14 HARV. NAT'L SEC. J. (forthcoming Dec. 2022)

A TRILEMMA FOR CROSS-BORDER CENTRAL BANK DIGITAL CURRENCIES

Giulia Fanti[*]

Today, most central banks worldwide are exploring some form of central bank digital currency (CBDC), or a digital form of central bank money accessible to the public.[1] There has been particular interest in cross-border CBDCs (also commonly called multi-CBDCs), which can be used to transfer assets from a CBDC ledger in one jurisdiction (typically one country) to another.

Important open questions surround how to *design* multi-CBDCs. For example, how should the system be architected? How should data flow? How should transactions be processed and settled? How should the system be governed?

In general, these questions remain open. Part of the challenge is that multi-CBDCs must satisfy many desired properties, which can sometimes interfere with one another. In this article, I discuss the tensions between three desired properties for cross-border CBDCs: security, privacy, and performance. I present a *trilemma*, which states that existing designs for multi-CBDCs do not achieve all three desired properties. I then illustrate how existing common designs for multi-CBDCs fail to achieve all three properties. However, I also argue that the limitations of current implementations are not fundamental. I believe that with proper cooperation and collaboration between stakeholders, these technical challenges can and will be circumvented, enabling secure, private, and performant cross-border CBDC transactions.

In the remainder of the article, I will assume that a cross-border CBDC would be built upon distributed ledger technology (DLT). The Bank of International Settlements defines DLT as "the protocols and supporting infrastructure that allow computers in different locations to propose and validate transactions and update records in a synchronised way across a network."[2] DLT is a natural design choice for multi-CBDCs, in which there is no central trusted party. Indeed, DLT has been the technology of choice in many early pilot multi-CBDC programs,[3] allowing independent domestic CBDC ledgers to be interlinked without requiring all CBDCs to interface on the same platform.

[*] Dr. Giulia Fanti is an Assistant Professor of Electrical and Computer Engineering at Carnegie Mellon University.

[1] Atlantic Council, "Central Bank Digital Currency Tracker", https://www.atlanticcouncil.org/cbdctracker/, Accessed on Aug 20, 2022; Board of Governors of the Federal Reserve System, "What is a Central Bank Digital Currency?", https://www.federalreserve.gov/faqs/what-is-a-central-bank-digital-currency.htm, Accessed on September 1, 2022

[2] BIS, "What is distributed ledger technology?", https://www.bis.org/publ/qtrpdf/r_qt1709y.htm, September 2017. Note that DLT is a superset of blockchain technology; that is, blockchains are a form of DLT, but all DLT solutions are not blockchains.

[3] Project Jura: Cross-border settlement using wholesale CBDC, https://www.bis.org/publ/othp44.pdf
Project Dunbar: International settlements using multi-CBDCs, https://www.bis.org/publ/othp47.pdf
Inthanon-LionRock to mBridge, https://www.bis.org/publ/othp40.pdf
Jasper-Ubin Design Paper: Enabling Cross-Border High Value Transfer Using Distributed Ledger Technologies, https://www.mas.gov.sg/-/media/Jasper-Ubin-Design-Paper.pdf?la=en&hash=EF5857437C4857373A9287CD86F56D0E7C46E7FF
STELLA – joint research project of the European Central Bank and the Bank of Japan, "Synchronised cross-border payments", https://www.ecb.europa.eu/paym/intro/publications/pdf/ecb.miptopical190604.en.pdf, June 2019

To my knowledge, every multi-CBDC pilot study to date has adopted an enterprise DLT solution. These enterprise DLT solutions are commercial software products that allow one or more organizations to maintain a DLT amongst themselves. For example, R3 has built a DLT platform called Corda, which has been used in several CBDC pilot studies. While these enterprise solutions are practical in many respects, they do not currently cover the full space of technical designs or properties one might envision in a multi-CBDC. Throughout the remainder of this article, I will discuss concrete examples of how pilot projects have used enterprise DLT offerings, and how these products' design choices and constraints affect the resulting multi-CBDC's system properties.

