Will colored coin extensibility throw a wrench into the automated information security costs of Bitcoin?

By Tim Swanson

Revised: May 29, 2014

The cost of securing the Bitcoin network for a given length of time is roughly equivalent to the value of the block reward over the same time. In economic terms this reads as MP = MC. In Bitcoin and most of its descendants, the labor force (called miners) are provided a hard-coded wage, a seigniorage subsidy called a block reward roughly every 8-10 minutes in consideration for their providing security and processing transactions. In return, this labor force provides the security in a method called "proof-of-work" – hashing through benign math work until it finds a special number, broadcasting that solution to the network (the other laborers) and, once a block is found, repeating the cycle once again.

Is there an economic flaw of proof-of-work as it relates to security? For instance, on most cryptocurrency chains the asset value of the chain has to be proportional to the proof-of-work otherwise this could lead to an economic incentive to attack the chain. Compounding this issue are new financial instruments such as metacoins, colored coins and smart contracts that can be exchanged on the same chains and unquestionably increase the enterprise value of the chain, yet which do not proportionally incentivize security beyond the existing seigniorage subsidy.

Economically rational laborers will not spend more than the value of a bitcoin to extract the rents of that bitcoin. Because mining rewards were fixed with the genesis block in 2009 (providing a fixed income on a scheduled time table), and market participants are able to determine the percentage of the overall hashrate at a given time that their mining equipment represents, only relatively simple calculations are required to gauge the potential profitability of their mining activities.

In practice, laborers on the Bitcoin network must account for the capital costs of their hashing equipment, rent for the land, administrative overhead, taxes and increasingly important, the energy costs which can be very specific to their locality, depending on the equipment's geographic location. All of these costs are tallied against an inelastic wage which can only be attained if the hashing equipment they control is able to outcompete other such miners – it is a zero-sum game. And it can be scaled.

The Hashrate Wars

This subsequent escalation, dubbed a "hashrate war" (the competitive fight for ever increasing hashing equipment) created a technological S-curve that looks similar to the chart below:

