Demetri Kofinas: What's up everybody. Welcome to this week's episode of Hidden Forces with me, Demetri Kofinas. My guests for this episode are Vitalik Buterin and Vlad Zamfir. Vitalik needs little introduction. He is the founder and inventor of Ethereum, the first cryptocurrency enabled decentralized touring complete [00:00:30] machine ever created. Ethereum and Bitcoin combined to make up nearly 60% of the entire market cap of all cryptocurrencies.

Vlad is one of the Ethereum most prominent researchers and the leading figure in the development of Casper, a consensus mechanism that aims to enable the Ethereum platform to scale its existing architecture for use cases that reached beyond the networks current capacity, like decentralized car sharing applications, stock markets, and online games.

[00:01:00] In this conversation, we focus our attention on the Ethereum roadmap, specifically Casper, Plasma and the developer communities approach to shorten. I also give Vitalik and Vlad's reactions to the recent comments by ACC Director of corporation finance, William Hyndman, as well as their thoughts about governance models in open source crypto economies.

For more information about today's episode or if you want easy access to related programming in Blockchain and [00:01:30] cryptocurrencies, visit our website at hiddenforces.io, and subscribe to our free email list. You can follow us on Twitter, Facebook, and Instagram @hiddenforcespod for regular updates and audience feedback, including the latest information about future episodes, topics, and guests.

Now let's get right to this week's conversation. All right guys, welcome to Hidden Forces.

Vitalik Buterin: Hi. Hey Demetri. Great to be here.

Demetri Kofinas: It's great having you on.

Vlad Zamfir: Yeah. [00:02:00] Good to be here too.

Demetri Kofinas: So we were just talking about some libertarian history. You guys are fans of Doug Casey.

Vitalik Buterin: Yeah. Well, when I was in high school, my dad would send me Doug’s “Conversations with Casey” – those monthly newsletter – pretty much every month, and I remember this had to be somewhat convinced that the US dollar would actually hyperinflate, and the financial system would collapse and this actually, to some extent literally had me scared that I would basically not be able to survive unless [00:02:30] I try really hard to keep myself very useful and this motivated me to learn a bunch of stuff.

Demetri Kofinas: That's really funny because Doug's a funny guy, and you Vlad, you were looking for Doug when you found capital account.

Vlad Zamfir: Well, I mean I was – I kind of followed Doug on YouTube. I thought he was among the better among the like libertarian sphere that I used to kind of follow on the Internet. I think a lot of people in cryptocurrency still are in that sphere. I think over time I've kind of drifted a little bit away and become a little bit more maybe realistic and also absurdist.

Vitalik Buterin: I'm [00:03:00] not sure the libertarians sphere, like as it existed in the late 2000 sort of exists today. I feel it's perfricated in a bunch of directions. There is some that kind of moved into the cryptocurrency space. Some that, in my opinion, very unfortunately, turned into alt-right people. There's all of these different subtribes.

I mean there was obviously a lot of people who believe in the same values, but the strategy is the way that they interpreted the [00:03:30] financial crisis, the way that they interpreted cryptocurrency, and the way that they interpreted Trump and all of those things. I think movements will always exist at a time and place and then sort of evolve as the time and place changes.

Demetri Kofinas: That is true. That is interesting. It is interesting, like you said, how it's fragmented into subcultures. How initially the cryptocurrency/blockchain space – I mean I guess it wasn't even conceptually thought of as Blockchain – was heavily influenced by anarchism and libertarian ideals and how that sort of changed. Certainly, [00:04:00] I think that's much more the case with Bitcoin than it is with Ethereum.

Vlad, before we start, you got to tell our audience how you and I first met, actually, virtually met because it wasn't in person. I'll tell you from my perspective. get on Facebook. This was months after I had been aware of your work. I had been reading about Casper, and I get on my Facebook page, and I see a message from Vlad Zamfir, so I'm confused here, how it is that I'm getting a text message from Vlad and it turned out this was from a couple of years ago when I had published a story, and you had messaged [00:04:30] me, and you used to watch Capital Account, so you had learned about what I was doing from Capital Account, and you liked that show.

Vlad Zamfir: Yeah, it was my favorite show back in the day. I still remember being really upset when you cancelled it.

Demetri Kofinas: That's amazing. Thinking about all the work that you've been doing with Casper and proof-of-stake, and the amount to which economic modeling goes into thinking about that. The fact that you used to watch that show was very humbling.

Vlad Zamfir: I was into global finance and stuff and that's actually the kind of reason I got into Bitcoin in the first place because I kind thought, "Oh look, the global financial system is [00:05:00] really fucked up and Bitcoin is going to save the world." I kind of stopped following media after I got into Blockchain. But definitely, you know ---

Demetri Kofinas: Anyway, that was a selfish plug for Capital Account and what I'd done there, because I was so proud of it and it made me so happy to learn that, and just really cool point of history. Guys, I want to get right into this because there's so many things that we want to discuss. You just flew in from Singapore. How was that?

Vitalik Buterin: It was two long flights. [00:05:30] I mean, we travel around the world all the time. We're used to it.

Demetri Kofinas: Anything interesting to report? Have you guys been doing any consulting work for the government in Singapore? I know that they have been looking at the possibly implementing blockchain for their local currency.

Vitalik Buterin: I've definitely talked to various people in the Singapore government internally on token related issue is and other things. But so far it's been fairly informal.

Vlad Zamfir: I'd say like, very preliminary.

Demetri Kofinas: All right. [00:06:00] Before we get into the two kinds of areas that I want to discuss, which I think are most interesting. Certainly to me they are most interesting, and I think they are to the audience and they're certainly points of confusion for me as well, and they are basically the domain of architecture – the technology in other words – and the domain of governance. But before we do that, I do want to ask you, because it is timely that we're recording this on Sunday, June 17th, a few days ago whenever it was that the ACC director of corporate [00:06:30] finance, William Hyndman came out with a statement around a crypto currencies and securities. Do you feel like his statement provided any clarity for you and if so, what was the impact?

