Reasons for why centralization has occurred and potential solutions

By Tim Swanson

Many adopters mistakenly think that mining is cheap, free or miraculously perpetual. It is none of the above, over the next year alone at $650 a token or $2.3 million a day, nearly $850 million in capital and operating costs will be spent securing the network and likely 2-4 times that amount due to externalized costs. Earlier this spring a meme was repeated, that these collective costs – what was then estimated at $600 million that was irreversibly spent securing the network – was actually a good thing. Yet that is what happens in all industries with depreciating capital stock. Automobile companies could, but do not brag in commercials that they burnt $1 billion in resources and capital (coal, steel, alloys) to build their new engines – that would look conceited and vain.

Consequently, because it costs real money and investors want to recoup their costs, mining will gravitate towards solutions that provide a reliable rate of return and this ultimately leads to industrial scale mining in centralized geographic regions.

What miners are faced with is the following: the more lottery tickets (or scratch-off puzzles) that they can obtain, the more chances at winning a block as miners are continuously incrementing the nonce in hopes to get the "lucky number."

What is this lucky number? Meni Rosenfeld described it thusly, "The miner plays with the nonce to get a block, up to a point. Since nonce is a 32-bit integer which only allows for 4B values, eventually it will need to ask the server (whether locally or on a pool) for a new merkle root to work on (where things like the extra nonce have been changed)."

The graph below illustrates the choices that investors have in mining in what is otherwise a zero-sum scenario (MV=MC). Isn't there some kind of upperbound limit to operating costs (energy) in the mining process? No, because the auto adjusting difficulty rating scales with the hashrate. Recall that in Satoshi's original FAQ: "When Bitcoins start having real exchange value, the competition for coin creation will drive the price of electricity needed for generating a coin close to the value of the coin."). Dave Hudson recently ran a Monte Carlo simulation 10 million times and found that because of the Poisson process you would need to be gambler to want to bet on those odds; in contrast investors want stable, reliable flow:

