



OPEN

Analysis of distributed ledger technologies for industrial manufacturing

Lam Duc Nguyen^{1,5}✉, Arne Bröring^{2,5}, Massimo Pizzol³ & Petar Popovski⁴

In recent years, industrial manufacturing has undergone massive technological changes that embrace digitalization and automation towards the vision of intelligent manufacturing plants. With the aim of maximizing efficiency and profitability in production, an important goal is to enable flexible manufacturing, both, for the customer (desiring more individualized products) and for the manufacturer (to adjust to market demands). Manufacturing-as-a-service can support this through manufacturing plants that are used by different tenants who utilize the machines in the plant, which are offered by different providers. To enable such pay-per-use business models, Distributed Ledger Technology (DLT) is a viable option to establish decentralized trust and traceability. Thus, in this paper, we study potential DLT technologies for efficient and intelligent integration of DLT-based solutions in manufacturing environments. We propose a general framework to adapt DLT in manufacturing, and then we introduce the use case of *shared manufacturing*, which we utilize to study the communication and computation efficiency of selected DLTs in resource-constrained wireless IoT networks.

Industrial Internet of Things (IIoT) is a recent concept that gained traction with the emergence of the wireless 5G technology and it is already exhibiting a great impact within the manufacturing domain¹. The general trend is to embrace digitalization and automation towards manufacturing plants that act as cyber-physical systems. This results in an increasing number of smart devices with sensors and actuators that are being integrated in industrial automation processes. In parallel, local edge computing infrastructures are being built up in manufacturing plants, which provide resources for advanced computing and henceforth the basis for next generation IIoT applications². The key economic driver behind this technological evolution is the increase in the production flexibility. This allows for smaller lot sizes and more individualized products for customers. These trends are supported by business models, such as manufacturing-as-a-service, where manufacturing facilities are utilized more flexibly by numerous tenants who utilize the machines in the plant, which are offered by different providers. These economic forces drive manufacturing plants towards an increase in technological complexity and require improvements in system reliability, intelligence, and trustworthiness during operation³. Especially the opening of the manufacturing plant's ecosystem to a diverse set of involved parties poses many challenges for manufacturing enterprises to satisfy the trust requirements of multi-partner collaboration⁴.

Distributed Ledger Technology (DLT) can be used to address those trust and privacy challenges in the manufacturing environment of the future, e.g., to transparently store machinery's usage data as a basis for pay-per-use business models on the manufacturing shop floor. A DLT is a distributed ledger of transactions—rather than being kept in a single, centralized location, the information is held by all the nodes of a network⁵. In general, all these network nodes have copies of the same ledger. This removes the need for a third-party to assure that rules are being implemented correctly, instead, this is implicitly done through a decentralized system. Although the most widely known instance of DLT is Blockchain, and, specifically, the most popular decentralized digital currency based on Blockchain is Bitcoin, the transactions on a DLT do not have to be financial. In essence, a transaction simply represents a change in state for whichever data point the DLT's stakeholders want to track. DLTs are driven by consensus: when a node or a *DLT-client* initiates a transaction, its details are broadcast to the entire network, checked by other nodes and accepted if there is consensus. *DLT-clients* are considered as lightweight devices which have limited resources and just initiate transactions, as well as transmit transactions to *DLT-managers* to validate. When the DLT takes the form of a Blockchain, once a transaction has been validated, it is bundled with other transactions into a block of data. Each block is secured via a cryptographic algorithm.

¹Software Systems Research Group, CSIRO Data61, Sydney 2015, Australia. ²Siemens AG, 81739 Munich, Germany. ³Department of Planning, Aalborg University, 9220 Aalborg, Denmark. ⁴Connectivity Section, Aalborg University, 9220 Aalborg, Denmark. ⁵These authors contributed equally: Lam Duc Nguyen and Arne Bröring. ✉email: lam.nguyen@data61.csiro.au

This results in a unique signature for each block known as a hash. These blocks are then ordered sequentially into a chain of blocks, with each block also containing the previous block's hash⁶. This makes it extremely difficult to tamper with a block, as altering a single piece of data would result in a different hash value, making it evident to the DLT's users and causing the transaction and block to be rejected.

In short, DLTs allow the storage of transactions in immutable records and every record is distributed across many nodes. Thus, security in DLTs comes from the decentralized operation, but also from the use of strong public-key cryptography and cryptographic hashes⁵. The key benefits of the integration of DLTs into manufacturing systems are: (1) auditable and guaranteed immutability and transparency for stored data (e.g., machine usage data, sensing data about machine conditions, or logs about user/technician engagements), (2) no need for a third party to assure the rules between the different parties in the manufacturing ecosystem are met, (3) enabling security and privacy of information in manufacturing networks in conjunction with others techniques e.g, private transactions and private channels, which is urgently needed as more than 25% of cyber attacks will involve IoT⁷.

To showcase these benefits and to have a realistic use case as an example for our studies, we implement in this work the DLT-based application of *shared manufacturing*, which relates to the economic driver of 'flexible production'. Specifically, a robot arm, as part of a production cell with multiple machines, is offered by a provider, who allows different tenants of the plant the usage of the robot arm, while expecting a usage fee. In this application, the DLT is required to capture the usage times of the robot arm through the various tenants, which is then the basis for a correct billing and payment for usage time. Some parts of this process can be done automatically with smart contracts. These involve two entities turning a business contract into code that recognizes actions on the DLT. For example, a smart contract might recognize that a rental of a machine from "provider A" to "customer B" on a certain date for a specific time period should be for a specific price⁸. This simplifies processes that take significant time to check. This structure gives DLT participants confidence in their transaction without the need to trust each other. Nor do they need to agree on a trusted third party to make sure they're both following the rules. Because the ledger of transactions is consensus-based and distributed, records stored in it cannot be erased or changed.

To be able to implement the above described application and reap the described benefits, the system designs of current manufacturing plants need to be adjusted to be able to accommodate the operation of a DLT and overcome certain limitations: First, today's computation infrastructures of industrial manufacturing plants are typically designed as centralized systems, where cloud services perform data aggregation and analysis⁹. While the manufacturing infrastructure comprises a multitude of IoT devices and sensors that collect data and have only little computing power, the gathering and processing of data in a centralized cloud service may lead to network overload and single points of failure¹⁰. To setup a DLT network in such an environment, a sufficient amount of local computing capacity¹¹ needs to be available and, potentially, edge computing facilities can be integrated in the computation infrastructure. Furthermore, industrial communication systems have been traditionally designed for reliable operation in a noisy factory environment, employing mainly wired and proprietary communication technologies to connect sensors, actuators, and controllers. Nevertheless, with the emergence of IIoT, future factories will increasingly rely on diverse communication technologies, including wireless standards, to ensure reliability, interoperability, and remote operation and control of production processes through the Internet. These wireless links are potentially less reliable and are more constrained, which needs to be considered, when operating a DLT network. From these limitations regarding the system infrastructure, we derive the key *research question* of this work: "What is the computation and communication overhead that results from the operation of a DLT network in a manufacturing environment?" The answer to this question will be critical to understand for future research on applications of DLT in manufacturing, as well as for practitioners who want to deploy a DLT network in a manufacturing plant.

