All About Project Hamilton - June 9, 2022

Federal Reserve Bank of Atlanta *Talk About Payments* Webinar

All About Project Hamilton

June 9, 2022

Transcript

David Lott: Welcome, everyone. Thank you for taking the time to join us today in our ongoing *Talk About Payments* webinar series. My name is Dave Lott, and I'm with the Retail Payments Risk Forum here at the Federal Reserve Bank of Atlanta.

We're going to call a little bit of an audible today, in that the plan was that we were going to be joined by Jim Cunha and Bob Bench. Unfortunately, Jim just a few minutes ago had a family emergency come up. He hopes to be able to join the webinar a little bit later, but we're going to let Bob take care of everything. Bob is an assistant vice president, and I think it would be fair to regard Bob as the "in the weeds" technologist behind Project Hamilton.

Before I turn the presentation over to Bob, and hopefully Jim, I wanted to mention some things about the webinar. First of all, to provide a disclaimer for myself and Bob and Jim, any opinions expressed here are those of our individual selves and not necessarily those of the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of Boston, or the Board of Governors.

This webinar is being recorded and will be available on the Retail Payments Risk Forum website next week. A copy of the presentation that's going to be given will be available as well. We ask that you submit your questions through the Q&A panel on your screen there. We have planned on having plenty of time to address those questions.

For those of you that have been following the Fed's activities with regard to central bank digital currency, you'll know that there are two separate but parallel paths: Project Hamilton, which Bob is going to cover in detail, as well as the Board of Governors' policy issues around central bank digital currency and monetary policy. Bob's going to touch on those as well. With that, I'm going to turn the presentation over to Bob.

Robert Bench: Thank you, David. I appreciate that. I'm going to have a very difficult job trying to replace Jim Cunha, my boss, but I'll do my best here. I want to read you our disclaimers that David stated. My discussion today represents my views alone and not necessarily those of the Federal Reserve Bank of Boston, Federal Reserve System, or the Federal Reserve Board of Governors. I think, importantly, Project Hamilton should not be seen as indicative of a preference by the Federal Reserve System or the Board of Governors for or against a central bank digital currency. That is exclusively in the purview of the Federal Reserve Board of Governors.

I think also, importantly, one disclaimer I always like to make at the beginning, in case I forget to say it many times, is that Project Hamilton is technology research, and research alone. This is not an effort to build a production-grade system. This is an effort to better understand technology, and to better provide data to policymakers and leadership, and in the open source to understand how this could work, and to also understand how it shouldn't work, for example.

All those disclaimers aside, I'm going to discuss what I'm going to talk about today. What we're going to talk about today is, generally speaking, Project Hamilton and retail central bank digital currency. First, why are we doing this, and what's the context for the why? Second, what did we do here? What kind of requirements did we come up with, and what did we actually do? What was our day-to-day effort? Then finally, what were the results of that effort in our phase one? I'll briefly discuss where we want to go after the completion of phase one this year, and then importantly, what's yet to be done? What are the myriad policy issues out there that may be necessary to do technological research around, to better understand what kind of data could support a multitude of policy options?

That's the crux of the discussion today. What I hope to do is provide a clear picture of the why, the what, and the how on Project Hamilton and our efforts towards retail central bank digital currency, and then enable you all to ask a lot of questions, because I think that's where the most learning will come from.

This is a slide that Jim likes to use. We just generally discuss types of money. Money, generally speaking for most persons, is private money. When we think about "private money" right now, it comes in a lot of forms. Almost all private money right now is effectively financial institution funds. This is either commercial bank money, which is the money that shows up in your bank account, maybe it's private money in the shape of money market funds in the securities world, or maybe it's your Venmo account or your PayPal account. These are all forms of private money that tend at some point to be anchored in commercial bank money.

That's the big chunk. What Jim has here is a couple of new types of money that have emerged over the last 15 years that we like to think about when we think about Project Hamilton and how we learn. One is just generally called cryptocurrencies, which really was founded at this point almost 14 years ago with the invention of Bitcoin. With the invention of Bitcoin, what we think about is an alchemy of existing technologies that were put together for the first time to develop a truly peer-to-peer digital cash system. We've tried to learn as much as we can from this space, and all the lessons that have been learned over the last 15 years in this industry, to understand: Are there any learnings here that can assist our

mission at the Federal Reserve of building a more efficient, more secure, and more inclusive payments system?

Next, Jim breaks out stablecoins. Stablecoins are, effectively, usually dollar values that ride on top of Level 1 blockchains and are executed via secure contracts. The most notable ones are Tether and USDC. A full disclaimer: I formerly worked at Circle and helped build USDC in the private sector.

