The Road to Accountable and Dependable Manufacturing

Jan Pennekamp
Communication and Distributed Systems, RWTH Aachen University, Aachen, Germany

Roman Matzutt
Communication and Distributed Systems, RWTH Aachen University, Aachen, Germany

Salil S. Kanhere
School of Computer Science and Engineering, University of New South Wales, Sydney, Australia

Jens Hiller
Communication and Distributed Systems, RWTH Aachen University, Aachen, Germany

Klaus Wehrle
Communication and Distributed Systems, RWTH Aachen University, Aachen, Germany

***Abstract*—In manufacturing, advances from the IoT foster the vision of a highly dynamic and interconnected Industrial IoT. However, business-driven use cases mandate different levels of security, privacy, accountability, and verifiability alike. Blockchain technology addresses these requirements and thereby enables previously unforeseen collaborations. The authors emphasize the need for active research at the intersection of IoT, CPS, and blockchain.**

MANUFACTURING is expected to significantly benefit from recent advances in the areas of Internet of Things (IoT) and Cyber-Physical Systems (CPS). Particular development directions include establishing highly-dynamic business relations and creating interconnected production environments, even for short-lived collaborations, through increasing degrees of automation based on (sensor) data [1]. Concepts of the Industrial IoT (IIoT) or Internet of Production (IoP) [2] explicitly target to implement these improvements.

Research mainly evolves around three existing pillars (P0–P2): **(P0)** CPS and site-related improvements (↻) with limited external influences, **(P1)** extended data sharing along the supply chain (⇄), e.g., to reduce the bullwhip effect, and **(P2)** secure industrial collaborations across supply chains (↑↓), e.g., to reduce ramp-up costs. To achieve **P1** and **P2** not only with today's (established) long-term trust but also in settings with dynamically evolving and flexible short-term relationships, we identify a new research pillar **(P3)** that enables accountable and dependable dataflows for stakeholders without any trusted or previous relationships (❖). In this article, we focus on the research pillars P1–P3 that consider multiple stakeholders in collaborative processes.

Such industry-driven settings mandate special needs that traditional solutions in the IoT cannot satisfy. These aspects encompass improved accountability and verifiability to deal with uncertainty concerning the origin [3] and reliability of data [4], but also security and privacy requirements have to be considered as information leakage can have tremendous consequences in highly competitive environments [2]. We envision that the consequent integration of blockchain technology provides these desired features by design. Its tamperproofness offers verifiability and reliability once information has been recorded on

the blockchain. Similarly, blockchains are decentralized and thus well-suited for securing interactions among mutually distrustful parties. Finally, the extensible nature of blockchain technology enables scalability features, such as sidechains or sharding [5], as needed for solutions across different use cases and domains.

Given that research at the intersection of IIoT and blockchain is still in its infancy, we identify three key research areas. We discuss *blockchain-specific research questions* for the industrial setting, which mainly evolve around the general scalability of proposed solutions and the privacy of participants. Similarly, we identify a lack of manufacturing-specific solutions that integrate blockchains to improve accountability in this domain. We discuss *scenario-driven research directions* that close this gap and realize fast, versatile, accountable, and dependable manufacturing enabled by blockchains. Furthermore, we discuss arising *socio-economic challenges*. Particularly, new legal frameworks will need to take into account the increased usage of external data, potentially in safety-critical applications. First and foremost, however, we want to raise awareness on how to establish trust into the authenticity and correctness of data on the blockchain as a foundation for interorganizational data sharing within the IIoT.

MOTIVATION & POTENTIALS

Manufacturing is expected to compile vast amounts of process and product data in the near future [2]. Consequentially, we have to deal with associated big data challenges that stand out due to virtually infinite volumes of available sensor data and the increased need for high-frequency sensing [1]. However, big data also provides opportunities when properly extracting its encapsulated knowledge [1]. Regarding manufacturing, this potential has previously been neglected for lack of globally available process information, and even data sharing along the supply chain was limited. Figure 1 illustrates the data sharing along (⇄) and across (↑↓) supply chains, which we detail hereafter based on two fine blanking lines.

Information Sharing along Supply Chains (P1 ⇄)

Traditionally, supply chain data sharing was driven by large companies dictating their requirements to all suppliers. In this setting, information was collected in data sinks accessible by single (large) players [4], e.g., automotive manufacturers. Furthermore, due to privacy concerns, data is usually shielded from external stakeholders, for example, even rather insensitive information, such as delivery schedules or shipment tracking, is retained locally. Today, additional data is only shared under the promise of large financial impacts despite production data being expected to improve manufacturers' productivity and overall product quality [2].

This situation is unsatisfactory as it fails to address several desired aspects. Especially regarding legislation, today's landscape cannot reliably provide (long-term) verifiability of relevant information [6], e.g., provenance data for parts in the aerospace industry or associated maintenance protocols. Although additional processes are often in place, counterfeit or non-fair trade products are, occasionally, still entering legitimate supply chains [7]. To improve the reliability of (received) data, we envision technical solutions that minimize the room for manipulations and provide an efficiently verifiable certification for each individual product. Furthermore, a unified approach could improve governmental oversight, which is especially desirable for safety-critical products or food chains [8].

Another insufficiency stems from the lacking identifiability of root causes of manufacturing or product failures [6]. Currently, accountability is mostly limited to contractually-bound stakeholders. If not explicitly contractually negotiated, individual untrusted suppliers may remain passive or even behave adversely for their benefits, e.g., when covering up incidents. Simultaneously, missing feedback to estimate the lifetime or the fit of a product, which both might depend on the application, hinders the implementation of improvements. To overcome such limits, accessible production and usage data can provide insights [2].