The use of DLT in manufacturing has received attention from both academia and industry because of its promise for easing supply chain and manufacturing operation management problems due to its advantages in transparency, traceability, and security. In industry, Bosch increasingly connects their products to the IIoT in order to directly participate in the digital economy. The goal is to build an *Economy of Things*, which will be based on DLT¹². Another example is a concrete solution by Siemens, which enables their *MindSphere* IIoT platform to track products of the food and beverage industry transparently throughout their entire life cycle based on DLT¹³. *MindSphere* exploits all useful information before forwarding only a crucial subset to the distributed ledger. The DLT then makes sure the collected data is safe and transparently accessible to everyone who is part of the ecosystem. In academia, Li et al.¹⁴ introduced a distributed P2P system that improves the security and scalability of the cloud-based manufacturing platform based on DLT. Danzi et al.¹⁵ analyze the communication aspects in terms of delay and overhead between IoT devices and Blockchain network. The authors demonstrate that, if the statistics of account updates and the channel state are known, the lightweight IoT clients can construct a list of events of interest that provides a predictable average communication cost. In addition, a survey¹⁶ about performance of different Blockchains is conducted, but the work mainly focuses on theoretical aspects, and lacks a detailed analysis in specific application areas such as manufacturing. Fu et al.¹⁷ presented an innovative environmentally sustainable DLT-energized strategy for the fashion apparel manufacturing industry. Yu et al.¹⁸ proposed a DLT-based service composition architecture for manufacturing. In general, the public DLT-based applications are characterized by the distinctive metric of computational trust.

In order to be able to answer our research question, we extend state-of-the-art through the following research contributions:

- General analysis of different DLTs and their capabilities when used in industrial manufacturing environments;
- System design for DLT-based IIoT manufacturing systems that can integrate and adapt multiple features and components;

	Hyper. fabric ¹⁹	Quorum ²⁰	Ethereum ²¹	IOTA ²²	Solana ²³
<i>DLT type</i>	Private	Private	Public/private	Public/private	Public/private
<i>Goals</i>	Open DLT framework	Open, based on Ethereum	Broad ecosystem	Lightweight	High scalability
<i>Application</i>	Enterprise DLT	Enterprise DLT	DApps	IoT	DApps
<i>Governance</i>	Linux Foundation	ConsenSys	ETH Foundation	IOTA Foundation	Solana
<i>Cryptocurrency</i>	N/A	N/A	Ether (ETH)	MIOTA	SOL
<i>Consensus</i>	Pluggable	Voting protocol	PoW	Tangle	PoH
<i>Smart contract</i>	nodejs, go, java	Solidity	Solidity	Solidity, Go, Rust	Rust
<i>Throughput</i>	~ 2000 tps	~ 100 tps	~ 100 tps	1000–1500 tps	~ 1400 tps
<i>Latency</i>	~ 250 ms	~ 414 ms	~ 2150 ms	~ 258 ms	~ 500 ms

Table 1. Comparison of different enterprise DLT platforms.

- Evaluation of communication and computation overhead of different DLTs in resource-constrained IoT networks. This benchmark of different DLTs for manufacturing scenarios will help interested parties to understand the trade-offs in DLT-based systems.

The remainder of this paper is organized as follows. In the next section, we present the results of this study. First, we present a general analysis of five different DLT platforms. Second, we introduce a system design for using DLT in industrial manufacturing. Third, we implement the shared manufacturing use case and perform a performance evaluation of the five different DLTs. Finally, we discuss our findings and indicate avenues for future research.

Results

In this section, we first study five different DLT platforms, then we propose a general framework to integrate DLT in manufacturing. Finally, we implement the use case of shared manufacturing and conduct the evaluation.

Analysis of DLTs for industrial manufacturing. Although a large number of DLTs are available, within the scope of which work we have selected five representative DLT platforms that are either already used or appear as most promising for manufacturing environments: Hyperledger Fabric, Ethereum, Quorum, Solana, and IOTA. The overview comparison of these DLTs is shown in Table 1.

Each DLT can be categorized as public, private, or hybrid, where the latter one can support features of both public and private ones. DLTs allow any user to pseudo-anonymously join the DLT network and do not restrict the rights of the nodes on the network. We are investigating in this paper the public DLTs Ethereum²¹, IOTA²², and Solana²³. However, for the implementation of our use case within a manufacturing plant such public DLTs are used in a private deployment by installing local networks. In contrast, private DLTs restrict access to their network to certain nodes and may also restrict the rights of nodes on the network. In this paper, we are investigating the private DLT platforms Hyperledger Fabric¹⁹ and Quorum²⁰. The identities of the users of a private DLT are known to the other users of that private DLT. In a Hybrid DLT, every transaction can happen quickly in its own private chain and commits to the public chain only happen as and when necessary, e.g., when public verification is required. This provides the immutable trust from the Blockchain as well as the scaling from private DLTs. Layer 2 solutions and side-chains²⁴ are variations of this concept.

Besides their type, the five DLTs have different goals and applications in focus. Hyperledger Fabric and Quorum are both aiming to offer an open foundation for new components to build a broad ecosystem that supports enterprises with various functionalities to deploy their own private DLT. Ethereum has a large community of developers and already an established ecosystem that focuses on decentralized applications (DApps), e.g., for decentralized finance. Solana follows a similar application focus, while aiming for higher scalability than Ethereum. IOTA's focus is on IoT applications and therefore aims to support DLT participants with a small footprint.

IIoT applications in the manufacturing environment will involve many stakeholders with different roles, functionalities, and information with access rules, identities and security factors. An important factor to provide security is the support to validate transactions generated by participating nodes. While Hyperledger Fabric and Quorum are suited solely for private setups, Ethereum, IOTA and Solana are designed for public networks, but can also be configured for private purposes. In terms of security and confidentiality, public networks can show certain advantages over private ones, especially if they are able to provide transparency and distributed storage. Besides, the more users a public DLT has, the more secure it is. However, for enterprise use (i.e., also for typical manufacturing scenarios) public DLTs are not ideal as companies deal with highly sensitive data and cannot allow arbitrary users to join their network. Also, private DLTs provide very low or no fees for validation and a faster consensus process. However, a private DLT can be altered by its owners, making it more vulnerable to hacking²⁵. Besides, only Hyperledger Fabric supports data confidentially by default; this is done via in-band encryption and guarantees the privacy of data by creating private channels (e.g., to setup for departments within an organization). Therefore, Hyperledger Fabric allows for authorization with trusted Certificate Authority per channel. These features are vital in a trusted IoT system for enterprises.