Vitalik Buterin: It definitely was perceived by the market as a providing clarity or that at least some cryptocurrencies are kind of safe in some sense and we have peace in our time in the crypto industry. The extent to which the situation [00:07:00] is going to evolve over the next few years is of course still up in the air. These are nonbinding statements and they totally have the ability to reinterpret it if they want to and you're like – there's going to be a change of administration in two and a half, six and a half, whatever years. So, there are a whole bunch of unknown variables. But, like, the way the markets interpreted it is definitely at least, I feel, as a kind of reduction in uncertainty.

Vlad Zamfir: I thought it was encouraging to see that [00:07:30] they kind of understood the difference between a token and the agreements between the parties involved and that the ... as opposed to just imagining that if someone does a token sale, then no matter what happens in the future, that token will always be classified as a security.

Demetri Kofinas: Do you think that's less of a concern for the two of you and for the core developers than it is for the distributed app developers and people that are sort of entrepreneurs in this space trying to build and raise capital?

Vitalik Buterin: [00:08:00] Possibly, yeah. Because there's different categories of coins and there’s base layer protocol, cryptocurrencies like Bitcoin, Ether, Zcash and so forth. That's one category of thing. But I feel right now the barrier to entry in terms of being able to create a new base layer that actually gets used by a lot of people are just going up and up because people are expecting ultra-fast virtual machines. They're expecting [00:08:30] scalability. They're expecting privacy. You can't just make a clone of bitcoin, crank up some parameters and push it out.

But there's also this other category of cryptocurrencies that's more like application coins and these application coins tend to be ... not always but sometimes a more centralized in nature. They're also designed around particular products and particular systems and often enough from an economic and [00:09:00] possibly even regulatory standpoint, they lean closer to being securities and that is a kind of different subcategory of cryptocurrencies that has grown a lot over the last year.

I mean the SEC messages definitely don't give a kind of unconditional safe harbor to that category. They absolutely don't, but even still this is an area that people ... there is substantial demand to participate in and it's kind of evolving over time.

Demetri Kofinas: So you feel [00:09:30] pretty optimistic about the direction?

Vitalik Buterin: Optimistic about what. Because there is optimism versus pessimism about regulatory action. There is optimism versus pessimism about whether or not it'll work and work in the sense of becoming big, work in that sense of being useful to society…

Demetri Kofinas: Not optimistic about the technology, but optimistic about the way in which governments, specifically the United States, but more broadly global governments are going to, maybe not embrace, but sort of make room for the emergence of things like Ethereum.

Vitalik Buterin: [00:10:00] On the base layer side. I think it's definitely seeing fairly green lights on multiple fronts and not just in United States. I believe Korea is considering unblocking certain types of cryptocurrency trading though I haven't been following recently too closely. But at the same I do feel government agency is basically ... are going to be scared of basically having cryptocurrency [00:10:30] be a back door through which the entire securities regulation apparatus gets unraveled and so they would want to try to make some kind of barrier that says like, "You can't use this as some backdoor way of basically issuing company shares."

I think ... obviously, base layer protocol cryptocurrencies are very far from company shares but there definitely are cryptocurrencies that are much closer to that. There's going to be back and forth in the middle there.

Demetri Kofinas: You got any thoughts on that Vlad or should we go right into the meat and potatoes of the conversation?

Vlad Zamfir: Well, I'm [00:11:00] not a lawyer, but one of the things that I did learn that I found kind of interesting is that the definition of security has to do basically with this contract between the buyer and the seller and it's not natively to do with the technology that underpins things. I am optimistic that everyone will continue learning about the technology.

Demetri Kofinas: I think it's a situation where we just need to get some clarity around it, because the language of these contracts is important and the language of the way that the money is raised…to your point. So, we can begin either with the technology. [00:11:30] I'm inclined to go on the tech side because we can talk about governance too and I want to get to that towards the end, but let's start with the architecture and the roadmap.

It is very challenging for me to wrap my head around the solutions and the roadmap for Ethereum because of the anarchic nature of the way in which the conversations happen. So, there's you, there's Vlad, there's Joseph, there are a few people, right? And there are three main components to the solution, to the problem of scale that I've heard discussed. One is POS, [00:12:00] specifically Casper, FFG and CBC. CBC, as I understand it is being led by you Vlad. Is that right?

Vlad Zamfir: Mm-hmm (affirmative).

Demetri Kofinas: FFG is something that you and who else have worked on Vitalik?

Vitalik Buterin: I'm the primary kind of creator, but then there's also Danny Ryan who helped a lot on the implementation and Virgil Griffith helped to co-author the paper last year.

Demetri Kofinas: FFG is the friendly finality gadget. Is that a kind of compromise, a transient, temporary compromise between [00:12:30] proof-of-work and proof-of-stake? Is that compromise driven by the politics of wanting to appease the miners? Or is it about –

Vitalik Buterin: It's important to distinguish here between kind of two ideas that sometimes get called Casper FFG. One of them is specifically the way that the Casper FFG algorithm works and the reason why Casper FFG is called a finality gadget. So like, the word “gadget” instead of “algorithm” is that Casper FFG as described and envisioned [00:13:00] in the paper is sort of deliberately incomplete in some sense, right?

It handles coming to consensus, but it does not handle what practitioners and academics call leader election. Basically, leader election or a proposal that general category of things is basically a matter of coming up with proposals for, "let's agree on this." Then the second step as well, if someone proposes “let’s agree on this,” does the entire set of participants in [00:13:30] the network actually come together and say, “yes, we agree."

So Casper FFG handles the “yes, we agree” or “no, we don't agree” part, but it explicitly leaves a kind of open hole where you can plug in any separate algorithm for proposing things.

Demetri Kofinas: For who decides who the leader is that proposes the next block? Is that what you're describing?

Vitalik Buterin: Basically, yeah. In the original kind of vision, the leader would start off being the existing proof-of-work chain. The proof-of-work chain, instead of being the kind of full consensus, [00:14:00] would just sort of become a suggester and the proof-of-stake would become a ratifier.

Demetri Kofinas: But you're saying the suggester is the actual longest chain, or is it actually the minor who won the--

Vitalik Buterin: Well, the suggester is the chain acting as a kind of virtual entity consisting of all the miners, if that makes any sense at all.

