Networked Carbon Markets: Permissionless Innovation with Distributed Ledgers?

White Paper
by

Adrian Jackson, Ashley Lloyd, Justin Macinante, Markus Hüwener

Networked Carbon Markets:
Permissionless Innovation with Distributed Ledgers?

Adrian Jackson, Ashley Lloyd, Justin Macinante, Markus Hüwener

Overview

A principal objective of global efforts to avoid dangerous climate change impacts is mitigation of greenhouse gas (GHG) emissions. Putting a price on carbon is one obvious way to do this, whether by taxing the activity or by creating a market in which the cost of carbon (i.e., the related GHG emission) is internalised.

Carbon markets and carbon trading are, thus, key components in the climate change mitigation response. Whilst local, national or regional carbon trading schemes exist, and more are being put in place, connecting these markets has the potential to allow a more integrated, efficient, and globally consistent price on carbon which will engender greater confidence in the market, investment and, ultimately, help foster new technology development through climate finance.

The Paris Agreement directly supports the idea of international trading between carbon markets by recognising that Parties may engage voluntarily in cooperative approaches that allow internationally transferred mitigation outcomes (ITMOs) to contribute towards Nationally Determined Contributions (NDCs). Such engagements should promote sustainable development, ensure environmental integrity and transparency, including governance, and apply robust accounting, in particular in the avoidance of double counting.

The possibility of double counting also points to the practical challenges faced by many multinational companies and businesses operating in multiple jurisdictions with diverse carbon markets: a difficulty in reconciling their overall carbon footprint (and exposures) across all aspects of their business. Such entities would benefit from the connection of the diverse carbon markets they are operating in, and the ability to trade carbon assets across those carbon markets.

The challenge is that each individual carbon market has its own legal and regulatory framework and its own rules for assigning and accounting for the carbon units it trades. This presents significant legal and political hurdles that militate against a single carbon market being created within which all carbon trading might take place.

To progress from the present situation, a first step would be to connect two (or, at least, a small number of) carbon markets with bilateral agreements on how the markets will interoperate and collaborate. However, negotiating such arrangements is slow. Moving beyond such limited linking would be even more difficult and time-consuming, and hard to achieve with a top-down, negotiated approach. Divergence between the size and nature of respective economies, GHG emissions profiles, population sizes, and other factors, naturally make governments nervous of external influences and jealous of sovereignty when negotiating such arrangements.

This encourages the situation where to enable the connection of multiple markets, which are in essence *data infrastructures*, the alternative 'bottom-up' solution is to provide a system where carbon may be traded between a range of markets without forcing legal and regulatory homogeneous standardisation and conformance on those markets. The goal sought is a distributed infrastructure that is not controlled by a single organisation or entity and in which individual organisations (carbon markets in this context) can join or leave as they see appropriate to their own circumstances.

It is these multiple, dynamic, layers of authority, characterised by economic, systemic and geographic heterogeneity that require any market mechanism to be supported by a distributed data infrastructure with multiple layers of permission. One candidate technology is the 'Distributed Ledger' (DL), which directly supports multiple layers of permission for different users, even though it is widely implemented as an 'unpermissioned' system, such as the one used by the BitCoin crypto currency. Whilst DL technology does not provide any significant functionality that could not be achieved with established distributed database technology and custom application development, the combination of a distributed database with public/private key encryption and a decentralised infrastructure does provide the potential for some innovative solutions to data sharing or transaction management application areas and a good first-order match to the emerging requirements for an interoperable carbon market infrastructure.

A number of general DL implementations already exist that could form a diverse set of potential solutions. However, to meet the objectives of the Paris Agreement, a solution needs to be found that facilitates adoption by sufficient numbers of users and capacity for diffusion across geographical borders, throughout dynamic hierarchies of stakeholders, to form a global-scale distributed infrastructure. The established literature on innovation and technology diffusion gives guidance on this issue, but no ability to predict success.

The purpose of this White Paper, therefore, is to outline the most important questions identified in relation to the connecting of these carbon markets through application of DL technologies, and outline the authors' current thoughts on those questions. It is designed to facilitate discussion in the wider community, not as a final position on such questions. This is a starting point for designing a system that could change the face of carbon markets.

Table of Contents

Why networked carbon markets?

Networking of carbon markets (NCM) is an approach to connecting diverse and heterogeneous carbon markets, not by making them adapt and reach conformity with each other, but rather by recognising the differences between them and placing a value on those differences. NCM has the potential to enable significant climate change mitigation by making the trading of and accounting for carbon a global concern. This should lead to both a sustainable price for carbon across a range of carbon markets, incentives for investing in carbon mitigation projects and manufacturing in carbon efficient ways, and much more transparent auditing/accountability for carbon emission mitigation in global carbon markets, thereby meeting the requirements of the Paris Agreement for such transfers in a secure and cost-efficient environment.

