Evolving the Industrial Internet of Things: The Advent of Secure Collaborations

Jan Pennekamp
Communication and Distributed Systems
RWTH Aachen University, Germany
pennekamp@comsys.rwth-aachen.de

***Abstract*—The Industrial Internet of Things (IIoT) leads to increasingly-interconnected industrial processes and environments, which, in turn, result in stakeholders collecting a plethora of information. Even though the global sharing of information and industrial collaborations in the IIoT promise significant improvements concerning productivity, sustainability, and product quality, among others, the majority of stakeholders is hesitant to implement them due to confidentiality and reliability concerns. However, strong technical guarantees could convince them of the contrary. Thus, to address these concerns, our interdisciplinary efforts focus on establishing and realizing secure industrial collaborations in the IIoT. By applying private computing, we are indeed able to reliably secure collaborations that not only scale to industry-sized applications but also allow for use case-specific confidentiality guarantees. Hence, improvements that follow from industrial collaborations with (strong) technical guarantees are within reach, even when dealing with cautious stakeholders. Still, until we can fully exploit these benefits, several challenges remain, primarily regarding collaboration management, introduced overhead, interoperability, and universality of proposed protocols.**

***Index Terms*—security; privacy; private computing; reliability**

I. Introduction

The advent of the Internet of Things (IoT) and networked Cyber-Physical Systems (CPSs) allows for the collection of vast amounts of business and production data [1]. For the first time, we now have the opportunity to facilitate an Internet-like knowledge exchange in the Industrial IoT (IIoT) [2], i.e., a large, distributed network to globally share information, even among mutually-distrusting stakeholders [3]. Specifically, the research cluster "Internet of Production" [4], [5] stipulates a corresponding vision to improve innovation, productivity, sustainability, and product quality, among others.

To turn this vision into reality, several dimensions, including operational security and legal ramifications, have to be addressed jointly to lastingly evolve the IIoT [6]. However, most importantly, we identify information security to be essential for this upcoming evolution because many stakeholders still hesitate to participate in any knowledge sharing due to confidentiality concerns. They act cautiously because their sensitive information constitutes their competitive advantage. Consequently, research must secure corresponding information flows to provide stakeholders with (strong) technical guarantees.

We thus introduce the notion of "secure collaborations" to enable grouping (future) activities, related to properly securing information flows in the IIoT, under a common term:

> *A secure collaboration considers the sensitivity of exchanged or shared information (i.e., knowledge) as part of global dataflows to account for the confidentiality needs of stakeholders in a networked IIoT.*

That being said, secure collaborations go beyond simple data sharing. On a more conceptual level, we can refer to this collective term as a managed and secure overlay approach to the networked IIoT. Exceeding this superficial abstraction, the utility of collaborations specifically profits from the novel availability of sensors, information, and compute resources.

Our work focuses on the confidentiality and reliability concerns of involved stakeholders in the IIoT to address the outlined research gap, namely, providing technical guarantees for secure collaborations. Down the road, further steps are needed to fully close the gap because, today, we largely lack generic approaches and suitable protocols to generalize our contributions to all sorts of use cases, deployments, and settings. Hence, we postulate the need to pursue the development of blueprints on how to design, set up, and manage secure collaborations as well as their underlying technical framework.

To systematically tackle the first step, we consider different types of collaborations, most importantly, (i) along and (ii) across supply chains. The former term refers to exchanging information up- and downstream of established supply chains and is also known as vertical collaborations; the latter covers the flow of information and knowledge between different (independent) supply chains and closely corresponds to horizontal collaborations [7]. Sourcing this diversity in our covered settings allows us to eventually draw conclusions on how to evolve processes, protocols, and organizations to establish the aforementioned Internet-like knowledge exchange in the IIoT.

