Analysis of Nakamoto Consensus

Ling Ren

University of Illinois Urbana-Champaign
renling@illinois.edu

1 Introduction

The famed Bitcoin white paper presented an unconventional (at the time) Byzantine fault tolerant consensus algorithm that is now known as the Nakamoto consensus [4]. Nakamoto consensus centers around the proof-of-work (PoW) mechanism and the "longest-chain-win" rule. It is extremely simple and can be described very succinctly: at any time, an honest node adopts the longest PoW chain to its knowledge and attempts to mine a new block that extends this longest chain; a block is committed when buried sufficiently deep in the chain. Such a simple algorithm deserves a simple analysis, which is what this paper aims to provide.

2 Model and Overview

We assume the readers are familiar with how Nakamoto consensus works and we review its basics only to introduce notations. Transactions in Nakamoto consensus are batched into blocks. Each block is linked to a unique predecessor block via PoW, thus forming a PoW chain. A block's height is its predecessor block's height plus one. Upon adopting a new longest chain, either through mining or by receiving from other nodes, a node broadcasts and mines on top of the new longest chain. Ties can be broken arbitrarily.

Mining in Nakamoto consensus is modeled by Poisson processes as done in the Bitcoin white paper. A Poisson process with rate λ is denoted as $\{N(t;\lambda), t \geq 0\}$. The number of blocks mined within a time interval (t_1, t_2) is independent of other non-overlapping internals (mining is memoryless). It follows a Poisson distribution with parameter $\lambda' = \lambda(t_2 - t_1)$, i.e., $\Pr[N(t_2) - N(t_1) = k] = p(k;\lambda') = \frac{e^{-\lambda'}\lambda'^k}{k!}$. Let α and β be the collective mining rate of honest nodes and malicious nodes, respectively. If a block is mined by an honest (resp. malicious) node, we call it an honest (resp. malicious) block. This paper makes an ideal assumption that the mechanism of mining difficulty adjustment keeps α and β stable.

We will prove the traditional safety and liveness properties.

- **Safety**. Honest nodes will not adopt different blocks at the same height.
- **Liveness** Every transaction is eventually committed by honest nodes.

The liveness property of Nakamoto consensus is analyzed as two separate parts in the literature [2, 5] as *chain growth* and *chain quality*. Together, they state that honest blocks keep making into the longest chain, and hence keep committing new transactions.

If a group of nodes have zero communication delay between them, then they can extend a chain at their collective mining rate. We assume this is the case with malicious nodes. Between honest nodes, however, we assume a known bounded communication delay of Δ. With such a delay, honest nodes extend the chain at a rate slower than their collective mining rate, because blocks mined less than Δ time apart may not extend one another. The core of the proof is to analyze the mining rate loss due to communication delay. We will show that the effective honest mining rate is at least $g\alpha$ for liveness and at least $g^2\alpha$ for safety where $g = e^{-\alpha\Delta} < 1$. Thus, Nakamoto's protocol solves consensus if the effective honest mining rate is noticeably larger than the malicious mining rate, i.e., if $g^2\alpha > (1+\delta)\beta$ for some constant $\delta > 0$. We remark that if $\Delta \ll 1/\alpha$, i.e., the communication delay is much smaller than the expected block interval, then $g \approx g^2 \approx 1$ and the above condition becomes the "honest majority" assumption.

Δ	upper bound on communication delay
α	collective honest mining rate
β	collective malicious mining rate
g	$= e^{-\alpha\Delta}$, a discount factor of honest mining rate due to communication delay

Table 1: Notation

3 Proofs

3.1 Preliminary

The following Poisson tail bound and Chernoff tail bound will be frequently used. It is not surprising that the two bounds have almost identical forms because the Poisson distribution is a limiting case of the binomial distribution. We defer the proof of Lemma 1 to appendix and use $0 < \delta < 1$ for the rest of the paper.

