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, and 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 block, either through mining or receiving from other nodes, a node adopts, broadcasts, and mines on top of the new longest chain (ties are 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), and 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 miners and malicious miners, respectively. If a block is mined by an honest (resp. malicious) miner, 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.

If a group of miners have zero communication delay between them, then they can extend a chain at the rate of their collective mining rate. We assume this is the case with malicious miners. Between honest miners, however, we assume a bounded communication delay of Δ. Namely, whenever an honest node has a new block (by either mining it or receiving it from others), it takes up to Δ time for the block to reach all other honest miners. With such a delay, honest miners extend the chain at a rate slower than their collective mining rate, because blocks mined less than Δ time apart may not all make it to the longest chain. We will show that the effective collective honest mining rate is at least $\gamma = g\alpha$ where $g = e^{-\alpha\Delta}$. In particular, honest miners can extend a chain at a rate of at least γ. Furthermore, we show that Nakamoto consensus achieves safety (defined soon) if the effective honest mining rate is noticeably larger than half of the total mining rate, i.e., if $\gamma > (1+\delta)(\alpha+\beta)/2$. Note that this implies $\beta < \alpha$.

We will prove that Nakamoto consensus satisfies the traditional safety and liveness properties.

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

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

Table 1: Notation

Since Nakamoto consensus is a randomized protocol, its guarantees are probabilistic. We will show that the probability they are violated decreases exponentially with time. The liveness property of Nakamoto consensus is analyzed as two separate parts in the literature [2, 5]: *chain growth*, the longest chain grows at a "reasonable" speed, and *chain quality*, in the longest chain, there is a "reasonable" fraction of honest blocks. These two parts combined state that honest blocks keep making into the longest chain, and hence keep committing new transactions.

3 Proofs

3.1 Preliminary

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

Lemma 1. *Let* $F(k;\lambda) = \sum_{j=0}^{k} p(j;\lambda)$ *and* $\bar{F}(k;\lambda) = \sum_{j=k}^{\infty} p(j;\lambda)$. *Let* $k_1 = \lfloor(1-\delta)\lambda\rfloor$ *and* $k_2 = \lceil(1+\delta)\lambda\rceil$ *where* $0 < \delta < 1$. *Then* $F(k_1;\lambda) < e^{-\Omega(\delta^2\lambda)}$ *and* $\bar{F}(k_2;\lambda) < e^{-\Omega(\delta^2\lambda)}$.

3.2 Liveness

To capture the loss of collective mining rate due to communication delays, we introduce the notion of *tailgaters*. If two honest blocks are mined less than Δ time apart, we call the latter block a tailgater; otherwise, we call the latter block a non-tailgater.

Lemma 2. *Honest non-tailgaters have different heights.*

Proof. Suppose for contradiction that two distinct honest blocks B and B' have the same height and neither is a tailgater. Without loss of generality, suppose B is mined first. B reaches all honest miners within Δ time, which is before B' is mined (otherwise B' tailgates B). Upon receiving B, an honest miner will attempt to extend B and will never mine a block at the same height as B from then on. □

Lemma 3. *Let* $\gamma = g\alpha$ *where* $g = e^{-\alpha\Delta}$. *In a time interval of duration* t, *the number of honest non-tailgators follows a Poisson distribution with parameter* γt.

Proof. Recall that interarrival times in a Poisson process follow independent exponential distributions with the same parameter. Thus, each honest block has a probability $g = e^{-\alpha\Delta}$ of being a non-tailgater, independent of other honest blocks. Let N and N^* be the number of honest blocks and honest non-tailgators mined during this time interval, respectively.

$$\begin{aligned}\Pr[N^* = k] &= \sum_{i=k}^{\infty} \Pr[N = i]\binom{i}{k} g^k (1-g)^{i-k} = \sum_{i=k}^{\infty} e^{-\alpha t}\frac{(\alpha t)^i}{i!}\frac{i! g^k (1-g)^{i-k}}{k!(i-k)!} \\ &= e^{-\alpha t}\frac{(g\alpha t)^k}{k!}\sum_{i=k}^{\infty}\frac{[\alpha t(1-g)]^{i-k}}{(i-k)!} = e^{-\alpha t}\frac{(g\alpha t)^k}{k!}e^{-\alpha t(1-g)} = e^{-\gamma t}\frac{(\gamma t)^k}{k!}.\end{aligned}$$

□

Remark. g is the discount factor for the collective honest mining rate. The loss is due to communication delays up to Δ. If Δ is small compared to the expected time to mine a block (which is $1/\alpha$), then $g \approx 1$.

