Dataflow Challenges in an *Internet* of Production: A Security & Privacy Perspective

Jan Pennekamp*, Martin Henze†, Simo Schmidt‡, Philipp Niemietz‡, Marcel Fey‡, Daniel Trauth‡, Thomas Bergs‡, Christian Brecher‡, Klaus Wehrle*

*Communication and Distributed Systems, RWTH Aachen University, Germany · †Cyber Analysis & Defense, Fraunhofer FKIE, Germany · ‡Machine Tools and Production Engineering, RWTH Aachen University, Germany

{pennekamp, wehrle}@comsys.rwth-aachen.de · martin.henze@fkie.fraunhofer.de

{s.schmidt, p.niemietz, m.fey, d.trauth, t.bergs, c.brecher}@wzl.rwth-aachen.de

ABSTRACT

The Internet of Production (IoP) envisions the interconnection of previously isolated CPS in the area of manufacturing across institutional boundaries to realize benefits such as increased profit margins and product quality as well as reduced product development costs and time to market. This interconnection of CPS will lead to a plethora of new dataflows, especially between (partially) distrusting entities. In this paper, we identify and illustrate these envisioned inter-organizational dataflows and the participating entities alongside two real-world use cases from the production domain: a fine blanking line and a connected job shop.

Our analysis allows us to identify distinct security and privacy demands and challenges for these new dataflows. As a foundation to address the resulting requirements, we provide a survey of promising technical building blocks to secure inter-organizational dataflows in an IoP and propose next steps for future research. Consequently, we move an important step forward to overcome security and privacy concerns as an obstacle for realizing the promised potentials in an Internet of Production.

CCS CONCEPTS

• **Security and privacy** → **Privacy-preserving protocols**; • **Computer systems organization** → *Embedded and cyber-physical systems*; • **Information systems** → Enterprise information systems.

KEYWORDS

Internet of Production; Dataflows; Information Security

ACM Reference Format:

Jan Pennekamp, Martin Henze, Simo Schmidt, Philipp Niemietz, Marcel Fey, Daniel Trauth, Thomas Bergs, Christian Brecher, and Klaus Wehrle. 2019. Dataflow Challenges in an *Internet* of Production: A Security & Privacy Perspective. In *ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC'19), November 11, 2019, London, United Kingdom.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3338499.3357357

CPS-SPC'19, November 11, 2019, London, United Kingdom

ACM ISBN 978-1-4503-6831-5/19/11…$15.00

https://doi.org/10.1145/3338499.3357357

1 INTRODUCTION

Cyber-physical systems (CPS) allow the interconnection of a variety of physical objects in a computer network, e.g., enabling computer systems to (remotely) execute control over entities in the physical world. Through this interconnection of previously isolated systems, CPS provide numerous advantages in the control and operation of complex systems, especially as they are extremely adaptable and versatile, thus contributing to increased efficiency. Recent advances in the field of CPS resulted in a dramatic increase in the deployment of sensors, actuators, control processing units, as well as communication devices communicating amongst each other and with the Internet, fostering a thing-to-thing [26] or device-to-device communication without (necessarily) humans in the loop.

To implement additional improvements in the area of manufacturing, the Internet of Production (IoP) [33, 48] takes the idea of CPS one step further to afford collaboration between different manufacturing processes to establish a "production-to-production" communication. Especially in the context of manufacturing and production, these advantages of CPS go hand in hand with security and privacy concerns [17]. Consequently, manufacturing companies are understandably cautious when sharing, transferring, and storing their valuable production details and process information. To ensure a large-scale deployment of an IoP, this new kind of production-to-production communication thus first requires a profound security and privacy dataflow analysis. In this paper, we refer to privacy and trade secrets interchangeably to improve the readability. This decision is reasonable as trade secrets are essentially confidential process or product information that need to be protected from privacy-invasive applications or entities. Hence, our terminology covers the business privacy of a single stakeholder in a CPS rather than the traditional notion of human privacy.

While research in the area of traditional CPS is mainly concerned with the view of a single stakeholder [37, 54], i.e., how to integrate and enable networking and networked control into a local environment, these advances are insufficiently covered wrt. an IoP because they usually neglect systems that handle inter-organizational information where multiple stakeholders process information. In traditional CPS, nodes mainly communicate with a trusted entity, e.g., a gateway or directly with a web or cloud server. However, in an IoP, communication of valuable data is expected to also take place with external, potentially untrusted entities, e.g., along supply chains, or even with competitors [48]. Hence, the control over data and the influence on the CPS is shifted from a single entity

to a set of entities with different levels of impact. Consequentially, depending on the involved entities, the posed requirements of the CPS and the respective dataflows can change significantly when compared to typical scenarios that CPS cover. Ultimately, this difference mandates a separate investigation of information leakage and sharing for each of these parts beyond the efforts for addressing security and privacy challenges of CPS in general [27, 56].

In this paper, we specifically study the security and privacy challenges resulting from an increasing interconnection of (previously isolated) CPS across different organizations as advocated by the IoP. To this end, we identify the resulting novel dataflows in manufacturing and production processes between (potentially distrusting) entities, which are necessary to achieve extensive advances in today's production landscape as envisioned by the IoP [48], effectively utilizing inter-organizational data even across different domains. Based on these dataflows, we can derive challenges wrt. security and privacy that need to be overcome to turn the vision of an IoP into reality and hence allow every participant to benefit from overall production process improvements, an increased flexibility to quickly react on change requests, and reduced product research cycles. More specifically, our contributions are as follows:

(1) We identify a comprehensive list of inter-organizational dataflows in an Internet of Production based on real-world use cases from the manufacturing and production sector.
(2) Based on each distinctive group of dataflows, we first highlight security aspects before we categorize these security and privacy challenges into three categories: *authenticity of information*, *scope of data access*, and *anonymity*. These aspects are potential obstacles when realizing the full vision of an Internet of Production, especially considering the envisioned inter-organizational cross-domain collaborations.
(3) To pave the way for an effective and successful implementation of an Internet of Production, we conduct a survey of building blocks that potentially allow the industrial stakeholders to tackle the identified challenges to enable security- and privacy-respecting dataflows.

2 INTERNET OF PRODUCTION

As a foundation to analyze dataflow challenges in CPS and specifically an Internet of Production wrt. security and privacy, we first provide an introduction into the concept of the IoP before we present details on general information security considerations.

The Internet of Production (IoP) describes a vision that enables manufacturers of any product to utilize domain knowledge due to advances in the IoT and the groundwork of interconnected CPS [48] which enable the measurement and extraction of massive amounts of data related to a specific manufactured product and its production process [20]. This data can be used to improve the CPS itself as well as the product and the production process, e.g., by using this data to create digital twins (fully representing manufactured products) [69] or the new concept of digital shadows (a simple, real-time capable abstraction of a product) [59]. The ultimate vision of an IoP is to leverage this information to create a platform for global inter-organizational collaborations across companies to combine data gathered in the different stages of any product, i.e., during development, production, and customer usage [65]. With information

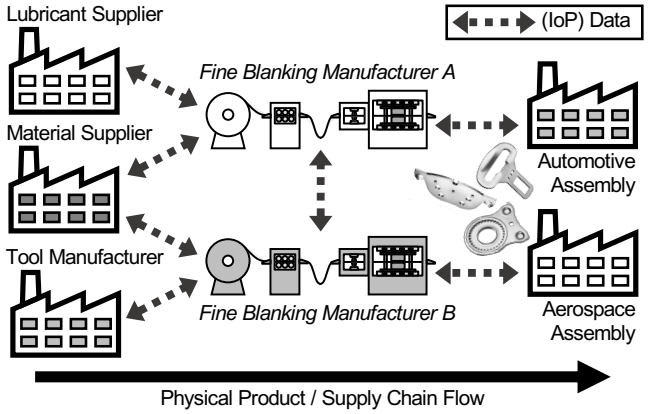

Figure 1: Exemplary IoP dataflows that bring benefits to the local CPS by integrating information from external stakeholders. Data provided by suppliers (here: lubricants, material, and tools) can be utilized by both manufacturers who themselves exchange production data with subsequent assembly stations. Similarly, both manufacturers exchange production process information on their CPS (here: fine blanking lines), the processed material, and their interplay.

automatically shared in real time by all involved entities, companies are able to minimize production interruptions independent of the responsible source, improving their local CPS. Currently, CPS usually rely on local information only, effectively sealing knowledge in stakeholder-specific data silos that are not connected to external parties at all, rendering any automated data exchange infeasible. The establishment of automated inter-organizational dataflows helps to maximize the benefits envisioned by an IoP.

