
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Nikander, Pekka; Autiosalo, Juuso; Paavolainen, Santeri
Interledger for the Industrial Internet of Things

Published in:
2019 IEEE 17th International Conference on Industrial Informatics (INDIN)

DOI:
[10.1109/INDIN41052.2019.8972167](https://doi.org/10.1109/INDIN41052.2019.8972167)

Published: 30/01/2020

Document Version
Peer reviewed version

Please cite the original version:
Nikander, P., Autiosalo, J., & Paavolainen, S. (2020). Interledger for the Industrial Internet of Things. In *2019 IEEE 17th International Conference on Industrial Informatics (INDIN)* (pp. 908-915). (IEEE International Conference on Industrial Informatics). IEEE. <https://doi.org/10.1109/INDIN41052.2019.8972167>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Interledger for the Industrial Internet of Things

Pekka Nikander
School of Electrical Engineering
Aalto University, Finland
Email: pekka.nikander@aalto.fi

Juuso Autiosalo
School of Engineering
Aalto University, Finland
Email: juuso.autiosalo@aalto.fi

Santeri Paavolainen
School of Electrical Engineering
Aalto University, Finland
Email: santeri.paavolainen@aalto.fi

Abstract—The upsurge of Industrial Internet of Things is forcing industrial information systems to enable less hierarchical information flow. The connections between humans, devices, and their digital twins are growing in numbers, creating a need for new kind of security and trust solutions. To address these needs, industries are applying distributed ledger technologies, aka blockchains. A significant number of use cases have been studied in the sectors of logistics, energy markets, smart grid security, and food safety, with frequently reported benefits in transparency, reduced costs, and disintermediation. However, distributed ledger technologies have challenges with transaction throughput, latency, and resource requirements, which render the technology unusable in many cases, particularly with constrained Internet of Things devices.

To overcome these challenges within the Industrial Internet of Things, we suggest a set of *interledger approaches* that enable trusted information exchange across different ledgers and constrained devices. With these approaches, the technically most suitable ledger technology can be selected for each use case while simultaneously enjoying the benefits of the most widespread ledger implementations. We present state of the art for distributed ledger technologies to support the use of interledger approaches in industrial settings.

Index Terms—Distributed Ledger Technology; Blockchain; Internet-of-Things Devices; Industrial Internet

I. INTRODUCTION

The term *Industrial Internet* was introduced by General Electric in 2012 [1], originally denoting the the combination of intelligent machines and advanced analytics, supporting more intelligent design, operations, and maintenance. Today the meaning has shifted somewhat, mostly referring to the industries moving from closed intranet-based ICT architectures to more open, Internet-based ones, while deploying Internet of Things (IoT) technologies.

Within this movement, there is a growing interest towards utilizing distributed ledger technologies (DLTs), aka blockchains, to induce collaboration across companies. For example, in the energy sector alone, there are some 140 projects that utilize such technologies [2]. However, there are few if any successfully commercialized industrial deployments where multiple companies collaborate across a distributed ledger. This is, of course, partially explained by the immaturity of the technology in general. However, there are other concerns, the most often mentioned of which is companies keeping their data tightly controlled and being unwilling to share it, due to fearing their data being used against them or due to the difficulty of defining a price for the data. This may be attributed to the still poorly understood non- or *anti-rival*

nature of data [3], [4]. The most often cited technical concern is the trade-off between the ledger scalability in terms of users vs. transactions commit delay.

In this paper, we introduce the idea of *interledger transactions* for the Industrial Internet and especially for the Industrial IoT (IIoT), showing how such transactions be applied to address the scalability, data control, and resource consumption concerns. Furthermore, we surmise that the interledger approach could be also used to alleviate the concerns related to the immaturity of the technology, by allowing changes in the underlying technology choices to be evolved. We refer to this latter ability as *ledger agility*, alluding the structural similarity with *crypto agility* [5].

The rest of this paper is structured as follows. First, in Section II we provide background information on industrial internet, distributed ledgers and typical deployment models. In Section III a number of industrial use cases are described, followed by Section IV on technological challenges of using DLTs in the industrial internet context. A few potential solutions are discussed in Section V. Conclusions are presented in Section VI.

II. BACKGROUND

A. Industrial IoT Security

Industrial IoT (IIoT) is facing security concerns from both fundamentally vulnerable legacy applications and inherently insecure IoT devices. Firstly, industries have traditionally built their core functions to rely on internal networks that are isolated from the public Internet. Therefore, many of the industrial applications provide open access to sensitive information and machine control interfaces. Partly due to this, the Industrial Internet landscape is still strongly divided into silos that can not be easily made accessible from the public Internet.

Secondly, lower-end consumer-class IoT devices are often connected straight to the public Internet, for the sake of ease of installation and simplified usability. However, due to low security, huge farms of IoT devices have been harnessed into launching distributed attacks [6]. Hence, in the IIoT, the systems as wholes have to be protected from both the legacy vulnerabilities and the low-end-device issues.

These two concerns are now being brought together in a multi-level architecture style for cyber-physical systems, leveraging intelligent machines and their twins [7].

B. Digital Twins

Digital twins are becoming an integral part of the IIoT architecture. With roots in many different fields and under various terms [8], the term has had many definitions [9]. However, while details vary, the basic idea remains the same: a digital twin is a *virtual entity with a one-to-one relationship to a real world object*. Today, each twin is typically custom built to serve a use case and therefore they look very different from each other.

Ideally, a digital twin is used to convey all communication between a physical device and the external world. In other words, an ideal digital twin is a 24/7 one-stop-shop for the best available information about the physical counterpart it represents. From the organizational security point of view, sharing data about real world objects will be easier to control when the datasets are arranged to well defined digital twins. Furthermore, use of DLTs in industrial manufacturing context is proposed to be performed by a specific software component [10] that can be included in the digital twin.

