




South African Reserve Bank

PROJECT KHOKHA

*Exploring the use of distributed ledger technology
for interbank payments settlement in South Africa*



FOREWORD

*Francois Groepe, Deputy Governor,
South African Reserve Bank*

Our goal with Project Khokha is to contribute to the global initiatives which assess the application and use cases of distributed ledger technology (DLT) through this collaborative effort piloted by the South African Reserve Bank (SARB) together with the national banking community.

The role of banking institutions, central banks and financial market infrastructures, both domestically and internationally, is being disrupted at a rapid pace by innovative and emerging technologies. Instead of ignoring these developments, the South African financial services industry has risen to the challenge of learning, growing and understanding the implications and opportunities that they present in a South African context. We recognise that there is a cost to impeding innovation, and we are keen to support innovation where it could give rise to greater efficiencies without undermining the soundness and stability of the broader financial system.

The SARB initiated Project Khokha in the latter part of 2017, with the project team consisting of seven banking industry participants, a technical service provider (ConsenSys), and consulting practice PricewaterhouseCoopers Inc. The results of this project are captured in this paper, which describes the key objectives, approach, technical testing results and insights gained from this proof of concept.


The scope of Project Khokha was to trial interbank wholesale settlement using DLT. Several astute insights were attained from taking a robust approach to the trial. A significant contribution to the global DLT body

of knowledge was made, as it is thought that this was the first time that the Istanbul Byzantine Fault Tolerance consensus mechanism and Pedersen commitments for confidentiality had been used with Quorum. Furthermore, the DLT nodes were operated under a variety of deployment models (on-premise, on-premise virtual machine, and cloud) and across distributed sites while processing the current South African real-time gross settlement system's high-value payments transaction volumes within a two-hour window.

An important outcome from this project is the realisation that there is more with which to deepen and expand our knowledge. It is evident that significant benefits will be realised by leveraging collaboration across an industry – and even more so when global communities work together.

Given the speed and disruptive nature of digital innovation, it is undisputable that regulators will need to be more proactive and more responsive. The practice of regulation following innovation thus also needs to change. Regulators cannot control the future, but, through the establishment of appropriate and responsive regulatory frameworks, we can contribute positively to the betterment of our societies, including their economic well-being.

We as the SARB are committed to contributing to efforts aimed at leveraging innovation and technology to improve efficiencies while promoting the safety and soundness of the broader financial system. And although Project Khokha is but one initiative in this regard, it is nonetheless an important step on this journey. We invite you to be part of the journey.



“OUR GOAL WITH PROJECT KHOKHA IS TO CONTRIBUTE TO THE GLOBAL INITIATIVES WHICH ASSESS THE APPLICATION AND USE CASES OF DISTRIBUTED LEDGER TECHNOLOGY (DLT) THROUGH **THIS COLLABORATIVE EFFORT PILOTED BY THE SOUTH AFRICAN RESERVE BANK (SARB) TOGETHER WITH THE NATIONAL BANKING COMMUNITY.**”

Francois Groepe, Deputy Governor

CONTENTS

06

06 EXECUTIVE SUMMARY

09 INTRODUCTION

10

CHAPTER 1 | THE CONTEXT AND BACKGROUND OF PROJECT KHOKHA

11 AN EVOLVING TECHNOLOGY
AND GROWING EXPERIMENTATION

18 THE SOUTH AFRICAN RESERVE BANK
AND THE APPROACH TO FINTECH

20 THE SOUTH AFRICAN REAL-TIME
GROSS SETTLEMENT SYSTEM

22 THE RATIONALE AND KEY OBJECTIVES
OF PROJECT KHOKHA

24

CHAPTER 2 | THE STRUCTURE OF PROJECT KHOKHA

25 PARTICIPANTS AND ROLES

26 OBJECTIVES, SCOPE
AND EXCLUSIONS

32 TECHNOLOGY SELECTION

33 THE APPROACH

35 HYPOTHESIS TESTED

36

CHAPTER 3 | RESULTS

37 PROJECT SUCCESS FACTORS

38 THE OUTCOME OF HYPOTHESIS TESTS

40 TECHNICAL FEASIBILITY

41 SCALING

44 THE CONSENSUS MECHANISM

44 CONFIDENTIALITY

46

**CHAPTER 4 | COMPARISON
TO SIMILAR PROJECTS**

52

**CHAPTER 5 | CONCLUSION
AND FUTURE WORK**

53 INNOVATION AND COLLABORATION

54 THE SUITABILITY OF DISTRIBUTED
LEDGER TECHNOLOGY

58 IMPLEMENTATION CONSIDERATIONS

60 POTENTIAL FUTURE DIRECTION

64

CHAPTER 6 | APPENDICES

65 DETAIL ON CENTRAL-BANK DLT PROJECTS

73 DETAIL ON PROJECT KHOKHA
TESTING ITERATIONS

75 LIST OF REFERENCES

77 LIST OF ABBREVIATIONS

78 ACKNOWLEDGEMENTS

EXECUTIVE SUMMARY

Project Khokha is a collaborative project led by the SARB and involving a consortium of South African settlement banks as well as technical and support partners. The goal of the project was to build a proof-of-concept (PoC) wholesale payment system for interbank settlement using a tokenised South African rand on distributed ledger technology (DLT). The project commenced in January 2018 and ran for 14 weeks. It is the first project initiated by the recently constituted Fintech Unit in the SARB.

Project Khokha builds on the initiatives previously undertaken by other central banks, including those of Brazil, Canada, Europe, Japan and Singapore. These projects have typically been rapid builds of PoCs exploring the use of DLT in a banking context. This report identifies three 'waves' of this type of work. The first wave consists of simple PoCs built using Ethereum. The second wave takes this work and investigates the interconnected issues of scalability, resilience, confidentiality and finality. The third wave, which is just starting, diverges into areas such as securities settlement and cross-border transactions.

Project Khokha forms part of the second wave and was designed to provide a realistic test of a DLT-based wholesale payments system. In particular, it investigates whether confidentiality could be achieved at scale, i.e. with production-level volumes, and whether multiple node types, each configured by a participating bank, could be accommodated.

The project was built on Quorum, using Istanbul Byzantine Fault Tolerance (IBFT), Pedersen commitments and range proofs to deliver on the combination of scalability, resilience, confidentiality and finality. This is thought to be the first time that a PoC has used Pedersen commitments and range proofs on a Quorum network using IBFT. Each bank was responsible for configuring its own node on the network, and the seven banks used a mixture of physical machines, on-premise virtual machines and cloud-based virtual machines.

The results show that the typical daily volume of the South African payments system could be processed in less than two hours with full confidentiality of transactions and settlement finality. This was done using ISO 20022 standard messages, propagated within two seconds, across a network of geographically distributed nodes, with distributed consensus providing the requisite resilience. The SARB was able to view the detail of all the transactions to allow for regulatory oversight.

This project has laid the foundations for future collaborative work – essential in the blockchain context – and has fulfilled its objective of providing useful insights to all participants. This report notes that there are many issues to consider before the decision to take a DLT-based system into production can be taken. Some of these issues relate to the practicalities of implementation, but also to legal and regulatory factors and to the broader economic impact. The future direction will also be influenced by further development of the technology and by central banks and other regulators globally continuing to contribute to this field of knowledge.

The SARB anticipates continuing work in this area and expects to continue contributing to the body of work in DLT-based systems.



KHOKHA IS THE
ISIZULU WORD
MEANING **'PAY'**.





INTRODUCTION

Project Khokha was formally initiated by the SARB early in 2018 to experiment with DLT as a mechanism for managing wholesale payments between settlement banks. Khokha is the isiZulu word meaning 'pay'. The project was a collaborative effort initiated by the SARB with the goal of building on the body of knowledge of DLT through a PoC.

The initiative involved numerous partners, including a consortium of banks (Absa, Capitec, Discovery Bank, FirstRand, Investec, Nedbank and Standard Bank) together with ConsenSys as the technical partner and PricewaterhouseCoopers Inc. (PwC) as the support partner. Planning started late in 2017, with execution running for 14 weeks from January to April 2018. The project was managed by the Fintech Unit within the SARB, and it is expected that the SARB will contribute to innovative and collaborative initiatives in future in the fintech arena.

01

THE CONTEXT AND BACKGROUND OF **PROJECT KHOKHA**

AN EVOLVING TECHNOLOGY AND GROWING EXPERIMENTATION

Distributed ledger technology is developing rapidly and its use is growing.

With the publication of the seminal paper in 2008, the pseudonymous Satoshi Nakamoto introduced bitcoin to the world. The blockchain technology that underpins this crypto-currency has since been developed in many industries, with financial services leading the way. The bitcoin blockchain is public and decentralised, and therefore needs a high degree of cryptographic security in order to support a 'trustless' environment. The potential applications of DLT to uses in areas such as payments, identity and trade finance, and many other transaction types, mean that financial institutions around the world have been experimenting with this technology. Many banks and other financial services organisations have publicised their work with DLT and we have seen a number of consortiums forming to collaborate in this area.

Private crypto-currencies are still a relatively novel concept, therefore different interpretations and terminology have been used in the public domain. Some global standard-setting bodies, such as the Financial Stability Board, prefer the term 'crypto-assets' or 'crypto-tokens'. This report uses both of these terms as reference to these financial technology innovations.

***Distributed ledgers** are a family of technologies where participating nodes in a network each update their copy of a common set of data. **Blockchains** are a type of distributed ledgers where transactions are organised into consecutive groups, or blocks, that are cryptographically secured.*

The trade-off for security in blockchains like bitcoins is the proof-of-work consensus mechanism which effectively limits the speed at which transactions can be processed and has well-publicised implications for energy usage and the concentration of groups of miners who perform the transaction verification. In permissioned, or private, blockchain environments, the implicit trust that the members of the network have in each other means that simpler consensus mechanisms can be used, and these distributed ledgers form the basis of the systems that financial institutions globally are experimenting with. DLT platforms developed specifically for financial services use include Corda, Hyperledger Fabric, Quorum and Ripple. Quorum is based on Ethereum, and is the technology used by this project.

The process of understanding and leveraging new technology tends to move through stages as the technology becomes more broadly and deeply accepted by its users. The first uses to which we put new technologies tend to be focused around the things that we perceive them to be replacing 'a faster horse'. The next phase entails a better understanding of the technology, it becomes cheaper and more abundant and supporting structures are created. In this stage, new business models emerge that make full use of the technology. Cars, for example, moved from being toys of the rich to spawning new industries in commerce and leisure. Finally, the technology becomes ubiquitous and almost invisible as it becomes part of everyday life.

DLT has followed this pattern so far in that early applications are typified by replicating things that are already achieved with previously existing technology. This approach of doing things better often enables, or leads directly to, experiments with doing things that were impossible with the previous generation of technology. Financial services are currently leading the move into this stage with DLT, enabling new business models which are the most exciting aspect of this sphere of activity. In the early stages of a new technology, the necessary skills to implement it are scarce and its capabilities are not well understood. It can often take a while for enabling technologies and structures to emerge before adoption takes off and this is what we are seeing with the current developments in DLT.

A key feature of these schemes is their asymmetry: the work must be moderately hard (but feasible) on the requester's side but easy to check for the service provider.

A proof-of-work protocol, or function, is an economic measure to deter denial-of-service attacks and other service abuses by requiring some work from the service requester, usually meaning processing time by a computer.



Several central banks globally are experimenting with distributed ledger technology.

Many commercial banks around the world have been among the leaders in experimenting with blockchain technology. Inevitably, these experiments started with small teams and then grew, with loan syndication or trade finance being typical use cases. Cross-border remittances are another use case in which banks have used crypto-assets inside a trusted network to transfer value locally and internationally.

Like commercial banks, central banks are cognisant of the potential implications of fintech, and are creating the capabilities and structures to understand it better, often by setting up specialist units and/or the so-called 'sandboxes' where companies can experiment, within limits, in a regulatory 'safe space'.

For central banks evaluating blockchain and DLT, the wholesale payments process has been the most common type of experiment undertaken at any scale, and several central banks have already carried out, or are currently engaged in, experiments in wholesale payments systems based on DLT. It is into this group that Project Khokha fits and the results of this work will add to the global knowledge base that has come from central banks publishing the results of their work.

Table 1 summarises some of the central-bank initiatives that have been made public to date. While central banks other than those listed have done work with DLT and crypto-currency, they have not all published their results, or they have provided only high-level summaries. This group includes the Riksbank of Sweden, which is investigating an 'e-krona' as a complement to cash, and the Bank of England (BoE).



Table 1: Summary of central-bank initiatives

Project	Results published	Main focus	Technology used
Brazil (Central Bank of Brazil) 			
Phase 1	August 2017	Identify use cases and build a working prototype for a select use case.	Ethereum
Phase 2	August 2017	Look to build 'realistic' functionality for a real-time gross settlement (RTGS) system proof of concept (PoC) on distributed ledger technology (DLT), and analyse DLT platforms.	Corda Hyperledger Fabric Quorum
Canada (Project Jasper) 			
Jasper – Phase 1	September 2017	Create a wholesale interbank payments settlement (RTGS) PoC.	Ethereum
Jasper – Phase 2	September 2017	Expand on the scalability and flexibility of the original PoC, including building in liquidity saving mechanisms (LSMs).	Corda
Europe and Japan (Project Stella) 			
Stella – Phase 1	September 2017	Build an RTGS PoC with DLT, including LSMs, and investigate the performance of the system.	Hyperledger Fabric
Stella – Phase 2	March 2018	Build a delivery versus payment (DvP) PoC using a number of different DLT platforms.	Corda Elements Hyperledger Fabric
Singapore (Project Ubin) 			
Ubin – Phase 1	March 2017	Create an RTGS PoC using DLT.	Ethereum
Ubin – Phase 2	November 2017	Expand on the original PoC by incorporating LSMs, and compare DLT platforms.	Corda Hyperledger Fabric Quorum
South Africa (Project Khokha) 			
Khokha	June 2018	Build an RTGS PoC on DLT, expanding on privacy and scalability in a 'realistic' environment.	Quorum

The BoE has conducted a number of different studies, publishing only high-level results. In June 2016, it was announced that the BoE was conducting a DLT PoC leveraging Ethereum and considering gross settlement and transfer of value. The PoC assessed scalability, security, privacy, interoperability and sustainability. Since then, the BoE has done other studies, including a PoC with Ripple to explore how the use of Ripple Connect and the Interledger protocol may be able to lower settlement risk and improve the speed and efficiency of cross-border payments. A further study involved an academic exercise that looked at ways to maintain privacy over DLT while still allowing a regulatory body to have a view of the data.

South Africa's banking industry is collaborating around blockchain.

South Africa's financial services industry is innovative and sophisticated, with the soundness of banks ranked 37th out of 137 countries by the World Economic Forum's Global Competitiveness Report¹ and the world's third-highest life insurance penetration². It is worth noting that the domestic financial sector weathered the global financial crisis of 2008 well due to the country's sound macroeconomic fundamentals and a robust financial regulatory framework.