The idea here is that that kind of peer-to-peer cash system that was invented under Bitcoin has certain functionalities, primarily through Ethereum, that allow people to move what are claimed to be dollar equivalents over Level 1 public blockchains. These are increasingly being discussed in policy circles as a new form of monetary innovation.

Finally, CBDCs. Central bank digital currencies are, effectively, digital representations of public money. CBDC is one of the first new public monies that has come to bear in a long time, and the idea here is, "Should we take the lessons learned from the private money world, whether it's the advances in digital money with PayPal or Venmo, or the advances in crypto, and should we take those lessons learned and put that in the public money world? And should, and can, public money advance the same way that the private money has, from a technology standpoint?" That is part of what we do with central bank digital currencies.

One question is the why. Why do we do CBDC research? I think that's a common question that we get on Project Hamilton. Jim has listed a lot of contributing factors, or some context, here of the why. China certainly has built a central bank digital currency that is now used by over a quarter billion people and appears to be very similar to the private sector initiatives by Alipay and WeChat Pay that built the largest digital payment system in the world, which is the Chinese mobile payment world.

Certainly Facebook's use of Libra brought a lot of attention to digital money. The commercial banking sector has begun issuing their own stablecoins, as well as the tech industry. Crypto continues to be very interesting, but its volatility also remains interesting. The legislative bodies, particularly Congress, have begun getting more active in debates about central bank digital currencies.

The much higher-level discussion is, technologies tend to happen, and technology emerges, and it has begun happening in the private money world. Whether or not to do or issue a CBDC is absolutely a policy question. I think what is unquestionably important is that we understand the technology and understand whether the technology is appropriate for the mission of the Federal Reserve.

That is why we do it, at the end of the day, to understand: How is technology emerging? What tools do technology currently offer those who are responsible for managing payments systems? Should those technologies be used, or can we get better data on the implication of those technologies on the Federal Reserve's mission in payments?

What's in context? What's happening out there around the globe in retail central bank digital currency? As I stated earlier, China is incredibly important. They have been working on a central bank digital currency for the better half of a decade and have launched that with a quarter billion people with many billions of yuan in play. Their product is called the e-CNY, and as I mentioned it looks and works and feels a lot similar to the most popular mobile apps in payments in the world, Alipay, WeChat Pay.

However, other countries are getting involved. The Bahamas has launched theirs, as has the Eastern Caribbean Central Bank. Many countries remain exploring the possibilities, but I think the ultimate takeaway with the global activity in retail central bank digital currency is that every country has their own requirements. This is not a one-size-fits-all technology. When I met with the Eastern Caribbean Central Bank over three years ago, probably two and a half years ago, they had a hard problem to solve: moving cash island to island. A central bank that had numerous jurisdictions over dozens and dozens of islands meant that moving cash was very expensive and very slow, and for a population that was fairly impoverished. They needed a digital cash that moved money, only faster than a boat and cheaper than $100, which was the fixed fee.

You always need to think, "What's the purpose?" In Sweden, cash is going away, so they need to figure out how to build a digital cash. In Germany or Japan, privacy is incredibly important due to their historic issues with state activity and surveillance. Every country has its own model. For America, we're already thinking about our requirements. We need to think about, "How do you handle the largest currency in the world? What type of performance implications are there?" A lot is happening globally, but importantly, every country has to focus on their own requirements for their own economy.

What did we do? How did we approach this problem? We know that money has developed from a technology standpoint, certainly since the advent of mobile phones, and the invention of fintechs, but then very drastically since the invention of Bitcoin in 2008. Taking all of that information was what we were trying to do. When we titled Project Hamilton, the idea was: How do we take all this technology that's been built, particularly over the last 14 years, and understand is there anything there that is helpful to our mission?

The first thing we had to do was find people who were good at this stuff, who had done it before, who had researched it, built it, managed it, and shipped it. Fortunately, at the Federal Reserve Bank of Boston, we have a company in town that has done this, with USDC. We hired people from Circle, which launched USDC, and that's where I came from. Importantly, Boston also has, right across the river in Cambridge, the world's preeminent research university in digital currencies.

The Federal Reserve Bank of Boston signed a multi-year contract with Massachusetts Institute of Technology's Digital Currency Initiative to understand this. We put together the best researchers in the world, this technology, with persons who built and designed and shipped and managed a stablecoin, as well as the leading thinkers inside the Federal Reserve System, who have managed money at the stakes of the Federal Reserve for a long time. That was a very unique combination of persons, and we were very fortunate to bring them together on this project.

Our goal was to create a hypothetical research platform for retail central bank digital currency. We wanted multiple use cases, but we thought that it would have to be something primarily from a performance standpoint that could handle peer-to-peer payments, personal commerce payments, commerce-to-person payments, government-to-business payments, any type of way that money could move, we wanted this to support.