Demetri Kofinas: Where does the leader who actually chooses the next block come from?

Vitalik Buterin: Sure. The leader is basically implicitly the miner that creates the check points, [00:14:30] that creates the block that the gadget ends up actually coming to agreements on.

Demetri Kofinas: Okay. So that happens before this implementation of the consensus protocol or mechanism is put in play.

Vitalik Buterin: Well, proof-of-work and FFG both kind of run in parallel and they're both running constantly, forever, but generally, yes. Step one, the proof-of-work chain grows to some height, at which point it proposes a new block and step two that block gets confirmed by the proof stake layer.

Demetri Kofinas: Okay. All right. That's [00:15:00] a good sort of general overview. Vlad what about CBC. Now this is, I think, am I correct in understanding that when people think about Casper and sort of the grandest visions of Casper and proof-of-stake and moving away from proof-of-work, what they're really talking about is correct-by-construction. Casper, CBC, what you've been working on for the last number of years?

Vlad Zamfir: I think firstly people think about proof-of-stake and this idea of using security deposits for proof-of-stake in this kind of blockchain-like consensus protocol. And then [00:15:30] somehow the finality gadget and the CBC, Casper family of protocols are in some sense distributed systems solutions and they are going to be associated with security deposits and economic incentives. But Casper research really started in the economics side of the equation.

Then we moved over time to – after we do a bunch of thinking about what is economic security – how do we make sure that [inaudible] follow the protocol, we have to actually say like, "Oh, what protocol is that we want to follow?" Then we go and build these [00:16:00] different distributed systems solutions. CBC Casper is a family of distributed systems protocols, of consensus protocols.

Demetri Kofinas: It's a family of consensus protocols.

Vlad Zamfir: It's a family of consensus and protocols, yeah.

Demetri Kofinas: You and I spoke about this recently, the one time that you and I had a chance to talk, and I got a sense from what you described that there are multiple different solutions. You apply different types of consensus to different types of problems? Let's see if you can try to explain that to me, because I didn't understand it when you spoke to me before.

I want to try to see if I can understand it now and [00:16:30] provide clarity to my audience where it's possible. Maybe you can just give us the grand vision for what Casper is and how you want to apply proof-of-stake and consensus in order to scale the network.

Vlad Zamfir: Sure. Well firstly, let me say that proof-of-stake by itself is just a solution for Sybil resistance and governing the nodes of the network so that they don't attack the system or if they do that it's expensive. It's not by itself a scalability solution. Although it helps in, for example, sampling nodes. [00:17:00] It's easier to sample nodes and proof-of-stake than proof-of-work and then we kind of need distributed systems that we're going to incentivize this proof of state.

Demetri Kofinas: When you say sample nodes. What do you mean?

Vlad Zamfir: Oh, for example, imagine that there's some shard which is a kind of a subset of the consensus protocol, and you need to make sure that the blocks there are valid and available. You might want to grab some nodes, tell them, "Go check it out," and it's easier to sample nodes and proof-of-stake and proof-of-work because the protocol has their deposits and you can [00:17:30] sample them according to their deposits.

Vitalik Buterin: I think one thing that might help us is a bit of an analogy. So, imagine you have a fairly big and complex country and you want to run this country with something as close as possible to a direct democracy. You're like, "People in this country happens to really completely not like politicians and they just wants to have the sort of direct will of the people decide everything as purely as they can."

One of the problems with this, I mean not the only problem, but one of them is that [00:18:00] basically there's so many possible laws that people or actions that they might want to take and vote on, that there's just not enough time for every person in the country to research all of them. If you try to just get everyone to vote for everything, then you'll basically just get random junk because no one would even know what they're voting for.

But one way that you could improve this is, imagine if for every bill you would basically randomly pick who had say 2000 citizens and it would be only [00:18:30] those 3000 citizens that would be able to vote and that vote would decide whether or not the bill passes. Right now because you're randomly picking 2000, statistically speaking, they're going to have roughly the same opinion as the entire country. At least if they can't diverge by more than a few percent. Or otherwise, it’s still, extremely unlikely.

But now for every bill, you only need 2000 people to spend time learning about it and deciding whether or not to support it instead of the entire group or that entire country. So, sharding is kind of applying that same principle [00:19:00]. Instead of checking bills you're basically just validating whether or not particular blocks are correct and follow the rules.

Vlad Zamfir: And, Proof-of-stake makes it easier because you have a list of citizens as opposed to only being able to tell indirectly through the work that they produced that they exist.

Demetri Kofinas: Well, I want to get into sharding and we'll get back into Casper. Sharding is another one. As far as what I've heard for sharding, I haven't heard anything unique to Ethereum. Just that you guys want to share the database [00:19:30] and please educate me on the details of it, but also, am I correct in assuming that it would be very difficult to do that effectively without finality?

In other words, is it possible to implement just a Sharding solution without actually having Casper?

Vlad Zamfir: Yeah, I mean there are proof-of-work sharding schemes out there, but they're less good by a margin because you can't sample the nodes as well. It's relatively – so, the proof-of-stake/proof-of-work side is relatively minor compared

[00:20:00] to the scaling advantage -- compared to like the distributed systems distinctions between these protocols.

Vitalik Buterin: Another thing that's probably worth pointing out is that there is Sharding in the sense of Sharding as in what database engineers have already been doing for 30 years and then they're Sharding in the sense of blockchain Sharding, and those two are kind of different problems. Blockchain Sharding does also involve some Sharding of the distributed systems or of the traditional database type.

[00:20:30] But the main difference is that when Amazon does Sharding, all of the shards are computers that are run by Amazon and so because all of the computers are run by Amazon, they are run by possibly just the same guy who is in charge of plugging all of the computers in and installing the software. Every computer can trust every computer. If you see a message from one computer, he can just assume that that actually is the result of the computation.

Here we're talking about a permissionless, open public blockchain, and so what you're doing is a [00:21:00] kind of sharding across administrative boundaries. You're basically taking work and you're splitting it up between users where the different users that you're splitting it between are totally different individuals, totally different people, totally possibly different companies and they don't necessarily have the interests of the system at heart.