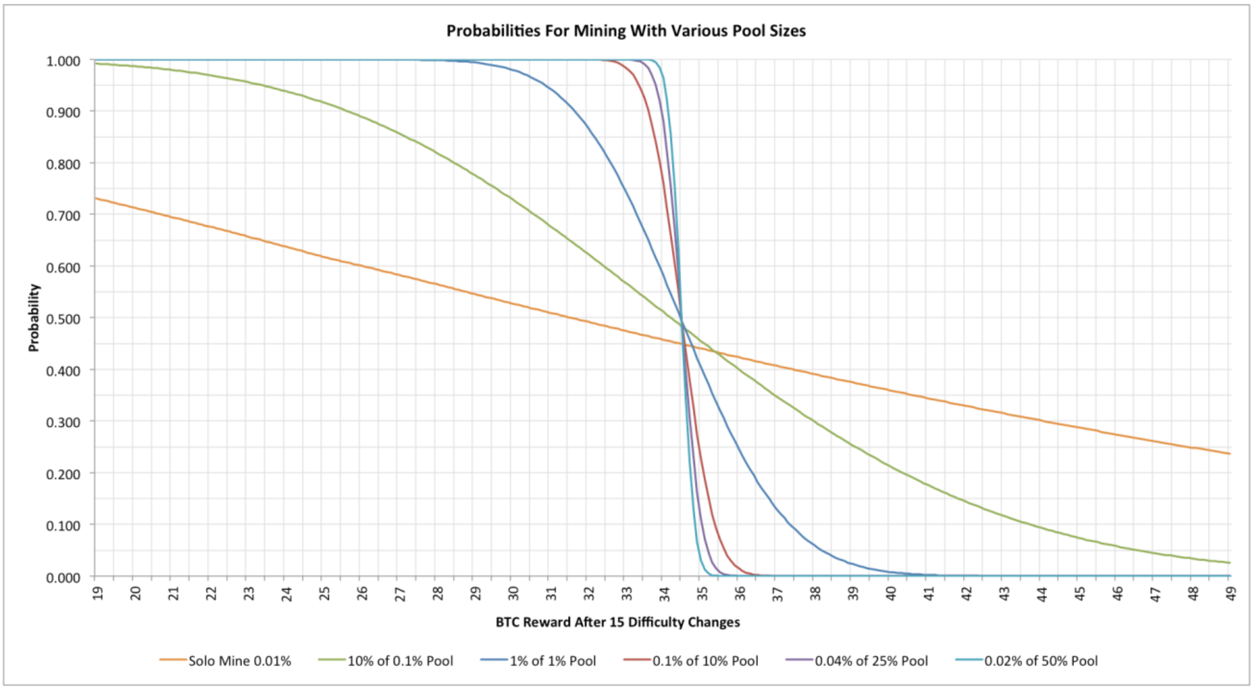

Thus as I have described before, there is an incentive to throw as much hashrate as possible to obtain the block reward before your competition does the same. Rather than repeating what has been discussed *ad infinitum*, below are several solutions:

- In his April interview (video) with *Money and Tech*, Mike Hearn explained the mining centralization issue and last week described a variety of problems and corresponding solutions to this hurdle.
- Peter Todd previously discussed this issue in a lengthy thread about "How a floating blocksize limit inevitably leads towards centralization." His solution is "Tree Chains" in *Let's Talk Bitcoin* Episode 104. Furthermore, in his interview with, *IamSatoshi* (video) when block sizes get larger, this will squeeze out more marginal players due to increased requirements (need gigabyte network throughput, terabyte hard drive platters, etc.). Someone has to pay to run a fully verifying node and need access to this level of technology (i.e., the local infrastructure has to support it).
- Two Phase Proof of Work (2P-PoW) by Ittay Eyal and Emin Sirer
- Proof of Activity: Extending Bitcoin's Proof of Work via Proof of Stake by Charlie Lee, Alex Mizrahi, Meni Rosenfeld and Iddo Bentov (as an aside, they identified three tragedy of the commons within the current protocol)
- Andrew Miller is a graduate student at the University of Maryland has at least one solution (Permacoin)
- Bitcoin Cooperative Proof-of-Stake by Stephen Reed
- Delegated Proof of Stake by Daniel Larimer
- Blockpad: Improved Proof-of-work function with descentralization incentives by Sergio Lerner
- Vitalik Buterin has some mining solutions related Ethereum, but will likely not be implemental for Bitcoin
- Greg Maxwell, a Bitcoin core developer has been discussing integrating a unique private key for each piece of hardware, soldered onto the physical hardware that is tamper resistant (not tamper proof) making it costly if destroyed. Bob could have all mining machines in one facility

but according to this design, they machines could be viewed as potentially decentralized with this quasi-TPM (Trusted Platform Module) device.

- Andrew Poelstra (andytoshi), has a paper on ASICs and decentralization noting that once you hit the thermodynamic limit of chip fabrication, the technology becomes commoditized and proliferates, potentially leading to decentralization. However this actually leads to global energy arbitrage, where miners move to the location with the cheapest energy and reliable internet access (e.g., Bitcoins: Made in China).

Perhaps the most novel approach is Proof-of-Idle (paper) by Tadge Dyrja: **[Video: https://www.youtube.com/watch?v=QN2TPeQ9mnA]**

However, irrespective of what solution is chosen it always boils down to this: what incentive do miners have to actually implement these? There is currently no incentive to implement new unprofitable code that removes the seigniorage subsidy because miners have sunk costs that have to be paid for. And there is no immediate incentive to upgrade to new software (majority of nodes are running older 0.8.5 and 0.8.6) so even if it was implemented in code, why would they upgrade? Similarly, even if proof-of-stake works (and thus far, all have led to centralization), there is no incentive for miners to use it (due to a lack of the subsidy) leading to a hard fork.

To compound this issue, there are vocal, influential members of the community effectively stonewalling efforts to discuss it – this includes those who are *not* involved in core development (the 10-15 guys consistently in #wizard IRC room), those whom have never mined before, and the largest segment: the ideological adopters who purge the community of skeptical discourse. In fact, bringing up criticism or skeptical points of view are continually met with vocal threats of "public shaming" by ideological groups – which stymies a free flowing dialogue of ideas. It cost GHash.io $90 million in hardware to achieve that level last month. A clever attacker would not need to brute force the ecosystem, but instead compromise network gear (with 0-day exploits), DDOS pools or "use a wrench":

