

Blockchain-Assisted Access for Federated Smart Grid Domains: Coupling and Features

Cristina Alcaraz^{1,2}, Juan E. Rubio¹, and Javier Lopez^{1,2}

¹Department of Computer Science, University of Malaga,
Campus de Teatinos s/n, 29071, Malaga, Spain

²ITIS Software, Ada Byron Research Building,
C/ Arquitecto Francisco Peñalosa 18, 29071, Malaga, Spain
{alcaraz,rubio,jlm}@lcc.uma.es

Abstract

Industry 4.0 technological expansion and the multiple accesses to the diverse Smart Grid domains (power networks, control systems, market, customer premises) entail the need to provide efficient interconnection mechanisms with connection from anywhere, at any time and in anyhow. However, this type of requirement should not only consist in imposing interoperability solutions between entities and domains, but also in searching the way to justify and trace connections (how, when, where, who) for future governance or auditing actions. This paper, therefore, presents a three layer-based interconnection architecture and several interconnection strategies, all of them adapting the traditional policy decision and enforcement approaches together with the blockchain technology to manage reliable and secure connections among entities, processes and critical resources. With this architecture in mind, the paper also analyzes the coupling level of the blockchain technology, and explores which interconnection strategy is more suitable for Smart Grid domains and their control systems.

Keywords: Smart Grid, Blockchain, Access Control, Technological Coupling, Industrial Internet of Things, Cyber-Physical Systems

1 Introduction

We are increasingly witnessing how the Smart Grid (SG) domains are adopting the new technologies to adapt the new industrial philosophy equivalent to Industry 4.0 [1]. The goal is now to (i) optimize and automate operational processes and (ii) efficiently produce and distribute energy according to real demand. The result is a complex environment based on multiple application domains composed of diverse stakeholders (power grids, control systems, providers, customers and market), as also outlined by the National Institute of Standards and Technology (NIST) in [2]. The interactions between stakeholders can be protected incorporating diverse security mechanisms [3]

and access control measures [4], and more even when the environment tends to be federated.

So far, the vast majority of these approaches follow the traditional Policy Decision Points (PDP) and Policy Enforcement Points (PEP) schemes [5, 6, 7, 8], originally introduced by the Internet Engineering Task Force (IETF) in [9]. But in this paper, we particularly explore the way of extending [5, 6] by bringing SG domains to the Industry 4.0 competitive advantages. Both works already anticipate the transition of the new Information Technologies (IT) towards the Operational Technologies (OT), while entities require gaining access to the diverse critical resources of the system such as: CPS (Cyber-Physical System) and/or IIoT (Industrial Internet of Thing) devices to verify states or lead Command and Control (C&C) actions in the field. As specified in [6, 5], access can be supported by IT technologies, in which different PEP interfaces connect to decentralized policy decision points to prove authentication and authorization tokens.

Figure 1 depicts with greyish background the two interconnection architectures established for access control in [5, 6], such that: One is based on PDP proxies [6] and the another one on a two-layer structure composed of cloud and fog PDP servers [5]. These two architectures lie the base of this research where the purpose is now to: (i) Expand the interconnection analysis with new models; and (ii) incorporate new security measures that guarantee a trustworthy governance in the entire consortium making use of the Distributed Ledger-based Technology (DLT). Through this technology, transparency and access traceability can be possible so that federated entities cannot only know the connections established in their respective resources but also the sequences of actions and responsibilities taken in such resources.

These decentralized solutions offer some benefits against the centralized alternatives for PDP-interconnected domains, since the latter can impose strong coordination restrictions between the many stakeholders involved. This includes issues like data ownership in cloud-based environments (i.e., delegating shared data in untrusted providers), their inability to cope with load balancing problems (i.e., bottlenecks) and the evident fact that they represent a single point of failure [10]. In turn, by integrating DLT-based solutions such as the blockchain, we ensure a horizontal integration of all entities under the same data management policy, which implies data replication with immutability warranty within a fault tolerance network. This is why we focus on these solutions for devising future access control measures.

Because of such data replication and immutability, enhanced traceability features are enabled for a distributed access control mechanism. These measures give the underlying system an attractive way to evaluate and quantify actions granted, and determine the real fulfillment of access control policies. This information can even feed the authorization components to improve the actions taken by their models (e.g., RBAC (Role-Based Access Control), PBAC (Policy-Based Access Control), ABAC (Attribute-Based Access Control) [4]). In this sense, authorization components would not only rely on the type of user, their roles or the characteristics of the context (as already detailed in [5, 6]), but also on “experience and past occurrences”. For that reason, this paper mainly focuses on the blockchain technology as represented in Figure 1. This type of technology allows the system to check and register large past evidence without entailing a centralized management and guaranteeing immutability and trans-

parency in the process [11, 12, 13, 14].




Figure 1. IT-based interconnection models for federated SG

Some related works have already shown the viability of the blockchain technology in the SG field [15] such as: Resilience through smart contracts [16]; grid management [17]; billing [18]; energy trading and supply [19, 20]; and sharing of resources [18]. But authorization-related works for the remote protection of SG resources are still insufficient, and especially when the environments are more and more federated. Alcarria *et al.* in [21] focus on a blockchain-based authorization system to facilitate access to consumer information and resource trading in communities; Zhou *et al.* present in [22] a cloud-assisted centralized authorization system with support in the blockchain technology, proposing a custom consensus algorithm; and Suciu *et al.* present in [23] a conceptual architecture for the integration of Attribute-based Access Control scheme to the SG with the use of Blockchain. But even so, these initial advancements do not evaluate whether the adaptation of blockchain technology can be really feasible for any SG network topology. To the best of authors' knowledge, this is the first paper that evaluates the integration of these elements with traditional authorization approaches and their impact on efficiency and control requirements. For these reasons, the contributions of this paper can be summarized as follows:

- Definition of a three-layer architecture that eases the accommodation of Blockchain technologies with the access models of an authorization framework in a SG scenario.
- Studio and discussion to assess whether traditional PDP and PEP-based authentication mechanisms can be assisted by a blockchain network [14] and hence fulfill different coupling conditions with respect to data provenance and traceability procedures.
- Concise analysis to qualitatively determine which of the Blockchain coupling models are effective for the interconnection of federated SG domains with tendency to IT-OT environments, in terms of efficiency and control requirements.