Commercial banks and other financial institutions in South Africa have been collaborating around blockchain since 2016. This collaboration has resulted in the creation of the South African Financial Blockchain Consortium (SAFBC) (safbc.co.za), recognised by *The Banker* magazine as 'Blockchain Technology Project of the Year' in 2017³. The SAFBC has launched a number

of initiatives, including issuing a syndicated loan via a shared blockchain network, for which the SARB participated as an observer. Members of the consortium represent a very high proportion of financial services activity in the country, including all the large banks. This has been of benefit to the current project, since many of the members have collaborated previously.

In August 2017, a Fintech Unit was formally established within the SARB, with the intention of exploring the implications of fintech innovation for the SARB and all the financial services in South Africa. Following discussions with members of the SAFBC, the South African Multiple Option Settlement (SAMOS) User Group (SUG) and the Monetary Authority of Singapore (MAS) late in 2017, the SARB embarked on its own DLT trial: Project Khokha.

¹ Schwab, K. and Sala-i-Martin, X. (2017)

² Swiss Re Institute (2017)

³ Macknight, J. (2017)

SOUTH AFRICA'S FINANCIAL SERVICES INDUSTRY IS INNOVATIVE AND SOPHISTICATED, WITH THE SOUNDNESS OF BANKS **RANKED 37TH** OUT OF 137 COUNTRIES BY THE WORLD ECONOMIC FORUM'S GLOBAL COMPETITIVENESS REPORT.



THE SOUTH AFRICAN RESERVE BANK AND THE APPROACH TO FINTECH

The SARB is responsible for price and financial stability.

As the central bank of the Republic of South Africa, the SARB's mandate is to achieve and maintain price stability in the interest of balanced and sustainable economic growth in South Africa. The achievement of price stability is quantified by the setting of an inflation target by government that serves as a yardstick against which price stability is measured. The achievement of price stability is underpinned by the stability of the wider financial system. The SARB is thus entrusted with the overarching monetary policy goal of maintaining price stability within a flexible inflation-targeting framework⁴. The disruptive potential of and the potential risk from fintech is widely recognised and, in common with many other central banks, the SARB has started to put structures in place to monitor and consider appropriate policy frameworks and responses to the changes that may result from fintech.

⁴ The enabling framework for the operations of the South African Reserve Bank is provided by sections 223 to 225 of the Constitution of the Republic of South Africa (1996), the South African Reserve Bank Act 90 of 1989, as amended (SARB Act), and the Regulations framed in terms of this SARB Act (<https://www.resbank.co.za/AboutUs/Mandate/Pages/Mandate-Home.aspx>).

The SARB has committed publically to three areas of exploration.

The SARB has recently established a Fintech Unit to assess the emergence of fintech in a structured and organised manner, and to consider its regulatory and strategic implications. On 13 February 2018, the Unit publicly announced its three initial initiatives, which include this project.

Figure 1: Fintech Unit initiatives



Project Khokha is a collaborative initiative using distributed ledger technology.

Project Khokha is seen as an initiative in collaborating for innovation, therefore both the process as well as the outcome of the project contribute to the SARB's goals. The decision was made to assess the use case for DLT in wholesale payments and interbank settlement and thus build on and extend the work done in other parts of the world. The SARB engaged ConsenSys as the technical partner on the project and worked with a consortium of banks made up of Absa, Capitec, Discovery Bank, FirstRand, Investec, Nedbank and Standard Bank.

PROJECT KHOKHA
IS SEEN AS AN
INITIATIVE IN
COLLABORATING
FOR INNOVATION.

THE SOUTH AFRICAN REAL-TIME GROSS SETTLEMENT SYSTEM

The SAMOS system was introduced on 9 March 1998. The system is an automated interbank settlement system provided by the SARB for banks to settle their obligations on an immediate real-time basis in central-bank money. SAMOS operates on a 24/7/365 basis and was designed to settle all large-value, wholesale and securities transactions as well as the interbank obligations resulting from retail payment clearing houses. Large-value payments are settled individually on a real-time gross settlement (RTGS) basis, while retail payments are settled as a batch on a deferred basis. The SAMOS system is linked to the various participant banks, clearing systems and operators. Participants in the settlement network include the SARB, commercial banks, registered branches of foreign institutions, mutual banks, cooperative banks and designated settlement systems. SAMOS is the only RTGS system in South Africa and is owned and managed by the SARB.

The NPSD needed an effective and safe mechanism for the exchange of money between transacting parties. The goal was to enable the safe and efficient transfer of money in the financial system. The RTGS system – initially developed between 1996 and 1998 by the SARB, the banking industry and local technology providers – aimed to provide a sophisticated, state-of-the-art system for high-value interbank transactions. SAMOS was officially introduced in March 1998 and brought domestic interbank settlement practices in line with international best practice, signalling a new era for payment practices in South Africa. SAMOS has been instrumental in reducing the systemic risk in South Africa and provides the facilities for banks to finally and irrevocably settle obligations within the South African payment system, settling more than 90% of the value of all interbank obligations directly.

History and purpose.

The SARB is ultimately responsible for the national payment system (NPS), with its National Payment System Department (NPSD) as the responsible management and oversight body. The NPSD supports the mission of the SARB by ensuring the overall effectiveness and integrity of the payment system.

The characteristics of a properly functioning payments system are:

- Enhance the stability of the financial system;
- Reduce transaction costs in the economy;
- Promote efficient use of financial resources;
- Improve financial market liquidity;
- Facilitate the conduct of monetary policy.

Settlement in South Africa.

The National Payment System Act 78 of 1998, as amended (NPS Act) has established that settlement may only be carried out using cash or through entries in the book of the SARB. This requires that the settlement system participants have an account at the SARB through which interbank settlement obligations are settled. These settlements are carried out on a pre-funded basis through the SAMOS system and help to reduce the risk of settlement failure by a participant in the payment system. An instruction to transfer funds through SAMOS will only be carried out if the bank issuing the transfer instruction has sufficient funds in its settlement account. This reduces risk in that banks are unable to build up exposure to each other in the SAMOS system. The funds transferred are final and irrevocable. The SAMOS system provides differing levels of functionality in order to meet both real-time and deferred settlement requirements.

The SARB has defined settlement as 'the final and irrevocable discharge of an obligation of one bank in favour of another bank, in central-bank assets or money'. Central-bank assets or money include banknotes, coins and credit balances held by the central bank.

THE RATIONALE AND KEY OBJECTIVES OF PROJECT KHOKHA

The project has built on previous central-bank projects.

Project Khokha aimed to build on the initiatives previously undertaken by global peers and to test some concepts in a South African context. The process of doing the work is almost as important as the output, in that it provides an opportunity to broaden the DLT skills base in the South African banking industry and also presents an opportunity to explore the type of collaborative innovation that is expected to become more common. The technology inherently requires the sharing of data between multiple parties, so this collaboration is critical in order to provide and maximise the benefits of DLT. The scope of the project was designed for the specific requirements of the local market and to complement previous work done in this area.

The scope of the project was to create a distributed ledger between participating banks for a wholesale payment system, backed by central-bank deposits, allowing participating banks to pledge, redeem and track balances of the tokenised rand on the ledger. The project aimed to assess the performance, scalability, privacy, resilience and finality of a DLT solution under conditions that were as realistic as possible, in that each bank was responsible for its own node and these nodes were distributed.

There are some specific areas of exploration.

Project Khokha was designed to simulate a ‘real-world’ trial of a DLT-based wholesale clearing and settlements system. The project therefore focuses on some practical aspects of using DLT to give practical experience to a broad group of participants. Significant time was invested early on in training participants to ensure that they were all well acquainted with the technology prior to mobilisation. The project testing schedule was designed to reflect a realistic level of volume while maintaining confidentiality of balances and transaction amounts. In addition, the banks were each responsible for setting up and managing their own node and private keys, which led to a mixture of on-premise and cloud-based nodes.

**PROJECT KHOKHA
WAS DESIGNED
TO SIMULATE A
REAL-WORLD TRIAL
OF DISTRIBUTED
LEDGER TECHNOLOGY.**



02

THE STRUCTURE OF **PROJECT KHOKHA**

PARTICIPANTS AND ROLES

The SARB assembled a consortium of South African banks.

In planning the project, a decision was taken that participants should be drawn from the SAMOS community as they are best positioned to consider the impact of DLT on wholesale payments. The consortium of banks involved in the project consisted of seven volunteers from the SUG, including the five biggest retail banks in the country. This process was also assisted by the industry's history of collaboration with the SARB. The SARB was the project owner, responsible for crafting the strategy and approach of Project Khokha. Absa and Standard Bank, referred to as the 'partner banks', were the first banks to be involved in the early design stages of the project. The remaining banks (Capitec, Discovery Bank, FirstRand, Investec and Nedbank) joined the distributed ledger subsequently, and all implemented their own nodes, having had varying degrees of technical involvement.

ConsenSys was brought in to build on the experience with the Monetary Authority of Singapore.

ConsenSys is a venture production studio, building decentralised applications and various developer and end-user tools for blockchain ecosystems, primarily focused on Ethereum. ConsenSys played a key role in Project Ubin, which had been initiated by the MAS. Its role as architect and technology partner on the Quorum work stream on Ubin's Phase 2 was leveraged in Project Khokha and enabled this project to be established quickly. ConsenSys's main responsibilities were to develop the solution and help the participants with any technical challenges.

PricewaterhouseCoopers Inc. was contracted to support Project Khokha.

PwC was contracted to support on Project Khokha, with a team of specialists across payments and blockchain not only to report on the findings of this PoC, but also to note the insights gained and bring together global research as well as possible future scenarios and impacts for consideration. PwC's team reviewed the project documentation and other central-bank white papers and conducted interviews with the key stakeholders across all the participants in the SARB, banks and ConsenSys.



OBJECTIVES, SCOPE AND EXCLUSIONS

The current wholesale payments system in South Africa is the South African Multiple Option Settlement system.

South Africa's RTGS system is SAMOS, as noted above. Large-value payments⁵ and payments requiring immediate finality are settled one by one on an RTGS basis, while retail payments are settled on a deferred nett settlement basis. SAMOS has 30 participants, including the SARB, commercial banks (registered under the Banks Act 94 of 1990), mutual banks (registered under the Mutual Banks Act 124 of 1993), cooperative banks (registered under the Cooperative Banks Act 40 of 2007) and designated settlement systems. The system operates 24 hours a day, 7 days a week, 365 days a year. The Society for Worldwide Interbank Financial Telecommunication (SWIFT) messaging network is used to transfer payment messages between the SAMOS system and its participants.

It should be noted that while much of Project Khokha was focused on understanding how DLT may be used in an RTGS environment, particularly around whether the performance is achievable under realistic conditions, the SARB does not intend to replace SAMOS with DLT at this juncture, nor does it intend to halt the current RTGS modernisation process. One objective of Project Khokha is to provide a better understanding of how SAMOS would integrate with a DLT system. The intention is not to consider changing the approach with the SAMOS replacement, but to provide input to that project.

⁵ Although all payments over R5 million must be settled directly through the South African Multiple Option Settlement (SAMOS) system, the majority of payments settled in SAMOS (by volume) is below R5 million, reflecting the wide definition of 'large-value payments'. The characteristics of such payments include specific markets (such as bonds, money and equity markets) but also wholesale business-to-business payments.

Figure 2: Project Khokha overview



OBJECTIVES

Develop and test a PoC that is built on a private (permissioned) Ethereum blockchain (Quorum) to create the tokenisation of the South African rand and to demonstrate that it can be transferred between commercial banks on a blockchain.

Extend the global body of knowledge on DLT-based RTGS systems.



GOALS

Create a distributed ledger between participating banks for a domestic payment system backed by central-bank deposits allowing participating banks to pledge, redeem and track balances on the distributed ledger.

Establish how DLT compares to the capabilities of existing technologies used by the SARB and participating commercial banks.



MEASURABLE GOALS

Transaction speed

70 000 transactions over 2 hours.

Block Propagation Times

95% propagation in <1s.
99% propagation in <2s.

Confidentiality

Fully confidential transactions ensuring appropriate privacy.



FOCUS

Performance

Ability to process required transactions.

Adherence to principles

Adherence to relevant PFMLs.

RTGS participation

Increased stakeholder participation.

Non-technical implications

Understand legal, regulatory, accounting implications.

ONE OBJECTIVE OF PROJECT KHOKHA IS TO PROVIDE A BETTER UNDERSTANDING OF HOW THE SOUTH AFRICAN MULTIPLE OPTION SETTLEMENT SYSTEM WOULD INTEGRATE WITH A **DISTRIBUTED LEDGER TECHNOLOGY SYSTEM.**

A key goal of the project was to establish a collaborative approach with participating banks.

Besides the technical objectives described below, a key objective of the project was to encourage collaboration within the industry and within the SARB in order to establish a framework of cooperation that would support future projects. Such collaboration is key, as the technology inherently requires the sharing of data. More importantly, the value of a distributed ledger benefits greatly from the network effect. As a PoC, there was no integration with the existing payment system or with banks' production systems. In addition, liquidity saving mechanisms (LSMs), queue handling and gridlock resolution were not tested.

Specific, measurable goals were set around the platform performance.

In order to establish how the DLT system compares to existing technologies and whether it can meet the performance requirements of a production environment, the following performance goals were set. These goals can be regarded as some of the requirements for a DLT solution to be able to operate in the South African context.

TRANSACTION PERFORMANCE

The DLT system should aim to achieve:

- 70 000 transactions per day, with a stretch goal to increase this to 200 000 transactions per day – to simulate typical SAMOS requirements and potential future growth; and
- 70 000 transactions in two hours – designed to simulate a requirement to process a full day's transactions in two hours in the event of a day's loss of processing.

