ASTRI

Whitepaper On Distributed Ledger Technology

11 Nov 2016

Commissioned by

HONG KONG MONETARY AUTHORITY
香港金融管理局

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.

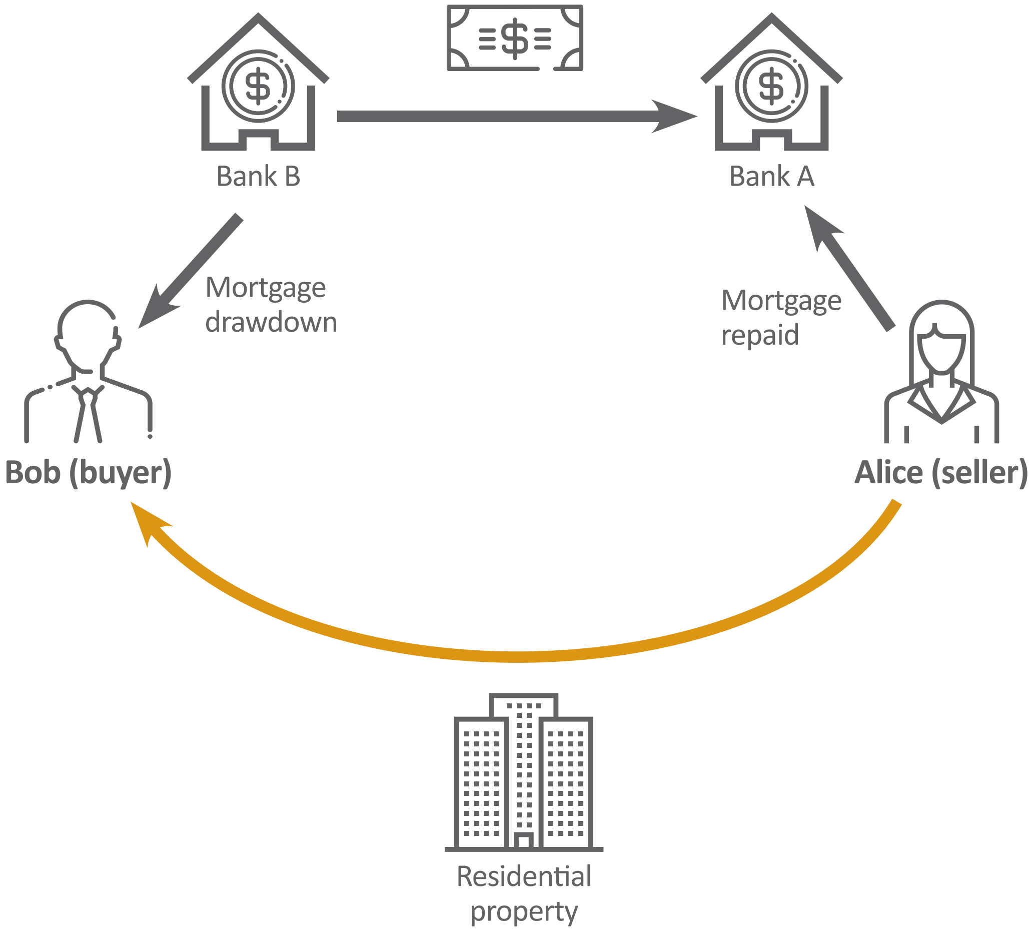

The property transaction

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.

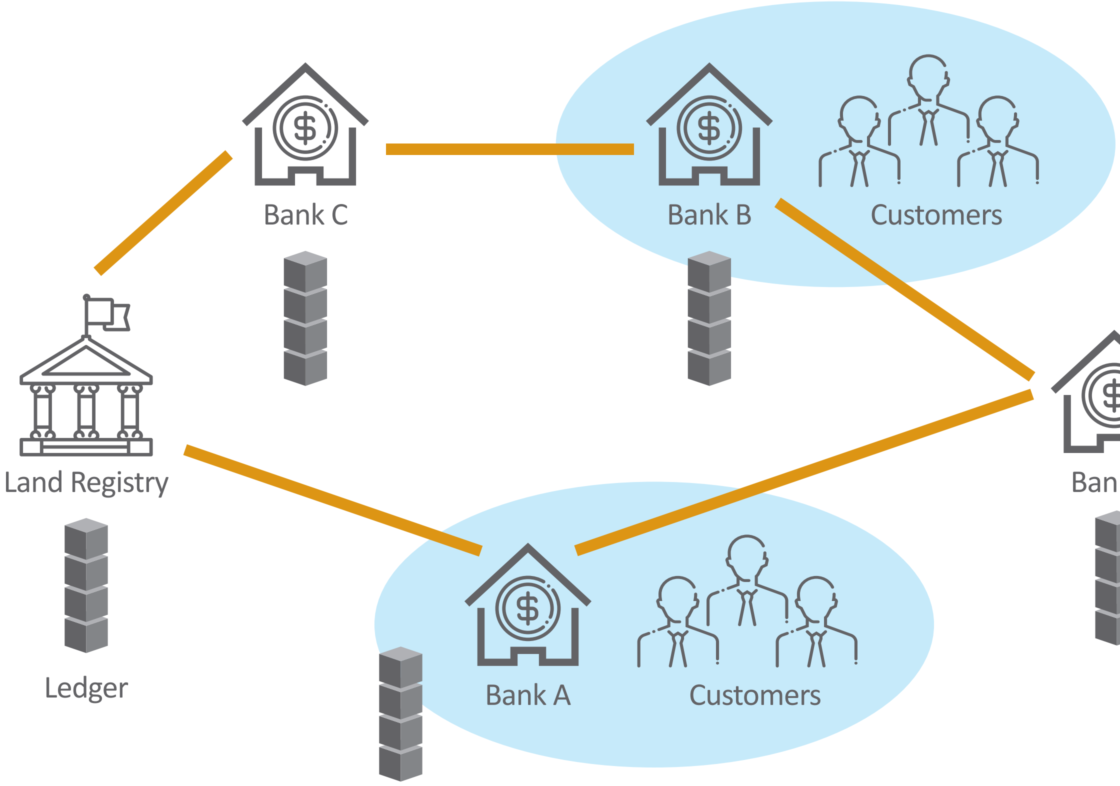

A DLT system in which the Land Registry and registered banks maintain replicated copies of the ledger.

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.

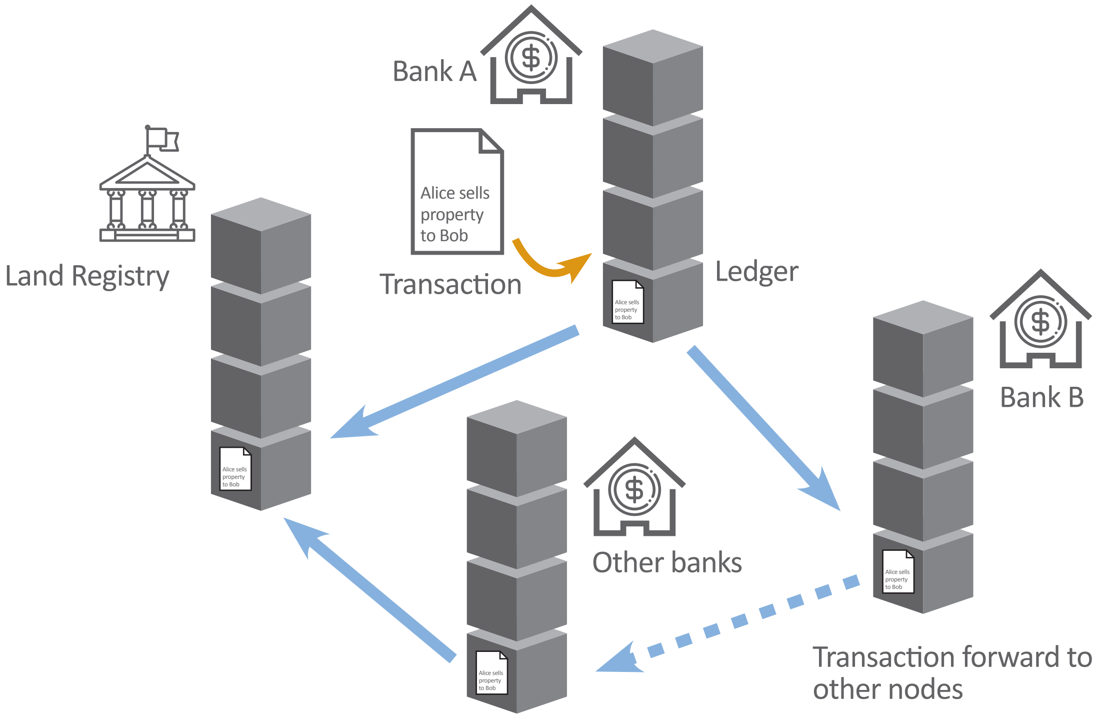
Bank A
Alice sells property to Bob
Transaction
Ledger
Land Registry
Bank B
Other banks
Transaction forward to other nodes

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).

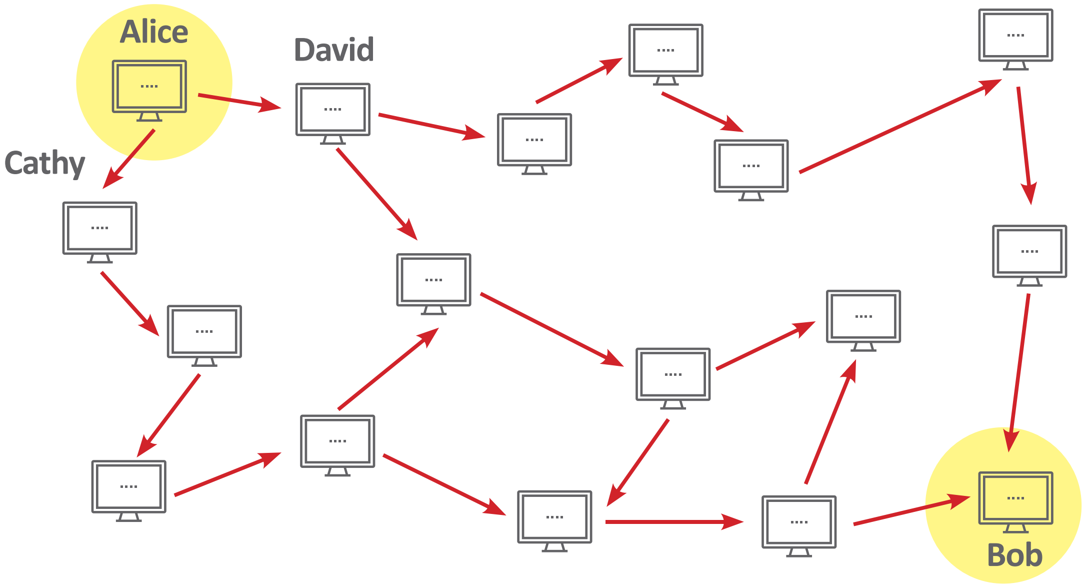

Example of a Peer-to-Peer Network (P2P)

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.

A Merkle tree is used to represent the set of transactions within a block. This enhances transaction searches and the verification of the chain of blocks, since each block may contain many transactions. Please refer to technical note (I) for further details.

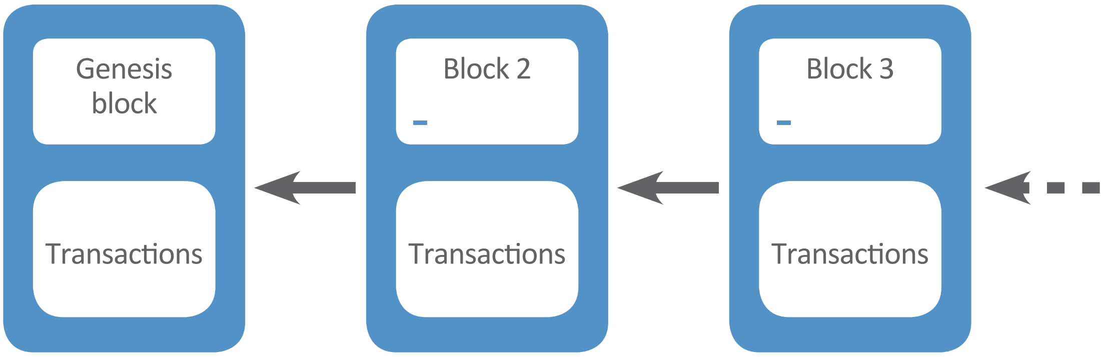

A simplified "chain of blocks" construction

The "blocks" of information: creation and addition

A key role of the miners, or the validating nodes, is to collect information about new transactions, put this information into a new "transaction block", and append it to the ledger.

An important feature of DLT is that this process of adding new transactions to the ledger needs to go through a consensus and validation mechanism, which helps prevent malicious parties from adding falsified transactions and blocks.

Let us look into how this is done in unpermissioned networks and permissioned networks respectively.

1. Unpermissioned networks

In an unpermissioned network, the miner must solve a mathematically difficult problem on the new block before being able to add the block to the ledger. The miner's solution is then forwarded to other miners and is verified by them before being accepted into their own copies of the ledger.

The mechanism for solving the difficult problem is known as "proof-of-work". A large amount of computational resources are needed to find the solution. Such resources may be understood in terms of power consumption or hardware resources, such as size of physical memory or special hardware logic. The first miner to succeed in solving the problem receives a reward in return for the efforts he or she has made. This behaviour, carried out in the expectation of receiving a reward, may be the reason why the process has come to be called "mining" and the parties involved "miners".

If more than one miner solves the mathematical problem for the same block, or generates a block with a different content, the differences are resolved through a consensus protocol, which is a set of rules allowing the parties involved to reach an agreement. Once the process of reaching an agreement is complete, only the block agreed upon is added to the ledger, after which all the replicated copies of ledgers in the network are identical again.

The process of reaching a consensus is also designed to prevent dishonest miners from adding blocks containing malicious transactions. If a dishonest miner attempts to alter the transaction history in such a way that he could wrongfully claim to own an asset which does not belong to him, the miner would have to replace a block in the middle of the ledger with a malicious block containing false transactions. However, this would then require the miner to re-create the rest of the blocks behind that block as well, or else his ledger would be shorter than the current one and would be automatically rejected by all other miners during the consensus process.

As explained above, if a miner was to attempt to re-create all such trailing blocks, he would have to solve all the mathematically difficult problems needed for creating those blocks. The longer the chain of blocks after it, the harder this task becomes. Given a situation where all miners are using large amounts of computational resources to carry out their mining, sometimes to their maximum capabilities, it would be difficult (if not impossible) for a dishonest miner to make this happen.

2. Permissioned networks

In a permissioned network, the validating nodes are trusted and therefore do not need to solve the mathematically difficult problem. Instead, they need only present proof of their trusted identity in order to participate in the consensus process. During the consensus process, each trusted validating node collects new transactions and exchanges its collection of these new transactions with other trusted validating nodes. Once the consensus process is complete, all the trusted validating nodes have the same new block containing the same set of new transactions.

Reaching final consensus on new transaction blocks

After generating a new block, the miners/validating nodes perform the consensus process to ensure that all the nodes see the same set of new blocks and also add them to the DLT network in the same order. As a result, all the replicated copies remain identical after the new blocks and transactions have been added.

The example of the property transaction given in chapter three is used in the diagram below to illustrate the technological design of DLT.

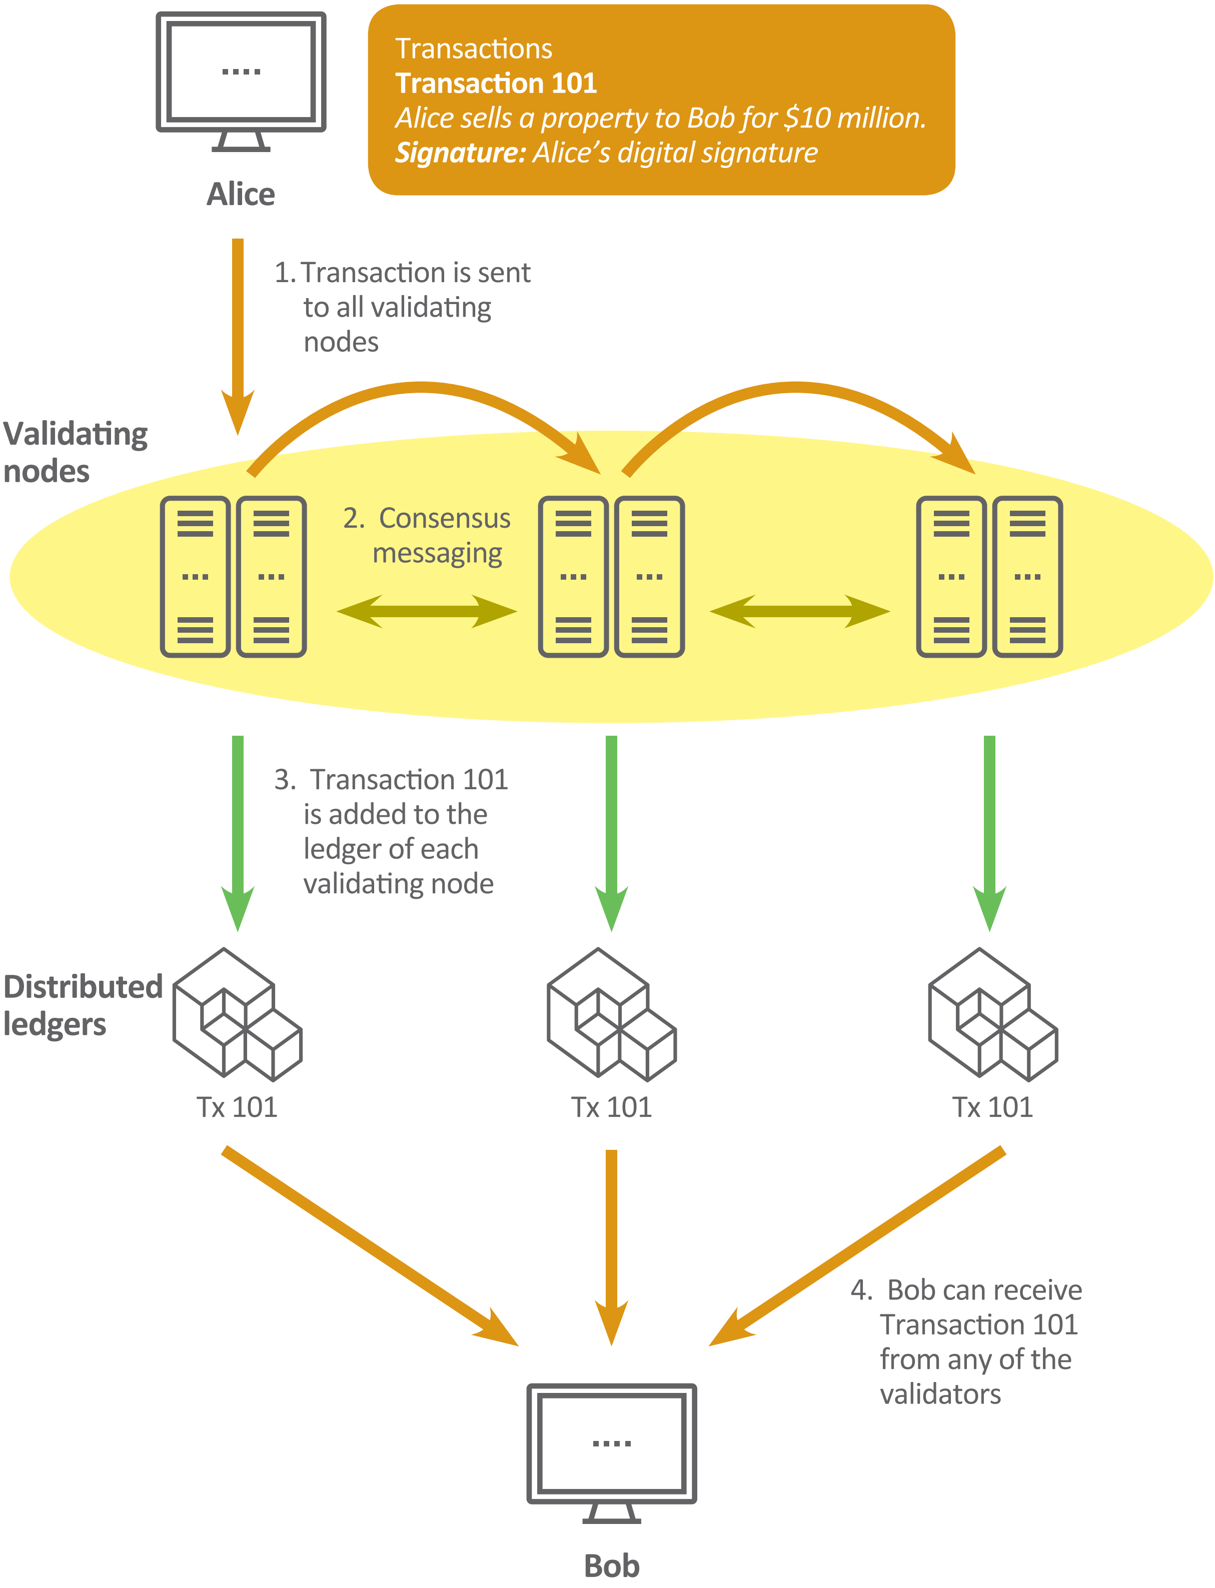
Transactions
Transaction 101
Alice sells a property to Bob for $10 million.
***Signature:** Alice's digital signature*
Alice
1. Transaction is sent to all validating nodes
Validating nodes
2. Consensus messaging
3. Transaction 101 is added to the ledger of each validating node
Distributed ledgers
Tx 101
Tx 101
Tx 101
4. Bob can receive Transaction 101 from any of the validators
Bob

Example: A DLT network where Alice signs and sends a contract to sell a $10 million property to Bob.

4.2 Protecting information integrity with cryptographic technology: a "micro" perspective

The DLT system utilises cryptography to protect the ledger data: this includes both data in transit over the network, and data in local storage. Data authenticity and confidentiality are two major requirements for DLT.

Hash

A DLT ledger contains a large number of blocks and a large amount of transaction data. Once the blocks and the transaction data have been stored in a ledger, these blocks and data are immutable, i.e. cannot be altered. At the macro level, the consensus process among validating nodes ensures that no malicious validating nodes can inject falsified transactions into a DLT network. At the micro level, hash technology is applied to detect any alteration made to the original data. This is a mathematical unidirectional function that summarises a piece of data, regardless of its size, as a piece of unique fixed-size short data called its "hash value". Any alteration to the data causes its hash value to change also. Hence, instead of having to inspect every piece of information in the lengthy data record in order to detect data alterations, one only needs to run the hash function on the data record to obtain a hash value. If the hash value has changed, then one can conclude that the data record has been altered.

The strength of the hash function is that it is close to impossible for anyone to determine a hash value in advance, and then to create a data record that will in turn produce the pre-determined hash value. A minor change in the original data record will result in major changes in the hash value. There are different hash algorithms in the technology world. A common example is SHA256, which reduces a data string of any size to a 256-bit number (see an illustration of the generation of an SHA256 hash below).

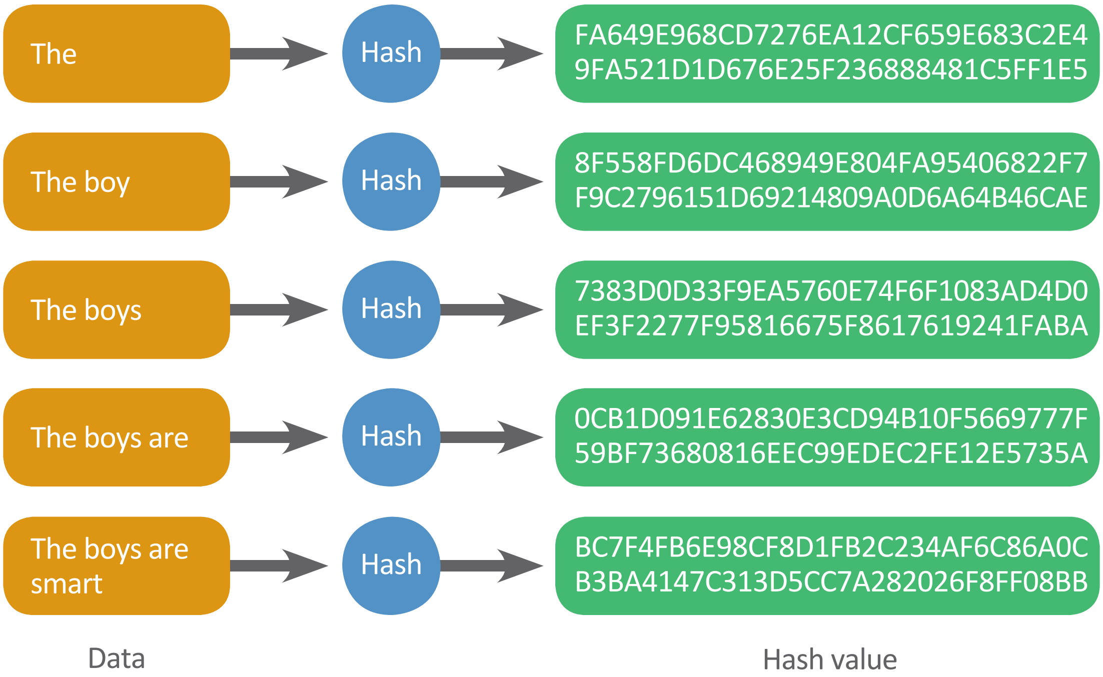
The
Hash
FA649E968CD7276EA12CF659E683C2E4
9FA521D1D676E25F236888481C5FF1E5
The boy
Hash
8F558FD6DC468949E804FA95406822F7
F9C2796151D69214809A0D6A64B46CAE
The boys
Hash
7383D0D33F9EA5760E74F6F1083AD4D0
EF3F2277F95816675F8617619241FABA
The boys are
Hash
0CB1D091E62830E3CD94B10F5669777F
59BF73680816EEC99EDEC2FE12E5735A
The boys are smart
Hash
BC7F4FB6E98CF8D1FB2C234AF6C86A0C
B3BA4147C313D5CC7A282026F8FF08BB
Data
Hash value

Example of the generation of a SHA256 hash for different character strings

As the hash operation reduces an arbitrarily large piece of data to a unique fixed-size short data string, it is often used as a unique identifier of the data itself. At the same time, it keeps the content of the data undisclosed. Not only is the unique ID quite useful for representing the data itself, it can also be used for other purposes. One example is for the proof-of-existence used in some DLT applications. A person who has a digital document containing confidential information can pass the document's hash value to a third party as a reference to the document and as a proof of its existence. When the third party is later provided with a copy of the digital document, that party can verify the genuineness of the digital document by verifying the hash value.