As part of this paper, we take a look at the resulting global dataflows, in particular, communication involving different organizations. Intentionally, we neither consider how to (company-) internally manage CPS and dataflows, i.e., when only a single stakeholder is involved, nor consider monetary flows which also take place when valuable information is globally exchanged.

2.1 Use Cases for an Internet of Production

To further illustrate inter-organizational dataflows and motivate the huge potential of large-scale collaboration of CPS in an Internet of Production, we now present two exemplary use cases based on real-world applications: fine blanking and a connected job shop.

2.1.1 Demanding Quality Requirements in Fine Blanking. Fine blanking is a precision forming process for manufacturing huge quantities of (ideally) identical work pieces, e.g., for the aerospace and automotive services [20, 35], thus reducing costs during production. To this end, fine blanking utilizes characteristic operating parameters that result in the production of parts with excellent quality of the sheared surface and geometric accuracy [78]. Despite its quality superior to traditional blanking or stamping processes [36], each produced work piece is not identical, while the process setup remains unchanged [72]. Due to the complex interplay of tool components, fluctuating material properties, changing environmental conditions, wear of active components in the fine blanking tool, and varying behavior of the machine tool itself, the understanding

of the behavior of the process has come to its limits [72]. Hence, no generic process setup exists requiring a manual setup for every upcoming product by a trial-and-error approach.

In an IoP-enabled fine blanking line [8], as illustrated in Figure 1, apart from delivering physical goods, the suppliers would also deliver digital product information which enables the manufacturers to adjust their CPSs according to the received material specification. Hence, adjustments before production can reduce the need to retrospectively polish the produced work piece to reach the desired (identical) quality and enlarge the tolerance of incoming material. Besides, the two manufacturers (here: automotive and aerospace fine blanking manufacturers) each operating their own fine blanking lines could benefit from directly exchanging process know-how to reduce scrap. Finally, both companies could share the properties of each fine-blanked component with their respective customer (assembly) to create awareness of minor quality deviations. This dataflow would allow the assembling companies to adjust their CPS accordingly, ultimately increasing the final product's quality without producing significant amounts of waste. At the same time, with data available along the supply chain, fault detection in case of a failure during the assembly could become more holistic, increasing the chance to detect and remedy its root cause.

To realize such an IoP-enabled fine blanking line, data of the different mentioned sources of influence as well as the resulting quality features of, ideally, each produced work piece need to be acquired, to effectively analyze the impact of outer influences onto the process. To assure quality guarantees as well as product estimates in an accountable and verifiable way, all fine-blanked components need to be trackable, which is yet not established in productive environments. Identifying each component requires markers applied onto the work pieces in an additional manufacturing step, or by laser printing that may also affect the physical properties of the work piece surface, both usually come with high economical costs [71]. Approaches to identify work pieces by their near to unique surface structure have been pursued, but yet not proven mature enough to be integrated in productive environments [50]. Establishing this trackability would enable customers of fine-blanked components to give componentwise feedback of the quality of each work piece during the assembly or usage in the end-product.

Given a fully implemented IoP, customer feedback for every component as well as information about the used material, respectively output and input of the process, would be available. In combination with the idea of globally connected CPS, this information detailing various production setups of fine blanking lines truly enables the development of data driven models to increase the understanding of the process leading to an autonomous process setup.

2.1.2 Connected Job Shops in Discrete Manufacturing. Discrete manufacturing is characterized by the production of individual products (like consumer goods, automotive parts and such) in units, where each product is produced separately [24]. Each producible part (i.e., the work piece) comes with its own set of quality requirements. Those have to be assured, commonly by intermittent quality measurements. In case of defective parts which are out-of-tolerance, batches of finished parts after the previous quality measurement have to be scrapped, adversely affecting production planning and potentially associated assemblies or other products.

Apart from those cascading effects, tracking down root causes for quality issues can be tedious, time-consuming and oftentimes difficult, as sufficient information about the complex machining process and involved components is lacking. In conjunction with a connected job shop, valuable data across the entire production cycle of each part could become available utilizing knowledge about all quality-influencing factors. Those factors include the machining tool data, history and condition, the work piece material, the raw geometry and required tolerances, a history of loads and positional deviations of all involved machine tools, or performance data of the used numerical control unit. Apart from efficiently troubleshooting quality issues, an IoP-enabled connected job shop would allow companies to optimize the availability and productivity of production lines across different production sites and even organizations.

The connected job shops in such a discrete manufacturing use case must rely on a supply chain for production and following distribution. For example, a company manufactures a CNC milling machine using received parts (drive and guide components, bearings, and milling spindle) of its supplier. Afterward, a product producer purchases the assembled milling machine and uses it to produce consumer products. While an external maintenance contractor is responsible for maintaining the availability of the manufacturing machine, another supplier provides and replenishes required tools, such as milling cutters, to both machine and product producer.

In a scenario without an IoP, both producers have to manually adjust their production according to the delivered parts and components as well as other influencing factors. For example, the real loads on machine components during machining are unknown due to a lack of detailed models and availability of comparison data. Similarly, accurate predictive maintenance is still in its infancy due to the complexity of failure mechanisms. Again, a large-scale data comparison would help to identify root causes more easily.

An IoP and its interconnectivity would allow the producers to automatically receive product and process parameters, allowing efficient machine configuration and better process control. Consequentially, the manufacturing processes could be automatically adjusted to improve the quality, to speed-up the process, or to reduce scrap production or machine downtime. Due to the reuse of data, the accuracy of such adaptations could be increased as well. Furthermore, based on the collected and processed information of the product producer, the machine producer would also be able to provide more sophisticated usage estimates (and guarantees) that adopt real-time or previous usage data for this manufactured machine to its customer. Finally, the production can be automatically adjusted for customer change requests based on historical knowledge derived from previous production processes and results extracted from digital models and simulations, i.e., the producers can offer Manufacturing-as-a-Service with limited interaction.

Overall, an Internet of Production would allow companies to create a holistic view of the connected job shop and its connections and dependencies along the supply chain, while reducing the need for human expert interaction. So far, issues regarding the ownership attribution and privacy challenges of these big amounts of data gathered by different stakeholders, i.e., the machine tool and machine manufacturers, machine component and tool suppliers or the end users, are the main issues which prevent these global optimization measures already today.

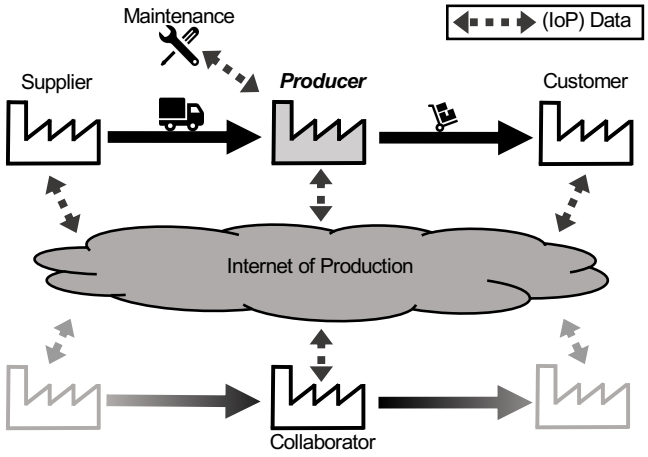

Figure 2: The vision of an Internet of Production (IoP) is that participating entities (along the supply chain and across the production landscape) collaborate across organizations to increase the productivity. Depending on the relationship of entities, different security and privacy challenges arise for each communication link and the associated dataflows.

3 A PRODUCTION LANDSCAPE WITH INTER-ORGANIZATIONAL DATAFLOWS

Based on our exemplary use cases, we can formally define a set of entities that potentially can collaborate within an Internet of Production (IoP) and derive the relationships between them. Our formalization will serve as a foundation for identifying (potentially security and privacy relevant) dataflows later on.

