

SOLUTION: SECURING SENSING IN SUPPLY CHAINS: OPPORTUNITIES, BUILDING BLOCKS, AND DESIGNS

Anonymous authors

Paper under double-blind review

Abstract

Supply chains increasingly develop toward complex networks with a growing number of actors. The lack of mutual trust in such networks results in challenges that are exacerbated by stringent requirements for shipping conditions or quality, and where actors may attempt to reduce costs or cover up incidents. In this paper, we develop and comprehensively study four scenarios that eventually lead to end-to-end-secured sensing in complex supply chains with many mutually distrusting actors while highlighting relevant pitfalls and challenges—details that are still missing in related work. Our designs ensure that sensed data is securely transmitted and stored, and can be verified by all parties. To prove practical feasibility, we evaluate the most elaborate design with regard to performance, cost, deployment, and also trust implications on the basis of prevalent (mis)use cases. Our work enables a notion of secure end-to-end sensing with minimal trust across the system stack, even for complex supply chain networks with opaque communication.

1 Introduction

Supply chain management involves provisioning, internal and external suppliers, vendors, logistics, bookkeeping, and billing [17,61]. Depending on the final product, supply chains can be highly complex as they include a copious number of actors and business interests, many interdependent production steps, and potentially high levels of variability and uncertainty [40,47]. Complex trust relationships in digital infrastructures—e.g., sensing and processing systems, opaque communication, and mobile networks—add to this.

While corresponding technical research on supply chain management largely and extensively focuses on the *flow* of information [17], advances in the Internet of Things (IoT), especially its *comprehensive sensing capabilities*, highlight the importance of considering the information’s *origin* and *acquisition processes*. These sensing-based approaches have drastically reshaped business processes, resulted in increasingly automated decision-making [11,48,63], and established new application domains, including industrial applications where IoT systems are used to monitor manufacturing processes [8], specifically in the context of intra-corporate deployments [34,47,63]. Apart from academia, IoT-oriented companies (e.g., upkeep [69], project44 [53], or roambee [55])

also seize the moment to widely deploy modern sensing technology in supply chains. Fittingly, recent research also investigates the use of such sensing equipment in the context of multi-stakeholder deployments that are characteristic of complex supply chain networks [11,29,62]. The corresponding large number of actors opens up important questions around *cyber security, dependability, trust relationships, authenticity, and accountability* to enable robust supply chains [46].

Pennekamp et al. [46] argue that end-to-end (E2E) secure means of supply chain monitoring can be implemented using Trusted Execution Environments (TEEs), tightly interwoven with immutable storage such as immutable databases based on Merkle trees [38] or public ledgers [37]. Specifically for perishable supply chains—cold chains in food supply or pharmaceuticals—E2E secure setups seem to introduce acceptable infrastructure costs, enable near real-time monitoring and detection of incidents, and minimize the required trust between the involved parties. This prior work superficially covers two aspects: (1) How to ensure that sensor data can be trusted, especially if multiple stakeholders sense, forward, process, and use the information, and (2) how and when distrusting stakeholders can trust and rely on each others’ (historic) sensor data. However, the authors leave crucial design details and relevant security considerations for future work, do not provide a discussion of deployment aspects, do not cover varying levels of trust relationships, and neglected performance and security evaluations [46].

This Paper. The overall complexity of modern supply chains is naturally reflected in design choices for supply chain monitoring systems, each of which involves different trust relationships and systems costs. The overall impact of design decisions when construing such a (secured) data processing pipeline is not well understood, and the definite applicability of possible solutions is unknown. With this paper, we intend to fill this research gap. Specifically, we design, investigate and compare several supply chain sensing scenarios concerning the previously unexplored security implications, i.e., tamperproofness, authenticity, and accountability, the resulting trust relationships, and implementation costs and performance. In our evaluation, we resort to five common use cases and four realistic misuse cases in supply chains to assess the technical readiness. We further provide a performance evaluation of the relevant components. As a result, we move the trust in processed data toward real-world requirements and explore the applicability of secure E2E sensing.

Contributions. The main contribution of this paper is a comprehensive design of four supply chain sensing and monitoring scenarios that involve different approaches to acquiring, processing, and storing sensed data based on the availability of trusted processing elements, management infrastructure for cryptographic credentials, and tamperproof storage solutions. More in detail:

- (a) We compile a universally-valid, EPCIS [21]-aligned list of desirable supply-chain-sensing use cases that expresses the respective sensor types, data volumes, sensing frequencies, and latency requirements.
- (b) We provide a list of realistic and threatening misuse cases that highlight the potential for actors to deceive others.
- (c) We design, investigate, and rate different evolutions for reliable end-to-end secure sensing scenarios that cover all sorts of supply chains. We derive these designs iteratively, which allows for precisely configured real-world deployments (individual needs and attacker model).
- (d) We evaluate and discuss the performance, security (misuse cases), and cost implications of the different designs, and place them in the context of real-world use.

We emphasize that any market-ready E2E-secured design should utilize TEE-based sensors and processing that eventually persists an immutable trust anchor. Such a design ensures trustworthy processing, long-term availability and authenticity of all data, even in dynamic settings, and with a minimal trusted computing base.

Beyond Supply Chains. We believe that our designs, analyses, and evaluation results are useful beyond the scope of supply-chain monitoring but can inform engineering decisions toward secure and trustworthy distributed many-stakeholder systems in the context of IoT, Cyber-Physical Systems, and for future critical infrastructures. Our prototypic open-source implementation and software development framework could be a valuable case study and starting point for future research to build upon and extend.

Open Science. Our paper’s evaluation artifacts will be open-sourced to ensure reproducibility and reusability.

2 Background: Technical Concepts

As a foundation for our work, we now give a concise overview of the two most-relevant, technical concepts.

Trusted Computing is a set of hardware security mechanisms to shield software on a device from untrusted access [30, 59]. Based on a hardware root of trust, software can isolate itself to become integrity and confidentiality protected from the surrounding untrusted operating system [30]. While different approaches can achieve this isolation, Trusted Execution Environments (TEEs) are a popular method that either provide separate trusted and untrusted environments, like ARM TrustZone [50], or that granularly shield specific memory regions, like Intel SGX [14] or Sancus [43]. In either

of these cases, TEEs typically enable isolated software to additionally *attest* itself to remote stakeholders. This attestation allows to bind an authenticated and encrypted communication channel to a specific component on the trusted computing device, giving the remote stakeholder the guarantee that they can verifiably and securely communicate with specific software on the device, which is useful for mutually distrusting actors. Related work also demonstrates attestation mechanisms between TEEs from different vendors [60, 71].

Concerning IoT and environmental sensing, some TEEs allow the direct (i.e., secure) control of peripherals by trusted software [30, 59]. This feature enables remotely attestable software to provide authentic and integrity-protected measurements that are independently verifiable by remote parties. In this work, we refer to them as *trusted sensors*.

Blockchain Technology orthogonally can support the reliable long-term storage of information in settings with mutually distrusting parties. As immutable and distributed append-only ledgers, blockchains utilize cryptography to irrevocably link information-containing *blocks* to form a chain [42]. By appending new blocks, altering or removing older blocks becomes computationally infeasible. Blockchains hence achieve *tamperproofness* and can guarantee the *existence* of data in distributed settings, replacing trusted third parties. These long-term guarantees make blockchains a valuable tool to enhance collaborations in (potentially) low-trust supply chains [5, 33, 66]. Despite its accountability features, scalability issues regarding computational and storage overheads remain open challenges [36, 77] and require careful consideration of the amount of data to store as well as the underlying consensus algorithm [58]. A common approach for minimizing both performance and storage overhead is to only persist *fingerprints* of the data on the blockchain [5]. Simultaneously, fingerprints also mitigate common (data) privacy concerns.

Over time, different variants of blockchains have emerged, namely *public* and *permissioned*. They offer different trade-offs regarding accountability, privacy, and scalability: While public blockchains are usually operated publicly without focusing on a single use case and offer a high degree of accountability, they particularly face scalability issues and require special considerations of privacy aspects. Permissioned (or *private*) blockchains usually offer better scalability because they are tailored to specific applications and operate with specialized consensus algorithms. However, their degree of accountability is arguably weaker than with public blockchains.

3 Supply Chain Sensing and Processing

Based on interactions with supply chain experts, we now briefly illustrate supply chains, the involved actors, as well as desirable sensing use cases while only presenting details that are relevant for our setting—the sensing in supply chains. Sourcing this overview, we then compile a list of goals for the reliable data collection and processing in supply chains.

3.1 Scenario Overview

Abstracted, supply chains primarily consist of a physical dimension (the flow of shipments, parcels, products, and paper-based documentation) and a digital one, which covers the exchange of information and data along the supply chain. Most importantly, supply chains are usually composed of multiple stakeholders where, especially over multiple hops, pre- and succeeding companies might not be trusted or even known to data-processing actors. Thus, the digital dimension requires security mechanisms to ensure a trustworthy and reliable flow of information across companies and stakeholders.

Involved Stakeholders. Various actors can sense, forward, and process information: (i) *production-related* commodity corporations, manufacturing companies, and machine suppliers; (ii) *logistics-related* shipping companies, customs authorities, and warehousing services; and (iii) *sales-related* distributors, retailers, and customers. Particularly, sales-related stakeholders are interested in the product’s origin, i.e., through its history of sensed data, whether it may be due to fair-trade, sustainability, or authenticity needs. Thus, within a supply chain, a multitude of (sensed) data can be demanded and provided by individual stakeholders at the same time.

Sensing and Monitoring. Information on a shipment is critical for today’s supply chain management. Broadly, we identify three groups of shipment information: (i) Tracking data, i.e., where the shipment is; (ii) monitoring data, i.e., what condition the shipment is in; and finally, (iii) information on the product itself (e.g., to swiftly adapt production processes). Gathering and handling this data is of utmost importance to the entire supply chain as it allows companies to manage their processes and schedule their operations.

Research Gap. From this brief description, we derive that the flow of information and its distributed sensing within supply chains is challenging and complex. Numerous threats during sensing, forwarding, processing, and (long-term) storage arise due to the large number of actors with potentially non-existing or little trust relationships. Unfortunately, today’s supply chains lack the technical means to satisfactorily address them. As a result, many actors still use inefficient processes (e.g., paper-based reports), do not collect desired information, or refrain from trusting received (remote) data.