Each DLT platform deploys a different consensus mechanism. Ethereum uses the Proof-of-Work (PoW) consensus that requires involved parties of a network to expend effort solving a mathematical puzzle to prevent anybody from gaming the system. PoW consensus consumes significant computing and energy resources, which is not suitable for resource-limited systems. Quorum, as an enterprise version of Ethereum, uses a voting-based consensus protocol. This consensus protocol achieves consensus on transactions and key network decisions by counting the number of votes cast by nodes on the networks and not consuming more energy for verification as compared to PoW. IOTA uses “little” PoW for preventing spamming attacks. In Ethereum, doing PoW is to receive the power to define the truth, the node with more power can solve the PoW faster and consume more energy. Meanwhile, IOTA use PoW with lower difficulty to prevent spamming and to allow transactions to be attached in the Tangle. The “little” PoW is still similar to original PoW but with lower difficulty. Hyperledger Fabric modularized the consensus part among distributed peers in an ordering service¹⁹, so that this platform allows users to choose their preferred algorithm, e.g., CFT (crash fault-tolerant) or BFT (byzantine fault-tolerant) ordering. Finally, Solana introduces a new consensus algorithm called Proof-of-History which allows timestamp field to be built into the blockchain itself instead of using values of timestamps as PoW DLTs. Whereas other DLTs require validators to talk to another in order to agree that time has passed, each Solana validator maintains its own clock by encoding the passage of time in a simple SHA-256, sequential-hashing Verifiable Delay Function (VDF)²⁶.

Table 1 specifies the smart contract programming language supported by the DLT. Smart contracts act as autonomous entities on the ledger that deterministically execute logic expressed as functions of the data that are written on the ledger. Therefore, smart contracts can be established to have automatic reactions from the DLT network to specific events. For example, in the use case of shared manufacturing, smart contracts can be used for restricting, tracking and payment for the usage of the rented machinery. The smart contract feature is currently supported by Ethereum, Solana and Hyperledger Fabric (called ‘Chaincode’). In IOTA, a smart contract is called Quobric which is now deprecated and the new beta version of smart contract is developing²⁷.

Furthermore, Table 1 states the performance characteristics of the DLTs. These values have been acquired both through our own experiments and the data available in the literature. These measures are of vital importance for IoT applications, particularly in manufacturing, where a large number of sensors may generate millions of data points per day. This requires high efficiency of the consensus mechanism, including the way in which transactions are processed by the peers, known as *endorsing peers* in Hyperledger Fabric, validators in Solana, and full nodes (peers) in Ethereum. Specifically, we have Solana, with 600 nodes and around 1000 validators. Currently Solana is hosting around 340 apps²⁸. Meanwhile, Ethereum has over 3000 Dapps running on its network. Regarding latency, the transaction confirmation time must be sufficiently short to avoid queuing in the DLT and to ensure consistency in the ledgers. The confirmation time of an Ethereum transaction in a public network is around 25 s in public networks. This value indicates that consensus over public networks may not be suitable for real-time IoT applications. However, other DLT platforms can achieve much lower confirmation times²⁹. Note that in Table 1 the transaction confirmation time is included in the end-to-end latency, which however does not account for the communication latency at the radio access networks.

Another important performance feature of the DLTs is related to the CPU usage and resulting energy consumption. The idea that DLT technologies and crypto-assets consume an excessive amount of electricity has been at the heart of recent discussions around this technology. The energy consumption of a DLT protocol should not be equated with its environmental footprint. Indeed, many use cases related to DLT technologies and crypto-assets may even contribute to improving the environmental footprint, in particular by using the surplus of decarbonised energy in certain geographical areas where the need for electricity is lower than the level of production³⁰. In the scope of this work, we study the carbon footprint of the different selected DLT platforms within the local testbed of the shared manufacturing use case.

The charge of fees to process the transactions, commonly known as *gas* is yet another factor to take into account to select the appropriate DLT. These may greatly increase the operational costs of the network, which negatively impacts the throughput of the DLT. On the one hand, transaction fees pose a problem in massive IIoT scenarios if the generation of a large number of transactions is essential. On the other hand, these fees may contribute to minimizing the amount of redundant transactions generated by the sensors, which in turn offloads the DLT. In industrial manufacturing domain, the required fee for generating transactions within a company or among some cooperative organizational setup may not be suitable. In addition, businesses have always required a reasonable degree of privacy as well as control over their networks, so that publishing the data on a Blockchain is not reasonable and potentially unsafe. Therefore, we consider only private Blockchains for the enterprise scenario.

System design for using DLT in industrial manufacturing. The proposed system design is described in Fig. 1. It comprises four key parts as described below.

DLT system. This component includes all modules to build various features of DLT technologies such as consensus, smart contract, data authorization, identity management, and peer-to-peer (P2P) communication. These components must ensure that every change to the ledger is reflected in all copies in seconds or minutes and provide mechanisms for the secure storage of the data generated by IoT devices and parameter configurations. There are numerous DLTs with different characteristics that may be beneficial for different target applications. The DLT nodes can be located everywhere and connected with base stations via the Internet.

Physical machines. This component consists of physical robots, machines, and IoT sensor devices which collect the data and publish to the distributed ledger for accounting or analyzing purposes.

