BRL Working Paper Series No. 12

Blockchain for Supply Chain: From Promise to Practice

Elias Strehle [1,*]

[1] Blockchain Research Lab, Colonnaden 72, 22303 Hamburg, Germany
[*] Correspondence: strehle@blockchainresearchlab.org

Published: 15 Jun 2020

Abstract: The potential of blockchain for the supply chain sector is widely acknowledged, but it is not always clear how this potential can be realized in specific use cases. Supply chain blockchains must fit into the wider context of complex business relationships, existing technical systems, and ever-changing requirements. Consequently, their design requires a holistic approach which appreciates the complex interplay between the technology and its context. This paper presents a corresponding framework for analyzing and designing blockchain systems in the supply chain sector. It outlines the benefits of blockchain technology, provides guidance on deriving requirements from the use case, and distills critical implementation features.

Keywords: Blockchain; Supply Chain

1 Introduction

During its early years, blockchain was often equated with cryptocurrency. Over time, however, the technology outgrew its first application. By now, the potential of blockchain as a source of trust and transparency is being explored for a multitude of use cases by a plethora of organizations, from tech start-ups to corporate consortia to public institutions.

The supply chain sector in particular appears to be a promising field for a technology built to operate in an environment of asymmetric information and complex incentive structures. Today's supply chains are characterized by a peculiar mixture of competition and antagonism on the one hand and cooperation and coordination on the other (Hoyt and Huq, 2000). Manufacturers and their suppliers are opponents in price negotiations but have the common interest that the final product can be sold at a competitive price. Manufacturers compete for the same customers but need to cooperate in the development of technical standards. Suppliers are torn between securing demand through tight integration with a single manufacturer and preserving a sufficient degree of independence.

Blockchain may be the missing ingredient to reproduce—and improve—this mixture of competition and cooperation in the digital world. It is designed to be operated collectively by many parties without any party relinquishing control. Pioneer projects are already demonstrating the viability of blockchain for supply chain use cases. Supply chain tracking, in particular, is emerging as a dominant use case (Rauchs et al., 2019, Figure 15).

Academic research on blockchain and supply chain is still in its infancy, with the vast majority of papers published during the last three years. One group of papers approaches the topic from a high level, including an early exploratory study among Finnish companies on the potential of blockchain for supply chains (Korpela et al., 2017) and extensive literature reviews by Wang et al. (2019) and Kshetri (2018). The latter also lists existing industry projects, as do Hewett et al. (2019) and Al Muhairi

et al. (2020). A second group of papers focuses on individual aspects or use cases. This includes a series of whitepapers by the World Economic Forum, each addressing a different aspect in the context of blockchain and supply chain: Digital identity (Jensen and Hewett, 2019), public and private blockchains (Hanebeck et al., 2019), data confidentiality (Flanagan et al., 2019), cybersecurity (Ogée et al., 2019), and interoperability (Pawczuk et al., 2020). Among papers on individual use cases, tracing—food tracing in particular—has received considerable attention. Kim and Laskowski (2018) propose an ontology data model for blockchain-based tracing systems. For a comprehensive literature review on food tracing in general see Section 2.2 in Hinkes and Peter (2020), on blockchain-based food tracing Tripoli and Schmidhuber (2018) and Pearson et al. (2019). Casino et al. (2019b) present a technical case study of a food tracing system based on Ethereum smart contracts. Specific traceability use cases discussed in the academic literature are provenance tracing in the Chinese dairy supply chain (Tian, 2018) and a tracking system for nuclear materials (Gasser, 2019).

What appears to be missing in the existing literature is a link between the generalist papers and the specialist papers. The broad interest in blockchain for supply chains is well-documented and the general advantages are frequently pointed out. At the other end of the spectrum, papers provide detailed accounts of individual aspects and selected use cases. There is little discussion, however, on how to bring the potential of blockchain to life within a given setting.

This paper therefore aims to contribute to the existing literature with a meso-level approach which is positioned between macro-level studies of blockchain technology and micro-level descriptions of selected use cases. It presents a framework for assessing potential blockchain use cases in the supply chain sector. The framework attempts to capture the complex relationship between the opportunities and requirements of supply chains on the one hand and the benefits and challenges of blockchain technology on the other.[1] This is particularly important because some features of blockchain are beneficial within the context of a public and open system such as Bitcoin but are a double-edged sword in supply chain use cases. Transparency and decentralization can make supply chains more efficient and predictable, but they can also threaten the competitive advantage of involved parties or collide with non-disclosure agreements and data protection laws. In this regard, the presented framework can help determine how blockchain systems in the supply chain fit into the wider context of existing systems, complex business relationships, and ever-changing requirements.

2 What is blockchain?

A blockchain is a distributed database.[2] It enables multiple parties to reach consensus on a set of records—the ledger—which they collectively create, maintain, and update. The records are made persistent by replicating them across multiple nodes (Rauchs et al., 2018, p. 24).

In comparison to other distributed databases, a blockchain is uniquely endowed to "understand" database transactions and the economic transactions they represent. It has a rich notion of validity that goes far beyond the type and uniqueness constraints of other databases (Greenspan, 2015b). Transactions are cryptographically signed by their originator, which allows an intuitive representation of economic ownership within the database. In addition, a blockchain can ensure that transactions follow economic rules (e.g. that they include a transaction fee) and can even initiate addi-

[1] This paper excludes purely financial use cases such as international cryptocurrency transfers or tokenized bonds. While these may also be beneficial to supply chain enterprises, the opportunities and challenges they present are not specific to the sector. For an overview of blockchain-based finance see e.g. Section 5.1 in Casino et al. (2019a).

