LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

The Bitcoin Mining Game: Open Review

Author: Nicolas Houy*†

Reviewers: Reviewer A, Reviewer B, Reviewer D

Abstract. The final version of the paper "The Bitcoin Mining Game" can be found in Ledger Vol. 1 (2016) 53-68, DOI 10.5915/LEDGER.2016.13. There were three reviewers who responded, none of whom have requested to waive their anonymity at present, and are thus listed as A, B and D. After initial review (1A), the author submitted a revised submission and response (1B). The revised submission was reviewed a second time by Reviewers A and B (2A) after which the assigned Ledger editor determined that the author had adequately addressed the reviewer concerns and asked the author for minor revisions which were carried out by the author, completing the peer- review process. Author's responses are in bullet form.

1A. Review, Initial Round

Reviewer A:

Dear Editor,

Thank you for the opportunity to review "The Bitcoin Mining Game," by Dr. Nicolas Houy. My assessment is that this paper makes several novel contributions and should be accepted for publication without hesitation. However, I believe there are two oversights—that if corrected—would alter some of the author's conclusions. Furthermore, I believe a new and important result is buried in the math and that the author should explicitly point it out. Therefore, I suggest that the paper be accepted subject to consideration of the revisions I suggest in this review.

In "The Bitcoin Mining Game," Houy analyzes the size of blocks a rational miner will produce, by balancing fees with orphaning risk. Although Rizun[1] considered a similar problem, there are important distinctions that make this work unique:

- The author examines the problem from a game theory perspective (rather than from a functional perspective).

- While Rizun analyzed in detail the case of many small miners (such that the "self-propagation" advantage was negligible), Houy considers in detail the case of two

†N. Houy (houy@gate.cnrs.fr) is an economics researcher at the Groupe d'Analyse et de Théorie Economique (GATE) at the University of Lyon, France.

*1PcPoNQ1YMihqumZzTzQirkGudZJB7EZ3E

miners, each with different hash powers where the self-propagation advantage is significant.

- Houy makes rigorous the "orphaning factor," whereas up to this point, prior work had used an approximation originally made by Andresen.[2]

Houy's work seems to confirm some of the findings of Rizun, including:

- That a fee market would exists without a block size limit.

- That the minimum fee density required for a (small) miner to profitably include a transaction is zRT^{-1}, where z is the propagation impedance (seconds/MB) and RT^{-1} is Bitcoin's inflation rate.

It also includes and intriguing new result:

- In the two-miner case, there exists a "game theoretic" limit to the maximum size of a block that a rational miner will produce *regardless of the fees paid*. In other words, even if the fees per byte become arbitrarily high, a miner controlling a fraction of the network hash power will not produce an arbitrarily large block.

This is a stronger result than what Rizun showed (he showed that the block size would be bounded for all *finite* fee densities).

I also think that author could derive some useful result concerning the self-propagation advantage of large miners, and I hope he considers doing so, as this too would be novel.

Later in this letter, I give a section-by-section review of Houy's paper. To begin, my "high-level" critiques are summarized in the following four points:

- The author uses the variable *c* to represent the cost of a transaction (its fee) and seems to treat is as some sort of protocol constant. In practice, fees are dynamic and adjusted by wallets to ensure confirmation within a reasonable time frame. I think it is necessary to re-frame this work by viewing this as a variable that dynamically responds to supply and demand. In other words, as the "market price" to get one's transaction mined.

- The author develops much of his theory from a "per transaction" view rather than from a "per kilobyte" view. Since kilobytes are ultimately what matter for block propagation, I think the work should be framed from that perspective. This would also serve to remove a few symbols (this paper already uses several different symbols and I worry that it will be difficult for a significant fraction of Ledger's readers to follow along).

- I feel a few of the symbols used for variables should be changed to make them more consistent with other work in the field.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

- The results in Section 4 are based on the propagation impedance calculated by Decker and Wattenhofer from 2013 (z = 80 sec/MB). More recent research indicates that the network connectivity is significantly better (e.g., Tradeblock study from 2015 suggested 17 sec/MB,3 as does G. Andrew Stone's recent paper).

The remainder of this letter offers feedback of the paper, section by section.

Section 2. Model

In this section, the author begins to formalize his model. My suggestions are as follows:

- To align better with the work of D. Pinna and Rizun, I suggest the author use the symbol h_i for the ith miner's relative computation power (rather than α_i).

- The statement "then, the mining Poisson process of the mining process" is awkward (paragraph 1) and should be rephrased.