A TRILEMMA FOR CROSS-BORDER CBDCS

Computer scientists sometimes describe the technical tradeoffs of a system in terms of a *trilemma*: a set of three properties that cannot all be satisfied at once. For example, Vitalik Buterin, the creator of the Ethereum smart contract platform, proposed a well-known (in the field) *blockchain trilemma*: it states that in general, a blockchain cannot satisfy more than two of the following three desired properties at once:[4]

1. Scalability: The blockchain can process and confirm many transactions per unit time.
2. Decentralization: The chain does not depend on a few centralized entities.
3. Security: The blockchain can withstand a large percentage of nodes behaving maliciously (e.g., trying to corrupt the state of the ledger).

This trilemma has primarily served as a call to action, helping to guide technical research to resolve these tensions. However, blockchains, particularly in the context of permissionless cryptocurrencies, have different requirements than a multi-CBDC. For example, decentralization is inherently less important in a multi-CBDC than it is in cryptocurrencies, which were initially proposed as a method for enabling decentralized payment systems that do not require users to trust any single party.[5] In contrast, CBDCs are inherently centralized, and a user's central bank is typically assumed to be trusted (to varying degrees).

Based on the requirements of multi-CBDCs and the properties of existing multi-CBDC solutions, I propose a different trilemma. It is my view that *existing* multi-CBDC solutions can achieve, at most, two of the following three properties at a time:

1. **Security**: Do the ledgers of uncompromised parties (e.g., banks) remain consistent and correct even if some parties in the system are compromised (either internally or through third-party malicious agents)? Even if end users trust their own banks, a counterparty's bank could be compromised. In this case, to resolve disputes, there must be a mechanism

[4] Buterin, Vitalik, "Why sharding is great: demystifying the technical properties", https://vitalik.ca/general/2021/04/07/sharding.html, 2021

[5] A more decentralized blockchain is often viewed as less susceptible to corruption—and generally superior—in the cryptocurrency community. For example, see:
Conway, Luke, "Measuring Decentralization: Is Your Crypto Decentralized?", BlockWorks, 2022.

for resolving conflicts. This definition of security is narrow, and does not include many other facets, such as smart contract security, wallet key management, or system availability.[6] It is most closely related to the concept of *integrity*, which is often viewed as a sub-category of the security of computer systems.[7] However, I use this definition because it is, in my view, a prerequisite for other types of security. If a multi-CBDC cannot ensure ledger consistency, then there is no point to building a smart contract platform on top of it.

2. **Privacy**[8]: Is transaction data visible to the parties that need to see it for regulatory compliance (transparency) while remaining invisible to parties that have no need to see it (privacy)? In a multi-CBDC, transparency and privacy have security implications in a broader sense. A privacy-conscious CBDC can have inherent security benefits by not concentrating valuable data in one place.[9] Moreover, transparency requirements regarding anti-money laundering, counter proliferation financing, and combating the financing of terrorism allow regulatory oversight bodies to combat practices that have (inter)national security implications.[10] Today, there is little consensus on what is the right balance between privacy and transparency; these choices depend heavily on cultural norms and governmental postures.[11] While many countries have stated in writing that privacy is a central concern surrounding the deployment of CBDCs,[12] it remains unclear whether these concerns will materialize into designs that shield user financial data from central banks in the way that cash does.

3. **Scalability**: Can the system achieve performance metrics of transaction throughput (transactions per second) and latency (time to confirmation) needed to support international trade?

[6] Giulia Fanti, Kari Kostiainen, William Howlett, Josh Lipsky, Ole Moehr, John Paul Schnapper-Casteras, and Josephine Wolff, "Missing Key: The challenge of cybersecurity and central bank digital currency", https://www.atlanticcouncil.org/in-depth-research-reports/report/missing-key/, 2022

[7] Russell, Deborah, Debby Russell, G. T. Gangemi, Sr Gangemi, and G. T. Gangemi Sr. *Computer security basics*. " O'Reilly Media, Inc.", 1991.

