Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

PRepChain: A versatile privacy-preserving reputation system for dynamic supply chain environments

Jan Pennekamp[a,*,1], Lennart Bader[b,1], Emildeon Thevaraj[a,c], Stefanie Berninger[c], Martin Perau[c], Tobias Schröer[c], Wolfgang Boos[c], Salil S. Kanhere[d], Klaus Wehrle[a]

[a] *Communication and Distributed Systems, RWTH Aachen University, Ahornstraße 55, Aachen, 52074, NRW, Germany*
[b] *Cyber Analysis & Defense, Fraunhofer FKIE, Fraunhoferstraße 20, Wachtberg, 53343, NRW, Germany*
[c] *Institute for Industrial Management at RWTH Aachen University, Campus-Boulevard 55, Aachen, 52074, NRW, Germany*
[d] *School of Computer Science and Engineering, University of New South Wales, Engineering Rd, Kensington, 2052, NSW, Australia*

ARTICLE INFO

Keywords:
Confidentiality
Unlinkability
Subjective and objective ratings
Homomorphic encryption
Trust

ABSTRACT

Despite their significant added value in the context of consumer-oriented e-commerce, reputation systems have seen limited adoption in other business settings and models these days. Yet, reliable reputation scores are essential in such settings for easing the establishment of new business relationships—an aspect that is particularly crucial in dynamic supply chain environments, where business partners change frequently. Existing approaches, however, usually target other application domains and fall short in addressing the specific challenges of dynamic supply chains—especially with respect to reliability (incl. availability) and privacy preservation (incl. confidentiality). To close this research gap and to support novel directions in this important research area, we propose PRepChain, our highly-configurable approach that leverages fully homomorphic encryption and distributed competences to provide businesses with a versatile reputation-enriched ecosystem. PRepChain is specifically designed to operate in dynamic environments by also offering a trade-off between data availability and confidentiality guarantees. We make contributions in four primary directions: (i) It offers performant privacy preservation even in large-scale settings, (ii) ensures availability of computed reputation scores, (iii) seamlessly integrates with existing supply chain information systems, and (iv) in addition to subjective reputation scores, it also supports reliably-calculated, i.e., objective, ones, thereby strengthening the reliability of third-party-sourced information. Our evaluation of PRepChain documents its performance—based on a real-world use case—, security, and privacy preservation, hence, its applicability. We conclude that it is indeed destined for practical deployments in modern supply networks.

1. Introduction

Due to advances in digitalization and developments such as the Industrial Internet of Things (IIoT), traditional product-oriented supply chains are rapidly evolving into digitized and interconnected supply networks. In this context, sophisticated, data-driven supply chain information systems (SCISs) introduce amenities for businesses and consumers alike [1–3]. Relatedly, this accessibility of additional data also eases the otherwise complex establishment of new business relationships since businesses have a better foundation for their decision-making. The growing demand for individualized products and related small-batch production as well as more frequent disruptions in established supply networks amplify this trend [4–7]. These circumstances lead to dynamic business environments [8,9].

Despite these developments, the full-fledged implications for next-generation supply chain management approaches are still largely unknown to date [10]. Certainly, trust is a vital aspect of (successful) business relations [11]: Given the extensive accessibility of information, new relationships are more frequently bootstrapped using reputation (scores) [9]. Surprisingly, while supply chain information systems and collaboration have evolved extensively in the recent past [12,13], reputation systems for supply chain environments have not [14]. Existing technical solutions often lack important features or are tailored to

[*] Corresponding author.

E-mail addresses: pennekamp@comsys.rwth-aachen.de (J. Pennekamp), lennart.bader@fkie.fraunhofer.de (L. Bader), thevaraj@comsys.rwth-aachen.de (E. Thevaraj), stefanie.berninger@fir.rwth-aachen.de (S. Berninger), martin.perau@fir.rwth-aachen.de (M. Perau), tobias.schroeer@fir.rwth-aachen.de (T. Schröer), wolfgang.boos@fir.rwth-aachen.de (W. Boos), salil.kanhere@unsw.edu.au (S.S. Kanhere), wehrle@comsys.rwth-aachen.de (K. Wehrle).

[1] J. Pennekamp and L. Bader contributed equally to this work.

https://doi.org/10.1016/j.future.2025.108024
Received 20 December 2024; Received in revised form 7 July 2025; Accepted 13 July 2025
Available online 11 August 2025

other domains [15–18]. Even modern data ecosystem architectures that primarily build on organizational security [19,20] still require significant evolution to provide sufficient confidentiality and privacy for reputation systems in practice.

We thus identify a pressing need for *reliable, privacy-preserving reputation systems* that are explicitly designed for dynamic supply chain environments. Such systems can improve resilience, support data-driven compliance with emerging regulations (e.g., the proposed European Supply Chain Act [21]), and foster reliability across organizations—especially when deriving reputation from objective (IIoT) measurements and information taken from SCISs. Accordingly, any supported aggregation mechanisms and models need to provide sufficient flexibility for privacy-preserving processing in these dynamic environments.

Contributions. Our primary contributions are as follows. With PRepChain—**P**rivate and Reliable **Rep**utation for Dynamic Supply **Chain**s—we are the first to present a sophisticated, distributed, multi-agent reputation system that comes with technical guarantees to tackle the challenging requirements of dynamic supply chain environments. Its supported key features in dynamic environments are:

- support for objective and subjective ratings,
- integration with existing SCISs,
- (weighted) joining of voter and votee ratings,
- adaptive aggregation of ratings into reputation, and
- optional data availability guarantees,

while maintaining a privacy-preserving and scalable operation. To underline PRepChain's configurability, we discuss several influences that impact the utility of reputation systems well beyond our own design. Finally, our highly-configurable approach fits to various (modern) use cases and settings with appropriate performance, security, and privacy preservation while also promising a scalable operation.

Open Science. We open-source our prototypical implementation of PRepChain to foster future research [22].

Organization. The remainder of this paper is structured as follows. First, in Section 2, we introduce the core concepts for this work. Second, in Section 3, (i)we outline the scenario along with related work, and (ii)based on this information, we then derive the research gap and compile a list universally-applicable goals. After presenting our design, PRepChain, in Section 4, we extensively evaluate its performance and security in Section 5. Finally, in Section 6, we discuss PRepChain's limitations and utility before concluding the paper in Section 7.

2. Background

To set the stage for the remainder of this paper, we first outline our understanding of supply chain information systems, supply networks, and reputation systems in Section 2.1 to 2.3, respectively. This foundation is critical for understanding the characteristic requirements of reputation systems in (volatile) supply chains. Subsequently, in Section 2.4, we complement this foundation with a description of fully homomorphic encryption (FHE), a technical building block that is key for our design.

2.1. Supply chain information systems

In this paper, we refer to supply chain information systems (SCIS) to describe (parts of) business application systems that assist companies in efficiently managing their supply chains. Such systems enable the collection, processing, and analysis of data necessary for planning, monitoring, and controlling multifaceted activities within the supply chain [23]. These activities include aspects such as procurement, production, inventory management, distribution, and customer service. SCISs help companies create transparency in their supply chains, improve communication among the various stakeholders, and make informed decisions [24].

In practice, a SCIS is usually not a single proprietary system; rather, it is a compilation of relevant supply chain information that is available in specific operational application systems, such as enterprise resource planning (ERP), manufacturing execution system (MES), Supplier Relationship Management (SRM), etc. [23,24]. This wide-ranged information serves as the foundation for planning, maintaining, and developing the supply chain [23]. Traditionally, the ERP system has been the primary source of information and the system where data is being stored, aggregated, and analyzed. However, nowadays, specialized business information (BI)-tools increasingly complement it to provide cross-system supply chain information [23]. Similarly, research also proposed several information-system designs that simultaneously manage, analyze, and store information of multiple businesses [1,25] to eventually improve higher-layer applications such as supply chain management (SCM) [8].

In the context of this work, the connection between reputation system and SCIS, therefore, refers to the interface to the corresponding business application system in which data relevant for determining reputation scores is available.

2.2. Supply networks

For this paper, we follow the discussion by Braziotis et al. [26] when defining the scope of supply chains and supply networks. Supply networks are essential for any (manufacturing) business to enable value creation. At a high level, they are associated with the tasks of production and exchange of goods, as well as the management of the related information and financial flows. More specific tasks of SCM include the design of supply networks, including supplier selection, collaborative planning, tracking, tracing, and sharing of product information [8]. As companies operate in increasingly volatile environments, collaboration and supply networks are subject to constant change and evolution [27].

For example, Rolf et al. [28] identify uncertain markets, disruptions, and fluctuations in demand and supply as main drivers for dynamic network design. Although long-term partnerships remain a crucial element in fostering effective collaboration along the supply chain, more agile approaches to supply network design are needed. In this context, digitalization of supply chain management is regarded as a key enabler [27], facilitating the rapid reconfiguration of networks and thereby leading to long-term resilience.

Especially in dynamic environments, finding reputable and trustworthy business partners is a major challenge [8,9]. A commonly-chosen attempt to address this challenge is the use of reputation systems that provide insights into the past performance of potential suppliers (and customers). They are based on the data taken from SCISs, and derived metrics such as on-time delivery or quality performance of suppliers support the data-based decision-making. However, cross-enterprise reputation systems are scarce in supply networks, limiting the availability of reliable, objective reputation. Moreover, enterprise-internally, subjective information that feeds reputation scores often predominates [29].

Next, we thus discuss the diversity of reputation systems while considering their fit for supply chains.

2.3. Reputation systems

Reputation systems constitute a distinct type of information systems (potentially multi-agent), which serve and maintain reputation scores of participating entities, e.g., businesses. In addition to consumers rating businesses (e.g., on Amazon), other reputation systems target settings where businesses rate each other [30]. Gurtler and Goldberg [30] also distinguish different forms of reputation directionality (who rates whom). Given the broad range of chosen terms in literature, we briefly introduce our understanding next.

We refer to a single vote as rating r. Voters submit such ratings to influence the reputation score S (the aggregation of submitted ratings)

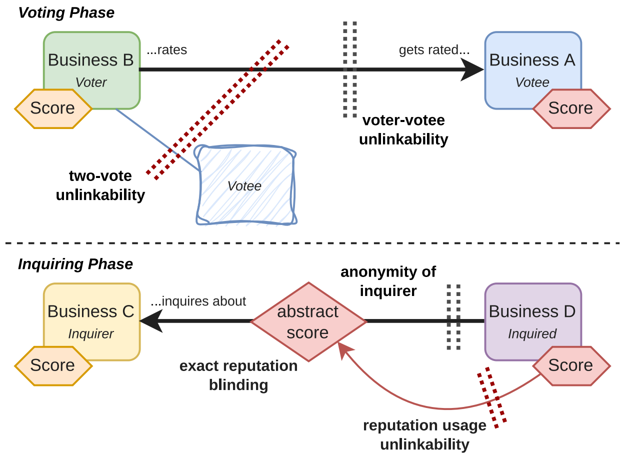

Fig. 1. In supply networks, reputation systems should guarantee anonymity and unlinkability for voters and votees, as well as inquiring and inquired businesses.

of another participant; in this context, the votee. To later retrieve the reputation score of a business, i.e., the inquired, the inquirer queries the reputation system. We visualize the different entities and their roles in Fig. 1.