Thus, we identify the need to provide actors with a secure E2E sensing design to allow them to reliably (trustworthy and timely) detect undesirable delivery statuses or environmental conditions of their shipments and other issues, even when untrusted stakeholders reported otherwise. In fact, we even plan for *distrust* between stakeholders, where different stakeholders have no prior reason to trust one another. Long-term availability of data would further improve the actors’ chances for accountability and (data) verifiability, e.g., to unequivocally assign blame in case of disputes. Hence, closing this gap is especially beneficial for issues that are not immediately noticeable, like a temporarily interrupted cold chain.




Figure 1: Overview of the logical actors: Sensing parties own sensors and forward data to requesting parties who may share it with accessing parties on-demand for later use.

3.2 Logical Actors in Supply Chains

When analyzing the flow of sensed data, we identify three classes of logical actors in our setting: sensing, requesting, and accessing parties. We illustrate their relation in Figure 1. Conceptually, each supply chain stakeholder can take the roles of multiple logical actors simultaneously.

First, a *sensing party* owns and deploys the sensors in use. It also makes sure to provide the sensed information to the requesting party. The sensing party is usually a shipment provider during transit, but for shipments with sensitive, expensive, or fragile cargo, a customer might also request the inclusion of its own sensors. Likewise, warehousing departments can act as sensing parties when utilizing smart readers to process incoming or outgoing shipments.

Second, *requesting parties* are the intended, original recipients of sensed data, e.g., customers requesting details on their purchase and the transit of their goods. The requesting party may also be the sensing party (for documentation and benchmarking purposes). Still, sensed information is usually relevant to only one party. However, if shipment providers group cargo by different customers in a single container, all customers might be interested, e.g., in maintaining the cold chain. Focusing on potential data availability needs, the requesting party is responsible for the long-term storage.

Third, additional *accessing parties* might be interested in the sensed information at a later point in time. Accessing parties can be virtually any stakeholders that are concerned with the supply chain, from production companies, over suppliers, retailers, and governmental agencies, to end customers. In this case, the sensed information is originally shared by the requesting party, and if no direct business relationship between the accessing and requesting party exists, the information must pass multiple hops. For example, following an accident, data might only become relevant after years.

3.3 Typical Sensing in Supply Chains

Based on the electronic product code information services (EPCIS) industry standard [9, 21], we can identify a number of relevant sensing use cases in supply chains. The corresponding EPCIS data model captures the dimensions of what, when, where, and why [21], the latter being irrelevant in our scenario. The recent successor, which is intended for state-of-the-art supply chain data interoperability, further includes the

dimension “how” [22], matching the research gap of providing companies with reliable (sensed) data.

3.3.1 Use Cases Overview

When looking at different sensing applications within supply chains, we identify five general use cases with increasing (technical) complexity and processing requirements. They are not limited to a single granularity (e.g., shipments only). Instead, they can either be applied on a shipment, parcel, or product level according to the specific use case needs. By default, each measurement includes the dimensions *what* (i.e., the focus of the measurement) and *when* (i.e., timing information). Thus, we focus on *where* and *how* instead.

Status Tracking. To track the status of a physical flow, requesting parties are interested in corresponding status changes. For example, when handling a parcel, such as moving it from a container to a warehouse or changing the mode of transportation (e.g., from truck to aircraft), this information, as well as its location (*where*), must be recorded. This tracking can be achieved using stationary RFID readers or BLE beacons (*how*). To the general public, this use case is well-known from the last-mile tracking of consumer parcels.

Location Tracking. If requesting parties also want to know the approximate locations of products, they demand periodic or real-time updates of the respective locations. Thus, location sensors (*how*) must reliably sense this information (*where*). In the context of consumer parcels, this sensed location data is nowadays frequently available as part of last-mile deliveries.

Integrity Monitoring. Moreover, companies might also be interested in the (physical) integrity of their shipments (*where*). For example, when dealing with pharmaceuticals, detailed documentation might even be required by law [18]. Potential violations cannot only result in monetary damages but also in harm to humans (e.g., food poisoning). Thus, precisely capturing such data and maintaining access logs are important aspects. To this end, sensing parties can deploy secure (smart) locks and other surveillance sensors (*how*).

Condition Monitoring. Extending this previous use case to a continual or real-time yet reliable monitoring (*where*) can be equally relevant. Most prominently, inspecting compliance with the temperature requirements is essential, e.g., to identify spoiled goods in cold chains. Likewise, manufacturers of sensitive products might define constraints for shipment environments, such as humidity or impacts. Thus, deployed sensors must reliably provide this information (*how*).

Visual Monitoring. When considering very valuable products or livestock in transit, video-based monitoring (*where* and *how*) is a use case as well. Depending on the setting, the corresponding image or video feed might only be transmitted after specific triggers, e.g., opening after unlocking a smart lock or when exceeding a specific noise threshold. Thus, the exact needs vary significantly depending on the setting.

In real-world deployments, stakeholders commonly rely

Use Case	Sensor	Payload	Measure. Frequency	Time Criticality
Status Track.	Smart Reader	<1– <100 KiB	Triggered	Minutes
Location Track.	e.g., GPS	<1 KiB	<1 records/min	Minutes
Integrity Mon.	Smart Lock	<1 KiB	Triggered	Hours
Condition Mon.	Various	<1–<10 KiB	<6 records/min	Hours
Visual Mon.	Camera	>10 KiB	Variable	Variable

Table 1: Technical overview and equipment of our use cases.

on a combination of these use cases. Depending on the exact setting, several types of sensors with different densities must be deployed: Location, temperature, humidity, air pressure (altitude), light, shock (impact), acceleration, tilt, or weight sensors, as well as smart locks and scanners (readers).

3.3.2 Technical Perspective

In light of the associated computational burden of our use cases, we also consider payload sizes and sensing frequencies (i.e., the processing bandwidth). Latency is of interest to guarantee a timely handling of sensed status or condition changes. In Table 1, we provide an overview of these aspects.

Overall, most use cases have moderate needs when sensing relevant information. The exact payload size depends on the sensor in use, as well as the size of added context information. For the first use case, status tracking, payload sizes can range from 96 bit for the most common type of RFID tags, up to 100 KiB or more for specialized RFID tags with extended user data [72]. Similarly, location tracking and integrity monitoring can be realized in less than 1 KiB, most of the time. For condition monitoring, different sensor types introduce varying payload sizes, with simple measurements covering several bytes. In contrast, visual imagery monitoring or video feeds may lead to more excessive needs (far greater than tens of kilobytes), depending on the required image quality and resolution, as well as available compression methods.

The individual sensing frequency for condition monitoring can be comparably high (i.e., several measurements per minute). However, the respective latency requirements are usually not demanding because the condition monitoring can often not be acted upon immediately during the transit of a shipment but is instead intended to be an authoritative reference. Essentially, while the information gathered by status and location tracking may be needed within minutes of gathering, integrity and condition monitoring only produce data that has to become available within hours since the data gathering. Data aggregation at the source could help to lower the amount of data to transmit, especially if only data outliers are of interest to requesting parties. Thus, the continual transmission of large payloads is unlikely in most settings.

3.4 Misuse Cases: Credibility Threats

With multiple supply chain actors, we need to consider several misuse cases in real-world settings that may all be used to deceive the requesting or an accessing party. All misuse cases build on one or multiple of the following actions.

Data Tampering. Parties might have an incentive to directly manipulate data during transmission for various reasons, e.g., shipment providers trying to cover up shipment treatment deficiencies. This desire may emerge in case of accidents that should not reflect negatively on a shipping company or once requesting parties try to deceive accessing parties.

Data Hiding. If direct data tampering is not possible, simply hiding the existence of data or of a specific range of data may be equally desirable. The lack of data may not be surprising to the victim and could easily be blamed on unreliable technology or environmental events such as power outages.

Data Injection. A malicious party could attempt to insert forged information, i.e., data that originates from unauthentic and unrelated sources or is made up entirely. Similar to the previous misuse cases, such actions could be useful to deceive actors and convince them to accept the shipment conditions.

3.5 Design Goals for Secure E2E Sensing

Based on these threats (misuse cases) and in line with prior work [46], we now compile five pressing (technical) design goals for end-to-end-secured sensing (in supply chains).

G1: Tamperproofness. Sensor data must be verifiably untampered when being assessed by the requesting or an accessing party at any point in time. This property is vital to ensure that sensor data can be relied upon by all parties.

G2: Authenticity. The design must ensure that sensed data verifiably originates from authentic sensors. Thereby, malicious actors are prevented from (retroactively) forging data. This goal covers both sensing and requesting parties.

G3: Completeness. To truly enable E2E-secured sensing, recipients (i.e., the requesting and all accessing parties) must be able to verify that the data they receive is complete. Apart from checking for data completeness from a single sensor, recipients need to confirm that the measurements of all relevant (i.e., deployed and expected) sensors are available.

G4: Affordability. Given the overhead of any technical solution, both the (one-time) costs for additional hardware or hardware upgrades and the associated operating costs should be kept to a minimum. Consequently, new designs should be careful to (i) avoid performing computation-intensive tasks and (ii) not introduce excessive duty cycles during (regular) operation. Otherwise, real-world deployments are unrealistic on low-cost IoT devices, preventing widespread adoption.

G5: Latency Agnosticism. Any solution must be agnostic to any network latencies or network disruptions (offline periods) experienced by the sensing nodes while generally supporting frequent live updates to infrequent batch uploads.

These goals capture the actors' reliability and security needs. Fulfilling them would greatly improve the basis for the decision-making of all involved stakeholders in a supply chain. However, we still need to consider relevant pitfalls.

