}\text{ violates }(\tau,\mathcal{H})\text{-security}) \le A\exp\left(-(a/2^{1-\epsilon})\tau^{1-\varepsilon}\right) + 2B(|\mathcal{H}|+\tau)\exp(-(b/2)\tau). \tag{28}$$

Let a' and b' be positive constants such that $a' < a/2^{1-\epsilon}$ and $b' < b/2$. Let $\bar{\tau}$ be so large that $A\exp\left(-[(a/2^{1-\epsilon}) - a']\tau^{1-\varepsilon}\right) \le 1$ and $2B\exp(-[(b/2)-b']\tau) \le 1$ for all $\tau \ge \bar{\tau}$. Then for $\tau \ge \bar{\tau}$

$$\mathbb{P}(\mathsf{tx}\text{ violates }(\tau,\mathcal{H})\text{-security}) \le \exp\left(-a'\tau^{1-\varepsilon}\right) + (|\mathcal{H}|+\tau)\exp(-b'\tau) \tag{29}$$

Let b'' be a positive constant with $b'' < b$. Then, using the assumption $|\mathcal{H}| \ge 1$,

$$\begin{aligned}
(|\mathcal{H}|+\tau)\exp(-b'\tau) &= |\mathcal{H}|\exp(-b''\tau) + \tau\exp(-b'\tau) - |\mathcal{H}|(\exp(-b''\tau) - \exp(-b'\tau)) \\
&\le |\mathcal{H}|\exp(-b''\tau) + \tau\exp(-b'\tau) - (\exp(-b''\tau) - \exp(-b'\tau)) \\
&= |\mathcal{H}|\exp(-b''\tau) - (\exp(-b''\tau) - (1+\tau)\exp(-b'\tau)) \\
&\le |\mathcal{H}|\exp(-b''\tau) \quad \text{for all } \tau \text{ sufficiently large}
\end{aligned} \tag{30}$$

Combining (29) and (30) implies that there exists $\bar{\tau}'$ such that

$$\mathbb{P}(\mathsf{tx}\text{ violates }(\tau,\mathcal{H})\text{-security}) \le \exp\left(-a'\tau^{1-\varepsilon}\right) + \tau\exp(-b''\tau) \quad \text{for } \tau \ge \bar{\tau}' \tag{31}$$

Select positive constants a''' and b''' such that $a''' < a'$ and $b''' < b''$ and

$$1 \le \exp\left(-a'''\tau^{1-\varepsilon}\right) + \tau\exp(-b'''\tau) \quad \text{for } 0 \le \tau \le \bar{\tau}' \tag{32}$$

Combining (31) and (32) yields (27) for all $\tau \ge 0$. □

This concludes the proof of Theorem 3.2. □