Basically the sharding scheme has to be able to take that into account and be robust against the possibility of [00:21:30] some nodes being malicious some of the time. And it's that that is a kind of significant challenge and why you can't just take traditional sharding and just port it over you.

Demetri Kofinas: Let's stay with this a little bit. What you're pointing out here, if I understand correctly, is if you shard the database, if you basically partition it into smaller and smaller versions of itself – kind of like a fractal. You're taking one large database and you're breaking it up into much smaller versions of the same thing?

Vlad Zamfir: [00:22:00] It depends.

Vitalik Buterin: It depends.

Vlad Zamfir: [00:22:00] It could be.

Demetri Kofinas: All right, well, that's how I'm thinking about it. The challenge you face there is if you have X number of nodes, you're going to have a fewer number that are going to be available per shard, right? Because, you don't have the security of the entire network validating every transaction, you have each shard right? And then that shard has to reconcile with the entire network at some point.

Vitalik Buterin: It depends what you mean by “security of” because, technically it is true that if you have a million computers [00:22:30] in total and a thousand of them processing one shard and if those exact 1000 happened to be taken offline or attacking the network or whatever at the same time, then yes it is easier to attack the system and then the entire system has to go through some fairly complex recovery process.

But at the same time, that doesn't mean that if an attacker has a thousand computers that they can kind of direct those computers toward one particular shard and take it over. The [00:23:00] system doesn't let you choose which shard you get allocated to. It's kind of how in this hypothetical and direct democratic country, I suggested if you have a thousand neo nazis, they're not going to be able to kind of get together and say, "We all wants to be the ones that are voting on this particular bill." Like, "No, it's a lottery." It's randomly selected and chances are you're not going to have more than one of you voting on any bill at any time in the future.

Demetri Kofinas: When you talk about the sampling, what you're talking about is, does this also includes shuffling? You're talking about how you manage to keep these shards [00:23:30] fresh and pure so that there isn't an attempt by a malicious act on the network to concentrate malicious intent into a particular shard.

Vlad Zamfir: It's a defense against targeted attack.

Demetri Kofinas: That makes sense. That's one component. That's one challenge. That's one thing that you guys have to figure out, obviously if you're going to implement a sharding solution?

Vlad Zamfir: Mm-hmm (affirmative)-

Demetri Kofinas: And you feel pretty good about where you're at with that?

Vlad Zamfir: Yeah. I actually think at this point we're at the stage where we're very confident that our current suite of building blocks and our general way of arranging them totally in principle [00:24:00] works and what's left is basically implementation optimization, testing, narrowing down very fine details, and then in some cases less-fine details, but none of the things that we’re still narrowing down are kind of fundamental to the ability to shard in the first place. That's more about what's more efficient.

Demetri Kofinas: I want to also get to state channels and plasma, but before we do that, I want to stay here and I want to go deeper into where you're at in terms of this Casper implementation [00:24:30] and how confident you feel about your ability to reach finality and how – and humor me here – help me understand how you achieve finality. How you do that, technically speaking?

Vitalik Buterin: I feel like the CBC Casper family and the Friendly Finality Gadget family have slightly different solutions, but they're all kind of consensus protocols and consensus protocols all somehow fundamentally make irreversible decisions out of mutually exclusive alternatives and there's a kind of consensus safety proof [00:25:00] that shows for each of them that if there's not too many malicious nodes then they will all be able to do this in a safe manner.

Vlad Zamfir: I guess the slightly technical but still vaguely understandable explanation I could give people is ... you have this set of participants that are trying to come to consensus on things and if two thirds of the participants basically sign some message that says that they agree and then two thirds of the participants see that each of those [00:25:30] two third sees that two thirds of the others have signed the message. Then they can sign a message that says, "Yes, I see the two thirds of people are agreeing." And once two thirds of people do that, then basically those two rounds of agreements are kind of enough to achieve a kind of a lock on the message where the algorithm has no ability to sort of move away from that message anymore.

Demetri Kofinas: Is there some way for someone who's interested in understanding how this works to understand it conceptually without having to go through…what if they can understand the language of the mathematics and the white paper? [00:26:00] Is there some way for someone to understand conceptually how it is that you guys reached consensus? And, just to be clear with FFG in 20 minutes you have finality, right? And with Casper, with your construction Vlad, what are you looking to get?

Vlad Zamfir: Well, we can explore a big trade-off triangle and so, it depends on like what parametrization, but the promise is that we can achieve theoretically optimally low network overhead or low latency or large number of validators and kind of ... but there's a tradeoff between these. [00:26:30] We need to kind of choose a moderate amount of each or depending on-

Demetri Kofinas: Depending on use case?

Vlad Zamfir: Well, depending on kind of what people want out of the consensus protocol, which might depend on the use case. But for general purpose blockchain it’s going to be a little more subtle than that. But you know, I do want to say actually that in the CBC world we don't have this two thirds number and we don't insist on the rounds thing at all. We have instead some number, some kind of subjective number that different nodes decide on. Just like in bitcoin today, you decide how many confirmations to [00:27:00] wait for and the CBC Casper world you decide how much fault tolerance to wait for. So, if you want to wait for a block to be finalized with a third fault tolerance, that's fine. If you want to wait for it to be finalized with less or more, that's also something you can do. And the more fault tolerance you wait for the more likely that you're going to have to wait longer or that it might never happen.

And if you wait for less fault tolerance then you get less of it, but kind of fundamentally the CBC Casper Consensus family of protocols is not round based. We have this arbitrary DAG that everything is kind of built [00:27:30] on top.

Demetri Kofinas: This arbitrary DAG you said?

Vlad Zamfir: Yeah, I mean we have this message DAG which could look in any real shape as opposed to being round based. There's this notion that nodes make messages kind of in sync and that it's not going to be the one the validator makes a thousand messages while another one makes one.

Demetri Kofinas: How far along are you in your work on this?

Vlad Zamfir: It depends exactly what you mean. There's a lot of work in a [00:28:00] lot of different areas. If you're talking about deployment for a single consensus protocol, for example, a blockchain, it's basically the distributed systems research is almost entirely complete and we should have something running soon.

