Whitepaper On Distributed Ledger Technology

11 Nov 2016

Commissioned by

HONG KONG MONETARY AUTHORITY
香港金融管理局

1. Foreword . . . 3
2. Executive Summary . . . 5
3. Introduction to distributed ledger technology . . . 9
3.1 Basic building blocks and mode of operations . . . 9
3.2 A detailed walk-through illustration . . . 12
3.3 The Disruptive Properties . . . 16
3.4 Smart contracts . . . 17
3.5 Conclusions . . . 17
4. Technology . . . 18
4.1 Protecting information integrity through design: a “macro” perspective . . . 18
4.2 Protecting information integrity with cryptographic technology: a “micro” perspective . . . 23
4.3 Smart Contracts . . . 27
5. Deployment . . . 28
5.1 Performance . . . 28
5.2 Interoperability . . . 31
6. DLT Platforms . . . 34
6.1 Bitcoin . . . 34
6.2 Ethereum . . . 35
6.3 Hyperledger . . . 35
6.4 Corda . . . 36
6.5 Ripple . . . 36
7. Governance . . . 37
7.1 Rule-Making Processes undertaken by Participants . . . 37
7.2 Defining roles and responsibilities . . . 38
7.3 Conflict Resolution . . . 41
8. Risk Management and Regulatory Compliance . . . 44
8.1 Operational Risks . . . 44
8.2 Identity Theft Risks . . . 44
8.3 Conduct Risks . . . 45
8.4 Regulatory Compliance . . . 45
9. Security and Privacy . . . 47
9.1 Security . . . 47
9.2 Privacy . . . 48
9.3 Challenges . . . 49
10. Legal considerations . . . 54
Distributed Ledger Technology Will Drive Legal Innovation (by Professor Carla L. Reyes) . . . 58

11. Proof-of-Concept Work . . . 59
11.1 Proof-of-Concept – Mortgage Loan Application . . . 59
11.1.1 Sub use case (1): The property valuation . . . 62
Review of the prototype integration experience (by the working group) . . . 65
11.1.2 Sub use case (2): Property ownership verification . . . 67
11.1.3 Sub use case (3): Mortgage count . . . 70
11.1.4 Conclusions – Benefits and challenges . . . 71
11.2 Proof-of-Concept – Trade Finance . . . 75
11.2.1 Smart contract in Open Account Trade . . . 79
11.2.2 Tracking the status of trade transactions: flows of Goods & Funds . . . 80
11.2.3 Information matching of Invoices with Purchase orders . . . 82
11.2.4 Benefits of and possible challenges to DLT . . . 83
11.3 Proof-of-Concept – Digital Identity Management . . . 84
Appendix 1 – Use case – Mortgage-Backed Securities (by R3) . . . 85
Appendix 2 – Blockchain for Trade Finance (by Vishal Batra, IBM Research, India) . . . 87
Technical note (I) – the Merkle tree . . . 90
Technical note (II) – Illustration: Proof-of-work mining using the SHA256 hash algorithm . . . 92
Technical note (III) – Considerations on the Deployment of Platforms and Applications . . . 93
References . . . 96

1. Foreword

To many, "fintech" is a term simply associated with the trendy banking or payment services they use via their smartphone apps, or at the virtual counters of their banks. Internet banking and mobile payment applications are certainly important areas in the application of fintech, but they are far from the only ones. Other technologies, from artificial intelligence to big data analytics to virtual reality, are pushing out the possible frontiers of fintech every day. These technologies could bring a sea change to banking and payment services. The subject of this report – distributed ledger technology (DLT) – is just one key example of this beginning to happen.

DLT is perhaps better known as "blockchain". It is essentially technology that supports networks of databases that enable participants to create, disseminate and store information in a secure and efficient manner. While database technologies are not new, what makes DLT special is that these networks of databases can operate smoothly and securely without necessarily being controlled and administered by a central party that is known and trusted by every participant.

The potential applications of DLT, as the fintech industry and many central banks and regulatory authorities soon found, are not limited to dealing in virtual currencies or commodities. The very fact that DLT allows information or records to be transferred and updated by network participants, and this to be done in a trustworthy, secure and efficient way, carries enormous potential. However, while the value proposition of DLT is gradually materialising, the use of DLT in financial services is also introducing new risks and giving rise to new legal and governance issues. These require in-depth study before its full potential can be realised. As a regulatory authority, we need to have a thorough understanding of the various governance, risk management and legal issues associated with DLT before its wider use begins in earnest.

The Research White Paper

In this connection, the Fintech Facilitation Office (FFO) of the Hong Kong Monetary Authority (HKMA) has commissioned the Hong Kong Applied Science and Technology Research Institute (ASTRI) to conduct a research project on the subject of DLT. The key objectives of this project are to carry out an open-minded and an in-depth examination of the technology (including an investigation into its potential, its risks, and its regulatory implications); and to identify possible applications of DLT to banking services by engaging in proof-of-concept work.

This white paper may be regarded as the first stage in this larger research project. It aims to provide the fintech industry in Hong Kong with a reasonably comprehensive study of the key features, benefits, risks and potential of DLT. It also includes the initial findings of the proof-of-concept work carried out on DLT applications in three areas: mortgage loan application, trade finance, and digital identity management. The next stage of this project will deliver more detailed findings from the proof-of-concept work, along with discussion on whether some of this work can be put into action. It will also address the regulatory implications of DLT, and the general control principles for DLT for the banking and payment industry. We plan to deliver the next set of results in the form of another white paper in the second half of 2017.