3.1 Participating Entities

Overall, we define five groups of entities which can be involved within a typical manufacturing or production process: *supplier*, *producer*, *collaborator*, *customer*, and *maintenance provider*. In Figure 2, we illustrate their embedding within the global production landscape from the point of view of a single producer. The upper row refers to a regular supply chain, where the producer receives goods from a supplier and manufactures a new product, component, or part. This company's output is then delivered to a customer, which can be either a merchant, an end customer, or another producer, i.e., the current company is a supplier wrt. the following entity in the supply chain. Furthermore, we define a maintenance entity, which directly interacts with other entities (e.g., suppliers, producer, or customers). In our illustration, we only include a single maintenance entity for simplicity. We mark the concept of an IoP, i.e., interconnecting different entities to establish previously non-existing inter-organizational dataflows, with a gray background and visualize communication links to the other entities (the collaborators) accordingly. Three of the introduced entities are also producers, i.e., the supplier, the collaborator, and depending on the type, even the customer. Overall, we end up with the following definitions of entities in our newly defined production landscape:

- **Supplier:** delivers materials or intermediate products to the producer (its collaborator along the supply chain).
- **Producer:** is our point of view of the production landscape (either a supplier, collaborator, or customer).
- **Collaborator:** acts as an end point of supply chain unrelated inter-organizational data exchanges.
- **Customer:** receives (intermediate) products from the producer (its collaborator backwards in the supply chain), can also receive a final product as an end customer.
- **Maintenance Provider:** directly interacts with other entities, i.e., its clients, to perform maintenance-related tasks.

3.2 Benefits of Industrial Collaboration

Following the introduction of all entities, we now take a look at the benefits of two entities collaborating. The respective dataflows are means to achieve improvements in the production domain that are otherwise either not possible or not as easily accomplishable. Given that the capabilities of entities vary significantly, we have to look into the different combination of entities individually. Hence, next, we first consider the traditional relationships of the producer along the supply chain before also referring to the benefits of collaboration between previously unaffiliated companies.

3.2.1 Supplier & Producer. Traditionally, the supplier delivers raw materials or intermediate products to the producer in a mainly unidirectional flow of data and information. With increased collaboration between entities in the production landscape, the supplier is expected to provide accountability for each delivered piece or batch of raw material that the company further processes, allowing the company to efficiently adapt its production. The supplier and the company are expected to cooperate more closely in the development of new products to utilize synergy effects based on insight of previous usage data. Overall, a closer collaboration of supplier and producer helps to reduce inaccuracies in the products (i.e., improve the quality), optimize the product development process, and cut down costs due to unexpected manufacturing process adjustments.

3.2.2 Maintenance Provider & Producer. Similarly, a collaboration between the maintenance provider and its clients (producers) might help to improve the response time by the maintenance provider. In a collaborative scenario, the maintenance provider should be able to minimize the client's downtimes by conducting predictive maintenance and shipping replacement parts on time solely based on insight into the running processes at the client. Besides, based on the usage information, the maintenance provider can predict usage estimates or offer remote repairs that further help both entities to schedule the production accordingly. Overall, both entities improve their product knowledge, which improves the efficiency (i.e., less unplanned situations and outages).

3.2.3 Producer & (End) Customer. Collaborations between the end customer and the producers are already well-established in the digital world. For example, tracking and usage information allows the company offering (digital) services to draw meaningful conclusions. With an interconnected production landscape, these collaboration-based advances will likely follow in the manufacturing sector. On the one hand, the customer is interested in receiving the best-suitable user-tailored product. On the other hand, the producer wants to minimize unnecessary expenses by only offering and supporting needed features based on the customer usage. These requirements call for a collaborative production environment with agile processes as well as a flexible product development.

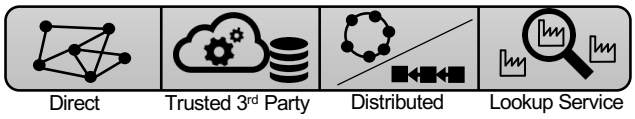

Figure 3: Different deployment models of a proposed IoP introduce varying challenges to start (new) collaborations.

3.2.4 Producer & Collaborator. As today's production landscape lacks collaborations between unaffiliated collaborators, the benefits of such concepts are particularly interesting. In general, any producer in an Internet of Production is a collaborator. The collaborators who can be part of any step of a supply chain are mainly driven by advantages through these computational advances and the occurring knowledge transfer. However, those collaborator relationships are not necessarily linked to a particular product or the supply chain. For example, both collaborators can simply process the same raw material or operate a machine by the same manufacturer (cf. Figure 1). So far, none of the potential is being exploited simply because these collaborations do not exist. As of today, the creation of these relationships between different collaborators remains unclear. Hence, it is a necessity to bootstrap an IoP to easily interconnect (previously unaffiliated) companies.

3.3 Realizing Inter-Organizational Collaborations in an Internet of Production

To realize these benefits, enabling collaborations between the different involved entities is essential. From a technical perspective, different architectures or design decisions can be taken for the different forms of collaboration [48], e.g., to account for different technical or security requirements, resulting in varying deployment and usage models. To capture all (potentially security critical) resulting dataflows, evaluating and considering all these different options for realizing an IoP is an important task. Overall, we identify four different approaches for realizing collaboration across organizations in an IoP (cf. Figure 3): *Direct communication*, through a *(trusted) third party*, in a *distributed network*, or a *lookup service*. While (more advanced) collaborations along the supply chain can easily be established due to the existing business relationship, additional mechanisms for a bootstrapping of new collaborations with unaffiliated entities are required to realize the full potential of an IoP. Otherwise, potentially valuable information remains accessible to a local entity only and is not shared and utilized globally.

Direct Communication. In a scenario with a direct communication between the different entities, the participating entities must only trust their current communication partner (who might already be a business partner) from a security and privacy perspective. Hence, the risks of data leakage itself are not as challenging as no information is centrally stored or processed. Therefore, we identify the bootstrapping of new collaborations as the main challenge because the different entities are initially unaware of additional entities besides their current partners. Given that no central registry listing all manufacturing companies, their deployed machines and processes, and their processed raw materials or intermediate products exists, a trivial solution for this challenge is missing.

(Trusted) Third Party. When looking for a simple solution to tackle the bootstrapping challenge of the direct communication scenario, relying on a (trusted) third party is a straightforward solution. In such a scenario, every entity connects to this central entity to share its data and process information. Consequentially, the matching of collaborators with each other can be easily facilitated by the central third party. The motivation for the third party is simple as more collaboration leads to more usage of the platform and therefore to an increase in revenue. Unfortunately, apart from introducing a single point of failure, this design is also more risky from a security perspective as all information is centrally stored (and potentially even processed), effectively creating a valuable target for data theft. Again, more collaborations also increase the amount of data available at the platform. Existing research in the domain of cloud computing shows the risks of such an approach for sensitive information [27, 67]. If the third party cannot be trusted, access to the information should be prevented, essentially creating an encrypted data storage platform that is unable to fulfill the bootstrapping needs as the required process information is inaccessible.

Distributed Network. To prevent a central platform from having access to all available information, a distributed platform could be established. Existing business concepts [6] could be transferred to the industrial domain. In contrast to a single third party, each entity must trust multiple nodes as communication and dataflows are split across various servers. Besides, bootstrapping and connecting to such a network introduces challenges when compared to the solution with a single server because no (dedicated) single entry point exists. Here, mechanisms are required to handle network joins and node failures. These aspects are well-known problems from research in the area of peer-to-peer networks [64].

Lookup Service. Companies can operate their own servers as an alternative design to dedicated distributed trusted nodes. In such a scenario, the central IoP would only serve as a lookup service that can be either operated with a trusted third party or in a distributed way. This decision is a trade-off between the mentioned bootstrapping challenges and the danger of introducing a single point of failure and high value target for data theft. In this scenario, the security requirements of the central IoP nodes are not as critical as in the previous designs because they only delegate data transmissions, i.e., security and privacy are enforced elsewhere. Instead, the entities must mainly trust the entities that the lookup service points to, which in turn is similar to the direct communication setting.

So far, we identified the actors in the production landscape, called entities, as well as benefits of collaboration in an industrial setting. Furthermore, we highlighted potential deployment models to create a large-scale IoP and to overcome bootstrapping issues. However, in the remainder of this paper, we focus on occurring dataflows in an IoP and consider IoP deployment challenges out of scope.

4 CHARACTERISTIC DATAFLOWS

As a foundation to identify security and privacy challenges in IoP deployments, we first need to identify the crucial potential dataflows that occur between the different entities within an Internet of Production. To this end, we now dissect the dataflows we identified in Section 3.2 and discuss resulting security considerations.