4 Evolving Reliable Sensing Concepts

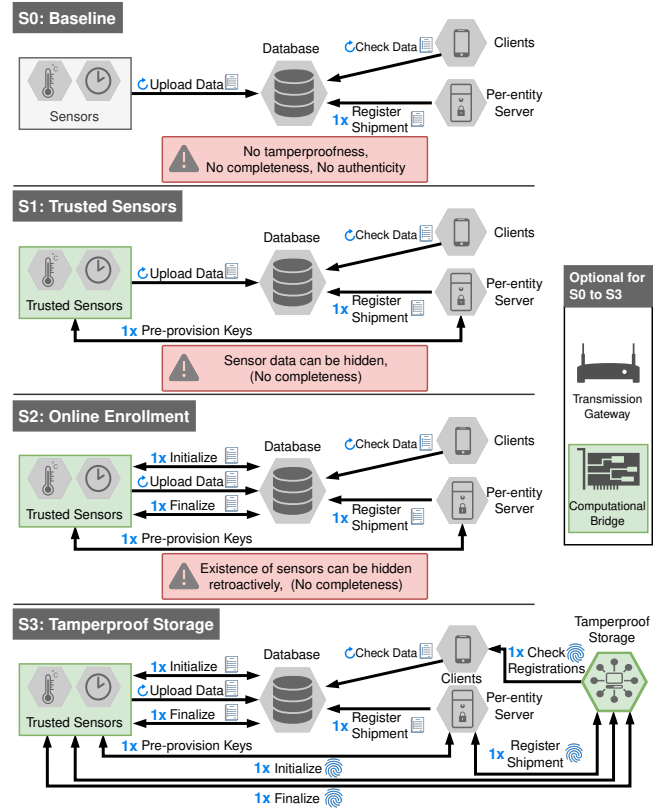


Figure 2: Comparison of supply chain sensing scenarios. The baseline (S0) does not provide trusted sensing. S1 introduces trusted sensors but lacks data completeness. S2's online phase addresses this issue but lacks sensor completeness. Finally, S3 includes a tamperproof storage for fingerprints.

To conflate potentially-diverging business interests of several stakeholders and the demand for increased transparency of shipment status, integrity, and product properties, specifically tailored yet flexible concepts are needed. Here, we propose our design that addresses the challenge of realizing secure E2E sensing in modern supply chains. Considering widely differing supply chain instances with differing needs, we derive our design based on a four-layered scenario model, which we illustrate in Figure 2. Our work bases on remote attestation and authenticated communication. As a high-level overview, we summarize the implications of each design stage on trust, security, and deployment aspects in Table 2.

We first introduce a traditional baseline sensing scenario in Section 4.1. In Section 4.2, we then outline and evolve our design decisions to secure this supply chain. Subsequently, in

	S0: Baseline	S1: Trusted Sensors	S2: Online Enrollment	S3: Tamperproof Storage
Design guarantees that...				
sensor data is not altered	–	✓	✓	✓
sensor data is not deleted	–	–	✓	✓
sensors cannot be hidden	–	–	–	✓
Deployment and Cost				
Communication latency	Configurable	Configurable	Configurable	Configurable
Sensor communication	Offline possible	Offline possible	Online phase required	Online phase required
Storage	Simple database (DB)	Simple DB	Simple DB	Simple DB plus tamperproof storage
Key management	–	Sensor keys	S1 + PKI for sensor receipts in online phase	Same as S2
Additional hardware costs	–	Trusted sensors	S1 + uplink hardware	Same as S2
Additional operational costs	–	–	Network uplink	S2 + tamperproof storage

Table 2: Comparison of the key design aspects in terms of security and deployment considerations for the scenarios in Figure 2.

Section 4.3, we briefly introduce optional components, and, in Section 4.4, we then elaborate on the technical details of the distilled, most-secure scenario. Finally, we discuss deployment challenges in Section 4.5. We defer a discussion on the security and limitations of secure E2E sensing to Section 6.

4.1 Baseline Scenario: (Insecure) Status Quo

Figure 2 illustrates a baseline scenario, which we denote as *S0*. Divided into arbitrarily many sensors and three remote stakeholders, the scenario centers around a storage database that is written to by sensing parties, e.g., shipment providers, to register shipments and their involved sensors. Each of these sensing parties is assumed to have its own server infrastructure, denoted as a per-entity server. The database may be maintained by an external party or by a shipment provider and should use, e.g., TLS certificates for authentication. Sensors registered to shipments regularly upload data to the database, which allows authenticated clients to retrieve the data and compare it to their expectations. In all discussed scenarios, clients are seen as an entity that accesses and verifies data, and can, as such, belong to any requesting or accessing party. At this level of abstraction, we focus on the trust a client has toward the other entities, and work toward an environment where the client has to maintain minimal trust in them.

Table 2 lists the guarantees given by the respective designs. The baseline scenario does not guarantee to clients that data has not been altered, deleted, or hidden. The overview shows that this, in practice, well-established baseline scenario only works in supply chains with already existing trust relationships. By being focused on cost-efficiency (G4), only a minimum of financial investments for sensor hardware and standard cloud computing costs is necessary.

However, this scenario is not suited for low-trust environments or for ensuring information integrity (G1–G3). Thus, in the following, we iterate on these trust assumptions to

Scen.	Validation Steps (additive)
	Perform the following steps for each sensor in the shipment:
S0	1. Check all sensor data for acceptable parameters
S1	2. Check integrity of all sensor data (through TEE attestation)
S2	3. Check the existence of start/end points of sensing periods 4. Check for data/measurement gaps in the sensing period
S3	5. Verify data completeness with tamperproof storage

Table 3: Steps to verify the validity in each sensing scenario.

reach a situation where the remaining assumptions are either unavoidable or can be accommodated by external means outside of the system design. Since such a reduction in trust necessarily comes with an increased burden on the client to verify a shipment’s validity, Table 3 keeps track of the steps a client needs to take. In *S0*, simple validation of the sensor data for acceptable parameters suffices.

4.2 Evolution Toward Trustworthy Sensing

As a first measure to guarantee data authenticity (G2), i.e., to ensure a verifiable information origin and prevent tampering with submitted data (G1), we extend the base scenario with *S1* to rely on *trusted sensors*. Attestation, as well as encryption and authentication of all communication, guarantee that involved stakeholders and external actors cannot tamper with data. In both *S0* and *S1*, sensor measurements can be cached offline (locally) until the shipment is completed, and can then be uploaded to the database as a whole dataset. Thus, the only effective change from *S0* to *S1* is the added hardware cost of employing trusted sensors, the added key management on the side of the sensing party, and the clients’ added computational effort to also verify the integrity and authenticity

of the data, ensuring that it originates from and is signed by the trusted sensors involved in the shipment. Nonetheless, without further measures, malicious stakeholders can still hide measurements (violating **G3**), either selectively or also a whole range of data starting at an arbitrary moment during shipment (*data hiding*).

To overcome these issues, the next scenario, **S2**, includes *online phases* at the beginning and the end of each shipment. During the first online phase, a sensor initializes the shipment by registering itself in the database to confirm its existence and relevance for the shipment. As soon as the shipment is completed, sensors finalize the shipment during the second online phase, indicating the number of measurements of this sensor during the shipment. While the initialization of the shipment signals a clear start of the shipment for future validation, the finalization ensures that no individual measurements are hidden, strengthening data completeness (**G3**). By involving the trusted sensors in an online phase, clients can verify that no data was lost in the database before sensing started, and that the sensing terminated only upon the final message in the database that also contains the number of measurements during this period. This improvement comes with the added requirement of a live network uplink during the start and end of a sensing period, in addition to the overhead of maintaining an infrastructure that allows the database and sensors to communicate directly with each other. Additionally, the start and end mechanisms have to be actively triggered from the outside, which can happen automatically when a sensor is turned on or through a message from a control server, i.e., as part of a handover procedure during shipping. We discuss all associated deployment considerations in Section 4.5.

While **S2** already greatly reduces the trust assumptions by the client, it still does not entirely prevent data hiding by removing all data of a sensor. A maliciously acting stakeholder could attempt to retroactively delete all records of a specific sensing from the database before the client accesses this information to filter out sensors with potentially encumbering measurements. Although this attempt might assume a high degree of criminal energy, **S3** takes this threat into account. Instead of solely relying on a potentially manipulable database, the respective fingerprints for shipment initialization and finalization are additionally stored on a tamperproof storage (e.g., a permissioned blockchain [5]). To enable a strict ordering of events, the shipment provider is required to also store a fingerprint of the shipment registration on the same tamperproof storage prior to the first sensor initialization. Thereby, we prevent retroactive hiding of sensors and ensure data inclusion for all client verifications. Although primarily relevant for ensuring data completeness (**G3**), the tamperproof storage further strengthens tamperproofness (**G1**) and authenticity (**G2**) in the presence of criminally acting parties. Since the tamperproof storage only stores fingerprints in the form of cryptographic hashes, its usage neither reduces privacy nor scalability. Next, we detail the technical aspects,

the information flow, and the respective implications on data trustworthiness and security of our final design.

4.3 Optional Components in E2E Sensing

In addition to the main components of our designs (sensors, per-entity server, and database), and for **S3**, also a tamper-proof storage, our design can further be extended with two optional and use-case-specific components.

A **transmission gateway** can be located in direct proximity of the (trusted) sensors to serve as a *cryptographically-passive* on-path relay that buffers data or simply serves as a network hub. It can be untrusted and does not require any trusted computing hardware. Overall, it can relieve (lightweight) sensors from the overhead of (i) supporting and managing (wireless) communication or (ii) accounting for sufficient buffer sizes (e.g., in use cases with longer offline periods).

In contrast, a **computational bridge** is equipped with trusted computing hardware to allow for *cryptographically-attested* on-path manipulations of the sensed data, e.g., to filter or aggregate measurements. Hence, they can also relieve lightweight sensors from complex (pre-)processing tasks. While the resulting computing requirements entail higher deployment costs in practice, computational bridges allow sensing actors to *reliably* reduce the amount of data that needs to be forwarded, stored, and processed without violating any of the introduced end-to-end security guarantees of our work.

4.4 Secure E2E Sensing: Technical Guarantees

S3 additionally stores the shipment registration as a fingerprint (cryptographic hash) in the tamperproof storage for two reasons. First, it creates a non-refutable link of used sensors to the shipment. Since the entry is signed by the shipment provider and integrity-protected by the storage, the sensors cannot be disassociated from the shipment anymore. Second, the entry in the storage serves, via relative ordering, as a timing marker for the sensors themselves when they register. Without this ordering, associating sensor data and shipment registration across different types of storage can be difficult. In this context, any sensor activation succeeding a shipment registration that includes this sensor is attached to the shipment, in contrast to requiring timekeeping for every sensor.