C. Distributed Ledger Technologies

A distributed ledger is basically an *append-only database* maintained by a set of computing nodes that communicate over the Internet¹ and jointly agree on storing the transactions submitted to the database by the database users. In many ledgers, the transactions are first packaged into so-called *blocks*, which are then *chained* together. These systems are the proper *blockchains*, including e.g. Bitcoin [11] and Ethereum [12], the best known distributed ledgers. However, the word *blockchain* is colloquially often used to denote all distributed ledger technologies.

The nodes responsible for serializing transactions into the ledger are often denoted as *miners*. Many ledgers also implement a *cryptocurrency system*. In some ledgers, such as Bitcoin and Ripple [13], the main function of the database is to transfer virtual funds. Other ledgers, such as Ethereum and Hyperledger Fabric [14], also include a virtual computing machine (VM), which is used to define business logic.

However, in almost all cases, the main function of the ledger is to act as a *trusted, shared database*. It is most often maintained by a large set of miners, who usually do not fully trust each other.

There are basically two defining features for distributed ledgers:

- The ledger database is *distributed*. There are tens to thousands of copies of the database. The miners implement a jointly agreed *consensus protocol*.
- The transactions *cannot be changed or removed* once they have been entered. Transactions are *chained* to each other with cryptographic fingerprints.

D. Types of ledgers

Distributed ledger technologies (DLTs) may be divided into two broad categories: *permissionless* and *permissioned*. De-

¹In the case of private or consortium ledgers, the node may communicate over a local or protected network instead of the Internet.

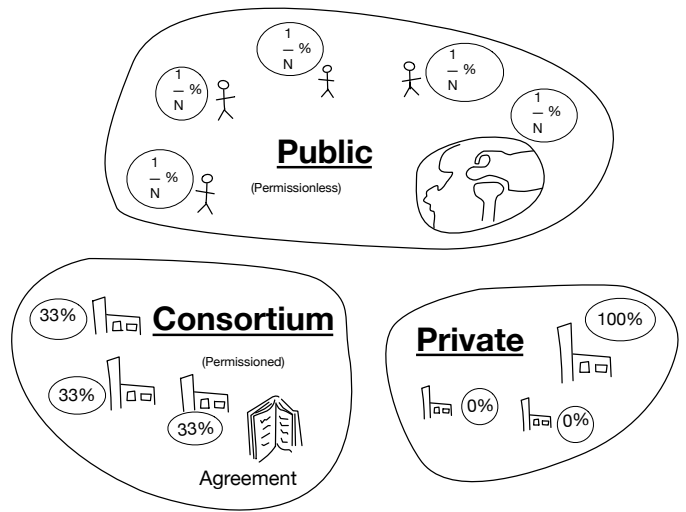


Figure 1. Governance differences between private, consortium and public ledgers. In private ledgers, the controlling authority is highly centralized even if the network itself is decentralized. In contrast, consortium ledgers have a limited set of authoritative nodes, but the overall authority is distributed without any single party having a majority rule. A public ledger is decentralized, with all parties having a fraction of the authority on the ledger (although not necessarily equally).

pending how a particular database is deployed, the installations may be further divided into *public*, *consortium*, and *private* ledgers.

In a *fully permissionless* DLT, anyone is able to join the DLT, taking up any role. Any node may opt to participate as a *validator*, as a *miner*, or as a *verifier*. In the majority of the blockchains there are no distinct validator and miner roles, but the miners are responsible for both functions. However, in some newer DLTs these roles are distinct.

In a *permissioned DLT*, the joining node needs to be authenticated and authorized to take up certain roles. For instance, a DLT could restrict the validator role to a predefined set of authorized nodes, but let anyone read the database. Other DLTs require authorization also for reading.

There is an essential difference between permissionless and permissioned ledger technologies: reaching consensus among a changing, dynamic group of untrusted strangers is fundamentally harder than reaching consensus within a well defined group. Roughly speaking, a permissionless and public ledger needs to be based on the so-called Nakamoto consensus [11]. These algorithms scale quite well in terms of the number of users, but they incur a large energy cost and are quite slow. In a permissioned or closed setting, on the other hand, redundant Byzantine consensus is sufficient while retaining resiliency against misbehaving parties [15]. Byzantine systems usually do not scale that well, but tend to be orders of magnitude faster and more energy efficient [16].

E. Deployment Scenarios: Public, Consortium, or Private

Considering the deployment scenarios, a fully *public ledger* is one where anyone can act in any role. By definition, they are based on some permissionless ledger technology,

such as Ethereum. A *consortium ledger* is one where certain roles are limited to the members of a consortium, which in turn consists of a number of distinct organizations or other entities. Technically, a consortium ledger may be built upon permissionless or permissioned technology. However, while many people build their initial concept systems on Ethereum, from the production point of view such practice doesn't make much sense, since the resulting closed Ethereum network either requires a lot of energy to safely commit the transactions or is prone to outside attacks by potential attackers with sufficient computing capacity. Hence, in general, it is better to base consortium ledgers on permissioned technology, such as Hyperledger Fabric. Finally, a *private ledger* is typically technically similar to a consortium ledger but enclosed inside a single organization [17].

We summarize some of the differences in Table I on the basis of ledger's throughput, latency, and resource requirements for some of the ledgers mentioned in the text. The data for each DLT is from the project's own website, documentation or wiki, unless otherwise specified.

III. INDUSTRIAL OPPORTUNITIES

In this Section, we cover a number of industry sectors, outlining the potential benefits. We first briefly present the existing literature, followed by a summary of potential benefits.

A. Examples of Industrial Use Cases

We consider the use of DLT in a number of industry sectors: logistics, energy markets, smart grid security, and food safety.