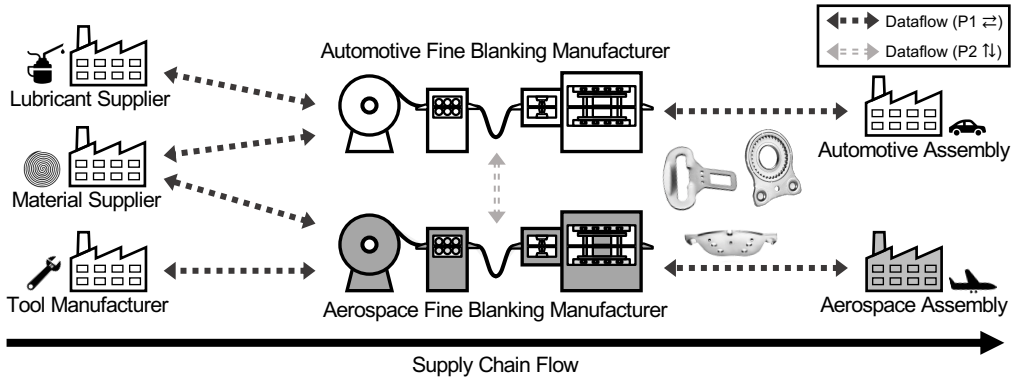

Figure 1: Manufacturing engulfs both dataflows along the supply chain (P1 ⇄)and across supply chains (P2 ⇅). Suppliers (here: for lubricants, material, and tools) support manufacturers who themselves provide subsequent assembly lines with production data. Similarly, manufacturers exchange process information (here: fine blanking lines), the processed material, and their interplay. A currently non-existing relationship between both assembling companies could be non-existent due to the untrusted environment (P3 ❖). We adapted the figure from our analysis of dataflows in an *Internet of Production* [2].

Foundations for Expanded Secure Industrial Collaboration Across Supply Chains (P2 ⇅)

In addition to the marginal data sharing along supply chains (P1 ⇄), data exchanges across supply chains (P2 ⇅)are basically non-existing in today's manufacturing landscape [1]. While manufacturers gather usage data from their customers (in centralized data silos), virtually no knowledge exchange happens between different operators of (identical) machines [2]. For example, experiences with used machine configurations or information about the (expectable) production quality can reveal interesting insights into newly configured manufacturing processes. Hence, all knowledge is retained locally without global availability, despite potentially tremendous benefits [2].

To improve productivity and to decrease costs, companies could, for instance, share ideal machine configurations for their workpieces, e.g., within their fine blanking line, without revealing all details to the machine supplier. Furthermore, this information exchange may reduce ramp-up times of new manufacturing processes by deriving machine parameters from readily available information (cf. Figure 1). Consequentially, non-competing companies can cooperate and jointly assemble a shared knowledge base in a give-and-take manner or offer their valuable data for sale. As of today, a lot of expected potential is still unexplored.

Ad-Hoc Relationships in Untrusted Environments (P3 ❖)

When considering relationships with previously unaffiliated and thus untrusted companies (P3 ❖), several additional use cases emerge. Along supply chains (P1 ⇄), identifying the ideal supplier for a component is simplified when the utilization of relationships among previously unaffiliated parties is improved. Similarly, exchanging information with companies in related domains across supply chains (P2 ⇅)is currently hindered by a lack of trust between the involved stakeholders. We expect that more use cases surface once the first steps towards secure industrial collaboration have been taken as businesses are naturally cautious when sharing sensitive and valuable details, especially production and product data [2]. Furthermore, we observe that currently no uniform standardization for data sharing exists, which especially hinders flexible relationships as company-specific adjustments are required for each new partner [4].

In the context of accountable and dependable manufacturing, we also have to address privacy and safety [9]. Appropriate means are not yet available, or they are not proven or tested in manufacturing [1]. A major milestone to establish trust can be achieved by providing accountability,

verifiability, and transparency for all actions and traded information. Consequentially, blockchains are a promising tool to establish trust in mutually distrustful manufacturing markets and to eventually allow for interorganizational data sharing and novel applications.

THE INFLUENCE OF BLOCKCHAINS

Blockchain systems have matured considerably since their introduction through Bitcoin in 2008. Initially created for the decentralized, yet secure, management of digital currency, the potential of blockchains for larger and more diverse tasks was quickly identified across academia and industry.

The State of Blockchain Integration

We now reiterate impactful milestones and applications of distributed ledger technology to assess its current level of integration into business processes and to identify areas where blockchains have been applied successfully.

Financial Origins Bitcoin paved the way for global financial transactions without banks as intermediaries. Besides inspiring numerous comparable cryptocurrencies, the banking sector also noticed the potential of blockchains to improve transactions between financial institutes. This development yielded major blockchain-based interbank networks, e.g., the Ripple payment and exchange network or JP Morgan's Interbank Information Network. Furthermore, blockchains promise to provide better, i.e., more direct, customer experience at lower costs due to more automated, disintermediated processes. Especially in scenarios where participants are known, and their majority is trusted, consortium blockchains are seen as key enablers for shaping new transaction processes in highly distributed applications, e.g., accounting in supply chains.

Digital Assets One of the first non-cryptocurrency applications of blockchains was the establishment of digital assets and notary services. While dedicated solutions, such as Namecoin, were launched quite early, numerous such services piggyback on existing blockchains, commonly Bitcoin [10]. Particularly, to transfer digital ownership of property, coupons, or stock-marketing shares through a cryptocurrency's blockchain, users can tie assets to blockchain transactions. Beyond that, notary services immutably attest the existence of documents by storing a cryptographic hash on a blockchain, a tamperproof identifier to which owners can subsequently refer to.