Theorem 4 (Chain growth)**.** *If an honest node adopts a chain of length ℓ_0 at time t_0, then at time $t_0+2\Delta+t$, every honest node adopts a chain of length at least $\ell_0 + (1-\delta)\gamma t$, except for $e^{-\Omega(\delta^2\gamma t)}$ probability.*

Proof. At time $t_0+\Delta$, every honest node adopts a chain of length at least ℓ_0. Due to Lemma 3 and Lemma 1, between time $t_0+\Delta$ and time $t_0+\Delta+t$, at least $(1-\delta)\gamma t$ honest non-tailgaters are mined except for the said probability. These non-tailgaters have different heights, all greater than ℓ_0, and they reach all honest nodes by time $t_0+2\Delta+t$, giving a chain of length at least $\ell_0+(1-\delta)\gamma t$ at every honest node. □

Theorem 5 (Chain quality)**.** *At time t, in the longest chain adopted among honest nodes, the fraction of honest blocks is at least $1-(1+\delta)\frac{\beta}{\gamma}$ except for $e^{-\Omega(\delta^2\beta t)}$ probability.*

Proof. Malicious blocks occur as a Poisson process with rate β. Due to Lemma 1, the number of malicious blocks $N_1(t) < (1+\delta')\beta t$ except for $e^{-\Omega(\delta'^2\beta t)}$ probability. Theorem 4 shows that the minimum chain length is $N_2(t) > (1-\delta'')\gamma t$ except for $e^{-\Omega(\delta''^2\gamma t)}$ probability. The honest fraction is smallest if all malicious blocks make it to the longest chain. By picking $\delta' = \delta'' = \delta/4$, and observing that $\frac{1+\delta/4}{1-\delta/4} < 1+\delta$, the fraction of honest blocks is $\frac{N_2(t)-N_1(t)}{N_2(t)} > 1-\frac{1+\delta'}{1-\delta''}\cdot\frac{\beta}{\gamma} > 1-(1+\delta)\frac{\beta}{\gamma}$, except for $e^{-\Omega(\delta^2\beta t)}$ probability. □

3.3 Safety

Lemma 6. *If a block B^* is mined at time t^* and buried k blocks deep at time t^*+t, then $t > \frac{k}{(1+\delta)(\alpha+\beta)}$ 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^*. 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 $1-\sum_{i=0}^{k-1}\frac{e^{-\lambda t}(\lambda t)^i}{i!}$. Let $l = k/(1+\delta)$ and $t_1 = l/\lambda$. We have $\Pr[t \le t_1] = 1-\sum_{i=0}^{k-1}\frac{e^{-l}l^i}{i!} = \sum_{i=k}^{\infty}\frac{e^{-l}l^i}{i!}$. In the last step, note that the summands equal Poisson distribution probability density with parameter l. Since $k=(1+\delta)l$ is an integer, by Lemma 1, $\Pr[t \le t_1] < e^{-\Omega(-\delta^2 l)} = e^{-\Omega(-\delta^2 k)}$. □

Theorem 7. *Suppose $\gamma > (1+\delta)(\alpha+\beta)/2$. After an honest node adopts a chain that buries a block B^* by κ blocks deep, no honest node will adopt a chain that does not extend B^* except for $e^{-\Omega(\kappa)}$ probability.*

Proof. Let t^* be the time B^* is mined. Let t^*+t be the first time that some honest node adopts a chain that buries B^* by κ blocks deep. By Lemma 6, $t > \frac{\kappa}{(1+\delta)(\alpha+\beta)}$ except for $e^{-\Omega(\delta^2\kappa)}$ probability.