[8] This category could be more accurately (but less tersely) called "Data access control", as it includes both privacy and transparency.

[9] Giulia Fanti, Kari Kostiainen, William Howlett, Josh Lipsky, Ole Moehr, John Paul Schnapper-Casteras, and Josephine Wolff, "Missing Key: The challenge of cybersecurity and central bank digital currency", https://www.atlanticcouncil.org/in-depth-research-reports/report/missing-key/, 2022

[10] Laurinaitis M, Štitilis D, Verenius E (2021) Implementation of the personal data minimization principle in financial institutions: Lithuania's case. Journal of Money Laundering Control, 24(4), 664-680, Emerald
NCA (2016) Guidance on submitting better quality Suspicious Transaction Reports (STRs). Available at: https://www.clc-uk.org/wp-content/uploads/2018/01/Guidance-on-Submitting-Better-Quality-STRs.pdf

[11] Allen, Sarah, Srđjan Čapkun, Ittay Eyal, Giulia Fanti, Bryan A. Ford, James Grimmelmann, Ari Juels et al. *Design choices for central bank digital currency: Policy and technical considerations*. No. w27634. National Bureau of Economic Research, 2020.

[12] Board of Governors of the Federal Reserve System, "Money and Payments: The U.S.Dollar in the Age of Digital Transformation", January 2022
Bank of England, "Central Bank Digital Currency: Opportunities, challenges and design", Discussion paper, March 2020

In this article, I aim to explain the reasoning behind this apparent multi-CBDC trilemma and suggest what would be needed to resolve it. I will next justify the trilemma by discussing how to achieve each pair of properties above, and why the remaining third property cannot be satisfied using current solutions.

Option 1: Independent ledgers: Scalability and privacy, but not security

A naive and simple design for a cross-border CBDC is akin to what is done today. Namely, a cross-border payment would be routed over a series of one or more correspondent banks, each of which performs services like foreign exchange and compliance checks. Ledgers would be updated pairwise at each intermediate financial institution without running explicit synchronization or consensus protocols. The main difference between this design and today's cross-border payment system (e.g., via the correspondent banking network) is that routing would be automated, rather than requiring the (often manual) compliance checks that occur today.

Scalability

This design is scalable, in the sense that it would be able to meet the throughput and latency requirements of today's cross-border payments system. In fact, by automating compliance checking and transaction processing, this simple design could already eliminate several of the latency bottlenecks in today's cross-border payment ecosystem. These bottlenecks can arise from various sources, including (but not limited to) manual compliance checks and requirements that ledgers can only be updated during local working hours.

Privacy/Transparency

The design is also private, in the sense that only the payer, payee, and intermediary banks need to see transaction details. At the same time, intermediary financial institutions can collect and share data about transaction participants to comply with local regulations. Such data can be transmitted to the relevant intermediaries as the transaction is passed to its destination.

Security

This design is *not* secure in the sense of my definition above. If a sender, Alice, sends a payment to a receiver, Bob, and Bob's receiving ledger is compromised, the two ledgers can diverge. In this case, the multi-CBDC is no longer consistent. If Alice and Bob try to transact with a recipient, Charlie, in a third jurisdiction, Charlie will be unable to verify the correctness of either ledger, and therefore cannot verify transaction validity.

