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 [3]. 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.

We will prove that Nakamoto consensus guarantees safety and liveness with overwhelming probability.

– **Safety**. Honest nodes do not commit different blocks at the same height.

– **Liveness** Every transaction is eventually committed by honest nodes.

The liveness property of Nakamoto consensus is separated into two parts in the literature [1, 4] as *chain growth* and *chain quality*. Together, they state that honest blocks keep making into the longest chain, and hence honest nodes keep committing new transactions.

We make the following assumptions: (1) mining is memoryless, (2) the global mining rate does not change, and (3) communication delay is upper bounded by Δ.

Let α and β be the collective mining rate of honest nodes and malicious nodes, respectively. A block mined by an honest (resp. malicious) node is called an honest (resp. malicious) block. By assumptions (1) and (2), the number of honest (resp. malicious) blocks mined is a Poisson process with rate α (resp. β).

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, the (bounded) communication delay will lead to a (bounded) loss in their collective mining rate. The core of the proof is to bound the honest mining rate loss. 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 well-known "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

In a Poisson process with rate λ, the number of event arrivals (blocks mined) in a time interval (t_1, t_2) is independent of other non-overlapping internals and follows a Poisson distribution with rate $\lambda(t_2 - t_1)$. The interarrival times in a Poisson process follow independent exponential distributions with the same rate parameter λ, whose cumulative distribution function is $\Pr[T > t] = e^{-\lambda t}$. Let S_k be the arrival time of the k-th event. S_k is the sum of k i.i.d. exponential distributions with rate λ, which is an Erlang distribution.

The following tail bounds will be frequently used. It is not surprising that the Poisson tail bound is almost identical to the Chernoff bound, because the Poisson distribution is a limiting case of the binomial distribution. In fact, its proof is also very similar to the proof of the Chernoff bound. We defer the proofs to the appendix. We use $0 < \delta < 1$ throughout the paper.

Lemma 1 (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 $\Pr[X < (1-\delta)\mu] < e^{-\Omega(\delta^2\mu)}$.*

Lemma 2 (Poisson tail)**.** *Let X be a Poisson random variable with rate μ (which is also its expectation). For $0 < \delta < 1$, $\Pr[X > (1+\delta)\mu] < e^{-\Omega(\delta^2\mu)}$ and $\Pr[X < (1-\delta)\mu] < e^{-\Omega(\delta^2\mu)}$.*

Lemma 3 (Erlang tail)**.** *Let S_k be the arrival time of the k-th event in a Poisson process with rate λ. For $0 < \delta < 1$, $\Pr[S_k < \frac{k}{(1+\delta)\lambda}] < e^{-\Omega(\delta^2 k)}$ and $\Pr[S_k > \frac{k}{(1-\delta)\lambda}] < e^{-\Omega(\delta^2 k)}$.*

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 "tailgates", 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 4 (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 [4]) does not tailgate and is not tailgated. Note that these notions apply to honest blocks only. The next two lemmas establish key properties of non-tailgaters and loners.

Lemma 5. *(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 6. *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. Pick $\delta_1 \in (\frac{\delta}{3}, \frac{2\delta}{3})$ such that $n = (1-\delta_1)\alpha t$ is an integer. For sufficiently large t, such an δ_1 exists. Now consider the first $n+1$ honest blocks mined since the beginning of the time interval. Number these blocks $1, 2, \ldots, n+1$. Define block 0 to be the last honest block mined before the interval.

Let A_0 be the event that $S_n \leq t$, i.e., at least n honest blocks are mined in the interval. By Lemma 3, $\Pr[\overline{A_0}] = \Pr[S_n > t] = \Pr[S_n > \frac{n}{(1-\delta_1)\alpha}] < e^{-\Omega(\delta_1^2 n)} = e^{-\Omega(\delta^2 \alpha t)}$.

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. Define $X = \sum_{i=1}^{n} X_i$ and $Y = \sum_{i=1}^{n} Y_i$.

Let A_X be the event that $X \geq (1-\delta)g\alpha t$. Recall that interarrival times in a Poisson process follow i.i.d. exponential distributions with the same rate parameter. Thus, $\Pr[X_i = 1] = e^{-\alpha\Delta} = g$, independent of each other. Let $\delta_2 = \delta - \delta_1$; we have $\delta_2 \in (\frac{\delta}{3}, \frac{2\delta}{3})$ and $(1-\delta_2)(1-\delta_1) > (1-\delta)$. Thus, $\Pr[\overline{A_X}] = \Pr[X < (1-\delta)g\alpha t] \leq \Pr[X < (1-\delta_2)gn] < e^{-\Omega(\delta_2^2 gn)} = e^{-\Omega(\delta^2 g\alpha t)}$, where the last inequality is by Chernoff.

