Blockchains in 2016: status quo and scaling challenges

JP Aumasson, Kudelski Security
Philipp Jovanovic, Decentralized and Distributed Systems lab, EPFL

In this report we review some of the latest developments in blockchain technologies, and focus on a new technique developed at EPFL to scale blockchains and to address critical issues with Bitcoin.

Blockchain 101

We'll start with a recap of the basics of blockchain and cryptocurrencies. You can find many online resources explaining these concepts in more detail, for example on https://bitsonblocks.net/ or https://bitcoin.it/.

At its core, a blockchain is a peculiar kind of database, namely a digital ledger. In summary, a blockchain is

1. **A chain**, or a sequential list of blocks of data, where blocks represent transactions added to the database. To encode many transactions in a single block, most blockchains use Merkle trees to represent the transactions of many fingerprints as a single one.
2. **Immutable, append-only,** meaning you can add new blocks to the chain but you cannot remove or modify previous blocks. That is, you can only add items to a blockchain—except in private blockchains that authorize it.
3. **Distributed** among multiple machines rather than stored on a single machine as a traditional database. The computing effort necessary to include transactions is also distributed among a set of machines, the miners.
4. **Decentralized,** as blockchains are peer-to-peer systems, as opposed to classical server-clients systems. Blockchains remove single points of failure, and a blockchain-based system will work as long as a majority of participants behave.
5. **Secured,** as blockchains use cryptography to ensure that items being added cannot be modified. For example, transactions are cryptographically signed, and cryptographic puzzles (for blockchains that use one) help to prevent double spending.

Thanks to decentralization, blockchains solve the problem of trusted third-parties: Instead of having to trust a single entity such as a bank, payment network, or governmental organization, a blockchain decentralizes trust to a multitude of parties.

Cryptocurrencies

Bitcoin[1] and most cryptocurrencies work as follows:

- Accounts are called **wallets**, and are not bound to any person or organization as they are in traditional banking. That is what makes Bitcoin pseudonymous, but unfortunately, that's what also makes Bitcoin easy to steal. A wallet identifier is for example `14T2Mqc6TQ81nKks1ZUu7PjvJwtP4zXkZ9`. You can manage your own wallet yourself or delegate the burden to an online wallet service.
- Blocks within Bitcoin's blockchain encode **transactions**, or statements that wallet X transferred an amount of Y bitcoin to wallet Z. The Bitcoin blockchain is therefore the whole history of who paid whom.
- To **issue a payment** transaction, it is sent first to the Bitcoin network for validation. Computers verify the cryptographic signature created by the payer, a proof that the payer does agree to transfer money to a given wallet. Once validated, the network runs a consensus protocol to actually add the transaction to the blockchain.
- In order to prevent **double spending**, or use of the same bitcoins twice, Bitcoin makes it hard to create new chains of blocks. Hardness is guaranteed through a **proof-of-work**, or a computational puzzle that takes a large number of operations to solve prior to integrating a new transaction block into the blockchain.
- The process of solving proofs-of-work is called **mining**, and computers or parties carrying out that task are called **miners**. For each mined block, miners are rewarded with a certain number of bitcoins.

[1] *When capitalized "Bitcoin" refers to the whole system whereas lowercase "bitcoin" refers to one unit of the currency.*

- Bitcoins can be traded on **exchange platforms**, such as Coinbase or Kraken. Exchange platforms let you buy and sell bitcoin in one of the main currencies, such as USD or EUR, and will charge you per-transaction fees as classical trading platforms.

The Bitcoin network validates a new block every ten minutes, where blocks are, at most, one megabyte long. On average, a block encodes more than a thousand transactions at the time of writing. Bitcoin supports up to seven transactions per second, or approximately 4200 in a single block.