Paper Organization. This dissertation [6] digest is structured as follows and covers five years of research on network and service management, designing communication and information-sharing protocols, and research methodology. First, in Section II, we introduce the IIoT in more detail along with the required technical foundation from the area of private computing that allows us to reliably secure collaborations. Subsequently, in Section III, we detail our primary research contributions, separated by their type of collaboration, i.e., along and across supply chains. Moving on, in Section IV, we then discuss important key takeaways for the management and operation of secure collaborations in an evolved IIoT. Finally, in Section V, we elaborate on our abstract methodology to bet-

 DOI: 10.1109/NOMS59830.2024.10575325

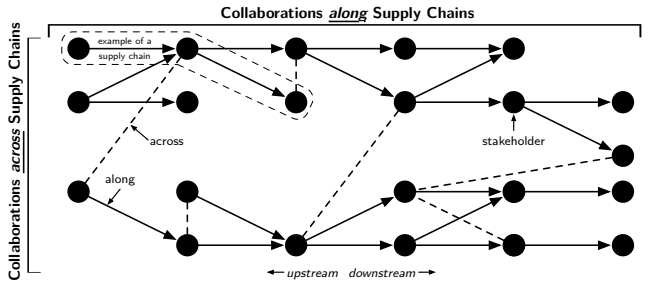

Fig. 1. Graphs can suitably model supply chains and embedded collaborations.

ter structuring interdisciplinary (security) research, e.g., in the IIoT, before concluding this dissertation digest in Section VI.

II. Background, Motivation, and Challenges

As preliminaries for the remainder of this digest, we now look at the relevant building blocks in more detail. To this end, we first put supply chains (and collaborations along and across) into perspective with respect to the IIoT. Additionally, we introduce private computing on a high level to raise conceptual challenges when applying it to secure collaborations.

A. Formalizing Supply Chains for Application in the IIoT

Given the diversity of actors and organizations in the IIoT, several entities have relevance for secure collaborations [8]. Depending on the specific information flow and setting, they can take different roles, e.g., supplier, manufacturer, customer (consumer), or collaborator. Here, the established trust and the individual confidentiality needs largely depend on the relationship between these entities. Especially in volatile settings with frequently changing business partners, introducing trust is a significant challenge [9], and thus, calls for technical means that could replace complex (paper-based) legal contracts.

Since CPS- and production site-specific improvements cannot make up for the expected benefits of globally-exchanged knowledge, information increasingly flows across mutually-distrusting organizations [3]. Conceptually, we can thus distinguish two overarching types of collaborations in the IIoT: (1) along supply chains or (2) across supply chains. Under this umbrella, participating entities can then focus on pursuing various use cases [7]: collaborative planning, designing supply chains, tracking, tracing, integrating critical infrastructures, or sharing product information. On a formal level, we can neatly express supply chain relationships in directed acyclic graphs (DAGs), as we illustrate for two networks in Figure 1. Once secure industrial collaborations are more widely in practical use, we expect that the list of (traditional) use cases will further evolve and expand in terms of scope, scale, and impact.

To study the feasibility of our proposed approaches, we evaluated our designs for securing collaborations extensively using several real-world use cases: For details on the use cases, (i) product composition and product properties, (ii) operation and procurement of machine tools, (iii) internal and external company benchmarking, and (iv) sharing and exchanging production parameters, we refer to the dissertation [6, Chapter 3].

B. The Boon and Bane of Private Computing

Moving toward the technological foundation of secure collaborations, holistically, we would have to consider operational security, network security, and information security [6], [10]–[12]. However, we examine the latter due to our focus on exchanging information and knowledge as well as its significance for successful collaborations in an evolved IIoT. Specifically, we consider dealing with malicious-but-cautious adversaries [13]. Assuming this attacker model is reasonable because involved organizations have an incentive to cheat, but they also depend on their public reputation and are bound to specific legislation.

When going for technical guarantees in the area of information security, private computing [14] is a fitting candidate since it covers the intersection of privacy-preserving computations and confidential computing, i.e., we are open to both software- and hardware-based approaches. In the following, we give a few examples that allowed us to realize reliable guarantees when designing novel protocols for secure collaborations, namely, trusted execution environments (TEEs), homomorphic encryption (HE), attribute-based encryption (ABE), private set intersection (PSI), and oblivious transfers (OTs), among others. Moreover, we are open to augmenting this security foundation with other conceptual approaches like blockchain technology or federated learning as long as they support us in securely advancing the exchange of knowledge. Based on our building block survey [6], [8], we have to conclude (and confirm) that, so far, no building block is able to simultaneously provide all relevant security aspects [8], i.e., authenticity of information, scope of data access, and anonymity. Thus, until the available building blocks have extensively evolved, research must resort to picking and combining them to craft secure collaborations for specific use cases and settings.

