REDBELLY

AMA Transcript

6 December 2022

REDBELLY

Technical AMA 06.12.2022

Tuesday, Dec 06, 2022 8:03PM AEDT • 43:51

You can find the full recording on our Discord Channel within the Post-AMA-Info sub-channel

SUMMARY KEYWORDS

blockchain, consensus, transaction, network, participants, vincent, questions, users, validators, miners, problem, node, block, data, outcome, chain, governance, typically, machines, run

00:09 [Leighlah Ashmore]
Hello, everyone. Thanks to everyone that has joined us today for Redbelly's first ever technical Ask Me Anything with our founder and CTO Professor Vincent Gramoli. By way of introduction, I'm Leighlah Ashmore, and I'm Head of Partnerships and Community and I have the pleasure of moderating today's session.

So to set the scene, the format for the AMA today will be broken into three sections. Section 1, I'll be interviewing Vincent so that we can hear more about him and Redbelly and the journey up until now. Section 2, we will be answering some of the questions that have been raised directly from our community. And a big thanks to everybody that submitted questions. We've selected 5, and we'll answer those today. And those that submitted the questions will each receive one of the 5 x $25 USD prizes, and the final section will be an open Q&A forum.

Please remember that this is a Technical AMA. And the questions are for Vincent only. Please feel free to add your questions into Telegram and we'll attempt to work through them as much as time permits today. But rest assured, if we haven't answered your questions today, and other questions raised within the community, we will post our responses and we are recording today's session. And so we'll post that recording as well. So, without further ado, let me welcome Professor Vincent Gramoli.

02:11 [Prof Vincent Gramoli]
Hi Leighlah

02:13 [Leighlah Ashmore]
Hi, Vincent. Thanks for joining. So can we possibly start with an introduction to yourself and maybe a bit of an overview or background into Redbelly, please?

02:25 [Prof Vincent Gramoli]
Yeah, sure. So, I'm Vincent, Founder and CTO of Redbelly Network, and my background is Academia. I've been doing research and teaching for quite some time, I've been working, in particular, on the consensus problem, which is fundamental to the distributed computing discipline. And that dates back even before blockchain started with Bitcoin. And so I've been doing that, working on this consensus problem for the last 15 years or so at different places, working at INRIA in France, but also at Cornell University in the US, before moving to the Swiss Federal Institute of Technology in Lausanne, Switzerland. And finally, I ended up getting this professor position at the University of Sydney 10 years ago. And when I decided to come to Australia, I had the chance to teach very good students at the University of Sydney and I started working with the Commonwealth Science Industry Research Organisation, CSIRO, here in Australia. And we were looking at blockchain technologies.

That was back around 2015-2016, when we started working on trying to understand how the existing classic blockchains were working. And because we suspected that they had to solve an interesting problem that is consensus, on which we had been working for quite some time, and that we knew was difficult to solve. And then, after some time, we discussed with industry partners, including Arcery, which is a consortium of financial institutions, but also CBA, a large bank here in Australia. They were both using Ethereum in their private testnet and lab environment. And then we realised that there were possibilities to do double spending in their configuration settings, because of the way they were deploying Ethereum in their private networks. And that served as a very important motivation for us to try to come up with an alternative solution that would be targeting high value assets with deterministic guarantees, and that's what will lead to the publication of the consensus algorithm back in 2018, but I guess we will talk about that a bit later.

05:05 [Leighlah Ashmore]
Great. Thanks, Vincent. And so can you explain the role of CSIRO and University of Sydney and what that looks like today?