In practice, different types of aggregation modes exist [31], i.e., the process of updating a reputation score once a new rating has been received. These aggregation modes range from simple calculations of the mean over purely monotonic updates to more sophisticated Bayesian models, offering varying flexibility for the expressibility of reputation scores. Depending on the processing of a rating r, models can be either voter-agnostic or voter-conscious (considering the voter's reputation score as well) [30]. Further, reputation scores are usually limited to a specific range or set of values, which might also introduce restricted visibility or limited durability [31]. Lastly, the majority of systems feature liveliness [31], i.e., scores do not reach a final state and remain updatable.

Especially in business settings, confidentiality and privacy are critical when dealing with reputation systems, especially, for practical, real-world deployments. So-called privacy-preserving reputation systems attempt to address corresponding requirements; the most important properties being [30,31]: (i) *voter–votee unlinkability*, (ii) *two-vote unlinkability*, (iii) *anonymity* of votee, voter, and inquirer, (iv) *reputation–usage unlinkability*, i.e., not strictly linking entities with their reputation score, and (v) *exact reputation blinding*, i.e., hiding exact reputation scores (only providing rough orientation). Fig. 1 concisely captures them in the context of different roles and phases of a reputation system.

Depending on the setting and threat model (more in Section 3.2), operators and participants face risks with different severity and importance [14,31]: At large, they range from, e.g., entity collusion, data manipulation, reputation tracking, over Sybil attacks, ballot stuffing, bad-mouthing, whitewashing, oscillation, and random ratings to free riding. Given their broad range and differences, each threat must be considered individually and oftentimes requires a dedicated mitigation strategy.

In this light, Gurtler and Goldberg [30] argue that addressing all desirable properties of reputation systems at once is challenging while calling for follow-up research. Hasan et al. [31] second this statement by arguing that rating confidentiality and user privacy are conflicting goals. Thus, in this paper, we explore which confidentiality and privacy guarantees are achievable in reputation systems that have been designed specifically for supply chain environments.

2.4. Preliminary: Homomorphic encryption

More foundationally, we now introduce homomorphic encryption since our proposed design, PRepChain, relies on this technical building block to ensure the confidentiality of handled ratings and reputation scores.

Fully Homomorphic Encryption (FHE) relies on homomorphic properties of special cryptosystems that allow for operations on the ciphertext to also be reflected in the "underlying" plaintext [32]: That is, FHE enables computations directly on encrypted data (without the need for decryption). Cloud computing and machine learning particularly contribute to the utilization of FHE. With FHE schemes differing in the supported operations, data types (e.g., Booleans, integers, or approximated reals), and computational overhead [32]. Which scheme to choose ultimately depends on the use case requirements at hand.

We suggest to rely on CKKS [33] in PRepChain to enable comparably accurate and performant computations on rational numbers. This way, reputation scores remain flexible and expressive, i.e., offer great utility for higher-layer applications. With FHE, competences in data handling can be distributed while preserving information privacy and confidentiality, i.e., different entities can perform operations on encrypted information without gaining additional insights, offering desirable properties for reputation systems.

3. Scenario and desired properties

Now, we first discuss related work in Section 3.1 before outlining the threat model that we consider as part of our research in Section 3.2. Based on the content presented so far, in Section 3.3, we then outline the research gap and express four crucial design properties of reputation systems for practical use in supply chains.

3.1. Related work: Supply chain reputation

Establishing novel business relationships highly depends on trust—a well-known fact to businesses and research—such that the importance of a business' reputation grows in the light of flexible and short-lived supply networks [9]. Apart from various studies on e-commerce, e.g., Hendrikx et al. [34],Shi et al. [35],Zhou et al. [36], recent work also looks into the value and pitfalls of reputation systems in business-to-business settings [9,37,38] as well as reputation-aware supplier selection [39]. Generally, we observe that several surveys [30,31] study the multitude of general-purpose reputation systems, focusing on conceptual differences and privacy preservation. We refer to Table 1 in [30,31], and [14], respectively, for details on technical foundations, privacy, and trust requirements. These works emphasize desirable properties and showcase the challenges of simultaneously fulfilling them. However, barely any research considers the specific challenges that supply chain environments introduce [14], despite a general acknowledgment of challenges such as rating fairness and objectivity as well as the required trust into a reputation platform [40].

In this light, recent, previously unsurveyed, work [41,42] also explores the benefits of monetizing participation in reputation systems using blockchain technology. While relying on blockchain technology removes the requirement for trust in the reputation platform provider [40] and further mitigates risks regarding data loss and fraud [36], its distributed and open nature amplifies challenges regarding, e.g., scalability, privacy, and information availability. Different from a blockchain-based realization, PrivBox [43] addresses these challenges for e-commerce by utilizing homomorphic cryptography to protect individual ratings in a distributed setting while still relying on a centralized marketplace to authorize and preserve (partially encrypted) ratings.

Attentively, Bader et al. [14] point out that even state-of-the-art (privacy-preserving) reputation systems do not adequately account for these unique needs modern supply networks introduce (i.e., privacy and availability, among others). We thus conclude that additional (interdisciplinary) effort is needed to bring mostly theoretical advances from computer science to practical applications in supply chain environments for future use.

3.2. Threat model in the targeted business setting

Just like related work in the area [14], in our research, we consider malicious-but-cautious attackers [44]. Under this threat model, entities can misbehave in every possible way as long as they do not leave any evidence of their misbehavior, permitting local protocol deviations. Likewise, entities may also collude as long as they do not leave any (publicly-verifiable) trace of such a collusion. We assess that this choice is reasonable since all involved businesses are driven by the strive for profits while still being bound by legislation and to specific jurisdictions. This way, we have to consider a broad range of realistic threats while ruling out threats that are unrealistic in business settings like apparent collusion.

3.3. Research gap and corresponding goals

While general-purpose reputation systems introduce several setting-independent (confidentiality) requirements, our considered scenario of offering privacy preservation in dynamic supply chain environments adds complementary needs and constraints. We first surveyed related work and subsequently discussed the situation and our findings with supply chain experts. This approach allowed us to explicitly tailor our list of properties to dynamic supply chain environments. For referenceability, we grouped them into four key properties, as we detail in the following.

P1 : Genuine Operation. Any reasonable design must (reliably) provide the key functionalities of a reputation system (cf. Section 2.3). Only authentic scores and ratings should be processed, handled, and returned by the reputation system to make its operation genuine. On a technical level, this need boils down to a form of authorization (including access control) as well as integrity-protected information flows between the involved entities.

P2 : Secure Operation. This second property slightly conflicts **P1** because it mandates that submitted ratings and scores, as well as the relationships of participants, are not publicly accessible, seemingly impairing a genuine operation. Thus, appropriate measures are needed in practical designs. More in detail:

P2a: Privacy. Especially business environments mandate that the participants' behavior and their relationships remain private to protect business secrets and allow them to maintain their competitive advantage. Consequently, any system must provide unlinkability and anonymity features while simultaneously ensuring privacy of relationship as well as unlinkability of ratings (cf. Fig. 1).

P2b: Confidentiality. Extending the previous thoughts, ensuring confidentiality of ratings and scores is also key for any deployment in business settings since they capture highly-sensitive and private information.

P3 : Accurate Operation. Closely related to **P1**, and in conflict with **P2** (building blocks for confidentiality and privacy-preservation might negatively impact the precision of computed results), is the need for accurate operation. That is, ratings and aggregated scores should be as accurate as needed for the respective use case. Otherwise, they lose their value for higher-layer decision-making, e.g., for SCM.

P4: Applicability. This property is mostly unique to the presented motivation (cf. Section 1): First, any design must be compatible with dynamic supply chain environments and their frequent changes of business relationships, which also includes the risk of unresponsive entities (e.g., non-complying or defunct businesses). In this regard, reputation scores should remain available to ensure a reliable and dependable operation at all times. Second, again considering the setting, any proposed system has to scale to the size of modern supply networks—an aspect with increasing importance, given the expected dynamism in such environments. Third, the setting stresses the need for strictly authentic ratings (reliability), emphasizing **P1**. Finally, to ensure compliance with different deployments and their individual needs, the handled scores and the supported aggregation functions must offer sufficient granularity. Jointly, these aspects contribute to the property of applicability, which has not been considered by related work in full (cf. Section 3.1).

Real-time Requirements. The first two aspects of **P4** further raise the question about real-time requirements concerning the processing of ratings and the inquiry of reputation scores. After discussing with domain experts, we concluded our considered scenario does not come with such tight requirements: Given that reputation scores consists of a large number of ratings in business settings, a single rating is unlikely to extensively impact the aggregated reputation score. This aspect aligns with related work, targeting other domains, since they frequently implement threshold-based rating aggregations, which introduces a certain delay in processed ratings by design. Hence, delays in the voting process are acceptable. Moreover, in business settings, operational decisions regarding sourcing as well as other long-term decisions are rarely made at the last minute. Businesses carefully consider their options, i.e., they have a certain leeway until they have to reach a decision. Thus, the inquiring process does not mandate tight processing requirements either.

Implicitly, and after discussing with domain experts, we are further convinced that meaningful approaches should satisfy these properties largely by providing technical guarantees since "organizational security" [20] is unfit for use in the defined threat model (cf. Section 3.2) as well as for any setting with mutually distrusting (business) entities.

Research Gap. Based on our literature review on existing reputation systems and the supply chain-specific needs for technical enforcement of system properties, we find a lack of *reliable* and *privacy-preserving* reputation systems that are applicable to the flexible environments of supply chains while further providing technical guarantees at real-world scales. Existing research for other domains offers various approaches for achieving the desired properties. However, the unique *combination* of requirements and desired properties reveals the research gap for achieving reliable business reputations in dynamic supply chain environments.

Takeaway. *The manifested lack of (convincing) reputation systems for supply chains (cf. Section 3.1) paired with the setting-specific requirements (**P4**) reveal a research gap despite the importance of having the corresponding functionality available (cf. Section 1).*

4. PRepChain: Private and reliable reputation for dynamic supply chains

To close the outlined research gap, we designed PRepChain, our FHE-based approach to provide a versatile privacy-preserving reputation system, which is specifically tailored to use in dynamic supply chain environments (**P4**), irrespective of the exact requirements. This novel multi-agent information system further accounts for sophisticated privacy and confidentiality requirements (**P2**) without neglecting well-known desires like a reliable operation (**P1**).

To introduce PRepChain's key functionality step by step, we first present a design overview in Section 4.1. In Section 4.2, we then outline the involved entities and their responsibilities before discussing the system's voting and inquiring processes in Section 4.3. Finally, we summarize how PRepChain advances the state of the art in Section 4.4.