Since the vision of globally exchanging knowledge in the IIoT is a recent idea and secure collaborations are not yet in large-scale use (e.g., to benefit from IIoT-related improvements), the suitability of private computing to secure industry-sized applications is mostly unknown. Therefore, we systematically studied this challenging research gap for both types of collaborations as part of the aforementioned dissertation.

III. Secure Collaborations in the Industrial IoT

After presenting the proposed technical foundations of and motivation for secure collaborations, we now introduce the dissertation's primary technical contributions in more detail.

A. Secure Collaborations Along Supply Chains ($\rightleftarrows$)

For the first type of collaboration, which involved organizations are arguably less concerned about, we distinguish two settings. First, we look at securely exchanging information within an established volatile supply chain network. Second, we pick up the volatility and present stakeholders with protocols to reduce the disclosure of sensitive information when establishing business relationships with new, previously unknown partners, i.e., the during initial stages of procurement.

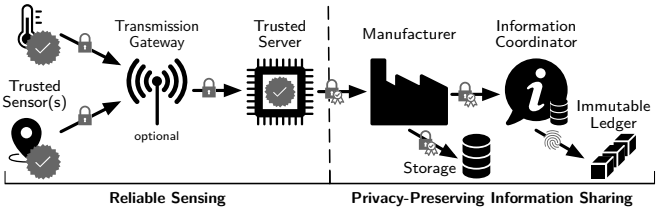

Fig. 2. Reliable information processing consists of two connected angles.

1) A Reliable Information Processing Pipeline: The challenge of establishing a secure and reliable exchange of information along supply chains consists of two angles: (1) challenges related to reliable sensing even in otherwise untrusted environments, i.e., ensuring the authenticity and correctness of processed data, and (2) privacy-preserving information flows between mutually-distrusting entities, also over multiple hops.

We thus look at both angles individually while ensuring their overall compatibility, as we visualize in Figure 2.

Sensing. To improve the reliability and trustworthiness of sensed information, we present the notion and concept of *reliable end-to-end (E2E) sensing* [15], [16]. In this context, we are the first to rely on (i) trusted sensors, which verifiably secure the sensing by utilizing TEEs, and (ii) subsequent processing of sensed data, using only TEE-backed infrastructure. As a result, our approach provides verifiable technical guarantees regarding the authenticity and correctness of processed data, even if it was initially sensed, processed, and forwarded in otherwise untrusted remote environments. In connection with supply chain information systems that ensure verifiable and accountable information retrieval (cf. second angle on information *sharing* below), these guarantees even hold long-term. Our evaluation underlines the feasibility of our proposed approach in terms of performance, costs, and security for various application areas, ranging from status and location tracking to integrity, condition, and visual monitoring.

Sharing. Given the lack of supply chain information systems that (i) allow for privacy-preserving sharing of information over multiple hops while (ii) also supporting volatile and complex supply chain networks, we developed *PrivAccI-Chain* [17], [18], our design for establishing accountable and confidential information flows along supply chains. Specifically, we rely on ABE and its concept of policies to model and implement fine-granular access control even for settings with flexible and highly-dynamic business relationships. As a result, the exact (intended) recipients do not have to be known when encrypting the information, and PrivAccIChain does not even require any later involvement of data-providing entities.

The use of (encrypted) tracing references neatly enables multi-hop information flows and further allows for efficient traversal of entire supply chain networks. This traversal is especially beneficial for the use case of tracing, both downstream (→) and upstream (←). Based on our evaluation of two real-world applications (one focusing on the supply chain of a real-world urban vehicle [18]), we conclude that PrivAccIChain's performance is satisfactory for large-scale deployments.

Next, we intend to fuse both angles to improve the reliability of supply chain reputation systems while ensuring confidentiality [19]. Likewise, future work should rigorously survey which mechanisms for key management and key exchange are most suitable for large-scale use in an evolved IIoT.

2) Privacy-Preserving Purchase Inquiries: While conducting our research, jointly with practitioners, we noticed that establishing new business relationships as part of procurement is still severely hindered by confidentiality concerns stemming from the need to share sensitive information upfront. Since an evolving IIoT will increasingly require stakeholders to bootstrap business relationships between mutually-distrusting organizations, our research [20], [21] on privacy-preserving purchase inquiries (the first step during procurement, cf. [21, Appendix A]) is paramount. To the best of our knowledge, we are the first to address this research gap.