Besides being useful for summarising data and detecting data alteration, the hash function has been innovatively applied by DLT for other purposes. Unpermissioned DLT networks use it to perform proof-of-work mining. When creating a new block containing a set of transactions, a miner runs an indefinite number of iterations of the hash operation. In each iteration, the miner selects a unique number called nonce, combines it with the block's content, and runs the hash operation on the combination. The miner then checks to see whether or not the resulting hash value matches the bit pattern specified by the DLT network. If it does not match, that means the miner has not found a suitable nonce and so has to start another round of the hash operation using a new nonce. When a suitable nonce is finally found, the miner can then legitimately declare that a new block has successfully been created and the nonce can also be presented as proof.

Mining is a computationally intensive process because the process requires very significant hash efforts for a miner to find a nonce that will produce a matching hash value pattern. Miners in an unpermissioned DLT compete with each other to be the first one to find the right nonce. The first one to do so normally receives a reward in the DLT system. Technical readers should refer to Technical Note (II) for a detailed illustration.

Symmetric Key Encryption

In symmetric key cryptography, a single key is used to encrypt and to decrypt data. This allows a person to encrypt his or her data and prevent others from learning its content. As long as that person does not reveal the key to others, no one else can decrypt the data.

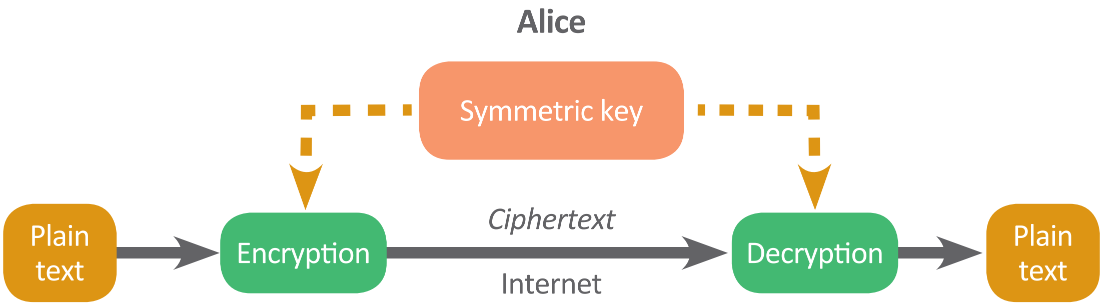

Illustration: Alice uses a symmetric key to encrypt and decrypt data

In the diagram above, a symmetric key is used to encrypt a plain text into ciphertext. The ciphertext can then be securely sent through a public channel such as the Internet. The receiving side uses the same key to decrypt the ciphertext back to plain text.

A popular example of symmetric key cryptography is Advanced Encryption Standard (AES). One common usage of AES is for on-line transactions and communications, where it protects the confidentiality of data exchanged over a public network. There are different AES key sizes, with longer key lengths for better protection. A longer key length also makes it harder for outsiders to guess. Currently, many secure online banking transactions are protected by AES cryptography. During each online banking session, a new AES key is established to encrypt subsequent network communications after the bank web server and the client have authenticated each other using asymmetric key cryptography.

Asymmetric Key Encryption

An asymmetric key is also commonly used for data encryption. It allows two parties who may not want to share a symmetric key to send ciphertext to each other.

Unlike symmetric key cryptography, where the same key is used for encryption and decryption, asymmetric key cryptography uses a pair of keys: a public key and a private key. The public key is used to encrypt the data, while the private key is used to decrypt the data. For example, if a person wants to use asymmetric key cryptography to protect the data exchange between himself and other parties, a pair of keys (i.e. the public and private keys) is created for this purpose. The public key is then sent to the other parties, while the private key is kept by the person in question. When a third party wants to send data to this person, the third party makes use of the public key to encrypt the data. As the recipient possesses the private key, he is the only person capable of decrypting the encrypted data sent by the third party (see the illustration in the diagram below).

The power of asymmetric key cryptography is that it allows two strangers to exchange confidential data on the public network without worrying about data security, and without sharing a single key.

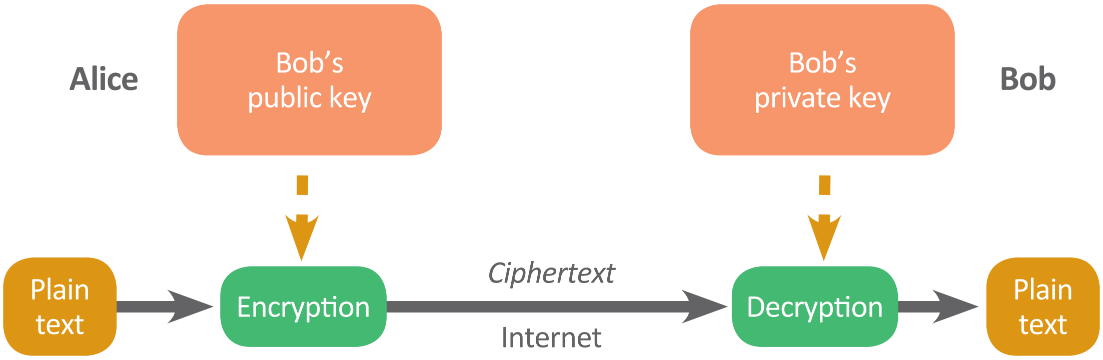

Illustration: Alice passes encrypted data to Bob using Bob's public key

RSA is one of the most widely used asymmetric key cryptography systems. Its key length usually starts at 2048 bits, and goes higher. It is commonly used in setting up secure connections between clients and trusted servers. An example is a client connecting to a bank's web server, where confidential data is transmitted from the client to the bank's web server. The client encrypts the data using the bank's public key and sends the resulting ciphertext to the bank's web server. On receiving the ciphertext, the bank's web server then decrypts the ciphertext and turns it into plain text. An unauthorised third party who attempted to intercept the encrypted data during its transmission over the Internet would find it difficult (if not impossible) to decrypt the ciphertext and access the confidential data.

Digital Signature Technology

Digital signatures are based on asymmetric key cryptography and are used in DLT to certify the authenticity of transactions, i.e. to show that a person is the true owner of an indicated digital identity. When a person creates and sends a DLT transaction, the transaction must also bear that person's digital signature.

In digital signature technology, a person creates a pair of keys: a public key and a private key. The person releases the public key to the public, while keeping the private key under tight security. When the person sends a transaction to a third party, the content of the transaction is signed using the person's private key. The recipient of the transaction can then use the public key to verify that the digital signature attached to the confidential transaction is the person's true digital signature. Anyone else trying to sign the transaction using a different private key will be easily detected as bogus.

A digital signature provides a secure means for conducting online transactions and transmitting certified digital documentation over the Internet. For example, e-Cheques bearing the payer's and the bank's digital signatures, implemented by the banking sector in Hong Kong in 2015, can be sent over the public network without any concerns about tampering by an unauthorised party. In the same way, DLT uses digital signatures to verify the authenticity of transactions over the network.

4.3 Smart Contracts

A smart contract is a DLT feature where a contract is executed by DLT itself. A contract is an agreement between multiple parties. The contract terms are specified in computer language or code that can be executed by DLT validating nodes. DLT validating nodes impartially execute the smart contract based on the specified terms.

A typical smart contract consists of four major components:

- **Contract terms** – a set of pre-defined terms and execution conditions agreed by all relevant parties;
- **Event** – one or a series of specific events that can trigger the transaction, which should be carefully defined in the contract;
- **Execution** – the transfer of value between the contract signing parties when the transaction is triggered; and
- **Settlement** – the settlement of both on-chain assets and off-chain assets.

As in normal contracts, a smart contract needs to be drafted by relevant parties. To cause DLT to execute the smart contract, the contract terms may specify requirements such as the receipt of digital signatures from the parties involved and the occurrence of certain events or conditions.

The creation and execution of smart contracts in DLT do not require the involvement of a middle man. Two parties can simply enter a smart contract into a DLT network, with the smart contract being executed in precise accordance with its terms and conditions. This brings a number of benefits, including fast and automated execution, possible lower transaction costs, and non-ambiguity of contract terms.

Smart contracts can be drafted for many purposes, such as (for example) to handle escrow services, project funding pledges, public company corporate actions, and investor voting systems.

A multi-signature feature is commonly used when a smart contract requires endorsement from multiple parties in order for a deal to be closed. A requirement like this means that the smart contract terms state the identities of all the required signers, and the DLT network will not execute the smart contract unless the required digital signatures are all present. Asymmetric key cryptography is used to support the multi-signature requirement, and the smart contract specifies the required signatures by stating the public key information of the required signers. Alternatively, the smart contract may specify a set of authorised signatures and at the same time state the minimum number of signatures required to trigger the execution of the contract.

5. Deployment

The ease and effectiveness with which DLT can be deployed requires a careful consideration of the various characteristics of this innovative technology. This section looks at the deployment of DLT platforms from the perspective of both performance and interoperability.

5.1 Performance

DLT performance is an important factor in deciding whether and how DLT can be effectively applied for different financial services. In applications with a high volume of activities, such as fund transfer systems, securities trading systems and post-trade infrastructures, performance is a major critical success factor.

The performance of a DLT platform can be measured in the following terms.

a. Block Size Limit & Block Generation Frequency

The throughput of a DLT network is based on the number of transactions the network can process in a given time period. The processing of a transaction involves validating it and adding it to a block within the DLT network. DLT platforms normally generate blocks periodically. The block size limit usually puts a ceiling or limit to the number of transactions a block may carry. DLT networks can seek to increase throughput either by increasing the frequency of block generation, or by raising the block size limit.

Block Size Limits

Many DLT platforms impose a limit on the block size, for various reasons. The limit is specified in the platform rules and policies.

Formation of an unpermissioned DLT network is usually started by a group of developers, who are normally the ones who set the initial policies and rules for the DLT network. Because most of the time such a network is open source, subsequent updates to rules are normally open to public suggestions. These suggestions, if accepted by the group of developers, will be added to the set of rules governing the DLT network.

DLT platforms that place a size limit on the block of a ledger may use different parameters to specify the limit.

Some DLT platforms set the block size limit directly in terms of the maximum number of bytes it can accommodate. Other platforms do not limit the byte count, but instead use other parameters. Once the target set by the parameters has been reached, no more transactions can be added to the current block. At this point the block is immediately processed, validated, and added to the DLT ledgers. Examples of such parameters could include the degree of urgency of the transactions, or the computation overheads of the validating transactions.

For example, Ethereum is a DLT platform that limits the block size in terms of a "gas" limit.[2] Processing of a transaction consumes a certain amount of gas. Each Ethereum block declares the limit of the total amount of the gas that can be consumed by all the transactions inside the block. This limit is not a fixed value and may change from block to block; it is determined by the miners. However, Ethereum exerts control on the allowable rate of change in gas limit between successive blocks.

Block Generation Frequency

Some DLT platforms set a fixed frequency for block generation. Others have a dynamic block generation time.

For example, Bitcoin generates one block every 10 minutes, while Ethereum currently targets block generation at a median period of every 13 seconds.

Permissioned DLT networks may find reasons to implement a dynamic block generation time. For example, the presence of urgent transactions may warrant their immediate validation and incorporation into the DLT ledgers.

b. Transaction Confirmation Time

This refers to the time delay required before a DLT transaction can be considered as confirmed. A transaction is considered confirmed if it has been added to a block within the DLT ledgers and has survived the consensus process among the validating nodes in the whole DLT network for at least a certain amount of time. After a block has been added to the DLT ledgers, the length of history of this block can usually be measured by the number of subsequent new blocks added after it. The more the number of trailing blocks after it, the more likely that it will not be revoked later. Unpermissioned DLT networks usually have a longer confirmation time because the consensus process, (i.e. proof-of-work mining) takes a longer time to perform. Some permissioned DLT networks reduce the confirmation time by applying both special network architecture and consensus algorithms, which reduce both the time needed for the confirmation process and the risk of block revocation.

How many block confirmations are needed for a transaction to be considered as confirmed?

When a block has been created by a validating node and added to the ledger, it is possible that another validating node has created a competing block at almost the same time. Since the validating nodes are connected through the P2P network and experience network communication latency, some time may pass before the two validating nodes realise the collision. If one block wins, then some of the transactions in the losing block may be revoked. Hence, after a block is added to the ledger, the transacting parties are advised to wait multiple block confirmation times, or mining periods, before considering the transactions in the block to be secured and not subject to subsequent revocation. During each mining period, validating nodes run a new round of the consensus process to confirm the addition of a new block to the ledger. If the block manages to stay inside a ledger after multiple rounds of mining periods, this means that no block collision has been detected by the validating nodes within that extended period.

Non-technical Considerations on Setting Performance

It might seem preferable to have a high performance DLT network. However, there are reasons why some DLT platforms might prefer to keep performance below a certain performance ceiling. For example, the block generation rate on a DLT platform handling native cryptocurrencies is directly related to the rate at which new coins are generated, since the availability of new coins can affect the price of the cryptocurrency. Such platforms therefore deliberately keep the block generation rate under a protocol-defined limit.

Considerations on the Computational Capability of Validating Nodes

When a DLT platform increases its block size or block generation frequency, more transactions can be processed. The burden of processing the higher volume of transactions in a timely manner falls on the validating nodes. Payment transactions require the validating nodes to verify the source of payment, while smart contract transactions require them to run the smart contract scripts.

An unpermissioned DLT platform needs to take this into consideration when planning to raise its performance. If the new performance criteria require a significant increase in the computational power of the validating nodes, the DLT platform must ensure that this does not push out many weaker validators and thus leave the validating actions to a few powerful (and potentially dishonest) validators. If this happens, the DLT network will be exposed to the 51% risk, or the risk that dishonest validators gain a majority so that they can control the platform and disrupt its regular functions.

Confirmation Time and Network Architecture

DLT generally runs on P2P networks, which can cover a vast area. A DLT system involves a replicated database that is maintained collectively by a set of validating nodes. Reaching consensus among the validators requires the exchange of information between all validating nodes in the DLT network. However, packet exchanges may experience network delays and occasionally interruptions. Consensus algorithms are designed to take these issues into consideration, and deliver reliable transaction confirmation methods. This has an impact on the length of the confirmation time. If the confirmation time period is set too short, this increases the chance of ledger forking (where the copies of the ledger distributed among the nodes have conflicting contents), thus creating a higher chance of a previously confirmed transaction being unconfirmed or revoked. A P2P network may also sometimes experience fragmentation, where a fragment of the network is temporarily separated from the rest of the network. If transactions are processed and blocks are being created by validating nodes in different fragments of the DLT network, these nodes need to reconcile the differences in their DLT ledgers once the full DLT network is restored.

If performance is to be increased, especially in a permissioned network, consensus algorithms and network architecture are among the factors to be considered. However, while the confirmation time may need to be longer than that required for centralised databases, a DLT network delivers other features and additional efficiencies that make it an attractive technology.

5.2 Interoperability

Interoperability between DLT networks

DLT platforms usually run in their own network domain. DLT platforms are designed to solve different problems arising in the applications they are intended to support. As DLT networks become more widely adopted and their applications diversify, the number of new DLT platforms is also continuing to increase. Two recent examples are Corda[3] and Hyperledger[4]. Each has its own specific characteristics, such as unique consensus algorithms and features that are specialised for different applications.

Currently, a technology called "pegged sidechain" enables multiple DLT ledgers to run side by side,[5] and assets to be transferred between different DLT ledgers and platforms. As more DLT ledgers with different specialties are joined together, a global ecology of heterogeneous DLT networks is gradually forming. Attaching DLT ledgers together, and transferring assets between them, creates both technical and security challenges. Technology must enable DLT ledgers to interface with each other cleanly and smoothly. Each DLT ledger also needs to ensure that any misbehaviour in its neighbouring ledgers will not move across the interface and affect it.

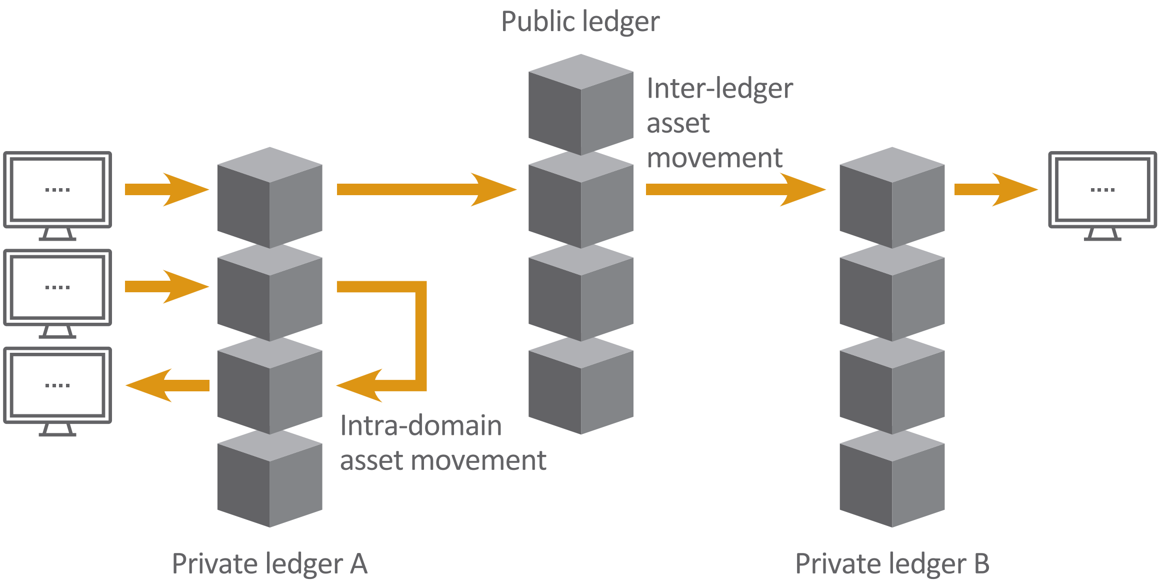

Illustration of the moving of asset (a) within a ledger and (b) between ledgers

Changing DLT Policies, & Creating a New Ledger

Policies and rules of consensus are clearly defined for each DLT network. Attempts to participate by validating nodes with non-conforming behaviour are rejected. However, some DLT allow a group of administrators to agree to create another DLT ledger for a different kind of use and with a new set of rules, and then move assets between DLT ledgers. For example, some DLT platforms define ledger rules and policies with a set of parameters that may be modified by the ledger administrators. The administrator may then set up ledgers with different policies by assigning different parameter settings to them.

Interoperability between DLT networks and non-DLT systems

As an emerging technology, DLT has begun to be applied significantly across many areas of activity. However, interoperability will be important for DLT to be able to integrate or interface with the existing range of applications.

The complex financial world, for example, already has a comprehensive set of multifaceted financial transactions systems, which work together to perform many different kinds of financial transactions. While some of these may gradually migrate to a DLT base, that process is likely to be gradual. This means there will be a mixture of legacy financial systems and newly emerging DLT systems working together at the same time, creating the need for an interfacing mechanism between legacy systems and DLT systems.

This interoperability must cover both control paths and data paths. Different interface methods may be needed to suit different non-DLT systems. In receiving data from a non-DLT system, for example, the interface design will need to decide whether the data is to be pulled onto the DLT platform by the validating nodes, or pushed onto it by the non-DLT system.

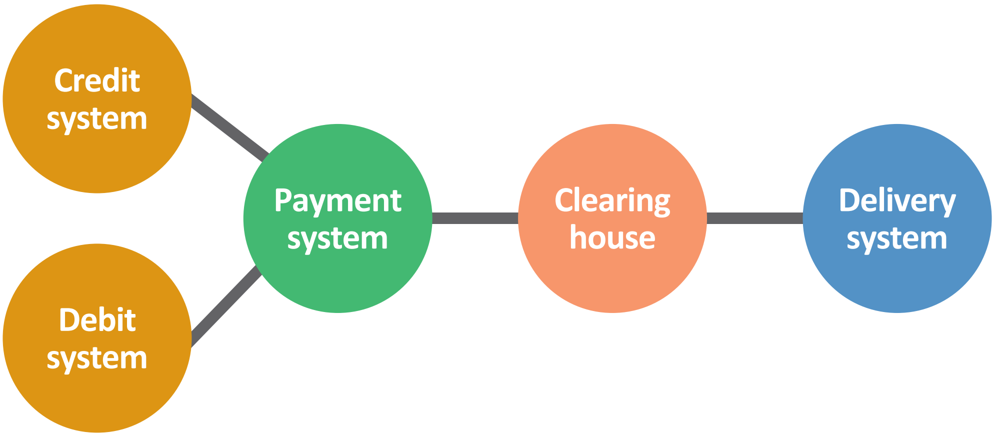

a) A simplified block diagram of a legacy payment system

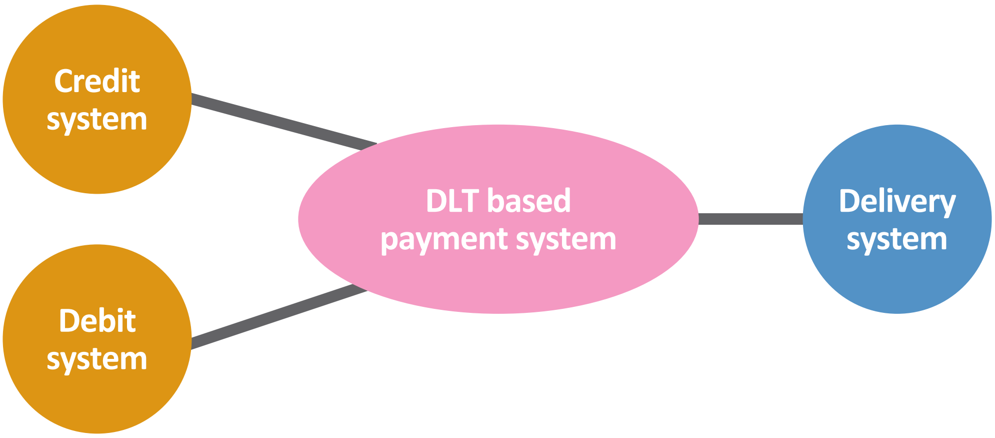

b) A simplified block diagram of a hybrid DLT and legacy payment system

Choice between Data Pull and Data Push

The choice between data push/pull needs to be based on a set of application criteria, including data freshness and data responsiveness.

A pull data mode is suitable when the validating node wants to take the active role of initiating data transfer from a non-DLT system. This allows it to gain data access whenever it needs the data, and avoid the nuisance of receiving data it does not need. For example, suppose the required data is an asset price provided by a non-DLT system, where the price is continually changing. The validating node does not need to know all the changes in the asset price during the day, but only the price at a particular moment specified by a smart contract. In such a case, the data pull model is more suitable. However, since the data has first to pass the barrier between the non-DLT system and the DLT network, certain delays will be experienced before the data is acquired. Care has to be taken when multiple validating nodes in the DLT system are acquiring the same piece of external data. There is a need to make sure that the data is consistent even if there are network delays.

A push data mode is suitable when the validating nodes want to be informed of certain data whenever it becomes available, with the data being stored locally for subsequent access as necessary. It is also suitable if the data does not change after its generation. In this case, the validator can gain immediate access to this locally stored data and does not need to worry about whether the data is obsolete. A more complex interface may use a hybrid of both the pull and push data model.

Interoperability is clearly an increasingly important issue to be addressed as more DLT platforms are developed and are required to interface with traditional platforms.

6. DLT Platforms

Bitcoin marked the introduction of blockchain-based DLT. Since then, there has been a rapid evolution in the design of DLT platforms. Platforms with varied features and characteristics have emerged on which developers can build different applications.

DLT platforms can be divided into two main categories: unpermissioned and permissioned. In unpermissioned platforms, the ledger is maintained by collaborative action among nodes in the public network, and is accessible to anyone. In a permissioned platform, participation is restricted to member nodes only: the ledger is maintained by authorised nodes and is accessible to registered members only. Permissioned platforms enable fast transaction validation to take place, offer enhanced privacy, and at the same time take less energy to operate.

DLT platforms in each category also differ among themselves, each having its own unique features. For instance, some are designed for specific types of applications and others for general use. This section briefly introduces some of the most popular DLT platforms. Reader should refer to Technical note (III) in this document, “Considerations on the Deployment of Platforms and Applications,” for further information.

6.1 Bitcoin

Bitcoin is an unpermissioned DLT introduced by Satoshi Nakamoto in 2008 as a digital asset and payment system. It supports a native cryptocurrency called bitcoin.

According to Satoshi’s thesis “Bitcoin: A Peer-to-Peer Electronic Cash System”[6], Bitcoin has the following characteristics:

1. Two parties are able to transact directly without the need for a trusted third party
2. Transactions are non-reversible
3. Double-spending is impossible

Anyone can join the Bitcoin operation. Each user creates a “wallet” with a unique address that identifies it in the network. The address comes with a pair of cryptographic keys for signing and verifying transactions. Transactions are sent and received with reference to these addresses.

Every full participant in Bitcoin keeps a full history of all transactions. Although transactions in the Bitcoin ledger are not encrypted, user pseudo-anonymity is preserved because only the wallet addresses of users are exposed.

Transactions sent by users are verified by mining nodes. Mining nodes continually collect verified transactions into new blocks, doing which requires solving mathematically difficult problems. The first mining node that succeeds in solving the problem wins, and receives a reward. The new block is then added to the blockchain and propagated to the network.

6.2 Ethereum

Ethereum, launched in July 2015, is an unpermissioned DLT platform for decentralised applications[7]. Like other unpermissioned blockchain platforms, the Ethereum blockchain is maintained by all connected nodes in the public network. Transaction verification and block creation is performed by validating nodes, also known as miners. Newly created blocks are then propagated to the network.

Ethereum serves as a flexible platform for anyone wishing to program blockchain applications on the Ethereum Virtual Machine (EVM).[8] EVM can be thought of as a decentralised machine containing a number of different objects known as "accounts", which are capable of executing code. There are two types of accounts: Externally Owned Accounts (EOAs) and Contract Accounts. EOAs are accessed externally with private key verification, and Contract Accounts are governed by their internal codes. Smart contracts are computer codes added to transactions by users and sent to the ledger for deployment. Smart contract execution may be triggered by additional user-sent transactions. Sending a transaction to the account in EVM requires a small fee payment known as "gas". The fee is a reward to the miners to compensate them for their costs in hardware and electricity in performing transaction validation and block creation.