PROPAGATION TIMES FOR TRANSACTIONS

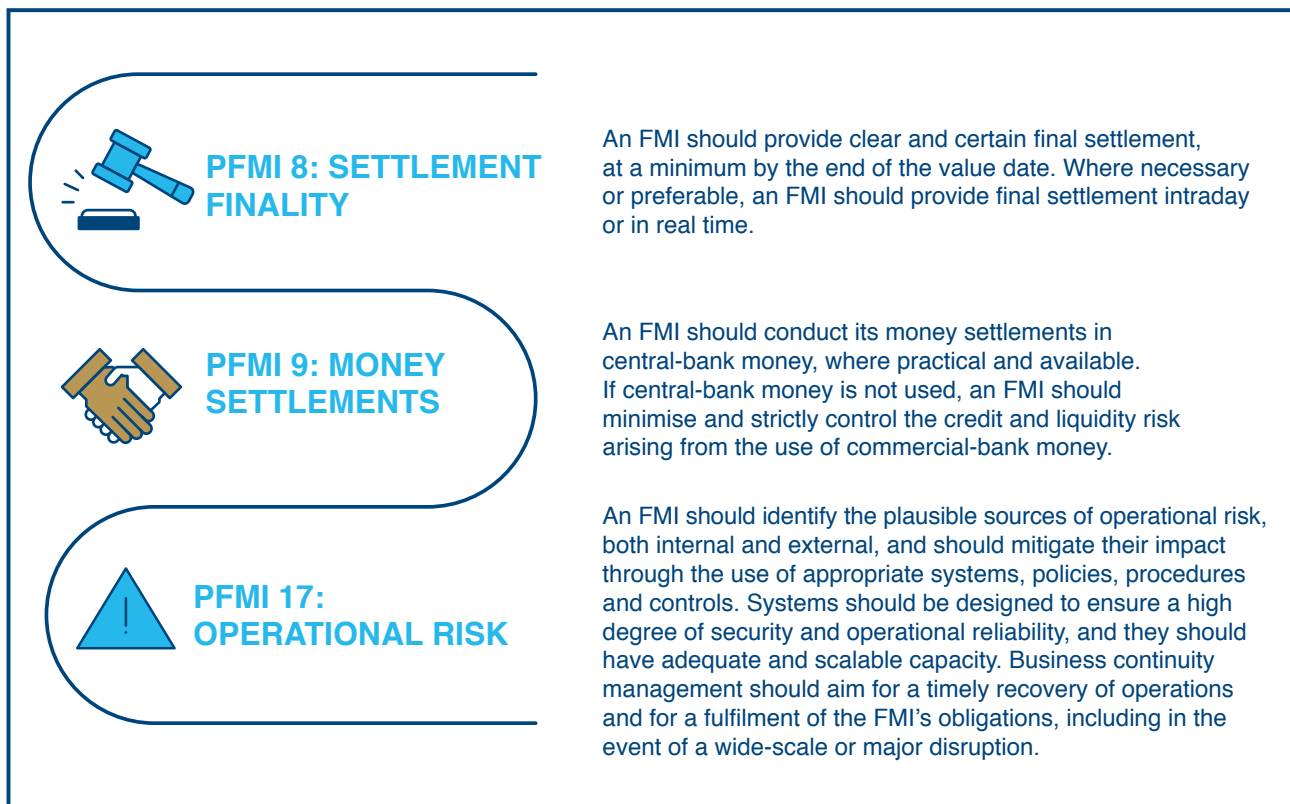
In order to enable transaction performance, the vast majority of propagation times for transactions must be under one second. More specifically, there should be a 95% confidence level that blocks (containing transactions) are propagated through the entire network in less than one second and a 99% confidence level that blocks are propagated through the entire network within two seconds.

CONFIDENTIALITY

Transaction details must be fully confidential, i.e. Bank C should not be able to see the details of a transaction between Bank A and Bank B, thus ensuring the appropriate level of privacy to all participants. However, the SARB, as regulator, should be able to see all the transactions.

In addition to the targets noted above, there is a requirement to adhere to the globally recognised principles encapsulated in the Principles for Financial Market Infrastructures (PFMIs). This project aims to specifically highlight the ability of a distributed ledger to adhere to the following:

- PFMI 8: Settlement finality
- PFMI 9: Money settlements
- PFMI 17: Operational risk

Figure 3: The Principles for Financial Market Infrastructures evaluated for Project Khokha























The above-mentioned PFMI were the only principles considered for this PoC since they specifically relate to the transaction element of an RTGS system. Principles that apply to production-level financial market infrastructures (FMIs), such as those relating to governance and legal aspects, were excluded.








Project Khokha has added to the central-bank distributed ledger technology knowledge base.

Table 2⁶ highlights the scopes of work publicly released by other central banks. The typical approach has been to start by building a relatively simple PoC (usually in wholesale payments) and then to expand functionality in further phases – typically looking at performance testing, privacy and LSMs. Some projects have also made comparisons of different technology platforms for the same use case.

⁶ The icon key appears below the table.

Table 2: Scope of work of published central-bank initiatives

Project	Phase	Scope	Themes
Brazil (Central Bank of Brazil)	1	Identify use cases for the central bank using distributed ledger technology (DLT). Identify an appropriate DLT platform and build a minimal proof of concept (PoC).	 
	2	Analyse competing blockchain platforms using the selected use case as a benchmark. Address the privacy issues identified in the previous phase.	  
Canada (Project Jasper)	1	Build a wholesale interbank settlement capability on an Ethereum DLT platform, and demonstrate its ability to exchange value in the form of a central-bank-issued digital settlement asset. Evaluate the performance of this platform, using the Principles for Financial Market Infrastructures (PFMIs) as a framework.	
	2	Evaluate the scalability and flexibility of DLT by moving to an alternative technology platform (Corda) and by continuing to build in more of the functionality observed in today's interbank settlement solutions, most notably the liquidity savings mechanisms (LSMs). Provide a data-driven simulation exercise with operational data sets to evaluate the platform and LSM performance.	  
Europe and Japan (Project Stella)	1	Build a real-time gross settlement (RTGS) system on DLT, including LSM functionalities. Assess whether the specific functionalities of existing payment systems could be safely and efficiently run in a DLT application, focusing on hands-on testing.	  
	2	Explore ways in which delivery versus payment (DvP) can be conceptually designed and technically achieved in a DLT environment.	 
Singapore (Project Ubin)	1	Build a PoC for domestic payments for interbank obligations on a distributed ledger, denominated in balances backed by a central bank. Identify the non-technical implications of moving this into a production environment.	
	2	Assess the potential implications of deploying DLT for specific RTGS functionalities by focusing on LSMs. Understand how RTGS privacy can be ensured on DLT. Compare alternative DLT platforms.	   
South Africa (Project Khokha)		Build an RTGS system on DLT, with tokens backed by funds held in the central bank. Investigate privacy solutions while simultaneously achieving the required throughput. Perform tests in a 'realistic' environment by having each participant run their own node under a variety of deployment models in different locations.	  

	Key
	RTGS <i>The project involved the creation of an RTGS PoC on DLT.</i>
	RTGS+ <i>The project expanded on the original RTGS PoC, typically looking to implement more functionality.</i>
	LSMs <i>The project looked at implementing an LSM within the PoC.</i>
	Privacy <i>The project assessed different methods to ensure privacy of transactions.</i>
	Technology comparison <i>The project assessed different DLT platforms in order to determine suitability for the use case.</i>
	Distributed nodes <i>The project assessed the effect of distributed nodes on the performance of the system.</i>
	Delivery versus payment <i>The project implemented a DLT PoC for delivery versus payment (DvP).</i>

Project Khokha is therefore similar to other experiments in that its first phase is aimed at building a wholesale payment system PoC. However, it is important to note that it brings new ways to achieve privacy while meeting required transaction volumes – balancing a key trade-off in DLT design. Each participant ran a full node with different deployment models and in different locations, creating a relatively realistic testing environment.

TECHNOLOGY SELECTION

The selection of Quorum was pragmatic to enable the project to mobilise rapidly.

The purpose of the project was not to compare DLT technologies, but to create a PoC that explored the applicability of DLT in the RTGS space while building on the global body of knowledge. A number of other central banks have used Quorum, and as such – for the purpose of this study – it was a pragmatic choice, particularly given ConsenSys's technical skills and experience from working on Project Ubin. The choice of ConsenSys and Quorum was brought about through collaborative discussion with, for instance, MAS, as well as through the monitoring of what other central banks have done and collaborative learning through participation in international bodies.

The key features of Quorum are relevant to this work.

While the purpose of the project was not to provide an evaluation of the strengths and/or weaknesses of Quorum, nor to provide a comparison with other enterprise DLTs, an overview of the key features of Quorum is provided here for completeness.

Quorum is an enterprise-focused version of Ethereum that was co-developed by JPMorgan in partnership with EthLab. It is open-source and, like Ethereum, licensed according to the GNU Operating System general public licence (GPL) and lesser general public licence (LGPL) – which was done to facilitate and encourage collaboration and innovation. Quorum is built upon Go Ethereum, the base code for one of the most commonly used clients on the Ethereum blockchain. It operates very similarly to Ethereum, but with four major distinctions: network and peer permissions management, increased transaction and contract privacy, voting-based consensus mechanisms and higher performance. There is intent for it to update as Ethereum evolves, as they are tightly coupled due to the base code being the same foundation.

At a high level, Quorum can be considered to be a version of Ethereum for private consortiums, with three key pillars⁷.

Figure 4: The three key pillars of Quorum



PRIVACY & TRANSPARENCY

Both transaction-level privacy and network-wide transparency are supported. It is customisable to requirements. For this project, Whisper peer-to-peer messaging, Pedersen commitments and range proofs were the mechanisms used to enable privacy.



PERFORMANCE & THROUGHPUT

Quorum has been designed to achieve realistic throughputs associated with the financial services industry. This is partially enabled by including options for consensus mechanisms. The mechanism used here was Istanbul Byzantine Fault Tolerance (IBFT).



PERMISSION & GOVERNANCE

Quorum supports blockchains around permissioned groups of participants, with transaction validation and block creation distributed throughout the network.

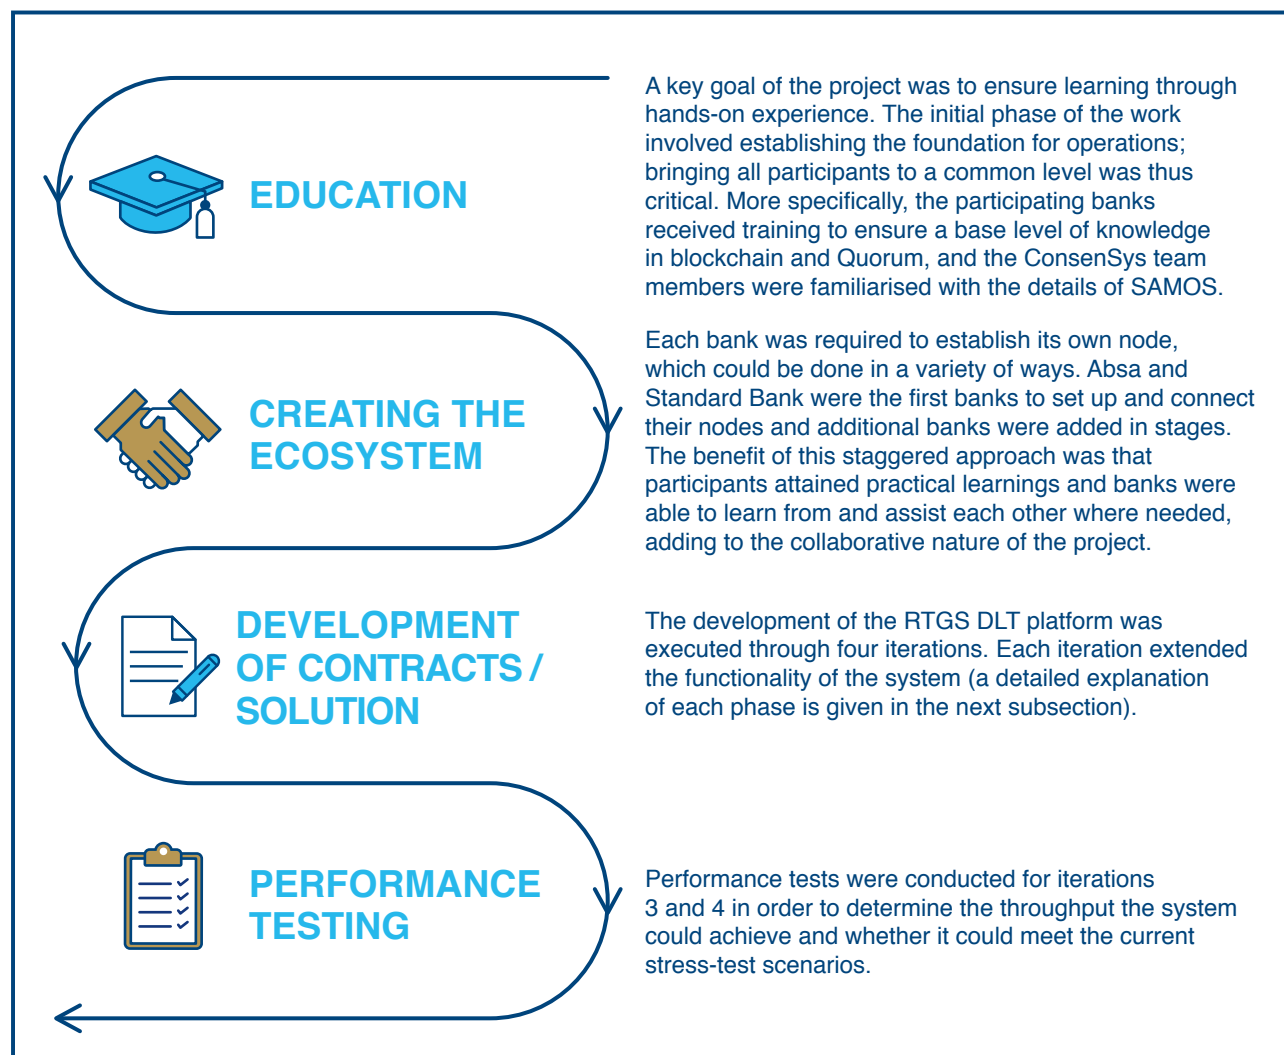
⁷ JPMorgan (2018)

THE APPROACH

Project Khokha comprised four broad phases.

The project ran over a 14-week period, and a collaborative, agile process was followed in order to meet the project deadlines. The project was mobilised with fundamentals training and a joint application design (JAD) session during which the scope, objectives and timelines were established. Clear project governance was established, including brief coordination meetings that were held three times a week, with an in-depth technical meeting occurring weekly as part of the design and troubleshooting.

Figure 5: The four phases of Project Khokha



The project followed an agile approach, building the proof of concept in four iterations.

The solution was developed in four iterations, with each iteration building on the functionality of the previous one.

- At a high level, iteration 1 involved enabling two banks to transfer tokens, as well as a minting function (restricted to the SARB node).
- Iteration 2 was closer to the current SAMOS operation, with the SARB approving the payment but all transactions visible to the whole network.
- Iteration 3 involved shielding the amounts of transactions and balances using Pedersen commitments, with the SARB given access to the openings of the commitment, allowing it to verify and approve the payment. Encrypted peer-to-peer ‘Whisper channels’ were used at setup to share encryption keys between pairs of nodes and with the SARB node.
- The final iteration increased the resilience of the system as this version used range proofs as well as the Pedersen commitments. The range proofs allow the participant nodes to verify that the amount being sent and the balance after sending are both positive, without revealing the balances or amounts. The SARB node still has the openings to the Pedersen commitments and therefore it has full visibility of transactions, but because the other nodes can verify transactions via the range proofs, the SARB node no longer needs to perform this role.

All iterations ran on a network using IBFT as the consensus mechanism.

Additionally, under iterations 3 and 4, Know Your Customer (KYC) and anti-money laundering (AML) payment information could be shared between Bank A, Bank B and the Financial Intelligence Centre (FIC) by enabling them to decrypt the relevant parts of the transaction. This would require separating the encryption of customer data from the transaction details so that the SARB could monitor liquidity and the FIC could monitor KYC and AML data.