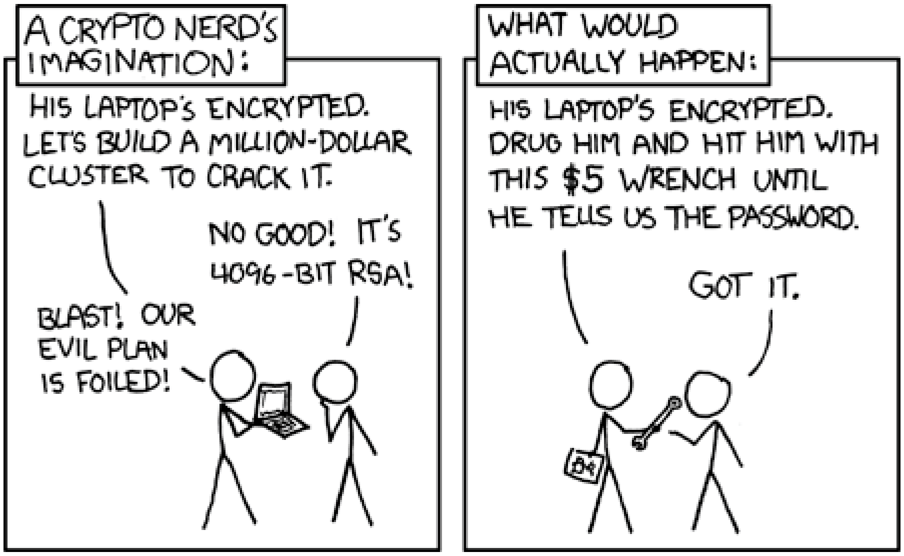
A CRYPTO NERD'S IMAGINATION:
HIS LAPTOP'S ENCRYPTED. LET'S BUILD A MILLION-DOLLAR CLUSTER TO CRACK IT.
NO GOOD! IT'S 4096-BIT RSA!
BLAST! OUR EVIL PLAN IS FOILED!
WHAT WOULD ACTUALLY HAPPEN:
HIS LAPTOP'S ENCRYPTED. DRUG HIM AND HIT HIM WITH THIS $5 WRENCH UNTIL HE TELLS US THE PASSWORD.
GOT IT.

Other solutions and hurdles

- Change the hashing algorithm from SHA256. However scrypt (which is used in Litecoin and Dogecoin) is no longer a deterrent as shown with the large supply of scrypt-based ASICs now available for commercial purchase. Other choices include: Scrypt-N, Scrypt Jane, Groestel (Grøestl), Keccak, Quark, X11, X13 (note: X11 and X13 are a cauldron of hashing algorithms).
- Getblocktemplate BIP 23 from Luke-Jr (which Hearn discussed as well), however there is no straight forward incentive mechanism for mining pools to use this today.
- Blacklisting, whitelisting and redlisting of pools that propagate certain blocks. This is a controversial issue that was debated back in April because of BitUndo (see also this thread on the Bitcoin developer mailing list).
- Change the Poisson process in the code, but then it is no longer random, see Dave Hudson's article: "Hash Rate Headaches"
- Change the difficulty reset period to another arbitrary time (instead of every 2016 blocks). The automatically readjusting difficulty rating reinforces the zero-sum of hashing (e.g., exergy is consumed linearly, MV=MC) yet any other time period would likely lead to similar result, albeit protracted (or contracted).
- Lastly, who is going to pay for and test the code? This is a public good's problem. Jeremy Allaire CEO of Circle recently challenged the developers to "step up" and create a more inclusive process for development and simultaneously explains how investors (venture funding) will secure the network. Yet investors understandably desire consistent reliable return-on-investment, this creates an incentive to mine at the large pool – to cut down on variance and orphan rates. However this still does not answer the question: who will pay for all of the code?

Is centralization a real issue?

Greg Maxwell created an attacker success probability calculator:

- 40% of hashrate, successful probability of ~50%
- 49% of hashrate, successful probability of ~96%
- 51% of hashrate, successful probability of 100%

I spoke to several other experts and below are their insights on this matter.

