

# Corda: A distributed ledger

Mike Hearn

November 25, 2016

Version 0.4

*Confidential: For R3 DLG members only*

## Abstract

A decentralised database with minimal trust between nodes would allow for the creation of a global ledger. Such a ledger would have many useful applications in finance, trade, supply chain tracking and more. We present Corda, a decentralised global database, and describe in detail how it achieves the goal of providing a platform for decentralised app development. We elaborate on the high level description provided in the paper *Corda: An introduction*<sup>1</sup> and provide a detailed technical overview, but assume no prior knowledge of the platform.

THIS DOCUMENT DESCRIBES THE CORDA DESIGN AS INTENDED. THE REFERENCE IMPLEMENTATION DOES NOT IMPLEMENT EVERYTHING DESCRIBED WITHIN AT THIS TIME.

# Contents

|          |                                                                |           |
|----------|----------------------------------------------------------------|-----------|
| <b>1</b> | <b>Introduction</b>                                            | <b>4</b>  |
| <b>2</b> | <b>Overview</b>                                                | <b>6</b>  |
| <b>3</b> | <b>The peer to peer network</b>                                | <b>7</b>  |
| 3.1      | Network overview . . . . .                                     | 7         |
| 3.2      | Identity and the permissioning service . . . . .               | 7         |
| 3.3      | The network map . . . . .                                      | 8         |
| 3.4      | Message delivery . . . . .                                     | 8         |
| 3.5      | Serialization, sessioning, deduplication and signing . . . . . | 9         |
| <b>4</b> | <b>Flow framework</b>                                          | <b>10</b> |
| 4.1      | Overview . . . . .                                             | 10        |
| 4.2      | Data visibility and dependency resolution . . . . .            | 12        |
| <b>5</b> | <b>Data model</b>                                              | <b>14</b> |
| 5.1      | Transaction structure . . . . .                                | 14        |
| 5.2      | Composite keys . . . . .                                       | 15        |
| 5.3      | Timestamps . . . . .                                           | 16        |
| 5.4      | Attachments and contract bytecodes . . . . .                   | 17        |
| 5.5      | Hard forks, specifications and dispute resolution . . . . .    | 19        |
| 5.6      | Identity lookups . . . . .                                     | 20        |
| 5.7      | Oracles and tear-offs . . . . .                                | 21        |
| 5.8      | Encumbrances . . . . .                                         | 23        |
| 5.9      | Contract constraints . . . . .                                 | 24        |
| 5.10     | Event scheduling . . . . .                                     | 25        |
| <b>6</b> | <b>Common financial constructs</b>                             | <b>25</b> |
| 6.1      | Assets . . . . .                                               | 25        |
| 6.2      | Obligations . . . . .                                          | 26        |
| 6.3      | Market infrastructure . . . . .                                | 28        |
| <b>7</b> | <b>Notaries and consensus</b>                                  | <b>29</b> |
| 7.1      | Comparison to Nakamoto block chains . . . . .                  | 29        |
| 7.2      | Algorithmic agility . . . . .                                  | 30        |
| 7.3      | Validating and non-validating notaries . . . . .               | 31        |
| 7.4      | Merging networks . . . . .                                     | 32        |
| 7.5      | Guaranteed data distribution . . . . .                         | 33        |
| <b>8</b> | <b>The vault</b>                                               | <b>34</b> |
| 8.1      | Direct SQL access . . . . .                                    | 35        |
| 8.2      | Key randomisation . . . . .                                    | 36        |
| <b>9</b> | <b>Domain specific languages</b>                               | <b>37</b> |
| 9.1      | Clauses . . . . .                                              | 37        |

|           |                                            |           |
|-----------|--------------------------------------------|-----------|
| 9.2       | Combinator libraries . . . . .             | 37        |
| 9.3       | Formally verifiable languages . . . . .    | 39        |
| 9.4       | Projectional editing . . . . .             | 39        |
| <b>10</b> | <b>Secure signing devices</b>              | <b>39</b> |
| 10.1      | Background . . . . .                       | 39        |
| 10.2      | Confusion attacks . . . . .                | 40        |
| 10.3      | Transaction summaries . . . . .            | 41        |
| 10.4      | Identity substitution . . . . .            | 42        |
| 10.5      | Multi-lingual support . . . . .            | 42        |
| <b>11</b> | <b>Client RPC and reactive collections</b> | <b>43</b> |
| <b>12</b> | <b>Data distribution groups</b>            | <b>44</b> |
| <b>13</b> | <b>Deterministic JVM</b>                   | <b>46</b> |
| <b>14</b> | <b>Scalability</b>                         | <b>49</b> |
| <b>15</b> | <b>Privacy</b>                             | <b>51</b> |
| <b>16</b> | <b>Conclusion</b>                          | <b>53</b> |
| <b>17</b> | <b>Acknowledgements</b>                    | <b>54</b> |
|           | <b>Bibliography</b>                        | <b>54</b> |

# 1 Introduction

In many industries significant effort is needed to keep organisation specific databases in sync with each other. In the financial sector the effort of keeping different databases synchronised, reconciling them to ensure they actually are synchronised and resolving the ‘breaks’ that occur when they are not represents a significant fraction of the total work a bank actually does!

Why not just use a shared relational database? This would certainly solve a lot of problems using only existing technology, but it would also raise more questions than answers:

- Who would run this database? Where would we find a sufficient supply of angels to own it?
- In which countries would it be hosted? What would stop that country abusing the mountain of sensitive information it would have?
- What if it were hacked?
- Can you actually scale a relational database to fit the entire financial system?
- What happens if The Financial System™ needs to go down for maintenance?
- What kind of nightmarish IT bureaucracy would guard changes to the database schemas?
- How would you manage access control?

We can imagine many other questions. A decentralised database attempts to answer them.

In this paper we differentiate between a *decentralised* database and a *distributed* database. A distributed database like BigTable<sup>2</sup> scales to large datasets and transaction volumes by spreading the data over many computers. However it is assumed that the computers in question are all run by a single homogenous organisation and that the nodes comprising the database all trust each other not to misbehave or leak data. In a decentralised database, such as the one underpinning Bitcoin<sup>3</sup>, the nodes make much weaker trust assumptions and actively cross-check each other’s work. Such databases trade performance and usability for security and global acceptance.

*Corda* is a decentralised database platform with the following novel features:

- New transaction types can be defined using JVM<sup>4</sup> bytecode.
- Transactions may execute in parallel, on different nodes, without either node being aware of the other’s transactions.

- Nodes are arranged in an authenticated peer to peer network. All communication is direct.
- There is no block chain<sup>3</sup>. Transaction races are deconflicted using pluggable *notaries*. A single Corda network may contain multiple notaries that provide their guarantees using a variety of different algorithms. Thus Corda is not tied to any particular consensus algorithm. (§7)
- Data is shared on a need-to-know basis. Nodes provide the dependency graph of a transaction they are sending to another node on demand, but there is no global broadcast of *all* transactions.
- Bytecode-to-bytecode transpilation is used to allow complex, multi-step transaction building protocols called *flows* to be modelled as blocking code. The code is transformed into an asynchronous state machine, with checkpoints written to the node's backing database when messages are sent and received. A node may potentially have millions of flows active at once and they may last days, across node restarts and even upgrades. Flows expose progress information to node administrators and users and may interact with people as well as other nodes. A Flow library is provided to enable developers to re-use common Flow types such as notarisation, membership broadcast and so on.
- The data model allows for arbitrary object graphs to be stored in the ledger. These graphs are called *states* and are the atomic unit of data.
- Nodes are backed by a relational database and data placed in the ledger can be queried using SQL as well as joined with private tables, thanks to slots in the state definitions that are reserved for join keys.
- The platform provides a rich type system for the representation of things like dates, currencies, legal entities and financial entities such as cash, issuance, deals and so on.
- States can declare a relational mapping and can be queried using SQL.
- Integration with existing systems is considered from the start. The network can support rapid bulk data imports from other database systems without placing load on the network. Events on the ledger are exposed via an embedded JMS compatible message broker.
- States can declare scheduled events. For example a bond state may declare an automatic transition to an "in default" state if it is not repaid in time.

Corda follows a general philosophy of reusing existing proven software systems and infrastructure where possible. Comparisons with Bitcoin and Ethereum will be provided throughout.

## 2 Overview

Corda is a platform for the writing of “CorDapps”: applications that extend the global database with new capabilities. Such apps define new data types, new inter-node protocol flows and the “smart contracts” that determine allowed changes.

What is a smart contract? That depends on the model of computation we are talking about. There are two competing computational models used in decentralised databases: the virtual computer model and the UTXO model. The virtual computer model is used by Ethereum<sup>5</sup>. It models the database as the in-memory state of a global computer with a single thread of execution determined by the block chain. In the UTXO model, as used in Bitcoin, the database is a set of immutable rows keyed by (`hash:output index`). Transactions define outputs that append new rows and inputs which consume existing rows. The term “smart contract” has a different meaning in each model. A deeper discussion of the tradeoffs and terminology in the different approaches can be found in the Corda introductory paper<sup>1</sup>.

We use the UTXO model and as a result our transactions are structurally similar to Bitcoin transactions: they have inputs, outputs and signatures. Unlike Bitcoin, Corda database rows can contain arbitrary data, not just a value field. Because the data consumed and added by transactions is not necessarily a set of key/value pairs, we don’t talk about rows but rather *states*. Like Bitcoin, Corda states are associated with bytecode programs that must accept a transaction for it to be valid, but unlike Bitcoin, a transaction must satisfy the programs for both the input and output states at once. *Issuance transactions* may append new states to the database without consuming any existing states but unlike in Bitcoin these transactions are not special and may be created at any time, by anyone.

In contrast to both Bitcoin and Ethereum, Corda does not order transactions using a block chain and by implication does not use miners or proof-of-work. Instead each state points to a *notary*, which is a service that guarantees it will sign a transaction only if all the input states are un-consumed. A transaction is not allowed to consume states controlled by multiple notaries and thus there is never any need for two-phase commit between notaries. If a combination of states would cross notaries then a special transaction type is used to move them onto a single notary first. See §7 for more information.

The Corda transaction format has various other features which are described in later sections.

## 3 The peer to peer network

### 3.1 Network overview

A Corda network consists of the following components:

- Nodes, communicating using AMQP/1.0 over TLS. Nodes use a relational database for data storage.
- A permissioning service that automates the process of provisioning TLS certificates.
- A network map service that publishes information about nodes on the network.
- One or more notary services. A notary may itself be distributed over multiple nodes.
- Zero or more oracle services. An oracle is a well known service that signs transactions if they state a fact and that fact is considered to be true. They may also optionally also provide the facts. This is how the ledger can be connected to the real world, despite being fully deterministic.

A purely in-memory implementation of the messaging subsystem is provided which can inject simulated latency between nodes and visualise communications between them. This can be useful for debugging, testing and educational purposes.

Oracles and notaries are covered in later sections.

### 3.2 Identity and the permissioning service

Unlike Bitcoin and Ethereum, Corda is designed for semi-private networks in which admission requires obtaining an identity signed by a root authority. This assumption is pervasive - the flow API provides messaging in terms of identities, with routing and delivery to underlying nodes being handled automatically. There is no global broadcast at any point.

This ‘identity’ does not have to be a legal or true identity. In the same way that an email address is a globally unique pseudonym that is ultimately rooted by the top of the DNS hierarchy, so too can a Corda network work with arbitrary self-selected usernames. The permissioning service can implement any policy it likes as long as the identities it signs are globally unique. Thus an entirely anonymous Corda network is possible if a suitable IP obfuscation system like Tor<sup>6</sup> is also used.

Whilst simple string identities are likely sufficient for some networks, the financial industry typically requires some level of *know your customer* checking, and differentiation between different legal entities, branches and desks that may

share the same brand name. Corda reuses the standard PKIX infrastructure for connecting public keys to identities and thus names are actually X.500 names. When a single string is sufficient the *common name* field can be used alone, similar to the web PKI. In more complex deployments the additional structure X.500 provides may be useful to differentiate between entities with the same name. For example there are at least five different companies called *American Bank* and in the past there may have been more than 40 independent banks with that name.

More complex notions of identity that may attest to many time-varying attributes are not handled at this layer of the system: the base identity is always just an X.500 name. Note that even though messaging is always identified, transactions themselves may still contain anonymous public keys.

### 3.3 The network map

Every network requires a network map service, which may itself be composed of multiple cooperating nodes. This is similar to Tor's concept of *directory authorities*. The network map publishes the IP addresses through which every node on the network can be reached, along with the identity certificates of those nodes and the services they provide. On receiving a connection, nodes check that the connecting node is in the network map.

The network map abstracts the underlying IP addresses of the nodes from more useful business concepts like identities and services. Each participant on the network, called a *party*, publishes one or more IP addresses in the network map. Equivalent domain names may be helpful for debugging but are not required. User interfaces and APIs always work in terms of identities - there is thus no equivalent to Bitcoin's notion of an address (hashed public key), and user-facing applications rely on auto-completion and search rather than QRcodes to identify a logical recipient.

It is possible to subscribe to network map changes and registering with the map is the first thing a node does at startup. Nodes may optionally advertise their nearest city for load balancing and network visualisation purposes.

The map is a document that may be cached and distributed throughout the network. The map is therefore not required to be highly available: if the map service becomes unreachable new nodes may not join the network and existing nodes may not change their advertised service set, but otherwise things continue as normal.

### 3.4 Message delivery

The network is structurally similar to the email network. Nodes are expected to be long lived but may depart temporarily due to crashes, connectivity interrup-



tions or maintenance. Messages are written to disk and delivery is retried until the remote node has acknowledged a message, at which point it is expected to have either reliably stored the message or processed it completely. Connections between nodes are built and torn down as needed: there is no assumption of constant connectivity. An ideal network would be entirely flat with high quality connectivity between all nodes, but Corda recognises that this is not always compatible with common network setups and thus the message routing component of a node can be separated from the rest and run outside the firewall. Being outside the firewall or in the firewall's 'de-militarised zone' (DMZ) is required to ensure that nodes can connect to anyone on the network, and be connected to in turn. In this way a node can be split into multiple sub-services that do not have duplex connectivity yet can still take part in the network as first class citizens. Additionally, a single node may have multiple advertised IP addresses.