Carbon markets[1] are highly regulated and vary in scope, design, implementation, and rules and standards across the globe. They are not free markets: regulators in each carbon market control supply of tradable emission units, as well as set requirements for carbon accounting. Therefore, any system to network such markets needs to be able to cope with these variations in policy and approach and still enable trading between participants.

Furthermore, given that carbon emissions and climate change are global issues, a solution is required that is accessible to developing nations as well as developed nations, without barriers to adoption such as heavy investment in bespoke software or hardware for nations with limited resources, but still allowing participation on an equal footing with rich, developed economies.

What types/levels of market are we looking to support?

Whilst the overall goal is to allow the networking of a range of individual carbon markets to enable the international transfers of mitigation outcomes, it is necessary to recognise that there are multiple levels where carbon trading or carbon accounting and tracking may occur.

Corporate Level

Firstly, individual corporations may need systems to allow them to undertake carbon accounting across their various jurisdictions, locations, and sub-organisations. Businesses will currently be doing some form of tracking of this data, be it with spreadsheets that are e-mailed between participants, or bespoke applications/data infrastructures. Different jurisdictions may require different accounting/tracking procedures, meaning a single spreadsheet or application may not be able to do this accounting across all those jurisdictions. There are also likely to be oversight and auditing requirements within the corporation that will need to either approve actions on the data, or at least be able to see some of the data held in the system.

Single Emissions Trading System (ETS) Level

Secondly, individual carbon markets (ETSs), be they local area, single country, or regional, require systems to organise and administer their operations. This includes assigning and distributing emission units to companies, allowing for trading of emission units between organisations (subject to the rules of the ETS), and allowing companies to surrender/report emission units for obligations when required (recorded through ETS administrator-maintained registries). Such ETSs are likely to have strict timetables, with emission units and obligations time limited. In such markets a range of different organisations should be able to participate. Traditionally, this has been mainly the remit of large, multi-national, corporations, but for climate mitigation to truly be a success it should be able to account for much smaller organisations as well[2], including those without the funds to purchase and operate costly software/computing infrastructure.

[1] In this paper, the authors are focusing in the first instance on the connecting of government-instigated emission trading schemes (ETSs), as a first step towards the ultimate objective of being able to connect more generally a broader range of mitigation actions/carbon pricing measures (e.g., including carbon taxes, renewable energy certificate schemes and so on). Note also that the expressions 'market', 'carbon market' and 'ETS' or 'ETS market' are used interchangeably throughout the paper to mean the same..
[2] Or even for individuals to trade, if such is approved by the relevant jurisdiction.

Networked Carbon Markets

Finally, the aim is to network these individual ETSs, to allow trading between the different jurisdictions' markets. Given the varied natures of ETSs, it is unlikely that direct, unrestricted, trading between all markets is possible. The varying timescale and rules of different schemes, the auditing, environmental, and accountability standards of the different markets, and other differences, mean there are large variations between markets. As such, networking is based on the idea that the ETSs, designed and implemented to achieve mitigation of GHG emissions to the atmosphere, will generate mitigation outcomes upon which it is possible to place a measureable value. How such a mitigation value (MV) is arrived at, while integral to the concept of NCM, is not part of the scope of this paper.[3] Suffice it to say, for the purpose of this paper, that establishing a defined mitigation value (MV) methodology will eventually facilitate trading between ETS markets based on conversion rates, derived from the respective MVs of the counterparties (or, at least, the MVs of the emission units from their respective ETSs).

Global agreement across all jurisdictions to be networked is unlikely to be feasible, so a mechanism is going to be required that recognises individual markets are currently operating separately, but should they so decide, provides for how they can interact with other markets. The system will need to be able to support a range of different interactions between markets controlling how/when/whether they interact with each other. This may even include options to allow individual organisations to decide how to trade with organisations in other markets and decide their own trading price and contract conditions.

Flexible membership is also important: individual jurisdictions should be able to join or leave the networked trading arrangements, as they deem appropriate to their individual circumstances. The ability to join once a system is already operational, or leave after being part of such a system, is key in supporting both a growing trading system and the Paris Agreement intent of Parties engaging in cooperative approaches voluntarily.

Ideally, the intention is to create a single system, or at least a single set of tools (a *software infrastructure*) that can enable all three levels we have described above, ranging from providing the functionality a single organisation needs to manage its carbon assets/exposure, to a platform that can network multiple carbon markets. Whilst there is no technical necessity to have a single software infrastructure for all these levels of use, such an infrastructure would provide the potential for benefits, from enabling easy adoption and sharing of data, to full system auditing and tracking of carbon to ensure accountability and discourage fraud in the system. Further, as will be discussed in more detail, the technical issues that need to be addressed for such a system inside an organisation are applicable at the networked market level and vice versa. Therefore, a technical solution for a single organisation should support participation in a network of carbon markets comprised of organisations using the same system.