The remainder of the paper is organized as follows: Section 2 introduces the interconnection models for federated environments, composed of multiple actors, processes and technologies. Section 3 introduces a three layer-based interconnection architecture composed of three main data provenance phases following the ProvChain model presented in [24], to later establish in Section 4 the interconnection and coupling conditions for IT-OT networks. Section 5 widely discusses the real applicability of the blockchain technology in the different interconnection schemes, and Section 6 highlights the future challenges for the fourth control industry in this context.

2 Interconnection models for federated SG scenarios

Considering the traditional interconnection systems [25], and the PEP/PDP-based authorization frameworks for energy Industry 4.0 [5, 6], it is now possible to classify access models according to the type of access and the capacity to manage such an access. These models are also illustrated in Figure 1 and detailed as follows:

Centralized access: Interconnection is led by centralized powerful devices working at GHz (with two or more microprocessors) with support for large data warehouses and with the capacity to establish authorized connections. Within this category, we stress, among others, the role of dedicated PDP proxies or cloud-computing PDP servers by offering storage and processing services of large data volume, as well as interoperability services [26].

Decentralized access: Authorization services are spread out in multiple decentralized PDP nodes [27], with the technical capacity to collect and process data from a SG domain/area and share the information with other domains of the federated network. We understand by technical capacity those PDP devices working in terms of GHz with enough storage resources (e.g., decentralized fog servers deployed one per domain [28]) or those PDPs with limited but with sufficient capacities to process large data volume without interfering in the operational tasks.

Hybrid access: Authorization services based on hierarchical architectures capable of decentralizing access per domain/area while large access-related data volume can be concentrated on one or several PDP devices within the hierarchy. Local data

can be managed by decentralized PDP nodes, such as fog servers, and global data can be computed by centralized nodes in the cloud to understand general states. This way of managing data certainly helps the system provide a better understanding of the health state of a context (a domain, several domains or the entire system), and determine access according to such a context. [5].

Here we discard distributed access, mainly because the vast majority of CPS or IIoT devices present significant HW/SW constraints ($\sim 13\text{MHz}$ - 200MHz with 256 bytes-64MB RAM, 8KB- 32MB flash memory and 16KB-256KB EEPROM such as RTU (Remote Terminal Units), ~ 4 - 32MHz , 4-512KB RAM, 48-192KB ROM such as industrial sensors, or ~ 8 - 50MHz , 4-32KB RAM and 32-512KB flash memory in the case of smart meters [29]) to compute operational tasks and incorporate minimal protection services for security [30]. Any new security layer may significantly increase complexities that may impact on operational tasks. For example, if access control measures should be considered per node, then mechanisms for event management and accountability should also be contemplated to determine the degree of “responsibility” in the own access and per domain [6, 31]. Through event management, access requests, context states and actions in the field can be recorded to establish decisions accordingly. Their contents could include information about: (i) The kind of access to a specific domain and resource, (ii) who has requested the access and its permissions, (iii) when and from where it has been requested, as well as (iv) the actions carried out in the field. Many of these aspects are well underlined in [32]. The authors carefully analyze how the diverse interconnection modes can impact on the operational performance at substations, concluding that the best access modes should be, for now, via front-end.

So far, we have explored different interconnection architectures in these environments, which manage diverse access types and heterogeneous data. Yet, it is also necessary to find out how these architectures impact on the control performance, data treatment and its traceability, as well as the authorization in real time, which are critical for SG scenarios. For this reason, the remainder of this paper focuses on analyzing the coupling level of the new access modes assisted by current cutting-edge technologies such as the blockchain technology. To do this, two *Coupling Conditions* (CC) are established:

- **CC1:** Where to store access registers to later contrast them with security policies; and, in the worst case, to allow the system in a near future to adjust parameters, review policies or establish responsibilities. This condition implicitly leads to consider how to store access registers to make sure data replicability in a critical and shared infrastructure without impacting on the underlying system.
- **CC2:** Where to configure the traceability services and related security services without colliding on the operational performance.

The following section expands the interconnection models to include the blockchain, and identifies the main properties of this technology to evaluate its suitability with the access control purpose in SG scenarios and taking into account the control requirements given in [33].

3 Three layers-based interconnection architecture

In order to address the two Coupling Conditions for reliable interoperability in critical contexts and introduce the analysis about the effectiveness of the blockchain technology for federated SG environments, an interconnection architecture is proposed in this section. Basically, this architecture supports three interconnected network layers, capable of adapting the traditional capacities of the PDP technology with the blockchain technology to cover many of the expectations commented above:

Layer 1 (L1): Physical network layer, generally composed of multiple OT devices capable of perceiving states of the system and protecting it according to these states. In this sense, only authorized stakeholders from **L1** may interact with these devices to change states and deviate behavior through specific C&C actions.