More details on the iterations are given in Appendices.

A zero-knowledge proof (ZKP) is a protocol via which one party can demonstrate to another that they know a certain value without disclosing that value.

A range proof is a type of ZKP that enables one party to prove that a secret value (such as one contained in a Pedersen commitment) is within a certain range in this case, only positive numbers.

A Pedersen commitment is a cryptographic primitive that allows a party to commit to a secret numerical value, while maintaining the ability to do mathematical operations on the secret value. Commitment schemes are designed so that a party cannot change the value after they have committed to it – in other words commitment schemes are binding.

Istanbul Byzantine Fault Tolerance (IBFT) is a consensus mechanism with a manageable validator set. Validators propose and vote on new blocks, and as long as no more than a third of them are faulty, instant finality is achieved which means that no forks or invalid blocks can occur.

HYPOTHESES TESTED

The technical goals are expressed formally as hypotheses.

The project set out to test a number of key hypotheses which are depicted in figure 6:

Figure 6: The key hypotheses of Project Khokha



03

RESULTS

RESULTS

Project success factors.

A number of factors contributed to the success of this project, including:






- The establishment of communities of learning within the regulators and industry created conducive learning environments, reflecting a willingness to learn and explore in South African financial services.
- The SARB's interest to learn more about DLT, not only as a regulator but also as a participant in the process, was important.
- Open communication and a culture of collaboration, facilitated by previous banking relationships and involvement in the SAFBC and SUG, were vital. Participants reported that all the parties strived towards ensuring that the project was a success and sought to help each other and learn.
- Collaboration was further enabled by the use of agile processes and supporting communication tools.
- Rapid value realisation and being able to see the results of what had been achieved within the 14-week period reinforced the spirit of collaboration and momentum within the project.
- Key stakeholder involvement aided the achievement of success as the right stakeholders were involved at the right time across all aspects of the project. Significant effort was made at the beginning of the project to ensure that the appropriate people (a mixture of business and technical staff) were nominated by each participant.
- Training and alignment on mobilisation enabled each participant to start with a common body of knowledge.
- The involvement of the technical partner ConsenSys, including team members with experience from Project Ubin, was a key factor in mobilising quickly, leveraging the insights gained.
- The phased execution (banks were brought onto the network sequentially) and the on-boarding approach for participating banks helped to overcome common challenges and to gain further insight from each other's experiences.
- The report development process enabled independent reflection on the project objectives and outcomes within a global context.



THE OUTCOME OF HYPOTHESIS TESTS

Table 3 depicts the hypotheses tested in this PoC and the produced results.

Table 3: Outcomes of hypotheses tested


Hypothesis	Outcome	See page
 Real-time gross settlement transactions can be executed using standard payment message formats (leveraging the ISO 20022 format).	Proven	p.38
 Messages can be processed at sufficient scale in line with the current system processing times.	Proven	p.41
 Blocks should be propagated within one second to a 95% confidence level and within two seconds to a 99% confidence level.	Proven	p.38
 The confidentiality and privacy of transactions between commercial banks is maintained.	Proven	p.44
 The visibility of the system for the SARB is sufficient for oversight and operational management.	Proven	p.38

The PoC was designed in compliance with ISO 20022, demonstrating that the first hypothesis above is correct. With reference to the last hypothesis above, the focus of the PoC was not on operational management and NPS oversight was not part of the PoC. Regardless, the SARB node gave the SARB full visibility of all the transactions, so the visibility of the system may be considered sufficient.

The final results were positive, indicating that the Quorum platform can deliver the performance required, matching and even exceeding the current performance criteria. The current testing, without full optimisation, indicates twice the current SAMOS system’s performance for transactions per second (tx/s), with confidentiality and

settlement finality being enabled as part of the design. System latency was within the targeted times for the writing of blocks, with 99% of the blocks propagated through the network in 1s and 100% in 1.25s, as can be seen in Figures 7 and 8.

Figures 7 and 8: Block propagation times



The Project Khokha blockchain solution is handling the transactional element of RTGS and integrating back into the testing versions of existing systems to complete full RTGS operational requirements. More details on the specific results per category are provided above.

TECHNICAL FEASIBILITY

The Project Khokha solution has shown that it is technically feasible to process wholesale payments between participants with full visibility to the central bank from an operational perspective.

A file size of 70 000 transactions requires a rate of 9.72 tx/s in order to complete processing within two hours. Full details of the performance testing are discussed in section 5.4, but both the third and the fourth iteration achieved rates that satisfied the design requirement. The final iteration – using a combination of IBFT, range proofs and Pedersen commitments – removed the need for the SARB to approve each transaction, and allowed processing at 12.86 tx/s, thus still achieving the design requirement despite the additional computation demands. Note that the node machine specifications were upgraded for the final test, so iterations 3 and 4 are not directly comparable.

The technical feasibility of integrating with other systems within the SARB and the commercial banks has not been tested. Nor has some of the existing RTGS system's functionality, including liquidity management, been tested. These technical elements were not considered to be required as part of the PoC, but they would require investigation and further PoCs to understand the complete end-to-end feasibility of the solution. The current design, when taken to a production-ready state by integrating with participants' back-office systems and putting operational controls in place, could offer a potentially robust backup solution to the RTGS system for transaction message processing.



SCALING

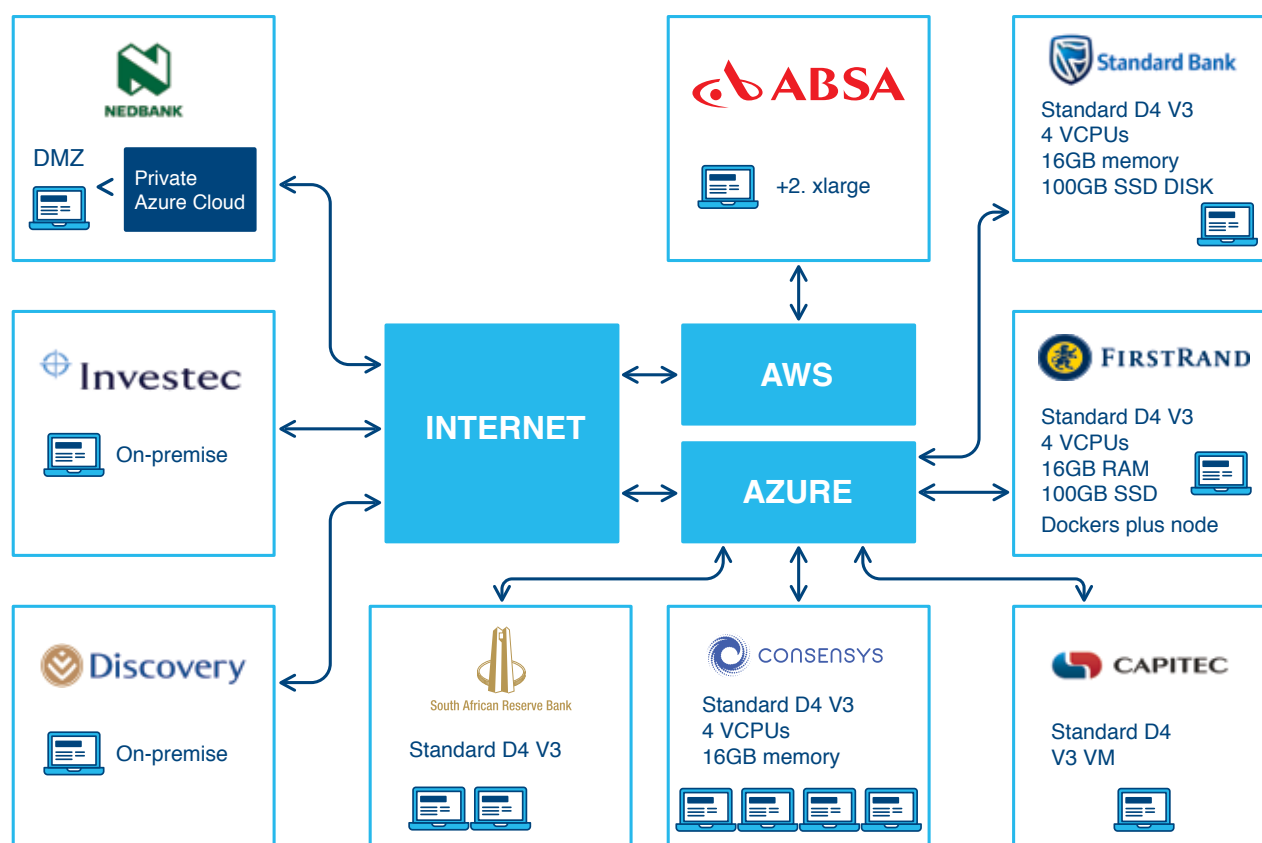
Scalability depends on a number of factors, including the confidentiality and consensus mechanisms dealt with below, as well as the node architecture. Some of the issues relating to node architecture are discussed below.

NODE ARCHITECTURE

The Project Khokha design document detailed the minimum specification for the network nodes. The design was for an Ubuntu server, which could be a physical machine, an on-premise virtual machine, or a virtual machine in the cloud. All three of these permutations were employed on the project (refer to Figure 9 for more detail), which meant that the PoC was conducted in an environment that attempted to emulate a real-world scenario closely.

Even though the Quorum node had not yet been deployed as a Docker container, one of the participants used Docker to build a containerised solution, utilising components already placed in Docker containers. The design document described a baseline machine based on a simple Microsoft Azure Virtual Machine. There is considerable scope to upgrade the architecture, including the number of cores and memory as well as the virtual storage of the virtual machine. An additional enhancement could be the deployment of containers for services like the user interface or dedicated physical Linux hardware.

Figure 9: Project Khokha node architecture



Performance tests

The details of the performance tests for iterations 3 and 4 are shown below. Note that the number of nodes shown is for the participant banks, and that it excludes the ConsenSys and SARB nodes shown above, which were on the network for all the transactions. The specifications above were upgraded for the second and third run in iteration 4.

ITERATION 3 RESULTS

In iteration 3 the transactions are confidential and the SARB node approves all the transactions after checks.

Table 4: Results for the testing of iteration 3

Iteration 3			
	Test run 1	Test run 2	Test run 3
Test type	First benchmark test on iteration 3	Second benchmark test on iteration 3	Third benchmark test on iteration 3
Date	3 April 2018	6 April 2018	12 April 2018
Nodes (banks)	4	5	7
Total transactions	70 000	90 000	90 000
Duration	1h 14m 21s	1h 15m 51s	1h 16m 00s
Processing speed	15.69 tx/s	19.78 tx/s	19.74 tx/s
Node specifications	Four cores, 16GB RAM	Four cores, 16GB RAM	Four cores, 16GB RAM

ITERATION 4 RESULTS

Iteration 4 performed fully distributed, confidential transactions. The SARB had no operational involvement in approving the transactions, although it still had sight of all the transactions. The practical import of this is that the SARB node could be down and the banks would still be able to pay each other.

Each transaction needed to generate range proofs to show that the amount being transferred was positive and that the balance of the sending bank was still positive after the transaction. These proofs have to be verified by all the nodes on the network, which introduces more computation per transaction. The iteration 4 tests were therefore run on a separate network because the specification of the virtual machines on the blockchain needed to be improved to accommodate the additional processing required. For comparative purposes, a test was run using the original lower-specification machines before the specification of the machines was improved for the two benchmark tests in iteration 4. Between the first and second benchmark tests for this iteration, the team was able to identify some bottlenecks and optimise the code as well as the blockchain configuration for the final test run.

Table 5: Results for the testing of iteration 4

Iteration 4			
	Test run 1	Test run 2	Test run 3
Test type	Comparative benchmark test on iteration 4	First benchmark test on iteration 4	Second benchmark test on iteration 4
Date	20 April 2018	20 April 2018	20 April 2018
Nodes (banks)	6	6	6
Total transactions	10 000	70 000	70 000
Duration	38m 23s	1h 54m 46s	1h 30m 41s
Processing speed	4.34 tx/s	10.71 tx/s	12 tx/s
Node specifications	4 cores, 16GB RAM	16 cores, 64GB RAM	16 cores, 64GB RAM

THE CONSENSUS MECHANISM

This is thought to be the first time that a PoC based on Quorum has used IBFT. The critical features of this implementation are its performance and scalability.

Overview of Istanbul Byzantine Fault Tolerance.

The IBFT algorithm is a state machine replication algorithm that follows these steps:

- The proposer multicasts the block proposal to the validators;
- The validators agree on the block and broadcast their decision to others; and
- Each validator waits for $2F+1$ commits from different validators with the same result before inserting the block into the blockchain.

The algorithm maintains the order of transactions and ensures data consistency and integrity. A network of N validators can tolerate F faulty nodes, where $N = 3F + 1$. Practically, this equates to being able to maintain integrity should up to 33% of nodes act maliciously. The impact of faulty nodes was not specifically tested as part of this project.

CONFIDENTIALITY

The Project Khokha Quorum solution uses Whisper for private messaging, Pedersen commitments and range proofs (in the final iteration). This ensures that the solution provides robust confidentiality while enabling the transaction throughput required.

Whisper messaging is used at start-up.

Whisper is a secure peer-to-peer secure messaging system built into Ethereum. At the start-up of the network, each bank exchanges encryption keys with every other bank via this channel and these keys are then used in transactions between those banks. The SARB has a copy of all the keys so that it can view all the transactions, i.e. the Whisper messages are shared with the SARB.

Pedersen commitments enable confidentiality.

The implementation of Pedersen commitments in the Project Khokha solution has shown to be effective in ensuring full confidentiality of transactions. This form of cryptographic algorithm was initially proposed in 1991 and is part of the technique used in the crypto-currency monero to ensure privacy. Pedersen commitments allow the provider of the proof – in this case the sender of the payment – to commit to a value without being able to change it once the commitment is provided. Pedersen commitments are perfectly hiding and computationally binding under the discrete logarithm assumption. The specific implementation in Project Khokha is the first time that this has been used with Quorum, and it has delivered significant performance improvements over other confidentiality mechanisms.

THE IMPLEMENTATION
OF PEDERSEN COMMITMENTS
IN THE PROJECT KHOKHA
SOLUTION HAS SHOWN
**TO BE EFFECTIVE
IN ENSURING FULL
CONFIDENTIALITY
OF TRANSACTIONS.**

04

COMPARISON TO SIMILAR PROJECTS

Several central banks have published the results of their DLT experiments. In this section, we review the Central Bank of Brazil's PoC as well as those for Projects Jasper Stella and Ubin.

A number of themes are apparent:

- The projects rapidly deliver PoCs, often with collaboration across an industry.
- The projects are building up a global body of knowledge around DLT and its uses for central banks.
- Specific areas of focus include building knowledge within scalability, privacy and resilience in the RTGS space.

Central banks are using agile delivery methods to learn about and implement distributed ledger technology proofs of concept.

