

Decentralised Currencies Are Probably  
Impossible  
But Let's At Least Make Them Efficient

Ben Laurie  
(ben@links.org)

Tue Jul 05 09:22:21 2011 +0100 (19:483a5782de61)

## 1 Abstract

Lately, there's been a good deal of excitement about Bitcoin[1], an (allegedly) decentralised currency, based on proof-of-work.

I explore the limitations and costs of Bitcoin and introduce an efficient alternative.

Both Bitcoin and my alternative proposal suffer from a problem for which there is no known solution: creating consensus in a group with open, changing membership. But at least my proposal fails in an energy efficient way, unlike Bitcoin.

## 2 What is a Currency?

A currency consists of a finite pool of tokens, each representing some amount of "value"<sup>1</sup>. Let's call these tokens *coins*.

Each coin is in the possession of exactly one participant in the scheme at any one time, and it is possible to transfer coins between participants. Let's call these participants *purses*<sup>2</sup>.

The number and identity of coins and purses may vary over time, but at any particular time there must be agreement about which coins exist and which purses they are in<sup>3</sup>.

## 3 Agreement

In traditional fiat currencies, agreement on the question of which coins exist is achieved by fiat, as their name suggests: some central authority issues coins and is the final arbiter of the validity of coins<sup>4</sup>. The question of which coin is in which purse is settled by the simple fact that the coins are material - they can only be in one place at a time.

However, it has also long been possible to represent coins in a purely notional way, such as when I deposit money in a bank account. Very often this money has never existed as actual physical currency. But nevertheless its allocation to the appropriate purse is handled by a central authority through a hierarchy of delegated powers (for example, the Bank of England allows recognised banks to represent money as mere annotations in bookkeeping systems, and banks, in turn, allow me to write cheques, which cause these annotations to change).

But we are thinking about decentralised systems. In this case, there can be no central authority to defer to. Instead we must have agreement (or consensus) amongst some group. Group consensus is a well-studied problem and can be

---

<sup>1</sup>I will not attempt to define "value" and will instead rely on your intuitions, since the meaning of "value" is not important to the technical discussion. For example, you could think of each token as representing £1.

<sup>2</sup>Of course, the purse is probably owned by someone - you, me or a bank.

<sup>3</sup>In practice there may be temporary periods of uncertainty, which is OK so long as agreement is eventually reached. If periods of uncertainty are extensive, then the currency is unlikely to be very usable, but still fits technically within my definition.

<sup>4</sup>In practice a certain level of forgery is usually tolerated by the authority.

arrived at in many ways, but in essence all solutions are the same: consensus is arrived at when some sufficient number of members of the group agree, where “sufficient” means enough such that, under the rules of consensus, whatever they are, no number of dissenting opinions would cause a change in the agreement. For example, we could say that consensus is arrived at when more than half the members agree, and this would work, since the remaining members cannot change the consensus<sup>5</sup>.

To match this to the notion of “decentralised” (i.e. lacking central authority), the consensus group must be, at least, all participants in the currency. This does not present any real problem when that group is known. For example, it would be possible to define the group as “all people currently in the United States” – where the currency would be something akin to the US Dollar. Assuming the majority decide to behave honestly (as seems likely, after all, that is what happens now), then they should have no difficulty in forming consensus on who has how much money at what time<sup>6</sup>.

However, the most general notion of decentralisation does not admit such restrictions. After all, in some sense, placing any such restriction simply pushes the central authority back a layer: instead of controlling the currency, the authority controls membership of the consensus group.

A system like this must allow any entity to participate, and to join and leave the scheme at will. And here lies the problem. If you can never know who is in the scheme (bear in mind that knowing who is in is also a consensus problem!), then you can never get agreement.

## 4 Bitcoin Agreement

Now that we understand the core problem, namely that of agreement, we can quite easily understand Bitcoin’s solution to the problem.

Bitcoin defines the consensus group as “all the computing power in existence”, and requires participants to prove their possession of whatever fraction of this power they care to spend on Bitcoin by using it to produce proof-of-work tokens.

And once we state the problem like this, we can quite clearly see the flaw. Until at least half of the computing power in existence is actually used to produce Bitcoins, we cannot know that we have consensus!

If, for example, 1% of the total power available<sup>7</sup> is used to produce Bitcoins at present (in fact, the amount is far less than that), then at any point someone could come along with a further 1.1% of the total power and use this to define their own consensus<sup>8</sup>, thus invalidating all the work, *and all the money*, of the initial group, and instead take possession of the entire currency for themselves.

---

<sup>5</sup>Note that a central authority is just a special case of this where we define consensus to mean “whatever the central authority thinks”.