Layer 2 (L2): Interconnection layer, composed of dedicated PDP servers capable of authenticating and evaluating any access request to **L1**, and according to a specific interconnection model (centralized, decentralized or hybrid). To do this, a set of factors are considered: ID and role of subject (*subject*), type of demanded resource (*resource*), type of actions to carry out in the resource (*action*) at a certain moment in time, and the type of actions taken by the subject in the past (*pastActions*); such that: $\langle \text{subject} \rangle := \langle \text{subjectID} \rangle \langle \text{role}_{\text{primary}} \rangle \langle \text{role}_{\text{secondary}} \rangle$; $\langle \text{resource} \rangle := \langle \text{resourceID} \rangle \langle \text{domain} \rangle \langle \text{infrastructureOwnerID} \rangle \langle \text{contextState} \rangle$; $\langle \text{action} \rangle := \langle \text{C\&C} \rangle \langle \text{timestamp} \rangle$; and $\langle \text{pastAction} \rangle := \{ (\langle \text{subject} \rangle, \langle \text{resource} \rangle, \langle \text{action} \rangle)^* \}$. Many of these aspects have already been formalized in [6] under restrictive RBAC-ABAC rules to allow or deny access in SG environments, where $\langle \text{contextState} \rangle$ defines the availability level of the assets demanded.

Layer 3 (L3): Distributed ledger layer, in which multiple distributed IIoT devices are in charge of sending, receiving, storing and validating transactions of type $(\langle \text{subject} \rangle, \langle \text{resource} \rangle, \langle \text{action} \rangle)$, which are associated with the granted access in **L2**. As a result, a common but distributed database is created to enable the synchronization of immutable but linkable data between all partners within the federation. This technology aids an authorization system to transparently declare which state of the database is considered as valid over time, allowing to later create security services that help improve the quality of access to **L1** and the governance in the entire federated system (**L2**).

To implement **L3**, multiple blockchain architectures and platforms have already been classified for its application to power grids [15, 11], which are discussed in the following. These have been used in this area as a transparent, tamper-proof and secure system that enables a plethora of business applications, ranging from P2P energy trading in microgrids [34] to record keeping systems with privacy protection [35]. After all, the blockchain consists of a shared and distributed database that contains a continuously expanding log of transactions (i.e., access registers in our case) sorted in chronological order, which are aggregated and stored into larger structures called blocks. These are subsequently signed and cryptographically linked to previous blocks, hence forming

the ‘Blockchain’. When combined with smart contracts (i.e., user-defined programs executed in the ledger), it enables an accurate traceability of events between the different devices and partners, ensuring the veracity of data while also removing the need of intermediaries.

When it comes to data ownership and visibility, a Blockchain can be public (also known as permissionless) or private (permissioned). If we are dealing with a federated SG scenario and distributed ledgers are public, then all parties are granted access to read past transactions. Hence, to preserve confidentiality and privacy, this may require novel ways to protect sensitive information which would make energy consumption and access registers not traceable to individual users. For this reason, only architectures of type “consortium” are considered in this paper, where a set of partners are allowed to collaboratively manage the ledger. These permissioned blockchain schemes oblige those federation partners to be identified and authorized prior to participate in network operations, which reduces dramatically the number of nodes compared to public blockchains [36]. The latter are based on a *Proof of Work* (PoW) consensus between the partners, such as the Nakamoto algorithm (where the consensus depends on demonstrating the resource consumption implied by solving a complex mathematical problem), since there is no previous trust assigned to the rest of peers within the network (that in turn offer a higher peer-to-peer scalability). In contrast to them, permissioned blockchains allow the deployment of more efficient consensus algorithms featuring a higher transaction capacity [37], based on the assumption of an increased trust between their participants for internal business operations.

Examples of widely used consensus algorithms for permissioned blockchains include the *Proof of Stake* protocol (PoS), where the network participants pledge their crypto actives (e.g., their tokens) and wait to be probability selected to add new blocks instead of competing with others, in such a way that validators with large stakes will be chosen more often, achieving a high performance. A related algorithm is *Proof of authority* (PoA), where a pre-selected subset of participants are elected as authorities that put their reputation at stake, so that new blocks are generated when a majority is reached by them. Another semi-centralized approach is also applied by Raft, which assumes the presence of a leader to propose new blocks that are confirmed by the rest of followers within the network, resulting in a faster block time. In Raft, that leader is automatically elected after a period of time if the previous one fails, hence becoming a Crash Fault Tolerant (CFT) consensus engine. Lastly, the Byzantine Fault Tolerant (BFT) variants of consensus are opposed to this principle, where participants in the private network do not assume honesty toward each other. The main difference from CFT algorithms like Raft is that while followers in Raft blindly trust their leader, each block in PBFT requires multiple rounds of voting to arrive at a mutual agreement, which achieves a higher protection as long as the number of peers is kept reasonably low (due to the overhead introduced by messages exchanged across the network).

These consensus algorithms can be implemented and flexibly configured on the software applied to the application context. Among the plethora of platform alternatives for permissioned blockchains that are increasingly used in nowadays business scenarios, we can highlight these three examples, whose main features are summarized in Table 1:

- Hyperledger Fabric: Promoted as a cross-Industry pluggable framework. This infrastructure provides a modular architecture that enables a configurable consensus between the consortium members and the execution of flexible smart contracts [38].
- Quorum: This is an Ethereum-based distributed ledger protocol that especially focuses on providing private transactions and contracts, besides the possibility of integrating alternative consensus mechanisms to ensure a higher performance [39].
- Corda: It is an open-source platform that also enhances privacy while offering fine-grained access control to digital records. Corda also enables the consortium members to transact directly using smart contracts, hence removing the burden of costly business transactions [40].

The provisioning and deployment of these platforms can be conducted on premises or by means of Blockchain-as-a-Service (BaaS) providers to reduce costs and enhance the scalability of resources [41]. Altogether, the particular selection of the consensus algorithm and the implemented platform is influenced by the scalability, security and performance requirements of the network, as well as both functional and non-functional aspects that fall out of the scope of this paper. Regardless of this selection, this way of injecting, filtering, managing and storing authorization assets still adds the need to determine how the interactions between layers ($\mathbf{L1} \rightarrow \mathbf{L2}$ and $\mathbf{L2} \leftrightarrow \mathbf{L3}$), can be effective without impacting on the operational processes. Therefore, it is still required raising the two previous issues: **CC1** and **CC2**.