Central banks have taken a cautious and prudent approach to DLT, assessing key opportunities and risks and implementing PoCs to validate assumptions and gain insights regarding the technology. Typically, these projects have been run by small teams within the central bank, looking at spearheading innovation, or at least the exploration of DLT, in the organisation and assessing the impact on the industry. These projects have been delivered within relatively short timelines, showing value quickly by leveraging some form of agile delivery methodology.

Finally, these projects have often involved a variety of stakeholders and have incorporated relevant industry participation, where applicable. A shared ledger inherently requires active participation by the relevant players, so collaboration is a key ingredient for success in these initiatives.

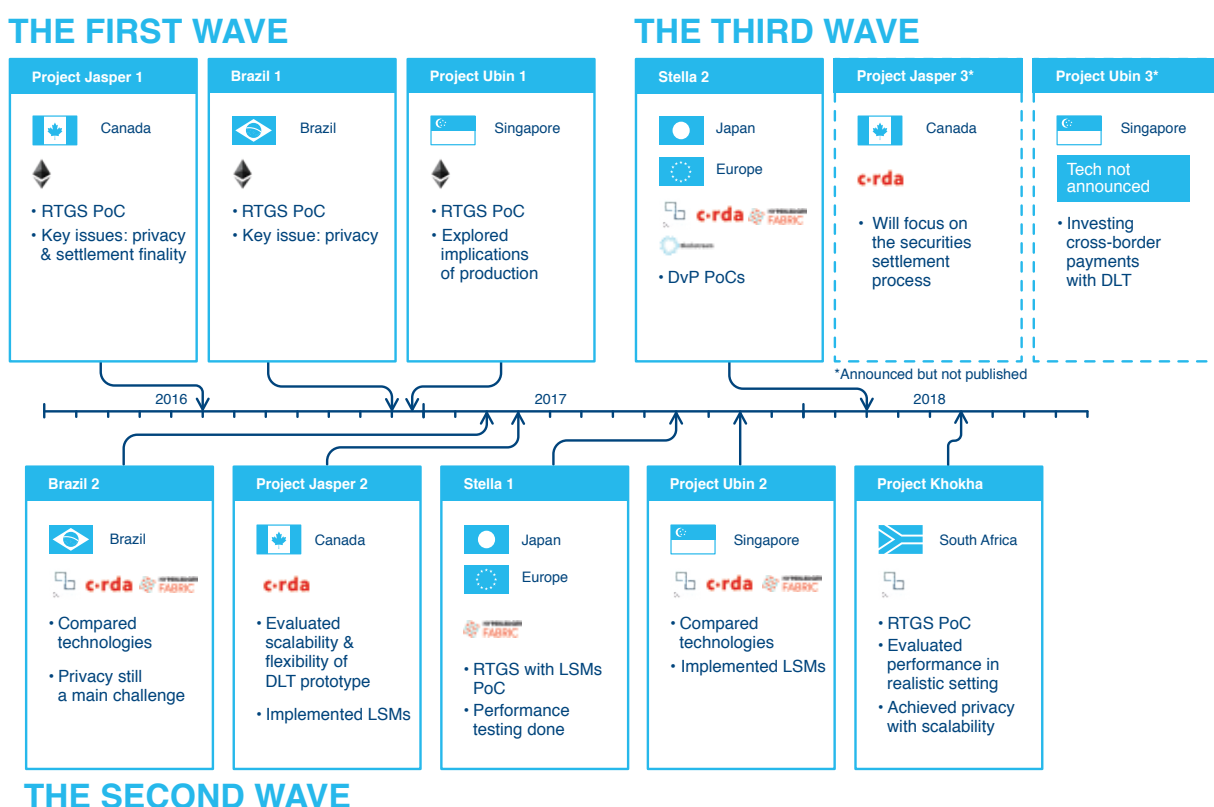
In the above aspects, Project Khokha has been managed in a similar way to its global peers.

The global body of knowledge has grown rapidly, with projects building upon each other.

Figure 10 summarises specific projects and highlights the key technologies and focus areas as piloted through the various PoCs published to date. We have grouped the projects into different waves of central-bank activity.



Figure 10: Three waves of central-bank distributed ledger technology experiments



The first wave

The first wave represents the creation of relatively simple PoCs that have all used Ethereum-based systems. Many of the findings from these projects indicated that there could be benefits to DLT, particularly around the resilience of the systems, but that many of the features of Ethereum, such as proof of work and the visibility of transactions by default, are ill-suited for private, enterprise blockchains. This is especially the case in the wholesale settlement scenario where transaction speed, finality of settlement and confidentiality are critical.

The second wave

Following the first wave of DLT projects, a second wave emerged that sought to build on the insights from the first, often leveraging different enterprise platforms. The projects in the second wave have made progress in terms of scalability, privacy and many of the key issues identified previously. Essentially, this work has taken a more thorough look at how this technology could be implemented in a production setting.

The third wave

Currently, a third wave of projects is emerging, with central banks moving beyond the replication of existing systems and looking at how DLT can enable new models. Phase 2 of Project Stella could mark the beginning of this wave, by considering how DLT can enable new models of delivery versus payment (DvP). The project has 'proven that cross-ledger DvP could function even without any

connection between individual ledgers, a novelty which does not exist in today's set-up'⁸. In particular, it seems likely that the use of DLT for new ways of achieving cross-border payments will be a key theme of upcoming research, with the MAS announcing that future phases of Project Ubin will consider new ways of conducting cross-border payments⁹. Commercial banks, like the Royal Bank of Scotland¹⁰, have also shown interest in this use case, highlighting that cross-currency international payments have no central authority and that, as such, DLT might be a suitable technology, particularly given the claims that the current system that is 'costly, inefficient, and not transparent'¹¹.

In this context, Project Khokha falls in the second wave, ultimately adding to the body of knowledge in this sphere in two key respects:

1. Each bank has deployed its own node, operating under a variety of different deployment models (on-premise, on-premise virtual machine and cloud) and across distributed sites. The operation of this system is thus more robust than similar nodes being run in the same environment. The fact that the performance of the system is able to accommodate current volumes brings much more weight to the viability of this technology as an alternative to centralised infrastructures in the wholesale payments space.
2. Confidentiality has been a key issue within wholesale payments PoCs, typically forcing trade-offs in finality and speed. Several different architectures have been tested to deal with confidentiality. The approach taken during Project Khokha made use of new methods to achieve confidentiality while still achieving operational-level transaction volumes and allowing a distributed consensus mechanism.

⁸ Kobayakawa, S., Kawada, Y., Kobayashi, A., Bullmann, D., Bachmann, A., Humbert, C., Mögelin, S. and Galano, G. (2018)

⁹ <http://www.mas.gov.sg/Singapore-Financial-Centre/Smart-Financial-Centre/Project-Ubin.aspx>.

¹⁰ Creer, D., Crook, R., Hornsby, M., Avalis, N., Simpson, M., Weisfeld, N., Wyeth, B. and Zieliński, I. (2016)

¹¹ *Ibid.*



The global body of knowledge around scalability, privacy, resilience, and finality has grown significantly.

This second wave of central banks creating RTGS platforms using DLT has typically considered four key factors: scalability, privacy, resilience and finality. Below is a broad overview of what the DLT solutions have achieved.

SCALABILITY

In general, projects have sought to address scalability and transaction processing concerns that arose in the first wave of projects. By moving away from a proof-of-work mechanism, DLT has been able to execute transactions at a rate that could accommodate the RTGS volumes seen today. However, this scalability has typically come at a cost – often through limiting privacy and/or the resilience of the system.

PRIVACY

The exploration of different DLT platforms has resulted in a number of different architectures or designs to enable privacy. A core focus of the second wave of central-bank activity has been to investigate how best to achieve this privacy. Typically, we have seen the inclusion of channels which broadcast messages or transactions only to the applicable parties. Overall, studies have shown that privacy concerns around using DLT in an RTGS environment can be mitigated in a number of different ways.

RESILIENCE

Resilience is a feature of blockchain and DLT systems, in that full nodes hold a copy of all the transactions, so that if a node fails, it can replicate the full data set when it comes back online. Where the consensus process is distributed, there is also resilience to nodes failing, since the consensus does not rely on a single node. In a private or permissioned blockchain, where the consensus process depends on a single node, e.g. a notary node, the single point of failure is reintroduced and this element of resilience is lost. Furthermore, the use of channels, where the full details of transactions are not broadcast globally, could be an issue in terms of recovery in the event of failure.

FINALITY

In general, the second wave of central-bank activity has used DLT platforms that ensure finality. The early experiments in DLT showed that proof of work is deterministic and that theoretically settlement finality is probabilistic and never final. Subsequent experiments


have looked at resolving this typically through the introduction of a central authority providing a final ‘stamp of approval’ in the form of a notary node (Jasper, Ubin Phase 2), the use of a Practical Byzantine Fault Tolerance (PBFT) consensus mechanism with Hyperledger Fabric, or the use of Raft or IBFT in Quorum.

TRADE-OFFS

These characteristics should not be viewed in isolation, as the design decisions of creating a DLT solution typically involve trade-offs between these elements. It is worth noting that while the trade-offs are present today, there is significant development underway seeking to address these issues. A simple example of this would be to consider the first wave of experiments, which repeatedly discussed the issue of settlement finality. Today, a number of alternative consensus mechanisms have been developed, so that achieving finality is no longer an issue, since proof of work, and the associated probabilistic finality, is not used. It is expected that many of the other issues with DLT systems today will be addressed in future.

However, some of the reports have highlighted trade-offs in both privacy and resilience for scalability. Jasper Phase 2, Stella Phase 1 and Ubin Phase 2 highlighted some of these trade-offs. More specifically:

- Jasper Phase 2 used Corda, which brings privacy and transaction speed but presents some challenges to resilience most notably the introduction of a notary node as a single point of failure.
- With Stella Phase 1, transaction speeds comparable to today’s RTGS systems were achieved, but without addressing privacy.
- Lastly, Ubin Phase 2 provides interesting observations regarding the different DLT platforms:
 - The Corda prototype illustrated a similar trade-off to that of Project Jasper.
 - The Hyperledger Fabric prototype achieved the necessary throughput but introduced a single point of failure. Perhaps most interestingly, the use of channels for privacy presents a scalability challenge, as the management of these channels becomes increasingly complex as the number of participants grows.
 - For the Quorum prototype, the way that zero-knowledge proofs (ZKPs) were achieved significantly increased transaction times (to approximately six seconds overall), with implications for scalability.

Figure 11: The three waves of distributed ledger technology projects

05

CONCLUSION AND FUTURE WORK

INNOVATION AND COLLABORATION

The existing industry groups helped to lay the foundation for collaboration on Project Khokha.


Project Khokha builds on previous South African financial services industry collaboration. The fact that some of the project team members had already worked together, both at the SAFBC and the SUG, provided the foundation on which the successful collaborative approach of Project Khokha could be built.

The collaboration on Project Khokha has been successful.

Many participants in the project have remarked on how well the collaborative approach has worked in resolving issues quickly and supporting other participants to make swift progress. This bodes well both for future initiatives that develop from Project Khokha and also for future collaborative innovation in other areas. There has been a global trend in recent years of fintech companies and financial institutions shifting from an adversarial to a more collaborative approach. With the advent of open banking and similar movements, the ability to collaborate effectively will become increasingly important.

Other initiatives fostering innovation and collaboration are being set up.

Another collaborative initiative has been set up in which regulators and policymakers have engaged with industry to develop key considerations and a more harmonised approach to fintech-driven innovations. This initiative is known as the Intergovernmental Fintech Working Group (IFWG) and includes four financial regulators, namely the FIC, the Financial Sector Conduct Authority (FSCA), National Treasury and the SARB. The purpose of the initiative is to identify the risks and benefits of financial innovation driven by fintech in order for regulators and policymakers to develop appropriate policies and implement effective frameworks that allow for responsible innovation.



THE SUITABILITY OF DISTRIBUTED LEDGER TECHNOLOGY

Project Khokha has achieved its stated objectives.

Project Khokha has proven that DLT can perform the basic transaction function of an RTGS wholesale payment system. The SARB has adopted the PFMI as the key risk management standard for systemically important FMIs in South Africa, so it is appropriate to apply these to this work. Project Khokha was primarily a PoC implementation and as a result excluded a number of PFMI from the evaluation of the initiative. The governance and legal aspects that related to production-ready FMIs were not included in the project and remain an area of focus should DLT be considered for future production-ready environments.

The project considered a limited set of PFMI, namely PFMI 8 (Settlement finality), PFMI 9 (Money settlements) and PFMI 17 (Operational risk).

The success of the project in meeting these PFMI may be assessed as follows:

- **Settlement finality (PFMI 8)**
There are two aspects to settlement finality: operational finality and legal finality. Operational finality is achieved via the IBFT consensus mechanism described above, in that it is not possible to revoke transactions once they have been committed to the ledger and the point at which funds are transferred from one bank's wallet address to another is the point of settlement. Legal settlement is defined in the NPS Act as follows: "Settlement is effected in money or by means of entries passed through the SARB settlement system or a designated settlement system". For the settlement to be considered legally binding therefore requires the SARB to adjust the legal framework, assuming that the tokenised rand is recognised for these purposes. Establishing this clarity would be a part of the process if a DLT system were put into production.
- **Money settlement (PFMI 9)**
The principle of the Project Khokha system is that banks' DLT balances are pre-funded with the SARB. Final settlement is done in a tokenised rand, which is a claim against the SARB. This would satisfy the requirements of PFMI 9 for a production environment.

Production-ready real-time gross settlement?

- Operational risk (PFMI 17)

The core operational-risk considerations in wholesale payment systems relate to resilience, security and scalability. A detailed assessment of each of these operational risks is not possible in the case of Project Khokha as the system is not at production level. Managing operational risk is intimately linked to security and business continuity and the reduction of risk in these areas requires the implementation of coordinated end-to-end procedures. Operational risks in payment and settlement systems arise through technical failure, problems at the premises of the operator and/or incorrect procedures that threaten the clearing and settlement process. The knowledge around the use of DLT-based systems in wholesale settlements is relatively limited given the infancy of the technology. Additional insight needs to be garnered into these risks and the potential impact that they could have on a payment system. One of the requirements for participation in the SAMOS system is to have adequate business continuity planning and sufficient disaster recovery procedure facilities in place. These requirements would be no different for a DLT system. While Project Khokha has been successful in demonstrating the ability of DLT to perform transactions as required in the wholesale settlement environment in South Africa, the overall design and architecture of the system would need to be further evaluated to understand and mitigate against any risks that could impact on the overall integrity, security and availability of the network.

To address the full set of relevant PFMI, additional work would need to be done to achieve a complete RTGS solution. This would entail building additional capabilities into the system, such as the ability for banks to obtain credit and to reject or queue payments, and a more detailed assessment of the operational and legal issues relating to the other PFMI that are relevant to RTGS.