Summary

This simple design bears some important similarities to the designs that have been adopted by nearly every multi-CBDC pilot to date. Today, most multi-CBDC pilots rely on enterprise DLT products, which allow users to specify certain transactions as *private*. A private transaction is typically only exposed in plaintext to the payer, the payee, and a small number of specialized nodes called *validators,* which confirm the validity of a transaction (e.g., are there sufficient funds). A key observation is that for these special private transactions, transactions are sometimes validated by very few validators (even just one). This is the case in Corda, a DLT solution that has been used by several multi-CBDC pilots.[13] In Corda, there is a custom consensus protocol that checks for invalid transactions. However, it does not algorithmically reconcile cases when one or more ledgers is arbitrarily compromised. In other DLT offerings, private (encrypted) transactions are not externally validated at all, and are only maintained unencrypted in the payer's and the payee's ledgers. This is the case for Quorum, which has been used in Project Jura.[14]

The limited validation of transactions in these systems is necessitated by the practicalities of private transactions. Since the transactions cannot be widely disseminated—at least not in unencrypted form—they also cannot be validated to the same degree as public transactions. More specifically, since transactions are not shared (in plaintext) with validators, validators are unable to run so-called *Byzantine-fault tolerant consensus protocols*—algorithms that establish a consistent ledger ordering even in the presence of misbehaving participants. These algorithms require at least a minimum number of validators, and are therefore incompatible (to varying degrees) with existing privacy measures in enterprise DLT solutions. This prevents the system from satisfying a basic security guarantee.

Option 2: Global consensus on unencrypted data: Security and scalability, but not privacy

To resolve the security vulnerability from the previous section, a multi-CBDC could choose to broadcast unencrypted transactions to all validators, and have this set of validators run a Byzantine Fault Tolerant protocol. The main difference between this design (Option 2) and the previous design (Option 1) is that all transactions in Option 2 are passed to the entire set of validating nodes. The validating nodes would then conduct a consensus protocol to agree on the ledger state.

Security

Because this design runs a Byzantine Fault Tolerant algorithm to validate transactions, this design is secure against compromised or misbehaving ledgers or validators. Of course, a design can have other security flaws, but in terms of my definition for this article, this design is secure.

Scalability

This design can be scalable, depending on the implementation. If the set of validating nodes is small (e.g., fewer than 20 nodes), the additional communication and computational overhead of

[13] Project Jura: Cross-border settlement using wholesale CBDC, https://www.bis.org/publ/othp44.pdf
Project Dunbar: International settlements using multi-CBDCs, https://www.bis.org/publ/othp47.pdf
Inthanon-LionRock to mBridge, https://www.bis.org/publ/othp40.pdf
Jasper-Ubin Design Paper: Enabling Cross-Border High Value Transfer Using Distributed Ledger Technologies, https://www.mas.gov.sg/-/media/Jasper-Ubin-Design-Paper.pdf?la=en&hash=EF5857437C4857373A9287CD86F56D0E7C46E7FF
"Transactions", Corda Documentation, https://docs.r3.com/en/platform/corda/4.8/open-source/key-concepts-transactions.html

[14] Project Jura: Cross-border settlement using wholesale CBDC, https://www.bis.org/publ/othp44.pdf,
"Private transaction lifecycle", GoQuorum Documentation, https://consensys.net/docs/goquorum/en/stable/concepts/privacy/private-transaction-lifecycle/

running a consensus protocol is manageable.[15] Indeed, such consensus protocols (with low numbers of validators) were the cornerstone of prior proposals for privately-run digital currencies.[16]

However, as the number of validators grows, the efficiency of consensus protocols decreases rapidly.[17] This is a well-known and longstanding problem in the computer science community.[18] Indeed, one of the major technical insights of Bitcoin was to propose a consensus protocol that can scale to thousands of validators without requiring advance knowledge of their identities.[19]

In the context of a multi-CBDC, this raises an important question: who should run validator nodes? If the multi-CBDC is run as a single global ledger (as in Project Dunbar[20]), then each domestic CBDC may want to contribute some validating nodes to the global system. However, if there are hundreds of validators (one per country), each validating all transactions, this will quickly lead to serious scalability bottlenecks, inherently limiting how equitable or distributed a multi-CBDC ledger can be.

