Saudi Central Bank and Central Bank of the U.A.E. Joint Digital Currency and Distributed Ledger Project

البنك المركزي السعودي
SAMA
Saudi Central Bank

مصرف الإمارات العربية المتحدة المركزي
CENTRAL BANK OF THE U.A.E.

The views expressed in this report are those of the authors and not necessarily the views of the Saudi Central Bank and Central Bank of the U.A.E.
This report cannot be used to present Saudi Central Bank and Central Bank of the U.A.E. viewpoint.
The report is also concerned with proving the technical feasibility of joint digital currencies only for Aber project scope. It is not a recommendation to any digital currencies.

Contents

Acknowledgement

We would like to express our gratitude to the six participating commercial banks in the two countries, which came together to execute this project; and the various teams from each of our central banks. We also extend our gratitude to IBM our partner in success.

Project Aber: Final Report

Foreword

Final Report Project Aber

Foreword

The Saudi Central Bank (SAMA) and Central Bank of United Arab Emirates (CBUAE) are pleased to present this report on Project Aber: Joint Digital Currency and Distributed Ledger proof of concept Project.

This report reflects our journey in the project that aimed to contribute in the body of knowledge in CBDC and DLT technologies filed.

In the light of the new experiments and researches both central bank led this project as an innovative driven initiative. The initiative sought to explore whether distributed ledger technology could enable cross-border payments between the two countries to be reimagined: using a new, dual-issued digital currency as a unit of settlement between commercial banks in the two countries and domestically. The name Aber was selected because, as the Arabic word, for “crossing boundaries”, it both captures the cross-border nature of the project as well as our hope that it would also cross boundaries in terms of the use of the technology.

Over the course of a year, use cases were designed, implemented, and operated; with the solution, results, and key lessons learned documented in this report.

The project confirmed that Distributed Ledger Technology (DLT) can provide central banks with the ability to reimagine both domestic and cross-border payment systems in new ways. We believe that it represents a significant contribution to the body of knowledge in this field and lays the foundation for future work that we plan to explore in the future.

We are pleased by the promising results, insights, and learnings described in this report and trust that they will benefit the central banking community and broader financial ecosystem in visualizing the potential of this new technology to transform the GCC financial markets and indeed our industry.

Executive Summary

Executive Summary

Project Aber was an initiative launched by the central banks of Saudi Arabia and United Arab Emirates to explore the viability of a single dual-issued digital currency as an instrument of domestic and cross-border settlement between the two countries.

The high-level objectives of the initiative were:

- To explore, experiment, and gain a deeper understanding of distributed ledger technology (DLT) and analyse its maturity;
- To explore an alternative DLT-based cross-border payment solution that can overcome inefficiencies in existing cross-border interbank payment approaches;
- To understand and experiment with the dual issuance of a central-bank digital currency;
- Benchmark findings against those of other central banks.

The project was structured into three distinct phases or use cases:

- Use case one to explore cross-border settlement between the two central banks;
- Use case two to explore domestic settlement between three commercial banks in each country
- Use case three to explore cross-border transactions between the commercial banks using the digital currency.

Following an extensive assessment of the current payment systems, prior work in the application of DLT in the field, and informed by the current state of the art in DLT technologies, a number of key principles were agreed by all participants to guide the execution of the project.

Firstly, commercial banks must be active participants, running local nodes on the network and engaging the fullest from a technical and business perspective throughout the lifecycle of the project. This was to ensure that the employees of both central and commercial banks would benefit from the acquisition of knowledge around this new technology and also so that the project could be better informed as to challenges, risks, or improvements, from a commercial bank perspective, that would need to be addressed if the full value of the technology was to be realized in this context.

Secondly, real money would be used in the project. This was important because it forced greater consideration of the non-functional aspects, such as security, that

would need to be addressed moving forward; and also how the system would interact with existing payment systems, such as the domestic RTGS system.

Thirdly, rather than simply replicate the way in which conventional payment systems work, the project sought to explore how such systems can leverage the unique characteristics of DLT to drive greater levels of distribution. By doing that, it sought to develop a system that was more resilient to single points of failure.

The project confirmed that a cross-border dual issued currency was technically viable and that it was possible to design a distributed payment system that offers the two countries significant improvement over centralized payment systems in terms of architectural resilience. The key requirements that were identified were all met, including complex requirements around privacy and decentralization, as well as requirements related to mitigating economics risks, such as central bank visibility of money supply and traceability of issued currency. The performance objectives that were originally set for the project were exceeded, proving that DLT technologies could offer high levels of performance whilst not compromising safety or privacy.
As such, the project has confirmed the viability of DLT as a mechanism for both domestic and cross-border settlement and confirmed the technical viability of a single digital currency issued by both central banks. The project has also identified further areas that need to be explored in the future if the approach of a single digital currency is to be implemented: key amongst these are the need to understand impacts to the monetary policy of participating states and to address, in particular, the means by which interest is calculated and disbursed to the commercial banks in each jurisdiction and how this can be applied with a single digital currency.

In terms of future work, there are many directions that this project can evolve towards. Firstly, it could provide the basis for a backup to domestic and regional RTGS; providing a more distributed and potentially resilient alternative to the centralized systems that are implemented or being implemented today.
Secondly, by offering a DLT-based payments rails, there is the possibility to expand to Delivery versus Payment (DvP) scenarios such as using the Aber network as a means of settlement for other forms of transaction, such as the sale of bonds or other dematerialized assets. Thirdly, there is the possibility of extending it geographically to include regional or other international central banks or linking heterogeneous networks together.

In summary, the project was successful in meeting its objectives, demonstrated possible incremental benefits of this new approach to payments, identified important

lessons learned that can benefit other central banks exploring the field, and has identified several areas of future expansion that can be considered by either the participants in this project or other central banks. As such, we believe this project has made a material impact on industry understanding of the field and is a substantial contribution to the body of knowledge in how the emerging technology of DLT can be applied to cross-border and domestic payments.

Chapter 1:
Introduction

Background

Saudi Central Bank (SAMA) and Central Bank of the UAE (CBUAE) announced a joint digital currency initiative in January 2019. It was named Project Aber, which literally translates to 'one who crosses boundaries' highlighting the cross-border focus of the project.

The unique aspects of Aber

The following key objectives were outlined before the start of the project:

- To explore, experiment, and gain a deeper understanding of distributed ledger technology and analyse its maturity;
- To explore an alternative DLT based cross border payment solution that can overcome inefficiencies in existing cross-border interbank payment approaches;
- To understand and experiment how to dually issue a central-bank digital currency;
- Benchmark findings against those of other central banks.

The central banks laid out several principles which formed the basis for deriving the objectives and requirements of the project. Whilst there have been other central bank digital currency experiments around the world, the unique elements below made Project Aber a "first of a kind".

Active commercial bank participation

An important aspect of the project was active business and technical involvement of commercial banks from a very early stage. A total of six commercial banks (three from each jurisdiction) participated in all phases of the project. This allowed an opportunity for commercial banks -- key stakeholders of any wholesale CBDC project — to have a first-hand experience of using and operating a distributed ledger based interbank payment solution. The deployment architecture was highly distributed and provided complete flexibility to banks in hosting their environments. The graphic in Figure 1 shows the six commercial banks that participated in Project Aber.

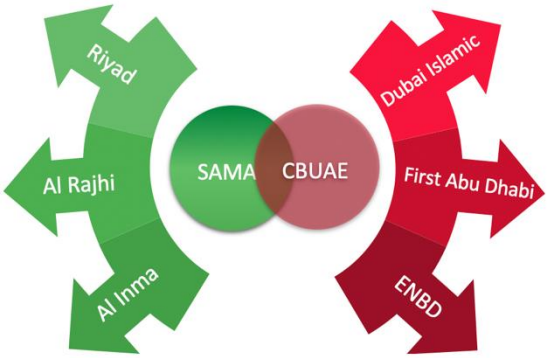

Figure 1: Stakeholders of the Aber Ecosystem

Simplify Cross Border Payments

The vision of Project Aber was to create a central bank digital currency that could be used for settlement of cross-border payment obligations between commercial banks. Such an instrument could be especially useful in a region like the Gulf Coordination Council (GCC) states where there is substantial intra-region trade and movement of citizens and residents. This could address an inefficiency in the existing correspondent banking-based payment systems that often results in delays and required commercial banks to maintain substantial Nostro accounts with their correspondent banks. This has been characterized as "trapped liquidity" in some studies [McKinsey, 2016] as, for some commercial banks, there is an opportunity cost in maintaining these balances and it creates material compliance overheads. In contrast, in Project Aber, the movement of funds would be done in real time, without the need for the commercial banks to have correspondent bank Nostro account in each country.

While there have been other CBDC experiments involving two central banks, the requirement of having a single network and single digital currency for settlement of cross-border payments was a unique aspect of Project Aber.

Real Money in Operations Phase

Another unique aspect was the use of "real money" in the pilot project. This was achieved by commercial banks pledging real money from the deposits that they held with the central bank; using these funds to then fund their digital currency accounts on the distributed ledger. These funds would then be subsequently returned to the commercial banks' accounts once the project had completed.

Having a three month operations phase and the use of real money had the following tangible benefits:

- Use of real money motivated both commercial and central banks to think about how the digital currency will be managed in their books and which core banking systems would be impacted if such a system was to be implemented in production and scaled. This led to a series of lessons learned and observations that would have been unlikely to have been made if real money had not been pledged and used;
- Having three months of operations phase, including several knowledge transfer sessions for banks, provided business and IT operations teams of the banks with invaluable hands-on experience in managing a DLT based payment system;
- The project team received valuable feedback from the banks on additional features and tooling that would be required or would be useful in managing and using a CBDC of this type in the future;

The experience gained and lessons learned from the operations phase is expected to help streamline issues in the eventual rollout of a DLT based interbank payment system as an alternative to traditional RTGS and international wire transfers;

Maximum Decentralization

One of the key mandates was to design a solution that is decentralized to the maximum extent possible, the motivations for which were:

- A truly decentralized solution would increase the architectural resilience of a mission critical financial system by mitigating the impacts of single points of failure and risk;
- The objective is also well aligned with the cross-border focus of the project, ensuring equal authority and participation from banks in both jurisdictions;
- A solution that does not realize this objective would not benefit from one of the distinguishing and unique qualities of blockchain technology;

Chapter 2

Use Cases & Project Phases

Use cases

The project addressed three high level use cases and was executed in three phases. The three use cases were as follows:

- Use Case 1 (UC1): Payment between central banks
- Use Case 2 (UC2): Domestic Payments between Commercial banks
- Use Case 3 (UC3): Cross-border Payments between Commercial banks

As seen in Figure 2, Use Case 1 involved creating a shared ledger between the two central banks on which various digital currency transactions could take place.
Use Cases 2 and 3 were focused on inter-commercial bank payments. Use Case 2 addressed domestic payments between commercial banks in each country; while Use Case 3 extended this to allow cross-border transactions between commercial banks. Since settlement was on an individual payment basis, these use cases imitated the functionality of a national RTGS or cross-border RTGS systems.
While central banks were not directly involved in the payment flow in these use cases, they were the creators and destroyers of digital currency, as well as, auditors for all transactions. Use Case 3 also required settlement between central banks in some situations. This was done using the shared ledger between central banks. Thus Use Case 1 could be seen as a building block for Use Case 3.

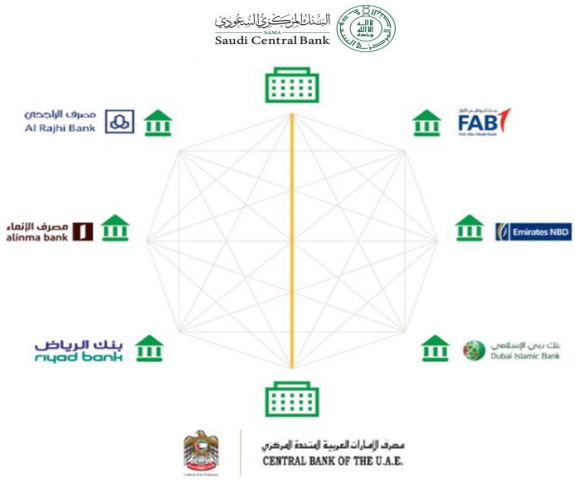

(a) Use case 1

(b) Use case 2

Figure 2: The three primary use cases for Project Aber (a) UC1, (b) UC2, (c) UC3

Project Execution

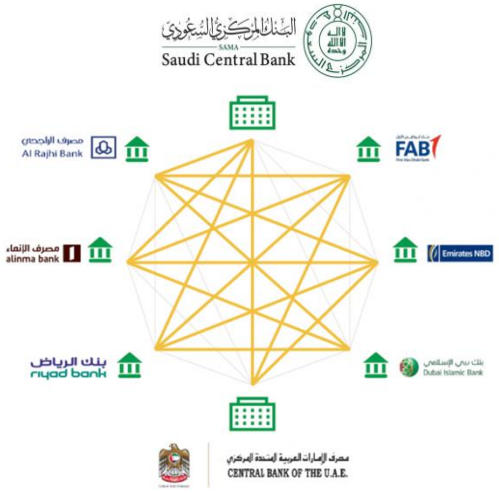

(c) Use case 3

Figure 3: Timeline of Project Aber

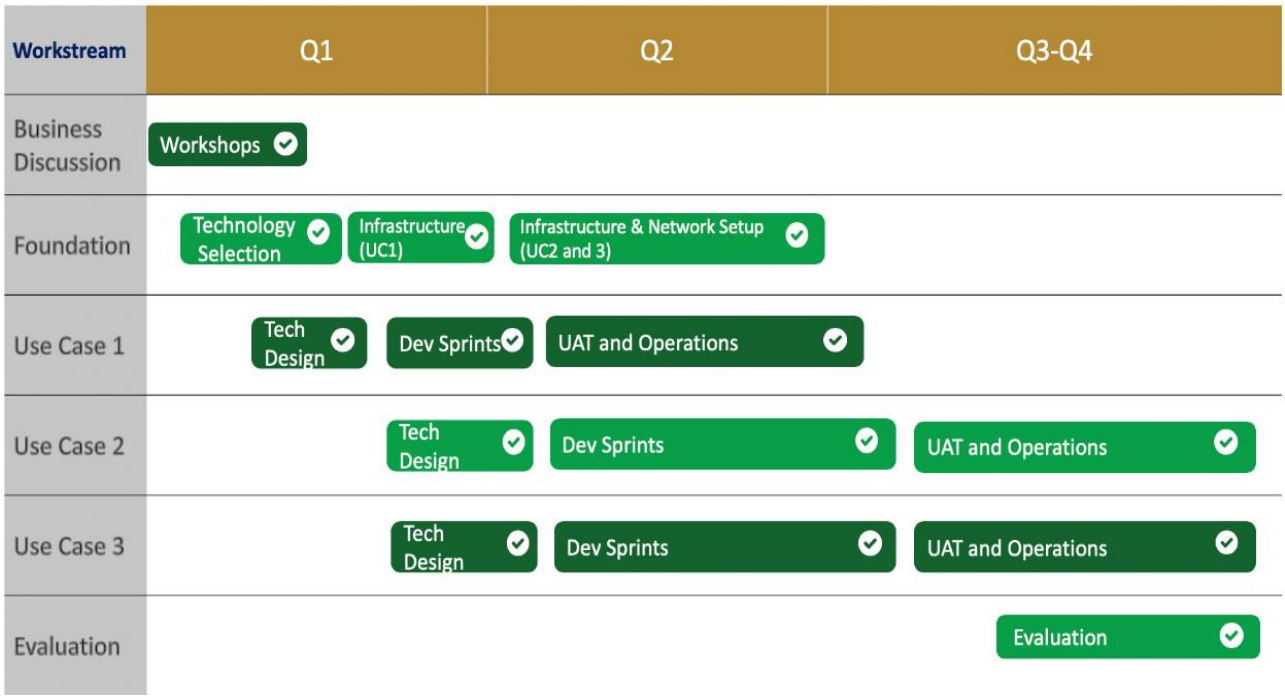

The project was carried out in agile sprints with design, development, and operations phases of the three use cases overlapping with each other as shown in Figure 3. The project began with blockchain seminars focused on general blockchain knowledge transfer followed by focused sessions on the various aspects of CBDCs, such as the different types, prior work, benefits, and risks. This laid the foundation for a series of business workshops that were focused on detailed requirements gathering, and involved active participation from different stakeholders across the central and commercial banks.

The first major milestone was the full implementation of Use Case 1 which was delivered in Apr 2019. It was followed by a 3 month operations phase. In parallel, the development of Use Cases 2 and 3 was completed in July 2019. This was followed by a four month operations and evaluation phase.

Chapter 3

Prior CBDC work

Defining Central Bank Digital Currency

Prior to exploring the different CBDC projects, it is important to first define what a CBDC is – and what it is not. In classifying money, the Bank of International Settlements (BIS) established a framework or taxonomy of money that, whilst based on a Venn diagram, is colloquially called the "Money Flower" [Beck, 2017].