05:14 [Prof Vincent Gramoli]
Yep. So while we were working with CSIRO and University of Sydney, and because we had designed this consensus protocol, we thought, okay, let's try to build a blockchain together. We're using this consensus protocol that was now published as a scientific paper. And then we ended up, you know, doing all sorts of experiments. We ran it in a very large scale environment, involving Amazon Web Services, availability zones around the world, we deployed it in 10 countries over four continents at a time. And then, the funny thing is that AWS felt that we were doing a denial-of-service attack against their infrastructure, because we were using pretty much all the availability zones we had access to. And then we had to call them because they shut down our accounts. And, then when they realised that we were researchers doing experiments of probably the largest blockchain deployment, on cloud infrastructure, they decided to do a PR with us and we ended up having this news article back in 2018, with CSIRO and AWS and from then on, we decided to, and I think that was from what I recall, an initiative from CSIRO wanted to patent the technology. The inventions that were part of this Redbelly

blockchain, given the good results that we obtained. Like at the time, we we're having peak performance of around 660,000 transactions per second, which was quite impressive, but that was in the UTXO model. And, then they decided to patent this, to file a patent application at least. And that was back in 2018. And three years later, the University of Sydney, CSIRO and Block8, created a joint venture here in Sydney to commercialise this technology into what will become Redbelly Network.

07:30 [Leighlah Ashmore]
Great, thank you. So Vincent, can you give us an overview of what you see as the competitive advantages of Redbelly?

07:42 [Prof Vincent Gramoli]
Yeah, there are several of them. So we could start with the inventions that we patented back in 2018. So, a long time ago, but this essentially combines two inventions. The first one is that we took a collaborative approach. If you look at classic blockchains, they typically take a competitive approach. And by that, I mean, when you run a blockchain network in the classic way, you have several machines that will potentially propose different blocks for the same index, right? This is what the miners would typically do, right? If two miners solve the proof of work roughly at the same time in Ethereum or the Bitcoin network, then they will typically propose this block as a candidate for the index that is available in the chain. And then the classic blockchain will have to select one of these blocks, right? This is the way it works in classic blockchains, and this is competitive in the sense that you might have a lot of miners or just a few, it doesn't really matter. At the end, you only pick one of the blocks among the ones that were proposed. And so you only have one winner. And that's what I mean by competition. You have one winner at the end, which is the miner that managed to have his block appended. Our strategy was different because we realised that the consensus was not designed. And even the definition, the formal definition of the consensus problem itself, usually states that one of the values has to be decided among all the proposed values, which is exactly what Ethereum and Bitcoin are doing. And we thought, okay, we really want to scale. Of course, consensus was defined for small networks, right? Originally. But if we want to scale and have a very large network of participants, and have a lot of transactions being committed at the same time, then we should better have a collaborative approach. And by that, I mean that you typically would like to leverage the resources of all the participants. So if you have four validators or miners, it doesn't really matter. But as soon as you start having hundreds of validators or miners, if you can combine blocks coming from all these different validators, and miners, that are potentially, you know, having disjointed sets of transactions, then what you can do is you can commit to orders of magnitude more transactions than what you would do with a competitive approach. So that's the first innovation that was patented.

The second one is around verification. So there is something that is extremely costly, and because we're, you know, we were working distributed system, we were not really aware of the cost of signature verification at the time, so we were surprised to realise that it was slowing down our blockchain technology. So we had to come up with a novelty, which was the idea of what we call sharding, the verification but it's a bit of a misnomer, in the sense that it's, it's not the sharded blockchain. But what we decided to do is to reduce the amount of work that will be done by each validator in terms of verifying transactions. So in classic blockchain, every miner will typically validate every single transaction, potentially twice. Here, what we do is we tried to minimise the number of validations per node per transaction. So these are the two innovations that we had at the beginning. But of course, we have, we have continued the research, we came up with this definition of accountability within the consensus protocol, which is, we believe is interesting as well, when you want to transfer high value assets, because it raises the level of security that you have in your system through incentives. The idea is to be able to generate undeniable proof of fraud on the fly, while your consensus protocol is running. And the idea is extremely simple. It's like forcing all the participants of the consensus to sign their messages. And if you detect that they're lying, because they want to influence the outcome of the consensus in order to double spend, then you would typically create a proof of fraud, which is the concatenation of two signed messages that are conflicting with each other. And then this allows you to, you know, kick out the malicious participants or to slash their coins, for example. So accountability was a big thing.