The reference implementation provides this functionality using the Apache Artemis message broker, through which it obtains journalling, load balancing, flow control, high availability clustering, streaming of messages too large to fit in RAM and many other useful features. The network uses the *AMQP/1.0*<sup>7</sup> protocol which is a widely implemented binary messaging standard, combined with TLS to secure messages in transit and authenticate the endpoints.

### 3.5 Serialization, sessioning, deduplication and signing

All messages are encoded using a compact binary format. Each message has a UUID set in an AMQP header which is used as a deduplication key, thus accidentally redelivered messages will be ignored.

Messages may also have an associated organising 64-bit *session ID*. Note that this is distinct from the AMQP notion of a session. Sessions can be long lived and persist across node restarts and network outages. They exist in order to group messages that are part of a *flow*, described in more detail below.

Messages that are successfully processed by a node generate a signed acknowledgement message called a 'receipt'. Note that this is distinct from the unsigned acknowledgements that live at the AMQP level and which simply flag that a message was successfully downloaded over the wire. A receipt may be generated some time after the message is processed in the case where acknowledgements are being batched to amortise signing overhead, and the receipt identifies the message by the hash of its content. The purpose of the receipts is to give a node undeniable evidence that a counterparty received a notification that would stand up later in a dispute mediation process. Corda does not attempt to support deniable messaging.

## 4 Flow framework

### 4.1 Overview

It is common in decentralised ledger systems for complex multi-party protocols to be needed. The Bitcoin payment channel protocol<sup>8</sup> involves two parties putting money into a multi-signature pot, then iterating with your counterparty a shared transaction that spends that pot, with extra transactions used for the case where one party or the other fails to terminate properly. Such protocols typically involve reliable private message passing, checkpointing to disk, signing of transactions, interaction with the p2p network, reporting progress to the user, maintaining a complex state machine with timeouts and error cases, and possibly interaction with internal systems on either side. All this can become quite involved. The implementation of Bitcoin payment channels in the bitcoinj library is approximately 9000 lines of Java, very little of which involves cryptography.

As another example, the core Bitcoin protocol only allows you to append transactions to the ledger. Transmitting other information that might be useful such as a text message, refund address, identity information and so on is not supported and must be handled in some other way - typically by wrapping the raw ledger transaction bytes in a larger message that adds the desired metadata and giving responsibility for broadcasting the embedded transaction to the recipient, as in Bitcoin's BIP 70<sup>9</sup>.

In Corda transaction data is not globally broadcast. Instead it is transmitted to the relevant parties only when they need to see it. Moreover even quite simple use cases - like sending cash - may involve a multi-step negotiation between counterparties and the involvement of a third party such as a notary. Additional information that isn't put into the ledger is considered essential, as opposed to nice-to-have. Thus unlike traditional block chain systems in which the primary form of communication is global broadcast, in Corda *all* communication takes the form of small multi-party sub-protocols called flows.

The flow framework presents a programming model that looks to the developer as if they have the ability to run millions of long lived threads which can survive node restarts, and even node upgrades. APIs are provided to send and receive object graphs to and from other identities on the network, embed sub-flows, and report progress to observers. In this way business logic can be expressed at a very high level, with the details of making it reliable and efficient abstracted away. This is achieved with the following components.

**Just-in-time state machine compiler.** Code that is written in a blocking manner typically cannot be stopped and transparently restarted later. The first time a flow's `call` method is invoked a bytecode-to-bytecode transformation occurs that rewrites the classes into a form that implements a resumable state

machine. These state machines are sometimes called fibers or coroutines, and the transformation engine Corda uses (Quasar) is capable of rewriting code arbitrarily deep in the stack on the fly. The developer may thus break his or her logic into multiple methods and classes, use loops, and generally structure their program as if it were executing in a single blocking thread. There's only a small list of things they should not do: sleeping, directly accessing the network APIs, or doing other tasks that might block outside of the framework.

**Transparent checkpointing.** When a flow wishes to wait for a message from another party (or input from a human being) the underlying stack frames are suspended onto the heap, then crawled and serialized into the node's underlying relational database using an object serialization framework. The written objects are prefixed with small schema definitions that allow some measure of portability across changes to the layout of objects, although portability across changes to the stack layout is left for future work. Flows are resumed and suspended on demand, meaning it is feasible to have far more flows active at once than would fit in memory. The checkpointing process is atomic with changes to local storage and acknowledgement of network messages.

**Identity to IP address mapping.** Flows are written in terms of identities. The framework takes care of routing messages to the right IP address for a given identity, following movements that may take place whilst the flow is active and handling load balancing for multi-homed parties as appropriate.

**A library of subflows.** Flows can invoke sub-flows, and a library of flows is provided to automate common tasks like notarising a transaction or atomically swapping ownership of two assets.

**Progress reporting.** Flows can provide a progress tracker that indicates which step they are up to. Steps can have human-meaningful labels, along with other tagged data like a progress bar. Progress trackers are hierarchical and steps can have sub-trackers for invoked sub-flows.

**Flow hospital.** Flows can pause if they throw exceptions or explicitly request human assistance. A flow that has stopped appears in the *flow hospital* where the node's administrator may decide to kill the flow or provide it with a solution. The ability to request manual solutions is useful for cases where the other side isn't sure why you are contacting them, for example, the specified reason for sending a payment is not recognised, or when the asset used for a payment is not considered acceptable.

Flows are named using reverse DNS notation and several are defined by the base protocol. Note that the framework is not required to implement the wire protocols, it is just a development aid.

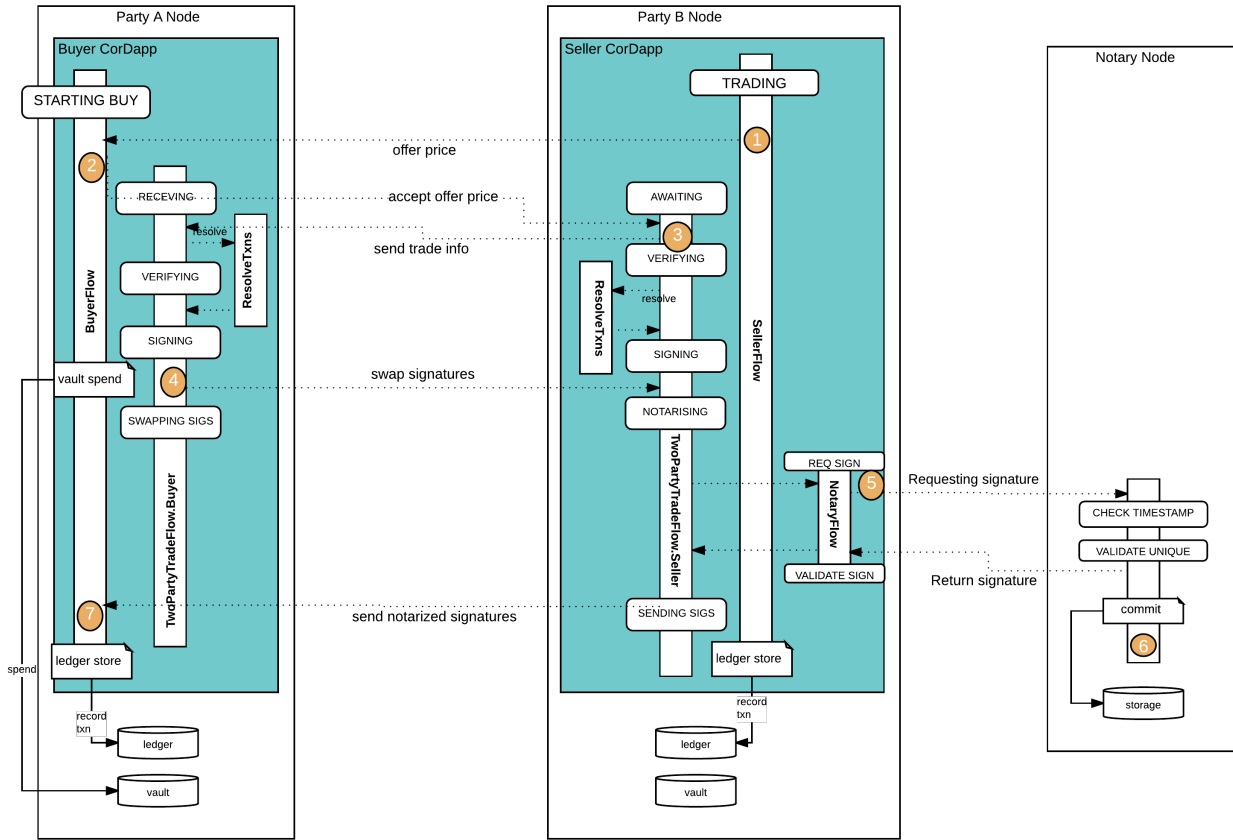


Figure 1: A diagram showing the two party trading flow with notarisation

## 4.2 Data visibility and dependency resolution

When a transaction is presented to a node as part of a flow it may need to be checked. Simply sending you a message saying that I am paying you £1000 is

only useful if you are sure I own the money I’m using to pay you. Checking transaction validity is the responsibility of the **ResolveTransactions** flow. This flow performs a breadth-first search over the transaction graph, downloading any missing transactions into local storage and validating them. The search bottoms out at the issuance transactions. A transaction is not considered valid if any of its transitive dependencies are invalid.

It is required that a node be able to present the entire dependency graph for a transaction it is asking another node to accept. Thus there is never any confusion about where to find transaction data. Because transactions are always communicated inside a flow, and flows embed the resolution flow, the necessary dependencies are fetched and checked automatically from the correct peer. Transactions propagate around the network lazily and there is no need for distributed hash tables.

This approach has several consequences. One is that transactions that move highly liquid assets like cash may end up becoming a part of a very long chain of transactions. The act of resolving the tip of such a graph can involve many round-trips and thus take some time to fully complete. How quickly a Corda network can send payments is thus difficult to characterise: it depends heavily on usage and distance between nodes. Whilst nodes could pre-push transactions in anticipation of them being fetched anyway, such optimisations are left for future work.

Whilst this system is simpler than creating rigid data partitions and clearly provides better privacy than global broadcast, in the absence of additional privacy measures it is nonetheless still difficult to reason about who may get to see transaction data. This uncertainty is mitigated by several factors.

**Small-subgraph transactions.** Some uses of the ledger do not involve widely circulated asset states. For example, two institutions that wish to keep their view of a particular deal synchronised but who are making related payments off-ledger may use transactions that never go outside the involved parties. A discussion of on-ledger vs off-ledger cash can be found in a later section.

**Transaction privacy techniques.** Corda supports a variety of transaction data hiding techniques. For example, public keys can be randomised to make it difficult to link transactions to an identity. “Tear-offs” (§5.7) allow some parts of a transaction to be presented without the others. In future versions of the system secure hardware and/or zero knowledge proofs could be used to convince a party of the validity of a transaction without revealing the underlying data.

**State re-issuance.** In cases where a state represents an asset that is backed by a particular issuer, and the issuer is trusted to behave atomically even when the

ledger isn't forcing atomicity, the state can simply be 'exited' from the ledger and then re-issued. Because there are no links between the exit and reissue transactions this shortens the chain. In practice most issuers of highly liquid assets are already trusted with far more sensitive tasks than reliably issuing pairs of signed data structures, so this approach is unlikely to be an issue.

## 5 Data model

### 5.1 Transaction structure

States are the atomic unit of information in Corda. They are never altered: they are either current ('unspent') or consumed ('spent') and hence no longer valid. Transactions consume zero or more states (inputs) and create zero or more new states (outputs). Because states cannot exist outside of the transactions that created them, any state whether consumed or not can be identified by the identifier of the creating transaction and the index of the state in the outputs list.

Transactions consist of the following components:

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Input references</b> | These are ( <b>hash</b> , <b>output index</b> ) pairs that point to the states a transaction is consuming.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Output states</b>    | Each state specifies the notary for the new state, the contract(s) that define its allowed transition functions and finally the data itself.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Attachments</b>      | Transactions specify an ordered list of zip file hashes. Each zip file may contain code, data, certificates or supporting documentation for the transaction. Contract code has access to the contents of the attachments when checking the transaction for validity.                                                                                                                                                                                                                                                                               |
| <b>Commands</b>         | There may be multiple allowed output states from any given input state. For instance an asset can be moved to a new owner on the ledger, or issued, or exited from the ledger if the asset has been redeemed by the owner and no longer needs to be tracked. A command is essentially a parameter to the contract that specifies more information than is obtainable from examination of the states by themselves (e.g. data from an oracle service). Each command has an associated list of public keys. Like states, commands are object graphs. |
| <b>Signatures</b>       | The set of required signatures is equal to the union of the commands' public keys.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

|           |                                                                                                                                                     |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Type      | Transactions can either be normal or notary-changing. The validation rules for each are different.                                                  |
| Timestamp | When present, a timestamp defines a time range in which the transaction is considered to have occurred. This is discussed in more detail below.     |
| Summaries | Textual summaries of what the transaction does, checked by the involved smart contracts. This field is useful for secure signing devices (see §10). |

Signatures are appended to the end of a transaction and transactions are identified by the hash used for signing, so signature malleability is not a problem. There is never a need to identify a transaction including its accompanying signatures by hash. Signatures can be both checked and generated in parallel, and they are not directly exposed to contract code. Instead contracts check that the set of public keys specified by a command is appropriate, knowing that the transaction will not be valid unless every key listed in every command has a matching signature. Public key structures are themselves opaque. In this way algorithmic agility is retained: new signature algorithms can be deployed without adjusting the code of the smart contracts themselves.