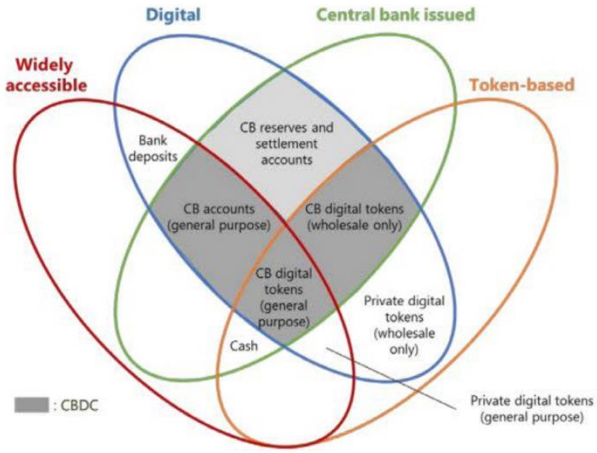

Figure 4: Money flower from BIS report on Central Bank Cryptocurrencies. Source: [Beck, 2017]

The Money Flower describes four key properties of money: the issuer, the form, the accessibility and the technology.

The issuer refers to who has issued the money. It could be the central bank or, in the case of cryptocurrencies or other forms of tokenized value, it could be some private sector entity such as a commercial bank or a firm. In the context of CBDC, we are, of course, referring to central bank-issued forms of money.

The form is the physical form that the money takes. It could be physical, in the case of cash, or it could be digital, such as reserve accounts or bank deposits.

The accessibility refers to who is able to access the particular form of money and this represents an important bifurcation. Some money is accessible only to the central bank and a select number of commercial banks whereas other forms of money, like cash or bank deposits, is available to a broad section of the population and hence is considered generally accessible. In the context of CBDCs, we therefore distinguish similarly between those that are available to the general public (known as retail CBDCs or generally available CBDCs) versus those that are only available to a select number of parties or wholesale CBDCs. In the case of Project Aber, we were focused on developing a CBDC that would be used as an instrument of settlement by participating commercial banks and would not be made available to the general public: hence, it considered a wholesale CBDC.

The final property of money relates to the technology. There are some forms of digital money that are token-based and some that are account-based. The key distinguishing factor is that token-based approaches represent money as a set of tokens that are held by their owner (similar, in many respects, to physical cash) and can sometimes be bearer instruments (also similar to cash). Account-based approaches hold balances in accounts like a bank account and therefore do not have similar characteristics such as cash but more closely resemble bank accounts. Retail CBDCs that seek to simulate or replace cash have tended towards token-based approaches or a hybrid, whereas wholesale CBDCs have tended to use account-based approaches.

There have been numerous papers written by central banks exploring the different practical and theoretical aspects of both wholesale and retail CBDCs. In the case of retail CBDCs, whilst there is a recognition of the role they can play as an alternative to cash in societies where cash use is declining, such as Sweden, there are significant risks that have been raised by multiple central banks around, for example, the disintermediating impact on the banking system if the general public gets access to central bank money and the risk that such access could introduce in facilitating "bank runs" at times of stress in the banking system. As such, most central banks have focused instead on exploring wholesale CBDC use cases.

Potential Benefits

The first potential benefit is that wholesale CBDCs on DLT can improve the architectural resilience of the payment infrastructure in a country through their relative decentralization. Conventional systems are often prone to single points of failure, either from a technology perspective or an organizational one, such that, if the system goes down or fails, there can be substantial adverse economic impacts. DLT holds out the promise of having decentralized systems where the failure of a single node will not materially impact the payments infrastructure of a country.

The second benefit relates also to resilience but wholesale CBDC on DLT can lead to improvements in security. By using wholesale CDBCs as a backup, for example, to RTGS in a country, the overall security of the payments system can be improved. The reason is that DLT is a fundamentally different technology so attack vectors used to compromise the core system cannot be reused for the backup.

A third benefit is that DLT-based solutions can be designed to make on-boarding of non-bank participants, such as FinTechs and other entities, easier and lower cost. This can support the objectives of some jurisdictions to improve diversity in access to payment systems.

A fourth benefit is that, as Blockchains are adopted as the basis for the tokenization of different types of security or asset, such as bonds and debentures, there is a need to solve the Delivery versus Payment (DvP) problem for these assets. The ability to provide a payment rails that can be used for the atomic swap of assets for CBDC is an important benefit, particularly as the adoption of blockchain becomes more mainstream. It is this benefit that underpins some of the recent central bank projects that have focused on these DvP or Payment vs Payment (PvP) aspects.

Evolution of CBDC

Central banks around the world have been evaluating distributed ledger technology and CBDC for wholesale payments. There have been several proof of concepts leveraging from each other's experience. The evolution of CBDC has gone through the following stages:

Explore: These were early experiments done in 2016-2017 timeframe. The idea was to get acquainted with distributed ledger technology, its complexity, and its new programming model. Enterprise blockchain technology was in its infancy, so many of these experiments started with public Blockchains. Use cases were built around simple asset transfers. Several practical concerns such as privacy or correctness guarantees were generally ignored as the focus was primarily on understanding the unique aspects and potential of the technologies.

Evaluate: This stage started once several permissioned DLTs had started demonstrating the promise of production readiness. This allowed central banks to investigate designs that addressed privacy and finality in a distributed architecture. Performance and scalability of the networks was also examined during this stage.

Envision: In this stage, the use cases moved beyond simple domestic payments to those where blockchain demonstrated promise compared to traditional systems. Examples of dominant use cases in this stage were: cross border payments, interoperability between multiple ledgers or technologies, tokenization of multiple asset types, and atomic swaps.

Expedite: As DLT platforms have matured and successful CBDC experiments helped make the technology more viable for use within banking contexts, central banks have started joint exercises with commercial banks in order to build confidence and gain acceptance from key stakeholders.

Figure 5 shows a chronology of some of the major CBDC experiments around the world. The timeline also shows how projects have had the opportunity to leverage and extend endeavours from other banks as well as their own. The projects are colour coded, with the lightest shade representing Explore stage and the darkest shade

representing Expedite stage. Project Aber demonstrates characteristics of Evaluate, Envision and Expedite stages of CBDC evolution.
Some of the central banks leading the innovation agenda in this space are: Bank of Canada, European central bank (ECB), Bank of Japan (BoJ), Monetary Authority of Singapore (MAS), South African Reserve Bank, and now SAMA and CBUAE.

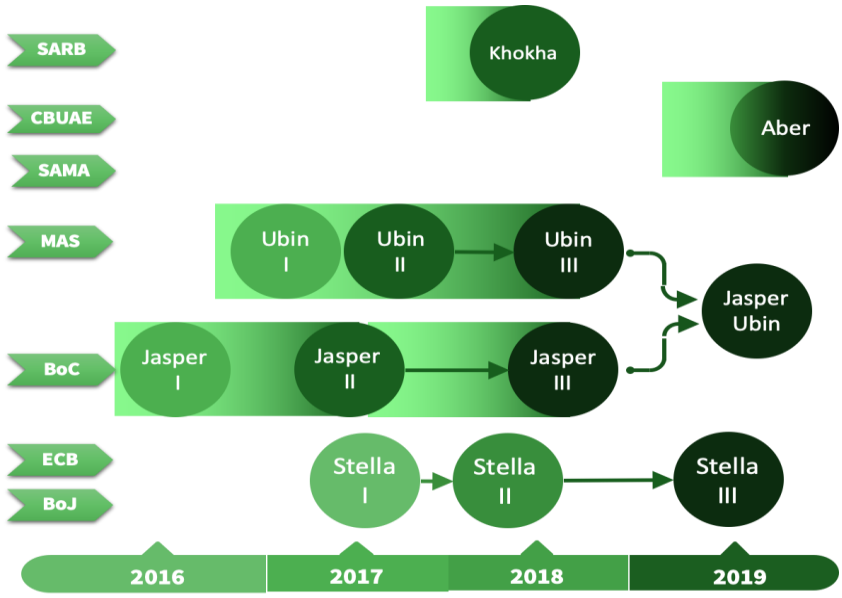

Figure 5: A timeline of wholesale CBDC projects

Central bank Proof of Concepts: An Overview

In this section an overview of objectives and findings of some of the previous central bank PoCs are discussed. The summary of each project has been sourced from its cited project report. This will help set the context for some of the key motivations and principles that underpinned Project Aber.

Project Jasper

The project represents a collaborative effort between Payments Canada, its member financial institutions, the Bank of Canada, and other market participants
Its aim was to investigate and understand how DLT could transform the future of payments in Canada. The project has delivered three phases as detailed below:

Jasper Phase 1

Phase 1's objective was to experiment with central bank-issued digital receipts for deposited currency to support settlement on a DLT platform. Phase 1 was delivered on the Etherium platform and used a 'proof of work' consensus protocol.
The solution delivered the identified functionality successfully in a non-production setting. However, there was a concern cited in terms of efficiency as well as a number of gaps identified [BoC, 2017]. For example, there wasn't a clarity on transaction finality (due to the use of proof of work), there were some operational risks shared, and participant data privacy wasn't supported as participants had full visibility into the central bank ledger.

Jasper Phase 2

Phase 2 was launched in late 2016 (report released in September 2017) and intended to assess and evaluate the scalability and flexibility of DLT by moving to an alternative technology platform and to add more functions to the existing interbank settlement systems. As such, Jasper Phase 2 was built on a permissioned DLT platform from R3 called Corda [R3, 2019] which used 'notary node' as an alternative consensus model to the proof of work used in Phase 1. The solution also supported liquidity saving mechanism (LSM) using a centralized queue that settles batches of queued payments on a net basis. This promotes funding efficiency and enables a smoother intraday flow of payments. Phase 2 achieved the intended functionalities, such as user privacy, processing and high-volume transactions within an acceptable window. However, platform operational resiliency wasn't fully achieved due to some required activities to be performed by centralized notary node which can create a single point of failure [JAS].

Jasper Phase 3

Phase 3's objective was to explore the ability of DLT to transform payments and securities settlement in Canada. The Proof of Concept (POC) intended to provide delivery vs. payment (DVP) settlement by integrating securities and cash ledgers. The PoC was, like the previous phase, built on R3's Corda platform. To summarize the result, payments and security settlement were successfully implemented and demonstrated using DLT, however, efficiency gains and cost saving from using the DLT solution could not be determined from this experiment as broader scope should also be considered [BoC, 2018].

Project Stella

In December 2016, the Bank of Japan (BoJ) and European central bank (ECB) jointly announced project Stella: a research project to study and understand the use of DLT

as the basis for financial market infrastructure. Project Stella has had three phases so far; completed in September 2017, March 2018 and June 2019.

Stella Phase 1

Phase 1's goal was to determine whether a DLT could efficiently and safely run the liquidity saving mechanism (LSM) of BoJ and ECB's RTGS systems. The project was built on Hyperledger Fabric [HLF, 2019] platform and ran two types of smart contracts. The first one without queuing and offsetting features and the second offering LSM based on queuing and offsetting of their RTGS systems to process the payments. As a result of the findings, the project confirmed that DLT could meet the desired safety and performance needs. However, when number of requests were above average traffic levels, the system performance suffered. The conclusion was that, whilst some functional characteristics could be met, DLT was not yet mature to run large-scale systems such as RTGS due to non-functional limitations on performance and scaling [ECB,BoJ, 2017].

Stella Phase 2

Stella Phase 2's goal was to explore delivery versus payment (DVP) for securities and specifically how the delivery of securities against cash payments could be conceptually designed and operated using DLT technology. Prototypes were built using three DLT platforms: R3 Corda, Elements, and Hyperledger Fabric. Phase 2 successfully delivered DvP using DLT platforms however more exploration was suggested on safety, efficiency and legal aspects [ECB, BoJ, 2018].

Stella Phase 3

The objective of phase 3 was to explore whether cross-border payments could be improved with respect to interoperability, efficiency, and safety aspects using DLT technology. As part of this, they wanted to study the Interledger Payments (ILP) [Thomas, 2015] protocol that synchronizes payments between different types of ledgers (a centralized ledger and a DLT ledger, DLT ledgers, and centralized ledgers). Hyperledger Fabric protocol was used in this phase. The conclusion of this project was that using synchronized payments and locking funds in a payment chain could indeed improve the safety in cross-border payment transactions however legal, cost-benefit analysis and technology maturity aspects need to be looked at before considering implementing this new method of cross-border payments [ECB, BoJ,2019].

Project Ubin

In late 2016, the Monetary Authority of Singapore (MAS) commenced a collaborative project with financial institutions and technology providers to explore the use of

Distributed Ledger Technology (DLT) for clearing and settlement of payments and securities. This initiative was known as Project Ubin.

Ubin Phase 1

In Phase 1, the objective was to study the feasibility of using a central-bank-issued digital currency for interbank payments. A digital form of the Singaporean Dollar, backed by central bank money, was tokenized on Etherium based distributed ledger [MAS, Mar 2017]. The project proved the technical possibility of tokenisation using DLT.

Ubin Phase 2

Phase 2 commenced in July 2017 with the goal of studying the potential implications of deploying DLT for specific RTGS functions. In particular, the following functions were considered as they were determined to be key factors that DLT would need to meet in order to be viable as an RTGS alternative or alternative to subset of RTGS functions:

- Digitization of payments;
- Decentralized processing;
- Privacy of transactions;
- Settlement Finality;
- Liquidity Saving Mechanisms;

Three DLT platforms were evaluated in parallel: R3 Corda, Hyperledger Fabric, and Quorum (a JP Morgan created private version of the Etherium protocol). Solutions based on the three platforms were designed for a common set of functionalities, with the goal of conducting an objective (but non-quantitative) assessment of each of the three technologies and their ability to meet the requirements of the overall initiative. To summarize the findings, the Corda based solution had some privacy limitations (as noted in the findings section of [MAS, Nov 2017]). While the Hyperledger Fabric (HLF) based solution had good privacy properties, both HLF and Corda relied on some central services for ensuring safety (correctness). The Quorum-based approach used zero-knowledge proof techniques for privacy but suffered from higher latencies [MAS, Nov 2017] than Corda and Fabric.

Jasper-Ubin

Bank of Canada (BoC) and Monetary Authority of Singapore (MAS) collaborated on Jasper Ubin Project and announced the results in May 2019 [BoC,MAS, 2019]. The main objective of this project was to experiment in the use of distributed ledger technology for cross-border high value payments. They used different DLT networks in different

jurisdictions without involving single trusted entity with the intention of increasing the efficiency and reducing risks for cross-border payments. The use case built a cross-border payments solution between Canada and Singapore that involved cross-currency (CAD and SGD) and cross platform (using Corda for Canada and Quorum for Singapore) DLT protocols for atomic transactions. They used a technique called Hash Time-Locked Contracts (HTLC) to connect the two networks and allow Payment vs. Payments (PVP) settlement without the involvement of a trusted third party acting as an intermediary for settlement. To summarize the findings, the project successfully delivered the desired objective by performing atomic transactions between two different DLT networks - Quorum network in Singapore and Corda network in Canada - using HTLC. However, the solution is subject to single point of failure due to the involvement of intermediaries in all transactions, also, further exploration on using HLTC for more than two networks would be required.

Project Khokha

In late 2017, the South African Reserve Bank (SARB) commenced a collaboration with seven banks, technology and consulting providers to explore the use of distributed ledger technology (DLT) for interbank settlement in South Africa.

In early 2018, SARB launched Project Khokha with the objective of simulating a "real-world" trial of distributed ledger technology-based interbank clearing and settlements system modelled on South Africa's real-time gross settlement (RTGS) framework. The project sought to assess the following functions of a DLT solution: performance, scalability, privacy, resilience, and finality. Whereas other central bank projects had focused in large part on the functional aspects of a DLT-based payments system or currency, Khoka was focused on proving the ability of DLT to meet non-functional aspects, notably scale and performance. SARB decided to use Quorum as a DLT platform to build the project. The Khokha project findings confirm that the Quorum platform delivered the required performance and exceeded the performance criteria in some cases. Also, the confidentiality and privacy of transactions between commercial banks were achieved. Some areas were not considered for evaluation such as risks related to integrity, security and availability of the network, in addition to legal and regulatory factors. As per their report, these should be assessed and evaluated further before implementing such a system [SARB, 2018].

Summary

Figure 6 is a summary of focus areas and technology platforms for the advanced CBDC projects. This helps put Project Aber in context relative to what has preceded it.

The focus on real money and a co-executed operations phase between network participants and technology provider was a significant step in the direction of expediting usage of CBDC for interbank payments. While other central banks have also explored cross-border payments, the major difference was in Aber's dually-issued single digital currency approach and use of real money.