4.1 Supplier & Producer

Within an IoP, we identify a number of dataflows between suppliers and producers. The vision is that any supplier shares properties of

the supplied item along with expected usage properties to enable reliable adjustments of the running process. This supplied item can be a part or component, a tool, or even a production machine. Hence, the extent of digital information attached to this item might vary significantly. For example, a supplier of fine-blanked parts can store production properties of its production in a distributed ledger. In a proof-of-concept, data collected from a fine blanking machine is already published to the IOTA tangle [68]. These part details are accessible to the supplied company to provide accountability through a tamper-proof distributed storage. In the opposite direction, the company can transfer details about the expected usage requirements for any ordered item. This information helps the supplier to only deliver items that are actually usable by the company. For example, a lower hardness of a non-structurally used metal piece might help to reduce the machine's wear without having an influence on the final product. These details, however, are sensitive as they can reveal process secrets of the producer to the supplier.

Machine Supplier. When taking a look at more specific suppliers, we identify that a machine supplier might receive even more data. For example, the producer could share details about the machine downtimes as well as about the used components of the machine to the machine supplier with the goal of improving reliability. A use case example for this scenario is Feintool, a company manufacturing fine blanking machines (cf. Section 2.1.1), which offers a service to monitor the machine's functionality [23]. In this scenario, the machine submits anonymized, aggregated usage information to the machine supplier every ten minutes. In return, Feintool is able to share estimates on machine downtimes and various condition changes to the producer. This approach can be extended into a Manufacturing-as-a-Service business model where the company only rents capacity on a machine from the supplier. Then, the machine supplier is responsible for all maintenance-related tasks and for keeping the machine ready for operation. In this scenario, the machine supplier has direct access to the production process, a potentially very sensitive aspect of production. Apart from theoretically being able to reverse-engineer aspects of classified manufacturing processes, traceability of the productive and non-productive times of individual production machines can become an issue.

Tool Supplier. Similar observations regarding the amount of received data hold for tool suppliers providing, for example, cutters, cutting inserts, or grinding discs. The goal of sharing data between a producer and a tool supplier is to reduce downtimes by accessing production-specific data to provide tools in time. However, in contrast to the machine supplier, the tool supplier might not have direct access to the production process. Hence, his access to production parameters is limited. Nonetheless, specific process parameters or wear characteristics can also reveal information on the running process and the handled material (e.g., aerospace-grade aluminum for military vehicles), allowing the supplier to obtain knowledge about manufactured products and the utilization of the site.

Security Perspective. We identify two main challenges when dissecting the dataflows between supplier and producer. First, transferred information, i.e., production values, might allow the reverse-engineering of particular aspects of the production process and can result in a loss of intellectual property or business secrets. This challenge holds for both entities: Supplier knowledge can also inadvertently flow from the producer to the supplier's competitors. This situation is typical, if (unmetered) readable data access exists and is usually dealt with through contracts as both entities have a business relationship. Concerning the second challenge, the involved entities could deliberately deliver incorrect values to achieve a (monetary) advantage in the relationship. The situation intensifies when data is further propagated along the supply chain. Currently, companies are forced to at least occasionally verify the provided information for delivered goods. However, not all process parameters and product properties can be checked without destroying the work piece, hindering the verification of promised properties.

While we focus on information flows in this paper, an extensive survey on supply chain security [73] provides interesting insights wrt. physical aspects. Advances in the IoT [1] can complement their findings to improve the physical security along the supply chain.

4.2 Maintenance Provider & Producer

Wrt. dataflows, the maintenance provider is similar to a supplier, as the maintenance provider directly interacts with the producer (cf. Section 3.2.2). Hence, the dataflows, such as usage values in one direction and usage estimates in the other direction, are very similar and, consequentially, the security perspective is comparable as well. A common approach is that maintenance providers directly establish collaborations with producers to offer their services, i.e., without the involvement of the original manufacturer of the machine. Consequently, the maintenance provider receives valuable information not only about the producer but potentially also about the manufacturer of the maintained machine. In the other direction, firmware updates or configuration recommendations might be passed along from the machine manufacturer or the machine supplier via the maintenance provider to the producer to limit the number of entities that are involved with the machine.

Security Perspective. The maintenance provider can gain valuable insights into the processes of the producer due to its direct access. However, business relationships of the maintenance provider with direct competitors of the producer are an even more severe risk. Hence, contracting an external maintenance might result in (unintentionally) transferring knowledge to direct competitors who are clients of the same maintenance provider. Furthermore, on the one hand, the producer is interested to check the authenticity of manufacturer updates or configuration settings that are passed along from the maintenance provider to minimize the risk of deploying malicious code and parameters. On the other hand, the maintenance provider does not want to be liable for any damages.

4.3 Producer & (End) Customer

For the relationship between the customer and the producer, we identify the following dataflows within an IoP. First, the customer can receive maintenance recommendations and firmware updates from the producer (if provided) to improve the product's availability and productivity. Second, the customer shares her usage requirements along with usage values to send feedback to the manufacturer. This kind of dataflow is identical to the relationship of a supplier and a producer (cf. Section 3.2.1) as the point of view defines the role within the supply chain (cf. Section 3.1), i.e., the customer is basically another company that purchases manufactured products from another (previous) supplier. Hence, in the following,

we mainly investigate a customer who is either a merchant or an end customer as this situation varies from a security perspective.

From a business perspective, the producer can offer its end customers discounts or other benefits for sharing privacy-sensitive information as the customer can only gain less sensitive (and interesting) information through dataflows from the producer. Here, the incentive for establishing a collaboration resides on the site of the producer. This information can, for example, support the producer to link customer satisfaction as well as wear with particular (integrated) subcomponents of a product, i.e., identifying a single supplier that is responsible for abnormal, well or disappointing behavior. Furthermore, detailed usage data can help the producer to provide improved support to its customers increasing both its knowledge about the product and the customer's satisfaction. For example, within discrete manufacturing (cf. Section 2.1.2), some machine manufacturers provide a process ramp-up service, where they support the customer in finding stable process parameters for new machining processes in exchange for knowledge about the products being machined on their machines. Their motivation is to further increase their process parametrization expertise and secure their business interest and customer loyalty.

Security Perspective. A significant threat for end customers for such dataflows is the risk of tracking and surveillance based on usage values [76]. Without proper anonymization or aggregation of data, entities of the supply chain might be able to identify customers based on the data being passed backwards. This situation especially holds for products with only a few buyers. The privacy of these customers is particularly at risk. However, the producer is also put at risk by such dataflows as the information they are receiving from their end customers can also be manipulated. Given the low number of customers, this usage data can have a significant impact on the decision-making of product development or manufacturing. For example, they might falsely adjust their product based on the incorrect usage data which can ultimately have a negative performance on future batches of the product.

4.4 Producer & Collaborator

An IoP deployment enables the collaboration of a producer with other companies outside the existing supply chain, e.g., competitors or companies utilizing the same machines, tools, or components. This new unique type of relationship introduces particularly interesting dataflows. Data that is being exchanged between these entities can vary, e.g., both entities receive parts, components, or material from the same supplier and therefore, they have an incentive to exchange knowledge about how to process the received items in the most efficient way (cf. Figure 1). In a different scenario, both entities could operate machines by the same manufacturer. To utilize these machines efficiently, they have an incentive to rely on the experiences by the respective other entity, e.g., which workload is beneficial for the machine wear or sharing the key performance indicators proving what output rate can be achieved. In an anonymized form, this "collaboration" is even imaginable between competitors if both companies expect improvements following their participation. Then again, dataflows in this context could also occur before a real business relationship is established. On the one hand, one company could (anonymously) check whether a particular component can be manufactured by the collaborator and for what costs without revealing the exact specification. On the other hand, multiple collaborators could unite to establish a syndicated procurement of components or material. Overall, these dataflows allow significant advances in the production domain simply by making information available across stakeholder boundaries.

Security Perspective. Dataflows which are part of this category of relationship are very important for two reasons. First, they are crucial for the success of an IoP as they promise the largest advances in production technology: Without inter-organizational (cross-domain) collaborations most potential of an IoP remains unutilized. Second, due to the flexibility of collaborating entities, questions of trust and accountability are especially challenging. Currently, long-lasting relationships of (existing) business partners do not call for sophisticated security and privacy evaluations as most rights and duties are covered by long established business contracts anyway. However, in an IoP, these relationships are more dynamic and possibly only short-lived, resulting in a frequent linking of previously unaffiliated entities. Hence, the aspects of authenticity, confidentiality, and correctness must be separately analyzed with caution to only guarantee desirable and expected behavior for all involved entities. Finally, companies must act with caution especially when they collaborate with other anonymous entities. Here, security measures must be in place to prevent misuse of an IoP and to protect honest companies that want to collaborate.