It is noted that, whilst this discussion is of a single software *infrastructure*, such a focus does not necessarily mandate a single software *solution*. This paper *is not* proposing a bespoke, closed-source, commercial solution that requires all participants to purchase the same software and hardware to participate in the market. Indeed, the discussion is broadened to include consideration of the impact of design considerations on adoption and diffusion. In doing so, the analysis recognises that to meet the targets set by the Paris Agreement, any proposed *solution* needs to diffuse across countries and stakeholder groups very quickly as well as support market innovations that will draw more investment into the operation of those markets. Hence, what is proposed is to design a data infrastructure that can enable all of the functionality that has been outlined as being desirable, and create a reference implementation that can be used, without precluding others implementing their own solutions and still participating in the system in general.

Furthermore, it is noted that it would be possible to create a solution for the requirements outlined without using distributed ledger technologies (DLTs). However, the core benefit that DLTs promote is the ability to ensure access to the data in the network is distributed and available to all participants (subject to any access restrictions deemed necessary in carbon markets/carbon trading[4]) providing equal participation in the market. Also, other functionality, such as auditability and visibility of transactions, can easily be enabled and

[3] The derivation of methodologies for arriving at MV values is the subject of research by other groups, including some under the purview of the World Bank.

[4] As we will discuss later in the paper, we are not necessarily advocating fully visible, fully unpermissioned, ledgers. There may be good reasons that access to or permissions to change some data may be restricted in the system.

extended/augmented as required, ensuring that carbon market(s) can be made as transparent, trusted, and liquid as possible.

What is distributed ledger technology (DLT)?

DLT covers a wide range of potential functionality, from completely decentralised, permissionless systems in which anyone can participate and be a part of the data infrastructure, to permissioned and controlled systems which are very similar to current data networks controlled by a single or small number of entities. As such, it is necessary to be able to understand the variation in DLT, the choices that can be made when designing a DLT based system, and the impacts these choices would have on the functionality provided by the resultant system.

The following features (*Table 1*) are recognised as key for any system to be considered a DLT:

Table 1: Design features defining a Distributed Ledger

Feature	Description
Distributed infrastructure	The system must be composed of multiple entities, or nodes, that each have a full copy of the ledger. This implies that nodes can be added and removed from the system without affecting the overall operation of the system or the integrity of the data stored in the ledger.
Encrypted transactions	The ability of participants to use public/private key encryption to authorise/assign/interact with transactions in the system.
Global consensus	The nodes in the system must have a mechanism for coming to agreement about what is a valid addition to the ledger and what is the current state of the ledger. This means that entries to the ledger can be added by different nodes and the ledgers held by the nodes in the system will be updated to reflect the changes made at individual nodes. This also implies that there is a system for deciding whether an entry, or set of entries, that has been added to one or more of the copies of the ledger is valid and reject entries that are not valid.
Immutability	The ledger is accumulative: once entries have been accepted into the ledger it should not be possible to go back and amend them. This is achievable by including some type of *hashing*[5] in the ledger, generating a key from the data in an entry in the ledger and incorporating that key in the subsequent entries. This process makes it extremely hard to alter entries in the ledger without that alteration being noticed as any change to the data in the ledger will change the key (or hash) generated from that data. An audit of the ledger (passing back through the ledger, calculating the hashes for each entry and comparing them with the hash stored in the subsequent entries) will quickly highlight if entries have been changed. Immutability becomes more problematic in the distributed infrastructure as it is possible for different versions of the ledger to exist (remembering that each node has a version of the ledger and they may be in different states of being up-to-date). This is overcome by the global consensus model of the system, but this does need careful consideration to ensure that immutability is not compromised. For instance, if a node appears with a ledger that diverges from the majority of other ledgers many entries ago, should this be accepted as the new ledger? If so, it would change entries already added to many existing ledgers. There are mechanisms for addressing these issues which must be considered when designing the immutability and consensus models of the DLT.

[5] A 'hash' is a one-way mathematical function that summarises a piece of data, regardless of its size, as a piece of unique, fixed-size, short data. The hash function turns data into a key of random characters called a 'hash'.

The following features (Table 2) are recognised as configurable[6] for a DLT depending on the application it is being used for and the environment in which it exists:

Table 2: Design features of DLTs that are configurable.