[2] Some authors use the term distributed ledger technology (DLT) instead of blockchain, or define blockchain as a subset of DLTs. This paper makes no distinction between the concepts and uses the more familiar term blockchain.

Minimal requirements for blockchain use cases:

- Multiple parties need to reach consensus on shared data
- Transactions need to be checked for validity on a technical and economic level
- The use case is clearly delineated and well understood

Figure 1: Minimal requirements for blockchain use cases.

tional programmatically-executed transactions—known as smart contracts—based on complex business logic (e.g. paying commission to a party that facilitated the transaction).

A blockchain's deep understanding of data is supplemented by two additional capabilities. First, each party with a full view of the ledger is able to independently validate transactions and the integrity of the system. Second, the system resists attempts to unilaterally change past data, and each attempt produces tamper evidence which is easily detected by others (Rauchs et al., 2018, p. 24).

In other words, a blockchain enables its users to collectively define and enforce a comprehensive set of technical and economic rules for transactions, leading to strong security and consistency guarantees. However, these tend to come at the cost of efficiency, confidentiality, and flexibility (Strehle, 2020). Blockchain is therefore not an all-purpose tool. Other technologies specialize on one concern but aim to provide it for a wide range of applications. A classical data warehouse, for instance, is built around the single concern of storing data but aims to address this concern for many applications at the same time. A blockchain, by contrast, addresses many concerns at once—storing and sharing data, validating and initiating transactions—but is most useful when tailored to a single, well-defined application.

These general properties of blockchain imply a number of necessary, though not sufficient, conditions which a supply chain use case should satisfy to benefit from a blockchain (cf. Figure 1). First, the use case should involve multiple parties that need to reach consensus on shared data. Second, it should benefit from the blockchain's rich notion of validity, for example because it is important to know who "owns" data, smart contracts are used to boost efficiency, or a high degree of tamper resistance is required. Third, there should be a clearly delineated and well understood use case which can be represented through the users, objects, protocols, and smart contracts of a blockchain system.

Supply chains are likely to yield a large number of use cases which satisfy these requirements. They involve many parties that must exchange data but cannot always assume that the data they receive is valid in a technical and economic sense.

3 Benefits of blockchain

It is important to identify how a given use case will benefit from a blockchain. Most of the benefits of a blockchain are not emergent properties—they do not simply materialize because a blockchain is used. Indeed, some of the benefits can quickly evaporate as a result of unfortunate design decisions. This section describes potential benefits of using a blockchain and describes the design decisions which impact them. Figure 2 provides a summary.

The benefits in this section should be seen as ideals which guide the development of a blockchain system. In supply chain use cases, they will likely characterize some, but not all, aspects of the system. For example, a system may decentralize the maintenance of shared data records but keep other parts of the system, such as the onboarding of new users, centralized.

Potential benefits of blockchain:

- Decentralization
  - Disintermediation
  - Dilution of economic power
  - Fluctuating membership
- Transparency
  - Consensus on data and business logic
  - Tamper evidence
- Security
  - Preventive and detective security measures
  - Security within the system
- Catalyst
  - Increased cooperation
  - Automation
  - Standardization
  - Natural fit for use case
  - Marketing device

Figure 2: Potential benefits of blockchain.

Decentralization

One of the most frequently discussed aspects of blockchain is its potential for decentralization. Strong security measures and the redundancy of system components make it possible to distribute control.

In the most extreme case, a blockchain system can enable complete *disintermediation*. While a disintermediated system still relies on a group of facilitators (such as the miners on Bitcoin), no single facilitator performs a critical role. In other words, every single member of the system can be removed without affecting the functionality of the system. Bitcoin is the most prominent example of a disintermediated blockchain system.

Decentralization can also be understood in a weaker sense as the *dilution of economic power*. Even if a single intermediary performs a critical function in the system, the users retain control over the data and the system specification. This allows users to foster competition for the intermediary's position. Consequently, the intermediary has less economic power and thus less power to charge high service fees.

In addition, decentralization implies that blockchains are prepared to handle *fluctuating membership*. They make it easy to keep a system usable and useful when new users join and old users leave. In a blockchain system, the nodes of new members are automatically synchronized. Existing members can typically leave the system without bringing the system to a halt. This means for instance that a blockchain system can outlive its founding members, which might be particularly interesting in supply chains with a high degree of fluctuation.

Transparency

Many aspects of blockchain make it easy to ensure a high level of transparency. A blockchain allows users to reach *consensus on data and business logic*. Full nodes can independently audit the system and guarantee that code in smart contracts is executed without interference. In addition, blockchains provide strong *tamper evidence*: members who violate the roles, by accident or on purpose, will usually leave behind traces which are visible to others.

Transparency, however, is a commitment. If confidentiality is a key concern, the benefits from transparency are likely to be limited. Within supply chains, transparency can therefore be a double-edged sword. It can increase supply chain integration and efficiency but may endanger the competitive advantage or collide with nondisclosure obligations of some parties.

Nonetheless, where transparency is pursued, blockchains make it uniquely easy to get a complete view of the system, its code and its data. Transparency can ensure that users "trust" the blockchain and through it the parties they interact with.

Security

To establish the security potential of blockchain, it is helpful to make two distinctions. First, there is security within the system and security against outside attacks. Second, one can categorize security measures as preventive, detective, and corrective. Preventive security measures discourage a violation before it occurs. Detective security measures identify a violation and provide information about it. Corrective security measures fix a violation after it occurred and return the system to a normal state.

The tamper evidence and tamper resistance of blockchain provides exceptional *preventive and detective security measures which ensure security within the system*. The rules on a blockchain are strict, difficult to break and almost impossible to break unnoticed. Use cases where security concerns of these kinds exist can benefit from the "trust machine" blockchain.