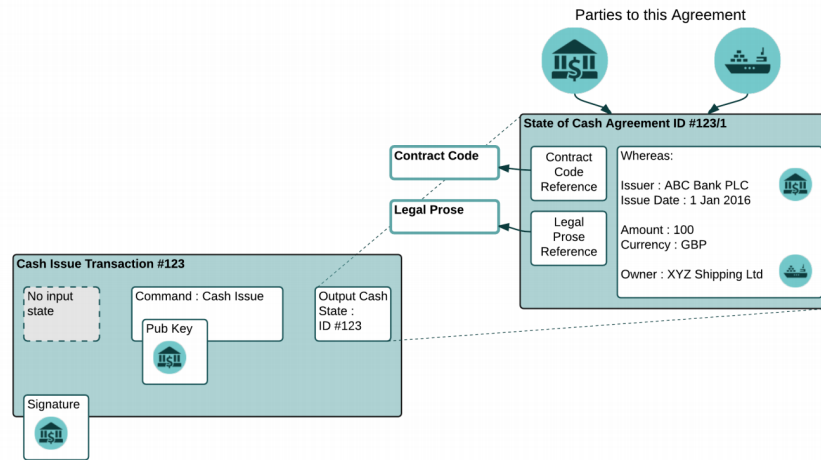
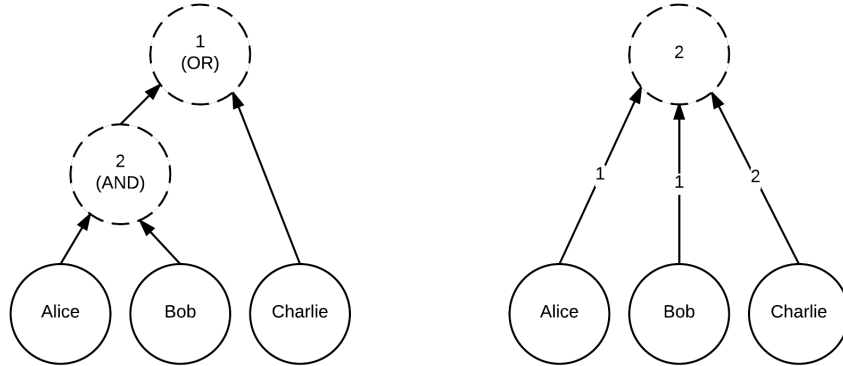


Figure 2: An example of a cash issuance transaction

## 5.2 Composite keys

The term “public key” in the description above actually refers to a *composite key*. Composite keys are trees in which leaves are regular cryptographic public keys with an accompanying algorithm identifiers. Nodes in the tree specify

both the weights of each child and a threshold weight that must be met. The validity of a set of signatures can be determined by walking the tree bottom-up, summing the weights of the keys that have a valid signature and comparing against the threshold. By using weights and thresholds a variety of conditions can be encoded, including boolean formulas with AND and OR.



Composite keys are useful in multiple scenarios. For example, assets can be placed under the control of a 2-of-2 composite key where one leaf key is owned by a user, and the other by an independent risk analysis system. The risk analysis system refuses to sign if the transaction seems suspicious, like if too much value has been transferred in too short a time window. Another example involves encoding corporate structures into the key, allowing a CFO to sign a large transaction alone but his subordinates are required to work together. Composite keys are also useful for notaries. Each participant in a distributed notary is represented by a leaf, and the threshold is set such that some participants can be offline or refusing to sign yet the signature of the group is still valid.

Whilst there are threshold signature schemes in the literature that allow composite keys and signatures to be produced mathematically, we choose the less space efficient explicit form in order to allow a mixture of keys using different algorithms. In this way old algorithms can be phased out and new algorithms phased in without requiring all participants in a group to upgrade simultaneously.

### 5.3 Timestamps

Transaction timestamps specify a `[start, end]` time window within which the transaction is asserted to have occurred. Timestamps are expressed as windows because in a distributed system there is no true time, only a large number of desynchronised clocks. This is not only implied by the laws of physics but also



by the nature of shared transactions - especially if the signing of a transaction requires multiple human authorisations, the process of constructing a joint transaction could take hours or even days.

It is important to note that the purpose of a transaction timestamp is to communicate the transaction's position on the timeline to the smart contract code for the enforcement of contractual logic. Whilst such timestamps may also be used for other purposes, such as regulatory reporting or ordering of events in a user interface, there is no requirement to use them like that and locally observed timestamps may sometimes be preferable even if they will not exactly match the time observed by other parties. Alternatively if a precise point on the timeline is required and it must also be agreed by multiple parties, the midpoint of the time window may be used by convention. Even though this may not precisely align to any particular action (like a keystroke or verbal agreement) it is often useful nonetheless.

Timestamp windows may be open ended in order to communicate that the transaction occurred before a certain time or after a certain time, but how much before or after is unimportant. This can be used in a similar way to Bitcoin's `nLockTime` transaction field, which specifies a *happens-after* constraint.

Timestamps are checked and enforced by notary services. As the participants in a notary service will themselves not have precisely aligned clocks, whether a transaction is considered valid or not at the moment it is submitted to a notary may be unpredictable if submission occurs right on a boundary of the given window. However, from the perspective of all other observers the notary's signature is decisive: if the signature is present, the transaction is assumed to have occurred within that time.

**Reference clocks.** In order to allow for relatively tight time windows to be used when transactions are fully under the control of a single party, notaries are expected to be synchronised to the atomic clocks at the US Naval Observatory. Accurate feeds of this clock can be obtained from GPS satellites. Note that Corda uses the Java timeline<sup>10</sup> which is UTC with leap seconds spread over the last 1000 seconds of the day, thus each day always has exactly 86400 seconds. Care should be taken to ensure that changes in the GPS leap second counter are correctly smeared in order to stay synchronised with Java time. When setting a transaction time window care must be taken to account for network propagation delays between the user and the notary service, and messaging within the notary service.

## 5.4 Attachments and contract bytecodes

Transactions may have a number of *attachments*, identified by the hash of the file. Attachments are stored and transmitted separately to transaction data and

are fetched by the standard resolution flow only when the attachment has not previously been seen before.

Attachments are always zip files<sup>11</sup> and cannot be referred to individually by contract code. The files within the zips are collapsed together into a single logical file system, with overlapping files being resolved in favour of the first mentioned. Not coincidentally, this is the mechanism used by Java classpaths.

Smart contracts in Corda are defined using JVM bytecode as specified in “*The Java Virtual Machine Specification SE 8 Edition*”<sup>4</sup>, with some small differences that are described in a later section. A contract is simply a class that implements the **Contract** interface, which in turn exposes a single function called **verify**. The verify function is passed a transaction and either throws an exception if the transaction is considered to be invalid, or returns with no result if the transaction is valid. The set of verify functions to use is the union of the contracts specified by each state (which may be expressed as constraints, see §5.9). Embedding the JVM specification in the Corda specification enables developers to write code in a variety of languages, use well developed toolchains, and to reuse code already authored in Java or other JVM compatible languages.

The Java standards also specify a comprehensive type system for expressing common business data. Time and calendar handling is provided by an implementation of the JSR 310 specification, decimal calculations can be performed either using portable (**strictfp**) floating point arithmetic or the provided bignum library, and so on. These libraries have been carefully engineered by the business Java community over a period of many years and it makes sense to build on this investment.

Contract bytecode also defines the states themselves, which may be arbitrary object graphs. Because JVM classes are not a convenient form to work with from non-JVM platforms the allowed types are restricted and a standardised binary encoding scheme is provided. States may label their properties with a small set of standardised annotations. These can be useful for controlling how states are serialised to JSON and XML (using JSR 367 and JSR 222 respectively), for expressing static validation constraints (JSR 349) and for controlling how states are inserted into relational databases (JSR 338). This feature is discussed later.

Attachments may also contain data files that support the contract code. These may be in the same zip as the bytecode files, or in a different zip that must be provided for the transaction to be valid. Examples of such data files might include currency definitions, timezone data and public holiday calendars. Any public data may be referenced in this way. Attachments are intended for data on the ledger that many parties may wish to reuse over and over again. Data files are accessed by contract code using the same APIs as any file on the classpath would be accessed. The platform imposes some restrictions on what kinds of data can be included in attachments along with size limits, to avoid people placing inappropriate files on the global ledger (videos, PowerPoints etc).

Note that the creator of a transaction gets to choose which files are attached. Therefore, it is typical that states place constraints on the data they're willing to accept. Attachments *provide* data but do not *authenticate* it, so if there's a risk of someone providing bad data to gain an economic advantage there must be a constraints mechanism to prevent that from happening. This is rooted at the contract constraints encoded in the states themselves: a state can not only name a class that implements the `Contract` interface but also place constraints on the zip/jar file that provides it. That constraint can in turn be used to ensure that the contract checks the authenticity of the data - either by checking the hash of the data directly, or by requiring the data to be signed by some trusted third party.

## 5.5 Hard forks, specifications and dispute resolution

Decentralised ledger systems often differ in their underlying political ideology as well as their technical choices. The Ethereum project originally promised “unstoppable apps” which would implement “code as law”. After a prominent smart contract was hacked<sup>12</sup>, an argument took place over whether what had occurred could be described as a hack at all given the lack of any non-code specification of what the program was meant to do. The disagreement eventually led to a split in the community.

As Corda contracts are simply zip files, it is easy to include a PDF or other documents describing what a contract is meant to actually do. There is no requirement to use this mechanism, and there is no requirement that these documents have any legal weight. However in financial use cases it's expected that they would be legal contracts that take precedence over the software implementations in case of disagreement.

It is technically possible to write a contract that cannot be upgraded. If such a contract governed an asset that existed only on the ledger, like a cryptocurrency, then that would provide an approximation of “code as law”. We leave discussion of the wisdom of this concept to political scientists and reddit.

**Platform logging** There is no direct equivalent in Corda of a block chain “hard fork”, so the only solution to discarding buggy or fraudulent transaction chains would be to mutually agree out of band to discard an entire transaction subgraph. As there is no global visibility either this mutual agreement would not need to encompass all network participants: only those who may have received and processed such transactions. The flip side of lacking global visibility is that there is no single point that records who exactly has seen which transactions. Determining the set of entities that'd have to agree to discard a subgraph means correlating node activity logs. Corda nodes log sufficient information to ensure this correlation can take place. The platform defines a flow to assist with this, which can be used by anyone. A tool is provided that generates an “investigation

request” and sends it to a seed node. The flow signals to the node administrator that a decision is required, and sufficient information is transmitted to the node to try and convince the administrator to take part (e.g. a signed court order). If the administrator accepts the request through the node explorer interface, the next hops in the transaction chain are returned. In this way the tool can semi-automatically crawl the network to find all parties that would be affected by a proposed rollback. The platform does not take a position on what types of transaction rollback are justified and provides only minimal support for implementing rollbacks beyond locating the parties that would have to agree.

Once involved parties are identified there are at least two strategies for editing the ledger. One is to extend the transaction chain with new transactions that simply correct the database to match the intended reality. For this to be possible the smart contract must have been written to allow arbitrary changes outside its normal business logic when a sufficient threshold of signatures is present. This strategy is simple and makes the most sense when the number of parties involved in a state is small and parties have no incentive to leave bad information in the ledger. For asset states that are the result of theft or fraud the only party involved in a state may resist attempts to patch things up in this way, as they may be able to benefit in the real world from the time lag between the ledger becoming inaccurate and it catching up with reality. In this case a more complex approach can be used in which the involved parties minus the uncooperative party agree to mark the relevant states as no longer consumed/spent. This is essentially a limited form of database rollback.

## 5.6 Identity lookups

In all block chain inspired systems there exists a tension between wanting to know who you are dealing with and not wanting others to know. A standard technique is to use randomised public keys in the shared data, and keep the knowledge of the identity that key maps to private. For instance, it is considered good practice to generate a fresh key for every received payment. This technique exploits the fact that verifying the integrity of the ledger does not require knowing exactly who took part in the transactions, only that they followed the agreed upon rules of the system.

Platforms such as Bitcoin and Ethereum have relatively ad-hoc mechanisms for linking identities and keys. Typically it is the user’s responsibility to manually label public keys in their wallet software using knowledge gleaned from websites, shop signs and so on. Because these mechanisms are ad hoc and tedious many users don’t bother, which can make it hard to figure out where money went later. It also complicates the deployment of secure signing devices and risk analysis engines. Bitcoin has BIP 70<sup>9</sup> which specifies a way of signing a “payment request” using X.509 certificates linked to the web PKI, giving a cryptographically secured and standardised way of knowing who you are dealing

with. Identities in this system are the same as used in the web PKI: a domain name, email address or EV (extended validation) organisation name.

Corda takes this concept further. States may define fields of type **Party**, which encapsulates an identity and a public key. When a state is deserialised from a transaction in its raw form, the identity field of the **Party** object is null and only the public (composite) key is present. If a transaction is deserialised in conjunction with X.509 certificate chains linking the transient public keys to long term identity keys the identity field is set. In this way a single data representation can be used for both the anonymised case, such as when validating dependencies of a transaction, and the identified case, such as when trading directly with a counterparty. Trading flows incorporate sub-flows to transmit certificates for the keys used, which are then stored in the local database. However the transaction resolution flow does not transmit such data, keeping the transactions in the chain of custody pseudonymous.

**Deterministic key derivation** Corda allows for but does not mandate the use of deterministic key derivation schemes such as BIP 32<sup>13</sup>. The infrastructure does not assume any mathematical relationship between public keys because some cryptographic schemes are not compatible with such systems. Thus we take the efficiency hit of always linking transient public keys to longer term keys with X.509 certificates.

## 5.7 Oracles and tear-offs

It is sometimes convenient to reveal a small part of a transaction to a counterparty in a way that allows them to check the signatures and sign it themselves. A typical use case for this is an *oracle*, defined as a network service that is trusted to sign transactions containing statements about the world outside the ledger only if the statements are true.