Feature	**Description**
Permissioning	DLT systems have traditionally been viewed as permissionless systems, in which anyone can participate and be part of the system. No one entity, organisation, or system participant is in charge or has any more authority than any other participant. However, this is not necessarily a fundamental property of distributed ledgers. The authors would consider designing ledgers that do have different types of participants, different levels of authorisation, should that be required or beneficial for the proposed application of DLT.
Proof of work	One of the key features of permissionless crypto currency systems, such as BitCoin, is the ability to decide when a transaction (addition to the ledger) is valid or not. With no permissions in the system there is no authority to consult on validity, therefore a computational system for ensuring validity is required. This is tied up with the global consensus feature, which *proof of work* (also known as *mining*) schemes use to both control the availability of entries in the ledger (coins in the crypto currency) and ensure that the selection of which node can add an entry next is random in the system.
Smart Contracts	When entries are added to the ledger it is also possible to automatically perform some actions associated with that entry or ledger. A *smart contract* is simply a set of operations to be performed automatically once the contract has been fulfilled.[7] In the context of DLTs, fulfilling the contract is likely to be the same as adding an entry into the ledger, although smart contracts could rely upon multiple ledger entries to be fulfilled.
Settlement, exchanges, or payment systems	Whilst DLTs can be used as crypto currencies and can also support many business processes, the actual transferral of money between parties, or settlement for physical assets, need not be part of the DLT. Indeed, there are numerous BitCoin exchanges that operate to allow transfer between BitCoin and Fiat currencies. These exchanges use the BitCoin DLT but are not part of the DLT functionality itself.

Designing for Use Cases

The distinction between functionality that defines a DLT versus those features that are configurable provides scope for designing DLTs that are quite specific to anticipated Use Cases and User Groups.

Permissionless systems, for example, are open to all for participation. Anyone may become a ledger node (i.e. hold a copy of the ledger) and add valid entries. The converse of this is a permissioned system, where only authorised entities can hold a copy of the ledger or participate in transactions. However, it may also be possible for a range of permissions to be applied to the system, i.e. permission may be required to become a ledger node but individuals may interact with ledger nodes without permissions; or anyone may be able to become a ledger node but may require permissions to add entries to the ledger; or there may be different permissions for adding entries compared to viewing entries; and so on.

To give this more structure from a design perspective, there are two different types of permissions that determine a participant's scope of interactions with the ledger:

- Hosting the DLT (i.e. having a copy of the ledger, being a node in the system)
- Interacting with the ledger (i.e. viewing, or adding to, the ledger)

[6] Configurable: Potentially present or not in the system. Also, potentially different levels or types of this functionality present in the system.
[7] Or in a more legalistic sense, it might be seen as the transactional terms and conditions embedded in computer code which allow automatic execution of the relevant transaction once precise conformity with those terms and conditions has been established.

Whilst it may seem strange to separate out these two types of permissions (after all does it make sense to be a node in the system if one cannot add or view transactions?) there may be use cases when entries can only be added to a ledger following external approval (e.g., in a regulated environment), or where visibility of entries in the ledger is restricted (e.g., if commercially sensitive data is present).

Some of these use cases would stretch the bounds of the DLT definition, since it could be asked, if the ledger entries are not viewable by all participants, whether the system is truly a DLT? Nevertheless, it is considered important that all potential configurations are canvassed.

Using a single ledger for carbon markets?

One option would be to create a single distributed ledger that contains the emission units to be assigned, traded, and surrendered in all the ETS markets that will be involved in the network. This has the benefit of providing a single ledger for all trades, meaning that emission units traded between markets can easily be tracked and audited. However, the requirements to support different operational modes and timescale for individual ETS markets (and their respective emission units) make such a system unwieldy.

A single level system would require multiple regulator (or administrator) nodes in the system, with all able to issue and accept surrender of emission units. They would also potentially have the ability to authorise or block transactions of emission units issued by them or traded with organisations being regulated by them, which has the potential to either significantly reduce the responsiveness of such a DLT/market, or make the system very complex.

More critically, how different administrator nodes assigned emission units to compliance entities, the differences in the respective mitigation values between ETS markets, and the fact that each jurisdiction would need to continue maintaining and operating its own domestic registry (that is, its part of the ledger), make a single market solution very hard to conceive. The translation of value between emission units issued by different administrator nodes but in the same ledger would require significant functionality in the systems surrounding the DLTs.

Finally, having a single level DLT would make it hard for ETS markets to join or leave the network. It would require full migration of operations into the DLT for a market prior to joining, and full withdrawal of operations from the DLT into something different when an individual ETS wanted to leave the network. Therefore, whilst a single level DLT provides the requisite functionality for a market to operate, it is not compatible with the evolution of that market, and hence a broader view of the relationship between design and user adoption is required.

Designing for Diffusion

To have meaningful impact on practice, any new solution has to be adopted by user groups that extract value, share knowledge about use, invest in support, and thereby help to accelerate adoption by others. This enables the solution to diffuse across its intended markets, growing in scale and added value. Having designed a DLT for a specific set of Use Cases, it should be recognised that new users will bring new contexts and modalities that might not be predicted, but do need to be anticipated if the system is to achieve its design objectives.