4.1. Design overview

PRepChain builds on several core principles to achieve the goal of proposing a supply chain-oriented reputation system. First, it utilizes fully homomorphic encryption and a distribution of competences across different entities (Fig. 2), separating ratings and scores from the key material, to satisfy the need for confidentiality (**P2a**). Likewise, to offer the desired level of privacy (**P2b**), namely voter–vote unlinkability and two-vote unlinkability, PRepChain relies on the use of pseudonyms and (distributed) *aggregation engines* to process submitted ratings. Pseudonyms prevent the tracking of rating and inquiring

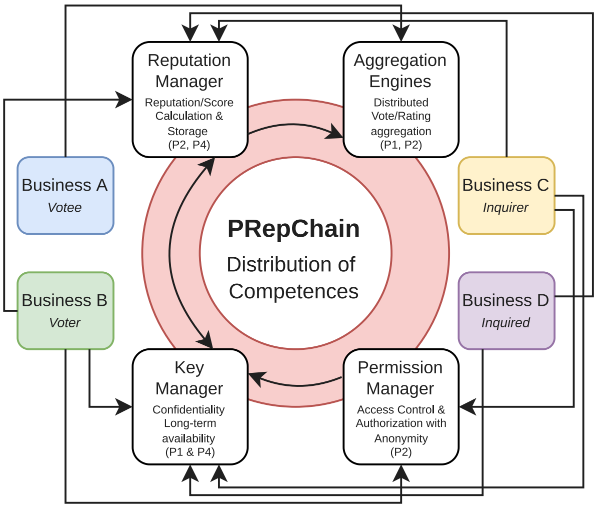
Reputation Manager
Reputation/Score Calculation & Storage
(P2, P4)
Aggregation Engines
Distributed Vote/Rating aggregation
(P1, P2)
Business A
Votee
Business B
Voter
Business C
Inquirer
Business D
Inquired
PRepChain
Distribution of Competences
Key Manager
Confidentiality Long-term availability
(P1 & P4)
Permission Manager
Access Control & Authorization with Anonymity
(P2)

Fig. 2. PRepChain follows the concept of distribution of competences to assure reliability, security, and accountability of information. Thus, businesses have to interact with multiple independent entities during the voting and inquiring processes.

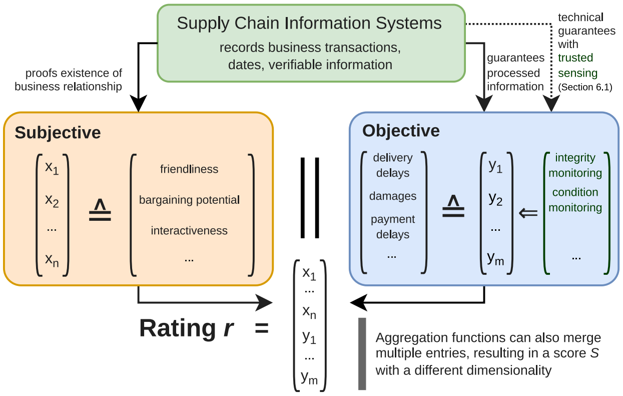
Supply Chain Information Systems
records business transactions, dates, verifiable information
proofs existence of business relationship
guarantees processed information
technical guarantees with trusted sensing (Section 6.1)
Subjective
x_1
x_2
...
x_n
≙
friendliness
bargaining potential
interactiveness
...
Objective
delivery delays
damages
payment delays
...
≙
y_1
y_2
...
y_m
⇐
integrity monitoring
condition monitoring
...
Rating r =
x_1
...
x_n
y_1
...
y_m
Aggregation functions can also merge multiple entries, resulting in a score S with a different dimensionality

Fig. 3. The integration with existing supply chain information systems enables PRepChain to efficiently consider subjective and objective information for ratings. Objective rating aspects can be based on contract data, payment proofs as well as information from trusted sensing [45,46].

patterns. They are generated by the *pseudonym manager*, which simultaneously implements a form of access control. They represent an anonymous-yet-entity-bound variant of *tickets* [30] or *tokens* [31]. Specifically, PRepChain further sources information from existing SCISs that capture business interactions to only permit (authorize) ratings for actual business transactions (**P1**). We further stipulate a threshold-based aggregation at the *reputation manager* to prevent reputation tracking, i.e., the score S^E is updated once k-many new ratings have arrived.

Altogether, PRepChain features a hybrid architecture that utilizes conceptually centralized reputation and key managers to safeguard (encrypted) reputation scores while ensuring their availability (**P4**). Moreover, a centralized reputation manager ensures consistency of the reputation scores, i.e., we achieve global visibility, which is favorable in business settings.

Finally, our voter-conscious design offers significant flexibility in terms of the applied aggregation model: (i) how to "join" voter and votee ratings at the aggregation engine and (ii) how to update (retrievable) reputation scores at the reputation manager once the threshold of k ratings has been reached. The former also provides us with the ability to weigh ratings depending on the reputation of voter and votee. Due to this flexibility, PRepChain can be configured to operate as simplex, half- or full-duplex [30] reputation system. Overall, its aggregation model is only constrained by the underlying FHE cryptosystem; i.e., in theory, every operation is computable [32]. Further, scores are commonly non-monotonic if not altered by a specifically tailored model. Additionally, since PRepChain implements a vector-based approach (range of $\mathbb{Q}$), we offer support for objective and subjective ratings as well as liveliness per feature (Fig. 3).

In Section 5, we discuss our evaluation results of PRepChain's performance and security to stress its conformance with **P2** and **P4**. Likewise, given its reliance on FHE, we will further carefully examine the precision of calculated computations to ensure **P3**. For now, we focus on PRepChain's design in more detail.

4.2. Involved entities and their roles

Moving on, we introduce the different entities in our design as well as their responsibilities. We sort them by their order of appearance when a voter submits a rating.

Participant(s). Depending on the current task at hand, businesses take different roles when interacting with PRepChain: Regarding the voting process, a participant is either a voter or votee. The former not only submits the encrypted rating r_r^E but also forwards the encrypted rating r^E that has been processed by the aggregation engine to the reputation manager. By submitting r_e^E, the votee may also contribute to this processed rating. Then, as part of the inquiring process, we distinguish between inquirer and inquired. During operation, these entities interact with the other entities as needed (cf. Section 4.3).

Participants may leverage different approaches, including objective and subjective ones, to collect or sense data that is used to come up with ratings. Given that this data is usually pre-processed in one way or another, i.e., the voting process is decoupled from the data collection, constrained devices do not have to interact directly with PRepChain. Instead, the voting process (just like the inquiring process) is expected to be processed through a server under the control of the respective participant. Thus, we consider the support of lightweight sensors or edge devices to be experimental matters that exceed the requirements of typical deployments. Regardless, advances in the area of trusted sensing [46] show that a lightweight processing of (objective) data is possible, with the concept being compatible to PRepChain.

Pseudonym Manager. This entity is crucial for supporting access control and authorization in PRepChain, and thus, it is likely supposed to be linked to existing SCISs. This way, it can assess whether a pseudonym should be handed out for a specific business relationship or transaction (authorization). In real-world deployments, it may further confirm the authenticity of businesses (i.e., participants) by interacting with responsible government entities, e.g., using data from company registers. We consider its exact realization out of the scope of this paper because it does not have implications for the remainder of the design.

Key Manager. As outlined in Section 4.1, we rely on a distribution of competences. That is, PRepChain separates the key material from encrypted data. Hence, the key manager is essential to ensure confidentiality. Then again, handling parts of the key material (incl. certain decryption keys) provides the desired level of availability, which is needed in light of rapidly developing and changing supply chain environments. PRepChain's availability guarantees can be improved by weakening the security property (cf. Section 6.2). In this case, the key manager also possesses the FHE decryption keys. Irrespective of this configuration, during the inquiring process, the key manager is further crucial to prevent the reputation manager from decrypting the relayed scores (cf. Section 4.3.2). In our design, it is conceptually centralized, i.e., in practice, its responsibilities for specific entities may also be partitioned across multiple independent key managers.

Aggregation Engine(s). This entity joins the encrypted voter's and votee's ratings (r_r^E and r_e^E) into a single encrypted rating r^E and thus ensures unlinkability of ratings, i.e., it shields their relationship from the other entities. PRepChain's security guarantees build on the existence of multiple independent aggregation engines to prevent a single aggregation engine from observing all relationships as well as having access to all ratings. Voters can arbitrarily choose different aggregation engines. Since all ratings are encrypted under FHE, the aggregation engines can aggregate the provided values while being oblivious of their values and the result, ensuring the required confidentiality.

Reputation Manager. The conceptually-centralized reputation manager (which also allows for partitioned responsibilities) ensures availability of the participants' reputation scores. Due to the lack of decryption keys, it is oblivious of the stored reputation scores, i.e., it cannot track any scores (observing the frequency and patterns of newly arriving ratings is possible, though). During operation, the reputation manager performs the threshold-triggered processing of k ratings t_i^E into the "long-term" score S^E. Moreover, as part of the inquiring process, it further serves as a relaying proxy to separate inquirer and inquired.

Jointly, these entities shape the information flows of PRepChain, as we outline in the next subsection.

4.3. Core functionality: Sequential steps

In reputation systems, we distinguish between two main processes. First, in Section 4.3.1, we detail how a voter submits a rating that is eventually persisted in the votee's reputation score. Afterward, in Section 4.3.2, we outline the order of steps when an inquirer retrieves the reputation score of an inquired when using PRepChain.

4.3.1. Voting process

The voting process in PRepChain is triggered by a voter who is interested in submitting her rating of another business, i.e., a votee. By design, each rating and reputation score is a finite vector, containing both objective and subjective entries (for both, the number of entries is chosen from $\mathbb{N}_0$). Hence, depending on the use case, configuring only objective or only subjective entries is supported as well. However, the latter choice would sacrifice the benefit of PRepChain supporting objective ratings and reputation scores.

① Voter acquires a pseudonym PDM (i.e., access token) from the pseudonym manager

② Voter retrieves votee-specific FHE public key k_e^E from the key manager (the votee generates a new key pair if it is the first rating for him and submits the public key to the key manager)

③ Voter encrypts her rating r_r to r_r^E using k_e^E

④ Voter retrieves the votee's current reputation score S_e^E from the responsible reputation manager (if a reputation score is already available)

⑤ Voter forwards $\langle \mathrm{PDM}, r_r^E, S_e^E \rangle$ to an aggregation engine of her choice for further processing
After these steps, the selected aggregation engine continues the process by handling the submitted rating.