The reward is paid in the Ethereum native cryptocurrency, 'Ether'. It can be traded or used as a medium of payment on cryptocurrency exchanges, such as the Coinbase wallet system. Ether has been developed as a vehicle to facilitate the operation of peer-to-peer smart contracts.

6.3 Hyperledger

Hyperledger is a permissioned blockchain platform[9]. It is open source, and represents a collaborative effort among businesses and industries to advance the way business transactions are conducted.

Hyperledger has no native cryptocurrency. Participation is restricted to members only. Business transactions are coded in smart contracts which may be written in different computer languages. Transactions are encrypted and may be viewed in plain only by parties that have permission. A membership management service and a transaction cryptography management service are provided. Validating nodes verify transactions and execute smart contracts submitted to the ledger. Unlike unpermissioned DLT platforms, validating nodes in Hyperledger do not receive rewards for their efforts.

Hyperledger supports multiple consensus algorithms through its algorithm plug-in interface. Validating nodes apply the Practical Byzantine Fault Tolerance (PBFT) algorithm by default to reach consensus on the state of the blockchain.

6.4 Corda

Corda is a permissioned DLT system developed by R3, an innovation firm that leads a consortium partnership with a group of global financial institutions[10]. Its development was inspired by the concept of blockchain. However, it differs from many other DLT platforms in certain respects. Firstly, its ledger entries are recorded in structures that do not follow the blockchain pattern. Secondly, it is mainly designed for managing, recording, and executing financial agreements between businesses. The DLT operation and the business agreement execution operation are carried out in an environment that is safely governed, easily auditable, and regulatory compliant. Thirdly, sharing of individual ledger data is limited to parties with a legitimate need to know and see the data within an agreement[11].

Its consensus mechanism is also unique. Instead of implementing consensus on a system level, consensus is implemented on individual deals. Different kinds of consensus are supported. However, transactions are validated only by designated validating nodes.

In practice, businesses in a transaction may develop applications which draft business agreements in the form of smart contracts. Corda serves as the platform where the smart contracts can be executed. The agreement is validated and recorded in the ledger, then executed in Corda's secure environment by designated validating nodes.

Participation in the Corda operation is available to registered members only, in order to preserve the confidentiality of business agreements and ensure their safe execution. Corda does not have a native cryptocurrency.

6.5 Ripple

Ripple is a permissioned DLT specifically developed for financial transactions[12]. Its technology is designed as an improvement over current direct bank-to-bank payment systems. It enables banks to send real-time international payments across multiple networks. DLT enables Ripple to settle payments instantly, unlike traditional settlements which can take days to complete. This is done by the instant recording of payment instructions and their execution within the DLT system.

Execution of payment instructions in DLT lowers the cost of banks and operators, while at the same time shortening the settlement period. Another benefit is the complete traceability of the transactions, making the operation auditable.

Consensus is based on a protocol called InterLedger Protocol (ILP), which is based on the Practical Byzantine Fault Tolerance algorithm. This protocol runs between different bank ledger systems and across national boundaries[13].

Since banks are highly regulated, Ripples is designed to match bank infrastructure and practice by supporting risk management, the meeting of compliance requirements, and the preservation of privacy. The confidentiality of transactions is preserved through the use of cryptographic technologies.

Ripple has a native digital asset known as XRP. It is not required for performing bank payment transactions. However, it is positioned to create competitive foreign exchange markets for cross-border payments.

7. Governance

A DLT network, whether permissioned or unpermissioned, needs to be under some form of governance to ensure its proper operation. The governance structure may start as a set of operating rules established by the original designers. These rules are then ingrained in the design of the DLT software and embedded in its operational protocol. The rules aim to govern and control the behaviour and operation of the DLT network and of its participants.

7.1 Rule-Making Processes undertaken by Participants

Formation of an Initial Set of Rules

An unpermissioned DLT network is usually formed by a group of developers, who are also normally the ones who set the policies and rules. As such a network is open-sourced most of the time, subsequent updates to rules are also open to public suggestions. The suggestions, if accepted by the group of developers, are added to the existing set of rules for the DLT network.

The formation of a permissioned DLT network may involve multiple authorised parties. As founding parties, they either are directly involved in the initial rule-making process, or else delegate this task to others. One example is R3Cev, a consortium of multiple financial institutions.[14] Together, these financial institutions define the set of rules for the distributed and replicated ledgers for various financial services.

Rule Updates

In most cases, an unpermissioned DLT network has no central authority. It runs on a fixed set of written rules that are built into the DLT software. The advantage is that no one can bend the rules, or force others to do the same. This sits well with the absence-of-trust nature of an unpermissioned DLT network. Making new rules or changing existing rules requires modifications to the software. As DLT gains greater acceptance and is used more widely by the general population, the need for new features and bug fixes inevitably arises. This requires adequate change control procedures, which allow for adding new rules or modifying existing ones. The process begins with achieving consensus for the changes from the users and developers. The developers then modify the software. Finally, it is up to the miners as to whether or not they migrate to the new/updated software, i.e. the modified DLT network.

The length of this process depends on multiple factors. Among them, urgency and level of controversy are major ones. For urgent bug fixes, gaining consensus among users and developers is usually not difficult. Because the miners are committed to the safety, soundness and smooth operation of the DLT network, they are usually more than willing to adopt new software to fix any problem that has arisen.

In some situations a new rule may be proposed that does not gain support from the entire group, including miners and users. In such cases, the process will have both supporters and opponents and resolution could become difficult.

In a case like this, the original DLT network may split into two different networks. One group may end up adopting the new software, while the other group insists on using the original software. However, it is impossible to predict such an outcome, and the democratic process among the users and the miners/validating nodes will determine how the situation will evolve. For example, if the preferences of miners/validating nodes are driven by the size of the transaction fees they might receive, they will choose the network that offers the largest volume of transactions. It is also up to users to decide which network they prefer to use.

In a permissioned DLT network, the rule-making process is similar to that of an unpermissioned one – it also requires software modification. However, as the circle of validating nodes is restricted to registered or authorised members only, the process is generally easier and more straightforward. One reason for this is that the validators, in general, share common interests and have common goals. In cases of disagreement, it also is easier to get the stakeholders, who have common interests, together to address the issue. The matter may become even simpler if the DLT network is under the administration of a central trusted party.

In the case of a permissioned DLT network, it is a good idea for the founding parties to set up an oversight and monitoring function as part of the governance structure. This helps ensure ongoing compliance with the rules and the conflict resolution process, and facilitates the proper administration of access controls over the DLT network.

7.2 Defining roles and responsibilities

Different groups of people are involved in DLT networks. They can be grouped according to their functions as follows:

Developers

Developers are responsible for implementing the DLT protocol, including setting up policies and putting rules in place. The composition of the developers affects how this process operates. In an open source DLT network, developers can be divided into two groups: core developers and non-core developers. Anyone can contribute their software to the network. Their contribution will be examined by a core group of developers to verify its quality and check its features. The core group decides whether or not to accept it.

The general public or non-developers also have a role to play in DLT software development. Although they do not develop software, their opinions on the new features or changes have a bearing on the decisions of the core developers when changes to the software are required.

Before being accepted, any modification to the software should be subject to the change control process and to a quality examination. In addition, developers should ensure that the changed software is adequately tested and is accepted by users. Importantly, proper systems documentation must be maintained and kept updated to ensure that ongoing systems maintenance and support can be carried out smoothly.

Validating nodes

Validating nodes have the role of validating transactions and adding them to the DLT ledger. In addition, unpermissioned DLT validating nodes perform computationally intensive block mining operations to create blocks of validated transactions that are eligible to be added to the DLT ledger. Validating nodes engage in a consensus process among themselves to maintain consistency among the replicated copies of the ledger. Validating nodes are essential to the operation of the DLT network, as they commit significant levels of resources to delivering the performance and security required by the DLT platforms. While parties in permissioned DLT networks are willing to make such investments, validating nodes in unpermissioned DLT networks (also known as miners) seek to receive rewards in return for having committed significant computational resources to the process.

Users

Users of a DLT network conduct transactions in the network, for example by making payments, entering contracts with other parties, transferring asset ownership, and performing other activities. Though not strictly seen as one of their roles, users in fact play an essential functional role: they keep the DLT network alive and operational. It is users who send transactions to the DLT platform, giving the validating nodes transactions to work on.

In the case of an unpermissioned DLT network, the role of the users is normally one of rewarding the miners with transaction fees. Without transaction fees, miners in an unpermissioned DLT network may have few incentives to commit their computational resources to the operation.

Membership manager and cryptographic key manager

A permissioned DLT network has two additional roles: controlling access to the network, and controlling the content of each transaction. These two roles require a membership manager and a cryptographic key manager. A membership manager controls the admission of members to the network. If transactions in the network are encrypted, the cryptographic manager, together with the membership manager, determines which users will be granted right of access to the transactions in plain form.

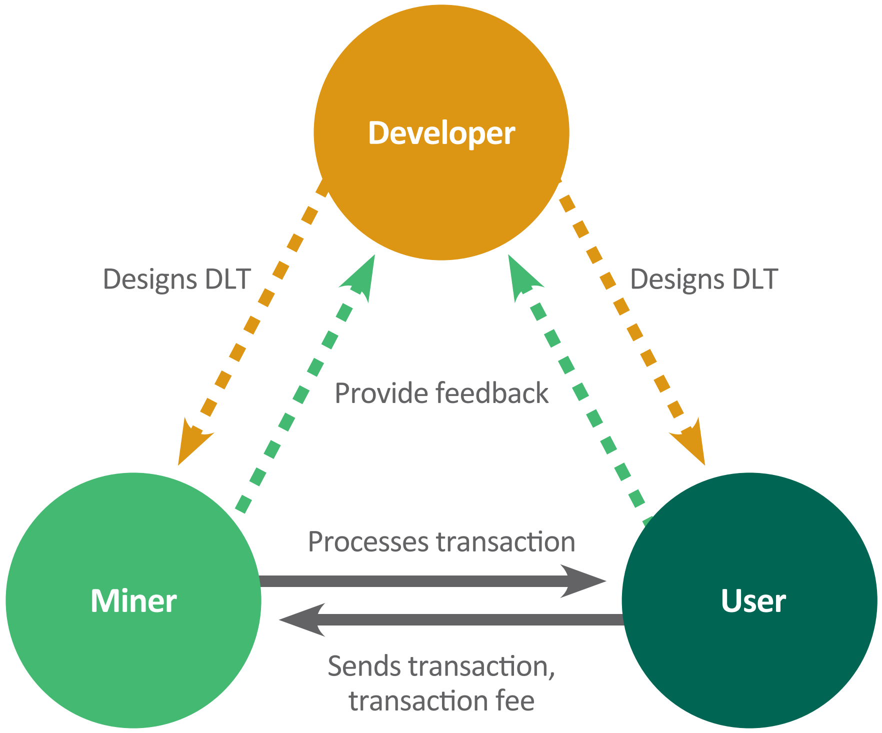

Governance structure and role players in an unpermissioned DLT network

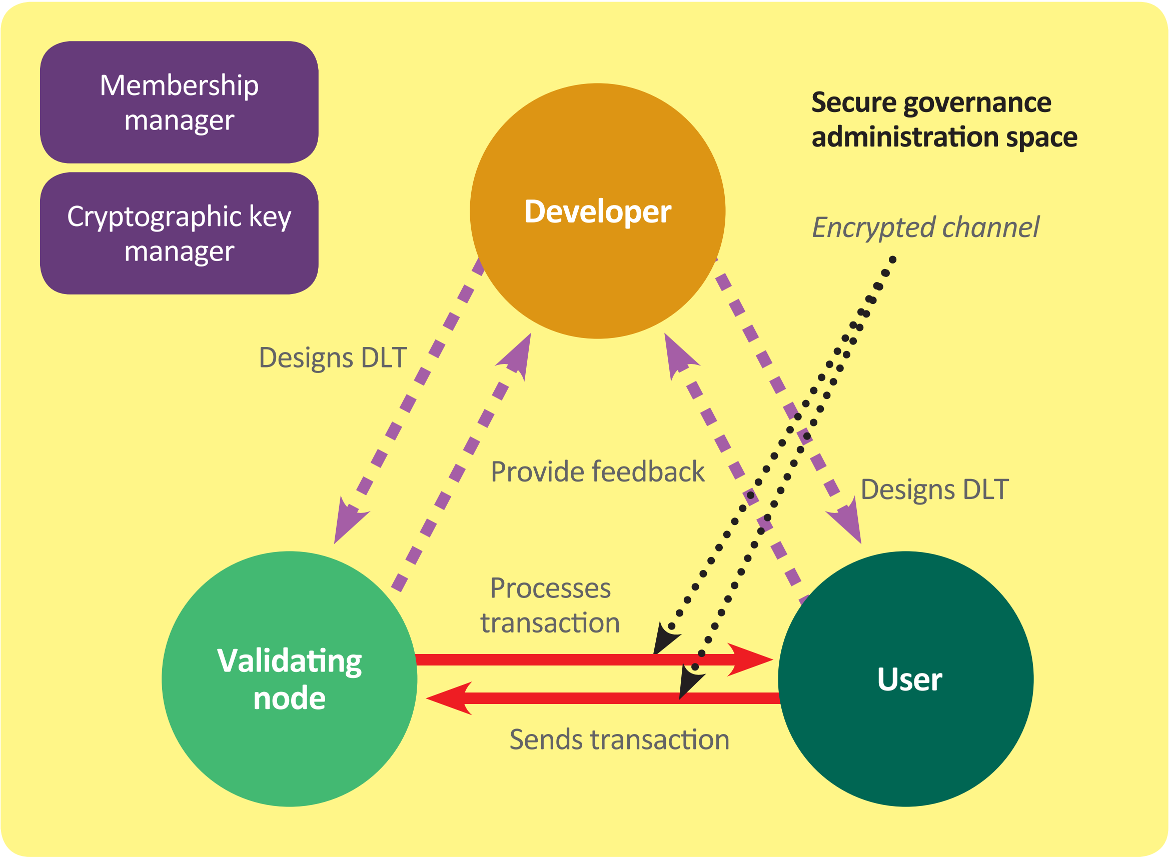

A common governance structure of, and roles in, a permissioned DLT network

7.3 Conflict Resolution

When a new transaction is sent to the DLT P2P network, the transaction is checked and processed by the validating nodes spread across the P2P network. This validation is done based on the rules set in the DLT platform. Once the validating process is complete, the transaction is incorporated by the validating nodes into a block, which is then linked to the ledger.

The validation is done in strict conformance to the DLT rules. Rules are also set to resolve conflicts that may arise due to the nature of a DLT network. However, other kinds of conflicts may also arise that require external intervention.

Transaction Validation

The validation process includes checking for incorrect transaction syntax and for violation against the state of the DLT network. The transaction syntax check normally includes an inspection to ensure the presence of mandatory information, the authenticity of the digital signature, and the validity of any ledger account addresses. An example of a violation against the state of the DLT network is where an asset, which the ledger indicates has previously been spent already, is double-spent by a transaction.

For example, in a cryptocurrency payment transaction, a check is done to confirm the existence of the source of the stated paying funds, which must be recorded in the ledger. Next, a check is made to discover whether the funds are still available, or have already been spent. The transaction is rejected if any non-conformity is noted.

Conflicts Automatically Resolved by Ledgers

As a DLT system is one of replicated databases managed by multiple validating nodes which are physically separated, situations can arise where the validating nodes see contradicting snapshots of the current state of the system. For example, a validating node may see a transaction that other validating nodes do not see. The DLT protocols should have well-defined rules for resolving such conflicts.

If two miners hold contradicting copies of the ledgers, they apply the consensus algorithm to resolve the differences. On unpermissioned DLT platforms employing proof-of-work mining, the rule for selecting between two contradicting but valid ledgers always involves selecting the ledger of higher difficulty. As a ledger is actually a chain of transaction blocks, a ledger has a higher difficulty if the chain of transaction blocks is longer than the chain in the other ledger. If both chains are of identical difficulty, the resolution can still be reached, but it takes longer. This is because while one camp of miners is holding one copy of the ledger and the other camp of miners is holding a different one, any miner in any camp that first successfully mines the next block will then hold the latest and longest chain. This newly extended chain will be propagated to both camps. Since this newly extended chain is now of the higher difficulty, it will be accepted by all.

This is demonstrated in the diagrams below.

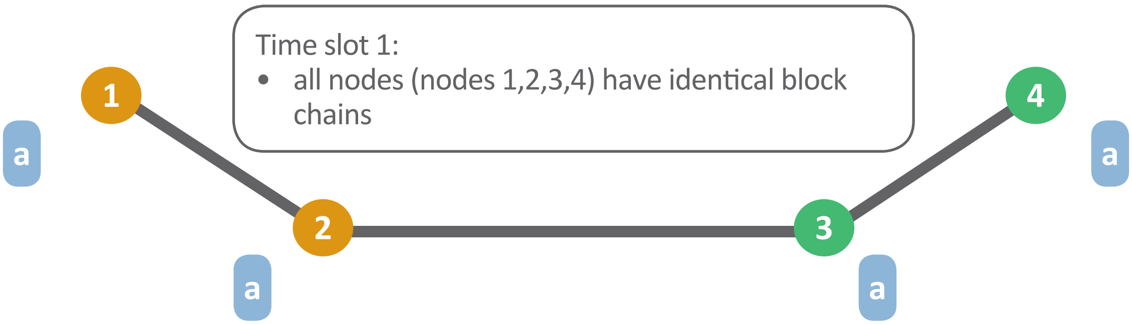

Time slot 1:

- all nodes (nodes 1,2,3,4) have identical block chains

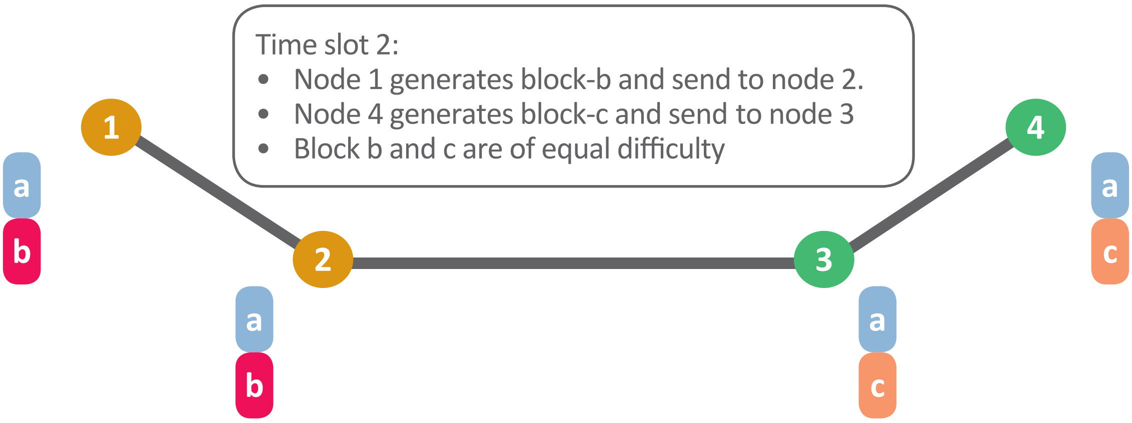

Time slot 2:

- Node 1 generates block-b and send to node 2.
- Node 4 generates block-c and send to node 3
- Block b and c are of equal difficulty

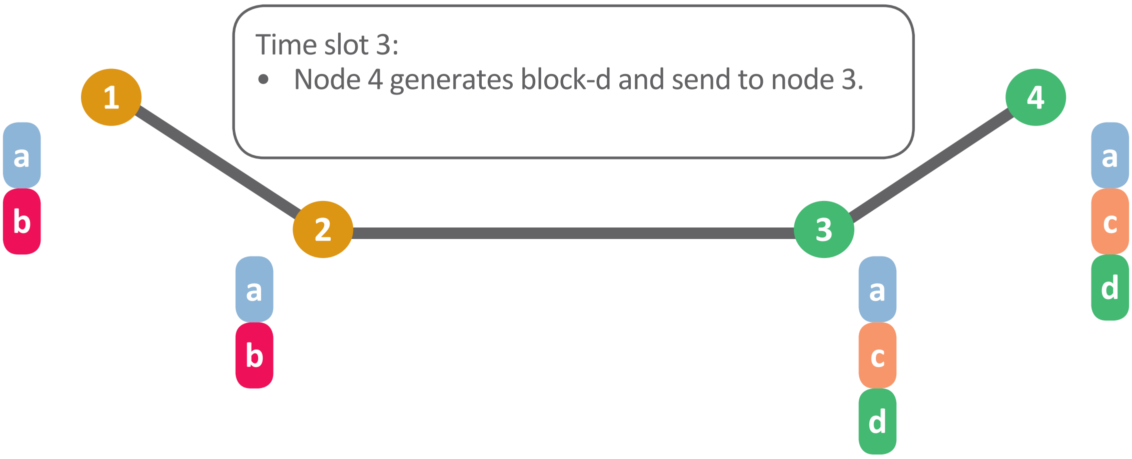

Time slot 3:

- Node 4 generates block-d and send to node 3.

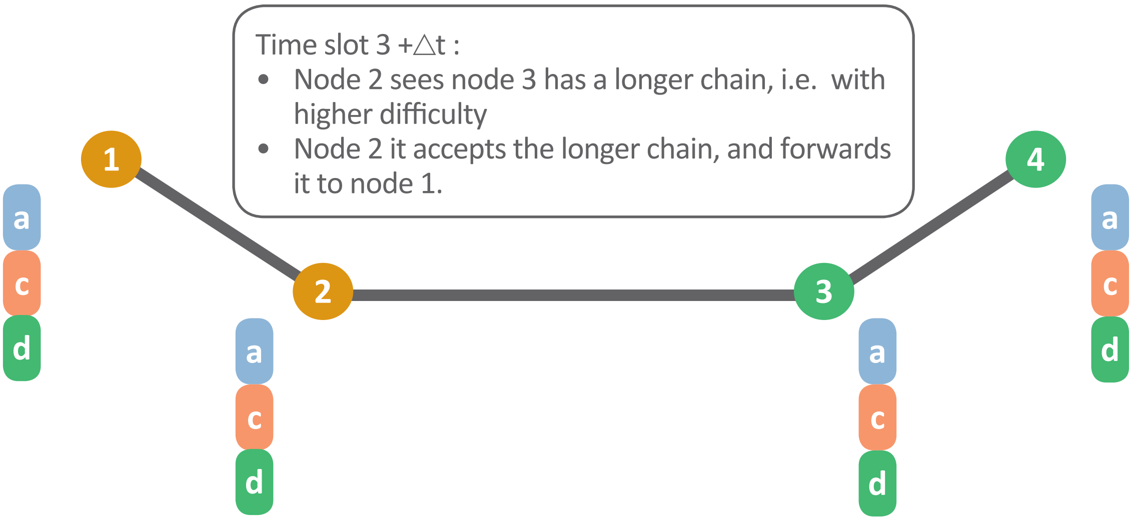

Time slot 3 +△t :

- Node 2 sees node 3 has a longer chain, i.e. with higher difficulty
- Node 2 it accepts the longer chain, and forwards it to node 1.

In this respect, it is very important to ensure that dishonest miners do not hold the majority of the mining power. Otherwise, these miners will manage always to be the first to successfully solve the proof-of-work problem, and generate the longest chains of transaction blocks. This is an example of the well-known 51% risk.

Permissioned networks, such as Ripple[15] and Hyperledger, employ different consensus algorithms to resolve ledger conflicts. In general, these algorithms solve the Byzantine Fault Tolerance problem. The Byzantine Fault Tolerance problem refers to a situation in which a group of Byzantine generals are in different locations, and need to find a way to communicate a common plan: either attack or retreat. It is imperative for them to come to a common agreement. The algorithms are designed to formulate methods for reaching agreement in such a situation.

Conflicts Requiring External Intervention

DLT is designed to handle events that occur as part of the nature of a distributed ledger, such as transient inconsistency due to network delays. However, even if a DLT network is correctly implemented and the network is in perfect condition, it is still possible for other events to occur that may demand further resolution by human intervention.

One example comes from Bitcoin. The current Bitcoin block size limit is 1 Megabyte. Statistics indicate that a block typically contains 2,000 transactions, each of around 500 bytes. As the Bitcoin user base has expanded, the number of transactions has increased. Some foresee a need to lift the block size limit in order to accommodate the rising number of transactions. However, not all stakeholders see a bigger block size as being to their benefit. For example, miners stationed in locations with lower network bandwidth might consider themselves at a disadvantage because their network cannot cope with the increased volume of transaction traffic.[16] If and when a resolution is reached, its adoption must also be carried out by human intervention: i.e. by changing the Bitcoin software and installing the new software in the Bitcoin nodes.

In most DLT networks, there is no central authority to manage, administer and control the operation of the network. Such an operating model is very different from conventional system infrastructures and operating models. A new kind of governance framework and structure is certainly needed, and should be put in place to oversee and manage the DLT network on an ongoing basis. Further study will be conducted in the next stage of this research project with a view to identifying possible regulatory implications, as well as formulating a set of control principles for the regulatory community. This will help prepare for the wider adoption of DLT in financial services.

8. Risk Management and Regulatory Compliance

Some proposed DLT platforms would be deployed and operated over the public network, and financial transactions may be transmitted over the network and processed by these DLT platforms. The possible risks involved in operating such platforms should not be underestimated. These may include operational risks, cyber attacks, and money laundering issues. This chapter aims to identify a list a possible risks and consider their implications for DLT development. Like the governance issues set out in the previous chapter, further study of the risk management issues will also be carried out in the next stage of this research project, with an aim of formulating a set of control principles that will address the risks described below.

8.1 Operational Risks

Malicious validating nodes

The DLT consensus algorithms should have the capability to detect, identify and exclude malicious validating nodes from the DLT network. For example, a virus-infected validating node may inject blocks containing false or unauthorised transactions into the DLT network. Any spike of attacks attempting to change the transaction history or inject fictitious transactions should be able to be promptly detected and identified by other validating nodes, and be adequately dealt with by the consensus algorithms.

Network problems and attacks

DLT nodes residing on the DLT network may be subject to various network breakdown problems and malicious attacks. While network problems such as congestion may cause communication delays, malicious attacks such as a Distributed Denial of Service Attack (DDoS) or a targeted cyber attack may bring down some validating nodes. Under normal circumstances, the DLT network topology is designed to be able to survive any temporary network disruption problems. The presence of multiple validating nodes also enables the network to survive DDoS attacks targeted at selected validating nodes. However, there are other forms of DDoS that are capable of attacking all the validating nodes, by targeting specific validation functions and making use of the network to spread the attack. As the types of DDoS attacks vary, new preventive and remedial counter-attack measures may be required to fend them off.