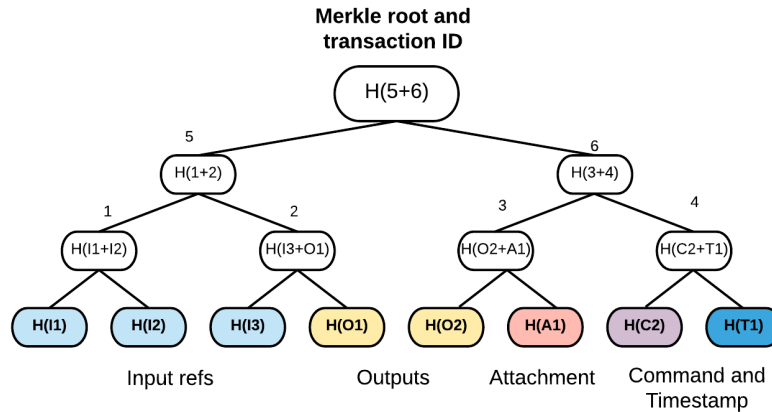
Here are some example statements an oracle might check:

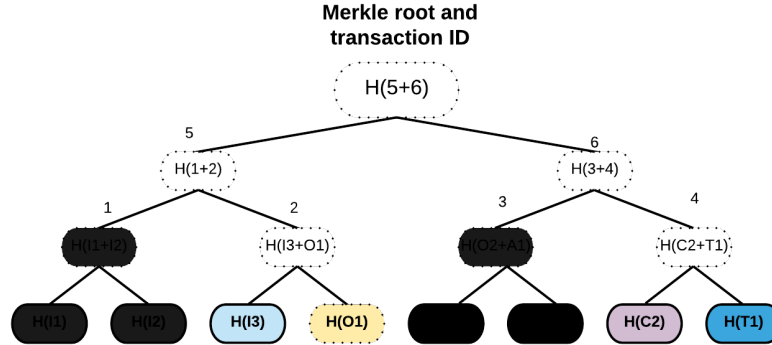
- The price of a stock at a particular moment was X.
- An agreed upon interest rate at a particular moment was Y.
- If a specific organisation has declared bankruptcy.
- Weather conditions in a particular place at a particular time.

It is worth asking why a smart contract cannot simply fetch this information from some internet server itself: why do we insist on this notion of an oracle. The reason is that all calculations on the ledger must be deterministic. Everyone must be able to check the validity of a transaction and arrive at exactly the same answer, at any time (including years into the future), on any kind of computer. If a smart contract could do things like read the system clock or fetch arbitrary web pages then it would be possible for some computers to conclude a transaction

was valid, whilst others concluded it was not (e.g. if the remote server had gone offline). Solving this problem means all the data needed to check the transaction must be in the ledger, which in turn implies that we must accept the point of view of some specific observer. That way there can be no disagreement about what happened.

One way to implement oracles would be to have them sign a small data structure which is then embedded somewhere in a transaction (in a state or command). We take a different approach in which oracles sign the entire transaction, and data the oracle doesn't need to see is "torn off" before the transaction is sent. This is done by structuring the transaction as a Merkle hash tree so that the hash used for the signing operation is the root. By presenting a counterparty with the data elements that are needed along with the Merkle branches linking them to the root hash, that counterparty can sign the entire transaction whilst only being able to see some of it. Additionally, if the counterparty needs to be convinced that some third party has already signed the transaction, that is also straightforward. Typically an oracle will be presented with the Merkle branches for the command or state that contains the data, and the timestamp field, and nothing else. The resulting signature contains flag bits indicating which parts of the structure were presented for signing to avoid a single signature covering more than expected.





There are a couple of reasons to take this more indirect approach. One is to keep a single signature checking code path. By ensuring there is only one place in a transaction where signatures may be found, algorithmic agility and parallel/batch verification are easy to implement. When a signature may be found in any arbitrary location in a transaction's data structures, and where verification may be controlled by the contract code itself (as in Bitcoin), it becomes harder to maximise signature checking efficiency. As signature checks are often one of the slowest parts of a block chain system, it is desirable to preserve these capabilities.

Another reason is to provide oracles with a business model. If oracles just signed statements and nothing else then it would be difficult to run an oracle in which there are only a small number of potential statements, but determining their truth is very expensive. People could share the signed statements and reuse them in many different transactions, meaning the cost of issuing the initial signatures would have to be very high, perhaps unworkably high. Because oracles sign specific transactions, not specific statements, an oracle that is charging for its services can amortise the cost of determining the truth of a statement over many users who cannot then share the signature itself (because it covers a one-time-use structure by definition).

## 5.8 Encumbrances

Each state in a transaction specifies a contract (boolean function) that is invoked with the entire transaction as input. All contracts must accept in order for the transaction to be considered valid. Sometimes we would like to compose the behaviours of multiple different contracts. Consider the notion of a "time lock" - a restriction on a state that prevents it being modified (i.e. sold) until a certain time. This is a general piece of logic that could apply to many kinds of assets. Whilst such logic could be implemented in a library and then called from every

contract that might want to benefit from it, that requires all contract authors to think ahead and include the functionality. It would be better if we could mandate that the time lock logic ran along side the contract that governs the locked state.

Consider an asset that is supposed to remain frozen until a time is reached. Encumbrances allow a state to specify another state that must be present in any transaction that consumes it. For example, a time lock contract can define a state that contains the time at which the lock expires, and a simple contract that just compares that time against the transaction timestamp. The asset state can be included in a spend-to-self transaction that doesn't change the ownership of the asset but does include a time lock state in the outputs. Now if the asset state is used, the time lock state must also be used, and that triggers the execution of the time lock contract.

Encumbered states can only point to one encumbrance state, but that state can itself point to another and so on, resulting in a chain of encumbrances all of which must be satisfied.

An encumbrance state must be present in the same transaction as the encumbered state, as states refer to each other by index alone.

## 5.9 Contract constraints

The easiest way of tying states to the contract code that defines them is by hash. This works for very simple and stable programs, but more complicated contracts may need to be upgraded. In this case it may be preferable for states to refer to contracts by the identity of the signer. Because contracts are stored in zip files, and because a Java Archive (JAR) file is just a zip with some extra files inside, it is possible to use the standard JAR signing infrastructure to identify the source of contract code. Simple constraints such as “any contract of this name signed by these keys” allow for some upgrade flexibility, at the cost of increased exposure to rogue contract developers. Requiring combinations of signatures helps reduce the risk of a rogue or hacked developer publishing a bad contract version, at the cost of increased difficulty in releasing new versions. State creators may also specify third parties they wish to review contract code. Regardless of which set of tradeoffs is chosen, the framework can accomodate them.

A contract constraint may use a composite key of the type described in §5.2. The standard JAR signing protocol allows for multiple signatures from different private keys, thus being able to satisfy composite keys. The allowed signing algorithms are `SHA256withRSA` and `SHA256withECDSA`. Note that the cryptographic algorithms used for code signing may not always be the same as those used for transaction signing, as for code signing we place initial focus on being able to re-use the infrastructure.



## 5.10 Event scheduling

State classes may request flows to be started at given times. When a state is considered relevant by the vault and the implementing CorDapp is installed and whitelisted by the administrator (e.g. in the config file), the node may react to the passage of time by starting new interactions with other nodes, people, or internal systems. As financial contracts often have a notion of time in them this feature can be useful for many kinds of state transitions, for example, expiry of an option contract, management of a default event, re-fixing of an interest rate swap and so on.

To request scheduled events, a state may implement the `SchedulableState` interface and then return a request from the `nextScheduledActivity` function. The state will be queried when it is committed to the vault and the scheduler will ensure the relevant flow is started at the right time.

## 6 Common financial constructs

### 6.1 Assets

A ledger that cannot record the ownership of assets is not very useful. We define a set of classes that model asset-like behaviour and provide some platform contracts to ensure interoperable notions of cash and obligations.

We define the notion of an `OwnableState`, implemented as an interface which any state may conform to. Ownable states are required to have an `owner` field which is a composite key (see §5.2). This is utilised by generic code in the vault (see §8) to manipulate ownable states.

From `OwnableState` we derive a `FungibleAsset` concept to represent assets of measurable quantity, in which units are sufficiently similar to be represented together in a single ledger state. Making that concrete, pound notes are a fungible asset: regardless of whether you represent £10 as a single £10 note or two notes of £5 each the total value is the same. Other kinds of fungible asset could be barrels of Brent Oil (but not all kinds of crude oil worldwide, because oil comes in different grades which are not interchangeable), litres of clean water, kilograms of bananas, units of a stock and so on.

When cash is represented on a digital ledger an additional complication can arise: for national “fiat” currencies the ledger merely records an entity that has a liability which may be redeemed for some other form (physical currency, a wire transfer via some other ledger system, etc). This means that two ledger entries of £1000 may *not* be entirely fungible because all the entries really represent is a claim on an issuer, which - if it is not a central bank - may go bankrupt. Even assuming defaults never happen, the data representing where an asset may be

redeemed must be tracked through the chain of custody, so ‘exiting’ the asset from the ledger and thus claiming physical ownership can be done.

The Corda type system supports the encoding of this complexity. The **Amount<T>** type defines an integer quantity of some token. This type does not support fractional quantities so when used to represent national currencies the quantity must be measured in pennies, with sub-penny amount requiring the use of some other type. The token can be represented by any type. A common token type to use is **Issued<T>**, which defines a token issued by some party. It encapsulates what the asset is, who issued it, and an opaque reference field that is not parsed by the platform - it is intended to help the issuer keep track of e.g. an account number, the location where the asset can be found in storage, etc.

## 6.2 Obligations

It is common in finance to be paid with an IOU rather than hard cash (note that in this section ‘hard cash’ means a balance with the central bank). This is frequently done to minimise the amount of cash on hand when trading institutions have some degree of trust in each other: if you make a payment to a counterparty that you know will soon be making a payment back to you as part of some other deal, then there is an incentive to simply note the fact that you owe the other institution and then ‘net out’ these obligations at a later time, either bilaterally or multilaterally. Netting is a process by which a set of gross obligations is replaced by an economically-equivalent set where eligible offsetting obligations have been elided. The process is conceptually similar to trade compression, whereby a set of trades between two or more parties are replaced with an economically similar, but simpler, set. The final output is the amount of money that needs to actually be transferred.

Corda models a nettable obligation with the **Obligation** contract, which is a subclass of **FungibleAsset**. Obligations have a lifecycle and can express constraints on the on-ledger assets used for settlement. The contract allows not only for trading and fungibility of obligations but also bi-lateral and multi-lateral netting.

It is important to note here that netting calculations can get very complex and the financial industry contains firms that compete on the quality of their netting algorithms. The **Obligation** contract provides methods to calculate simple bi-lateral nettings, and verify the correctness of both bi and multi-lateral nettings. For very large, complex multi-lateral nettings it is expected that institutions would use pre-existing netting implementations.

Netting is usually done when markets are closed. This is because it is hard to calculate nettings and settle up concurrently with the trading positions changing. The problem can be seen as analogous to garbage collection in a managed runtime: compacting the heap requires the running program to be stopped so the contents of the heap can be rewritten. If a group of trading institutions

wish to implement a checked form of ‘market close’ then they can use an encumbrance (see §5.8) to prevent an obligation being changed during certain hours, as determined by the clocks of the notaries (see §5.3).

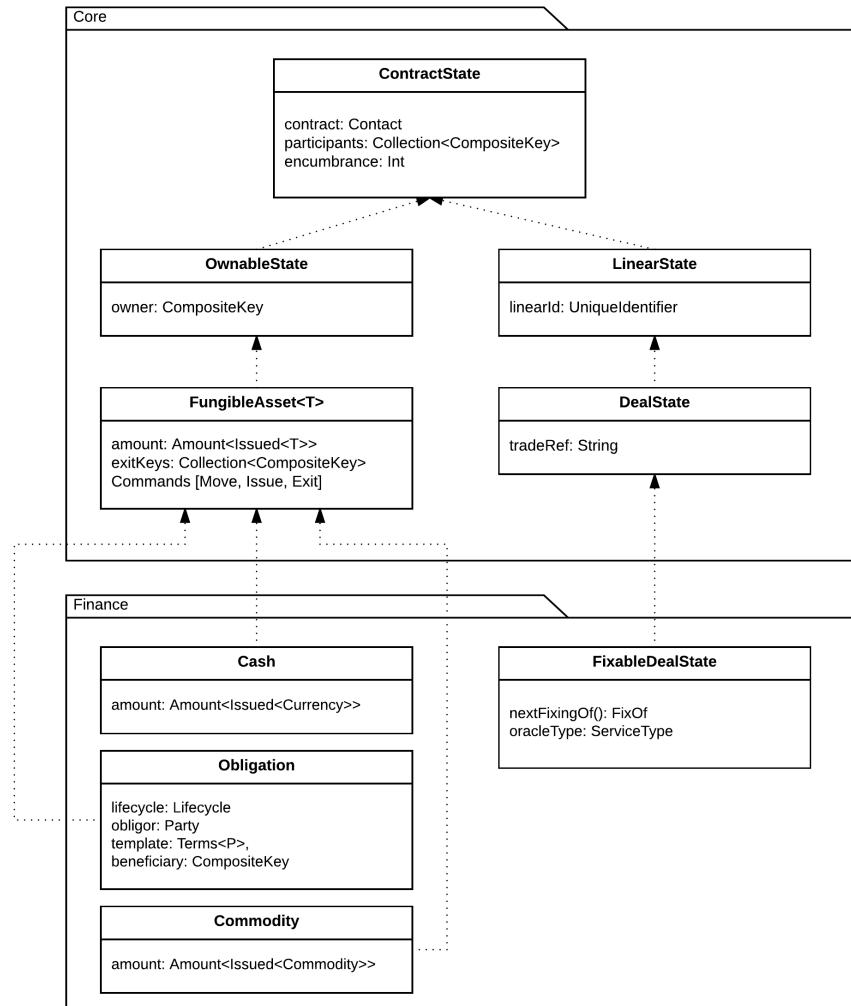


Figure 3: Class hierarchy diagram showing the relationships between different state types

### 6.3 Market infrastructure

Trade is the lifeblood of the economy. A distributed ledger needs to provide a vibrant platform on which trading may take place. However, the decentralised nature of such a network makes it difficult to build competitive market infrastructure on top of it, especially for highly liquid assets like securities. Markets typically provide features like a low latency order book, integrated regulatory compliance, price feeds and other things that benefit from a central meeting point.