There are now many cryptocurrencies other than Bitcoin, such as Litecoin, Monero, Dash, or Dogecoin, with slightly different properties than the Bitcoin system. For example, Litecoin uses a different proof-of-work mechanism than Bitcoin in order to facilitate mining on off-the-shelf computers, thanks to the hardware-unfriendly scrypt algorithm. But many **altcoins** only address niche use cases or are only used for speculative purposes.

Decentralized Applications

The idea of using a blockchain to track transactions can be generalized to transactions other than "wallet X pays Y bitcoins to wallet Z". This idea is at the core of Ethereum (https://www.ethereum.org/), launched in July 2015. With a market capitalization of about $900 million (ten times less than Bitcoin), Ethereum is by far the most important blockchain application after Bitcoin.

Like Bitcoin, Ethereum maintains a blockchain wherein new blocks are added through a proof-of-work and where computers are rewarded for solving **computation-** and **memory-expensive** operations. Rewards come in a unit called **Ether**. Like bitcoin, Ether can be bought and sold against normal currencies—at the time of writing, one Ether is worth $11.

The big difference from Bitcoin is that Ethereum transactions are not restricted only to payment contracts, but can include arbitrary operations under arbitrary conditions, which are then called **smart contracts**. For example, you can create a **bet application,** wherein two parties place a bet based on the outcome of some event, and after the event the winner receives their due amount. Only the rules of the bet are committed to the blockchain, and the actual bets and rewards are created by interacting with the blockchain, just as one would interact with a computer program. More generally, Ethereum's blockchain can be used to host software applications implementing any business logic. Such **decentralized applications** allow their clients to interact with each other without resorting to a trusted third party, thereby eliminating a single point of failure. Individuals and organizations could even create their own cryptocurrency (or token of value) running on top of Ethereum's blockchain.

More technical details on Ethereum can be found in the corresponding whitepaper[2].

Challenges with Blockchains

Operational Risks

The greatest risk with Bitcoin—aside from financial risks—is the loss or theft of an account. Lose your 256-bit key, and there is no way to recover the bitcoins in your account. For example, in 2013 a British man threw out a hard drive disk containing the key to unlock his 7500 bitcoins, worth $7.5 million at the time, which are (alas for him) now lost forever.

Online Bitcoin trading services also offer wallet management services—they store your private key for you, just as a bank would store your cash. This is a way to transfer the risk from the individual to a Bitcoin company. But hacking a remote computer and stealing a 256-bit value is easier (and usually less risky for the thieves) than robbing a bank and running away with bags of cash. For example, in July 2011, the MyBitcoin wallet platform lost 51% of the funds it held on behalf of its members, and in February 2014 the Bitcoin exchange platform Mt Gox got robbed of the equivalent of $350M worth of members' assets. A 2016 study even claims that between 2009 and 2015, 33% of Bitcoin exchanges had been hacked.

Likewise, you can lose the key to your Ethereum account. But Ethereum's complexity introduces additional risks: If a software application is insecure, its decentralized version will be insecure too. Worse, attackers may exploit vulnerabilities in Ethereum itself and compromise a part of the network's assets. For example, in June 2016 an unknown attacker exploited a flaw in the design of The Decentralized Autonomous Organization (The DAO), an Ethereum application that acts as a venture capital fund. The attacker managed to transfer $60M worth of Ether from The DAO investors to his own accounts.

Anonymity

While Bitcoin accounts (wallets), unlike traditional bank accounts, are anonymous, transactions (who paid whom, and how much) are public, so that anyone can observe what payments are made. This makes transactions publicly verifiable, and prevents fraud. However, anyone can **trace** the transactions of a given account and observe who's doing business with whom. To avoid this limitation, researchers created a so-called **privacy-preserving** version of Bitcoin called Zcash (https://z.cash), launched in fall 2016.

Zcash makes Bitcoin transactions anonymous, and therefore untraceable, by hiding the identifier of the payer. Zcash does so by using **zero-knowledge proofs**, cryptographic mechanisms that allow you to prove the knowledge of some information without revealing it. In

[2] https://github.com/ethereum/wiki/wiki/White-Paper

Zcash, this information is the ownership of **some** Zcash coins. Details are available in the original Zerocash paper[3].