But if you're talking about sharding or if you're talking about the proof-of-stake stuff, it's a little bit earlier stage, but sharding for me is not inside the CBC Casper Framework and it was kind of related to this question you asked, but we kind of got derailed about how different [00:28:30] consensus protocols can decide on different types of values.

Traditionally people kind of imagined that consensus protocols have to create an order of events, but that's not really fundamentally true at all. Consensus protocols are much more broad family of protocols that can decide.

Demetri Kofinas: Well, you don't need to have ordering. But are you able to do ordering?

Vlad Zamfir: Yeah. Yeah. So when I said like, "Oh, the blockchain thing is pretty much done." That's like-

Demetri Kofinas: So you would use directed acyclic graph within each block?

Vlad Zamfir: No, all of the CBC consensus protocols have a directed acyclic graph [00:29:00] of messages, whether it's agreeing on a bit or an integer or a blockchain. The data structures that we're agreeing on is independent of the message structure. From the point of view of the round or not.

Demetri Kofinas: What would that do for you in terms of ordering? In terms of ordering transactions? [crosstalk 00:29:13]

Vlad Zamfir: If I were to order transactions, then I would have to either come to consensus on an ordering of transactions. I can come to consensus on a blockchain, which concludes transactions. There's a number of ways to order things.

Demetri Kofinas: Seems challenging.

Vlad Zamfir: It's not really. [crosstalk 00:29:27]

Vitalik Buterin: I think that the core message that's important [00:29:30] to get at is that on the FFG side, code has been written, code has been tested, and it's a matter like, there was even a test net running for some period of time between a few nodes and on the CBC side Vlad can talk more but in general, the fundamental kind of theoretical ability to establish the guarantees that we have been talking about and we want to establish as like totally there 100%.

When we're optimizing DAG [00:30:00] structures and message passing, this is fairly far off in optimization land. That's sort of no [inaudible]. If you want to have something minimal than it actually in general isn't really all that difficult. I mean for the kind of moderately math inclined among you, I would encourage people to just read the Casper FFG paper. It's only as about as long as the bitcoin white paper and it's got pretty good diagrams and graphs.

Vlad Zamfir: When I say it's not too hard, I mean it's not a big difference between the different consensus protocols in the CBC family tend [00:30:30] to replicate a bit or an integer or an order.

Demetri Kofinas: What are the biggest challenges for you for both of you and trying to scale the system? What do you feel are the biggest hurdles that you're working hardest to overcome?

Vitalik Buterin: At this point, it's just raw research and development effort I would say. It still takes quite a bit of effort to bring an algorithm from the core set of ideas to [00:31:00] something that has all of the little tiny details that make sure it works and then make sure it satisfies all different properties. And make sure that you can actually join and leave the proof-of-stake system and all of that and obviously we need people to develop it. Obviously you need people to audit it, do all of the security work and there is still a fairly long pipeline both in terms of all of these research optimizations and in terms of actually building the thing.

Vlad Zamfir: But I would say that overall the research risk is actually very low and the difficulty, [00:31:30] the research isn't ... for the most part too hard. The thing that's much harder is finding qualified people who are willing to work on it and are willing to drop whatever they're doing to do it.

Vitalik Buterin: So just to reiterate once again there's a lot of research left, but research risk is low, right? At this point we're not in the stage of figuring out whether or not, say, heavier than air flight is fundamentally possible. It's about optimizing the airplane.

Demetri Kofinas: I'm conscious [00:32:00] about time. I want to discuss plasma. Give me a sort of overview of what that is and what the roadmap for that is as well. Which Vlad is not a big fan of I think.

Vlad Zamfir: Kind of. It depends what you mean exactly.

Vitalik Buterin: Plasma is basically a way to create second layer chains, that sort of plug into a lower level of chain, like Ethereum and what this allows [00:32:30] you to do is it allows you to create applications that derive their security from the security of more decentralized kind of lower level base chain, but at the same time we get higher performance, more throughput, lower fees and better salability from chains that have either a smaller number of nodes or possibly are just entirely centralized servers.

You could have a plasma chain that's between a thousand nodes, between one hundred nodes, or you could [00:33:00] even have a plasma chain which is just run by one server. Even just to take the most extreme case where you have a plasma chain that's run by one server, the benefit is basically that you can kind of send coins into a plasma chain contract and the plasma chain contract is a smart contract that lives on the Ethereum chain and then you would be able to basically trade and move those tokens around, do things with the tokens inside of the plasma chain and in this case, because [00:33:30] it's a server or will be just extremely fast and extremely cheap, but regardless of what the server does, even if a server shuts down, even if the server turns out to be malicious, even if the server gets hacked, if the server tries to do anything that kind of violates the rules of the system, then you will always be able to kind of appeal to the plasma chain contract and get your coins out on the main chain.

Demetri Kofinas: That's a are really good question, because that's what I was kinda thinking about. Let me see if I get this straight. The way I understand channels is, let's say I go to the Corner [00:34:00] Bodega every day and I get a bagel, and I know that I'm going to get about this many bagels a month and I have about this much money in my account and I reconcile that every day or I pay for that. I have this open channel between me and the vendor, the Bagel provider, whatever. At the end of that month I have a certain amount that I reconcile, but then if that's --

Vitalik Buterin: Channels are slightly different from plasma, but I'll try to give an analogy for channels too. Here's the analogy. Well let's suppose that [00:34:30] the only way to make a payment is to basically use checks and cash checks into a bank, but the bank is not a very nice bank and it charges say a $5 fee for cashing checks.

Here's what you do. First of all, the first time you go and buy a Bagel, you write to the Bagel owner a check and you -- say a Bagel costs $3. The check is for $3.

The second time you go and buy a Bagel, you ask the Bagel owner for the first check and you basically simultaneously rip up the first check [00:35:00] and you give them a second check and the second check is for $6.

The next time you wrap up the $6 check, give them a check for $9 and then you keep on repeating this process and eventually whenever you ended up moving out of town, the bagel owner sees that. Okay, this looks like it's finally the last check. It's gotten, I don't know, a $1723 and cashes it at once and pay's only a single $5 to the bank and everyone except for the bank is happy.