The Corda data model allows for integration of the ledger with existing markets and exchanges. A sell order for an asset that exists on-ledger can have a *partially signed transaction* attached to it. A partial signature is a signature that allows the signed data to be changed in controlled ways after signing. Partial signatures are directly equivalent to Bitcoin's **SIGHASH** flags and work in the same way - signatures contain metadata describing which parts of the transaction are covered. Normally all of a transaction would be covered, but using this metadata it is possible to create a signature that only covers some inputs and outputs, whilst allowing more to be added later.

This feature is intended for integration of the ledger with the order books of markets and exchanges. Consider a stock exchange. A buy order can be submitted along with a partially signed transaction that signs a cash input state and an output state representing some quantity of the stock owned by the buyer. By itself this transaction is invalid, as the cash does not appear in the outputs list and there is no input for the stock. A sell order can be combined with a mirror-image partially signed transaction that has a stock state as the input and a cash state as the output. When the two orders cross on the order book, the exchange itself can take the two partially signed transactions and merge them together, creating a valid transaction that it then notarises and distributes to both buyer and seller. In this way trading and settlement become atomic, with the ownership of assets on the ledger being synchronised with the view of market participants. Note that in this design the distributed ledger itself is *not* a marketplace, and does not handle distribution or matching of orders. Rather, it focuses on management of the pre- and post- trade lifecycles.

**Central counterparties.** In many markets, central infrastructures such as clearing houses (also known as Central Counterparties, or CCPs) and Central Securities Depositories (CSD) have been created. They provide governance, rules definition and enforcement, risk management and shared data and processing services. The partial data visibility, flexible transaction verification logic and pluggable notary design means Corda could be a particularly good fit for future distributed ledger services contemplated by CCPs and CSDs.

## 7 Notaries and consensus

Corda does not organise time into blocks. This is sometimes considered strange, given that it can be described as a block chain system or ‘block chain inspired’. Instead a Corda network has one or more notary services which provide transaction ordering and timestamping services, thus abstracting the role miners play in other systems into a pluggable component.

Notaries are expected to be composed of multiple mutually distrusting parties who use a standard consensus algorithm. Notaries are identified by and sign with composite public keys (§5.2) that conceptually follow the Interledger Crypto-Conditions specification<sup>14</sup>. Note that whilst it would be conventional to use a BFT algorithm for a notary service, there is no requirement to do so and in cases where the legal system is sufficient to ensure protocol compliance a higher performance algorithm like Raft<sup>15</sup> may be used. Because multiple notaries can co-exist a single network may provide a single global BFT notary for general use and region-specific Raft notaries for lower latency trading within a unified regulatory area, for example London or New York.

Notaries accept transactions submitted to them for processing and either return a signature over the transaction, or a rejection error that states that a double spend has occurred. The presence of a notary signature from the state’s chosen notary indicates transaction finality. An app developer triggers notarisation by invoking the `Finality` flow on the transaction once all other necessary signatures have been gathered. Once the finality flow returns successfully, the transaction can be considered committed to the database.

### 7.1 Comparison to Nakamoto block chains

Bitcoin organises the timeline into a chain of blocks, with each block pointing to a previous block the miner has chosen to build upon. Blocks also contain a rough timestamp. Miners can choose to try and extend the block chain from any previous block, but are incentivised to build on the most recently announced block by the fact that other nodes in the system only recognise a block if it’s a part of the chain with the most accumulated proof-of-work. As each block contains a reward of newly issued bitcoins, an unrecognised block represents a loss and a recognised block typically represents a profit.

Bitcoin uses proof-of-work because it has a design goal of allowing an unlimited number of identityless parties to join and leave the network at will, whilst simultaneously making it hard to execute Sybil attacks (attacks in which one party creates multiple identities to gain undue influence over the network). This is an appropriate design to use for a peer to peer network formed of volunteers who can’t/won’t commit to any long term relationships up front, and in which identity verification is not done. Using proof-of-work then leads naturally to a requirement to quantise the timeline into chunks, due to the probabilistic

nature of searching for a proof. The chunks must then be ordered relative to each other. The incentive to build on the most recently announced proof of work is in tension with the reality that it takes time for a proof to circulate around the network. This means it is desirable that proofs are produced at a rate that is slow enough that very few are circulating at the same time. Given that transactions are likely to be produced at a higher rate than this, it implies a need for the proofs to consolidate multiple transactions. Hence the need for blocks.

A Corda network is email-like in the sense that nodes have long term stable identities, of which they can prove ownership of to others. Sybil attacks are blocked by the network entry process. This allows us to discard proof-of-work along with its multiple unfortunate downsides:

- Energy consumption is excessively high for such a simple task, being comparable at the time of writing to the electricity consumption of an entire city<sup>16</sup>. At a time when humanity needs to use less energy rather than more this is ecologically undesirable.
- High energy consumption forces concentration of mining power in regions with cheap or free electricity. This results in unpredictable geopolitical complexities that many users would rather do without.
- Identityless participants mean all transactions must be broadcast to all network nodes, as there's no reliable way to know who the miners are. This worsens privacy.
- The algorithm does not provide finality, only a probabilistic approximation, which is a poor fit for existing business and legal assumptions.<sup>17</sup>
- It is theoretically possible for large numbers of miners or even all miners to drop out simultaneously without any protocol commitments being violated.

Once proof-of-work is disposed of there is no longer any need to quantise the timeline into blocks because there is no longer any need to slow the publication of conflict resolution proposals, and because the parties asserting the correctness of the ordering are known ahead of time regular signatures are sufficient.

## 7.2 Algorithmic agility

Consensus algorithms are a hot area of research and new algorithms are frequently developed that improve upon the state of the art. Unlike most distributed ledger systems Corda does not tightly integrate one specific approach. This is not only to support upgrades as new algorithms are developed, but also to reflect the fact that different tradeoffs may make sense for different situations and networks.

As a simple example, a notary that uses Raft between nodes that are all within the same city will provide extremely good performance and latency, at the cost of being more exposed to malicious attacks or errors by whichever node has been elected leader. In situations where the members making up a distributed notary service are all large, regulated institutions that are not expected to try and corrupt the ledger in their own favour trading off security to gain performance may make sense. In other situations where existing legal or trust relationships are less robust, slower but byzantine fault tolerant algorithms like BFT-SMaRT<sup>18</sup> may be preferable. Alternatively, hardware security features like Intel SGX® may be used to convert non-BFT algorithms into a more trusted form using remote attestation and hardware protection.

Being able to support multiple notaries in the same network has other advantages:

- It is possible to phase out notaries (i.e. sets of participants) that no longer wish to provide that service by migrating states.
- The scalability of the system can be increased by bringing online new notaries that run in parallel. As long as access to the ledger has some locality (i.e. states aren't constantly being migrated between notaries) this allows for the scalability limits of common consensus algorithms or node hardware to be worked around.
- In some but not all cases, regulatory constraints on data propagation can be respected by having jurisdictionally specific notaries. This would not work well when two jurisdictions have mutually incompatible constraints or for assets that may frequently travel around the world, but it can work when using the ledger to track the state of deals or other facts that are inherently region specific.
- Notaries can compete on their availability and performance.
- Users can pick between *validating* and *non-validating* notaries. See below.
- In some models for how these technologies will be adopted, it is possible that issuers of assets will find it convenient to 'self-notarise' transactions that pertain to assets they have issued and this necessarily requires support for multiple notaries in the same network. Such a model is likely to be a transitional state, not least because such a model is inherently limited in the range of operations that can be supported.
- Separate networks can start independent and be merged together later (see below).

### 7.3 Validating and non-validating notaries

Validating notaries resolve and fully check transactions they are asked to de-conflict. Thus in the degenerate case of a network with just a single notary

and without the use of any privacy features, they gain full visibility into every transaction. Non-validating notaries assume transaction validity and do not request transaction data or their dependencies beyond the list of states consumed. With such a notary it is possible for the ledger to become ‘wedged’, as anyone who knows the hash and index of a state may consume it without checks. If the cause of the problem is accidental, the incorrect data can be presented to a non-validating notary to convince it to roll back the commit, but if the error is malicious then states controlled by such a notary may become permanently corrupted.

It is therefore possible for users to select their preferred point on a privacy/security spectrum for each state individually depending on how they expect the data to be used. When the states are unlikely to live long or propagate far and the only entities who will learn their transaction hashes are somewhat trustworthy, the user may select to keep the data from the notary. For liquid assets a validating notary should always be used to prevent value destruction and theft if the transaction identifiers leak.

## 7.4 Merging networks

Because there is no single block chain it becomes possible to merge two independent networks together by simply establishing two-way connectivity between their nodes then configuring each side to trust each other’s notaries and certificate authorities.

This ability may seem pointless: isn’t the goal of a decentralised ledger to have a single global database for everyone? It is, but a practical route to reaching this end state is still required. It is often the case that organisations perceived by consumers as being a single company are in fact many different entities cross-licensing branding, striking deals with each other and doing internal trades with each other. This sort of setup can occur for regulatory reasons, tax reasons, due to a history of mergers or just through a sheer masochistic love of paperwork. Very large companies can therefore experience all the same synchronisation problems a decentralised ledger is intended to fix but purely within the bounds of that organisation. In this situation the main problem to tackle is not malicious actors but rather heterogenous IT departments, varying software development practices, unlinked user directories and so on. Such organisations can benefit from gaining experience with the technology internally and cleaning up their own internal views of the world before tackling the larger problem of synchronising with the wider world as well.

When merging networks, both sides must trust that each other’s notaries have never signed double spends. When merging an organisation-private network into the global ledger it should be possible to simply rely on incentives to provide this guarantee: there is no point in a company double spending against itself. However, if more evidence is desired, a standalone notary could be run



against a hardware security module with audit logging enabled. The notary itself would simply use a private database and run on a single machine, with the logs exported to the people running a global network for asynchronous post-hoc verification.

## 7.5 Guaranteed data distribution

In any global consensus system the user is faced with the question of whether they have the latest state of the database. Programmers working with block chains often make the simplifying assumption that because there is no formal map of miner locations and thus transactions are distributed to miners via broadcast, that they can listen to the stream of broadcasts and learn if they have the latest data. Alas, nothing stops someone privately providing a miner who has a known location with a transaction that they agree not to broadcast. The first time the rest of the network finds out about this transaction is when a block containing it is broadcast. When used to do double spending fraud this type of attack is known as a Finney Attack<sup>19</sup>. Proof-of-work based systems rely on aligned incentives to discourage such attacks: to quote the Bitcoin white paper, “He ought to find it more profitable to play by the rules ... than to undermine the system and the validity of his own wealth.” In practice this approach appears to work well enough most of the time, given that miners typically do not accept privately submitted transactions.

In a system without global broadcast things are very different: the notary clusters *must* accept transactions directly and there is no mechanism to ensure that everyone sees that the transaction is occurring. Sometimes this doesn’t matter: most transactions are irrelevant for you and having to download them just wastes resources. But occasionally you do wish to become aware that the ledger state has been changed by someone else. A simple example is an option contract in which you wish to expire the option unless the counterparty has already exercised it. Then exercising the option must not require the seller to sign off on it, as it may be advantageous for the seller to refuse if it would cause them to lose money. Whilst the seller would discover if the buyer had exercised the option when they attempted to expire it, due to the notary informing them that their expiry transaction was a double spend, it is preferable to find out immediately.

The obvious way to implement this is to give notaries the responsibility for ensuring all interested parties find out about a transaction. However, this would require the notaries to know who the involved parties actually are, which would create an undesirable privacy leak. It would also place extra network load on the notaries who would frequently be sending transaction data to parties that may already have it, or may simply not care. In many cases there may be no requirement for the notary to act as a trusted third party for data distribution purposes, as game-theoretic assumptions or legal assurances are sufficiently

strong that peers can be trusted to deliver transaction data as part of their regular flows.

To solve this, app developers can choose whether to request transaction distribution by the notary or not. This works by simply piggybacking on the standard identity lookup flows (see §5.6). If a node wishes to be informed by the notary when a state is consumed, it can send the certificates linking the random keys in the state to the notary cluster, which then stores it in the local databases as per usual. Once the notary cluster has committed the transaction, key identities are looked up and any which resolve successfully are sent copies of the transaction. In normal operation the notary is not provided with the certificates linking the random keys to the long term identity keys and thus does not know who is involved with the operation (assuming source IP address obfuscation is in use, see §15).

## 8 The vault

In any block chain based system most nodes have a wallet, or as we call it, a vault.

The vault contains data extracted from the ledger that is considered *relevant* to the node's owner, stored in a form that can be easily queried and worked with. It also contains private key material that is needed to sign transactions consuming states in the vault. Like with a cryptocurrency wallet, the Corda vault understands how to create transactions that send value to someone else by combining asset states and possibly adding a change output that makes the values balance. This process is usually referred to as 'coin selection'. Coin selection can be a complex process. In Corda there are no per transaction network fees, which is a significant source of complexity in other systems. However transactions must respect the fungibility rules in order to ensure that the issuer and reference data is preserved as the assets pass from hand to hand.

Advanced vault implementations may also perform splitting and merging of states in the background. The purpose of this is to increase the amount of transaction creation parallelism supported. Because signing a transaction may involve human intervention (see §10) and thus may take a significant amount of time, it can become important to be able to create multiple transactions in parallel. The vault must manage state 'soft locks' to prevent multiple transactions trying to use the same output simultaneously. Violation of a soft lock would result in a double spend being created and rejected by the notary. If a vault were to contain the entire cash balance of a user in just one state, there could only be a single transaction being constructed at once and this could impose unacceptable operational overheads on an organisation. By automatically creating send-to-self transactions that split the big state into multiple smaller states, the number of transactions that can be created in parallel is increased. Alter-

natively many tiny states may need to be consolidated into a smaller number of more valuable states in order to avoid hitting transaction size limits. Finally, in some cases the vault may send asset states back to the issuer for re-issuance, thus pruning long transaction chains and improving privacy.

The vault is also responsible for managing scheduled events requested by node-relevant states when the implementing app has been installed (see §5.10).