Specifically, we propose multiple designs [21] with slightly differing confidentiality guarantees that ensure two-way privacy for purchase inquiries for different settings. In addition to an intuitive PSI-based design, called *PPI*, that utilizes two computational phases (product then price) to process a purchase inquiry, we also introduce two HE-based designs, *HPI* and *cHPI*, which build on computations on encrypted inputs. Our designs only handle bilateral inquiries, i.e., information is never exposed publicly or to uninvolved parties. However, as certainly required, stakeholders can still trigger concurrent protocol runs with every organization they are interested in.

Our evaluation further underlines that our designs are suitable to securely realize this novel example of industrial collaborations. While our protocols' security builds on well-established building blocks and their attested security, they also ensure the confidentiality needs of both buyers and sellers (PPI with minor deductions). We also demonstrate their real-world feasibility based on two real-world applications. Besides, they also scale well with the number of potential sellers that should be considered as the protocol runs are independent of each other, i.e., buyers can trigger as many runs as needed and computationally supported at once. Hence, our work provides stakeholders with sufficient flexibility when establishing new business relationships in an evolved IIoT.

After this presentation of our two technical contributions for reliably securing collaborations along supply chains (and within volatile supply chain networks), we next shift our focus to exchanging knowledge across (independent) supply chains.

B. Secure Collaborations Across Supply Chains (↑↓)

Setting the stage for collaborations across supply chains by providing reliable technical guarantees is an important but challenging endeavor since they have rarely been studied so far. The primary reason is that stakeholders are simply too concerned to be involved because their sensitive information is mostly exchanged with unknown (untrusted) entities, potentially even their competitors. With our research, we intend to change this undesired momentum by demonstrating the feasibility of *secure* collaborations with varying degrees of invasiveness (on the management and operation of businesses).

1) Privacy-Preserving Comparisons: As an example of a collaboration with few direct implications for the participating businesses, we looked at comparisons of business data, i.e., no external information is directly fed back into (local) processes. In particular, we considered industrial benchmarking, an activity that is frequently performed these days, however, without sufficiently considering all confidentiality requirements. Specifically, related work largely fails to protect the underlying algorithm that is used to compute the benchmark [22], [23]. This essential piece is valuable to the operator of the benchmark since it constitutes its competitive advantage.

With this contribution, we can outline the design space (in terms of building blocks from private computing) that is available when composing protocols for secure collaborations. Specifically, we realize the same task with two diametrical strains, namely, HE as software-based and TEEs as hardware-based foundation. The high-level protocol does not differ between both designs, *SW-PCB* and *HW-PCB* [23], highlighting the interchangeability. However, we still need to account for building block-specific limitations, for example, in terms of the precision of computations when using HE in SW-PCB.

Our evaluation [23] shows that both strains are readily available to secure benchmarks in real-world deployments. Thus, when rolling out corresponding secure collaborations, the key question is which conceptual technology should serve as the root of trust, i.e., trusted hardware (a TEE) or an HE scheme, mainly because the remaining properties do not prohibit practical realizations. Both designs fulfill the performance requirements, with HW-PCB computationally outperforming SW-PCB. While HW-PCB's accurate computations promise quick and precise results, SW-PCB is easier to deploy as it is designed for untrusted hardware. The exact realization (design) then likely depends on the availability of a TEE and the willingness to build on its associated security assumptions (e.g., trusting the underlying security concept, the vendors, and remote attestation). Otherwise, HE-based implementations also promise secure and practical benchmarks for the IIoT.

Having this non-invasive example in mind, we continue with a slightly more invasive application in the following.

2) Privacy-Preserving Matchings: In contrast to the previous example, for this contribution, we worked on the challenge to match information in a privacy-preserving way, even when dealing with fuzzy queries, i.e., identifiers. Practitioners motivated us to tackle this example of a secure collaboration because they lacked adequate training data to apply their transfer learning on, even though, in theory, sufficient information is available across different organizations [24].