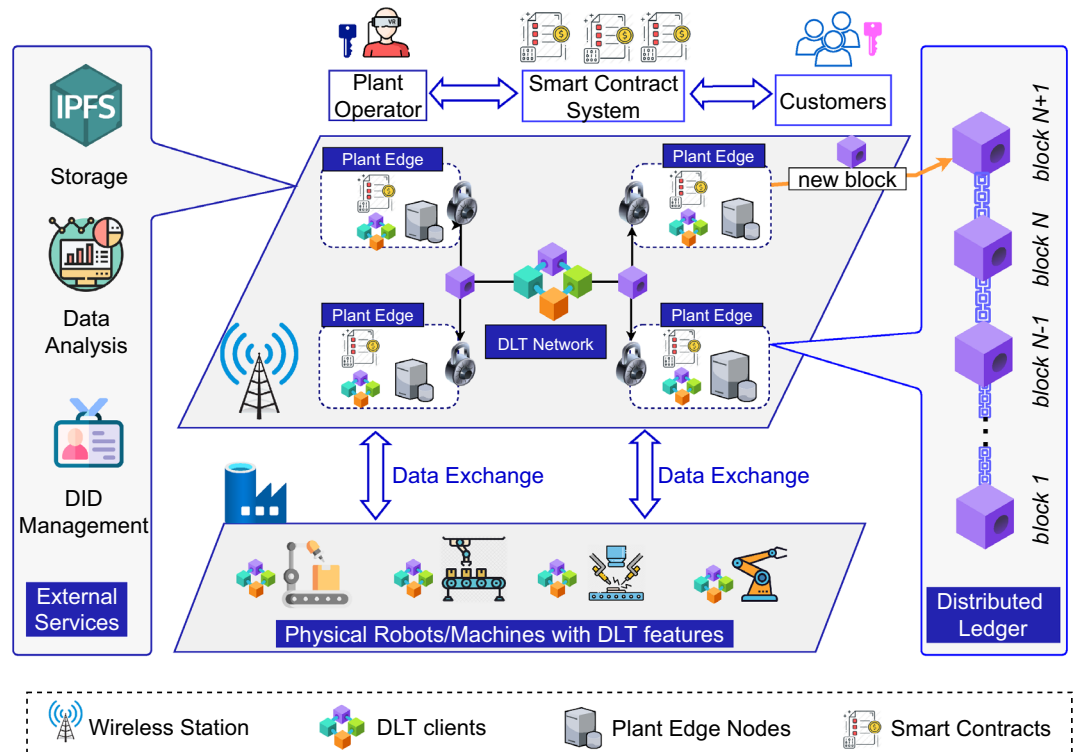


Figure 1. Overview of the system design.

Plant edge system. Even though DLT-based solutions offer significant countermeasures to secure data from tampering and support the distributed nature of the IoT, the massive amount of generated data from sensors and the high energy consumption required to verify transactions make these procedures unsuitable to execute directly on resource-limited IoT devices. Instead, edge servers with high computation resources can be used to handle real-time applications and to further increase the degree of privacy (e.g., through cloud computing)³¹. The edge network is a potential entity to cooperate with the DLT network in computationally heavy tasks and return the estimation results (e.g., from solving proof-of-work (PoW) puzzles, hashing or algorithm encryption) to the DLT network for verification.

External services. The devices of the manufacturing environment are typically resource-constrained with limited storage space and low computation capacity. Hence, external infrastructure which operates on the edge may be incorporated to provide external services, such as storage and computing. For example, the Interplanetary File System (IPFS) is a distributed file storage system that can store data generated from IoT networks and return a hash to the ledger based on the content of the data. Since the ledger cannot handle and store the massive amount of manufacturing data collected by the sensors, machines, and robots, the service provided by the IPFS is a vital component. The default configuration of IPFS connects to the global distributed network. In some cases regarding to privacy and confidentiality, a private IPFS network is preferred over connecting to the public IPFS network. In our scenario, we prefer to configure IPFS privately in a local cluster. Second, we introduce the application of payment channels to sharing manufacturing use case because of its natural advantages. In specific, a Payment Channel is a process where customers can make multiple transfers with e.g., plant operator, without sending a transaction to the DLT. Once the final transaction occurs between the participants the recipient can claim their funds by submitting one final transaction to the Smart Contract on the ledger. This allows both parties to avoid fees involved with multiple transactions. Smart Contracts can be an agreement about the rental time, specific tasks between customers and plant operators, or smart contracts created at the beginning of the process of payments. In addition, a Digital Identity Management (DID) could be added to support managing identity of participant devices in a distributed manner.

Performance evaluation of DLTs in a shared manufacturing use case. In this section, we analyze the application of the Shared Manufacturing use case and study its performance. The application uses DLT to automate the management of rentals of industrial robots, where the manufacturing plant operators and their customers can make agreements without third parties and the associated delay.

Along with data sharing³³ and vehicle sharing³⁴, the machine sharing concept in industry manufacturing has been recently identified as a key innovation for implementing the next industrial evolutionary step¹⁸. Open and shared manufacturing factories are composed of a number of industrial robots and other production machines that can be rented by customers. The advantage over traditional manufacturing plants is that such plants can have a higher workload and less idle periods, which in turn can make the production cheaper. Therefore,

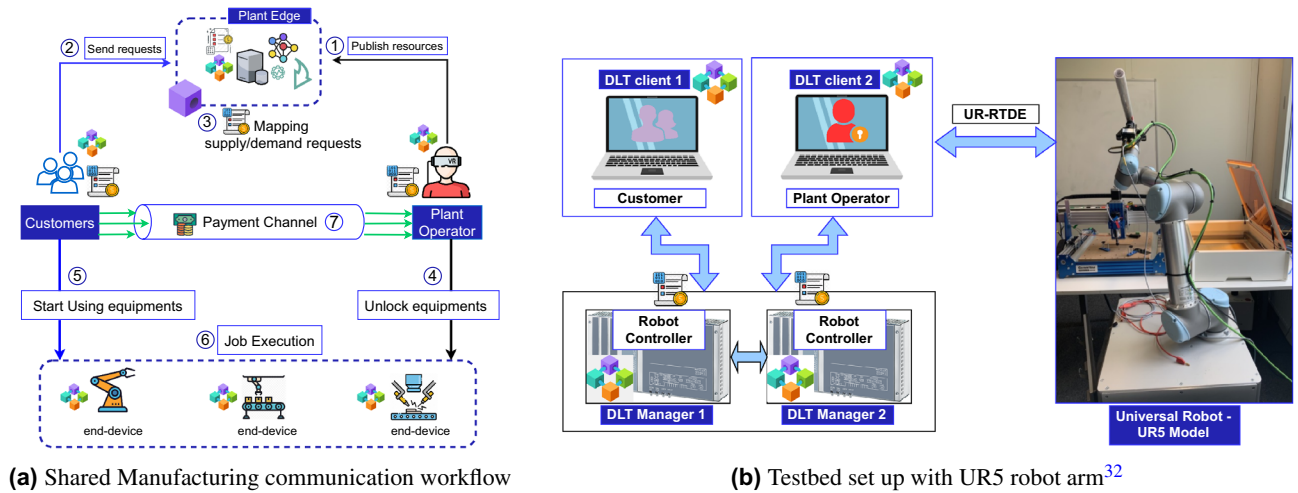


Figure 2. System model of shared manufacturing and local test-bed setup.