In terms of engineering, Zcash is a major piece of work. When it launches, it will likely attract attention of miners (who hope to make a good investment) and from vendors, including dark markets (who hope to benefit from its untraceability).

Consensus

One of the main distinguishing features between different blockchain systems is the way consensus is achieved in the network. This includes which peers belong to the so-called **consensus group**, the set of nodes that is allowed to make decisions, and how new candidates gain access to that group. There are two approaches:

- **Permissioned** or **federated** blockchain systems have a **closed** consensus group. This means that candidates who wish to participate in the decision-making process have to request access and proof that they meet certain (predefined) requirements. Only if a quorum of the consensus group members approves the application, the contender gains access and is allowed to take part in the consensus mechanism. Hyperledger, Ripple, and Stellar are representatives of this category.
- **Permissionless** blockchain systems allow basically anyone to participate in the consensus mechanism. Their consensus group is therefore called **open**. The right to make decisions (i.e. append blocks to the blockchain) is granted by solving (difficult) computational puzzles and presenting a valid solution, the previously mentioned **proof-of-work**, to the peers of the network. Bitcoin, Ethereum, and Zcash are representatives of this category.

Mining

Peer-to-peer networks generally face the problem of so-called **Sybil attacks**, in which adversaries create large amounts of pseudonymous identities to obtain disproportionately large influence. While the very nature of the closed consensus group in permissioned blockchains prevents this kind of malicious behaviour, permissionless systems with their open membership need a mechanism to protect themselves against the above attacks.

Permissionless blockchains use **mining** for protection, which usually refers to the process of solving (hard) computational puzzles. **Proof-of-work** (PoW) is the outcome of a successful mining process and, although is hard to create, is easy to verify.

Proof-of-Work

Bitcoin's approach to mining is the twofold evaluation of the SHA-256 cryptographic hash function on a newly formed block's header and checking if the hash has a certain number of leading zeroes. If not, then a counter included in the block is increased and the above procedure is repeated. If the hash however has the above shape, then the combination of hash and block forms a valid proof-of-work. However, this particular approach has its drawbacks, e.g. consumption of massive amounts of energy for large-scale mining.

Different blockchain systems use different proof-of-work algorithms, with different requirements. For example, Ethereum uses **Ethash**[4], a memory-hard hash that needs a 1GB table to compute the hash (a technique that has been called "ROM-hard", though in this case the table is stored in RAM), but verifying the hash needs only minimal memory. This low-memory verification is the main difference with password hash functions such as Argon2 or scrypt.

Zcash's proof-of-work is **Equihash**[5], an algorithm by the designers of Argon2. Like Ethash, Equihash is hard to compute but easy to verify—such functions can be seen as NP-hard languages, and are sometimes called **asymmetric** PoWs. Equihash it totally different from Ethash though. Computing an Equihash is equivalent to solving a variant of the **generalized birthday problem**, namely finding N values $x_1, x_2, .., x_N$ and a nonce V such that

$$H(I \,||\, V \,||\, x_1) \oplus H(I \,||\, V \,||\, x_2) \oplus \ ... \ \oplus H(I \,||\, V \,||\, x_N) = 0$$

where H is a hash function, and with some additional difficulty conditions.

Proof-of-Stake

An alternative to proof-of-work is **proof-of-stake** (PoS), also called **virtual mining**. Proof-of-stake is a generalization of proof-of-work: Instead of trading compute resources (the mining effort) for coins (the mining reward), proof-of-stake trades other assets than computation, or accepts risks. For example, **proof-of-deposit** requires miners to prove that they moved some of their coins into a bond deposit locked for a certain period of time, while **proof-of-burn** requires miner to pay a certain amount of cryptocurrency, by burning it, that is, transferring it to an unspendable wallet, equivalent to `/dev/null`.