This white paper could not have been completed without the help of many people. The experts at ASTRI have played a key role as authors and project manager. My special thanks also go to the many industry experts who have contributed thematic articles discussing very specific and relevant issues surrounding the use of DLT. I am grateful for the efforts of participating banks and other industry players, including HSBC, Standard Chartered Bank, the Bank of China (Hong Kong), the Hang Seng Bank, and the Bank of East Asia. They have all been extremely helpful in sharing their experience, insights, and honest assessments in relation to the use of DLT in their businesses.

All these efforts, I believe, will not only make this research project a unique contribution to the debates taking place around the world on the future development of DLT, but will also form a solid basis on which the HKMA and our banking sector can deliberate on how best to put this technology to use.

Howard Lee
Senior Executive Director, Hong Kong Monetary Authority

2. Executive Summary

Financial Technology (fintech), a term often used to refer to newly emerging digital technologies adopted in the finance industry, is said to disrupt traditional banking models to bring about greater convenience and efficiency for consumers of financial services, as well as offering the possibility of reducing risk and lowering the cost of operations for financial service providers. A number of major financial markets, including those of London, New York, Hong Kong and Singapore, have tried to establish a sustainable fintech ecosystem and attract fintech talents in order to maintain their competitiveness.

To support the development of the fintech industry and maintain Hong Kong's status as a leading international financial centre, the Hong Kong Monetary Authority (HKMA) set up a Fintech Facilitation Office (FFO) in March 2016. As part of its major functions, FFO is tasked with initiating banking and payment industry research into the potential application of novel fintech solutions which could have a significant impact on banking and payment services.

2.1 The Purpose of this White Paper

Against this background, the HKMA, through the FFO, has commissioned the Hong Kong Applied Science and Technology Research Institute (ASTRI) to conduct a research project on the subject of a much talked-about fintech topic-distributed ledger technology (DLT), of which a well-known example is "blockchain". A key objective of this project is to undertake an open-minded and in-depth examination of the technology, and identify its potential and its risks. This white paper is the first deliverable of this research project. It aims to provide the fintech industry in Hong Kong with a reasonably comprehensive study of the key features, benefits, potential and risks associated with DLT.

2.2 Distributed Ledger Technology

DLT is built upon a series of networks of databases that allow participants to create, disseminate and store information in an efficient and secure manner. These networks of databases can operate smoothly and securely without the need for any central party or central administrator that every participant knows and trusts. At the same time, these networks make constantly available for examination a full audit trail of information history, which can be traced back to the moment when a piece of information was created for the first time. Furthermore, unauthorised changes to the information and its history are very difficult, if not impossible, to make. In other words, DLT operations are designed in such a way that information stored and communicated through the networks has a high level of trustworthiness, and every participant in the network can get simultaneous access to a common view of the information.

Structurally speaking, a blockchain may be considered as a series of blocks of information that are securely chained together. Any given digital record of an asset, be it a copy of the title deeds of a bricks-and-mortar property or a virtual commodity, can be stored in a block. New blocks are formed whenever participants create a piece of new information or change an existing piece of information about an asset, for example by entering transaction records, changes of status, new market prices, or new owners. All blocks newly formed after the first block are securely chained to the previous one, thus ensuring their authenticity and creating a trustworthy audit trail. In fact, one of the earlier uses of DLT was in the area of virtual commodities (e.g. Bitcoin), for which change in the ownership of a commodity is recorded in the blockchain.

DLT design clearly enjoys advantages over some traditional technologies. However, such evolving technology also brings possible risks if issues of governance, deployment, risk management, and regulatory compliance, along with the legal implications (as set out below), are not adequately taken into account.

2.3 Governance

Despite its decentralised approach, DLT still requires a set of common rules by which all participants operate in order to ensure its accuracy and trustworthiness. The decentralised model poses some challenges when there is a need to make changes to or update the rules, because such changes need to be agreed upon and accepted by all participants in order for DLT to function consistently.

A governance framework is therefore important for the implementation and sustainable operation of DLT. This framework needs to take into consideration oversight and monitoring functions, rule setting, and acceptance and change control management.

2.4 DLT Platforms

There are two main categories of DLT platforms: unpermissioned and permissioned. The former type is maintained by public nodes, and is accessible to anyone. Bitcoin is a well-known example of an unpermissioned platform. It has operated as a digital asset and payment system since 2008, and its example has significantly accelerated the development of DLT platform design. The latter type, which has as an example the Corda platform, only involves authorised nodes and thus facilitates faster, more secure and more cost-effective transactions. DLT platforms in each category have their own unique features. Some are designed for specific types of applications, and others for general use. For instance, the sharing of individual ledger data in Corda is limited to parties with a legitimate need to know, which is not the case on other platforms.

2.5 Deployment

The ability of a new technology to succeed depends on whether it can be practically implemented. In the case of DLT, the performance and computing resources required for transaction processing, validation and fraud detection have an impact on which financial services it can best be applied to. In addition, the efforts needed to ensure interoperability between different DLT networks, between ledgers within the same network, and with other non-DLT systems, should not be underestimated, and also need to be considered carefully before deployment.