⑥ Aggregation engine retrieves the voter's score S_r^E from the reputation manager using the pseudonym

⑦ *Optional Step:* Votee may share an (objective) self-rating r_e^E with the aggregation engine

⑧ Aggregation engine joins both ratings r_r^E and r_e^E into a single rating r^E; the exact details depend on the aggregation model; the reputation scores S_r^E and S_e^E may also serve as weights in this step

⑨ Aggregation engine signs the result (r^E) and sends it to the voter

⑩ Voter has received proof of the processed rating and forwards it to the reputation manager

⑪ Reputation manager treats r^E as t_i^E and aggregates it into S_e^E once the threshold of k new ratings has been reached

This processing by the reputation manager concludes the voting process of a rating in PRepChain.

4.3.2. Inquiring process

This process is triggered once an inquirer is interested in obtaining the reputation score of a business, i.e., the inquired business. Again, all scores are vectors and thus feature multiple entries.

① Inquirer acquires a pseudonym PDM from the pseudonym manager

② Inquirer requests the inquired's reputation score S_d^E from the reputation manager using PDM for anonymous authentication

③ Reputation manager requests the encrypted FHE secret key $E_h(k_d^E)$ from inquired
To appropriately protect the secret key during transit, PRepChain encrypts k_d^E by utilizing a fresh temporary key pair h that is served by the key manager.

④ Inquired retrieves a fresh encryption key h from the key manager

⑤ Inquired encrypts her secret key k_d^E using the encryption key h to obtain $E_h(k_d^E)$

⑥ Inquired sends $E_h(k_d^E)$ to the reputation manager

⑦ Reputation manager shares this encrypted key $E_h(k_d^E)$ and reputation score S_d^E with the inquirer
This relaying approach ensures the separation of inquirer and inquired (privacy of relationship).

⑧ Inquirer requests the corresponding decryption key for h from the key manager

⑨ Inquirer decrypts the inquired's FHE secret key using the requested key to obtain k_d^E

⑩ Inquirer decrypts the inquired's reputation score S_d^E using k_d^E

After completing this process sequence, the inquirer has access to the inquired's current reputation score.

4.4. Advancing the state of the art

When comparing our design, PRepChain, with the state of the art, we identify several aspects that advance the field.

First, while privacy preservation is a well-known aspect in literature [30,31] (Section 3.1), we are the first to propose a privacy-preserving reputation system that is tailored to dynamic supply chain environments and their specific properties (Section 3.3). Second, PRepChain exercises great care in ensuring the availability of reputation scores—an aspect related work did not consider isolatedly. Prior work rather distinguished between centralized and decentralized approaches without considering the availability of reputation scores. Overall, our design does not store them locally with the participants; instead, the conceptually-centralized reputation manager records them. In Section 6.2, we later discuss how the mode of deployment (PRepChain supports two) influences the availability of reputation scores in scenarios where businesses might also be "leaving" the ecosystem.

Third, due to our focus on supply chain environments, we further integrate existing SCISs into the regular operation of our design. This way, PRepChain can extract reliable information about business relationships or transactions, allowing the pseudonym manager to check whether authorization for voting should be granted or not. Lastly, we account for differences in the reliability of processed ratings by distinguishing objective (e.g., by using trusted sensing [46]) and subjective information (sources). To the best of our knowledge, these last two aspects have not yet been considered by related work at all, i.e., given these specific features, this paper is advancing the field beyond the proposed design.

The design choices we made to accommodate joining voter and votee ratings and weighing ratings during aggregation into reputation scores required us to utilize FHE, as we have hinted at in Section 2.4. Less powerful homomorphic schemes like partially homomorphic encryption (PHE) [32] do not support a sufficient number of operations. As a result, a PHE-based PRepChain (cf. Section 6.2) could not support the aforementioned expressive weighting of ratings at different stages of the voting process. Similarly, other design variants would introduce other disadvantages. Secure multiparty computation-based designs, including those protocols that build on threshold cryptography, would introduce significant drawbacks in terms of the data availability concerning both ratings and reputation scores alike paired with increased protocol complexity, potentially even reducing the observed performance and scalability. However, given the focus on dynamic supply chain environments, data availability is paramount. Thus, for the moment, we see benefits in pursuing a comparably easy-to-comprehend design like PRepChain. Nonetheless, future work may explore other avenues for satisfying the outlined properties (cf. Section 3.3); our research insights may come in handy in this endeavor.

Takeaway. *By distributing competences among entities, using pseudonyms, and relying on FHE, PRepChain considers the privacy and confidentiality needs of businesses (**P2**). Its ticket-based approach, as well as the support for objective and subjective ratings, contribute to a genuine operation (**P1**). PRepChain is specifically suited for application in dynamic supply chain environments (**P4**) since relevant data is available at conceptually-centralized entities.*

5. Evaluating PRepChain

Complementing the design, we now focus on the practical feasibility of PRepChain, primarily from a technical point of view. To this end, we give details on our proof-of-concept realization in Section 5.2 and detail our experimental setup in Section 5.2. Subsequently, in Section 5.3, we show that PRepChain is suitable for practical use in supply chain environments (**P3** and **P4**). Moreover, in Section 5.4, we discuss its confidentiality and privacy guarantees (**P2**). Hence, we first focus on the technical (security) guarantees we can achieve and which overhead PRepChain introduces. Finally, in Section 5.5, we complement this evaluation by assessing the (improved) situation in a real-world setting.

5.1. Implementation

To evaluate PRepChain in detail, we implemented a prototype in Python, which is publicly available [22]. This implementation is a means to assess our design's feasibility and is not intended to serve as the foundation for a production-ready deployment, i.e., we refrained from optimizing the architecture and performance of our prototypical realization.

Implementation. We rely on MongoDB [47] as data storage and use the FHE cryptosystem CKKS [33], specifically through the Python library Pyfhel [48], with the default parameters, which offers an API to SEAL [49]. For improved convenience, we encode transmitted data and ciphertexts in JSON. In the querying process, our implementation uses an efficient hybrid encryption construction based on Fernet [50] (symmetric) and RSA [51] (asymmetric).

Our prototype implements three main components. The first component automates the rating and query processes, handling input for voters, votees, and ratings. The second component simulates user interaction through an interface, processing user actions via RESTful API calls. The third component implements the remaining entities using Flask [52] and RESTful APIs. In our prototype, the votee runs on a Flask server while the voter interacts with it via API calls. A dedicated script initializes the different components and prepares the required databases, ensuring that new operations start from scratch. Our implementation is currently compatible with Linux (Ubuntu 22.04) and Windows (10), promising deployability, flexibility, and scalability.

Deployment. We realized PRepChain with good usability in mind. This intention includes providing a simple setup. Any deployment of PRepChain only involves configuring the database, setting up the Flask server, and securely realizing the key management to prepare for the intended encryption. However, in real-world deployment, integration with existing SCISs is also required. This aspect results in company–individual configurations and adaptations, adding minor complexity. In the long run, we expect that standardization and interoperability approaches will improve the situation. As we detail in Section 5.3, the overhead introduced by PRepChain is moderate. Moreover, our implementation does not come with specific hardware requirements.

When compared to an insecure design, FHE introduces significant computational overhead, leading to prolonged execution times and higher CPU load [53]. Comparisons of state-of-the-art FHE libraries [54–56] indicate minor potential for optimizations since they highlight the benefits of selecting the most suitable FHE implementation. We leave corresponding considerations for future work and now focus on assessing PRepChain's general feasibility instead.

5.2. Experimental setup

For our assessment of PRepChain, we also evaluate its performance on a dedicated machine against a real-world use case from industry, looking into its potential for deployment.

Experimental Setup. We evaluate PRepChain's performance on a single server (`Intel Xeon E5-2620v2` with 16 GB of memory). Specifically, all entities run on said server. In practice, their operation would be distributed across different organizations. We repeat each measurement 20 times and further provide 99% confidence intervals in the following.

Real-World Use Case. We also conduct an evaluation that utilizes a real-world use case. Specifically, our real-world-oriented evaluation sources data from a manufacturing company with different granularity. It comprises data from a SCIS as well as sensor data from the shopfloor.

The company is a contract manufacturer with a frequently changing production program. Consequently, the supply network is very dynamic, with frequently changing customer and supplier relationships. The data under (technical) evaluation comprises ERP data that is created during business transactions such as incoming and outgoing goods bookings, purchase orders, and invoices. This broad data allows reputation scores to be generated for up- and downstream material, financial, and information flows. From the supply network perspective, the company can represent both the customer and supplier side, i.e., to imitate a real-world deployment. Furthermore, raw data from the shopfloor originates from sensors of a band saw, allowing an evaluation of the actual production process. From a business perspective, the retrieved reputation scores (for, e.g., on-time delivery or quality performance) can eventually be applied to support data-driven decisions for supplier selection and performance measurement in dynamic supply chain settings.

Given our focus on the technical feasibility in this paper, we now discuss this dimension in more detail. The corresponding performance measurements later serve as the foundation for discussing the business dimension, particularly its value for SCM and decision-making. Altogether this real-world use case provides us with 8 unique (reputation) features, i.e., $n + m = 8$ (cf. Fig. 3), for our evaluation.

5.3. Performance and accuracy evaluation

Moving on, we first measure the computational performance of both voting and inquiring processes (Sections 5.3.1 and 5.3.2) to give insight into PRepChain feasibility. Moreover, we present the network demand and storage requirements for each entity (cf. Table 1) to complement the runtime measurements. Afterward, in Section 5.3.3, we discuss PRepChain's scalability and encryption-induced precision.

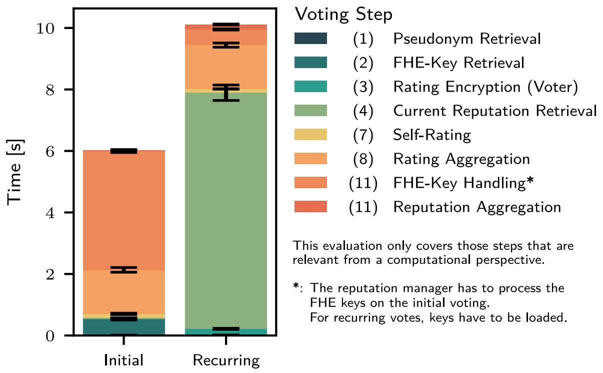

Fig. 4. In the simpler case (initial rating), the aggregation at the reputation manager accounts for most of the processing. In the recurring case, with existing reputation scores, handling these scores adds complexity. Still, the processing concludes quickly.

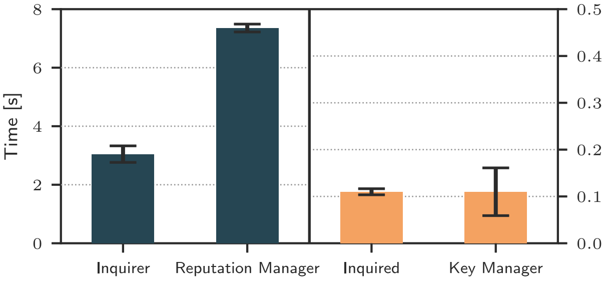

Fig. 5. Computationally, the inquiring process mostly burdens the inquirer and reputation manager.