8.2 Identity Theft Risks

A person within a DLT network proves his or her identity by presenting a digital signature. Since a private key is needed to create the digital signature used to identify oneself and to prove ownership of assets, the owner of the assets must adequately safeguard the private key to protect the assets. If a private key is stolen, the perpetrator can act as the owner and is thus able to change the ownership of assets.

It is the owner's responsibility to take measures to ensure that the private key is kept safely and securely. The private key should also be encrypted, and access restricted to authorised users only. Normally, "cold storage" provides better protection for this than online key storage does. In cold storage, private keys are stored offline. An example of cold storage is the Trezor device, which is a security token in which private keys are stored. When a digitised transaction needs to be signed, the digitised transaction is forwarded to the security token for signing.

While the loss of a normal user's private key may not have any direct connection with the integrity, safety and soundness of the DLT platform, it is possible that the loss could lead to unauthorised transfers of the ownership of assets, which in turn could significantly undermine the confidence of DLT platform users.

8.3 Conduct Risks

Money laundering

Illegally gained fiat currency may be converted by currency exchange into cryptocurrency, and a series of transactions then conducted through a DLT platform to hide the source of the money. A DLT platform may thus be used for money laundering activities and the anonymous transfer of assets if adequate money laundering controls are not in place, such as Know Your Clients (KYC) rules and transaction monitoring controls over the DLT platform.

Sales of illegal drugs and contraband

Criminals may use DLT platforms as payment gateways for illegal drugs and smuggled items, or as illegal asset exchange platforms. Again, KYC measures and transaction monitoring controls are crucial for DLT platforms, especially if these platforms allow participation by public users.

Receipt of ransom payments

Recent years have seen the rise of ransomware computer viruses which infect many personal and business computers. Hackers sometimes demand that victims deposit ransom payments into the hackers' DLT "wallets", which are pseudo-anonymous.

It is prudent to anticipate continued DLT misconduct in the future. An effective risk management system and procedures should be in place so that any problem in the application will not spread to the DLT platform and undermine its stability, or affect public confidence in the integrity of the platform. Likewise, any attack on the DLT platform should not jeopardise the safety of the DLT assets of users.

Adequate preventive measures and detection measures should be implemented to minimise the chance of abuse and the risk of operational disruptions. Given ongoing advances in cryptography, this will prove a continuing challenge, as users can hide the details of their transactions using cryptography.

DLT platform operators should implement tools for logging and analysing system activities that can detect suspicious activities. DLT platforms should also develop better reporting and incident response systems. In addition, a robust business continuity management framework and related arrangements should be put in place to handle any serious disruptions to DLT platforms.

8.4 Regulatory Compliance

Financial and banking stability and consumer protection are the key objectives of all regulatory authorities. This is why the financial and banking sectors are heavily regulated. The growing interest in and increasing adoption of DLT for financial and banking services is obviously increasing pressure on the regulatory authorities.

In general, DLT compliance with regulatory requirements remains an area that remains unexplored or has received little in-depth investigation. To date, regulatory authorities have issued little in the way of regulatory guidance or control principles. In addition, the decentralised and cross-border nature of certain proposed DLT platforms makes the regulatory issues even more complicated. This leads to questions about which activities should be regulated, how activities should be regulated, and by whom they should be regulated. Although it might seem most straightforward if regulators were simply to adopt a traditional regulatory approach (given their technology neutral stance), it remains unclear whether some of the key risks and legal issues associated with DLT set out in this white paper can be adequately dealt with using such an approach.

Regulatory compliance is certainly an area to keep an eye on, and one which requires further study before any formal regulatory requirements are implemented. This white paper aims to conduct an in-depth review and examination of DLT and set out the possible risks and issues involved in its use. It is thus helping build a good foundation for the next stage of this research project, which will include efforts to determine sound regulatory guidance and control principles for the regulatory community.

9. Security and Privacy

9.1 Security

Unpermissioned DLT Networks

As set out in previous chapters, in unpermissioned DLT networks the assumption is that no miners are trusted. To protect the integrity of an unpermissioned DLT network, mining mechanisms are designed to ensure that no miner is capable of cheating or taking advantages of other miners by creating fictitious blocks and transactions. These mechanisms include proof-of-work mining, discussed in detail in chapter four.

Similarly, no user of an unpermissioned DLT network is trusted. To prevent a user from double-spending assets recorded in the unpermissioned DLT network, miners check the asset transaction history and compare their copies of the DLT ledger. This comparison, commonly known as the consensus process, is one of the mechanisms that enable miners to detect double-spending by dishonest users of an unpermissioned DLT network.

Although ownership of an asset is visible in an unpermissioned DLT network because transactions are seen by all, a digital signature of the asset owner is required before the asset can be traded and ownership transferred. This applies not only to asset trading, but to all transactions conducted within an unpermissioned DLT network.

To guard against any kind of failure or misappropriation, or the disappearance of copies of the DLT ledgers, replicated copies are kept and maintained by physically separate miners. This ensures that (unless all the replicated copies are destroyed or amended) the unpermissioned DLT network and its sets of recorded transactions remain intact.

Permissioned DLT Networks

Like unpermissioned DLT networks, permissioned DLT networks also use cryptography for record protection, but in an even more extensive manner.

Only authorised parties are allowed to join a permissioned DLT network. This reduces the risk of a dishonest validating node mishandling the operation of the permissioned DLT network and misappropriating transaction records within the network. Honest nodes gain access to the permissioned DLT network only after passing security authentication, such as through the validation of a registered member's digital signature.

Further protection is provided by having all validating nodes run a consensus process to ensure they all agree on the transactions to be added to the DLT ledger. This ensures the health, integrity and consistency of the replicated copies of the DLT ledger. It also helps isolate validating nodes that are acting abnormally due to technical failures or other problems.

9.2 Privacy

Unpermissioned DLT Networks

User identities are pseudo-anonymous since only a limited amount of personal data and identity information is stored in unpermissioned DLT networks. Although transactions in an unpermissioned DLT network are usually unencrypted, they can only be traced to DLT wallet addresses, which are numerical strings associated with wallets but which do not reveal the wallet's private (secret) key, and do not expose much wallet owner information. This helps preserve privacy and hide the identity of the owner.

Users can further protect the privacy of their transaction history by utilising different DLT wallet addresses to conduct individual transactions.[17] This makes it very difficult for other parties to trace the owner's identity, or to identify the owners of the transactions and assets stored in the unpermissioned DLT network.

Permissioned DLT Networks

Transactions in a permissioned DLT network are encrypted to prevent unauthorised parties from reading them. Network communication among validating nodes and users is also encrypted to make data unintelligible even when intercepted during data transmission over the permissioned DLT network. Data privacy can be further protected by storing confidential data in a restricted location, with the DLT ledgers merely storing the hash of this data.

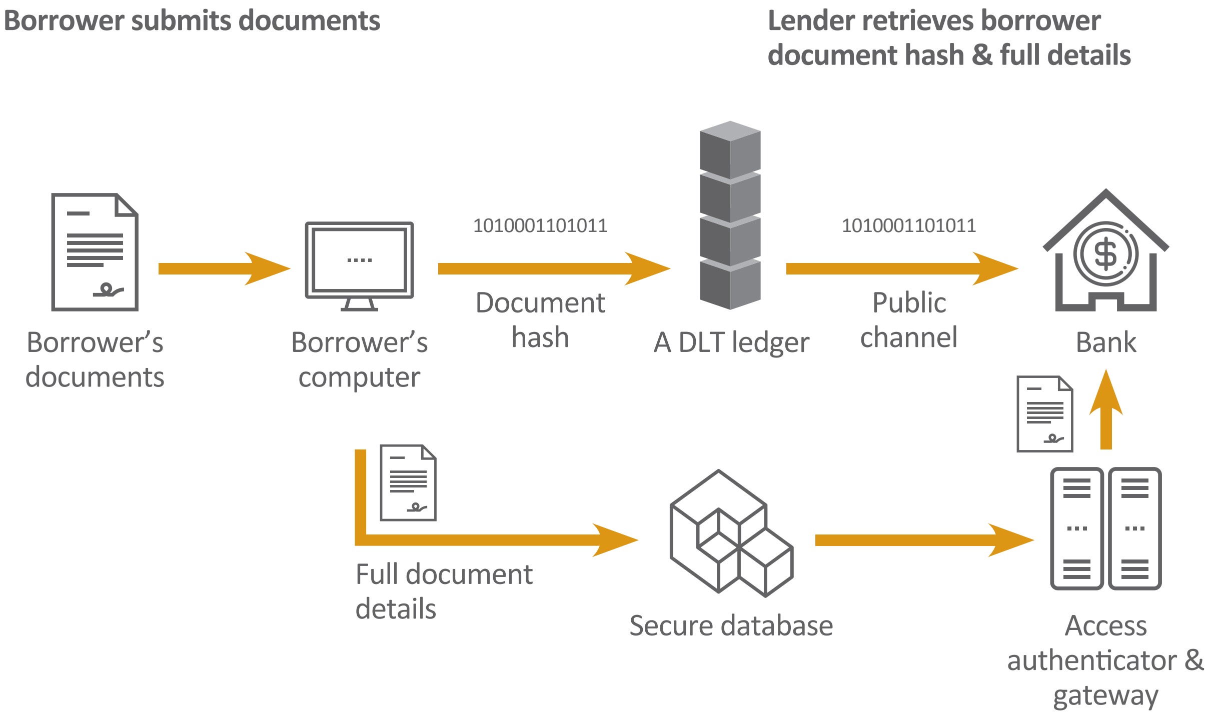

Illustration showing the storing of document hash in a DLT ledger

Although only registered members can participate in a permissioned DLT network, members' identities are not necessarily revealed in the transactions. To keep a member's activity history private, a DLT platform may assign each member a transaction certificate when a new transaction is conducted. This means that different transactions initiated by a member may bear different transaction credentials. As a result, other members can only relate a transaction to a particular credential, and cannot infer the identity of the member from it. Some permissioned DLT networks utilise this approach to protect the privacy of members and their activity history by issuing pseudonymous member credentials in the form of transaction certificates.[18] Only authorised auditors are granted permission to access membership information and are able to relate transaction certificates to members' identities.

9.3 Challenges

Unpermissioned DLT Networks

Potential Attacks by a Dishonest Majority

The integrity of an unpermissioned DLT network relies on the assumption that the majority of miners, who possess most of the mining capability, are honest in mining and in maintaining the network. If there is a chance that the majority of the miners are colluding together, or that a few dishonest miners possess most of the mining capability, these miners would be capable of compromising the integrity of the unpermissioned DLT network.

DDoS (Distributed Denial of Service) Attacks

While the multiplicity of miners and the high connectivity of the P2P network make it difficult for an unpermissioned DLT network to be successfully attacked by a DDoS so that it is brought down completely, this kind of attack can still slow down the performance of the DLT network, or temporarily fragment the DLT network. This exposes the need for high network and operation resilience to guard against such attacks.

Theft of Wallet Keys & Other Risks

The assets recorded in an unpermissioned DLT network are at risk if their private keys are stolen. Multiple digital signatures may be used to implement an escrow service that strengthens security controls, say by requiring not only the owner's digital signature but also a digital signature from a trusted third party in the DLT network when signing the transaction.

However, theft of private keys remains a threat. There have been occasions when private keys have been stolen from cryptocurrency exchanges and crypto-wallets, thefts which had the potential to bring down the operation of some cryptocurrency exchanges should heavy financial losses and a consequent plummeting of user confidence have occurred.

Incidents of Theft

According to some news reports, both Bitcoin and Ethereum have experienced incidents of theft. However no evidence has been found to indicate that these incidents were caused by the technological design of DLT or by DLT protocol problems. Rather, some incidents were carried out by hackers who took advantage of bugs or vulnerabilities in the software developed by users for keeping their private keys safe.

One example based on a news report relates to Bitfinex, an exchange company that was attacked by a hacker. The hacker managed to breach the company's system and take 119,756 Bitcoins from its user accounts. The company had implemented a system that stored users' Bitcoin deposits online and made it easier for its users to access their Bitcoin accounts. The company partnered with BitGo in implementing a multi-signature protection scheme to prevent unauthorised access to users' accounts. Despite all these protection measures, the perpetrator managed to find a way through the system and steal the Bitcoins from users' accounts.[19]

The privacy challenge

In addition, there is always a challenge and in some cases a dilemma for an unpermissioned DLT network when faced with the need to protect user privacy while also being required to assist law enforcement authorities. Government authorities want to be able to detect illegal cryptocurrency activities such as money laundering, tax evasion, drug trafficking, and ransom payments. On the other hand, users may be concerned that details of their crypto-currency activities, and hence of their personal data and lifestyles, could be exposed to unauthorised parties. The challenge of dealing with the conflict between the need for users' privacy and the right for authorities to know about and monitor activities for various legitimate purposes cannot be underestimated.

Permissioned DLT Networks

Security Breaches

Attacks can come in many forms, such as DDoS and the implanting of viruses and malicious software, so preventive and detection measures should be implemented. The risk of DDoS attacks may be mitigated by monitoring and detection tools, as well as by network equipment that can detect and divert DDoS packets to a quarantined location for analysis and appropriate treatment. System and communication software should be examined and audited to remove any possible vulnerability to DDOS attacks. It is imperative for the system to be protected by anti-virus software, and all software running on the system should go through a stringent change control management process before installation.

Effective Administration

A permissioned DLT network needs to implement an effective administration methodology that allows administrators to monitor, configure, and control the DLT network to preserve its integrity and smooth operation. In the event of a virus attack, network disruption, or other incidents, the network administration team should be empowered to reverse the damage and bring the network back into normal operation. More importantly, a robust business continuity plan (BCP) should be in place to address any contingencies that may arise.[20,21] In addition, getting all the validating nodes to agree on the BCP and any contingency actions to be taken on the DLT network is a huge challenge, so sufficient time and resources should be devoted to this issue when formulating a permissioned DLT network.

(a) Network fragmentation issues

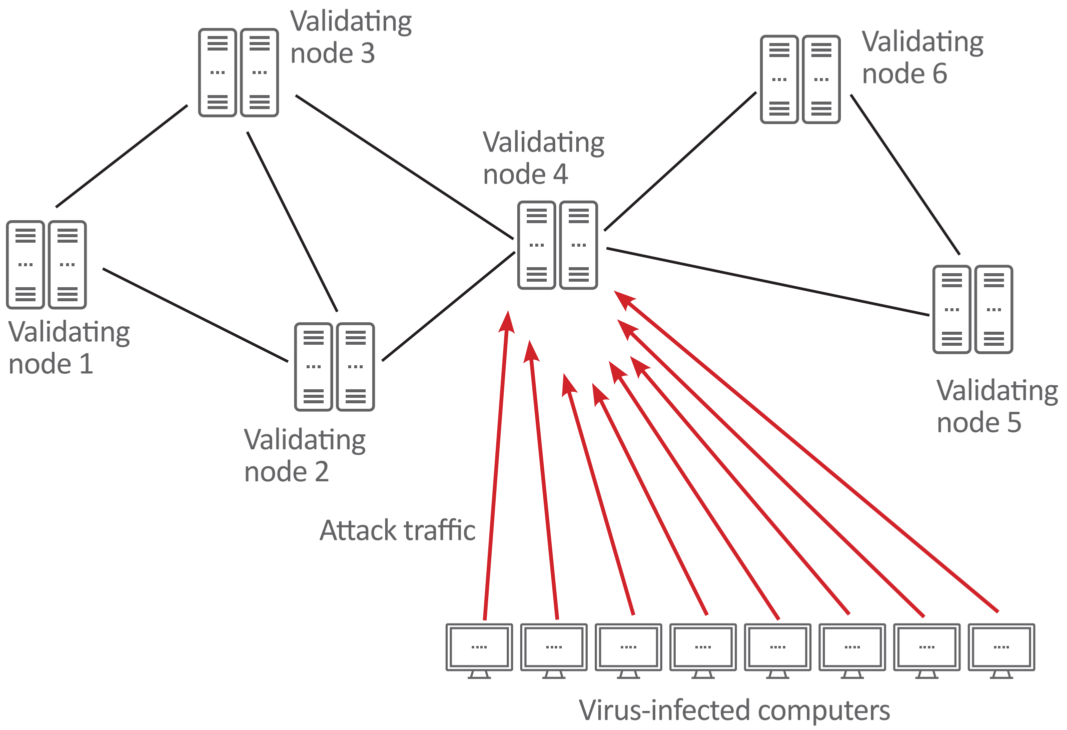

Illustration: a DDoS on a single validating node

In the simplified DLT network system shown in the diagram above, the P2P network is vulnerable at node 4. If node 4 comes under attack and is no longer able to forward packets between nodes on its left and right sides, the P2P network will be split into two network fragments:

Fragment 1: includes nodes 1, 2 and 3
Fragment 2: includes nodes 5 and 6

Although this is a simplified illustration of a DLT network, it demonstrates how important it is for a permissioned DLT network to provide sufficient connectivity and redundancy among the validating nodes to avoid network fragmentation occurring.

(b) Network performance issues

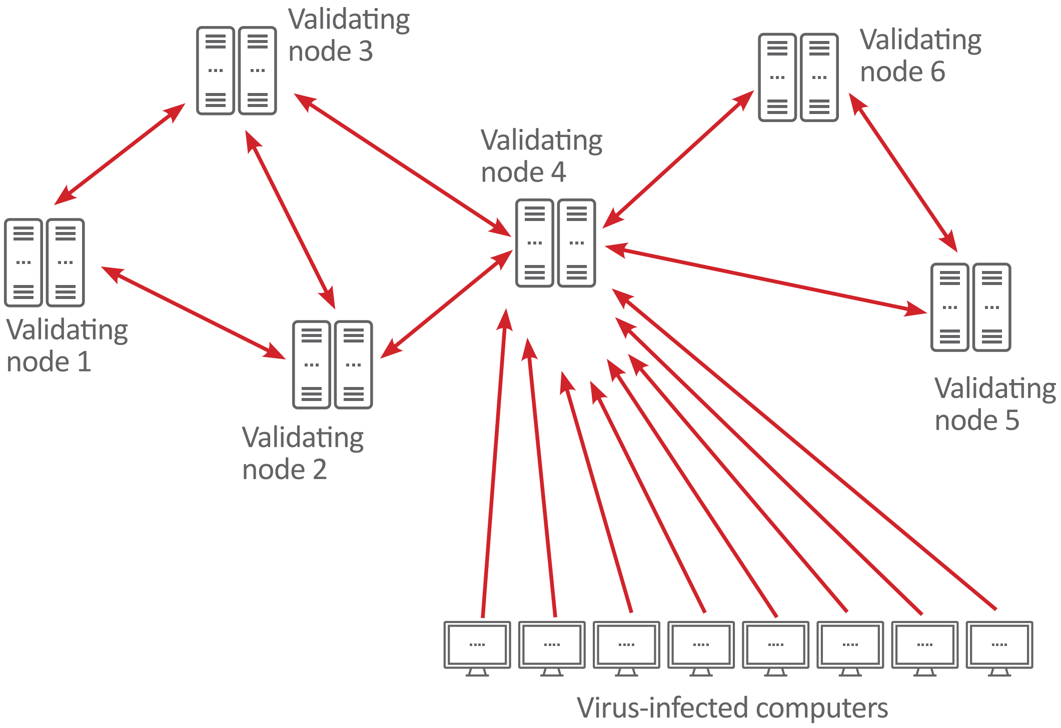

Illustration: a DDoS on all validating nodes

In the illustration above, the DDoS attack sends legitimate transactions to validating node 4. Instead of bringing down node 4, the transactions contain instructions that cause all validating nodes to devote a large amount of resources to processing and verifying them. Since the transactions are legitimate, validating node 4 will both process the transactions and pass them on to other validating nodes for the same process to be conducted. This can result in the performance of the entire network suffering, with significant degradation occurring due to the heavy computational resources required for this process.

The two examples above show the possible impact of DDoS on DLT network operations. It is therefore important to implement adequate preventive, monitoring and remedial measures to guard against such attacks on the DLT network.

The privacy challenge

Sensitive transaction data in a permissioned DLT network often needs to be treated with confidentiality, and access to such data should be restricted to authorised parties only. Strong data encryption and adequate access controls are essential for protecting a permissioned DLT network. A permissioned DLT network should be carefully designed, with the security and data protection requirements being embedded at the system design stage and thoroughly scrutinised before implementation.

A node may serve the purpose of storing a ledger or a database. As new applications evolve for processing data in the DLT ledgers, the need will also arise for new cryptography technologies to ensure adequate data protection. For instance, a need may arise for carrying out a new kind of big data analysis on the records stored in a DLT ledger. To ensure the safety and efficiency of the DLT network, it would be desirable to be able to perform big data analysis on such records in an encrypted form. Also, advances in computing power (e.g. quantum computing) might one day render current encryption technology breakable. New cryptographic technology will always be needed over time to provide improved protection.

In addressing the challenge of balancing data privacy and data transparency so that authorities can carry out certain actions, permissioned DLT networks are obviously a better option. A permissioned DLT network can be designed in a way that meets data privacy requirements while at the same time complying with the requirements of law enforcement agencies and regulatory authorities. For example, since the members of a permissioned DLT are trusted and known parties, the accepted operating rules can specify the need for them to comply with local laws and regulatory requirements.

10. Legal considerations

Clearly, DLT has introduced a brand new way for businesses to operate in relation to business partners and customers. Not only does it have the potential to reduce errors and opportunities for fraud due to its unique distributed nature and the way transactions are recorded in the ledger, it also provides a higher level of data accessibility and resiliency. While DLT has the potential to bring benefits to many organisations that are still in the early phases of exploring or adopting it, its unconventional and unique features could give rise to a number of legal challenges relevant to its eventual adoptability.

As a discussion of the legal considerations relating to DLT would require input from and in-depth study by lawyers and other legal experts, this white paper does not attempt to provide any detailed legal guidance in relation to DLT. Rather, it aims to point out the possible implications of a number of potential legal issues. A more detailed legal study of DLT is needed and should be carried out by legal experts. ASTRI plans to engage interested legal experts to take part in this study at its next stage, in order to provide the banking industry with a higher level of legal guidance.

Legal implications in relation to data privacy

A DLT network may involve the storage of personal data, as for example in the fictional property transaction between Alice and Bob described in chapter three where personal particulars were recorded in the DLT ledger. The handling of personal data raises major issues relating to legal and compliance requirements. If handling of personal data is carried out in Hong Kong, the requirements of the Personal Data Privacy Ordinance (PD(P)O) become relevant.

The PD(P)O imposes a number of major principles to ensure that the personal data collected by institutions is properly stored, kept no longer than necessary, and used only for the purpose for which it is collected. Data users are also required to take practical steps to safeguard this personal data from unauthorised use, and to make their personal data policies and practices known to the public.

Given the decentralised and distributed nature of DLT, customers involved should be made fully aware that their personal data is being shared among all the participating parties of the DLT platform. A proper governance structure should be set up among the participating parties, delivering proper notification to customers through an agreed mechanism. The DLT platform may also be required to come up with a set of agreed rules and policies to be followed by all participating parties and made known to customers.

Even if personal data is permitted to be shared among the participating parties of the DLT platform, proper governance should be in place that defines and agrees the purpose of the collected data, and ensures that this data is only used for well-defined purposes. In cases where the participating parties wish to use the collected personal data for a new purpose, consent should be obtained from customers, and such consents properly indicated according to the ledger record access policy incorporated in the DLT platform.

One of the major benefits of DLT is its immutability, meaning the data cannot be altered or deleted. However, this may contravene requirements that the data retention period for personal data should be defined, and the data should be deleted or purged when it is no longer in use or when the customer requests it.

One potential way of resolving this issue would be to add an extra layer of encryption for each data with a key. Instead of purging the whole record, only the key to the specific data need be deleted. However, further thinking regarding key management should be undertaken, and testing should be performed to verify the practicalities. In addition, proper legal advice is needed in some of these areas before any implementation is begun.

The issues discussed above relating to privacy concerns pose an added layer of complexity for unpermissioned DLT networks. Since there is no central administrator or authority that manages the network, no one takes responsibility for any risk. This is of course one of DLT's unique features, but it does bring with it some significant challenges.

Litigation and legal disputes

Whether a DLT network is permissioned or unpermissioned, a defective code or program bug may cause damage or financial loss to specific participants. The decentralised nature of DLT makes it difficult for recourse to be sought in a court. The network does not take the form of a single company, but of a group of participants connected together for the purpose of carrying out certain commercial activities. There may not be any central administrator or authority to take responsibility for defective operational design or for the misbehaviour of participants. In a worst-case situation, the network may be like a gentleman's club, where members join at their own risk and are not protected by any company laws.

To illustrate this, let us use the example from chapter three of Alice selling her property to Bob in an unpermissioned DLT network. Suppose that, due to a defective code, Alice receives her payment a few days late, even though the property title has been successfully transferred to Bob on the agreed date. Because of this, Alice cannot settle the outstanding mortgage loan payment and gets a heavy late charge penalty from her bank as a result. It would be difficult for Alice to locate any individual who is responsible and can be held accountable for her financial loss, since there is no central administrator of a DLT platform, especially in unpermissioned DLT networks. The situation would become even more complicated if the identities of buyers and sellers were pseudonymous.

Rules and conditions in code

Traditionally, every legal or legally binding transaction involves certain legal documents. In the case of the property transaction between Alice and Bob, these would include a formal sale and purchase agreement for the transaction, a title deed for the change of ownership of the property, and a mortgage agreement between Bank B and Bob.

In the future world of DLT, all contracts may be written in computer code (i.e. smart contracts). This raises the question as to whether legal societies and law firms are aware of and prepared for this. Should a dispute arise, lawyers would be required to review the computer code as well. Paper contracts are clearly written in an easily readable way, and are therefore understandable by all parties involved. However, contracts that are written in code involve the terms and conditions being defined in code logic, which can be hard to understand for untrained eyes. Therefore, having DLT contract drafting and analysis tools available would be a useful step towards ensuring the accuracy and completeness of contracts, and of the terms and conditions within these contracts.

Compliance with laws and regulations

As with a number of the issues discussed above, the absence of a central administrator (especially for an unpermissioned DLT network) makes it difficult to carry out certain basic maintenance activities for system operations, as well as to achieve general compliance with various laws and regulations (e.g. those relating to data privacy, anti-money laundering requirements, etc.). Additionally, cross-border transactions or connections with cross-border DLT platforms raise further major issues relating to the applicability and enforceability of laws.