1) *Logistics*: Logistics of physical items involves integration of material and information flows together with various processes, and includes multiple stakeholders with different needs and goals. Typically, the stakeholders do not fully trust each other. According to Korpela et al. [18], many-to-many integration models, including open source DLTs, are most cost-effective in logistics, offering data security and cost-effective transactions. DLTs have been suggested to be used to track the origin and/or enhance the supply chain of e.g. pharmaceuticals [19], [20], manufacturing [21], diamonds [22], biorefining [23], ship building [23], container transport [24], and consumer delivery optimization [25]. In particular, Gallay et al. describe a peer-to-peer platform for decentralized logistics, where an IBM Bluemix [26] / Hyperledger Fabric² blockchain is used as a shared ledger, together with other components [23].

The potential benefits include enhanced transparency, consumer and small company empowerment, and improved sustainability and reduced environmental impact resulting from increased transparency and consumer choice [23], [25], [27]–[29].

²IBM Bluemix and Hyperledger Fabric are roughly speaking two different marketing names for the same technology.

2) *Energy Markets*: The energy sector is facing growing challenges in its transition to renewable energy sources (RES), as energy production is becoming less centralized. The grid needs to be updated to support reverse energy flows, and the increasing RES production implies increasing needs for electricity storage and real-time source capacity adjustments. As a result, there is a growing need to use digitalization to help balancing the supply and need.

Andoni et al. have provided so far the most comprehensive academic review on blockchains and DLT in the energy sector, with over 140 energy sector blockchain projects studied [2]. According to them, DLTs are considered to potentially 1) improve the efficiency, 2) accelerate the development, 3) provide innovation, and 4) improve current practises in the energy sector. DLTs could provide solutions across the so-called energy trilemma: reduce costs, improve security, and promote sustainability. DLTs could play a role in e.g. billing, trading, grid automation, grid management, identity management, sharing of resources, creating competition, increasing transparency, and creating data markets for third parties.

3) *Smart Grid Security*: Mylrea and Gourisetti [30], [31] note that the power grid lacks the necessary security and resilience measures to prevent cyber-attacks on distributed energy resources (DER), including solar panels and wind. This also extends to the control systems beyond the (smart) electricity meter.

Dong et al. propose a very generic DLT based smart grid cyber-physical infrastructure model which combines elements from IoT, cloud computing, and DLT [32]. They propose DLTs to be used for grid data protection, smart meter data aggregation, and direct load control. Interestingly, they give a conceptual example of a hierarchically organized ledger structure. In [33], the same research group report early simulation experiments on a benchmark system consisting of 54 generators, 118 nodes and 186 branches, basically measuring the probability of an attacker succeeding in confusing the system state as a function of the number of sensors the attacker needs to manipulate.

Ebrahimi et al. [34] report the findings from the First International Cyber Security Workshop in Smart Energy Systems, noting that DLTs could be used to enhance the self-defensive capabilities of the energy systems against cyber attacks. Kim and Huh [35] propose using a blockchain called Rainbowchain to the smart grid and energy exchange.

4) *Food Safety*: To our knowledge, Tian [36] was the first to propose using DLTs, together with RFIDs, to enhance food safety, followed soon by e.g. Ge et al. [37], Lin et al. [38], Tse et al. [39], Caro et al. [40], Galvez et al. [41], Kamath [42], and Yiannas [43]. The basic idea is simple enough: increasing trust through the immutability and transparency offered by a public or consortium ledger. Furthermore, most of the papers were at least partly motivated by the recent food scandals, including the 2006 melamine milk scandal in China [44] and the 2013 horse meat scandal in Europe [45], [46]. Some of the more recent works, such as Galvez et al. [41], focus more on the safety aspects, while many others are more focused on

the overall traceability.

From the present paper view, perhaps of most interest is the preliminary experimental work with Ethereum and Hyperledger Sawtooth by Caro et al. [40]. In particular, they show how “both implementations have different properties and capabilities that need to be considered”, referring to aspects such as latency, scalability in terms of network size, maturity, and operational costs.

B. Potential Benefits

Summarising the commonalities across the four industry sectors, there are basically three frequently mentioned potential benefits: **disintermediation**, **transparency**, and **reduced costs**. However, it should be noted that there appears to be a lot of uncertainties related to the perceived cost structures.

Disintermediation may be considered as a specialized form of decentralization, referring to the disappearance of intermediators or trusted third parties. DLTs are quite generally expected to lead to market structures where there is no trusted market maker, like in today’s digital platforms, but where the market is based on a DLT database recording the requests, bids, deals, and even the events related to contractual obligations, including payments. However, this process is perceived to take some time, largely due to the perceived immaturity of the technology and associated legal and regulatory barriers.

Another commonly cited benefit is *transparency*. Most DLTs allow all participating users to perceive and inspect all transactions entered to the database, thereby becoming aware of not only their own obligations and relationships, but also of those that their business partners have with third parties. Such increased visibility is likely to lead better understanding of the financial situation of all of the companies in the field, both by regulators and by the competitors. This, in turn, has the potential of leading to increasing trust and reversing the current erosion of social capital.

Finally, while many works note the uncertainty related to costs or even the higher operational costs of many DLTs compared to the traditional distributed databases, there are also strong expectations about *reduced cost* for transacting, contracting, enforcement, and compliance. These expectations are usually not directly perceived as immediate benefits from deploying DLTs, but from their longer term effects in terms of reducing the search and transaction costs in one hand and contract agreement and enforcement in the other hand, both due to the increased transparency. It is also expected that the costs related to compliance enforcement may be reduced, apparently again mostly due to the increased transparency.

However, we surmise that the increased transparency alone is unlikely to lead to the discussed potential benefits. Instead, there is a need for additional measures and technologies, such as regulation endorsing new governance models, introducing new compensation methods, and enabling smart behaviour by integrating ledgers with physical actuation and measurement using cryptographically strong, easily verifiable manners.

