Michael Parisi Presicce (mjp9zn)
CS 4501 Crypto Currency Cabal Fall 2015
Project Report

1. the motivation for your project

The problem itself is that Bitcoin currently has a hard limit of 7 transactions per second, and given the rate at which Bitcoin is growing, this limit will be reached soon. Thus, a solution is required that would allow the Bitcoin network to grow alongside its demand and past this boundary.

The inspiration for my project came from the presentation given in class by Nick Skelsey and Alex Kuck. Specifically, it came from their discussion of the problem Bitcoin is facing with respect to scaling and some possible solutions that would modify Bitcoin rather than shift to something else. Throughout their lecture I wondered about each of the presented options and how they would be implemented and, in particular, about what the implications would be should each option be chosen. That is, what would be the result once Bitcoin underwent the necessary changes and how different would this be from the original Bitcoin. This question was the basis for my project and the reason why it is entitled rather openly as Bitcoin Block Size Options.

Preliminary research quickly led me to the discovery that numerous competing interests exist within Bitcoin. Despite all groups having a vested interest in the success of Bitcoin, many have differing opinions on what success means and how that success ought to be achieved. Thus, the sights of my project shifted to an attempt to find alternative possibilities that might satisfy all parties involved, or at least enumerate certain qualities that are desirable in whichever solution is realized.

2. background, including a description of related work (with references)

The background for my project is heavily intertwined with the project itself, as one of my intentions was to decide which properties of current solutions are undesirable, and construct an alternative knowing what not to include.

BIP 100

BIP 100 is the earliest BIP that tries to respond to Bitcoin's scaling problem with a solution that involves modifying the block size in some way. It was created and written-up by Jeff Garzik and currently holds about 66% of the miners' votes out of the 90% it requests to become active. The main idea of this BIP is to allow the miners to decide the maximum block size through votes given in each block's scriptSig. Specifically, this proposal would get rid of the present 1MB cap, keep the historical 32MB cap, and introduce a new floating limit. The floating limit would be recalculated every 12000 blocks (~3 months) to be the minimum value of all the votes after the top 20% and bottom 20% are cut off.

Reference: [1], [2]

BIP 101

BIP 101 is one of the main changes that BitcoinXT hopes to implement for the Bitcoin network. This proposal intends to increase the max block size in a scheduled manner and at a predictable rate. It would first increase the 1MB limit to 8MB then double every two years for 20 years, finishing with a maximum of 8192MB. It plans to begin once January 11, 2016 has passed and 750 out of the past 1000 bocks found are created by miners using BIP 101 (BitcoinXT) software.

Reference: [3], [4], [5]

BIP 102

BIP 102 is a temporary solution also proposed by Jeff Garzik. It is exceedingly simple and would just increase the current 1MB cap to 2MB. The purpose of this BIP is to give the Bitcoin community more time to determine which solution they wish to implement. One small issue with this proposal is that it included a date by when it was to be instated, and that date has already passed without the proposal being accepted. Though this is the case, this BIP still stands as an option to consider.

Reference: [6]

BIP 103

This proposal by Pieter Wuille would systematically increase the block size by about 4.4% approximately every 97 days, which is equivalent to 17.7% every year. These values are based off numbers in a Cisco report on bandwidth growth, which Wuille considers to be the bottleneck of the Bitcoin network. BIP 103 is intended to reduce the effect of large blocks on the Bitcoin network, which will be discussed later in this section, as well as minimize the strain on all nodes, allowing for greater decentralization in that respect.

Reference: [7], [8], [9]

BIP 105

BIP 105 is similar to BIP 100 in that the miners get to decide what the block size should be through the blocks they mine. This proposal was created by BtcDrak in order to address the issue that large miners, whose blocks would be more prevalent and thus would get more votes in this system, benefit from larger blocks, as I will elaborate upon later, and thus would always aim for the highest maximum block size and gain an advantage. BIP 105 intends to stifle this by inflicting a cost upon the miners who want to vote for an increase to the block size. This cost would be in the form of a higher difficulty in the block they mine. It also tabulates votes differently, selecting the median rather than the lowest of the middle 60%.

Reference: [10]

BIP 106

Written by Upal Chakraborty, BIP 106 is the first proposal that attempts to respond automatically to the amount of space used in recent block. In essence, if the size of 50% of the blocks in the past difficulty cycle are above 90% of the current max block size, the max block size doubles. Otherwise, if the size of 90% of the blocks found in the past difficulty cycle are below 50% of the max block size, the max block size halves. If neither of these cases are true, the max block size remains the same. To address a number of arguments, it was later edited to propose a second solution which would determine the max block size based on the last block size calculation as well as the transaction fees collected by miners.

