FTC FinTech Forum: Artificial Intelligence and Blockchain
March 9, 2017
Peter Van Valkenburgh
Transcript

PETER VAN VALKENBURGH: Thank you to the FTC and thank you to Berkeley for having me to talk today. I have the unenviable task of doing some level setting for blockchain. I don't really like the word "blockchain" personally, because the word "blockchain" to me is a lot like the word "vehicle."

And no one ever says to you in a friendly conversation or a policy conversation like, "Hey, how do you feel about vehicle?" Or, "Hey, we can fix this problem with vehicle." And we might talk in the abstract about "vehicle technology", but I think even that would be quite strange. We should probably be having a conversation about, "Well, there's a bus" or "You could take an Uber."

And blockchain is actually the same. The word is very broad. There is no "the blockchain", any more than there is "the vehicle." And "blockchain technology" as a bundle of things is a really broad category.

So there's one thing that we definitely know fits inside of this category, and that's Bitcoin-- the Bitcoin network and the software. And that's something I can say with pretty good certainty, because the term blockchain was actually originally invented to describe the list of transactions that all the computers connected to the Bitcoin network help compile and keep locally in order to describe the movement of bitcoins on this peer-to-peer network.

But since Bitcoin's release in 2009 as a running network, and it's invention in late 2008, and a white paper, there have been several follow-on innovations from individuals, companies, non-profits, for-profits, which have fallen into this category, "blockchain technology," usually because they've borrowed some of the technological aspects from the original Bitcoin software or from research done surrounding that software. But what do all of these things really have in common, aside from the fact that they all borrowed something from Bitcoin? What can we actually say about blockchain technology-- this grand sounding phrase-- that's usually or almost always true?

Well, all blockchain technologies have these three essential components-- peer-to-peer networks, consensus mechanisms, and blockchains, which I mean hash-linked data structures. Now, you might be saying, all right, Peter why did they call it "blockchain technology" if there's these three things? And I think that's actually going to come down to branding, probably.

All of these things existed before Bitcoin. For example, we had hash-linked data structures that existed. We had consensus mechanisms like Paxos, which was developed in the 1980s. We've had peer-to-peer networks, obviously-- BitTorrent, probably, being the one most people are familiar with.

And from a branding standpoint-- "Well, this is piracy," "This is something eggheads in Berkeley talk about," and, "Wooo! Well, I haven't seen that before."-- that's a good catchy name for that. It's a little weird, but it's better than "cryptography," which sounds like it happens in the basement of a church or something. So we'll go with it.

But really, these three things are important to think about in total, because that actually helps demystify it to a lot of people who will discuss this for policy purposes, or for personal purposes because they're interested. Because blockchain is just one part of it. And I can say that because we can take these three essential components and turn them into a sentence that actually describes blockchain technology in a much more useful way.

So blockchain technology always works like this. Some of my-- there we go, enumerated lists aren't working here. So blockchain technology always does this-- connected computers reach agreement over shared data. If it's a block chain technology, it should do that. It should allow connected computers to reach an agreement over shared data.

And each part of that sentence corresponds to one of these technological aspects. How are the computers connected? You don't interact with somebody on the Bitcoin network by going through an intermediary necessarily, so it's not like using Facebook. You're actually interacting peer-to-peer, like you would if you were using BitTorrent.

How did the computers reach agreement? Now, this is a consensus mechanism. This is where longer scary terms like Practical Byzantine Fault Tolerance come in. But really, the long and short of it is, we need to have software that develops rules for all of the connected computers that facilitate agreement, that make sure that-- even if some of the computers go offline, if some of them are messed with in order to try and inject fraudulent data into the consensus that's trying to be achieved-- that it's ignored. Rules that facilitate agreement in software.

And finally, what are they reaching agreement over? It's this shared data. Now, why do we call it a blockchain if it's just shared data? Why wouldn't we just call it a database? It's because we employ some cryptographic techniques to make the data that we're reaching agreement over easier to verify. So we have Merkle trees that push little atomistic data into blocks, and blocks which are chained together to create a timeline of things that happened.

Now, what this means is, if you change any data one point down the line, the math won't work out. The cryptographic hashes won't work out anymore. So it really just makes the data very verifiable. And that is part of what makes a consensus algorithm work, because the computers all across the network can rapidly verify and validate the data.

But anyway, the long and short of it is that all blockchain technology does this-- it allows connected computers to reach agreement over shared data. All of these do that thing. And Bitcoin did that thing and still does that thing today.