Lemma 1 (Poisson tail)**.** *Define* $F(k;\lambda) = \sum_{j=0}^{\lfloor k \rfloor} p(j;\lambda)$ *and* $\bar{F}(k;\lambda) = \sum_{j=\lceil k \rceil}^{\infty} p(j;\lambda)$*. For* $0 < \delta < 1$*,* $F((1-\delta)\lambda;\lambda) < e^{-\Omega(\delta^2\lambda)}$ *and* $\bar{F}((1+\delta)\lambda;\lambda) < e^{-\Omega(\delta^2\lambda)}$*.*

Lemma 2 (Chernoff)**.** *Let* $X = \sum_{i=1}^{n} X_i$ *be the sum of* n *independent Boolean random variables and* μ *be the expectation of* X*. For* $0 < \delta < 1$*,* $\Pr[X > (1+\delta)\mu] < e^{-\Omega(\delta^2\mu)}$ *and* $\Pr[X < (1-\delta)\mu] < e^{-\Omega(\delta^2\mu)}$

3.2 Non-tailgaters and Loners

Let us put all honest blocks on a time axis based on when they are mined. An honest block is essentially "wasted" if it "tailgaters", i.e., mined too closely after another block. On the other hand, honest blocks that do not tailgate contribute to the liveness and safety of Nakamoto consensus.

Definition 3 (Non-tailgaters and loners)**.** *Suppose an honest block* B *is mined at time* t*. If no other honest block is mined between time* $t - \Delta$ *and* t*, then* B *is a non-tailgater (otherwise,* B *is a tailgater). If no other honest block is mined between time* $t - \Delta$ *and* $t + \Delta$*, then* B *is a loner.*

In other words, a *non-tailgater* is mined more than Δ time after the previous honest block. A *loner* (called a "convergence opportunity" in [5]) does not tailgate and is not tailgated. Note that these notions apply to honest blocks only. The next two lemmas establish useful properties of non-tailgaters and loners.

Lemma 4. *(i) Non-tailgaters have different heights. (ii) A loner is the only honest block at its height.*

Proof. It suffices to show that if two honest blocks do not tailgate one another, then they have different heights. Let the two blocks be B and B'. Without loss of generality, assume B is mined first. B reaches all honest nodes within Δ time, which is before B' is mined (otherwise B' tailgates B). Upon receiving B, an honest node will attempt to extend B and will only mine at a height greater than B. □

Lemma 5. *During a time interval of duration* t*, (i) at least* $(1-\delta)g\alpha t$ *non-tailgaters are mined except for* $e^{-\Omega(\delta^2 g\alpha t)}$ *probability, and (ii) at least* $(1-\delta)g^2\alpha t$ *loners are mined except for* $e^{-\Omega(\delta^2 g^2\alpha t)}$ *probability.*

Proof. Since honest blocks follow a Poisson process of rate α, by Lemma 1, there are $N > (1-\frac{\delta}{2})\alpha t$ honest blocks mined during the interval, except for $e^{-\Omega(\delta^2\alpha t)}$ probability. Number these blocks $1, 2, \ldots, N$. For convenience, let us define block 0 to be the last honest block mined before the interval and block $N+1$ to be the first honest block mined after the interval. Let $X_i = 1$ if the i-th honest block mined is a non-tailgater, and 0 otherwise. Let $Y_i = 1$ if the i-th honest block mined is a loner, and 0 otherwise. The number of non-tailgaters is $X = \sum_{i=1}^{N} X_i$. The number of loners is $Y = \sum_{i=1}^{N} Y_i$.

Recall that interarrival times in a Poisson process follow independent exponential distributions with the same parameter α. Thus, $\Pr[X_i = 1] = e^{-\alpha\Delta} = g$, independent of each other. By Chernoff, $X > (1-\frac{\delta}{2})gN > (1-\frac{\delta}{2})^2 g\alpha t > (1-\delta)g\alpha t$, except for $e^{-\Omega(\delta^2 g\alpha t)}$ probability.

$Y_i = X_i X_{i+1}$, so $\Pr[Y_i = 1] = g^2$. However, Y_i and Y_{i+1} are dependent (both depend on X_{i+1}), so we cannot directly invoke Chernoff. Luckily, Y_i and Y_{i+2} are independent. Thus, Y can be broken up into two summations of independent Boolean random variables $Y = \sum_{\text{odd}} Y_i + \sum_{\text{even}} Y_i$. Applying Chernoff on the two summations separately completes the proof of (ii). □

3.3 Liveness

Theorem 6 (Chain growth)**.** *At time t, the longest chain adopted among honest nodes has length at least $(1-\delta)g\alpha t$ except for $e^{-\Omega(\delta^2 g\alpha t)}$ probability.*