	Hyperledger Fabric	Quorum	Corda
Governance	Linux Foundation	Ethereum Developers and JP Morgan Chase	R3 Consortium
Initial release	2016	2016	2016
Consensus Algorithm	Pluggable Framework (SOLO, Kafka, Raft)	Majority Voting	Asynchronous Byzantine Fault Tolerance
Programming Language for Smart Contracts	Javascript, Go, Java	Solidity	Kotlin, Java
Advantages	Modularity, private communications	Public and private transactions, high transaction processing speed	Data privacy, performance

Table 1. Examples and main features of some of the most used permissioned blockchain platforms

3.1 CC1: Access data management and provenance for SG scenarios

Following the three interconnection models described in Section 2, we distinguish in this section three main access-related data management modes for **L2** (also depicted in

Figure 2) so as to reduce latency. We cannot forget that any critical-safety system has to be under control and accessible in due course. Therefore:

Centralized data management: Access instances are processed, interpreted and stored by a single central PDP service. This service can be integrated in a dedicated PDP proxy or a cloud-coupled PDP server with the ability to handle multiple interconnection and security services.

Decentralized data management: Decision-making is completely decentralized throughout PDP proxies or fog-coupled PDP servers.

Hybrid data management: It focuses on heterogeneous architectures [42]. Decentralized systems can individually manage access per area, thereby avoiding bottlenecks and single failure points. However, the access information should later be replicated in a centralized node to offer a greater overview of the real state of the entire system, characterizing anomalous events originated at one or more domains.




Figure 2. Access data management models and auditing

Unfortunately, data scalability considerably shoots when transactions of **L3** are copied in PDPs of **L2**, even it could lead throughput penalties when large transaction volumes have to be processed [43]. This impact can be reduced if blockchain-based systems apply lightweight structures with space to store a predefined number of transactions per block or specific techniques to simplify the register in the entire chain as detailed [12]. If so, then it is possible to defend that the computational cost invested in the own access verification process with data extracted from **L3** can be beneficial for protection of critical domains and resources. In this sense, the verification processes can be more effective and suitable for the decision making in the field, and subject to extra information of type: $\{(<\text{subject}>, <\text{resource}>, <\text{action}>)*\}$.

This redundancy level also brings itself the need to manage provenance measures for traceability. Data provenance is defined by Buneman *et al.* in [44] as “the process of tracing and recording the origins of data and its movement between databases”,

and it is translated as *the way to distribute and store large data flows at the diverse PDP points and how these data can be linked each other*. Traditional provenance techniques [45, 46, 12] are mainly based on traditional asymmetric cryptography schemes and probabilistic and coding methods, such as: Encryption approaches, digital signature, message authentication codes, hash functions, privacy schemes (searchable data encryption, proxy encryption), watermarking, arithmetic coding, Bayesian methods, Bloom filters and timing, such as [47, 48]. However, the own nature of blockchain in **L3** already addresses itself this issue by implicitly establishing the property of data lineage within its own chain. Each transaction is part of a block, and each block is part of an immutable and linkable sequence of blocks; implicitly fostering traceability and subsequent security services.

The first related works on blockchain-based data provenance date from 2017, such as [24, 49, 50]. The first two, i.e., [24, 49], are focused on cloud-computing and [50] on preserving privacy by implementing a randomized voting system where any deviation is punished by a monetary penalty using smart contracts. However, and although there is not sufficient related work on this research field linked to data provenance for large federated environments, there is a special attraction and novelty in the use of this technology [51, 52]. As stated in [24, 12], the technology itself presents multiple benefits for traceability and integrity, where lineage is based on sequences of transactions where the identity of the entities remains anonymized by the use of hashes; i.e., $H(\text{subjectID})$.

3.2 CC2: Additional security services for SG scenarios

According to Souali *et al.* in [53], traceability can be defined as “*the ability to keep a detailed history of all activities and changes that a particular object can undergo throughout its entire life cycle, taking into account the different relationships that may appear*”, and under specifications, laws or standards such as the EU General Data Protection Regulation (2016/679) [54]. However, in federated environments, this level of traceability also entails the need to manage complete and immutable traces that help the system guarantee tamper resistance, governance, auditing and accountability. According to ENISA (European Union Agency for Network and Information Security) in [14], auditing processes consists in “*showing a status of a given system or organization in a given point in time*”, which can be extracted from the traceability itself and from the continuous monitoring. Both properties, traceability and auditability, are then ideal criteria to: (i) Eventually show the granted access modes to the different domains and critical elements of a SG; and (ii) explain: *how, when, where* and *who* invoked access. This way of tracing actions also helps to establish accountability, which is defined by the Cambridge dictionary in [55] as “*the fact of being responsible for what you do and able to give a satisfactory reason for it, or the degree to which this happens*”.

In order to establish the precise responsibilities in a SG context in terms of traceability and accountability, we can take into account these definitions and recognized standards like the IEC-62351 [56]. This is a reference framework in industrial networks and power systems, that provides guidelines for introducing different security services concerning data and communications. The standard is composed of eleven parts, where part 8 is especially applied to control access mechanisms. It defines a set

of roles that can be part of this new concept of Industry 4.0, such as operators, administrators, engineers and “auditors”, so that each one is assigned with clear responsibilities in form of specific rights that allow them to read, write, report and configure different assets within the organization. As for auditors, denoted as SECAUD in IEC-62351-8, they are in charge of verifying the well performance of the underlying infrastructure, the applications of which can be extended to PDP nodes [5] to ensure the correct application of the authorization policies with the verification of access registers (stored as transactions in **L3**) at all times. Taking advantage of this, Figure 3 outlines the actions that virtual auditors could contribute in the new architecture. In this case, each auditor would be also responsible of locally capturing information from **L3** to **L2** in order to update the databases of their allocated PDP; and with this to foster reliable decision-making during the access verification processes, as also has been stated in Section 3.1.