5 SECURITY & PRIVACY CHALLENGES

Based on the IoP-specific dataflows that we derived in the previous section, we can now identify and group together the main security- and privacy-related challenges. To this end, we rely on the well-established general information security concepts of confidentiality, integrity, and availability (CIA) [74] as well as authentication, authorization, and accountability (AAA) [74]. Importantly, we deliberately focus on security and privacy challenges of dataflows in the following and do not specifically consider the already addressed orthogonal problem of network security in CPS [27].

The relevance for each dataflow differs depending on the involved entities, the transmitted data, the deployment model of an IoP, and the company's preference wrt. malicious entities and attackers. Besides, some parties might have an aversion against inter-organizational collaborations, hence, their need for perceived security and privacy is higher than the need of parties with an open-minded view. To address these concerns, we identify three higher-order categories of security and privacy challenges in an IoP: *authenticity of information*, *scope of data access*, and *anonymity*.

Authenticity of Information. The first category covers all aspects related to the correctness and origin of the data. In inter-organizational relationships, the *authenticity* of information is a vital aspect because an entity utilizing this information must be sure that the data is reliable. Otherwise, adjusting the machine parameters accordingly can incur significant damages or jobs might be scheduled that have no actual buyer. Similarly, *integrity* protection of dataflows should be in place to prevent any tampering of the data. Once the authenticity of data is ensured, establishing *accountability* is the next challenge. This aspect can introduce liabilities for entities in an Internet of Production. Depending on the exchanged information, accountability can only affect two directly

involved entities or even a longer chain. For example, a machine manufacturer guarantees the validity of its usage estimates to the machine operator, hence, he should be accountable. Along the supply chain, end-to-end guarantees are imaginable as well. Here, a producer of brake pad carriers attests to the end customer of a car that this fine-blanked component endures the car's lifetime. A significant challenge with accountability is to realize a linking between the physical object and its digital information because attaching a physical identifier, such as a barcode, RFID tag, or black light marker, might not always be an option. Besides, to prevent abuse and forgery, such identifiers should be tamperproof. Otherwise, the task of detecting knock-offs or re-labeled components is a virtually impossible endeavor.

Auditing capabilities are required for all dataflows to allow for a verification of processes and data exchanges. Basically, this goal overlaps with the accountability aspect because companies should be able to prove which interactions they initiated anyway. *Immutability* and *referenceability* enable stakeholders to make sure that the transmitted information can be located at a later point again (e.g., for verification). Hence, all data processing should be designed with these requirements in mind. These aspects integrate nicely with the vision of an IoP because companies are expected to improve their processes by applying collected (past) knowledge. Hence, information must be retrievable and available anyway.

Scope of Data Access. While the first category mainly dealt with the trustworthiness and reliability of information, the second category comprises different challenges related to the access of data. In the context of production, most information is valuable because it contains details about the production process, registered patents, or created intellectual properties. Consequentially, all entities have a large incentive to retain their knowledge locally. *Confidentiality* mechanisms and reducing the number of dataflows as well as the extent of dataflows to a minimum help to realize the vision of implementing inter-organizational dataflows even in conservative industrial environments. A reduced *granularity* of information (e.g., through aggregation or anonymization) can help to rule out the danger of process reverse-engineering based on shared data. Monetary compensation helps if no production data is available for an exchange to still implement benefits in mostly one-sided relationships (e.g., between a company and its end customers).

The second large part of this category deals with challenges related to data access. Proper *authentication* should ensure that no information is leaked to unintended parties, i.e., requiring each entity to authenticate themselves. Furthermore, *authorization* must be granted as well to obtain access to information. While these aspects are insignificant for a dataflow only concerning two entities, the challenges are more demanding when taking into account that data is expected to be forwarded along the supply chain. Therefore, options regarding *data control* should be evaluated carefully. In particular, the list of authorized entities must be expandable. Questions wrt. data access can also affect intermediate entities in an IoP. For example, depending on the deployment model, (external) cloud services also have access to the data, and they might even perform intermediate processing, effectively having access to confidential information. Consequentially, their capabilities must be properly defined. As preventing unauthorized data forwarding (from any entity) is a technically nearly impossible task, most regulation is likely based on contracts. To counter misuse, the previously mentioned auditability helps in identifying data leaks once information was transferred in an undesired way.

Anonymity. The third category that we identified deals with the anonymity of participants in an IoP. While direct partners along the supply chain know each other, multi-hop knowledge might not be required or desired. Companies might not have an incentive to reveal their network of suppliers or maintenance providers. Regardless, their actions should still be covered by *identifiability*, providing a unique reference for each action and entity to achieve accountability. In addition, *untrackability* must be taken into account to prevent that side-channel information or communication patterns can de-anonymize participating companies. Overall, an IoP requires anonymity mechanisms in place to also support new areas of dataflows that complement the existing traditional flow of information in production. Especially parties with an aversion of sharing information might not collaborate otherwise.

Apart from these three major categories of dataflow challenges, we also identified three minor aspects that we consider out of scope for this paper as they are more business related: (i) The mentioned bootstrapping issues (cf. Section 3.3) that also need established trust between participating entities, (ii) as well as everything that evolves around the need of achieving availability, reliability, and resilience of an IoP architecture and information, and (iii) open questions wrt. determining the value of shared information between collaborators.

6 SECURITY & PRIVACY BUILDING BLOCKS

To address the pressing security and privacy challenges in an Internet of Production, we now present a comprehensive set of (technical) building blocks that cover distinct subsets. Given that we can identify clusters within these building blocks, we group them into five larger groups that loosely target similar challenges. In particular, we categorize building blocks for security and privacy in an IoP as follows: (i) *data security* covers building blocks that mainly deal with the access to data, (ii) *data processing* concerns technologies which aim to conceal information during computation, (iii) *proving support* deals with mechanisms to establish authenticity of information, (iv) *platform capabilities* incorporate building blocks that realize strict rules for all participants, and (v) *external measures* contain supporting concepts that facilitate the creation of an IoP while not primarily focusing on security aspects.

In the following, we introduce these categories and the individual building blocks they encompass, before deriving takeaways based on the current research state of our identified dataflow challenges. We present a high-level overview of the different building blocks and the security and privacy challenges they address in Table 1.

Data Security. We grouped building blocks with a strong focus on the access to data, i.e., providing confidentiality, into this category. Here, the most basic form to achieve confidentiality is to rely on *encryption* [7]. While regular encryption has no drawbacks wrt. the other challenges we defined, it also lacks a feature to dynamically update the number of entities that are allowed to access the information without leaking the used key or still sharing the content with removed entities (even when data is updated at a later point). In an IoP, relationships are more short-lived and thus, access must be granted in a flexible manner to changing entities. To

Table 1: A mapping between our surveyed building blocks (y-axis) and our categorization of security and privacy challenges (x-axis) shows that no single one fits all solution exists. Depending on the security goal, the applicability of the different building blocks also varies significantly (from ++, over + and +/- to - and --). No entry denotes that no direct impact is notable.

	authenticity	integrity	accountability	auditing	immutability & referenceability	confidentiality	granularity	authentication	authorization	data control	identifiability	untrackability
	Authenticity of Information					**Scope of Data Access**					**Anonymity**	
Data Security												
Encryption [7]						+ +						
Data Usage Control [52]			+	+		+ +	+	+	+	+		
Secret Sharing [60]	+		+			+				+ +	+	+
Data Processing												
Secure Offloading [10]	+					+ +			+	+ +	+	+
Secure Computation [43]			- -	-	-	+ +			+ +			
Anonymization [62]	-	-	-		- -	+	+ +				+	+
Proving Support												
Digital Fingerprints [71]	+ +		+ +		+							- -
Digital Signatures [55]	+ +	+ +	+ +	+								- -
Distributed Ledgers [44]	+ +	+ +	+ +	+ +	+				+			+/-
Version Control [40]		+	+	+ +	+ +				+		+	-
Platform Capabilities												
Access Control [57]				+		+	+	+ +	+ +			- -
Policies [32]				+		+	+ +	+	+ +	+/-		
Smart Contracts [75]	+	+	+ +	+	+			+	+			+/-
Trusted Computing [58]	+ +	+ +	+	+		+ +	+	+	+	+ +	+	+
External Measures												
Data Markets [5]		+ +				+/-	+	+ +	+ +	+	+	-
Legal Contracts [4]	+ +		+ +	+	+	-	+ +	+ +	+	+ +	+	- -
Smart Payments [34]				+ +	+					+	+/-	+/-

improve usability, Ma et al. [41] proposed an enhanced encryption scheme especially targeted for the industrial context that is able to make encrypted information searchable. Traditionally, systems processing solely encrypted data must rely on an additional indexing schemes to support search queries based on this extra information.