Table I
PROPERTIES OF SELECTED LEDGERS

DLT	Throughput	Latency
<i>Technologies for public or permissionless ledgers</i>		
Bitcoin	7 tps	>10 minutes
Ethereum	25 tps	>15 seconds
Ripple	50 000 tps	seconds
<i>Technologies for private / consortium or permissioned ledgers</i>		
Hyperledger Fabric	>1 000 tps	a few seconds
R3 Corda	15-2000+ tps	seconds–minutes

IV. TECHNOLOGY CHALLENGES

There are several challenges facing the widespread adoption of DLTs for industrial use. Some of these problems relate to the security, governance, performance, and reliability of distributed ledgers. A common theme currently is that these concerns are often solved by some, but not all, generally available DLTs. We may expect that as the technology matures, these concerns will be eventually addressed.

At present, the easiest use case for distributed ledgers is in the private, “internal” use by a single organization. Here, the organization makes the selection of which distributed ledgers to use. While there may be some concerns left, it is likely that most requirements can be met. However, if there is any need to interact with public or permissionless ledgers, the business needs, such as reaching the maximum potential customer audience with minimal investments, narrow the choice of available ledgers to only the few with the largest market share (for example, Bitcoin and Ethereum).

After the concerns such as the market share, the biggest technological problems facing an organization on the choice of a ledger are throughput, latencies, and resource requirements. For many industrial applications, throughput or latency may limit the applicability of the ledger, while resource requirements limit the capability of industrial devices to become active participants in the ledger. These are described in detail in the following sections.

A. Throughput

The common measurement of distributed ledger throughput is its *transactions per second* rate (TPS). It is also often misleading, as for an accurate comparison one would need to take the size of transactions into consideration. In ledgers that are used only for value transfers across accounts, a transaction is of fixed size, or varies typically with only a small range. However, in ledgers with more flexible evaluation models, e.g. incorporating smart contracts, the “size” of a transaction is more variable. For example, a fully utilized Ethereum block may contain anything from over three hundred transactions of minimal size to just a single, computationally “large” transaction³ — thus, the TPS for the public Ethereum network can vary between 0.07 and 25. This may be compared to the throughput of a value-transfer based public ledger,

³Based on the Ethereum network settings as of early March 2019.

such as Ripple, which claims to be able to attain a rate of 50 000 transactions per second [13].

While there are developments and technologies to increase the throughput of slower public distributed ledgers, such as payment channels [47], [48], and sharding [49], [50], one still must confront the restrictions that a low, or highly variable TPS rate of a distributed ledger imposes on industrial use cases. A smaller transaction rate can lead to higher contention and competition by users to complete transactions, leading to increased transaction costs and processing times. This effect has been observed multiple times in Ethereum, for example, when sudden interest in an ICO [51] or a game [52] has led to a shortage of the available transaction processing capacity.

Overall, if interaction with a lower throughput ledger is required, this restricts feasible industrial usage scenarios. Pragmatically, any use case that requires a high guarantee on timely transaction completion, or with a high rate of transactions, is practically eliminated.

B. Latency

The first, and economically still the largest distributed ledger, Bitcoin, has a *block time*, e.g. the average time between completed blocks and thus the fastest time a transaction can be included in the ledger, of 10 minutes. To guard against double-spend attacks, more than one block should be observed before a transaction should be considered as committed. The exact number of blocks to wait depends on risk factors, such as the value of the transaction and the capabilities of the assumed adversary [53]. The general trend, since Bitcoin, has been towards ledgers with lower latencies. For example, Ethereum's block time is 15 seconds on the average, and Ripple claims to have transaction commitment delay of a few seconds [13].

Thus, while the throughput of the distributed ledger may be sufficient for an use case, the latency may still prove to be an obstacle. If interaction with Bitcoin is required, and any latency less than half an hour is required, then one must instead look into the use of payment channels [47], [48]. Another complication is that when a multi-step operation is performed across different organizations using a smart contract, the total latency is multiplied by the number of distinct transaction steps that are required for the operation. Thus, if a use case requires multiple sequential transactions to be completed within a few tens of seconds, many distributed ledgers currently in widespread use are immediately ruled out.

C. Resource Requirements

Participating in a distributed ledger can impose heavy requirements on the available resources. These include, in varying degree on different ledgers, requirements on the CPU, memory, storage and network. For any kind of industrial device either completely or partially dependent on battery power, the cumulative effect of CPU, storage and network use on the battery life is also an important consideration.

Note that in all cases discussed here, we assume that the resource requirements apply to a node that is not an active participant in the consensus algorithm (e.g. not a miner, which usually is the most resource-intensive role on a DLT).

1) *CPU*: A node that validates all of the state changes on the ledger does require processing proportional to the actual transaction rate (throughput) on the network. Generally, public DLTs supporting smart contracts have the greatest CPU requirement, and private ledgers with only value transactions require the least.

2) *Storage*: For ledgers with a large state and no method of intermediate state pruning (such as in Vault [54]), the storage requirements can grow large. For example, even if transaction histories are omitted and only the set of unspent transaction outputs in Bitcoin is serialized⁴, even this is three gigabytes in size [55]. Similarly, the storage requirements of an Ethereum node varies, but even in minimal configuration is tens of gigabytes [56]. Thus, the consideration of storage space available on the node is a severe constraint on a DLT database.

3) *Network*: The network requirements of a node can vary significantly between *bootstrapping* and steady-state operations [54]. During bootstrap, e.g. during the node setup on the network, the amount of data that needs to be fetched by the device can vary from relatively small (a few hundred megabytes) to hundreds of gigabytes depending on the DLT. During the steady-state operation, the frequency of updates on the network has a strong effect on the network requirements and the latency of a node with respect to the network [57].

D. Other Challenges

A large number of many other concerns can be attributed to the relative newness of the research field, including lack of *established* standards in the area⁵. Similarly, there are no established benchmarks for measuring DLT performance, and any benchmarks that exist are strictly limited to a narrow set of DLTs. This is in large contrast to established benchmarks that exist, for example, in the field of databases (TPC-C benchmarks). For industrial uses, the licensing model of the technology is an issue — while non-viral open-source licenses are usually acceptable for integration to industrial systems, the availability and continuity of development and support must be considered for any system with a long expected lifetime. Furthermore, there is currently a definite competence gap, with relatively few people with good understanding of DLTs, making competence acquisition and retention a challenge.