Once a shipment is registered, sensors can be activated and start with their online phase to initialize the sensing. This process again serves two purposes: First, the sensor uploads and verifies the upload of a non-refutable and signed statement that it is now active and can start sensing. This prevents retroactive claims by the shipment provider that the sensor was never active and thus never produced any data. Second, the sensor can provide additional metadata in this initialization message, such as starting conditions or even relative IDs that it uses as start markers in its sensing data. Due to these purposes, the sensor requires a response from the storage that

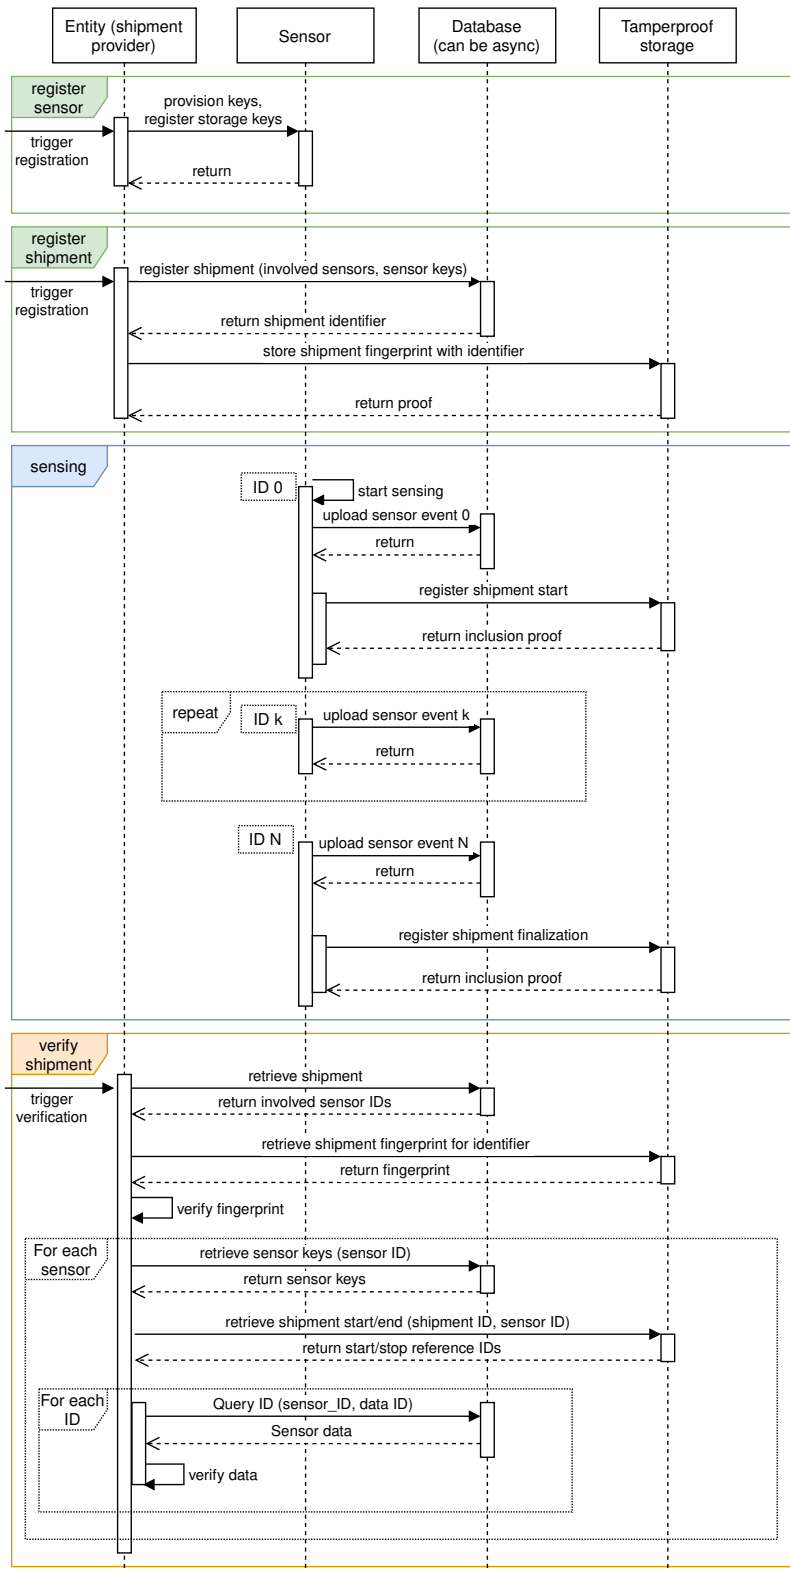


Figure 3: Sequence diagram for S3. The figure shows the complete sequence diagram of S3. After independent phases to register multiple sensors and a shipment, sensing can start after an online phase. A similar online phase is required to finish the sensing, requiring both beginning and ending proofs to verify the shipment.

its entry was included in the database, which requires local verification before the sensor starts operating. In Section 4.5, we discuss how a check of this inclusion proof impacts the practical performance requirements of the involved sensors.

After initializing the shipment and securing it in the tamperproof storage, the sensors start their operation and upload sensor data to the database, using an internal counter to keep track of the number of measurements made. By design, every sensor data that is stored in the tamperproof storage is also signed to prevent data integrity attacks during data upload. The final measurement counter is then used in the finalization of the shipment to announce and persist it reliably. The last step is again part of an online phase to ensure that the final measurement is also captured in the fingerprint that will be persisted in the tamperproof storage, i.e., the sensing has been terminated gracefully and covers all data.

As a result, clients can (i) retrieve the shipment registration for metadata, (ii) retrieve data of each sensor, and (iii) verify the authenticity and integrity of all received data to ensure it originated from this sensor. Figure 3 illustrates a complete sequence diagram of this verification process.

4.5 Crucial Deployment Challenges

Across these scenarios, various deployment challenges arise: We now discuss the associated key management, nuances for existing trust relationships with the storage provider, checking of inclusion proofs, and the support for optional components.

Key Management. For all scenarios using a trusted sensor (i.e., all but S0), the deployed sensors have to directly communicate with one or multiple storages. To do so securely, sensors and storage have to authenticate each other, which is usually realized through some form of multi-level public key infrastructure (PKI). Sensing parties are expected to pre-deploy devices with keys that exist within this PKI, which can be used by the storage to authenticate the sensor.

Lightweight architectures (e.g., Sancus [43]) may not be able to utilize PKI cryptography, i.e., they only support symmetric key operations. This constraint, however, is simply a deployment concern, as sensing parties can also distribute symmetric keys that are rotated after each shipment, for example, by communicating them to the storage for a specific shipment. Alternatively, a *computational bridge* (cf. Section 4.3) in the form of a more powerful TEE can serve as an on-path processing node between the (trusted) sensor and the database. Since such bridges may also support lightweight sensors by performing the inclusion proof checking, our prototype implementation utilizes one computational bridge to accurately assess the performance impact of such a design.

Trust in the Storage Provider. We envision an immutable ledger for the tamperproof storage in S3. However, its exact realization can vary based on the considered trust model and established trust relationships. For example, if a single trusted entity exists, they could maintain a singular trusted database

that also serves as tamperproof storage. As such, our design S3 may, in some situations, be equivalent to S2 if the storage database can be seen as a trustworthy alternative.

Related work provides certain database functionality from inside a TEE [45,56], which could be another method to establish trust in the database, removing the need for a tamperproof storage. If the provider running this service can then be verifiably prevented from performing rollback attacks on the data, these concepts are a viable alternative to an immutable ledger and make a TEE-backed storage practically tamperproof.

Inclusion Proof. S3 places additional work on the sensors to check and verify the inclusion proof received by the tamperproof storage. Depending on the nature and time constraints of this check, it can introduce non-negligible overhead for lightweight sensors. Additionally, depending on the used tamperproof storage, the online part of the communication with the sensor may suffer from undesirable delays.

We see two approaches to deal with these challenges. First, having an online phase immediately before sensing is not necessarily essential as long as the sensor performs the verification at some point. It could start sensing in the meantime. Hence, a delay in communication or computation may be acceptable until the inclusion proof has been verified. The sensor could already sense and upload data to the storage while the inclusion check is running in the background, which would allow to tentatively start a shipment.

If such a time overhead is infeasible, e.g., for constrained devices or short shipment durations, a TEE-based computational bridge between the sensor and the tamperproof storage could be deployed instead, as explained before. In this case, the sensor would directly communicate with the powerful TEE of the computational bridge and receive the confirmation (guarantee) that the inclusion proof has been verified on the tamperproof storage—without unacceptable delay.

Gateway to Buffer Data. While we have already discussed a computational bridge, lightweight sensors may additionally benefit from a passive, scalable on-path *transmission gateway* (cf. Section 4.3) that buffers data or serves as a network hub (as also suggested by [46]). Particularly, long-lasting shipments without permanent network connectivity may exceed the storage limitations of constrained sensor nodes, requiring a dedicated gateway between the (trusted) sensors and the database that receives all data from nearby sensors and uploads their data once network connectivity is restored.

5 Evaluation of Secure E2E Sensing

The discussed scenarios each introduce stronger technical sensing and processing guarantees. S1 adds trusted sensors to the baseline scenario to achieve tamperproofness and authenticity (G1-G2), and is improved upon by S2 and S3, which also tackle the issue of data and sensor hiding (G3). In the following, we evaluate our designs with regard to performance, deployment, and cost considerations.




Figure 4: Computation and transmission times of data from the trusted sensor to the bridge for varying payload sizes.

5.1 Performance Evaluation

Evaluating designs that focus on end-to-end-secured sensing and data processing boils down to the performance of the individual components. Accordingly, we now discuss the performance capabilities and requirements on the two core computational pillars of the discussed designs: (i) The sensing and processing equipment and (ii) the tamperproof storage, here, an immutable ledger. All reported performance evaluations are part of our artifacts.

