GARY GENSLER: Hello, everybody. Good weekend? Everybody staying dry, I hope. So we're going to dive back in to blockchain, money. And we didn't lose too many people when we did cryptographic hash functions and digital signatures last week. So I thank you for all being back.

So what were today's study questions again? What's the Byzantine Generals problem? Anybody want to tell me what the Byzantine Generals problem is? Ben.

AUDIENCE: So, the Byzantine General problem is this sort of general mathematical puzzle. And basically, what it is is, how do you coordinate actors when they may be an actor who's not acting in the best interests of the group-- how do you sort of get a good actor on that then?

GARY GENSLER: So there-- it might be how to coordinate when somebody is not acting in good faith as a malicious actor. But it also might be just somebody that doesn't get the communication. Somebody that-- there's a thought, whether it's malicious or not.

How we doing, Alene? Yeah.

How does proof of work and mining in Bitcoin address it? We're going to walk through this, and I'm going to give you my sense of it. But does anybody want to give a short version-- prepared for a-- Brodush.

AUDIENCE: So the proposed way solves it in a probabilistic way, rather than a deterministic way, using amount of CPU power to solve a problem of certain complexity to prove that one consensus has been reached by a majority of the participants.

GARY GENSLER: All right. So what Brodush said is it's probabilistic instead of deterministic. That you use CPU power to form some consensus. I think that's what you said. I'm going to walk through this in more detail. But does anybody want to give another shot at it?

AUDIENCE: I could say so when a transaction happens, it is posted. And then miners get in, and that they try to compete in solving this. And whoever gets it first, then he claims the award.

I would also say it requires a lot of CPU processing power. I'm quoting the [INAUDIBLE]. So it needs to be powered. Big computers. So there is a question around [INAUDIBLE].

GARY GENSLER: So remind me your first name?

AUDIENCE: Riham.

GARY GENSLER: Riham.

AUDIENCE: Yes.

GARY GENSLER: Riham says it's about people called miners or computers called miners, which we'll talk about today, using computer power again. But transactions were part of that as well as to how it comes together.

I'm not going to torture you and ask a bunch more but, Addy?

AUDIENCE: I think one of the other important ideas is that even though solving the problem is really hard, validating that the solution is correct is easy. So what is ensured is that even though miners are thinking about the computing power to solve [INAUDIBLE], it can propagate it [INAUDIBLE] and then validate it.

GARY GENSLER: So a key point that Addy-- Addy raises is that once somebody solves the puzzle, it can be propagated across a network. And then others on the network can validate it. So what we'll talk about a little bit later is that it's a hard puzzle to solve, but an easy puzzle to verify.

And this is an important asymmetry in essence, that there's a lot of resources to solve a puzzle, but once knowing the answer, there's very limited resources to verify that it's the right answer. If it was not asymmetric, it would not work as well.

So it's a key part of the design of many cryptographic things, but particularly, Adam Back's sort of novelty in the 1990s of proof of work.

I don't know if that's just a relaxed bit for you, or you have your hands up, Derek.

AUDIENCE: Ah, no. That's--

GARY GENSLER: That's just a relaxed pose. I like it.

We're going to talk about other consensus protocols as well. Proof of work is not the only

consensus protocol, but how to address ourselves to Byzantine fault tolerance. And then some of the economic incentives, so we'll talk about the native currency. What's the native currency of the first blockchain application?

AUDIENCE: Bitcoin.

GARY GENSLER: Bitcoin. How many of you in this room have owned a Bitcoin at some point in time? So you've all owned the native currency that helps a blockchain.

You got the opportunity to read a paper that a group of us here at MIT co-authored. I hope some of you actually were able to download it and-- good. Apparently, even when you write something, it still goes behind a copyright wall. So I'm glad to know that it was actually available.

And then back to the National Institute of Standards and Technology paper as well-- this time the next chapters. And then a paper from about 25 years ago on the Byzantine Generals problem itself.

So what are we going to do today? We're going to go back through the design. We're going to talk about consensus through proof of work

Bitcoin mining-- it's important. It's relevant. But it's kind of some fun facts about that as well. The native currency-- of course, Bitcoin of the first. But there's now at least 1600 different native currencies.

What does it mean to have a network? And why do networks matter particularly for blockchain? Some of the other consensus protocols, and then just wrap up.

So just going back to the review, and this is what we talked about a bit on Thursday. But I think it's relevant to just kind of bring it back.

I found when I was first learning this, it's hard to keep all the moving parts. Remember, there's-- that graphic, you'll see all semester. But it's all the different blocks. And it's append only.

What does append only mean? Andrew. Why does that matter, this word append only that we talked about last Thursday?

AUDIENCE: Yeah. So that it's immutable. It cannot be changed.

GARY GENSLER: Right. So it's immutable. Now, of course, because cryptography maybe can be broken, but we use the word immutable that it cannot be changed, except for maybe as Alene so-- I like that that's in his book. You all have to see this. This is very clever. You know, a little-- how to do a flag.

Maybe it's immutable except for one out of 10 to the 40th times it could be broken or something.

AUDIENCE: Can I interject and actually say that it's not that good, actually. So technically, Bitcoin and all of these permissionless cryptocurrencies, one way to attack them is to mine-- to get a lot of mining power and mine.

But another way to attack them is to just take control over the actual network, like the internet. So if you're an internet service provider, or if you're China, you can actually fork Bitcoin with zero mining power by just controlling the network.

So there's actually an assumption behind how this thing works, which is that the network works. Because everybody sees the messages.

GARY GENSLER: I agree with you. But there's also some assumptions-- let's say that China, or any state actor, chose to fork one of these. If it's considerably less than a majority--

AUDIENCE: No. Zero mining. You can fork with zero mining. You won't get any mining charge, if you control the network.

GARY GENSLER: If you control the worldwide network, or just the network--

AUDIENCE: Not the mining network. The actual internet.

GARY GENSLER: I understand. Are you talking about the worldwide internet, or you're talking about one country's internet?