Process Automation *Smart contracts* [5] realize the automated execution of transactions once the blockchain's state satisfies their one-time programmable conditions. This tamperproof programmability allows for transparent automation of global processes. While Ethereum popularized blockchain-based smart contracts, business applications are commonly built using consortium blockchains, e.g., created through *Hyperledger Fabric* or the Ethereum-compatible *Quorum*. Beyond the banking sector, insurers process insurance claims without human interaction through smart contracts. An increased demand for blockchain-based process automation sparked the creation of Blockchain-as-a-Service solutions, e.g., offered by Microsoft Azure, IBM, and Amazon Web Services. These services lower the barrier for creating blockchain-backed architectures, but also introduce an infrastructure provider as a new *centralized* entity.

Internet of Things Advances in process automation proliferated the vision of coupling autonomous IoT devices with blockchains. The main advantages of blockchain-based IoT infrastructures lie in the immutable and decentralized IoT-based sensing of physical environments in conjunction with the accountable recording of actuation events. If seized well, these capabilities can significantly simplify applications for smart cities, e.g., smart microgrids [11] or vehicular networks [12]. Here, blockchains aid trust management and access control to sensed data alike.

Supply Chain Blockchains may be used as an architectural pillar for reshaping supply chains [13], [7], [6], [14], especially due to improved financial transactions, asset management, process automation, and data management. However, smooth integration is still lacking [9]. TrustChain [8] or ProductChain [3] al-

ready tackle important issues of supply chain deployments, such as reputation-based trust management among suppliers and provenance tracking for customers. Still, holistic, all-encompassing approaches to improve supply chains based on distributed ledgers are yet to come.

Useful Properties for Diverse Applications

Even today's limited integration of blockchain technology into business processes highlights that distributed ledgers have proved to provide *valuable foundations* for various domains, applications, and use cases. Particularly, we highlight that blockchain technology provides desirable contributions to flexible collaborations and especially to applications involving supply chains. First, the *decentralized* nature of blockchain applications suits the highly distributed and heterogeneous environments created by collaborating companies and supply chains. Second, blockchains can provide data *integrity* and *verifiability* even if collaborators are partially distrusting each other. As part of this process, recorded data is kept on a *tamperproof* ledger. Finally, established measures to keep track of digital assets and to prevent double-spending enable the public, transparent *traceability* of products or their components. However, the decentralization and immutability of blockchains creates issues that were not present in traditional business processes. Next, we thus dive into resulting challenges that, once tackled, will help realize suitable full-stack solutions for improving business processes via distributed ledgers.

OPEN RESEARCH AREAS

We identify three layers of open research areas that we illustrate in Figure 2: **(L1)** yet unaddressed challenges for the use of blockchain technology in manufacturing, **(L2)** new opportunities for a fast, versatile, accountable, and dependable manufacturing enabled by blockchains, i.e., scenario-driven challenges, and **(L3)** socio-economic challenges stemming from immutably recorded production data and highly flexible cross-company collaborations. We consider these layers to be highly relevant when shaping the future of interconnected manufacturing.

Open Blockchain-Inherent Challenges (L1)

As groundwork for more scenario-specific research, we identify blockchain-induced research areas that surface when relying on blockchains for accountable and dependable manufacturing.

Scalability Permissionless blockchains traditionally struggle with limited scalability in terms of transaction throughput, transaction latency, and storage requirements. For instance, Bitcoin famously has a low transaction rate of only 3.5 transactions per second as its 10-minute inter-block delay requires users to wait for an hour to safely accept payments [5]. Even though consortium blockchains can utilize more efficient consensus algorithms [15], recording large numbers of events on-chain still remains challenging. Solutions may aggregate multiple events into single or few (on-chain) transactions, similar to micropayment channels that boost transaction throughputs in today's cryptocurrencies. Furthermore, applying *sharding schemes* [5] to consortium blockchains may improve their transaction throughput as these schemes target to partition the network and to distribute the responsibility for transaction processing.

Another scalability issue is the ever-increasing storage requirement to operate blockchains. For instance, heavily-utilized blockchains today accumulate hundreds of Gigabytes of historical data. This problem is aggravated in the context of supply chain applications once suppliers are required to tie their reports for other contractors immutably to the blockchain. Pruning strategies have been proposed to unburden blockchain nodes from storing historic transaction data that has become obsolete meanwhile [16]. However, applications relying on blockchain-extrinsic data cannot immediately seize this potential since what constitutes obsolete data has to be defined on a per-application basis. Again, also partitioning data storage across the network with sharding schemes can reduce per-node storage requirements. Overall, future research needs to assess the need for long-term data availability to allow for efficient and scalable solutions.