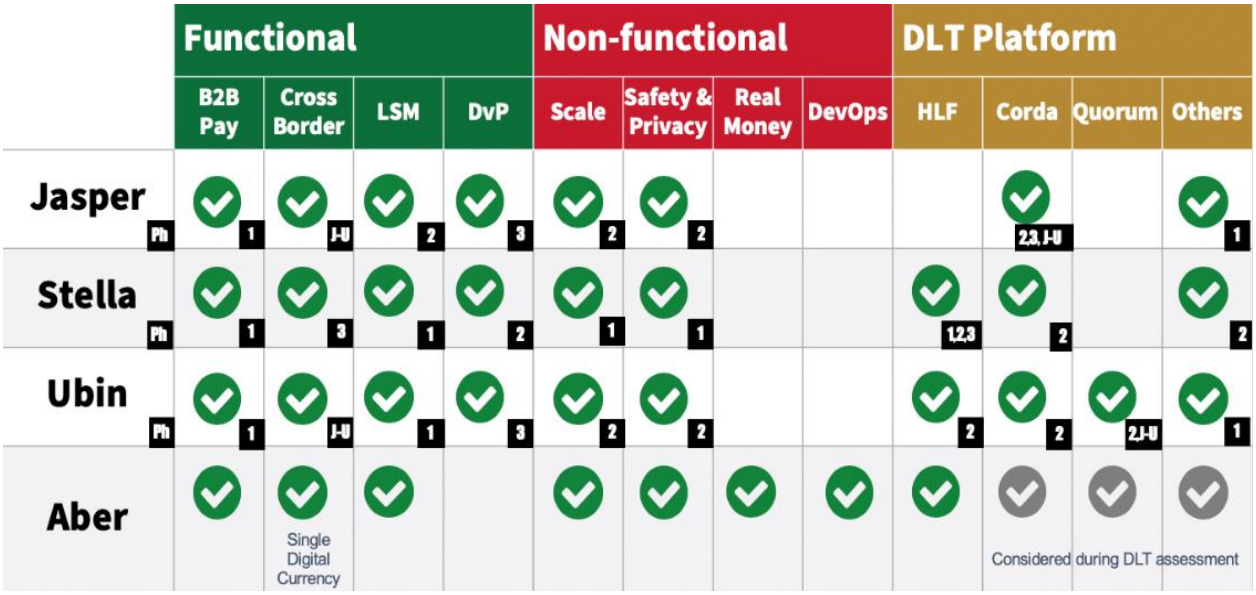

Figure 6: Focus areas and technology platforms in CBDC projects.

Chapter 4

Aber: Business Requirements

Digital Currency Lifecycle

Figure 7 shows the lifecycle of digital currency from issuance to destruction. Typical lifecycle of digital currency is described below:

1. commercial bank will pledge cash collateral in an account held by central bank;
2. central bank converts cash collateral to generate new digital currency;
3. central bank funds the newly created currency in the commercial bank's account on the ledger;
4. The commercial bank transfers the new currency to an account belonging to the Counterparty on the ledger;
5. The Counterparty redeems the currency for cash collateral via the central bank in its jurisdiction;
6. The central bank destroys the created currency for this transaction.

While all of the digital currency goes through this cycle, step 4 can be repeated as many times as needed as the digital currency is exchanged between participants. Finally, the issued currency is destroyed as part of a redemption request in which it is converted, by the central bank, back to cash and deposited back with the commercial bank.

Business Requirements

Business requirements apply to one or more steps in the lifecycle:

1. System should be decentralized to maximum extent possible: The system should enable commercial banks to settle with each other even in cases where the central bank is unavailable or disconnected from the network. Functions that are purely central bank functions, such as issuance, pledging and redemption, would not be available but the core functions of domestic and cross-border settlement should be available. Only counterparties to the payment transaction need to be online for the payment to be settled. The rationale behind this is to enable the system to offer a higher level of architectural resilience than traditional centralized systems that depend on the availability of centralized services and thus avoid a single point of failure.

   Applies to: Transfer (step 4)

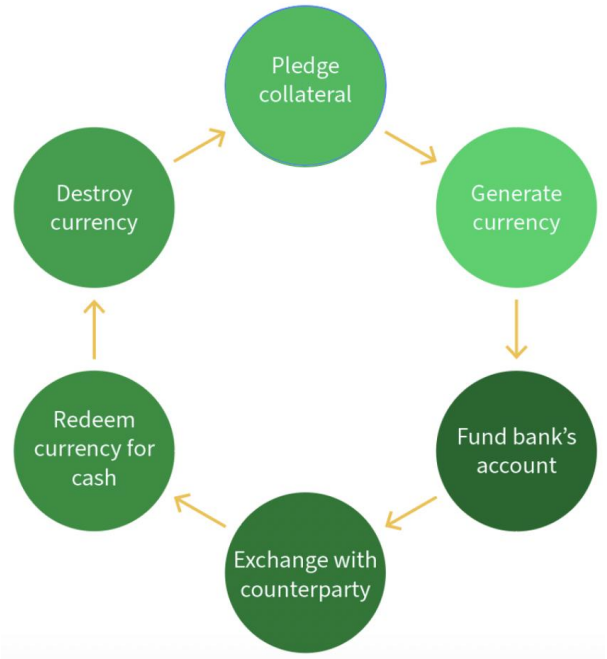

Figure 7: The digital currency lifecycle

2. The Currency shall be pegged: Given that both SAR and AED are pegged currencies, the digital currency that will be used as the Medium of Exchange for the project shall be pegged.

Applies to: Generate, Redeem (Steps 2, 5)

3. The Exchange Rate or Conversion Rate for fiat to Digital Currency shall be fixed: The rate at which the commercial bank will receive the new currency should be fixed for the duration of the project. This was a reasonable assumption given both Saudi Arabia Riyal (SAR) and UAE Dirhams (AED) are pegged.

Applies to: Pledge, Generate, Redeem, and Destroy (Steps 1, 2, 5, 6)

4. Same currency shall be used for domestic and cross-border usage: The same digital currency will be transferrable and redeemable both domestically and cross- border. This increases the utility of the currency as a medium of exchange.

Applies to: Transfer, Redeem, and Destroy (Steps 4, 5, 6)

5. Central banks should have full visibility of money supply: Each central bank will be able to see the total amount generated and issued by the other central bank. Each bank will be issuing a currency that could be used/redeemed in the other jurisdiction, representing a potential liability for both banks. It is thus important, that both central banks have visibility to all digital currency issued in the network.

Applies to: Fund (Step 3)

6. Only central banks can issue currency: This is to ensure that all digital currency used in the system is backed by central banks and can be used to replace fiat payments between banks.

Applies to: Pledge, Generate (Steps 1, 2)

7. Resilience to gridlock: The system should be resilient to gridlock scenarios where transactions are submitted but there is insufficient liquidity in the system to fulfil at that point in time. The system should incorporate appropriate algorithmic approaches to solve this problem in a decentralized way, such as multi-lateral netting or other liquidity saving mechanisms.

Applies to: Transfer (Step 4)

8. Cross Border Issued Currency: Currency issued by a central bank should remain liability of the issuing bank, regardless of the jurisdiction of redemption. The solution should thus support settlement between central banks, when redeeming cross-border issued currency. Having the currency as a claim on the central bank that issued it, simplifies management of the risk around cross-border impacts on money supply.

Applies to: Redeem, Destroy (Steps 5, 6)

Non-Functional Requirements

Security and Privacy

The solution should support Hardware Security Module (HSM) technologies for safeguarding key storage.

- The solution should support Public Key Infrastructure (PKI) based blockchain permissioning: only principals having identities issued by one or more verified certificate authorities will have access to the network.
- Communication Security: The messages exchanged between blockchain peers should be encrypted using transport level encryption.
- The solution should define a point after which settlement can be considered final and irrevocable.
- The solution should be able to provide safety against the types of attacks generally referred to as double spending. It should be resilient against any other type of forgery, concealment, and non-repudiation.

- Only relevant parties can see transactions and balances: Commercial banks should see the transactions in which they are a party but should not be able to see other transactions in the network. They should not know the account balances of the other commercial banks. Central banks, on the other hand, should not have visibility to domestic transactions and ledger balances from other jurisdiction. This requirement is discussed in more detail **under Key Design Considerations**.

Auditability

- All currency should be linked to the central bank who issued it: The Digital Currency should be traceable to the issuing central bank. At any stage in the currency lifecycle, it should be clear which authority issued the currency. Each central bank should be able to see, in real time, the current amount of currency by issuer in the system and, upon presentation by a commercial bank for redemption, know who issued the digital currency and from where the funds for redemption should be sourced. The central banks should have full visibility to any monetary activity involving a bank in their respective jurisdiction. This includes both domestic and cross-border transactions. In addition, the central bank should have full view of the account balances of banks within its jurisdiction.

Scalability

- The system should have the ability to support a large number of network participants: While the pilot will have a maximum of eight participants, the solution itself must be able to scale to large number of banks (e.g. 50-100 commercial banks and up to six central banks). The architecture should be able to scale horizontally with the number of nodes in the network.

Chapter 5

Key Design Principles

Prior to the development phase, there were a series of discussions between the business and IT teams in the stakeholders to determine a set of principles that would guide the successive phases of the project, minimize risks, and ensure that the project would be focused on what was important in the context of the objectives set by the central banks.

Minimize Business Impact to As-Is Systems

Impact to production systems and business processes should be minimized to the maximum extent possible.

As a limited pilot, albeit with real value being exchanged, there was an early risk identified that any changes to existing systems or processes would introduce potentially prohibitive disruption and delays. As such, whilst it was important that the system interact, in some ways, with existing systems, such as Real-Time Gross Settlement (RTGS) in the two central banks and commercial bank's core banking systems, it was also important that such integrations or interactions didn't require any change to these systems. As such, a key design principle was to ensure that there were no requirements introduced that would drive changes to existing systems or create risks that would need to be considered by the owners of those systems. This meant that, whilst APIs would be built and made available for optional system-level integrations, there should be a web interface that participants from the different banks could use to interact with the system thus mitigating the need for any significant process changes or system modifications.

Decentralize with Safety

The system should focus on a higher degree of decentralization than is currently present in the existing payments infrastructure in the two countries.

A key design principle was that the system should go beyond just replicating the current "business model" on a distributed ledger. Instead, it should take advantage of DLTs potential to allow far more distributed technical and business architectures and should aspire to achieve the highest degrees of decentralisation whilst ensuring that privacy, security, performance and other requirements are met. For example, whereas current systems are subject to a

single point of failure due to centralisation, the Aber network should be resilient to such points of failure and allow transactions to be executed without the direct support or involvement of the central banks. Whilst some functions, such as currency issuance and redemption, would implicitly require the involvement of the central banks, the system should be designed such that, in the event of the central bank being unavailable, commercial banks could still transact. This would represent a significant improvement over current systems and more fully take advantage of DLT's unique characteristics. At the same time, it is important that privacy, security, and finality are not compromised as a result of this design principle.

Scalable Design

The system should be designed with the future possibility of expansion and rollout in mind. Whilst a limited pilot, it is important that the system should be designed in a way that it is possible to extend it and move forward into production or at-scale deployments.

Firstly, the system should be able to scale to accommodate the probable transaction volumes involved in at at-scale deployment. It should not require a fundamental redesign or rework of the architecture in order to achieve this.

Secondly, the system should be designed in such a way that other participants can be added in the future and without needing redesign or rework. This could be new commercial banks within the existing jurisdictions, or it could be expanding to new jurisdictions and adding additional central banks.

Thirdly, there are assumptions that will be made for the purposes of the Pilot, such as the use of fixed foreign exchange rates between the Digital Currency and the UAE Dirham and Saudi Riyal. These assumptions should be represented in the system in such a way that, in the future, they can be changed without the need for a complete redesign.

Fourth, future use cases would be all or a subset of those domestic use cases that use RTGS and cross-border use cases; therefore, it's important that the system can accommodate these in the future, such as enabling the use of SWIFT messaging formats or similar.

Fifth, it is known that different jurisdictions have different economic policies and practices with respect to the central banks and their commercial banks. Key amongst these is the payment of interest for overnight deposits by the

commercial banks with the central bank. E.g. in the case of UAE and KSA, KSA may pay interest while UAE may not. The system may not address these differences now but should, particularly with regards to the calculation and payment of interest, be able to accommodate this in the future.

Incentivize Use of Digital Currency

Design decisions should prioritize utility of the currency to the commercial banks. When designing the system, consideration and prioritisation should be given to requirements that will maximize the utility and usefulness of the system to the commercial banks; meaning that it should offer them value that exceeds what they can currently get from their existing payment systems and traditional approaches to cross-border payments, such as the correspondent banking system. Many requirements, such as using the same currency for domestic and cross-border or allowing the inclusion of metadata or the introduction of Liquidity Saving Mechanisms, were formulated based on this; and, through discussions with the commercial banks, there was identification of further value that the system could offer if expanded further, such as being used as the basis for trade finance (i.e. Letter of Credit) between the two countries.

Secure and Private by Design

System should be designed with equivalent levels of privacy as exist today. In a conventional system, the central bank has visibility over all the accounts of the commercial banks who hold accounts with them but each commercial bank has no visibility into the balances of their counterparties. As a design principle, commercial banks should not have visibility into each other's balances and the central bank should not have visibility unless one of the commercial banks in the exchange falls under their jurisdiction. Likewise, non-participating banks should not have visibility into a specific transaction. In general, the system should not require a material compromise in privacy.

Chapter 6

Technology Assessment

Methodology

A critical activity during Phase 1 was technology assessment, i.e. evaluation of potential DLT platforms for the pilot and deciding on the platform (aka the DLT protocol) to be used for the implementation phase.

The assessment was performed in two stages. First, a qualification criteria was used to come up with a set of candidate DLT technologies/protocols. Second, an assessment criteria were used to evaluate the shortlisted technologies. Rather than applying assessment criteria on the technologies directly, it was done in context of the cross-border interbank payment problem. The objective was to assess the ability of different technologies to satisfy (current and future) Project Aber requirements. The evaluation approach is depicted in Figure 8.

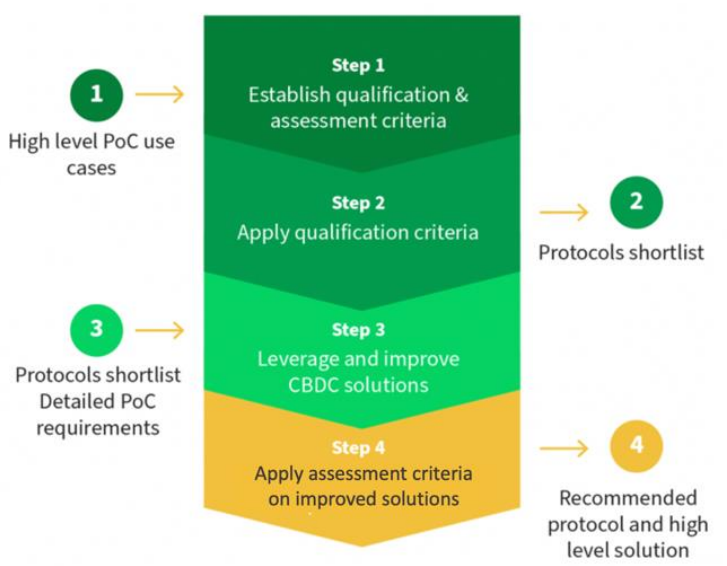

Figure 8: The technology assessment process for Aber

The qualification criteria was simply a minimum capability set derived from the high level expectations and Aber use cases. The key capabilities were:

- Support for smart contracts;
- Built-in support for permissioning and privacy in the DLT platform;
- Support for consensus with finality property;
- Well known in financial industry;
- Use in previous CBDC projects (preferred).

- Note that public blockchain protocols such as Ripple and Stellar, which are often positioned for cross-border remittance use cases, were ruled out because of the obvious need for permissioning and privacy for an interbank payment use case (which these protocols didn't support). Application of the qualification criteria resulted in shortlisting of the following three DLT platforms or protocols:
- R3 Corda
- Hyperledger Fabric (HLF)
- JPMC Quorum

Assessment Criteria

A comprehensive assessment criteria was used to evaluate the protocols. Rather than being generic, the assessment was very specific to Aber requirements. Instead of applying the criteria directly to technologies, a high level solution overview for each of the target platforms was drawn up and then evaluated against the criteria.

Solutions used in previous central bank experiments, served as starting points. The solutions were then enhanced to address observations from previous PoCs and to incorporate requirements that were specific to Aber.

The analysis evaluated effectiveness over several areas that could be considered mutually exclusive. These were termed as dimensions of the evaluation. Further, to be comprehensive and detailed, several facets within a dimension, were considered. Each facet was characterized by multiple levels representing least to most desirable attributes of that feature. The facets were weighted according to their relative importance and levels were assigned scores between 10 and 100.

The following dimensions were considered:

1. **Decentralization**: The primary benefit of using blockchain technology is to avoid a single point of trust or failure. This was considered especially important for the project due to the multiple jurisdictions involved in a cross-border payment system. Multiple levels corresponding to degree of decentralization (e.g. whether central banks need to authorize each and every transaction) and safety (e.g. whether decentralized was achieved at the cost of correctness or payment finality) were determined and assessed.

Another facet considered was whether the solution needs a trusted setup phase.