Finally, almost all ledger technologies today have severe limitations in terms of crypto-agility [5]. Each and every cryptographic algorithm has a finite lifetime. Some of them are more resilient than others, extending their lifetime by e.g. increasing the key length, but it is seldom longer than 30 years. For the so-called hash functions, the secure lifetime has usually been less than 10 years [58].

Most modern cryptographic protocols have been designed to be crypto-agile, meaning that the underlying cryptographic functions can be changed without the protocol itself being affected. Not so for blockchains. Today's public blockchains

⁴The set of unspent transactions in Bitcoin needs to be known by nodes to prevent double-spend attacks.

⁵There are standardization efforts, such as ISO/TC 307 or ILP from W3C.

and many other DLTs do not support changing the underlying algorithms. Furthermore, it is very unlikely that they would even in the near future, due to their inflexible governance models [59].

V. POTENTIAL SOLUTIONS

We separate potential solutions of how to improve usability of DLTs in industrial context, with special focus of integrating IoT devices to a ledger, into two categories: *direct integration models* and *indirect integration models*, discussed below.

A. Direct Integration Models

In direct integration models the desired functionality of the target ledger is directly integrated into the service system, potentially through intermediary nodes. Therefore, for an IoT device operating on a DLT, we identify four possible integration models, based on and slightly paraphrased from earlier work by Özyılmaz and Yurdakul [60]:

1) *Full Validating Node*: The device may itself be a full validating node, if it has sufficient resources to operate on the DLT directly. Note that the assumption is that for most public ledgers this is not feasible. However, this may be a feasible approach for private and consortium ledgers where the ledger technology can be selected more freely.

2) *Thin Client*: The device may operate as a thin client, and use so-called “light protocols” to access the DLT. The light protocols that are used in public blockchains often can provide only probabilistic security assurances, while on permissioned ledgers, if there is an explicit trust model that the device can use to compute ledger validity efficiently, the device may attain higher security guarantees even as a thin client. Note that in the general thin client scenario, the device has no direct trust relationship with the light protocol server node it is communicating with (in contrast to *trusted gateway* below).

3) *Trusted Gateway*: The device may have access to a trusted gateway. In this scenario, the device uses a ledger protocol (such as light client protocol) so that it still retains understanding of the DLT itself. The trusted gateway is assumed to be a validating node, able to provide high security assurances of the ledger state to the device. The trust relationship may be explicit (such as a commercial service contract) or implicit (same owner). The trusted gateway thus provides a trusted interface to the ledger, but it does not assume the role of the device.

4) *Trusted Service*: The device may also relegate all DLT operations to a trusted service, thereby allowing the service to assume the role of the device and operate on its behalf. The device itself does not retain any explicit knowledge of the DLT anymore. The device-to-service interface can be specifically tailored to the needs of the device and service, allowing low resource consumption. However, from the device’s point of view, it is completely reliant on the trusted service to uphold the security model regarding DLT operations.

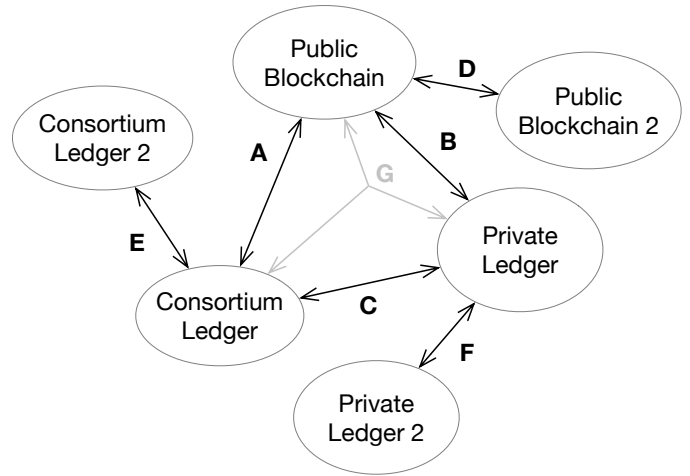


Figure 2. Interledger operations across differently governed ledgers. Cases A–F describe operations across two ledgers, but it is possible to construct interledger operations that operate across more than two ledgers simultaneously, as in case G.

B. An Indirect Approach: Interledger

The other approach is to integrate the device to a target DLT indirectly, *across multiple ledgers*, where the device integrates to some other ledger than the target ledger, and interledger operations are used to integrate across these ledgers. Here the device integrates directly (in the most secure manner) on the ledger that can now be selected to be compatible with the device’s resource constraints. The operations that must be performed on another ledger (such as on a public blockchain) are then bridged across the ledgers using an interledger gateway.

Several techniques exist, and are being researched, on how to perform interledger operations securely with full guarantees of transaction completion across multiple ledgers. For example, [61] defines six different categories of interledger techniques: atomic cross-chain transactions, sidechains, bridging, ledger-of-ledgers, payment channels, and layered value transfer protocols. In other cases, the interledger operations may carry financial risks, but the protocol is designed to make any nonconformance auditable [62].

The use of interledger operations allows many different types of DLTs to be joined together, as shown in Fig. 2. One common need is to transfer value (payments) across different ledgers, which has been demonstrated to be feasible across multiple public and consortium ledgers [63] (cases A and D in Fig. 2).

Interledger operations across different consortiums can occur when two consortiums operating in the same business area collaborate. For example, competing logistic consortia could potentially collaborate on product safety issues, using interledger operations to establish a non-repudiable auditing trail without devolving confidential business information. Similarly, collaboration across industries facilitated through interledger operations are possible, such as integrating sensor information from cargo logistics across to just-in-time man-

ufacturing (case E). Similarly, organizations partaking in a consortia may wish to use an internal ledger for their industrial device integration (e.g. use a ledger suitable for constrained devices), or even across organizations' private ledgers (cases C and F).