## 8.1 Direct SQL access

A distributed ledger is ultimately just a shared database, albeit one with some unique features. The following features are therefore highly desirable for improving the productivity of app developers:

- Ability to store private data linked to the semi-public data in the ledger.
- Ability to query the ledger data using widely understood tools like SQL.
- Ability to perform joins between entirely app-private data (like customer notes) and ledger data.
- Ability to define relational constraints and triggers on the underlying tables.
- Ability to do queries at particular points in time e.g. midnight last night.
- Re-use of industry standard and highly optimised database engines.
- Independence from any particular database engine, without walling off too many useful features.

Conda states are defined using a subset of the JVM bytecode language which includes annotations. The vault recognises annotations from the *Java Persistence Architecture* (JPA) specification defined in JSR 338<sup>20</sup>. These annotations define how a class maps to a relational table schema including which member is the primary key, what SQL types to map the fields to and so on. When a transaction is submitted to the vault by a flow, the vault finds states it considers relevant (i.e. which contains a key owned by the node) and the relevant Cordapp has been installed into the node as a plugin, the states are fed through an object relational mapper which generates SQL UPDATE and INSERT statements. Note that data is not deleted when states are consumed, however a join can be performed with a dedicated metadata table to eliminate consumed states from the dataset. This allows data to be queried at a point in time, with rows being evicted to historical tables using external tools.

Nodes come with an embedded database engine out of the box, but may also be configured to point to a separate RDBMS. The node stores not only state data but also all node working data in the database, including flow checkpoints. Thus the state of a node and all communications it is engaged in can be backed up by simply backing up the database itself. The JPA annotations are independent

of any particular database engine or SQL dialect and thus states cannot use any proprietary column types or other features, however, because the ORM is only used on the write paths users are free to connect to the backing database directly and issue SQL queries that utilise any features of their chosen database engine that they like. They can also create their own tables and create merged views of the underlying data for end user applications, as long as they don't impose any constraints that would prevent the node from syncing the database with the actual contents of the ledger.

States are arbitrary object graphs. Whilst nothing stops a state from containing multiple classes intended for different tables, it is typical that the relational representation will not be a direct translation of the object-graph representation. States are queried by the vault for the ORM mapped class to use, which will often skip ledger-specific data that's irrelevant to the user like opaque public keys and may expand single fields like an `Amount<Issued<Currency>>` type into multiple database columns.

It's worth noting here that although the vault only responds to JPA annotations it is often useful for states to be annotated in other ways, for instance to customise its mapping to XML/JSON, or to impose validation constraints<sup>21</sup>. These annotations won't affect the behaviour of the node directly but may be useful when working with states in surrounding software.

## 8.2 Key randomisation

A standard privacy technique in block chain systems is the use of randomised unlinkable public keys to stand in for actual verified identities. Ownership of these pseudonyms may be revealed to a counterparty using a simple interactive protocol in which Alice selects a random nonce ('number used once') and sends it to Bob, who then signs the nonce with the private key corresponding to the public key he is proving ownership of.

Generating fresh keys for each new deal or asset transfer rapidly results in many private keys being created. These keys must all be backed up and kept safe, which poses a significant management problem when done at scale. The canonical way to resolve this problem is through the use of deterministic key derivation, as pioneered by the Bitcoin community in BIP 32 'Hierarchical Deterministic Wallets'<sup>13</sup>. Deterministic key derivation allows all private key material needed to be derived from a single, small pool of entropy (e.g. a carefully protected and backed up 128 bits of random data). More importantly, when the full BIP 32 technique is used in combination with an elliptic curve that supports it, public keys may also be deterministically derived *without* access to the underlying private key material. This allows devices to provide fresh public keys to counterparties without being able to sign with those keys, enabling better security along with operational efficiencies.

Corda does not place any constraints on the mathematical properties of the

digital signature algorithms parties use. However, implementations are recommended to use hierarchical deterministic key derivation when possible.

## 9 Domain specific languages

### 9.1 Clauses

When writing a smart contract, many desired features and patterns crop up repeatedly. For example it is expected that all production quality asset contracts would want the following features:

- Issuance and exit transactions.
- Movement transactions (reassignment of ownership).
- Fungibility management (see §6).
- Support for upgrading to new versions of the contract.

Many of these seemingly simple features have obscure edge cases. One example is a need to prevent the creation of asset states that contain zero or negative quantities of the asset. Another is to ensure that states are summed for fungibility purposes without accidentally assuming that the transaction only moves one type of asset at once. Rather than expect contract developers to reimplement these pieces of low level logic the Corda standard library provides *clauses*, a small class library that implement reusable pieces of contract logic. A contract writer may create their own clauses and then pass the set of contract clauses together to a library function that interprets them.

### 9.2 Combinator libraries

Domain specific languages for the expression of financial contracts are a popular area of research. A seminal work is ‘Composing contracts’ by Peyton-Jones, Seward and Eber [PJSE2000<sup>22</sup>] in which financial contracts are modelled with a small library of Haskell combinators. These models can then be used for valuation of the underlying deals. Block chain systems use the term ‘contract’ in a slightly different sense to how PJSE do but the underlying concepts can be adapted to our context as well. The platform provides an experimental *universal contract* that builds on the language extension features of the Kotlin programming language. To avoid linguistic confusion it refers to the combined code/data bundle as an ‘arrangement’ rather than a contract. A European FX option expressed in this language looks like this:

---

```
fun fx_swap(expiry: String, notional: BigDecimal,
            strike: BigDecimal, foreignCurrency: Currency,
```

```

    domesticCurrency: Currency,
    partyA: Party, partyB: Party) =

    arrange {
      actions {
        (partyA or partyB) may {
          "execute".givenThat(after(expiry)) {
            swap(
              partyA, notional * strike,
              domesticCurrency, partyB,
              notional, foreignCurrency
            )
          }
        }
      }
    }

    arrange {
      actions {
        acmeCorp may {
          "exercise".givenThat(before("2017-09-01")) {
            fx_swap("2017-09-01", 1.M, 1.2.bd, EUR, USD,
              acmeCorp, highStreetBank)
          }
        }

        (acmeCorp or highStreetBank) may {
          "expire".anytime {
            zero
          }
        }
      }
    }
  }

```

---

The programmer may define arbitrary ‘actions’ along with constraints on when the actions may be invoked. The **zero** token indicates the termination of the deal.

As can be seen, this DSL combines both *what* is allowed and deal-specific data like *when* and *how much* is allowed, therefore blurring the distinction the core model has between code and data. It builds on prior work to enable not only valuation/cash flow calculations, but also direct enforcement of the contract’s logic at the database level as well.

### 9.3 Formally verifiable languages

Corda contracts can be upgraded. However, given the coordination problems inherent in convincing many participants in a large network to accept a new version of a contract, a frequently cited desire is for formally verifiable languages to be used to try and guarantee the correctness of the implementations.

We do not attempt to tackle this problem ourselves. However, because Corda focuses on deterministic execution of any JVM bytecode, formally verifiable languages that target this instruction set are usable for the expression of smart contracts. A good example of this is the Whyle language by Dr David Pearce<sup>23</sup>, which checks program-integrated proofs at compile time. By building on industry-standard platforms, we gain access to cutting edge research from the computer science community outside of the distributed systems world.

### 9.4 Projectional editing

Custom languages and type systems for the expression of contract logic can be naturally combined with *projectional editing*, in which source code is not edited textually but rather by a structure aware editor<sup>24</sup>. Such languages can consist not only of traditional grammar-driven text oriented structures but also diagrams, tables and recursive compositions of them together. Given the frequent occurrence of data tables and English-oriented nature of many financial contracts, a dedicated environment for the construction of smart contract logic may be appreciated by the users.

## 10 Secure signing devices

### 10.1 Background

A common feature of digital financial systems and block chain-type systems in particular is the use of secure client-side hardware to hold private keys and perform signing operations with them. Combined with a zero tolerance approach to transaction rollbacks, this is one of the ways they reduce overheads: by attempting to ensure that transaction authorisation is robust and secure, and thus that signatures are reliable.

Many banks have rolled out CAP (chip authentication program) readers to consumers which allow logins to online banking using a challenge/response protocol to a smartcard. The user is expected to type in the right codes and copy the responses back to the computer by hand. These devices are cheap, but tend to have small, unreliable, low resolution screens and can be subject to confusion attacks if there is malware on the PC, e.g. if the malware convinces the user they are performing a login challenge whereas in fact they are authorising a

payment to a new account. The primary advantage is that the signing key is held in a robust and cheap smart card, so the device can be replaced without replacing the key.

The state-of-the-art in this space are devices like the TREZOR<sup>25</sup> by Satoshi Labs or the Ledger Blue. These were developed by and for the Bitcoin community. They are more expensive than CAP readers and feature better screens and USB connections to eliminate typing. Advanced devices like the Ledger Blue support NFC and Bluetooth as well. These devices differ from CAP readers in another key respect: instead of signing arbitrary, small challenge numbers, they actually understand the native transaction format of the network to which they're specialised and parse the transaction to figure out the message to present to the user, who then confirms that they wish to perform the action printed on the screen by simply pressing a button. The transaction is then signed internally before being passed back to the PC via the USB/NFC/Bluetooth connection.

This setup means that rather than having a small device that authorises to a powerful server (which controls all your assets), the device itself controls the assets. As there is no smartcard equivalent the private key can be exported off the device by writing it down in the form of "wallet words": 12 random words derived from the contents of the key. Because elliptic curve private keys are small (256 bits), this is not as tedious as it would be with the much larger RSA keys the financial industry is typically using.

There are clear benefits to having signing keys be kept on personal, employee-controlled devices only, with the organisation's node not having any ability to sign for transactions itself:

- If the node is hacked by a malicious intruder or bad insider they cannot steal assets, modify agreements, or do anything else that requires human approval, because they don't have access to the signing keys. There is no single point of failure from a key management perspective.
- It's more clear who signed off on a particular action - the signatures prove which devices were used to sign off on an action. There can't be any back doors or administrator tools which can create transactions on behalf of someone else.
- Devices that integrate fingerprint readers and other biometric authentication could further increase trust by making it harder for employees to share/swap devices. A smartphone or tablet could be also used as a transaction authenticator.

## 10.2 Confusion attacks

The biggest problem facing anyone wanting to integrate smart signing devices into a distributed ledger system is how the device processes transactions. For



Bitcoin it's straightforward for devices to process transactions directly because their format is very small and simple (in theory - in practice a fixable quirk of the Bitcoin protocol actually significantly complicates how these devices must work). Thus turning a Bitcoin transaction into a human meaningful confirmation screen is quite easy:

Confirm payment of 1.23 BTC to 1AbCd0123456.....

This confirmation message is susceptible to confusion attacks because the opaque payment address is unpredictable. A sufficiently smart virus/attacker could have swapped out a legitimate address of a legitimate counterparty you are expecting to pay with one of their own, thus you'd pay the right amount to the wrong place. The same problem can affect financial authenticators that verify IBANs and other account numbers: the user's source of the IBAN may be an email or website they are viewing through the compromised machine. The BIP 70<sup>9</sup> protocol was designed to address this attack by allowing a certificate chain to be presented that linked a target key with a stable, human meaningful and verified identity.

For a generic ledger we are faced with the additional problem that transactions may be of many different types, including new types created after the device was manufactured. Thus creating a succinct confirmation message inside the device would become an ever-changing problem requiring frequent firmware updates. As firmware upgrades are a potential weak point in any secure hardware scheme, it would be ideal to minimise their number.

### 10.3 Transaction summaries

To solve this problem we add a top level summaries field to the transaction format (joining inputs, outputs, commands, attachments etc). This new top level field is a list of strings. Smart contracts get a new responsibility. They are expected to generate an English message describing what the transaction is doing, and then check that it is present in the transaction. The platform ensures no unexpected messages are present. The field is a list of strings rather than a single string because a transaction may do multiple things simultaneously in advanced use cases.

Because the calculation of the confirmation message has now been moved to the smart contract itself, and is a part of the transaction, the transaction can be sent to the signing device: all it needs to do is extract the messages and print them to the screen with YES/NO buttons available to decide whether to sign or not. Because the device's signature covers the messages, and the messages are checked by the contract based on the machine readable data in the states, we can know that the message was correct and legitimate.

The design above is simple but has the issue that large amounts of data are sent to the device which it doesn't need. As it's common for signing devices to have

constrained memory, it would be unfortunate if the complexity of a transaction ended up being limited by the RAM available in the users' signing devices. To solve this we can use the tear-offs mechanism (see §5.7) to present only the summaries and the Merkle branch connecting them to the root. The device can then sign the entire transaction contents having seen only the textual summaries, knowing that the states will trigger the contracts which will trigger the summary checks, thus the signature covers the machine-understandable version of the transaction as well.

Note, we assume here that contracts are not themselves malicious. Whilst a malicious user could construct a contract that generated misleading messages, for a user to see states in their vault and work with them requires the accompanying CorDapp to be loaded into the node as a plugin and thus whitelisted. There is never a case where the user may be asked to sign a transaction involving contracts they have not previously approved, even though the node may execute such contracts as part of verifying transaction dependencies.

## 10.4 Identity substitution

Contract code only works with opaque representations of public keys. Because transactions in a chain of custody may need to be anonymised, it isn't possible for a contract to access identity information from inside the sandbox. Therefore it cannot generate a complete message that includes human meaningful identity names even if the node itself does have this information.

To solve this the transaction is provided to the device along with the X.509 certificate chains linking the pseudonymous public keys to the long term identity certificates, which for transactions involving the user should always be available (as they by definition know who their trading counterparties are). The device can verify those certificate chains to build up a mapping of index to human readable name. The messages placed inside a transaction may contain numeric indexes of the public keys required by the commands using backslash syntax, and the device must perform the message substitution before rendering. Care must be taken to ensure that the X.500 names issued to network participants do not contain text chosen to deliberately confuse users, e.g. names that contain quote marks, partial instructions, special symbols and so on. This can be enforced at the network permissioning level.

## 10.5 Multi-lingual support

The contract is expected to generate a human readable version of the transaction. This should be in English, by convention. In theory, we could define the transaction format to support messages in different languages, and if the contract supported that the right language could then be picked by the signing device. However, care must be taken to ensure that the message the user sees

in alternative languages is correctly translated and not subject to ambiguity or confusion, as otherwise exploitable confusion attacks may arise.