Proof. Follows from Lemma 4(i) and 5(i). □

Theorem 7 (Chain quality)**.** *At time t, in the longest chain adopted among honest nodes, the fraction of honest blocks is at least $1-(1+\delta)\frac{\beta}{g\alpha}$ except for $e^{-\Omega(\delta^2\beta t)}$ probability.*

Proof. The honest fraction is smallest if all malicious blocks make it to the longest chain. At time t, except for the said probability, the number of malicious blocks $N_1(t) < (1+\delta/4)\beta t$ by Lemma 1, the chain length $N_2(t) > (1-\delta/4)g\alpha t$ by Theorem 6, and the fraction of honest blocks $1-\frac{N_1(t)}{N_2(t)} > 1-\frac{(1+\delta/4)\beta}{(1-\delta/4)g\alpha} > 1-(1+\delta)\frac{\beta}{g\alpha}$. (For the last step, note that $\frac{1+\delta/4}{1-\delta/4} < 1+\delta$.) □

3.4 Safety

Theorem 8 (Safety)**.** *Let B^* and B^{**} be two distinct blocks at the same height. If $g^2\alpha > (1+\delta)\beta$, then once an honest node adopts a chain that buries B^* by k blocks deep, no honest node will adopt a chain that buries B^{**} by k blocks deep, except for $e^{-\Omega(\delta^2 k)}$ probability.*

Proof. Let t_1 be the first time after B^* is buried k deep that some honest node adopts a chain that buries B^{**} by k blocks deep. Thus, right before t_1, an honest node (potentially the same one) adopts a chain that extends B^*. Let these two diverging chains end at B_1 and B_1' respectively. Let Block B_0' be the last common ancestor of the two diverging chains. Let B_0 be the most recent *honest* ancestor of B_0', and let t_0 be the time it is mined. Let h_0, h_0', h_1, h_1' be the height of blocks B_0, B_0', B_1, B_1', respectively. Note that if B_0' is mined by an honest node, then $B_0 = B_0'$ and $h_0 = h_0'$; otherwise, the blocks between B_0 (excluded) and B_0' (included) are all malicious blocks. Without loss of generality, assume $h_1 \le h_1'$. Figure 1 illustrate the scenario, which is similar to the one used in Garay et al. [2].

Let Z be the set of malicious blocks mined between time t_0 and t_1, a duration of t. Let Y be the set of loners mined between time $t_0+\Delta$ and $t_1-\Delta$, a duration of $t-2\Delta$. Let $t = t_1 - t_0$.

Lemma 9. $|Z| \ge |Y|$.

Proof. We first show that every loner $y \in Y$ has height $h \in (h_0, h_1]$. By time $t_0+\Delta$, all honest nodes have received B_0 with height h_0 and will never again mine on height h_0 or lower. If any loner at height $h > h_1$ has been mined before time $t_1 - \Delta$, then at t_1, no honest node will adopt a chain ending at B_1 with h_1.

Now, we prove the lemma by pairing every $y \in Y$ with a distinct malicious block in Z as follows. If y has height $h \in (h_0, h_0']$, then it is paired with the height-h block that buries B_0, which is a malicious block by the definition of B_0. If y has height $h \in (h_0' h_1]$, then it is paired with the height-h block on the diverging chain, which must be a malicious block due to Lemma 4(ii). In either case, the paired malicious block belongs to Z because it extends B_0 (mined at t_0) and is known by some honest node before t_1. Lastly, these malicious blocks are distinct because they have distinct heights (the loners have distinct heights). □

Lemma 10. *If we ignore all unlikely events, then $|Z| < |Y|$.*

Proof. If we ignore all unlikely events, then $|Y| > (1-\delta_1)g^2\alpha(t-2\Delta)$ by Lemma 5(ii) and $|Z| < (1+\delta_2)\beta t$ by Lemma 1. When k is sufficiently large (the exact condition can be derived from Lemma 11), $t - 2\Delta > \frac{t}{1+\delta_3}$. Picking $\delta_1 = \delta_2 = \delta_3 = \delta/8$, $|Y| > (1-\delta_1)g^2\alpha(t-2\Delta) > \frac{1-\delta_1}{1+\delta_3} g^2\alpha t > \frac{1+\delta_2}{1+\delta} g^2\alpha t > (1+\delta_2)\beta t > |Z|$. □