Efficiency Wide-spread adoption of blockchain technology in supply chains necessitates an efficient operation of the

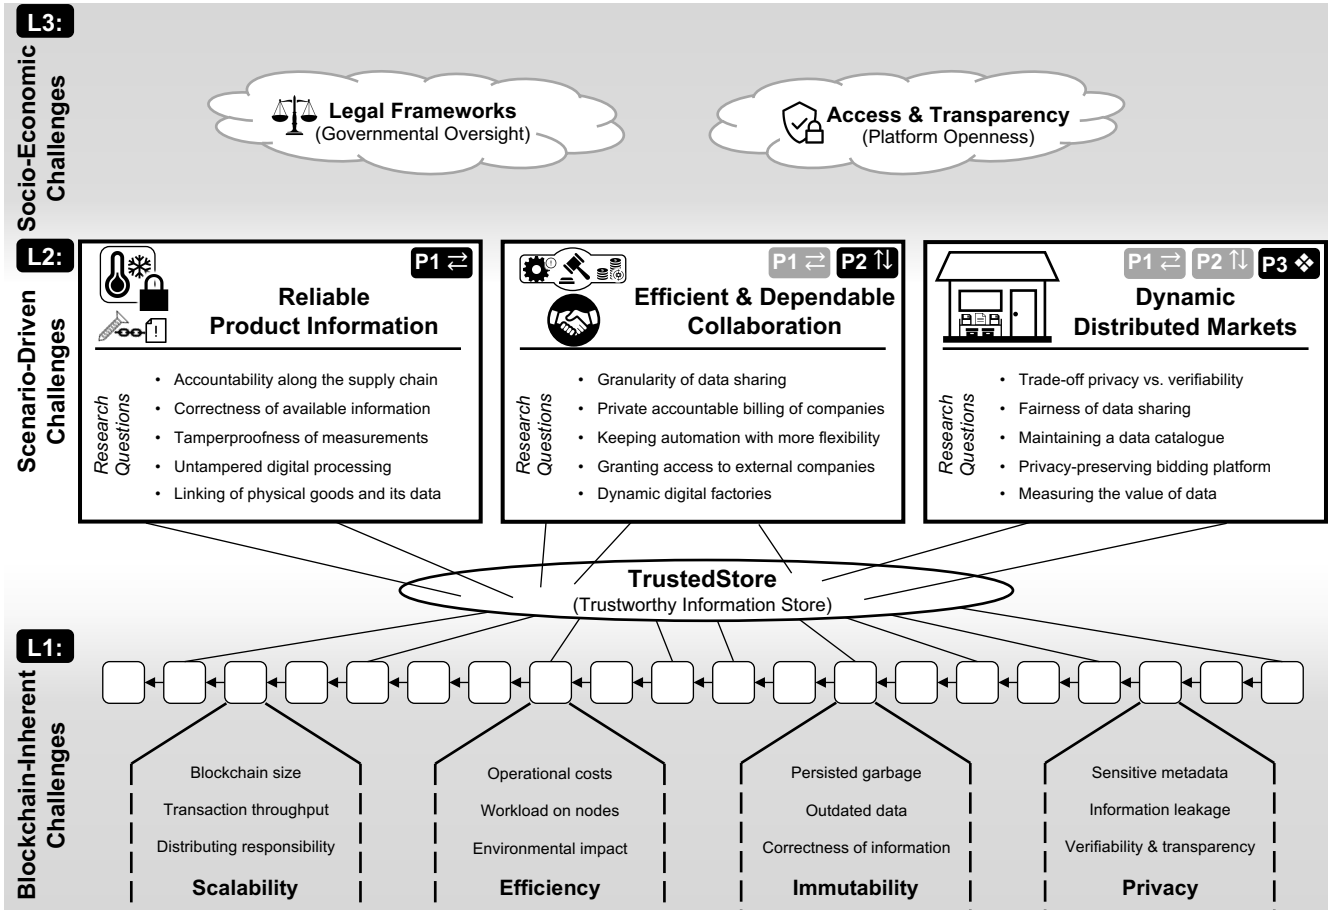

Figure 2: We group research towards accountable and dependable manufacturing into three layers.

L1: Blockchain-inherent challenges that concern the properties of blockchain technology which is expected to serve as an underlying key component of our envisioned TrustedStore.

L2: Scenario-driven challenges that can be grouped into three main research directions that each focus on a specific research pillar, i.e., *along supply chains* (P1 ⇄), *across supply chains* (P2 ↑↓), and *situations with insufficient trust between stakeholders* (P3 ❖).

L3: Socio-economic challenges that have an impact on underlying collaborations and improvements.

To offer viable solutions for accountable and dependable manufacturing, research must consider and tackle all layers and their individual research challenges.

infrastructure. To this end, any proposed architecture must take the deployment and operation costs into account, with a special focus on computing overhead for securely keeping data on-chain. Improvements in efficiency mainly originate from more fundamental lines of research, e.g., advances in authentication, distributed consensus, or secure communication. Yet, a proper integration of these advances into a full blockchain-based architecture is mandatory to seize this potential for efficient data management and to not undermine any requirements of the overall system. The main bottleneck of traditional blockchains is the redundant execution of various tasks, such as verifying digital signatures or maintaining a local state [5]. This redundancy not only increases costs but also creates a potentially avoidable environmental impact. Solutions, such as sidechains or sharding [5], that distribute the workload without lowering security guarantees will help to reduce the operating costs. While these concepts are primarily being researched for public settings, the envisioned high-frequency

utilization and large volumes of data call for similar developments for consortium blockchains.

Immutability Recording events immutably despite the presence of adversaries eager to alter history is arguably the blockchain's key achievement. Thus, storing non-financial, application-specific data on-chain or referencing such data through on-chain fingerprints, has become a frequent proposition [10]. However, this immutability has also proved to create further issues than only impacting the long-term scalability of blockchains, e.g., distributing and storing unwanted blockchain data can cause legal liability [16]. While the prevalence of known identities within consortium blockchain mitigates such risks, different stakeholders may nevertheless be in conflict about the value of recorded data, e.g., whether data is outdated or when unknown raw data formats pollute the shared storage. Overall, the quality of recorded information becomes more important as participants should be able to rely on data that is recorded by other parties that exhibit varying individual levels of trust. Today, a link between a physical (product) property and its digital data is missing, which limits the consensus algorithms' ability to verify claimed events before persisting them on-chain, e.g., sensor readings from inaccessible, remote environments. Correcting identified errors is trivially possible by overwriting data in a new transaction, but implies a more complex transaction processing by all parties. Hence, further research is required to explore the trade-off between data availability and data utility as well as data verifiability and efficient corrections.

