Blockchain
Hype or Hope?

RADIA PERLMAN

Radia Perlman's work has had a profound impact on how computer networks function today, enabling huge networks, like the Internet, to be robust, scalable, and largely self-managing. She has also made important contributions in network security, assured delete, key management for data at rest encryption, DDoS defense, and user authentication. She is currently a Fellow at Dell EMC, and has taught as adjunct faculty at University of Washington, Harvard, and MIT. She wrote the textbook *Interconnections* and co-wrote the textbook *Network Security*. She holds over 100 issued patents and has received numerous awards, including induction into the Inventor Hall of Fame, lifetime achievement awards from ACM's SIGCOMM and USENIX, election to the National Academy of Engineering, induction into the Internet Hall of Fame, and an honorary doctorate from KTH. She has a PhD in computer science from MIT.
radia@alum.mit.edu

In this article, I describe the technology behind Bitcoin's blockchain, and its scalability, security, and robustness. Most of what is written about "blockchain technology" talks about how it will revolutionize all sorts of applications without contrasting it with alternative solutions. To complicate matters, there are all sorts of proposed variants of the original blockchain (the technology behind Bitcoin), making the definition of "blockchain technology" very unclear. I explain how Bitcoin's blockchain technology works, along with its performance implications.

A lot has been written about "blockchain technology" recently, but most of it talks about how it "is being investigated" for various applications and how it is a revolution in computing that will change the world [1]. It is not that easy to discover, from these sorts of articles, how the technology works or what its true properties are. These articles treat "blockchain" as a sort of black box that stores and retrieves data, with certain properties:

- Append-only log
- "Immutable"
- No central point of control

The term *blockchain* was introduced as the name of the technology that powers Bitcoin. Given that Bitcoin's technology is widely deployed and unlikely to change very dramatically, it is possible to describe how it works and what its scalability, robustness, and security properties are. It is not clear how much this system can be modified and still be called *blockchain technology*. Therefore, with the term blockchain technology being less and less well-defined, I will not attempt to describe the properties of every variant proposed, and for the rest of this article, when I say "blockchain," I am referring to Bitcoin's blockchain.

Description of Blockchain

In this section I'll give an overview of the Bitcoin blockchain technology.

Bitcoin

Bitcoin was introduced to the world in a 2008 article [2] and, shortly thereafter, was released as open source software. The concepts are described in the paper, but the details are defined by the implementation. The open source community in control of the software may make changes, but the more widely deployed it is, the more difficult it is to make incompatible changes.

The design goal of Bitcoin was to create a currency that could not be controlled by any government or any known organizations. This design is intended to foil the ability of governments to do things like:

- Enforce tax laws
- Follow a money trail
- Prohibit payments to certain countries or organizations
- Inhibit criminals from anonymously collecting ransom money

These may or may not be desirable goals for a currency, but I will examine the performance implications of a design with these goals, and whether applications other than cryptocurrency really benefit from a design without known entities at the center.

The basic concepts behind blockchain:

- A large (thousands) community of anonymous entities called "miners" collectively agree upon the history of transactions, in an append-only data structure known as "the ledger."
- Users of Bitcoin are not identified with names, but rather, with public keys, and a user is allowed (even encouraged) to change public keys often, to make transactions more anonymous.
- The ledger contains a list of every Bitcoin transaction since Bitcoin was invented.
- A transaction records that a public key X pays a certain amount of Bitcoin to public key Y.
- In order to add transactions to the ledger, a miner must validate the transactions and compute a valid block containing them.
- A valid block contains a hash of the previous block in the blockchain, a set of new valid transactions, and a random number chosen so that the hash of the block meets certain conditions. A valid block is, by design, just hard enough to compute that the collective compute power of the miner community will find a new block at some cadence (about every 10 minutes).
- The miner who is lucky enough to be the first to find the next valid block is awarded with some amount of Bitcoin.

Now I will describe these steps in more detail.

Format of the Ledger

Each block in the blockchain contains the hash of the previous block, a nonce (a random number), the public key of the lucky miner who was the first to find a valid next block, and valid transactions that have not yet been recorded in the ledger (Figure 1).

Transactions

The information in transactions looks like this:

A transaction (with hash T1) consists of the payer (public key X) signing away all of the Bitcoins that X had been paid in some previous transaction (with hash T2).