2. **Privacy**: Achieving decentralization with safety is hard, but privacy requirements make this problem even more challenging. Different privacy levels were defined depending on what is protected from unauthorized access e.g. transaction data (amount), endpoints (sender, receiver), liquidity positions, time of payment etc. The protection mechanism and support for auditing despite privacy restrictions were also considered.
3. **Scalability:** In this dimension, several non-functional characteristics such as latency of payments and variation of latency and throughput with increasing number of nodes (banks), were considered.
4. **Solution Complexity**: This dimension captures the complexity and the effort required in building and supporting a wholesale CBDC solution based on the evaluated technology. This considered several factors such as overall complexity, level of reuse of existing CBDC designs, and complexity of managing the solution.
5. **Security**: The means by which and extent to which the underlying technology supports permissioning was considered. There was a preference for PKI-based permissioning. Support for pluggable consensus was another factor considered, for instance, from the point of view of supporting byzantine fault tolerance (i.e. resilience in face of arbitrary and malicious behaviour). Built-in support for HSM standards was also considered a favourable capability.
6. **Production Readiness**: In this dimension, the readiness of the solution for eventual deployment into a production environment and its use in running real-world production transactions was considered. Like security, this dimension was largely dependent on the underlying technology. Various factors such as use of production ready components, operability, supportability and availability of production deployments based on the technology, were considered.
7. **Long term viability**: This was again a technology dependent criteria. Factors, or facets considered were: availability of skills, diversity of vendors involved in the project, extensibility, open standards and open governance.

As a sample, a graphical representation of the privacy dimension from the technology assessment report, appears in Figure 9.

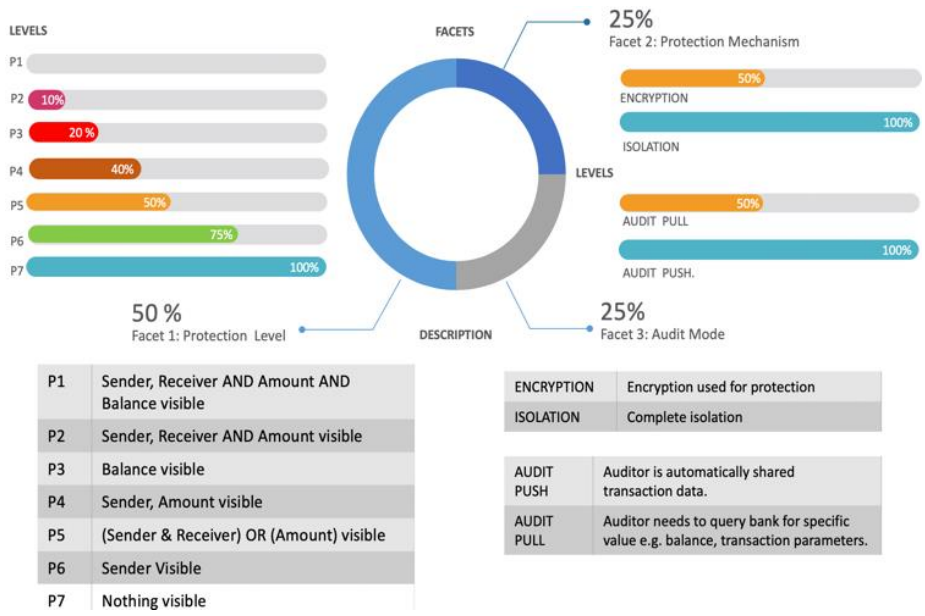

P1	Sender, Receiver AND Amount AND Balance visible
P2	Sender, Receiver AND Amount visible
P3	Balance visible
P4	Sender, Amount visible
P5	(Sender & Receiver) OR (Amount) visible
P6	Sender Visible
P7	Nothing visible

ENCRYPTION	Encryption used for protection
ISOLATION	Complete isolation

AUDIT PUSH	Auditor is automatically shared transaction data.
AUDIT PULL	Auditor needs to query bank for specific value e.g. balance, transaction parameters.

Figure 9: Rating criteria, for "Privacy and Visibility" dimension, showing relative weights and levels of the various factors considered.

Assessment Result

Some of the conclusions from our assessment were as follows:

- HLF and Corda based solutions, by virtue of scoring higher in the 'production readiness' dimension, were considered more enterprise ready than Quorum based ones;
- The approaches that use zero-knowledge proofs for validation do not scale very well with the number of network participants due to being computationally expensive;
- Solutions that use built-in privacy features of a platform should be preferred to reduce solution complexity and reduce risk in trying to create one's own crypto-schemes or similar.
- The HLF-based solution was the only one that concurrently satisfied the privacy, decentralization and safety requirements of the pilot.

The final assessment of the solutions across the dimensions appears in Figure 10.

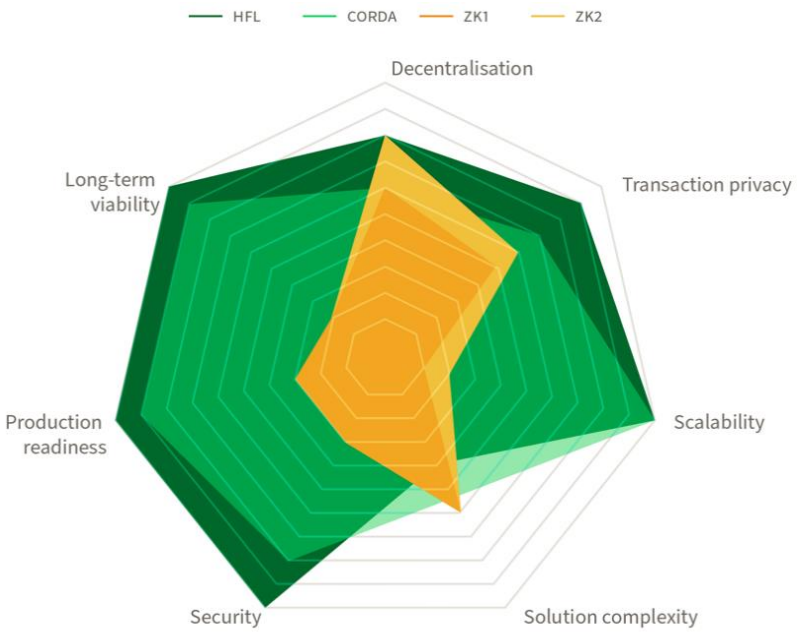

Figure 10: Comparative analysis of different DLT platforms such as Hyperledger Fabric, Corda and Quorum for Aber solution.

Four candidate solutions were considered: one each based on Hyperledger Fabric (HLF) and Corda, while two were based on Quorum (ZK1, ZK2). ZK1 was similar to the Ubin Phase 2 solution using Quorum [MAS, Nov 2017], while ZK2 was based on and used techniques similar to Project Khokha [SARB, 2018]. Based on this assessment and observations above, Hyperledger Fabric was selected as the technology for the implementation phase.

It is important to note that the purpose of Project Aber was not to prove a specific technology but rather it was to test the feasibility of achieving the business objectives or outcomes using a class of technology (DLT). The selection of HLF was a point in time decision. Given the rate and pace of change in the DLT space, any successive phase or project should consider the state of the art in blockchain protocols at that time and make the most appropriate selection.

Hyperledger Fabric Overview

Hyperledger Fabric (HLF) [HLF, 2019] is a permissioned blockchain technology, with a pluggable architecture that allows “plug and play” of important components such as consensus and membership service.

It allows two types of peers: endorsers and committers. Execution of smart contract or chain code is limited to endorsing nodes while committers only maintain the ledger.

In addition, there is a fault tolerant ordering service that delivers sequence of transactions grouped into blocks to the peers. This ensures that all non-faulty peers agree on the transactions and their order.

Privacy Features

Membership service provides an abstraction which can be used to support multiple permissioning mechanisms. The most commonly used is a PKI based membership service provider. However, privacy preserving techniques such as identity mixer can also be plugged in. This allows clients to sign transactions in such a way that their network membership can be verified, without revealing their real identity to the rest of the network [HLF, 2019]. Fabric provides a certificate authority out of the box, and supports cryptographic standard PKCS#11 which simplifies integration with HSMs for key protection.

Channels represent the primary mechanism for providing privacy in a Fabric based network. Ledgers exist in the scope of a channel. A ledger can be shared across the entire network or privatized to include only a specific set of participants. An organization can participate in multiple ledgers.

Another privacy feature is private data collections. This feature provides privacy within a channel, such that private transaction data and private state is visible only to a subset of organizations in a channel. However, the existence of the transaction and hash of private data and state can be seen by all members.

Collections are used when transactions (and ledger) must be shared among a set of organizations, but when only a subset of those organizations should have access to some (or all) of the data within a transaction. Additionally, since private data is disseminated peer-to-peer rather than via blocks, private data collections are also useful when transaction data must be kept confidential from ordering service nodes.

Transaction Flow

Hyperledger Fabric uses a distinctive transaction flow that goes through the steps of endorsement, ordering and validation. This design significantly improves scalability compared to other models based on state-machine replication. The key is to decouple two computationally heavy tasks — chain code execution and transaction ordering — so that they can be executed on separate nodes.

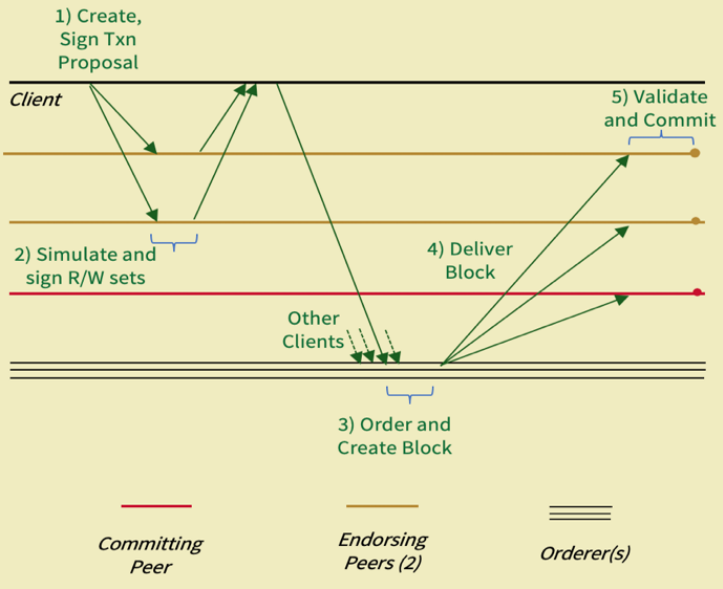

Figure 11: A representation of the Endorse-Order-Validate transaction flow of Hyperledger Fabric

The chain code is installed on peers and instantiated on a channel. The endorsement policy determines the peers that must endorse the transaction, for it to be considered valid. The transaction flow of Hyperledger Fabric is shown in Figure 11. The transaction goes through several steps before it gets committed to the ledger. These steps are outlined below:

Endorse: The client digitally signs the transaction proposal and sends it to the endorsing peers (as per the endorsement policy), which simulate execution of the smart contract (chain code) associated with the transaction, digitally sign the read/write sets (key value pairs that were read or would be written by the execution) and return to the client.

Order: The client consolidates the responses and sends the transaction to the ordering service which is responsible for ordering transactions from multiple clients, minting blocks and broadcasting it to the peers.

Validate and Commit: All peers that maintain ledger (i.e. endorsing as well as committing peers) validate transaction by verifying its endorsement policy and by checking whether the read sets are still valid. The transaction once validated as described above is committed to the ledger.

Chapter 7

Solution overview

Key design considerations

Hyperledger Fabric (HLF) is a general purpose blockchain platform. It does not assume a specific domain or industry or asset type. Thus, to implement an interbank payment using HLF, a payment protocol was designed and implemented specifically for Aber requirements. This is referred to as the Aber Protocol in this paper.

Since Aber is essentially an interbank payment system, its requirements were naturally derived from the features of traditional RTGS systems, but extended to cover cross-border payments. What was obviously different from traditional implementation though was the need for decentralization. This single requirement had a direct implication on several non-functional aspects such as privacy, security and scalability. In this section, these requirements are reviewed, existing solutions and drawbacks are considered, in order to motivate the need for a new privacy architecture.

Decentralization

The system is expected to allow payment and settlement between commercial banks to continue without the central bank(s) nodes being available. This suggest a T-shape transaction model as shown in Figure 12. This differs from the previous attempts at CBDC which have largely targeted a Y-Shape transaction model. While the T-shape model may not satisfy all of the basis for RTGS system, the key property of irrevocable transactions can be achieved as long as transaction validity requires a quorum of participants that does not necessarily include the central bank(s).

Privacy and Visibility

The privacy and visibility requirements can be summarized by the graphic in Figure 13

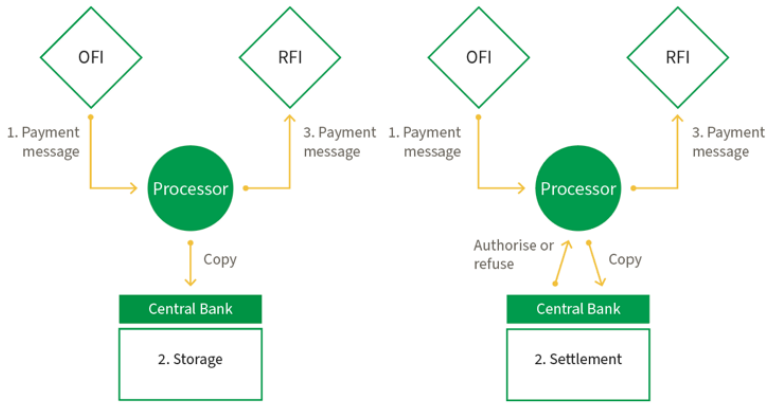

Figure 12: Role of central bank in a T-shape model (left) compared to a more traditional Y-shape model.

In the picture, green and red nodes represent KSA and UAE commercial banks respectively while arrows represent transactions between banks. The large circles represent the sphere of visibility for SAMA and CBUAE, while small circles represent the same for a commercial bank in each jurisdiction.

Commercial banks have complete knowledge of their own node and the transactions they participate in, but do not know about other nodes (i.e. their overall DLT balances) or transactions between other nodes.

Central banks have visibility to everything regarding the nodes and transactions in their jurisdiction. In addition, they also have knowledge of all cross-border transactions. However, central banks are not supposed to know about overall DLT balances and domestic transactions of participants in the other jurisdiction. While in general, central banks should not have visibility into the other jurisdiction, there is one exception to this rule. Both central banks must have full visibility of money supply, including the digital currency that was issued in the other country. The reason for this is that any digital currency issued in one jurisdiction could, theoretically, end up transferred to the other jurisdiction and subsequently presented for redemption.

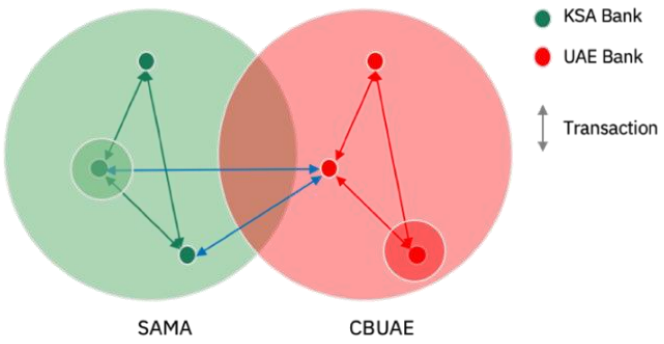

Figure 13: Sphere of visibility in the Aber Solution

Safety

The safety property in a distributed system is a guarantee that something bad will never happen. In the context of decentralized payment systems, bad usually refers to the possibility of double spend. Simply speaking, this means that a participant should not be able to use the same digital currency in two valid payments. The primary goal of DLT based payment systems is thus to prevent double spend in the absence of a central trusted authority.

When discussing decentralization, a need of having a quorum of participants (not necessarily including central banks) that is sufficient to validate a payment, was discussed. This quorum should be a large group in order to achieve safety. However, this directly conflicts with the privacy requirement which restricts commercial banks to transactions directly involving them. Providing safety of payments, while preserving privacy of transactions and balances was one of the major considerations in the design of Aber protocol.

Scalability

For an alternative cross-border RTGS system, it is not only important to execute transactions with privacy and safety, but also provide high transaction throughput.

Blockchains are generally associated with poor transaction throughput and scalability. This is primarily because consensus with the safety properties described above, is hard to achieve in a network with a large number of nodes. As number of nodes in the network increase, the demand of transactions in the network goes up, but the capacity to service the transaction reduces. Thus, to start with, the consensus problem – something that all blockchain networks need to contend with – may seem conflicting with the goal of building scalable payment systems.

Public blockchain networks, such as Bitcoin and Etherium, where the throughput problem is particularly severe are looking at parallelism through sharding and

layer two networks, where computation is moved off-chain, using innovations such as lightning network [Poon, 2016].
A particularly challenging task in Aber protocol design was managing the conflict between its privacy, parallelism and safety features such that neither of the important non-functional aspects were compromised.