Privacy Tightly related to the individual data value for different stakeholders involved in the consortium blockchain is the notion of data privacy, which applies not only to traditional privacy, e.g., storing and trading customer data, but to information leakage in general [16]. On the one hand, blockchains may disclose sensitive business secrets [13], such as capabilities of production machines or process details, e.g., required temperatures or metal alloys, both directly and indirectly. On the other hand, meta-information such as the frequency of transactions between two collaborators or key performance indicators may be inferred, putting affected parties at a disadvantage against competitors, e.g., during price negotiations or when company acquisition is imminent. A key challenge for sustainable consortium blockchains will be carefully gauging the desired level of point-to-point collaborations and consequently tackling arising trust barriers through both trust and data management.

Scenario-Driven Research Directions (L2)

On top of the blockchain-inherent challenges, further research directions may lead to a fast, versatile, accountable, and dependable blockchain-backed manufacturing (cf. Figure 2). Research into (i) *reliable product information* will ensure the availability of high-quality data alongside all production steps of a supply chain (P1 ⇄), ranging from tamperproof sensing to secure blockchain storage. Based on this reliable, high-quality information more (ii) *efficient and dependable collaborations* can form in the future that will increasingly affect dataflows across supply chains (P2 ↑↓). Ultimately, (iii) *dynamic distributed markets* allow for flexible sharing of data and advertising services, especially when stakeholders without any trusted or previous relationships intend to collaborate (P3 ❖). This way, collaborators can efficiently foster fast, versatile, and dependable business relations.

Reliable Product Information Today, large-scale production and supply chains (P1 ⇄) are opaque regarding processes and the origin of processed goods [4]. Consequentially, failure root causes and other issues cannot be tracked down efficiently, creating massive administrative overheads [6], [14], e.g., hampering legal investigations, causing over-dimensioned product recalls, or an inefficient lookup of compatible spare parts for repairs or assembling bigger workpieces. Similarly, feeding back information from mid-term or long-term field experience into manufacturing processes for improvements is hard [2].

To overcome these limitations, manufacturing needs a *reliably accessible*, *tamperproof* information store that links *clearly identifiable* products to their physical state in a *verifiable* manner. For example, the transportation of fresh produce, which must uphold a mandated cold chain, requires the container's temperature to be con-

tinually monitored such that tricking sensors is infeasible [8].

First, this process requires measures to *achieve a tamperproof gathering of physical-state information*. Here, we identify tailored machine learning mechanisms for anomaly detection as promising research area. Such a machine learning algorithm can base on the following data: (i) Using multiple sensors allows for cross-checking gathered data, e.g., sensors redundantly monitoring the container from different vantage points can increase tamper resilience as already subtle monitoring inconsistencies could unveil manipulations. (ii) Similarly, different sensor types and measuring methods further increase the range for sensing correlation to detect anomalies regarding the coherence of real-world physical effects. As sensor nodes cheapen and allow for long-lasting battery-based operation, these solutions are also becoming increasingly economically viable. (iii) Further, high sampling rates also improve tamper resilience, as more readings are available to identify inconsistencies. Overall, the gathered data provides promising input for a machine learning-based anomaly detection.

Still, storing these large amounts of raw data (i–iii) in globally replicated tamperproof storages such as the blockchain remains challenging. Instead, we envision a combination of *mid-term local storages* maintained by companies and a *long-term distributed information store*. In this deployment model, companies store their raw production data locally and signal its availability on-chain via fingerprints. Further, the blockchain stores (small-sized) insights that result from analyses of the locally stored raw data. Likewise, this storage happens in a certified manner, overall creating a *trustworthy information store*, which we refer to as *TrustedStore*. To ensure that companies fully preserve raw data locally, certified service providers (verifiers) periodically check if local stores match with the TrustedStore, so that misbehavior can be detected in a timely manner and appropriately acted upon (legally). As the amount of data renders full-blown checks impracticable from remote locations and on-site checks involve high costs, they have to happen only rarely. In between, verifiers remotely request data for randomly selected fingerprints to frequently, yet economically, check for data availability. Alternatively, companies store raw data in globally distributed certified data stores and prove such storage to the TrustedStore. Overall, decoupling the storage of *large amounts of raw data* from *derived insights and key properties* ensures the immutability and availability of rich raw data while keeping reasonable loads for globally maintained infrastructures.

Second, a *tamperproof digital processing of gathered data* ensures that original sensor readings enter the blockchain-backed TrustedStore correctly. This way, data can be collected even from untrusted or hostile environments, e.g., to realize new collaborations without sufficient trust levels. *Tamperproof sensors* can provide this form of dependable data gathering and processing [17]. Such devices combine traditional sensors, e.g., RFID scanners, or temperature or humidity sensors [18], with trusted computing mechanisms, such as hardware security modules (HSMs). These security-enhanced sensors are able to immediately hand over data to HSMs for processing, thereby minimizing the attack surface for tampering. Ultimately, the HSM uploads the sensor readings to the local storage and stores their fingerprints on the TrustedStore. From this point on, the reliably-sensed data is persisted immutably.