2.6 Risk Management and Regulatory Compliance

Any introduction of new technology inevitably introduces new types of risk, and DLT is no different. Even when an asset owned by a participant is protected by the participant's digital certificate, and no changes can be made to the information without the correct digital signature, certain traditional cybersecurity issues still apply to DLT. For example, denial of access attacks and other cyber attacks may still be launched against DLT in an effort to cause its operation to fail.

Due to the anonymous nature of participants in some DLT applications (in particular Bitcoin), DLT is sometimes seen as being associated with issues of money laundering and the sale of illegal goods, and as supporting the ransomware payment model. Although these issues may largely be addressed when DLT is implemented in a "permissioned" network (which only authorised and authenticated participants may join), this kind of solution still needs to be examined in detail.

Regardless of whether a DLT platform is operating in "permissioned" mode or not, there are personal data privacy issues that need to be addressed when information concerning individuals is stored in DLT. For example, as information stored in DLT cannot be altered or deleted once added, any application will need to address how to comply with the data protection principles of accuracy and an individual's right of correcting data. In addition, some DLT applications may be implemented across various jurisdictions without a single entity responsible for their running, so issues relating to cross-border data flow, legal enforceability, liability, dispute resolution, discovery and extraterritorial reach need to be addressed too.

Although at this stage it is not the focus of this white paper to explore fully all these complex legal and regulatory issues, the paper identifies these potential issues and calls for further study in these areas in the next stage of this research project, with likely input from the legal community.

2.7 Potential Applications

As the fintech industry and many central banks and regulatory authorities agree, DLT has a wide range of potential applicability to many banking and payment services, such as cryptocurrencies, post-trade settlements, record checking and management, and cross-border fund transfers. The fact that DLT allows information or records to be transferred and updated by participants of the networks (who can be total strangers to one another) and this to be done in a highly trustworthy, secure and efficient way carries enormous potential. DLT is even more appealing as a possible replacement for existing processes in which important information needs to be communicated and stored in a highly secure manner, but which are currently largely manual, labour-intensive and paper-based.

2.8 Proof-of-Concept Work

During the preparation of this white paper, a number of banks and industry players have participated in this project by looking more closely at "proof-of-concept" work involving the use of DLT in actual banking businesses. The following three areas have been identified at this stage in which DLT could play a useful role –

1. Mortgage Loan Applications: Banks need fast and accurate information about the estimated value of a property in order to make good credit decisions. However, communication between banks, law firms and valuation firms remains a largely paper-based and (sometimes) error-prone process. A DLT network that connects these participants could therefore be helpful for them, for example by enabling them to confidently share copies of digitised valuation reports and legal documents or even transfer titles, thus reducing the time and cost of transactions.

2. Trade Finance: This is another key banking business which involves paper-intensive processes. With digitised documents, DLT could help improve the efficiency and accuracy of the workflow by making the entire transaction history and its collateral information more transparent. More importantly, it could help reduce the risk of fraud through the use of forged documents and the double or multiple presentations of invoices.

3. Digital Identity Management: Existing "know your customer" (KYC) requirements and customer authentication processes are very manually intensive, and require significant resources from banks to ensure regulatory compliance. In addition, the current manually intensive procedures can be inconvenient for customers and lead to undesirable user experiences. In this connection, an attempt is made here, through the performance of proof-of-concept work on a DLT network, to implement a digital identity management platform that could automate some of the KYC requirements and the customer authentication process.

This proof-of-concept work has been developing at different rates. Most advanced is work on the mortgage loan application, for which testing is now being carried out. More details about the mortgage loan application proof-of-concept work, including a detailed operating model and a discussion of possible issues to be resolved, are given later in this white paper.

2.9 Other Reference Materials

Throughout this white paper, expert opinions from academia and industry have been included to provide a range of different details and viewpoints on the potential benefits and challenges relating to DLT and its applications.

2.10 Ways Forward

More in-depth findings from the proof-of-concept work, and discussion about whether some of this work can be put into practice, will be the focus of the next phase of this project, which is expected to be delivered in the form of another white paper in the second half of 2017. The second white paper will also cover the regulatory implications of DLT, and explore general control principles for DLT for the banking and payment industries.

Ultimately, it is hoped that this white paper will provide a better understanding of DLT, and of how its potential applications could benefit both customers and banks: customers, by providing better banking services, and banks, by helping them provide services of greater security, quality and efficiency. This in turn, it is hoped, will assist in maintaining and improving the stability of Hong Kong's banking and financial sectors.

3. Introduction to distributed ledger technology

DLT is a protocol for building a replicated and shared record ledger system. Such a system may be used to record a wide range of items, such as asset ownership, asset transfer transactions, and contract agreements. While its ledger function is similar to that of a conventional paper-based or electronic-based ledger system, its capabilities go much further. It provides a new way of constructing a secure record system that offers stakeholders more transparency, and that encourages member participation in its operations.

DLT came onto the scene with the introduction of Bitcoin. Bitcoin, one of the best-known applications of DLT, was introduced by the pseudonymous Satoshi Nakamoto in 2008[1]. Bitcoin is a digital asset that utilises a peer-to-peer payment system, enabling transactions to take place between users directly without the need of a centralised authority to control and administer the system. All the transactions are verified by network nodes, and recorded in a globally distributed database.