Some early versions of interledger technologies are already in production use, such as the Wanchain 3.0 Bitcoin Bridge to Ethereum [64].

A number of new use cases for DLTs can be conceived by the interledger approach, as it enables leveraging the benefits of private, consortium and public ledgers in one application. For example, a network of digital twins could have its own lightweight private ledger to which the devices are also connected. This private ledger is used for quick and reliable transfer of information between the devices and their digital twins. The digital twins can be responsible for the heavier roles of the ledger, whereas devices just utilize the reliability. With interledgers, this private ledger can be tied to the rest of the world, for example thereby enabling payments, through any consortium and/or public ledgers.

Finally, we surmise that our interledger approach could be used to introduce crypto-agility to DLTs through *ledger agility*. That is, instead of making a specific blockchain, such as the Bitcoin, crypto agile, we suggest making the interledger system *independent* of the underlying ledgers. Once a new and better DLT is introduced, it may be added to the system through the interledger mechanisms. Similarly, if a vulnerability is found in an existing ledger, it may be phased out of the system.

Overall, however, interledger operations have currently inherent limitations, and are most constrained by the latencies of associated ledgers. These may be sometimes alleviated through the use of sidechains, off-chain cryptographic proofs, and operating under optimistic assumptions, but substantial open questions still remain.

VI. CONCLUSION

There is a growing interest towards distributed ledger based solutions within various sectors of the industry, with about 140 distinct research projects in the energy sector alone [2] and probably several thousand altogether. However, despite of this quite tremendous interest, there are few if any projects that have actually led to commercially viable, market ready solutions. The reasons for this situation appear to be manifold, including issues related to business models, the anti-rival nature of data, immaturity of technology, lack of established standards, and e.g. lack of enforcing regulation.

In this paper, we have suggested using the so-called interledger approaches to solve or alleviate some of the problems that are currently slowing the adaption of the technology. In particular, the use of interledger methods allows organizations to mitigate technology and vendor lock-in risks, maximizing their capability to integrate into existing ledger networks, and offer flexibility regarding new developments in the rapidly developing field of distributed ledger technologies.

REFERENCES

- [1] P. C. Evans and M. Annunziata, "Industrial internet: Pushing the boundaries," 2012.
- [2] M. Andoni, V. Robu, D. Flynn, S. Abram, D. Geach, D. Jenkins, P. McCallum, and A. Peacock, "Blockchain technology in the energy sector: A systematic review of challenges and opportunities," *Renewable and Sustainable Energy Reviews*, vol. 100, pp. 143–174, 2019.
- [3] Anonymized, "Anonymized," in *Proceedings of the 30th European Conference of the International Telecommunications Society*, Helsinki: International telecommunications Society, 2019-06-16/2019-06-29.
- [4] C. Jones and C. Tonetti, "Nonrivalry and the Economics of Data," in *Society for Economic Dynamics 2018 Meeting Papers*, vol. 477, 2018.
- [5] *Crypto agility*, in *Wikipedia*, Page Version ID: 887668579, 2019-03-14. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Crypto_agility&oldid=887668579 (visited on 2019-03-15).
- [6] C. Kolias, G. Kambourakis, A. Stavrou, and J. Voas, "DDoS in the IoT: Mirai and other botnets," *Computer*, vol. 50, no. 7, pp. 80–84, 2017.
- [7] J. Lee, B. Bagheri, and H.-A. Kao, "A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 3, pp. 18–23, 2015-01-01, ISSN: 2213-8463. DOI: 10.1016/j.mfglet.2014.12.001.
- [8] J. Autiosalo, "Platform for industrial internet and digital twin focused education, research, and innovation: Ilmatar the overhead crane," in *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*, 2018-02, pp. 241–244. DOI: 10.1109/WF-IoT.2018.8355217.
- [9] E. Negri, L. Fumagalli, and M. Macchi, "A Review of the Roles of Digital Twin in CPS-based Production Systems," *Procedia Manufacturing*, 27th International Conference on Flexible Automation and Intelligent Manufacturing, FAIM2017, 27-30 June 2017, Modena, Italy, vol. 11, pp. 939–948, 2017-01-01, ISSN: 2351-9789. DOI: 10.1016/j.promfg.2017.07.198.
- [10] T. Kobzan, A. Biendarra, S. Schriegel, T. Herbst, T. Müller, and J. Jasperneite, "Utilizing Blockchain Technology in Industrial Manufacturing with the help of Network Simulation," in *2018 IEEE 16th International Conference on Industrial Informatics (INDIN)*, 2018-07, pp. 152–159. DOI: 10.1109/INDIN.2018.8472011.
- [11] S. Nakamoto, *Bitcoin: A peer-to-peer electronic cash system*, 2008. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>.
- [12] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum Project Yellow Paper*, vol. 151, 2014.
- [13] Ripple - One Frictionless Experience To Send Money Globally, [Online]. Available: <https://ripple.com/> (visited on 2019-03-20).
- [14] Hyperledger Fabric, [Online]. Available: <https://www.hyperledger.org/projects/fabric> (visited on 2019-03-20).
- [15] P. L. Aublin, S. B. Mokhtar, and V. Quéma, "RBFT: Redundant Byzantine Fault Tolerance," in *2013 IEEE 33rd International Conference on Distributed Computing Systems*, 2013-07, pp. 297–306. DOI: 10.1109/ICDCS.2013.53.
- [16] C. Cachin and M. Vukolić, "Blockchain consensus protocols in the wild," *arXiv preprint arXiv:1707.01873*, 2017.
- [17] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, "An overview of blockchain technology: Architecture, consensus, and future trends," in *2017 IEEE International Congress on Big Data (BigData Congress)*, IEEE, 2017, pp. 557–564.
- [18] K. Korpela, J. Hallikas, and T. Dahlberg, "Digital supply chain transformation toward blockchain integration," in *Proceedings of the 50th Hawaii International Conference on System Sciences*, 2017.
- [19] S. Apte and N. Petrovsky, "Will blockchain technology revolutionize excipient supply chain management?" *Journal of Excipients and Food Chemicals*, vol. 7, no. 3, p. 910, 2016.
- [20] R. Casado-Vara, A. González-Briones, J. Prieto, and J. M. Corchado, "Smart contract for monitoring and control of logistics activities: Pharmaceutical utilities case study," in *The 13th International Conference on Soft Computing Models in Industrial and Environmental Applications*, Springer, 2018, pp. 509–517.
- [21] S. A. Abeyratne and R. P. Monfared, "Blockchain ready manufacturing supply chain using distributed ledger," 2016.
- [22] S. Underwood, "Blockchain beyond bitcoin," *Communications of the ACM*, vol. 59, no. 11, pp. 15–17, 2016.
- [23] O. Gallay, K. Korpela, N. Tapio, and J. K. Nurminen, "A peer-to-peer platform for decentralized logistics," in *Proceedings of the Hamburg International Conference of Logistics (HICL)*, epubli, 2017, pp. 19–34.