Assuming mechanisms for tamperproof sensing and blockchain inclusion, we finally must clearly link these readings to the respective physical products, e.g., via camera tracking, RFID tags, imprints, or other markings. Importantly, this identification must also be tamperproof, using suitable mechanisms as described before.

In summary, this research will yield a reliably accessible, tamperproof TrustedStore for production data to establish *accountability along any supply chain*. Beyond aiding legal investigation, managing product recalls, and optimizing parts utilization, this TrustedStore can further serve as a medium to foster collaborations among well-known and novel companies alike.

Efficient and Dependable Collaboration

Established business relations with trust in place can increase their efficiency with a dependable TrustedStore. This claim especially holds for dataflows across supply chains (P2 ↑↓) that could improve the productivity in manufacturing qual-

ity [2]. Additionally, sharing workpiece data, production machine schedules, and states in a timely manner enables close collaborations, accumulating companies into *digital factories* with production efficiencies similar to single, multi-factory companies. Rich information flows allow for a cross-company allocation of machine time and flexible handling of process deviations [2], e.g., by automatically reallocating machine capacity in case of delays. Here, the TrustedStore enables trustworthy tracking methods for workpieces along the full (multi-factory) supply chain. As a result, problems can easily be tracked, and clearly assigned responsibilities motivate participants to comply with their obligations. Most basically, this information allows for detecting infringements early on, e.g., misconfiguration or maintenance backlogs.

Beyond supply chain management, TrustedStores simplify the billing of goods or machine usage (Manufacturing-as-a-Service) [2]. Especially with production environments shifting from generic mass production to individual products, companies require verifiable and highly automated payment processes to keep administrative burdens at a reasonable level. Even pay-as-you-go contracts for cost-efficient machine usage in adaptive production are conceivable where customers pay only for the resources and energy required to create the requested (potentially low-quantity) workpieces. Thereby, high degrees of automation enable manufacturers to maintain a high utilization as multiple customers can share single machines with almost no downtime.

Managing data from mid-term and long-term field experience on the TrustedStore promises further benefits. In contrast to the previously discussed less sensitive product data, the process data considered here is more valuable and, thus, must be protected accordingly. Nowadays, information on product life cycles, required maintenance intervals, or production quality variations is exclusively accessible to the manufacturer. Using the TrustedStore, such data becomes accessible to current and prospective machine users alike (cf. Figure 1). Here, the TrustedStore provides evidence of data correctness. Data of individual machines further facilitates reselling as prior usage and output quality become assessable.

Research has to answer questions on the required granularity of sharing data to achieve these envisioned benefits. As business secrets are potentially at risk when providing information to external, partially trusted collaborators [2], companies have to make informed decisions when trading off efficiency and profit for data privacy.

Dynamic Distributed Markets Ultimately, we envision (distributed and transparent) blockchain-based bidding platforms that realize fast, versatile, yet dependable markets for *goods*, *services* (e.g., machine rentals), and *configuration knowledge*, especially fostering collaborations between–previously unknown and potentially untrusted–business partners (P3 ❖). Today's business relations typically evolve over long periods and trust builds up slowly or is enforced through complex contracts. Blockchains can largely substitute social trust through technical guarantees and thus foster the establishment of new business relations. Furthermore, a distributed TrustedStore allows for efficient automation, e.g., the allocation of machine time, achieving high utilization even in adaptive manufacturing processes. Consequentially, manufacturers can generate profit even from short-time business relations for single workpieces, which would otherwise be uneconomical and incur high risks.

Customers can search for the best-matching offer and benefit from reasonable prices due to increased market competition. Especially smaller manufacturers can profit from low-barrier market access to appeal to customers and business partners and easily increase (domain) knowledge through the TrustedStore.

However, the realization of these distributed markets faces a big challenge, i.e., the potential disclosure of business secrets. For example, big companies could exploit the TrustedStore's information to suppress competitors, e.g., by engaging in well-informed price dumping. Thus, a fundamental research question is how to match business partners based on desired capabilities and quality-guarantees without requiring manufacturers to reveal too sensitive information up front. Promising building blocks for such a *privacy-preserving catalog* are known from privacy-preserving computing. However, they require extensive research to fit the desired scenario of privacy-preserving bidding platforms for manufacturing.

Such mechanisms must realize *fair data sharing*, i.e., participants must not obtain detailed information about other participants, especially competitors, without providing said information themselves. To this end, mechanisms to *assess the value of data* can provide measures to rate-limit or charge participants with extraordinary usage patterns.

Socio-Economic Challenges (L3)

Beyond the outlined technical measures to realize accountable and dependable manufacturing, we also briefly discuss overarching socio-economic challenges (cf. Figure 2).

Legal Frameworks Legislation currently fails to cover blockchain-based smart contracts and analyses have to show whether general rules suffice to enable the envisioned business relations. Especially when considering global supply chains, also different legal frameworks and multi-national agreements must be taken into account. To realize the desired accountability, legal frameworks must further clarify the responsibility for the accuracy of information in a TrustedStore. An exemplary question is whether manufacturers should be responsible only for the data they provide or whether they should also be responsible for consistency checks on the received data.

In terms of privacy, all systems must comply with local as well as multi-national rules for data privacy, such as the GDPR, including the right to erasure of previously recorded data. Thus, an extensive analysis has to show which data is safe to be stored on-chain, and systems must prevent the inclusion of data that falls under the right to be forgotten or provide mechanisms for data removal without undermining the desired goals.

Furthermore, several third-party services that use the available data are conceivable, e.g., utilizing individual usage data to offer improved maintenance for all customers. To this end, legal frameworks have to clarify who owns the data on the blockchain and who is allowed to process which data in which way. Similar questions also arise for any derived knowledge.