In a supply chain, such use cases can emerge as a result of trust problems between or within companies, but also in complex environments where it is difficult to exclude human or technical error.

Catalyst

In addition to its direct advantages in the realms of decentralization, transparency, and security, blockchains "can act as a powerful catalyst to [...] encourage entities to rethink existing infrastructure and business processes, and [...] solve political roadblocks in the creation of industry utilities and drive behavioural change in industries that are generally resistant to change" (Rauchs et al., 2019, p. 18).

Even where the transparency of a blockchain is not strictly necessary, it can provide enterprises with the confidence which is necessary for *increased cooperation* even when business interests are not fully aligned.

Like every other technical system, blockchain has the potential to increase the level of *automation*. Smart contracts in particular have the potential to efficiently handle business processes which occur often and can be represented as computer code. This can reduce the number of errors and save money. Blockchain systems are particularly suited to automate processes which involve multiple parties and would otherwise require a matching of information from various decoupled systems.

A close relative of automation is *standardization*. Creating a shared system requires a reconciliation of how business objects and processes are defined and implemented. As a system that enables

Design considerations for supply chain blockchains:

- Objects
  - Exogenous data and oracle problems
  - Level of object abstraction: class vs. instance database
  - Level of object granularity: batches, products, parts, orders, shipments, etc.
- Users
  - Permissions for user actions: validating, writing, reading
  - Transparency vs. confidentiality
  - Effective stakeholder transparency
  - End-to-end security and usability
  - Timeliness of information: asynchronous vs. synchronous

Figure 3: Design considerations for supply chain blockchains.

the codification of business logic and invites data sharing among multiple parties, it provides strong incentives to cooperate and agree on standardized data formats and processes.

In addition, for some use cases where other technologies would also be satisfactory, blockchain might simply be a *natural fit* (Greenspan, 2015b). It enables a robust peer-to-peer database with multiple writers, provides strong tools for detecting and handling conflicting transactions, has a strong concept of who owns what, and makes it possible to codify and enforce business logic directly in the database. Especially when the use case benefits from an integrated digital currency, it might be much easier to implement with a blockchain than with any other technology.[3]

Last and least, blockchain has also proved effective as a *marketing device* to call attention to necessary and useful changes. Rauchs et al. (2019) refer to cases where blockchain is only used for this reason as "Blockchain as an excuse (BaaE)" (p. 18).

4 Use case and requirements

Designing a blockchain system requires a thorough understanding of the use case it will serve. Few aspects of a blockchain system can be developed in isolation. A step-by-step approach—write scope document, select blockchain implementation, define data structures, implement business logic, run tests, go live, onboard users—is unlikely to be successful. Merz (2019) compares the process of designing a blockchain to repeatedly riding an elevator: From business processes down to technical details up to economic implications and down again.

The benefits of blockchain are brought to life in relation to a specific and well-defined use case. Mapping a use case to a comprehensive description of the blockchain system can be challenging. This section outlines some design considerations which are critical for supply chain blockchains. An overview can be found in Figure 3.

On the highest level of abstraction, the system consists of objects and users. The objects are represented by data records; they are modified through transactions by users and smart contracts. Mapping the real-world use case to a list of blockchain objects and users provides clarity on how the benefits

[3] Notice that it is in no way necessary for enterprises to create an actual cryptocurrency which replaces other means of payment. The currency tokens on their blockchain can simply be used as an accounting tool, e.g. to keep track of micropayments, and be offset against real-world currency on a regular basis.

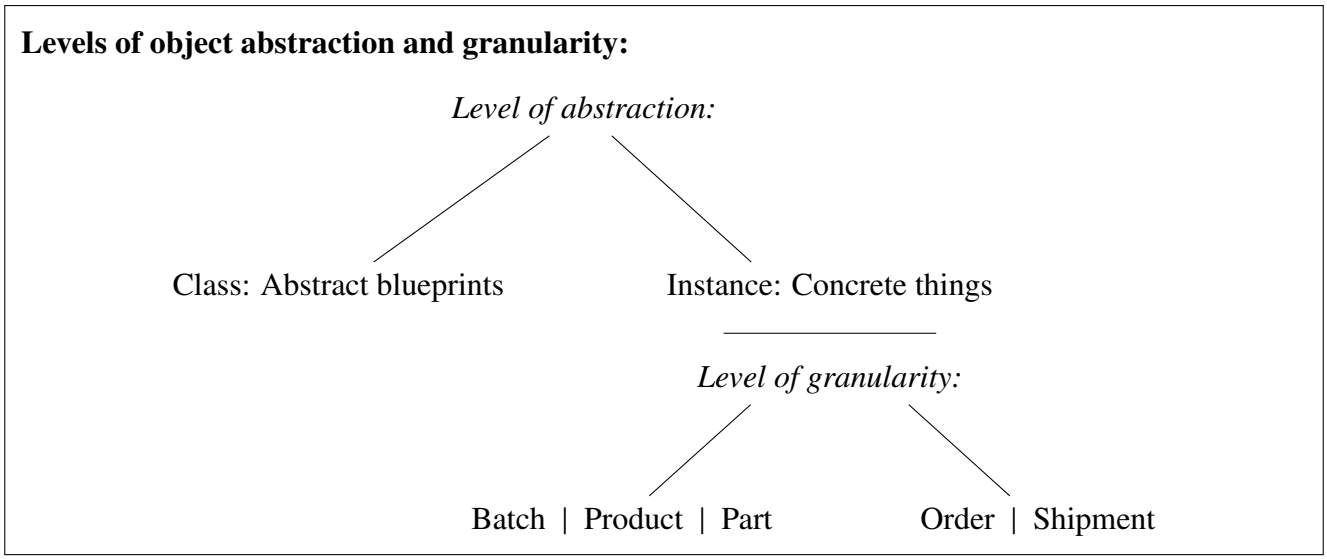

Figure 4: Levels of object abstraction and granularity.

of the system will materialize and which requirements it must satisfy.