	DLT-manger-1	DLT-manger-2	DLT-client 1	DLT-client 2
Devices	Siemens Microbox	Laptop	Rapsberry Pi 3+	Rasperry Pi 3+
RAM	4 GB	8 GB	1 GB	1 GB
Connectivity	Ethernet	Ethernet	Wifi	Ethernet
Capacity	Intel(R) Core i7-351UE CPU @ 1.70 GHz x4 GHz	Intel(R) Core(TM) i7-8550U CPU @ 1.80 GHz 1.9 GHz	Quad Core 64 bit ARM cortex at 1.2 GHz	Quad Core 64 bit ARM cortex at 1.2 GHz

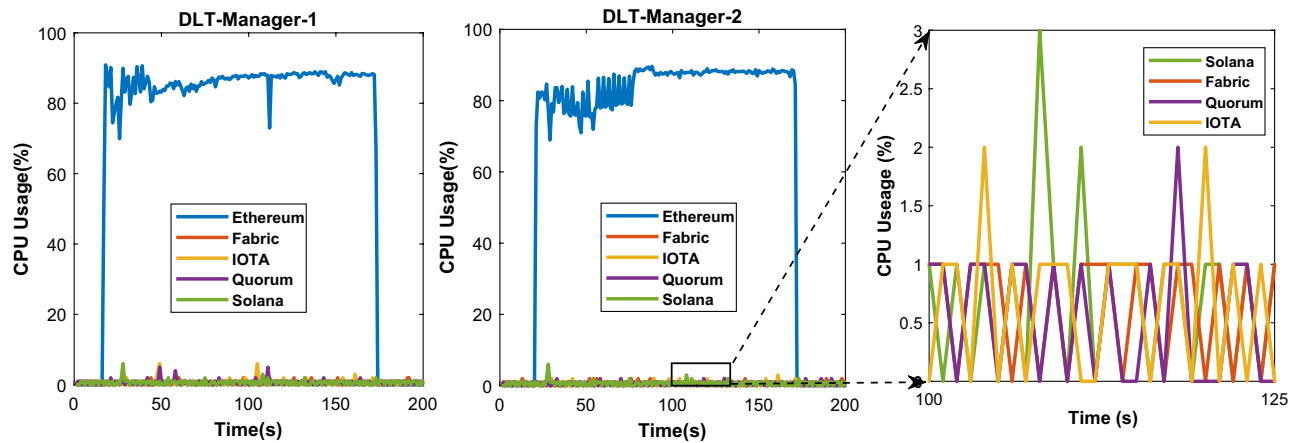
Table 2. Testbed settings.

production tasks need to be efficiently allocated on the available machine resources under consideration of system performance.

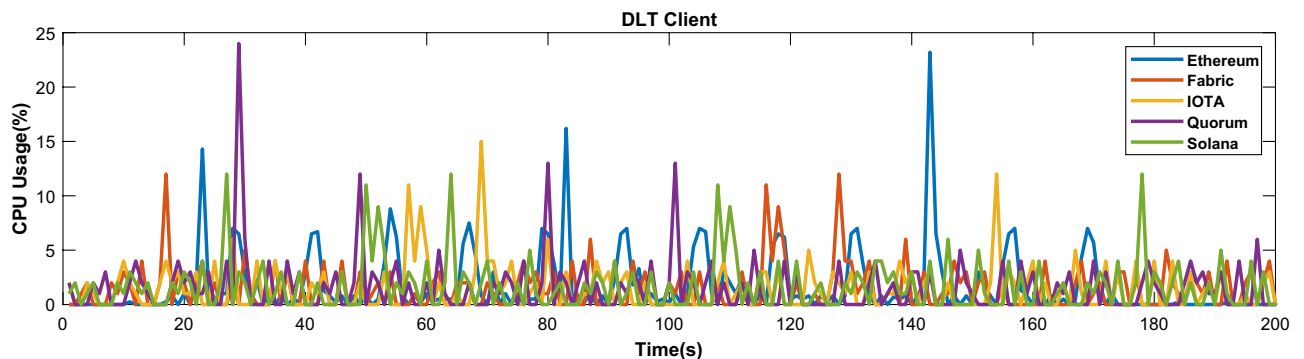
Our shared manufacturing application scenario is described in Fig. 2a. As an initial step (1), the plant operator of a factory publishes the list of machine resources which are available for rent. Thereby, each machine has a unique ID and described capabilities to perform specific jobs. A *manufacturing marketplace* running on a DLT-based network can be implemented in such an environment to offer access to those machine descriptions. In the DLT-based manufacturing network, smart contracts are running to receive requests from customers rent machines (step 2) and match them to resources offered by the plant operator (step 3). In addition, the rules and agreements, e.g., about the rent period, specific tasks, or payment methods between plant operator and customers are pre-defined in the smart contracts and executed autonomously. The customers can check the list of available machines published by the plant operator, and if the customers have a relevant job coming up, they can request the suitable machines via smart contracts. This is the first difference between the DLT-based and non-DLT shared manufacturing system. In a non-DLT based system³⁵, a plant operator and customers could not work directly by exchanging messages without the guarantee about the trust of contracts as well as payment. This guarantee requires a third-party to complete the deal. After DLT-based smart contracts executed and mapped the requests from plant operator and customers, the plant operator account will unlock automatically the available machines (step 4) and assign the control of the machines to customers. Then, the customers can start control and program the machines for their jobs (step 5), which are then executing these jobs (step 6). Compared to standard shared manufacturing, the second innovation in DLT-based systems is that we implemented a *layer 2 payment channel*³⁶ between the plant operator and customer for micro-payments (step 7).

To study the communication and computation overhead resulting from DLT in manufacturing systems, we have implemented the above described *shared manufacturing* application in a private setup as shown in Fig. 2b. The setup involves the DLT components *DLT-manager 1 and 2* and *DLT-clients*. The *DLT-manager* has a high computation capacity and enough storage for a full ledger with all the information and data. The *DLT-clients* are lightweight and are limited in terms of computation and resources. The *DLT-clients* can query and access the data from the ledger without downloading the full chain of blocks. The *DLT-managers* are implemented in two different equipments: as a Siemens Microbox³² and a Macbook Pro. The *DLT-clients* are implemented in Raspberry Pi 3+. The specifications of these devices are found in Table 2. *DLT-manager-1* and *DLT-manager-2* can be connected via Ethernet, and communicate with DLT-clients via local WiFi. The communication method can be extended to other long-range communication or global internet depending on specific scenarios. The distributed ledger is deployed in the DLT Managers. We have implemented five types of DLTs, namely Ethereum, Quorum, IOTA, Hyperledger Fabric, and Solana.

During our evaluation, the DLT-client sends 10 transactions per second to the ledger, which is hosted by the DLT-manager-1 and -2. The reason why we set 10 transactions per second is that traffic from UR5 robots to a



(a) Energy Consumption of *DLT-managers* running full DLT features in separated scenarios.



(b) Energy consumption of *DLT-clients* which transmit data to *DLT-managers* and synchronize data when needed.