AUDIENCE: Let's say you're in China. There's 50% of Chinese miners in China. I forbid these Chinese miners to broadcast with blocks. They find a block that goes here, and they find another block that goes here, it goes here, and it goes there. The rest of the world will find the block that goes here. I have a fork. It goes here and here, and I have two forks. They don't see each other.

GARY GENSLER: So we're going to talk about forks a little later. Can we hold Alene's point until then?

AUDIENCE: Sure. Sure.

GARY GENSLER: And then I'm going to share with you what Satoshi Nakamoto wrote about this very issue back in 2010.

AUDIENCE: And he was writing back where you have mining power. He was assuming the network works. Again, just to clarify.

GARY GENSLER: Alene's raising a point as to whether somebody captures part of the internet. And if the internet itself, by capturing part of the internet, you fork the blockchain.

And what I was just-- said I would hold for later, but instead I'll cover now, is this question was raised in an email exchange with whomever Satoshi Nakamoto was back in around 2010. And his answer to Alene's question that I'm just helping share with you all, is that as long as that part of the internet that was walled off was less than a majority, and in fact, if it was China, because that was in the example even eight years ago, it would be considerably less than 50%, that within a reasonable amount of time, maybe it would take a few hours, but within a reasonable amount of time, one chain would be where the majority of the mining power was. And that it could take a while, but the other one would probably stop, that people would stop investing electricity and CPU time within China, because they would realize some way. Now, that was the theory at least.

AUDIENCE: And just to add onto that, something to realize is like if you-- I mean, like, you said we'll go into forks. But just for people who don't know, like if you go in one fork, then anything after that you go in all forks. So like if you lose connection to the main Bitcoin network, you'll still have that, as long as you have your private keys. So if somebody in China realizes that they're on the wrong chain, it's not like you've lost your actual bitcoin.

GARY GENSLER: That's correct. They won't lose it up to, if I can use the term, prior to the fork.

AUDIENCE: Correct.

GARY GENSLER: Was a question? Brodush.

AUDIENCE: So just to add to Alene's point, essentially. So the way he said--

GARY GENSLER: Since you're in the back of the room, speak up.

AUDIENCE: Yeah. So to add to Alene's point actually. So there's an assumption, underlying assumption, for the-- that the real problem in the context of blockchain is-- you have the assumption that the network can actually verify that the-- what is being-- [INAUDIBLE] what the network is actually valid information. So that is kind of an underlying assumption.

If the network is contaminated, then the premise on which the problem is being solved, or the protocol that is being given here as a solution to the problem, is actually not valid if the network is contaminated. So that way, it is indeed a underlying [INAUDIBLE].

GARY GENSLER: So I think, because I didn't pick up every word, you're just saying that there is an underlying assumption that the network protocol, the communication protocol of the internet, is not compromised or walled off. But also, that it's working.

It creates a database. We've talked about it through hash functions and digital signatures, and then consensus. So what were the technical features?

I thought about it a little bit from our last class to help just thinking through in three buckets-- the cryptography and timestamping that we talked about last Thursday; what we're going to talk about today, the decentralized consensus protocols and the network, of course, and the native currency; and then lastly, transaction script that we're going to talk about this coming Thursday.

Now, it's not just three buckets because it's three lectures. But it's three buckets because they have something to do with each other. The cryptography, which is at the core of cryptocurrencies and blockchains, and is the core of a lot of things on the internet today, the consensus mechanism, and then the transaction script itself.

Cryptography, as we've talked about-- communications in the presence of adversaries-- also, a form of ways to make commitments and secure computation.

Hash functions, if you recall what we talked about. What's the key of a hash function? Here. Joaquin.

AUDIENCE: The key of hash function?

GARY GENSLER: One-- the elevator pitch. You just have to make sure that your sibling knows you. What's that?

AUDIENCE: If you have two, the private and the public key.

GARY GENSLER: All right. That's a good-- that's cryptography, but not a hash function.

AUDIENCE: OK.

AUDIENCE: It's a fingerprint of a fixed length of any amount of data.

GARY GENSLER: I like that. One way data compression-- a crossword puzzle. Anybody here do the *New York Times* crossword puzzle on a mobile app? Good. So if you do a Wednesday's *New York Times* crossword puzzle, does it tell you whether you're correct on Wednesday? Or does it not tell you you're correct? When does it-- I don't do the *New York Times* crossword puzzles. But--

AUDIENCE: Erin.

GARY GENSLER: Erin.

AUDIENCE: I'm actually-- I usually do ones in the past. But I think it will tell you maybe either that day or the next day, or maybe [INAUDIBLE].

GARY GENSLER: Stephanie.

AUDIENCE: So it tells you as soon as you finish the puzzle whether or not you have any errors. But it won't-- but you can't actually check what the errors are unless you want to invalidate your streak for that day. So basically, you get a streak every time you [INAUDIBLE].

GARY GENSLER: I don't actually know if the *New York Times* use hash functions. But they could. They could, because they could stick the whole entire crossword puzzle into a hash, and it's a commitment scheme. And remember, if you change even one thing in the input data, the hash will come out differently.

So the *New York Times* could use a hash function, so that Stephanie could find out, right? Because you can only push to see if it's correct when you finish the whole. And it either tells you you have it or not, right?

So I'm just-- I'm bringing it to real life that a hash function-- just think of the *New York Times* crossword puzzle. And if you don't remember, ask Stephanie.

We talked about append-only logs. And recall that in blockchain, in Bitcoin, there is a bunch of

information in the head of the block. And that which is in the head of the block is put together like the *New York Times* crossword puzzle. And we have a chain of blocks.

Most of the data, though, is stored efficiently in something called a Merkle tree. Again, it uses a whole lot of hash functions. And so it's a way to be efficient, but it's also a way to secure the data.

So now, we're going to get back to your favorite thing-- digital signatures. So what's a digital signature do?

AUDIENCE: You can prove that you're signing something like a transaction with your private key. And the other person on the outside could prove that you are the one that signed it with your public key.

GARY GENSLER: Perfect. So it guards against tampering and impersonation. I didn't go through this last Thursday, but think of digital signatures two different ways-- a digital signature that you use without a hash, and in Bitcoin and blockchain, often it's actually-- it's combined with a hash.