The study of the diffusion of technology is a well established field and relevant to the design of a DLT-based network infrastructure. The configurable design criteria discussed above that are offered as a justification for rejecting a single level DLT design include: complexity, responsiveness, and flexibility. These qualities are known to impact the adoption and diffusion of any new technology, and were first expressed by Everett Rogers[8] in 1962 as five characteristics associated with the speed of adoption: *relative advantage* over existing technologies; *observability* of those benefits; *compatibility* with existing needs and use patterns; relative *complexity* of the innovation and the degree to which it has *trialability*, where availability for experimentation is positively correlated with adoption speed.

Considering the configurable design criteria of *Permission* it is reasonable to posit that each of the qualities associated with adoption and diffusion may have some dependence on degree of permission (Table 3):

---

[8] Rogers, E. M. (1962) Diffusion of Innovations. (1st ed.) New York: The Free Press of Glencoe.

Table 3: Potential relationships between DLT design and Diffusion: Focus on Permission.

Diffusion Characteristic	**Potential relationship with Degree of Permission**
Relative Advantage	Unpermissioned DLTs have applications, such as cryptocurrencies that have no flat currency analogue, and hence a significant relative advantage in that application area. Fully permissioned DLTs may operate with a central authority and offer an alternative to traditional systems where the locus of relative advantage is more constrained. *Relative advantage might therefore be seen as inversely correlated with degree of permission, but highly market specific.*
Observability and Trialability	Unpermissioned DLTs may be used by anyone without restriction whilst fully permissioned DLTs require user identification and authentication that defines the type of access that is available to the user. To this extent, a potential user may experiment with the unpermissioned DLT and hence directly observe the purported advantages. Fully permissioned DLTs however require user engagement before full experimentation can occur. *Observability and Trialabilty might therefore be seen as inversely correlated with degree of permission.*
Complexity	Complexity: the complexity that is 'observable and trialable' by the user can be constrained by the degree of permission: strongly permissioned systems define the user access rights and hence constrain the range of features available to a user. Unpermissioned systems in contrast, not only require technical expertise but also significant computing resources to engage with fully, representing the difference between using a cryptocurrency versus cyptocurrency ('bitcoin') mining. *Complexity might therefore be seen as inversely correlated with degree of permission.*
Compatibility	Compatibility: in markets where applications for DLTs are seen as improvements over legacy applications, the match between existing user needs and DLTs is easier to establish for fully permissioned systems. In terms of impact on diffusion, compatibility will be strongly market specific – whilst it is possible that an unpermissioned DLT will diffuse quickly across a global user base, so too will a fully permissioned system that is a clear upgrade path from existing systems for an established user base. *Compatibility might therefore be seen as positively correlated with degree of permission but highly market specific.*

The associations above highlight the importance of the link between design decisions and the potential for a technology to diffuse (Diffusion Potential), being necessary but not sufficient pre-conditions. The associations also establish links to other literatures on the diffusion of information and communications technologies, of which two are highlighted here: (i) Open Systems development and Open Innovation where, for example, the degree of Permission is associated both with 'Openness' and potential for others to innovate at the 'edge' of these technologies; and (ii) that diffusion is generally modelled as a nonlinear process, with the generalised Bass diffusion model (1969)[9] used to explain Roger's empirical observations across a range of technologies. If it is assumed that a contestable market opportunity exists – as no new product is likely to diffuse if it cannot demonstrate relative advantage and/or compatibility with user needs – then Diffusion Potential has an exponential relationship with Degree of Permission. This is illustrated in Figure 1, along with the assumed inverse relationship between Degree of Permission and Openness.

[9] Bass, Frank M., A NEW PRODUCT GROWTH FOR MODEL CONSUMER DURABLES, January 1969, Management Science, Volume 15, Number 5, pp. 215-227.

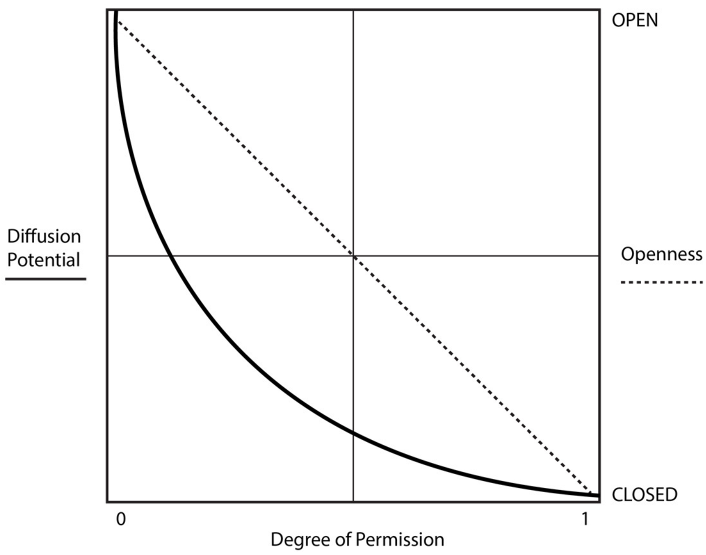

Figure 1: Diffusion potential as a function of Degree of Permission (see Table 3.)

