Cryptocurrencies: principles and challenges
What are they for? How do they work?

Arthur Breitman
Tezos

[special issue of *Réalités Industrielles,* february 2019]

Abstract:
The origins of cryptocurrencies are both political and technological. To understand their purpose and usefulness, we will present these two origins by looking first at their ideological roots, then by briefly describing the cryptographic and information technologies used in rolling out blockchain solutions.

From a political project to a technological one

To understand the technical foundations of cryptocurrencies, we must first understand their political and social roots. Originally, cryptocurrencies were inspired by libertarian and "cypherpunk" thinking, with libertarians attempting to establish (or in fact restore) the separation between government and currency, while cypherpunks endeavour to use cryptography to defend privacy.

These communities of thought focused on currency because it lies at the heart of all economic activity and is therefore central to most human endeavours. Currency is involved in trade, in contracts and investments, and in the relationship between citizens and the government. Even family or religious structures, which are non-commercial by their very nature, cannot completely avoid currency. Therefore, controlling the currency inherently means controlling society. This means controlling the economy, first of all, through monetary policy, even though the creation of independent central banks was a major step towards curbing the historical abuses of the right to mint money. More recently, controlling the currency has become a way to control individuals thanks to the digitalisation of payments and the transformation of the banking system into an electronic surveillance mechanism.

This power of control is an asset and a safeguard for the public authorities, especially for fighting crime or collecting taxes. However, it also presents a significant risk for civil liberties.

English version: Translation Centre of the French economy and finance ministries

On the one hand, securing data on a wide scale often proves to be unexpectedly difficult. Electronic systems frequently fall prey to attacks by criminal organisations or to economic espionage by foreign governments. Nowadays, a French company can no longer count on confidentiality for its transactions, and even more so, for its suppliers, clients, trips by its executives, etc.

This argument is reminiscent of the "crypto wars" of the 1990s, which resulted in liberalisation of encryption technologies in the US. In 1993, as industrial groups showed growing interest in encryption, the NSA developed a chipset known as the Clipper Chip. This chipset gave civilians access to encryption technologies that had previously been reserved for the armed forces. As a compromise, the chipset openly included a "backdoor" that allowed intelligence services and law enforcement agencies to decode messages. In less than a year, cryptographer Matt Blaze[1] cracked the algorithm, showing that the backdoor was actually wide open. US authorities eventually had to concede the obvious fact that mathematics is neutral and does not distinguish between users' intentions. Attempting to curb technology is in vain or done at the expense of everyone's security. Confidentiality is either absolute or it isn't; there is no "backdoor".

On the other hand, considering the very high stakes, it would not seem cautious or reasonable to presume that all public authorities are benevolent. This is not a matter of being paranoid or believing in conspiracy theories; it is simply an application of the principle of precaution. This circumspect approach has very respectable historical precedents, such as the ratification of the US Constitution or the Declaration of the Rights of Man and of the Citizen of 1789.

To take a less daunting example, the fact (as mentioned above) that central banks are by law independent clearly draws on lessons from history to remove governments' discretionary power to issue currency. Yet aside from the matter of monetary policy, there must also be safeguards against the risk of totalitarianism.

In 2017, more than half the world's population was living under an authoritarian regime; 44% were living in a dictatorship. These modern regimes have considerable technological resources. The Chinese government, equipped with the tools to monitor all electronic payments, is currently developing Orwellian surveillance systems that give citizens ratings based on their consumption habits, their social ties or their political opinions. Venezuela currently uses the same technology. What impact would total surveillance of exchanges have during ethnic cleansing? These are genuine risks that cannot be ignored even in western countries. The survivors of the Vel' d'Hiv mass arrest of Jews in Paris in 1942 can still bear witness and they urge us never to forget.

[1] M. Blaze, "Protocol Failure in the Escrowed Encryption Standard", *Proceedings of the 2nd ACM Conference on Computer and Communications Security*, pp. 59-67.

The political and economic risk associated with the monetary system is not – to paraphrase Soviet dissenter Aleksandr Solzhenitsyn – the evil design of dark-souled men. Instead, it reflects a basic technological limitation. Historically, and until quite recently, transfers of money at a distance involved either the transport of banknotes and coins, or interbank credit.

To protect individuals' civic liberties, cryptocurrencies offer an alternative system in which cash actually becomes electronic. Cryptocurrencies make remote payments possible without resorting to credit, and therefore without resorting to the banking system and the governmental apparatus needed to enforce the payment of debts. The design of these cryptocurrencies is based on decades of research in cryptography and distributed computing.

The digital signature: the cornerstone for cryptocurrencies

In 1976, Whitfield Diffie and Martin Hellman described the concept of a digital signature and broadened the application scope for cryptography beyond coding and decoding. Coding is generally used to preserve the confidentiality of messages; digital signatures attest to their authenticity. These signatures are unforgeable, unalterable and irrevocable. In fact, they prove that the signer holds a private digital key that is only known to the signer. From smart cards to websites, these digital signatures are omnipresent nowadays.