We also felt that we should be technology agnostic. There are multiple ways to move money securely, quickly, and resiliently, and we wanted to explore all of them, understanding that our expertise primarily lay in traditional, distributed architectures in parallelized computing.

We began this project in earnest in the summer of 2020, and we released our phase one in February 2022. Importantly, as Jim bolded here, there has been no decision to move to pilot or production. This project remains, as intended, research only.

What were our requirements? Like I said, every country has their own requirements. To do things well, you need to pick a couple of things and do them well. We really focused on what we called performance, resiliency, security, and flexibility. For performance, we had two main goals. One was, we wanted a minimum of 100,000 transactions a second. We wanted this to be future-looking. Most real-time payments systems right now are in the "three or four digits" transactions per second. We wanted in the "at least six digits" transactions per second, understanding that our mentality here is this is something that we want to understand the future of money, not necessarily the present of money.

We thought speed was critical. By speed, we meant full clearing and settling in under five seconds. Most real-time payments systems today clear and settle in around 45 seconds, and oftentimes those are netted, which leads to interbank balance liabilities. One of the ideas here is that as payment volumes increase dramatically, in a world where potentially payments are automated and become machine-based payments as opposed to person-based payments, the payment speeds and payment volumes are increased and thus settlement will need to be quicker.

We believe resiliency is important, understanding that the United States is a very complex place from a geography standpoint. We would have to understand that on the East Coast hurricanes happen, on the West Coast earthquakes happen, and in the Midwest tornadoes happen. How do you make sure that a farmer in Ohio is able to sell his crop if there's a hurricane on the East Coast that's taken down most of the East Coast cloud infrastructure? We try to ensure that this is a system that no matter whether there was an earthquake in California and a hurricane on the East Coast, that that farmer can still go to market and sell their goods.

Privacy was critical. What we did, from an architectural standpoint, is understanding that it's very easy to build compliance into systems later on. What is very hard to do is build in privacy if the system takes in a lot of data at the core. Our current model takes in almost no data at the core architecture, other than what's necessary to process transactions. We also thought that that was critically important from a systems efficiency standpoint, as well as a cybersecurity standpoint.

Importantly, what we did for all of these areas is try to create an incredibly flexible system. Again, our work is research-only for a theoretical central bank digital currency. However, we want to understand that if there's certain policy options that policy colleagues would like us to test out to better understand their own research, and importantly, if there are other stakeholders that have interest in policy options that we could build on top of our system, test those options out, and provide reliable data about those policy options. Flexibility was critical.

That's phase one. That's what we did, finishing up this winter. Phase two, which we've been working on now for about six months, is understanding what are those policy options that we find to be so pressing that we should examine? Or, alternatively, are there functionalities in money that have emerged since the advent of Bitcoin that are particularly interesting to us? Some of those are very critical, such as new ways to do privacy, such as zero-knowledge proofs. Programmability, adding smart contract functionality to our system, understanding whether there are ways to do interoperability across digital assets that may make certain frictions in the payments system less. A classic example there is what are called "atomic swaps of tokens."

Also, increasing accessibility in offline payments. America is a big place. A lot of America doesn't have internet or Wi-Fi, so can you still use a research version of digital cash? Can that still be transferred in the absence of that? That's another area of technology we'd love to explore from a research standpoint.

What did we do? I think the most important thing we talk about is, what did we actually get out of this? Importantly, what we did is, again, what we didn't do. We did not build a production system. A production system requires significant hardening to ensure that it can work reliably time and time again. We did build a pretty effective system. One, we built from scratch. This was not pulling down from someone else's work. We built the entire system from scratch with our own engineers in C++. We built a flexible model with minimal data in the system. The system parallelizes as much data as possible to increase efficiency. Importantly, what we did is we built a system that was centralized as opposed to decentralized. I think that's one important lesson learned, was that a system run by the Federal Reserve would need to be centralized. You can get a lot more efficiency than the traditional decentralized systems you'll find in, say, Bitcoin or Ethereum.

What were our data points? Accordingly, we built two systems. We built one that was an order processing engine. That order processing engine topped out at about 170,000 transactions a second. We think we might be able to get it a little bit higher, say, 200 to 250,000. That's fast, right? That's still about 10 times Visa on a good day, but it wasn't fast enough. We thought we could get a lot faster.

We built a second system, called the two-phase commit (2PC) system, which achieved linear scalability. We reached 1.7 million transactions a second, about 100 times faster than Visa, with 99 percent of finality in under one second. As I mentioned, most real-time payment systems do it in about 45 seconds. We did full settlement finality in under one second.