Key solution concepts

Hybrid transaction model

There are two types of transaction accounting models that are typically used in DLT technologies. There are Bitcoin-style models that based on tracking unspent transaction outputs and termed as the UTXO model. Etherium, Ripple, Stellar use a model where account balances are explicitly recorded on-ledger. In permissioned DLT technologies, Corda subscribes to the UTXO model; whereas Hyperledger Fabric does not commit to a particular model, but most implementations use an account-based model.
We reviewed the solutions and observations of prior CBDC experiments from central banks that use these accounting models. Due to privacy reasons, in either case, the main ledger needs to be partitioned into multiple sub-ledgers, one for each pair of commercial banks in the network. Since transactions within a sub-ledger are private to its members, to guarantee safety, a solution must address the potential double spend problem arising out of digital currency moving across ledgers.
This problem of providing safety with privacy in a decentralized setting has not been completely solved in previous CBDC approaches. In past solutions, we have seen at least one of the three properties: decentralization, safety or privacy being compromised.
The UTXO approach, in the absence of a global ledger, requires the receiver to execute a validation check requiring access to the chains of transactions leading up to any of the inputs of the current transaction. This issue has been termed as walking the chain in Corda documentation [R3, 2019] and has obvious privacy implications. Confidential identities is a partial solution to the issue, which at least protects the transacting identities from third party disclosure in most cases. However, transaction amounts are still exposed to third parties.
The account-based approaches in previous CBDC projects have typically provided privacy at the cost of either decentralization or safety. This results in either a centralized solution or one in which double spend can be detected but not prevented.

Aber uses a transaction model which is a hybrid between an account and token (UTXO) model. There are two types of digital currency transfer transactions: within a privacy group (sub-ledger) and across privacy groups. A UTXO model is used for accounting movement of digital currency across sub-ledgers. Once the digital currency is moved into a ledger, it can be exchanged peer-to-peer using an account model. This allows us to decouple the privacy and safety concerns and handle them in two different types of ledgers, using different transaction models.

Channel structure

Aber uses a unique approach for solving the privacy and safety issues. The channel membership is shown in Figure 14 for all the channels in which a participant (bank B1) from KSA is involved in. There are three types of channels used by the solution:

Bilateral Channels: These are the peer-to-peer channels between each pair of commercial banks. It uses an account-based transaction model. The central bank of each peer also participates in the channel. Thus bilateral channel membership is either 3 or 4 depending on whether it is domestic or cross-border. However, endorsement policy of chain code (the term for smart contract used in HLF documentation) installed in this channel requires only the two commercial banks to endorse a transaction. This allows a payment to be confirmed even when one or both central banks are unavailable.

Private Channel: A private ledger between a commercial bank and its central bank. This is used for making private requests such as issue and redeem requests. Chain code installed in this type of channel requires endorsement from both participants.

Primary Channel: All the banks (commercial as well as central) participate in this channel. This channel uses a UTXO based transaction model. Endorsement policy for chain code installed on this channel requires at least 5 (out of the total 8) organizations to endorse transactions with at least 2 of them from each jurisdiction.

Transaction privacy and visibility requirements discussed under **Non Functional Requirements** are addressed by the bilateral channels. Privacy of issue and redeem operations is managed by private channels. The primary channel is responsible for association of digital currency with bilateral channels. It does so while ensuring the following properties:

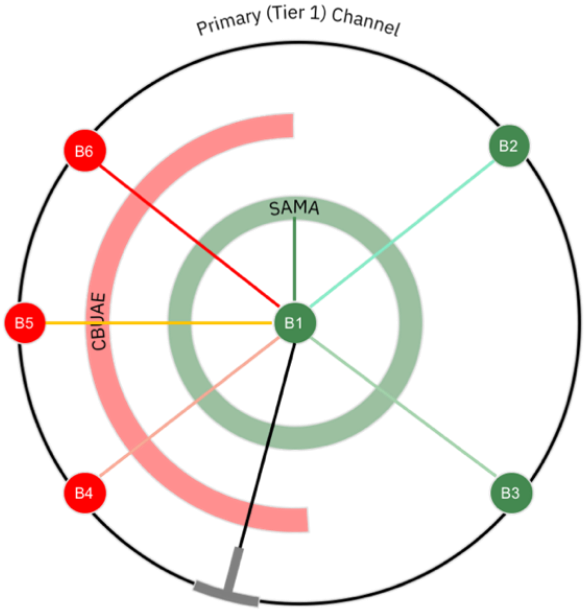

Figure 14 Channel structure for Aber Solution

1. At a given point in time, a token of digital currency can be associated with at most one bilateral channel.
2. Movement of a token of digital currency (DC) from one bilateral channel to another (or its redemption), requires agreement from both commercial banks in the channel.
3. The association between digital currency and owning bilateral channel is kept secret from the network, i.e. it is only known to the participants of the bilateral channel.

The table below is a summary of different types of ledgers used in the solution.

Ledger	Purpose, # instances and membership
Primary	Issuance and anonymous tracking of slices of DC
	1, CBs and all commercial banks
Bilateral	Peer to peer payments
	15, commercial bank pair and their CB(s)
Private	Ledger for issue and redeem requests
	6, commercial bank and its CB

Pseudonyms and Slices

Digital Currency is issued by central banks through a special "issue" transaction on the primary ledger. When commercial banks request funds to be issued by central banks, rather than tagging the currency against their enrolment public key, which is associated with their real identity, it gets tagged to a pair of pseudo identities, identified by one time use public keys or addresses called pseudonyms.

More specifically, each issued fund is tied to two public keys in this manner. Normally these public keys belong to the two commercial banks in a bilateral channel. This effectively locks the funds to a bilateral channel and requires agreement (consent) from both members for it to be moved. The transaction for moving the currency between two channels or from a channel to central bank is termed as inter channel or slice (introduced below) transfer.

Pseudonyms conceal the identity of the channel and hence the bank to which digital currency is issued. However, as per Aber requirements, the amount of digital currency also needs to be protected. This is important for both issue and slice transfer transactions on the primary channel.

To protect the amounts, digital currency is issued in fixed units called slices. E.g. a slice could be configured as 1000 digital currency. All transactions in the primary ledger (i.e. slice issue and transfer) use a single slice as the denomination. Combined with the fact that all transactions use pseudonyms rather than channel or bank identifiers, means that the primary ledger nodes only see transactions of fixed amounts going to unknown addresses, thus revealing no information on issuance, redemption or redistribution of funds to individual banks.

Apart from providing secrecy, a slice is also the unit of accounting in the distributed ledger. The slice size should be carefully controlled in an implementation. If the slice size is too small, the overhead of tracking the digital currency across bilateral ledgers will be high. Alternatively, if the slice size is too large, it will reduce granularity of issuance and redeem operations and make it harder for commercial banks to precisely match availability of digital currency with its requirement.

It should be noted, that the notion of slice only limits the granularity of issuance and redemption requests and slice transfer transactions. It does not in any way impose lower or upper limits of a payment transaction which takes place between commercial banks in bilateral channels. In an actual implementation, the size of slice would depend on the average transaction size in the network. Typically, it should be tens of times the average transaction size.

Consent

When a slice is locked to a bilateral channel, it cannot be moved without the agreement of both the parties. While this rules out double spending funds in another channel without counterparty knowledge, it also means that funds could be stuck if counterparty goes offline or becomes unresponsive. Moreover, at any point in its lifecycle, the two parties in the bilateral channel in which it is currently associated, must agree on their shares of the slice.

To avoid locking of funds or ambiguity in ownership of slices, the notion of consent is defined. Explicit agreement called consent is exchanged between the two parties to agree on the share. A consent for a share **X** of a slice from party **A** to **B**, means that **A** agrees to **B** owning share **X** of the slice. When the share of party **B** in the slice reduces from **X** to **Y** where **X > Y**, then **B** must surrender the consent for X shared by **A** earlier and obtain another consent from **A** for share **Y**. The surrendering of consent, issued from bank **A** to **B**, by bank **B** is termed as consent revocation.

Overview of Aber Protocol

Aber is designed based on the concepts described above. The three digital currency lifecycle operations: Issue, Transfer and Redeem are orchestrated through a workflow involving the three types of channels described above. In addition, there are two other important processes and algorithms implemented as part of Aber which simplify and improve liquidity management. These are called Automated Fund Management and Gridlock Resolution. Collectively, these five workflows represent a majority of Protocol Aber.

Issue Workflow

Issue Workflow is the process of a commercial bank requesting its central bank to issue a specific amount of digital currency and the central bank issuing it. The high level workflow to implement this business process using the design is described here.

A commercial bank (e.g. Bank **A**) decides to make an issuance request. Digital currency has to be requested (and issued) in multiple of slices in this solution. Unless the automated fund management feature is enabled, Bank **A** must also decide on a bilateral channel(s), where this currency is to be issued. In a simple case, let's say Bank **A** wants a single slice to be issued in its bilateral channel with Bank **B**. The workflow doesn't depend on whether **B** is from the same jurisdiction or not. The process can be described as follows, as also illustrated in Figure 15:

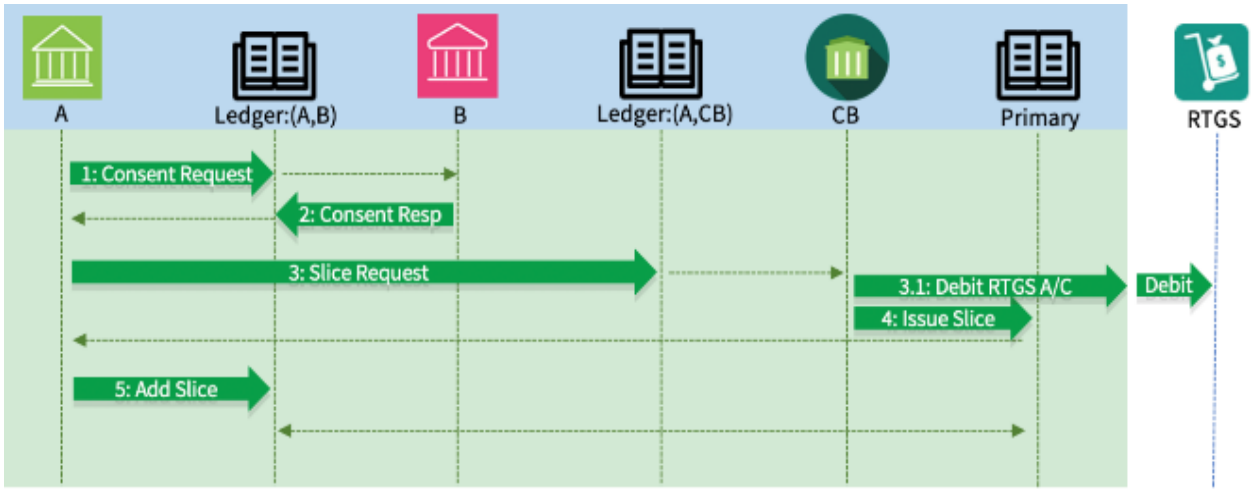

Figure 15 Aber Protocol: Issue Workflow

- (Steps 1-2) Bank **A** and **B** generate a pseudonym (one time address) each against which this currency will be issued. Bank B also signs and shares a consent agreeing to hundred percent ownership by A's address. This workflow takes place between DLT adapters of A and B, through the (**A,B**) bilateral ledger.
- (Step 3) Bank **A** makes a request to central bank (**CB**) to issue the slice against the two addresses. This request is made in private using the private (**A,CB**) ledger.
- (Step 4) Central bank issues the slice in the primary channel locked to the two one-time use addresses.
- (Step 5) Bank **A** requests addition of the slice to its bilateral channel. The chain code checks existence of the slice and its owning addresses in the primary ledger before adding the slice.

For real money transactions, the same process was followed, except before issuing the slice in step 4, the commercial bank pledged money in an Aber account at the central bank, dedicated for this purpose.
As per the principles described earlier, there was no direct integration with core banking systems; and a manual process was in place. Reports exported from the DLT solution were used to update core banking.

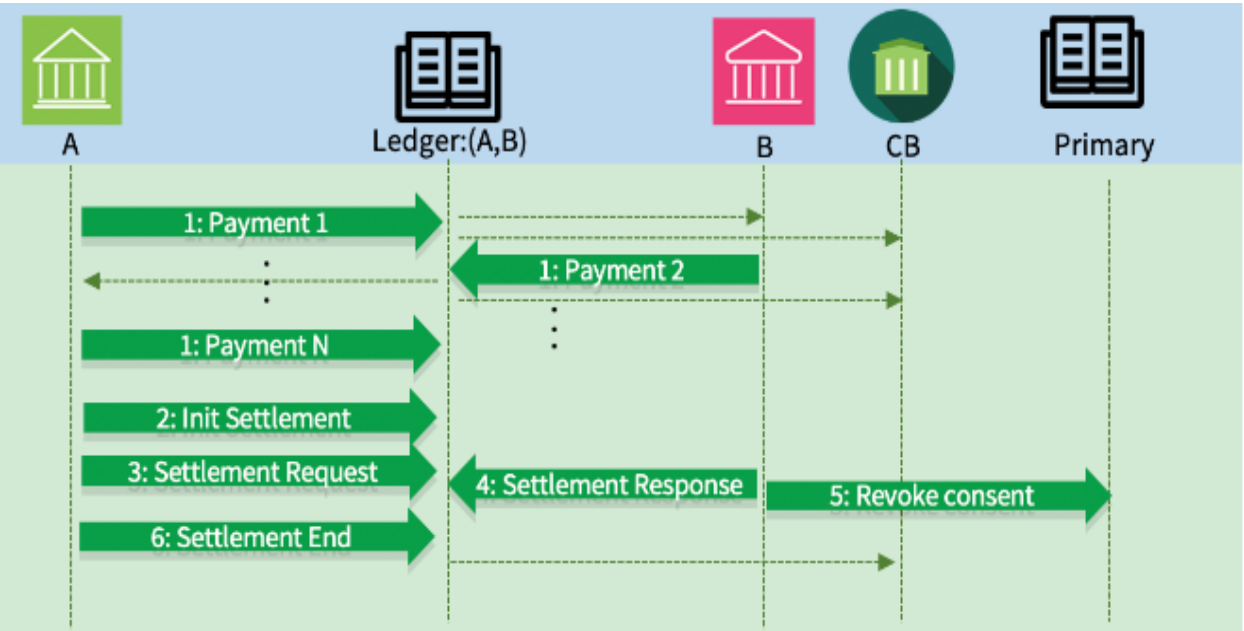

Figure 16: Aber Protocol: Transfer Workflow

Transfer Workflow

Transfer workflow is the process of a commercial bank requesting to make a payment to a counterparty with which it shares a bilateral channel. Note that this is different from the slice transfer operation that moves slice sized

chunks of digital currency across channels. In the description below, the two banks making bilateral payments to each other are **A** and **B**. The process works as follows, and illustrated in Figure 16:

- (Step 1) Banks **A** and **B** make payments to each other by invoking the chain code on the (**A, B**) bilateral ledger. In case of insufficient balance, payments get processed in an on-chain queue. Bilateral netting was implemented, so before adding a payment to a queue, netting opportunities with payments queued on the other side are explored. Representing the new ownership was not exchanged.
- (Step 2) At some point, the settlement cycle gets triggered. This could be after every transaction, based on time, no. of payments or movement in positions/balances -- all measured since the last settlement.
- (Steps 3-4) **A** and **B** exchange consents with each other based on the new shares of slices. Note that at any point of time, at most one slice in the channel is partially owned.
- (Steps 5-6) Previously issued consents for those slices need to be invalidated. Consent issued by **A**, needs to be invalidated by **B**. This operation is performed on the primary ledger. If the netting is in favor or **A** (i.e. **A**'s balance is supposed to increase because of the netting), then only **B** needs to revoke successful, either party can request ending the settlement cycle.

The frequency of the settlement cycle is a configuration parameter in the solution. Having a large settlement cycle means that payments take longer to confirm but is more efficient since consent/revocation need to be exchanged/executed for every transaction.

Redeem Workflow

The redeem workflow is used by a commercial bank to request digital currency to be destroyed and funds to be returned to their core-banking account outside the DLT. For transactions involving real money, the Aber accounts maintained by the central banks were used for this purpose. As mentioned before, there

was no direct integration with core banking and a manual process was in place, where reports extracted from the DLT network were shared with the core banking team. The redeem process in Aber protocol can be described as follows (refer to Figure 17):

Bank **A** decides to redeem digital currency. Unless using the automated fund management feature, it must decide the bilateral channel and the amount. The amount must be multiple of a slice and there must be at least those number of slices available and fully owned by **A** in the channel. The rest of this process description assumes that **A** wants to redeem a single slice in channel (**A, B**) for which it has a 100% ownership consent from **B**.