So as Joaquin just went through, you can have a private key that you sign something with, the sender's public key, and a signature, and exactly that. But it's also able to do it where you have a hash as well. You take all the data, all the message, and you've put it into it with a signature.

And this is a little bit too complex, and it was last Thursday's lecture, but it's important to know that what blockchain is basically doing, most blockchains do, is they take a lot of information, a transaction for instance, hash that information. And why do we hash it again? Kelly?

AUDIENCE: To protect it from other users of the network, sort of like we talked about Alice and Bob, and how one has to be aware of each other's key, and then back verify the incoming message.

GARY GENSLER: Right. And it also compresses some of the data. But it's a commitment scheme. It's like this is it. This is actually the *New York Times* crossword puzzle that answers all the questions. So usually it first hashes it, meaning it's a commitment. And then put a digital signature on it.

And there was one last thing we talked about last week. What are Bitcoin addresses? Isabella, can you tell me when a Bitcoin address is?

AUDIENCE: Umm. Is that what-- like, I guess tells you where the Bitcoins being sent.

GARY GENSLER: So it tells you where Bitcoin is being sent. Ben, you want to help out a little.

AUDIENCE: So it's the public hash, public key? Public hash of the--

GARY GENSLER: It's close. So it's basically that-- between Isabelle and Ben, you've got it. It's basically how any of the native currency-- Bitcoin-- can be identified. But it is a public key with a couple extra hashes, and a little bit other fancy footwork to make it compressed and smaller. But it is literally what you can send Bitcoins to.

So it's determined by the public key, but it's not identical to. And I found a fancy little chart to define it. A private key leads to a public key through some form. And in Bitcoin, it's called elliptic curve multiplication. But there are other forms of public and private keys.

The public key, then it gets hashed. And then it goes through a code that makes it shorter, which is the Bitcoin address. Part of the reason it was hashed, and part of the reason it goes through that extra code, is to make it even more secure. It's not the only reason. It also compresses it a bit more. But those of you who have ever owned Bitcoin, you have a wallet. And the wallet keeps those Bitcoin addresses.

All right. So now let's talk about decentralized networks, the topic of today. Any questions about the review for last Thursday? I know it was quick. Alon?

AUDIENCE: I have a question about the double hash part. Does that mean that it's now less feasible to be, like, less immutable? Because if you take a 24-digit hash, and you contract it to a 4-digit hash, there's fewer options.

GARY GENSLER: I think that you're-- the question is, is if the output of a hash function is shorter, is it possibly more breakable? I think mathematically, that might be correct.

However, this actually goes through two hashes-- one, which is this mechanism called SHA-256. And the other one, I'm going to mispronounce, but down to 160. So I think because it's going through two different hashes, the answer is it's even harder to break both. Does that--

AUDIENCE: That makes sense.

GARY GENSLER: Any other questions about the review? No. Please. Derek.

AUDIENCE: So you said hashing the public key makes it more secure. I'm just wondering, because the public key is for the public. So what is the-- where does the added security come from?

GARY GENSLER: So the only thing that you're actually showing is a Bitcoin address. Until later-- and we'll talk a lot about this on Thursday-- when you actually do a transaction, you have to then disclose your public key.

So initially, the storage is around Bitcoin addresses. And some will advise-- and it's why many wallets do this-- that you should never use the same public key twice. Though, numerous people do in blockchains. But to be most secure, you would constantly be creating new public key/private key pairs. And once you've used it, move on and get a new set of keys. Got it?

AUDIENCE: Yeah.

GARY GENSLER: So distributed networks-- we talked about Byzantine Generals problem. So I found some Byzantine generals. They want to all attack that castle. Or what if only three of them do, and two of them say retreat? That's the visual. That's the problem.

The only way to win in this mathematical game theory, a paper that was written some 25 years ago, is if they all said attack, or all said retreat. But the same thing sort of came to computers.

And the core thing about a permissionless system is there is no central authority. And if there's no central authority, how does a distributed network, like the distributed set of generals, come to some agreement? Do we attack? Do we retreat?

Well, it's based on a consensus protocol and a native currency. That's the key innovation of Satoshi Nakamoto, is to pull it all together. But it was built on the backs of other people. Adam Back, in 1997, he proposed a way to address email spam and other types of computer problems called denial of service attacks.

Now, it ultimately wasn't used. I mean, he proposed it. It was used for a short while, and then it wasn't subsequently used. But it's important to understand that the proof of work in the middle of Bitcoin was created 11 or 12 years before the Bitcoin paper. And the key was basically require a bunch of computational work using hash functions.

And so the email, or the header of the email-- this is Adam Back's, not Bitcoin. But the email, or the header of the email, went into the hash function, creates a hash.

But the difficulty of finding whether it's confirmed was was it in a certain range of hashes? And he did that by the quote, "leading zeros." Does anybody want to guess why he did it this way?

Or Alene's just going to tell us probably and not guess, but-- who hasn't spoken yet?

Emily, you want to try it out?

AUDIENCE: I'm not totally sure.

GARY GENSLER: Daniel?

AUDIENCE: I mean, I guess just, like, preserve some privacy around the emails?

GARY GENSLER: Well, so it definitely preserved privacy. But he was trying to put some computational work. Every email that would be sent would take one to two or three seconds of computational work. That was in his original paper. It would take a few seconds.

AUDIENCE: The fact that we-- earlier we were talking about the fact that we need a way in which we have to make the puzzles difficult to solve, but easy to validate. This is exactly how the whole thing is accomplished, by setting the hash into a fixed characteristic, like leading-- a number of leading zeros, what you get is to modify a small piece of the whole information, and try and try until you get that specific hash. And that makes it really computational intensive, but validating is just running one hashing function. So--

GARY GENSLER: Do you have a--

AUDIENCE: I mean, it comes back to what we talked about last time with the nonces. You need to try out a bunch of different random numbers in order to get the right number of leading zeros. And it could be-- I don't know. I think it could be like leading anything. But he chose zero because it's nice. But you need to try to get the numbers in order to get the right number.