Both are highly resilient, with multi-tiered replicated architectures that are distributed across what we call three AZs, or three "cloud

environments" across the United States. If we want to take out our East Coast AZ and our West Coast AZ, Ohio keeps turning on and doing the job just fine and we don't lose any data. That's really important.

We use cryptographic techniques such as private and public key pairing. Importantly, what we did here that makes us different from almost all other modern payments systems is we used the UTXO model, or the unspent transaction output model. This makes the nature of our form of money an object-oriented form of money, as opposed to a data point on a traditional relational database. We think that there are critical values here when it comes to security, resiliency, double-spend proofing, counterfeit proofing, and cybersecurity reasons. That was a major finding of ours. The idea here is that this platform we built can handle any number of design or policy choices.

Again, before I start this slide, what did we ultimately learn? We learned that some things from the post-Bitcoin digital money era could be useful for central bank digital currency. We learned that the UTXO data format for money, money as an object, not as a piece of data on a digital database, makes money move more securely, we think. It works in a highly efficient money format. I think there was some skepticism about whether UTXO models could move money quickly.

We importantly learned that an ordering monetary system, whether it's a blockchain or a traditional payments system that orders payments, leads to a performance bottleneck. There is a way to remove transaction ordering but enable transactions to not conflict with one other and happen at considerable scale, with considerable efficiency. Our two-phase commit model, importantly, will remove transaction ordering and receive linear scalability in our transaction space. That was really important.

What we also learned was that some of these things from crypto don't make sense. Importantly, the whole idea of Bitcoin as a peer-to-peer cash system assumed a malicious actor, and assumed that you had to do, for example, proof of work mining to solve that problem.

With a central bank, you are assuming a central actor. Removing that mistrusted actor assumption in Bitcoin and removing the proof of work mining system, or even a proof of stake mining system and replacing that with what we call a "Raft consensus system," enables you to move money even faster. So you get that same level of security, you get that same level of efficiency, but you get considerably stronger performance in the speed and settlement times. Those are two major functions.

Importantly, we are a small team of central bankers and software researchers…to understand, how can we get smarter in learning from others? This is not an RFP, or a request for procurement-type system. What it really is, is understanding: Is there incredible work being done by other parts of the world, that we can learn from to better understand how money can work? That's something we're continually working on, to identify how smart companies can help us get smart.

One question we get is, how could retail CBDC work? I won't go too deep in this, but the idea is that this could work a lot like cash, or it could work a lot like other forms of payments. You can have person-to-person, including offline person-to-person. You pay me $5, and we click phones, and that could work. Consumers can pay for things. This happens currently in China, where you can go to your local Walmart and use e-CNY and buy goods. Business-to-consumer, right? People can pay their staff in CBDC. Cross-border, this is a particular area of stress. Moving money cross-border is very slow and very expensive, and this is always seen as one of the most important things affecting CBDC.

What we think about is, how could that be done better? And there are certainly use cases here where retail CBDCs could be used to either move a single currency cross-border, or actually enable currencies to swap automatically. And then business-to-business. How do businesses exchange money with each other? Is there a use case for a retail CBDC?

We think about micropayments in the future. That's big in IoT. Again, as I discussed, our requirements are thinking about 10, 20, 30 years down the line of computers making payments. If micropayments are a thing, you may need to have a lot more transactions, for example, to read an article on the internet. If you have a native browser attached to a digital payment mechanism, maybe every page you look at in the *New York Times* is a couple of cents. We want to think about how much performance would we need to be able to handle that type of activity.

Then, it was always financial inclusion. It's something that is always a goal in the Federal Reserve System. It's a goal of ours, but we know it's one that continues to be fleeting. How can we understand how modern technology can get more persons involved in the financial system, and make the financial system less expensive for more persons? That's always an outstanding public policy goal.

This is a favorite of Jim's. Us being a Boston team, between the Boston Fed and MIT, Jim's favorite line is "Don't pave the cow paths." What we're really trying to do is understand, can we think differently? Are there ways to think differently about money that will lead us to a much more efficient world of money than we have right now? Again, as anyone who's been traveling Boston who doesn't know these roads like the back of their hand, you can oftentimes find yourself walking in circles.

Other things that are happening in this space right now. I'll discuss them, but I won't comment. Importantly, we are one team in the Federal Reserve, thinking about this. There are many others. Very importantly, our leadership on the Federal Reserve Board of Governors issued a public consultation paper that sought input. That input closed in May, and then Governor [Lael] Brainard spoke about her thoughts on central bank digital currency and digital money in the prior weeks. Certainly, some of the things that have been identified by our Board of Governors are questions about the appropriate level of privacy for money, a preference for financial intermediation in central bank digital currencies, and then questions about ensuring that the government can do its necessary jobs with digital money, such as law enforcement and national security.

