Unlocking Secure Industrial Collaborations through Privacy-Preserving Computation

by Jan Pennekamp (RWTH Aachen University), Martin Henze (Fraunhofer FKIE) and Klaus Wehrle (RWTH Aachen University and Fraunhofer FKIE)

In industrial settings, significant process improvements can be achieved when utilising and sharing information across stakeholders. However, traditionally conservative companies impose significant confidentiality requirements for any (external) data processing. We discuss how privacy-preserving computation can unlock secure and private collaborations even in such competitive environments.

Recent developments demonstrate the value data science can have for industries. A prime example is the research cluster "Internet of Production" [L1], which aims to turn data into value throughout the entire product lifecycle, i.e., production, development, and usage. The cluster, which was established in 2019, brings together more than 200 engineers and computer scientists from more than 35 institutes at RWTH Aachen University and the Fraunhofer Society. Its key vision is to interconnect companies with the aim of exchanging knowledge and know-how globally (Figure 1), i.e., advancing use cases within and across domains to establish reliable, cost-efficient, sustainable, and accountable production. Not surprisingly, the involved industrial stakeholders mandate strict confidentiality concerning their data as they fear a loss of control [1]. To address these concerns, privacy-preserving computation with its diverse building blocks, such as homomorphic encryption (HE), private set intersection (PSI), or oblivious transfers (OTs), can act as a key enabler. Here, industrial settings provide unique challenges and opportunities compared to traditional privacy-preserving computation: While demanding strict confidentiality and scalability around data volumes and data rates, industrial settings can benefit from publicly known stakeholders, which depend on their reputation to conduct business, easing the identification and sanctioning of misbehaviour.

A research roadmap to unlock secure industrial collaborations

In this work, we report on our research roadmap for realising secure industrial collaborations, consisting of three research directions (Figure 2), sorted by increasing complexity. First, privacy-preserving comparisons allow companies to identify unrealised potential without an immediate feedback loop into existing processes. Extending on this idea, privacy-preserving matching provides a mechanism to retrieve information to directly improve local production, while still requiring manual interaction. Finally, privacy-preserving machine learning promises to feed newly derived knowledge directly into running production processes without manual interaction.

Privacy-preserving comparisons

A prominent application of privacy-preserving comparisons is company benchmarking, i.e., comparing business performance among companies. Studying real-world requirements for such benchmarks in industrial settings, we identified two key challenges [2]: First, meaningful benchmarks require complex and hierarchical computations of key performance indicators, imposing a significant burden for privacy-preserving computation. For example, in a benchmark for the injection moulding sector, one performance indicator, measuring the overall effectiveness of manufacturing equipment, covers 23 inputs and 83 calculations (23x addition/subtraction, 27x multiplication, 25x division, 8x minimum). Second, given these complex calculations, the benchmark algorithm itself becomes a valuable asset that warrants protection.

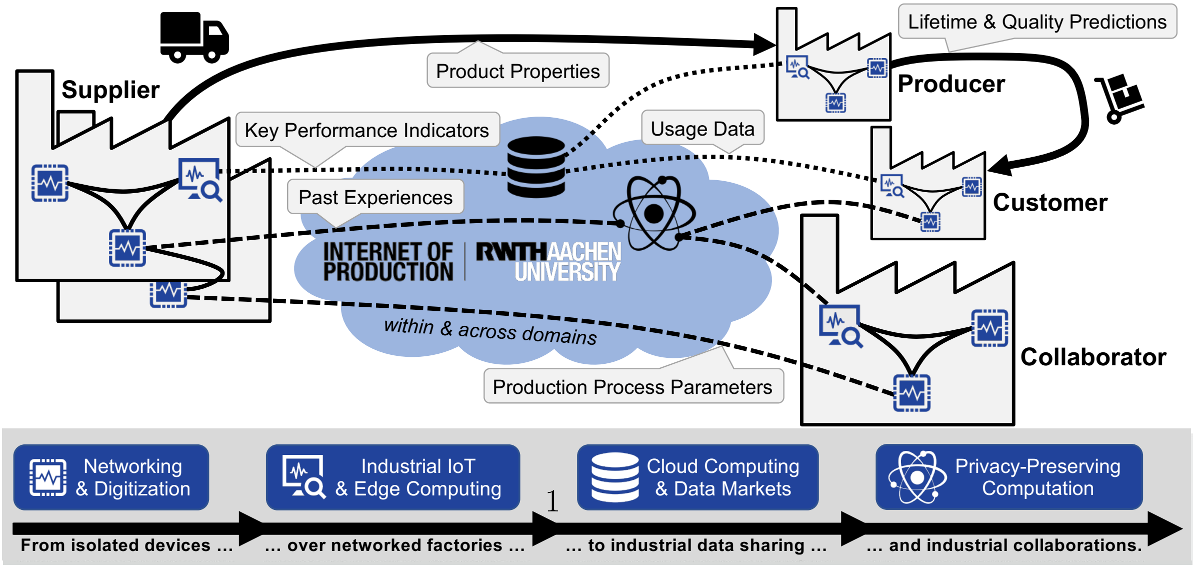
Supplier
Product Properties
Lifetime & Quality Predictions
Producer
Key Performance Indicators
Usage Data
Past Experiences
Customer
INTERNET OF PRODUCTION
RWTH AACHEN UNIVERSITY
within & across domains
Collaborator
Production Process Parameters
Networking & Digitization
Industrial IoT & Edge Computing
Cloud Computing & Data Markets
Privacy-Preserving Computation
From isolated devices ...
... over networked factories ...
... to industrial data sharing ...
... and industrial collaborations.