GARY GENSLER: And in the email circumstance, his thought was it will take two or three seconds for anyone sending an email to do this proof of work. But it will take a nanosecond or less to confirm it. But if you were sending spam, and you had a computer to send millions of spam times two or three seconds apiece, that would be too much for the spammer. That was it.

So any one person-- any person sending one email, it won't be too bad. Anyone sending millions of emails, it would be lousy. And so that's why this concept was in the midst of emails. And it could be efficiently proved.

So back to blockchain, the innovation was basically, how do we do this with a chain? How do we do this with a chain set of works? And remember, Stuart Haber, that whole thing about the

blockchain and what's in the *New York Times* was that chain of information. But here, why don't we do a proof of work between the chain?

And I found a little graphic. But the SHA-256, that's the formula which is used to hash the header-- the previous hash, the transaction hash, a time stamp, and a nonce. Can you find a hash that has a certain number of leading zeros? This was the key innovation. In a sense, or maybe Satoshi Nakamoto was just taking Adam Back's email proof of work.

Remember the reading for last week of blocks? This is colored green because each of them have hashes that in this case have leading, if I'm right-- is this leading four zeros? Leading four zeros. What if we change one thing? What's going to happen?

Is it Alfa?

AUDIENCE: Yeah.

GARY GENSLER: What happens if we change one thing?

AUDIENCE: The hash should change completely.

GARY GENSLER: The hash will change completely. So what happened? What did we change? Here, I'll go back and forth.

AUDIENCE: You changed the color.

GARY GENSLER: What's that? Zan?

AUDIENCE: You changed the coinbase transaction, so that $100, I guess in this example, went to you instead of [INAUDIBLE].

GARY GENSLER: Well, why shouldn't I be able to get $100 for free?

AUDIENCE: Well, in this example, I guess the coinbase is for the miner, right? So there's one transaction dedicated for whoever validates the block, gets, right now, 12 and 1/2 Bitcoin. And so you add that in addition to all the other transactions that you're validating. But in this case, you're not actually this miner. So you shouldn't be getting that much.

GARY GENSLER: So the little 18-minute video that was assigned for last Thursday, I just went in, and I was trying to-- I was trying to get-- I was trying to get the money for me. And it invalidated the rest of the chain.

And that's really-- that's the sort of innovation or genius is, is if you try to go into a former block, whether it's the last block or a block 100,000 blocks away, and change one little whisker of information, or one letter on that crossword puzzle, it's going to change the entire blockchain.

And I bring it back to the crossword puzzles, or a whisker on a cat. It's just any little bit of information. So an innovation about hash functions became, and an innovation about timestamp blocks, all of a sudden came together with this proof of work innovation.

So now to the chains themselves. The consensus of blockchains-- and many people would say that, in fact, the reality of blockchains is only the longest chain is the one that other miners, other people, will build upon. As I understand, though, it's not written into the base computational code. It's really just a consensus that comes about.

This is an example-- the purple block and the black blocks. The purple blocks are kind of stale blocks. They were mined. They were computationally solved a proof of work. But nobody mined on top of them. And if somebody doesn't mine on top of them, then eventually they're ignored.

Some people call them orphan blocks. But I'll call them stale blocks, because they were actually created. But the information that's in them is kind of worthless. It's not needed.

In the actual Bitcoin technology, this happens from time to time. But it hasn't happened in over a year. The technology, it's-- and you can look on various websites to find this out. Probably at the maximum, the longest stale chain goes out to two or three blocks. But it's very, very rare.

So back to Alene's question of what if China carved off and had the presumption-- it may not work-- but the presumption is, let's say, China is the purple blocks. Because China's walled off its entire network.

The presumption is there'll be some communication outside of the network. It might be on television. It might be by courier, that the Chinese miners would know that they're not in the majority, and they would stop expending electricity to even mine in that circumstance. Because whether it's a few hours or a few weeks or a few months, they know that their expenditures would be worthless. Tom.

AUDIENCE: In these stale blocks, these forked blocks, are the miners receiving Bitcoins?

GARY GENSLER: So in the purple blocks, there will be, if I can go back to this, there will be a coinbase transaction. But it will be worthless. Because it's a coinbase transaction in a block that's not on the main block. And it won't be usable later.

But in Bitcoin itself, there's software that says you cannot use a coinbase output for 100 blocks. It's written right into the base code. And it has been since the beginning.

AUDIENCE: So can you not verify though that-- so in the situation where network is walled off. People start mining on top of this segmented block. Would they not realize that their Bitcoins are invalid for 100 blocks, save for some external knowledge.

GARY GENSLER: Yes. Let me just go back to the-- sorry-- the chain. You're saying if-- Tom's question is, is what if the purple side chain goes on for 100 blocks. What happens? And in fact, we have circumstances of that.

Bitcoin has split between Bitcoin and Bitcoin Cash. It was called a hard fork last year. And for a moment, let's call the purple chain Bitcoin Cash. It's not only gone on for 100 blocks, it's now gone on for tens of thousands of blocks. It is now its own native currency. Within that community that purple blockchain is so long now, that people have found value in that. And it is its own native currency.

And the reason I share that is to Tom's question of, well, what if China was walled off so long, it's plausible-- unlikely, but plausible-- that there would become some value and call it the Bitcoin China blockchain versus the Bitcoin global blockchain. It would be-- what is money, but a social consensus?

AUDIENCE: How does this society work? I mean, is it based on supply and demand, the amount of forks out there? And then the other question that I have. Who decides it? So if it's supply and demand, is it community? And what is the form of reward?

GARY GENSLER: The question is, is the reward-- I'm going to hold a part of it for a little bit when we talk about native currency. But the reward is, in nearly every blockchain is a new native currency of that blockchain. Bitcoin for Bitcoin. ETH, or E-T-H, for the Ethereum. XRP. For each blockchain there's a native currency.

Who-- the second question was, who decides it? It's generally, but not always, hard programmed into the first release of that blockchain.