Reference: [11], [12], [13], [14], [15]

Adam Back's Compromise

Like BIP 102, this proposal is a temporary measure. It would immediately increase the max block size to 2MB. Then it would increase to 4MB after 2 years, and 8MB 2 years after that. As stated, this is a provisional proposal based on the principle that there is no way to know where technology will be or how it will advance.

Reference: [16]

Flex cap

Flex cap is another proposal that is not a BIP, but has nonetheless garnered attention within the Bitcoin community as a potential solution. It was thought up by Gregory Maxwell and is another twist on giving miners control over the block size maximum. In the flex cap system, each individual miner decides for himself whether to increase his max block size with a penalty of a higher difficulty or decrease his max block size and be allowed a lower difficulty. The change in difficulty would vary depending on how much larger the miner wants their block to be, with the current working proposal being quadratic with respect to the difference in size. There would also be a limit in place to ensure no miner deviates too far from the current maximum.

Reference: [17], [18]

No Limit

In Scaling Bitcoin 2, a paper and presentation by Peter R argued for the removal of the block size limit entirely without any other changes. The rationale behind this was that whenever a miner wished to include an additional transaction in a block, there is an increase in the risk of that block orphaning due to the additional time it would take for that block to propagate across the network. This risk can be used to create a supply and demand graph, whereby the miner can determine which transactions to include for maximum profit.

Reference: [19]

Large Block Size

There has been much discussion on the potentially negative effects of allowing large blocks on the network under any of the above methods or any other method that may increase the maximum block size. This concern comes from the amount of time it takes for blocks to propagate across the network and how that time increases exponentially with the size of the block. Consider a miner that has just found a block. It sends the new block to the nodes it is connected to, and while they verify it and send it to their neighbors, the miner that found the block is already working on finding the next one. If that miner is small and possesses little hashing power, this is not much of an issue, as they have a low chance of finding the block in that time. If the miner has a lot of hashing power, this becomes more of a problem, as they have a greater chance of finding the next block before the other miners even realize that they should be looking for it. Small versus large blocks has a very similar effect as small versus large miners. Small blocks propagate quickly and give the miner that found the block very little in the way of a head start, but as the blocks grow larger, so does the time it takes for the block to reach the rest of the network. In this way, the advantage that large miners get is compounded.

The threat that this poses is not simply that they large miners will have an easier time finding blocks, but that this may eventually lead to centralization of the Bitcoin network with only a few large miners and no small ones simply because there is no environment that can sustain them. No one in the community denies that larger blocks are beneficial to larger miners, but there is some variance on how much they do.

Reference: [20], [21], [22], [23], [24]

3. Explanation of what you did

The first thing I did as part of my project was to discover and accumulate any and all problems with the individual proposals. This was done primarily in order to understand why there hadn’t been any decision on the part of the community on which proposal to accept. Once I had the problems associated with the proposals, I attempted to derive the broader qualities that make the proposals unsuitable. Additionally, I grouped the proposals on whether they showed inherent favoritism for users, miners, or nodes. In the process of assigning group preference to each proposal, I enumerated the specific needs of miners, users, and nodes. Finally, I used the information I had gained to construct a new proposal that possessed few or none of the aforementioned qualities.

4. your results

First in this section I will discuss the needs of the various parties in Bitcoin, then I will list each proposal from earlier in this report and identify its faults and shortcomings. I would like to add that this paper is not meant to show that none of the current proposals would function, but rather that they do have issues that can be addressed.

Needs of miners

The need of miners is another piece of background that ought to be addressed in this section. Most miners are a part of the Bitcoin network in order to make profit. This is currently accomplished mostly through generation of new coins, but the intention is for it to be eventually managed by transaction fees. Therefore, whatever solution is chosen must keep this this in mind and be conducive to a healthy transaction fee market.

Needs of users

The needs of the users counteract those of the miners when it comes to transaction fees. While the miners require fees in order to operate, the users obviously want to pay as little as possible to get their transactions validated and included in the block. While I emphasized paying as little as possible, it is also important that there is space in the blocks for the users' transactions.