Figure 1: Privacy-preserving computation is a promising technology to unlock secure industrial collaborations, i.e., an exchange of knowledge that goes beyond simple data sharing, while still considering the confidentiality needs of companies.

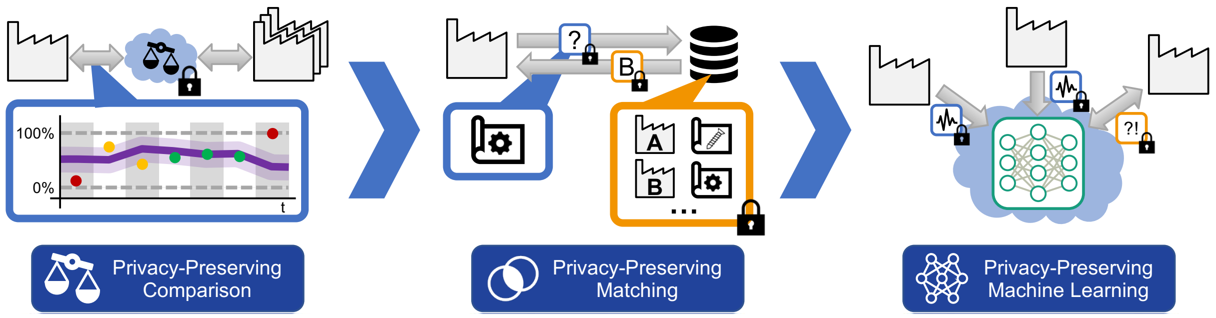
100%
0%
t
Privacy-Preserving Comparison
?
B
A
B
...
Privacy-Preserving Matching
?!
Privacy-Preserving Machine Learning

Figure 2: The increasing complexity of comparisons, matching, and machine learning also challenges the development of suitable privacy-preserving computation solutions, especially in industrial settings.

Addressing these challenges, we developed a secure solution for company benchmarking using HE that preserves not only the privacy of participating companies but also protects the benchmarking algorithm. We evaluate our approach by repeating a real-world benchmarking in the injection moulding industry, covering 48 distinct performance indicators calculated out of hundreds of input values. A runtime of 8.7 minutes per company and an average deviation of 0.16% compared to plaintext calculations underline the real-world applicability of our approach.

Privacy-preserving matching

Moving one step further in secure industrial collaboration, privacy-preserving matching directly impacts production processes, e.g., when commissioning and configuring new production lines. Traditionally, companies use empirical testing to identify machine parameters, which is costly and time-consuming. As others might already operate similar production lines, reusing this knowledge is a sensible and sustainable approach, if realised securely. To achieve this goal and thus allow companies to securely exchange information that can, e.g., be used to configure production sites, we developed two approaches with different privacy trade-offs [3]. By combining established building blocks (OTs, PSIs, and Bloom filters), we show the potential of industry-tailored privacy-preserving computation. To evaluate our approach, we (i) realise a process parameter retrieval for injection moulding to reduce ramp-up phases and (ii) exchange machine tool parameters to improve the machine settings for individual workpieces. Our evaluation shows that our approach meets today's real-world privacy and processing requirements. Thus, privacy-preserving computation can enable the secure exchange of sensitive industrial information, even in competitive environments.

Privacy-preserving machine learning

Finally, privacy-preserving machine learning provides a more tightly integrated knowledge sharing, directly feeding new (and improved) information into local industrial processes. For example, in high-pressure die casting, machine learning-based quality prediction allows defects to be discovered even when in-situ methods are not applicable. Here, technical advances promise to utilise a larger set of input data across stakeholders, i.e., using federated learning, and thus improve predictions. However, thoughtless decisions can result in unfounded high scrap rates, while in other settings, feedback loops can even result in physical harm, e.g., when coordinating line-less mobile assembly systems. Likewise, privacy-preserving computation must ensure that no sensitive information is leaked (indirectly), e.g., through dataset inference or reconstruction attacks. Our ongoing work aims to enhance the industry's decision-making and feedback loops by securely utilising external knowledge and data.

Conclusion

Privacy-preserving computation indeed promises to unlock sophisticated secure industrial collaborations. In the future, the relevance of such collaboration will further increase, as the goals of confidentiality and sustainability will complement today's dominant factors of costs and product quality. Until then, we need to address several research challenges [L2] to reliably and securely realise the vision of globally-interconnected production. These challenges are not limited to privacy-preserving knowledge exchange, but also include device and network security, among others.

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

Links:

[L1] https://www.iop.rwth-aachen.de

[L2] https://www.comsys.rwth-aachen.de/research/industrial-internet-of-things

References:

[1] J. Pennekamp et al.: "Dataflow Challenges in an Internet of Production: A Security & Privacy Perspective", ACM CPS-SPC, p. 27-38, 2019.

[2] J. Pennekamp et al.: "Revisiting the Privacy Needs of Real-World Applicable Company Benchmarking", WAHC, p. 31-44, 2020.

[3] J. Pennekamp et al.: "Privacy-Preserving Production Process Parameter Exchange", ACSAC, p. 510-525, 2020.

Please contact:

Jan Pennekamp
RWTH Aachen University, Germany
jan.pennekamp@comsys.rwth-aachen.de

Martin Henze
Fraunhofer FKIE, Germany
martin.henze@fkie.fraunhofer.de