- I like how the author relates the problem to the number of transactions; however, I think it is better to formalize the model using the number of bytes. Transactions can be of all different sizes and the propagation delay depends on bytes, not transactions. For example, a transaction four times as big adds four times the marginal orphaning risk. I suggest the symbol Q for the quantity of bytes. Viewing the problem from the perspective of bytes also aligns better with other work on this topic.

- After reframing in terms of bytes (or MB), I would then then define the propagation impedance as $\tau(Q) = zQ,\ z > 0.$ (That is, I would not use the k and x variables).

- With this reframing, set $c \rightarrow \rho$ and $x \rightarrow Q$ in the last sentence of this section, where ρ is the fee per byte (rather than the fee per transaction).

Section 2.1. Mining Payoffs

In this section, the author derives an equation for the miner's expected revenue as a function of the size of the block propagated, taking into account orphaning risk. While Rizun used Andresen's approximation, Houy attempts to solve the problem rigorously, and this explains the added complexity with the Π-product in Eq. (1).

My feeling is that the readers will become lost at this point, which is a shame because this derivation is foundational to the work that follows. I suggest that the author try to explain with a few words why this Π-product is needed so that readers can gain some intuition about the process even if they cannot follow the math. Further, I suggest the author try to explain in words what the terms Q and A represent (this Q should be renamed if the author accepts my suggestion to reframe this paper in terms of the quantity of bytes included in the block).

Again, I would formalize in terms of quantity of bytes rather than transactions. I think Eq. (1) would then become

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

$$\left[\prod_{j\in N, t+z(Q_i-Q_j)\geq 0} e^{-h_j T^{-1}\left(t+z(Q_i-Q_j)\right)}\right] h_i T^{-1} dt$$

and Eq. (3) would become

$$\Pi_i(\vec{Q}) = (R + \rho Q_i) P_i(\vec{Q}).$$

Lastly, the "**1**" notation took me a while to figure out and should probably be explained.

Proposition 1

Propositions 1.1 and 1.2 seem odd to me (although 1.3 and 1.4 seem clear and correct). How can the expectation value of the ith miner's revenue increase as the jth miner includes more transactions in general (which includes the case where the jth miner also increases his expected revenue)? My intuition tells me that if the jth miner increases his expected revenue, then this must decrease the expected revenue of the ith miner. Perhaps my intuition is misleading me, but could the author speak to this?

Lastly, I was not familiar with the "\{i}" notation.

Lemma 1

Reframe with $c \to \rho$ and $x \to Q$.

Section 2.2. The Bitcoin mining game

In this section, the author introduces "the Bitcoin mining game" from the game theory perspective. Because Ledger's audience is interdisciplinary, the author should explain "the usual meaning for the notation x_{-i}" rather than assuming the reader already knows. Furthermore, Ledger discourages multi-letter variables (so that multiplication can be implied between any two adjacent letter). The author should replace BR and NE with either the full words ("best response" and "Nash equilibrium") or with a single-letter symbol and a subscript.

Possible error: Should $\vec{x}$ have a subscript i in Eq. (4)?

Section 3. The two-miners case

In this section, the author applies the model he developed in Section 2 to the case of two miners.

The author should explain the Π^+ and Π^- notation.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

Proposition 5

Here, the author solves the asymmetric two-miner case. This is the most general case solved analytically in the paper.

The first equation describes the conditions under which a miner will produce an empty block. Re-writing this equation in terms of the symbols I suggested gives

$$\frac{T}{z(1-h_1)} - \frac{R}{\rho} \leq 0.$$

Solving this equation for the critical fee density below which a miner will produce empty blocks yields

$$\rho \leq z(1-h)RT^{-1}.$$

Note that this is precisely the same result as Rizun in the limit as h → 0 (Rizun considered the many-miner case where all miners have small hash power relative to the network total).

The second equation describes the size of the miner's block for non-empty blocks. Re-writing in terms of the symbols I suggest gives

$$Q = \frac{T}{z(1-h_1)} - \frac{R}{\rho}$$

Solving this equation for the fee density required for a miner to produce a block of size Q and expanding in a power series yields

$$\rho = \frac{z(1-h)R}{T-(1-h)zQ}$$

To compare to Rizun, allow $h \to 0$ and expand in a power series about $Q = 0$

$$\rho = zRT^{-1} + z^2RQT^{-2} + z^3RQ^2T^{-3} + \cdots.$$

Comparing to Rizun's Eq. (10):

$$\rho = zRT^{-1}e^{\frac{zQ}{T}}$$

$$= zRT^{-1} + z^2RQT^{-2} + \frac{1}{2}z^3RQ^2T^{-3} + \cdots.$$

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

Both the constant and linear terms in the two power series expansions are identical (as $h \to 0$). In plain language: both models predict the same marginal cost of transactions for "small blocks," and the same initial rate of increase for that marginal cost, despite approaching the problem for very different perspectives. This agreement should be noted.