To characterize these actions, our architecture is based on the ProvChain model specified in [24]. The model comprises a set of actions categorized into three data provenance phases: (i) Data collection from a determined network (**L2** $\rightarrow$ **L3**), (ii) access verification through the blockchain technology (**L3**), and (iii) transaction download for auditing tasks (**L3** $\rightarrow$ **L2**). These phases correspond with the processes of the approach and data treatment, whereas the layers correspond with the interconnection between networks. Namely:

Phase 1 (P1): Access-related data collection from the interconnection network (**L2**) to be later injected to the blockchain network. Particularly, this phase transforms the access data in metadata containing specific attributes for traceability. Each metadata generated in layer **L2** (see Figure 4) is transferred to **L3** so that this last can create a transaction of class: $\langle \text{transaction} \rangle := \{ \langle \text{subject} \rangle, \langle \text{resource} \rangle, \langle \text{action} \rangle \}$. Basically, this phase corresponds to the data transference from **L2** $\rightarrow$ **L3**.

Phase 2 (P2): Access verification and lineage through the blockchain technology [57, 12]. Due to the critical restrictions of the underlying infrastructure, the blockchain-based system must be based on lightweight platforms working under efficient consensus protocols, as explained in Section 3.

Phase 3 (P3): Databases updates configured in PDP nodes by (i) periodically capturing transactions from **L3** [24] to assess the policy enforcement, and (ii) validating this information to detect possible tamperings. To do this, auditors could calculate the Merkle tree root (based on a set of linked hashing operations [24]) and comparing it with the Merkle tree root value included in each block downloaded (see Figure 3).

This way of connecting **L2** and **L3** makes that PDPs serve as intermediary interfaces among layers, and auditors SECAUD as software inspectors capable of offering a clearer picture of the activities carried out in the field. But to meet the design precondition **CC2**, this type of connection should not, in turn, collide with the interconnection tasks in which actions of SW auditors should then run in background. In this way,




Figure 3. Data provenance phases in our approach: **P1**, **P2** and **P3**

decision-making would be based on historical information handled in parallel under the condition of granting PDPs the right to ignore or interrupt the feedback at any time, thereby preventing any computational penalty. Here, connection times to **L3**, and processing and storage of large transaction volumes in **L2** should carefully be questioned.

4 Interconnection and coupling conditions: IT-OT networks

So far we know that accessibility to OT domains of an SG environment, in which multiple stakeholders can request PEP access, is already possible [6] through the deployment of a three-layer-based architecture (see Section 3). However, this eagerness to accommodate the blockchain technology and show its replicability and traceability, in turn, forces us to analyze whether the technology itself can really coexist in control domains taking into account the two Coupling Conditions: **CC1** and **CC2** (see Section 2). To do this, it is fundamental first to explore if this new technology can confluence in OT environments without infringing on the five control requirements [33]:

Real-time performance: Any interconnection point needs to process multiple and concurrent accesses in optimal times (**CC1**), reducing any implicit *computational*, *storage* and *communication overhead* carried out by other security services for traceability and auditing (**CC2**). In this sense, *responsiveness* of auditors SECAUD to read or write in databases – either in (relational, NoSQL,

triplestore and graph) databases, log files, tables or in any virtual data warehouse – should be composed of straightforward and lightweight procedures that do not interfere on the operational processes. For example, making use of generic ontology languages, such as RDF and OWL [58], simple messages for the exchange and storage, such as XML [59] or metadata structures containing *fine-grained* information based on auditable attributes.

Sustainability: Interconnection systems and their auditors have to continuously be updated to ensure its validity for a long period of time. This requirement is directly linked (i) to *maintainability* property, which is, in turn, related to upgrade HW/SW components, including data warehouses, and (ii) *HW scalability* to allow the inclusion of new components within the system such as new PDP devices with support for large databases. Within scalability it is also important to highlight the data scalability according the storage space, the depth and stages of ancestry established not only for the PDP systems, but also the blockchain [43]. In this regard, the aim is to keep transaction rate per second as high as possible, based on the number of peers and the consensus algorithm implemented in the network.

Dependability: Three-layers-based interconnection systems in trustworthy federated networks should be fault-tolerant to causal or regular internal faults, making sure *availability* and *reliability* of resources and data in all time. The former can be addressable through redundant measures and data reliability through validity mechanisms to establish accuracy and quality conditions. However, as our study is focused on the decision-making for access, this reliability also has to be related to the consistency degree with the reality, using, for example, past experience such as $\langle \text{pastAction} \rangle := \{ \{ \langle \text{subject} \rangle, \langle \text{resource} \rangle, \langle \text{action} \rangle \} \}$.

Survivability: Interconnection systems and their auditors must be able to cope (deliberate or unforeseen) with malicious actions in a timely manner, protecting (i) *confidentiality*, *integrity*, *availability* of access records and (ii) *privacy* of the entities of the entire consortium [45]. This also implies the introduction of privacy-preserving techniques to conduct load monitoring and billing procedures in the SG and the custody of aggregated data from consumers, according to the EU regulation [54]. Cryptography, authentication, non-repudiation and unforgeability schemes, robust information and network infrastructures, prevention and response, trust models, regulatory frameworks and security policies are hence indispensable measures for construction of secure interconnected networks. Apart from this, it is also essential to consider the *traceability* of the data itself. In this regard, attribution (i.e., the data copyright through unique identifications or non-repudiation), simplicity of the data (or metadata) for the provenance, and its linkage in terms of data scalability, are keys to track the lineage and the origin of an event [60] – associated in this case with access information and data scalability. As stated in Section 3.2, through this traceability it is possible to establish other security measures such as auditing and accountability.