In order for the transaction T1 to be valid,

- There must be a prior transaction with hash T2, in which X was the payee of the amount of Bitcoin being paid in transaction T1.
- The signature on T1 must properly validate, using public key X.
- There must be no other transaction in the ledger in which X has already spent the proceeds of T2.

There are extra details. For example, notice in the third line of Figure 2 (the transaction with hash x17), A is signing over to C the results of the transaction with hash x15, in which X received 74.92 Bitcoins. But A is only paying 74.21 in transaction x17, even though in transaction x15, A had received 74.92. The difference (74.92 – 74.21) is a transaction fee, paid to the miner who adds a block to the blockchain that contains transaction x17. This rewards the miner for including this transaction in the new block.

The Hash

The mining community imposes conditions on the hash of a valid block. These conditions are designed to be just difficult enough to meet, that it will take the community about 10 minutes to find a block with the appropriate hash.

A good cryptographic hash is like a random number. Given random input, it should have probability 0.5 that the first bit in the hash will be 0, or probability 0.25 that the first two bits would both be 0. The method that blockchain uses to adjust the difficulty of computing the hash is to have a maximum value that the hash must have. Currently, the maximum value of the hash has about 70 leading zeroes. That means that for any block, the probability of its hash having 70 leading 0s is $1/(2^{70})$. Using a brute force search, and the collective compute power of the mining community, it takes about 10 minutes for at least one miner to find a block with a small enough hash. If blocks are found too quickly, then the maximum hash value is adjusted to be smaller. If blocks are found too slowly, then the maximum hash value is adjusted to be larger.

Traditional Integrity Checks vs. Blockchain Hash

Traditional public key cryptography creates digital signatures that can be efficiently computed, if and only if the signer knows

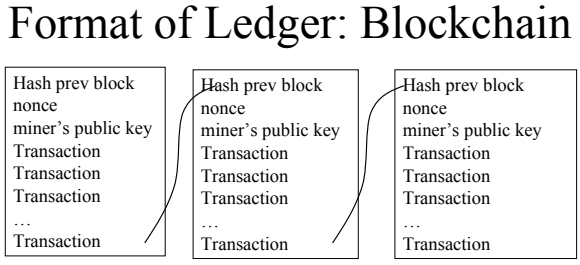

Figure 1

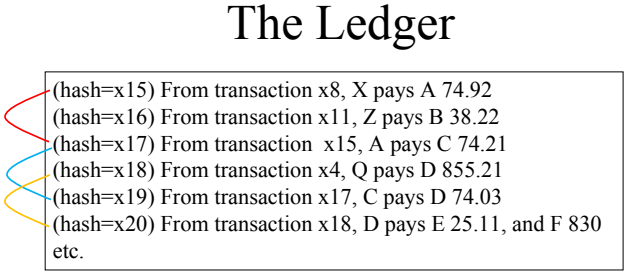

Figure 2

a secret known as the *private key*. The signature can be verified by anyone with knowledge of the associated *public key*. And an essential component of any public key system is that there will be some way of making sure that the public key is well-known.

With a traditional public key system, the cryptography ensures that there is an enormous gap between the computation needed for someone with knowledge of the private key to generate a signature, and someone without knowledge of the key to forge a signature. With RSA, the computation necessary to generate a signature (knowing the private key) is a small power of the length of the key (between 2 and 3). In contrast, brute force breaking of a key is almost exponential in the length of the key. So, for instance, for a 1024 bit RSA key, it is about 2^{63} times more expensive to forge a signature than to generate one. Increasing the key size increases the gap between forging and generating signatures. If an RSA key were increased from 1024 bits to 2048 bits, the gap becomes about 2^{94} times more expensive to forge rather than generate a signature.

Since it's hard to imagine these huge numbers, another way to say it is that signing with RSA 1024 takes about a millisecond on a typical CPU, and signing with RSA 2048 might take 6 milliseconds on the same CPU. However, breaking RSA 1024 takes about as much computation as all the Bitcoin miners do in an hour. Breaking RSA 2048 takes about as much computation as all the Bitcoin miners would do if they continued at the present rate for a million years.

The startling aspect of the Bitcoin hash is that it is *equally* difficult for the community of miners to compute a hash as for someone to forge a hash. This means that the security of Bitcoin depends on the assumption that no entity or collection of entities can amass as much compute power as the Bitcoin mining community. This is a very surprising assumption. It would indeed be easy for a nation-state to amass more compute power than the Bitcoin community.