Needs of nodes

The needs of the nodes are rather simple. Nodes are supporters of the Bitcoin network unable to make profit nor transact with a Bitcoin address. Thus, they are relatively unaffected by this problem as a whole. A couple aspects of their functioning are worth discussing, however, which are the desires for low bandwidth costs, meaning smaller blocks, and a predictable growth rate. Paying for bandwidth is the main operating cost of a node, so anything that increases the amount of bandwidth required by the node is unappreciated. Having a predictable growth rate for the max block size allows the nodes to know when they should schedule an upgrade or get out of the network due to them being unable to fulfil the demands of the network. Both of these needs that affect nodes are in the interest of decentralization.

BIP 100

The fundamental issue with BIP 100 is the amount of power it grants to the miners. Currently, mining in Bitcoin is not heavily decentralized, with 4 miners controlling over 75% of hashing power, and 7 controlling 92%. Given this, those 7 miners are in full control of the block size. They may choose to increase it to its maximum for their own benefit, or any of them with 21% hashing power may decide to keep the max block size to a minimum completely uninhibited and cause transaction fees to spike. In other words, this proposal introduces a new 21% attack and there is no guarantee or even incentive for the miners to keep the block size in line with demand.

Tags: miner control, bad for users

BIP 101

BIP 101 avoids any party having ultimate control over the block size through scheduling of upgrades at specific times. This, however, raises two important issues. First, the values for the max block size to be set to are arbitrary, as is its growth rate. They are numbers that were chosen because they sound safe, and are unrelated to the demands of the network or needs of any party. By the time it has completed its growth, the max block size will be at 8GB, and there is nothing that suggests the number of transactions would grow alongside it. Additionally, it is highly likely that the number of generated Bitcoin per block would have reached 0.78125, which when coupled with the effectively nonexistent transaction fee market would make mining unsustainable.

Tags: arbitrary cap, nonresponsive, bad for miners

BIP 102

BIP 102 is quite simple, but even it has reasons why it cannot be considered a solution. Primarily, it is the fact that it is temporary. BIP 102 is just pushing back the problem to a later date. Furthermore, like BIP 101, the 2MB proposed cap is arbitrary.

Tags: temporary, arbitrary cap

BIP 103

This proposal is unique amongst the other proposals in this paper, as it mainly caters to the needs of nodes instead of users or miners. It focuses on keeping the maximum on what the network can physically handle. This, however, addresses neither the needs of users nor that of miners adequately. If the demand for space in a block increases at a rate greater than 17.7% a year, there will not be enough space in the blocks for all users' transactions. Otherwise, if the rate is less than 17.7% per year, which is far more likely, there will be no incentive for transaction fees, and miners will suffer.

Tags: bad for miners, bad for users

BIP 105

BIP 105 addresses the large block issue when it comes to having miners in full control of the max block size, but it ignores the small block issue. Unlike BIP 100, it has no 21% attack, as it selects from the median, but it inherently incentivizes keeping the block small, as that has no difficulty penalty. Not only would miners want to keep the blocks small, so they do not have to mine at a higher difficulty, but also the blocks that do not vote for an increase would be more prevalent in the block chain by virtue of the fact that there is a higher chance to find them with the same hashing power. Additionally, having blocks mined at various difficulties will affect the difficulty calculations for the next cycle. If the calculations choose to ignore the higher difficulty faced by some miners, then the overall and average time to find blocks will decrease. If the calculations take the higher difficulty blocks into account, the amount of time it takes to find blocks will remain the same, but only until miners change their votes, either because the block size changed successfully, or because they intended to game the system for a lower difficulty.

Tags: miner control, affects difficulty calculations

BIP 106

BIP 106 is the first proposal that attempts to produce a max block size in response to activity on the network, which is something that cannot be done with scheduling nor when giving any particular party control. This principle is something I emulate in my alterative solution. BIP 106, however, is only responsive to a certain extend due to the arbitrary caps that it implements. Since it changes the max block size by a factor of 2, any transaction fee market that may have developed in the time leading up to the change is destroyed.

In BIP 106's second proposal, this is accounted for, and the max block size is changed by a ratio of the previous average block size over the current one. It also takes transaction fees into account to prevent spamming with useless transactions. This brings a completely new variable into the equation which introduces its own problems. Primarily, it fails to take the exchange rate of Bitcoin into consideration and assumes value will remain consistent between cycles, which history has shown is not necessarily the case. Another small detail is the use of an average over median is assessing the size of blocks.

Tags: bad for miners, arbitrary caps

Adam Back's Compromise

Very much like BIP 102, this is a temporary measure and this cannot be considered a true solution to the problem as it just passes the problem to a later date. Much like other proposals, it imposes an arbitrary cap that sounds right.