Safety-critical: PDPs systems must also be able to safeguard proofs through optimized *replication* approaches (but without adding extra overloads) to address




Figure 4. Technological coupling: **CM-1**, **CM-2** and **CM-3**

those unplanned events that might result in drastic physical damages in databases [33]. However, this level of redundancy also implies *data transparency* in terms of location (where is the data?), so as to determine its real setting in each replication and its use within the federated network [61]: e.g., in a DLT network, everyone maintains a copy of the transaction.

5 Blockchain coupling analysis for federated SG domains

To determine the coupling level of the blockchain network, it is now necessary to analyze whether the inclusion of a third layer into the architecture defined in Section 3 can make sense for future SG environments; or it may be optimized to 2-layers. To do this, three kinds of coupling models are addressed in the remainder of this paper and taking into account: The three interconnection models stated in Section 2, the three network layers (IT-OT) specified in Section 3 together with the three deployment phases of Section 3.2, and the three data management systems of Section 3.1. Particularly, these three Coupling Models (CM) are as follows (see Figure 4):

- **CM-1:** A centralized PDP in **L2** with an uncoupled blockchain network in **L3** where **L2** launches the processes **P1** and **P3**, and **L3** executes **P2**.
- **CM-2:** Decentralized PDPs in **L2** with an uncoupled blockchain network in **L3**, where **P1-P2-P3** are established as **CM-1**.
- **CM-3:** Decentralized PDPs in **L2** based on a coupled blockchain network in **L2** where this latter contemplates the three data provenance phases **P1**, **P2** and **P3**.

As can be noted, our analysis omits the coupling in centralized networks since the own blockchain technology demands decentralization by itself [12], and hence, it is incongruous with the application scenario.

5.1 CM-1: Centralized interconnection and uncoupled blockchain

From an interconnection point of view, the number of connections per domain is expected to grow in Industry 4.0 [26]. This increase implicitly leads to a significant number of transactions to be checked and registered by **L3** as illustrated in Figure 4 (but only considering the existence of one PDP), and to later be downloaded by the only one auditor SECAUD in **L2** [12]. This feature may be critical for centralized systems. Its susceptibility to generate bottlenecks when multiple devices simultaneously interact, may seriously cause a deterioration on the operational performance or a denial of service during access to **L1** in the worst scenario. This situation may even isolate the decision point, hampering the new access requests or the actions of the auditor SECAUD to update the data warehouses in **L2**.

Indeed, the expected connectivity of the fourth industrial revolution may imply frequent downloads of blockchain receipts from **L3**, requiring a greater computational requirement to validate Merkle tree root values with respect to the chain of transactions already stored locally. This action may even originate a certain penalty on the own computation of the PDP, which should in turn provide authorization services and access on-demand. This hit due to **P3** may also overflow the storage space of the PDP if lightweight database management procedures, either simple data structures or external resources (e.g., cloud-computing or servers), are not strictly applied. **P1** and **P2**, to the contrary, do not intercede on the computation and storage requirements of the PDP as the blockchain network is completely uncoupled from **L2**, and **P1** is more centered on the communication processes between **L2** and **L3**. However, all the transactions towards or from **L3** may add an extra overhead in **L2** if the number of transactions in **L2** and **L3** is considered, which might collapse the connections to the PDP and inherently block the control to **L1**.

Upgrading of resources (either HW/SW of PDPs or databases) in centralized systems can temporarily disable all the communication with the blockchain network (corresponding to **P1** and **P3**), the SECAUD activities, and the operational actions in the field. Depending on updating period, this fact may queue a large number of transactions, likely generating unnecessary computational additional costs both in **L3** and **L2**. In contrast, the distributed nature of **L3** adds the capacity to establish a more gradual maintainability as specified in [32]. Here, their devices can be upgraded following a gradual procedure without involving the overall disconnection of the entire blockchain system. Therefore, centralization leads to simple points of failure, and consequences that might be irreparable or inadmissible from the control standpoint.

Continuing with the maintainability, the deployment of a blockchain in **L3** for only one PDP can become quite costly, and does not provide any benefit beyond a simple data centralization, due to the presence of a single domain of trust. However, the price to pay for a secondary network that guarantees the implicit features of the technology, such as immutability in a federated environment, authenticity and redundancy, can be quite attractive and brings profitable features for the underlying system. The environment is complex and shared by diverse entities, making necessary that: (i) The information of the entire consortium remains protected against possible tamperers launched by insiders; (ii) the technological uncoupling for data provenance in **P2** avoids to sacrifice the interconnection in **L2**; and (iii) data redundancy and its availability become two

implantation requirements to benefit the safety-critical conditions. Related to this latter, the way to manage the data in **P2** inherently adds data transparency both in **L2** and **L3**, letting know in all times where and how the access information is processed and stored (thanks in part to the data provenance). Indeed, as **L3** transactions are copied in **L2** without breaking the lineage of the data and their relationships, any member of the consortium with reading permission can access to the data and establish traceability, auditing and accountability. Here, data simplicity depends on the consortium, which is responsible for pre-defining the base structure of the metadata and its attributes (e.g., $\langle \text{transaction} \rangle := \{ \langle \text{subject} \rangle, \langle \text{resource} \rangle, \langle \text{action} \rangle \}$) to later be managed by entities of **L2** and **L3**.

