

How many bitcoins does it cost to maintain the Bitcoin network?

Tim Swanson

Let us be quite clear: if Bitcoin was a cheaper or more efficient transaction method, for-profit organizations such as large payment processors would have forked it long ago and would likely already be using it internally in order to shore up their margins. They do not because it is not cheaper, in fact, it is significantly more expensive to maintain than any of a number alternative centralized methods (e.g., running MongoDB on [a Pi server](#)).

The bottom line to them is that the marginal value in these centralized solutions has to be greater than the cost of maintaining it ($MV > MC$) otherwise none of the companies would be able to generate a profit. As described below, Bitcoin currently does not fulfill that characteristic.

And this error – that Bitcoin's network is cheap or free – is an oft repeated theme in the Bitcoin community and shares a common root in the 'seen versus unseen' in the aggregate transaction costs of the network.

For instance, Gavin Andresen recently [wrote an article](#) on speeding up block propagation:

People claiming that "Bitcoin Doesn't Scale" are theoretically correct: you still need $O(n)$ bandwidth and CPU to fully validate n transactions-per-second.

Someday, when Bitcoin is the number 2 payment network in the world, we might have to start worrying about that. Here are a couple of back-of-the-envelope calculations that show that we should be able to scale up to $n=15,000$ transactions per second before running into that $O(n)$ bandwidth limit.

For perspective, the number 1 payment network in the world today (Visa) handles about 212 million transactions per day; 2,500 transactions per second on average. Their peak processing capacity, needed on the busiest shopping days, is reported to be 40,000 tps.

My home Internet connection is getting about 90 megabits download bandwidth per second right now. An average Bitcoin transaction is about 2,000 bits, so my current consumer-level Internet connection could download 45,000 transactions per second, over ten times average Visa transaction volume.

While it is nice to know that I *could* run a full node handling more-than-Visa-scale transaction volume from my house, running a dedicated machine in a data center somewhere makes more sense. 15,000 250-byte transactions per second works out to about 7 terabytes of bandwidth per month. One of my hosting providers charges \$20

per month for a virtual private server with 8 TB of bandwidth per month-- or \$240 per year to handle MasterCard-level transaction volume today (August 2014).

Andresen's solution to the propagation scaling issue is innovative and seems valid at first glance. But it only factors in propagation costs and not the actual securing or transactional processing services provided by miners. Hashing is not free. As described copiously in [Chapter 3](#), the network does not run on goodwill – quite the opposite. And consequently the real fee, the real cost to use the Bitcoin network [fluctuates](#) between \$30 - \$50 per transaction. These costs are not directly seen by users. Additional coins minted are in fact inflation, inflation which devalues all existing coins – thus these transactions come at the cost of an inflation or dilution tax (e.g., without mining there would be no new coins but also no transactions).

This economic description contrasts with others space including Antonis Polemitis [who disagrees](#) with the ways to measure and calculate the actual costs of the network, contending that these are “self-correcting and of no concern to a BTC user” and “don't underestimate the fact that anyone can enter mining just through \$\$\$.”

These are both false. Contra Polemitis, the network costs real money to run that users will eventually have to face once subsidies disappear. For comparison, Visa externalizes those costs in less than 1% of the fees (see Richard Brown's new [piece](#) on that).

The way to measure the actual costs is as follows: in a competitive open market the marginal revenue (or value) from a good or service trends towards the cost of that same good or service ($MV=MC$). We empirically see this time and again in virtually every market segment with low barriers to entry. If that value or revenue is removed, the labor force goes elsewhere. Bitcoin is no different.

It costs one bitcoin to make one bitcoin?

Thus the real question should be, in a perfectly competitive marketplace, how much of a bitcoin does it cost to make a bitcoin?

For the sake of argument, let us assume that bitcoin is the unit of account and all costs (machines, property, administrative overhead, electricity) are also denominated in bitcoin as well.

For example, miners (the labor force) are continually competing in a process of ‘burning’ one type of good (bitcoin) to make the same good (bitcoin). That is to say, they continually have to expend value somewhere with the goal of receiving an equal or larger amount of value in return.