Traditionally, a payment transaction requires a centralised authority acting as an intermediary (e.g. a bank or clearing house) to prove that the payer or paying bank has sufficient money, as well as to process the money and transfer it between accounts. Unlike these conventional transactions, Bitcoin distributes the DLT maintenance responsibility to the whole network. So long as the validating nodes verify transactions and publish them to the ledgers, each transaction is added to the distributed database, rendering its status as validated. Hence, no centralised authority is required to manage, control or authorise transactions between participants.

3.1 Basic building blocks and mode of operations

I. What is a "ledger"?

A ledger, according to the Oxford Dictionary, is "a book or other collection of financial accounts". Ledgers have existed for thousands of years, first arising when people started trading goods and services and needed to keep records of transactions. Today, the conventional ledger system is often a centralised system that is maintained inside the information system infrastructure of an organisation.

One example of a modern-day "ledger" familiar to most people is a bank account record, in which every debit or credit transaction of a bank customer is maintained. Importantly, bank customers trust their banks to have the capability of maintaining their banking records (i.e. the bank account information in the ledger) safely and securely.

II. What is a "distributed ledger"?

A "distributed ledger" system, unlike a conventional ledger system, is collectively maintained by all the participants of that system, rather than by one central party (e.g. a bank or a clearing house). Each participant is considered to be a "node" of the distributed ledger system. Essentially the nodes are the computers of individual participants, which each contain a complete set of transaction records. Together, the nodes participate in building and maintaining the distributed ledger. Since a "local" copy of the same ledger is maintained and developed in each node (instead of being centrally controlled and administered by a certain party), the system is known as a "distributed" ledger system.

(a) How is a distributed ledger updated?

Similarly to a conventional ledger, a distributed ledger is updated whenever a transaction takes place. However, instead of the previous record being overwritten (as in a conventional ledger), the transaction information is exchanged between nodes (for example, between two system participants) and added as a new ledger entry.

In the absence of a trusted central party, the process of updating the distributed ledger relies on a process for achieving consensus among the nodes (or "distributed consensus") regarding all new information added to the ledger. Achieving "distributed consensus" in turn requires two important processes to take place: validation of each transaction, and the "broadcast" of the validated result to all the other nodes of the distributed ledger.

- "Validation" – The nodes together determine whether or not new entries in a transaction block are valid, as well as whether the transaction block can be admitted to the ledger. Specifically, the participants (the nodes) are required to perform validation of every transaction in the block to ensure that its contents are legitimate. For example, they must verify that the sender of a transaction is the true owner of the asset being sold. For transactions containing a contract execution instruction, validating nodes will also execute the instruction that has been received and confirmed by the consensus process.

- "Broadcast & Consensus" – This is the process that enables validating nodes to reach a consistent view of the new entry in the distributed ledger. It begins when a validating node has validated one or more transactions and initiates the process of adding them to the ledger. The validating node first broadcasts information about the new block to the other validating nodes. The other validating nodes may have also validated the same set or different sets of transactions, but the consensus process allows them to communicate among themselves and agree on a common set of validated transactions to be added to the ledger.

(b) Mining – an important but resource-intensive task in the validation process

As mentioned above, there is a need to achieve distributed consensus in an "open" form of a distributed ledger, i.e. one in which anyone can contribute data to the ledger and no one can claim control as the central trusted authority (commonly known as an "unpermissioned ledger" or a "permissionless ledger").

One important way of doing this is by the so-called "proof-of-work mining" process. This process involves all the validating nodes competing to perform a computationally demanding calculation. The first node to solve the computational problem then helps to build a transaction block.

The "mining" process, however, has at least two problems:

- First, the mining process requires significant computing resources to perform the calculations. Therefore, participants need to have incentives for investing the resources needed to participate in the mining activity and, ultimately, to maintain the ledger. (For example, in the case of Bitcoin, the incentive for being the first node to solve the algorithm problem (and therefore successfully build a transaction block) is a "reward" of a certain number of Bitcoins.)

- Second, the mining process normally takes time for complex problems to be solved. While a validating node is busy doing demanding computations on a set of validated transactions in an attempt to add them to the block, it cannot in parallel process another set of new transactions into another new block. Hence, the mining process slows down the prompt processing of transactions.

(c) Fashioning a distributed ledger with a central trusted party

In view of the problems arising from the mining process in an unpermissioned ledger, another type of distributed ledger has been developed.

Commonly known as a "permissioned" ledger or "private ledger", this type of ledger may be owned, controlled and managed by a central trusted party or a group of participants in the form of a consortium. Only trusted or vetted participants are allowed to participate in the control and maintenance of permissioned ledgers. Distributed identical copies of ledgers are kept by all participants. This more controlled sharing of ledgers among registered or authorised participants can be used to support an industry-level record system that keeps track of asset ownership, the movement of confidential documents, the status of settlements, and other transactions.

An important advantage of permissioned ledgers over unpermissioned ones is that the validation process does not involve the computationally intensive mining process, which consumes large amounts of both electrical energy and computing resources. The validating nodes simply check the validity of a transaction without needing to perform the mining task. This means the ledgers can be updated in a much faster and more energy-efficient manner. By restricting ledger management to trusted participants only and by reducing the number of labour-intensive and duplicated processes required, permissioned ledgers also benefit from a relatively lower risk of cyber attacks and security breaches, as well as lower operating costs (e.g. fewer computing resources are required).