Access and Transparency Before realizing immutable TrustedStores, research must work out the access requirements for different entities and the corresponding trade-off between verifiability and privacy. On the one hand, broad access to information increases transparency such that customers can obtain information more easily. Research must reveal which information is necessary, e.g., to alleviate the required trust from today's slowly forming business relations via technical measures to ease collaboration without pre-established trust. Legal entities may further demand access, e.g., to discover cartels.

On the other hand, information stored on a (semi-)public blockchain must not subvert privacy-legislation. Specifically, granting broad access to information may put business secrets and privacy at risk. Furthermore, reasonable freedom of action for market participants must be maintained. For example, adequate measures must prevent customers from exploiting the knowledge of a participant's low machine utilization to achieve an uneconomic price. In the end, socio-economic research must develop guidelines for blockchain-based platforms that do not only optimize cost but lead to a healthy ecosystem with incentives for high quality, economically healthy companies, and employee well-being.

■ **FUTURE** manufacturing will be driven by exciting advances stemming from the combination of IoT and blockchain technology to implement a dependable and accountable ecosystem. We identified relevant future use cases for both supply chain-related and unrelated aspects that should significantly improve the utilization of manufacturing data (cf. Figure 1). In particular, research must address open challenges on different layers, ranging from system-specific blockchain questions to overarching socio-economic challenges (cf. Figure 2). Regardless, we believe that most effort must be invested in scenario-driven tasks to enable trustworthy information stores, i.e., *TrustedStores*, in competitive, business-driven, and potentially distrustful industry environments. Fortunately, smaller advances are already achievable in increments, and as such first changes should be realizable in the near future.

ACKNOWLEDGMENT

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2023 Internet of Production – 390621612.

REFERENCES

1. A. Kusiak, “Smart manufacturing must embrace big data,” *Nature*, vol. 544, no. 7648, pp. 23–25, 2017.
2. J. Pennekamp, M. Henze, S. Schmidt, P. Niemietz, M. Fey, D. Trauth, T. Bergs, C. Brecher, and K. Wehrle, “Dataflow Challenges in an Internet of Production: A Security & Privacy Perspective,” in *Proceedings of the ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC ’19)*. ACM, 2019, pp. 27–38.
3. S. Malik, S. S. Kanhere, and R. Jurdak, “ProductChain: Scalable Blockchain Framework to Support Provenance in Supply Chains,” in *2018 IEEE 17th International Symposium on Network Computing and Applications (NCA)*. IEEE, 2018, pp. 20:1–20:10.
4. L. Gleim, J. Pennekamp, M. Liebenberg, M. Buchsbaum, P. Niemietz, S. Knape, A. Epple, S. Storms, D. Trauth, T. Bergs, C. Brecher, S. Decker, G. Lakemeyer, and K. Wehrle, “FactDAG: Formalizing Data Interoperability in an Internet of Production,” *IEEE Internet of Things Journal*, vol. 7, no. 4, pp. 3243–3253, 2020.
5. R. Matzutt, B. Kalde, J. Pennekamp, D. Arthur, M. Henze, T. Bergs, and K. Wehrle, “How to Securely Prune Bitcoin’s Blockchain,” in *2020 IFIP Networking Conference (Networking)*. IEEE, 2020, pp. 298–306.
6. S. Wang, D. Li, Y. Zhang, and J. Chen, “Smart Contract-Based Product Traceability System in the Supply Chain Scenario,” *IEEE Access*, vol. 7, pp. 115 122–115 133, 2019.
7. M. Montecchi, K. Plangger, and M. Etter, “It’s real, trust me! Establishing supply chain provenance using blockchain,” *Business Horizons*, vol. 62, no. 3, pp. 283–293, 2019.
8. S. Malik, V. Dedeoglu, S. S. Kanhere, and R. Jurdak, “TrustChain: Trust Management in Blockchain and IoT supported Supply Chains,” in *2019 International Conference on Blockchain (Blockchain)*. IEEE, 2019, pp. 184–193.
9. 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 (HICSS)*. AIS, 2017, pp. 4182–4191.
10. M. Bartoletti and L. Pompianu, “An analysis of Bitcoin OP_RETURN metadata,” in *International Conference on Financial Cryptography and Data Security (FC)*. Springer, 2017, pp. 218–230.
11. E. Mengelkamp, J. Gärttner, K. Rock, S. Kessler, L. Orsini, and C. Weinhardt, “Designing microgrid energy markets: A case study: The Brooklyn Microgrid,” *Applied Energy*, vol. 210, pp. 870–880, 2018.
12. Z. Yang, K. Yang, L. Lei, K. Zheng, and V. C. Leung, “Blockchain-Based Decentralized Trust Management in Vehicular Networks,” *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 1495–1505, 2019.
13. Q. Lin, H. Wang, X. Pei, and J. Wang, “Food Safety Traceability System Based on Blockchain and EPCIS,” *IEEE Access*, vol. 7, pp. 20 698–20 707, 2019.
14. J. Pennekamp, L. Bader, R. Matzutt, P. Niemietz, D. Trauth, M. Henze, T. Bergs, and K. Wehrle, “Private Multi-Hop Accountability for Supply Chains,” in *2020 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2020, pp. 1–7.
15. C. Cachin and M. Vukolić, “Blockchain Consensus Protocols in the Wild,” in *31st International Symposium on Distributed Computing (DISC 2017)*. Schloss Dagstuhl, 2017, pp. 1:1–1:16.
16. R. Matzutt, J. Hiller, M. Henze, J. H. Ziegeldorf, D. Müllmann, O. Hohlfeld, and K. Wehrle, “A Quantitative Analysis of the Impact of Arbitrary Blockchain Content on Bitcoin,” in *International Conference on Financial Cryptography and Data Security (FC)*. Springer, 2018, pp. 420–438.
17. J. Pennekamp, F. Alder, R. Matzutt, J. T. Mühlberg, F. Piessens, and K. Wehrle, “Secure End-to-End Sensing in Supply Chains,” in *2020 IEEE Conference on Communications and Network Security (CNS)*. IEEE, 2020, pp. 1–6.
18. 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)*, 2016, pp. 8:1–8:6.