And we also worked on the privacy, this is something that is extremely recent. We have come up with ensuring that the data of the users of our blockchain technology will remain private, and will not be shared with the rest of the system. Through something that is called verifiable secret sharing, this is an old technique that comes from cryptography, but we have applied it in the context of data privacy. And, now we're really focused on the product that is around the identity concept, because we believe that it's very important to have users that have custody of their data. And, and so we're innovating in this area, and the privacy is one step towards that direction.

13:24 [Leighlah Ashmore]
Thanks for that, Vincent. I think we'll be touching on identity a little bit later as well. So look, I know you've been travelling a lot recently. Can you share more details on what you've been doing during your travels and the conferences that you've attended, please?

13:42 [Prof Vincent Gramoli]
Yes, so, I went to a few places. I went to Salerno in Italy, to present at The Principles of Distributed Computing. This is one of the largest or predominant distributed computing conferences, where, you know, researchers in distributed computing will meet on a regular basis, like yearly basis. And I went there to present the formal verification of our blockchain consensus. This is joint work with people from Europe that are experts in verification, I'm typically not an expert in verification. So, but we teamed up because we wanted to verify a distributed algorithm.

And so, we ended up presenting the model checking of our consensus algorithm, this is also something that I presented in Augusta, Georgia in the US recently. Essentially the idea is that we came up with a mathematical representation of our consensus protocol, that was written as a threshold automata. And we also, defined or specified the consensus problem in linear temporal logic. And this allowed us to feed a model checker, which is a software, not only was the problem that we were trying to solve the consensus problem, but also our potential solution. And the role of the model checker was to check that in every possible execution or algorithm, then the properties of the consensus would be fulfilled. And the model checker output was that it was correct. And so that was very interesting outcome, because most of the time these consensus algorithms are too complex to be formally verified, especially when you look at the blockchain consensus, there might be some very simple consensus algorithm, from the old days that could be formally verified, but usually when you take more than one that is quite efficient, the number of states that you have in your threshold automata prevents you from verifying the property that it ensures.

So, we managed to do that, thanks to some decomposition that we did around the threshold automata to reduce the space. But in the end, it's very comforting because we know that in consensus algorithms, you have a lot of mistakes, and for example, when we published our consensus algorithm paper, there were nine pages of handwritten proofs. And of course, as humans, you know, we always make errors. And so it's very reassuring to know that there is a software that did this automatic verification. We're not excluding the risk of adding errors, because we, you know, to make sure you need to verify the architecture on which you run the model checker, you need to verify the model checker to make sure that there are no errors. But seeing that the model checker that was implemented by verification experts returned that the outcome was successful is very reassuring for us and for the level of security of a blockchain.

So that was one thing, another thing that we mentioned was, that I presented was at the conference organised by Stellar. I think it's the Foundation of Stellar, that organised this conference. So they invited me to talk about TPS. So TPS is the acronym for Transaction Per Second, it has been used a lot in the blockchain space, to refer to various things. So I took a scientific approach, trying to define which kind of TPS was interesting. That's the throughput. It's not the sending rate, for example, when you talk about blockchain, but also it's not the only metric that is of interest, right? It's also very important to combine the throughput with latency, for example, because you can reach very high throughput, but if you have to wait a year for a single transaction to be committed, then the service is, or the system is not very useful.

And then I presented the work that we've done recently with the Swiss Federal Institute of Technology. In particular, a student was coming from Technical University of Munich, that was around a benchmark that we did: so a lot of people were involved from the University of Sydney EPFL. And, so what we did is we developed this benchmark, in order to test the performance, in terms of throughput and latency, of the modern blockchains on the same ground, and so we took these blockchains: Algorand, Solana, Avalanche, Quorum, Ethereum. And that's pretty much it, I guess, Diem, which is the new Libra, or what's called Libra before now it's Diem to test them on the same ground using decentralised applications with traces taken from centralised application like exchanges, mobility service, games, and so on, so forth. And, we really deployed them in the exact same configuration settings across different scales, just to make sure that we could compare the performances. And so I presented that and also, a student of mine at the University of Sydney went a bit further and used the benchmark to evaluate

Redbelly blockchain as well, which gave, you know, incredible results when we did the experiment at the end. So that was very interesting outcome. We presented all these results at the conference. And that's, that's pretty much it.