Another building block to achieve data security is *data usage control* [52] which allows distributing decisions regarding data access to multiple parties. Hence, this approach fulfills all aspects of the challenge wrt. the scope of data access, i.e., limiting the access to information for external stakeholders in an IoP. With the correct set of policies, logging functionality to achieve accountability can be integrated as well. However, so far, this technique is more a (theoretical) concept than an established functional system.

Finally, *secret sharing* [47, 60, 63] allows data sharing with multiple entities in a confidential way. To reveal the information, a subset of the entities must collaborate to reconstruct the original information allowing a certain degree of data control. Hence, apart from computational overhead, its applicability might be limited in a dynamic environment such as an IoP where entities often change. Regardless, Zhou and Chao [79] show an application in the Internet of Things to establish a security architecture. In a more static context, Cyran [13] uses secret sharing in another domain (healthcare) with strict confidentiality requirements. Overall, such an approach could help to overcome today's trust issues of companies.

Data Processing. The category of data processing covers approaches that try to hide information during computations from unintended recipients, i.e., they extend the concept of simply limiting access to data to approaches that can also operate on or with data in a secure manner. In particular, we identify three larger building blocks in this category: secure offloading [10] (operating directly on ciphertext), secure computation [43] (jointly computing a function without revealing individual inputs), and anonymization [62] (a collection of one way functions to anonymize data).

The specific implementations of *secure offloading* support different complexity of computations (e.g., homomorphic [19, 70] and order-preserving encryption [2]). They have in common that encrypted data is sent to another party who performs calculations on the ciphertexts without inferring the content. Afterward, the entities with the correct key are able to decrypt the resulting ciphertext to obtain the result. Such an approach enables stakeholders to rely on (untrusted) cloud services for computation without the fear of leaking information [80], i.e., confidentiality and data control are preserved. Furthermore, it allows stakeholders to offload their computation anonymously because no conclusions about the data owner can be drawn. The production domain is a logical applicant as companies operate with large amounts of process data.

Approaches in the area of *secure computation*, such as secure multi-party computation [38], oblivious transfer [53], and zero-knowledge proofs [22], provide protocols between multiple (distrusting) stakeholders to either jointly compute a result or to exchange information or secrets obliviously. Hence, they are particularly suitable for dataflows between previously unaffiliated collaborators. Recent work even shows the possibility of privacy-preserving database lookups without a trusted third party [14] to reduce any leakage. Unfortunately, being oblivious reduces the accountability and referenceability of this approach significantly because the individually provided inputs are only locally available and hence, no (external) verification is possible without cooperation.

Third, *anonymization* approaches, such as k-anonymity [66], differential privacy [16], data aggregation [25, 61], and noise [15, 21],

allow entities to protect sensitive information by aggregating it with other data points or by altering their precision. Then, they can collaborate with other entities in an IoP without leaking valuable process information. While these techniques also limit the accountability and authenticity of information, they also allow stakeholders to participate anonymously as no single data point can be traced back to a single entity. For example, recent work has shown that consumer usage data can be properly anonymized [29, 30], a direction that is likewise promising in an industrial context, e.g., to enable anonymous comparisons of the efficiency of production processes across manufacturers [48]. However, companies should take into account that dataflows can already reveal relationships between different entities based on communication patterns only [31].

Proving Support. The approaches presented so far particularly deal with the challenge of controlling access to information. Now, we look into a group of building blocks that specifically enable stakeholders to verify the authenticity of information. The respective approaches range from proving physical aspects of a work piece, i.e., digital fingerprints [50, 71], to providing evidence for the origin and correctness of digital information (e.g., digital signatures [55], distributed ledgers [44], and version control [40]). While different in scope, these approaches have in common that their ability to attest the authenticity and integrity of information contradicts the desire of stakeholders to remain untrackable. *Digital fingerprints* of physical products, i.e., having a unique digital identifier of a work piece or product available, are difficult to realize in an industrial context because attaching a barcode or a unique identifier to a manufactured product is not always possible. Consequentially, new solutions are required to reliably link a specific product to its digital information to prove its authenticity and to remain accountable.

Nowadays, *digital signatures* are commonly used in the context of the Internet to provide authenticity and this concept can be extended easily to an IoP to provide similar verifiability there. To improve auditing and immutability capabilities of this traditional solution, *distributed ledgers* have proven to be a suitable approach. Blockchain [45, 46] as well as the IOTA tangle [51] allow establishing a persistent record of information and past dataflows, being a good fit in an environment where multi-hop traceability (along the supply chain) is a strict requirement. Similarly, *version control systems*, such as Git, are also suitable to track changes of data and to allow audits. These properties are required when dealing with a global knowledge system like an Internet of Production. However, in contrast to distributed ledgers, they are not tamper-proof. Moreover, current version control systems might not support industry-specific data formats without adjustments or overhead.

Platform Capabilities. Apart from the technical building blocks encountered so far, we can also make use of centrally deployed mechanisms that define and enforce rules for an IoP. On the one hand, the traditional idea of *access control* [57] can help to restrict the scope of data access by setting rules for all individual entities. However, such restrictions are only possible if the participating entities can be tracked. Here, approaches from the context of the Internet of Things [39] can be transferred to the industrial sector given that the overall attack vectors are similar. On the other hand, *policies* [32] directly attached to the data can offer similar flexibility [28] because usage or access constraints are directly attached to the data. Instead of defining access rules for each entity, policies constrain the scenarios where and how a specific piece of information can be used, i.e., they are independent of the entity processing the data giving some control to the data owner, i.e.,, the company.

From a different perspective, the concept of *smart contracts* [11, 18, 75] links the idea of blockchain with concepts of automated contracts. Consequentially, apart from proving the authenticity of information, smart contracts are also able to enforce the scope of data access to a certain extent. Previous work already showed a suitable prototype for the IoT [77], however, the impact of an IoP with flexible relationships on this design remains to be seen. A significant challenge is to determine the mode of operation.

Finally, *trusted computing* realizes an isolated enclave where guarantees about the running code and thereby about data accesses can be made [58]. Examples in this area include ARM TrustZone [3] and Intel SGX [12]. In the industrial context, this technique allows fulfilling most of our security challenges as it was developed to provide a secure area in insecure environments. However, the incoming and outgoing dataflows still must be carefully analyzed wrt. the derived challenges once they leave the secure environment. Besides, it contradicts interoperability and might result in a vendor lock-in because a specific trusted computing solution must be chosen. Furthermore, examples have shown in the past that security issues cannot be mitigated in a simple manner [9], i.e., they might require new hardware instead of (more) simple software patches. A requirement that is unlikely to be realistic for an IoP. Regardless, recent work [49] showed the feasibility in an industrial context while only partially addressing the required security aspects.

External Measures. The last category of building blocks contains supporting approaches that might help to realize an IoP without primarily focusing on the security of dataflows. They offer suitable approaches for inter-organization collaborations. Therefore, they affect different parts of our defined challenges while the aspects of all other categories are usually focused on specific areas.

To monetize the value of sensitive information (in the industrial context), (distributed) *data markets* enable all participating collaborators to sell and buy access to data. Besides mediating access to data, such a central data market can also ensure authentication and authorization, i.e., fulfill aspects wrt. data access. Depending on the exact implementation, confidentiality can also be ensured if data is only shared in an encrypted format. Recent examples, such as the International Data Space [5] or other data markets [42], have shown that centralized concepts to securely share information are feasible even in larger contexts. However, such a centralized approach shifts a lot of power to this market place which is in turn a valuable target for attackers of industry data and thus, might prevent a wide-spread and accepted application in an IoP.