3.2 A detailed walk-through illustration

A transaction in a conventional centralised database or system

To better understand the properties of a DLT network, we illustrate below its potential application for handling a property transaction including the subsequent change of property title.

Suppose Alice is selling a flat to Bob for HK$ 10 million. The key events which occur with regard to the change of property title are as follows: After signing a sale and purchase agreement, Bob obtains a mortgage from Bank B. Bank B then transfers the funds on behalf of Bob to Bank A on the agreed date. Following that, the change of the ownership title of the property is submitted for registration at the Land Registry.

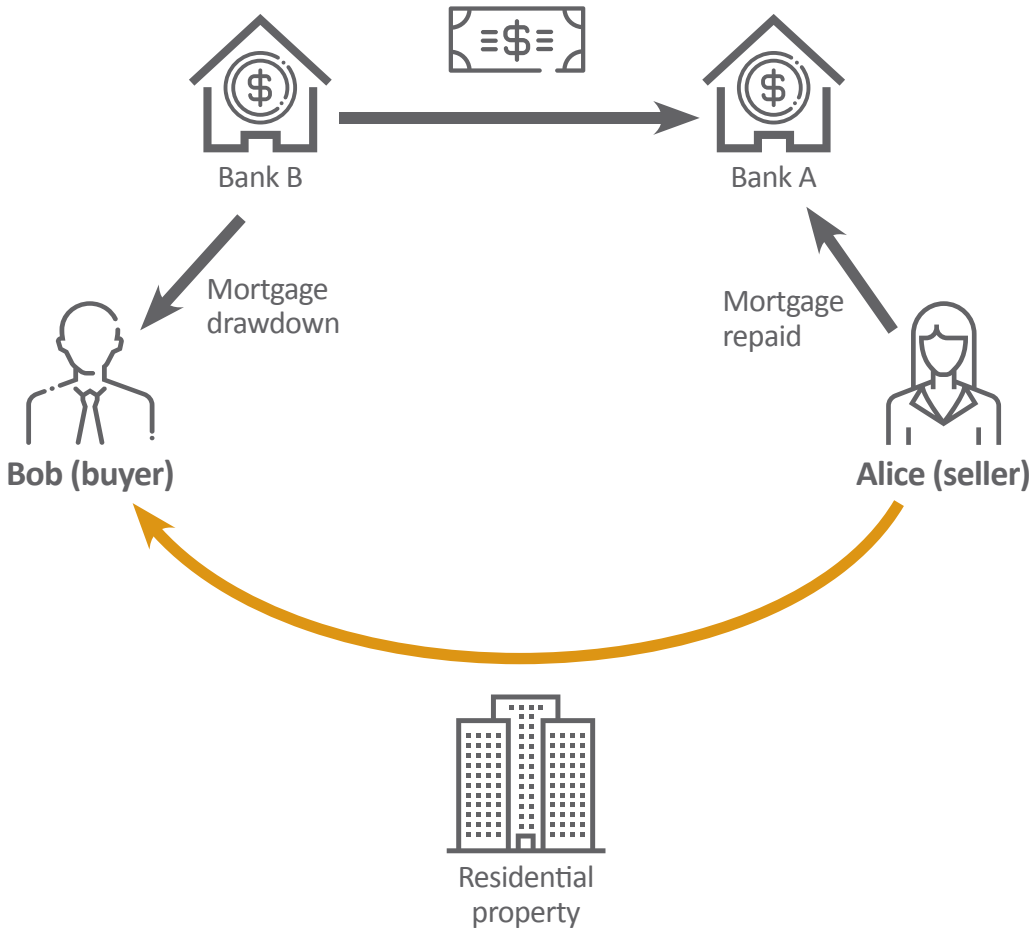

The property transaction

A real-life example would involve more parties, such as a surveyor to provide a property valuation assessment to Bank B before Bob could obtain the mortgage, and a solicitor to handle all legal documentation.

The above example is therefore a simplified but typical property transaction which is based on trust in reliable centralised authorities – specifically, a government department that assures us that the legal title of a property has been transferred, and a bank that confirms that our money has been transferred to the designated account.

DLT provides an alternative way of building trust through transparency and consensus, with participants engaging in a cooperative consensus process to construct the distributed ledger, and record and verify every entry in the ledger.

How would the transaction work in DLT?

At present, the Land Registry holds transactional data and other relevant information about the flat. This is all stored in a conventional centralised database, which can be accessed by different parties (e.g. Alice, Bob and the banks) for a fee.

In future, it may be possible for this process to take place in a DLT network designed to keep track of property ownership entitlement, and in which the transactional data and other relevant information about the flat is stored.

To keep things simple, we will assume here that a permissioned DLT network is used, since it is likely that only a limited group of trusted participants will join such a network (e.g. the Land Registry and the banks). In such a network, these participants will be the nodes.

Bank A (as a node) will create a transaction record containing a set of information (e.g. the personal particulars of Alice and Bob, the transaction date, the address and the price of the property), along with the digital signature of the seller (needed for signing the electronic record).

The digital signature is crucial for the transaction. It is a mathematical scheme adopted to prove two core elements of the transaction: the sender's authenticity, and the integrity of the information. The technology employed here is asymmetric cryptography, which provides the required level of security for creating and sending the transaction. This concept will be discussed further in chapter four.