Whilst it is clear from the design considerations applying to a single DLT for carbon markets that design needs to support the evolution of a market infrastructure, 'designing for diffusion' illustrates just how nonlinear the diffusion process is, and hence the importance of incorporating this into the design of any potential solution that will meet the timetable defined by the Paris Agreement.

To better understand the barriers to diffusion, a more finely grained look at market structure and in particular User and Stakeholder groups is required.

DLT/Carbon-Matrix

The DLT/Carbon Matrix (Table 4) incorporates the linear and nonlinear elements of design. Applications at the level of corporate accounting and internal transactions over different ETS will form the basic use case for any further implementation at the next level. The first level already carries the essential components for the following levels and diffusion requirement. If the functional requirements of each use case are not met then there is effectively no solution to diffuse, but as design decisions move from the left to the right, the network effects and diffusion impacts become stronger. These decisions are explored in the following text.

Table 4: DLT/Carbon Matrix

DLT/Carbon-Matrix	Ledger Use Case	Transaction Mechanism	Permissioning
Corporate Level	1.	2.	3.
Single ETS	1.	2.	3.
NCM	1.	2.	3.

Ledgers (Matrix 1.)

What system can we have if we don't have a single ledger?

The alternative to a single level structure is to have multiple different ledgers for the different ETS markets or modes of operation required. This would enable separate ETS markets to run their own ledgers, manage their internal market operations as required by local laws and regulations, but inter-connect with other markets to undertake networked trading.

The disadvantage of a multi-level approach is loss of the ability to have a single emission unit that can be used in all markets. However, as already discussed, the single emission unit across all markets introduces a significant amount of complexity into the ledger and its functionality, so separating out markets and allowing them to define and regulate emission units in their markets as the authorities see fit removes this complexity. Even more so, the idea of a single emission unit capable of being used across all participating jurisdictions runs counter to jurisdictions preserving their sovereignty and maintaining the flexibility to join or leave the network as they see fit in light of their own circumstances. This does not exclude the possibility of introducing a unit of exchange, to facilitate the transaction mechanism, however this subject is not considered further at this time (it will be the subject of a further sub-paper, by the authors).

The challenge with a multi-level approach is how to enable the networking of distinctive markets, the trading between ETSs, which a single level design would facilitate. The proposed solution is to use multiple levels of DLT to link together different markets or operational environments, that is: **Federated Distributed Ledgers**. Such a system could be used to create a distributed ledger (i) within a single organisation to manage its own carbon trading data, (ii) to operate a single carbon market, and (iii) to create a network between carbon markets. The same technology is used to operate different markets with different requirements, but can interoperate and allow movement of data between the different levels, as outlined in Figure 2.