19:46 [Leighlah Ashmore]
Excellent. And actually, if you check out some of our community channels, we have posted some of the presentation decks and videos from some of these conferences as well. May not be the best quality but hopefully the underlying message gets through. So I'm just mindful of time here as well, Vincent. I've got a question here about “can you tell us more about some of the recent developments in Redbelly, and any up and coming milestones you have for the technical roadmap”?

20:16 [Prof Vincent Gramoli]
Yep. So I'm gonna go through that quickly. I mentioned it briefly. But yeah, we have achieved this privacy outcome in the last quarter, around verifiable Secret Sharing implementation within our distributed Redbelly network. This happened to be quite successful in the sense that we have good performances, and the users can use this to recover their assets, even though they have lost their wallet, and nobody else has the complete information that they used to have. So that was a nice outcome that goes into the direction of Identity. But we have another milestone coming up, we've been working hard this quarter to come up with a Testnet. And that will be ready in January. Along with the TGE, which is the Token Generation Event, that will also happen in January. We have the plan of using the Devnet for inviting people to do some tests. And that's going to happen in April. And finally, we're working very hard to release our Mainnet, which will happen first of July, next year.

21:38 [Leighlah Ashmore]
Excellent. Is it worth taking this time to share details of the growth within the team as well at this time?

21:48 [Prof Vincent Gramoli]
Yes. So yeah, we doubled in size recently, we moved, I think, from 23 engineers to something around 50+ engineers, so we have doubled in size. This is a side effect of our successful fundraising. And we got some investments despite the bear market. So we have been lucky enough to be able to grow in these tough times.

22:28 [Leighlah Ashmore]
Excellent. Alright. Well, that's it for Section 1, we should move into Section 2, which is the questions from the community. And we've selected five. So let's start with the first question from Air Catch Dropper. “Can you please explain more about your consensus algorithm DBFT? And what's the reason behind you choosing this consensus”?

22:56 [Prof Vincent Gramoli]
Yeah. So that's a very good question. In the sense that it's a key differentiator of our blockchain compared to other blockchains. And I think the term democratic, I've explained it, right. It's democratic in the sense that it's, it combines two proposals coming from different validators. And more precisely, if you look at most of the practical,

Byzantine fault tolerant consensus algorithms, they rely on the leader, right, and it's no surprise that they typically select one block out of many or one value out of many that were proposed. It's because the leader, the protocol works in such a way that the leader will try to impose its value to the rest of the network, right? That's the way it works. And if the leader is incorrect, and doesn't manage to impose its value, then another leader will be elected. And again, you know, will retry, this new leader will try to impose its value to the rest of the system.

So it's very competitive. And it's because it's leader based. Now, the leader pattern is also a bit of a problem, if you want to scale really worldwide. The problem is that the leader will typically have to gather messages coming from all the participants, and will have to typically send messages to all the participants. So you can imagine that as I increase the number of participants to hundreds, if not 1000s, then I obtain the network bottleneck on the network interface of the leader machine. This is something that is well known. And because our democratic consensus algorithm doesn't have any leader, then we don't have this bottleneck, right? So instead, we're using as many routes as you might have, between any pair of computers that are participating in the consensus. And that's the way you better balance the load across all the routes, and you better exploit the bandwidth.

Yeah, and the last thing is that DBFT is also deterministic. And that's a big differentiator with a blockchain that would have a probabilistic outcome, right? So if you run for long enough a blockchain, then you may expect that if it has an epsilon chance of failing each time you append a block, then eventually, you know that it's going to fail with probability 1. We took a different approach because we wanted to get high value assets. And here, we wanted to make sure that we have deterministic guarantees. So, that is also a difference.