What could a malicious set of miners, with more compute power than the honest Bitcoin miners do? They could discriminate against certain transactions, refusing to ever record them in the ledger. They could compute an alternate ledger, where transactions they had previously spent were not recorded anymore, and then they could double-spend.

And not only is the security assumption highly questionable, since it is hard to believe that the community of honest miners has cornered the market on all computation power on the planet, but it means that the computation required by the honest miners is mind-bogglingly huge.

What Would Motivate Someone to Be a Miner?

Miners have to do a lot of computation if they ever hope to be rewarded with any Bitcoins. Currently, the miner community earns about 2 million US dollars every day. And reports are that this barely covers the amount they are spending on electricity. That amount of electricity is estimated to be equal to what a nuclear power plant generates per day, about 500 megawatts.

So any application of this technology must somehow generate revenue for the miners.

Other Costs

It is also necessary to store the entire ledger so that transactions can be checked for validity. Currently, the ledger is about 100 GB and is stored in thousands of places around the network. Also, there is a huge amount of network bandwidth to broadcast transactions and new blocks to all the Bitcoin nodes, as well as to be able to download the entire ledger to any node that is joining the community.

What Is Novel about Blockchain?

If "blockchain" is truly a revolution in computing, there must be something about it that did not exist before. What could it be?

Is It Having a "Ledger"?

Blockchain's "ledger" is an append-only log that needs to be kept in its entirety, and needs to be world-readable and world-writable. Very few applications really want these properties. Much more flexible databases have of course existed for a long time.

Is It Replicating the Data?

Blockchain highly replicates the ledger so that it will not easily get lost. Obviously, the more locations in which something is stored, the less likely it is that it will become permanently lost. Large public clouds tend to store data in perhaps six places, carefully chosen to be located in different locations so that a natural disaster in one location will not wipe out all copies of the data. If any copy is lost, the public cloud quickly replicates the data to new locations to replace the ones that have lost the data. In contrast, blockchain stores the ledger in thousands of locations.

To store something in N places requires N times as much storage, as well as network bandwidth to communicate the data to all the places. What is the optimal number of locations? It is unlikely that the extra redundancy of thousands vs. six merits the storage cost and network bandwidth for replication. And despite how many copies are kept, there have been many clones of Bitcoin that eventually failed due to lack of interest, and all of the copies then were lost, because there is no obligation for a node in a blockchain system to maintain the data.

Is It Being "Immutable"?

The term *immutable* means the data cannot be modified. The term "immutable ledger" isn't quite true. The data can certainly be modified, but the assumption is that there is an integrity

check that can be used to detect whether the data has been modified. Blockchain did not invent the concept of an integrity check, just the concept of a horrendously expensive-to-compute integrity check. Traditional cryptography has long known about easy-to-compute integrity checks that are computationally infeasible to forge.

Furthermore, the ledger in blockchain is not actually immutable. Forks can occur, starting from, say, block *N*, where multiple different subsequent blocks *N+1* and further might be found. The hope is that this situation would be resolved quickly, because a miner seeing two different valid chains will only accept the longer one. However, a fork can persist for a long time if there were an Internet partition, or if the gossip network connecting the miners got partitioned, due to some highly connected node going down, perhaps. Also, if there were any incompatibility in code, such that a transaction looked valid in one version of the code and invalid in a different version, then the miners running different versions will ignore each other's chains. This situation actually occurred in 2013. If blockchain were truly decentralized, then this situation would be permanent. However, there are a few people who really are paying attention and in charge, and after the fork in 2013, they decided which version of the blockchain should live.

Is It Being Decentralized?

The concept of having a ledger agreed upon by consensus of thousands of anonymous entities, none of which can be held responsible or be shut down by some malevolent government, is fairly unique. However, most applications would not require or even want this property. And, as demonstrated by the Bitcoin community's reaction to forks, there really are a few people who are in charge who can control the system, by, for example, making a decision on which fork should be chosen.

The concept of general distributed databases is very old. For instance, this is a survey paper about the state of such systems from 1981 [3]. Such systems are more complicated than Blockchain, because they handle things like having multiple nodes simultaneously attempting to update the same location and atomic transactions. In contrast, Blockchain is an append-only log.