Vlad Zamfir: There are some similarities between channels [00:35:30] and plasma as I always like to point out namely that you're doing a bunch of updates, some state off chain and if there is a failure mode, if the customer just stop showing up or if the server goes offline forever, then you go to the chain and you can settle based on the last kind of state that you have so that the chain as the recourse for when things go badly or when the transaction is kind of over right is shared between plasma and channels.

I think [00:36:00] of plasma as a non-unanimous channel that is regularly notarized on the main chain. And, but the reason why this gets kind of controversial is because people are used to channels and bitcoin and think of channels as being always unanimous. But if you drop the requirement that everyone is participating in the channel signs off on every state change, then it becomes very much like plasma. Especially if you also notarized at the main chain regularly so that you don't have a competition between different inconsistent versions of events.

Demetri Kofinas: [00:36:30] These channels are basically like, you open up a separate accounting ledger, and you engage in a bunch of transactions and then you reconcile the last transaction to the ledger when you're done so that you save the trouble of having to do all of those on chain. What happens if there is a dispute at the end of that? This is what we're discussing.

Let's say there's a dispute at the end of that, it turns out that the guy says, "Actually, [00:37:00] you never paid me anything. And this is what I'm going to go back and say." How does that get resolved? Whereas I feel like it wouldn't be a problem if it were all happening on chain.

Vitalik Buterin: Imagine if just to go back to the example with the checks, imagine if every check had a two numbers on it. One of these numbers is kind of a number that refers to a kind of transaction ID. The second number is [00:37:30] a sequence number and suppose that it's illegal for the bank to cash multiple checks where the first number is the same but for the second number, basically what the bank has to do is the bank has whenever it receives a check it has to wait for some period of time, and if someone provides a check which is signed by the person who's paying that has a higher number than the bank replaces that check with the one with the higher number.

Basically the idea is [00:38:00] that if, let's say you personally have in your possession, a check where the second of four or some channel idea where the second number on the check, the sequence number is say 47 and you personally know that you… Or let's say this check has $400 on a going to you. Let's say you personally know that you never signed a check that has a number higher than 47. Then you know that if you wanted to at any time you could take [00:38:30] this check with number 47 submitted to the bank and the bank would go through it's like standard period of waiting for other checks. But because you never signed a check and you personally know you never signed a check with a number higher than 47, you know that nobody will be able to challenge you. And so you know that the bank would have to give you the $400.

Basically because you know how you participated in the game and because this check has a kind of higher number than all the other checks that you've signed and then you know that if you try to [00:39:00] cash you'll be able to win it. You can act as though you already have the money even though what you have so far is just a check.

Demetri Kofinas: High level here. Where are we in terms of scaling the Ethereum network and when do you guys think that developers will be able to build things like a decentralized Uber on your platform?

Vitalik Buterin: In terms of scaling, there's multiple, I think applications using state channels that are live on Ethereum [00:39:30] may not already. I know there's been a lot of research over the last half year and development over the last half year or so on making channels viable and work, so that number is going to increase fairly quickly. The plasma side is a bit behind channels because it's technology that was established but it is rapidly getting there.

Sharding, we expect to take longer because sharding is a kind of a more fundamental change to the actual base of the Ethereum blockchain. It's not like a layer to a thing that anyone can just go and spin up by themselves.

Vlad Zamfir: [00:40:00] Yeah. And it doesn't just benefit from layer one security. The way that plasma and state channels do. Plasma and state channels kind of free ride on layer one security for their failure modes. Whereas sharding-

Demetri Kofinas: We're talking about the dispute resolution for example.

Vlad Zamfir: For example. Yeah. Whereas that it has to be done in the protocol. There's a kind of-

Demetri Kofinas: Because the same protocol that the main Ethereum network is using is what the shard would be using?

Vlad Zamfir: That's the end goal.

Demetri Kofinas: For consensus.

Vlad Zamfir: That's the end goal. In the interim, what we're likely to have-

Demetri Kofinas: It's a more fundamental change, [00:40:30] we're talking about a much more fundamental change

Vlad Zamfir: Or like a more kind of very different problem in terms of distributed systems.

Demetri Kofinas: Are there any other systems out there that you look at how they do this and you think that there's something that you can draw on or you can learn from that you can implement with what you're doing with Ethereum that's useful? What are your thoughts about that? Is anyone doing anything interesting that you feel you can borrow from that you can learn from that you can implement in Ethereum?

Vitalik Buterin: I think a lot of the 3.0 projects [00:41:00] and putting "3.0" in quotes, I think going in a direction that Vlad and I both feel as one that is seriously misguided in a lot of ways because I feel there's a lot of people that go into the blockchain space and see, "Oh, there is a governance problem and the governance problem needs to be solved.? Well, how do you solve the governance problem with the mechanism? Well, the mechanism has to be decentralized? Well, what is a decentralized governance mechanism? [00:41:30] Well, basically voting.

Demetri Kofinas: You're talking about on-chain governance?

Vitalik Buterin: Yeah, yeah, exactly right, and the problem is that even though on-chain governance often feels kind of structured and clean and decisive, it has a whole bunch of fairly serious issues in pretty much all of the implementations of it that we've seen and even though to a lot of untrained eyes, it seems like progress. I think in both of our opinions, it’s [00:42:00] actually regress and potentially fairly dangerous for ... basically it reduces our ability to be confident that the system’s will be able to maintain any particular kind of guarantees that we care about going into the long-term future.

Demetri Kofinas: There's no failsafe system-wide in other words, in a sense?

Vitalik Buterin: Well, the problem is more a matter of that ... it's the users are not represented.

Vlad Zamfir: The failsafe that normally is there in blockchain governance is that people who run nodes decide whether to [00:42:30] install upgrades so that if the Dev team is coerced to install an update then it will be rejected by the network, but the on-chain governance removes the necessity for node operators from participating in governance by creating a default that where it just follows the chain governance kind of solution.

Node operators are disenfranchised from participating in governance by on-chain blockchain governance and that basically means that the [00:43:00] on-chain governance mechanism will be empowered and effectively will own the blockchain. Whereas like today, the blockchain doesn't really have any owners. It's actually quite...