[3] http://zerocash-project.org/media/pdf/zerocash-extended-20140518.pdf

[4] https://github.com/ethereum/wiki/wiki/Ethash
https://github.com/ethereum/wiki/wiki/Ethash-Design-Rationale
https://github.com/LeastAuthority/ethereum-analyses/blob/master/PoW.md

[5] https://www.internetsociety.org/sites/default/files/blogs-media/equihash-asymmetric-proof-of-work-based-generalized-birthday-problem.pdf

Connecting Payment Networks

Blockchains are often seen as a cheaper book-keeping solution than existing financial networks. Blockchain technology has the potential to dramatically **cut operational costs**, at least in principle. For example, a 2015 white paper sponsored by Santander Bank argued that "The first major application [of distributed ledgers] is being seen in payments, especially across borders," and that smart contracts "will lead to a wide variety of potential uses in securities, syndicated lending, trade finance, swaps, derivatives or wherever counterparty risk arises."

The most advanced blockchain-based cross-border payment technology is **Ripple**, a protocol and software that enables banks "to directly transact with each other without the need for a central counterparty or correspondent" (https://ripple.com/). After Bitcoin and Ethereum, Ripple has the third largest market capitalization among blockchain-based technologies (more than $200M at the time of writing). Banks experimenting with Ripple include Santander, UBS, UniCredit, ReiseBank, CIBC, National Bank of Abu Dhabi (NBAD), and ATB Financial.

Worth noting as well is **R3** (http://r3cev.com/), a consortium of financial institutions that started in 2015 with Barclays, BBVA, Commonwealth Bank of Australia, Credit Suisse, Goldman Sachs, J.P. Morgan, Royal Bank of Scotland, State Street, and UBS. Other institutions later joined to support R3 in the creation of a distributed ledger technology matching the needs of its members' financial services.

To be useful, blockchain-based payment networks need to be able to interact with each other, and to support payments with non-blockchain networks that keep a ledger of transactions. Here ledger is understood in a broad sense, meaning a system that tracks accounts and balances. Ripple created the **Interledger**[6] protocol (https://interledger.org/). Interledger defines a high-level, ledger-agnostic protocol to enable cross-ledger payments, while ensuring correctness and preventing fraud. Interledger is integrated in Ripple, but is still considered experimental and not yet used by major banks or networks.

Focus: A Solution to Scaling Blockchains

Bitcoin faces some fundamental security and performance hurdles that prevent large-scale adoption: In comparison to classic payment providers, such as VISA or PayPal, which are able to perform hundreds to thousands of transactions (tx) per second with very low transaction confirmation latencies, Bitcoin's current transaction throughput is limited to about 7 tx/sec and its consensus mechanism requires users to wait tens of minutes for transaction commitment, and even then, only provides probabilistic consistency guarantees. This means that inconsistencies (**forks**) might occur when different miners find new blocks independently and at about the same time, splitting the peers' views on the blockchain. Fork-resolution regularly destroys large numbers of transactions, sometimes even hours after their initial submission, thereby wasting all the computational power spent on the orphaned branch. As a consequence, Bitcoin's peer-to-peer network establishes a consistent view on the distributed ledger only eventually. The probabilistic consistency of Bitcoin's consensus mechanism is also one of the reasons that the cryptocurrency is susceptible to all kinds of attacks, such as selfish mining and double spending. One way to mitigate many of these problems, is to eliminate Bitcoin's sloppy fork-resolution mechanism and adopt **strong consistency**, a more proactive approach that offers the following important benefits:

- All miners agree on the validity of the blocks right away, without wasting computational power to resolve forks.
- Clients don't need to wait for extended periods to be certain that a submitted transaction is committed; as soon as it appears in the blockchain, the transaction can be considered confirmed.
- Once a block has been appended to the blockchain, it stays there forever (as long as there is an honest majority of miners). This property is also often referred to as **forward security**.