AUDIENCE: Yeah. So back to the question on Bitcoin and Bitcoin Cash. If the Bitcoin Cash is the purple line that we see on the chart here, does that mean-- because it's shorter, compared to the block-- Bitcoin chain. Does that mean that under the assumption of the majority consensus, the value of which is essentially zero?

GARY GENSLER: So the question is is if it's shorter, does it--

AUDIENCE: So there's no validity in that chain. And therefore, the value becomes zero.

GARY GENSLER: So there's two-- using this chart just as an example. There's two ways-- there was-- the main point of this chart was to say that the black chain, as represented in black, is the main chain. And that is where the social consensus will stay. That's where the consensus is. And generally speaking, the stale blocks don't mean anything, and the stale blocks go away.

Occasionally, there is something called a hard fork, where the social consensus continues to maintain. And I was using this chart as a rough answer to Tom's earlier question about Bitcoin Cash. And if the purple chain kept going for thousands of blocks, and there was a social consensus to keep both chains going, you'd start to see separate currencies, as you've seen with Bitcoin Cash. Does that help?

So I was using a graphic to answer a separate question. I'm going to take two more questions, and then go to native currencies. I haven't heard from Daniel yet.

AUDIENCE: So my question is similar to the mining. So if your transaction is on one of the blocks, does that transaction become void, so to speak?

GARY GENSLER: It's not so much void. It's just-- it's meaningless. It's-- yes, in a sense, effectively, it's void.

AUDIENCE: So I guess would somebody-- if you initiated that transaction, would you be aware of that and reinitiate it?

GARY GENSLER: Very good question. The transactions will still be in what's called the memory pool of anybody who's mining on the main chain.

So transactions-- which we'll talk a lot about on Thursday, this coming lecture-- go in through the network. They're propagated through the network to the entire node network. In Bitcoin, there's about 10,000 nodes. And they will receive those coins and those proposed transactions. So anything on the purple chain will still be in the other chain's memory pool.

One more, and then I'm going to--

AUDIENCE: In the case of the hash--

GARY GENSLER: Your first name is?

AUDIENCE: Iash.

GARY GENSLER: Iash.

AUDIENCE: In the case of the hard fork, so between Bitcoin and Bitcoin Cash, what are the differences? And what about the differences in value between those two?

GARY GENSLER: That's a much longer question. The question is, is what are the differences of value between Bitcoin and Bitcoin Cash? And though, I think Bitcoin is trading around $6,300, and Bitcoin Cash is--

AUDIENCE: It's about $435.

GARY GENSLER: --$435. Thank you, Zan. That gives you the monetary difference of about 15 to 1. But it would take more conversations about why that happened and background and so forth.

So let me talk about the difficulty factor. So proof of work, at least in Bitcoin's case, has a difficulty factor with regard to these leading zeros in the hash. And Satoshi Nakamoto said, let's change that every 10 minutes. Let's ensure that every block comes on average every 10 minutes.

And to do that, define how many leading zeros there needs to be. And it adjusts about every two weeks. Every blockchain can be different. It doesn't have to adjust every two weeks. This is just what Bitcoin did. This is what Nakamoto did to maintain an average of 10 minutes.

So what has happened? Currently it takes 18 leading zeros. And because this is in a 60-- it's in a hexadecimal character system. Every decimal is-- what's that?

AUDIENCE: Four bits. So it's 64 leading zeros in bits, and 18 in hexadecimal. Is it?

GARY GENSLER: So it's 2 to the 64th.

AUDIENCE: But the probability of finding a block is 1 over 2 to the minuses.

GARY GENSLER: So what Alene just said was that it's a very small chance of finding a block, because this is the equivalent of 18 leading zeros-- so that's more than 64. It's 18 times 4.

AUDIENCE: Oh. Yeah. I know. I'm sorry. I can't do arithmetic.

GARY GENSLER: Yeah. PhD in computer science, but can't do arithmetic. So this is the most recent block I grabbed off the blockchain this morning. And it has 18 leading zeros, and then all those other digits. That's block number 541,974. 18 leading zeros.

The genesis block, the very first block in January of 2009, had 10 leading zeros. But the requirement that Satoshi Nakamoto actually put into the computer code was you only needed eight leading zeros. So the probabilities have gone way up. So let me take it off of fancy numbers like that, and just say this is the actual Bitcoin mining difficulties on a logarithmic scale. Because if it weren't logarithmic, you couldn't really read it.

The difficulty was set at one. This is all scaled to how difficult was it for the first year and a half of mining in 2009 and early 2010, one. And now, it is at one trillion. It's actually more than one trillion because it's logarithmic. It's at about seven trillion. It is currently seven trillion times harder to find the answer to the puzzle than it was in 2009. And that's because there's a lot of computers trying to hash all of this stuff.

AUDIENCE: So is that where the--

GARY GENSLER: Kelly.

AUDIENCE: --the collectors in the pools of mining nodes work to be able to achieve this at a more efficient rate?

GARY GENSLER: Correct. Correct. And the hash rate is now somewhere around 50-- it's not terahashes. I'm trying to remember what the-- what's that?

AUDIENCE: Hexahash.

GARY GENSLER: 50 hexahash per second, which is like 1,000 trillion hashes. Because a terahash is a trillion hashes. Zan.

AUDIENCE: I think it's worth noting, though, it didn't scale linearly as, like, number of computers got on the network. It's also the hardware has gotten incredibly more sophisticated. So it's not that you can just assume there's 15 trillion number of people that are mining Bitcoin. It's just the same

people that are doing it better.

GARY GENSLER: I can't ask for a better setup than that. Bitcoin mining evolution-- did you see my slides?

AUDIENCE: I just read your mind.

GARY GENSLER: So what what's the evolution? So it started with central processing units. And CPUs-- and I'm not sure my numbers are accurate, because I might be using CPU power today, and not CPU power in 2009. Apologies for those who know CPU power better than mine. You could do about 2 to 20 million hashes a second on a CPU properly geared, apparently.

They didn't last that long. By 2010, some folks figured out there was something faster, and it was called a graphics processing unit. We all use GPUs all day long, because that's what gives us all our quick graphics if you live stream something on your laptop. And graphics processing units, somebody figured out you can use that, and you could hash faster.