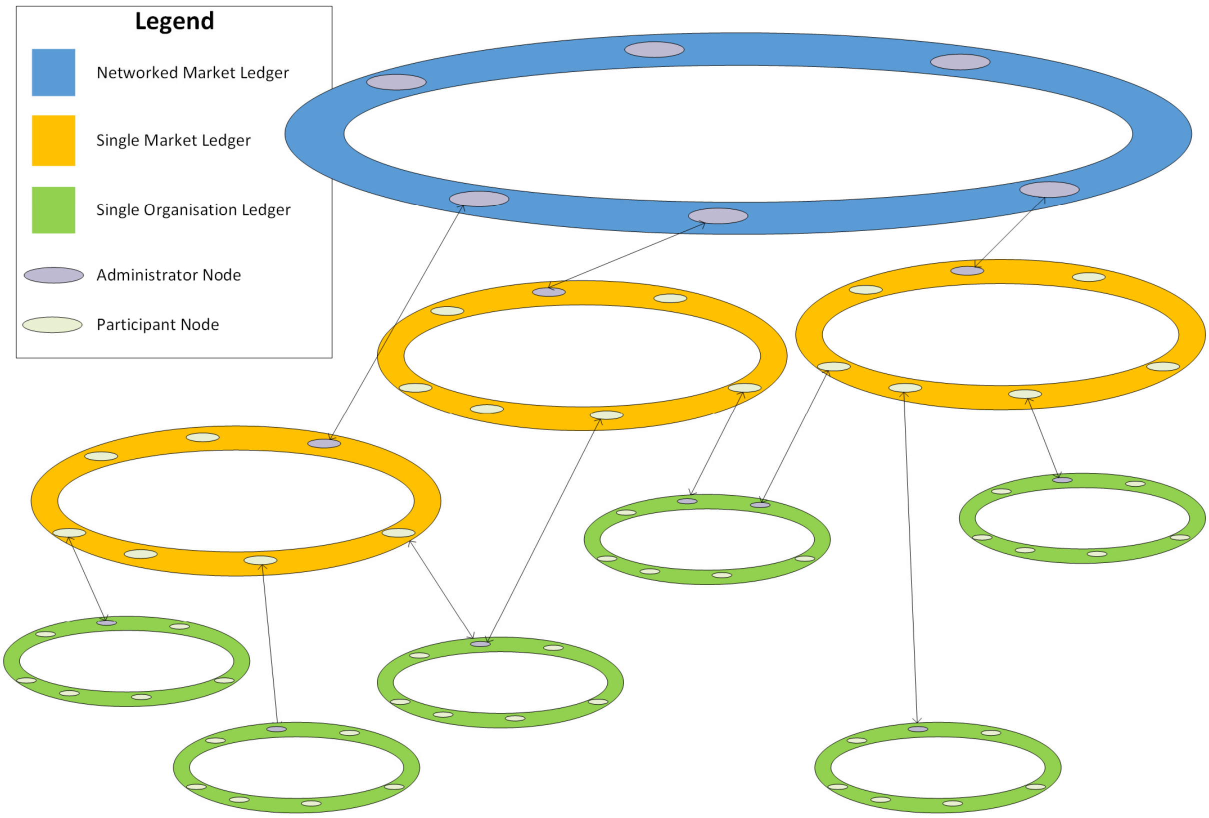

Figure 2: Multi-level Distributed Ledger for Networked Markets

Connection of ledgers is not a new concept. Indeed, side chains[10], inter ledger protocol and connection technology, such as that used by Ripple[11], are actively being developed, tested, and used. We are proposing using an approach broadly similar to the *side chain* model, using intermediate ledgers (also known as *side chains*) to undertake the federation of our primary market ledgers, enabling many primary ledgers (or *parent chains*) to trade/transfer resources.

The Network Market Ledger is the connection (*side chain*) between Single Market Ledgers for trading between these separate markets. Likewise, the Single Organisation Ledger is the connection (*side chain*) between Single Market Ledgers for accounting and tracking of assets across the markets an organisation is involved in. However, we need to enable such functionality without large waiting periods or delays between operations.

[10] This paper has a good explanation of side chains and the two-way peg mechanism: https://blockstream.com/sidechains.pdf
[11] https://ripple.com/technology/

These intermediate ledgers provide a place to both undertake conversion of emission units between ETS markets, and ensure resources can be securely traded with both parties having visibilities of the resources to be traded and their availability. This can be achieved using two-way peg like mechanisms through the intermediate ledgers, although research is required in the safest way to integrate such mechanisms into DLTs without significantly impacting ledger update rates and ensuring secure integration with ledger divergence reconciliation mechanisms[12].

Also, as is evident in Figure 2, our concept requires different types of participants in the overall ledgers, with administrator nodes able to interoperate with ledgers at levels higher or lower in the hierarchy. However, this is in keeping with the general requirements, where the ability is required to have different types of participants in the ledger for different operations (e.g., the regulator or administrator who can create and allocate emission units).

The attractiveness of this multi-level approach is that the DLT functionality can be used independently of other operations, for instance, a company can adopt the solution to manage its carbon assets even if the local market is not using such a solution. A local market can be created using this solution without requiring organisations participating the market to use the solution for local accounting. An individual market can use this approach without networking with other markets, or even before other markets are available. Because the different levels are separated, organisations or markets can change their approach or technology without impacting other markets or other involvement in these systems, allowing the system to diffuse easily.

The goal is to create a set of software tools that are sufficiently configurable to create and participate in any of the levels outlined, but do not mandate any particular level for their use. However, as the figure is illustrative it does not fully represent the required distributed infrastructure. With an operational solution, each of the nodes holds a full copy of the ledger and communicates, in a peer-to-peer fashion, with other nodes in the system.

Transaction Mechanism (Matrix 2.)

Are Smart Contracts fit for purpose?

Any system designed for carbon market trading and administration must be flexible enough to enable a range of metadata to be stored with emission units, or a range of configurable smart contracts to be linked to individual emission units. This is necessary to ensure emission units can be identified as being from a particular ETS, as part of a particular market cycle, for instance, and then not used in future market cycles (unless for instance, the jurisdiction allows banking, in which case this would be built into the metadata), or it may be required to ensure emission unit conversion rates and prices can be decided and applied automatically in trades.

However, direct inclusion of smart contracts in the emission unit has been highlighted as problematic for current distributed ledgers (see The DAO hack[13]). Having fully functional smart contacts that are part of the emission unit itself, places strict requirements on the code implementing those contracts to be correct and secure, so there is no way to change the contract once it has been added to the ledger.