The Project Khokha PoC was not intended to replace the SAMOS system, and production readiness would require additional work on integrating with bank systems and addressing other issues. The full functionality of SAMOS includes liquidity management and credit extension, and it should be noted that these elements have not been included as part of this PoC.

The key element to note is that DLT opens up a new environment of learning and new scenarios that can be explored for the entire industry, not only for wholesale payments. There are numerous use cases that can be explored, and the fundamental element to consider is how this could contribute to the economic well-being of societies. A further step would therefore be to identify what these use cases are and to explore them further to build on this work.

The benefits and risks framework of the Bank for International Settlements is useful for analysis.

Over and above the PFMLs, the framework published by the Bank for International Settlements (BIS) for evaluating the benefits and risks associated with a DLT implementation is relevant here. This framework can be used for evaluating the suitability of DLT-based systems. It covers four areas:

- Functionality/scope of the system;
- Efficiency implications;
- Safety implications; and
- Broader financial market implications.

The BIS framework is specifically targeted at central banks and other entities interested in analysing and reviewing DLT applications, with the key objective of understanding their use cases together with the identification of opportunities and risks. This framework is outlined below; it highlights some of the key considerations that relate to the suitability of DLT in wholesale settlements.

Functionality and scope.

The functionality of DLT is well suited to the wholesale payments use case, as several central banks have demonstrated. The retail use case of DLT has not been explored at this stage, and there are numerous factors to take into consideration because, although it gives rise to greater efficiencies and hence lower frictional costs, in doing so, it should not impact on the safety and soundness of the payment system. It should also be noted that not all the scenarios and functionality of DLT have been tested within a production environment yet, and this raises numerous questions that need to be explored further. In particular, issues relating to regulatory compliance and reporting, cross-border payments, securities and retail offer ample scope for further work.

Implications for efficiency.

Efficiency includes both speed of operations and market efficiency. Project Khokha has shown that a DLT system can sustain process speeds that enable it to cope with the typical transaction volumes of the operational environment. In the case of Project Khokha, the DLT system was able to process the current daily volumes of SAMOS within a two-hour period. There would be technology and resource costs to implementing a new DLT system, but these must be offset against the additional functionality that a DLT system would bring. A DLT system also brings a higher degree of resilience, in that the reliance on a single point of failure is removed and each node additionally acts as a backup of the ledger. In that regard, it has the effect of redistributing some of the costs of resilience among the participants, which, in a live environment in South Africa, could remove the costs to the SARB of running a backup system for SAMOS. The levels of costs incurred and saved for the various parties in this situation are beyond the scope of this study.

Implications for safety.

The term ‘safety’ is used by the BIS to cover risks relating to legal and governance issues as well as operational risks and data protection. The issue of the legal status of smart contracts has been raised in relation to DLT systems, but the smart contracts used in Project Khokha are simply carrying out the processing rules. The legal issues relate to the status of the payment on the system, which needs a more detailed review in relation to the legal status of the tokenised rand before it is converted back into money on the SARB’s systems. The governance of the DLT system in a live environment would need to reside with the SARB due to its regulatory oversight mandate, regardless of the system used.

The cybersecurity of the DLT system is critical in implementation; the relative novelty of DLT systems will require additional exploration to fully understand the cybersecurity implications. The participants would also need to determine how to manage and protect their private keys. The protection of confidential data in the system has been a key focus of Project Khokha and the work with Pedersen commitments and range proofs has shown that it is possible to provide for real-world confidentiality requirements on a DLT system. There is an additional benefit, in that the details of transactions may be selectively revealed to regulators, as required, by giving them access to the system, which could make regulatory reporting more efficient. Operationally, Project Khokha has also demonstrated how finality of settlement can be achieved using a DLT system.

Broader financial market implications.

The broader implications of the implementation of DLT systems relate primarily to financial markets. The advent of DLT payment systems brings additional complexity to the ecosystem, but also brings advantages. Securities settlement is one such example that requires multiple system linkages as the systems that manage the transfer of value are separate to those that manage the transfer and settlement of securities. If the asset and the currency were on the same blockchain, then DLT brings potential for great improvements.

As different actors in global financial markets introduce DLT systems, the potential complexity in interoperability and the potential for new business models both increase. One of the key challenges that exist in relation to implementation considerations lies in the ability of each of the selected systems to communicate – not only with each other, but also with existing legacy systems. One of the ways in which this can be addressed is through the introduction of technical interoperability standards as these would provide a base layer of connectivity that would help to standardise and lower implementation and integration costs. One of the key benefits of such initiatives would be the ability of financial systems to bring in network scale effects and efficiencies.

The nature of DLT may result in the disintermediation of certain functions and entities within the financial market ecosystem, which might affect the competitive balance in financial markets. A disruption of such a nature could have a material impact on the balance within the ecosystem and as a result lead to various implications for financial market architecture. An area of uncertainty is in the introduction of new non-bank players that are not accounted for in many regulatory regimes. This may drive the regulation of market activities rather than market entities.

IMPLEMENTATION CONSIDERATIONS

There are several issues to resolve before distributed ledger technology systems will be ready for production.

There are a number of considerations that require further exploration before determining whether DLT systems are suitable for use in wholesale interbank settlement. Project Khokha offers unique practical insights into some of the challenges faced in setting up a distributed system, with the participants utilising different architectures in their testing environments. This is only the starting point.

The implementation of a production-ready RTGS system requires a deeper understanding of complexities that were not within the scope of this project. Key considerations that need to be addressed include the evaluation of supporting frameworks and other systems that integrate with the RTGS system, as well as the legal, regulatory and compliance factors, and the further complexities of a full RTGS system noted above.

The macroeconomic considerations of distributed-ledger-technology-based real-time gross settlement.

The BoE has conducted a theoretical evaluation of the macroeconomic considerations relating to a central-bank-issued digital currency (CBDC) and published its findings in a working paper.¹² One of the main challenges the report highlights is that no monetary regime currently exists in which CBDC has been fully adopted, which makes the practical assessment of DLT difficult to achieve. This was largely due to, at the time of publication in 2016, the lack of available technology to make such systems feasible and resilient. In addition to this, there had been very little evidence that could help in understanding the costs and benefits of transitioning to such a system. This also increases the difficulty of evaluating the various ways in which monetary policy could be implemented. As a result, the approaches used to date have been largely theoretical, but they do provide some insight into the challenges and impact that DLT presents.

The economic impact of DLT was not assessed in Project Khokha but, based on the results of the findings provided in the working paper compiled by the BoE, it is reasonable to assume that the impact of such a DLT system would be felt across multiple areas. The scope of a CBDC initiative is significantly wider than that of Project Khokha and remains a separate topic for consideration. It does, however, provide an indication of the complexity involved in determining the full-scale macroeconomic considerations of a production-ready RTGS DLT system.

¹² Barrdear, J. and Kumhof, M. (2016)

THERE ARE A NUMBER
OF CONSIDERATIONS
THAT REQUIRE FURTHER
EXPLORATION BEFORE
**DETERMINING WHETHER
DLT SYSTEMS ARE SUITABLE
FOR USE IN WHOLESALE
INTERBANK SETTLEMENT.**

POTENTIAL FUTURE DIRECTION

The insights gained from this work, and the momentum generated by the project, prompt the question: “What comes next?” While other central banks are experimenting in this area, DLT technology and the business models built around it are evolving rapidly, and there are national and regional implications should a DLT system be put into operation. A fully live DLT-based payments system is not currently planned in South Africa. This section explores some of the work that would need to be done to understand the potential implications if such a decision were considered.

Blockchain technology continues to evolve, and so technology decisions are difficult.

Several blockchain systems are being developed that are focused on financial service applications and problems such as privacy, performance and resilience; use-specific operational issues are dealt with in different ways by different systems. These factors contribute to the difficulty in identifying the best route forward as far as technology selection and future architecture are concerned. As ever, there are trade-offs in choosing one system over another, not least of which are the potential network effects from similar choices by peers, although these would be mitigated as interoperability tools are developed.

A system based on distributed ledger technology opens up other avenues of exploration.

As described in the discussion of the third wave above, broadening the scope of DLT use opens up multiple possibilities which central banks are starting to explore. Understanding the technology is the first challenge, then collaborating, then going into production. If one starts from the point where money is tokenised (as is done in this project) and then represented on a DLT system, then this system can be developed to enable other uses beyond wholesale settlement. Examples include the exchange of tokenised money for other tokenised assets, like bonds or securities. These may exist on the same system or on different blockchains. Such a system has considerable implications for other industries, like asset management. Access to the payments system could also be widened to include large companies, small companies, and even retail customers. It would need to be established whether these parties have direct access via their own nodes or as clients of existing node owners and, if they operate nodes, then the role that those nodes play in the consensus process. As the system becomes broader, a wider variety of trade finance applications would be possible, with additional payment experiments such as DvP an obvious next step. Several central banks have now reached similar stages with a DLT wholesale payments system. If they coordinate their future experiments, the learning process can be sped up through cooperative cross-border projects or through ensuring that they do not duplicate work unnecessarily.

As several central banks create similar systems, the possibility arises of linking DLT networks across national borders. Assuming that the technological challenges can be overcome – and this is the stated goal of some proposed projects, such as the future phases of Project Ubin – this could enable cross-border payments. The process would require a link between the two countries’ DLT systems, either by connecting them or by using a linkage outside the DLT system. One option is that the same organisation has nodes on both blockchains and creates the link between the two, thus creating a path for payments.

A preferable option might be to have the two central banks connect, enabling a payment to flow without exposing details to any commercial bank. If the two systems are in different currencies, then the conversion would be an added factor to deal with. Conceptually, cross-border connections between central banks could be thought of as a macro version of in-country DLT systems. The global coordination of central banks that this enables could conceivably facilitate monetary policy and other fiscal coordination across borders, although this is currently a long way off. As a first step, the SARB could participate in cross-border DLT projects with other central banks globally or drive a Southern African Development Community cross-border experiment.

If full-scale retail access to central-bank accounts is allowed and a CBDC is created, then the shape, architecture and impact on the banking system would require careful review. These types of issues have been explored by others, and their potential implications have been cited. The possibilities mentioned include smart money and enabling real-time tax collection and government spending. From a macro perspective, it has been suggested that the introduction of a CBDC may bring economic and monetary policy benefits^{13,14}. CBDC requires careful policy consideration and research.

¹³ Ehsani, F. (2016)

¹⁴ Barrdear, J. and Kumhof, M. (2016)

Taking a distributed ledger technology system into production would have multiple implications beyond the core technology.

While DLT is not yet mature, the technology has been the subject of a number of projects carried out by central banks. The issues adjacent to these new systems will vary on a case-by-case basis, and will include integration with other systems, the fundamentals of cybersecurity, the management of private keys by the parties running nodes, and the legal and accounting aspects of incorporating DLT into current structures.

This work would potentially be time-consuming, with a large number of parties involved, and would be a key component of any implementation programme. The regulatory and legislative elements are particularly complex. A DLT-based system could possibly enable more efficient regulation that is done in real time via an observer node for the regulator. At the same time, in a distributed system, there are elements of control that a regulator like the SARB may have to consider relinquishing in order to achieve some of the other benefits (such as distributed consensus), which would improve resilience by removing the single point of failure. It has been argued that the regulatory challenges are much harder to resolve than the technological ones.

One logical first step for adoption, which has been mooted by the Central Bank of Brazil, is to install a DLT-based system as a backup to the 'normal' payments infrastructure. Taking into account that this would require integration and other work noted above, this idea might provide the impetus to build a production-level system without taking the significant step of going fully live with a DLT system. It could be argued that this would require as much development as taking a system live, unless the disaster recovery version has less onerous requirements for integration with other systems and the other issues noted in the preceding paragraphs. The only reason for not going live with a DLT system would then be a value judgement based on the relative qualities of the systems.

The potentially broad impact of a distributed-ledger-technology-based system needs to be considered.

The business case for implementing DLT systems is usually not made on a like-for-like replacement basis. In the South African context, for example, the case to do a straight replacement for the SAMOS wholesale payments system with a DLT equivalent would probably not be clear-cut. What is more interesting is the broader impact of a DLT system and the kinds of things it enables as applications and use cases are built around it, such as those noted above. A logical next step from Project Khokha might therefore be an economic impact analysis of the implementation of DLT under a number of scenarios. This could take into account the various options for development, their interaction and their full financial impact. Such a study would need to consider varying degrees of DLT usage in the economy and would be an important parallel piece of work to further experiments with the technology.

WHAT IS MORE INTERESTING
IS THE BROADER IMPACT
OF A **DLT SYSTEM AND
THE KINDS OF THINGS
IT ENABLES AS
APPLICATIONS.**

06

APPENDICES

Detail on central-bank distributed ledger technology projects.

Table 6 provides an overview of key insights gained from initiatives previously undertaken by other central banks.

Table 6: Breakdown of central-bank projects

	Approach	Key insights gained
Brazil (Central Bank of Brazil – Phase 1)	Identify uses for DLT inside the Central Bank of Brazil. Elect one use case and platform for prototyping. Produce a minimal PoC.	<p>SCALABILITY Scalability was not a key focus of this report.</p> <p>PRIVACY The innate data transparency of Ethereum infringes on privacy requirements. While a privacy solution was created, there were several drawbacks (e.g. smart contracts could not see all the data, and the solution lacked forward secrecy).</p> <p>RESILIENCE A private Ethereum blockchain using proof of work is inherently resilient. The report noted that, to enable privacy gaps, a trusted node could be used to ensure privacy amongst participants, but at the cost of this resilience.</p> <p>FINALITY The use of proof of work does not bring settlement finality.</p> <p>OTHER Four key use cases were identified: digital identity, a local-currency payments system, agreement on reciprocal payments and credits and an alternative system for transactions settlement. The last-mentioned use case was selected, looking at an alternative system for transactions settlement which would be able to immediately replace the core functionalities of the main Brazilian RTGS system in case of a full collapse.</p>
Brazil (Central Bank of Brazil – Phase 2)	<p>Analyse competing blockchain platforms using the selected use case as a benchmark.</p> <p>As privacy was a key issue in Phase 1, the main goal of Phase 2 was to seek a platform which allowed the development of an RTGS system that agrees with all requirements.</p>	<p>SCALABILITY Scalability was not a focus of this study.</p> <p>PRIVACY The Hyperledger Fabric 0.6 prototype had the same limitations as Ethereum with privacy. There is a need for an additional off-chain layer to create private transactions. The report noted the promise of channels in the upcoming Hyperledger Fabric 1.0. Furthermore, node administrators have access to all the information in the blockchain.</p> <p>The Corda prototype broadly met the required privacy requirements, given that there is no globally shared blockchain and transactions occur in a peer-to-peer manner.</p> <p>The Quorum prototype created private transactions that were encrypted and addressed to specific nodes on the network. Nodes not privy to the contents receive hashes of the encrypted private transaction data.</p>