5.3.1. *Voting process*

First, we evaluate all relevant steps of the voting process that might involve computationally expensive operations. In Fig. 4, we detail the best- and worst-case situations: For the best case (*initial*), neither voter nor votee have any previous reputation, reducing the number of processing steps. For the worst case (*recurring*), their previous reputations is being processed as part of the voting. The complexity of the initial case is dominated by the initial processing of FHE keys as part of Step ⑪. For subsequent ratings (recurring), retrieving the votee's reputation score (Step ④) adds significant processing. Since the duration of the voting process is below 11 s, we conclude that PRepChain performs adequately for real-world use, where individual transactions usually exceed minutes. Thus, in real-world deployments, ratings are being submitted less frequently.

5.3.2. *Inquiring process*

In contrast to the voting process, the complexity of the inquiring process is dominated by the operations performed by the inquirer and the reputation manager (Fig. 5), barely adding load to the inquired and key manager. Specifically, the aggregation of reputation scores (Step ②) and relaying key material (Step ⑦) burden the reputation manager. Caching aggregated results for subsequent queries could even reduce the load on the reputation manager.

Based on our evaluation of a real-world use case from industry, we conclude that PRepChain's performance is well-suited for practical deployments, irrespective of the exact content that is captured within the reputation scores.

5.3.3. *General consideration*

Following the evaluation of PRepChain's computational performance, we now investigate and discuss its scalability and induced overheads, i.e., networking and storage requirements, as well as the FHE-achievable precision.

Table 1
PRepChain's Protocol-induced Overheads.

Entity	Aspect	Voting Phase [MB]		Inquiring Phase [MB]	
Voter/Inquirer	Network	↑ 3.99	↓ 41.41	↑ 0.00	↓ 3.83
	Storage	No storage used		No storage used	
Votee/Inquired	Network	↑ 1.02	↓ 1.18	↑ 0.31	↓ 0.31
	Storage	No storage used		No storage used	
Permission Man.	Network	↑ 0.00	↓ 0.00	↑ 0.00	↓ 0.00
	Storage	No storage used		No storage used	
Key Manager	Network	↑ 35.81	↓ 0.29	↑ 0.311	↓ 0.29
	Storage	37.33		No storage used	
Aggregation Eng.	Network	↑ 1.27	↓ 5.16	Not involved	
	Storage	No storage used		— " —	
Reputation Man.	Network	↑ 3.83	↓ 1.35	↑ 4.14	↓ 0.09
	Storage	1.28		No storage used	

Scalability and Overheads. Hasan et al. [31] previously identified network and storage as relevant performance metrics. Accordingly, we summarize corresponding measurements for both processes in Table 1. By design, the reputation manager is the only entity with (permanent) storage needs. Specifically, evaluation keys concerning the FHE encryption are responsible for the majority, whereas the rating score itself is only a fraction of the storage overhead. Our evaluation purposely disregards all networking constraints. To still outline the expected burden on the network, Table 1 gives specific numbers. In modern networks, these numbers are not an issue, even when setting up a deployment with significantly larger ciphertext sizes (the FHE ciphertext size depends on the required precision and the configured security level; cf. the upcoming paragraph on precision).

PRepChain ensures great scalability by design, given that all influencing factors scale at most linearly. That is, the number of voters and votees does not have an influence on the processing of individual ratings or reputation scores. Naturally, an increasing number of ratings and captured reputation scores also increase the processing time and storage requirements. Overall, any observed overhead in PRepChain is independent of the number of participating businesses, submitted ratings, and reputation scores processed, promising an appropriate scalability for real-world use. This conclusion thus also holds for large supply networks with many businesses and business transactions.

Precision. FHE enables PRepChain to obliviously preserve the content of individual ratings, at the expense of (storage) overhead (cf. Table 1), i.e., FHE greatly contributes to the desired confidentiality guarantees (**P2b**). Now, to assess **P3** with our configuration and real-world evaluation data (values from $\mathbb{Z}$), we chained FHE operations on corresponding CKKS ciphertexts. Even after more than 8000 operations, the results were as accurate as with plaintexts, despite the chosen scheme which only approximates numbers. While FHE computations on values from $\mathbb{Q}$ generally affect the precision, the FHE cryptosystem and its configuration influence it as well, allowing for trading off performance (overhead) and precision as needed when deploying PRepChain. As such, it is well-suited across settings, independent of the number of (chained) computations on a ciphertext.

5.4. *Security discussion*

In reputation systems for supply chains, illegitimate insertions, manipulations or deletions of reputation scores and ratings for one or multiple businesses are desirable attack vectors for malicious participants. To mitigate these threats, the security (**P2**) of PRepChain relies on the distribution of competences, redundancy of important entities, cryptography for privacy and access control, and the information security provided by underlying SCISs. Further, businesses cannot participate anonymously such that detected malicious actions entail juristic consequences and PRepChain respects the malicious-but-cautious threat model on a technical level.

The concept of distributed competences requires entities to collude to achieve illegitimate access to PRepChain. Since such collusion is always visible to multiple entities, a malicious-but-cautious attacker (cf. Section 3.2) would abstain from such an attempt, especially as participants are registered with a unique business identifier, such that malicious actions could entail juristic consequences. The unique identifiers for each business registered at the pseudonym manager also prevent businesses from *whitewashing* their reputation by re-entering the system or from performing sophisticated Sybil attacks. The usage of pseudonyms for anonymous authorization further prevents entities from linking ratings to votees or multiple ratings to each other, ensuring the desirable *unlinkability* property. Since the pseudonym manager has no access to the submitted ratings and the reputation manager has no access to the votee's identity, breaking these unlinkability properties requires collusion of at least these two entities. For environments with strong unlinkability requirements, multiple layers of pseudonyms can be used to flexibly tune the required number of colluding entities. Due to the redundancy of critical entities, data loss—either by accidental or malicious deletion—requires collusion and is thus prevented by technical means, which extends to attacks that aim to manipulate existing information. Using multiple replicas for data exacerbates such attempts at the cost of increased storage requirements and (insignificant) communication overhead.

The pseudonym manager—without further measures—has the capability to impersonate other entities within the system as it provides both anonymization and access control to all participants. Hence, it can issue a pseudonym to itself for attempting a Sybil attack. While using this "ticket" is detected by the reputation manager, it is unable to detect whether this usage is legitimate or not. Multiple instances of the pseudonym manager can counteract this threat: Whenever a pseudonym is used for authentication, the reputation manager has to verify its legitimacy against *all* pseudonym managers. Using a majority vote for this verification further circumvents the risk of a single pseudonym manager willingly blocking authorization requests and provides redundancy. Since the reputation manager has no decryption capabilities, illegitimately revealing stored information always requires a conspirator as well. Hence, a malicious-but-cautious attacker would abandon such an attempt.

While encryption, distributed competences, redundancy, and real-world identification of businesses protect information that already entered the system (countering, e.g., reputation tracking) and is exchanged between entities, multiple attack tactics aim to submit arbitrary information, i.e., ratings, from the beginning, to manipulate reputation scores to match their desires. Here, randomized or illegitimate ratings for different businesses (bad-mouthing, ballot stuffing, random ratings), as well as illegitimate ratings for their own reputation (Sybil attacks, oscillation), are potential attack vectors. The combination of access control (voting authorization), business identifiers (preventing, e.g., Sybil attacks), and the integration with SCISs allow PRepChain to limit the possibility of undetected reputation manipulation. Further, including concepts such as end-to-end-secured sensing [46] for calculating objective ratings further complicates such manipulation attempts. To further enhance the trustworthiness of submitted ratings, a *verification engine* can be integrated into PRepChain to oversee the reputation calculation process. This entity then moderates between the reputation manager, the voter, and the aggregation engines to verify signatures of provided information and signs the results of computations to attest their correctness, improving the reliability of ratings at the cost of increased communication and performance overhead. Orthogonal to ensuring the correctness of ratings, preventing excessive ratings and potentially revoking incorrect or malicious ratings can increase the resilience of PRepChain against such attacks. While rate-limiting the votee's or voter's rating capabilities is an easy-to-adoptable feature, supporting revocable or degradable ratings introduces unique challenges (cf. Section 6.1).

5.5. Assessing PRepChain in a real-world setting

Deploying PRepChain in real-world supply networks yields significant practical benefits, particularly in dynamic and complex environments. The following conclusions build upon (informal) discussions with domain experts who are familiar with reputation in supply networks. At a strategic level, PRepChain enables objective, data-driven assessments of the current supply chain performance by leveraging (reliable) data from SCIS. This approach enhances transparency regarding the supply network's performance and reduces dependence on individual employees' subjective, often inaccessible knowledge. While objective data serves as the foundation, PRepChain further integrates subjective stakeholder ratings. Optionally, their relevance and timeliness can be fed into the verification engine for verification purposes. This way, PRepChain accounts for the differing needs of supply networks and is thus applicable in various situations.

Introduction PRepChain into a business operation allows for continuously identifying weak points in the supply network, such as unreliable suppliers, on an operational level. Afterward, practical measures to increase supply chain performance can be agreed upon with the affected stakeholders, or they can (collaboratively) undertake an agile reconfiguration of the supply network. Consequently, PRepChain is especially valuable in supplier and supply chain management operations. In highly-dynamic environments, such as the contract manufacturer of our real-world use case, it supports the rapid formation of completely new, trusted partner networks. Due to the availability of a global reputation system that collects, processes, and maintains reputation scores of multiple businesses, PRepChain facilitates supplier pre-selection and negotiation processes for material suppliers as well as logistics partners based on (optionally-verified) ratings. In practice, this feature promises to reduce manual, subjective research efforts, supplement internal historical data with reliable, external evaluations, and minimize time spent on inter-company information exchanges.

Finally, the integration of trusted sensing capabilities (i.e., objective data) enables the reliable and timely detection of disruptions and misbehavior within the supply network, e.g., the interruption of a cold chain, at the stakeholder level. This aspect is particularly of interest since it addresses the lack of trust in dynamic supply chain environments with mutually distrusting entities. The availability of such information not only accelerates root cause identification, enhancing quality assurance and liability management but also contributes to the long-term elimination of systemic issues, thereby improving overall supply chain performance and collaboration—a desirable outcome for any operation.

PRepChain prospectively introduces diverse and measurable benefits when being deployed in real-world settings.

Takeaway. *Despite the utilization of cryptography with the distribution of competences to provide both privacy (**P2a**) and confidentiality (**P2b**), PRepChain excels with good performance, scalability, and precision (**P3** and **P4**). The combination of reliable and accountable information stemming from sourced SCISs as well as redundant entities ensure a genuine operation (**P1**). Future extensions, such as the integration of verification engine(s), promise to further harden and improve PRepChain, even beyond our considered malicious-but-cautious threat model.*

6. PRepChain's utility for supply chains

After assessing the performance and security of PRepChain, we now move slightly to the business perspective. First, in Section 6.1, we outline (technical) limitations of our design that may impact deployments. Subsequently, we discuss general real-world implications (Section 6.2).