$A_0 \cap A_X$ is the event that "at least n honest blocks are mined in the interval and at least $(1-\delta)g\alpha t$ among the first n of them are non-tailgaters", which implies part (i) of the lemma statement. Thus, the probability that part (i) does not hold is at most $\Pr[\overline{A_0 \cap A_X}] = \Pr[\overline{A_0} \cup \overline{A_X}] < e^{-\Omega(\delta^2 g\alpha t)}$.

Let A_Y be the event that $Y \geq (1-\delta)g^2\alpha t$. $Y_i = X_i X_{i+1}$, so $\Pr[Y_i = 1] = g^2$. Y_i and Y_{i+1} are dependent (both depend on X_{i+1}), so we cannot directly invoke Chernoff bounds. 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 bounds on the two summations separately yields $\Pr[\overline{A_Y}] = \Pr[Y < (1-\delta)g^2\alpha t] < e^{-\Omega(\delta^2 g^2 \alpha t)}$. The rest of the proof for part (ii) is identical to the proof of part (i). □

3.3 Liveness

Theorem 7 (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 5(i) and 6(i). □

Theorem 8 (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 $N_1(t) < (1+\delta/3)\beta t$ by Lemma 2, the chain length $N_2(t) > (1-\delta/3)g\alpha t$ by Theorem 7, and the fraction of honest blocks $1 - \frac{N_1(t)}{N_2(t)} > 1 - \frac{(1+\delta/3)\beta}{(1-\delta/3)g\alpha} > 1 - (1+\delta)\frac{\beta}{g\alpha}$. For the last step, note that $\frac{1+\delta/3}{1-\delta/3} < 1+\delta$. □

3.4 Safety

Theorem 9 (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 and B_0 be the last *honest* common ancestor of the two diverging chains. Note that if B_0' is mined by an honest node, then B_0 and B_0' are the same block; otherwise, the blocks between B_0 (excluded) and B_0' (included) are all malicious blocks. Let h_0, h_0', h_1, h_1' be the height of blocks B_0, B_0', B_1, B_1', respectively. Let t_0 be the time B_0 is mined. Without loss of generality, assume $h_1 \leq h_1'$. Figure 1 illustrate the scenario, which is similar to the one studied in [1]

Next, we derive a contradiction in the following two lemmas to finish the proof. Let $t = t_1 - t_0$. 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 $t - 2\Delta$.

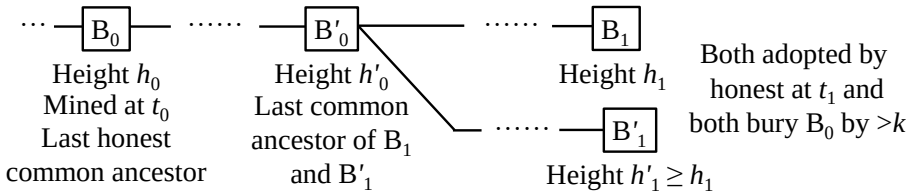

Figure 1: The scenario analyzed in the proof of Theorem 9.

Lemma 10. $|Z| \geq |Y|$.

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 5(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 11. $|Z| < |Y|$ *except for* $e^{-\Omega(\delta^2\beta t)}$ *probability.*

Proof. If we ignore all unlikely events, then $|Y| > (1-\delta_1)g^2\alpha(t-2\Delta)$ by Lemma 6(ii) and $|Z| < (1+\delta_2)\beta t$ by Lemma 2. When k is sufficiently large (the exact condition is not hard to derive), $t - 2\Delta > \frac{t}{1+\delta_3}$. Picking $\delta_1 = \delta_2 = \delta_3 = \delta/5$, $|Y| > (1-\delta_1)g^2\alpha(t-2\Delta) > \frac{1-\delta_1}{1+\delta_3}g^2\alpha t > \frac{1+\delta_2}{1+\delta}g^2\alpha t > (1+\delta_2)\beta t > |Z|$. The unlikely events involve Chernoff or Poisson tail bounds and happen with $e^{-\Omega(\delta^2\beta t)}$ probability (note $g^2\alpha > (1+\delta)\beta$). □

A safety violation can occur only if the contradiction between Lemma 10 and 11 is avoided, which happens with $e^{-\Omega(\delta^2\beta t)}$ probability. As the last step of the proof, we need to convert this probability to $e^{-\Omega(\delta^2 k)}$. To this end, note that t is larger than the $(2k)$-th arrival time in a Poisson process with rate $\lambda = \alpha + \beta$. By Lemma 3, except for $e^{-\Omega(\delta^2 k)}$ probability, $t > \frac{k}{(1+\delta)(\alpha+\beta)}$, in which case $e^{-\Omega(\delta^2\beta t)}$ is bounded by $e^{-\Omega(\delta^2 k)}$. □