Unfortunately, surveying related work revealed that prior approaches are not able to satisfy the (confidentiality) needs of both data-providing and data-querying entities. To mitigate this situation, we proposed a modular concept [6] to realize a privacy-preserving exchange platform for the IIoT. This concept utilizes Bloom filters, PSIs, and OTs, depending on the exact configuration (*BPE*, *PPE*, or *OPE*). As a result, individual deployments can be tailored to address use case-specific confidentiality, performance, and scalability needs.

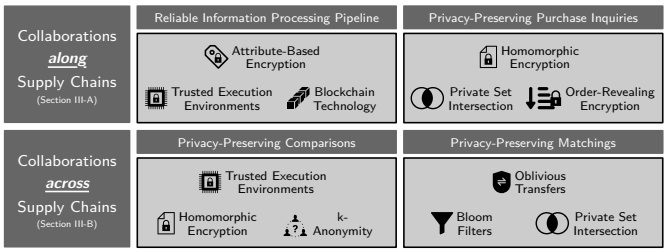

Fig. 3. The careful selection of appropriate building blocks is paramount.

The evaluation of these configurations shows that the runtime to offload records to the exchange platform is negligible. Our performance evaluations of synthetic inputs and two real-world applications (each with multiple queries) demonstrate the feasibility of BPE for these settings. In contrast, the computationally more demanding concepts, PPE and OPE, only scale to smaller settings. By realizing the potentially-sensitive computations locally at querying entities, we account for their confidentiality needs. While the Bloom filter-based matching in BPE slightly violates the data providers' confidentiality desires, the PSI-based matching in PPE and OPE reliably addresses this drawback. Importantly, this work is not limited to a specific domain or type of data that should be exchanged across supply chains. The protocols only require a suitable and globally agreed-upon indexing scheme for the information.

This application concludes the presentation of the dissertation's primary technical contributions, where we successfully covered collaborations along and across supply chains. Based on the presented examples, we can follow that designing secure, IIoT-focused protocols that scale to industry-sized applications is possible, even if we mandate the configuration of use case-specific *technical* confidentiality guarantees by sourcing current building blocks from private computing.

IV. Takeaways for Service Management and Operation in the Evolving IIoT

We note that related work in the area of "secure collaborations" is rather sparse. Relevant large-scale initiatives, such as Alice [25], Gaia-X [26], or IDS [27], still mostly focus on organizational security [28], [29] instead of ensuring strong technical guarantees (which we consider to be crucial in light of the likelihood of dealing with malicious-but-cautious adversaries, cf. Section II-B). Regardless, based on the dissertation's findings, we can already derive several important takeaways.

Takeaways. Based on our contributions, we identify two key results. First, we can implement secure collaborations even in settings with malicious-but-cautious adversaries that reliably improve the status quo for stakeholders in the IIoT. Second, as we also summarize in Figure 3, we have to carefully select appropriate technical building blocks according to the needs of the use case at hand. This selection is especially challenging when having to balance opposing desires, for example, the trade-off between transparency and privacy preservation [18]. As a result, we build on a bouquet of building blocks, which, for this reason, covers software- and hardware-based concepts.

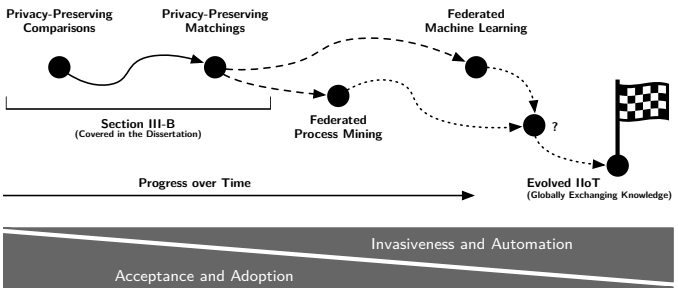

Fig. 4. The level of automation related to collaborations is likely to increase.

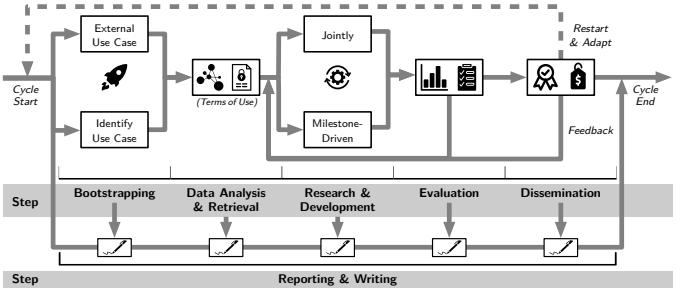

Fig. 5. This process cycle is intended to strengthen interdisciplinary research.