By creating an identification process for electronic systems that can be mathematically verified, these signatures paved the way for digital currency. However, attempts to build decentralised monetary systems based on digital signatures have run into the double spend problem. The rest of this paper sketches out certain key factors for designing these systems.

The double spend problem: a stumbling block

To illustrate this point, let's make a naïve attempt to build a simplified system. Let's assume that, from the outset, Alice holds all the money in the system, i.e. 1,000 doubloons.[2] Alice spends this amount by signing two digital cheques: one for 700 doubloons payable to Oscar, and the other for 300 doubloons payable to Bernard. The validity of these transactions is verified with the authenticity of Alice's signature and the fact that $700 + 300 = 1{,}000$. Let's assume that Oscar decides to spend the 700 doubloons that he received from Alice. Oscar signs a transaction that transfers 700 doubloons to Carole, but he also signs another transaction to transfer the same amount to Bernard. These two transactions are incompatible, but to realise this, Carole and Bernard must both be aware of the transaction that the other one received.

[2] Here, we have chosen the term "doubloon" because it is neutral and has a certain old-fashioned charm. However, the reader may just as easily replace it with bitcoin, euro, etc.

Historically, under the gold standard, this kind of problem never arose. The laws of physics governed accounts: nothing is lost, nothing is created. For fiduciary moneys, the account mechanism is typically a hierarchy. At the top, a central bank holds an account ledger for its clients (the banking institutions). In the case of decentralised currencies, shared knowledge of all transactions by all participants is what guarantees that the accounts are kept properly.

From a technical standpoint, all participants must agree on the order of transactions. After Carole has accepted Oscar's transaction, the system must reject the conflicting payment that Oscar attempts to make to Bernard. This therefore assumes that all participants acknowledge that one of these transactions was made before the other. The order chosen is not, in and of itself, very important, but it must be indisputable.

In distributed computing theory, this problem is known as the "consensus problem". It models a set of processes that should arrive at an agreement, in a finite period of time, as to the contents of a log of transactions. One particularly challenging variant of the consensus problem entails creating consensus protocols in cases involving unreliable participants known as "Byzantine generals". These Byzantine generals act as they please, without necessarily following the rules of the protocol. They can also corrupt the network by slowing down the circulation of messages. Therefore, the honest participants must reach a consensus despite the presence of these dishonest participants. The problem was described for the first time in an article entitled "The Byzantine Generals Problem",[3] which, taking a more general perspective, demonstrates that the problem is solvable if – and only if – less than one-third of participants are Byzantine generals.

This approach allows account-keeping duties to be distributed, but it is based on selecting an unchanging set of participants. It is therefore not suited to a large-scale decentralised network, which by definition must be open to everyone. Opening up the network is particularly problematic when there are Byzantine generals. Indeed, in an open, anonymous network, it is easy for an attacker to pretend to be a number of different identities and to use this trick to thwart the consensus. This is known as a "Sybil attack".

[3] L. Lamport, R. Shostak and M. Pease (1982), "The Byzantine Generals Problem", *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, July.

Proof of work

In 2008, Bitcoin[4] proposed an unconventional approach. Participation in the consensus was based not on the notion of identity, but by proving consumption of computing power. The technique, known as "proof of work", was originally introduced by cryptographer Adam Back[5] to curb spam emails. It is based on the principle of partial cryptographic hash inversions. Participation in the Bitcoin consensus mechanism is therefore not measured in terms of distinct "entities", but instead in terms of computing power. This approach not only safeguards against Sybil attacks, but also lends itself to an economic mechanism that both rewards honest participation in the protocol and punishes Byzantine behaviour. The network therefore tolerates fairly amoral participants who, in pursuing their own personal interests, contribute to network security by "mining" new blocks created by proof of work. Aside from its security-related properties, the proof-of-work approach enables the initial, anonymous and unbiased distribution of bitcoins to these "miners" for each block they create.

What's important to understand is that the computations carried out for the proof of work are not intrinsically useful. They do not determine the validity of transactions or update a database. Proof of work is only used to prove that actual resources – in this case, energy – have been irrevocably used.

Despite its obvious advantages, the proof-of-work approach is not immune to criticism. Cryptographer Ben Laurie has stated that, to ensure security, proof of work must account for half the total computing power worldwide.[6] This argument is taken to the extreme, but it is true that over the past few years, the power dedicated to ensuring that the network operates properly has taken on considerable proportions of several gigawatts. In addition, the arguments in favour of proof of work – its decentralised aspect, in particular – have been challenged both by actual practice and by game theory analyses.