And then all of a sudden, hobbyists started to wire the GPUs together. And they could figure out a way to get between 20 million hashes to 300 million hashes a second. I'm told that even today you could maybe get up closer to a thousand million hashes a second, or a billion hashes a second, if you did a GPU rig. But that's yesteryear on Bitcoin.

Now there's something called an application-specific integrated circuit, an ASIC. Just think about a circuit that the only thing the circuit does is create hashes. In fact, the circuit is wired-- I use the word wired, because I'm old enough to remember wires. But it's-- the circuit board is manufactured in a way that all it does is the SHA-256 hash function to Bitcoin mine.

And the first ASICs, which are dedicated circuit boards to do this mining, came out in 2013. And even since then, they have moved up the scale. The most expensive that sells for about $3,000 or $4,000 in ASIC could do 16 terahashes per second, or at least that's what it's rated for if you go on Amazon and try to buy it. And you could do that. But you'd be competing with something that looks like this.

A modern map mining factory for Bitcoin has thousands of ASICs. They have water cooling systems to keep it cooled down. And they're probably buying their electricity for less than 3 kilowatts, $0.03 per kilowatt. And they might even be paying off the local government officials, and not even paying the electricity company, and just bribing to get their electricity. Emily.

AUDIENCE: This might be a dumb question, but is there an economic opportunity cost of using all this

processing power just for mining Bitcoin? Like, is there a more efficient allocation of that processing capability in terms of like more-- for a more stable economic usage?

GARY GENSLER: There are certainly trade-offs here. And the aggregate electricity for all of Bitcoin mining, now that it's seven trillion times harder than it was in 2009, has been compared to the electricity use of countries like Ireland, on the way to the electricity use of countries like Denmark, I think. It's somewhere between Ireland and Denmark. See ya, Larry. Alon.

AUDIENCE: Well, add to that the cost is-- let's assume it's in dollars or whatever currency, and the reward is in Bitcoin, the volatility of Bitcoin makes it hard to answer that question. Because you don't know if there's an economic value for you, because you don't know what will happen to Bitcoin.

GARY GENSLER: So I said in our first class, I'm neither a blockchain maximalist or a blockchain minimalist. And you all will have a chance through this course to form your own views. But one of the debates is, all right, Emily's [INAUDIBLE]. Is this a good use of economic-- a good use of resources?

But I would note that all strong currencies, strong monies, for centuries have had something to limit the supply. And so now we're doing it electronically and through this mining. That doesn't mean it's the best use. I'm just saying it's another way.

Extracting gold out of the ground is very hard. And in the 19th century, to have big vault doors and security guards with rifles was a way to insure it. And one could even say that having central banks takes cost.

So I think of it as a trade-off of how you ensure a currency as a harder currency to create. But it doesn't mean that proof of work is the best way, which is, of course, then, the setup to the question of, are there other ways to do consensus?

So one other thing is all of this hashing, how is it distributed? And this, I pulled off the internet this morning. You can see these statistics every day. Proof of work and mining has formed mining pools.

And these mining pools come together for simple economic reason that it's so unlikely to solve the riddle, solve the puzzle of mining, that if you can only solve it once a year, or maybe even once every 10 years, you weren't going to invest in mining. So mining pools started around 2010 to smooth out the revenue.

So if Amanda doesn't want to get it once every 10 years, she might say, well, why doesn't all

the 80 people in this room-- you might, Amanda, I don't know-- say, why don't we all form a pool, and we'll all going to use our laptops. And now this is still 2009 or 2010, when you could mine Bitcoin on your laptop.

But we could say, why don't we all do that together? And then all of us could say, well, Amanda, that's a bright idea, but could you create the Merkle root for us? Could you do some other things so that our computer doesn't have to do all that other fancy stuff?

So then Amanda might say, well, I want to charge all of you a little bit. How about if I charge you 1% of the take? And Amanda would call herself a mining pool operator.

That's what's happened, is basically the economics of mining have clumped around mining pool operators. And the standard fees range from 1% to 3%. That the mining pool operator provides a number of services to the miners themselves, and those services are compensated, as I say, somewhere between 1% and 3% of the returns. But mostly, it smooths out the economics for all the miners. It does some other things as well, but that's the primary.

AUDIENCE: So I also had one question. You laid out the difficulty, as mining has become increasing difficult. And the cost of the electricity-- the break even point has actually become lower and lower in terms of the electricity cost.

So with that in mind, over the next couple of years, if that's the case, people will start to lose incentive in keep doing the mining. And once that happens-- once that happens, [INAUDIBLE].

GARY GENSLER: Well, it could go either way. As we said-- just Bitcoin, this is-- Bitcoin adjusts the difficulty of mining every two weeks. So if there's fewer people mining, the difficulty will go down.

And if you remember, I said you had to have 18 leading zeros. It might go back to 17 leading zeros or 16 leading zeros. And every two weeks, it adjusts based upon the prior 2016 blocks.

Did it average 10 minutes? If it averaged, for instance, 14 minutes, then it will lower the difficulty. If it averages six minutes, it will increase the difficulty. Kelly.

AUDIENCE: So is this where the proof of stake comes in? Does BTC always get their 19% because they have the largest stake in the system?

GARY GENSLER: So Kelly is asking whether this is where proof of stake comes in. What is proof of stake? Anybody who read the Coindesk article? See, you when you hide your first name, I can just

call you US Air Force. It's true. Bo's name card says US Air Force.

AUDIENCE: That's true. So proof of stake is the coins-- all the coins are already dispersed onto the network, and the verification allocation is allocated based on--

GARY GENSLER: So proof of stake is an alternative consensus mechanism. And Bo described it well. But Kelly, it's not related to this chart here. This is all proof of work. BTC has 19% of the hash rate. That means that they literally have about-- if the total hash power on Bitcoin is 60 hexa-- do I have the word--

AUDIENCE: It starts with an e-x-- exa.

GARY GENSLER: Exahash-- then 19% of that, or about 12.

AUDIENCE: That's what they're attempting to do. That's not the stake that they already have.