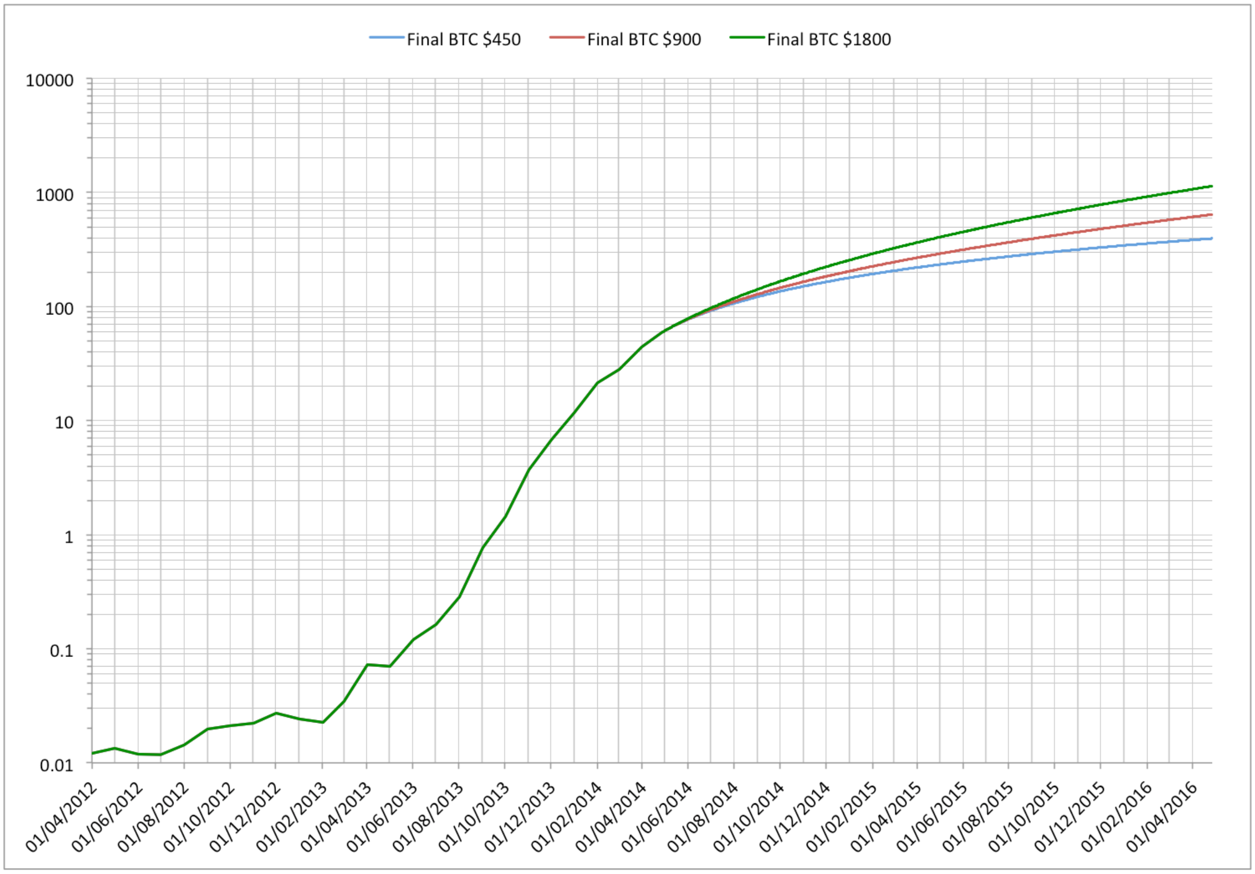

The vertical axis in the chart above is logarithmic and illustrates the hashing rate (showing that it will slow down once ASICs hit fabrication node limitations). The horizontal axis projects two years into the future (see also Bespoke Silicon).

Ignoring all of the various issues related to public goods challenges and game theory (such as "selfish mining"), this system has served the bootstrapping phase with relative ease. If it continues to expand at its current rate, the hardware side could potentially become commoditized in the next 3-4 years whereupon a miner's competitive advantage will solely lay in energy arbitrage. In fact, Satoshi Nakamoto, the pseudonymous creator of the protocol foresaw this noting in the original FAQ that "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."

Thus the relationship between enterprise value and hashing power has been known for some time.

A challenge however, presents itself when this seigniorage subsidy is halved, a structural feature of most cryptocurrencies. With Bitcoin, every 4 years (or every 210,000 blocks) the subsidy is reduced by 50%. This is equivalent to the miners – the labor force – being told they would receive a 50% pay cut. While this issue typically remains hidden and muted when token values appreciate and rise, in the long run continual halvings discentivize laborers from providing security and utility to the network. There have been several "cryptocurrencies" whose labor force fled after their profitability period was over – most notably with Auroracoin – and as a consequence the network was left insecure and vulnerable to double-spending attacks (called a 51% attack).

One such popular token that is currently facing this dilemma is Dogecoin, which is losing 20-30% of its security force every 2 months. While there are potential solutions Dogecoin developers could adopt, incorporate or migrate to, because Dogecoin is still relatively young it has the flexibility of moving towards a different security mechanism. This issue has the potential to become systemic – and thus more difficult to address – in other digital currency ecosystems.

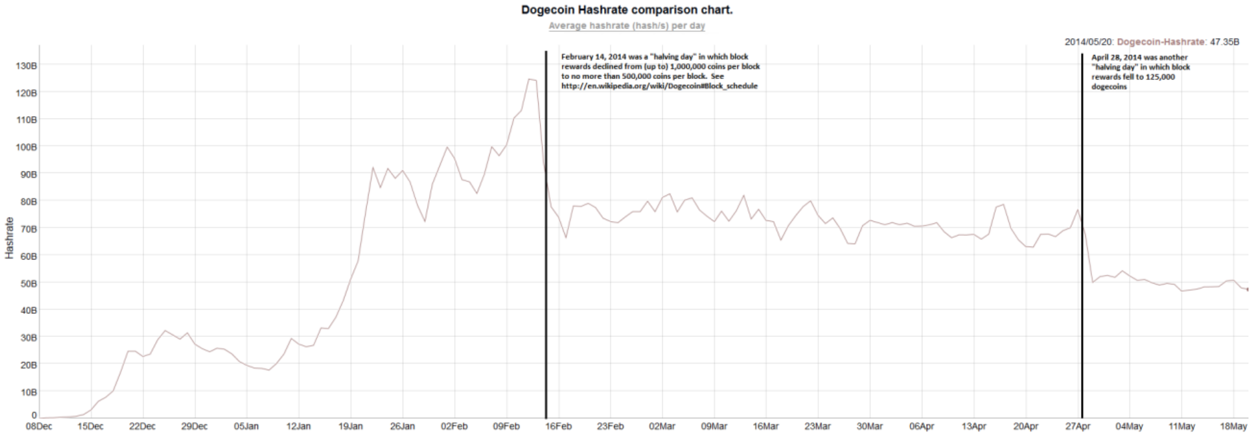

Are there any other areas of asymmetric, unbalanced security?