Robert Sams, founder of Cryptonomics:

> Choose-your-own-difficulty which goes something like this. A miner can choose what difficulty he mines at, and the reward is some non-linear function of difficulty chosen. This will allow people with inferior hardware to mine some coins, even though they'll be paying more in electricity for them than the market rate. I think people will do that, as virgin coins have anonymity value. This scheme would likely lead to MC > MV, which is good... mining will no longer be profitable (you can't "sell" virgin coins and retain their anonymity value).
>
> To my knowledge, this approach hasn't been explored in detail by anyone (including myself). But I have a gut feeling that it's promising. The essence of the idea is that the coinbase is actually more valuable than coins with a history, but it's a value that isn't tradable. If you make it feasible for people to mine some coin in a reasonable period of time, people will even if the mining costs are greater than the market value of the coin. So if, for example, all the guys buying drugs and naughty stuff acquire their coin by mining under this scheme (feasible... your commodity

> hardware may get you .1 coin in a couple of weeks), you could have mining economics that make it unfeasible for anyone to mine on scale, anyone who has to sell coin to pay for electricity bills."

Jonathan Levin, co-founder of Coinometrics:

> One really important point is to ensure that any new solution does not make things too botnet friendly.
>
> Another simple thing about this is that it is unsurprising that the bitcoin network got into this mess as it is economically rational to join the biggest pool. Minimises variance and ceteris paribus reduce orphans increasing expected return per hash. The other point is that there is still hardware bottlenecks so designing the theoretically most robust system may fail due to market imperfections. Implicitly in many arguments I hear about mining people assume perfect competition. Do we need to remind people what are the necessary conditions for perfect competition? Perfect information, equal access to markets, zero transportation costs, many players .... this is clearly not going to be a perfectly competitive decentralised market but it certainly should not favour inherently the big players.

Dave Babbitt an interdisciplinary graduate student at Northwestern University:

> Centralization wouldn't have been a surprise if they modeled bitcoin before they launched it. (As I keep on saying to myself while looking at the huge number of hours required to get it done.) Efforts to formally model the Bitcoin economy didn't start picking up steam until February of this year. But the cross-disciplinary field of Agent-Based Modeling (ABM) was mature enough in 2007 to do just that. The Bitcoin economy could have been modeled with readily-available software and clusters. Even certain equation-based models would've validly predicted the centralization problems we are having with Bitcoin. Devs are using phrases like "you don't need to model the web to design TCP/IP" to justify not worrying about the economic aspects of their design. But just as Kleinrock, Baran, Davies, and Licklider modeled the packet net before Kahn and Cerf designed TCP/IP, so core developers should have modeled the currency and banking aspects of their design.

Sergio Lerner, an independent security researcher at Certimix:

> The only way to give a theoretical solution to the mining centralization problem is by forcing miners to use real identities, and people vote/trust on those. This is because with anonymous mining all miners could be controlled by a single party. Having real identities implies legal liabilities and users trust, which in turn implies centralization (institutions, pool, companies) to reduce personal risks and provide higher trust. So it's a paradox. Decentralization looks more like Ripple paradigm than Bitcoin paradigm.
>
> Some argue proof-of-stake of hybrid system can have better decentralization incentives. All methods I've analyzed are inherently more complex and have many security problems than simple proof-of-work. So I expect decentralization comes on the form of a proof-of-work mining that practically (not theoretically) has deterrents against centralization; scrypt with a high memory footprint does it. Also see my LIMIO protocol as an innovative way of descentralization in addition to the Blockpad proof-of-work already mentioned.

In conclusion, there will likely be dozens perhaps hundreds of other proposals and experiments in the coming months and years, each with their own pros and cons. For instance, one potential issue highlighted by Sams' proposed approach is that the block reward is programmed to decrease and get smaller. Simultaneously it cannot be known as to whether or not that the dollar value of the reward is going to get larger (the two are not causally linked). If mining moves to individuals who do not mind mining at a loss in the quest for an anonymous, "virgin" coinbase that does not have a history, then perhaps this loss-bearing activity can continue for years.

Another ongoing issue will be botnets. In the beginning Satoshi Nakamoto assumed that botnets were actually a good thing because they might reduce spam based botnets – yet it is clear that they simply externalize the costs onto other parts of the economy and squeeze out marginal participants.

In the end, despite the multitude of avenues presented above, proof-of-work may simply not be a viable solution as a trustless means for arriving at a consensus in a distributed manner.