Let t_1 be the first time after t^*+t that some honest node adopts a chain that does not extend B^*. Thus, right before t_1, an honest node (potentially the same one) adopts a chain that extends B^*. Let Block B_0 be the least common ancestor of the two diverging chains. Let B_0' be the most recent *honest* block that is an ancestor of B_0 and let them be k_0 blocks apart. Let t_0 (resp. t_0') be the earlier time when B_0 (resp. B_0') is adopted by some honest node. Note that if B_0 is mined by an honest node, then $B_0' = B_0$, $k_0 = 0$, and $t_0' = t_0$; otherwise, the blocks between B_0' (excluded) and B_0 (included) are all mined by malicious nodes. The chronology of the scenario is illustrated in Figure 1.

Suppose B_0 is buried k_1 blocks deep by the shorter of the two diverging chains. This chain is adopted by some honest node at time t_1 or right before t_1. Invoke Theorem 4 from t_0 to t_1, and from t_0' to t_1, we have that, except for negligible probability,

$$k_1 > (1-\delta')\gamma(t_1-t_0-2\Delta), \quad k_1+k_0 > (1-\delta')\gamma(t_1-t_0'-2\Delta). \tag{1}$$

At time t_1, the number of blocks extending B_0' is at most $N(t_1-t_0';\beta)+N(t_1-t_0;\alpha)$: these blocks must be mined either (1) by malicious nodes after t_0', or (2) by honest nodes after t_0. Given the existence of the two diverging chains, there are at least $2k+k_0$ such blocks. However, we next show that there will be less

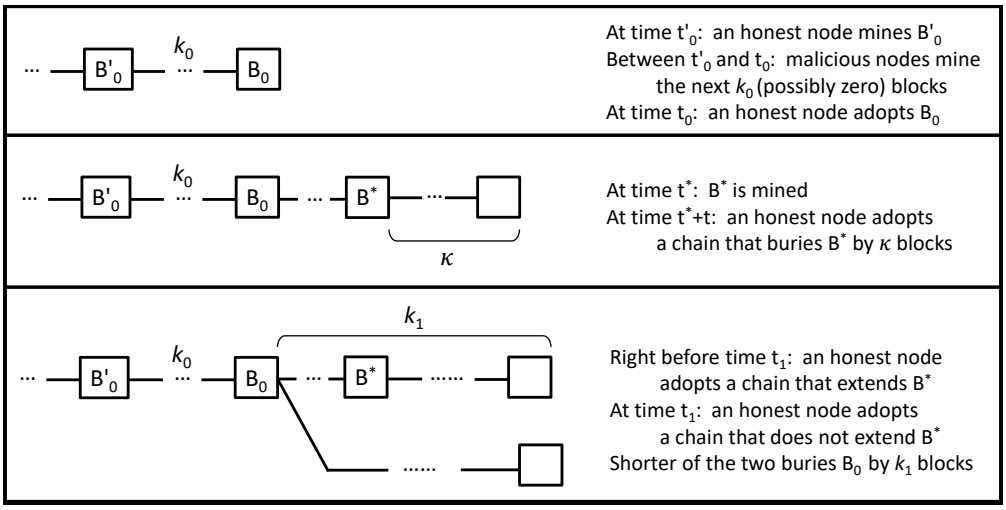

Figure 1: A chronology of the events involved in the proof of Theorem 7.

than $2k + k_0$ such blocks — unless some unlikely bad event happens. Note that $t_1 - t_0 > t > \frac{\kappa}{(1+\delta)(\alpha+\beta)}$. When κ is sufficiently large (the exact condition is easy to get), we have $2\Delta < (\delta/4)(t_1 - t_0 - 2\Delta)$, and