Colored coins, metacoins, smart contracts and user-created assets are buzzwords trumpeted by many cryptocurrency enthusiasts this past year. I even wrote a short book about these groundbreaking possibilities. Considerable publicity has been dedicated to new functionality which promises to expand the extensibility of cryptoprotocols to go beyond tracking ledger entries for just one specific blockchain-managed asset (a coin) and allows users to instead "colored" tokens to represent cars, houses, commodities, stocks, bonds and other financial instruments and wares. For example, there are several colored coin projects currently in beta that allow users to take a fraction of a bitcoin, such as 0.001 BTC and "color" it "blue" (or any other arbitrary color) which represents say, a specific make and model of an automobile like a 2010 Camry LE. The user can then transfer that asset, the title of the Camry, along a cryptoledger (such as the Bitcoin network) to other individuals. Instead of having to transfer tens, hundreds or thousands of bitcoins in exchange for a good or service, users can instead exchange and manage entire asset classes in a trustless, relatively decentralized framework.

However, in this model the labor force providing security has no incentive to consume more capital or create additional hashrate just because the market value of colored coins is in excess of the uncolored value (since the value of miners' new coins will be solely based on uncolored exchange value). Just because social conventions on the edges of the network add value perceptions to the network, based on the current code, miners do not automatically receive any additional value for providing that security.

So we should ask: does this raise the risk of a double-spend? Perhaps, because more hashrate is required for a proof-of-work blockchain with additional color value transactions on the chain. Yet, there is no automatic mechanism to do reward this additional labor leading a (remote) possibility of having to remove some Script's altogether. Script is the built-in scripting language used for creating and customizing transactions.

The gap between mining value and enterprise value

For instance, assuming this colored coin technology works and is adopted by 1,000 people the following scenario could take place. The total market value of a block reward (currently 25 bitcoins) is roughly $12,500 (or $500 per bitcoin), thus *ceteris paribus* the labor force is only spending $12,500 every 10 minutes to secure the blockchain (in practice it is a lot more, there are several exceptions). One such exception is the expectation of token value appreciation – that is to say that if Bob the miner believes that a bitcoin's value is $1000, but the price is currently $500, Bob is still willing to expend up to $1000 for mining each bitcoin, discounted by his internal calculation for the probability that bitcoin will rise to that price. However, if colored coins are adopted and used via the built-in scripting methods, there is potential for a seemingly unlimited amount of assets to be traded on the Bitcoin network. If these several thousand colored coin users add additional value, this creates an incentive for attackers to attack the network through colored coin-based double-spending attacks.

For example, where each of these 10,000 users places the title of a 2010 Camry each valued at $10,000 that would theoretically add $100 million in value that the network is transferring, but for which miners are not being proportionally rewarded or paid to secure those assets. As a consequence, over time as tens of thousands of assets – and functionality – are added to the network, the gap between mining reward value and enterprise value widens which creates a vulnerability, an economic incentive for criminals to use hashrate to attack the network. A rogue attacker could sell an asset and build a competing tree (consensus in Bitcoin is based on whatever is the longest tree of blocks). After a successful 51% attack, the rogue attacker could then broadcast a fake chain built without the corresponding asset, having switched it out thus effectively double-spending. And if the total value that the network is transacting is at least twice as much as bitcoin value is, then there is a financial incentive for rogue participants to attack the network. The impact of a successful attack involves a lot of speculation and will likely fill continue to provide researchers many more volumes of conjecture and modeling.

Money for nothing

This scenario raises the question: what then is the potential divergence in value between bitcoin the currency and bitcoin the network (which can transfer and protect other data)? This issue only presents itself now as, previously, only bitcoins – and no other apps, assets or instruments – existed on the network. This gives rise to a coordination problem because miners would have to also keep track of the color, keep track of the exchanges the color is being traded on, and keep track of the settlement price (if there is such a thing) so that they could adequately gauge market clearing prices and readjust the coinbase reward every 10 minutes. Again, even if this coordination problem is solved the seigniorage reward does not increase – the current fixed income does not reflect the actual value being transacted on the network. So colored coins on a fully decentralized network could end up on an *undersecured* network of their own making with the only solution: recode the block rewards based on the value of the color and this presents a number of technical and social engineering challenges. In some ways this issue is related to the hypothetical economic disconnection between blacklisted and whitelisted tokens (due to Coin Validation) – a blacklisted token would be sold for less than what a whitelisted token would sell for.