25:26 [Leighlah Ashmore]
Excellent. Another question that was raised by Fukuyama Satoshi, “If Redbelly is fixing transaction fees, then can you explain how you're planning to prevent spamming please”?

25:40 [Prof Vincent Gramoli]
Yeah. So that's also very good. Very good thing. Very good question. And it's true that if you have completely fixed fees, then you expose yourself to flooding attacks or denial of service attacks. The idea here is that we have a mechanism that it can trigger when we detect the denial of service attack, right. So I think there is no solution against denial of service attack, to be honest, right. So there is always a way to spam a service beyond its capacity. But to mitigate it, at least, what we try to do is to accept transactions at a fixed fee. And if we see that there is an abnormal behaviour, I don't know, you know, millions of transaction per second, things like this that are extremely high compared to what you would expect normally, then what you immediately start doing is you increase the price of each transaction, which means that most of the transaction will be immediately dropped. And then the attacker wants to really flood the system will have to pay a very high price, which disincentivizes him from doing that. And then all these transactions that were associated with regular fees and are that are coming from correct, clients will

be in queues. And then the backlog will have to re-execute when the flooding attack will stop. So what we will see is that there will be some, maybe, increased latency during this period of flooding attacks. But we believe that in the end, it's better than crashing the entire system, or raising the fees for the current participants. And eventually their transaction with the regular fees will be persisted.

27:38 [Leighlah Ashmore]
Excellent. Now this is one of my favourites. And this came from Till'D'End, "In light of recent data breaches and hacks on Australian enterprises, although not exclusive to, will Redbelly hold personal data and business data, and how will you protect and secure the data from future breaches? I know you touched on this briefly. But can you spend a bit more time on that, please?

28:06 [Prof Vincent Gramoli]
Yes, yes. So it's clear that we need Redbelly blockchain to hold some data. But we will try to hold this data in an encrypted form, right. So we will not have personal data in clear text available on the network. And this is of paramount importance because we want to, of course, to incentivise users to use our network. But also we want to liberate the decentralisation of the blockchain that is an inherent property of our technology. You're probably referring to the Optus data breach or the Medibank data breach that happened recently. So the problem with this situation is that Optus is a centralised organisation, and as soon as it collects a lot of data from a big set of users, like millions of users, then they become a honeypot. Because a hacker can intrude in a single system, the Optus system, and actually make use of this data and sell it. In a decentralised system like ours, we don't do that, we don't centralise the information. So, there is no such honeypot and we will make sure that, you know, this information will not be available at the validator nodes. Instead what will happen is that the user, because they use a mobile device, for example, and or its personal device on the network, and then what would happen is it would keep this personal information like Driver's Licence on his own devices. And he will not share it. But it's true that there will be some pieces of information that will be private. But this is in order for the user to recover its assets. So for example, if you don't have a solution to recover your private key and you lose your wallet or your mobile device, then you lose all your assets, right. And that's also a big issue. So the middle ground solution that we have found is that we're going to have users (and that's done automatically by the wallet right now), chopping their private key into chunks, and then distribute the chunks on the machines of the network. The nice thing is that it's a mathematical process. That is called verifiable Secret Sharing. It's actually the Shamir Secret Sharing Scheme. And that will guarantee that none of the nodes of the network will have enough information to retrieve, to get close to retrieving the secret, at all, right? And only the user who can get sufficiently many of these chunks back will be able to recover its private key, the secret, in its original form. And so you see that you're distributing some information. But this information cannot allow anyone to retrieve the secret by hacking even a fraction of the entire network. So it will cost a lot of money for an attacker to try to steal users' data, because you will have to attack every individual mobile device of all the users.

31:29 [Leighlah Ashmore]
Excellent. Sorry, Vincent. I'm being pushed around for time a little bit here. So if I can jump straight into the next question, “Will Redbelly have a multi chain function?” And that question came from Heinz Elman.