The work described below shows how to implement strong consistency in Bitcoin by introducing **ByzCoin**, a novel Byzantine consensus protocol. Results from experimental evaluation indicate that ByzCoin-improved Bitcoin can increase its throughput by two orders of magnitude.

This research piece on scaling blockchains originally appeared in the Usenix Security 2016 paper Enhancing Bitcoin Security and Performance with Strong Consistency via Collective Signing[7]*, by Kokoris Kogias, Jovanovic, Gailly, Khoffi, Gasser, and Ford (EPFL, Switzerland).*

ByzCoin

ByzCoin is a novel scalable Byzantine fault-tolerant (BFT) consensus protocol that provides strong consistency, while scaling to processing throughputs of hundreds of transactions per second, among hundreds to thousands of decentralized miners. ByzCoin uses an adapted version of the well-studied Practical Byzantine Fault Tolerance (PBFT) algorithm and introduces four key improvements over Bitcoin:

- ByzCoin's improved PBFT-like consensus mechanism commits Bitcoin transactions irreversibly within seconds.

[6] https://interledger.org/interledger.pdf

[7] https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_kokoris-kogias.pdf

- ByzCoin preserves Bitcoin's open-membership property by dynamically forming hash power-proportionate consensus groups that represent recently successful block miners.
- ByzCoin uses communication trees to further optimize transaction commitments and verification under normal operations, while guaranteeing safety and liveness under Byzantine faults.
- ByzCoin decouples the election of a new leader from transaction verification, an approach inspired by Bitcoin-NG, that enables ByzCoin's transaction throughput to further increase.

Together, all these optimizations enable ByzCoin to achieve throughputs higher than PayPal currently handles, and to provide low confirmation latencies. Another benefit of ByzCoin's fast transaction commitment, ranging from a few seconds, up to at most one or two minutes after submission, is the mitigation of double-spending and selfish mining attacks.

Design

An overview on ByzCoin's design is in Figure 1 below:

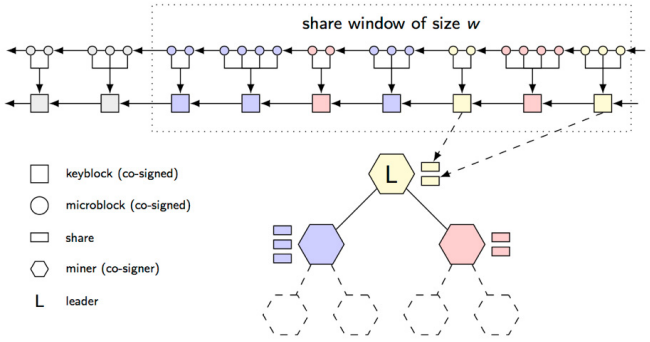

Figure 1: ByzCoin overview

The lower part of Figure 1 shows ByzCoin's consensus group that is comprised of recently successful block miners and uses a PBFT-like mechanism to reach consensus. Instead of PBFT's MAC-based authentication, which has quadratic communication and computation complexity, ByzCoin uses **CoSi** (https://github.com/dedis/cosi), a distributed protocol that utilizes aggregated Schnorr signatures and tree-based communication to make large-scale, decentralized, collective signing practical. CoSi reduces communication complexity from quadratic to logarithmic, enables third-party verifiability, and signature verification in constant-time complexity.

Another feature that ByzCoin adopts from PBFT is the role of the consensus group leader whose task is to bundle transactions into blocks and initiate new signing rounds. All actions taken by the leader have to be approved by a two-thirds supermajority of the consensus group members, which effectively leads to strong consistency and all of its benefits discussed earlier. In case the leader misbehaves, the miners in ByzCoin's consensus group can start a voting round and dismiss the Byzantine node but only if, again, a two-thirds supermajority approves. The requirement of the two-thirds supermajority for decision making comes from Byzantine agreement theory that permits at most f malicious/faulty nodes among a total of 3f + 1 nodes.