Table 2
Classification of PRepChain according to a prior comparison framework by Hasan et al. [31, Table 1].

Architecture	Set/Range	Granularity	Set/Range	Liveliness	Visibility	Durability	Non-Mono.	Aggr. Model
H	$\mathbb{Q}^n$	S	$\mathbb{Q}^n$	●	G	●	●	Open

6.1. Limitations

PRepChain provides a sophisticated design for dynamic supply chain environments with a particular focus on information privacy and confidentiality without the need for a trusted third party. This dedicated focus introduces different limitations for its applicability and deployment. As part of this paper, we focused on the design of a versatile privacy-preserving protocol to demonstrate its capabilities and hint at its added value (cf. Section 6.2). As a result, we consider PRepChain's embedding in business workflows as well as realizing production-grade deployments as out of scope.

Considering the deployment requirements in realistic scenarios, the number of required independent entities can be a limiting factor, i.e., PRepChain primarily targets larger deployments with at least 100 participants. This requirement is in line with the real-time needs expressed in **P4**.

In the design context, we identify three primary limitations: First, PRepChain does not incorporate *verifiable* homomorphic encryption, limiting the degree to which it can guarantee usable, correct, and sound computations without external verification. In Section 7.2, we detail respective promising mitigation strategies and future research directions. Second, the computational overhead of the rating process grows linearly with the number of rating entries (i.e., votable aspects). While we assess this (unavoidable) overhead to be reasonable, the resulting computational and operating must be considered when selecting the desired number of vector entries. Finally, revoking a submitted rating is a challenging and potentially error-prone procedure. The reputation manager can delete submitted ratings to enhance privacy after they have been used for the aggregation. Hence, removing a rating requires the provision of a *valid* copy of the original rating to subtract it from the overall score. A carefully-adapted design could potentially get rid of this constraint, likely at the expense of added complexity, entailing degraded performance, and/or storage overhead.

Additionally, from a deployment perspective, realizing compatibility between and integrating PRepChain and the multitude of different SCISs (cf. Section 5.1) is another step future work has to tackle (cf. Section 7.2). Beyond standardizing data inputs and data-quality criteria, we expect this step to mostly constitute an implementational challenge rather than a research-specific one. Accordingly, in this paper, we focused on conceptually outlining the benefits of their interaction to showcase the benefits of our design.

Next, we go into more detail on PRepChain's properties to highlight their respective impact and added value.

6.2. Discussion: Impact and added value

The technical guarantees and operational features for supply networks provided by PRepChain cover several aspects of added value, as we discuss hereinafter.

Applicability. The central property of a reputation system is its applicability (**P4**) to its field of application, i.e., supply chain environments. Here, PRepChain supports these dynamic environments by distributing competences and reducing the required interactivity of businesses by offloading storage and processing of information to dedicated, independent entities. Further, by combining subjective and objective ratings as well as requiring (pseudonymous) authentication, PRepChain offers a reliable rating process. By distinguishing aspects such as friendliness and delivery delays during the voting process, our design also achieves a flexible granularity for reputation derivation and aggregation. Lastly, as highlighted in our evaluation (Section 5.3), it provides adequate performance, scalability, and security for dynamic, real-world supply networks.

Extensibility. To further extend the features and guarantees offered by PRepChain, its native interaction with existing SCISs allows for refining the information sources for objective ratings. For instance, the integration of trusted sensing [46] is supported by default and could offer valuable, guaranteed-reliable information for businesses' reputations in various real-world deployments. The previously mentioned verification engine (cf. Section 5.4) could also be extended to verify the integrity and correctness of the information from trusted sensing and the information systems in general, ensuring that this information correctly impacts a business' reputation.

Mode of Deployment. By nature, confidentiality (**P2**) and availability of information (**P4**) are slightly conflicting properties in the scenario of a business leaving the system due to the handling of the required key material. PRepChain ensures the availability and confidentiality of the ratings and scores by storing them FHE-encrypted at the reputation manager. However, PRepChain also supports an adjusted role of the key manager, i.e., an alternative mode of deployment. In the default scenario, which we outlined in this paper, confidentiality is valued higher than availability: The keys required for decrypting a business' reputation are always held by the business itself. Once it leaves the ecosystem, no further access to the private key is possible, severely impairing the availability of its reputation scores.

Alternatively, the handling of private keys can also be shifted to the key manager(s), ensuring the availability of all required information at the loss of control over the key material by the affected businesses. In this case, the storage needs as well as the networking load of the key manager(s) would slightly increase (cf. Table 2). However, we consider this overhead as marginal, i.e., this adapted PRepChain would still be practical and performant in real-world deployments.

Operational Trade-offs. We designed PRepChain in a way so that it comes with an adjustable configuration, allowing for setting-specific deployments. Overall, these operational trade-offs are as follows. First, as we have just discussed, the mode of deployment primarily impacts the availability of reputation scores. However, it also affects the messages that are being sent during the inquiring process. Second, integrating votees into the rating process is optional in PRepChain. While we see significant benefits in incorporating the perspective of both voters and votees, certain deployments might not depend on it, providing room for fewer messages and reduced processing. Third, the threshold k, which defines when new ratings are being added to the reputation score, is adjustable, impacting the timeliness, achievable confidentiality guarantees, and processing overhead. Lastly, as our performance evaluation highlights, voting and inquiry processes both scale to realistic supply chain settings. We are confident that the benefits and technical guarantees PRepChain provides make up for its deployment overhead and cost of operation. Given that PRepChain is only conceptually centralized, adding computing hardware and resources is supported through both horizontal (scaling out) and vertical (scaling up) scaling, offering flexibility.

A Constrained PRepChain Variant. As we have briefly hinted at in Section 4.4, conceptually, we could replace the FHE cryptosystem with a PHE cryptosystem like Paillier [57]. In this case, we would lose access to several operations on encrypted ratings. In particular, this constrained variant of PRepChain cannot support the voter–votee joining by the aggregation engine (⑧ of the voting process) as well as the weighted aggregation into reputation scores by the reputation manager (⑪ of the voting process), taking away significant benefits of PRepChain. These changes do not impact the inquiring process, but they impact the information that ratings and reputation scores can capture, i.e., ultimately, the reputation score's expressiveness. However, in scenario where this functionality is not needed, changing the cryptosystem would promise a significant performance overhead by reducing processing times and ciphertext sizes while promising accurate computations.

Systematic Classification. Apart from PRepChain's exceptional applicability to dynamic supply chain environments, we can further assess it in terms of "traditional" reputation system properties. Specifically, we apply the comparison framework by Hasan et al. [31], as we summarize in Table 2. As a ticket-based approach, PRepChain belongs to the category *token-based systems*. Its intuitive vector-based approach (cf. Fig. 3), which handles values in the range of $\mathbb{Q}$, and the support for flexible aggregation models are unique. Notably, despite this feature richness, PRepChain does not sacrifice other desired properties, e.g., durability or liveliness, either.

Practical Implications. Having a reliable, privacy-preserving reputation system in place allows companies to discover trustworthy business partners more easily and reliably. The prevailing business practice of relying solely on the historical data available in company's own SCIS or generally available but non-verified data when making network decisions is significantly expanded by the information provided by our objective reputation scores. Consequently, PRepChain supports the flexible and rapid evolution of supply networks. This benefit also increases business resilience and even supports the establishment of new business models and associated partnerships in the context of sustainability and circular economy.

As we have discussed, PRepChain offers significant added value for businesses in dynamic supply chain environments (while still being compatible with non-dynamic settings) as it enables data-based decision-making. Its flexibility further ensures that real-world use is possible across diverse use cases and settings.

7. Conclusion and takeaways

In this section, we highlight relevant takeaways from our research by summarizing our contributions in Section 7.1. Subsequently, we outline research directions for future work that follow from our work and findings in Section 7.2.

7.1. Conclusion

Dynamic supply chain environments introduce specific requirements, partly due to their unique setting, for reputation systems. Specifically, corresponding solutions need to ensure a genuine, secure, and accurate operation while also accounting for the applicability to the setting at hand (cf. Section 3.3). To the best of our understanding, research did not yet appropriately consider these (crucial) aspects.

Given this unexpected research gap (many different reputation systems have been proposed for other domains), we developed PRepChain, our approach to offering a versatile and privacy-preserving reputation system that is tailored toward practical use in modern supply chain environments. To achieve this applicability, we design PRepChain as a distributed, multi-agent information system (targeting reliability) that utilizes FHE to offer the desired confidentiality guarantees. With this design, we primarily advance the state of the art in four directions (cf. Section 4.4): (i) accounting for reliable privacy preservation in large-scale supply chain settings, (ii) ensuring (long-term) availability of reputation scores even in dynamic environments, (iii) sourcing information from existing SCISs by default, and (iv) distinguishing and processing both objective and subjective ratings. Our evaluation, which covers a real-world use case, and our detailed security discussion confirm the compliance with the intended system properties, providing valuable information for deployments and next-generation SCM approaches.

7.2. Future work

In addition to PRepChain addressing the identified research gap (Section 3.3) and bringing several advances to the field, it also paves the way for follow-up research. In particular, we can distinguish between aspects that would evolve our design beyond the scope of this paper and more general research directions in the context of reputation systems, particularly in (dynamic) supply chain environments.

Evolving PRepChain. Concerning the technical dimension, the focus of this paper, exploring the possibility of enhancing PRepChain's access control by utilizing attribute-based encryption [58] or by integrating it more closely with existing SCISs might be worthwhile. Likewise, studying whether multi-key FHE [59] or secret sharing/threshold cryptography [60] are able to introduce additional benefits is a promising next step. Since PRepChain cannot verify made computations, trusted execution environments or zero-knowledge proofs could potentially add further technical guarantees [20]. In the same direction, developments in the context of FHE [61,62] might introduce schemes and features that are worthwhile integrating into our design.

Orthogonal Research Directions. The technical contributions of PRepChain can have significant implications for other domains. In particular, business experts should assess the impact on SCM and how future business decisions are being made. Moreover, studying these influences also touches upon concepts that are known from game theory. Hence, future work (still) has to cover a lot of different angles before bringing sophisticated reputation systems to broad use. As a first step, we would like to explore PRepChain in a couple of more (dynamic) real-world deployments and study how businesses interact with the system in general, especially in light of the added value it introduces for their decision-making in higher-layer applications. Moreover, defining a pre-defined set of metrics that can be retrieved from SCIS for the objective reputation score and relevant assessment categories for the subjective reputation score could be a decisive step to enable practical implementation. In an effort to move toward real-world deployments, Berninger et al. [63] already proposed a conceptual framework for joining the business and technical dimensions when it comes to reputation in supply networks. Future work should bring this framework into (evaluated) practice.