- [24] N. Hackius and M. Petersen, "Blockchain in logistics and supply chain: Trick or treat?" In *Proceedings of the Hamburg International Conference of Logistics (HICL)*, epubli, 2017, pp. 3–18.
- [25] R. Polim, Q. Hu, and S. Kumara, "Blockchain in megacity logistics," in *IIE Annual Conference. Proceedings*, Institute of Industrial and Systems Engineers (IISE), 2017, pp. 1589–1594.
- [26] K. Kobylinski, J. Bennett, N. Seto, G. Lo, and F. Tucci, "Enterprise application development in the cloud with IBM Bluemix," in *Proceedings of 24th Annual International Conference on Computer Science and Software Engineering*, IBM Corp., 2014, pp. 276–279.
- [27] A. Badzar, "Blockchain for securing sustainable transport contracts and supply chain transparency-an explorative study of blockchain technology in logistics," 2016.
- [28] H. M. Kim and M. Laskowski, "Toward an ontology-driven blockchain design for supply-chain provenance," *Intelligent Systems in Accounting, Finance and Management*, vol. 25, no. 1, pp. 18–27, 2018, ISSN: 1099-1174. DOI: 10.1002/isaf.1424.
- [29] M. Dobrovnik, D. Herold, E. Fürst, and S. Kummer, "Blockchain for and in Logistics: What to Adopt and Where to Start," *Logistics*, vol. 2, no. 3, p. 18, 2018.
- [30] M. Mylrea and S. N. G. Gourisetti, "Blockchain: A path to grid modernization and cyber resiliency," in *2017 North American Power Symposium (NAPS)*, IEEE, 2017, pp. 1–5.
- [31] —, "Blockchain for smart grid resilience: Exchanging distributed energy at speed, scale and security," in *2017 Resilience Week (RWS)*, IEEE, 2017, pp. 18–23.
- [32] Z. Y. Dong, F. Luo, and G. Liang, "Blockchain: A secure, decentralized, trusted cyber infrastructure solution for future energy systems," *Journal of Modern Power Systems and Clean Energy*, vol. 6, no. 5, pp. 958–967, 2018-07-06.
- [33] G. Liang, S. R. Weller, F. Luo, J. Zhao, and Z. Y. Dong, "Distributed blockchain-based data protection framework for modern power systems against cyber attacks," *IEEE Transactions on Smart Grid*, 2018.
- [34] R. Ebrahimi, C. Morisset, H. Patsios, and Z. Pourmirza, "Report on Smart Energy Systems and Cyber Security," 2018.
- [35] S.-K. Kim and J.-H. Huh, "A study on the improvement of smart grid security performance and blockchain smart grid perspective," *Energies*, vol. 11, no. 8, p. 1973, 2018.
- [36] F. Tian, "An agri-food supply chain traceability system for China based on RFID & blockchain technology," in *2016 13th International Conference on Service Systems and Service Management (ICSSSM)*, IEEE, 2016, pp. 1–6.
- [37] L. Ge, C. Brewster, J. Spek, A. Smeenk, J. Top, F. van Diepen, B. Klaase, C. Graumans, and M. d. R. de Wildt, *Blockchain for Agriculture and Food: Findings from the Pilot Study*, 2017-112. Wageningen Economic Research, 2017.
- [38] Y.-P. Lin, J. Petway, J. Anthony, H. Mukhtar, S.-W. Liao, C.-F. Chou, and Y.-F. Ho, "Blockchain: The evolutionary next step for ICT E-agriculture," *Environments*, vol. 4, no. 3, p. 50, 2017.
- [39] D. Tse, B. Zhang, Y. Yang, C. Cheng, and H. Mu, "Blockchain application in food supply information security," in *2017 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, IEEE, 2017, pp. 1357–1361.
- [40] M. P. Caro, M. S. Ali, M. Vecchio, and R. Giaffreda, "Blockchain-based traceability in Agri-Food supply chain management: A practical implementation," in *2018 IoT Vertical and Topical Summit on Agriculture-Tuscany (IOT Tuscany)*, IEEE, 2018, pp. 1–4.
- [41] J. F. Galvez, J. C. Mejuto, and J. Simal-Gandara, "Future challenges on the use of blockchain for food traceability analysis," *TrAC Trends in Analytical Chemistry*, 2018.
- [42] R. Kamath, "Food traceability on blockchain: Walmart's pork and mango pilots with IBM," *The JBBA*, vol. 1, no. 1, p. 3712, 2018.
- [43] F. Yiannas, "A New Era of Food Transparency Powered by Blockchain," *Innovations: Technology, Governance, Globalization*, vol. 12, no. 1-2, pp. 46–56, 2018.
- [44] C. Xiu and K. K. Klein, "Melamine in milk products in China: Examining the factors that led to deliberate use of the contaminant," *Food Policy*, vol. 35, no. 5, pp. 463–470, 2010.
- [45] J. Premanandh, "Horse meat scandal—A wake-up call for regulatory authorities," *Food control*, vol. 34, no. 2, pp. 568–569, 2013.
- [46] *2013 horse meat scandal*, in *Wikipedia*, Page Version ID: 882673924, 2019-02-10. [Online]. Available: https://en.wikipedia.org/w/index.php?title=2013_horse_meat_scandal&oldid=882673924 (visited on 2019-03-15).