31:44 [Prof Vincent Gramoli]
Yes, so I'm not sure about the definition of multi-chain here. But I think our sharding plan, it relates to this idea of multi-chain. In some sense, what you have in Ethereum 2 is, you know, beacon chain with what we can call shard chains. This is similar to what we're developing, in the sense that we will have a main chain, and then we will be able to spawn shard chains. The only difference is that we do it in a dynamic way. So whenever you have some participants who maybe require some privacy or don't want to store their entire data, or want to constrain their data in a specific jurisdiction, for example, they will have to spawn a shard chain with individual machines in this jurisdiction, and then they will store this data exclusively on their machines. But this shard chain will be linked to the main chain, because if these participants want to do that they need first to deposit some assets that they can use in the shard chain. And after some time, if they want, they can close the shard chain. It's also good for scalability and to mitigate denial-of-service attacks.

33:07 [Leighlah Ashmore]
Yep. No, that's good. Alright. The fifth question that wins a prize is from dodcrypt, “What are the specs to run a node and with the potential to have 600,000 TPS latency becomes an issue. How is this accounted for in the protocol and the node selection?”

33:29 [Prof Vincent Gramoli]
Yeah, so the specs of the node. At the moment we have run Redbelly blockchain on machines with not a lot of resources, right. So we're not really exploiting special Intel instruction. Or, like Solana, for example, we don't have these requirements, which run the blockchain on what we would call normal machines. But for the sake of, you know, having enough resources, we recommend the node operators to run. I think it's eight vCPUs with 16GB of RAM on each validator node.

34:14 [Leighlah Ashmore]
You are interested in running a node, we'll be releasing the Node Runner Program also in January. So please register your interest.

34:24 [Prof Vincent Gramoli]
Yeah, and I’m not sure who this is related to the latency being an issue. I mean, from what we could, we could observe, you know, we run this with these type of machines. We ran a very large experiment across five different continents where we ran decentralised application with the NASDAQ workload, which is exceptionally highly demanding at 9am Wall Street time, and couldn't see any problem with our blockchain. That's the only blockchain we tested that could commit all the transactions. So we're quite confident that this will not impact, the latency much regardless of the application we're considering.

35:08 [Leighlah Ashmore]
Excellent. Well, they're the five community questions. Thank you for raising them. And we'll be in touch with regards to the prizes. Let's move into Section 3, which is the open Q&A. So if you have some additional questions that we haven't answered, please put them in the chat and Vincent will select some and answer them now.

35:33 [Prof Vincent Gramoli]
So I don't see any questions. But I know that we have received other questions as well. I have private questions that people are asking whether I'm from Sydney University. I guess I answered that already. It's a different channel.

Oh, yeah, "Hacking is not a big deal at this moment, recently FTX got the rug and hacked also and few blockchain projects also face same situation. So what's your plan for this type of situation?"

Yeah. So I think the FTX bankruptcy or dramatic losses is due to the way FTX was centralised and managed. We take a different approach. We believe that a proper governance that is fully decentralised and carefully decentralised should be the one that will dictate how the assets or how the network will be running, right. So the governance will typically have a right to say about the software upgrade. So if we want to change the protocol, then the governance will kick in. It's a very different approach from what Ethereum is doing. For example, if you want to or classic blockchain, if you want to upgrade Ethereum, then you would typically hard fork right, because you will always have miners who will decide to keep the old version, the way we do the upgrade here is we first have to agree, like as a validator of the Redbelly Network, you'd have to agree that there might come a day where the governance will say that we have to move to version 2. So you agree to this plan when you run Redbelly blockchain on your machine. And so what will happen is that either your software will stop working, or your software will upgrade to this new version independently of your will, right. So there won't be any hard fork each time we will do an upgrade, because it will just be part of the protocol and it will be included in the smart contract and the governance will invoke this reconfiguration automatically if sufficiently many votes approve this upgrade. So we really want to have this decentralised governance to make sure that we don't have a misuse of the power of a single person, for example.