Transaction broadcast and validation by network nodes

At this point, the transaction (with the signature appended) is broadcast by Bank A to all other nodes (i.e. the Land Registry and the other participating banks) so that the transaction can be validated by any one of them.

As each node holds a local copy of the ledger that contains a complete set of the historical transactional data records of that property, the node is able to look into its own chain and record history to check the validity of the transaction, i.e. to ascertain whether Alice genuinely holds the title to the property.

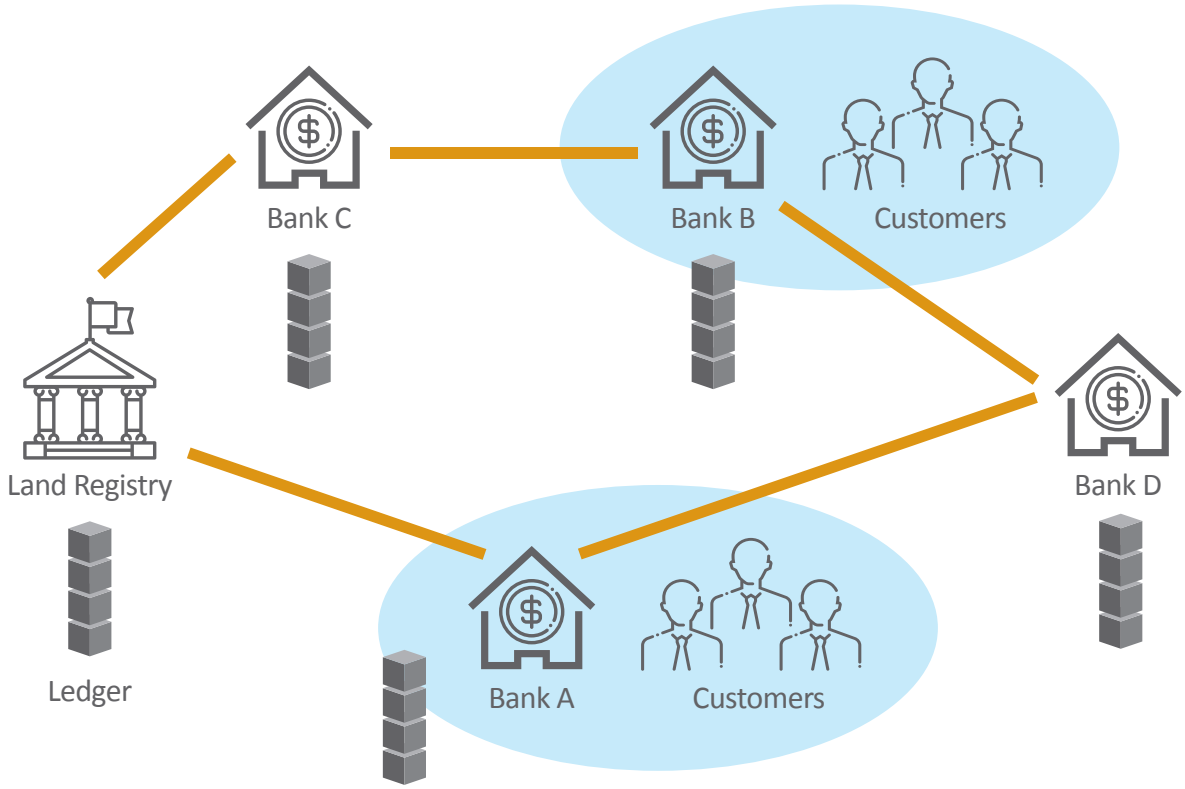

A DLT system in which the Land Registry and registered banks maintain replicated copies of the ledger.

While this particular transaction is being validated by nodes of the DLT network, it is likely there will be other transactions taking place at the same time. For example, Cathy may be buying a property from David, and settling payment to Bank D via Bank C.

Nodes will then group all the newly created transactions which have not yet been recorded into the permissioned ledger, including the transaction created by Alice, and compile these transactions into a "transaction block".

New transaction block

Once the transaction block has been compiled by a validating node, it then broadcasts it to all other nodes within the DLT network. The other validating nodes may have also validated the same set or different sets of transactions, and the consensus process allows them to communicate among themselves and agree on a common set of validated transactions to be added to the ledger. Then the process starts again.