5.1.1 Sensing and Processing Equipment

The selected sensing equipment—microcontrollers, storage, and peripherals—mostly depends on the expected throughput of data and the complexity of the processing to be performed on the sensor node. Specifically, we distinguish between extremely lightweight 16-bit equipment, such as Sancus processors [43] running at 8 MHz, and substantially more powerful 32-bit nodes based on, e.g., ARM Cortex with TrustZone.

To demonstrate the feasibility of our designs, we created a Sancus-based prototype implementation of S3. We show that even under these limiting resource constraints, sensors can address the requirements of S3 and easily handle, i.e., sense and process, most payloads, as compiled in Table 1 (except for video mon.). As we have discussed in Section 4.5, very lightweight sensors can utilize computational bridges to perform computationally intensive operations.

For this reason, the sensors in our prototype utilize an Intel SGX enclave that directly communicates with the Sancus enclaves, and which further serves as a computational bridge between the sensor and the cloud components. In our testbed, we thus simulate the technical requirements of S3 over 100 runs by attaching a XuLA StickIt! [74] board running Sancus 2.1 [57] to an Intel Core i3-7100U running Intel SGX on Ubuntu 22.04.1 LTS, with 16 GB of RAM. Since network conditions are highly use case-specific, this testbed setup abstracts away network delay. Moreover, we do not fix a delay introduced for communicating with the tamperproof storage but perform all computations that would be required for real-world deployments. These operations include the computation of fingerprints for the tamperproof storage (cf.

Section 5.1.3) and signing the data to be sent to the database with a pre-deployed public key certificate. Moreover, the Sancus and Intel SGX enclaves are mutually attested and all on-path communication is encrypted and authenticated.

Thus, this evaluation focuses on the conceptual performance of both hardware components to estimate the feasibility of deploying the most lightweight sensors available.

Starting and finalizing the sensing introduces a computational overhead of 130 ms; 83 ms of this time are consumed by the sensor, and 46 ms are spent by the computational bridge. These numbers exclude the time to verify data inclusion in the tamperproof storage as that depends on deployment specifics but include all necessary steps to upload the data. In Figure 4, we further detail the individual runtimes of the sensor and computational bridge for common payloads in supply chains between 1 and 10 KiB. The calculated 99 % confidence intervals illustrate only minor deviations over all runs. Each run is triggered externally, simulating a manual sensing request to the sensor. After sensing and encrypting the data, the sensor sends it to the bridge, which re-encrypts the data and signs the fingerprint with pre-deployed PKI keys. This setup allows us to evaluate the longest end-to-end timings and, as such, also serves as an over-approximation of any sensing that could periodically be triggered by the sensor itself, i.e., which is not triggered by an external component or party.

Our evaluation shows that even payloads of up to 10 KiB can be sent (and processed) by the sensor roughly 6 times a minute, where the sensor makes up for the majority of computation time and the Intel SGX-based bridge only takes between 55 and 150 ms. Only large payloads, such as high-resolution visual monitoring or specialized RFID tags that contain large unique identifiers of up to 100 KiB (Table 1), are infeasible to be processed on lightweight sensors. In these cases, more powerful microcontrollers should be considered for operation, which may lead to increased deployment costs.

5.1.2 Optional Gateway Equipment

In Section 4.5, we have discussed the use of a cryptographically passive transmission gateway between the deployed sensor and any cloud infrastructure. Depending on the use case, these gateways either purely serve as a networking hub or also cache data during offline periods. In the first case, the gateway simply requires minor storage for the cache. In the second case, the gateway only needs to buffer data until a network connection is restored, which may take anything from minutes to weeks in the case of long-distance shipments. However, for both cases, the observed performance is defined only by the number of sensors, the data size, and the sampling rate of attached peripherals. Considering the previously discussed use cases, each use case would only require a data storage ranging from $1 \text{ KiB} \cdot 60 \cdot 24 = 1440 \text{ KiB} \approx 2 \text{ MiB}$ per day per sensor for small use cases like location tracking and up to $10 \text{ KiB} \cdot 6 \cdot 60 \cdot 24 = 86400 \text{ KiB} \approx 100 \text{ MiB}$ per day per

sensor for more extensive use cases. With visual monitoring, the data requirements can become arbitrarily large, which would have to be addressed per use case by the gateway.

Overall, gateways can be scaled to serve multiple sensors at a time, leading to latency agnosticism (G5). Thus, we exclude passive networking and storage gateways from our performance evaluation, as their computational requirements are minuscule, and their deployment costs mostly depend on the necessary buffer storage and used communication medium.

5.1.3 Database & Tamperproof Storage

Considering potential performance bottlenecks after the sensing, we ascertain that all relevant components allow for horizontal scalability: First, despite conceptually being a single entity, the per-entity server (cf. Figure 2) can be implemented by arbitrarily many servers to match the required performance capabilities. Second and similarly, each actor can utilize multiple clients for data verification depending on their specific demands. Third, database systems supporting vertical and horizontal scalability are well-established and openly available [5, 54]. Thus, all required components allow for horizontal scalability and do not constitute a performance bottleneck.

Given that all information is persisted in a conventional database, storing fingerprints and signatures of said information is sufficient to ensure tamperproofness. Hence, the tamperproof storage solely serves for information integrity and consistency purposes in our design. Since the tamperproof storage can be implemented as an immutable ledger (cf. Section 6.1), we evaluate the performance of a (private) Quorum blockchain [13] with a proof-of-authority consensus [15]. To this end, we deploy four Quorum nodes on the aforementioned server and prepare transactions on the same server. We derive that fingerprints, along with associated metadata, e.g., shipment or sensor IDs, result in a payload of 124 B. Along with the transaction overhead, including headers and signatures, a single fingerprint per transaction hence requires 267 B. Preparing a transaction (TX) with a single fingerprint—including the calculation of signatures—only takes 20 ms (on average over 1000 runs) on our lightweight server. Hence, we are able to prepare 50 TX/s, which greatly exceeds our needs. In a real-world setting, the corresponding functionality would likely even run on a more powerful (per-entity) server.

Baliga et al. [6] have shown that submitting prepared transactions to a Quorum blockchain can very well result in a throughput of 740 TX/s, which suffices for the use cases presented in Section 3.3: Each shipment only requires a single registration fingerprint along with 2 fingerprints (initialization, finalization) per involved sensor. Irrespective of any time constraints, merging multiple fingerprints in a single transaction, utilizing meta-fingerprints that cover multiple shipment events, sidechains, and sharding are concepts to further scale such a system performance- and storage-wise [4, 5, 76].

This performance evaluation underlines the feasibility and

scalability of our E2E sensing: computationally-wise, corresponding solutions are appropriate for real-world use.

5.2 Hardware Deployment & Cost Estimates

Goods are typically shipped in crates, with several crates fitting into a standard container. Monitoring equipment can be installed in crates or containers, depending on the required granularity of monitoring and the trust relationships between crate owners and shipping companies. To map the use cases and requirements from Table 1, less than five sensors would be required per crate, which can, in most cases, be operated even by a single lightweight Sancus processor. Realistically, such a setup, including basic sensors, can be built for approximately 10 € per crate utilizing off-the-shelf IoT sensors. Slightly more expensive sensors may be necessary if more demanding use cases such as visual monitoring are required. Concerning computational hardware, Sancus processors are not commercially available, while TrustZone-enabled ARM microcontrollers are offered for under 100 €, which would realistically place the cost for the more lightweight Sancus processors in the tens of euros, if they would be commercially deployed. Container and crate equipment would rely on extended storage and processing capabilities in a data center or in the cloud [23]. Cloud-deployed TEE infrastructure, which could serve as the computational bridge, is commercially available at marginally higher prices than today’s commonly used cloud infrastructures [25].

Importantly, existing equipment from today’s supply chain monitoring systems, specifically sensors and often also processing components, may be reused following our approach. That is, the use of TEEs puts no specific requirements on individual sensors, while many more recently purchased microcontrollers may readily support TrustZone functionality [50]. Techniques to seamlessly integrate and reuse sensors in such scenarios are commercially available [3, 75], and research to provide strong security on low-end processors that is orthogonal to TEEs [31] might provide viable solutions to extend the lifespan of existing equipment. For TEE-based equipment, re-deployment and re-attestation of sensing and processing software provide strong guarantees of system integrity even after a potential runtime compromise, thereby further extending the lifespan of a deployed setup. Specifically, regarding equipment costs and operational expenses, our approach thus satisfies G4 (affordability) and keeps costs minimal.

Importantly, deployments need to account for extended shipment periods of several weeks to months, potentially without external connectivity and also without external power supply. For these periods, optional gateway equipment (cf. Section 4.3) can provide the necessary data storage, up to several GiB for each connected sensor. Batteries in a crate or container also require projections and planning of worst-case scenarios, e.g., two to three months for door-to-door shipping between China and Europe. Ultra-low power equipment, such

as Sancus-based sensors, would typically consume less than 1 mA when active and only a few μA in sleep mode¹. Thus, these lightweight sensors can potentially operate for several months on a conventional 2 Ah AA battery cell.

ARM-based sensors and gateway equipment, in particular devices that provide permanent wireless connectivity, might consume several hundred mA and necessitate battery capacities of around 1000 Ah to operate without interruption over a period of two months. Specialized equipment based on TrustZone-enabled low-power processors, such as the ARM Cortex-M23 and specialized low-power WiFi, might reduce this need substantially. Note that our approach introduces very limited additional processing and communication overheads in comparison with less trustworthy monitoring solutions. Therefore, our proposal for TEE-based end-to-end security in supply chains will only marginally increase the power consumption of remote sensing equipment in shipments.

6 Assessing the Concept’s Impact

While Section 5 confirms the general technical applicability and financial feasibility of our design, we now discuss its capabilities and limitations. Particularly, we revisit the design goals (Section 3.5) and assess whether and how our design achieves them regarding the varying requirements arising from different use cases (cf. Section 3.3).