Open Challenges. We further identify several remaining aspects that could hinder the wide dissemination of secure collaborations in the wild. Specifically, these aspects are related to embedding secure collaborations and their protocols in the IIoT. That is, stronger focus should be put on *collaboration management*, e.g., how businesses set them up, how they can discover fitting collaborators, and how to globally deploy new protocols for new use cases, among other things. Additionally, more resources should be allocated to investigate the *overhead* that secure collaborations introduce, not only in terms of computational performance but also in terms of costs that accumulate due to their provision, management, and operation. To eventually come to a situation where collaboration protocols are universal, i.e., independent of specific use cases, *interoperability* between different designs and underlying technical frameworks is needed. While the FactDAG model [30] might be a candidate to address this challenge for data, in a truly evolved IIoT, we still need a semantic model that concisely represents and captures available knowledge across all stakeholders to make it accessible. Lifting the FactDAG model to such a solution could be a worthwhile approach.

Application Evolution within Reach. As a next step, secure collaborations are likely to further evolve from comparisons over matching to sophisticated federated settings, which include machine learning and process mining (Figure 4). These advances are particularly interesting to increasingly exploit collaborations across supply chains. However, corresponding applications are likely to implicitly feed knowledge directly into local processes, further impacting the collaborations' invasiveness. Such autonomous operations are thus challenged by the threat of adversarial behavior, most prominently in light of safety and environmental issues. Consequently, research must come up with designs that reliably deal with (i) untrusted entities who have an incentive to misbehave, (ii) potentially unauthentic (untrustworthy) information, and (iii) the computational overhead of introducing additional "layers" of security.

The Need for a Blueprint. As hinted at before, today's secure collaborations are not yet of universal nature, i.e., we mostly have to tailor new protocols for different use cases. Thus, we are in need of a blueprint on how to design, set up, and manage secure collaborations as well as their underlying technical framework, including the selection, configuration, and development of building blocks, in an evolved IIoT.

V. Research Methodology

In addition to our technical contributions, we also made an effort to conserve and formalize our experience in this interdisciplinary environment by deriving an abstract research methodology for global use [31]. This abstract process cycle is meant to support researchers who deal with security challenges in somewhat applied interdisciplinary research. Especially recent developments like the IoT and networked CPSs can serve as a catalyst for significant (applied) innovation. By illustrating our experience, we are able to provide a realistic overview of typical challenges and pitfalls in such environments.

Apart from supporting research related to the aforementioned blueprint for secure collaborations (cf. Section IV), our goals for this methodological contribution are twofold. First, we want to ease the challenges of and reservations against interdisciplinary collaborations. Second, we hope to contribute to bootstrapping additional security research that also considers applications in other domains to further boost research on building blocks from private computing. Eventually, we could even be able to derive flexible building blocks that can easily be re-used across use cases, settings, and domains.

As we visualize in Figure 5, except for the accompanying "Reporting & Writing" step, all other steps for conducting interdisciplinary research build upon each other (and can be revisited as needed). Once a full cycle has been completed, another use case can be tackled while also incorporating recent experience. As such, every iteration also influences future challenges and potentially contributes to their resolutions. Finally, we also looked into the connection between our methodology and research data management [31, Section V] to further strengthen the acceptance and impact of our work.

VI. Conclusion and the Road Ahead

For this digest, we extracted the main findings and takeaways from a recent dissertation [6]. While the dissertation also features various interdisciplinary aspects and discussions with relevance to both academia and industry from several domains, this digest is written for computer scientists with a strong interest in the evolution of the Industrial IoT. That is, its contributions allow us to answer the information security-focused research question "*How can we enable secure industrial collaborations in real-world settings?*" for the first time.

Our work has shown that today's building blocks from private computing are suitable to reliably realize secure collaborations in the IIoT, even in settings with strong confidentiality requirements, as exemplified by collaborations across supply chains. Due to our evaluation of diverse real-world use cases, we further argue that cleverly-designed collaborations, already today, also scale to the needs of businesses in industry.