<sup>6</sup>Of course, I assume a very modern world in which everyone has devices acting on their behalf which are connected to the Internet, busy forming this consensus.

<sup>7</sup>Strictly, I mean energy rather than power, since Bitcoin actually, in effect, sums power over time.

<sup>8</sup>By forking history right back to the first block, and producing a hash chain that is longer than the current consensus.

That is, in the pure Bitcoin model, where “longest chain wins” is the only rule, an adversary with more power than you can always come along at some point with a longer chain than the one you thought was longest, and effectively unwind all coin generation and transactions back to the first mined block.

Even worse, it is clear that arriving at the equilibrium state for Bitcoin is incredibly expensive: half of all the computing power in existence must be burnt, in perpetuity, maintaining agreement about the current state of the currency.

It also unknowable: we can never be sure that we actually are burning half of all the power in existence, because we do not know how much power exists.

## 5 Checkpointed Bitcoin

The Bitcoin developers have recognised this problem, and so they do not, in fact, run the pure Bitcoin protocol. Instead they introduced periodic *checkpoints*, which are moments at which a snapshot of the state of the Bitcoin currency is taken<sup>9</sup>. Once these snapshots are established, it is not permitted to go back to before the snapshot and revise history.

Note that there is something of an ad hoc balancing act in taking these snapshots: you can't take one that's very recent indeed, because you can't be sure whether the snapshot you are seeing is really the consensus, or whether you are seeing a branch off the “main” chain that is going to be invalidated by a consensus you happen to not be part of at the moment. The longer you wait, the more certain you are that you have seen consensus; but you are never totally certain.

The exact mechanism by which these snapshots are established is not important, so long as one thing is understood: they cannot be established by universal consensus under the rules of Bitcoin, or they would suffer from the same problems discussed above - it would always be possible that they might be later invalidated.

So, the mechanism by which these checkpoints are established must use some other form of consensus, and one which is robust over time. This implies, if you believe what I have said so far, that the protocol for checkpoints must operate through consensus in a known group<sup>10</sup>, *or* that the problems of group consensus I have outlined have been solved by some mechanism other than Bitcoin (and one that does not require large amounts of computing power).

## 6 Basis For An Efficient Solution

So now we have either an existence proof of an efficient solution, or a proof that Bitcoin doesn't work.

If Bitcoin is, indeed, using a known consensus group, then it has, after all, a central authority (that consensus group), and is not, therefore, a decentralised currency.

---

<sup>9</sup>The nature of Bitcoin makes this a very cheap operation - essentially one has to remember a single, relatively small, number.

<sup>10</sup>In practice, I believe this group is the Bitcoin developers.

If, on the other hand, Bitcoin has somehow solved this problem without a centralised authority, then we can use whatever mechanism it uses for agreeing checkpoints as the basis for an efficient solution. Let's label this magical protocol *efficient unbounded agreement*.

## 7 An Efficient Solution

Optimistically assuming that Bitcoin is, in fact, a decentralised currency, let us proceed. We can now use efficient unbounded agreement as a building block to create an efficient currency.

First of all, we need to issue coins. Bitcoin's model is that a new coin should come into existence roughly once every ten minutes. So, let's follow that model. Once every ten minutes, using efficient unbounded agreement, we agree that a new coin is issued - these could be numbered sequentially.

Next we have to agree who should get the coin. This is not particularly hard. First we use efficient unbounded agreement to number the current participants<sup>11</sup> sequentially. We then use it to agree a consensus random number. This could be done, for example, by agreeing a commitment for each participant, and then revealing the value they committed to, adding them all together and taking the modulo of that total, which would randomly designate a participant.

Once initial allocation has been dealt with, transactions are trivial - we simply agree to them with efficient unbounded agreement.

## 8 Conclusion

Of course, it is far more likely that Bitcoin has not solved the core problem and is therefore not a decentralised currency.

But if it has, I have shown that we could instead save a lot of energy by using an efficient protocol.

Alternatively, we could conclude that whilst Bitcoin is not strictly decentralised, it is as good an approximation as we can get. However, we must appreciate that this approximation relies on a certain level of honest behaviour from certain parties, and trust in those parties. If we have such behaviour and trust, why not leverage them in an efficient protocol, instead of burning CPU on proof-of-work?

## 9 Acknowledgements

I would like to thank Brian Warner, Zooko, Leigh Honeywell and Ray Daniels for useful comments and corrections.

---

<sup>11</sup>By "current" I mean just the participants in the current consensus - i.e. the set of all instances of the software that are currently running and in communication with each other.

## References

- [1] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://www.bitcoin.org/bitcoin.pdf>.