Another important part is the leader election mechanism, which brings us to the next component of ByzCoin's design, its blockchain (upper part of Figure 1). ByzCoin's blockchain is divided into two interdependent sub-chains: one for **keyblocks** and one for **microblocks**. We present these two types of block below, and refer to the research paper for further details.

Keyblocks

Keyblocks are used to manage ByzCoin's consensus group membership. These blocks are generated by the miners through proof-of-work roughly every 10 minutes, as in Bitcoin, and are collectively signed by ByzCoin's consensus group. Keyblocks form a regular blockchain. A miner who successfully mines a new keyblock is rewarded with a consensus group share, a so-called **proof-of-membership**, thereby gaining entry into the consensus group if they are not already a member, and becomes the next group leader.

A fixed-size sliding window mechanism constitutes the total number of available shares: Any share beyond the current window expires and miners who no longer hold any valid shares drop out of the consensus group. The number of valid shares in the possession of a miner reflects his voting power within the consensus group, when committing transactions. Moreover, this number determines the portion of coins a miner receives as a reward, when a new keyblock is found. In other words, ByzCoin rewards not only the node that mines a new keyblock but instead splits, proportionate to the valid shares each miner holds, the produced coins among all miners of the consensus group.

ByzCoin also uses that technique to split transaction costs as a reward, once no more coins can be mined. The proof-of-membership approach ensures liveness, as dormant miners are removed from the consensus group and the share-proportionate rewards further incite all miners to remain active and contribute to the progress of the system.

Microblocks

Microblocks, unlike keyblocks, contain transactions. They are proposed by the current leader and do not require proof-of-work, and therefore are committed much more frequently by the consensus group.

Each microblock contains, in addition to the list of transactions, a hash of the last microblock. This ensures total ordering, as well as a hash of the leader's keyblock to identify the era the microblock belongs to. Even though microblocks are created by the consensus group leader, ByzCoin's witness-mechanism deters leaders from misbehaving (such as mounting double-spend attacks),

because any misconduct would be immediately detected by the other group members, which in turn can trigger a view change (thereby removing the malicious node).

Experimental Results

The authors wrote a prototype of ByzCoin, available on GitHub as part of EPFL's cothority project (https://github.com/dedis/cothority), and conducted thorough experiments, measuring transaction confirmation-latency and throughput. They experimented with consensus group sizes between 144 and 1008 nodes, which corresponds to a window of successful keyblock miners ranging from the previous days up to the previous weeks. Figure 2 shows ByzCoin's throughput in comparison to other systems. The data for the simulations is based on actual transactions from a portion of the Bitcoin blockchain.

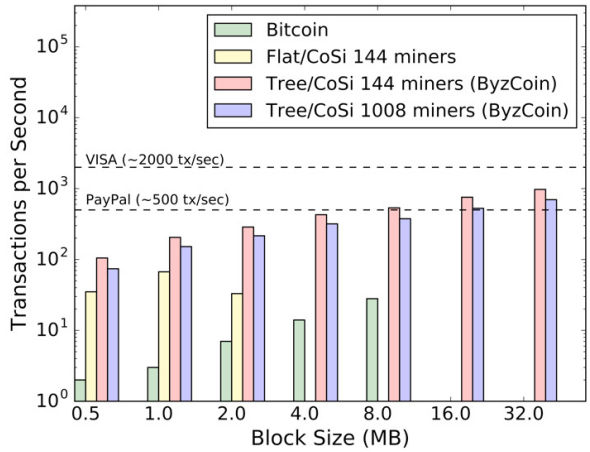

Figure 2: ByzCoin's throughput

The authors report an average latency around 90 seconds for 32 MB blocks (~66000 tx) and a consensus group size of 144 members, as an example. For this configuration, ByzCoin's throughput (~700 tx/sec) outperforms PayPal's. For a more elaborate discussion of ByzCoin's performance evaluation we refer you to the research paper.