- (Steps 1-3) Bank **A** first removes the slice from the bilateral

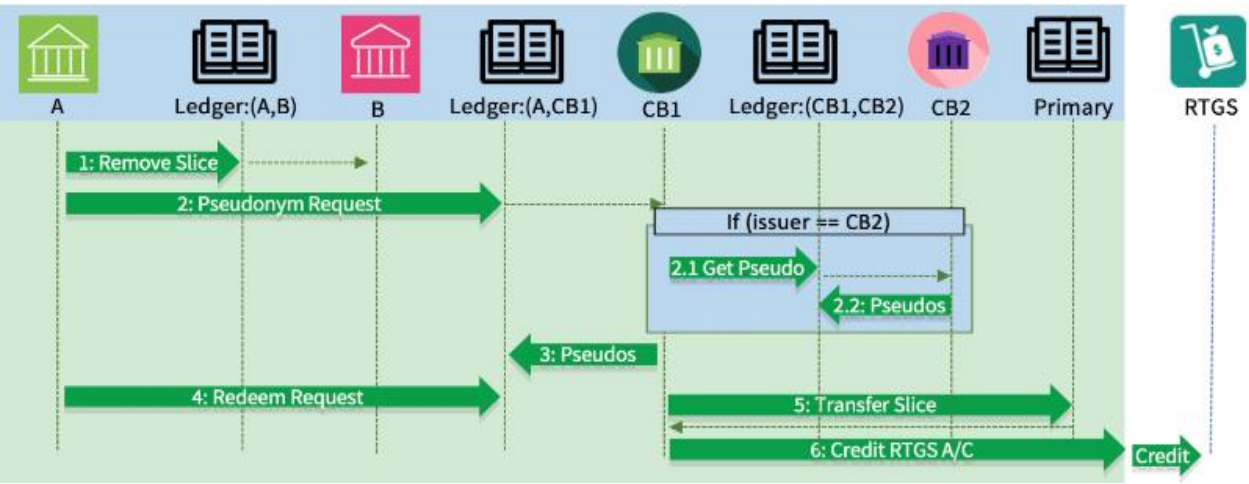

Figure 17:Aber Protocol: Redeem Workflow

channel (**A, B**). It requests two pseudonym addresses from central bank, as the control of the slice is to be returned its central bank (**CB1**). This request is made on the (**A, CB1**) channel. If the slice being redeemed was issued by the other central bank (**CB2**), then **CB1** gets the addresses from **CB2**. This exchange happens over the (**CB1,CB2**) channel.

- (Step 4) Bank **A** prepares a slice transfer request and shares with **CB1** on the private ledger (**A, CB**1). The request is for moving the slice to the pair of addresses returned by **CB1**.
- (Step 5) **CB1** executes the slice transfer operation in the primary ledger.
- (Step 6) On successful notification of this, **CB1** makes a request for returning the funds by crediting the commercial bank account in a core-banking system as mentioned above.

- (Step 7, optional) In the event that the digital currency was issued by **CB2**, then on receiving successful notification of slice transfer, **CB2** initiates a payment to **CB1** in the (**CB1**, **CB2**) channel.

Evaluation

Advanced Liquidity Management

Besides the basic digital currency lifecycle flows described above, some other advanced flows were implemented as part of Aber protocol. Two flows under the liquidity management category are described here:

Gridlock Resolution

The multilateral netting algorithm was based on the Ubin Phase 2 approach for the Fabric prototype where a distributed version of EAF2 algorithm was implemented [MAS, Nov 2017]. The enhancements introduced were:

- In the settlement phase, explicit creation and destruction of currency in bilateral ledgers was avoided. This allowed all issued currency (slices) to be (anonymously) tracked by all participants in the primary ledger. Instead, a new assignment to the slice shares that satisfied the settlement requirement was computed.

- Support for cross-border multilateral netting while preserving jurisdiction based visibility constraints was introduced, i.e. a foreign central bank should not be able to view domestic transactions even during gridlock settlement. Several options including privacy preserving encryption techniques were considered. However, the decision was to go for a computationally simpler approach where multilateral netting was performed in two stages: domestic and cross-border, with the output of the first stage being used in the second stage.

Automated Fund Management

Earlier in this section, the slice transfer operation that can be used to move a slice between channels, was discussed. This could be used to provide end user complete control over managing liquidity in each of a bank's channels. However, this would have required significant operational effort and inherent lag caused by a manual activity, thus making it less efficient. This problem was solved in Aber, through the Automated Fund Management feature. It works by predicting

liquidity deficit for the next cycle, based on demand in previous N cycles, and status of incoming and outgoing queues. This prediction is done for every channel, the bank is involved in. The deficit is used in two ways:

- At the time of issuance, it is used to automatically allocate requested slices to candidate ledgers belonging to the bank. Thus channels with higher liquidity deficit are allocated proportionally more slices
- It is also used to automatically rebalance funds from channels with excess liquidity to those in deficit. The channels are categorized with one of the liquidity levels: sufficient, donor, recipient. The automatic rebalancing feature then invokes slice transfer operations to carry out the actual donor to recipient channels.

Chapter 8

Aber Evaluation

This section presents an assessment of the Aber design and implementation against some of the key project requirements and financial market principles. Both functional and non-functional aspects have been considered. The technical evaluation also provides an insight into some of the design decisions.

Decentralization

Project Aber sought to push the limits when it comes to decentralizing payments in a permissioned blockchain setting. The decentralization directive under business requirements states:

Only counterparties to the payment transaction need to be online for the payment to be settled. The rationale behind this is to enable the system to offer a higher level of architectural resilience than traditional centralized systems that depend on the availability of the central bank and to avoid single point of failure.

In terms of the Aber Protocol, this translates into the following two conditions:

I. Transfer workflow should be possible to perform without availability of central bank(s).
II. Slice transfer operation must be possible to perform without involvement of central bank.

The former condition ensures that commercial banks can agree on payments bilaterally without involving a trusted third party. The latter condition ensures that a commercial bank can move its funds freely between the channels in which it participates without a dependency on the central banks or another party.

The transfer workflow involves transactions on the bilateral channel and primary channel (Figure 16). Bilateral channel operations do not require central banks endorsement, as per the endorsement policy specified under channel structure. Consent revocation is the only action that requires the primary channel, where again the endorsement policy does not require central banks. This is because, as discussed under channel structure, only 5 out of 8 participants are needed to endorse transactions on the primary channel.

The slice transfer operation, even though it changes the association of a slice from one bilateral channel to another, is also executed on the primary channel. Thus, due to the same argument as above, it does not necessarily require a central bank endorsement. Thus condition (ii) is also satisfied.

One other aspect related to decentralization is the distribution of the ordering service. For the pilot, a single ordering service was used for all the channels and its deployment was distributed between the two central banks. In a real production implementation, a more distributed deployment of the service,

including at least some of the commercial banks, is recommended. The privacy implications of doing this are considered later in this evaluation.

Privacy and Visibility

Transaction

We first consider privacy and visibility of the payment transaction. Transaction privacy requirement dictates that if a payment takes place between two commercial banks, then a third party commercial bank should have no knowledge of the following:

- The transaction amounts;
- Sender or the recipient bank identities;

Conversely, the visibility requirement states that the above information should be auditable (only) by the central banks who are the governing authorities for the transacting commercial banks. This means that domestic payments must be auditable by the central bank in that jurisdiction, while cross-border payments must be auditable by both central banks.

The bilateral channel design in the Aber solution is closely aligned with this privacy and visibility requirement. A bilateral channel between two commercial banks has the two banks and their central banks as members. This ensures isolation from third party commercial banks but at the same time satisfies the visibility requirements for central banks.

To be completely compliant with transaction privacy requirements, one also need to consider privacy from the ordering service nodes. There are two compliant deployments of the ordering service:

- A separate ordering service for each channel in the solution, such that only the nodes participating in the channel host the ordering nodes;
- A single multisite ordering service for all channels (possibly including both central banks and some or all commercial banks). However, each bilateral channel having a single private collection that is used for all its transactions;

In the latter case, amounts can be part of private transaction data and the nodes comprising the ordering service only sees hashes of the private data and private state.

Account Balances

Account balances should remain private to the commercial banks and their central banks. Account balance of a commercial bank in the DLT is the sum of its balances across all its bilateral channels. Since the central bank of the commercial bank's jurisdiction is the only entity having visibility to all of a commercial banks' channels, the overall position of commercial bank is known only to itself and its central bank.

Note that even though commercial banks are aware of each other's balance in their common bilateral channel, this does not amount to a privacy leak. In a large network, this balance cannot be used to predict overall balance of the counterparty (which is an aggregate of all that counterparty's bilateral channels with all the other banks in the network).

Another partial leakage of account balances could be possible through the slice transfer operation if the identity of the bank transferring the slice gets disclosed to other commercial banks. Slice transfer occurs in two situations. During automated fund management, it is used to move a slice between a bank's channels. It is also used to return a slice to the central bank during the redeem flow. As evident from Figure 17, in the latter case, the actual slice movement operation is executed by the central bank, thus the identity of commercial bank is not revealed. In the former case though, if the commercial bank were to perform this operation itself, then a partial revelation of its digital currency ownership will take place.

In the implementation, this leakage is avoided using a very simple solution. The bank interested in executing the slice transfer operation prepares the transaction and shares it with the other commercial bank in the source bilateral channel, which then executes it on the primary channel. In a network where banks have relationship with every other bank, this arrangement reveals zero information about the requesting commercial bank. In the worst case, the frequency of these transactions can be used to predict the level of imbalance that exists in a commercial bank's channels, but is not helpful in predicting the overall liquidity.

Safety

Safety properties in distributed systems imply that certain undesirable system states can never be reached. In the context of a decentralized payment system, the bad state is usually allowing some kind of double spend to occur due to a lack of trusted third party that reviews all transactions. Another safety consideration is that collusion between a small number of participants should

not be able to subvert the entire network. Defences against several types of double spending attacks are discussed below.

Unauthorized Slice Movement Scenarios

Removing a slice without counterparty knowledge

An honest participant should first remove a slice from a bilateral channel, before transferring it into another channel or redeeming it. However, a dishonest participant (let's say Bank A) can skip this step, so the other participant (bank B) does not know. It is also possible that the resulting notification of slice transfer operation was delayed or missed altogether by B. In this case, A could try to use this slice in a payment to B, even though it has been already moved and could be legitimately consumed, elsewhere.

However, such an attack is foiled because of a property guaranteed by the Aber protocol implementation: a consent can be consumed or revoked at most once. This also means that a consumed consent cannot be revoked and a revoked consent cannot be consumed. For a slice to be redeemed or transferred to another channel, a slice transfer transaction must have been executed in which a 100% consent for the slice from B would have been consumed. When the same slice is now used within the bilateral channel with B, e.g. for making a payment to B, the settlement will require revocation of the same consent, on the primary channel. This will fail because the consent had already been consumed. Thus B will not accept such a payment that has been made using a tainted slice.

Moving a partially owned slice out of the channel

Let's now consider the situation in which **A** and **B** have a partially owned slice in their common bilateral channel and **A** attempts to move it out of the channel. **A** does not have a valid 100% consent from **B**, but may have a consent received previously. However, this consent must currently stand revoked since **B** owns part of the slice now. Once again, the consent property guarantee described above becomes applicable. If **A** tries to use a revoked consent to move the slice out of the channel, the slice transfer transaction would fail on primary channel because a revoked consent cannot be consumed.

Moving the same slice to two different channels

An attempt to move a slice to two different channels with two different pseudonyms will also fail because only the first one of these transactions will succeed on the primary channel. The second will fail because the ownership (pseudonym pair corresponding to the slice) would have changed after the first movement. This will make the second transfer invalid as the consent was issued against one of the older pseudonyms.

PFMI Analysis

The Bank of International Settlements (BIS) Principles for Financial Market Infrastructures (PFMI) [CPSS, 2012] are international standards for payment systems, central securities depositories, securities settlement systems, central counterparties and trade repositories. The PFMI are designed to help ensure the safety, efficiency and resilience of these infrastructures supporting global financial markets — so their full, timely and consistent implementation is fundamental.

In this section, an evaluation of the solution against the PFMI principles is presented. The analysis focuses on those aspects that would likely be impacted by use of the use of DLT. However, at this stage, legal framework (e.g. any legal implication arising from tokenization, distributed ledger) and network governance are considered out of scope. Table 1 lists all the principles by category, highlighting those that have been considered and the reasons for some categories not being considered.

CATEGORY	PRINCIPLES	COMMENTS
General Organization	Legal basis, Governance, Framework for comprehensive risk management	Governance and legal not in scope
Credit and Liquidity Risk Management	**Credit Risk**, Margin, Collateral, **Liquidity Risk**	Considered
Settlement	**Finality**, **Money Settlements**, Physical Deliveries	Considered
Central Securities depositories and exchange-of-value settlement systems	Central security depositories, Exchange of value settlement system	Not applicable
Default Management	Participant default rules and procedures Segregation and Portability	Governance not in scope
General business and operational risk management	General business risk, custody and investment risks, **Operational Risk**	Considered
Access	Access and participation requirements, Tiered participation arrangements, FMI links	Not impacted by DLT
Efficiency	**Efficiency and effectiveness**. Communication procedures and standards	Considered
Transparency	Disclosure of rules, Key procedures and market data, Disclosure of market data by trade repositories	Not applicable

Table 1: PFMI principles by category. The principles considered in this evaluation appear in bold font.

Principle 4: Credit Risk

An FMI should effectively measure, monitor, and manage its credit exposures to participants and those arising from its payment, clearing, and settlement processes.

Since the digital currency is issued by central banks, there is no risk of default to participants that accept the currency. Note that as per business requirements, redemption of digital currency always happens against the local central bank, regardless of the issuer. Since central bank is already trusted by commercial banks for their deposits, it does not introduce a new risk. However, the actual liability of the currency is with the issuing central bank, so it does require some level of trust between the two central banks; but the transparency afforded by DLT into the money supply and the redeem process of cross-border issued digital currency makes it easier to achieve this trust.

Principle 7: Liquidity Risk

An FMI should maintain sufficient liquid resources in all relevant currencies to effect same-day and, where appropriate, intraday and multiday settlement of payment obligations with a high degree of confidence.

The only implication on liquidity with the use of distributed ledger is that LSMs are harder to implement in a distributed system. While LSMs have been implemented in previous CBDC projects, having multiple jurisdictions introduces additional challenges. The Liquidity Management section describes the challenges addressed by the solution.

Principle 8: Settlement Finality.

This is the property most likely to be impacted when moving to a decentralized RTGS implementation without a central authority for authorizing payments. This is thus the most important principle to be considered for a DLT based implementation.

The principle states:

An FMI should provide clear and certain final settlement, at a minimum by the end of the value date. Where necessary or preferable, an FMI should provide final settlement intraday or in real time... An FMI should clearly define the point after which unsettled payments, transfer instructions, or other obligations may not be revoked by a participant.

Payments between commercial banks take place in bilateral channels in the Aber solution. The point of finality for a payment from Bank A to Bank B in a bilateral channel can be precisely defined:

A payment from A to B can be considered final and irrevocable if and only if a settlement cycle (see Transfer Flow) including that transaction has been completed successfully between the two participants (i.e. the end settlement has been committed to the bilateral ledger).

Since at this point, the settlement cycle ID along with transaction IDs and transactions have been recorded in the immutable blockchain ledger, and both participants have a copy of the ledger, the transaction can be considered final. The network banks could be made to sign a contract acknowledging the possession of such transaction records by the counterparty as unconditional and irrevocable transfer of asset. The discussion on safety earlier in this section are relevant in this respect. In particular, the defences discussed against double spend ensure the following:

- At any given point in time, a slice could be associated with at most one bilateral channel. Thus it is not possible that the slice could have been simultaneously committed for any other payments, without the knowledge of both A and B.
- For each of the slices impacted by a settlement cycle, complete ownership history from its creation (slice issuance) to the point it came into A and B's possession and the sharing of slice between the two, is immutably recorded in the ledgers.

Principle 9: Money Settlements

An FMI should conduct its money settlements in central bank money where practical and available.

This is by definition true for a CBDC based payment system. In Project Aber, while there is a single digital currency, there are two issuers. Thus, commercial banks do accept payments that are not necessarily settled in their local central bank issued currency. However, from a commercial bank perspective, this distinction does not matter, since they are always guaranteed redemption of the currency against the local central bank.

Principle 17: Operational Risk

An FMI should identify the plausible sources of operational risk, both internal and external, and mitigate their impact through the use of appropriate systems, policies, procedures, and controls. Systems should be designed to ensure a high

degree of security and operational reliability and should have adequate, scalable capacity.

One of the distinct advantages of using DLT for interbank payments is the much higher architectural resilience when compared to other traditional implementations. Each ledger entry is replicated across multiple parties and physical locations; and the network has consensus about transactions and their relative order. Moreover, the design principle "Decentralize with Safety" allows payments to be made between commercial banks without any of the central banks having to be online. This eliminates a single point of failure in traditional interbank payment systems, thereby reducing the operational risk.

It should be noted that Safety and Privacy, two of the most important aspects for operational risk evaluation have been considered as separate topics earlier in this section. Thus, they are not discussed here.

Principle 21: Efficiency and Effectiveness

An FMI should be efficient and effective in meeting the requirements of its participants and the markets it serves.

The collective resources required in a distributed ledger based interbank payment solution are expected to be more than a traditional centralized solution. This is primarily due to multiple copies of the ledger that need to be synchronized using a communication intensive consensus process. However, Aber was designed to be a practical solution and uses an underlying DLT technology that scales well. The analysis presented in the next section demonstrates that Aber performance is indeed quite acceptable and that it can be easily ramped up to meet the needs of a cross-border payment network.

Performance Analysis

At the end of operations phase, a performance assessment was carried out. No hardware upgrade was performed before the performance study, thus servers with modest CPU/RAM (generally dual CPU and 4/8 GB RAM virtual machines) were used. The environment was hybrid with some banks opting for on-premise, while other banks opting for cloud based deployments. The data centre for a bank was within the geographical limits of its home country. Dynamic Multipoint VPN (DMVPN) was used to provide both hub-spoke and spoke-spoke connectivity. One of the central banks was used as the VPN hub.