In the mortgage loan application proof-of-concept illustrating the proof-of-concept work in chapter ten, a number of real legal issues have been identified which may require changes to the law in order to make the DLT model work in the real world. To date, the bulk of effort and resources have been spent in proofing the feasibility of the technology, and there has been an insufficient focus on the legal implications.

International developments

ASTRI understands from the HKMA that the central bank community and regulatory authorities are fully aware of the importance of the legal implications to the further development of DLT. In this connection, a number of working groups on digital innovation focusing on DLT have been formed to examine various of these issues. Set out below are some of the major legal issues that have been identified and discussed by the HKMA. Some of these are similar to ASTRI's observations as set out in this chapter.

Applicable law

- What is the applicable law for the DLT arrangement? How is this law enforceable, particularly in decentralised, unrestricted, or cross-border arrangements?

- What is the legal basis of the arrangement's rules, procedures and processes?

- To what extent are there valid, binding and enforceable contractual arrangements established between participants, including in cross-border scenarios?

- If an asset is constituted in digital form on the DLT ledger, what is the legal basis of the asset in digital form?

Rights and obligations

- What are the rights and obligations of the participants and the arrangement? How are these rights and obligations enforced? Are the rights and obligations known and understood by all participants?

- Under what conditions can the rights and obligations of the participants be challenged? What is the mechanism for changing the rights and obligations in the arrangement?

System accountability and responsibility

- Who is liable if something goes wrong with the arrangement? If there is no central administrator that is liable, what is the dispute mechanism for liability issues?

Enforceability

- How does the arrangement set and enforce its rules, procedures and contracts, including who is involved in the setting and enforcement? How does the arrangement enforce its rules, procedures, and contracts across borders, where relevant?

- How effective are the arrangement's tools to enforce its rules, procedures, and contracts in a timely manner? Would these tools be held up in a court of law? What would be the practical legal, financial, and business impact if these tools were not effective or not valid in a court of law?

These legal issues are not straightforward and cannot be resolved hastily. More research work, study and input from legal experts is required to ensure that they are adequately dealt with. This chapter ends with a paper titled "Distributed Ledger Technology Will Drive Legal Innovation", contributed by Professor Carla L. Reyes. In it, Professor Reyes shares her observations on the disruptions that DLT brings not only to the application of current laws in the U.S., but also to the very foundational elements in the way governments conceive of the creation, implementation and enforcement of law.

Distributed Ledger Technology Will Drive Legal Innovation (by Professor Carla L. Reyes)

Innovation. This word is top of mind for entrepreneurs, regulators and academics alike. Although this is not unusual in and of itself, this new round of innovation discussion arguably centers more heavily than prior rounds on the disruptive impact of technology on the law. Various countries and international organizations have expressed a growing awareness of the need for regulatory innovation to match the recent technology-driven period of industry innovation. In the United States, various federal and state law making and regulatory entities are specifically considering the impact of technological innovation on their mandates, methods and goals. Industry, for its part, appears to welcome the opportunity to influence a shift in regulatory thinking. For example, when the United States Department of Treasury Office of the Comptroller of the Currency issued a paper entitled Supporting Responsible Innovation in the Federal Banking System: An OCC Perspective, and called for public response, at least sixty-two (62) individuals and entities offered comments.

Although other disruptive technologies are among the innovations with which regulators currently wrestle, distributed ledger technologies lie squarely in the center of many regulatory discussions. As with the discussion on innovation generally, the distributed ledger industry appears generally interested in engaging regulators and governments as they craft legal responses and determine enforcement priorities. For example, the United States Department of Health and Human Services received over seventy (70) responses to its call for proposals to use distributed ledger technologies to improve health information technology and health-related research. This sense of openness to collaboration should be fostered, particularly as the regulatory emphasis shifts away from the payments applications of decentralized ledger technologies (e.g., bitcoin as a payment mechanism) and explores other applications of the technology.

To date, regulation of bitcoin and other cryptocurrencies as a payment mechanism has been heavily influenced by traditional payments regulation, with little flexibility or regulatory innovation to accommodate the unique nature of decentralized ledger technologies. The result, in the United States at least, is an industry faced with competing categorizations of digital assets traded on distributed ledgers. The United States Department of Treasury treats cryptocurrencies as "value that substitutes for currency," and therefore subject to money transmission regulation depending upon the business activity at issue. Meanwhile, the United States Internal Revenue Service categorizes cryptocurrencies as property. The Commodities Futures Trading Commission undertook enforcement actions against several industry actors, revealing that, depending on its use, cryptocurrencies may be subject to commodities regulations. State judges and state legislatures have also voiced divergent view points as to the categorization of the digital assets traded on distributed ledgers. These same federal and state governance bodies are now shifting attention to other use cases for distributed ledger technology, including efficient record-keeping for legal issues related to provenance, securities trading, real property records, privacy of records, and improving efficiencies in government processes. As they do so, an open dialogue with industry may pave the way for more novel regulatory approaches to both the technology and its use cases, and other substantive areas of law. In this way, distributed ledger technology disrupts not only the application of current laws, but also the very foundational elements of the way governments conceive of the creation, implementation and enforcement of law. Such legal disruption will only benefit the governed, whether industry or consumer, if both the governed and those charged with governing jointly rise to the challenge of legal innovation.

11. Proof-of-Concept Work

As commissioned by the HKMA, ASTRI is working with a number of banks and related industry players in Hong Kong, including HSBC, Standard Chartered Bank, the Bank of China (Hong Kong), Hang Seng Bank, and the Bank of East Asia, to explore the feasibility of applying DLT to a few sub use cases in the banking industry. Following discussions with the participating banks, the HKMA and ASTRI decided to conduct proof-of-concept work in three areas: mortgage loan application, trade finance, and digital identity management. For each area, an industry working group has been established to discuss the proof-of-concept plan, formulate the scope and design of the proof-of-concept work, and carry out the proof-of-concept work in the HKMA-ASTRI Innovation Hub.

As of the end of October 2016, a DLT prototype for carrying out property valuations for use in the mortgage loan application had been built to the final stage, and testing carried out. With regard to the other two areas (trade finance and digital identity management), the proof-of-concept work is relatively more complex and requires further study to finalise the best operating models. Although the progress of each sub use case varies, we would like to take this opportunity to share their current statuses in terms of the DLT solutions proposed, the lessons learned, the experience of implementing them, along with benefits and recommendations. Since all three sub use cases require more time for the proof-of-concept work to be completed, further updates on each one will be provided in the next stage of this research project, sometime next year.

In the appendix, we have also included a number of sharing contributed by R3[6] and IBM addressing the application of DLT to mortgage-backed securities (MBS) and trade finance, to provide different angles on the application of DLT.

11.1 Proof-of-Concept – Mortgage Loan Application

Introduction

A mortgage is an essential financing tool that provides capital resources to a mortgage applicant (mortgagor) for the purpose of purchasing a property, or obtaining financing by using the property as collateral. As mortgage loans normally involve a large loan amount and a long tenor, banks are obligated to perform sufficient due diligence when processing mortgage loan applications.

The current mortgage loan application process is time-consuming, laborious, manually-intensive and paper-based. It also requires the participation of multiple parties, such as the mortgage applicant (the property owner), banks (sources of financing), surveyors (for property valuation assessment), solicitors (to handle all legal materials), a credit bureau (i.e. TransUnion), and the Land Registry (which maintains an up-to-date list of property title ownership).

Together with five participating banks, the HKMA and ASTRI have formed a mortgage DLT working group to explore the possibility of applying DLT to the mortgage loan application process.

Mortgage loan application process

The current process of obtaining a mortgage loan from a bank is illustrated below:

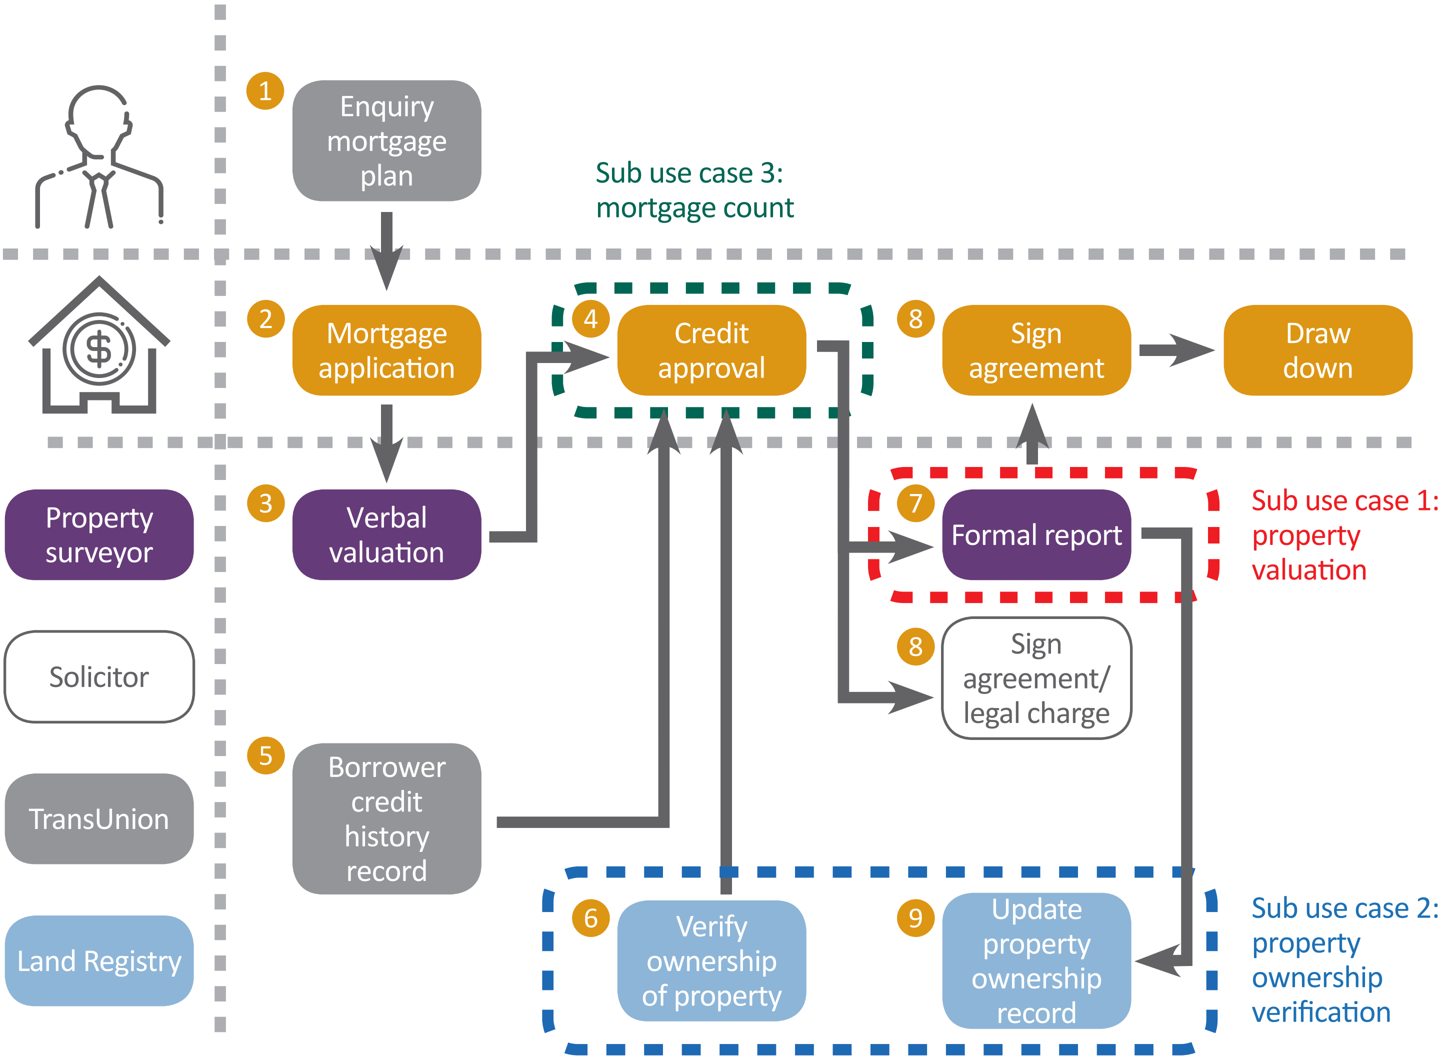
1
Enquiry mortgage plan
Sub use case 3: mortgage count
2
Mortgage application
4
Credit approval
8
Sign agreement
Draw down
Property surveyor
3
Verbal valuation
7
Formal report
Sub use case 1: property valuation
Solicitor
8
Sign agreement/ legal charge
TransUnion
5
Borrower credit history record
Land Registry
6
Verify ownership of property
9
Update property ownership record
Sub use case 2: property ownership verification

Process description:

1. When a buyer and a seller reach an agreement to transact a property and sign a sale and purchase agreement (S&P agreement), the buyer may approach more than one bank for mortgage plans.

2. Banks ask surveyors for an initial property valuation, which is used to provide the estimated mortgage amount. When the buyer has decided which bank to submit the mortgage loan application to, the application process begins. To start with, the buyer provides required information to the bank, including bank statements, income proof, the S&P agreement, and other existing loan information, as well as giving consent for a credit report to be obtained from TransUnion, a consumer credit reporting company.

3. The bank asks a surveyor for a verbal valuation result (via email or fax) to determine the mortgage loan amount.

4. The bank starts the credit approval process based on the information submitted by the buyer and other external parties.

5. A TransUnion report is requested by the bank to check the historical credit records of the buyer, including any negative credit information and the buyer's mortgage count (i.e. the number of the buyer's outstanding mortgages and personal loans). The bank uses this information to make its credit decision.

6. The bank does a property land search at the Land Registry in order to verify the ownership stated in the S&P agreement.

7. Before any legal documents are signed, the bank asks the surveyor to mail the finalised valuation report by post, and uses this report to verify the property status with the credit approval result.

8. The bank makes an offer to the buyer and notifies the solicitor to arrange for the signing of the mortgage loan agreement and the mortgage deed.

9. After all legal documents are finalised, the solicitor sends the signed mortgage deeds to the bank so that the bank can arrange for the drawdown, as well as sending a notice to the Land Registry to enable it to update the title deeds.

The three major areas within the mortgage loan application process

As illustrated in the previous section, the full mortgage loan application process involves a number of time-consuming steps, and some sub-processes are also needed to engage a number of different parties. The working group has identified three areas where the application of DLT could potentially improve the current mortgage business workflow. These three areas are **property valuation, property ownership verification**, and **mortgage count.**

To simplify the study, the working group has made three assumptions as outlined below:

- The DLT network on which banks can share mortgage-related data may contain sensitive personal data. In view of this, the working group believes the DLT network should be operated in private mode. Only authorised parties would be allowed access to the DLT network.

- The mortgage loan application process involves multiple parties, such as solicitors, the credit bureau, and the Land Registry. Getting the participation of these parties may involve setting up commercial arrangements and even making changes to the legal framework (which will be discussed later in this chapter). In order to simplify the study, we have assumed that banks are the only participants in the DLT network, unless otherwise specified.

- As sub use cases 2 and 3 of the mortgage loan application require a complete set of information from the banking industry, to illustrate the benefits it is assumed that all banks are participating in the DLT network.

11.1.1 Sub use case 1: The property valuation

A property valuation report evaluates the market value of the property in question. When a property buyer obtains a mortgage loan from a bank, the bank requires a valuation report from a surveyor. This acts as a safeguard to ensure that the mortgage loan granted is an adequate percentage of the market valuation of the property, and that the bank will be able to recover the outstanding amount if the mortgage is unpaid. In the current business process, banks are obligated to obtain and properly keep the physical valuation reports for two years.

To streamline the process of obtaining a valuation report from the surveyors, ASTRI has developed a DLT prototype to digitise the manual and paper-based process involved. The DLT prototype is able to capture the key data and hash of the digitised valuation reports to the distributed ledgers and store it there, thus increasing the efficiency and security of the process.

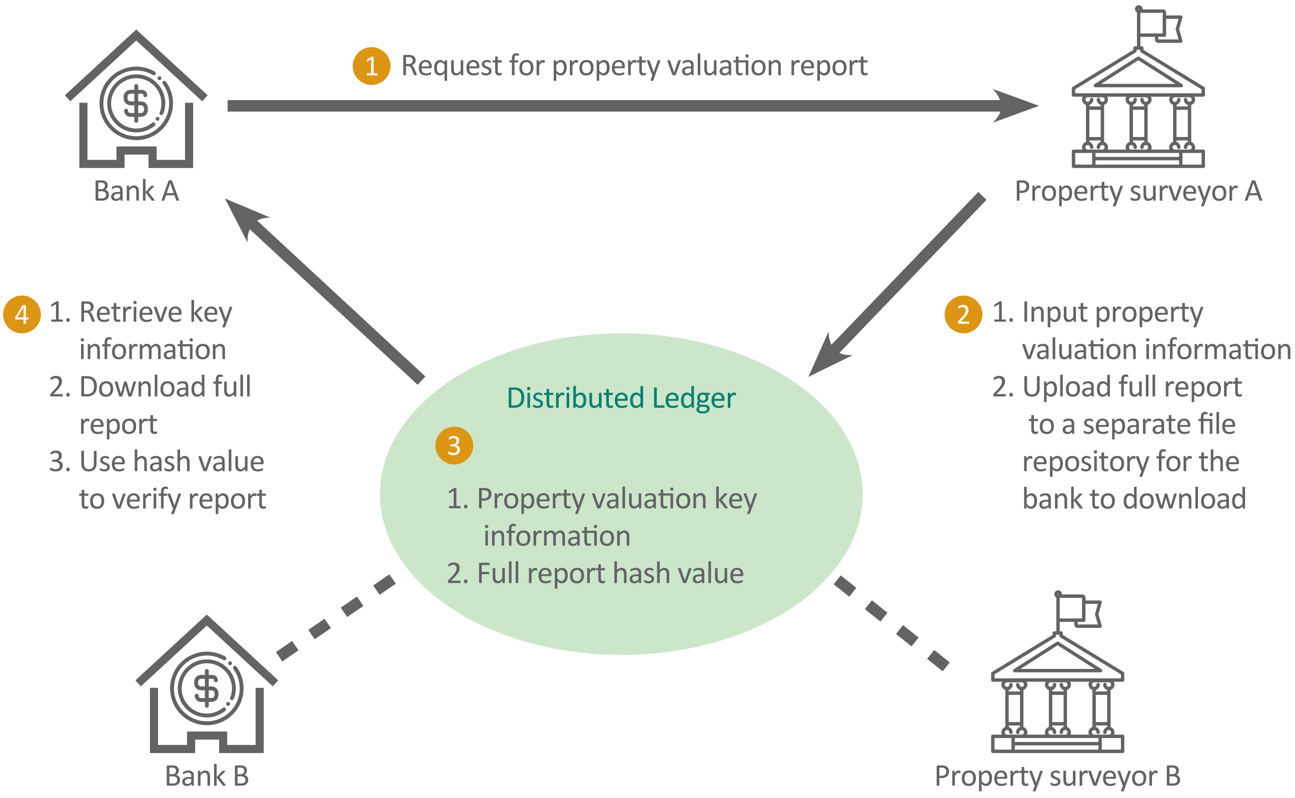
1 Request for property valuation report
Bank A
Property surveyor A
4 1. Retrieve key information
2. Download full report
3. Use hash value to verify report
2 1. Input property valuation information
2. Upload full report to a separate file repository for the bank to download
Distributed Ledger
3
1. Property valuation key information
2. Full report hash value
Bank B
Property surveyor B

The property valuation DLT workflow

Process description:

1. For a mortgage loan application, Bank A requests a property valuation report from Property Surveyor A via email or over the phone.

2. Property Surveyor A prepares the valuation result and the full report. The surveyor then inputs the result and uploads the report via a user interface to the distributed ledgers.

3. A network node generates a full report hash value and stores the most significant property valuation information (including the property address, price and area, etc.) on the distributed ledgers.

4. Bank A retrieves the valuation information from the network and uses the hash value to verify the report.

Current process inefficiencies:

- Physical valuation reports are required
- There is no version control over digitised valuation reports, and it is difficult to validate digital copies of the reports

Improvements achieved by applying DLT:

- Valuation reports can be digitised
- Data formats across the industry can be standardised and process integrity can be ensured
- Increased efficiency of the mortgage application process
- Improved security protection of the valuation reports. Making unauthorised changes to the reports is difficult
- Physical storage space can be saved

The proof-of-concept development has involved building a user interface enabling surveyors to receive valuation requests from banks, and to submit important data from the digitised valuation reports to the DLT network. Key attributes agreed by the majority of the banks in the working group are: the market value, the property address, the property reference number, the area, the month and year of completion, the valuation date, the full report hash value, the name of the valuation company, and the property type.

Upon completion of the approval procedure by the surveyor, this data is submitted to the DLT network where it is shared among participating banks.

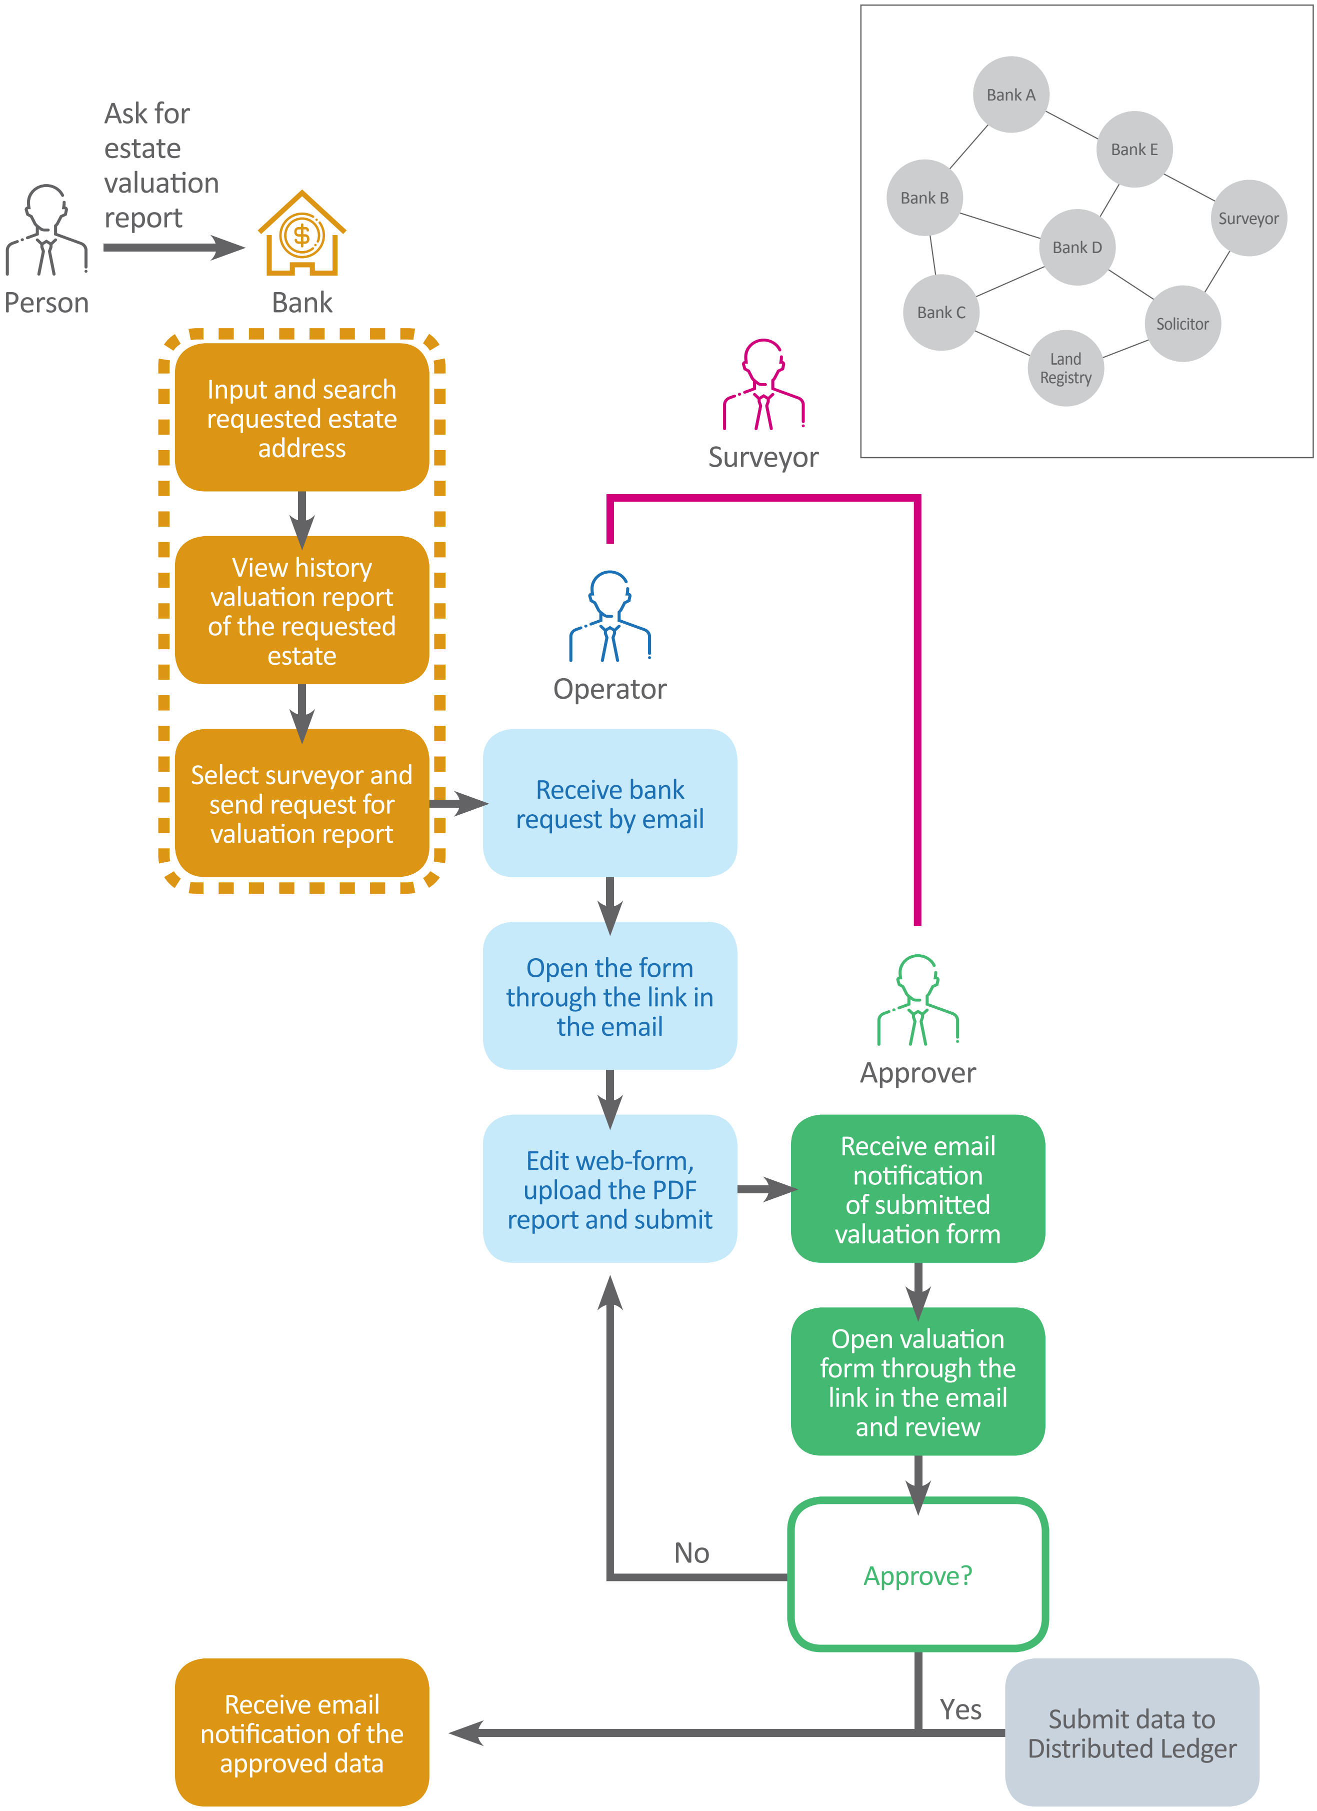
Bank A
Bank E
Bank B
Surveyor
Bank D
Bank C
Solicitor
Land Registry
Person
Ask for estate valuation report
Bank
Input and search requested estate address
View history valuation report of the requested estate
Select surveyor and send request for valuation report
Surveyor
Operator
Receive bank request by email
Open the form through the link in the email
Edit web-form, upload the PDF report and submit
Approver
Receive email notification of submitted valuation form
Open valuation form through the link in the email and review
Approve?
No
Yes
Receive email notification of the approved data
Submit data to Distributed Ledger

The detailed workflow before a property valuation report is submitted to the DLT network

Review of the prototype integration experience (by the working group)

According to the practical experience of one of the members of the working group, the following are the key benefits of adopting DLT for the property valuation process:

- ***Operation costs reduced by eliminating paper-copy report transfer and storage***
- ***Streamlined process, enhancing customer satisfaction***
- ***Operation risk minimised due to the new technology***
- ***Standardised mortgage process established within the financial industry***
- ***"Green Finance" implemented with paperless operations***

Property valuation involves non-sensitive and publicly accessible information, so it is ideal as a sub use case for testing this new technology. Furthermore, as property valuation involves only two types of participants, banks and surveyors, features of permission-based DLT can be efficiently tested without the need to involve too many different parties.