Other issues out there…not to go in too many other areas…monetary policy is something that our economists think about greatly. I am not an economist, so I leave that there. Then again, financial inclusion.

Some of the things we think about particularly is, how can we make this a platform for innovation? When we talk about a flexible core processing engine for a central bank digital currency, we think about are there ways that smart people can build on top of this, whether it's the private sector, whether it's the public sector, wanting to add functionalities? How can a research core processing engine lead to more learnings about how money can work? This is something we hope to build on in the future.

Importantly, all of Project Hamilton is open source. It's available in GitHub right now. If anyone on this call wants to contribute to the codebase, to criticize the codebase, to add comments to the codebase, it's all there in public. One, because we believe that innovation comes from the most unexpected places. Two, that money is a question of trust and you should understand the code behind your money, or at least, your research form of money.

Everything I've discussed to date has been about what we call "general-purpose central bank digital currency." Central bank digital currency in a research form that could be used by any party. Another branch of research that is underway, not by our team, but by other parts of the Federal Reserve System and abroad, is wholesale central bank digital currency, or limited-purpose central bank digital currency. This is central bank digital currency that is limited to certain parties, by definition. We don't necessarily think, from a research standpoint, that the core processor of a wholesale central bank digital currency needs to be any different than a retail central bank digital currency. Frankly, we think there's probably a lot of efficiencies to be gained there if they use the same core architecture.

However, there are a lot of efficiencies that can be gained in wholesale. For example, wholesale means your participants are defined, so the whole question of privacy and identity, and the rights around privacy and identity, go away in the sense that these tend to be corporate entities who have certain licenses, normally. You can probably make the system a lot more efficient because they'll be possibly less participants, and they will all be identified.

I think that's a really interesting thing for us in wholesale CBDC, is you could probably make it even a lot more efficient than retail CBDC from a research standpoint. There's a lot going on there. Importantly, the Federal Reserve Bank of New York's Project Cedar is looking into understanding wholesale CBDC from a research standpoint. Singapore, Canada, Japan, and the UK have worked on wholesale CBDCs. Again, the main question for wholesale CBDCs is, who are the economies that need to think about their capital markets deeply? You see places with deep capital markets, like Singapore, and Canada, and the UK, and Japan, and the US, wanting to look into this greatly because capital markets are so important to those economies.

I just covered a lot. I think what I hope to do is be clear about why we did our research in retail central bank digital currency, primarily because technology is always moving forward, and we need to understand how money might move forward as well. What did we do? We built a core processing engine from scratch, by our own engineers and the engineers at MIT DCI. What did we end up building? We built two core processors, one that makes money move at 1.7 million transactions a second with less than a second of finality.

We learned that there's a lot to be learned from the crypto space, in digital money. You can actually make it a lot more efficient if you centralize it, for the purposes of a central bank. What I want to emphasize is, there's a lot more to do. Our work is research only, and not meant for production, but we've added a little bit of data to that conversation, and we hope that's been valuable. There's a lot more work to be done. There's a lot more data to derive from research. I look forward to the questions, particularly if anyone else is doing this research, I look forward to learning from you.

That's all I have for you, David. I'm happy to take questions.

Lott: Thank you so much, Bob. I appreciate you covering all of this. We do have a number of questions that have already been submitted. Just to remind people to use the Q&A panel in order to submit those questions for us. The first one may be a little bit of a softball. It'll give you a chance to catch your breath here a little bit. I'm sure there's got to be an interesting story behind how Project Hamilton got its name. Can you give us a little bit of background on that?

Bench: Yes, you bet. When we were thinking about doing this project, it was definitely new for the Federal Reserve to do an open-source software build from scratch with software engineers. The idea of building a digital dollar, even in research form, was a big idea. We started researching, what are the other big government technology projects that have happened in the past? I read all about the creation of the internet and ARPANET that was done by the government, that led to the modern internet and the trillions of dollars of private sector wealth. We learned about Oppenheimer and how the nuclear bomb was built by the government in order to handle the stress of the Second World War.

Then we also of course learned about the moon landing, and how the government partnered with academia to get a moon mission to happen. Who we were mostly focusing on weren't the politicians or the policymakers. We were focusing on, who was that woman or man putting the code in the codebase, or putting the uranium in the shell? A woman that came up was Margaret Hamilton, who was a software engineer at MIT working on the lunar project. There's no one cooler in my mind than Buzz Aldrin, and those guys who first went up in that plane, but someone pretty darn cool as well is Margaret Hamilton, who built all the software that got them up there, and got them to land safely, and got them home safely.

We were thinking about who we name this project after, and we were thinking about one of those folks, and then I was listening to the Hamilton soundtrack, and that hit it for me, and we went from there. Alexander Hamilton was a pretty cool person in central banking history, and a great play was made about him, but this is a lot more about Margaret than it was about Alexander.