Figure 3. Computation overhead of each network component of the 5 studied DLTs namely Ethereum, Hyperledger Fabric, IOTA, Quorum, and Solana.

server is usually low at around 1–5 updates per hour or per day, depending on the specific scenario. Therefore, we stress the network with 10 TPS to test the efficiency of the system. In addition, we assume that the time required for completing a task with a robot could be a day or a week, so the robot could update the task progress and sensing information in a period of minutes or hours is sufficient. By recording the CPU usage percentage of the DLT-specific processes, we have observed the computation overhead in each case of the 5 selected DLT platforms. Looking at the DLT Managers, we have found Ethereum as an outlier, as it requires by far the most computation time of around 85% of CPU as shown in Fig. 3a due to the usage of the Proof of Work (PoW) for the consensus and verification process. The *non-PoW* DLTs, Solana, Hyperledger Fabric and Quorum, require in our private network setting only around 1–3% CPU usage in both *DLT Manager 1* and *DLT Manager 2*. Similarly, the IOTA platform uses *PoW* only rarely in order to prevent spam attacks, so the CPU usage of the DLT Managers is relatively low, similar to Hyperledger Fabric and Quorum. On the DLT Client component, the CPU usage is primarily the generation and transmission of transactions, so that these DLTs require around 5–10% CPU usage of the Raspberry Pi as shown in Fig. 3b.

Figure 4a,b show the communication overhead of the five different DLTs in our shared manufacturing setup. The Hyperledger Fabric produces more network traffic than the others on the DLT Managers. The reason is that the network architecture of Hyperledger Fabric is optimized for an enterprise environment with high security requirements, where the raw data need to be formatted for signed transaction proposals, then going through the complex endorsement and validation process, before attachment to the Blockchain. This process introduces more communication overhead. IOTA produces the lowest traffic on the DLT Managers thanks to the design based on the Tangle²². Specifically, the interconnected Tangle infrastructure does not require total verification across the whole ledger. Instead, all parties are verifying simultaneously and, as a result, the energy and time required to complete transactions are shortened. In addition, Tangle's verification process purports to ensure that there are no duplicate transactions that would lead to double-spending. On the *DLT-clients*, the communication overhead is mainly coming from publishing the collected data via formatted transactions to DLT managers. In specific, a single transaction in IOTA consists of 2673 trytes which is equivalent to 1589 bytes, if encoded, a Ethereum transaction includes around 109 bytes of header, and no limited metadata, Hyperledger Fabric transaction sizes depends on the type of transactions, for example, 3.06 kB for spend and 4.33 kB for mint¹⁹. The overhead of a Solana transaction includes 64 bytes signature and maximum 1232 bytes for given metadata.

In order to analyze the scalability of the integration of Blockchains in manufacturing, we simulate Blockchain networks with a varying number of *DLT managers* and *DLT clients*, from 4 to 16 nodes. The number of

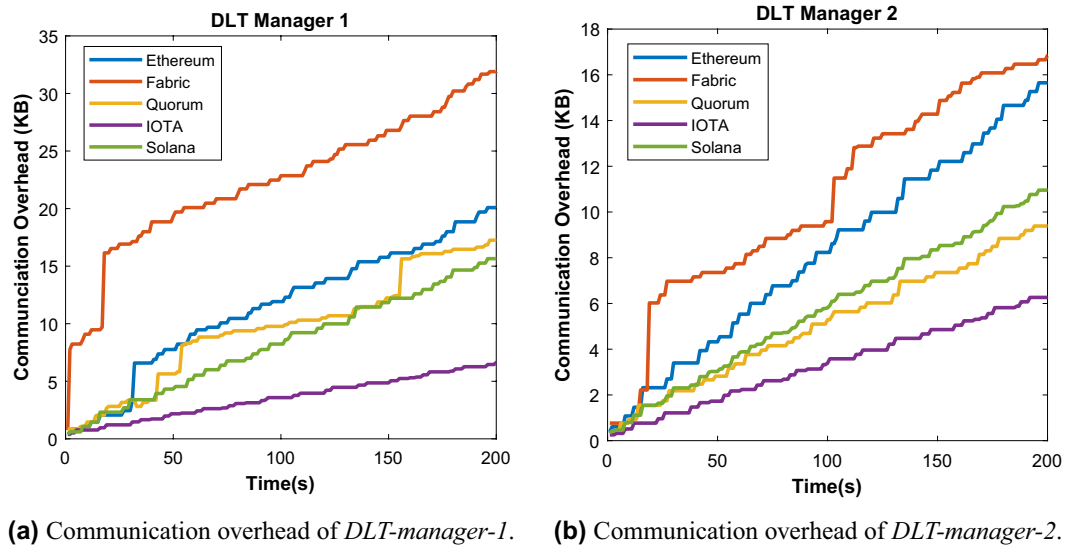


Figure 4. Communication overhead comparison of the 5 studied DLTs namely Ethereum, Hyperledger Fabric, IOTA, Quorum, and Solana.

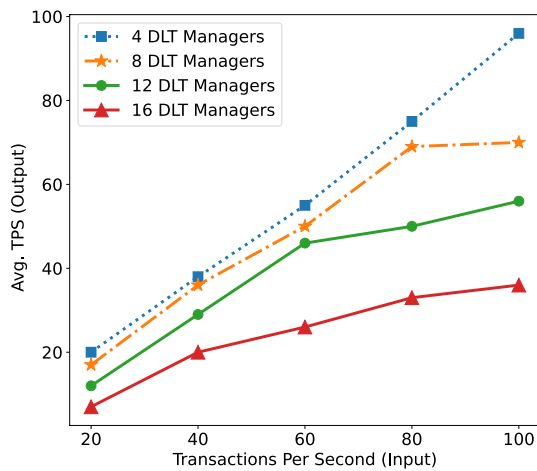


Figure 5. Impact of number of DLT-managers on the average transactions per second.

input transactions per second is generated from 20 TPS to 100 TPS. In this experiment, we choose to simulate the Ethereum Blockchain network. The simulation is extended from the framework³⁷ of Blockchain for Swarm Robotics. Specifically, both *DLT-manager* and *DLT-client* is executed in separated Docker Containers. The simulated robots are considered as light Blockchain clients and periodically published data to the managers for validation. The results in Fig. 5 show that the more the number of *DLT-managers* increases, the more the number of transactions validated per second decreases. The reason is that the more number of validators leading to the more number of transactions to be validated and exchanged among managers, as well as adding the system some delays in validating transactions. We observe that increasing the number of participants raises the challenges e.g. scalability, throughput, latency. Hence, private Blockchains are suitable for industrial IoT manufacturing environments. However, hybrid DLTs which provides flexibility on data visibility without compromising security also could be immense potential.