As for data security in terms of integrity, availability and confidentiality, it relies on the security policies imposed by the consortium. Depending on the access level to the data itself, the protection of the communication channel and the handle of the transactions in **L2** or **L3** (bouncing on **P1** and **P3**), may add new security risks. Any deliberated action may entail information leaks, data manipulation or falsification, as stated in [62]; but thanks to the immutability, these actions can be penalized. The penalization can be very variable; for example, the definitive expulsion from the consortium (which can be submitted to a majority vote between peers) or the variation of a reputation system in the space of $[0 - 100\%]$ to further limit the trust level. As stated in [63], an initial reputation of 100% can be assumed, but its value can vary according to $f(x_i) : x_i \pm y$, such that x_i corresponds to the current reputation of an entity_{*i*} and y refers to the increment/decrement value. As this procedure generally focuses on controlling actions, we adapt it to mitigate insider actions [64, 65]. In this case, not only incorrect actions should be penalized but also correct actions should be rewarded; and both cases should be part of a governance plan to clarify when to apply y as increase/decrement value.

Last but not least, several authors have already underlined in [12, 62] that transactional privacy is not always guaranteed, regardless the security level established. The transactions in our context are managed by resources belonging to the consortium. If security policies allow read permissions without a suitable use of cryptographic measures both in **L2** and **L3**, members of the consortium might derive and reveal identities and actions. Note that all these security and privacy issues are equally applicable for any interconnection system, either with a coupled or an uncoupled blockchain system.

5.2 CM-2: Decentralized interconnection and uncoupled blockchain

Decentralization is a solution that can mitigate some of the **CM-1** problems by simply delegating authorization services to several PDP proxies in **L2**. With this, SG entities would be able to establish more fluid connections to critical resources without risk to cause major problems that may lead to computational overheads and/or bottlenecks. But as the number of connections is expected to grow with the advent of Industry 4.0, what it is very probable that the number of transactions to/from **L3** may also be significant for the next industrial generation. This fact may even affect the overhead on the IT layers corresponding to **L2** and **L3**. Likewise, continued downloads of blockchain receipts from **L3** may negatively impact on the computation and communication of the PDPs. SW auditors require in this case to gain access to the blockchain system

to continuously process and validate transactions (**P1**), and update databases accordingly (**P3**). A way to relax the overhead, it would be to establish policies that do not impose frequent monitoring processes, but sufficient to keep access records updated in the diverse policy decision points.

At this point, the delegation of PDP actions is key to keep up the operational tasks in **L1** regardless of the conditions of provenance processes (**P1**, **P2** and **P3**) and the PDP nodes in **L2**. Indeed, PDP decentralization equally helps guarantee resource availability in **L1**. If a PDP is not able to manage access at a certain time, its action is relegated in another PDP, which is also in charge of managing access requests to **L1** and access data to **L3** through **P1**. Depending on the transaction volume compiled and sent to the blockchain space, topics related to ‘data’ scalability should also be addressed [12], mainly to anticipate and concrete the optimal storage space for **L2** and **L3**. In our case, we have specified lightweight transaction structures containing simple metadata of class $\{ \langle \text{subject} \rangle, \langle \text{resource} \rangle, \langle \text{action} \rangle \}^*$, the lineage of which has to be locally maintained by the diverse PDPs to guarantee a suitable past evidence-based access. Aligned with the scalability, the uncoupling of network layers (IT-OT) certainly benefits the inclusion of new PDPs or operational technologies without impacting on operational processes. Any new PDP device in **L2** could transparently connect to **L1** and **L3** without requiring modifying the existing architecture; offering transparency and efficiency in the communication processes from **L2** to **L1**.

This efficiency is not so appreciated when multiple replicates of a same transaction are distributed both in **L2** and **L3**, entailing complexity for its maintenance but benefiting in turn the safety-critical of the underlying system. All the PDPs must be able to download blockchain receipts to have an exact copy in their databases, offering data redundancy and data availability to face extreme situations. If, in addition, updating policies are regulated and established by the consortium, this capacity further promotes data transparency and accessibility to this information. In this way, if the downloads to **L2** are irregular (**P3**), it is quite probable that the decision-making of the PDPs are not so realistic and reliable since part of the past experience is missed. Therefore, the consortium must establish specific synchronization policies to determine when to proceed with the updates of data warehouses. Nonetheless, this synchronization does not have to be launched concurrently. It can be executed in close times to guarantee updated sequences in the all PDPs, and avoid bottlenecks in **L3** if several auditors SECAUD try to download of the last blockchain receipts at the same time.

Although the complexity of the **L2** system grows significantly, the upgrade HW/SW of PDPs and their associated data warehouse is less aggressive. The maintenance can be done from a gradual standpoint, allowing PDPs to temporally delegate their actions to other active PDPs whereas software components and resources are updated. Additionally, the Software Defined Networking (SDN) technology can be leveraged at this point, to dynamically manage and optimize network resources between **L2** and **L3** [66]. Dependability in the three layers is also guaranteed since the own decentralization aids to keep data copies in both **L2** and **L3**, and delegated access when PDPs present casual faults.

5.3 CM-3: Decentralized interconnection and coupled blockchain

Unlike **CM-2**, **CM-3** allows PDPs to not only process the authorization service but also to distribute and validate any access transaction considering a consensus protocol. The benefit of this model is the automatic synchronization of authorization policies between the participants within the consortium, which hence removes the need to manually execute **P3** periodically. For the same reason, the own policy rules could be implemented by means of Smart Contracts, which would enable an interoperability between authorization and blockchain operations [67]. This way of concentrating diverse actions at a same critical point may bring about extra overhead and delays that may penalize the responsiveness of the interconnection and access to critical resources in **L1**. PDPs not only present the ability to determine the connection in the field, but the capacity to verify, together with the rest of PDPs, the validity of a transaction for data provenance and traceability. This double functionality notably aggravates the management of PEP requests [6], mainly because part of the computation, the storage and the communication have to be reserved to: (i) Validate the pending transactions, (ii) interact with the rest of PDPs for its verification, and (iii) compete to add a block to the chain in a distributed environment. Depending on the consensus protocol and its added difficulty for the competition, the computational overheads can become very variable, hampering any access required for the control in the field.