Lott: Very interesting. You mentioned just a few minutes ago your partnership with MIT. Can you go into that in a little bit more detail in terms of the roles and responsibility of the Federal Reserve Boston group versus the MIT group that you were working with?

Bench: Yes, sure. One, I can't speak more highly about the Digital Currency Initiative and their leadership, primarily that of Dr. [Neha] Narula. The DCI has been working on this problem since 2015, longer than any other research university in the world. Importantly, they've been doing so in a very unconflicted manner, an unbiased manner. There have literally been billions of dollars to be made if you're really good at this software over the last 14 years. This has been a group that has been steadfastly focused, almost in a monastic state, of making sure this technology works and not worrying about shilling a coin somewhere.

As a public servant, you want to make sure that you're doing work that's critical with the parties that you feel most comfortable that work in the public interest. One, these folks have been working hard maintaining some of the most important infrastructure in digital money over the last seven years. They've also been doing so in a way that has enabled great comfort, as a public servant, that you are all working on the same page.

The way it works day-to-day is our team is currently nine persons. Seven of those persons are engineers who work in this space, so the Boston Fed has seven engineers who we all hired out of industry or academia to work for us very recently. MIT has their own research engineers, who are some of the leading experts in the world in questions of distributed computing, privacy, high-performance parallelizable systems. We all work together in a codebase, and we manage it like a start-up. It's run like a start-up, understanding that our stakeholders are ultimately the US public, and doing so with that in mind.

Lott: Thank you. Along that same thing, you mentioned Project Cedar that the New York Fed is involved in. Can you talk a little bit more about the work that you all are doing in terms of relationship to that project, as well as anything going on at the Board level?

Bench: Sure. Like I said, there's a lot of work happening around the Federal Reserve System by a lot of smart folks. Project Cedar is led by the BIS Innovation Center at the Federal Reserve Bank of New York, led by a great guy named Per von Zelowitz, and what they're trying to do is understand what are the requirements, or what are the things they need to think about, for the obligations of New York, which is the world's leading capital market center? How can they learn from this type of technology that could have an implication on their duties in New York, in the keystone of the world's capital markets, definitely the world dollar markets?

They're looking at that deeply and building partnerships on their own, and also working collaboratively, not only with the Federal Reserve System, but with the global BIS Innovation Hub. They have a fantastic view, not only of what's happening in New York's markets and how that might help them learn from the digital space, but also the collective learning from the BIS Innovation Hub. That's a huge win for the Federal Reserve System, having not only that opportunity within the BIS, but then their leadership is fantastic.

The Federal Reserve Board has what's called the TechLab, and the Federal Reserve Board TechLab is doing their own research on central bank digital currency and related digital currency areas, and they're building a great team. The Federal Reserve Board also has a policy section that is deeply studying all the policy issues around this area. I think one thing I didn't realize until I joined the Federal Reserve System is how many smart folks there are. Not just smart folks out of PhD programs, but smart folks who've been working in these payments systems for a long time. They come together in the Federal Reserve System so learning here is like learning from a fire hose. The Board particularly has those PhDs and those experts coming together to help guide our policymakers and our decision-makers in how they inform either the legislative branch or the executive branch on the decisions that may or may not need to be made regarding digital money.

That's happening there. We have a great team up in Minneapolis that's working on standards, as I'm sure you have a lot of payment experts on here. This all needs to go through standards bodies, and we have a couple of experts up in Minnesota that are deeply understanding of what type of standards changes would be needed for digital money. There's more and more coming out every day from across the system, and we're very excited to see it.

Lott: Great. One more question along the same theme again. You had a number of slides there talking about the activities that other central banks were doing with regard to digital currency. Has the Federal Reserve collaborated with any of those that you can discuss?

Bench: As far as I understand, no. There may be stuff going on that is just not in my wheelhouse. Importantly, MIT has entered into research collaborations with the Bank of England, as well as with the Bank of Canada. What you do have is likely a lot of shared learnings between the Bank of England and the Bank of Canada. At MIT, that may be built upon the learnings we had with Project Hamilton.

I think a lot of countries are learning together on this. You see the Bank of International Settlements doing a lot of leadership that is well informed by our own teams, as well as the other global central banking bodies that the Federal Reserve has prominent roles in, and that leadership usually comes from the Federal Reserve Board of Governors, and the Federal Reserve Bank of New York. I think those parties are very active in those global organizations and are leading there very well.

Lott: Great. We're going to move into your wheelhouse now and ask you some more of the more technology questions. You had mentioned that one of your codebases could process over a million transactions per second. The question is, if you launched into production, is that what you would expect to see?