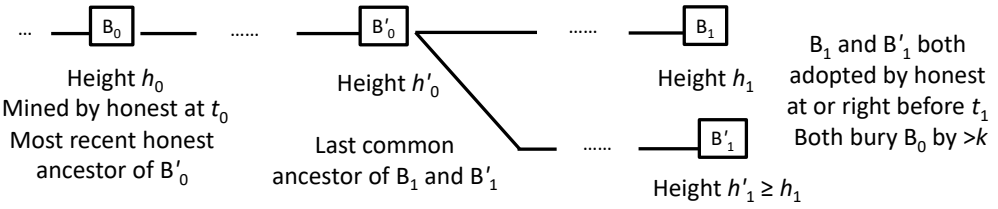

Figure 1: A chronology of the events involved in the proof of Theorem 8.

To avoid the contradiction between Lemma 9 and 10, some unlikely event involving Poisson or Chernoff tail bounds must happen. These probabilities are bounded by $e^{-\Omega(\delta^2\beta t)}$ (note that $\beta < g^2\alpha$). The following lemma helps converts the probabilities to use k and completes the proof.

Lemma 11. *$t > \frac{k}{(1+\delta)(\alpha+\beta)}$ except for $e^{-\Omega(\delta^2 k)}$ probability.*

Proof. t is smallest if all mined blocks (honest or malicious) form a chain that buries B_0. In this case, t is the sum of k interarrival times of a Poisson process with rate $\lambda = \alpha + \beta$, which follow i.i.d. exponential distributions. Sum of k i.i.d exponential distributions is an Erlang distribution, whose cumulative distribution function is $1 - \sum_{i=0}^{k-1} \frac{e^{-\lambda t}(\lambda t)^i}{i!}$, which equals the Poisson upper tail $\bar{F}(k; \lambda t)$ defined in Lemma 1. Plugging in $t = \frac{k}{(1+\delta)(\alpha+\beta)}$ and denote $l = k/(1+\delta)$, we have $\Pr[t \leq l/\lambda] = \bar{F}((1+\delta)l; l) < e^{-\Omega(\delta^2 k)}$ by Lemma 1. □

With Lemma 11, $e^{-\Omega(\delta^2\beta t)}$ is bounded by $e^{-\Omega(\delta^2\kappa)}$ because $\beta t > \frac{\beta k}{(\alpha+\beta)(1+\delta)} = \Omega(k)$. □

4 Remarks

Prior work. The first rigorous analysis of Nakamoto consensus is by Garay et al. [2], and they used the standard lock-step synchrony model. The lock-step model is a clean theoretical model but is not practical because it assumes that nodes have perfectly synchronized rounds and messages can only be sent at round boundaries. Pass et al. [5] extended the analysis to the non-lock-step synchrony model with significant added complexity. Pass and Shi [6] later simplified the analysis but it is still much longer and more involved than this paper. Kiffer et al. [3] used Markov chains to tighten the safety condition in [5]; their result has a non-closed form.

This paper also adopts the non-lock-step synchrony model and further removes the artificial notion of rounds by working with continuous time. We clarify that some papers [5, 3] incorrectly called non-lock-step synchrony "asynchroy" or "partial synchrony". Those two terms are well established in the literature and they describe much weaker models that Nakamoto consensus cannot handle. For example, the partial synchrony model assumes the communication bound Δ is "unknown", meaning that the adversary can choose Δ after all the other protocol parameters are fixed. Then, it can easily pick a very large Δ so that β far exceeds the effective honest mining rate $g\alpha = e^{-\alpha\Delta}\alpha$.

Source of simplicity. Simplicity of this paper results from many sources. We list a few.

1. We use continuous time rather than discrete, i.e., Poisson processes rather than Bernoulli trials. This simplifies various quantities and formulae. The Poisson model is by no means new; it is the model used in Nakamoto's original white paper [4]. Given that hash operations take very little time and the winning probability is very low, the Poisson model is well justified.
2. We directly abstract mining as an ideal lottery while Pass et al. [5] spent significant efforts proving this claim using the random oracle model. Because random oracles are also ideal tools that do not exist in reality, we do not find the extra complexity worthwhile.