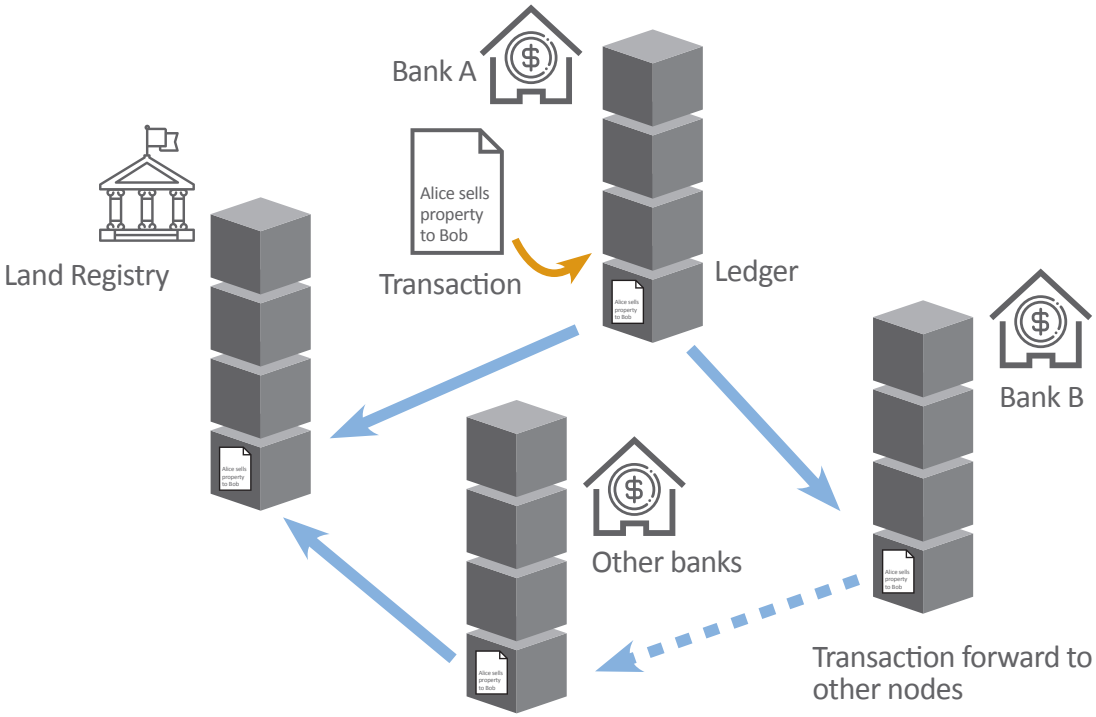

A DLT system in which Bank A sends an "Alice sells property to Bob" transaction to other nodes, which then append the transaction to their copies of the ledger.

The settlement of the cash payment

What about payment? This depends on the design of the DLT network.

Since the DLT network is simply a distributed ledger system containing data records, no physical assets are transferred. Only new records or entries are added to the DLT network. Payment for the property transaction can therefore either be settled outside the DLT network (i.e. through a separate payment system), or through the same DLT network. In the latter case, the payment will be recorded as a movement of digital currency between the participants. In other words, the property title transfer and the related payment can in theory be done in real time on a "delivery versus payment" basis.

3.3 The Disruptive Properties

DLT is tamper-proof

The distributed ledger system is well known for its tamper-proof nature, which in turn facilitates the building of trust among participants with regard to the integrity of the system. Its tamper-proof nature is achieved by two elements: a proof system, and cryptographic technologies.

Conceptually, DLT is a chain of blocks. Each block contains a set of entries known as transactions. New entries are collected into a new block by validating nodes which then add the new block to the chain. More blocks are created as more new entries are collected, and the chain grows in length.

DLT ensures that this chain is very difficult to tamper with. Any attempt to modify the chain requires the perpetrator to present proof of the authenticity of the modification. Generating such proof involves performing non-trivial cryptographic operations which are both lengthy and costly. In addition, fabricated proofs can be easily detected by other validating nodes within the DLT network.

The blocks inside the chain are connected together with links that are built with a hash function. The link of a block is integrally tied to its content. Any attempt to change the content of a block in the chain causes the value of its hash link to be changed also. This breaks the chain, with the remaining chain being much shorter than the original one. It is a situation that can be immediately detected by other participants, who will invariably reject the changes and continue using the original chain.

A hash function is a one-way mathematical function that turns data into a trunk of random characters called hash. Changes to the data, no matter how slight, change the hash dramatically and in an unpredictable way. It then becomes impossible to derive the original data from the hash.

DLT is immutable and transparent

The distributed ledger system is transparent in that all transactions are public, traceable, and permanently stored in the DLT network. While a private DLT network may add access restrictions to transactions, it preserves the feature of stakeholders having common access to their set of common transactions.

The moment anyone starts trading on the DLT system, a history of all their transactions starts to be logged in the system. This history is permanently recorded, unalterable, and accessible either to the public or to stakeholders. This high level of transparency and reliability is an important factor in building trust in the integrity of the network.

3.4 Smart contracts

As the technology has evolved, 'smart contracts' have emerged that have added further versatility to DLT. Participants are allowed to enter self-drafted agreements (i.e. smart contracts) and embed them in the records of the DLT network. Such contracts are developed in computer code, enabling DLT to execute them automatically and in precise conformity with the contract terms. Triggering events can be designed and built into the smart contracts to activate certain actions when specified events occur or certain data is received. A typical example involves a payment being triggered when a specified date is reached.

This study identifies a number of possible sub use cases for proof-of-concept work, and the smart contract concept will be included in some of these sub use cases to help us gain a better understanding of the potential of smart contracts.

3.5 Conclusions

A DLT system clearly has the potential to bring new opportunities and efficiencies to the banking and payment industries, based on the key strengths set out in this chapter. These include the capability of establishing trust in a distributed system, efficiency in broadcasting information in a speedy and secure way, the ability to achieve complete traceability of records and transactions, the possibility of lowering operation costs, and the potential for high resiliency. However, before concluding that DLT represents the best solution for all banking and payments issues, more work needs to be done to determine whether the existing DLT is mature enough to fulfil the requirements of the financial community, and to ascertain what key attributes or requirements DLT needs to incorporate in order to be widely and comfortably adopted by the banking and payment industries.