GARY GENSLER: Correct. Most Bitcoin miners sell their coins. So the coins that are created-- the coins that are created each year are sold into the broad community. Very few miners hold onto their coins for great lengths of time.

I mean, they might for days. That might for-- they might keep some for all sorts of reasons. And as I truly believe, but can't factually prove, a number of the biggest mining pools or miners are in places where they're doing illicit activity. They're getting their electricity for less than what it's really costing on the grid by bad actors.

But nonetheless, they have a choice whether to sell their coins or keep their coins. Got a question over here, and then I want to keep going.

AUDIENCE: If the mining industry is like so formalized, like you've so many pools, what prevents all the pools from coming together and saying that, let's just solve lesser puzzles so that the value becomes lower, and then the charge is much lower for each and every one?

GARY GENSLER: So the question is, is what happens if the mining pools collude and come together, either, as [INAUDIBLE] says, to, let's say, well, why should we have so much mining capacity? Let's, as a cartel-- like OPEC, the oil cartel-- say we should constrain supply and so forth.

I think what constrains that is it's an open system. But it's possible. It's plausible. I think the bigger question, and there's been numerous academic papers around this, is what happens if the mining pools come together and try to do what's called a 51% attack, and try to take over

the blockchain? And that's a more interesting challenge. And we'll talk about that throughout the semester. It hasn't happened as of yet.

Let me talk about the native currencies. Native currency helps do all this. What Nakamoto said, is it was an incentive system. There was an incentive system, but it was also a peer-to-peer way to create a new money.

And embedded in most blockchains, not 100% of them, there is something, I put quotes around it, monetary policy, in essence, that limits the supply of the currency. Not every blockchain has this, but the vast majority do. And when we start talking about initial coin offerings, you'll find some that don't.

But Bitcoin limits it. And I'm just going to say what it is. It's created in a coinbase transaction in each block. It was initially 50 Bitcoins per block. But now, because it's halved every 210,000 blocks, it's just 12 and one half Bitcoins. That's the number of Bitcoins you earn each time if you mine a block that's approximately $75,000 US dollars in value today, give or take, or $80,000, roughly, to mine a block.

The inflation rate for Bitcoin right now is 4.1%. So think-- for any of you that have taken monetary policy courses or financial courses that talk about the Federal Reserve, Bitcoin is growing about 4% a year right now. But it halves every 210,000 blocks. So the inflation rate will go down to 2%, and then later to about 1%, and later about a half a percent. And it caps around the year 2040.

So whether Satoshi was one person or a team of people, back in 2008 they put in place a monetary policy that is hard-coded into the computer base code, and is supposedly going to be there forever to cap Bitcoin at 21 million Bitcoin.

I'm going to throw up Ether just because it's an alternative. Currently, that mines three ETH per block. And the inflation rate's about 7 and 1/2%. It's a different stage of development, different inflation rate.

But there has been a proposal recently to literally-- it was a proposal by the programmers, we should really lower the inflation rate. And if it's accepted, it will be adopted in November.

The fees in Ethereum are largely paid in something called Gas. Gas is just a small unit of Ethereum. What's the small unit of Bitcoin?

AUDIENCE: Satoshi.

GARY GENSLER: Satoshi. So Gas and Satoshi are very similar. Brodush.

AUDIENCE: I think the limit of 21 million has switched to 2140, not 2040.

GARY GENSLER: It's not 2040?

AUDIENCE: It's 2140.

GARY GENSLER: Yes. Wait. You think it's 2140?

AUDIENCE: [INAUDIBLE].

GARY GENSLER: What's that?

AUDIENCE: I also [INAUDIBLE].

GARY GENSLER: All right. So I typed poorly. Aviva, did you have a question? No. Alene.

AUDIENCE: So this is-- Bitcoin is a deflationary currency?

AUDIENCE: Yes.

GARY GENSLER: Well, Alene says is Bitcoin a deflationary currency, depending upon your use of that word. But others have written that Bitcoin is a deflationary currency because it's not growing. If the economy is growing at x percent, and x percent is bigger than how Bitcoin is growing, that would sort of define it as deflationary.

I would just note that those who are fond-- and there's a lot of academic literature. And this goes back decades, if not centuries. For hard currencies, where monetary policy is absolutely formula-driven, whether it's the Taylor rule or other rules-- rule-based monetary policy would have a fondness for what you could maybe put in computer code.

Those who think that humans should be involved, and many people think there is a need for some human involvement, would say this is too hard-coded, and you'd want something where you can modify it and change it. And that would be dangerous in times of war, in times of stress, in times of economic peril. Or like the 2008 crisis, that this would make a crisis worse.

And so the academic literature and the real life reality of the last couple of hundred years of hard currencies, hard monetary policies versus human involvement and some judgment, is

kind of an interesting debate that goes right in the middle of all of this.

AUDIENCE: [INAUDIBLE] question. Would you consider this as being currency? Or is it an asset-- it's a class of asset that has some aspect of a currency.

GARY GENSLER: So the question is, is this is a real currency, or is it just an asset that has aspects of a currency? Ben, what do you think?

AUDIENCE: So there were three roles of-- well, three ways that you could define a currency. It was the unit of account, stored value, and medium of exchange. So I guess this has all of those attributes. So you could call it a currency.

GARY GENSLER: You're saying if it has all of those.

AUDIENCE: It does have those attributes.

GARY GENSLER: Oh, it does. All right. So Ben thinks yes. How many people are with Ben as of September 18, 2018, not 2118? So Isabella's there. Zan, Joaquin? How many people-- Tom. How many people think not? I don't-- I think that you all get to decide yourself.

Mark Carney, who's the governor of the Bank of England, gave a speech earlier in the year, which is assigned in a later class. And he says, I don't think we should call them cryptocurrencies. They should be called cryptoassets. They're not yet evidencing all three of these.

So that's Mark Carney, who I have great respect for. But there are others that say, no, it's evidencing enough. I would say this, though. It's plausible that they will-- and this-- if you take nothing else from the class, it's plausible in my view that they could provide.