Final Remark. With PRepChain, we neatly demonstrate that reasonable, practical privacy preservation in reputation systems—a novel application for business use—is achievable with today's concepts and technical building blocks, regardless of the need for availability and reliability, even in sophisticated dynamic supply chain environments.

Research data

We are committed to open science and as such we have open-sourced our prototype of PRepChain on GitHub [22].

CRediT authorship contribution statement

Jan Pennekamp: Writing – review & editing, Writing – original draft, Visualization, Supervision, Project administration, Methodology, Conceptualization. **Lennart Bader:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Methodology, Conceptualization. **Emildeon Thevaraj:** Writing – review & editing, Software, Investigation, Data curation. **Stefanie Berninger:** Writing – review & editing, Writing – original draft, Data curation. **Martin Perau:** Writing – review & editing. **Tobias Schröer:** Writing – review & editing. **Wolfgang Boos:** Writing – review & editing. **Salil S. Kanhere:** Writing – review & editing, Supervision. **Klaus Wehrle:** Writing – review & editing, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

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

Data availability

Due to the sensitive nature of our real-world use case data, we cannot disclose the manufacturing company's name nor share the original data used for deriving the reputation scores. Our implementation can be used with arbitrary data. Accordingly, we include exemplary data on GitHub [22].

References

[1] P. Gonczol, P. Katsikouli, L. Herskind, N. Dragoni, Blockchain implementations and use cases for supply chains-a survey, IEEE Access 8 (2020) 11856–11871, http://dx.doi.org/10.1109/ACCESS.2020.2964880.

[2] G.S. Ramachandran, S. Malik, S. Pal, A. Dorri, V. Dedeoglu, S. Kanhere, R. Jurdak, Blockchain in supply chain: Opportunities and design considerations, in: Handbook on Blockchain, vol. 194, Springer, 2022, pp. 541–576, http://dx.doi.org/10.1007/978-3-031-07535-3{_}17.

[3] E.I. Vazquez Melendez, P. Bergey, B. Smith, Blockchain technology for supply chain provenance: increasing supply chain efficiency and consumer trust, Supply Chain Manag. 29 (4) (2024) 706–730, http://dx.doi.org/10.1108/SCM-08-2023-0383.

[4] M. Linnartz, U. Motz, T. Schröer, V. Stich, K. Müller, C. Greb, Increasing resilience in procurement in the context of the internet of production, J. Prod. Syst. Logist. 1 (2021) (2021) http://dx.doi.org/10.15488/11350.

[5] H. Aboutorab, O.K. Hussain, M. Saberi, F.K. Hussain, A reinforcement learning-based framework for disruption risk identification in supply chains, Future Gener. Comput. Syst. 126 (2022) 110–122, http://dx.doi.org/10.1016/j.future.2021.08.004.

[6] P. Brauner, M. Dalibor, M. Jarke, I. Kunze, I. Koren, G. Lakemeyer, M. Liebenberg, J. Michael, J. Pennekamp, C. Quix, B. Rumpe, W. van der Aalst, K. Wehrle, A. Wortmann, M. Ziefle, A computer science perspective on digital transformation in production, ACM Trans. Internet Things 3 (2) (2022) http://dx.doi.org/10.1145/3502265.

[7] G. Dimitrakopoulos, P. Varga, T. Gutt, G. Schneider, H. Ehm, A. Hoess, M. Tauber, K. Karathanasopoulou, A. Lackner, J. Delsing, Industry 5.0: Research areas and challenges with artificial intelligence and human acceptance, IEEE Ind. Electron. Mag. 18 (4) (2024) 43–54, http://dx.doi.org/10.1109/MIE.2024.3387068.

[8] J. Pennekamp, R. Matzutt, C. Klinkmüller, L. Bader, M. Serror, E. Wagner, S. Malik, M. Spiß, J. Rahn, T. Gürpinar, E. Vlad, S.J.J. Leemans, S.S. Kanhere, V. Stich, K. Wehrle, An interdisciplinary survey on information flows in supply chains, ACM Comput. Surv. 56 (2) (2024) http://dx.doi.org/10.1145/3606693.

[9] S. Hemmrich, J. Schäfer, P. Hansmeier, D. Beverungen, The value of reputation systems in business contexts – a qualitative study taking the view of buyers, in: Proceedings of the 57th Hawaii International Conference on System Sciences, HICSS'24, AIS, 2024, pp. 4383–4392.

[10] E. Hofmann, H. Sternberg, H. Chen, A. Pflaum, G. Prockl, Supply chain management and industry 4.0: conducting research in the digital age, Int. J. Phys. Distrib. Logist. Manage. 49 (10) (2019) 945–955, http://dx.doi.org/10.1108/IJPDLM-11-2019-399.

[11] D. Ford, L.-E. Gadde, H. Hakansson, I. Snehota, Managing Business Relationships, third ed., Wiley, 2011.

[12] M. Ben-Daya, E. Hassini, Z. Bahroun, Internet of things and supply chain management: a literature review, Int. J. Prod. Res. 57 (15–16) (2019) 4719–4742, http://dx.doi.org/10.1080/00207543.2017.1402140.

[13] A. Leckel, M. Linnartz, Towards the internet of production–how to increase data sharing for successful supply chain collaboration, J. Prod. Syst. Logist. 3 (2023) http://dx.doi.org/10.15488/13409.

[14] L. Bader, J. Pennekamp, E. Thevaraj, M. Spiß, S.S. Kanhere, K. Wehrle, 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), 593, Springer, 2023, pp. 464–475, http://dx.doi.org/10.1007/978-3-031-63989-0_24.

[15] S. Malik, V. Dedeoglu, S.S. Kanhere, R. Jurdak, PrivChain: Provenance and privacy preservation in blockchain enabled supply chains, in: Proceedings of the 2022 IEEE International Conference on Blockchain (Blockchain '22), IEEE, 2022, pp. 157–166, http://dx.doi.org/10.1109/Blockchain55522.2022.00030.

[16] J. Pennekamp, J. Lohmöller, E. Vlad, J. Loos, N. Rodemann, P. Sapel, I.B. Fink, S. Schmitz, C. Hopmann, M. Jarke, G. Schuh, K. Wehrle, M. Henze, 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, pp. 489–505, http://dx.doi.org/10.1007/978-3-031-34560-9_29.

[17] A. Sarfaraz, R.K. Chakrabortty, D.L. Essam, AccessChain: An access control framework to protect data access in blockchain enabled supply chain, Future Gener. Comput. Syst. 148 (2023) 380–394, http://dx.doi.org/10.1016/j.future.2023.06.009.

[18] A. Mosteiro-Sanchez, M. Barcelo, J. Astorga, A. Urbieta, End to end secure data exchange in value chains with dynamic policy updates, Future Gener. Comput. Syst. 158 (2024) 333–345, http://dx.doi.org/10.1016/j.future.2024.04.053.

[19] T. Goertler, M. Papert, I. Fischer, D. Reich, N. Werner, I can see clearly now: A bibliometric exploration of digital platforms in supply chain management, in: Proceedings of the 57th Hawaii International Conference on System Sciences, HICSS'24, AIS, 2024, pp. 4995–5002.

[20] J. Lohmöller, J. Pennekamp, R. Matzutt, C.V. Schneider, E. Vlad, C. Trautwein, K. Wehrle, The unresolved need for dependable guarantees on security, sovereignty, and trust in data ecosystems, Data Knowl. Eng. 151 (2024) http://dx.doi.org/10.1016/j.datak.2024.102301.

[21] G. Felbermayr, K. Friesenbichler, M. Gerschberger, P. Klimek, B. Meyer, Designing EU supply chain regulation, Intereconomics 59 (1) (2024) 28–34, http://dx.doi.org/10.2478/ie-2024-0007.

[22] J. Pennekamp, L. Bader, E. Thevaraj, S. Berninger, M. Perau, T. Schröer, W. Boos, S.S. Kanhere, K. Wehrle, PRepChain:: A versatile privacy-preserving reputation system for dynamic supply chain environments, 2025, https://github.com/COMSYS/PRepChain.

[23] H.-H. Wiendahl, A. Kluth, T. Schröer, J. Janßen, M. Linnartz, R. Kipp, Supply chain management 2023, third ed., 2023, http://dx.doi.org/10.25716/thm-297, Aachener Marktspiegel Business Software.

[24] G. Schuh, R. Anderl, R. Dumitrescu, A. Krüger, M. ten Hompel, Industrie 4.0 Maturity Index: Managing the Digital Transformation of Companies, Technical Report, acatech STUDIE, 2017.

[25] L. Bader, J. Pennekamp, R. Matzutt, D. Hedderich, M. Kowalski, V. Lücken, K. Wehrle, Blockchain-based privacy preservation for supply chains supporting lightweight multi-hop information accountability, Inf. Process. Manage. 58 (3) (2021) http://dx.doi.org/10.1016/j.ipm.2021.102529.

[26] C. Braziotis, M. Bourlakis, H. Rogers, J. Tannock, Supply chains and supply networks: distinctions and overlaps, Supply Chain Manag.: Int. J. 18 (6) (2013) 644–652, http://dx.doi.org/10.1108/SCM-07-2012-0260.

[27] A. Brown, BDO: 50% of manufacturers plan to secure backup suppliers in 2021, 2021, https://www.supplychaindive.com/news/BDO-manufacturing-survey-reshore-supplier-technology/593409/. (Accessed 7 October 2024).

[28] B. Rolf, V. Klementzki, S. Lang, I. Jackson, S. Trojahn, T. Reggelin, A scoping review on dynamic networks in supply chains, IFAC- Pap. 56 (2) (2023) 203–214, http://dx.doi.org/10.1016/j.ifacol.2023.10.1570.

[29] L. Chang, Y. Ouzrout, A. Nongaillard, A. Bouras, Z. Jiliu, Multi-criteria decision making based on trust and reputation in supply chain, Int. J. Prod. Econ. 147 (Part B) (2014) 362–372, http://dx.doi.org/10.1016/j.ijpe.2013.04.014.

[30] S. Gurtler, I. Goldberg, Sok: Privacy-preserving reputation systems, Proc. Priv. Enhancing Technol. 2021 (1) (2021) 107–127, http://dx.doi.org/10.2478/popets-2021-0007.

[31] O. Hasan, L. Brunie, E. Bertino, Privacy-preserving reputation systems based on blockchain and other cryptographic building blocks: A survey, ACM Comput. Surv. 55 (2) (2022) http://dx.doi.org/10.1145/3490236.

[32] A. Acar, H. Aksu, A.S. Uluagac, M. Conti, A survey on homomorphic encryption schemes: Theory and implementation, ACM Comput. Surv. 51 (4) (2018) 1–35, http://dx.doi.org/10.1145/3214303.

[33] J.H. Cheon, A. Kim, M. Kim, Y. Song, Homomorphic encryption for arithmetic of approximate numbers, in: Proceedings of the 23rd International Conference on the Theory and Application of Cryptology and Information Security, ASIACRYPT'17, 10624, Springer, 2017, pp. 409–437, http://dx.doi.org/10.1007/978-3-319-70694-8_15.