Looking specifically at how Bitcoin does it might help unpack this construct of this idea a little bit better. So what are the peer-to-peer-- what are the peers, what are the connected computers in Bitcoin? Well, if I ran a Bitcoin software wallet on my phone-- and in essence, my phone is one of the connected computers that's part of a system that's coming to agreement over Bitcoin transactions. Now, that would mean that if I generated an address using this phone at which I could be paid, this is the only phone that actually has the credentials. This is the part of the network that matters to me.

Now, that is something that generates risk potentially for the consumer if they're not a sophisticated consumer or if the user interface hasn't been well-designed. So plenty of people will not actually necessarily be a node on the network in any capacity, but rather use a company like Coinbase to be the node on the network for them. In that sense, they go to a website, looks like online banking. But what Coinbase is in the business of doing is running these computers that connect to the Bitcoin network.

Reaching agreement-- Bitcoin as an example. The consensus mechanism in Bitcoin is much more complicated than this very simplified example. And of course, it's in computer code, it's not in human-readable language. But these two rules, paraphrased, give you a decent understanding of what these rules for agreement actually look like in a blockchain technology. So for Bitcoin, we have at least these two and a few others.

First, nobody can send bitcoins that they've not first received from somebody else. This is just basically a rule against counterfeiting. You need to reference a transaction where you are the recipient in order to send funds on. And what are you referencing? You are referencing the shared data we all have.

And then the second rule-- every 10 minutes or so one of the connected computers will be selected to choose the order of valid transactions for that period. Now, why do we have this? Because if you have a bunch of computers networked together across the entire world, they have a tendency to get out of sync. Some of them might hear an order of transactions come in that looks like "A, B, C." Some of them might hear "C, B, A." We need a provably fair and open mechanism for picking one computer on the network to state the authoritative order of valid transactions for that period.

Now, note I said "valid transactions." Because if you tried instating your authoritative order because you got selected as the winner for that period, you tried to inject some fraudulent transactions where you just made up bitcoins, you'd be violating rule 1. So you see now how these rules make it easy for all the computers in the Bitcoin network to facilitate an agreement over something fairly simple-- in this case, a list of transactions between the nodes on the network.

Finally, the shared data. This is what we're all coming to agreement on in the Bitcoin context. It's a list of transactions. This is again highly simplified. Instead of names we would have what are called Bitcoin addresses, which are, essentially, random but unique numbers, which can be used as pseudonyms and have matching private keys for authorization purposes. And of course, there's a whole bunch more transactions than this. There's millions and millions of transactions that have happened since 2009 when the network went live.

So now you might be saying, all right, well, we're here to talk about more than just picked one. What about these other blockchain technologies? All of them do this-- they allow connected computers to reach agreement over shared data.

Now, what can differ? The data can differ. So rather than it being a list of transactions, we could have a list of identity credentials. So it could be something like a bearer token that someone carries around, that allows them to prove that they have a certain credit score, rather than proving that they can send someone some bitcoins. It could be a bearer token of sorts that allows somebody with a device-- once, and once only-- to vote for a candidate, and then evidence of how that person with that device voted. It could be Internet of Things related, so it could be-- who has permissions to open this door, turn on this smart light bulb in this home, and who doesn't.

Because this is again something we want to have social consensus over. In the physical world, we have window shades and actual physical doorways. In the digital world, how do we delimit, deliminate this space that is mine, that I should have dominion and control over, versus the space that someone else has, especially when my smart light bulb might actually be phoning home to a server somewhere in the middle of America?

Records of securities transactions. So it could be actual peer-to-peer transfer of an instrument that represents a share of Apple stock, or something like that. It could be property records for deeds. It could be interbank settlement records. It could be records about the provision of digital goods and automatic payments to remunerate people for providing those digital goods. It could be all sorts of things.

And then the other thing that could differ from Bitcoin is aspects of the consensus mechanisms. What are the rules and what are the design choices made in designing that consensus mechanism? The first big question we could ask here for consumer protection purposes-- I'm just thinking about the architectures of these systems-- is this an open network that's like the internet or is this a closed network that's like a company intranet? Who is allowed to partake in the consensus? Who is allowed to help reach agreement over the shared data? Is it open to everyone-- as you'll find in, say, Bitcoin, Zcash, or the Ethereum network-- or is it only open to a closed set of previously identified and credentialed nodes, which will usually belong to business entities, I would say?

So this is, for example, what R3CEV has worked on with their Corda platform, which would be a way of getting a bunch of banks to agree on certain shared data-- similar to what Symbiont's built, what Axoni's built, and a number of other companies in this space. Now, the difference really is about a pseudonymous system where anyone can join-- just like IP addresses or pseudonyms, and anyone can get one-- versus a fully identified system where only a few people can join, but they need to have identities provided by a central authority.