You could have digital currency that does not have a central authority. I mean, I think that innovation is there. Whether you call it a cryptocurrency or cryptoasset at this point in time, I leave to all of you. There won't be a right answer to that on any paper you submit. You can use whichever term you think fits your thinking.

I want to talk about the network. We only have 10 minutes. But the network is important. And a lot of times when you talk about Bitcoin and blockchain, folks aren't going to talk about it. But I want to quickly hit eight or nine players on the network.

There are full nodes. A full node is a group-- is a computer, I should say-- that stores the full

blockchain, and is able to validate all transactions. It doesn't have to. It's a volunteer thing. But it can validate all transactions.

A pruning node-- you're not going to read a lot about it, but I just have it. It prunes transactions once they've validated, and they have a certain age. They're saying, all of those early transactions, we're not going to focus.

There's been probably six times the number of transactions that have happened compared to the actual extant transactions right now. So all the transactions that have ever happened, five, six of them have already been used. They're not around anymore. Why do we have to lug it around in our data set? So you could have a pruning node.

Lightweight nodes, which if any of you have a Bitcoin wallet or any other wallet, you probably have a lightweight-- some form of lightweight node, or what's called an SPV node. It stores just those blockchain headers, rather than all this detail underneath-- a lot less storage. But a lightweight node has to rely on the full nodes for verification, because the lightweight node's not going to be doing that on it's own.

Miners-- we talked about miners. I want to just mention, miners don't have to be full nodes Amanda, you're running this mining operation for the whole class. And we're paying you 1% to 3%. But you're Sloan, right? You're probably charging at the high end, right?

So do you think that Amanda, as a mining pool operator's, operating a full node? Tom?

AUDIENCE: Yes.

GARY GENSLER: You're saying it tentatively. But Tom, you're paying Amanda 2% of your fees. Don't you want to make sure she's validating everything?

AUDIENCE: Just lost a whole percent.

GARY GENSLER: What's that?

AUDIENCE: She thinks 3%.

GARY GENSLER: Oh. She thinks 3%.

AUDIENCE: Yes.

GARY GENSLER: But any of you who are just miners, remember Amanda's our pool operator. Anybody who's-- Andrew, do you think you need a full node if Amanda--

AUDIENCE: No.

GARY GENSLER: No. So a lot of miners are not operating full nodes. They've got all those racks of ASICs. They're running all their electricity, and they're paying Amanda to check on it. Don't let them down. Alene?

AUDIENCE: I think this is a terrible thing. Because in principle, we have 20, 30 mining pools, which means you have 20, 30 computers which validates the newly proposed block. And the thing that everybody wants to believe is that these systems are decentralized, and you have thousands of contributors.

GARY GENSLER: So I might start calling you Nouriel Roubini, but I won't. But Nouriel Roubini, who's an economist-- he's sometimes called Dr. Doom. And he likes that phrase, because he caused downturns in the markets.

There's a later reading, and I might have even done a video of-- Roubini has this view, that it's not decentralized, and mining pools are an Achilles heel of the system.

But the full nodes, the 10,000 nodes, actually still do validation. And there's an interesting social construct where there's a lot of nodes doing, in essence, noncompensated work validating transactions beyond Amanda. And any miner-- Andrew could validate if he wants to. So there's a lot of free riding that goes along, and the economics of free riding.

And then there's wallets, which probably 30 or 40 of you have on your computer somewhere. They store and view and send all the transactions. But importantly, also create the key pairs.

So a lot to cover. There's one that's not a node itself at all. It's called the mempool, or the memory pool. And we're going to talk about this more on Thursday. But the memory pool stores all the unconfirmed, but yet, already validated.

So they've been validated by somebody. A transaction goes out into the network. A full node validates it, and it's put in a memory pool. And Amanda grabs the memory pool, and sends it out to everybody in this class in a block. And then we mine it.

We're not going to chat much about these right now. But we talked about-- there was a

Coindesk article. How many of you actually skimmed it, looked at it? Did it mean anything to you, or just meant that there's some alternative? What's that, Prya?

AUDIENCE: Just that there's some alternative.

GARY GENSLER: There's some alternative. That's what it meant to me the first time I looked into this about six months ago. But I want to just mention what the alternatives come down to.

They generally randomized or delegate the selection. So rather than saying any one of 10,000 nodes can prove that this works, they use various mathematical means-- randomized or delegated. And sometimes they do a little bit of delegation and randomization to pick who's going to validate the next block.

It all comes down to who is picking the next block. Is it by Adam Back's sort of, as Satoshi Nakamoto put out there, we'll call it Nakamoto consensus, proof of work? And you'll have a paper for Thursday that talks about Nakamoto Consensus, the Clark paper. But-- or is there some other randomized, delegated way to do it?

In some of them, they have a second check. If there's a delegated person to validate something, they put a second check in there that there's another group that officiates and says whether it's correct. And so there's proof of stake, which is based on the stake you have in the underlying currency. There's proof of activity, which is kind of a hybrid of proof of work and proof of stake.

Proof of burn-- are you willing to give up coins? Proof of capacity-- do you have storage capacity? And you might have a tiered system.

The major permissionless blockchains all use proof of work. And the reason is nobody's really solved-- all of these other alternatives, no one's really solved for a couple of problems in them. But they usually find a way to be more efficient through delegation and randomization, and might have a backup set of checks on it.

DASH and NEO will say that they use proof of stake. But they're actually kind of using some form of masternodes or set up professional nodes. But DASH and NEO are kind of, I think, the 13th and 15th largest market value cryptos, which means everything else kind of--

And Ripple doesn't-- Ripple's really almost like a permission system rather than permissionless. I mean, they would say they're permissionless. But it's a confirmed set of

nodes in the node system.

So that's it for today. We're going to do transactions on Thursday. I moved that study question. And then I'm going to ask you to read through the Clark paper, which is really the academic pedigree. Where is this built on? What's the background? But I think it's a good way to bring it all together.

And remember on Thursday, you all come in to answer this question-- your own view as to who Satoshi Nakamoto is. There is no right answer. But if MIT's Blockchain and Money class can answer that, I'm sure that we'll get a write-up somewhere about it.