[34] F. Hendrikx, K. Bubendorfer, R. Chard, Reputation systems: A survey and taxonomy, J. Parallel Distrib. Comput. 75 (2015) 184–197, http://dx.doi.org/10.1016/j.jpdc.2014.08.004.

[35] Z. Shi, K. Srinivasan, K. Zhang, Design of platform reputation systems: Optimal information disclosure, Mark. Sci. 42 (3) (2023) 500–520, http://dx.doi.org/10.1287/mksc.2022.1392.

[36] Z. Zhou, M. Wang, C.-N. Yang, Z. Fu, X. Sun, Q.M.J. Wu, Blockchain-based decentralized reputation system in E-commerce environment, Future Gener. Comput. Syst. 124 (2021) 155–167, http://dx.doi.org/10.1016/j.future.2021.05.035.

[37] B. Liu, T. Ju, J. Lu, H.K. Chan, Hide away from implication: potential environmental reputation spillover and strategic concealment of supply chain partners' identities, Int. J. Oper. Prod. Manage. 44 (9) (2024) 1595–1620, http://dx.doi.org/10.1108/IJOPM-08-2023-0649.

[38] D. von Berlepsch, F. Lemke, M. Gorton, The importance of corporate reputation for sustainable supply chains: A systematic literature review, bibliometric mapping, and research agenda, J. Bus. Ethics 189 (1) (2024) 9–34, http://dx.doi.org/10.1007/s10551-022-05268-x.

[39] X. Xu, J. Gu, H. Yan, W. Liu, L. Qi, X. Zhou, Reputation-aware supplier assessment for blockchain-enabled supply chain in industry 4.0, IEEE Trans. Ind. Inform. 19 (4) (2023) 5485–5494, http://dx.doi.org/10.1109/TII.2022.3190380.

[40] M. Möhlmann, T. Teubner, A. Graul, Leveraging trust on sharing economy platforms: Reputation systems, blockchain technology and cryptocurrencies, in: Handbook of the Sharing Economy, Edward Elgar Publishing, 2019, pp. 290–302, http://dx.doi.org/10.4337/9781788110549.00033.

[41] S. Hemmrich, Business reputation systems based on blockchain technology—A risky advance, in: Proceedings of the 31st European Conference on Information Systems, ECIS'23, AIS, 2023.

[42] Ö. Dogan, H. Karacan, A blockchain-based E-commerce reputation system built with verifiable credentials, IEEE Access 11 (2023) 47080–47097, http://dx.doi.org/10.1109/ACCESS.2023.3274707.

[43] M.A. Azad, S. Bag, F. Hao, PrivBox: Verifiable decentralized reputation system for online marketplaces, Future Gener. Comput. Syst. 89 (2018) 44–57, http://dx.doi.org/10.1016/j.future.2018.05.069.

[44] 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, http://dx.doi.org/10.14722/ndss.2014.23379.

[45] J. Pennekamp, F. Alder, R. Matzutt, J.T. Mühlberg, F. Piessens, K. Wehrle, Secure end-to-end sensing in supply chains, in: Proceedings of the 2020 IEEE Conference on Communications and Network Security, CNS'20, IEEE, 2020, http://dx.doi.org/10.1109/CNS48642.2020.9162337, Proceedings of the 5th International Workshop on Cyber-Physical Systems Security (CPS-Sec '20).

[46] J. Pennekamp, F. Alder, L. Bader, G. Scopelliti, K. Wehrle, J.T. Mühlberg, Securing sensing in supply chains: Opportunities, building blocks, and designs, IEEE Access 12 (2024) 9350–9368, http://dx.doi.org/10.1109/ACCESS.2024.3350778.

[47] MongoDB Inc., Mongodb, 2009, https://www.mongodb.com.

[48] A. Ibarrondo, Pyfhel, 2017, https://github.com/ibarrond/Pyfhel.

[49] I. Microsoft, Microsoft SEAL, 2018, https://github.com/Microsoft/SEAL.

[50] cryptography.io, Fernet (symmetric encryption), 2014, https://cryptography.io/en/latest/fernet/.

[51] R. Rivest, A. Shamir, L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, Commun. ACM 21 (2) (1978) 120–126, http://dx.doi.org/10.1145/359340.359342.

[52] A. Ronacher, Flask, 2010, https://palletsprojects.com/p/flask/.

[53] V. Sidorov, E.Y.F. Wei, W.K. Ng, Comprehensive performance analysis of homomorphic cryptosystems for practical data processing, 2022, http://dx.doi.org/10.48550/arXiv.2202.02960, arXiv:2202.02960.

[54] L. Jiang, L. Ju, Fhebench: Benchmarking fully homomorphic encryption schemes, 2022, http://dx.doi.org/10.48550/arXiv.2203.00728, arXiv:2203.00728.

[55] J. Takeshita, N. Koirala, C. McKechney, T. Jung, Heprofiler: An in-depth profiler of approximate homomorphic encryption libraries, 2024, Cryptology ePrint Archive 2024/1059.

[56] T. Rahman, A.M.I.M. Osmani, M.S. Rahman, M.M.A. Shibly, S. Islam, Benchmarking fully homomorphic encryption libraries in IoT devices, in: Proceedings of the 11th International Conference on Networking, Systems, and Security (NSysS '24), ACM, 2024, pp. 16–23, http://dx.doi.org/10.1145/3704522.3704546.

[57] P. Paillier, Public-key cryptosystems based on composite degree residuosity classes, in: Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques, EUROCRYPT'99, 1592, Springer, 1999, pp. 223–238, http://dx.doi.org/10.1007/3-540-48910-X_16.

[58] J. Bethencourt, A. Sahai, B. Waters, Ciphertext-policy attribute-based encryption, in: Proceedings of the 2007 IEEE Symposium on Security and Privacy (SP '07), IEEE, 2007, pp. 321–334, http://dx.doi.org/10.1109/SP.2007.11.

[59] M. Clear, C. McGoldrick, Multi-identity and multi-key leveled FHE from learning with errors, CRYPTO'15, in: Proceedings of 35th Annual Cryptology Conference, vol. 9216, Springer, 2015, pp. 630–656, http://dx.doi.org/10.1007/978-3-662-48000-7_31.

[60] M. Stadler, Publicly verifiable secret sharing, in: Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques, EUROCRYPT'96, 1070, Springer, 1996, pp. 190–199, http://dx.doi.org/10.1007/3-540-68339-9_17.

[61] A. Viand, P. Jattke, A. Hithnawi, Sok: Fully homomorphic encryption compilers, in: Proceedings of the 2021 IEEE Symposium on Security and Privacy (SP '21), IEEE, 2021, pp. 1092–1108, http://dx.doi.org/10.1109/SP40001.2021.00068.

[62] A. Viand, C. Knabenhans, A. Hithnawi, Verifiable fully homomorphic encryption, 2023, http://dx.doi.org/10.48550/arXiv.2301.07041, arXiv:2301.07041.

[63] S. Berninger, S.-Y. Kim, J. Piel, M. Perau, S. Geisler, F. Piller, K. Wehrle, J. Pennekamp, Privacy-aware supply chain ratings: Interdisciplinary research on collaborative supply chain management, in: Proceedings of the 7th Conference on Production Systems and Logistics, CPSL'25, publish-Ing., 2025, pp. 343–354, http://dx.doi.org/10.15488/18879.

Jan Pennekamp received the B.Sc., M.Sc., and PhD degrees in Computer Science from RWTH Aachen University. He is a researcher at the Chair of Communication and Distributed Systems (COMSYS) at RWTH Aachen University. His research focuses on security & privacy aspects in the Industrial Internet of Things (IIoT). In particular, his special interests include privacy-enhancing technologies, the design of privacy-preserving protocols, and secure computations as well as their application.

Lennart Bader received the B.Sc. and M.Sc. degrees in Computer Science from RWTH Aachen University. He is a researcher at the Cyber Analysis & Defense (CA&D) department at Fraunhofer Institute for Communication, Information Processing and Ergonomics FKIE where he is a member of the Secure Production & Energy Networks group. Accordingly, his research primarily covers the security in industrial networks with a specific focus on energy networks.

Emildeon Thevaraj, M.Sc., studied computer science at RWTH Aachen University, Germany, and Charles University in Prague, Czech Republic. He has been working in the Production Management group at FIR at RWTH Aachen University, Germany, since 2019.

Stefanie Berninger, M.Sc., studied Logistics and Transport Management at the University of Gothenburg in Sweden and Stellenbosch University in South Africa. She has been working in the Supply Chain Management group at FIR at RWTH Aachen University since 2021. Her research focuses on supply-chain-management and circular economy and is part of the Cluster of Excellence "Internet of Production".

Martin Perau, M.Sc., studied production engineering and economics at RWTH Aachen University and Aalto University in Finland. He has been working at FIR at RWTH Aachen University since 2022 and has headed the Supply Chain Management department since 2024.

Tobias Schröer, M.Sc., studied industrial engineering at the Technical University of Clausthal and has been working at the FIR at RWTH Aachen University since 2016. He is currently Head of Production Management at the FIR at RWTH Aachen University.

Prof. Dr. Wolfgang Boos, MBA studied mechanical engineering at RWTH Aachen University after completing his training as a tool mechanic. Following his studies, he completed his doctorate at the Machine Tool Laboratory (WZL) and became managing senior engineer of the Chair of Production Systems at the WZL. He has been an adjunct professor there since 2018 and works as a lecturer in the field of sustainable production. Prof. Dr. Boos has been Managing Director of FIR e.V. at RWTH Aachen University since January 2023.

Salil S. Kanhere is a Professor of Computer Science and Engineering at UNSW Sydney, Australia. He received his M.S. and Ph.D. degrees from Drexel University, Philadelphia, USA. His research interests span the Internet of Things, cybersecurity, pervasive computing, blockchain, and machine learning. He received the Friedrich Wilhelm Bessel Research Award (2020) and the Humboldt Research Fellowship (2014), from the Alexander von Humboldt Foundation in Germany. He is a Distinguished Member of the ACM and was an ACM Distinguished Speaker from 2019–2021. He is a Senior Member of the IEEE and an IEEE Computer Society Distinguished Visitor.

Klaus Wehrle received the Diploma (equiv. M.Sc.) and PhD degrees from University of Karlsruhe (now KIT). Since 2010, he is a full professor and the head of the Chair of Communication and Distributed Systems (COMSYS) at RWTH Aachen University. His research interests include (but are not limited to) engineering of networking protocols, (formal) methods for protocol engineering and network analysis, reliable communication software, and operating system issues of networking. He is a member of IEEE, ACM, VDE, GI/ITG-Fachgruppe KuVS, and ACATECH. Before joining RWTH Aachen University in 2006, he was a postdoctoral researcher at the International Computer Science Institute (ICSI).