	Approach	Key learnings
Brazil (Central Bank of Brazil – Phase 2)	<p>Analyse competing blockchain platforms using the selected use case as a benchmark.</p> <p>As privacy was a key issue in Phase 1, the main goal of Phase 2 was to seek a platform which allowed the development of an RTGS system that agrees with all requirements.</p>	<p>RESILIENCE</p> <p>For the Hyperledger Fabric prototype, a node needed to be enrolled using certificates provided by a membership service, which may present a single point of failure. The use of PBFT as a consensus mechanism increases the resilience of the system, as consensus does not rely on a single, set node. Interestingly, the Hyperledger Fabric 0.6 environment supports smart contracts with fewer coding restrictions than Ethereum, which means that it could be possible to write non-deterministic code that could present a risk to the system.</p> <p>In Corda, there is no public blockchain and no private channels; each transaction occurs in a peer-to-peer manner. Without a disaster recovery feature built into the platform, the process of retrieving data from the network would be cumbersome and risky. Furthermore, the notary function presents a single point of failure.</p> <p>The Quorum prototype is regarded as resilient, particularly as it benefits from (at the time) two years of Ethereum code being tested publically.</p> <p>FINALITY</p> <p>Settlement finality was achieved as a result of the consensus mechanism for Hyperledger Fabric (PBFT) and Corda (a notary node).</p> <p>The Quorum prototype utilised QuorumChain. It is not considered suitable for transactions that require settlement finality, as forking is possible.</p> <p>OTHER</p> <p>Ring signatures or ZKP techniques are promising and may be suitable for future work. Upcoming features of DLT platforms, such as channels in Hyperledger Fabric, were also cited as being able to address privacy issues in better ways.</p>

	Approach	Key learnings
Canada (Project Jasper – Phase 1)	Build a PoC as an experimental review of DLTs applicability to wholesale interbank payments settlement.	<p>SCALABILITY</p> <p>The proof-of-work consensus mechanism raised concerns around the scalability of the system, and could not provide the throughput required as volumes increased.</p> <p>PRIVACY</p> <p>The solution, built on Ethereum, provided full visibility of the ledger to all the participants, which does not meet data privacy requirements.</p> <p>RESILIENCE</p> <p>The solution was built on Ethereum, with all the participants having a copy of the distributed ledger and all being able to validate transactions. As a result, the study found that a DLT-based payment platform could provide more cost-effective resilience by having no single point of failure. However, the study noted that additional technology components such as private key, identity and system access management, are based on centralised models and need to be made highly available and backed up for disaster recovery.</p> <p>FINALITY</p> <p>With a proof-of-work consensus mechanism, there is doubt as to whether settlement finality is ever achieved, as it is probabilistic, i.e. there is always a small probability that the payment could be reversed; while settlement becomes increasingly certain over time, it never reaches a point of being completely irrevocable.</p>
Canada (Project Jasper – Phase 2)	<p>Further evaluate the scalability and flexibility of DLT by building a Corda PoC.</p> <p>Expand the exchange capabilities to include both atomic and deferred net settlement.</p> <p>Continue to build in more of the functionality observed in today's interbank settlement solutions (i.e. LSMs).</p> <p>Provide a data-driven simulation exercise with operational data sets to evaluate platform and LSM performance.</p>	<p>SCALABILITY</p> <p>The Corda platform greatly increased the transaction speed in comparison to the previous phase.</p> <p>A simulation was performed around a 'low-volume day' (26 000 transactions with a value of US\$104.5 billion) and a 'high-volume day' (37 000 transactions representing US\$227.9 billion). Roughly 61 payments per minute occurred in the simulation, while only 49 payments per minute were experienced by the Large Value Transfer System (the Canadian RTGS system) on its observed high-volume day of 53 000 transactions in 2016. The simulation testing illustrated that current-day average volumes could be processed within an acceptable window.</p> <p>PRIVACY</p> <p>The consensus protocol of Corda not only achieves more expedited transaction processing; it also supports privacy of information among participating financial institutions. The platform partitions data such that each participant's proprietary ledger reflects only its transaction activity, with only the notary and supervisory nodes maintaining the full content of the shared database.</p>

	Approach	Key learnings
Canada (Project Jasper – Phase 2)	<p>Further evaluate the scalability and flexibility of DLT by building a Corda PoC.</p> <p>Expand the exchange capabilities to include both atomic and deferred net settlement.</p> <p>Continue to build in more of the functionality observed in today's interbank settlement solutions (i.e. LSMs).</p> <p>Provide a data-driven simulation exercise with operational data sets to evaluate platform and LSM performance.</p>	<p>RESILIENCE</p> <p>Each participant would need to invest in data replication and archiving to ensure business continuity and the recreation of the ledger. They would need to back up all the private data, including private keys and customer account information.</p> <p>The notary and supervisory nodes are needed for consensus, thus both represent single points of failure – and both would need to maintain high-availability systems. One solution would be to run several notary node clusters.</p> <p>Corda nodes can queue pending requests to other nodes in the event that a node is forced offline. While the transactions will be processed when this node comes back online, the impact of an offline node on the ecosystem could be substantial.</p> <p>Incorporating a central queue or LSM component would require consideration of the need for high availability, as this would introduce a single point of failure.</p> <p>FINALITY</p> <p>To ensure legal settlement finality, Project Jasper was structured in such a way that the exchange of digital depository receipts between platform participants would be equivalent to a full and irrevocable transfer of the underlying claim on a central-bank deposit. The consensus mechanism used in this phase also provided greater settlement certainty.</p> <p>OTHER</p> <p>It is likely that a benefit of introducing a DLT interbank cash payment platform would be the possibility of end-to-end support of additional DLT use cases. This, however, would assume a certain level of interoperability with other systems.</p>

	Approach	Key learnings
Europe and Japan (Project Stella – Phase 1)	<p>Assess whether the functionalities of the existing payment systems could be safely and efficiently run in a DLT application, focusing on hands-on testing only.</p> <p>The PoC included LSMs (queuing, bilateral offsetting, and multilateral offsetting).</p> <p>Performance testing was done both with and without LSMs.</p>	<p>SCALABILITY</p> <p>The study found that DLT-based solutions could meet the performance needs of an RTGS system. The DLT prototype could process volumes comparable to the RTGS system in the euro area and Japan, with transactions processed in less than one second on average.</p> <p>However, when increasing the requests per second to 250, the trade-off between traffic and performance was not negligible.</p> <p>Generally, the tests proved the feasibility of implementing the processing logic of standard LSMs (queuing and bilateral offsetting) in a DLT environment.</p> <p>Specific performance testing found that DLT performance is affected by the network size and the distance between nodes:</p> <ul style="list-style-type: none"> • Increasing the number of nodes led to an increase in payment execution time. • The impact of the distance between nodes on performance was found to depend on the network configuration: provided the minimum number of nodes (quorum) required to achieve consensus was sufficiently close together, the effect of dispersion in the rest of the network on latency was limited. <p>PRIVACY</p> <p>While not explicitly stated, the study appears not to have incorporated privacy into the prototype. However, the report noted that sharing information among all the nodes could raise concerns regarding data privacy.</p> <p>RESILIENCE</p> <p>The report highlighted a number of important considerations for resilience.</p> <p>Tests showed that as long as the number of nodes required by the consensus algorithm (three out of four) was operational, the availability of the overall system was not affected by the failure of one node.</p> <p>A validating node was also able to recover in a short period of time (< 30s) for a range of plausible downtime scenarios.</p> <p>Critically, registering and authenticating participants is key to security, and Hyperledger Fabric ensures this through a certificate authority (CA). While transaction validation is distributed, the CA is centralised, and when it is not available, transactions are rejected.</p> <p>There were no difficulties in processing transactions with a correct format regardless of the percentage of incorrectly formatted messages.</p> <p>FINALITY</p> <p>The consensus mechanism (PBFT) provides settlement finality.</p>

	Approach	Key learnings
Europe and Japan (Project Stella – Phase 2)	<p>Explore how the settlement of two linked obligations, such as the delivery of securities against the payment of cash, could be conceptually designed and operated in an environment based on DLT.</p> <p>Importantly, the study examined ways in which DvP can be conceptually designed and technically achieved in a DLT environment drawing on existing DvP models as well as innovative solutions that are being discussed for distributed ledgers.</p>	<p>Stella Phase 2 has a different scope to the previous projects and does not use the lenses of scalability, privacy, resilience and finality. However, key findings are given.</p> <p>DvP can run in a DLT environment, but subject to the specificities of the different DLT platforms. DvP can be conceptually and technically designed in a DLT environment with cash and securities on the same ledger (a single-ledger DvP) or on separate ones (a cross-ledger DvP). The concrete design of a DvP, however, depends on the characteristics of the DLT platforms.</p> <p>DLT offers a new approach for achieving DvP between ledgers, which does not require any connection between ledgers. Cross-ledger DvPs could function even without any connection between individual ledgers, which is a novelty that does not exist in today's set-up¹⁵. Functionalities such as cross-chain atomic swaps have the potential to help ensure interoperability between ledgers (of either the same or different DLT platforms) without necessarily requiring connection and institutional arrangements between them.</p> <p>Depending on their concrete design, cross-ledger DvP arrangements on DLT may entail certain complexity and could give rise to additional challenges that would need to be addressed. Depending on the concrete design, this could impact on transaction speed and require the temporary blockage of liquidity (read the report for more details)¹⁶.</p> <p>The project also used a number of different DLT platforms, performing an assessment of the differences and benefits between them.</p>
Singapore (Project Ubin – Phase 1)	<p>Build a PoC for domestic payments for interbank obligations on a distributed ledger, denominated in balances backed by the central bank.</p> <p>Identify the non-technical implications of taking the PoC into production.</p>	<p>SCALABILITY Scalability was not a focus of this study.</p> <p>PRIVACY While details were not given, the PoC was run on an Ethereum-based blockchain, and thus – presumably – did not have privacy of transactions.</p> <p>RESILIENCE While details were not given, the PoC was run on an Ethereum-based blockchain, and thus – presumably – the system can be considered not to have single points of failure and thus as resilient.</p> <p>Furthermore, as banks fund their expected payments, there is virtually no liquidity risk in the distributed ledger. Thus, even the outage of the largest participant would not prevent the remaining participants from completing payments.</p> <p>FINALITY The report acknowledged the need for an appropriate legal structure to ensure that the transfer of the token is equivalent to a full and irreversible transfer of the underlying claim on the central bank's currency. Presumably, the prototype used a proof-of-work consensus mechanism and thus does not meet the requirements for settlement finality.</p>

¹⁵ Kobayakawa, S., Kawada, Y., Kobayashi, A., Bullmann, D., Bachmann, A., Humbert, C., Mögelin, S. and Galano, G. (2018)

¹⁶ Ibid.

	Approach	Key learnings
Singapore (Project Ubin – Phase 2)	<p>Assess the potential implications of deploying DLT for specific RTGS functionalities by focusing on LSMs.</p> <p>Understand how RTGS privacy can be ensured using DLT.</p> <p>Conduct an objective assessment on three blockchain platforms.</p>	<p>SCALABILITY</p> <p>Overall, Phase 2 showed that the key functions of an RTGS system – such as fund transfer, a queuing mechanism, and a gridlock resolution – can be achieved through different techniques and solution designs.</p> <p>For the Corda PoC, it was observed that, due to the fact that peers only see a subset of the entire ledger as opposed to ‘traditional’ DLT platforms which store and update the entire ledger at every peer, the scalability and performance issues commonly associated with ‘traditional’ DLT platforms were alleviated.</p> <p>For the Hyperledger Fabric prototype, the number of channels increased with every new participant, which increased the complexity of network and channel management. Furthermore, individual bilateral channels between pairs of transacting banks required bank operators to maintain funds in each channel, in addition to the total pledged fund from MAS. While solutions to this could be devised, it does present scaling concerns.</p> <p>For the Quorum prototype, it was observed that the current ZKP generation process takes approximately four seconds, with a total transaction processing time of five seconds for a fund transfer, which is a clear scalability concern.</p> <p>PRIVACY</p> <p>Privacy was enabled for each PoC using the platforms’ respective tools.</p> <p>Corda used its Unspent Transaction Model Output model and confidential identity. The ledger is distributed on a need-to-know basis instead of using a global broadcast method, and strengthens its privacy with an additional layer of confidential identities added to each transaction, whereby only the parties involved in a transaction can identify the participants, although the amounts transferred are not shielded (which could present privacy concerns).</p> <p>Hyperledger Fabric leveraged its channels design. Channels enable information to be shared between parties on a need-to-know basis, so that only the participants in the channel can see the transactions.</p> <p>Quorum used Constellation and ZKPs. Private transactions were sent directly to the recipients using Constellation, as encrypted blobs, where the rest of the network can only see a hash of the encrypted payload. ZKPs were used for managing the shielded salted balance of each participating bank, ensuring that the true balance position of a bank is only visible to itself.</p>

	Approach	Key learnings
Singapore (Project Ubin – Phase 2)	<p>Assess the potential implications of deploying DLT for specific RTGS functionalities by focusing on LSMs.</p> <p>Understand how RTGS privacy can be ensured using DLT.</p> <p>Conduct an objective assessment on three blockchain platforms.</p>	<p>RESILIENCE</p> <p>The project showed that key functions of an RTGS system – such as fund transfer, a queuing mechanism, and a gridlock resolution – can be achieved through a number of different platforms.</p> <p>Furthermore, decentralising the key functions of an RTGS system may not only mitigate the inherent risks of a centralised system, such as a single point of failure, but may also affirm the promised benefits of DLT, for example cryptographic security and immutability.</p> <p>For the Corda platform, the current prototype assumed a high availability of participants and therefore did not support the removal of any participant node during the gridlock resolution cycle, except for MAS.</p> <p>The design of the Hyperledger Fabric prototype entailed sending a transaction to all the peers in a channel for validation, presenting a single point of failure to the network (although this can be resolved with the implementation of a multiple-node ordering service).</p> <p>Should any node be disconnected from the Quorum prototype, the rest of the network would still be available, although no transactions with that node would be allowed. Furthermore, the transaction history is automatically synchronised when the node comes back online. Furthermore, the Raft consensus mechanism does not introduce a single point of failure for validating transactions.</p> <p>FINALITY</p> <p>Settlement finality was achieved for all the prototypes due to the consensus mechanisms used (Corda: notary node; Hyperledger Fabric: PBFT; and Quorum: Raft).</p>

Detail on Project Khokha testing iterations.

Further details for each iteration are given in Table 7 to supplement the description given on page 32.