- [47] C. Decker, R. Russell, and O. Osuntokun, "Eltoo: A simple layer2 protocol for bitcoin," *White paper: https://blockstream.com/eltoo.pdf*, 2018.
- [48] Y. Sompolinsky and A. Zohar, "Accelerating Bitcoin's Transaction Processing. Fast Money Grows on Trees, Not Chains," 881, 2013. [Online]. Available: <https://eprint.iacr.org/2013/881> (visited on 2017-09-06).
- [49] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A Secure Sharding Protocol For Open Blockchains", in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security - CCS'16*, Vienna, Austria: ACM Press, 2016, pp. 17–30, ISBN: 978-1-4503-4139-4. DOI: 10.1145/2976749.2978389.
- [50] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, "OmniLedger: A Secure, Scale-Out, Decentralized Ledger via Sharding," in *2018 IEEE Symposium on Security and Privacy (SP)*, 2018-05, pp. 583–598. DOI: 10.1109/SP.2018.000-5.
- [51] S. Gleiser. (2017-06-21). Ethereum Transactions Peaking As Ether Plunges, [Online]. Available: <https://bitcoinchaser.com/ethereum-transactions-problem> (visited on 2019-03-25).
- [52] M. Hrones. (2017-12-05). CryptoKitties Creates Massive Backlog on the Ethereum Network, [Online]. Available: <http://bitcoinist.com/cryptokitties-creates-massive-backlog-on-the-ethereum-network/> (visited on 2018-03-28).
- [53] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, "On the Security and Performance of Proof of Work Blockchains," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16, New York, NY, USA: ACM, 2016, pp. 3–16, ISBN: 978-1-4503-4139-4. DOI: 10.1145/2976749.2978341.
- [54] D. Leung, A. Suhl, Y. Gilad, and N. Zeldovich, "Vault: Fast Bootstrapping for the Algorand Cryptocurrency", p. 15,
- [55] S. Delgado-Segura, C. Pérez-Solà, G. Navarro-Arribas, and J. Herrera-Joancomartí, "Analysis of the Bitcoin UTXO Set", in *Financial Cryptography and Data Security*, A. Zohar, I. Eyal, V. Teague, J. Clark, A. Bracciali, F. Pintore, and M. Sala, Eds., ser. Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2019, pp. 78–91, ISBN: 978-3-662-58820-8.
- [56] A. Schoedon. (2017-11-29). The Ethereum-blockchain size will not exceed 1TB anytime soon., [Online]. Available: <https://dev.to/5chdn/the-ethereum-blockchain-size-will-not-exceed-1tb-anytime-soon-58a> (visited on 2018-01-11).
- [57] P. Danzi, A. E. Kalor, C. Stefanovic, and P. Popovski, "Analysis of the Communication Traffic for Blockchain Synchronization of IoT Devices," in *2018 IEEE International Conference on Communications (ICC)*, 2018-05, pp. 1–7. DOI: 10.1109/ICC.2018.8422485.
- [58] V. Aurora. (2017). Lifetimes of cryptographic hash functions, [Online]. Available: <http://valerieaurora.org/hash.html> (visited on 2019-03-30).
- [59] J. Mattila and T. Seppälä, "Distributed Governance in Multi-sided Platforms: A Conceptual Framework from Case: Bitcoin", in *Collaborative Value Co-Creation in the Platform Economy*, ser. Translational Systems Sciences, Springer, Singapore, 2018, pp. 183–205, ISBN: 978-981-10-8955-8. DOI: 10.1007/978-981-10-8956-5_10. [Online]. Available: https://link.springer.com/chapter/10.1007/978-981-10-8956-5_10 (visited on 2018-08-09).
- [60] K. R. Özyılmaz and A. Yurdakul, "Work-in-progress: Integrating low-power IoT devices to a blockchain-based infrastructure," in *2017 International Conference on Embedded Software (EMSOFT)*, 2017-10, pp. 1–2. DOI: 10.1145/3125503.3125628.
- [61] A. Karila, Y. Kortensniemi, D. Lagutin, P. Nikander, S. Paavola, N. Fotiou, G. C. Polyzos, V. A. Siris, and T. Zahariadis, "SOFIE: Secure Open Federation for Internet Everywhere", in *Proceedings 2018 Workshop on Decentralized IoT Security and Standards*, San Diego, CA: Internet Society, 2018, ISBN: 978-1-891562-51-8. DOI: 10.14722/diss.2018.23001.
- [62] Interledger Payments Community Group, [Online]. Available: <https://www.w3.org/community/interledger/> (visited on 2017-09-22).
- [63] M. del Castillo. (2017-06-02). Interoperability Boost: Ripple Sends Blockchain Transaction Across 7 Ledgers, [Online]. Available: <https://www.coindesk.com/interoperability-boost-ripple-sends-blockchain-transaction-across-7-different-ledgers> (visited on 2019-03-20).
- [64] D. Reecer. (2018-12-11). Wanchain 3.0 Launches Bitcoin Bridge to Ethereum, [Online]. Available: <https://medium.com/wanchain-foundation/wanchain-3-0-launch-bitcoin-ethereum-erc20-7cd504f25c0c> (visited on 2019-03-30).