The primary solution components are shown in Figure 18. Deployment was completely distributed so each participant had its own instance of the complete solution stack. This comprised of blockchain peers (two per organization), DLT adapter, off-chain database and a CBDC web portal. The off-chain database and the web portal were not central to this exercise. The focus of the performance

study was the REST API gateway (exposed by the DLT adapter) for integrating with the blockchain network.

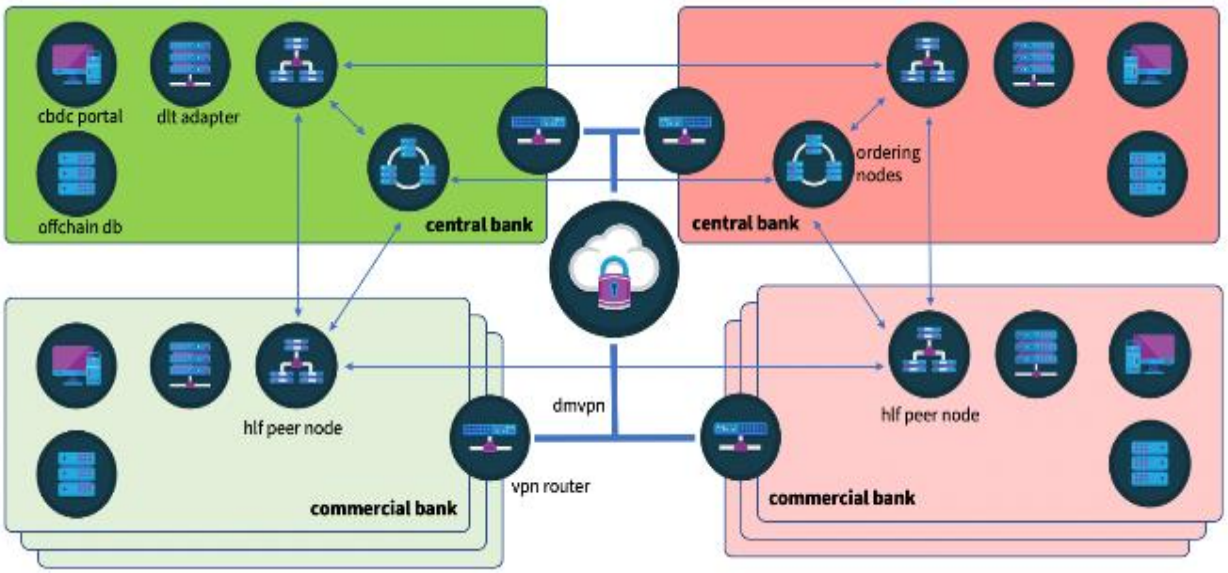

Figure 18: Simplified deployment view of Aber Solution. Hyperledger Fabric 1.4 LTS was used in the project.

Not all operations were equally important from a performance perspective. In a real deployment, digital currency issuance and redemption operations are expected to be much less frequent than its transfer. For this reason, in this report, we only cover the performance of digital currency transfer.

Tests were performed using the JMeter tool. One JMeter instance targeted payments on a single bilateral channel. The instance was executed at the sender bank and generated a sustained load. Measurements were started after a warm up time. The test interval varied from 5-30 minutes. Only transactions committed in the ledger by the end of the interval were considered for throughput calculation. Subsequent runs were carried out at greater load till saturation was reached. Besides throughput, latency was also recorded.

To study the impact of adding another interacting pair of banks, a separate JMeter instance targeting the bilateral channel between the banks was added. Since there were six commercial banks in the network, there were a maximum of (6¦2) or 15 bilateral channels that could be targeted.

The graph of Figure 19 was obtained by plotting total throughput across channels as well as average latency versus the number of channels. For a single channel, a throughput of roughly 15 transactions/sec (TPS) was observed. Moreover, the growth in throughput with number of channels was approximately linear. Thus a throughput in excess of 100 TPS could be easily

achieved for the 6 bank network, while using just 9 of the 15 available bilateral channels.
The graph also demonstrates an average latency of approximately 3 seconds that does not change appreciably with number of channels or the load. This is because the end to end latency was dominated by the multiple round trips performed between data centres spread across two countries over a VPN connection.

The above observations were made with a large settlement interval setting, so that the settlement cycle (steps 2-6 of Figure 16) was not invoked. We observed that the settlement cycle causes a fixed overhead of around 10 sec in our experiments, regardless of the frequency at which it is executed. The incoming payments during a settlement cycle are queued and processed after the cycle. Thus a settlement interval of 1 min will cause a
reduction of 16.7% (10/60) in throughput, while a 5 min interval will only cost 3.3% (10/300). However, settlement overhead does increase slightly with the number of slices impacted in the cycle. The 10 sec overhead for 2 slices roughly doubles to 20 seconds for 20 slices. This suggests why a very small value for slice size could be detrimental to settlement performance.

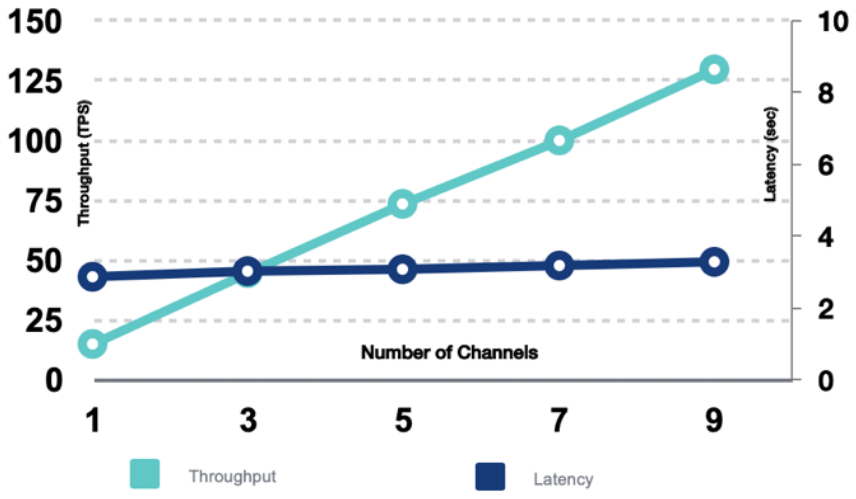

Figure 19: Variation of throughput and latency with number of channels in the network.

Chapter 9

Observations & Learnings

Observations and Learning's

Project Aber had some unique characteristics which presented several challenges but led to some significant lessons learned and observations. This section is a compilation of these experiences and lessons learned during Project Aber; which we expect to be a useful contribution to the community.

The challenges and lessons are grouped under three major categories: IT, Monetary and Business.

IT Challenges Network Connectivity is complex and requires time

The system required mesh connectivity between all participants meaning that each node had to have connectivity to all other nodes. This is contrast to non-DLT based solutions but it ensured higher degrees of architectural resilience, so no node became a single point of failure. This requirement, in the context of this project, was complex because of a number of reasons:

1. Nodes had to be able to communicate with each other multiple jurisdictions and countries. This meant, for example, that all banks in UAE had to have direct IP connectivity to all banks in Saudi Arabia. There was no current network that connected all these participants: the closest being a network between the two central banks but this would have made these central banks single points of failure and would have not fulfilled the objective of driving highest degrees of decentralization;
2. Given the need for nodes to be operated by each of the participating Central and commercial banks, they had different standards, technologies, network options and infrastructure in place and hence some degree of technical heterogeneity that made connecting them complex. Some had opted to use cloud providers which further complicated the matter;
3. There were dependencies on multiple telecom service providers to provide the networks who would take time in deploying such networks;

As a result, a decision was made to create a Dynamic Multipoint VPN (DMVPN) that would provide mesh connectivity to the participants. Participants would then configure their respective network devices to join this VPN and provide IP access to each node. This proved to be the right technical and architectural decision but, still, there were lessons that should be considered for future projects of similar scope or requirement:

1. Start early in discussing the network requirement with the networking and security teams in each participating bank, coordinating joint discussion on how the network connectivity should be designed and implemented;
2. Since the central banks own the project, it is critical for their networking and security teams to own the discussions and lead joint decision-making;
3. Network teams' familiarity with DMVPN configurations is important, particularly when using third party hosting or cloud vendors, so consideration should be given to this; There was a requirement for significant testing and trouble-shooting of port matrices to ensure that there was connectivity between all the different nodes. This required time and coordination effort so should be factored into all deployment timelines.

IT Security should be involved at early stage.

Security requirements go hand-in-hand with networking. The security requirements for the project were complex due to the following reasons:

- The project has to comply with stringent security requirements laid out by both the central banks as well as the individual security requirements of the different commercial banks;
- It will test payments using blockchain which may have its own security requirements that go beyond what are typically required for such systems;
- Multiple options/technologies are available to meet security requirements and different approaches may be used already by different banks which meant aligning on common approaches was difficult;
- The security requirements are likely to become even more stringent once a rollout decision is taken, thus the pilot should demonstrate design being extensible to accommodate long-term production deployment;

The key lesson was to discuss the security requirements early and, as with networking, work with the different teams across the banks. There should be pragmatism as this is a limited pilot and the security approaches should be specific to that, but ensure that any future challenges or issues, relative to security, are surfaced so that they can be addressed in case of an at-scale deployment later.

Software key management solutions should be used instead of HSMs but PKCS#11 support is key

Block chain technologies use public and private keys for its secure transaction processing. The private keys need to be kept very secure to ensure integrity of the transactions and data is maintained.

The preference is to have hardware security module (HSM) to secure private keys. However, as an alternative, software key management systems are also available along with HSM as a Service that can be used.

Based on the discussions with security teams of SAMA and CBUAE, it was decided to use software key management systems considering the following points:

- Pilot is done on limited, low-value transactions;
- All transactions are monitored;
- It is a time-bound short-term project with agreed timelines;

If the current project will be deployed on long-term basis, then a physical HSM should replace the software key management system. It is important to be flexible based on the current needs without compromising key security requirements in order to start quickly. It is possible to switch to a physical HSM at a later stage as long as all support the common PKCS#11 standard API for HSM integration.

Use of Open Source Certificate Authority (CA) and switch to alternatives for production.

Blockchain technologies use public and private keys for its secure transaction processing. The public and private keys have to be issued by a Certificate Authority (CA).

Implementation of CA for the project had a few options: dedicated CA server with online access, use existing CA server in SAMA / CBUAE that can provide offline access or open-source CA server.

The initial assumption from the banks was to use a commercial CA that can provide certificates but, since no web-based transactions are involved, there is little value and hence, after some discussions, a decision was made to use an open source CA. This was due to the fact that It was a limited pilot with a constrained timeline and because there were low value transactions executed during the pilot.

The design of the solution was, however, flexible enough to accommodate a dedicated CA to be hosted in SAMA and CBUAE if the limited pilot expands into a production system. As such, the key lesson learned was to be flexible in accommodating open source or licensed CA servers for the pilot since it can reduce a dependency on internal IT teams, reduces cost, enhances speed of delivery whilst still ensuring the requirements around key issuance and life cycle are met.

Cloud can simplify deployment but not necessarily faster

Pilot projects are generally short-term setup to test technologies and see if it will help meet organizational objectives. Some of them will be successful while others may require more testing or may get dropped. Such projects will need IT investments to setup the required environments. Procurement of such hardware and software can increase the costs while also taking more time to procure, deliver and install them.

Cloud provides an excellent mechanism to get required hardware and software on use and pay basis. The deployment time can be very quick and sometimes immediate. Once the pilot is completed, the hardware and software will be returned to the cloud service provider.

For this pilot, cloud option was considered and CBUAE decided to use this to provide the required environments for this project. SAMA decided to go with on-premise environments for the pilot as it provided an opportunity to SAMA IT teams to understand and learn by deploying new technology.

Although cloud was quick to setup the infrastructure, it was not faster than on-premise since SAMA and Saudi commercial banks already had the infrastructure available. If this had not been the case, then it is likely cloud would have been faster. It is expected that these hyperscale providers would be faster and could, in the future, employ available blockchain platforms as a service that could further accelerate setup times.

Testing is complex with multiple banks

Given nodes deployed into six commercial banks in two countries and two central banks, and the need for participants in each to interact with a distinct web application/API to execute transactions, there was significant complexity in testing the different scenarios that required coordination across these different

entities. Sufficient time should be allocated for such testing, consideration should be given to how communication is managed, and there should be project resources in each participating bank to support with trouble-shooting and guidance during the testing phase.

Simplify integration with existing systems

The project objective is to prove the viability of blockchain for a specific use case. In order to simplify the project, it was important to remove or limit the need for changes to existing systems or the introduction of complex activities that are outside the core scope and purpose of the project. As such, it was decided that integration with internal systems can be manually performed. The integration efforts can be high and extensive testing may be required. As such, a key lesson learned was to build the required APIs, so that the capability was ready, but then build a web front-end using the same API to mitigate the need for any system-level integrations. If the banks wanted to do a system-level integration, it would still be possible, but it would not be a prerequisite for the project. The updates to the core systems would be performed manually but still required the banks to determine necessary changes or impacts that would be useful later if a real integration was made.

Commercial banks have high degree of automation that make execution of a subset of cross border transactions challenging from an operational perspective.

It was originally thought that commercial banks could move all or a subset of transaction classes to the new system. On discussion with the commercial banks, it became clear that some have a high degree of automation and this would prove challenging since, in many cases, the systems would automatically trigger the necessary downstream steps and it is challenging to create exceptions. As such, a decision was made to create synthetic transactions that, although resembling real world transaction cases, would not be linked to actual real-world customers.

Monetary Policy

Currency rate should be kept static for project.

The FX rate between the two fiat currencies and the digital currency was considered for the Project and a decision was made to keep it fixed whilst

providing the ability for it to be changed in the future. The reason was that this would have introduced significant additional complexity outside of the core focus of the pilot.

Basket of currencies and rebalancing consideration

One of the initial ideas was that the digital currency would be linked to a basket of currencies that would be maintained by each central bank. For example, one digital currency would be linked to 50% of AED holdings and 50% of SAR holdings held in the respective central banks. The challenge with this concept was that, as other jurisdictions joined the network (if the project was taken further in the future), it would require the basket to be rebalanced and thus presented scaling challenges; and, if a non- pegged currency was to be introduced, it would similarly lead to challenges and complexity. As such, a decision was made to have each digital currency issued independently by each central bank and the commercial bank requesting the issuance to pledge using their local currency.

Currency should be traceable and immutably linked to issuing authority.

An early requirement was that that digital currency should be issuable by either central bank but must always be traceable back to the central bank that issued it even though it may have effectively moved cross-border multiple times. This was needed for general central bank supervisory purposes but also because, at the time of redemption, the digital currency issued by the local central bank would be redeemed from pledged funds whereas the digital currency issued by the foreign central bank would be redeemed using funds in their nostro account. The ability to know where a particular unit of currency originated was therefore a critical element.

Different Jurisdictions apply different interest rates.

In a cross-border context with a dual-issued digital currency, it became apparent that the differences in policy between the two jurisdictions would introduce some complexities if this was to be taken forward. Key amongst these was the difference in how each jurisdiction currently handles interest rates. As the digital currency was backed by pledged commercial bank funds with the respective central banks, this was an important observation because, whilst for one jurisdiction that didn't pay interest there was no impact, for jurisdictions that paid an overnight rate, there was a potential opportunity cost imposed on the commercial banks who held the digital currency overnight instead of

redeeming. Whilst interest wasn't implemented in the project, the options by which this could be handled were explored and some approaches documented:

1. Digital currency carries the interest of the issuing central bank. If digital currency carries the interest typically paid by the issuing central bank and given the interchangeability of the digital currency from a commercial bank perspective, this means that the commercial bank would only receive interest based on digital currency issued by central banks that pay interest and would not receive on others. This creates a distinction between the digital currency that undermines the objective of the project but also creates an economic incentive for commercial banks to engage in different arbitrage activities.
2. Digital currency pays an interest rate that is agreed by all participating central banks. This again exposes the system to arbitrage opportunities as some commercial banks may move funds into digital currency simplify to benefit from a rate that is higher than what is offered in their domestic central banks. This may be acceptable but is an important consideration if setting a universal rate for digital currency.
3. Digital currency pays a rate based on the jurisdiction in which it is held. In this scenario, the commercial bank always receives interest on their held digital currency based on the prevailing interest rate in that country – regardless of who issued the digital currency. This creates the right incentives and mitigates arbitrage opportunities, but further work is then needed on how these interest payments would be funded and by whom since some of the held digital currency will be issued by countries that pay no or less interest than the jurisdiction in which they are held

This is an area that needs further consideration if the system is to be taken into production but represents a significant opportunity for future research.

Full visibility of money supply between central banks.

One of the early lessons learned was that, in order for a dual issued digital currency to work, there should transparency between central banks on how much the other is issuing and what the total money supply is. This is because any currency issued by one central bank could theoretically end up in the second jurisdiction for eventual redemption which could lead to an assortment of

operational and policy impacts. As such, a decision was made that each central bank should have visibility of what the other is issuing and, in the future, limits could be mutually agreed e.g. based on balance of trade.

Interest need to be aligned with overnight rate to create right incentives.