Tags: temporary, arbitrary caps

Flex cap

Flex caps allows for different miners to have different difficulty levels based on their decisions. This introduces one large issue that cannot be overlooked, which is how the block chain will respond blocks of difficulties. The first case is that it takes the difficulty into consideration when deciding which chain is the longest. This means a large block mined at a higher difficulty will beat a smaller block mined at a lower difficulty if both are found at the same time. Consider the event that they are not both found at the same time, however. If a small block is found first, as is more likely the case, then should a miner mining a larger block stop and build off that chain or continue on his own, hoping he can find it before there is another addition to the current one? If the general response is to continue mining the same block, then the effective hashing power of all miners is reduced, as they would not all be working on the same block/ set of transactions. This would affect difficulty calculations in a similar manner as BIP 105 would. The other option for miners is to always build off the new block, which is the same as the case where the block chain does not take difficulty into account. This would make mining the smallest block possible optimal, as they are more likely to be found first and would constitute the longest chain. Additionally, this proposal gives miners full control over the max block size of the network.

Tags: miner control, affects difficulty calculations

No Limit

The No Limit proposal was argued just recently, and so far the main issue that I see with it is that it give no figures on approximately when and at what block size their supply and demand graphs would intersect. Given the increasing rate of propagation through the network, the point described in the paper may be a long way away. It is also possible that blocks will reach a size that can hold all transactions for the whole world before encountering this point. If this is the case, then there is no drive for transaction fees that would support mining.

Tags: bad for miners

Discussion of Qualities

One of the main trends I saw appear as I searched for problems relating to the individual proposals was whether or not they were responsive to demands of the network. It is evident that solutions that implement set schedules are unable to achieve this responsiveness. By definition, changes based on a schedule will change with time as the only factor, rather than block sizes, transaction fees, or decisions made by any part of the network.

This leads me to the second problem type that fails in a similar regard, which is any solution that relies on decisions made by one group. Unless we can trust that group to a) know exactly what the network needs and b) move the network in that direction, the decisions they make are not guaranteed to serve the needs of all groups. It is much more likely that they will favor the group making the decisions.

The one proposal above that has appropriate responsiveness without incorporating decisions by any group is BIP 106 which takes an algorithmic approach. Algorithmic approaches lack inherent favoritism for any particular group, as well.

The next quality worth discussing is the preference for any one group. The most major conflict between any groups on the Bitcoin network is that between miners and users. Miners want money, and eventually expect to be supported by users, while users wish to spend as little extra money as possible to get their transactions done. Thus, an equilibrium must be reached. Having miners decide in any way, shape, or form gives them the upper hand when determining a max block size. Likewise, any option that makes the block size unbounded or at least relatively so when compared to the amount of space being used goes directly in favor of users.

The last property to be considered is whether or not a solution affects difficulty calculations. Only two of the solutions above (BIP 105 and Flex cap) involve changing the difficulty in any way, but both had issues that must be resolves before any further consideration in this area may be given. Changing the difficulty affects the difficulty calculations and may lead to faster or slower block creation. It introduces an aspect of unpredictability and unreliability with confirmations.

My Proposal

My alternative solution to the max block size issue that Bitcoin is facing is as follows. The maximum block size would be a calculated by a function of previous blocks' sizes. Which set of previous blocks would be used by the function is yet to be determined, as it demands further calculation or a decision from the Bitcoin community as a whole. The set could vary in size from 144 (~24 hours), to 2016 (~2 weeks), to anything in between or beyond, depending on how much variance is wished upon the max block size. It could also be a set that is not contiguous with recent blocks, such as the range of 144 to 288 (24-48 hours) previous blocks. This would introduce a delay and make the network less responsive to demand, but may give extra time for transaction fee competition to emerge.

What the function previously specified would attempt to accomplish is to maintain blocks' fullness at a specified level. This level, like the set of previous blocks to be used by the function, has yet to be determined, as it requires a heavy amount of analysis and decisions from the community. The fullness is a representation of the amount of space used in a block relative to the current max block size. To get a sense of fullness of the selected set of blocks, the function would average that of each individual block, or give a certain weight to the more recent blocks (which is, again, a decision that requires more thought), or take the median of the set. Once the fullness has been calculated, the max block size would change by a factor of the desired fullness divided by the actual fullness. The desired fullness is intended to be a value between 0.0 and 1.0 that, when achieved, would produce an optimal transaction fee market. This value may be discovered though either observation of the instance when current blocks reach 1MB, or calculation by someone more experienced in economics.