$$\begin{aligned}
N(t_1 - t_0'; \beta) + N(t_1 - t_0; \alpha) &= N(t_1 - t_0' - 2\Delta; \beta) + N(2\Delta; \beta) + N(t_1 - t_0 - 2\Delta; \alpha) + N(2\Delta; \alpha) \\
&< (1 + \delta/4)[N(t_1 - t_0' - 2\Delta; \beta) + N(t_1 - t_0 - 2\Delta; \alpha)] \\
&< (1 + \delta/4)(1 + \delta'')[\beta(t_1 - t_0' - 2\Delta) + \alpha(t_1 - t_0 - 2\Delta)] && \text{(Lemma 1)} \\
&< \frac{(1 + \delta/4)(1 + \delta'')}{1 - \delta'} \cdot \frac{\beta(k_1 + k_0) + \alpha k_1}{\gamma} && \text{(Inequality 1)} \\
&= \frac{(1 + \delta/4)(1 + \delta'')}{1 - \delta'} \cdot \left(\frac{\alpha + \beta}{\gamma} k_1 + \frac{\beta}{\gamma} k_0 \right) \\
&< 2k_1 + k_0.
\end{aligned}$$

In the last step, we have chosen $\delta' = \delta'' = \delta/8$, and used the condition that $\gamma > (1 + \delta)(\alpha + \beta)/2$.

To avoid this contradiction, some unlikely event regarding the Poisson tail distribution must happen. These probabilities are bounded by $e^{-\Omega(\delta^2 \lambda(t_1 - t_0))}$ where λ is either α, β, or γ. All these are bounded by $e^{-\Omega(\delta^2 \beta t)} = e^{-\Omega(\frac{\delta^2 \beta \kappa}{(\alpha+\beta)(1+\delta)})} = e^{-\Omega(\kappa)}$. □

4 Remarks

Prior work. The first rigorous analysis of Nakamoto consensus is by Garay et al. [2], and they used the standard lock-step synchronous model. The lock-step model is a clean theoretical model but is not very practical because it assumes that nodes have perfectly synchronized rounds and message can only be sent at round boundaries. Pass et al. [5] extended the analysis to the non-lock-step synchronous 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. This paper also adopts the non-lock-step synchronous model and further removes the artificial notion of rounds by working with continuous time.

Note that Pass et al. called the non-lock-step synchronous model "asynchronous" and several other papers called it "partially synchronous". Both are misnomers. The well-established terms "asynchronous" and "partially synchronous" in the literature describe much weaker models (unbounded or unknown delays) and Nakamoto consensus can handle neither.

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. 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. Chain quality was analyzed over an arbitrary period of time in prior works [2, 5]. Such results can be shown in our framework using similar techniques but we opt for simplicity.
4. Pass et al. [5] relies on an insightful but more involved notion called "convergence opportunity", which we manage to avoid (more on this later).

Tightness. Theorem 4 is tight because malicious nodes can force this slow chain growth by simply remaining silent. Theorem 5 is tight given a simple variant of selfish mining [1, 2]. The mining rate condition in our Theorem 7, i.e., $g\alpha > (1+\delta)(\alpha+\beta)/2$, is not tight. Using convergence opportunities, the mining rate condition derived in Pass et al., when translated to the Poisson model, is $g^2\alpha > (1+\delta)\beta$, which is tighter when β is small compared to α. The mining rate condition is further tightened by Kiffer et al. [3] using an improved method of counting convergence opportunities. It should not be hard to adapt these methods to our model to obtain similar results. The relation between latency and failure probability (i.e., for a particular κ, what is the exact probability of safety violation) is quite loose in all existing analysis including this one. Hence, we still do not have a theoretically sound way to analyze the number of confirmations needed for a given failure probability in Nakamoto consensus. This remains an interesting future direction.

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 Bounds

Proof of Lemma 1. Recall Striling's inequality $k! < e\sqrt{k} \cdot (k/e)^k$. For $1 \le j \le k_1 = \lfloor (1-\delta)\lambda \rfloor$, we have $\frac{p(j-1;\lambda)}{p(j;\lambda)} = \frac{j}{\lambda} < 1-\delta$, and

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

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

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

Similar to the Chernoff bound proof, when $0 < \delta < 1$, both tail probabilities above are $e^{-\Omega(\delta^2\lambda)}$. □