Objects

The value of a blockchain system rarely derives directly from its objects. It derives from the attributes of the objects. A system is not useful because it references containers, it is useful because it stores the current location of each container. But one should not skip immediately to considering what attributes will be stored on the blockchain—location data, certificates of authenticity, prices, etc. Attributes pertain to objects, and it is critical to define the objects first.

A general challenge posed by supply chain use cases is that their blockchain objects reference *exogenous resources*, i.e. items or concepts which exist outside of the system (Rauchs et al., 2018, p. 48). This distinguishes them from early uses of blockchain technology, most notably cryptocurrency. The objects in early applications such as Bitcoin reference endogenous resources which have no independent existence outside of the system. Exogenous references expose a blockchain to the so-called oracle problem (Egberts, 2017). Oracles inform the blockchain system about the state of its resources and effect changes upon them. They can be simple sensors, complex machines, or people.

Erring input oracles can lead to a mismatch between reality and its representation in the system, erring output oracles can make the blockchain ineffectual. Blockchain designers should be aware of the oracle problem and anticipate its consequences. They must ensure that the system remains useful when oracles do not behave as expected. This is likely going to require a higher degree of flexibility than many blockchain technologies provide out of the box. As an example, consider a broken temperature sensor which wrongly reports that a shipment of food became too warm during transport and should be discarded. If the only problem was indeed the sensor, it must be possible to "tamper" with blockchain data and mark the shipment as okay.

One way to partially mitigate such problems is to introduce double checks, e.g. by using multiple sensors to measure the same fact or to compare input from automated devices with input from experts. Another way is to integrate a "reputation system" that gives less weight to devices which report nonsensical values. Williams and Peterson (2019) study incentive compatible reward systems for blockchain oracles.

An additional challenge for blockchain designers is to decide upon the *level of abstraction and granu-*

larity of objects, as illustrated in Figure 4. A useful distinction with respect to the level of abstraction is between *class* and *instance*. A class is an abstract concept which serves as a blueprint for concrete things. These things are called instances. For example, a vehicle model is a class, and each specific vehicle of that model is an instance of the class.[4] Both classes and instances can serve as references for blockchain objects.

Supply chain databases which reference classes are typically, and confusingly, referred to as product databases. An important subset of product databases are substance databases, which list some or all materials contained in consumer products. One example is the "International Material Data System" (IMDS), a comprehensive substance database for automobiles.[5] Another is the "database for Substances of Concern In articles as such, or in complex objects (Products)" (SCIP) which is currently being implemented at the mandate of the European Union.[6] Other types of product databases store product blueprints, operating and maintenance manuals, or answers to frequently asked questions. The latter has a proposed blockchain solution in the plastics industry (Licht et al., 2019). Product databases can also store sustainability data, e.g. for use in a Life Cycle Assessment (LCA).

Most proposals for blockchain systems, however, do not reference product classes but product instances. The former store general facts about abstract concepts, the latter make it possible to track differences between instances which belong to the same class. For example, a class database can store the average carbon dioxide emission associated with the production of a specific vehicle model; an instance database can store the exact carbon dioxide emission associated with each individual vehicle. Proposed use cases of blockchains with instance objects include tracing the provenance of food (e.g. Tripoli and Schmidhuber, 2018; Tian, 2018; Casino et al., 2019b), monitoring nuclear materials (Gasser, 2019), and controlling the temperature of sensitive trade items during transport (modum.io AG, 2017).

For instance databases, a further decision must be made with respect to the *level of granularity*. Does the blockchain track batches, products, or individual parts of the product? A second categorization, orthonogal to the first, is that of orders and shipments. Some use cases require the tracking of attributes that pertain to a collection of products, e.g. an order for five hundred computers. Which computer instances are used to fill the order is not relevant.

Determining the best resolution can be particularly challenging for uncountable goods like raw or bulk materials. What kind of blockchain object should store the temperature of milk in a tank truck—a truck object, a "truckload of milk" object, or multiple "liter of milk" objects? All approaches have their own advantages and disadvantages.

It is important to think about objects early in the blockchain design process because they define which attributes can, and must, be tracked. A class database cannot store the location of individual instances. At the same time, a higher level of granularity makes the system more complex. The aforementioned example of carbon tracking illustrates this. A class database requires a single emission value per vehicle model. An instance database requires one emission value per individual vehicle, which can only be realized with a high degree of automation and comprehensive coverage of all emission sources throughout the supply chain.

Class databases might therefore be the lower hanging fruit for supply chain blockchains. They tend to have low throughput since abstract product definitions are not updated very frequently. Their use cases also tend to involve relatively few users with similar skills and interests. Substance databases,

[4] Similar distinctions exist in various fields, for example the distinction between class and instantiation in object-oriented programming or the distinction between type and token in metaphysics (Wetzel, 2018).

[5] https://public.mdsystem.com/en/web/imds-public-pages/home.

[6] https://echa.europa.eu/de/scip-database.

for example, are maintained and used by engineering departments. In a blockchain solution, each engineering department can operate a blockchain node, mutual control can be assured and every user would be able to interpret tamper evidence or fix inconsistencies in the system.

For instance databases, the throughput is potentially much higher, especially when time varying features (such as the temperature of food during transport) must be tracked. The users are also likely to be more heterogeneous, for example when truck drivers and loaders need to enter information into the system.