With this target in mind, it is obvious how my proposal would satisfy both miners and users. The miners would have a healthy transaction fee market, supporting their costs, while the users are guaranteed enough room for their transactions in the blocks. Nodes are a different story completely, as this proposal does not address them. It relies on the fact that since the block size will be on the borderline of the demand of the network, what the network can handle is inconsequential. Any node unable to perform at the level the network requires is unneeded by the network.

One concern of algorithmic approaches to finding a solution is whether or not they can be gamed, by which I mean can a certain pattern of transacting or block creation influence the size of the blocks. While it may be true in theory for my alternative, gaming the function would require a continuous flow of transactions and thus be very expensive for an attacker.

Additionally, I have considered the possibility that there may be different expected transaction rates at different times due simply to time zones. This proposal could work with that and recalculate using the same parameters every 6 blocks or so (~1hour) as to ensure the desired fullness is maintained worldwide.

In conclusion, there are many factors to be considered when attempting to describe an ideal solution to the scaling problem the Bitcoin network is facing. There are a number of conflicting interests and viewpoints to consider beyond the safety and security of the Bitcoin network. My alternative solution attempts to fulfil all interests to a fair degree.

References

BIP 100

[1] http://gtf.org/garzik/bitcoin/BIP100-blocksizechangeproposal.pdf

[2] https://www.reddit.com/r/Bitcoin/comments/39kzyt/draft_bip_100_soft_fork_block_size_increase/

BIP 101

[3] https://forum.bitcoin.com/ama-ask-me-anything/i-m-mike-hearn-creator-of-lighthouse-bitcoinj-and-bitcoin-xt-ask-me-anything-t2207.html

[4] https://en.bitcoin.it/wiki/BIP_0101

[5] https://bitcoinxt.software/patches.html

BIP 102

[6] https://github.com/bitcoin/bips/blob/master/bip-0102.mediawiki

BIP 103

[7] https://github.com/bitcoin/bips/blob/master/bip-0103.mediawiki

[8] http://cointelegraph.com/news/115048/bitcoin-core-developer-proposes-177-yearly-block-size-growth-still-no-white-smoke

[9] https://www.mail-archive.com/bitcoin-dev@lists.linuxfoundation.org/msg00853.html

BIP 105

[10] https://github.com/bitcoin/bips/blob/master/bip-0105.mediawiki

BIP 106

[11] https://github.com/bitcoin/bips/blob/master/bip-0106.mediawiki#Proposal_1__Depending_only_on_previous_block_size_calculation
[12] http://lists.linuxfoundation.org/pipermail/bitcoin-dev/2015-August/010295.html
[13] http://upalc.com/maxblocksize.php
[14] https://bitcointalk.org/index.php?topic=1154536.0
[15] https://github.com/bitcoin/bips/pull/191/files

Adam Back’s Compromise

[16] https://www.reddit.com/r/Bitcoin/comments/3ihf2b/adam_back_2mb_now_4mb_in_2_years_8mb_in_4_years/

Flex cap

[17] https://www.reddit.com/r/Bitcoin/comments/359sdd/gregory_maxwells_take_on_increasing_the_block/
[18] https://bitcoinmagazine.com/articles/can-flexcaps-settle-bitcoin-s-block-size-dispute-1446747479

No Limit

[19] https://scalingbitcoin.org/papers/feemarket.pdf

Large Blocks

[20] http://arxiv.org/pdf/1507.06183v1.pdf
[21] https://www.reddit.com/r/Bitcoin/comments/3eawhi/new_paper_by_aviv_zohar_inventor_of_ghost_bitcoin/
[22] https://bitcoindebates.miraheze.org/wiki/Higher_block_propagation_latency_favors_large_miners
[23] http://gavinandresen.ninja/are-bigger-blocks-better-for-bigger-miners
[24] https://www.reddit.com/r/Bitcoin/comments/3j7nhd/bitfury_report_on_block_size_increase/

Other/All

[25] http://gavinandresen.ninja/bigger-blocks-another-way
[26] https://bitcointalk.org/index.php?topic=1216337.0
[27] https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2015-September/011031.html
[28] https://bitcoinmagazine.com/articles/everything-need-know-proposed-changes-bitcoin-block-size-cap-1440204699

[29] https://blockchain.info/pools
[30] https://www.reddit.com/r/Bitcoin/comments/3js809/merge_pull_request_191_from_upalchakrabortypatch1/