While we are the first to systematically pursue research in this rapidly evolving area, our findings also show that a lot of effort lies ahead of us. Specifically, we lack a universal approach that can serve us (and future work) as a blueprint for designing, setting up, and managing secure collaborations. At this point, we still have to meticulously select, adapt, and evaluate protocols and designs in light of each use case, the targeted setting, and the involved stakeholders. Regardless, we are convinced that our work can serve as a trigger and enable more developments. Once we have overcome the aforementioned technological issues and have addressed possible safety concerns and legal uncertainties, secure collaborations will pave the way for a more efficient distribution of knowledge as well as improved resilience for all stakeholders in the IIoT.

We look forward to eventually seeing a secure and full-fledged Internet-like knowledge exchange in the IIoT in place.

Acknowledgments

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2023 Internet of Production – 390621612. We are further grateful for Ike Kunze's feedback on this digest.

References

[1] E. Sisinni, A. Saifullah, S. Han, U. Jennehag, and M. Gidlund, "Industrial Internet of Things: Challenges, Opportunities, and Directions," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 11, 2018.

[2] H. Boyes, B. Hallaq, J. Cunningham, and T. Watson, "The industrial internet of things (IIoT): An analysis framework," *Computers in Industry*, vol. 101, 2018.

[3] F. Basso, S. D'Amours, M. Rönnqvist, and A. Weintraub, "A survey on obstacles and difficulties of practical implementation of horizontal collaboration in logistics," *International Transactions in Operational Research*, vol. 26, no. 3, 2019.

[4] J. Pennekamp, R. Glebke, M. Henze *et al.*, "Towards an Infrastructure Enabling the Internet of Production," in *Proceedings of the 2nd IEEE International Conference on Industrial Cyber Physical Systems (ICPS '19)*. IEEE, 2019.

[5] P. Brauner, M. Dalibor, M. Jarke *et al.*, "A Computer Science Perspective on Digital Transformation in Production," *ACM Transactions on Internet of Things*, vol. 3, no. 2, 2022.

[6] J. Pennekamp, "Secure Collaborations for the Industrial Internet of Things," Ph.D. dissertation, RWTH Aachen University, 2024.

[7] J. Pennekamp, R. Matzutt, C. Klinkmüller *et al.*, "An Interdisciplinary Survey on Information Flows in Supply Chains," *ACM Computing Surveys*, vol. 56, no. 2, 2024.

[8] J. Pennekamp, M. Henze, S. Schmidt *et al.*, "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.

[9] J. Pennekamp, R. Matzutt, S. S. Kanhere, J. Hiller, and K. Wehrle, "The Road to Accountable and Dependable Manufacturing," *Automation*, vol. 2, no. 3, 2021.

[10] J. Pennekamp, M. Dahlmanns, L. Gleim, S. Decker, and K. Wehrle, "Security Considerations for Collaborations in an Industrial IoT-based Lab of Labs," in *Proceedings of the 3rd IEEE Global Conference on Internet of Things (GCIoT '19)*. IEEE, 2019.

[11] M. Dahlmanns and K. Wehrle, "Protocol Security in the Industrial Internet of Things," in *Proceedings of the 2024 IEEE/IFIP Network Operations and Management Symposium (NOMS '24)*. IEEE, 2024.

[12] M. Dahlmanns, F. Heidenreich, J. Lohmöller *et al.*, "Unconsidered Installations: Discovering IoT Deployments in the IPv6 Internet," in *Proceedings of the 2024 IEEE/IFIP Network Operations and Management Symposium (NOMS '24)*. IEEE, 2024.

[13] M. D. Ryan, "Enhanced Certificate Transparency and End-to-end Encrypted Mail," in *Proceedings of the 21st Annual Network and Distributed System Security Symposium (NDSS '14)*. Internet Society, 2014.

[14] Confidential Computing Consortium, "A Technical Analysis of Confidential Computing," Tech. Rep. Version 1.3, 2022.

[15] J. Pennekamp, F. Alder, R. Matzutt *et al.*, "Secure End-to-End Sensing in Supply Chains," in *Proceedings of the 2020 IEEE Conference on Communications and Network Security (CNS '20)*. IEEE, 2020.