As contracts could be part of the emission unit, then this could be assumed to be a risk of participating in the system, with any mistakes within contracts becoming the risk of the issuing or participating organisation. However, even with this assumption, smart contracts that can modify their own state can still lead to issues with contract changes whilst those contracts are being fulfilled (it is possible to write contracts that can be updated by the owner, whilst the other party is fulfilling them, thus changing the contract rewards or conditions from those initially seen by the other party). This is especially problematic on distributed systems, such as DLTs, as ordering of transactions cannot be easily defined globally (although, of course, accepted transactions have a defined order).

[12] Because of the distributed nature of the ledger, different nodes in the system can have different versions of the ledger, which need to be reconciled at some point using the consensus mechanism. The use of interconnected ledgers may complicate this reconciliation, some further work is required to ensure it can be achieved without rolling back resources committed to remote ledgers or received from remote ledgers.

[13] https://blog.ethereum.org/2016/06/17/critical-update-re-dao-vulnerability/

Part of the work that needs to be done to enable NCM with DLT will be to design a safe way of executing contracts on the ledger. Indeed, to ensure emission units can be transferred between ledgers, automatic locking and releasing of emission units on market ledgers would be desirable, so participants from other markets can have confidence that double spending of emission units is not occurring.

How to enable consensus without mining?

Technologies like BitCoin or Ethereum use mining as a mechanism to ensure consensus can be achieved without any one entity being in control of the network and without participants being able to subvert the ledger. However, it has a significant computational overhead and therefore energy requirements that make it unsuitable as a technology that needs to scale globally to help mitigate climate change.

As this paper has outlined the need for a permissioned distributed ledger, it is advocating a system that does not undertake mining to create the emission units added to the ledger. However, there would still need to be a mechanism for ensuring consensus across the copies of the distributed ledger. Mining is used to enforce some level of randomness over the choice of which node in the system (which copy of the distributed ledger) gets to add its update to the ledger at any given point in the operation of the ledger. Without mining, there needs to be a method to provide consensus choice and ensure that no one node can hijack the ledger and maliciously alter it.

Permissions (Matrix 3.)

How can permissions be distributed?

This paper advocates a permissioned ledger. Indeed, one of the strengths of the system design proposed should be the ability to have different functionality, different roles, for individual nodes, or participants, in the system, but not through different software, purely through different permission levels outlined in the ledger. To be able to implement a permissioned ledger system and assign different roles to different participants, this paper is proposing the combining of two ledgers in our infrastructure; one ledger recording emission unit transactions, the other ledger storing permission transactions, using an Authorisation Unit (AU) to record permissions of organisations or individuals in the system.

Each node participating in the DLT will have a full copy of both ledgers, but what they can do with the ledgers, and how much data they can see, will depend on their entries in the permission ledger, their AU. Without an AU, they will only have basic access to the ledger. This may be read access to all the emission unit transactions, or may be simply being a passive node in the network, depending on how the DLT is configured.

The proposed solution may appear counter-intuitive for two reasons:

(i) A node may hold the whole ledger but not be able to read any of it! This functionality is enforceable using public key encryption and provided the relevant private keys are not compromised, varying levels of access to decryption of the ledger can still be implemented.
(ii) Who has authority to add entries to the authorisation ledger and how would it be initiated? For an unpermissioned ledger this would be unworkable, but in the proposed use cases there are market regulators and, as such, organisations in the ledger system that will have permissioning rights. These entities will be able to establish the market or the system, and they will be able to grant the authorisations required. They may also be able and willing to allow the market to evolve, perhaps by devolving some authorisation functionality to other participants in the ledger, for example a company may be able to set up more nodes in the system once they have their initial access granted.

Conclusions

This paper has considered the design criteria for a networked carbon market and argued that, whilst there are many technology alternatives, there is a very good match between the capabilities of distributed ledger technologies and the creation of a network of carbon markets that will help deliver the emissions reduction targets of the Paris Agreement.

Exploring use cases and DLT design flexibility indicates that design not only needs to be functionally capable to address the requirements of a viable market infrastructure, it also needs to address a complex and dynamic range of stakeholder requirements, so that the solutions can diffuse quickly to build the required scale on a timescale that is compatible with the deadlines for those Paris Agreement targets.

The paper proposes a solution based on two complementary permissioned ledgers with the flexibility to define user roles and privileges that will deliver the required functionality across different organisational and regulatory environments, based on a concept of Federated Distributed Ledgers. The solution proposed offers advantages such as reduced transaction costs by (a) making the transactions faster; and (b) eliminating intermediaries that both slow the transaction and charge fees. In addition, the intended design should afford potential participants ease of adoption with the objective of making the software (and, if applicable, hardware) as accessible as possible, so as to ensure a level playing field. Participants other than compliance entities will engage with the market also if there are opportunities for arbitrage. Therefore, it is incumbent to design a platform that facilitates normal market activity.

The paper suggests that the DLT technology is already mature enough to deliver a robust system, but also points to potential barriers to adoption and diffusion that require additional research.