Users

Most supply chain systems need to serve different stakeholders with different interests and varying levels of involvement and technical sophistication. Public blockchains like Bitcoin and Ethereum are open and permissionless, meaning that all users can view all data and perform all actions—validating, writing, and reading transactions. Most supply chain blockchains will need a more differentiated approach.

The task of identifying users can be approached from two directions. First, one can take a technical perspective and consider which user roles need permission to validate, write, and read some or all data on the blockchain. Second, one can take a business perspective and ask how the users will benefit from these permissions within the use case.

Validation serves two purposes. First, a blockchain needs a set of "full nodes" which validate all new transactions and therefore enforce the rules of the system. The full nodes ensure that only valid transactions are persisted in the ledger and report invalid transactions to the network. Consequently, they play a vital role in the system and are likely to require more attention and technical resources from their owners. For this reason, their operation is sometimes outsourced to a provider of "blockchain as a service." This may however be in conflict with the second purpose of validation: A full node is able to independently audit the system and ensure that it is in a valid state. In this respect, validation is the basis for trust. Users who are unwilling or unable to trust the system unless they can audit it on their own need to operate their own full node.

For blockchain designers, full nodes present a trade-off between transparency and confidentiality (Strehle, 2020). One can only validate what one can see. Techniques to hide data on a blockchain exist, e.g. by encrypting it or storing only its cryptographic hash, but these impact the blockchain's efficiency and typically weaken its security and consistency guarantees. Greenspan (2018) concludes: "In a general sense, the more information you want to hide on a blockchain, the heavier the computational burden you will pay to generate and verify transactions" (p. 62). In addition, the hidden part of the transaction cannot be validated. For example, if a user stores sensor data on the blockchain but encrypts the measured values, no business logic can be applied by validators to check that the values make sense. In the extreme case where all information on the blockchain is encrypted or hashed, the blockchain system becomes a timestamping service, a simple and occasionally useful tool which makes no use of the blockchain's concept of validity or the potential of smart contracts.

A key benefit of blockchain is its capacity to support multiple *writers* who submit new transactions to the system. The blockchain validators detect and resolve conflicts between transactions and ensure that all users agree on the same set of records. Since transactions are cryptographically signed, a blockchain makes it easy to represent economic ownership and ensure that only the owner of an asset can use it in a transaction.

As mentioned above, supply chain use cases may involve heterogeneous users, including some who are allowed to write new transactions but are not able to validate the system. Tracking applications

in particular may require write-only roles which can be provided to sensors, devices, or people with limited involvement in the blockchain system.

Some use cases may also benefit from including users who can only *read* selected data. One potential user group are output oracles which monitor the system state and become active when certain conditions are satisfied. For example, a barcode scanner may refuse to add an item to an order if the blockchain system has marked it as a counterfeit.

Another potential read-only group are people and organizations who are given access to the system in the interest of *stakeholder transparency*. These could be government authorities that use data on the blockchain to calculate taxes and duties, or customers who want to know where their food came from. In general, read access can be restricted to a subset of the information stored on the blockchain. One must be aware, however, that restricted read access is usually not sufficient to audit the system. This impacts the level of trust that it can provide. The blockchain itself does not create any trust—the ability to independently audit it does. A blockchain offers the tools to detect and report users who publish implausible or contradictory information. But only a validator can be sure that these tools are consistently applied. However, few supply chain systems will be able to give all stakeholders a full view of the data and business logic. In addition, it seems unlikely that every stakeholder who demands trustworthy information is willing to host and operate the full node required for validation. A more effective way of creating trust could be to rely on a trusted intermediary, such as an NGO, which is allowed to audit the entire system and provide selected information to other stakeholders.

Key challenges which come with a large number of heterogeneous users are *end-to-end security* and *usability*. It is important that the system provide both when some of the users are people with low technical proficiency and low involvement or automated devices. While blockchains tend to have a high level of internal security, issues like private key management are tedious and can quickly lead to security violations. Usability is also important to reduce the amount of incorrect or missing information, which can significantly reduce the power of blockchain by weakening the business logic that can be applied.

An aspect which is easy to miss from a technical perspective but shows itself clearly from a business perspective is the required *timeliness* of information. Here one can distinguish *asynchronous* and *synchronous* systems. In the context of instance databases, the former are roughly equivalent to tracing systems and the latter to (real-time) tracking systems. An asynchronous system does not guarantee—or even attempt—that information is written to the blockchain in real time. It can therefore observe the supply chain but is unable to intervene right away. That is, an asynchronous system registers data from the supply chain and allows the identification of deviations a posteriori, e.g. to trace back contaminated food from the supermarket to the farm. A synchronous system can potentially identify contaminated food and have it discarded before it arrives in the supermarket.

Synchronous systems may seem “better” but they require a much higher level of technical sophistication and supply chain integration. Events must immediately be recorded by devices or people. A high degree of interaction between the supply chain and the blockchain system can lead to unexpected feedback effects, attack vectors on the system, and high hardware cost. Even the most basic requirement of a synchronous system may not be satisfied in supply chain use cases: That a device must be connected to the internet when it needs to report data to the blockchain system. It is therefore important to clarify how much delay the system can tolerate and how the system will handle data that arrives too late or not at all.

Key considerations when selecting a blockchain implementation:

- Public vs. private
- Capability to define diverse assets with complex relationships
- Confidentiality
- Flexibility
- Advanced sharding capabilities
- Fine-grained user roles
- Transaction throughput

Figure 5: Key considerations when selecting a blockchain implementation.