A follow-up question that the community will likely debate is: Why wouldn't the value of a bitcoin increase as items of value are transferred on the blockchain via colored coins or another protocol, such that the miner's block rewards would adequately compensate the miners? According to Preston Byrne, a securitization attorney in London the answer to this is "that the value of bitcoin used in a colored coin transaction does not need to bear any relationship to the value of the associated asset – the network is being used to transmit information, and that information represents rights, and is the rights – not the token – which are valuable." If the price of bitcoin does not adequately incentivize the miners, then there will be a difference between value of a bitcoin and the network and then some entity will have to step in to compensate for that difference. Whether collective action is sufficient to provide this compensation is currently unknown but there are coordination problems inherent in this model that would make this difficult.

In contrast, the Ripple protocol, sidechains and perhaps even a proof-of-stake system could probably alleviate at least this specific concern. These alternative consensus mechanisms have one advantage to hash-based proof of work systems like Bitcoin, at least for the transfer of non-crypto value (i.e., colored coins). For instance, Ripple's distributed consensus mechanism allows users to exchange assets via gateways without needing to proportionally incentivize the security labor force. This is not necessarily an endorsement of this particular platform, rather it serves as examples of how it is immune to that particular attack vector.

Alternative approaches to network security

I reached out to several experts for their views on this issue. According to Robert Sams, founder of Krtyptonomic and Cryptonomics:

> One of the arguments against the double-spend and 51% attacks is that it needs to incorporate the effect a successful attack would have on the exchange rate. As coloured coins represent claims to assets whose value will often have no connection to the exchange rate, it potentially strengthens the attack vector of focusing a double spend on some large-value colour. But then, I've always thought the whole double-spend thing could be reduced significantly if both legs of the exchange were represented on a single tx (buyer's bitcoin and seller's coloured coin).
>
> The other issue concerns what colour really represents. The idea is that colour acts like a bearer asset, whoever possesses it owns it, just like bitcoin. But this raises the whole blacklisted coin question that you refer to in the paper. Is the issuer of colour (say, a company floating its equity on the blockchain) going to pay dividends to the holder of a coloured coin widely believed to have been acquired through a double-spend? With services like Coin Validation, you ruin fungibility of coins that way, so all coins need to be treated the same (easy to accomplish if, say, the zerocoin protocol were incorporated). But colour? The expectations are different here, I believe.
>
> On a practical level, I just don't see how psudo-anonymous colour would ever represent anything more than fringe assets. A registry of real identities mapping to the public keys would need to be kept by someone. This is certainly the case if you ever wanted these assets to be recognised by current law.

> But in a purely binary world where this is not the case, I would expect that colour issuers would "de-colour" coins it believed were acquired through double-spend, or maybe single bitcoin-vs-colour tx would make that whole attack vector irrelevant anyway. In which case, we're back to the question of what happens when the colour value of the blockchain greatly exceeds that of the bitcoin monetary base? Who knows, really depends on the details of the colour infrastructure. Could someone sell short the crypto equity market and launch a 51% attack? I guess, but then the attacker is left with a bunch of bitcoin whose value is...
>
> The more interesting question for me is this: what happens to colour "ownership" when the network comes under 51% control? Without a registry mapping real identities to public keys, a psudo-anonymous network of coloured assets on a network controlled by one guy is just junk, no longer represents anything (unless the 51% hasher is benevolent of course). Nobody can make a claim on the colour issuer's assets. So perhaps this is the real attack vector: a bunch of issuers get together (say, they're issuers of coloured coin bonds) to launch a 51% attack to extinguish their debts. If the value of that colour is much greater than cost of hashing 51% of the network, that attack vector seems to work.

In other words, while these new financial instruments could technically be exchanged in a trustless manner, the current protocol cannot automatically incentivize their protection or account for their enterprise value, the equivalent of using a mall security guard to protect Fort Knox. While miners may be able to protect against amateurish shoplifters or even unorganized cat burglars, once organized criminals calculate and realize that one “color” asset is worth the economic effort of attacking the vault they may try to do so. And because the blockchain is public and color assets could be known to the world-at-large, taking the Fort Knox analogy further, this would be like a mall cop standing in front of the contents of Fort Knox piled up on an open field (or behind a see-through glass vault). It is an attempt to guard the Crown jewels not in a fortress with armed guards, tanks and turrets, but with Paul Blart.

On this point, Jonathan Levin, co-founder of Coinometrics explained that:

> We don’t know how much proof of work is enough for the existing system and building financially valuable layers on top do not contribute any economic incentives to secure the network further. These incentives are fixed in terms of Bitcoin - which may lead to an interesting result where people who are dependent on coloured coin implementations hoard bitcoins to attempt to and increase the price of Bitcoin and thus provide incentives to miners.