Demetri Kofinas: That's a good way to transition into how governance works in Ethereum?

Vlad Zamfir: But wait, there's more about why other projects suck.

Demetri Kofinas: Let's talk about it. Let's talk about why they suck.

Vlad Zamfir: Basically to me there's a bunch of classes. One of them is proof-of-work. Any project that is based on proof-of-work is one that I think is kind of using very [00:43:30] suboptimal mechanism for doing the economics and the regulation of the protocol. I'm a proof-of-stake guy. I'm really into security deposit based proof-of-stake and if the state system you proposing doesn't have security deposits that I'm sure that the robustness-

Demetri Kofinas: Things that you could lose in other words if you act badly.

Vlad Zamfir: Yeah. Then definitely the equilibrium of the protocol if there even is one, isn't going to be as robust as it could be if you use security deposits.

Demetri Kofinas: Can I ask you something about that though, because that's a game theory approach, right? What you're saying is that you've got to have money at stake because if you don't have money at stake, you have nothing to [00:44:00] lose, but what if your nodes are infected and they're acting in a way that's not sort of in the best interest of the node were the node to be assumed to be a sort of self-interested homoeconomicus.

Vlad Zamfir: Yeah. Well, I also believe in kind of forgiveness, which is to say that, if a node exhibits a fault but it doesn't cause a failure, it's not an attack in the sense that it's not a successful attack.

Demetri Kofinas: This is like a New Testament POS.

Vlad Zamfir: Exactly. I used to be very old testament where I just burned down the whole city if there's a single center, but now [00:44:30] if there's just a few sins and they didn't cause any damage, then they don't need to be penalized quite as harshly.

Demetri Kofinas: This seems exponentially complicated because of the game theory aspects of it because of the fact that you're dealing with human beings and attack vectors that are sort of off the charts.

Vitalik Buterin: I would not say that it's exponentially more complicated than proof of work, maybe two times more complicated. And then again as I said, right, but the Casper FFG paper [00:45:00] is basically about as long as Satoshi's whitepaper.

Vlad Zamfir: Also thankfully we get to think of the robustness of our equilibriums independently of the incentives that happen outside. We can kind of maximize our security against the outside world by just pulling as hard as we can towards the protocol guarantees being satisfied, but that's just on the economic side. Then I also have lots of things I don't like that people do on the distributed system side.

I'm a huge fan of sharding. If the scaling solution proposed isn't [00:45:30] sharding then I'm not too happy with it and I have other relatively minor pet peeves I don't like in protocol fault tolerance thresholds. I think that's relatively-

Demetri Kofinas: What do you mean when you say that-

Vlad Zamfir: The classic thing is the one third, everyone's always talking about their one third fault tolerance and that how it's theoretically optimal.

Demetri Kofinas: Well, that's Hedera Hashgraph for example.

Vlad Zamfir: They're not alone.

Demetri Kofinas: And why don't you like that? What's your reasoning for that?

Vlad Zamfir: Because I think that it's completely possible to have consensus protocols with extra protocol fault tolerance thresholds, which means that the protocol [00:46:00] doesn't need to decide. The protocol developer doesn't need to decide, but the user can decide. I want it to tolerate one third equivocation faults and one third live faults. Or they could decide I want to tolerate 50% equivocation faults and 25% liveness faults as opposed to kind of blank boxing these things as Byzantine fault tolerant without saying what kinds of faults they are.

I think that basically, because people will think of Byzantine faults in this kind of generally abstract way as opposed to thinking more concretely about what exactly are the crimes that we're talking about, the fault tolerance profile that people imagine. [00:46:30] It's basically that one third is the maximum when really they're choosing a third, what they're saying is, "I want to tolerate a third for safety. And a third for liveness. Because that kind of maximizes my fault tolerance and in the worst case for either."

Whereas, I personally would be much happier tolerating 50% for safety and 25% liveness, but I don't want to make this decision as a protocol developer. I want the nodes to decide just the way bitcoin nodes decide how many confirmations do it.

Demetri Kofinas: How would they decide that?

Vlad Zamfir: There's just a couple of ways. One of them is, you just watch how much fault tolerance blocks end up getting [00:47:00] and then once you feel comfortable you go with it. The other is you can literally code it into your apps or plug it into the command line or to the user interface. For example, imagine if you're programming in exchange --

Demetri Kofinas: So you're saying the developers will be able decide what full tolerance they are willing to have.

Vlad Zamfir: Yeah, just like the exchange decides how many confirmations to wait for before they accept the deposit. They can decide how much full tolerance to require.

Demetri Kofinas: So what happens if you have a lot of developers that don't know what they're doing and they develop, let's say, they give too much leeway [00:47:30] and they end up let's say building multibillion dollar software on Ethereum, and then they miscalculate, for example.

Vitalik Buterin: That's a very interesting thing to say on the DAO fork anniversary.

Demetri Kofinas: Is today the DAO fork anniversary?

Vlad Zamfir: Yeah, two years.

Demetri Kofinas: That's amazing. Well, that's a good. Another good point of the importance of governance. What you guys did on the DAO fork. You had the ability to make a decision about what you want to do or not. That was the crisis control room. What was that like? That was like you out in the bunker. Like the day that the US government took out Osama Bin Laden. [crosstalk 00:48:00]

Vlad Zamfir: Well, we weren't in a room. We weren't in a bunker.

Vitalik Buterin: I was in Shanghai, a bunch of developers were in Europe. What continent were you in Vlad?

Vlad Zamfir: I was in Canada.

Vitalik Buterin: Okay. Yeah. See, we covered three continents.

Demetri Kofinas: You were virtually in the same bunker. You were in the same shard.

Vlad Zamfir: We were in many, many, many different channels with many, many different people talking about this. The threshold of attention that the Ethereum community paid during that time was very hard.

Demetri Kofinas: But to bring it back as it is a serious question. What you're describing here -- You're forfeiting control, [00:48:30] a great amount of control and agency over to the developers. Is that not a concern for the same reason?

Vlad Zamfir: To some extent, but I mean people today, do decide whether to choose one confirmation or six confirmations or zero and it's not a disaster even though like to me [crosstalk 00:48:50]. Zero conflict is sometimes a disaster.