5 Selecting a blockchain implementation

The expected benefits and requirements of the use case lay the foundation for selecting the most appropriate blockchain implementation. By now, a large number of implementations exists. They vary in technical focus, maturity, features, and usability. While this paper does not recommend one specific implementation, it outlines some requirements specific to supply chain use cases. Figure 5 summarizes them.

Public or private?

A fundamental decision is whether the system should use a public or a private blockchain. Public blockchains such as Bitcoin (Nakamoto, 2008) and Ethereum (Buterin, 2013) can be accessed and used by everyone. Private blockchains are specifically created for one application and can only be accessed by permissioned parties. They can be based on the same technology as public blockchains, but more often use a dedicated enterprise blockchain technology such as Hyperledger Fabric (Androulaki et al., 2018), MultiChain (Greenspan, 2015a), or Quorum (JPMorgan Chase, 2018).

Public blockchains can benefit from their large community and rich ecosystem. Instead of starting from scratch, designers who use a public blockchain can draw on already existing solutions for blockchain explorers, wallets, smart contracts, etc. The system itself is already live and secure, such that development and testing can begin immediately. This may be especially advantageous in a proof-of-concept phase or for small undertakings which lack the resources to set up an entire blockchain network on their own. Some use cases may also benefit from the cryptocurrency which comes with a public blockchain.

At the same time, public blockchains have downsides which can make deployment in a business context difficult or even impossible. One of the core problems is confidentiality. Every person with a computer and internet access can obtain a copy of all data on a public blockchain like Bitcoin or Ethereum, store it offline, and keep it forever. Even if data is encrypted, this poses a threat to the long-term security of data: Contemporary cryptographic algorithms are considered safe for at least the next years, but no legitimate forecast can be made that they will still be safe in ten or twenty years (Federal Office for Information Security, 2019, Section 6). This can pose a risk to companies and can even bring them in conflict with data protection laws.

One should also be aware that even if the content of transactions can be kept private, metadata analyses of the transactions can already provide valuable information. As an extreme example, suppose a company stores encrypted information about its products on a public blockchain and makes the

mistake of using a single blockchain address which creates one transaction per manufactured product. It is then trivial for external observers to derive how many products the company manufactured. Obfuscation is a countermeasure to metadata analysis—switch addresses often, batch data, publish fake transactions, et cetera. But each obfuscation technique makes the system more complicated and less flexible. While some degree of obfuscation may also be necessary on private blockchains, public blockchains can require so much of it that it turns from a technical detail into a major concern.

The lack of confidentiality and privacy on public blockchains can also make it more difficult for enterprises to retain a competitive advantage. On Ethereum, for example, the source code and the state history of all smart contracts are visible to everyone. But: "What makes a resource truly strategic—what gives it the capacity to be the basis for a sustained competitive advantage—is not ubiquity but scarcity. You only gain an edge over rivals by having or doing something that they can't have or do" (Carr, 2003, p. 6). If a smart contract is publicly visible and can be copied by everyone—who has an incentive to develop and maintain it?

A further disadvantage of public blockchains is the lack of control which the operators of a supply chain system have over them. They have almost no say in how the blockchain develops in the future, and the highly volatile transaction fees on popular blockchains like Bitcoin and Ethereum can pose a constant business risk.

In conclusion, public blockchains present blockchain designers with a less steep learning curve and lower initial investment, but can pose major obstacles in later development stages. They may however be a valid choice for systems where full transparency is unproblematic or even desired, particularly for "proof of existence" use cases (Hanebeck et al., 2019).

Critical features

Supply chain blockchains will typically cover different types of entities (e.g. different products) to which many things can happen (as opposed to a Bitcoin token which can only be minted and transferred). Often these entities will be interrelated, requiring mechanisms by which one thing can "become" another (e.g. metal ores being processed in a smelter) and how this affects their attributes. For example, a spoiled ingredient will fully spoil all meals that it went into, but the ingredient's carbon dioxide impact will be evenly split among the meals. It is therefore important that the blockchain implementation makes it easy to define *diverse assets with complex relationships*.

As pointed out above, most use cases will require a sufficient degree of *confidentiality*. Although not necessary, development can be greatly simplified if the blockchain system itself offers the possibility to share information privately with individual members. But recall that this inhibits the efficiency and security of the blockchain. Private data should be an additional feature, not the centerpiece of a blockchain system.

Another critical feature, deriving from the observations in the previous sections, is *flexibility*. Virtually all supply chain use cases are complex and will evolve over time. Furthermore, if the system covers exogenous resources, missing or wrong data are to be expected. It is important that the system anticipates this and is prepared to handle it in a way that keeps the system useful. For example, a tracked container might "disappear" in the system because the tracking device broke down. It might then "reappear" in another port once the device is fixed. While it is useful if the system detects and reports this, it must be possible to reinstate the container in the system without physically shipping it back to the port where the system lost track of it. Changes in business practice and regulation provide a further challenge for consistency, since a transaction might be valid on one date and invalid on another.

In addition, special cases are likely to occur. A company might want to ignore a mistake made by a business partner in the interest of maintaining a good relationship, or a court may order exceptional treatment of an entity (e.g. that an asset be handed back in spite of the rules of the system). This flexibility can be in contradiction to the tight security measures typically enforced by a blockchain.

While not a strict requirement, blockchains in a supply chain can benefit greatly from *sharding*. If the blockchain stores instance objects, its assets typically have a finite life. This means that entities can reach a final state such as "sold" or "discarded." Once an asset is in this state, it cannot switch to another state and it cannot be used in any new transactions. This provides a unique opportunity: Entities in their final state can be removed from the blockchain. They are no longer required for validating transactions and can be deleted or migrated into a more efficient long-term storage. This might make the problem of ever-increasing storage requirements and long-term data protection easier to solve than for cryptocurrency, where every token lives forever and must therefore be stored on the blockchain indefinitely.