However, terms after the second are different (Rizun's model is exponential whereas Houy's model is hyperbolic). It would be very interesting to know if this is due to the fact that Houy considers the two-miner case, whereas Rizun considers the many-miner case (where all miners are small), or whether the difference is due to the game-theory versus function frameworks the respective author used.

Important: A significant result is found in this section, and the author should draw more attention to it. Since the size of the block the miner will produce is

$$Q = \frac{T}{z(1-h_1)} - \frac{R}{\rho},$$

we can ask what happens in the limit as users bid for block space and the fees offered become arbitrarily large. Let's consider the "small miner case" ($h \to 0$):

$$Q_{\max} = \lim_{\rho \to \infty} \frac{T}{z} - \frac{R}{\rho} = Tz^{-1} = Q_c,$$

which is the network capacity block size as defined by G. Andrew Stone (the size of the block that would take 10 minutes to propagate). This is a fascinating result because it shows that the equilibrium block size is finite even if the fee density bid by users becomes arbitrarily high. What this shows is that no rational (small) miner would want to produce a single block greater than the network capacity!

***

In the last paragraph of Section 3, the author refers to the importance of checking the plausibility of the set of parameters for which miners do not include transactions. The fee density is not a network parameter but instead chosen by the wallet. It can respond naturally to supply and demand as required to ensure that the miners are always processing transactions. Again, this I think comes from the author viewing his "c" variable as some sort of protocol constant where all transaction pay this value, which does not reflect the reality of Bitcoin's fee market.

Section 4. The current case

In this section, the author puts numbers to his model and attempts to analyze current network conditions. I think this section needs to be re-written. The problem is that it uses the value of the propagation impedance from Decker and Wattenhofer (z = 80 sec/MB); however, more current research suggests that network propagation is much faster now (e.g., both the recent Tradeblock study and G. Andrew Stone's paper suggest ~17 sec/MB). Furthermore, this

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

section assumes the transaction fee is some sort of network constant (like I've mentioned before), while in reality wallets decide on the fee to pay in order to get the transaction confirmed.

I think the author should reframe this section by considering a range of propagation impedances (we don't really know what the correct value is—just that today it's probably close to 17 sec / MB). Also, rather than assuming a given fee density, the author should determine the fee density that would entice miners to produce non-empty blocks. Lastly, he should compare this to the average fees actually paid (see Rizun's fee market paper, Table 2 footnote (a)). I believe that by doing so, he will find that his work very closely matches that of Rizun.

Introduction and Conclusion

I suggest the author consider revising both of these sections after considering my comments at the beginning of this letter, and after re-interpreting his results (1) by using the faster propagation impedance (~17 sec/MB) and (2) considering fees to be a variable that wallets can adjust to ensure that a transaction gets confirmed (and not the "isStandard()" minimums for relay).

[1] Rizun P.R. (2015). "A transaction fee market exists without a block size limit". https://dl.dropboxusercontent.com/u/43331625/feemarket.pdf

[2] Andresen G. (2013). "Back-of-the-envelope calculations for marginal cost of transactions", https://gist.github.com/gavinandresen/5044482. Retrieved on 03/03/2014.

[3] Tradeblock (2015). "Bitcoin Network Capacity Analysis – Part 6: Data Propagation." https://tradeblock.com/blog/bitcoin-network-capacity-analysis-part-6-data-propagation

Reviewer B:

The paper analyses the incentives of a Bitcoin miner to include transactions in his blocks. The tradeoff is that more txs will result in higher fee collected if a block is found, while on the other hand slowing down the block's propagation and decreasing the chance a block will be accepted before competing ones. Since the expected payout for each miner depends on the actions of competing miners, this is modeled as a game rather than a single-player decision problem.

The main result is a characterization of the optimal number of txs to include, and the observation that in many real-world scenarios it is actually optimal not to include any transactions at all.

The results are interesting, as far as I know novel, and I could not find any material errors in the derivations. The paper would make a fine addition to the body of knowledge on mining strategies.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

I believe the significance of the result is somewhat overstated in the paper. A prime example is the sentence: "When it is the case [that the optimal strategy is to include 0 txs], Bitcoin can definitely not be used as a payment system". The author moves on to show that currently, it is in fact the case, so this statement contradicts the observational fact that Bitcoin is actually, in practice, being used as a payment system. So statements such as this should be weakened.