From the first phase of the mortgage loan application proof-of-concept work, the member experienced benefits from the technology in the area of generating trust among participating DLT members. Information was also able to be shared more efficiently and in a less costly manner compared with the existing manual or semi-manual paper-based operations of the property valuation process in the mortgage business. The next objective is to work on other areas of the process, so that DLT can be progressively applied more widely and collaboratively to the mortgage business. Below are four areas of discussion arising from the member's experience in implementing DLT.

1. Using the new system, and its potential for being extended

Suppose a customer, Cathy, applies for a new mortgage from Bank C. Bank C sends a request for a valuation report to a property surveyor. Meanwhile, Bank C sends another request to the DLT network to check whether any other mortgage has been taken out on the property for which Cathy is applying for the mortgage, or whether Cathy is still repaying any other outstanding mortgage. DLT makes the process more efficient and reliable compared to the conventional time-consuming, paper-based, error- and fraud-prone process, and thus facilitates the creation of a trusted business model. At the same time, it enriches the quality of the credit risk analysis carried out during the mortgage application process. Since Cathy has applied for a new mortgage from Bank C, this process provides additional credit information about Cathy which other banks are able to refer to (if needed) when performing additional credit risk assessments on Cathy at a later date.

The property valuation sub use case also provides a foundation for extending potential participation in the platform from banks and surveyors to solicitors and the Land Registry, who will play important roles in the sub use cases to be discussed in the remaining part of this chapter.

2. Technical and Operational Considerations

Cost is one of the major considerations when developing innovative technology such as DLT. In a DLT network, each participating player needs to build its own system and make it compatible with its own existing infrastructure and related policies, particularly in the area of network security. In addition, there is a lack of setup, operation and maintenance guidelines to help in cases such as client system installation, daily user and administrator operation procedures, and business contingency plans in certain system failure incidents. Although implementing a double-node design could mitigate the risk of failure, a system recovery procedure needs to be put in place. However, ASTRI has created a DLT operation monitor dashboard that provides an alert system which could help members deal with these contingencies. In order to better manage the network, the DLT community requires a dedicated party to take on this operation monitoring role, as well as to distribute the software program and oversee the future technical development of this DLT environment.

3. Measurement and Review of DLT platform

In order to ensure the stability and the performance of the DLT platform, ongoing work is required to measure the healthiness of nodes, the "friendliness" of individual members, interaction among members, the functional completeness of the system (particularly in relation to security and privacy), and the ease of maintaining the system, especially when it involves many different system and governance policies across different participants. Both qualitative and quantitative results (such as user experience, system latency, recovery lead time from system failure, etc.) would provide a more comprehensive way to measure the maturity of the technology in the financial industry.

4. Benefits of DLT for the community and society

A DLT platform offers completeness, reliability and transparency of its data set, thus providing a comprehensive data view that enables an authorised party to analyse information at different levels (e.g. according to geographic location, tiers of property value, frequency of property exchange, or seasonal trends and patterns). The results from this sort of analysis could provide the government with valuable information for important policy decisions. Moreover, when supplemented by economic data like interest rates and rental data, information relating to property transactions and mortgages could be used to further study the real estate industry and help the government draft related policy.

11.1.2 Sub use case 2: Property ownership verification

The information required for a bank to approve a mortgage loan application includes property ownership and mortgage count information. In the current process, this information is maintained in the Land Registry according to the Land Registry Ordinance, and extra procedures and processing time are needed both to obtain and to update it. According to the information provided by the banks participating in the working group, it can take up to 40 days from the date a property transaction is completed to the date when the ownership information is updated in the Land Registry. This relatively long period of time is due to the need for manual handling of a number of processes, including document handling and delivery, and the input of information into, and updating of, the computer system. This long lead time can create opportunities for fraud.

Under the arrangement proposed by the working group, banks would be able to store mortgage deeds and title deeds data on the distributed ledger, enabling participating banks to maintain up-to-date records of property ownership and mortgage information. The more banks that join the DLT network, the better the mortgage records will be in terms of accuracy and completeness.

Step 1: Uploading the deeds to the DLT network

Banks upload new mortgage and ownership information to the DLT network for each newly completed mortgage loan application:

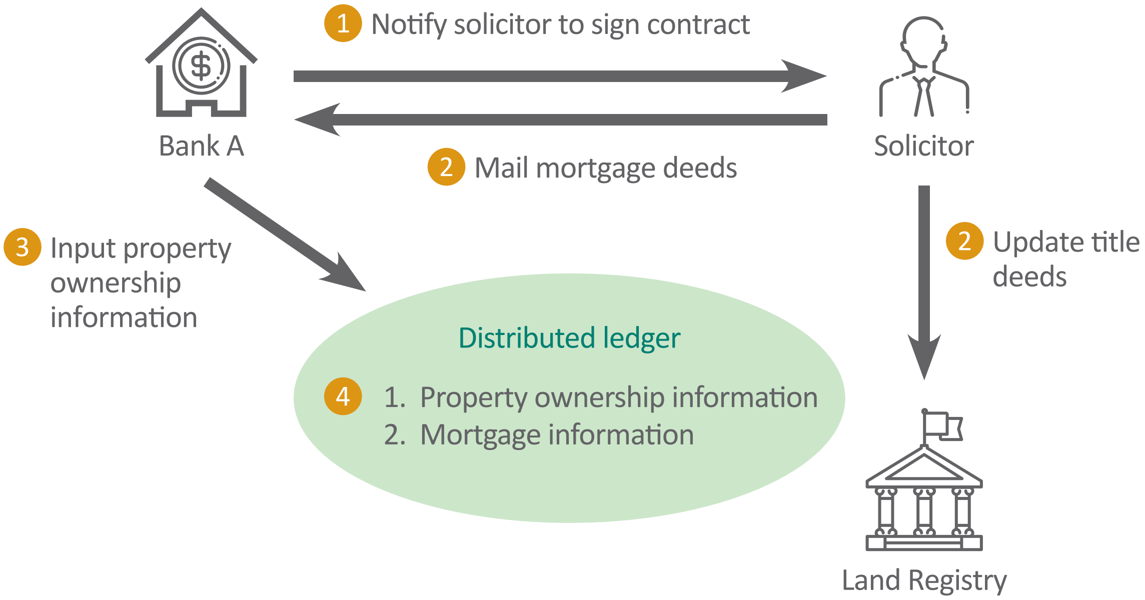

Process description:

1. Once the credit application is approved, the bank notifies the solicitor to sign the agreement.

2. The loan is paid to the solicitor on the title transfer date, the mortgage and title deeds are mailed to the bank by post, and the title deed information is provided to the Land Registry.

3. The bank inputs the property transfer information to the distributed ledgers.

4. The distributed ledgers store the property ownership information and mortgage information. Other banks can retrieve this property ownership information from the DLT network.

Step 2: Retrieving ownership information from the Land Registry if it does not exist on the DLT network

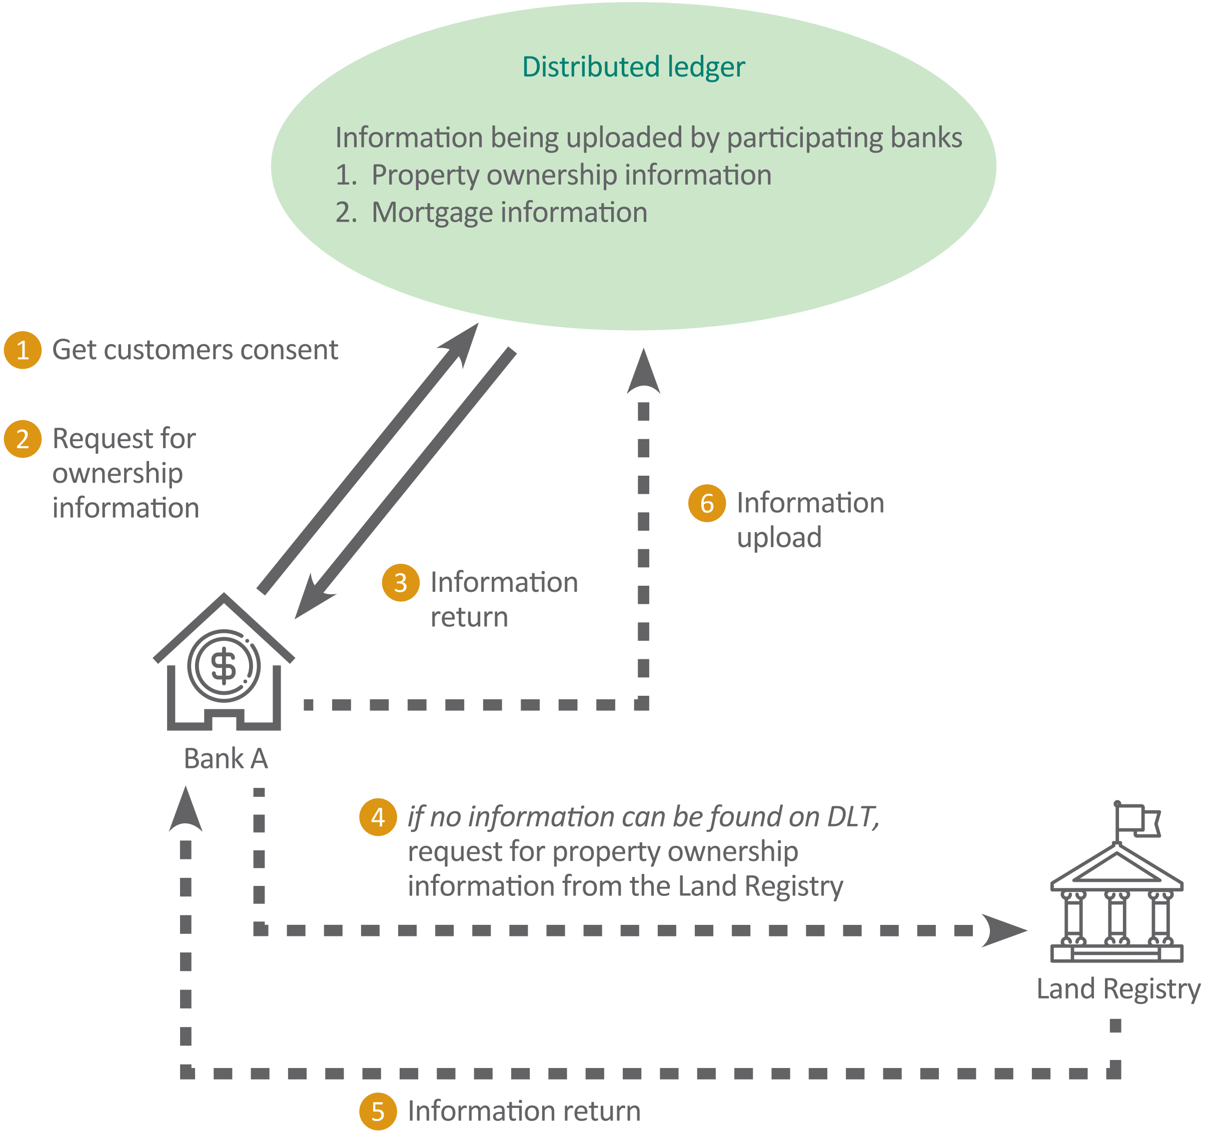
Distributed ledger
Information being uploaded by participating banks
1. Property ownership information
2. Mortgage information
1 Get customers consent
2 Request for ownership information
3 Information return
6 Information upload
Bank A
4 *if no information can be found on DLT,* request for property ownership information from the Land Registry
Land Registry
5 Information return

Ownership information retrieval DLT workflow

Process description:

1. Bank A obtains customer consent for a request for property ownership information

2. Bank A asks for property ownership information

3. Bank A retrieves the information from the DLT network

4. If no information can be found on the DLT network, Bank A asks for property ownership information from the Land Registry

5. Bank A retrieves the information from the Land Registry

6. Bank A uploads the updated property ownership information to the DLT network

Advantages

Since some time would be needed to gain a consensus from all existing customers by which they agreed to upload their existing mortgage and ownership information to the DLT network, it would be better to carry out this uploading each time a new application is made. This will mean that banks can start to benefit from the arrangement after only a short period of time.

Disadvantages

It may take banks years to update all their customers' mortgage and ownership information to the DLT network, as titles to properties do not change often. Ownership of some properties may not change for more than 30 years. Until all properties have been uploaded to the DLT network, some information will inevitably be missing, which may have an impact on the effectiveness of the DLT network.

Current process inefficiencies:

- It can take a long time for the Land Registry to update ownership information due to the complex manual processes required
- The time lag in updating information could lead to cases of fraud

Improvements achieved by applying DLT:

- DLT can facilitate the exchange of property ownership information among property owners, buyers, banks, and solicitors
- Efficiency in obtaining the latest property ownership information is increased
- Mortgage and ownership information is available on the DLT platform, which can help in implementing the mortgage count sub use case discussed in the next section

11.1.3 Sub use case 3: Mortgage count

As in sub use case 2, mortgage count information takes up to 3 working days to be obtained from a third party institution, one that acts as an agency requesting mortgage count information from all of the banks in Hong Kong. The Mortgage Working Group has suggested that mortgage deed information could be shared across a DLT network, thus improving the efficiency and accuracy of the business process.

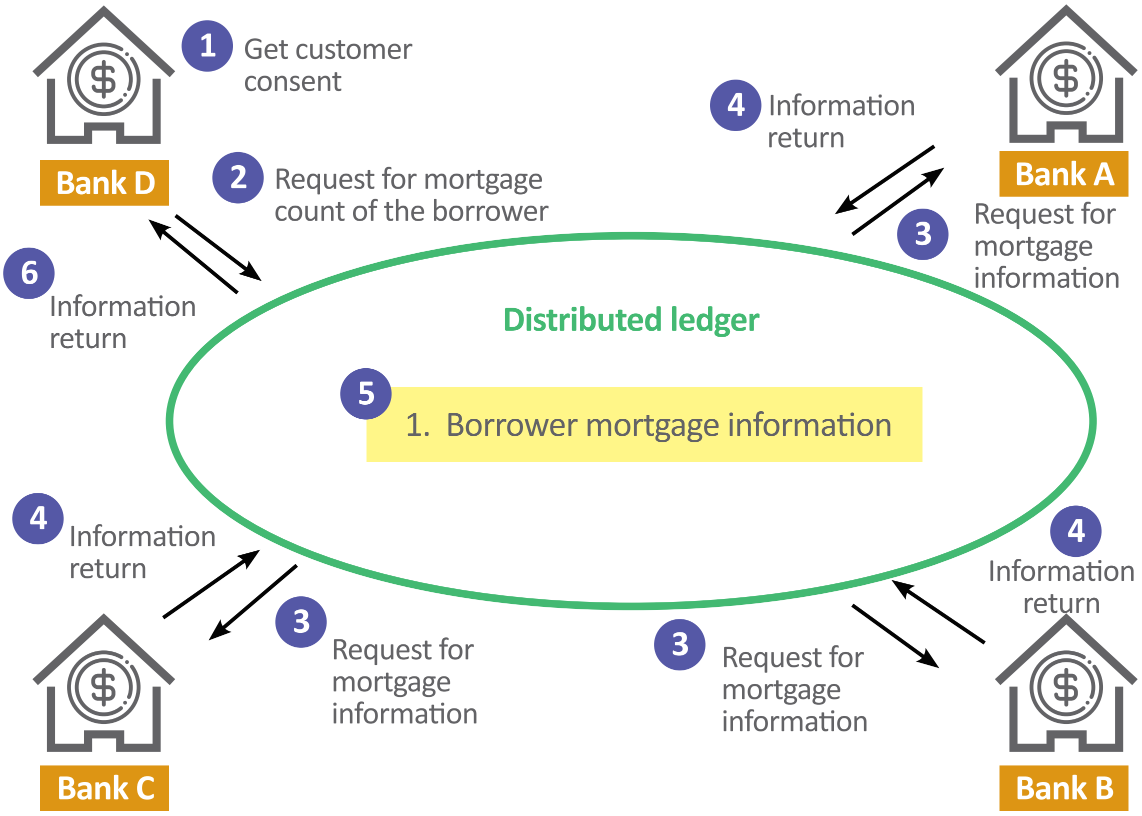
1 Get customer consent
Bank D
2 Request for mortgage count of the borrower
6 Information return
4 Information return
Bank A
3 Request for mortgage information
Distributed ledger
5
1. Borrower mortgage information
4 Information return
3 Request for mortgage information
Bank C
3 Request for mortgage information
4 Information return
Bank B

The mortgage count DLT workflow

Process description:

1. Bank D gains customer consent for a request for mortgage count information
2. Bank D asks for a mortgage count for the borrower
3. The DLT network sends the request to member banks through an application program interface developed for obtaining this information
4. Member banks send the information to the distributed ledgers through the interface
5. The borrower’s mortgage information is stored in the distributed ledgers
6. The requesting bank (i.e. Bank D) retrieves the information from the distributed ledgers, including the borrower’s mortgage indicator and corresponding property address(es), etc.

Current process inefficiencies:

- Long turnaround time (up to 3 working days) when requesting mortgage information from TransUnion

Improvements achieved by applying DLT:

- Participating banks are able to share the number of a customer's outstanding mortgage loans on the distributed ledgers
- The efficiency and accuracy of the entire process is enhanced

Enhancements in mortgage count checking

The title ownership of a property shows the rights of an individual or an organisation to manage a property, including selling it to others, collateralising it for financing (a mortgage), etc. Nowadays, banks rely on the Land Registry's records, in which the updating of any ownership transfer information can take up to 40 days to complete.

The adoption of DLT could help reduce the time lag associated with the title deed update by allowing banks to update the registration directly on the distributed ledgers. The updated records could thus be made available almost in real time for verification by banks within the mortgage application process.

11.1.4 Conclusions – Benefits and challenges

Benefits of DLT

A DLT platform provides an immutable, transparent and tamper-proof way for all stakeholders of instantly distributing digitised documents among all trusted parties. Adopting DLT for the mortgage loan application process could facilitate the digitising of paper documents and leave a complete audit trail for the mortgage loan application, result in a significant saving of manpower and storage space, as well as lowering the risk of manual errors. Transaction records can be placed in the DLT network within a very short period of time, thus dramatically reducing document collection and delivery times, from days to minutes. It also provides a transparent and chronological log for future investigation and auditing. In addition, digitising documents can reduce the need for manual data input and thus minimise the chance of human error. Moreover, the standardisation of the data format can help industry stakeholders enhance their record management procedures and streamline their business workflow.

This sub use case shows that DLT provides a great opportunity for standardising and digitising the paper documents involved as well as maintaining a complete audit trail for the mortgage business process. In the mortgage business process, there are six types of documents that have the potential to be digitised and stored in a DLT network. The benefits associated with digitising those documents are listed below:

Documents to be digitised	Benefits
Valuation reports	• The valuation reports and materials transmitted between banks and surveyors are standardised.
Instruction manuals and drawdown letters	• The data flow and input between banks and solicitors are standardised.
Mortgage deeds	• Verified true information about property owners, borrowers and guarantors can be shared. • Consistency and transparency of information is ensured. • Risk of mortgage fraud and manual errors is reduced. • Operating efficiency in accessing customer mortgage holdings is improved.
Title deeds	• Manpower required for physical filing and retrieval can be saved. • The drawdown process is automated. • Verification of property ownership becomes possible.
Credit records	• Credit assessments are facilitated. • Credit risk is reduced. • Intermediary costs are reduced.
Other land search documents (e.g. court orders)	• Important information that may affect valuation results and mortgage loan approvals can be obtained in a timely manner.

Challenges, and the next step forward

All three of the sub use cases discussed under the mortgage loan application proof-of-concept above have the potential to bring significant unarguable benefits to the industry. However, a number of issues pose challenges to the successful practical adoption of the technology in the real business world, especially in relation to the second and third sub use cases (property ownership verification and mortgage count). Below we discuss various major challenges identified by the Mortgage Working Group, and offer some possible steps for tackling them.

- **Electronic Transaction Ordinance**

  The Electronic Transaction Ordinance (ETO) (CAP 553), schedule 1 explicitly excludes any deeds, conveyances or other documents or instruments in writing, judgments, and lis pendens[22] referred to in the Land Registration Ordinance (CAP 128). This means that these documents must be in written form in order to be legally binding.

  The Mortgage Working Group needs to explore the possibility of amending the ETO (CAP 553) to cover deeds, conveyances, instrument judgments, and lis pendens that are stored and transacted digitally. The working group is making reference to the precedent case of the e-Cheque project, in which the HKMA worked with the Office of the Government Chief Information Officer (OGCIO) to amend the ETO (CAP 553) so that cheques were taken out of schedule 1 in the year 2014. Since then, e-Cheques have been able to be legally written by and transferred between banks and individuals.

- **Land Registry Ordinance and Property Conveyance Ordinance**

As with the ETO (CAP 553), the Land Registry Ordinance (CAP 128) only covers written forms of deeds, conveyances and other documents in relation to a property, and the Property Conveyance Ordinance (CAP 219) specifically states that related documents should be signed, sealed and delivered. The next step is to explore the possibility of amending the Land Registry Ordinance (CAP 128) to also include digital mortgages and title deeds, and the Property Conveyance Ordinance (CAP 219) to allow related documents to be handled electronically.

- **Participation of the Land Registry**

In the property ownership verification and mortgage count sub use cases described earlier, banks could only submit newly completed mortgage deeds to the DLT network, not existing ones. The Personal Data (Privacy) Ordinance (PD(P)O) (CAP 486) includes six Data Protection Principles (DPPs) which cover the lifecycle of a piece of personal data. One of these states that personal data must be used for the purpose for which the data is collected or for a directly related purpose, unless voluntary and explicit consent for a new purpose is obtained from the data subject. Therefore, the bulk upload of all existing mortgage deeds to the DLT network would involve a substantial consent-gaining exercise from all existing customers, which may not be feasible.

Until all property ownership and mortgage deed information has been uploaded to the DLT network one by one as described in the property ownership verification sub use case, banks will have to use alternative ways to collect information from other banks, rather than being able to retrieve the information directly from the DLT network.

Even if the Land Registry does not participate in the DLT network, banks can still contribute mortgage and ownership information to the network once a mortgage loan application has been approved. However, the database will never be complete because not every property transaction involves a mortgage loan.

To enable the mortgage count or fully functional mortgage loan applications on the DLT network, participation by the Land Registry is crucial. It is recommended that banks try to engage the Land Registry in studying the feasibility of this new technology.

- **Dynamic authorisation in DLT**

Although data is always available for all DLT nodes on the network, PD(P)O requirements mean that node controllers (e.g. banks) should not retrieve data on a network without a legitimate need. The mortgage count sub use case required a bank to obtain authorisation from the data subject before accessing the relevant information.

To comply with the PD(P)O requirements mentioned above, a dynamic authorisation mechanism is required. Our preliminary study indicates that dynamic authorisation could be enabled in DLT by making use of a digital identity. More details will be discussed in the second part of the white paper, to be published next year.

- **DLT searching capabilities**

With reference to the mortgage count sub use case, if all mortgage deeds were available on the DLT network, a bank could obtain the total number of mortgage loans owed by a single applicant simply by searching the DLT network. However, as described in an earlier chapter, records of transactions are stored chronologically along the chain of the ledger. As time goes by, and as more data is stored on the ledger, searching information in the ledger becomes more and more time-consuming.

Further research is required to ascertain possible ways of enabling efficient searching capabilities on DLT. Solutions might include building an index on top of DLT, or enhancing the DLT protocol to enable searching capabilities. ASTRI will continue to study the subject and provide further updates in the next stage of the white paper.