A less technical approach to restrict the scope of data access and to establish authenticity of information would be to rely on *legal contracts* [4]. They allow defining all kinds of requirements prior to the first initiated dataflow. However, such negotiations are not yet automated in any way and, therefore, might prove infeasible in a highly dynamic IoP. Regardless, the concept can be used to set a frame in which entities are willing to collaborate and then negotiate the exact parameters in an automated way. Even though such an implementation would also allow entities to define sanctions in case of misbehavior, monitoring their actions and identifying data leaks from a remote vantage point is extremely

challenging. Consequentially, this building block might only be applicable for dataflows in long-lasting business relationships.

To still facilitate automated data exchanges in an IoP, the different collaborators could also make use of *smart payments* [34] which allows them to automatically initiate data transfers once the recipient made a payment (for sensitive information). As this building block might be based on distributed ledger technology, it also supports auditing. However, it does not deal with securing the data and data access in any way, i.e., instead of securing existing dataflows, it makes new dynamic dataflows accessible.

Takeaways. Following our survey of building blocks, we can conclude that no one-fits-all solution that addresses the large variety of security and privacy challenges in an IoP is available. Instead, stakeholders must currently address their specific needs against malicious entities individually, limiting their participation to collaborations that match their standards, i.e., they are unable to participate globally. This finding especially holds for new (IoP-enabled) types of dataflows which challenge today's circumstances. Individual building blocks are only suitable for a small subset of the identified dataflow challenges. Especially, well-founded research in the direction of industrial needs of confidentiality and anonymity is still in its infancy, resulting in an insufficient coverage for real-world deployments. Furthermore, the transfer and application of established approaches from other domains to the industrial domain—while an interesting approach—has not been tackled adequately so far. A challenging aspect here is the scalability of approaches for future needs. Consequentially, available technologies that ensure security and privacy in an IoP are still mostly missing. Hence, further research to provide security and privacy in an IoP, especially considering novel inter-organizational dataflows, is highly-relevant.

In the near future, companies clearly have to define their needs wrt. secure industrial collaborations and its (new) dataflows.

7 CONCLUSION & FUTURE WORK

Interconnecting CPS of different stakeholders, as envisioned by an Internet of Production (IoP) to reduce product development costs and time to market as well as increase profit margins and general product quality, facilitates the need of analyzing resulting dataflows from a security and privacy perspective. To this end, we derive the entities (supplier, producer, collaborator, customer, and maintenance provider) that are part of such an IoP, identify characteristic dataflows, and analyze them wrt. their security and privacy perspective based on two real-world use cases (a fine blanking line and a discrete manufacturing connected job shop). Comparable dataflows to realize sophisticated inter-organizational collaboration are non-existent in today's production landscape, leading to novel, previously undiscovered, requirements.

Methodology. We identify three large classes of security and privacy challenges in an IoP: *authenticity of information*, *scope of data access*, and *anonymity*. To provide an insight into solutions which might tackle these challenges, we conducted a survey of promising (technical) building blocks to secure inter-organizational dataflows and rated these building blocks regarding their potential to cover the individual security and privacy preferences of companies. We obtained an impression of their current and future potential through a brief analysis of existing deployments.

Target Audience. Our survey presented in this paper is intended to serve as an overview about the current state of dataflows in an envisioned IoP from a security and privacy perspective to spark future research. The presented building blocks for security and privacy in an IoP and their coverage of the identified dataflow challenges clearly indicate that future research in this direction is necessary to achieve two important goals: First, the individual building blocks must be improved to work as a viable solution in large-scale production landscape of interconnected CPS. Second, the different orthogonal aspects of security and privacy building blocks visualize the need to also test and analyze the joint application of different building blocks to enable a coverage of large parts of the identified dataflow challenges. Otherwise, envisioned advances will flatline due to a lack of security and privacy. By advancing the respective building blocks, we will be able to overcome security and privacy concerns that hinder the adoption of an IoP and hence create an important foundation for actually realizing the improvements envisioned by the Internet of Production.

As a next step, research must trigger and encourage (production) companies to clearly communicate their decisive needs to enable a tailoring of security and privacy building blocks towards an IoP. Without significant advances in security research, new concepts of valuable process data, such as digital shadows [59], cannot be implemented securely and privacy-preserving, effectively hindering an adoption for inter-organizational use cases on a global scale.

Future Work. We envision the road towards an implemented Internet of Production to consist of different steps with individual milestones: (i) improving *existing industrial business relationships* through inter-organizational dataflows, (ii) integrating data of *non-competitors* into the local environment, and finally, (iii) turn into a production landscape that utilizes data even from direct *competitors*.

First, companies should explore these new dataflows in known settings, i.e., they can first experiment with (existing) trusted partners. Here, the risk of intellectual theft should be limited. Eventually, advances will allow them to establish new relationships more flexibly (e.g., to change suppliers based on specific requirements).

Second, the industry must be supported in integrating information of non-competitors (cf. Figure 1) to unlock additional data sources. The applicability of building blocks will determine the new advances in a setting where data leakage is increasingly critical.

Third, with the obtained knowledge, companies can strive towards an interconnected production landscape, relying on past experiences to prevent information leakage in this delicate endeavor. This step is the most challenging for research of technical building blocks as the setting imposes tight constraints wrt. data leakage.

Overall, the process of realizing high interconnectivity in today's production landscape will be both existing and challenging for security research, most likely, for at least a decade.

ACKNOWLEDGMENTS

The authors would like to thank the German Research Foundation (DFG) for the kind support within the Cluster of Excellence "Internet of Production" (IoP) under the project id 390621612. The authors are grateful for the fruitful discussions with Simon Knape during his time at the Laboratory for Machine Tools and Production Engineering (WZL) at RWTH Aachen University.

REFERENCES

[1] Mohamed Abdel-Basset et al. 2018. Internet of Things (IoT) and its impact on supply chain: A framework for building smart, secure and efficient systems. *Future Generation Computer Systems* 86.

[2] Rakesh Agrawal et al. 2004. Order preserving encryption for numeric data. In *ACM SIGMOD*.

[3] Thiago Alves and Don Felton. 2004. TrustZone: Integrated Hardware and Software Security.

[4] Alvaro Arenas and Michael Wilson. 2008. Contracts as Trust Substitutes in Collaborative Business. *Computer* 41, 7.

[5] International Data Spaces Association. 2019. IDS Reference Architecture Model.

[6] Lennart Bader et al. 2018. Smart Contract-based Car Insurance Policies. In *IEEE GC Wkshps*.

[7] Mihir Bellare et al. 2000. Public-Key Encryption in a Multi-user Setting: Security Proofs and Improvements. In *EUROCRYPT*.

[8] Thomas Bergs et al. 2019. Stamping Process Modelling in an Internet of Production. In *CIRP TESConf*.

[9] Guoxing Chen et al. 2018. SgxPectre Attacks: Stealing Intel Secrets from SGX Enclaves via Speculative Execution. *arXiv:1802.09085*.

[10] Xiaofeng Chen. 2016. Introduction to Secure Outsourcing Computation. *Synthesis Lectures on Information Security, Privacy, & Trust* 8, 2.

[11] Konstantinos Christidis and Michael Devetsikiotis. 2016. Blockchains and Smart Contracts for the Internet of Things. *IEEE Access* 4.

[12] Victor Costan and Srinivas Devadas. 2016. Intel SGX Explained. *IACR Cryptology ePrint Archive* 2016, 086.

[13] Marek A. Cyran. 2018. Blockchain as a Foundation for Sharing Healthcare Data. *Blockchain in Healthcare Today*.

[14] Markus Dahlmanns et al. 2019. Privacy-Preserving Remote Knowledge System. In *IEEE ICNP*.

[15] Cynthia Dwork et al. 2006. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *EUROCRYPT*.

[16] Cynthia Dwork et al. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9.

[17] Christian Esposito et al. 2016. Cloud Manufacturing: Security, Privacy, and Forensic Concerns. *IEEE CLOUD* 3, 4.

[18] Valentina Gatteschi et al. 2018. Blockchain and Smart Contracts for Insurance: Is the Technology Mature Enough? *Future Internet* 10, 2.

[19] Craig Gentry. 2009. Fully homomorphic encryption using ideal lattices. In *ACM STOC*, Vol. 9.

[20] René Glebke et al. 2019. A Case for Integrated Data Processing in Large-Scale Cyber-Physical Systems. In *HICSS*.