Vitalik Buterin: Yeah. I was about to say.

Demetri Kofinas: We're talking about the probabilistic consensus for bitcoin primarily.

Vlad Zamfir: In bitcoin. [00:49:00] Yeah, exactly.

Demetri Kofinas: Like how long do I wait? Depends on how valuable the good is.

Vlad Zamfir: Exactly. Yeah. People have heuristics that people develop and share for these things. It'll be the same. It's a matter of ... just like with the DAO hack. The reason why the DAO hack happened was because there was a developer norm not to have this reentrancy bugs that didn't exist at the time. Although today it really, really, really does. Where everyone kind of knows not to write that exact bug.

Vitalik Buterin: If all else [00:49:30] fails, just play the safe academic road and choose two thirds.

Vlad Zamfir: Not me. [crosstalk 00:49:35] No. Oh, sorry. So let's wait a second. I misunderstood. I thought he meant as the protocol developer. [crosstalk 00:49:42]

Demetri Kofinas: Let me ask you one last thing guys, because I did want to get into governance, we have a limited amount of time. So this is a good example which is that there isn't clear governance in the blockchain community. It is open source software. So you've got developers, you've got [00:50:00] miners, there are different stakeholders, you have Vitalik who is -- are you Vitalik, something more like Queen Elizabeth today? Are you more like Victoria? What is sort of your role in all of this? I've heard some people say, Vitalik the benevolent dictator. What do you think is the most accurate view? Are you a sort of a vote breaker in the middle? Like your tie breaker?

Vitalik Buterin: I guess like one thing that's very important to point out is that like, obviously my power isn't like zero because there's a whole bunch of people that listen to [00:50:30] me and all those like Twitter followers and so forth, but both of our power is almost entirely soft power. It's almost entirely based on basically the respect that the community has for us as a result of our track record of being able to produce good research and come up with good protocol ideas and if let's say for example, [00:51:00] one of us got captured by in some crazy third world governments that started trying to extort us and when I put back doors into the protocol like, sure, all right, the EIP, I'll write the code, the number of people who would listen to that is pretty close to zero.

Demetri Kofinas: Well, that's true. So the open source nature of this protocol prevents really bad things from happening in that sense, but it also makes it very difficult to move forward. Because you can get stuck in the mud, right? I mean, how do, for example, if you guys want [00:51:30] to move forward with Casper, let's say in a fully POS solution where you get rid of proof-of-work entirely, isn't that going to be a problem for the miners? Are they going to create problems in the implementation of that?

Vlad Zamfir: Basically there's a political process by which kind of the community makes up its mind about what to do and then what to do in political gridlock is kind of one of these questions, right? In the bitcoin world they have this norm against contentious hard forks and in Ethereum I don't think we really have that and I think that's ultimately the way that even if we can't agree, something will happen and I think [00:52:00] the Dow hard fork was a good example of a contentious hard fork and I personally have no qualms with doing a contentious hard fork if the miners don't like proof stake.

Vitalik Buterin: Another practical thing to point out is that I feel like basically almost the entire Ethereum community supports proof-of-stake and especially after the DAO Fork, like, basically, Ethereum classic ended up absorbing a lot of the proof-of-work fans. Chances are they are going to either stick with a proof-of-work or I could go with some rubber rose hybrid thing that [00:52:30]. I'm not exactly sure of the details of.

Demetri Kofinas: Let me ask you this last question. Are you hopeful about the ability to overcome the challenges that you guys are facing to scale? We've talked about some of them here today and if so, what do you think the timeline looks like ... again, to bring it back to the point about development on Ethereum, where a lot of these projects that have gotten funding for example, or are trying to get funding to build distributed applications on top of Ethereum, that they'll actually be able to do it. That you'll have the throughput [00:53:00] and that you'll have the low levels of latency.

Vitalik Buterin: I think like first of all, if you're your impatience and you want scalability, then like, choose channels and plasma, right? Because like channels are basically. Again, there are versions of specific use cases that are pretty much ready for use today and plasma implementations are making very rapid progress and we can expect them to be done with it within this year for at least some cases.

The kind of full dream [00:53:30] of baselayer, sharded proof-of-stake, especially with all of the different properties that we have will probably kind of come online in stages over the next few years and that's more difficult to predict.

Vlad Zamfir: For the most part. I feel very optimistic about tech governance. I'm much more pessimistic and concerned about governance with respect to irregular state transitions and things like reversing harms, preventing harms, and dealing with malicious applications. Those are kind of much more contentious and difficult [00:54:00] governance questions. Then will we be able to upgrade the tech.

Demetri Kofinas: Guys, I appreciate you coming in on such short notice. I wish we had more time and I wish you the best in your travels to ... Where are you going next?

Vlad Zamfir: Ithaca.

Demetri Kofinas: Do you have a home or you just like perpetually living out of a suitcase?

Vitalik Buterin: Perpetually living out of a 40 liter backpack.

Demetri Kofinas: Do you have to pay taxes because I mean you know, we talked about Doug Casey, and Doug doesn't have to pay taxes. He's an international man of mystery. He's not domiciled anywhere.

Vitalik Buterin: [00:54:30] I think cryptocurrency people tends to care particularly about capital gains and there's actually a number of fairly mainstream jurisdictions that don't even have to have taxes on capital gains, so it's a bit easier.

Demetri Kofinas: Is that where you have your Lambo. I'm just messing around, man. All right guys, thanks for coming on.

Vitalik Buterin: Yeah, cheers.

Vlad Zamfir: Cheers.

Demetri Kofinas: And that was my episode with Vitalik Buterin and Vlad Zamfir. I want to thank both of them for being on my program. [00:55:00] If you're a regular listener to this show, take a moment to review us on iTunes. Each review helps more people find the show and join our amazing community. Today's episode was produced by me and edited by Stylianos Nicolaou. For more episodes. You can check out our website HiddenForces.io. Join the conversation at Facebook, Twitter, and Instagram @hiddenforcespod or send me an email. [00:55:30] As always, thanks for listening. We'll see you next week.