I think I didn't answer the previous question. "The Redbelly Network is designed to focus on accountability. Accountability is enforced at a protocol level through a novel mechanism that constructs undeniable proofs of fraud -Polygraph; and at the functional level through an innovative identity layer that ensures all network participants are known." Could you please tell more about Polygraph and how exactly does it work?"

Yes, Polygraph is exactly the algorithm I was referring to. Which is extremely simple. In some sense. The biggest problem was to formally define the problem of accountable consensus, but the idea is, you force everyone to sign messages that could impact the outcome of the vote, or the outcome of the block decision. And you just ignore the messages that are not signed. Now, if someone lies and tries to double spend by trying to influence different nodes to choose different blocks that are conflicting, then it's very likely that some nodes will see the two messages that are in conflict and coming from the same person. And then it will concatenate them. And because there is a signature, this will consist of an undeniable proof of fraud that it can show to the rest of the network to kick out this malicious participant for example. There is a research paper that is published that is available on our website as well. There should be a preprint online.

39:43 [Prof Vincent Gramoli]
"What's the main difference between polygraph and a DAG?"

Yeah, so Polygraph is not a DAG in the sense that all the participants are running the consensus altogether. The main difference is that it solves the accountable consensus problem. So if everything goes fine, and you have less than a third of malicious participants, the coalition is not big enough to prevent you from reaching consensus, then you reach consensus. Now there is a case, in general networks where you cannot solve consensus, and you might expose yourself to double spending in a blockchain. And this is when the coalition exceeds a third of the entire set of participants. And when this happens, the polygraph guarantees that eventually all the …sorry, strictly more than a third of the malicious participants will be detected, right. So the proof of fraud will be generated, exchanged, and the malicious participants will simply be kicked out or their coins will be slashed. So it's not a DAG in the sense that you still have a total order of the transaction as opposed to a Directed Acyclic Graph, which is a DAG, where essentially, you have a partial order on the transactions. The good thing about the total order is that you take the last transaction and then you have the complete history of what happened. If you take a DAG, then you have to explore the tip of the DAG to retrieve the current state, which is, it takes a bit of time.

"Is it more of a block DAG like Kaspa, then?" I'm not familiar with Kaspa? Are you referring to Casper the Friendly Finality Gadget? If this is it, I think it's quite different from Casper, Casper from a paper I read on archive by Vitalic Buterin and others, I remember that they're not solving the same problem that we do like this deterministic consensus with liveness guarantees in a partially synchronous environment. They have different assumptions. And then what they need to have is a supermajority that will vote on the same branch or something similar.

"Ghostdag" Yeah, so yeah, Casper is probably based on ghost or one variant of Ghost. So yeah, that's exactly this algorithm, I guess. And, the way it works is that you have many branches, and then you have to vote for some branches. So you have a partial order, and eventually you hope that there will be convergence towards one branch. We don't do that. We just have one single branch all the time, which defines the total order among transactions.

42:42 [Leighlah Ashmore]
Excellent. Well, Vincent, it seems as though we're out of time for today. But I'd like to thank everybody for joining us. As a team, we're incredibly excited about the journey ahead, and we look forward to the continued support from the community. We will post

the details from today and any unanswered questions. But if you have some more, please feel free to put them in the channel.

I mentioned the Ambassador Program that we launched last week, and we've got the Node Operator launching in January. But if you have some specific areas of interest, then please write in the channels or send an email to info@redbelly.network.

But thank you Vincent, and the rest of the team for joining today and the community. Enjoy the rest of your day or your evening. And we wish you all a Merry Christmas and we look forward to 2023 when Redbelly will be launching to Mainnet with lots more exciting news to share.

43:43 [Prof Vincent Gramoli]
Thank you to the community and thank you Leighlah. Thank you all. Bye.

43:47 [Leighlah Ashmore]
Excellent. Thank you.