Other features will be critical for some use cases but less important for others. These include the ability to define various *user roles* with fine-grained permissions and a high *transaction throughput*.

6 Conclusion

Designing a supply chain blockchain requires a holistic approach. In comparison to other technologies, an effective blockchain system is highly customized to the use case it serves. Hewett et al. (2019) summarize: "Blockchain is a team sport—it requires collaboration. By its very nature, blockchain and distributed ledger technology make transformation from an isolated approach to end-to-end value-chain integration within fragmented and complex systems more attainable" (p. 19).

This presents blockchain designers with the challenge of acquiring and uniting expertise from multiple departments and companies. Not all elevator rides that blockchain designers need to make are metaphorical. At the same time, this inter-departmental and inter-company approach has the potential to lead to truly useful and seamless solutions which are able to unite all aspects of a business use case.

The framework in this paper provides a scaffolding for analyzing and designing supply chain blockchains. It attempts to cover the key aspects characterizing supply chain use cases and the interdependency between business requirements, technological potential, and economic trade-offs. The expected benefits from using blockchain should be determined early and reassessed throughout the design process. The use case dictates the requirements for the blockchain system, and the selected blockchain implementation influences how easily the use case can be represented.

A well-designed blockchain has the potential to cut through a supply chain's complexity. It can make supply chains more efficient, more transparent, and foster cooperation where it truly matters—for innovation, for social issues, for sustainability.

References

Al Muhairi, M., Termanowski, M., Balovnev, M., and Hewett, N. (2020). Inclusive deployment of blockchain: Case studies and learnings from the United Arab Emirates. Whitepaper, World Economic Forum. Retrieved May 29, 2020, from http://www3.weforum.org/docs/WEF_Inclusive_Deployment_of_Blockchain_Case_Studies_and_Learnings_from_the_United_Emirates.pdf.

Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., Manevich, Y., Muralidharan, S., Murthy, C., Nguyen, B., Sethi, M., Singh, G., Smith, K., Sorniotti, A., Stathakopoulou, C., Vukolić, M., Cocco, S. W., and Yellick, J. (2018). Hyperledger Fabric:

A distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference (EuroSys '18)*, pp. 1–15, Porto, PT. ACM Press. https://doi.org/10.1145/3190508.3190538.

Buterin, V. (2013). Ethereum whitepaper. Whitepaper, Ethereum Foundation. Retrieved June 3, 2020, from https://ethereum.org/whitepaper/.

Carr, N. G. (2003). IT doesn't matter. *Harvard Business Review*, (May 2003).

Casino, F., Dasaklis, T. K., and Patsakis, C. (2019a). A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and Informatics*, 36:55–81. https://doi.org/10.1016/j.tele.2018.11.006.

Casino, F., Kanakaris, V., Dasaklis, T. K., Moschuris, S., and Rachaniotis, N. P. (2019b). Modeling food supply chain traceability based on blockchain technology. *IFAC-PapersOnLine*, 52(13):2728–2733. https://doi.org/10.1016/j.ifacol.2019.11.620.

Egberts, A. (2017). The oracle problem: An analysis of how blockchain oracles undermine the advantages of decentralized ledger systems. Master's thesis, EBS Business School. https://doi.org/10.2139/ssrn.3382343.

Federal Office for Information Security (2019). Towards secure blockchains. Report, Federal Office for Information Security. Retrieved May 27, 2020, from https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Crypto/Secure_Blockchain.html.

Flanagan, A. J., Maclean, F., Sun, M., Hewett, N., and Liao, R. (2019). Inclusive deployment of blockchain for supply chains: Part 4 – Protecting your data. Whitepaper, World Economic Forum. Retrieved May 29, 2020, from http://www3.weforum.org/docs/WEF_Inclusive_Deployment_of_Blockchain_for_Supply_Chains_Part_4_Report.pdf.

Gasser, P. (2019). A distributed ledger technology (DLT) approach to monitoring UF6 cylinders: Lessons learned from TradeLens. In *Institute of Nuclear Materials Management 60th AnnualMeeting*, number LLNL-CONF-777061, Dalm Desert, CA, United States. Lawrence Livermore National Laboratory.

Greenspan, G. (2015a). MultiChain private blockchain: White paper. Whitepaper, Coin Sciences Ltd. Retrieved May 27, 2020, from https://www.multichain.com/download/MultiChain-White-Paper.pdf.

Greenspan, G. (2015b). Private blockchains are more than "just" shared databases. MultiChain Blog. Retrieved June 3, 2020, from https://www.multichain.com/blog/2015/10/private-blockchains-shared-databases/.

Greenspan, G. (2018). Where blockchains add real value. *Innovations: Technology, Governance, Globalization*, 12(1–2):58–69. https://doi.org/10.1162/inov_a_00267.

Hanebeck, H.-C., Hewett, N., and McKay, P. A. (2019). Inclusive deployment of blockchain for supply chains: Part 3 – Public or private blockchains – which one is right for you? Whitepaper, World Economic Forum. Retrieved May 29, 2020, from http://www3.weforum.org/docs/WEF_Inclusive_Deploymentof_Blockchain_for_Supply_Chains.pdf.

Hewett, N., Lehmacher, W., and Wang, Y. (2019). Inclusive deployment of blockchain for supply chains: Part 1 – Introduction. Whitepaper, World Economic Forum. Retrieved May 29, 2020, from http://www3.weforum.org/docs/WEF_Introduction_to_Blockchain_for_Supply_Chains.pdf.

Hinkes, C. and Peter, G. (2020). Traceability matters: A conceptual framework for deforestation-free supply chains applied to soy certification. *Sustainability Accounting, Management and Policy Journal*, ahead-of-print. https://doi.org/10.1108/sampj-04-2019-0145.