6.1 Trust & Security Discussion

In Section 5, we have already discussed G4 (affordability) and G5 (latency). Even in S3, the most complex scenario, appropriate latencies and data throughput is achievable while maintaining acceptable costs. Thus, we now focus on security guarantees, i.e., dedicated attacks by malicious parties, regardless of their involvement in the sensing and data processing. Instead of technical measures against fundamental attack vectors, we can also require a respective trust relationship between the involved parties (cf. Table 2).

Measurement Manipulation. In our context, data tampering corresponds to the manipulation of sensor measurements, e.g., to cover for issues during a shipment. The utilization of trusted sensors as of S1 with support for remote attestation enables unequivocal detection of such software manipulations, and manipulation of sensor measurements is always detectable in accordance with G1 (tamperproofness) and, at the same time, achieves G2 (authenticity).

Measurement Withholding. Data hiding can be attempted if direct data tampering is not possible, either by withholding or deleting sensor measurements. Since each sensor numbers

its measurements, any withholding is generally apparent as a gap in numberings. Our design further introduces online phases (as of S2) to ensure that (trusted) sensors report the total number of conducted measurements, allowing each party to reliably verify measurement completeness (G3).

Retroactive Data Removal and Manipulation. We specifically need to consider the hiding of complete sensors. In supply chain scenarios where parties cannot trust the database operator, our design S3 prevents retroactive data removal by recording fingerprints of essential sensor information on the tamperproof storage. These fingerprints serve as proof of existence and integrity protection since every deletion or manipulation of a data record from the database is noticeable, thus providing G1 (tamperproofness) and G3 (completeness).

Data Forging and Replaying. Finally, malicious parties could attempt to insert forged information, i.e., submit measurements that originate from unauthentic and unrelated sensors or are made up entirely (data injection). First, as our design involves a PKI for sensors, the origin and authenticity (G2) of submitted data are ensured, while unique and attested measurement IDs per sensor prevent replay attacks, as duplicated IDs would be detectable. Hence, retroactive inclusion of information would require the deletion of sensor registrations for shipments in combination with authenticating the forged data. Second, while the ledger prevents manipulations of existing data records (G1 and G3), the trusted sensors ensure that all signatures originating from the sensor are solely for data originating from that sensor and not from the outside.

Our design fulfills the outlined design goals (Section 3.5) while providing a tunable trade-off between data reliability features (G1-G3), deployment and operational costs (G4), and general performance (G5) based on existing (trust) relationships and use case-specific data reliability requirements.

6.2 Limitations with End-to-End Sensing

While E2E sensing shows potential for large-scale adoption, certain limitations need to be overcome. Moreover, organizations must deal with issues inherent to the mirroring of physical events in digital management systems. In the following, we specifically address physical attacks on deployed sensors, as well as malicious shipment providers that either cheat on their mapping of sensors to shipment or cheat by registering duplicate sets of sensors to hide incidents.

Direct Physical Attacks. (Trusted) sensors in a supply chain are naturally deployed in untrusted environments, thus being prone to physical attacks. Usually, when deploying embedded devices, physical attacks are out of scope, as for Sancus [43] and TrustZone [50]. Likewise, in this paper, we focus on the processing of information once it has been sensed by trusted sensors. Still, to broadly consider all threat vectors, sensing parties should account for such attacks, e.g., by using tamperproof physical isolation of the digital components or by physically hardening their equipment [2, 27]. We see these

¹Data on power consumption stems from commercial TI-MSP430 products [1, 65]. Sancus [43] is based on the openMSP430 core [19] but not a commercially available processor architecture, which, running on FPGA hardware for prototyping purposes, currently has incomparable power consumption characteristics.

mitigations as possible but non-trivial future work, i.e., they can augment our work and extend the security guarantees.

Sensor Registration Forgery. Trust into any reported sensor data can only be achieved if the underlying sensing devices are trustworthy. If an adversary can implant their own devices into the shipment from the beginning of the deployment process, then the existence of the whole chain of trust is in question. As a first step to mitigating sensor-related mistrust, we suggest only utilizing devices that have undergone some process to verify their legitimacy. This process would ideally include some general device certification, but in the absence of such, any verification process that is trusted by all stakeholders suffices. While this mitigation takes the first steps to ensure that no tampered or forged devices can be deployed by any stakeholder, some remaining issues persist: Most prominently, malicious shipment providers can register multiple sensor sets per shipment and ensure that, even if the first sensor set records issues with the shipment, the backup sensor set remains in a healthy (shielded) environment. Before handing off the shipment, a malicious shipment provider could then remove the faulty sensors to only report the shielded sensors' data. Similarly, malicious shipment providers can place sensors in a manipulated environment that differs from the intended one, i.e., conduct *physical sensor manipulation* [73]. To mitigate both issues, we suggest that future work investigates how to (semantically) verify measurements by sourcing environmental conditions of nearby sensors or different sensor types. Related work on consumer IoT already considers this research angle [7]. A similar approach could be used to also verify sensors in nearby or previous shipments.

Product Registration Forgery. Malicious shipment providers could attempt to cheat by being dishonest in mapping sensors to physically shipped products [73]. In this case, the shipment provider could pretend to provide sensor data to a shipment but then only attach the sensors to the shipping when the shipment arrives at the next destination. Here, we envision different mitigation strategies. First, companies could utilize sensors that make it exceedingly difficult to replace them unnoticeably. Depending on the type of shipment, this strategy might be costly due to the need for specialized equipment. Second, the linking of physical products to digital data could be strengthened through various means [24, 51, 70]. Modern marking approaches like molecular fingerprinting [10, 64] exemplify how to uniquely identify shipments. Embedding this unique identifier into the sensor metadata then allows actors to verify the binding between sensors and shipped products. We leave the corresponding feasibility and affordability studies to future work.

While these discussed aspects are relevant to keep in mind when deploying our proposed E2E sensing, they are also apparent in traditional sensing infrastructures, i.e., they are not specific to our work. Moreover, we argue that any real-world deployment would inherently require human decision-makers in the loop to make judgments and decisions based on

the output of the technical domain. This addition is not only necessary to handle any misbehavior but also to accommodate potential technical failures of the E2E-secured sensing.

6.3 Related Work

To the best of our knowledge, except for the work by Pennekamp et al. [46] (cf. Section 1), related work has not proposed comparable approaches to secure the processing of sensed information. Our research lies at the intersection of three topics: *supply chains*, *trusted computing*, and *blockchain technology*. Next, we discuss their respective influences and shortcomings concerning our proposed work.

Supply Chains. Business-oriented research in the context of supply chains and their management considers a multitude of research directions [11, 17, 61]. The closest overlap with computer science concerns the processing of data. Most prominently, related work puts great emphasis on tracking and tracing in supply chains [5, 20], with only a few approaches proposing end-to-end-secured information flows [41]. Yet, to date, barely any work considers the tamperproof and authentic sensing of data (G1 & G2) [46]. Simply deploying off-the-shelf IoT sensors is insufficient, especially in settings with business-oriented, distrusting stakeholders who might even have a (monetary) incentive to cheat. Thus, we require concepts to securely and reliably sense data in supply chains.

Trusted Computing. Generally, a wide range of trusted execution environments, both as research projects and as commercial products, is available [30, 59]. For trusted sensors, we identify Sancus [43] and TrustZone [50] as suitable TEEs, as both can be deployed for low costs (G4) while achieving the discussed performance requirements (G5). Other embedded security architectures [28, 44] may similarly suffice as long as they provide isolation and attestation primitives. Similarly, the utilized attestation protocol and implementation are equally interchangeable, and respective related work on improving remote attestation can be integrated [62].

Blockchain Technology. Supply chains are a prominent application area for blockchains: Related work frequently utilizes blockchain technology, most commonly when enabling or improving tracking and tracing of products or allowing for (origin) certification of products as well as counterfeit identification [5, 20, 32, 33]. While these approaches deal with the tamperproof storing and long-term availability, i.e., completeness of data, they are susceptible to “garbage-in, garbage-out” [52], as they *unreasonably assume* complete, authentic, and untampered (sensor) data as input. These shortcomings also hold for commercial products, such as the recently discontinued TradeLens [67], upkeep [69], project44 [53], or roambee [55]. Generally, they largely focus on ensuring interoperability and quality-of-operation improvements for businesses [26] while neglecting the reliable and authentic foundation of sensed data, which they indirectly build on.

Thus, we are convinced that a full-fledged end-to-end pro-

cessing design is both missing and essential. Without a specific focus on supply chains, other research intersecting the digital and physical world looks into the combination of IoT devices and blockchain technology without deploying specifically secured devices or sensors [16, 49]. As such, their work is complementary, and our design with trusted sensors can potentially benefit from corresponding advances, primarily through affordability and latency improvements (G4 & G5).

Conceptual Intersections. Several projects combine trusted computing and blockchain technology. For example, Microsoft CCF [56] realizes a replicated ledger inside a TEE. Other approaches explore TEEs for blockchain applications [12, 35, 68], while first projects rely on TEEs in real-world blockchain deployments [39]. Prior work [46, 73] already raised concerns about insecure sensors and unreliable data processing for supply chains and blockchains. To the best of our knowledge, a usable design utilizing specially secured sensors, i.e., trusted computing hardware, is still missing.

To conclude, we augment prior work by jointly applying trusted sensors, TEEs, and blockchain technology for practical use in supply chains to provide technical guarantees. Thereby, we reshape today’s trust boundaries in supply chains.

6.4 Universality of our Sensing Concepts

Even though the primary focus of our work is on the use of E2E sensing in supply chains, and we motivated our work accordingly, our design and the corresponding discussions are also relevant to other application areas. In particular, we can easily translate the foundation of our work to settings where mutually distrustful parties sense information in (remote) environments. More specifically, suitable areas are the application of shared inventory management, rental or parking services, digitized construction sites, and smart manufacturing. Especially the latter demands accurate processing of usage and state information due to the emergence of digital factories, where manufacturing equipment and raw material are shared between companies. Depending on the exact industry, accurate monitoring of tool wear is crucial to avoid significant damages (and, in turn, costs) to the production line. Consequently, we argue that our work and the presented findings have the potential to also impact applications that exceed the “simple” tracking and monitoring in supply chains.