[21] Satashu Goel and Rohit Negi. 2008. Guaranteeing Secrecy using Artificial Noise. *IEEE Trans. Wireless Commun.* 7, 6.

[22] Oded Goldreich et al. 1991. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM* 38, 3.

[23] Mitch Greeley. 2018 (accessed June 16, 2019). *Fine Blanking cuts into Industry 4.0*. https://blog.feintool.com/en/fine-blanking-cuts-into-industry-4-0/.

[24] Karl-Heinrich Grote and Erik K. Antonsson. 2009. *Springer Handbook of Mechanical Engineering*. Springer.

[25] Wenbo He et al. 2007. PDA: Privacy-Preserving Data Aggregation in Wireless Sensor Networks. In *IEEE INFOCOM*.

[26] Tobias Heer et al. 2011. Security Challenges in the IP-based Internet of Things. *Wirel. Pers. Commun.* 61, 3.

[27] Martin Henze et al. 2017. *Network Security and Privacy for Cyber-Physical Systems*.

[28] Martin Henze et al. 2016. CPPL: Compact Privacy Policy Language. In *WPES*.

[29] Martin Henze et al. 2017. Privacy-preserving Comparison of Cloud Exposure Induced by Mobile Apps. In *EAI MobiQuitous*.

[30] Martin Henze et al. 2017. CloudAnalyzer: Uncovering the Cloud Usage of Mobile Apps. In *EAI MobiQuitous*.

[31] Jens Hiller et al. 2019. Tailoring Onion Routing to the Internet of Things: Security and Privacy in Untrusted Environments. In *IEEE ICNP*.

[32] Karin Höne and Jan Harm Petrus Eloff. 2002. Information security policy — what do international information security standards say? *Computers & Security* 21, 5.

[33] Sabina Jeschke et al. 2017. *Industrial Internet of Things and Cyber Manufacturing Systems*.

[34] Merve Can Kus Khalilov and Albert Levi. 2018. A Survey on Anonymity and Privacy in Bitcoin-Like Digital Cash Systems. *IEEE Commun. Surveys Tuts.* 20, 3.

[35] Fritz Klocke and Aaron Kuchle. 2009. *Manufacturing Processes*. Springer.

[36] Kurt Lange. 1997. *Cold forming and fineblanking*. Wetzlar.

[37] Jay Lee et al. 2015. A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems. *Manuf. Lett.* 3.

[38] Yehida Lindell. 2005. *Secure Multiparty Computation for Privacy-Preserving Data Mining*.

[39] Jing Liu et al. 2012. Authentication and Access Control in the Internet of Things. In *ICDCS*.

[40] Jon Loeliger and Matthew McCullough. 2012. *Version Control with Git: Powerful tools and techniques for collaborative software development*. O'Reilly Media.

[41] Mimi Ma et al. 2018. Certificateless Searchable Public Key Encryption Scheme for Industrial Internet of Things. *IEEE Trans. Ind. Informat.* 14, 2.

[42] Roman Matzutt et al. 2017. myneData: Towards a Trusted and User-controlled Ecosystem for Sharing Personal Data. In *INFORMATIK*.

[43] Silvio Micali and Phillip Rogaway. 1991. Secure Computation. In *CRYPTO*.

[44] David C Mills et al. 2016. Distributed Ledger Technology in Payments, Clearing, and Settlement.

[45] Malte Möser et al. 2018. An Empirical Analysis of Traceability in the Monero Blockchain. *PoPETS* 2018, 3.

[46] Satoshi Nakamoto et al. 2008. Bitcoin: A peer-to-peer electronic cash system.

[47] Torben Pryds Pedersen. 1992. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In *CRYPTO*.

[48] Jan Pennekamp et al. 2019. Towards an Infrastructure Enabling the Internet of Production. In *IEEE ICPS*.

[49] Sandro Pinto et al. 2017. IIoTEED: An Enhanced, Trusted Execution Environment for Industrial IoT Edge Devices. *IEEE Internet Comput.* 21, 1.

[50] Stephen Pollard et al. 2018. Authentication of 3D Printed Parts using 3D Physical Signatures. In *NIP & Digital Fabrication Conference*, Vol. 2018.

[51] Serguei Popov. 2016. The tangle.

[52] Alexander Pretschner et al. 2006. Distributed usage control. In *Commun. ACM*.

[53] Michael O Rabin. 2005. How To Exchange Secrets with Oblivious Transfer. *IACR Cryptology ePrint Archive* 2005.

[54] Ragunathan Rajkumar et al. 2010. Cyber-physical systems: The next computing revolution. In *DAC*.

[55] Ronald Rivest et al. 1978. A Method for Obtaining Digital Signatures and Public-key Cryptosystems. *Commun. ACM* 21, 2.

[56] Ahmad-Reza Sadeghi et al. 2015. Security and Privacy Challenges in Industrial Internet of Things. In *DAC*.

[57] Ravi S. Sandhu and Pierangela Samarati. 1994. Access control: principle and practice. *IEEE Commun. Mag.* 32, 9.

[58] Nuno Santos et al. 2009. Towards trusted cloud computing. *HotCloud* 9, 9.

[59] Günther Schuh et al. 2018. The Digital Shadow of Services: A Reference Model for Comprehensive Data Collection in MRO Services of Machine Manufacturers. *Procedia CIRP* 73.

[60] Adi Shamir. 1979. How to share a secret. *Commun. ACM* 22, 11.

[61] Elaine Shi et al. 2011. Privacy-Preserving Aggregation of Time-Series Data. In *NDSS*.

[62] Jordi Soria-Comas and Josep Domingo-Ferrer. 2016. Big Data Privacy: Challenges to Privacy Principles and Models. *Data Science and Engineering* 1, 1.

[63] Markus Stadler. 1996. Publicly Verifiable Secret Sharing. In *EUROCRYPT*.

[64] Ralf Steinmetz and Klaus Wehrle. 2005. *Peer-to-Peer Systems and Applications*. Springer.

[65] Tim Stock and Günther Seliger. 2016. Opportunities of Sustainable Manufacturing in Industry 4.0. *Procedia CIRP* 40.

[66] Latanya Sweeney. 2002. k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzz.* 10, 5.

[67] Hassan Takabi et al. 2010. Security and Privacy Challenges in Cloud Computing Environments. *IEEE Security Privacy* 8, 6.

[68] Daniel Trauth et al. 2018 (accessed June 14, 2019). *Manufacturing Economy*. https://medium.com/industrial-iota-lab-aachen-wzl-of-rwth-aachen/manufacturing-economy-e541066889ee.

[69] Thomas H.-J. Uhlemann et al. 2017. The Digital Twin: Demonstrating the Potential of Real Time Data Acquisition in Production Systems. *Procedia Manuf.* 9.

[70] Marten Van Dijk et al. 2010. Fully Homomorphic Encryption over the Integers. In *EUROCRYPT*.

[71] Diana M. Segura Velandia et al. 2016. Towards industrial internet of things: Crankshaft monitoring, traceability and tracking using RFID. *Robot. Comput. Integr. Manuf.* 41.

[72] Herman Voigts et al. 2018. Dependencies of the die-roll height during fine blanking of case hardening steel 16MnCr5 without V-ring using a nesting strategy. *Int. J. Adv. Manuf. Technol.* 95, 5.

[73] Matthew Waller et al. 2008. Supply chain security: an overview and research agenda. *The International Journal of Logistics Management*.

[74] Michael E. Whitman and Herbert J. Mattord. 2011. *Principles of Information Security* (4th ed.). Course Technology Press.

[75] Gavin Wood et al. 2014. Ethereum: A secure decentralised generalised transaction ledger.

[76] Chunyong Yin et al. 2018. Location Privacy Protection Based on Differential Privacy Strategy for Big Data in Industrial Internet of Things. *IEEE Trans. Ind. Informat.* 14, 8.

[77] Yuanyu Zhang et al. 2018. Smart Contract-Based Access Control for the Internet of Things. *IEEE Internet Things J.*

[78] Qide Zheng et al. 2019. State-of-the-art and future challenge in fine-blanking technology. *Production Engineering* 13, 1.

[79] Liang Zhou and Han-Chieh Chao. 2011. Multimedia traffic security architecture for the internet of things. *IEEE Netw.* 25, 3.

[80] Jan Henrik Ziegeldorf et al. 2017. BLOOM: BLoom filter based Oblivious Outsourced Matchings. *BMC Medical Genomics* 10, Suppl 2.