It should also be noted that the engineers and those promoting extensibility such as colored coins do not see the technology as being limited in this way. If all colored coins can represent is 'fringe assets' then the level of interest in them would be minimal. Time will tell whether this is the case. Yet if Bob could decolor assets, in this scenario, an issuer of a colored coin has (inadvertently) granted itself the ability to delegitimize the bearer assets as easily as it created them. And arguably, decoloring does not offer Bob any added insurance that the coin has been fully redeemed, it is just an extra transaction at the end of the round trip to the issuer. That is an implicit negative for investors and users. This raises some concerns in the future, if a party had the ability to invalidate Bitcoin accounts based on their own criteria that the miners might gain an influence over the colored coins and may bias various aspects of the economy incentivized through some kind of backchannel payment. For instance, BitUndo is a new “double spending as a service” project that is trying to do just that, provide a way for users to send transactions to a mining pool in an attempt to reverse transactions something that has created a flurry of reactions in the community. In the end, colored coins ends up being expensive through imposed TX fees, and thus becomes less attractive to issuers and users.

According to Alex Mizrahi, lead developer of Chromawallet a colored coin project:

> It is true that currently block subsidy has a significant impact on network's security, but it is not meant to work this way in the long run.
>
> We'll go through 5 subsidy halvings in next 20 years, at that point block subsidy will be around 0.78 BTC. Reward miners get from fees is already on that scale (e.g. 0.134 BTC here) even though blocks aren't full yet.
>
> So transaction fees are going to play bigger role than subsidy. And value of those fees is linked to usefulness of transactions (i.e. value of those transactions) rather than to exchange rate.
>
> Colored coins increase incentive to attack, but they also increase usefulness of transactions, thus it isn't clear whether they will have negative or positive impact on network security.
>
> A couple other comments: "Script" is not required for colored coins, they work with very plain bitcoin transactions too. The incentive structure for bitcoin mining sucks from security perspective anyway, so I hope we'll eventually upgrade to a better protocol (e.g. including proof-of-stake) regardless of colored coin woes. And merged-mined sidechains will have even worse problems unless they are 'hardened' in some way.

I also contacted Jack Wang, co-founder of Bitfoo, a hosted wallet that was the first to implement proof-of-reserves. In his view:

> The security of the network depends on the aggregate hashing power. In one method of implementation, if Colored Coins could pay just one pool, say Eligius, extra to prioritize their transactions, but Eligius had only, say 25% of the network power, then the rest of the network could collectively decide to exclude the blocks that Eligius mined. This makes some sense to me since Eligius itself couldn't secure the network, yet is the only pool extracting the extra value out of Colored Coins. Colored Coins would need to distribute the extra rents to at least 50% of the network, and unless this lies within one pool then this is a danger to the Bitcoin network, but if it is 2 or more, this requires coordination and introduces potential holdout problems.
>
> A more natural way to implement this would be that colored coins users would pay higher transaction fees on their own so that any and all miners that included those transactions in their blocks would get more fees. But unless those fees are mandated by colored coins, what is the incentive for individual colored coins users to pay extra?

Towards a more functional future

While this is a speculative issue, what is knowable is that the economics behind it are math-based and built into these protocols. What is also known is that some proposed solutions should be easier to implement than others. For instance, Bitcoin developers could fork the code and create a proof-of-stake ledger proposed by Stephen Reed. Alternatively, because this new extensibility could create fungibility issues, a different – and admittedly impractical – solution might be for mining pools to utilize a trusted

Oracle data feed to colored coin exchanges and adjust mining rewards accordingly. Perhaps removing scripts entirely and relying on merge-mined sidechains, instead, could alleviate this potential pain point as well.

What is definitely known is that market participants have every incentive to keep miners mining. If fees are floated users will likely pay higher transaction fees if they do not want miners to go elsewhere. While speculative, colored coins users could become the biggest payer of transaction fees, though in practice, most users do not like paying any fee. Over the past several months this is an issue that Mastercoin and Counterparty developers have promoted: pay the miners higher fees for access to these new platforms because miners expect the value of these special transactions to go beyond the excess of bitcoin transactions. Miners could potentially auction block priority to these transactions over regular bitcoin transactions. One pool, Eligius, operated by Luke-Jr is already filtering out specific bitcoin transaction today. In conclusion, the interaction between second-generation blockchain technology and first-generation incentive mechanisms will continue to be thought-provoking. It is certainly an issue to keep one's eye on in the coming years.

Note: I would like to thank Preston Byrne, Petri Kajander and Taariq Lewis for their comments; and Joshua Zeidner for bringing this issue to my attention and for his extensive feedback.