If all that were needed was an append-only log, and an application (e.g., a consortium of banks) wished to collaborate on maintaining the log, a very simple solution would be to have an entry signed by any of the trusted parties in the consortium appended to the log. To handle Byzantine failures (where a minority of the entities in the consortium might become untrustworthy), the simple solution would be to require an entry to be signed by a majority of the consortium before it is appended to the log.

So the novel part of Blockchain is having a consortium of *unknown* entities maintain the ledger.

Blockchain vs. Traditional Solutions for Sample Applications

In this section we'll examine some applications that have been proposed as uses for blockchain and compare more traditional approaches. Since these systems are not actually deployed, it's not possible to completely predict the details of a blockchain-based approach, but we'll mention some issues.

DNS Names

Assigning DNS names is an interesting application. DNS is quite political. Which organization controls the names in a domain? What is the definition of a country? It might be tempting to "democratize" DNS names to first-come first-served, without any organization deciding who is allowed to have which name. With blockchain technology, we could do without any central organizations. And there is indeed a revenue stream for paying the miners, since people would still have to pay to rent a name.

However, people have come to assume that names have some meaning. They assume that the owner of the name *usenix.org* has some affiliation with the organization USENIX. And someone will still need to maintain the servers to map DNS names to IP addresses, along with all the other information stored in DNS.

So it would be preferable to have some mediation of names by a large, identifiable organization that could be held accountable if it misbehaved. And the current system is much less expensive than a blockchain system would be.

Health Records

When switching doctors, or when visiting several doctors with different specialties, it is important for them all to have access to your health records. However, is a universal, world-readable unstructured database with everyone's medical data the best answer? The sheer size of the database is daunting, especially when, as proposed by some blockchain enthusiasts, all medical devices attached to all people would report their readings into the blockchain. And this database would be stored in thousands of places.

Clearly with medical information, people will not want their information world-readable. Which leads to many questions that blockchain doesn't answer. Data must be encrypted. Who manages the keys? Who authorizes a new doctor you are meeting to see your records? What if you are in an accident? And, furthermore, who authorizes you, a doctor or a device, to write something about you in the blockchain?

With traditional technology, there would be a database stored with several trusted organizations, organized so that data for a particular patient could be quickly retrieved (rather than needing to have all the pieces found by searching through the blockchain). And even if encrypted, there would likely be access control on the data. And maintaining the database would be much less expensive if one organization, or a few large organizations, were using traditional digital signatures as an integrity check on the data.

Timestamping

One of the applications claimed for blockchain is the ability to prove that something happened before some time, because of where it appears in the blockchain. For instance, to prove you invented something, you could write a paper about it and store a hash of the paper on the blockchain.

However, there is much less expensive technology that can accomplish this. A trusted timestamping service can take a hash, append a timestamp, and sign it. Since this is such an inexpensive service, there could be hundreds or thousands of them. If Alice wants to be able to prove to Bob that something existed before some time, she needs to collect multiple signed copies to ensure that, when she needs to prove a timestamp to Bob, at least one of the timestampers she used is trusted by Bob. It is less expensive for everyone who wants this service to store their own signed copies than to store them publicly in a large blockchain.

Conclusion

Blockchain technology is extremely expensive in terms of computation, storage, and network bandwidth. With traditional technology, it is possible to replicate data, and public clouds are careful to do so. But there would be a handful of replicas; not thousands. Also, databases would be more structured than an append-only log combining information from all users and for many applications.

Most applications (such as financial ones) do want to have some collection of well-known organizations at the heart of the technology to mediate disputes and be held responsible if things go wrong. If it is distasteful to have a single organization in the center, it could be a consortium of several, and transactions could be considered valid only after a majority of the inner circle of organizations have signed the transaction. This would be immensely less expensive, and be a more natural trust model, than thousands of anonymous miners.

And traditional cryptographic integrity checks (digital signatures) by well-known organizations are practical and inexpensive.

References

[1] K. Torpey, "Why the Bitcoin Blockchain Is the Biggest Thing Since the Internet," *Bitcoin Magazine*, April 19, 2016: http://www.nasdaq.com/article/why-the-bitcoin-blockchain-is-the-biggest-thing-since-the-internet-cm608228.

[2] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," Bitcoin: https://bitcoin.org/bitcoin.pdf.

[3] P. A. Bernstein and N. Goodman, "Concurrency Control in Distributed Database Systems," *Computing Surveys,* vol. 13, no. 2, June 1981: https://people.eecs.berkeley.edu/~brewer/cs262/concurrency-distributed-databases.pdf.