The next thing to think about with rules and design choices is privacy versus transparency or auditability. And there is to some extent an eye on trade-off here, because if you're going to have a bunch of disparate nodes in the world agree on data, they are all going to need to have some ability to look at and audit that data to say that, yes, that's valid, we agree that that's part of the consensus. So data on the blockchain cannot be, as you might say, fully encrypted. It has to be partially transparent for all of the participants to be able to look at it and say, yes, that's a valid part of the consensus.

And really, Bitcoin, though famed to some extent for being an anonymous payment system, is quite far from a highly private system. Every transaction between two nodes on the network is evidenced on the blockchain. Yes, it's evidenced with this pseudonym. But if you can match a real human personality to a pseudonym, then you can actually see with a lot of robust evidence every transaction they've ever made using that pseudonym.

Now, R3CEV-- just to take another example-- with Corda has developed a platform that they say is a little bit better from a privacy perspective in some sense, because the data are relevant to the consensus. It's only shared with those for whom it is relevant. So if you're dealing with a contract between just two banks who might exist on the Corda system, the data relevant to those contracts is only shared with those two banks and verified by those two banks.

But you'll notice there is a trade-off here-- you now don't have independent and global verification of the data that they're agreeing on. The verifiers are then limited to this set that can see it. So it creates a kind of perimeter security for a small part of the blockchain technology or part of the blockchain network, rather than everything being aired and agreed upon in public-- starkly contrasted with Bitcoin where everything is effectively aired and agreed upon in public, where all the nodes can see it and verify that data.

And then there's Zcash, which is worth discussing because it's an interesting hybrid almost, to a certain extent. So Zcash is an open permissionless blockchain technology, just like Bitcoin where anybody can connect a computer by running free and open source software. And anyone can have that computer participate in the consensus. But transactions can actually be encrypted in that shared data, such that you will not be able to know the identities, even the pseudonymous identities, of the sender, the recipient or the amount sent.

Now, you might say, well, how is that possible? How can we know it's a valid transaction without knowing those three other details? And the way that Zcash actually creates that interesting arrangement is by using very novel cryptographic tools, called zero-knowledge proofs, to say that, yes, this is a transaction that did not invent any new bitcoins. We can say that with mathematical certainty by looking at the encrypted data, but we can't say anything else about it. It's very new, it's amazing technology, and it's really high science and something that's encouraging and very exciting to watch unfold, especially on an open network.

So finally, the last rule and design choice to think about these various blockchain technologies that will also have consumer protection implications is, where is the security in the system. The main thing that blockchain technology does-- probably more fundamental than anything to do with privacy or anything else-- is push security to the edge of the network.

What I mean by that is this-- if somebody in the crowd-- if I were to show a Bitcoin address I generated on my software while I'm on my phone right now, and somebody in the crowd was to snap a picture of that QR code and sent me some bitcoin-- we should have done this for real, so I can make some. No, I'm not hurting, so-- and sent me some bitcoin to that address, the only person in the world who can reverse that transaction, say, OK, good demonstration, Peter, now send me back my bitcoin. And I'd say, no, I don't want to. The only person in the world who can reverse that transaction is me.

The security on the network is at the edge. It's at the end user device. This is marked contrast to anything else that exists in the real world right now, where you have automated clearing houses, you have banks, where you could say, oh, hey, bank of America, sorry, I didn't mean to send that money, please reverse the transaction. Or where you have the credit card industry, which is basically operating for the express purpose, primarily, of making it easy to make transactions and also reverse them in the case of fraud. Those are centralized systems where the security and the power on the network exists at a data center somewhere or in a corporate business logic setting of some sort. This is a network where security exists on the edge.

Now, Bitcoin is all about the edge, and so, really, is Ethereum, and Zcash, and these open permissionless networks. Some of the block chain technologies-- that are either self styled as blockchain technologies or have borrowed some technology from Bitcoin and therefore call themselves blockchain technologies-- they're mostly just data centers. You might have a bunch of nodes, but you own them all, and they all sit in the same data center, and you have full control over them. It's not exactly clear how that changes the security calculation from a world where you gave all your data to Facebook. They can change your data, move your data, do whatever you want with it. What have you done that's exactly different?

So this is another really important way of looking at these systems. And whenever you're looking at a particular blockchain architecture, it's important to ask this question-- where is security, is it at the edge or is at the center? Sometimes this category is referred to as "immutable" versus "mutable." The Bitcoin blockchain is immutable, because the only way to change a transaction is to write a new transaction on top of it. And the only way to get somebody to do that, assuming they have the keys, is for them to use their end user device, their edge of the network device.

So just to sum up, connected computers reach agreement over shared data. This is what we're now going to talk about, and what we're going to talk about in the context of consumer protection, regulation, all of the issues that these technologies raise.