Table 7: Breakdown of Project Khokha's four iterations

Iteration	Description
1	<p>This was a straightforward transfer between two banks.</p> <p>The process entailed:</p> <ol style="list-style-type: none"> 1. Bank A commits a payment to Bank B by calling the 'approveFor' function on the 'ZAR smart contract'. This is a standard Ethereum Request for Comment (ERC20) function. 2. There are some simple checks in the contract to ensure that the amount being transferred is positive and that the post-transfer balance of Bank A is still positive. <p>All rand balances and payment values were visible to the whole network.</p>
2	<p>This was very similar to the current SAMOS process.</p> <p>Bank A effectively sent a payment instruction that the SARB executed.</p> <p>The process looked as follows:</p> <ol style="list-style-type: none"> 1. Bank A commits a payment to Bank B by calling the 'approveFor' method in the 'ZAR smart contract' with the amount and the destination address. 2. The SARB checks the values and the post-transfer balance. If both are positive, they transfer the funds using the 'transferFrom' method, which checks to see if Bank A has committed funds to the payment. <p>All rand balances and payment values were visible to the whole network.</p>
3	<p>The amounts and balances were shielded, with the network only seeing the Pedersen commitment for each transaction. The SARB still approved the transfer.</p> <p>The process entailed:</p> <ol style="list-style-type: none"> 1. Bank A calls for the 'approveFor' method which has a 'text' metadata field that will contain the encrypted versions of the amount and the blinding factor. 2. Bank A communicates the decryption key to Bank B and the SARB via a Whisper channel (performed during network set-up). 3. The SARB checks the transaction amount and the post-transfer balance of Bank A. If both the amount and the post-transfer balance are positive, the SARB transfers the funds using the 'transferFrom' method, which will check to see if Bank A has committed the funds to the payment and will update the balances of Bank A and Bank B by subtracting and adding the commitment submitted by Bank A from/to the balances of Bank A and Bank B respectively. 4. Once the balance commitments of Bank A and Bank B are updated, the transaction is settled. <p>The Rand balances and transaction amounts are only visible to the two transaction parties and the SARB.</p>

Iteration	Description
4	<p>The amounts and balances are shielded, visible only to the two parties to the transaction and the SARB. The network only sees the Pedersen commitment and the nodes on the network approve the transaction via range proofs.</p> <p>The process looked as follows:</p> <ol style="list-style-type: none"> 1. Bank A calls for the 'approveFor' method which has a 'text' metadata field that will contain the encrypted versions of the amount and the blinding factor. 2. Bank A communicates the decryption key to Bank B and the SARB via a Whisper channel (performed during network set-up). 3. The sender produces a range proof in addition to the Pedersen commitment, so that network participants can verify that the closing balance of the sender and the transaction value are both positive (in the range of 0 to 264-1). 4. Public balance updates are done via consensus between network participants. The SARB can see all the transaction amounts and balances, but would not need to approve or reject transactions as this is done by the network participants.

List of references

- Accenture (2017). Project Ubin Phase 2: Re-imagining interbank real-time gross settlement system using distributed ledger technologies. [online]. Accenture, pp. 1-64. Available at: https://www.accenture.com/t20171116T081243Z__w__/sg-en/_acnmedia/PDF-66/Accenture-Project-Ubin-Phase-2.pdf [accessed 24 April 2018].
- Bank for International Settlements (2017). 'Distributed ledger technology in payments, clearing and settlement'. Journal of Financial Market Infrastructures 6 (2/3), pp. 207-249. [online]. Available at: <https://www.bis.org/cpmi/publ/d157.pdf> [accessed 17 April 2018].
- Barrdear, J. and Kumhof, M. (2016). 'The macroeconomics of central-bank-issued digital currencies'. Social Science Research Network Electronic Journal. [online]. Available at: <https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2016/the-macroeconomics-of-central-bank-issued-digital-currencies.pdf?la=en&hash=341B602838707E5D6FC26884588C912A721B1DC1> [accessed 24 April 2018].
- Burgos, A., Filho, J., Soares, M. and De Almeida, R. (2017). Distributed ledger technical research in Central Bank of Brazil. [online]. Central Bank of Brazil, pp. 1-33. Available at: https://www.bcb.gov.br/htms/public/microcredito/Distributed_ledger_technical_research_in_Central_Bank_of_Brazil.pdf [accessed 24 April 2018].
- Chapman, J., Garratt, R., Hendry, S., McCormack, A. and McMahon, W. (2017). Project Jasper: are distributed wholesale payment systems feasible yet? [online]. Bank of Canada, pp. 1-44. Available at: <https://www.bankofcanada.ca/wp-content/uploads/2017/05/fsr-june-2017-chapman.pdf> [accessed 24 April 2018].
- Creer, D., Crook, R., Hornsby, M., Avalis, N., Simpson, M., Weisfeld, N., Wyeth, B. and Zieliński, I. (2016). Proving Ethereum for the clearing use case. [online]. Royal Bank of Scotland, pp. 1-10. Available at: <https://emerald-platform.gitlab.io/static/emeraldTechnicalPaper.pdf> [accessed 24 April 2018].
- Dalal, D., Yong, S. and Lewis, A. (2017). The future is here: Project Ubin: SGD on distributed ledger. [online]. Deloitte Singapore, pp. 1-44. Available at: <https://www2.deloitte.com/content/dam/Deloitte/sg/Documents/financial-services/sg-fsi-project-ubin-report.pdf> [accessed 24 April 2018].
- Driscoll, K., Hall, B., Sivencrona, H. and Zumsteg, P. (2003). Byzantine Fault Tolerance, from theory to reality. [online]. Available at: <https://www.cs.indiana.edu/classes/p545-sjoh/post/lec/fault-tolerance/Driscoll-Hall-Sivencrona-Xumsteg-03.pdf> [accessed 20 April 2018].
- Ehsani, F. (2016). The advent of crypto-banking. [online]. Rand Merchant Bank. Available at: http://www.foundry.co.za/wp-content/uploads/2017/09/The_Advent_of_Crypto_Banking.pdf [accessed 17 April 2018].
- Financial Stability Board (2017). Financial stability implications from fintech, supervisory and regulatory issues that merit authorities' attention.
- GitHub (2018). jpmorganchase/quorum. [online]. Available at: <https://github.com/jpmorganchase/quorum/wiki/Transaction-Processing> [accessed 20 April 2018].
- International Organization for Standardization (2018). Description of the payments message formats I ISO 20022. [online]. Available at: <https://www.iso20022.org/> [accessed 20 April 2018].
- JPMorgan.com (2018). Quorum I J.P. Morgan. [online]. Available at: <https://www.jpmorgan.com/country/US/en/Quorum> [accessed 17 April 2018].
- Kobayakawa, S., Kawada, Y., Kobayashi, A., Bullmann, D., Bachmann, A., Humbert, C., Mögelin, S. and Galano, G. (2018). Securities settlement systems: delivery versus payment in a distributed ledger environment. [online]. European Central Bank and Bank of Japan, pp. 1-52. Available at: https://www.ecb.europa.eu/pub/pdf/other/stella_project_report_march_2018.pdf [accessed 24 April 2018].
- Kobayakawa, S., Kawada, Y., Watanabe, A., Kobayashi, A., Bullmann, D., Chorley, F., Humbert, C., Leach, T. and Pinna, A. (2017). Payment systems: liquidity saving mechanisms in a distributed ledger environment. [online]. European Central Bank and Bank of Japan, pp. 1-23. Available at: https://www.ecb.europa.eu/pub/pdf/other/ecb.stella_project_report_september_2017.pdf [accessed 24 April 2018].
- Macknight, J. (2017). The Banker Technology Projects of the Year Awards 2017. Transactions and technology. [online]. The Banker. Available at: <http://www.thebanker.com/Transactions-Technology/The-Banker-Technology-Projects-of-the-Year-Awards-2017?ct=true#Blockchain> [accessed 18 April 2018].

- Metere, R. and Dong, C. (2017). Automated cryptographic analysis of the Pedersen commitment scheme. Cryptography and security. [online]. Available at: <https://arxiv.org/abs/1705.05897> [accessed 20 April 2018].
- Nakamoto, S. (2008). Bitcoin: a peer-to-peer electronic cash system. Available at: <https://bitcoin.org/bitcoin.pdf> [accessed 18 April 2018].
- National Payment System Act 78 of 1998, as amended.
- National Payment System Department (2008). Starter pack for participation within the national payment system (NPS). [online]. South African Reserve Bank. Available at: [https://www.resbank.co.za/RegulationAndSupervision/NationalPaymentSystem\(NPS\)/AccessToTheNPS/Documents/StarterPack.pdf](https://www.resbank.co.za/RegulationAndSupervision/NationalPaymentSystem(NPS)/AccessToTheNPS/Documents/StarterPack.pdf) [accessed 17 April 2018].
- National Payment System Department (2013). Position paper confirming the South African Reserve Bank's support of the Principles for Financial Market Infrastructures as published by the Committee on Payment and Settlement Systems and the Technical Committee of the International Organization of Securities Commissions. [online]. South African Reserve Bank. Available at: [https://www.resbank.co.za/RegulationAndSupervision/NationalPaymentSystem\(NPS\)/Legal/Documents/Position%20Paper/PPNPS01_2013PFMIs.pdf](https://www.resbank.co.za/RegulationAndSupervision/NationalPaymentSystem(NPS)/Legal/Documents/Position%20Paper/PPNPS01_2013PFMIs.pdf) [accessed 17 April 2018].
- National Treasury (2011). A safer financial sector to serve South Africa better.
- Schwab, K. and Sala-i-Martin, X. (2017). Global Competitiveness Report. [online]. World Economic Forum. Available at: <http://www3.weforum.org/docs/GCR2017-2018/05FullReport/TheGlobalCompetitivenessReport2017%E2%80%932018.pdf> [accessed 18 April 2018].
- Staib, D., Pain, D. and Puttaiah, M. SIGMA 3/2017 (2017). World insurance in 2016: the China growth engine steams ahead. [online]. Swiss Re Institute. Available at: http://www.swissre.com/library/publication-sigma/sigma_3_2017_en.html#inline [accessed 18 April 2018].
- Technical Committee of the International Organization of Securities Commissions (2012). Principles for Financial Market Infrastructures. [online]. Bank for International Settlements and the International Organization of Securities Commissions. Available at: <https://www.bis.org/cpmi/publ/d101a.pdf> [accessed 17 April 2018].
- Volker, W. (2013). Essential guide to payments: an overview of the services, regulation and inner workings of the South African national payment system.

List of abbreviations

AML	anti-money laundering	NPS	national payment system
AWS	Amazon Web Services	NPS Act	National Payment System Act 78 of 1998, as amended
BIS	Bank for International Settlements	NPSD	National Payment System Department
BoE	Bank of England	PBFT	Practical Byzantine Fault Tolerance
CA	certificate authority	PFMI	Principle for Financial Market Infrastructures
CBDC	central-bank digital currency	PoC	proof of concept
DLT	distributed ledger technology	PwC	PricewaterhouseCoopers Inc.
DMZ	demilitarised zone	RAM	random-access memory
DvP	delivery versus payment	RTGS	real-time gross settlement
ERC	Ethereum Request for Comment	SAFBC	South African Financial Blockchain Consortium
FIC	Financial Intelligence Centre	SAMOS	South African Multiple Option Settlement (system) (system)
fintech	financial technology	SARB	South African Reserve Bank
FMI	financial market infrastructure	SARB Act	South African Reserve Bank Act 90 of 1989, as amended
FSCA	Financial Sector Conduct Authority	SSD	solid-state drive
GB	gigabyte	SteerCo	Steering Committee
GNU	GNU Operating System	SUG	SAMOS User Group
GPL	GNU Public Licence	SWIFT	Society for Worldwide Interbank Financial Telecommunication
IBFT	Istanbul Byzantine Fault Tolerance	tx/s	transactions per second
IFWG	Intergovernmental Fintech Working Group	US	United States
ISO	International Organization for Standardization	VCPU	virtual central processing unit
JAD	Joint Application Development	VM	virtual machine
KYC	Know Your Customer	ZAR	South African rand
LGPL	GNU Lesser General Public Licence	ZKP	zero-knowledge proof
LSM	liquidity saving mechanism		
MAS	Monetary Authority of Singapore		

Acknowledgements

South African Reserve Bank

Executive Sponsor

Deputy Governor Francois Groepe

Fintech Unit

Anrich Daseman

Arif Ismail

Gerhard van Deventer

National Payment System Department

Edward Leach

Rhona Badenhorst

Tlalane Mokuane

Business Systems and Technology Department

Jan Mohotsi

Masego Jabosigo

Willem Pienaar

Financial Services Department

Charles Mayela

Nomwelase Skenjana

Thando Lekalakala

Technical partner

ConsenSys

Coenie Beyers

Ed Budd

Jason Smythe

Marc de Klerk

Peter Munnings

Ryno Beyers

Support partner

PwC

Chantal Maritz

David Baron

Jalal Ghiassi-Razavi

Paul Mitchell

Rodney Prescott

Participant banks

Absa

Marek Skocovsky

Romana Linkeova

Sean Mouton

Capitec

Andre Bouwer

Kevin Feng

Pierre Kotze

Discovery Bank

Chanel Kotze

Fred Malan

Gary Alter

Ghita Erling

Manie van Rooyen

Sheldon Lee

Sylvia van Nieuwenhuizen

FirstRand

Badi Sudhakaran

Bernard Carless

Brendon Tolan

Chris Tsimogiannis

Farzam Ehsani

Lucy Brickhill

Mandla Magagula

Investec

Chris Becker

Chris Kleb

David Aphane

Graham Nelson

Heinrich Dinkelman

Richard Williams

Nedbank

Charmaine Thiar

Francois Smit

Marijke Guest

Pierre Viljoen

Rene Jacobs

Reynhardt Cronje

Standard Bank

David Sobey

Ilze Prinsloo

Jeandre Engelbrecht

Olimpia da Rocha

Paresh Daya

Peter Train

Poovanthiren Moodley

Sugen Pillay

Credit to Shutterstock for
supplying the images on pages:
10, 13, 14, 17, 36, 40, 49 & 64.

The South African Reserve Bank acknowledges and thanks all the participants who have been part of this journey and have contributed to the success of Project Khokha. The progress achieved as well as the growth in the breadth and depth of skills must be noted, as this is to the benefit to the South African industry as well as to the global body of knowledge.

This is a journey in which business models are being disrupted, mandates are being blurred, behaviours are changing, and the future is unknown. The only way to navigate this journey is by positively contributing to the conversation as a global community and by working together to build the skills, understand the opportunities, and gain insight from the experiences to leverage innovation and maximise benefit to our societies.

This publication has been prepared for general guidance on matters of interest only and does not constitute professional advice. You should not act upon the information contained in this publication without obtaining specific professional advice. No representation or warranty (express or implied) is given as to the accuracy or completeness of the information contained in this publication, and, to the extent permitted by law, PricewaterhouseCoopers Inc., who assisted in authoring the publication, its members, employees and agents do not accept or assume any liability, responsibility or duty of care for any consequences of you or anyone else acting, or refraining to act, in reliance on the information contained in this publication or for any decision based on it.

© South African Reserve Bank

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without fully acknowledging Project Khokha of the South African Reserve Bank as the source. The contents of this publication are intended for general information only and are not intended to serve as financial or other advice. While every precaution is taken to ensure the accuracy of information, the South African Reserve Bank shall not be liable to any person for inaccurate information or opinions contained in this publication.



South African Reserve Bank

www.resbank.co.za