- **Readiness of the ecosystem**

To fully leverage the features of DLT, it is imperative that the wider ecosystem is in full readiness. For example, a smoothly operated blockchain mortgage system will require the participation of all relevant parties, such as surveyors, solicitors and the Land Registry, etc. We observe that some of the industry stakeholders are proactively adopting the emerging technology in their business processes. However, it will take time for the ecosystem to grow and become mature enough for the technology to fully take off.

- **Implications of the PD(P)O**

The implications of the PD(P)O will require further study. For example, while banks can contribute mortgage information to the DLT network, it is not entirely clear if consent is required from customers each time a bank wants to contribute ownership information. Given the immutable nature of information in the DLT network, the question of how the DLT network could meet PD(P)O requirements regarding the data retention period and individuals' rights of correction has yet to be resolved. Currently there is no official facility available to search for the number of properties owned by an individual. If the storing of mortgage data in a DLT network was implemented, it would become possible to search for ownership details. However, means of complying with the PD(P)O in relation to such mortgage data searches need to be explored.

11.2 Proof-of-Concept – Trade Finance

Introduction

Trade finance is an essential business tool. It is therefore an important service offered by banks, which act as intermediaries providing assurance and liquidity to both importers and exporters involved in global supply chains. However, the processes involved are without doubt labour-intensive and time-consuming because of the due diligence process, and heavy reliance on third-party paper documentation to reduce business risk. Various consortiums around the globe have carried out different types of proof-of-concept work using DLT in attempts to streamline the manual processes of import/export documentation, improve operational efficiency, reduce errors, and increase convenience for all parties through the digitisation of documents. The process also aims at making companies' working capital more predictable. Like other major financial markets, the banking sector in Hong Kong faces similar challenges.

In this connection, together with a trade service provider and the five participating banks, the HKMA and ASTRI have formed a DLT Working Group to explore the possibility of applying DLT to the trade finance process.

Open Account Trading

Traditionally, banks play an important role in international trade by providing working capital, financing manufacturing activities, preventing fraud, and guaranteeing payment for companies involved in trade transactions. Recent years have seen an emerging trend by which trading has transitioned from being carried out on a documentary credit basis to being on open account terms. An open account trade means that a seller delivers the goods to a buyer directly before any payment is due, and without relying on documentary credit issued by a bank. A major factor driving corporates to shift trade from documentary credit terms to open account terms has been the advances in technology which have made communication and the exchange of information between market participants over the Internet much easier.

Open account trading opens up opportunities for banks to offer a wide variety of financing solutions, such as invoice discounting, factoring, and buyer guaranteed financing. However, open account financing also means that banks are subject to a heightened risk of fraud and money laundering as compared with documentary trade, due to the lack of third-party documentation and the low visibility of the transaction status.

Key stakeholders in the trade finance ecosystem

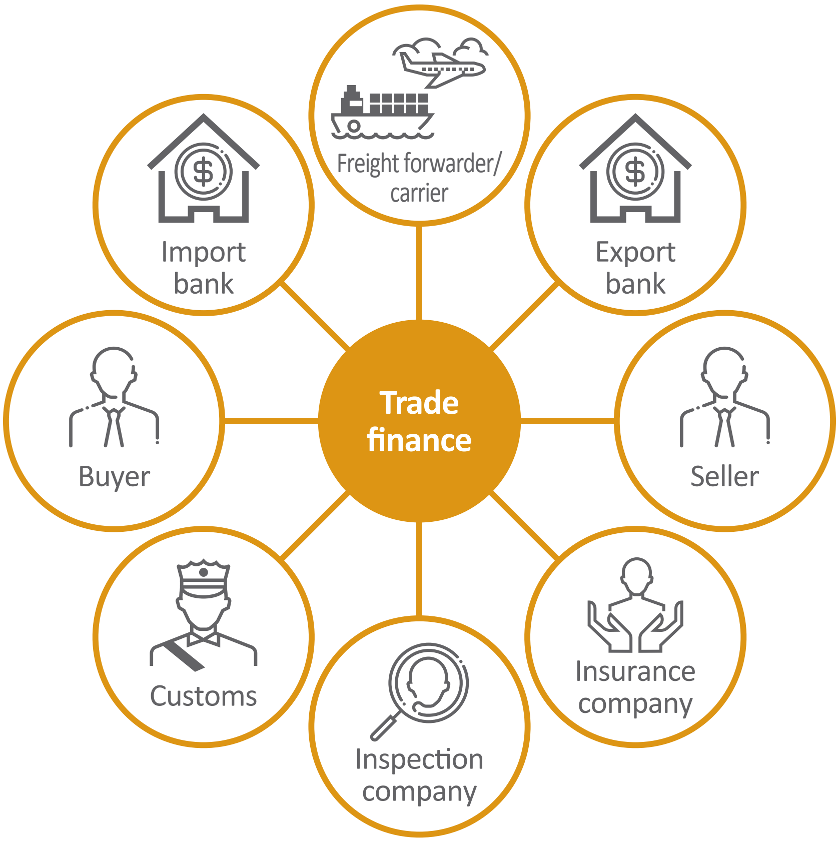

Current open account trading process

A simplified open account trading process of a sale from a supplier to a retailer is illustrated below:

Open Account Trade

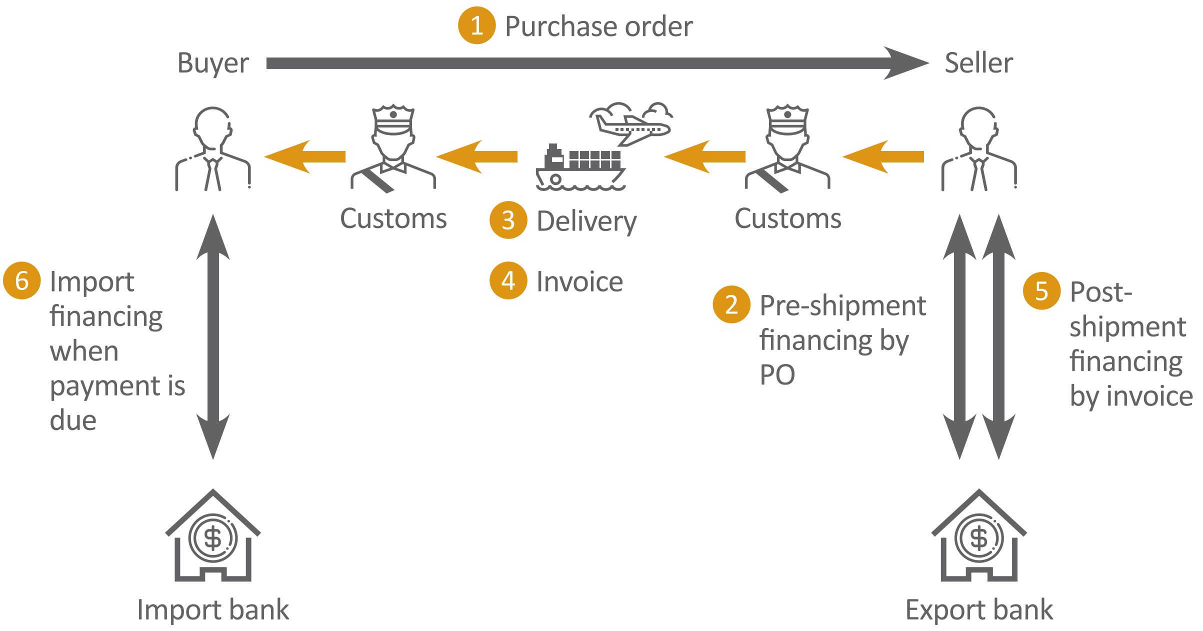

1 Purchase order
Buyer
Seller
Customs
3 Delivery
Customs
4 Invoice
6 Import financing when payment is due
2 Pre-shipment financing by PO
5 Post-shipment financing by invoice
Import bank
Export bank

Process description:

1. Buyer and seller agree on a trade transaction on open account terms at a specific date and time. A buyer (e.g. a retailer) creates a Purchase Order (PO) and sends it to a seller (e.g. a supplier) for confirmation.

2. The seller presents the PO to the export bank to request financing. The export bank determines pre-shipment financing subject to risks, such as that the PO is not confirmed or that financing is duplicated.

3. After the goods have been transported to the export terminal and inspected by customs, they are transported by freight to the import country. After customs inspection in the import country, the goods are delivered to the importer. Neither the import bank nor the export bank has access to the status of the goods delivery.

4. The seller provides invoices, bills of lading and other transport documents to the buyer via a document courier.

5. The seller asks the export bank for post-shipment financing by presenting the invoice. The export bank determines post-shipment financing for the seller, subject to the risk that duplicated financing and/or fraudulent transactions may have taken place.

6. The buyer accepts the invoice and asks the import bank for import financing. The import bank determines the import financing, subject to the risk that a fraudulent transaction, duplicated financing, or financing of the seller may have taken place.

The application of DLT to trade finance

Despite the shift from documentary credit to open account payment terms, trading parties continue to require banking services for financing, risk mitigation, and data transfer and matching. Banks need a holistic view of the contract terms of a trade transaction and the flow of goods between sellers and buyers in order to provide better value services in the supply chain, and to lower risk when performing their financing roles. The trade finance proof-of-concept aims to illustrate the ability of DLT to improve the transparency of trade transactions and facilitate the banks' financing services for customers in three major areas:
the use of smart contracts in open account trade, tracking of trade transaction statuses, and **the matching of invoices to POs.**

Documents to be put onto a DLT network

Based on the process described above, the following three types of documents are considered most critical for trade transactions and will be made available in the DLT network:

1. POs

2. Commercial Invoices: this proof-of-concept takes as its scope one commercial invoice raised for a PO

3. Transport Documents, which may include, Bills of Lading, House Air Waybills, and Sea Waybills.

Image files (in jpeg, gif, tif format) of other trade documents (such as packing lists and inspection reports) can be uploaded to the DLT platform or to a central repository. Although title transfer of transport documents could well be facilitated on DLT, it is not included in the current scope of the proof-of-concept work. The transport document on the DLT platform is simply used for information sharing, and provides information about the status of a goods shipment.

DLT features to be tested

The following four features of DLT will be demonstrated in the proof-of-concept work:

1. **Shared repository** – Multiple stakeholders in a trade transaction need to be able to view common information.

2. **Multiple writers** – Multiple stakeholders in a trade transaction take actions that need to be recorded and modified.

3. **Intermediaries (who add cost and complexity)** – Removal of a "central authority" record keeper or intermediaries has the potential to reduce costs (e.g. fees) and complexity (e.g. multiple reconciliation)

4. **Interactions are time sensitive** – Reducing delays has business benefits, such as reduced settlement risk and enhanced liquidity

11.2.1 Smart contract in Open Account Trade

A trade transaction is normally formalised by a PO. This is agreed on between and signed by a buyer and a seller, and is a legal document for executing the trade. A valid PO is often a key document setting out the trade terms, which make up essential information by which banks determine when and to what extent financing should be offered to the buyer and seller.

The Trade Finance Working Group has proposed the development of a smart contract template to record a purchase agreement between a buyer and a seller under open account trade terms. Based on the smart contract recorded on the DLT platform, the buyer, the seller and their banks can submit trade documents and gain access to transaction data subject to DLT platform specified permission. Banks can provide trade finance to customers more promptly according to 'triggers' based on the terms of the smart contract. As a result, transaction transparency is improved and the bank process for providing financing to customers becomes faster and more efficient.

Key DLT smart contract features to be built and tested:

1. When the buyer creates a PO, validated and accepted by the seller, a smart contract is formed and the PO is digitally signed by both the buyer and the seller.

2. Both the import bank and the export bank are permitted to read the PO. Financing triggers are generated when conditions under the smart contract are met. This allows the import bank and export bank to provide financing in a timely manner to the buyer and seller in the course of the trade transaction.

3. On the payment due date, the smart contract alerts the buyer and the import bank of the payment obligation, as well as reminding the seller and the export bank to settle the export financing.

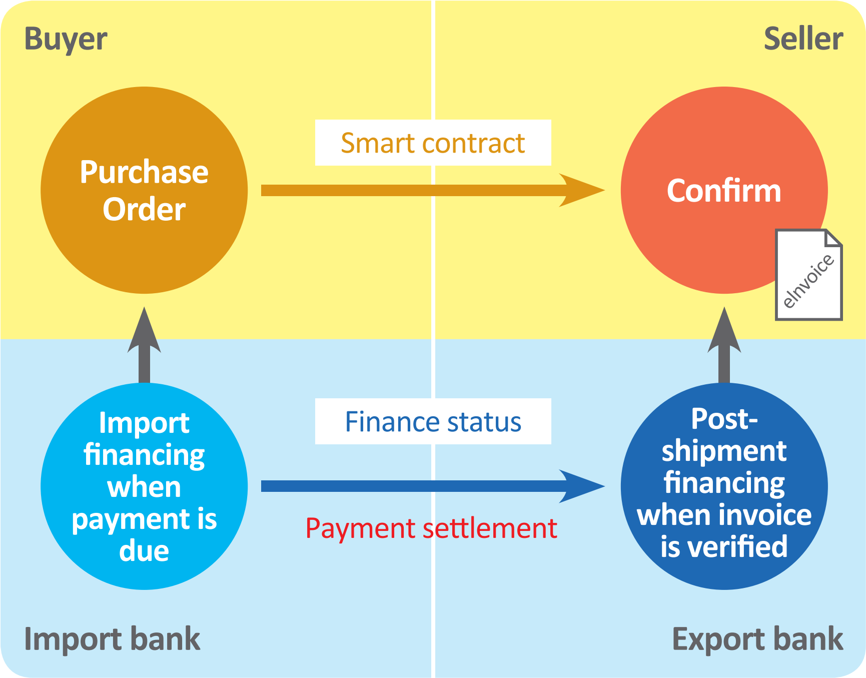

An example of the workflow of a smart contract

Current process inefficiencies:

- No standardised structure of a PO
- No way to verify the latest version of a PO
- Low visibility of contract terms and amendments to banks
- High error rate due to manual mistakes
- Inefficient and costly due to:
  - o Handling of paper-based documents
  - o Difficulty in authentication by banks
- Risk of fraud

Improvements achieved by applying DLT:

- Digitised and standardised PO
- Deployment of smart contract automates trade execution
- Facilitate cross-border trade finance by collaboration with overseas DLT projects
- Financing to customer is faster and enable stage financing along the supply chain
- Lower risk of financing and avoid double financing

11.2.2 Tracking the status of trade transactions: flows of Goods & Funds

Tracking the status of a trade transaction and the flow of goods is yet another challenge for open account financing, as the presentation of third-party documentation (such as transport documents) to the bank is not mandated, unlike with documentary credit.

The DLT solution proposed by the Trade Finance Working Group is to store and share key trade documents on the DLT network, so they are accessible by all stakeholders in the transaction. Moreover, a number of data feeds from the logistics service providers will be collected at key touch points to reveal the most current status of the flow of goods. The key information to be stored on the DLT network includes the goods delivery status reported by the seller, shipping information from logistics service providers, and the funding status from banks. This information will also be used for tracking the status of goods and funds. As a result, the visibility of the goods and the flow of funds in a transaction will be greatly enhanced, and the risk of fraudulent transactions or financing lowered.

Key trade transaction status tracking features to be built and tested:

1. Financing status

    - The financing status, i.e. pre-shipment financing, post-shipment financing for the seller, and import financing for the buyer will be tracked on the DLT ledgers and will be visible to all stakeholders in the transaction, helping prevent duplicated financing.

2. Delivery status

    - The delivery status of goods will be tracked using transport documents provided by the seller, and the actual "Cargo shipped" status will be based on the data feed received from the Freight Forwarder.

    - The delivery status will change to "Goods delivered" when the goods arrive at the destination port.

    - Additional tracking of the flow of goods using the data feed from the carrier and/or freight forwarder will lower the risk of fraudulent transactions.

3. Payment status

    - When the buyer accepts the invoice and the trade documents submitted by the seller, tracking of the payment status is updated, and an alert is generated to the seller, the import bank and the export bank.

    - The final payment by the buyer on the payment due date will be monitored.

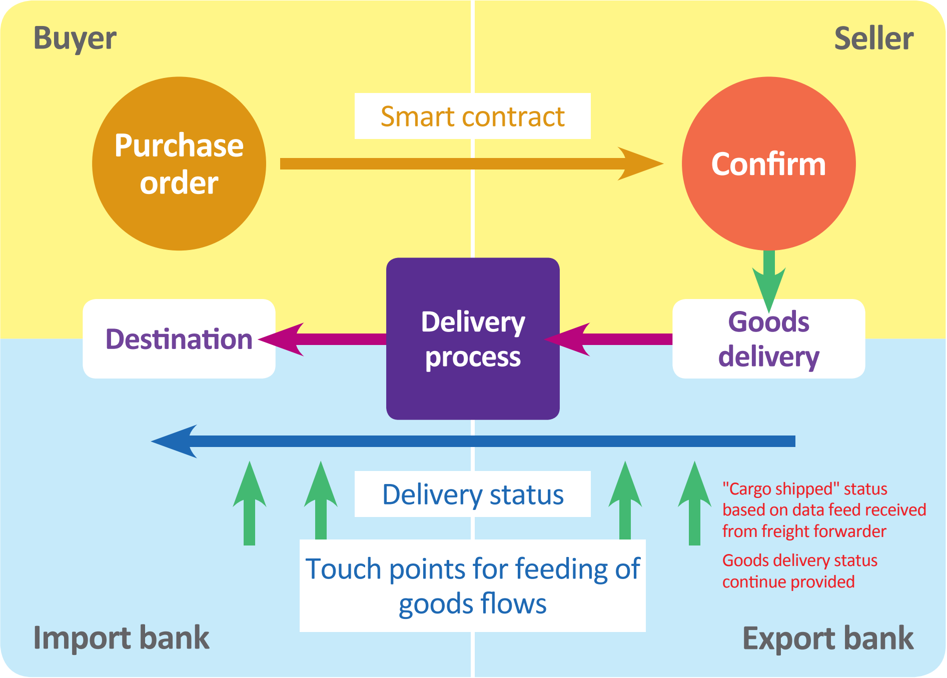
Buyer
Seller
Purchase order
Smart contract
Confirm
Destination
Delivery process
Goods delivery
Delivery status
"Cargo shipped" status based on data feed received from freight forwarder
Goods delivery status continue provided
Touch points for feeding of goods flows
Import bank
Export bank

An example of an update of the shipment delivery status

<u>Current process inefficiencies:</u>

- Difficult to track:
  - o Different stages along the Supply Chain (goods, document & fund flow)

<u>Improvements achieved by applying DLT:</u>

- Facilitates the ability of the buyer, seller, banks and forwarder to exchange trade documents, and update shipment, delivery and payment status
- Banks are able to verify invoices and other documents from third parties to reduce fraudulent financing; and able to review the goods flows and financing status (already financed amount and nature) of the trade transaction in evaluating a financing request of a customer
- Enhanced visibility of financing status of a PO/ invoice to reduce risk of double financing

11.2.3 Information matching of Invoices with Purchase orders

Most of the data required for preparing trade documents originates from the PO. Data inheritance from the PO to the invoice when issued can reduce the need to prepare invoices manually and the possible errors. Automatic matching of the invoice (and other trade documents) submitted by the seller with the PO can identify discrepancies and provide early alerts to the buyer and the financing banks.

The key features of invoice and PO matching to be built and tested:

1. When the seller creates an invoice against a PO, the DLT platform matches the line item details of the invoice and the transport documents with the PO.

2. Depending on the status of the matching results, either "clean" or "discrepant", post-shipment financing request will be triggered from the seller to the export bank, or an alert will be sent to all parties for reconciliation.

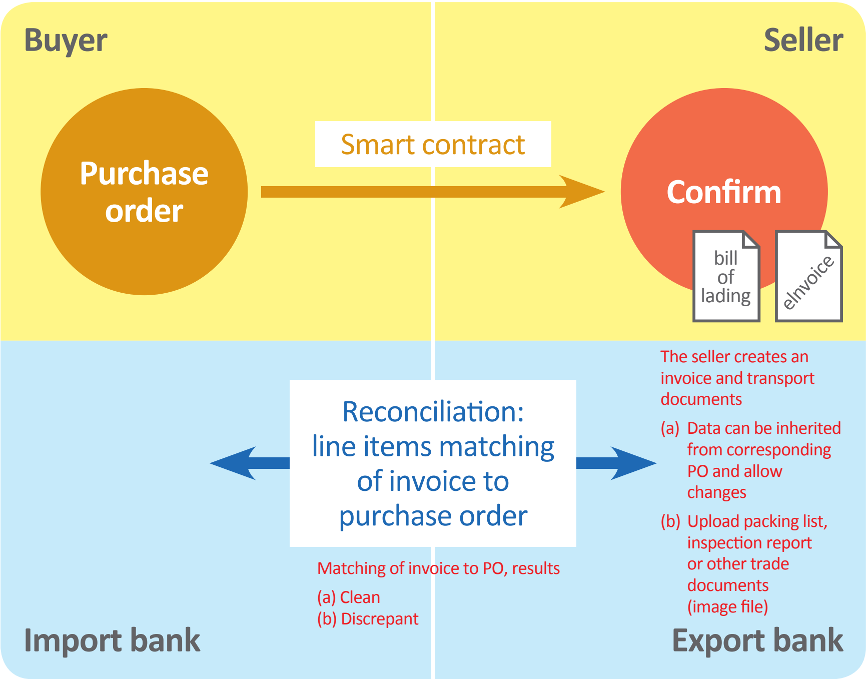
Buyer
Seller
Purchase order
Smart contract
Confirm
bill of lading
eInvoice
The seller creates an invoice and transport documents
(a) Data can be inherited from corresponding PO and allow changes
(b) Upload packing list, inspection report or other trade documents (image file)
Reconciliation: line items matching of invoice to purchase order
Matching of invoice to PO, results
(a) Clean
(b) Discrepant
Import bank
Export bank

An example of invoice and PO matching

Current process inefficiencies:

- Inconsistent and labour intensive due to manual preparation and reconciliation
- Possible double financing of single invoice or PO
- Risk of trade fraud due to fraudulent trade documents
- Extra effort needed to check outstanding PO items when partial shipment is allowed

Improvements achieved by applying DLT:

- Raises invoice to ledger, matches invoice details with PO
- Both buyer and seller can share the verified invoice with their banks to get post shipment financing (for seller), or import financing (for buyer)

11.2.4 Benefits of and possible challenges to DLT

Benefits

The tamper-proof, transparent and traceable nature of DLT allows relevant trade data to be digitised and stored on the DLT platform and shared with participating stakeholders on a real-time basis. Adopting DLT for the open account trade finance process could make the trading terms of purchase agreements more visible for all participating partners, and improve the transparency of open account trade transactions. This transparent system could also enable banks to provide better customer financing services. Trade data stored in the DLT network can be verified and cross-checked, thus lowering the risk of fraudulent financing. Last but not least, the standardising and digitising of trade data also helps in reducing reconciliation efforts and cutting down on manual errors.

Possible challenges

Although DLT provides a great opportunity for improving business and operation processes (as shown in the mortgage loan application proof-of-concept case), there remain a number of challenges that may hinder the successful application of the technology.

Managing the legal and regulatory issues (e.g. privacy issues) arising from data sharing within the DLT network, the non-standardised processes and procedures across the trade finance industry, and cross-border regulations, are all challenges that the working group needs to explore at the next stage.

As at the end of October 2016, the working group is moving towards the development of a prototype for testing the various DLT features. Development of this prototype has begun in the HKMA-ASTRI Fintech Innovation Hub. More details on the prototype and its testing will be reported in a later white paper, scheduled to be released in the second half of 2017.

11.3 Proof-of-Concept – Digital Identity Management

Introduction

Customers are asked to present personal identification when they open bank accounts, subscribe to banking services, or perform online transactions. Customers' assets are also tied to their personal identities. Identity theft can lead to financial loss and asset loss if personal identity information is not properly protected. In addition, other personal data (e.g. data relating to income, liabilities, default history, etc.) linked to one's personal identity constitutes useful information by which banks can find out about their customers' financial situations and credit history. Governments and regulatory bodies require financial businesses to collect and examine customer information, and to comply with KYC rules and anti-money laundering guidelines.

The growing number of regulatory requirements relating to KYC and money laundering has increased the incentives for financial institutions to find more cost-effective mechanisms for implementing these requirements. Also, demand for more convenient and user-friendly on-boarding processes for customers, especially for overseas customers, is increasing.

Estonia, in northern Europe, has started an e-Residency project in which a digital identity service is applied to banking and advisory services. An e-Residency ID links to the personal details of residents, and allows them to present their identity digitally and also sign digitally when using financial services.

Against this background, the HKMA and ASTRI have formed a Digital ID Working Group with the five participating banks to study the feasibility of applying DLT to digital identity management, with a view to addressing a series of challenges facing the banking sector in Hong Kong. This group aims to identify major issues and possible solutions, and also to ascertain whether DLT could provide a cost-effective solution for digital identity management.

Potential Benefits of DLT-based Digital Identity Management

DLT is a technology that has the potential for banks to share identity information in an effective and secure manner. Customers' records and documents can be digitised, updated and shared among banks through a DLT platform. Such an arrangement would have the following benefits:

- Customers would no longer be required to repeat the same processes and submit the same personal information to different banks for KYC purposes;
- The costs and resources needed for the identity verification process could decrease, as the information would be readily accessible and shared in the DLT ledgers;
- Checking of customer history could be carried out efficiently, as customers' information would be available in the DLT ledgers; and
- A better customer experience would result.

Working Group Status

The proof-of-concept work conducted by the Digital ID Working Group was still at a very early stage as of October 2016. Further study is required concerning the best way to gather and authenticate personal information from a reliable source, the authentication process for accessing such personal information, and the possible legal implications. The working group aims to report the details of the findings and recommendations relating to digital identity management at the next stage of this research project.

Appendix 1 – Use case – Mortgage-Backed Securities (by R3)

Another proposed use case that has been highlighted at conferences but requires future research relates to mortgage-backed securities (MBS).

A traditional MBS rests on the bank taking a position on the actual mortgages, performing analysis on them (e.g. creating pools/tranches), then issuing specific allocations to the market, pegged for specific tenor/risk baskets. This is a risky period for the MBS issuer. Prior to the credit crisis, auction notes were used to finance this bridging period. The drying up of the auction notes market was one of the main reasons for the fall of Bear Stearns.