Privacy
This design does not provide privacy. It broadcasts all transactions in plaintext to all validators. These validators could be run by domestic or international financial institutions, either in the payer's jurisdiction, the payee's jurisdiction, or a third-party jurisdiction. On the other hand, it enables full transparency for regulatory oversight.

Summary
This general design has been used to process public transactions in multi-CBDC pilots Project Jura and Inthanon-LionRock.[21] It is most commonly used in permissionless cryptocurrencies, such as Bitcoin and Ethereum. In such cryptocurrencies, this design provides only pseudonymity. However, in a multi-CBDC, it would very likely be coupled with identity verification requirements for Know Your Customer compliance. In that case, these designs would provide no privacy at all, but full transparency.

Option 3: Global consensus on encrypted data: Privacy and security, but not scalability

At face value, privacy and security (by my definitions) seem to be at odds. However, a remarkable technology from the cryptography community called zero-knowledge proofs can be

[15] Yin, M., Malkhi, D., Reiter, M.K., Gueta, G.G. and Abraham, I., 2018. HotStuff: BFT consensus in the lens of blockchain. *arXiv preprint arXiv:1803.05069*.

[16] The Diem Team, *DiemBFT v4: State Machine Replication in the Diem Blockchain*, https://developers.diem.com/docs/technical-papers/state-machine-replication-paper/, 2021

[17] Yin, M., Malkhi, D., Reiter, M.K., Gueta, G.G. and Abraham, I., 2018. HotStuff: BFT consensus in the lens of blockchain. *arXiv preprint arXiv:1803.05069*.
Castro, M. and Liskov, B., 1999, February. Practical byzantine fault tolerance. In *OsDI* (Vol. 99, No. 1999, pp. 173-186).

[18] PBFT complexity, scalability of BFT protocols

[19] Nakamoto, S., 2008. Bitcoin: A peer-to-peer electronic cash system.

[20] Project Dunbar: International settlements using multi-CBDCs, https://www.bis.org/publ/othp47.pdf

[21] Project Jura: Cross-border settlement using wholesale CBDC, https://www.bis.org/publ/othp44.pdf
Inthanon-LionRock to mBridge, https://www.bis.org/publ/othp40.pdf

used to circumvent this tension.[22] Zero-knowledge proofs (ZKPs) are cryptographic constructions that allow a prover (i.e., the transaction payer) to prove to a verifier (e.g., a validator) that some conditions hold over an encrypted quantity (e.g., that the transaction is valid and does not double-spend funds) without revealing any of the encrypted data to the verifier.[23] In theory, ZKPs can be used to prove arbitrary functions about an encrypted transaction; in practice, system designers have most successfully used ZKPs that are carefully tailored to specific functions and use cases, such as proving that a transaction spends only available funds.[24]

The final design template makes use of ZKPs to resolve the apparent tension between privacy and security. Under these designs, encrypted transactions are provided to all validators. The validators cannot decrypt transactions, but they can verify the validity of transactions in zero knowledge, even in the presence of Byzantine validators. This design is similar to Option 2, except all transactions are encrypted using ZKPs that are tailored to the validation and transparency requirements of the multi-CBDC.

Privacy
This design is private by design, because only the transaction payer and payee are able to see transaction details in plaintext. In cases where a transaction needs to be passed through intermediaries (e.g., for foreign exchange), the intermediaries may be able to decrypt transactions as well.

Security
This design can be made secure by having validators execute Byzantine Fault Tolerant protocols over the encrypted data. Such a design has been built and tested in production by the cryptocurrency Zcash.[25]

Scalability
Today's implementations of zero-knowledge ledgers suffer from scalability limitations. Specifically, the computational cost of using ZKPs, both for transaction creation and execution, is substantially higher than processing transactions unencrypted. In Zcash, the majority of transactions do not use ZKP-enabled privacy enhancements.[26] While we can only speculate about the reason for this, creating a shielded transaction in Zcash currently takes several seconds, which is at least an order of magnitude longer than it takes to create an unencrypted transaction in many existing cryptocurrencies.[27] These differences are likely to be exacerbated in a multi-