7 Conclusion

The growing complexity in supply chains comes with a need for extensive monitoring of goods and shipments and an increase in involved stakeholders, not all of whom trust each other. Hence, for such distributed settings, approaches for secure and reliable communication and sensing are needed. To alleviate this situation, we presented and discussed four designs that expand on the existing situation of trust in monitoring equipment and data, and that increasingly take adversarial

stakeholders into account. By utilizing trusted hardware at the sensor side and securely communicating the sensed and attested data into a database, which is backed by a tamperproof storage, our design achieves a high level of data integrity at minimal costs with real-world applicability. Our final design is, to a large extent, capable of reusing sensing equipment and builds on trusted execution primitives to establish and maintain trust over such equipment for an extended life span.

Overall, we show that realizing E2E-secured sensing is feasible and move the trust in processed data to the edge of the sensing while establishing trustworthy long-term availability of sensed data. These guarantees both apply to scenarios where shipments maintain mobile connectivity and transmit data continuously as well as to shipments that remain offline for large periods of the transit process, as is common in international freight handling. We provide an implementation and a comprehensive evaluation of the E2E-secured scenario, highlighting the feasibility, performance, and scalability aspects of all essential components. Our work fills a gap in literature and contributes a detailed discussion of design choices for trustworthy and robust supply-chain sensing that utilizes authenticated communication, a performance evaluation of crucial components necessary for real-world deployments, and a discussion of relevant pitfalls and challenges.

Acknowledgement

Paper under double-blind review.

Artifacts

The artifacts are currently under evaluation by the artifacts evaluation board.

References

- [1] Zack Albus, Adrian Valenzuela, and Marc Buccini. Ultra-low Power Comparison: MSP430F2x MCUs vs. Microchip XLP Tech Brief. White Paper SLAY015A, Texas Instruments, 2009.
- [2] Ross Anderson and Markus Kuhn. Low Cost Attacks on Tamper Resistant Devices. In *Proceedings of the 5th International Workshop on Security Protocols (Security Protocols '97)*, volume 1361, pages 125–136. Springer, 1997.
- [3] Emekcan Aras, Stéphane Delbruel, Fan Yang, Wouter Joosen, and Danny Hughes. A Low-Power Hardware Platform for Smart Environment as a Call for More Flexibility and Re-Usability. In *Proceedings of the 2019 International Conference on Embedded Wireless Systems and Networks (EWSN '19)*, pages 194–205. Junction Publishing, 2019.

- [4] Adam Back, Matt Corallo, Luke Dashjr, Mark Friedenbach, Gregory Maxwell, Andrew Miller, Andrew Poelstra, Jorge Timón, and Pieter Wuille. Enabling Blockchain Innovations with Pegged Sidechains. Technical report, Blockstream, 2014.
- [5] Lennart Bader, Jan Pennekamp, Roman Matzutt, David Hedderich, Markus Kowalski, Volker Lücken, and Klaus Wehrle. Blockchain-Based Privacy Preservation for Supply Chains Supporting Lightweight Multi-Hop Information Accountability. *Information Processing & Management*, 58(3), 2021.
- [6] Arati Baliga, I. Subhod, Pandurang Kamat, and Sidhartha Chatterjee. Performance Evaluation of the Quorum Blockchain Platform. arXiv:1809.03421, 2018.
- [7] Simon Birnbach, Simon Eberz, and Ivan Martinovic. Haunted House: Physical Smart Home Event Verification in the Presence of Compromised Sensors. *ACM Transactions on Internet of Things*, 3(3), 2022.
- [8] Philipp Brauner, Manuela Dalibor, Matthias Jarke, Ike Kunze, István Koren, Gerhard Lakemeyer, Martin Liebenberg, Judith Michael, Jan Pennekamp, Christoph Quix, Bernhard Rumpe, Wil van der Aalst, Klaus Wehrle, Andreas Wortmann, and Martina Ziefle. A Computer Science Perspective on Digital Transformation in Production. *ACM Transactions on Internet of Things*, 3(2), 2022.
- [9] Jaewook Byun, Sungpil Woo, and Daeyoung Kim. Efficient and privacy-enhanced object traceability based on unified and linked EPCIS events. *Computers in Industry*, 89:35–49, 2017.
- [10] Alice Capecchi, Daniel Probst, and Jean-Louis Reymond. One molecular fingerprint to rule them all: drugs, biomolecules, and the metabolome. *Journal of Cheminformatics*, 12(1), 2020.
- [11] Atanu Chaudhuri, Iskra Dukovska-Popovska, Nachappan Subramanian, Hing Kai Chan, and Ruibin Bai. Decision-making in cold chain logistics using data analytics: a literature review. *The International Journal of Logistics Management*, 29(3):839–861, 2018.
- [12] Raymond Cheng, Fan Zhang, Jernej Kos, Warren He, Nicholas Hynes, Noah Johnson, Ari Juels, Andrew Miller, and Dawn Song. Ekiden: A Platform for Confidentiality-Preserving, Trustworthy, and Performant Smart Contracts. In *Proceedings of the 2019 IEEE European Symposium on Security and Privacy (EuroS&P '19)*, pages 185–200. IEEE, 2019.
- [13] CONSENSYS. ConsenSys Quorum. <https://consensys.net/quorum/>, 2020.
- [14] Victor Costan and Srinivas Devadas. Intel SGX Explained. Cryptology ePrint Archive 2016/086, 2016.
- [15] Stefano De Angelis, Leonardo Aniello, Federico Lombardi, Andrea Margheri, and Vladimiro Sassone. PBFT vs Proof-of-Authority: Applying the CAP Theorem to Permissioned Blockchain. In *Proceedings of the 2nd Italian Conference on Cyber Security (ITASEC '18)*, volume 2058. CEUR Workshop Proceedings, 2017.
- [16] Volkan Dedeoglu, Raja Jurdak, Guntur D. Putra, Ali Dorri, and Salil S. Kanhere. A Trust Architecture for Blockchain in IoT. In *Proceedings of the 16th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQ-uitous '19)*, pages 190–199. ACM, 2019.
- [17] Hajar Fatorachian and Hadi Kazemi. Impact of Industry 4.0 on supply chain performance. *Production Planning & Control*, 32(1):63–81, 2021.
- [18] Food and Drug Administration. Standardization of Data and Documentation Practices for Product Tracing Guidance for Industry. Fda-2018-d-0688, Food and Drug Administration, 2018.
- [19] Olivier Girard. openMSP430 – a synthesizable 16bit microcontroller core written in Verilog. <https://opencores.org/projects/openmsp430>, 2009.
- [20] Peter Gonczol, Panagiota Katsikouli, Lasse Herskind, and Nicola Dragoni. Blockchain Implementations and Use Cases for Supply Chains-A Survey. *IEEE Access*, 8:11856–11871, 2020.
- [21] GS1 AISB. EPC Information Services (EPCIS) Standard. Technical Report Release 1.2, GS1, 2016.
- [22] GS1 AISBL. EPCIS Standard. Technical Report Release 2.0, Community Review Draft, GS1, 2021.
- [23] Martin Henze. The Quest for Secure and Privacy-preserving Cloud-based Industrial Cooperation. In *Proceedings of the 2020 IEEE Conference on Communications and Network Security (CNS '20)*. IEEE, 2020. Proceedings of the 6th International Workshop on Security and Privacy in the Cloud (SPC '20).
- [24] IBM Research. Changing the way the world works: IBM Research's "5 in 5". <https://research.ibm.com/blog/ibm-research-5-in-5-2018>, 2019 (accessed December 1, 2022).
- [25] Intel Corporation. Intel SGX Product Offerings. <https://www.intel.com/content/www/us/en/architecture-and-technology/sgx-product-offerings.html>, 2022 (accessed December 1, 2022).

- [26] Thomas Jensen, Jonas Hedman, and Stefan Henningsson. How TradeLens Delivers Business Value With Blockchain Technology. *MIS Quarterly Executive*, 18(4), 2019.
- [27] Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. In *Proceedings of the 19th Annual International Cryptology Conference (CRYPTO '99)*, volume 1666, pages 388–397. Springer, 1999.
- [28] Patrick Koeberl, Steffen Schulz, Ahmad-Reza Sadeghi, and Vijay Varadharajan. TrustLite: A Security Architecture for Tiny Embedded Devices. In *Proceedings of the 9th European Conference on Computer Systems (EuroSys '14)*. ACM, 2014.
- [29] Mario Lezoche, Jorge E. Hernandez, Maria del Mar Eva Alemany Díaz, Hervé Panetto, and Janusz Kacprzyk. Agri-food 4.0: A survey of the supply chains and technologies for the future agriculture. *Computers in Industry*, 117, 2020.
- [30] Pieter Maene, Johannes Götzfried, Ruan De Clercq, Tilo Müller, Felix Freiling, and Ingrid Verbauwhede. Hardware-Based Trusted Computing Architectures for Isolation and Attestation. *IEEE Transactions on Computers*, 67(3), 2017.
- [31] Jef Maerien, Sam Michiels, Danny Hughes, Christophe Huygens, and Wouter Joosen. SecLooCI: A comprehensive security middleware architecture for shared wireless sensor networks. *Ad Hoc Networks*, 25(Part A):141–169, 2015.
- [32] Sidra Malik, Volkan Dedeoglu, Salil S. Kanhere, and Raja Jurdak. TrustChain: Trust Management in Blockchain and IoT supported Supply Chains. In *Proceedings of the 2019 IEEE International Conference on Blockchain (Blockchain '19)*, pages 184–193. IEEE, 2019.
- [33] Sidra Malik, Salil S. Kanhere, and Raja Jurdak. ProductChain: Scalable Blockchain Framework to Support Provenance in Supply Chains. In *Proceedings of the 2018 IEEE 17th International Symposium on Network Computing and Applications (NCA '18)*. IEEE, 2018.
- [34] Leonardo Marques. Sustainable supply network management: A systematic literature review from a knowledge perspective. *International Journal of Productivity and Performance Management*, 68(6):1164–1190, 2019.
- [35] Sinisa Matetic, Karl Wüst, Moritz Schneider, Kari Kostianen, Ghassan Karame, and Srdjan Capkun. BITE: Bitcoin Lightweight Client Privacy using Trusted Execution. In *Proceedings of the 28th USENIX Security Symposium (SEC '19)*, pages 783–800. USENIX Association, 2019.
- [36] Roman Matzutt, Benedikt Kalde, Jan Pennekamp, Arthur Drichel, Martin Henze, and Klaus Wehrle. Coin-Prune: Shrinking Bitcoin’s Blockchain Retrospectively. *IEEE Transactions on Network and Service Management*, 18(3):3064–3078, 2021.
- [37] Roger Maull, Phil Godsiff, Catherine Mulligan, Alan Brown, and Beth Kewell. Distributed ledger technology: Applications and implications. *Strategic Change*, 26(5):481–489, 2017.
- [38] Ralph C. Merkle. A Digital Signature Based on a Conventional Encryption Function. In *Proceedings of the 7th Conference on the Theory and Applications of Cryptographic Techniques (CRYPTO '87)*, volume 293, pages 369–378. Springer, 1987.
- [39] MobileCoin. MobileCoin - Safe & easy payments at light-speed. <https://mobilecoin.com/>, 2017.
- [40] Teboho M. Mofokeng and Richard Chinomona. Supply chain partnership, supply chain collaboration and supply chain integration as the antecedents of supply chain performance. *South African Journal of Business Management*, 50(1), 2019.
- [41] Aintzane Mosteiro-Sanchez, Marc Barcelo, Jasone Astorga, and Aitor Urbietia. End to End Secure Data Exchange in Value Chains with Dynamic Policy Updates. arXiv:2201.06335, 2022.
- [42] Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press, 2016.
- [43] Job Noorman, Jo Van Bulck, Jan Tobias Mühlberg, Frank Piessens, Pieter Maene, Bart Preneel, Ingrid Verbauwhede, Johannes Götzfried, Tilo Müller, and Felix Freiling. Sancus 2.0: A Low-Cost Security Architecture for IoT Devices. *ACM Transactions on Privacy and Security*, 20(3), 2017.
- [44] Ivan De Oliveira Nunes, Karim Eldefrawy, Norrathep Rattanavipanon, Michael Steiner, and Gene Tsudik. VRASED: A Verified Hardware/Software Co-Design for Remote Attestation. In *Proceedings of the 28th USENIX Security Symposium (SEC '19)*, pages 1429–1446. USENIX Association, 2019.
- [45] Michael Paik, Jerónimo Irazábal, Dennis Zimmer, Michele Meloni, and Valentin Padurean. immudb: A Lightweight, Performant Immutable Database. Technical report, CodeNotary, 2020.