Bench: That's a great question. We don't know. We think, my hunch, is no because the significant hardening that would have to happen to the codebase would likely lead to absences of efficiency. The way I like to describe it is I'm a bit of a car person, and a car on the racetrack would never be allowed on the road because you've pulled out all the normal, traditional, commercial safety standards. You have certain roll bars and

things like that, but if you took your Nissan Sentra outside, and you pulled out the airbags, you pulled out the roll bars, and you pulled off the doors, you could probably go a lot faster.

We don't have a seatbelt, we don't have an airbag, we don't have roll bars, and all that kind of stuff that we need to make money safe and secure. Remember, at the end of the day the job of the Federal Reserve System, when it's money work, is make money safe and efficient. It's because the stakes are so high for the dollar, that my hunch is no, but that's why you do the research.

Lott: The next question is, have your tests considered use case scenarios most cryptos can't support, such as card-on-file or pre-auth[orization]? At the highest level, have you considered incorporating smart contracts within the architecture?

Bench: We have not considered those use case scenarios. We are researching incorporating smart contracts into our core processing engine to understand, one, what are those smart contracts capable of performing? Two, how much would those smart contracts degrade performance? And three, what security or resiliency issues arise from the inclusion of complexity into our core processing engine through those smart contracts?

Lott: There's a request to please elaborate on the consensus mechanism that was used.

Bench: Sure. We used what's called a Raft consensus mechanism. We do have validators of our core processing engine. Those validators are randomly assigned. For example, in a proof-of-work model, validators, meaning miners, compete to solve a puzzle. The miner who solves the puzzle is rewarded with Bitcoin or Ethereum. That mining or the solving of the puzzle by intense computation is what leads to the well-known energy consumption in proof-of-work standards. You not only need computation to run the computer, but you need cooling to keep the computer from overheating and working optimally. We got rid of that. That's one of the first things we usually talk about. Our randomly assigned validator model means that it uses no more energy than any other database.

Lott: Okay. There's a request: Could you talk a little bit more about the way "order independent version" works?

Bench: Sure. This is what we call our two-phase commit model. When we started our project, we assumed that we needed to have transaction ordering. We thought that for two reasons. One, for the traditional payments experts on the team, namely Jim Cunha, my boss, who has built and managed payments systems in the Federal Reserve System for almost 40 years, everything's been ordered. Things are ordered oftentimes to assume that there was enough money in account *X* to pay account *Y*, and you can audit it very easily.

Similarly, most of my team who comes from the cryptocurrency space believed in ordering because that's part of the definition of a blockchain. A blockchain is a chain of blocks in order. We assumed, based on our pedigree, that you should also continue to keep blocks in order. From the crypto space, the blocks are oftentimes in order, in order to ensure that you can avoid an attack on the system by malicious actors, or you can at least recover the system if there's a malicious attacker. When you have a centralized group, you don't need to worry about that, so maybe you don't need to worry about the order.

What we did is, instead of the transactions happening one by one and going through an ordering mechanism, they happened in parallel. So if Bob wants to pay David, and Jeff wants to pay Sally, those two transactions can happen at the same time as long as Bob's wallet has enough to pay David's wallet, and Jeff's wallet has enough to pay Sally's wallet. That can happen, and they don't conflict. Even if Bob wanted to pay David and Sally, that could also happen as long as Bob had enough value in his wallet to pay what he wanted to pay David and pay what he wanted to pay Sally.

The cool thing is, if you do that at speeds nearing one-and-a-half, two million transactions a second or more, they almost never conflict, particularly if they settle in under a second. What we found is that these transactions almost never conflict, and we built a whole testing system to monitor this. Even when they do, all transactions get canceled involving that wallet and they're rerun. When they're rerun, it fixes the problem.

The main idea there is that if you make the pipes wide enough and you make the transactions fast enough, not only from a transactions per second standpoint, but from a settlement standpoint, you don't get transactions that conflict, you don't get false transactions, you don't get the digital equivalent of check hiding. We find that really interesting from a computer science standpoint.

Lott: Okay. I think you touched upon this but let me ask the question as it was submitted to see if there's additional comments you have. The question is: In order to ensure interoperability and facilitate cross-border payments, do you anticipate the international standard-setting bodies, such as IOSCO [International Organization of Securities Commissions] and CPMI [Committee on Payments and Market Infrastructures], coming up with standards or guidance on technology to be used, key parameters, features, et cetera?

Bench: Those bodies are actively working on these types of things. I think there's an active question, in the policy circles, of if one country gets so far ahead in production of this thing, are they going to be the standard-setting group? This happened with 5G, right? Certainly they're active in it. As our team are technologists, we tend to think of interoperability when it comes to actual technological interoperability. One of the inspirations for this project was Linux. Almost every server in the world runs on Linux. Linux was an open-source software created in 1991 that remains today in the open source. There are multiple companies that have enterprise versions of Linux that update to this open-source software on a regular basis.