Based on the CPU usage and the utilized computing hardware, we determined the energy consumption of the five selected DLT platforms and computed the carbon footprint. We assume that electricity for running the computational operations is consumed and produced in Germany. As a measure of carbon intensity of the German energy mix, calculated in a life cycle perspective, we used data from the life cycle database ecoinvent v.3 cutoff system model³⁸. In particular, the dataset *Market for electricity, low voltage, DE* was chosen, which represents an average low-voltage energy mix for Germany. We obtained the life cycle impact of producing 1 kWh electricity according to this version of the database via the software *SimaPro* and using the default IPCC Global Warming Potential (GWP) method with a time horizon of 100 years³⁹. This resulted in a value of global warming impact of 0.540 kg CO₂-eq/kWh that represents the impact of all greenhouse gases emitted in the electricity production

Platform	Power of machine (kW)	Energy consumed on average (kWh)	Avg. CPU usage for blockchain operation (%)	Energy consumed for blockchain operation (kWh)	Greenhouse gas (GHG) emission in DE (kg CO ₂ -eq/kWh)**	GHG emission per blockchain operation (kg CO ₂ -eq)
Hyper.Fabric	0.06	0.06	0.625%	375×10^{-6}	0.540	203×10^{-6}
Ethereum	0.06	0.06	82.35%	$49,392 \times 10^{-6}$	0.540	$26,682 \times 10^{-6}$
Quorum	0.06	0.06	0.65%	390×10^{-6}	0.540	211×10^{-6}
IOTA	0.06	0.06	0.61%	366×10^{-6}	0.540	198×10^{-6}
Solana	0.06	0.06	0.61%	366×10^{-6}	0.540	198×10^{-6}

Table 3. Carbon FootPrint of private DLT testbed with 5 DLTs calculated in Germany market running per hour.

process and upstream activities in a life cycle perspective. This value was further used to calculate the total carbon footprint of the computation based on its energy requirements. The results are shown in Table 3. We observe that the annual CO₂ generated through PoW consensus is significantly higher than that of non-PoW Blockchains. The private Ethereum DLT produced around $26,692 \times 10^{-6}$ kg CO₂-eq/h, which is equivalent to the average of 4.3 charged smartphones⁴⁰. This compares to around 203×10^{-6} kg CO₂-eq/h, 211 kg CO₂-eq/h, and 198 kg CO₂-eq/h from Hyperledger Fabric, Quorum, and IOTA, and Solana respectively. Note that all results are extrapolated to the utilization of our private DLT setup for the shared manufacturing application over an operation day.

Discussion

We have seen that deploying a DLT in an industrial manufacturing environment allows to realize novel business cases. Choosing the right DLT platform for industrial use cases, such as the one we elaborated above, is challenging, as there are many options available. Therefore, we have conducted here an evaluation of five of the most popular and promising DLT platforms and proposed how to integrate those into a physical manufacturing system.

A clear observation is that Ethereum is an outlier in terms of CPU usage, due to its PoW consensus algorithm, which of course also results in high energy consumption. Therefore, we can conclude that Ethereum and other PoW DLTs should, in general, not be used in the envisioned manufacturing environments. In order to still be able to use many of Ethereum development tools, a plant operator can use Quorum, which is an enterprise version of Ethereum. Both Quorum and Hyperledger Fabric show a similar performance in our local evaluation regarding CPU usage. However, Hyperledger Fabric introduces higher communication overhead as compared to Quorum. Therefore, in an environment with communication restrictions, the operator could opt for Quorum out of these two by simply looking at slight performance advantages. IOTA, which is specifically designed for IoT networks, requires the lowest CPU and communication overhead and can hence be favoured by an operator that has strong requirements in this regard. However, IOTA's smart contract mechanism is still under development and also the tooling support is not as strong. Finally, Solana is a public Blockchain network with a focus on achieving high scalability. In our local network, Solana performed similarly to Quorum in terms of communication overhead as well as CPU usage.

The results measured from our local experiments can be considered as a benchmark regarding sustainability aspects in specific shared manufacturing use cases. In the scope of this research, we evaluated the greenhouse gas emission per Blockchain operation based on energy consumed by Blockchain activities. For example, IOTA foundation provided an energy benchmark for the IOTA network, which shows results that are similar to our experimental results⁴¹. In terms of Hyperledger Fabric, we have used Hyperledger Caliper for the benchmark evaluation⁴². Referring to prior research⁴³, the energy consumed by beyond-PoW blockchains, such as Polkadot⁴⁴, Cardano⁴⁵, or Hedera Hashgraph⁴⁶, is within a range that is similar to the energy consumed in our experimental setup.

Looking towards the future, we see many benefits for the use of DLT in manufacturing, enabling a broad range of use cases and business models. The vision at the horizon is a truly *collaborative industrial IoT* in which things (such as machines in a manufacturing plant) ubiquitously and automatically interact without intervention of humans. This is fuelled by the capability to autonomously make (micro-)payments. This would empower devices, e.g., to rent cloud server capacity for additional computational capacity when required, to pay directly to other devices for access to the Internet, or automatically pay for electricity consumed. The current payment systems are not well suited for massive-scale micro-transactions due to high transaction costs and limited capacity. This calls for a vision of payments between things, which will be a small per transaction, but autonomous and running efficiently at a massive scale. Besides, the research on governance, security aspects, and optimal consensus mechanism could be considered as a future work.

Data Availability

The datasets used and/or analysed during the current study available from the corresponding author on reasonable request.