- [46] Jan Pennekamp, Fritz Alder, Roman Matzutt, Jan Tobias Mühlberg, Frank Piessens, and Klaus 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. Proceedings of the 5th International Workshop on Cyber-Physical Systems Security (CPS-Sec '20).
- [47] Jan Pennekamp, Martin Henze, Simo Schmidt, Philipp Niemiets, Marcel Fey, Daniel Trauth, Thomas Bergs, Christian Brecher, and Klaus Wehrle. Dataflow Challenges in an *Internet* of Production: A Security & Privacy Perspective. In *Proceedings of the ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC '19)*, pages 27–38. ACM, 2019.
- [48] Lukas Pfahl and Claire Moxham. Achieving sustained competitive advantage by integrating ECR, RFID and visibility in retail supply chains: a conceptual framework. *Production Planning & Control*, 25(7):548–571, 2014.
- [49] Miguel Pincheira and Massimo Vecchio. Towards Trusted Data on Decentralized IoT Applications: Integrating Blockchain in Constrained Devices. In *Proceedings of the 2020 IEEE International Conference on Communications Workshops (ICC Workshops '20)*. IEEE, 2020. Proceedings of the 1st Workshop on Blockchain for IoT and Cyber-Physical Systems (BIoTCPS '20).
- [50] Sandro Pinto and Nuno Santos. Demystifying Arm TrustZone: A Comprehensive Survey. *ACM Computing Surveys*, 51(6), 2019.
- [51] Stephen Pollard, Guy Adams, Faisal Azhar, and Fraser Dickin. Authentication of 3D Printed Parts using 3D Physical Signatures. In *Proceedings of the NIP & Digital Fabrication Conference, Printing for Fabrication 2018*, pages 196–201. Society for Imaging Science and Technology, 2018.
- [52] Warwick Powell, Marcus Foth, Shoufeng Cao, and Valéri Natanelov. Garbage in garbage out: The precarious link between IoT and blockchain in food supply chains. *Journal of Industrial Information Integration*, 25, 2022.
- [53] project44. Advanced Supply Chain Visibility. <https://www.project44.com/>, 2014.
- [54] Sharvari Rautmare and Deepashree M. Bhalerao. MySQL and NoSQL database comparison for IoT application. In *Proceedings of the 2016 IEEE International Conference on Advances in Computer Applications (ICACA '16)*, pages 235–238. IEEE, 2016.
- [55] Roambee Corporation. Better Real-Time Supply Chain Visibility & Intelligence. <https://www.roambee.com/>, 2014.
- [56] Mark Russinovich, Edward Ashton, Christine Avanesians, Miguel Castro, Amaury Chamayou, Sylvan Clebsch, Manuel Costa, Cédric Fournet, Matthew Kerner, Sid Krishna, Julien Maffre, Thomas Moscibroda, Kartik Nayak, Olga Ohrimenko, Felix Schuster, Roy Schuster, Alex Shamis, Olga Vrousou, and Christoph M. Wintersteiger. CCF: A Framework for Building Confidential Verifiable Replicated Services. Technical report, Microsoft, 2019.
- [57] Sancus – GitHub. sancus-core. <https://github.com/sancus-tee/sancus-core/releases/tag/v2.1.0>, 2022.
- [58] Lakshmi Siva Sankar, M. Sindhu, and M. Sethumadhavan. Survey of consensus protocols on blockchain applications. In *Proceedings of the 2017 4th International Conference on Advanced Computing and Communication Systems (ICACCS '17)*. IEEE, 2017.
- [59] Moritz Schneider, Ramya Jayaram Masti, Shweta Shinde, Srdjan Capkun, and Ronald Perez. SoK: Hardware-supported Trusted Execution Environments. arXiv:2205.12742, 2022.
- [60] Gianluca Scopelliti, Sepideh Pouyanrad, Job Noorman, Fritz Alder, Christoph Baumann, Frank Piessens, and Jan Tobias Mühlberg. End-to-End Security for Distributed Event-Driven Enclave Applications on Heterogeneous TEEs. *ACM Transactions on Privacy and Security*, 2023.
- [61] Ashutosh Sheel and Vishnu Nath. Effect of blockchain technology adoption on supply chain adaptability, agility, alignment and performance. *Management Research Review*, 42(12):1353–1374, 2019.
- [62] Carlton Shepherd, Raja Naem Akram, and Konstantinos Markantonakis. Establishing Mutually Trusted Channels for Remote Sensing Devices with Trusted Execution Environments. In *Proceedings of the 12th International Conference on Availability, Reliability and Security (ARES '17)*. ACM, 2017.
- [63] Rajesh Kumar Singh, Pravin Kumar, and Mahesh Chand. Evaluation of supply chain coordination index in context to Industry 4.0 environment. *Benchmarking: An International Journal*, 28(5):1622–1637, 2021.
- [64] Anatoly P. Sobolev, Simone Circi, Donatella Capitani, Cinzia Ingallina, and Luisa Mannina. Molecular fingerprinting of food authenticity. *Current Opinion in Food Science*, 16:59–66, 2017.

- [65] Texas Instruments. MSP430 Ultra-low-power Micro-controllers. Product Bulletin SLAB034A, Texas Instruments, 1999.
- [66] Feng Tian. An agri-food supply chain traceability system for China based on RFID & blockchain technology. In *Proceedings of the 2016 13th International Conference on Service Systems and Service Management (ICSSSM '16)*. IEEE, 2016.
- [67] TradeLens. TradeLens. <https://www.tradelens.com/>, 2018.
- [68] Muoi Tran, Loi Luu, Min Suk Kang, Iddo Bentov, and Prateek Saxena. Obscuro: A Bitcoin Mixer using Trusted Execution Environments. In *Proceedings of the 34th Annual Computer Security Applications Conference (ACSAC '18)*, pages 692–701. ACM, 2018.
- [69] UpKeep. CMMS, EAM & IIoT Software by UpKeep Asset Operations Management. <https://www.upkeep.com/>, 2014.
- [70] Diana M. Segura Velandia, Navjot Kaur, William G. Whittow, Paul P. Conway, and Andrew A. West. Towards industrial internet of things: Crankshaft monitoring, traceability and tracking using RFID. *Robotics and Computer-Integrated Manufacturing*, 41:66–77, 2016.
- [71] VERAISON. Project VERAISON: VERificAtIon of atteStation. <https://github.com/veraison>, 2021.
- [72] Dong-Liang Wu, Wing W. Y. Ng, Daniel S. Yeung, and Hai-Lan Ding. A brief survey on current RFID applications. In *Proc. of the 2009 International Conference on Machine Learning and Cybernetics (ICMLC '09)*, volume 4, pages 2330–2335. IEEE, 2009.
- [73] Karl Wüst and Arthur Gervais. Do you need a Blockchain? In *Proc. of the 2018 Crypto Valley Conf. on Blockchain Technology (CVCBT '18)*, pages 45–54. IEEE, 2018.
- [74] XESS Corporation – GitHub. StickIt! Main Motherboard. <https://github.com/xesscorp/StickIt-MB>, 2015.
- [75] Fan Yang, Nelson Matthys, Rafael Bachiller, Sam Michiels, Wouter Joosen, and Danny Hughes. μ PnP: Plug and Play Peripherals for the Internet of Things. In *Proceedings of the 10th European Conference on Computer Systems (EuroSys '15)*. ACM, 2015.
- [76] Mahdi Zamani, Mahnush Movahedi, and Mariana Raykova. RapidChain: Scaling Blockchain via Full Sharding. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*, pages 931–948. ACM, 2018.
- [77] Zibin Zheng, Shaoan Xie, Hong-Ning Dai, Xiangping Chen, and Huaimin Wang. Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services*, 14(4):352–375, 2018.