4 Remarks

Prior work. The first rigorous analysis of Nakamoto consensus is by Garay et al. [1], 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. [4] extended the analysis to the non-lock-step synchrony model with significant added complexity. Pass and Shi [5] later simplified the analysis by removing the "no long withholding" lemma [4] but it is still much longer and more involved than this paper. Kiffer et al. [2] used Markov chains to tighten the safety condition; 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 [4, 2] 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 [3]. 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. [4] 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 [1, 4, 5] analyze chain growth and chain quality over all time intervals at all honest nodes. 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 using Markov chains [2].

Tightness. Whether the safety condition $g^2\alpha > (1+\delta)\beta$ is tight remains open. The best attack we know is still the simplest "private chain" attack described by Nakamoto [3] and it requires at least $\beta > g\alpha$. Another aspect that is quite loose in all existing analysis is the relation between latency and error probability, i.e., for a particular k, what is the probability of safety violation. Because of this, we still do not have a theoretically sound way to recommend the number of confirmations needed for a given error probability in Nakamoto consensus. These two aspects remain interesting future directions.

Acknowledgement. The author is grateful to Chen Feng, who pointed out a subtle mistake in a previous version of this paper and suggested the clever fix in the current proof of Lemma 6. The author thanks Jiantao Jiao, Kartik Nayak, and Elaine Shi for helpful discussions and feedback.

References

[1] 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.

[2] 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.

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

[4] 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.

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

A Proofs of Poisson and Erlang Tail Bounds

Lemma 2 (Poisson tail)**.** *Let X be a Poisson random variable with rate μ (which is also its expectation). For $0 < \delta < 1$, $\Pr[X > (1+\delta)\mu] < e^{-\Omega(\delta^2\mu)}$ and $\Pr[X < (1-\delta)\mu] < e^{-\Omega(\delta^2\mu)}$.*

Proof. Recall that the moment generating function of a Poisson random variable is $E[e^{tX}] = e^{(e^t-1)\mu}$. By Markov Inequality, we have

$$\Pr[X > (1+\delta)\mu] = \Pr[e^{tX} > e^{t(1+\delta)\mu}] < \frac{E[e^{tX}]}{e^{t(1+\delta)\mu}} = \frac{e^{(e^t-1)\mu}}{e^{t(1+\delta)\mu}}.$$

Setting $t = \ln(1+\delta)$, the right-hand side becomes $e^{[\delta-(1+\delta)\ln(1+\delta)]\mu}$. It is not hard to show (using derivatives) that $(1+\delta)\ln(1+\delta) \geq \delta + \delta^2/3$, which completes the proof of the upper tail bound. For the other side,

$$\Pr[X < (1-\delta)\mu] = \Pr[e^{-tX} > e^{-t(1-\delta)\mu}] < \frac{E[e^{-tX}]}{e^{-t(1+\delta)\mu}} = \frac{e^{(e^{-t}-1)\mu}}{e^{-t(1+\delta)\mu}}.$$

Setting $t = -\ln(1-\delta)$, the right-hand side becomes $e^{[-\delta+(1-\delta)\ln(1-\delta)]\mu}$. It is not hard to show that $(1-\delta)\ln(1-\delta) \geq -\delta + \delta^2/2$, which completes the proof. □

Lemma 3 (Erlang tail)**.** *Let S_k be the arrival time of the k-th event in a Poisson process. For $0 < \delta < 1$, $\Pr[S_k < \frac{k}{(1+\delta)\lambda}] < e^{-\Omega(\delta^2 k)}$ and $\Pr[S_k > \frac{k}{(1-\delta)\lambda}] < e^{-\Omega(\delta^2 k)}$.*

Proof. S_k follows an Erlang distribution. Its cumulative distribution function is $\Pr[S_k < t] = 1 - \sum_{i=0}^{k-1} \frac{e^{-\lambda t}(\lambda t)^i}{i!}$, which happens to be the upper tail distribution of a Poisson random variable with rate λt. Let $l = k/(1+\delta)$.

$$\Pr\left[S_k < \frac{k}{(1+\delta)\lambda}\right] = \Pr[S_k < l/\lambda] = \sum_{i=(1+\delta)l}^{\infty} \frac{e^{-l} l^i}{i!}.$$

By Lemma 2, the above probability is bounded by $e^{\Omega(-\delta^2 l)} < e^{\Omega(-\delta^2 k)}$. The other side is similar. □