For instance, an attack can consist of going backwards, in other words, rewriting the history of the blockchain. While the content of blocks cannot be modified, it is possible to pretend that these blocks were never produced or that other blocks were produced instead. Typically, such an attack can be successful only if more than 51% of miners participate because this "alternative" chain must be longer than the original chain in order to be considered legitimate. It is costly to take part in an attack that ultimately fails, but a strategically skilled attacker could bring together other miners by giving them a sort of "insurance policy" in case the attack fails, and promising them a small reward if it succeeds. In a model in which agents are amoral and blindly profit-seeking, the insurance guarantees that the attack will be successful and therefore costs the attacker nothing.

---

[4] S. Nakamoto (2008), *Bitcoin: A Peer-to-Peer Electronic Cash System*.

[5] A. Back (2002), "Hashcash - A Denial of Service Counter-Measure", *Technical Report*, August.

[6] B. Laurie, *Decentralised Currencies are Probably Impossible (But Let's at Least Make Them Efficient)*.

This example is not intended to suggest that the network is not viable, but to show that network security depends much more on participants' honesty than is sometimes implied.

Proof of stake

Another consensus approach – known as "proof of stake" – is actually older than Bitcoin.[7] It is currently gaining influence. The idea is to use the currency itself as a mechanism for thwarting Sybil attacks. Consensus participation is no longer based on the computational power consumed, but instead on the amount of currency held (the stake).

The flaw in this approach is that it is circular. The consensus must be secure in order to determine the rights to participate in the consensus. This circularity cannot be completely avoided and therefore the proof-of-stake approach cannot replicate all the security properties of the proof-of-work mechanism. This can be seen from another angle through a simple simulation.[8]

Let's assume, from a very general perspective, that blocks are not costly to create. This is one of the goals of proof of stake. In this case, nothing prevents dishonest participants from creating two parallel chains: one public, one secret. These participants can then, at any time, sell the currency they hold on the public chain and, in parallel, publish their secret chain. A new participant who discovers the system will see two chains: the authentic one and a fake one. There is no intrinsic property of the authentic chain that would enable it to be distinguished from the fake one. The duplicity of the dishonest participants can be revealed by comparing the two chains, raising the question of whether they should be punished. But how could they be punished, as they have nothing at stake (this problem is actually referred to as "nothing-at-stake").

This argument is correct, but although it is often presented as an insurmountable obstacle to proof of stake, it is not necessarily relevant. To begin with, most proof-of-stake approaches automatically freeze the funds of block-producing participants. If these funds are frozen for a month, for example, this means that during that lapse of time, it must be possible to guarantee that no "fake" chain published on the network diverges from the authentic chain. Otherwise, the dishonest participants who created this fake chain can be punished by destroying their funds. The security criterion therefore becomes the following: consensus participants must connect to the network at least once a month, while new entrants must verify the recent state of the blockchain. To do so, they can query merchants that accept the currency in question. Let's not forget that acceptance of a currency also reflects a human, social consensus. Blockchains are no exception to this rule of "weak subjectivity", regardless of whether they use proof of work or proof of stake.

---

[7] W. Dai, *B Money*, http://www.weidai.com/bmoney.txt

[8] A. Poelstra (2015), *On Stake and Consensus*.

Proof of stake offers unique features. In particular, it provides asymmetry for participants: honest creation of blocks is not very costly, whereas malicious failures to follow the protocol can be met with severe punishment. Thus, proof of stake reintroduces – to use Vitalik Buterin's terms – a form of asymmetry that is typical of cryptography and the cypherpunk movement, in which an attack is much more onerous than defence.[9] Given its security features and low cost, proof of stake is a different option for designing a cryptocurrency.

Following on from the success of Bitcoin and proof of work, proof of stake is currently enjoying fresh interest. Some projects take active ownership of the "subjective" features of proof of stake, for example, Tezos,[10] which also attempts to overcome the inherent governance challenges of proof-of-work models. (Author disclosure: I am personally very involved in the Tezos project.)

Regarded just a few years ago as an impossibility, proof of stake is now central to a new generation of projects, including: Tendermint, which is based on classic Byzantine fault-tolerant algos; Polkadot, which is pushing the boundaries of distributed computing by combining liveness and security; and Algorand, a blockchain designed by well-known cryptographer Silvio Micali, a recipient of both the Gödel Prize and the Turing Award.

Conclusion

Studying and designing cryptocurrencies is inevitably a multidisciplinary affair. The lion's share of this research involves the principles of distributed computing and cryptography, but it also calls on game theory, political and financial economics, and sociology. The ideological roots of cryptocurrencies are undeniable and can sometimes come into open conflict with certain governmental stakeholders, as was the case for the Internet in the 1990s. However, the well-informed will realise that this is an unavoidable disruptive innovation: a powerful idea whose time has come.

[9] V. Buterin (2016), *A Proof of Stake Design Philosophy*.
[10] L.M. Goodman (2014), *Tezos: A Self-Amending Crypto-Ledger*.