One of the key objectives was to maximize utility for both central banks and commercial banks. As such, the digital currency should not penalize commercial banks for holding it. It became apparent that, in Saudi Arabia, SAMA has an overnight rate paid to commercial banks who deposit with them at end of day. If the digital currency does not pay an equivalent rate, then there is an economic incentive for the commercial bank to redeem each day and deposit these funds to benefit from the said rate. For the purposes of the pilot, we did not introduce interest as there are other complexities in a cross-border context, but a key lesson learned for future projects, particularly if going live, is that unless interest is paid on the digital currency and equivalent to the overnight rate, there is a risk of commercial banks redeeming each day which obviously creates a burden on the central bank and increases costs as opposed to reducing them.

Business Challenges

Cut-off time for overnight rate interest calculation should be agreed.

An implication of paying interest on digital currency is that the time at which the interest calculation is made should be considered. There is a tension between

this concept and the desire that the DLT should enable the possibility of 24x7 cross-border payments.

The central bank should assure commercial banks of settlement finality.

There was some initial concern around the finality of settlement using the new system. In other words, that the counter-parties could hold each legally accountable for the transaction and there is legally-enforceable finality when a transaction was executed on the ledger. This is not a technical factor but a legal one and hence it was observed that, if the decision is made to take the system further, the central banks should issue a circular or directive clarifying this point to the participating banks.

Real value isn't the same as real customers and real transactions.

The objective of the pilot was to use real money which was initially thought to imply real customers and real transactions. Due to the complexity of carving out discrete transactions from production systems and other concerns related to the use of the pilot system for such transactions, such as legal liability to customers, it became apparent that it would be challenging and could delay the timeline for the pilot. As such, a decision was made to use real value and simulated customers and transactions. This was achieved by the commercial banks depositing/pledging actual funds with the central bank and equivalent amounts of digital currency created that would then be used to simulate defined transaction types between the two counter parties.

Use Nostro account for redemption.

A question arose around how redemption should be funded given the dual issued digital currency. The decision, which proved effective, would be the use of the central banks' Nostro account with the other central bank(s) to ultimately fund the redemption. This simplified things significantly, as it is already used for different types of transactions and settlement activity, but, if taken forward, consideration should be given to how this is managed to ensure sufficient liquidity. This could mean bilateral agreement between the central banks on how overdrafts can be handled as well as how each bank will replenish funds by acquiring the other currency.

Importance of limits.

As the project discussed the concept of the system with various stakeholders within the banks, it became clear that key to succeeding with the project was establishing limits. These limits would serve to manage risk and allow the project to proceed with real money. The limits were therefore set on the value of the digital currency, the total amount that could be issued, the total that could be redeemed, and the size of each transaction. These were daily limits and transaction-level limits. Through the application of these limits and agreement thereof with the different stakeholders in the Commercial and central banks, it was possible to secure their agreement to proceed with real money.

Chapter 10

Conclusion and Future Work

Conclusion

Project Aber was accomplished by a fruitful collaboration between all stakeholders, including valuable contributions from both the central banks, participating commercial banks and technology partner, reflecting the shared urgency to shape the application of DLT to prevail over existing pain points in cross-border transfers.

The path to success was interspersed with numerous challenges. During the project, several challenges were overcome and key milestones were achieved. A few contributions of Project Aber highlighting important achievements that were shared in this report.

The project was successful in achieving its key objectives which include using a new DLT based solution for real time cross-border interbank payments between commercial banks without the need to maintain and reconcile Nostro accounts with each other. This promises to address the inefficiency and costs that are inherent in existing cross-border payment mechanisms.

As envisioned by one of the objectives, the approach was based on a first of its kind, dually issued central bank digital currency. Additionally, the solution was designed after thorough analysis of previous CBDC projects. The project identified further areas that need to be explored in the future if the approach of a single digital currency is to be implemented.

The design principles were drafted such that the benefits of DLT could be accentuated. Foremost in this respect, was the need to decentralize to the maximum extent possible. Practically, this meant that the peer to peer payments between commercial banks could be executed without active involvement of central banks. Such decentralization, while addressing safety, finality and privacy concerns of commercial banks had not been achieved in a single payment solution before. This resulted in the design of the Aber Protocol. A significant part of project Aber was envisioning a new future for DLT based cross-border interbank payments

In Aber, we were able to make a number of significant changes to the status quo of cross-border payments. Firstly, by creating a common CBDC for use between jurisdictions, we removed the need for Commercial Banks to maintain nostro accounts with their various cross-border counter-parties. Secondly, we implemented a system that, whilst requiring the Central Banks to issue and

redeem CBDC, didn't require the Central Banks' systems to be running in order for cross-border payments to continue as this function was then distributed across all the participating Commercial Banks (whilst still ensuring that settlement finality, privacy, and other considerations are met). Thirdly, the role of the Central Banks evolved in that they maintained nostro accounts with each other that was then used to fund the redeemed CBDC (when that CBDC was originally issued by that particular Central Bank).

Another contribution of Aber protocol was its handling of the safety, privacy and audit ability implications of multiple jurisdictions. This added a completely new dimension to the interbank payments problem and the solution. We believe, Project Aber is a significant step forward in the world of dually issued single currency CBDC initiatives.

Suggested next steps

Just as Project Aber built on and extended prior work in this field, successive projects may build on the achievements of Project Aber in order to further expand the body of knowledge in Central Bank Digital Currencies.

Delivery versus Payment

Integration as block chain has seen significant adoption in different fields, particularly around the tokenization of different assets, such as bonds or debentures, there continues to be significant opportunity for improvement in how the Delivery versus Payment (DvP) challenged is addressed in a DLT-context. For example, if a debenture is being issued and traded on a debenture-oriented block chain, how can payment taken place between the counter-parties atomically and thus mitigate the inefficiencies that exist in many contemporary securities markets?

The challenge becomes more complicated when one considers that many of these different networks will exist on different blockchain platforms or protocols.

Developing approaches, frameworks, and assets focused on this integration and interoperability scenario would be an important area for future work; with the vision being that networks such as Aber can become a generic DLT-based payment rails for use alongside various purpose-built networks.

Expansion to multiple currencies

Although Aber sought to prove the viability of a single Digital Currency as an instrument of cross-border settlement, there is also potential to expand it accommodate multiple digital representations of fiat currencies. This would provide a Payment versus Payment (PvP) capability that could mitigate the settlement risk that exists particularly in these types of transactions and which organizations such as Continuous Linked Settlement (CLS) have sought to solve for a limited set of currencies. Given that Aber had a requirement to always ensure a currency was traceable to the authority that issued it, there is the potential to relatively easily adapt to multiple currency support so this could be area of future investigation. Interoperability with other CBDC networks as a corollary to the expansion to support multiple currencies, Aber could also integrate with the different CBDC networks or initiatives underway globally. This would present a number of technical challenges around integration and interoperability that would be important to solve and would support what is likely to be the future architecture for CBDC networks wherein regions or large countries will establish their own networks and will therefore seek to integrate with others.

Development of methods for calculating interest

As discussed in this paper, the project identified a broad set of considerations around policy when a single digital currency is being used cross-border. Key amongst these was the differential in the handling of interest across different countries and how this can be accommodated in such as CBDC solution. A range of high level options were identified but this is certainly an area for future work and research, particularly as there are features of DLT that could allow a reimagining of how interest is handled in such a system so as to create the right incentives for bank adoption and usage; and mitigate negative consequences such as the creation of arbitrage opportunities.

Straight-through processing

The interest question largely arises when there is a model wherein commercial banks acquire DC and then use it for a period of time before redemption. A possible alternative approach is to not allow the holding of the DC but rather, each time a transfer needs to happen, the issuance and redemption happens atomically as part of the exchange of value. This would require exploration of a number of other areas, such as how KYC, AML and similar, could work in a DLT-based cross-border network and also what standards need to be adopted in order for the network to be used for the different cross-border transactions. This was an area that was explored as part of the Aber project and the capability was developed to associate metadata, such as ISO20022 or SWIFT messages, with a transaction but this is certainly an area for future work.

Decentralized Gridlock Resolution Algorithms

As with Project Ubin, there is some degree of centralization in how gridlock is resolved where the Central Banks play an active role. An area for future research is to explore how gridlock resolution can be handled in a distributed manner which may involve the adaptation or creation of new algorithms that take advantage of the unique characteristics of DLT. Furthermore, there is potential for artificial intelligence or machine learning to be incorporated here alongside more decentralized approaches to liquidity saving.

Monetary Policy, Regulatory and Economic Considerations

At the technical level, Project Aber has addressed several impediments that have previously hindered the adoption of DLT technology in the CBDC space. However, turning the vision of DLT based CBDC into reality by considering the policy and regulatory aspects, is an area of future work. This includes decisions on monetary subjects such as exchange rates, interest rates that should be applied to the digital currency issued on DLT. The regulatory implications of tokenization and DLT based solutions (covering Aber as well as future scope) on existing legal framework and financial market infrastructure is another important area that requires comprehensive evaluation.

Finally, a comparative cost analysis of DLT based payments versus traditional cross-border fund transfer mechanisms should be carried out to make a compelling case for the technology.

Operational Roles of the Central Bank and Private Institutions in CBDC

Both wholesale CBDC initiatives, such as Aber, and retail CBDC initiatives raise important questions about the role of the Central Bank in the issuance of currency, management of the cross-border payment flows, and the eventual redemption of that currency. Just as we challenged the role of the Central Bank in being the enabler of cross-border payments by leveraging some of the unique characteristics of Distributed Ledger Technology, we can take this further in envisioning a CBDC that is not just a wholesale instrument but can be more generally utilized. In doing so, consideration needs to be given as to whether the general public should have direct access to a CBDC that is a claim on the Central Bank itself. A consequence of the Central Bank assuming direct relationship with the end user is that operational tasks, like KYC, would need to be performed by the Central Bank as well as the associated regulatory compliance burden. As such, it may be preferable to develop a tiered approach to CBDC issuance that resembles the current situation wherein Commercial Banks will have access to CBDCs that are claims on the Central Bank itself but Commercial Banks will themselves be able to issue CBDCs to the general public with the retail user only interacting with the Commercial Banks.

We believe that this is an important area of future study: particularly to explore the different monetary policy, structural and industry implications of the different technical possibilities and how these will lead to an evolution of the role of the Central Bank.

Current and emerging technologies allow us to reimagine how many functions are existed and the different roles of the different industry participants. Key amongst these are the role of the Central Bank and particularly its relationship with the retail sector and the Commercial Banks.

Abbreviations

AED	United Arab Emirates Dirham	**HSM**	Hardware Security Module
AML	Anti-Money Laundering	**HTLC**	Hash Time Locked Contract
API	Application Programming Interface	**ILP**	Inter Ledger Protocol
BIS	Bank of International Settlements	**KSA**	Kingdom of Saudi Arabia
BOC	Bank Of Canada	**KYC**	Know Your Customer
BOJ	Bank of Japan	**LSM**	Liquidity Saving Mechanism
CA	Certificate Authority	**MAS**	Monetary Authority of Singapore
CB	Central Bank	**PFMI**	Principles for Financial market infrastructures
CBDC	Central Bank Digital Currency	**PKCS#11**	Public-Key Cryptography Standard#11
CBUAE	Central Bank of the UAE	**PKI**	Public Key Infrastructure
CLS	Continuous Linked Settlement	**POC**	Proof of Concept
DC	Digital Currency	**PVP**	Payment vs. Payment
DLT	Distributed Ledger Technology	**RTGS**	Real Time Gross Settlement
DMVPN	Dynamic Multipoint VPN	**SAMA**	Saudi Central Bank
DVP	Delivery vs. Payment	**SAR**	Saudi Riyal
ECB	European Central Bank	**SARB**	South African Reserve Bank
FMI	Financial Market Infrastructure	**SWIFT**	Society for Worldwide Interbank Financial Telecommunication
FX	Foreign Exchange	**UC**	Use Case
GCC	Gulf Cooperation Council	**UTXO**	Unspent Transaction Output
HLF	Hyperledger Fabric	**VPN**	Virtual Private Network

References

List of References

1. Bech, 2017 Bech, M. L. and Garratt, R., Central Bank Cryptocurrencies (September 17, 2017). BIS Quarterly Review September 2017 [online]. Available at https://www.bis.org/publ/qtrpdf/r_qt1709f.htm [Last accessed, Dec 26, 2019]
2. CPSS, 2012 Committee on Payment and Settlement Systems (CPSS) and Technical Committee of the International Organization of Securities Commissions (IOSCO), "Principles for financial market infrastructures" (April 2012).
3. SARB, 2018 South African Reserve Bank, Project Khokha, Exploring the Use of Distributed Ledger Technology for Interbank Payment Settlements (Jun 2018) [online] Available at: https://www.resbank.co.za/Lists/News%20and%20Publications/Attachments/8491/SARB_ProjectKhokha%2020180605.pdf, [Last accessed, Dec 26, 2019]
4. ECB, BoJ, 2017 STELLA - joint research project of the European central bank and the Bank of Japan. Payment systems: liquidity saving mechanisms in a distributed ledger environment. (Sep 2017) [online]. Available: https://www.ecb.europa.eu/paym/intro/publications/pdf/ecb.miptopical190604.en.pdf [Last accessed, Dec 26, 2019]
5. ECB, BoJ, 2018 STELLA - joint research project of the European central bank and the Bank of Japan. Securities Settlement Systems: Delivery vs. Payment in a Distributed Ledger Environment. (Mar 2018) [online]. Available at: https://www.ecb.europa.eu/pub/pdf/other/stella_project_report_march_2018.pdf [Last accessed, Dec 26, 2019]
6. ECB, BoJ, 2019 STELLA - joint research project of the European central bank and the Bank of Japan. Synchronized Cross-Border Payments. (Jun 2019) [online]. Available at: https://www.ecb.europa.eu/paym/intro/publications/pdf/ ecb.miptopical190604.en.pdf [Last accessed, Dec 26, 2019]
7. MAS, Mar 2017 Project Ubin: SGD on Distributed Ledger, A report developed with the contributions of Bank of America Merrill Lynch, BCS Information Systems, Credit Suisse, DBS Bank, HSBC, J.P. Morgan, Mitsubishi UFJ Financial Group, OCBC Bank, R3, Singapore Exchange and UOB Bank (Mar 2017). [online] https://www.mas.gov.sg/-/media/MAS/ProjectUbin/Project-Ubin--SGD-on-Distributed-Ledger.pdf. [Last accessed, Dec 26, 2019]
8. MAS, Nov 2017 Project Ubin Phase 2: Re-imagining Interbank Real-Time Gross Settlement System Using Distributed Ledger Technologies (Nov 2017), [online] https://www.mas.gov.sg/-/media/MAS/ProjectUbin/Project-Ubin-Phase-2-Reimagining-RTGS.pdf [Last accessed, Dec 26, 2019]
9. BoC, 2017 Project Jasper, A Canadian Experiment with Distributed Ledger Technology for Domestic Interbank Payments Settlement (Sep 2017), [online] Available at: https://www.payments.ca/sites/default/files/29-Sep-17/jasper_report_eng.pdf [Last accessed, Dec 26, 2019]

10. BoC, 2018 Project Jasper Phase III, Securities Settlement using Distributed Ledger Technology, October 2018. [online] Available at: https://www.payments.ca/sites/default/files/jasper_phase_iii_whitepaper_final_0.pdf [Last accessed, Dec 26, 2019]
11. Poon, 2016 Poon J, Dryja T, The Bitcoin Lightning Network: Scalable Offchain Instant Payments, Jan 2016, [online] Available at: https://lightning.network/lightning-network-paper.pdf [Last accessed, Dec 26, 2019]
12. Thomas, 2015 Thomas S, Schwartz E, A protocol for Interledger Payments [online] Available at: https://interledger.org/ interledger.pdf [Last accessed, Dec 26, 2019]
13. BoC, MAS, 2019 Jasper – Ubin Design Paper, Enabling Cross-Border High Value Transfer Using Distributed Ledger Technologies (May 2019) [online] https://www.mas.gov.sg/-/media/Jasper-Ubin-Design-Paper.pdfs [Last accessed, Dec 26, 2019]
14. R3, 2019 R3 Corda master documentation [online]: Home: https://docs.corda.net, Key Concepts > Consensus: https://docs.corda.net/key-concepts-consensus.html [Last accessed: Dec 26, 2019]
15. HLF, 2019 Hyperledger Fabric documentation [online]: Home: https://hyperledger-fabric.readthedocs.io, MSP implementation with Identity Mixer: https://hyperledger-fabric.readthedocs.io/en/release-1.4/idemix.html [Last accessed: Dec 26, 2019]
16. McKinsey, 2016 Global Payments 2016: Strong fundamentals despite uncertain times [online]: https://www.mckinsey.com/~/media/McKinsey/Industries/Financial%20Services/Our%20Insights/A%20mixed%202015%20for%20the%20global%20payments%20industry/Global-Payments-2016.ashx [Last accessed: Dec 26, 201

In case of any questions, you may kindly write to us at the following:

Saudi Central Bank

E-mail: finsecdev@sama.gov.sa

Central Bank of the U.A.E

E-mail: *uaecbbod@cbuae.gov.ae*