The paper should have a more thorough discussion of the reasons miners might include transactions despite the model saying otherwise, and the importance of these in interpreting the results and bridging the gap between theory and practice - such as a stake in the long-term health of the Bitcoin network (mentioned in the conclusions), ideology, appearances (a pool that excludes txs, in blatant disregard to the well-being of Bitcoin, might be shunned) or the power of default.

Another aspect which should be discussed is the self-balancing nature of the system. If miners follow the strategy and start excluding txs, users who want their txs included will respond by increasing their fee until it reaches a level where miners will find it profitable to include them. So while some miners might want to change their strategy in light of the results, they pose little threat to the health of Bitcoin as a whole.

This ties into a small but significant calculation error in the paper. It says, "At the time this article is written, 0.001622 BC can be bought for about $4". Unless the paper was written at a future time when the price of a bitcoin is $2,500, the correct amount is $0.4. This changes the tone of the paper - for tx inclusion to be profitable, fees only need to rise to the level of $0.4, and not to the prohibitive level of $4 as could be inferred. Also, this figure will further reduce with technology advancement - since the value of k, the marginal propagation time per kB, scales with it.

The paper focuses on a model of the mining game where only a specific problem is brought to the front - which is fine, but it could use a "related work" section with mentions of literature that works with a different set of assumptions to analyze other instances where the "expected" behavior of miners is not optimal. This can help highlight the assumptions made in this paper, and pave the way for a more sophisticated model that considers multiple effects.

I'll conclude with a few small remarks:

1. In page 3, the propagation time is assumed linear, $k(x)=kx$. It is more intuitive, and consistent with the results of Decker and Wattenhofer, to consider an affine transformation with a constant term, $k(x)=kx-c$. The function is only used in a difference with itself, so the constant term cancels and is irrelevant - but I think this point should be clarified, in a footnote perhaps.

2. In the bottom of page 5, we have "as described in Equation 2". It should actually be "Equation 3".

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

3. The following is either a typo or a confusing piece of notation: Also in the bottom of page 3, it says $S_i=(\mathbb{R}^+)^N$. S_i is supposed to be the set of strategies for player i, so I would expect it to simply be R+. It is only the set of joint strategies that is equal to R+^N.

4. In page 6, section 3, the notation Π^+ is used without prior introduction. Perhaps this equation was supposed to define it, but it's not very clear.

5. Also there, we see $x_i

1B. Author's Response

First, let me thank the three reviewers for their very valuable comments. I am sure the revised version of the paper is of much better quality thanks to their remarks. I hope this revised version will meet their expectations.

Reviewer A:

I particularly thank Reviewer A for his thorough examination of the paper. I believe I integrated (almost) all his remarks in the revised version of the paper. I enumerate in the following list the only instances where I did not follow reviewer A's suggestion or went further than suggested.

- I did not try to explain functions A and B (with the new notation) as it is only calculation means to me. However, I tried to make the math expressions a bit more clear with some more "intuitive" explanation.

- Π+ and Π− functions are unnecessary in the main text in the revised version. They have are introduced only in the Appendix in the new version and explained there.

- Reviewer A made clear his/her point about the importance to compare my results with Rizun's. I fully agree with this point. I have introduced a new section in order to deal with this. In this new section, I try to analyze the relationship between Rizun's work and mine with a little bit less technicality than suggested by the reviewer. I also tried to answer reviewer A's question about the reasons of the differences between Rizun's study and mine in verbal terms. In a nutshell, the difference does not come from the number of miners but from the difference between the game-theoretical approach I propose vs the decision theoretic approach in Rizun's work. I hope the this is clear in the text. I am convinced that this section is an important qualitative improvement for the paper.

Reviewer B:

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

I believe all remarks made by reviewer B have taken into account. In particular, I tried to not overstate the results (and I agree with Reviewer's B critics on this point). Given the space limit for articles in Ledger and the amount of information that I would like to give in this article, I chose not to add a "related literature" section (but tried to be more complete in the citations).

Reviewer D:

I tried to make more clear the limiting assumptions in the model. I agree that my model suffers some degree of non realism. Like any model. I tried to be more explicit about these limitations in the manuscript and to add the suggested citations to mention other points of view. I did not understand the critics under the title "a unilateral deviation could increase its author's benefit up to 2%": Deviation in the manuscript is to be understood with its game theoretical definition.