What ultimately happens is that, in theory, the miners will spend no more than one bitcoin to extract one bitcoin as a reward for securing and providing transactional ability of the network – thus the cost of generating (or creating) one bitcoin will in the long-run equal the value of one

bitcoin. [Note: the spread between the two is called seigniorage, or as we refer to it in Bitcoin, a block reward.]

However, in practice there are, as Jonathan Levin has pointed out to me, a variety of disequilibrium's at play: primarily mining farms that receive the best, most efficient gear before anyone else. For them, it costs a small fraction of a bitcoin, to make a bitcoin. These are 'bumper coins' to them (like a [bumper crop](#) in agriculture). Yet, because mining is a zero-sum game, very few participants will ever reap similar rewards – in fact, most marginal participants do not because they are not the first ones to receive batches of the best equipment next. And again, contra Polemitis, the barrier to entry in August 2014 is not “just through \$\$\$” – it is large amounts of capital and the right connections to the next batch of equipment.

Consequently, with the Bitcoin network (and virtually all of its progeny) 99.8% of the income for the labor force comes from the seigniorage subsidy which if removed, the labor force leaves (it is actually 100% for almost all other networks because very few people pay transaction fees). These are real costs that few people are readily acknowledging and for whatever reason, assume that they do not exist.

Peter Todd, a Bitcoin developer, recognizes this issue, [stating](#):

Bitcoin is not an efficient payment system. It replaces what could be a single centrally managed datacenter with a vast army of miners turning electricity into heat, and thousands of copies of the central ledger. That's why the cost per transaction from the inflation subsidy is currently \$30USD. The one thing Bitcoin has going for it compared to existing systems is freedom from regulation, and it pays a heavy price to get that.

Robert Sams and Ray Dillinger have both independently written on this issue in the past.

Dillinger [equated](#) this hashing process equivalent to private money printers, “where people are spending money in an auction for the right to print money. Such an auction is more or less guaranteed to bring the costs of printing money right up to its value, which is an unnecessary (and unwanted) feature.”

Likewise, Sams [noted](#) eight months ago:

Miners (the peers who choose to do the hashing) will work on new blocks only when the expected value of the mining award exceeds the cost of electricity required to run the hashing hardware. There are no restrictions of entry to mining, and the equilibrating mechanism is the protocol's hashing difficulty. If the coin's exchange value increases, making mining profitable at current difficulty, more miners will join the hashing effort and because of this, after 2016 blocks the protocol will adjust the difficulty upward making expected value of mining = costs of mining again. The same process works in reverse in the scenario where exchange value decreases. In the creation of crypto coins, $MC = MP$.

[Note: MP stands for marginal productivity but can be used interchangeability with MR or MV.]

As I first [mentioned](#) in May, the ephemeral seigniorage spread miners (private money printers) see is continuously arbitrated away to where $MV=MC$. In contrast, that spread in a state seigniorage system is always $MV>MC$ (except for [pennies and nickels](#)...).

This is not an endorsement of either policy, but rather explains the actual economic motivations and phenomenon behind mining.

Conclusion

And at first glance, Proof of Idle seems to be a solution to this (see also Anton Bolotinsky's two short [critiques](#)) but it may be the case that the only solution to prevent $MV=MC$, that seems to be arising, is attempted collusion and cartelization in which the barriers to entry into mining eventually become prohibitive to new entrants; and as could be the case, as core developer are hired by mining pools and manufacturers, the core development code remains the same preventing a hard fork towards less energy intensive processes (a type of "regulatory capture").

Obviously this last part is speculative, but the underlying economic constraints have been known for some time and cannot be hand waived away. Nor does this mean that Bitcoin cannot be used for other purposes. As mentioned by Todd and argued in numerous publications by myself, it is not a competitive payment processor but it is, or at least pre-BINO was good at distributing trust. And incidentally, Polemitis has a [good post](#) on those potential use-cases.

In conclusion, the point of Bitcoin was to distribute trust. The word "trust" (or variation thereof) appears 11 times in the body of the whitepaper, once in a diagram and in one reference. That is what Satoshi Nakamoto was purportedly most passionate about. To distribute trust he had to design a network that axiomatically consumes more resources than a "trusted" centralized network. That's not to say it is "bad," but by design it cannot be greener or more efficient relative to running Mongo on a Pi box and still function in a trustless manner.