It is possible to imagine a number of scenarios if mortgages were held on a distributed ledger:

1. *If the mortgages were anonymised from an individual borrower perspective but the other data exposed on the ledger, it is possible issuers could model the mortgage-backed securities without holding them. This would allow the issuer to synthesise the pools without holding the mortgages, offer the MBS for sale, and essentially assemble the underlying securities as needed, without warehousing risk.*

2. *If the individual atomic cash flows from the individual mortgages were exposed on the ledger, and assignable, the MBS issuer could model the MBS out of the individual cash flows (e.g. coupon and tenor strips) and then match the specific cash flows they needed to the model out of a basket of cash flows, rather than from individual mortgages.*

3. *Perhaps the most radical "what if" scenario would involve creating a smart contract that (a) matched a lender to a model and (b) automatically matched the cash flows to investors directly on the other end via an MBS-like algorithm on the ledger, disintermediating traditional banking and creating a black box in place of the traditional mortgage lending primary market.*

However, as with the trade lifecycles of other securities, some of the challenges will be coordinating the parties involved in this process. These include:

- *The original mortgage issuer. The bank or mortgage lender who originates the mortgage*
- *Mortgage servicer. Collects payments from the mortgage holders*
- *Special purpose vehicles (SPVs). Contain the mortgages, issue loan notes to investors*
- *Note holders. Receive the mortgage cash flows based on some agreed formula*
- *Note custodians. Custodians holding SPV issued loan notes*
- *Collection account bank. Provides nostro accounts for customers to pay into*
- *Customer. Mortgage holder*

One proposed solution is to use distributed ledger technology to manage the data flows from a mortgage servicer to the finance team, and from the finance team to operations and the note custodian. The result would be no more need for reconciliations.

However, there remain some unanswered questions, such as the role a note custodian would play. Hypothetically, if the notes were issued on a distributed ledger then many of the administrative actions currently performed by the custodian could be done on the distributed ledger. Custodians also process payments. Presumably a distributed ledger would need cash-on-ledger to facilitate this – a situation which is a long way off. This also dovetails into the topic of central bank-issued digital currency, a topic of interest to many of our members.

Appendix 2 – Blockchain for Trade Finance (by Vishal Batra, IBM Research, India)

This paper describes an approach and method for leveraging blockchain technology for trade finance to mitigate risk and create an open marketplace for parties to reliably collaborate, mitigate risk and securely transact.

1. Blockchain

Blockchain is a novel peer-to-peer shared ledger technology without any central system or authority, that enables peers/parties to securely and reliably collaborate and transact on a desired set of records/digital assets while enforcing the business terms and conditions agreed upon by them. Each peer/party runs a blockchain node which connects with other blockchain nodes to create a network and each node is identical and maintains the exact same copy of records (on its local ledger) and uniformly applies (executes) the same set of business rules as defined by the parties-thus ensuring/ guaranteeing the same consistent outcome/final state of each transaction on every node of the network. The network, being private and permissioned, allows only authorized parties/peers to join the network and execute transactions with other parties. Further, the data is updated on the ledger only after reaching consensus among the collaborating parties on the network, making the ledger tamper proof and establishing trust and reliability in collaboration. The shared ledger also maintains the change history (previous value) of the records to keep the trail of all updates – a feature that enables auditing of transactions executed on blockchain. Figure 1 illustrates the concept of blockchain technology. Each party is shown to have the same set of immutable records on the ledger without a central system or authority.

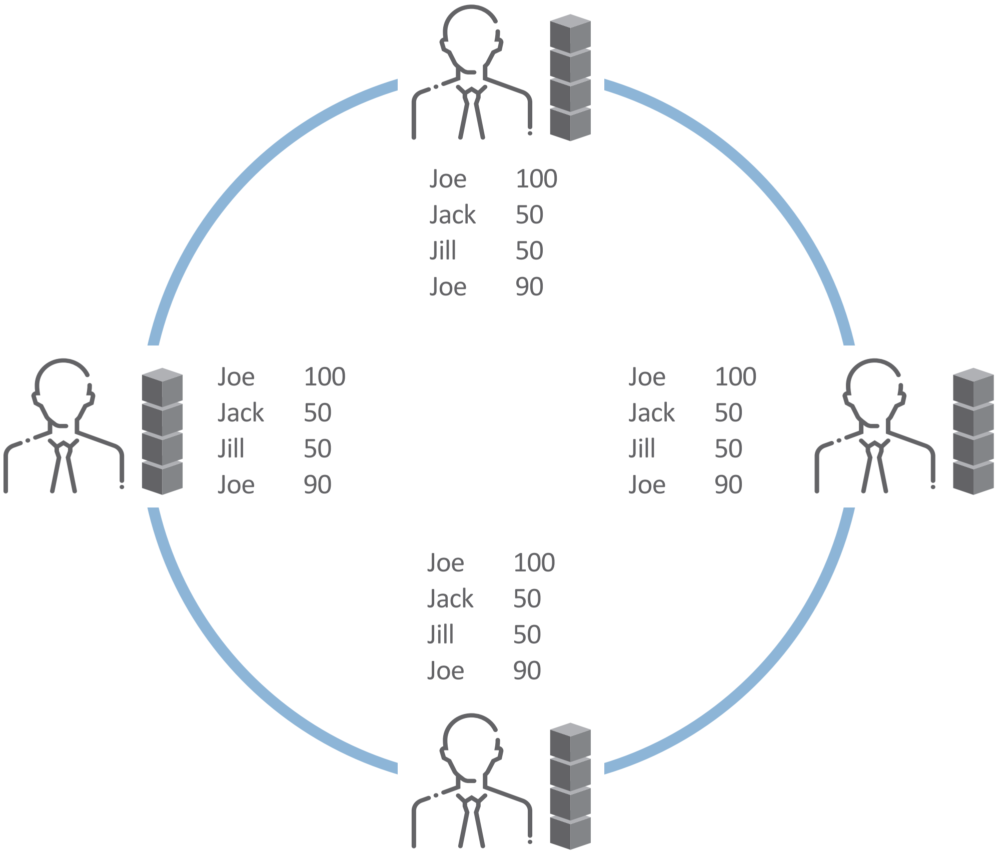

Figure 1. Blockchain Illustration

2. Trade Finance

International trade has tripled as a share of global GDP since 1945, generating an annual revenue of trade financing to about US$ 50 billion. Banks and financial institutions have traditionally played a crucial role in supporting global trade by providing Letter of Credit (LC) – an undertaking/promise given by the bank on behalf of the buyer (importer) to the seller (exporter), that, if the seller presents the complying documents to the bank as specified in the purchase agreement then the bank will make payment to the seller. The banks, thus, mitigate risk and provide the required guarantee to the exporter to reliably transact with a buyer who is often unknown and is in foreign jurisdiction. More recently, with increasing ease of cross-border electronic payments and global financial stability, legal uniformity and certainty governed by international trade treaties and laws for arbitration and settlement of trade disputes, Open Account based trade financing-a method used by trusted business partners by setting up their accounts with banks that are correspondent banks, has become a popular alternative to document-based trade practice using LC as long as global market conditions are stable. However, document-based trade financing using LC is preferred in volatile markets. The banks and financial institutions, and their systems and processes, therefore, have to flexibly and efficiently respond to market dynamics to meet the complex terms and conditions of trade contracts between buyer and seller. Blockchain provides the desired agility and flexibility for a secure, reliable and cost-effective trade finance network.

3. Blockchain for Trade Finance

A blockchain network is established between banks/financial institutions, buyers (importers) and sellers (exporter) and logistics companies, customs, etc. Each entity is an equal peer/party in the network and can easily collaborate with each other by defining the set of records to be shared among them, including LC and Export Documents – Invoice, Bill of Lading, Packing List, etc. and the set of transactions each party can undertake on these documents and its workflow by encoding the transaction rules in Smart Contract computer program which runs on blockchain. For example, the parties can specify that the importer first submits the LC Application on blockchain, which can then be reviewed by its bank and only on approval, the LC is generated from the given application and is presented to the exporter's bank for its review and approval. On approval by exporter's bank, the LC becomes irrevocable and binding among the parties and the exporter can reliably ship the order. The blockchain provides a common view and immutability on records and ensures only permissible transactions are executed thereon by the authorized party, thus eliminating any possibility of inconsistency and discrepancy on records and conflicts thereby.

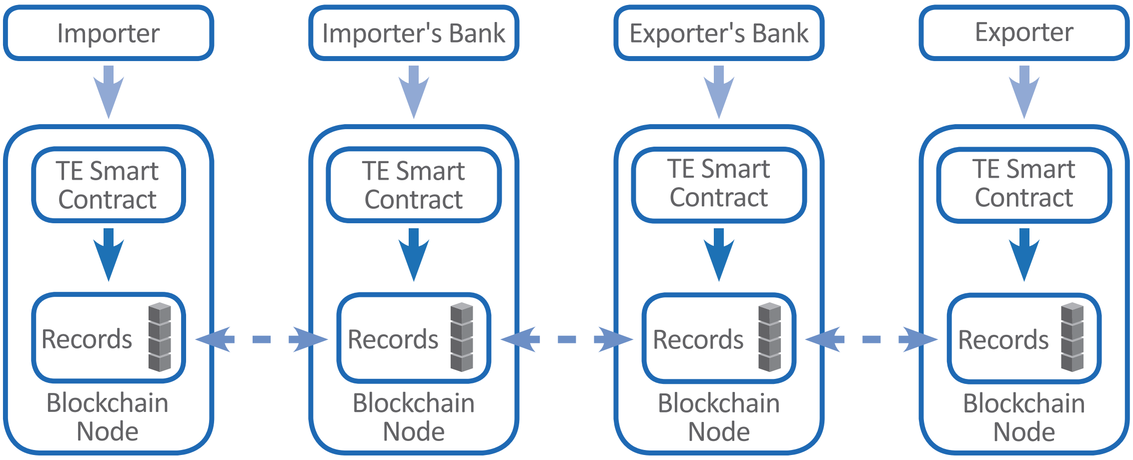

Figure 2. Trade Finance Blockchain Network

4. Trade Finance Marketplace

In the age of Internet and eCommerce, the buyers and sellers meet and negotiate online. In international trade, the buyers and sellers in different countries may have not met each other. The banks, therefore, mitigate the risk between unknown parties. The banks, however, have to setup channels and alliances with other banks in foreign geographies to mitigate risk and enable international trade. It is not cost-efficient for even big banks to have a leading presence in every market. Setting up such channels and alliances globally with other banks operating under different laws and ensuring regulatory compliance in local country is complex and cost intensive. The result is the banks form alliances and partnerships with only few leading/large banks in each geography. The smaller banks and financial institutions, therefore, have to route their trade financing through its correspondent bank which further submits the transaction to the leading bank in the geography to complete the end-to-end transaction. As the number of such intermediate banks increase, the overall transaction cost also increase for the importer and the exporter.

Blockchain technology can be employed to create a global trade finance marketplace and reduce the transaction cost by dynamically discovering the most cost-effective trust path between importer's bank and exporter's bank with least number of intermediaries. The blockchain-based marketplace allows banks and financial institutions of all sizes to join the global trade finance marketplace and offer its services. These entities connect with other entities on the network that they trust, thus defining a trust link between them. For each global trade, the blockchain discovers the most optimal trust path between the importer's bank and exporter's bank, thus eliminating the requirement for banks to form direct alliances and agreements with multiple banks in each geography. Each entity on the network publishes its terms of contract and fee. The blockchain protocol picks those entities in the trust path whose terms and conditions match with the other and also with the terms and conditions specified by the importer and exporter, while ensuring to keep the overall transaction cost to minimum. The entities, providing trust and liquidity for trade finance, act as market makers on global trade finance blockchain network.

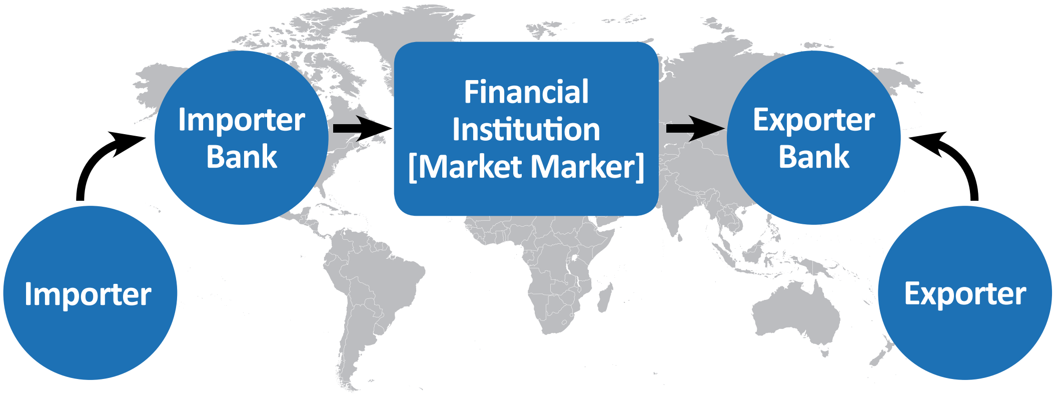

Figure 3. Trade Finance Marketplace on Blockchain Network

5. IBM Blockchain Technology and Solutions

Driven by business requirements and use cases, IBM and partners are implementing an enterprise-grade programmable, secure, permissioned and auditable blockchain fabric, which allows customized and extensible solutions on blockchain. The blockchain fabric – Hyperledger (https://www.hyperledger.org/), is an open source project under the Linux Foundation and is open for all. IBM also offers hosted Hyperledger nodes on secure SoftLayer cloud.

Technical note (I) – the Merkle tree

The Merkle tree is named after Ralph Merkle, who patented the concept of the hash tree. The tree is a binary tree of hashes. It provides an efficient way of summarising a set of leaf nodes. Each non-leaf node is labelled with the hash of its child nodes.

In DLT, a Merkle tree is used to summarise the transactions within a block. Each transaction is represented by a hash value. Transaction hashes are attached to the bottom of the Merkle tree as leaf nodes. The Merkle tree is then built upward from the leaf nodes until it reaches the top of the tree, where there is only one node known as a Merkle root. When using the SHA256 hash, the hashes of all nodes and of the Merkle root are 32 bytes long.

To prove the existence of a transaction within a block, one does not need to know the hashes of all the transactions within that block. To prove its existence within a block of N transactions, all that is required is the log2(N) number of hash values and the performing of an equal number of hash operations: one hash operation at each level, with the step repeated at the next level up until it reaches the root (as illustrated in the diagram below). If the calculated root hash value is the same as the known root value, then the existence of the transaction is proven. This significantly enhances the speed of verifying transactions while at the same time reducing the amount of block information needed to prove the existence of a particular transaction within the block (see diagram below).

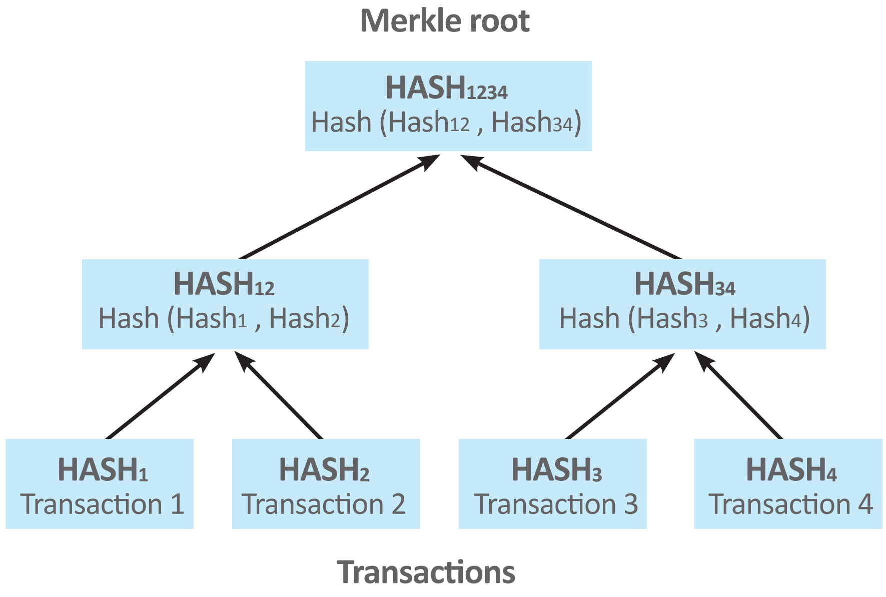

A Merkle tree of four transactions. Leaf nodes at the bottom are hash values of each transaction. Each pair of these hash values are then hashed again. The process is repeated until a single hash is obtained which is called the Merkle Root.
To prove the existence of transaction 4 in the block having a Merkle Root value of $HASH_{1234}$, one only needs to be provided with hash values of $HASH_3$ and $HASH_{12}$.

The Merkle tree is used to represent the transactions within a block. The calculated Merkle root is stored in a block's header.

Ethereum uses a tree structure called the Merkle-Patricia tree, or Patricia tree in short. This structure is used because Ethereum DLT needs not only to store immutable transaction history, but also the state information generated by contract execution, the representation of this state information is supported by the Patricia tree.

Technical note (II) – Illustration: Proof-of-work mining using the SHA256 hash algorithm

Block Data:
m0-rf9falgvoaiujrewmr89u12394=-19324132432490-132m51m4353ti3m-0223;2-0mpir,fbv0-p[da,;baewtr0-weryg;qwytgr2y0-apfd,a;gf0-afasodfasd9foas-o9fas0-fpaqw,reqk3=0aeg=0-as9df-a9=-g09sdhfios9ns9f=0-or4w,513259=0-9326t4-ow0f=-n0sfn-dno9=0-kr,v.a,fd v] fm= o-f[pae0fo 3-v8ap[3evl;dfva9ib[32]Fsbd9akreopbf,324mn1jvs

Target:
Find SHA256 hash of {Data + Nonce } such that leading 12 bits of the hash are all 0's.

Miner-A

Hash calculation result

Nonce	Hash
0:	736fcfa89db20b8f1...
1:	7fdc40dd3a03fe9b9...
2:	bcfbd8deff2d5f616...
3:	64ae0630a91821ea4...
4:	288bd40ab85e9216a...
5:	1984d6d9b141787cd...
6:	9fc9f88659a222716...
.	
122:	00cd22ac7ea221ddc...
.	
4000:	0792f2654d208824a....
.	
.	
.	
5203:	00053e801178649e4....

Miner-B

Hash calculation result

Nonce	Hash
0:	736fcfa89db20b8f1...
1:	7fdc40dd3a03fe9b9...
2:	bcfbd8deff2d5f616...
3:	64ae0630a91821ea4...
4:	288bd40ab85e9216a...
5:	1984d6d9b141787cd...
6:	9fc9f88659a222716...
.	
122:	00cd22ac7ea221ddc...
.	
4000:	0792f2654d208824a....

Proof-of-work

1. Miner-A and Miner-B both start the computation intensive repetitive hash operation to find the nonce that fulfils the target requirement
2. Miner-B has performed 5204 hash calculations and finds the nonce value of 5203, which fulfils the target requirement
3. At this point, Miner-B has only made 4001 calculations, and has not yet found a suitable nonce value.
4. Miner-A wins. He broadcasts his findings to the network and receives the proof-of-work reward.

Verification of proof-of-work

Instead of performing 5204 hash calculations, peers in the network only need to perform one SHA256 hash operation on {Data + "5203"} to obtain the hash value of 00053e801178649e4f....

Technical note (III) – Considerations on the Deployment of Platforms and Applications

Different DLT platforms have different features. Likewise, applications have differing requirements. This section identifies some major DLT features and cites examples of popular DLT platforms that support them. Such information should be taken into consideration when designing and deploying DLT applications. For instance, applications for general public participation will usually be deployed on unpermissioned DLT platforms. Applications for a smaller circle of peers are more suited to permissioned DLT platforms.

1. Participation of untrusted peers

DLT platforms may be deployed in private or public networks. Regardless of where they are deployed, participation in DLT transactions and operations is dictated by the design of the DLT platform. Unpermissioned DLT, also known as public DLT, allows anyone to participate as user or miner. In contrast, permissioned DLT, known as private DLT, requires membership control over participation.

Type	**Non-Trusted Peers Accepted (Unpermissioned DLT)**	**Trusted Peers Only (Permissioned DLT)**
Example	**Bitcoin**	**Ripple**
	Ethereum	**Hyperledger**
		Corda

2. Primary applications

Some DLT platforms are designed to support a wide range of applications, while others are tailored to specific kinds of applications. Both approaches have their advantages. The former provides greater flexibility, while the latter can provide higher efficiency and richer features for specific types of applications.

There are two major types of functionalities: immutable ledger recording and smart contracts. Each functionality may also have further specialisations.

Type	**Ledger Recording**	**Smart Contracts**
Example	**Bitcoin:** *payment*	**Bitcoin:** *Payment contract*
	Ethereum: *payment*	**Ethereum:** *generic applications*
	Ripple: *settlement*	**Hyperledger:** *generic applications*
		Corda: *financial applications*

3. Database structure

DLT started with Bitcoin, which built its distributed database on a chain of blocks containing transaction records. Some newer DLT platforms implement the database with other data structures for reasons such as performance, privacy, and better control.

Type	**Blockchain**	**Non-Blockchain**
Example	**Bitcoin**	**Corda**
	Ethereum	**Ripple**
	Hyperledger	

4. Consensus

Various consensus algorithms have been designed and incorporated into different DLT platforms. These algorithms are mainly of two types: (a) proof-based consensus, and (b) fault-tolerance consensus. Proof-based consensus algorithms are used in unpermissioned DLT networks, in which miners prove their trustworthiness in order to add new blocks to the blockchain. Validating nodes in permissioned DLT do not need to prove their trustworthiness, as they are pre-qualified. Instead, they engage in fault-tolerant consensus to ensure the consistency of all replicated copies of the ledger under their management.

Type	**Proof-based**	**Fault-Tolerant Consensus**
Example	**Bitcoin:** *SHA256 Hash Proof-of-Work*	**Hyperledger:** variety of consensus, including *PBPT (Practical Byzantine fault-tolerant consensus)*[23]
	Ethereum: *Ethash Hash Proof-of-work*	**Corda:** variety of *consensus on the level of individual deals rather than on a global level*[3]
		Ripple: *Byzantine, Altruistic, Rational (BAR) model consensus*[24]

5. Cryptocurrency

Some DLT platforms support native cryptocurrencies and hence can be used directly for payment transactions. Other DLT platforms do not, but instead support ledger recording and smart contract-based applications. Applications may also be developed for such platforms that would enable them to implement new cryptocurrencies.

Type	**Native Cryptocurrency**	**No Native Cryptocurrency**
Example	**Bitcoin:** *bitcoin (BTC)*	**Hyperledger**
	Ethereum: *ether (ETH)*	**Corda**
	Ripple: *Ripple (XRP)*	

6. Transparency

Unpermissioned DLT platforms provide full transparency to peers, who can access the complete ledger as well as the transactions it contains. Permissioned DLT platforms enforce various degrees of transaction access control. Access to transactions may be limited to participants involved in the transaction. Transactions may also be encrypted, with only relevant participants able to access the transaction in plain form. Some Permissioned DLT platforms may also restrict the validation and execution of smart contracts to a selected subset of the validating nodes.

Type	**Accessible to All**	**Limited Access by Permission**
Example	**Bitcoin**	**Hyperledger**
	Ethereum	**Corda**
		Ripple[25]

7. Script language Turing Completeness

Smart contracts in DLT platforms are written in computer languages. There are two types of languages: turing complete and non-turing complete. A turing complete language is highly flexible, supporting the drafting of complicated contract terms. A non-turing complete language places more controls over the contract complexity, and is better suited for specific kinds of contract applications.

Type	**Non-turing complete**	**Turing complete**
Example	**Bitcoin:** Script	**Ethereum:** *Solidity*
		Corda: *Java and others*
		Hyperledger: *Golang and others*

References

[1] See “Bitcoin: A Peer-to-Peer Electronic Cash System” at https://bitcoin.org/bitcoin.pdf

[2] See “Command line options” at https://github.com/ethereum/go-ethereum/wiki/Command-Line-Options

[3] See “Corda: An Introduction” at https://static1.squarespace.com/static/55f73743e4b051cfcc0b02cf/t/57bda319ebbd1acc9c030abd/1472045850269/corda-introductory-whitepaper-final.pdf

[4] See “Hyperledger Whitepaper” at http://www.the-blockchain.com/docs/Hyperledger%20Whitepaper.pdf

[5] See “Enabling Blockchain Innovations with Pegged Sidechains” at https://blockstream.com/sidechains.pdf

[6] See “Bitcoin: A Peer-to-Peer Electronic Cash System” at https://bitcoin.org/bitcoin.pdf

[7] See “Corporate website of Ethereum Foundation” at https://www.ethereum.org/

[8] See “Bitcoin Vs Ethereum: Driven by Different Purposes by Prableen Bajpai, CFA (ICFAI)” at http://www.investopedia.com/articles/investing/031416/bitcoin-vs-ethereum-driven-different-purposes.asp

[9] See “Hyperledger Whitepaper” at http://www.the-blockchain.com/docs/Hyperledger%20Whitepaper.pdf

[10] See “R3” at https://r3cev.com

[11] See “Introducing R3 Corda” at https://r3cev.com/blog/2016/4/4/introducing-r3-corda-a-distributed-ledger-designed-for-financial-services

[12] See “Ripple – Key Feature” at https://ripple.com/technology

[13] See “XRP Portal” at https://ripple.com/xrp-portal

[14] See “About R3” at http://r3cev.com/about/

[15] See “The Ripple Protocol Consensus Algorithm” at https://ripple.com/files/ripple_consensus_whitepaper.pdf

[16] See “What is the Bitcoin Block Size Debate and Why Does it Matter?” at http://www.coindesk.com/what-is-the-bitcoin-block-size-debate-and-why-does-it-matter/

[17] See “Protect your privacy” at https://bitcoin.org/en/protect-your-privacy

[18] See “Protocol Specification” at https://github.com/hyperledger/fabric/blob/master/docs/protocol-spec.md

[19] See “Bitcoin worth $72 million stolen from Bitfinex exchange in Hong Kong” at http://www.reuters.com/article/us-bitfinex-hacked-hongkong-idUSKCN10E0KP

[20] See “Contingency Plans” at https://en.bitcoin.it/wiki/Contingency_plans.

[21] See “Contingency Plan” at https://bitflyer.jp/en/contingency

[22] *Lis pendens* is a written notice that a lawsuit has been filed with respect to a property, concerning the ownership of the title or a claim in it

[23] See “Architecture of the Hyperledger Blockchain Fabric” at https://www.zurich.ibm.com/dccl/papers/cachin_dccl.pdf

[24] See “A Protocol for Interledger Payments” at https://interledger.org/interledger.pdf

[25] See “Key Features” at https://ripple.com/technology/