Generally speaking, if you're a server engineer at the Federal Reserve Bank of Atlanta, or you're a server engineer at the Reserve Bank of India, you're working on the same software. You could get on a plane and work on both, probably without much problem. That's how we think of interoperability, on our team. Can we make software that can universally work, from a research standpoint? Are there any learnings to be done

there? But certainly, those CPMIs and IOSCOs are happening, and the right people at the Federal Reserve are represented there. That's just not the role of the Federal Reserve Bank of Boston. We exclusively, on our team, do technology research.

Lott: Okay. I think your answer touches on, and let me paraphrase, this question. It basically asks, are there lawyers on your team?

Bench: That's a great question. I am a lawyer. I teach central bank digital currency at Harvard Law School. I teach financial regulation at Boston University Law School, and I was the associate general counsel at Circle, so I do a little bit of both. Money, and how money can work as a central bank digital currency. There are a lot of legal issues there. There are a ton of legal issues there, and I studied greatly outside of my role at the Federal Reserve. Importantly, we have very smart people studying it inside the Federal Reserve whose job it is to study it inside the Federal Reserve.

Also importantly, the president of the United States issued an executive order this winter asking these types of questions. What the president essentially asked for is a bunch of smart folks who work at the Federal Reserve, the Treasury, the National Security Council, and the Department of Justice, to come together and identify, not only the technology issues that we work on on our team, but also the issues around, "What is digital money?" For example, the president asked the Department of Justice, "Can you define who should be running a digital currency?" That's a great question. What does money look like in the digital currency, from a legal standpoint? What does AML/KYC [Anti-Money Laundering/Know Your Customer] look like as a digital currency? There is a lot of work to be done by a lot of smart lawyers to define a lot of this stuff, so that the right persons in government can work on it if they're asked to work on it.

Lott: Okay. One other question, kind of related to that. Let me point out that we do have links to the Board of Governor papers, the one that was released in January and then they just released another one last month really focusing on central bank digital currency and monetary policy. I think this question really relates to that. I'm not sure if you can comment on it, but I'll give you a shot on it. It says: Can you walk us through what you've found so far on how a CBDC might affect financial inclusion?

Bench: Our Federal Reserve Board of Governors has been very clear that that's something that's going to require a lot of research. The Federal Reserve Board of Governors, and the chair in her speech, prefers a central bank digital currency that is intermediated in some way, which I think is related to that. Financial inclusion is also a critical question. How do we make sure that money remains not only safe and efficient, but also inclusive? It's always been an overarching goal of the Federal Reserve and its payments systems, and this is a nut we have yet to fully crack.

We think about this in our work. For example, offline payments is something we want to think about. America is a big place. We're a government that may be asked to build a digital currency. You want to make sure that those folks in America who don't live in a place with internet access, or who may be outside of Wi-Fi, can use that technology, can use that public service, even if it's intermediated through a bank or a fintech.

That's something we think about. We think about, how can we make money easier to use? Like I said, there are really smart folks working on financial inclusion across the Federal Reserve and across government, and it's something we need to keep working on. We hope to get better at it, and we hope to research ways to get better at it on the CBDC side.

Lott: Okay. The final question, jumping back into the technology sector: Are you assuming a non-layered system, without channels, and payments and settlements happening in the same layer?

Bench: That's a great question. Our core processing engine enables transfer and settlement. Again, this is our research. What we haven't done in our research is actually build in layers to understand, what would an intermediary look like. How could that work with our core processing engine? That's work that has to be done, amongst a lot of other work that has to be done.

I think what we really try to focus on in the results of our research is the flexibility, in that our system could enable a multitude of layers, a multitude of intermediaries, based upon policymaker preferences or based upon market demand. What we tried to do is build something that was highly performant, that did not preclude optionality for either the markets or for policymakers. We hope, importantly, because our work is open-source code, those who are interested in testing what a layer might look like can pull the code down and build that. We think that's really interesting.

Lott: Great. Well, we're almost at the top of the hour. Bob, I want to thank you so much. I know that normally you and Jim do a tag team on this presentation, so I appreciate you going solo on this. I also want to thank all of the people that participated, particularly those that submitted questions.

Again, just as a reminder: we will have a recording of the webinar, as well as a transcript and the presentation, posted on our website here, hopefully in the next seven to 10 days. We'll be looking forward to having you all join us on our next webinar. Again, Bob, thank you so much. I appreciate it.

Bench: Thank you, David. Thank you all for being here.

Send questions about the webinar to David Lott at david.lott@atl.frb.org.

You may view previous *Talk About Payments* webinars.