[16] J. Pennekamp, F. Alder, L. Bader *et al.*, "Securing Sensing in Supply Chains: Opportunities, Building Blocks, and Designs," *IEEE Access*, vol. 12, 2024.

[17] J. Pennekamp, L. Bader, R. Matzutt *et al.*, "Private Multi-Hop Accountability for Supply Chains," in *Proceedings of the 2020 IEEE International Conference on Communications Workshops (ICC Workshops '20)*. IEEE, 2020.

[18] L. Bader, J. Pennekamp, R. Matzutt *et al.*, "Blockchain-Based Privacy Preservation for Supply Chains Supporting Lightweight Multi-Hop Information Accountability," *Information Processing & Management*, vol. 58, no. 3, 2021.

[19] L. Bader, J. Pennekamp, E. Thevaraj *et al.*, "Reputation Systems for Supply Chains: The Challenge of Achieving Privacy Preservation," in *Proceedings of the 20th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous '23)*. Springer, 2023.

[20] J. Pennekamp, F. Fuhrmann, M. Dahlmanns *et al.*, "Confidential Computing-Induced Privacy Benefits for the Bootstrapping of New Business Relationships," RWTH Aachen University, Tech. Rep. RWTH-2021-09499, 2021, Blitz Talk at the 2021 Cloud Computing Security Workshop (CCSW '21).

[21] J. Pennekamp, M. Dahlmanns, F. Fuhrmann *et al.*, "Offering Two-Way Privacy for Evolved Purchase Inquiries," *ACM Transactions on Internet Technology*, vol. 23, no. 4, 2023.

[22] J. Pennekamp, P. Sapel, I. B. Fink *et al.*, "Revisiting the Privacy Needs of Real-World Applicable Company Benchmarking," in *Proceedings of the 8th Workshop on Encrypted Computing & Applied Homomorphic Cryptography (WAHC '20)*. HomomorphicEncryption.org, 2020.

[23] J. Pennekamp, J. Lohmöller, E. Vlad *et al.*, "Designing Secure and Privacy-Preserving Information Systems for Industry Benchmarking," in *Proceedings of the 35th International Conference on Advanced Information Systems Engineering (CAiSE '23)*. Springer, 2023.

[24] J. Pennekamp, E. Buchholz, Y. Lockner *et al.*, "Privacy-Preserving Production Process Parameter Exchange," in *Proceedings of the 36th Annual Computer Security Applications Conference (ACSAC '20)*. ACM, 2020.

[25] S. Pan, D. Trentesaux, D. McFarlane *et al.*, "Digital interoperability in logistics and supply chain management: state-of-the-art and research avenues towards Physical Internet," *Computers in Industry*, vol. 128, 2021.

[26] A. Braud, G. Fromentoux, B. Radier, and O. Le Grand, "The Road to European Digital Sovereignty with Gaia-X and IDSA," *IEEE Network*, vol. 35, no. 2, 2021.

[27] B. Otto, S. Auer, J. Cirullies *et al.*, "Industrial Data Space: Digital Souvereignity over Data," Fraunhofer, White Paper, 2016.

[28] J. Lohmöller, J. Pennekamp, R. Matzutt, and K. Wehrle, "On the Need for Strong Sovereignty in Data Ecosystems," in *Proceedings of the 1st International Workshop on Data Ecosystems (DEco '22)*, vol. 3306, no. 51–63. CEUR Workshop Proceedings, 2022.

[29] J. Lohmöller, J. Pennekamp, R. Matzutt *et al.*, "The Unresolved Need for Dependable Guarantees on Security, Sovereignty, and Trust in Data Ecosystems," *Data & Knowledge Engineering*, 2024.

[30] L. Gleim, J. Pennekamp, M. Liebenberg *et al.*, "FactDAG: Formalizing Data Interoperability in an Internet of Production," *IEEE Internet of Things Journal*, vol. 7, no. 4, 2020.

[31] J. Pennekamp, E. Buchholz, M. Dahlmanns *et al.*, "Collaboration is not Evil: A Systematic Look at Security Research for Industrial Use," in *Proceedings of the Workshop on Learning from Authoritative Security Experiment Results (LASER '20)*. ACSAC, 2021.