Received: 23 February 2022; Accepted: 17 October 2022

Published online: 27 October 2022

References

- Kafle, B. *et al.* Topcon-technology options for cost efficient industrial manufacturing. *Solar Energy Mater. Solar Cells* **227**, 111100 (2021).
- Soret, B. *et al.* Learning, computing, and trustworthiness in intelligent IOT environments: Performance-energy tradeoffs. *IEEE Trans. Green Commun. Netw.* **6**, 629 (2021).
- Helu, M., Morris, K., Jung, K., Lyons, K. & Leong, S. Identifying performance assurance challenges for smart manufacturing. *Manuf. Lett.* **6**, 1–4 (2015).
- Al-Jaroodi, J. & Mohamed, N. Blockchain in industries: A survey. *IEEE Access* **7**, 36500–36515 (2019).
- Tschorsch, F. & Scheuermann, B. Bitcoin and beyond: A technical survey on decentralized digital currencies. *IEEE Commun. Surv. Tutor.* **18**, 2084–2123 (2016).
- Alrebd, N., Alabdulatif, A., Iwendi, C. & Lian, Z. Svbe: Searchable and verifiable blockchain-based electronic medical records system. *Sci. Rep.* **12**, 1–11 (2022).
- Top 10 security predictions (2016). <https://www.gartner.com/smarterwithgartner/top-10-security-predictions-2016>. Accessed on 12/06/2021.
- Nguyen, L. D., Pandey, S. R., Beatriz, S., Broering, A. & Popovski, P. A marketplace for trading AI models based on blockchain and incentives for IOT data. arXiv preprint [arXiv:2112.02870](https://arxiv.org/abs/2112.02870) (2021).
- Chen, B. *et al.* Edge computing in IOT-based manufacturing. *IEEE Commun. Mag.* **56**, 103–109 (2018).
- Hassan, W. H. *et al.* Current research on internet of things (IOT) security: A survey. *Comput. Netw.* **148**, 283–294 (2019).
- Shi, W., Cao, J., Zhang, Q., Li, Y. & Xu, L. Edge computing: Vision and challenges. *IEEE Internet Things J.* **3**, 637–646 (2016).
- Economy of things|Bosch global. <https://www.bosch.com/research/know-how/success-stories/economy-of-things-a-technology-and-business-evolution/>. Accessed on 12/06/2021.
- Blockchain IOT|exclusive content for the food and beverage industry|siemens global. <https://new.siemens.com/global/en/markets/food-beverage/exclusive-area/blockchain-iot.html>. Accessed on 12/06/2021.
- Li, Z., Barenji, A. V. & Huang, G. Q. Toward a blockchain cloud manufacturing system as a peer to peer distributed network platform. *Robot. Comput. Integr. Manuf.* **54**, 133–144 (2018).
- Danzi, P., Kalor, A. E., Stefanović, C. & Popovski, P. Delay and communication tradeoffs for blockchain systems with lightweight IOT clients. *IEEE Internet Things J.* **6**, 2354–2365. <https://doi.org/10.1109/JIOT.2019.2906615> (2019).
- Fan, C., Ghaemi, S., Khazaei, H. & Musilek, P. Performance evaluation of blockchain systems: A systematic survey. *IEEE Access* **8**, 126927–126950. <https://doi.org/10.1109/ACCESS.2020.3006078> (2020).
- Fu, B., Shu, Z. & Liu, X. Blockchain enhanced emission trading framework in fashion apparel manufacturing industry. *Sustainability* **10**, 1105 (2018).
- Yu, C., Zhang, L., Zhao, W. & Zhang, S. A blockchain-based service composition architecture in cloud manufacturing. *Int. J. Comput. Integr. Manuf.* **33**, 701–715 (2020).
- Androulaki, E. *et al.* Hyperledger fabric: A distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference*, 1–15 (2018).
- Baliga, A., Subhod, I., Kamat, P. & Chatterjee, S. Performance evaluation of the quorum blockchain platform. arXiv preprint [arXiv:1809.03421](https://arxiv.org/abs/1809.03421) (2018).
- Vujičić, D., Jagodić, D. & Randić, S. Blockchain technology, bitcoin, and Ethereum: A brief overview. In *2018 17th International Symposium Infoteh-Jahorina (Infoteh)*, 1–6 (IEEE, 2018).
- Popov, S. & Lu, Q. Iota: Feeless and free. IEEE Blockchain Technical Briefs (2019).
- Yakovenko, A. Solana: A new architecture for a high performance blockchain v0. 8.13. Whitepaper (2018).
- Sguanci, C., Spatafora, R. & Vergani, A. M. Layer 2 blockchain scaling: A survey. arXiv preprint [arXiv:2107.10881](https://arxiv.org/abs/2107.10881) (2021).
- Wüst, K. & Gervais, A. Do you need a blockchain? In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 45–54 (IEEE, 2018).
- Boneh, D., Bonneau, J., Bünz, B. & Fisch, B. Verifiable delay functions. In *Annual International Cryptology Conference*, 757–788 (Springer, 2018).
- Iota smart contracts beta release. <https://blog.iota.org/iota-smart-contracts-beta-release/>. Accessed on 05/04/2022.
- Homepage|solana docs. <https://docs.solana.com/>. Accessed on 01/14/2022.
- Xiao, Y., Zhang, N., Lou, W. & Hou, Y. T. A survey of distributed consensus protocols for blockchain networks. *IEEE Commun. Surv. Tutor.* **22**, 1432–1465 (2020).
- Ahl, A. *et al.* Exploring blockchain for the energy transition: Opportunities and challenges based on a case study in japan. *Renew. Sustain. Energy Rev.* **117**, 109488 (2020).
- Xiong, Z., Zhang, Y., Niyato, D., Wang, P. & Han, Z. When mobile blockchain meets edge computing. *IEEE Commun. Mag.* **56**, 33–39 (2018).
- Ur5 collaborative robot arm|flexible and lightweight cobot. <https://www.universal-robots.com/products/ur5-robot/>. Accessed on 01/26/2022.
- Nguyen, L. D., Leyva-Mayorga, I., Lewis, A. N. & Popovski, P. Modeling and analysis of data trading on blockchain-based market in IOT networks. *IEEE Internet Things J.* **8**, 6487–6497 (2021).
- Storch, D.-M., Timme, M. & Schröder, M. Incentive-driven transition to high ride-sharing adoption. *Nat. Commun.* **12**, 1–10 (2021).
- Renting machine made easy-mcpond. <https://mcpond.com/>. Accessed on 12/07/2021.
- Jourenko, M., Kurazumi, K., Larangeira, M. & Tanaka, K. Sok: A taxonomy for layer-2 scalability related protocols for cryptocurrencies. *IACR Cryptol. ePrint Arch.* **2019**, 352 (2019).
- Strobel, V., Castelló Ferrer, E. & Dorigo, M. Blockchain technology secures robot swarms: A comparison of consensus protocols and their resilience to byzantine robots. *Front. Robot. AI* **7**, 54 (2020).
- Wernet, G. *et al.* Theecoinvent database version 3 (part i): Overview and methodology. *Int. J. Life Cycle Assess.* **21**, 1218–1230 (2016).
- Stocker, T. F. *et al.* *Ipcc, 2013: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (2013).
- Greenhouse gas equivalencies calculator|us epa. <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator>. Accessed on 12/07/2021.
- Energy benchmarks for the iota network (chrysalis edition). <https://blog.iota.org/internal-energy-benchmarks-for-iota/>. Accessed on 02/11/2022.
- Hyperledger caliper-hyperledger foundation. <https://www.hyperledger.org/use/caliper>. Accessed on 02/11/2022.
- Platt, M. *et al.* Energy footprint of blockchain consensus mechanisms beyond proof-of-work. arXiv preprint [arXiv:2109.03667](https://arxiv.org/abs/2109.03667) (2021).
- Wood, G. Polkadot: Vision for a heterogeneous multi-chain framework. *White Paper* **21**, 2327–4662 (2016).
- Cardano|home. <https://cardano.org/>. Accessed on 02/11/2022.
- Baird, L., Harmon, M. & Madsen, P. Hedera: A governing council and public hashgraph network. *Trust Layer Internet, Whitepaper* **1**, 1–97 (2018).

Acknowledgements

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 957218 (Project IntelliOT), and in part supported by the Villum Investigator Grant "WATER" from the Velux Foundation, Denmark.

Author contributions

L.D.N and A.B. proposed the idea, implemented the system, and drafted the main manuscript. P.P revised manuscript and derived the structure for the research. M.P analyzed the sustainability aspect regarding CO₂ emission of the research. All authors contributed, edited, and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to L.D.N.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022