## 11 Client RPC and reactive collections

Any realistic deployment of a distributed ledger faces the issue of integration with an existing ecosystem of surrounding tools and processes. Ideally, programs that interact with the node will be loosely coupled, authenticated, robust against transient node outages and restarts, and speed differences (e.g. production of work being faster than completion of work) will be handled transparently.

To meet these needs, Corda nodes expose a simple RPC mechanism that has a couple of unusual features. The underlying transport is message queues (AMQP) and methods can return object graphs that contain Rx observables<sup>26</sup> which may in turn emit more observables.

It is a common pattern for RPCs to return a snapshot of some data structure, along with an observable that emits objects representing a delta on that data structure. The client library has functionality to reconstruct the snapshot + diffs into a observable collections of the type that can be bound directly to a JavaFX user interface. In this way, rendering data structures in the global ledger in a rich client app that stays fresh becomes a straightforward operation that requires minimal work from the developer: simply wiring the pieces together in a functional way is sufficient. Reactive transforms over these observable collections such as mappings, filterings, sortings and so on make it easy to build user interfaces in a functional programming style.

Because RPC transport takes place via the node's message queue broker, the framework automatically recovers from restarts of the node/node components, IP addresses changes on the client and similar interruptions to communication. Likewise, programs that need to live for a long time and survive restarts, upgrades and moves can request that observations be sent to a persistent queue. Backpressure and queue management is supplied by the broker. Additional capacity for processing RPCs can be added by attaching more RPC processors to the broker which load balances between them automatically.

It can be asked why Corda does not use the typical REST+JSON approach to communicating with the node. The reasons are:

- A preference for binary protocols over textual protocols, as text based protocols tend to be more susceptible to escaping and other buffer management problems that can lead to security issues.
- Message queue brokers provide significant amounts of infrastructure for building reliable apps which plain HTTP does not such as backpressure management, load balancing, queue browsing, management of speed differences and so on.

- REST based protocols have multiple conventions for streaming of results back to the client, none of which are ideal for the task.

Being able to connect live data structures directly to UI toolkits also contributes to the avoidance of XSS exploits, XSRF exploits and similar security problems based on losing track of buffer boundaries.

## 12 Data distribution groups

By default, distribution of transaction data is defined by app-provided flows (see §4). Flows specify when and to which peers transactions should be sent. Typically these destinations will be calculated based on the content of the states and the available identity lookup certificates, as the intended use case of financial data usually contains the identities of the relevant parties within it. Sometimes though, the set of parties that should receive data isn't known ahead of time and may change after a transaction has been created. For these cases partial data visibility is not a good fit and an alternative mechanism is needed.

A data distribution group (DDG) is created by generating a keypair and a self-signed certificate for it. Groups are identified internally by their public key and may be given string names in the certificate, but nothing in the software assumes the name is unique: it's intended only for human consumption and it may conflict with other independent groups. In case of conflict user interfaces disambiguate by appending a few characters of the base58 encoded public key to the name like so: "My popular group name (a4T)". As groups are not globally visible anyway, it is unlikely that conflicts will be common or require many code letters to deconflict, and some groups may not even be intended for human consumption at all.

Once a group is created other nodes can be invited to join it by using an invitation flow. Membership can be either read only or read/write. To add a node as read-only, the certificate i.e. pubkey alone is sent. To add a node as read/write the certificate and private key are sent. A future elaboration on the design may support giving each member a separate private key which would allow tracing who added transactions to a group, but this is left for future work. In either case the node records in its local database which other nodes it has invited to the group once they accept the invitation.

When the invite is received the target node runs the other side of the flow as normal, which may either automatically accept membership if it's configured to trust the inviting node, or send a message to a message queue for processing by an external system, or kick it up to a human administrator for approval. Invites to groups the node is already a member of are rejected. The accepting node also records which node invited it. So, there ends up being a two-way recorded relationship between inviter and invitee stored in their vaults. Finally the inviter side of the invitation flow pushes a list of all the transaction IDs that

exist in the group and the invitee side resolves all of them. The end result is that all the transactions that are in the group are sent to the new node (along with all dependencies).

Note that this initial download is potentially infinite if transactions are added to the group as fast or faster than the new node is downloading and checking them. Thus whilst it may be tempting to try and expose a notion of ‘doneness’ to the act of joining a group, it’s better to see the act of joining as happening at a specific point in time and the resultant flood of transaction data as an ongoing stream, rather than being like a traditional file download.

When a transaction is sent to the vault, it always undergoes a relevancy test, regardless of whether it is in a group or not (see §8). This test is extended to check also for the signatures of any groups the node is a member of. If there’s a match then the transaction’s states are all considered relevant. In addition, the vault looks up which nodes it invited to this group, and also which nodes invited it, removes any nodes that have recently sent us this transaction and then kicks off a `PropagateTransactionToGroup` flow with each of them. The other side of this flow checks if the transaction is already known, if not requests it, checks that it is indeed signed by the group in question, resolves it and then assuming success, sends it to the vault. In this way a transaction added by any member of the group propagates up and down the membership tree until all the members have seen it. Propagation is idempotent - if the vault has already seen a transaction before then it isn’t processed again.

The structure we have so far has some advantages and one big disadvantage. The advantages are:

- Simplicity The core data model is unchanged. Access control is handled using existing tools like signatures, certificates and flows.
- Privacy It is possible to join a group without the other members being aware that you have done so. It is possible to create groups without non-members knowing the group exists.
- Scalability Groups are not registered in any central directory. A group that exists between four parties imposes costs only on those four.
- Performance Groups can be created as fast as you can generate keypairs and invite other nodes to join you.
- Responsibility For every member of the group there is always a node that has a responsibility for sending you new data under the protocol (the inviting node). Unlike with Kademlia style distributed hash tables, or Bitcoin style global broadcast, you can never find yourself in a position where you didn’t receive data yet nobody has violated the protocol. There are no points at which you pick a random selection of nodes and politely ask them to do something for you, hoping that they’ll choose to stick around.

The big disadvantage is that it’s brittle. If you have a membership tree and a

node goes offline for a while, then propagation of data will split and back up in the outbound queues of the parents and children of the offline node until it comes back.

To strengthen groups we can add a new feature, membership broadcasts. Members of the group that have write access may choose to sign a membership announcement and propagate it through the tree. These announcements are recorded in the local database of each node in the group. Nodes may include these announced members when sending newly added transactions. This converts the membership tree to a graph that may contain cycles, but infinite propagation loops are not possible because nodes ignore announcements of new transactions/attachments they've already received. Whether a group prefers privacy or availability may be hinted in the certificate that defines it: if availability is preferred, this is a signal that members should always announce themselves (which would lead to a mesh).

The network map for a network defines the event horizon, the span of time that is allowed to elapse before an offline node is considered to be permanently gone. Once a peer has been offline for longer than the event horizon any nodes that invited it remove it from their local tables. If a node was invited to a group by a gone peer and there are no other nodes that announced their membership it can use, the node should post a message to a queue and/or notify the administrator, as it's now effectively been evicted from the group.

The resulting arrangement may appear similar to a gossip network. However the underlying membership tree structure remains. Thus when all nodes are online (or online enough) messages are guaranteed to propagate to everyone in the network. You can't get situations where a part of the group has become split from the rest without anyone being aware of that fact; an unlikely but possible occurrence in a gossip network. It also isn't like a distributed hash table where data isn't fully replicated, so we avoid situations where data has been added to the group but stops being available due to node outages. It is always possible to reason about the behaviour of the network and always possible to assign responsibility if something goes wrong.

Note that it is not possible to remove members after they have been added to a group. We could provide a remove announcement but it'd be advisory only: nothing stops nodes from ignoring it. It is also not possible to enumerate members of a group because there is no requirement to do a membership broadcast when you join and no way to enforce such a requirement.

## 13 Deterministic JVM

It is important that all nodes that process a transaction always agree on whether it is valid or not. Because transaction types are defined using JVM bytecode, this means the execution of that bytecode must be fully deterministic. Out of

the box a standard JVM is not fully deterministic, thus we must make some modifications in order to satisfy our requirements. Non-determinism could come from the following sources:

- Sources of external input e.g. the file system, network, system properties, clocks.
- Random number generators.
- Different decisions about when to terminate long running programs.
- `Object.hashCode()`, which is typically implemented either by returning a pointer address or by assigning the object a random number. This can surface as different iteration orders over hash maps and hash sets.
- Differences in hardware floating point arithmetic.
- Multi-threading.
- Differences in API implementations between nodes.
- Garbage collector callbacks.

To ensure that the contract verify function is fully pure even in the face of infinite loops we construct a new type of JVM sandbox. It utilises a bytecode static analysis and rewriting pass, along with a small JVM patch that allows the sandbox to control the behaviour of hashcode generation. Contract code is rewritten the first time it needs to be executed and then stored for future use.

The bytecode analysis and rewrite performs the following tasks:

- Inserts calls to an accounting object before expensive bytecodes. The goal of this rewrite is to deterministically terminate code that has run for an unacceptably long amount of time or used an unacceptable amount of memory. Expensive bytecodes include method invocation, allocation, backwards jumps and throwing exceptions.
- Prevents exception handlers from catching `Throwable`, `Error` or `ThreadDeath`.
- Adjusts constant pool references to relink the code against a ‘shadow’ JDK, which duplicates a subset of the regular JDK but inside a dedicated sandbox package. The shadow JDK is missing functionality that contract code shouldn’t have access to, such as file IO or external entropy.
- Sets the `strictfp` flag on all methods, which requires the JVM to do floating point arithmetic in a hardware independent fashion. Whilst we anticipate that floating point arithmetic is unlikely to feature in most smart contracts (big integer and big decimal libraries are available), it is available for those who want to use it.
- Forbids `invokedynamic` bytecode except in special cases, as the libraries that support this functionality have historically had security problems and

it is primarily needed only by scripting languages. Support for the specific lambda and string concatenation metafactories used by Java code itself are allowed.

- Forbids native methods.
- Forbids finalizers.

The cost instrumentation strategy used is a simple one: just counting bytecodes that are known to be expensive to execute. Method size is limited and jumps count towards the budget, so such a strategy is guaranteed to eventually terminate. However it is still possible to construct bytecode sequences by hand that take excessive amounts of time to execute. The cost instrumentation is designed to ensure that infinite loops are terminated and that if the cost of verifying a transaction becomes unexpectedly large (e.g. contains algorithms with complexity exponential in transaction size) that all nodes agree precisely on when to quit. It is *not* intended as a protection against denial of service attacks. If a node is sending you transactions that appear designed to simply waste your CPU time then simply blocking that node is sufficient to solve the problem, given the lack of global broadcast.

Opcodes budgets are separate per opcode type, so there is no unified cost model. Additionally the instrumentation is high overhead. A more sophisticated design would be to statically calculate bytecode costs as much as possible ahead of time, by instrumenting only the entry point of ‘accounting blocks’, i.e. runs of basic blocks that end with either a method return or a backwards jump. Because only an abstract cost matters (this is not a profiler tool) and because the limits are expected to be set relatively high, there is no need to instrument every basic block. Using the max of both sides of a branch is sufficient when neither branch target contains a backwards jump. This sort of design will be investigated if the per category opcode-at-a-time accounting turns out to be insufficient.

A further complexity comes from the need to constrain memory usage. The sandbox imposes a quota on bytes *allocated* rather than bytes *retained* in order to simplify the implementation. This strategy is unnecessarily harsh on smart contracts that churn large quantities of garbage yet have relatively small peak heap sizes and, again, it may be that in practice a more sophisticated strategy that integrates with the garbage collector is required in order to set quotas to a usefully generic level.

Control over `Object.hashCode()` takes the form of new JNI calls that allow the JVM’s thread local random number generator to be reseeded before execution begins. The seed is derived from the hash of the transaction being verified.

Finally, it is important to note that not just smart contract code is instrumented, but all code that it can transitively reach. In particular this means that the ‘shadow JDK’ is also instrumented and stored on disk ahead of time.



## 14 Scalability

Scalability of block chains and block chain inspired systems has been a constant topic of discussion since Nakamoto first proposed the technology in 2008. We make a variety of choices and tradeoffs that affect and ensure scalability. As most of the initial intended use cases do not involve very high levels of traffic, the reference implementation is not heavily optimised. However, the architecture allows for much greater levels of scalability to be achieved when desired.

**Partial visibility.** Nodes only encounter transactions if they are involved in some way, or if the transactions are dependencies of transactions that involve them in some way. This loosely connected design means that it is entirely possible for most nodes to never see most of the transaction graph, and thus they do not need to process it. This makes direct scaling comparisons with other distributed and decentralised database systems difficult, as they invariably measure performance in transactions/second. For Corda, as writes are lazily replicated on demand, it is difficult to quote a transactions/second figure for the whole network.

**Distributed node.** At the center of a Corda node is a message queue broker. Nodes are logically structured as a series of microservices and have the potential in future to be run on separate machines. For example, the embedded relational database can be swapped out for an external database that runs on dedicated hardware. Whilst a single flow cannot be parallelised, a node under heavy load would typically be running many flows in parallel. As flows access the network via the broker and local state via an ordinary database connection, more flow processing capacity could be added by just bringing online additional flow workers. This is likewise the case for RPC processing.

**Signatures outside the transactions.** Corda transaction identifiers are the root of a Merkle tree calculated over its contents excluding signatures. This has the downside that a signed and partially signed transaction cannot be distinguished by their canonical identifier, but means that signatures can easily be verified in parallel. Corda smart contracts are deliberately isolated from the underlying cryptography and are not able to request signature checks themselves: they are run *after* signature verification has taken place and don't execute at all if required signatures are missing. This ensures that signatures for a single transaction can be checked concurrently even though the smart contract code for that transaction is not parallelisable. (note that unlike some other systems, transactions involving the same contracts *can* be checked in parallel.)