2A. Review, Second Round

Reviewer A:

Dear Editor,

Thank you for allowing me to review the revisions made to "The Bitcoin Mining Game," by Dr. Nicolas Houy. In my opinion, the clarity of the presentation around the mathematical details has improved significantly with the changes to the notation and the more liberal use of endnotes. The new section comparing the author's results to Rizun's has both shown the agreement between to the two theories for small Q, and has made clear one of the novel advancements made by this current work: a "game theoretic" block size limit exists regardless of the fees offered by users.

I still have two criticisms, but I believe these can be dealt with between the author and the editor. I am recommending that Ledger ACCEPT this submission.

Criticism #1.

In my opinion, the paper is still unclear regarding the "standard fee" suggested by Bitcoin Core and the "average market fee" that transactions are empirically paying. For example, page 2, paragraph 3 reads:

"The variable reward is typically 10^-4 BTC per transaction today but it can also be considered as the price on the market for space in bocks."

However, between 1 April 2015 and 30 June 2015, the Blockchain grew by Q = 5,100 MB while the total fees earned by miners over the same period averaged 17.82 Ƀ/day × 91 days →

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

M = 1,621 Ƀ (source: blockchain.info). This shows that the average market price for block space was 1621 /5100 = 0.00032 BTC / kB or 0.0002 BTC per transaction (assuming 600 bytes). So the market fees (at least over this period of time) were twice as high as what the author assumed. This means that the miners' current behavior is much closer to being a Nash equilibrium than the paper suggests.

Criticism #2.

The author assumes a propagation impedance of 17 sec / MB based on results from the Tradeblock study and the recent paper by Andrew Stone. While I think these are the best estimates to use, in my opinion it is possible that the propagation impedance between hash power is faster than this (e.g., due to the Relay Network).

The author suggests that the miners are currently not operating at a Nash equilibrium, but he shows they would be if the propagation impedance were less (i.e., if miners were better connected than he assumed). I think the author should consider mentioning this as a second explanation.

Minor notes:

P2, Par 1: "We show that the Bitcoin miners are currently not playing strategies of a Nash Equilibrium for the typical fee." → too strong…could be see Criticisms #1 and #2.

P2, Par 3: "consensus is depending on its size" → "consensus is dependent on its size"

P2, Par 4: "block at the same date, t = 0" → "block at the same time, t = 0"

P6, second equation: (x1, x2) → (Q1, Q2)

P7, Table 1: here c = 10^-4 BTC (see criticism #1)

P7, Table 1: remove reference to Rizun re. 17 sec / MB (he used Tradeblock's estimate)

P8: missing units on R values.

P9, z = 0.0049 (missing units…I think it's best to say z = 4.9 sec / MB).

Table 2: Move to appendix? I don't find this very useful. Perhaps just include a pie chart?

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13

Reviewer B:

Hi Andrew,

Yes, my concerns were adequately addressed.

I do have a few more comments (though the paper is fine as it is):

1. One of my original remarks is that there was not enough discussion of the assumptions, and how they relate to other analyses with a different set of assumptions.

The author has made some improvements in this regard - and explained the lack of even more thoroughness by the space limitation, which is legitimate.

However, I still feel there are a few points where more can be said, and one in particular that comes to mind is the assumption that there is an infinite supply of transactions that all have the same fee/size ratio.

If we assume transactions are more heterogeneous - and that if there is indeed an infinitely long tail of transactions, they come at worse and worse fees - surely we will get more interesting market behavior. I don't, of course, expect the analysis of this to be in *this* paper, but a few more words can be said.

When transactions are not in infinite supply, miners should consider not only the current offering they can include in the current block, but also how inclusion of transactions can deplete the supply for future blocks. If a miner includes a tx now he cannot include it later, so miners have less incentive to be included.

In fact, including txs can also affect the long-term prospects of the fee market. If users see their txs are included easily, they will be less willing to pay high fees in the future, to the detriment of miners. This leads to interesting behavior that has been discussed at length (though perhaps not conclusively) as early as ~5 years ago – `https://bitcointalk.org/index.php?topic=6284.0`.

Of course, all these observations are in the direction of *less* incentives to include txs, so if the author's main point is that miners are currently including more txs than they should, they do not detract from it.

2. In the conclusion section, we have "As we said, for Bitcoin to be used as an efficient payment system...". I think - but am not sure - that this is actually a relic of a previous version of the paper, and that in the current incarnation this observation has not been explicitly made previously. If so, perhaps this sentence can be reworded to avoid confusion.

Ledger is published by the University Library System of the University of Pittsburgh as part of its D-Scribe Digital Publishing Program and is cosponsored by the University of Pittsburgh Press.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.13