Jan Pennekamp received his B.Sc. degree and M.Sc. degree in Computer Science from RWTH Aachen University with honors. He is a researcher at the Chair of Communication and Distributed Systems (COMSYS) at RWTH Aachen University, Germany. His research focuses on security & privacy aspects in the Industrial Internet of Things (IIoT). He is *IEEE Student Member*. Contact him at pennekamp@comsys.rwth-aachen.de.

Roman Matzutt received his B.Sc. degree and M.Sc. degree in Computer Science from RWTH Aachen University. He is a researcher at the Chair of Communication and Distributed Systems (COMSYS) at RWTH Aachen University, Germany. His research focuses on blockchain and its privacy impli-

cations. He is *IEEE Student Member*. Contact him at matzutt@comsys.rwth-aachen.de.

Salil S. Kanhere received his M.S. degree and Ph.D. degree from Drexel University in Philadelphia. He is a Professor of Computer Science and Engineering at UNSW Sydney, Australia. His research interests include Internet of Things, blockchain, pervasive computing, cybersecurity and applied machine learning. He is a *Senior Member of the IEEE and ACM* and an Humboldt Research Fellow. He serves as the Editor in Chief of the Ad Hoc Networks journal. Contact him at salil.kanhere@unsw.edu.au.

Jens Hiller received his B.Sc. degree and M.Sc. degree in Computer Science from RWTH Aachen University. He is a researcher at the Chair of Communication and Distributed Systems (COMSYS) at RWTH Aachen University, Germany. His research focuses on efficient secure communication in the Internet of Things. Contact him at hiller@comsys.rwth-aachen.de.

Klaus Wehrle received his Diploma (equiv. M.Sc.) and PhD degree from University of Karlsruhe (now KIT), both with honors. He is full professor at the Chair of Communication and Distributed Systems (COMSYS) at RWTH Aachen University, Germany. His research interests include network protocol engineering, methods for network analysis, and reliable communication. He is a *Member of IEEE and ACM*. Contact him at wehrle@comsys.rwth-aachen.de.

FURTHER READING

We provide references to further reading material related to this article for an overview into related work and today's relevant research challenges. In particular, our selected literature provides additional insights into challenges (1) and application areas (2–4) of blockchain technology as well as supply chain-specific research (5–6). Finally, we include literature on the envisioned *Internet of Production* (7) and associated challenges when processing big data (8).

1) **Blockchain challenges:**
Z. Zheng, S. Xie, H.-N. Dai, H. Wang and X. Chen, "Blockchain Challenges and Opportunities: A Survey," Inderscience, 2018, *International Journal of Web and Grid Services*, vol. 14, no. 4, p. 352–375.
2) **Financial blockchain applications:**
Y. Guo and C. Liang, "Blockchain application and outlook in the banking industry," Springer, 2016, *Financial Innovation*, vol. 2, no. 1, p. 24:1–24:12.
3) **Survey of blockchain-based applications:**
F. Casino, T. K. Dasaklis, and C. Patsakis, "A systematic literature review of blockchain-based applications: Current status, classification and open issues," Elsevier, 2019, *Telematics and Informatics*, vol. 36, pp. 55–81.
4) **Determining the suitability of blockchain:**
K. Wüst and A. Gervais, "Do you need a Blockchain?" in *Crypto Valley Conference on Blockchain Technology (CVCBT)*. IEEE, 2018, pp. 45–54.
5) **Blockchain-supported use cases in the context of supply chains:**
P. Gonczol, P. Katsikouli, L. Herskind, and N. Dragoni, "Blockchain Implementations and Use Cases for Supply Chains-A Survey," IEEE, 2020, *IEEE Access*, vol. 8, pp. 11 856–11 871.
6) **Research propositions for supply chains:**
A. Rejeb, J. G. Keogh, and H. Treiblmaier, "Leveraging the Internet of Things and Blockchain Technology in Supply Chain Management," MDPI, 2019, *Future Internet*, vol. 11, no. 7, pp. 1–22.
7) **Internet of Production:**
J. Pennekamp, R. Glebke, M. Henze, T. Meisen, C. Quix, R. Hai, L. Gleim, P. Niemietz, M. Rudack, S. Knape, A. Epple, D. Trauth, U. Vroomen, T. Bergs, C. Brecher, A. Bührig-Polaczek, M. Jarke, and K. Wehrle, "Towards an Infrastructure Enabling the Internet of Production," in *2019 IEEE International Conference on Industrial Cyber Physical Systems (ICPS)*. IEEE, 2019, pp. 31–37.
8) **Challenges in Big Data:**
A. Oussous, F.-Z. Benjelloun, A. A. Lahcen, and S. Belfkih, "Big Data technologies: A survey," Elsevier, 2018, *Journal of King Saud University-Computer and Information Sciences*, vol. 30, no. 4, pp. 431–448.