**Multiple notaries.** It is possible to increase scalability in some cases by bringing online additional notary clusters. Note that this only adds capacity if

the transaction graph has underlying exploitable structure (e.g. geographical biases), as a purely random transaction graph would end up constantly crossing notaries and the additional transactions to move states from one notary to another would negate the benefit. In real trading however the transaction graph is not random at all, and thus this approach may be helpful.

**Asset reissuance.** In the case where the issuer of an asset is both trustworthy and online, they may exit and re-issue an asset state back onto the ledger with a new reference field. This effectively truncates the dependency graph of that asset which both improves privacy and scalability, at the cost of losing atomicity (it is possible for the issuer to exit the asset but not re-issue it, either through incompetence or malice).

**Non-validating notaries.** The overhead of checking a transaction for validity before it is notarised is likely to be the main overhead for non-BFT notaries. In the case where raw throughput is more important than ledger integrity it is possible to use a non-validating notary. See §7.3.

The primary bottleneck in a Corda network is expected to be the notary clusters, especially for byzantine fault tolerant (BFT) clusters made up of mutually distrusting nodes. BFT clusters are likely to be slower partly because the underlying protocols are typically chatty and latency sensitive, and partly because the primary situation when using a BFT protocol is beneficial is when there is no shared legal system which can be used to resolve fraud or other disputes, i.e. when cluster participants are spread around the world and thus the speed of light becomes a major limiting factor.

The primary bottleneck in a Corda node is expected to be flow checkpointing, as this process involves walking the stack and heap then writing out the snapshotted state to stable storage. Both of these operations are computationally intensive. This may seem unexpected, as other platforms typically bottleneck on signature checking operations. It is worth noting though that the main reason other platforms do not bottleneck on checkpointing operations is that they typically don't provide any kind of app-level robustness services at all, and so the cost of checkpointing state (which must be paid eventually!) is accounted to the application developer rather than the platform. When a flow developer knows that a network communication is idempotent and thus can be replayed, they can opt out of the checkpointing process to gain throughput at the cost of additional wasted work if the flow needs to be evicted to disk. Note that checkpoints and transaction data can be stored in any NoSQL database (such as Cassandra), at the cost of a more complex backup strategy.

Due to partial visibility nodes check transaction graphs 'just in time' rather than as a steady stream of announcements by other participants. This complicates the question of how to measure the scalability of a Corda node. Other block chain systems quote performance as a constant rate of transactions per unit

time. However, our ‘unit time’ is not evenly distributed: being able to check 1000 transactions/sec is not necessarily good enough if on presentation of a valuable asset you need to check a transaction graph that consists of many more transactions and the user is expecting the transaction to show up instantly. Future versions of the platform may provide features that allow developers to smooth out the spikey nature of Corda transaction checking by, for example, pre-pushing transactions to a node when the developer knows they will soon request the data anyway.

## 15 Privacy

Privacy is not a standalone feature in the way that many other aspects described in this paper are, so this section summarises features described elsewhere. Corda exploits multiple techniques to improve user privacy over other distributed ledger systems:

**Partial data visibility.** Transactions are not globally broadcast as in many other systems.

**Transaction tear-offs.** Transactions are structured as Merkle trees, and may have individual subcomponents be revealed to parties who already know the Merkle root hash. Additionally, they may sign the transaction without being able to see all of it. See §5.7

**Key randomisation.** The vault generates and uses random keys that are unlinkable to an identity without the corresponding linkage certificate. See §8.

**Graph pruning.** Large transaction graphs that involve liquid assets can be ‘pruned’ by requesting the asset issuer to re-issue the asset onto the ledger with a new reference field. This operation is not atomic, but effectively unlinks the new version of the asset from the old, meaning that nodes won’t attempt to explore the original dependency graph during verification.

Corda has been designed with the future integration of additional privacy technologies in mind. Of all potential upgrades, three are particularly worth a mention.

**Secure hardware.** Although we narrow the scope of data propagation to only nodes that need to see that data, ‘need’ can still be an unintuitive concept in a decentralised database where often data is required only to perform security

checks. We have successfully experimented with running contract verification inside a secure enclave protected JVM using Intel SGX™, an implementation of the ‘trusted computing’ concept<sup>27</sup>. Secure hardware platforms allow computation to be performed in an undebuggable tamper-proof execution environment, for the software running inside that environment to derive encryption keys accessible only to that instance, and for the software to *remotely attest* to a third party over the internet that it is indeed running in the secure state. By having nodes remotely attest to each other that they are running smart contract verification logic inside an enclave it becomes possible for the dependencies of a transaction to be transmitted to a peer encrypted under an enclave key, thus allowing them to verify the dependencies using software they have audited themselves, but without being able to see the data on which it operates.

Secure hardware opens up the potential for a one-shot privacy model that would dramatically simplify the task of writing smart contracts. However, it does still require the sensitive data to be sent to the peer who may then attempt to attack the hardware or exploit side channels to extract business intelligence from inside the encrypted container.

**Mix networks.** Some nodes may be in the position of learning about transactions that aren’t directly related to trades they are doing, for example notaries or regulator nodes. Even when key randomisation is used these nodes can still learn valuable identity information by simply examining the source IP addresses or the authentication certificates of the nodes sending the data for notarisation. The traditional cryptographic solution to this problem is a *mix network*<sup>28</sup>. The most famous mix network is Tor, but a more appropriate design for Corda would be that of an anonymous remailer. In a mix network a message is repeatedly encrypted in an onion-like fashion using keys owned by a small set of randomly selected nodes. Each layer in the onion contains the address of the next ‘hop’. Once the message is delivered to the first hop, it decrypts it to reveal the next encrypted layer and forwards it onwards. The return path operates in a similar fashion. Adding a mix network to the Corda protocol would allow users to opt-in to a privacy upgrade, at the cost of higher latencies and more exposure to failed network nodes.

**Zero knowledge proofs.** The holy grail of privacy in decentralised database systems is the use of zero knowledge proofs to convince a peer that a transaction is valid, without revealing the contents of the transaction to them. Although these techniques are not yet practical for execution of general purpose smart contracts, enormous progress has been made in recent years and we have designed our data model on the assumption that we will one day wish to migrate to the use of *zero knowledge succinct non-interactive arguments of knowledge*<sup>29</sup> (‘zkSNARKs’). These algorithms allow for the calculation of a fixed-size mathematical proof that a program was correctly executed with a mix of public and private inputs. Programs can be expressed either directly as a system of

low-degree multivariate polynomials encoding an algebraic constraint system, or by execution on a simple simulated CPU ('vnTinyRAM') which is itself implemented as a large pre-computed set of constraints. Because the program is shared the combination of an agreed upon function (i.e. a smart contract) along with private input data is sufficient to verify correctness, as long as the prover's program may recursively verify other proofs, i.e. the proofs of the input transactions. The BCTV zkSNARK algorithms rely on recursive proof composition for the execution of vnTinyRAM opcodes, so this is not a problem. The most obvious integration with Corda would require tightly written assembly language versions of common smart contracts (e.g. cash) to be written by hand and aligned with the JVM versions. Less obvious but more powerful integrations would involve the addition of a vnTinyRAM backend to an ahead of time JVM bytecode compiler, such as Graal<sup>30</sup>, or a direct translation of Graal's graph based intermediate representation into systems of constraints. Direct translation of an SSA-form compiler IR to constraints would be best integrated with recent research into 'scalable probabilistically checkable proofs'<sup>31</sup>, and is an open research problem.

## 16 Conclusion

We have presented Corda, a decentralised database designed for the financial sector. It allows for a unified data set to be distributed amongst many mutually distrusting nodes, with smart contracts running on the JVM providing access control and schema definitions. A novel continuation-based persistence framework assists developers with coordinating the flow of data across the network. An identity management system ensures that parties always know who they are trading with. Notaries ensure algorithmic agility with respect to distributed consensus systems, and the system operates without mining or a block chain.

A standard type system is provided for the modelling of financial logic. The design considers security throughout: it supports the integration of secure signing devices for transaction authorisation, secure enclaves for transaction processing, composite keys for expressing complex authorisation policies, and is based on binary protocols with length-prefixed buffers throughout for the systematic avoidance of common buffer management exploits. Users may analyse ledger data relevant to them by issuing ordinary SQL queries against mature database engines, and may craft complex multi-party transactions with ease in programming languages that are already familiar to them.

Finally, the platform defines standard ways to integrate the global ledger with financial infrastructure like high performance markets and netting services.

## 17 Acknowledgements

The author would like to thank Richard Gendal Brown, James Carlyle, Shams Asari, Rick Parker, Andras Slemmer, Ross Nicoll, Andrius Dagys, Matthew Nesbit, Jose Coll, Katarzyna Streich, Clinton Alexander, Patrick Kuo, Richard Green, Ian Grigg, Mark Oldfield and Roger Willis for their insights and contributions to this design. We would also like to thank Sofus Mortesen for his work on the universal contract DSL, and the numerous architects and subject matter experts at financial institutions around the world who contributed their knowledge, requirements and ideas. Thanks also to the authors of the many frameworks, protocols and components we have built upon.

Finally, we would like to thank Satoshi Nakamoto. Without him none of it would have been possible.

## Bibliography

- [1] Brown, Carlyle, Grigg, Hearn. *Corda: An introduction*. <http://r3cev.com/s/corda-introductory-whitepaper-final.pdf>, 2016.
- [2] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. Bigtable: A distributed storage system for structured data. *ACM Trans. Comput. Syst.*, 26(2):4:1–4:26, June 2008.
- [3] Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. <https://bitcoin.org/bitcoin.pdf>, 2008.
- [4] Lindholm, Yellin, Bracha, & Buckley. *The Java Virtual Machine Specification Java SE 8 Edition*. <https://docs.oracle.com/javase/specs/jvms/se8/jvms8.pdf>, 2015.
- [5] Buterin et al. *A Next-Generation Smart Contract and Decentralized Application Platform*. <https://github.com/ethereum/wiki/wiki/%5BEnglish%5D-White-Paper>, 2014.
- [6] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13*, SSYM'04, pages 21–21, Berkeley, CA, USA, 2004. USENIX Association.
- [7] OASIS. Advanced message queuing protocol (amqp) version 1.0, 2012.
- [8] Mike Hearn. Bitcoin micropayment channels. <https://bitcoinj.github.io/working-with-micropayments>, 2014.
- [9] Mike Hearn, Gavin Andresen. Bitcoin payment protocol. <https://github.com/bitcoin/bips/blob/master/bip-0070.mediawiki>, 2013.

- [10] java.time.Instant documentation. <https://docs.oracle.com/javase/8/docs/api/java/time/Instant.html>, 2014.
- [11] PKWARE. Zip file format. <https://pkware.cachefly.net/webdocs/casestudies/APPNOTE.TXT>, 1989.
- [12] David Siegel. <http://www.coindesk.com/understanding-dao-hack-journalists/>, 2016.
- [13] Pieter Wuille. Hierarchical deterministic wallets. <https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki>, 2013.
- [14] Stefan Thomas. Crypto-conditions. <https://interledger.org/five-bells-condition/spec.html>, 2016.
- [15] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, USENIX ATC'14, pages 305–320, Berkeley, CA, USA, 2014. USENIX Association.
- [16] Christopher Malmo. <http://motherboard.vice.com/read/bitcoin-is-unsustainable>, 2015.
- [17] Tim Swanson. <http://tabbforum.com/opinions/settlement-risks-involving-public-blockchains>, 2016.
- [18] Alysson Bessani, João Sousa, and Eduardo E. P. Alchieri. State machine replication for the masses with bft-smart. In *Proceedings of the 2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, DSN '14, pages 355–362, Washington, DC, USA, 2014. IEEE Computer Society.
- [19] Hal Finney. Best practice for fast transaction acceptance - how high is the risk? <https://bitcointalk.org/index.php?topic=3441.msg48384#msg48384>.
- [20] Jsr 338: Java persistence api. [http://download.oracle.com/otn-pub/jcp/persistence-2\\_1-fr-eval-spec/JavaPersistence.pdf?AuthParam=1478095024\\_77b7362fd5bd185ebf8d2cd2a071a14d](http://download.oracle.com/otn-pub/jcp/persistence-2_1-fr-eval-spec/JavaPersistence.pdf?AuthParam=1478095024_77b7362fd5bd185ebf8d2cd2a071a14d), 2013.
- [21] Jsr 349: Bean validation constraints. <https://www.jcp.org/en/jsr/detail?id=349>, 2013.
- [22] Simon Peyton Jones, Jean-Marc Eber, and Julian Seward. Composing contracts: An adventure in financial engineering (functional pearl). *SIGPLAN Not.*, 35(9):280–292, September 2000.

- [23] David J. Pearce and Lindsay Groves. Designing a verifying compiler: Lessons learned from developing whiley. *Science of Computer Programming*, 113, Part 2:191 – 220, 2015. Formal Techniques for Safety-Critical Systems.
- [24] Markus Voelter and Sascha Lisson. Supporting diverse notations in mps’ projectional editor. In *Proceedings of the 2nd International Workshop on The Globalization of Modeling Languages co-located with ACM/IEEE 17th International Conference on Model Driven Engineering Languages and Systems, GEMOC@Models 2014, Valencia, - Spain, September 28, 2014.*, pages 7–16, 2014.
- [25] Bitcoin trezor device. <https://bitcointrezor.com/>, 2016.
- [26] Reactivex. <https://www.reactivex.io>, 2016.
- [27] C. Mitchell and Institution of Electrical Engineers. *Trusted Computing*. Computing and Networks Series. Institution of Engineering and Technology, 2005.
- [28] David L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–90, February 1981.
- [29] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 781–796, San Diego, CA, August 2014. USENIX Association.
- [30] Graal research compiler. <http://openjdk.java.net/projects/graal/>, 2016.
- [31] Eli Ben-Sasson, Iddo Ben-Tov, Alessandro Chiesa, Ariel Gabizon, Daniel Genkin, Matan Hamilis, Evgenya Pergament, Michael Riabzev, Mark Silberstein, Eran Tromer, and Madars Virza. Computational integrity with a public random string from quasi-linear pcps. Cryptology ePrint Archive, Report 2016/646, 2016. <http://eprint.iacr.org/2016/646>.