As DLT is still evolving, new and more innovative operating models continue to be introduced and tested. One example is R3's Corda DLT, which builds a distributed ledger without using the blockchain as its building block. Developments like this may offer additional DLT options in terms of providing potential operating models.

This paper will now go on to examine the technology associated with DLT, and identify possible issues associated with it. It will also present sub use cases to demonstrate the potential of DLT.

4. Technology

The previous chapter provided an overview of DLT. This chapter aims to deliver an in-depth study of the underlying technology and security design of DLT. In a distributed network, information is created, transmitted and stored in a way that renders unauthorised changes hard to make by, for example, a dishonest participant in the network or an unauthorised party. Its robustness in protecting the integrity of information is due (at a macro level) to its high-level design, and (at a micro level) to the detailed technicalities and specific security arrangements involved in managing and communicating information.

In the following paragraphs, we discuss how a distributed network protects information integrity; first from a "macro" point of view, by focusing on its general design, and then, at a "micro" level, by looking at how various technicalities and security arrangements work in certain key processes.

4.1 Protecting information integrity through design: a "macro" perspective

A distributed ledger is essentially a decentralised database in which information is replicated in multiple locations that are connected in the network. Each copy of the replicated database is managed and constructed by one of the participants.

These participants communicate actively based on a set of established procedures, or a "protocol", by which they (a) exchange transactions and (b) reach unanimous agreement that a transaction can be added to the ledger, and agree on the order in which the transaction is added. These participants are known as "miners" (in unpermissioned networks), or "validating nodes" (in permissioned networks).

A "distributed" network

The participants are connected through a distributed network. A common way of interconnecting participants is through a "peer-to-peer" network, known as P2P. A P2P network can be vast and cover a large geographical area. Each connected computer is called a "peer" node. A node joins the network by connecting to one of the well-known peers, whose information has been made public. It then learns about other peers through the information received from this well-known node. At the same time, other nodes learn about this new node.

A P2P network can cover a large area and include many peers. A single peer is connected to a few other peers, who themselves are connected to other peers in the network. When a peer sends a transaction to his or her peers, those peers forward the transaction to other peers, who in turn forward it to other peers again. In the end, all peers in the network receive the information relating to that transaction (see the diagram below for an illustration).

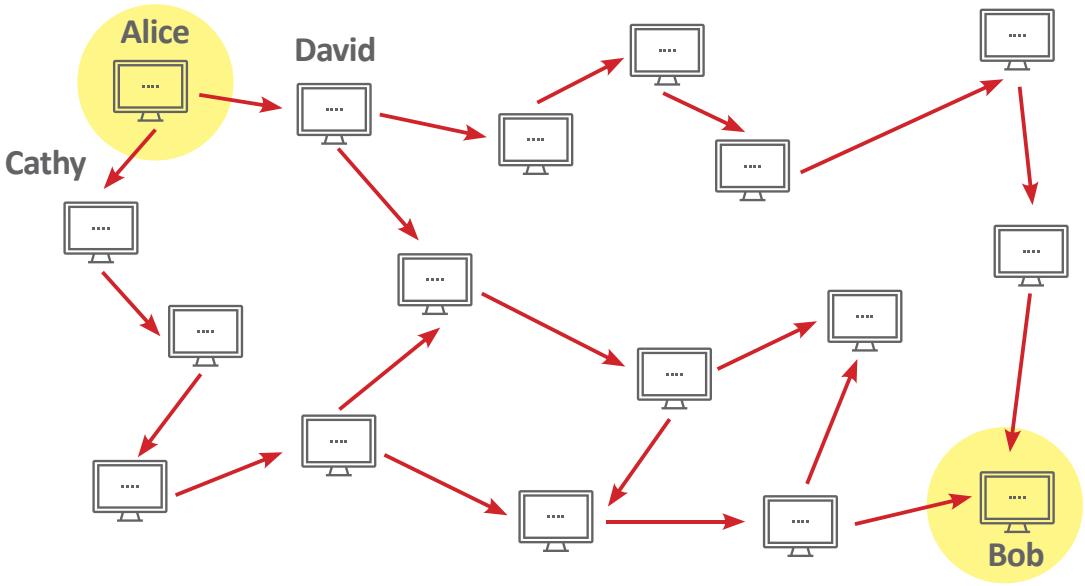

Example of a Peer-to-Peer Network (P2P)

1. Nodes are connected to a few other peers, who are connected to other peers again.
2. If Alice wants to send a message to all nodes in the network, she starts by sending it to her immediate peers, i.e. Cathy and David.
3. Cathy and David forward the message to their immediate peers

Illustration: Peer-to-Peer Network and Packet Forwarding

Data records in a chain of blocks

A ledger is organised as a chain of "blocks" of information. Each block contains a collection of transactions. New transactions are collected to form a new block, which is appended to the ledger. Hence, all transactions are immutably stored in that ledger, which is replicated among all validating nodes.

The first block in the ledger is called the "Genesis block". The next new block is appended to this block, thus forming a chain. As more new blocks are appended to the chain, the chain gets longer. Each block has a unique identifier and is usually represented by a "hash" value. Each block points to its immediate upstream block (see the simplified "chain of blocks" diagram below). A block's location is described according to its position in the ledger.