[22] I do not distinguish in this article between zero-knowledge proofs and zero-knowledge arguments, which differ in their technical definitions, but are used in similar ways.
[23] Feige, U., Fiat, A. and Shamir, A., 1988. Zero-knowledge proofs of identity. *Journal of cryptology*, *1*(2), pp.77-94.
[24] Sasson, E.B., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E. and Virza, M., 2014, May. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy* (pp. 459-474). IEEE.
[25] "How It Works", Zcash Documentation, https://z.cash/technology/
[26] Mike Dalton, "Zcash Privacy Back in Question after User Traces Shielded Transaction", Crypto Briefing, https://cryptobriefing.com/zcash-privacy-back-question-user-traces-shielded-transaction/, July 2020
Josh Olszewicz, "Zcash Price Analysis - Shielded addresses underutilized", Brave New Coin, https://bravenewcoin.com/insights/zcash-price-analysis-shielded-addresses-underutilized , September 2020
[27] ibid

CBDC, since the statements that would need to be proved in zero-knowledge would not just be limited to availability of funds, but would also need to encompass other regulatory compliance checks. In particular, they would need to expose enough information to enable both pre- and post-suspicion data sharing.

Summary

Today, running an entire multi-CBDC over encrypted data could incur unacceptable levels of performance overhead due to scalability issues in current ZKP implementations. However, these technologies are advancing rapidly. I believe that these constraints could be resolved in the next couple of years.

In a multi-CBDC setting, another important challenge is how to enable ZKPs to interact across ledgers. Today, cross-chain transactions are typically executed using a construction called a cross-chain atomic swap. This is a sequence of transactions that enable a party to send funds from one ledger to a receiver in another ledger (i.e., another domestic CBDC) without needing to trust a middleman. Typical cross-chain atomic swap constructions require the payer and payee to place transactions on one another's ledgers, and verify each other's transactions on the counterparty's ledger. However, in a cross-border CBDC design that provides privacy by encrypting ledgers, users would not have (plaintext) access to ledgers from other jurisdictions. Broadly, understanding how to build a multi-CBDC across multiple, encrypted ledgers is an open design question.

WHAT NEXT?

When a trilemma is proposed, there are typically two possibilities. The first is that the trilemma is true, and fundamental tradeoffs exist between the proposed quantities. In this case, it is impossible to satisfy all three properties at once. This can often be established through theoretical (mathematical) modeling and analysis.

The second possibility is that the trilemma is not actually fundamental and can, in principle, be broken through the development of new technologies. It is my belief that multi-CBDCs fall into the latter category. Today, such a system—that is, a multi-CBDC that is secure, private, and scalable—is within reach, but it will require new technological advances. These advances are also within reach; it is my opinion that if the appropriate technical requirements are clearly scoped and funded, the tools to meet those requirements would be developed in a matter of 2-3 years.

In my view, the most important precursor to breaking the multi-CBDC trilemma is to clearly define requirements and threat models. To the extent that this exercise has been done (at least publicly), it has been at a high level. I recommend outlining and *publicly* documenting these requirements at a much lower level of granularity and higher level of precision. For example, if a transaction is sent from a payer to a payee in different jurisdictions, and a validator in the payee's jurisdiction is compromised while the transaction is being settled, what are the tolerable outcomes? What happens if the compromised party changes in location, time, or severity of compromise? These questions should ideally be answered in a structured manner in a convening between stakeholders from different jurisdictions. Once multi-CBDC requirements are crisply documented and communicated to the broader technical and research communities, it is quite likely that we will see new designs emerge, as well as stronger, independent validation of current designs.

Regardless of the outcome, it is my belief that broader collaboration between central banks, private industry, nonprofits, academia, and end users is key for accelerating the resolution of the apparent trilemma that characterizes current designs of multi-CBDCs.