Hoyt, J. and Huq, F. (2000). From arms-length to collaborative relationships in the supply chain. *International Journal of Physical Distribution & Logistics Management*, 30(9):750–764. https://doi.org/10.1108/09600030010351453.

Jensen, H. H. and Hewett, N. (2019). Inclusive deployment of blockchain for supply chains: Part 2 – Trustworthy verification of digital identities. Whitepaper, World Economic Forum. Retrieved May 29, 2020, from http://www3.weforum.org/docs/WEF_Trustworthy_Verification_of_Digital_Identities_2019.pdf.

JPMorgan Chase (2018). Quorum whitepaper v0.2. Whitepaper, JPMorgan Chase. Retrieved May 27, 2020, from https://github.com/jpmorganchase/quorum/blob/master/docs/Quorum%20Whitepaper%20v0.2.pdf.

Kim, H. M. and Laskowski, M. (2018). Toward an ontology-driven blockchain design for supply-chain provenance. *Intelligent Systems in Accounting, Finance and Management*, 25(1):18–27. https://doi.org/10.1002/isaf.1424.

Korpela, K., Hallikas, J., and Dahlberg, T. (2017). Digital supply chain transformation toward blockchain integration. In *Proceedings of the 50th Hawaii International Conference on System Sciences (2017)*. Hawaii International Conference on System Sciences. https://doi.org/10.24251/hicss.2017.506.

Kshetri, N. (2018). Blockchain's roles in meeting key supply chain management objectives. *International Journal of Information Management*, 39:80–89. https://doi.org/10.1016/j.ijinfomgt.2017.12.005.

Licht, J., de Jong, T., Oudshoorn, K., and Pasotti, P. (2019). Circularise whitepaper. Whitepaper Version 2.3, Circularise. Retrieved May 27, 2020, from https://www.circularise.com/whitepaper.

Merz, M. (2019). *Blockchain for B2B Integration*. MM Publishing, first edition.

modum.io AG (2017). Modum whitepaper. Whitepaper V 1.0, modum.io AG. Retrieved May 27, 2020, from https://modum.io/sites/default/files/documents/2018-05/modum-whitepaper-v.-1.0.pdf.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Whitepaper. Retrieved May 27, 2020, from https://bitcoin.org/bitcoin.pdf.

Ogée, A., Furuya, S., and Hewett, N. (2019). Inclusive deployment of blockchain for supply chains: Part 5 – A framework for blockchain cybersecurity. Whitepaper, World Economic Forum. Retrieved May 29, 2020, from http://www3.weforum.org/docs/WEF_Inclusive_Deployment_of_Blockchain_for_Supply_Chains_Part_5.pdf.

Pawczuk, L., Nielsen, J. M., Kwan Hang, P. S., and Hewett, N. (2020). Inclusive deployment of blockchain for supply chains: Part 6 – A framework for blockchain interoperability. Whitepaper, World Economic Forum. Retrieved May 29, 2020, from http://www3.weforum.org/docs/WEF_A_Framework_for_Blockchain_Interoperability_2020.pdf.

Pearson, S., May, D., Leontidis, G., Swainson, M., Brewer, S., Bidaut, L., Frey, J. G., Parr, G., Maull, R., and Zisman, A. (2019). Are distributed ledger technologies the panacea for food traceability? *Global Food Security*, 20:145–149. https://doi.org/10.1016/j.gfs.2019.02.002.

Rauchs, M., Blandin, A., Bear, K., and McKeon, S. (2019). 2nd global enterprise blockchain benchmarking study. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3461765.

Rauchs, M., Glidden, A., Gordon, B., Pieters, G., Recanatini, M., Rostand, F., Vagneur, K., and Zhang, B. (2018). Distributed ledger technology systems: A conceptual framework. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3230013.

Strehle, E. (2020). A new perspective on enterprise blockchains. Working paper, Blockchain Research Lab. Forthcoming.

Tian, F. (2018). *An Information System for Food Safety Monitoring in Supply Chains Based on HACCP, Blockchain and Internet of Things*. PhD thesis, WU Vienna University of Economics and Business.

Tripoli, M. and Schmidhuber, J. (2018). Emerging opportunities for the application of blockchain in the agri-food industry. Report, Food and Agriculture Organization of the United Nations. Retrieved May 27, 2020, from http://www.fao.org/3/ca1335en/CA1335EN.pdf.

Wang, Y., Han, J. H., and Beynon-Davies, P. (2019). Understanding blockchain technology for future supply chains: A systematic literature review and research agenda. *Supply Chain Management: An International Journal*, 24(1):62–84. https://doi.org/10.1108/scm-03-2018-0148.

Wetzel, L. (2018). Types and tokens. In Zalta, E. N. (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2018 edition.

Williams, A. K. and Peterson, J. (2019). Decentralized common knowledge oracles. *Ledger*, 4. https://doi.org/10.5195/ledger.2019.166.

Declarations

Availability of data and materials
Not applicable.

Conflicts of interest
Not applicable.

Funding
Not applicable.

Acknowledgments
The author thanks Lennart Ante, Constantin Fischer, and Christopher Nigischer for reviewing an earlier version of the manuscript.

About the Blockchain Research Lab

The Blockchain Research Lab promotes independent science and research on blockchain technologies and the publication of the results in the form of scientific papers and contributions to conferences and other media. The BRL is a non-profit organization aiming, on the one hand, to further the general understanding of the blockchain technology and, on the other hand, to analyze the resulting challenges and opportunities as well as their socio-economic consequences.

www.blockchainresearchlab.org