3. Prior works [2, 5, 6] analyze chain growth and chain quality over all time intervals at all honest nodes in prior works. Such results can be easily shown in our framework using similar techniques. We opt for weaker but sufficient forms for simplicity.
4. We manage to prove the same or tighter results without needing the "no long withholding" lemma [5] or Markov chains [3].

Tightness. Theorem 6 is tight because malicious nodes can force this slow chain growth by simply remaining silent. Theorem 7 is tight given a simple variant of selfish mining [1, 2]. Tightness of the safety condition $g^2\alpha > (1+\delta)\beta$ is unknown. The best attack we know is still the simplest "private chain" attack described by Nakamoto [4] and it requires $\beta > g\alpha$. Another aspect that is quite loose in all existing analysis is the relation between latency and failure probability, i.e., for a particular κ, what is the exact probability of safety violation. We still do not have a theoretically sound way to analyze the number of confirmations needed for a given failure probability in Nakamoto consensus. These two aspects remain interesting future directions.

Acknowledgement. The author is grateful to Jiantao Jiao, Kartik Nayak and Elaine Shi for helpful discussion and feedback.

References

[1] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. *Communications of the ACM*, 61(7):95–102, 2018.

[2] 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.

[3] Lucianna Kiffer, Rajmohan Rajaraman, et al. A better method to analyze blockchain consistency. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 729–744. ACM, 2018.

[4] Satoshi Nakamoto et al. Bitcoin: A peer-to-peer electronic cash system. 2008.

[5] 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.

[6] Rafael Pass and Elaine Shi. Rethinking large-scale consensus. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 115–129. IEEE, 2017.

A Poisson Tail Bound and Chernoff Bound

Proof of Lemma 1. Recall Stirling's inequality $k! > (k/e)^k$ and Poisson distribution $p(k;\lambda) = \dfrac{e^{-\lambda}\lambda^k}{k!}$.

For $1 \le j \le k_1 = \lfloor (1-\delta)\lambda \rfloor$, we have $\dfrac{p(j-1;\lambda)}{p(j;\lambda)} = \dfrac{j}{\lambda} < 1-\delta$, and

$$\begin{aligned} F((1-\delta)\lambda;\lambda) = \sum_{j=0}^{k_1} p(j;\lambda) &< \frac{p(k_1;\lambda)}{1-(1-\delta)} = \frac{e^{-\lambda}\lambda^{k_1}}{k_1!}\cdot\frac{1}{\delta} < e^{-\lambda}\cdot\left(\frac{e\lambda}{k_1}\right)^{k_1}\cdot\frac{1}{\delta} \\ &< e^{-\lambda}\cdot\left(\frac{e}{1-\delta}\right)^{(1-\delta)\lambda}\cdot\frac{1}{\delta} && \left(\left(\tfrac{e\lambda}{k}\right)^k \text{ increases with } k \text{ when } k<\lambda\right) \\ &< \left[\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right]^{\lambda}\cdot\frac{1}{\delta} \end{aligned}$$

For $j \ge k_2 = \lceil (1+\delta)\lambda \rceil$, we have $\dfrac{p(j+1;\lambda)}{p(j;\lambda)} = \dfrac{\lambda}{j+1} < \dfrac{1}{1+\delta}$, and

$$\begin{aligned} \bar{F}((1+\delta)\lambda;\lambda) = \sum_{j=k_2}^{\infty} p(j;\lambda) &< \frac{p(k_2;\lambda)}{1-\frac{1}{1+\delta}} = \frac{e^{-\lambda}\lambda^{k_2}}{k_2!}\cdot\frac{1+\delta}{\delta} < e^{-\lambda}\cdot\left(\frac{e\lambda}{k_2}\right)^{k_2}\cdot\frac{1+\delta}{\delta} \\ &< e^{-\lambda}\cdot\left(\frac{e}{1+\delta}\right)^{(1+\delta)\lambda}\cdot\frac{1+\delta}{\delta} && \left(\left(\tfrac{e\lambda}{k}\right)^k \text{ decreases with } k \text{ when } k>\lambda\right) \\ &= \left[\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right]^{\lambda}\cdot\frac{1+\delta}{\delta} \end{aligned}$$

It is not hard to show that $(1+\delta)^{1+\delta} \ge \delta + \delta^2/3$ and $(1-\delta)^{1-\delta} \ge \delta + \delta^2/3$, which complete the proof. □