In addition to this, the complexity of the architecture and the risks multiply. If PDPs eventually stop working or presents anomalies, the probability of continuing operations decreases considerably as these PDPs need to delegate actions. Any new delegation may in turn entail an increase of activities in the PDPs, triggering a cascading effect. Moreover, to all these implications we also have to add the particular actions of the auditors for the writing and the reading of the databases, which might intensify the computation overhead of the PDPs. It is true that at this level, the labor of the auditors and the number of replicates simplify considerably, mainly because the blockchain space is part of the PDP network; but auditors SECAUD periodically need to track the databases from their own memory space to help authorization services make decisions according to past experience. On the other hand, and given that the blockchain space is integral in each PDP, auditors do not require establishing synchronization among them so as to get consistency in the access process as could occur in **CM-2**. Rather, the consistency is already implicit in the own consensus protocol. The rest of properties such as (PDP and data) scalability and maintainability is similar to **CM-2** since **CM-2** and **CM-3** are based on decentralized systems.

6 Final discussion and future challenges

Table 2 summarizes all the analyses done in Section 5, especially illustrating the influence of the three layers-based interconnection systems and contemplating the three coupling models (**CM-1**, **CM-2** and **CM-3**), equally defined in Section 5. To characterize this aspect, we also consider in this table the clout of the three data provenance phases of Section 3.2, the processes of which can also impact on the decision-making in **L2** and access to **L1**.

Table 2. Influence of the three layers-based interconnection strategies on the control

		Real-time performance					Sustainability			Depend.		Survivability						Sf.-crit.	
		Computational	Storage	Communication	Download (DB)	Read/Write (DB)	Maintainability	PDS scalability	Data scalability	Data consistency	Availability	Integrity	Confidentiality	Authen. & Author	Immutability	Privacy	Traceability	Redundancy	Transparency
CM-1	P1			-			-	-			-	7	7	7					
	P2						+				+	7	7	+	+	7	+	+	+
	P3			-	-		-	-		7	-	7	7	7					
	L1	*	*	*	*	*	*	*	*	*	*	*	*	*	+		+	+	
	L2	3,6	3,4,6	3	-	-	-	-	4,6	7	3	7	7	+	+	7	+	+	+
	L3	3,6	3,4,6				+		4,6		+	7	7	+	+	7	+	+	+

CM-2	P1			7								7	7	7					
	P2						+				+	7	7	+	+	7	+	+	+
	P3			7	7					7		7	7	7					
	L1	+	*	+	+	+	+	+	*	*	+	*	*	*	+		+	+	
	L2	3,6	3,4,6	3	-	-	+	+	4,6	7	+	7	7	+	+	7	+	+	+
	L3	3,6	3,4,6				+		4,6		+	7	7	+	+	7	+	+	+

CM-3	P1																		
	P2																		
	P3																		
	L1	-	-	-			+	+	*	*	+	*	*	*	+		+	+	
	L2	-	-	-		+	+	+	4,6	7	+	7	7	+	+	7	+	+	+
	L3																		

+: The data provenance phase is applicable AND the property does not impact on the control of L1.

-: The data provenance phase is applicable AND the property impacts on the control of L1.

*: Probability of impacting on access or on its mode in L1, and therefore on the control.

1: Depends on the efficiency of the SW auditor to verify blockchain receipts and update the data provenance database.

2: Limited to the requirements of the consortium and its trustworthy level such as the goodness of its members.

3: Depends on the number of connections.

4: Depends on the simplicity of the data itself, its structure defined for the metadata (e.g., links, use of hashes, IDs, etc.) and the value of its attributes.

5: Depends on the type of maintenance; concretely whether the HW/SW update process is carried out gradually.

6: Depends on the technology itself, and its technical capacities to process and store operations.

7: Depends on security policies and/or policies established by the consortium.

From this table, we underline that centralized PDP systems based on an uncoupled DLT (**CM-1**) or decentralized PDP systems based on a coupled blockchain (**CM-3**), are not suitable approaches for Industry 4.0. **CM-1**-based systems are clear single failure points where all the control exclusively relies on a unique access point; whereas decentralized systems require the deployment of more interconnection resources to distribute their PDP actions. This feature makes that decentralized systems are today considered as a good approach for authorization services, mainly because part of the actions can be delegated in other PDPs. Precisely, the related work [32] clearly states it, in which the full integration of IoT devices in **L2** to create completely distributed environments may intensify control issues and harm the business continuity [62]. HW/SW constraints of the majority IoT devices do not help progress in the provision of new interconnection strategies [32]. If in addition to this, other technologies (e.g., blockchain) are adopted

to modernize the actions of the control industry, the complexities and their associated problems may become much more significant as occurs with **CM-3** - see also Table 2. With this, we also underline that access in critical systems should work in optimal times, offering control services in all times as specified in [33]. Namely, any new technological coupling in OT networks should not hamper the actions in the field; rather they should offer support to additionally optimize the operational processes.

All of these aspects are also underlined in Table 2 by means of three colours: green (good for the control), orange (depends on other issues) and red (bad for the control). Each colour means the impact level on **L1**. As can be noted **CM-3** is indeed a bad approach for critical environments followed by **CM-1**. In contrast, **CM-2** presents better conditions for these types of environments except for storage space, which depends on: The number of connections, the simplicity of the data and its structure, as well as on the technical capacities of the resources. A way to reduce the bulky transactions managed by both **L3** and **L2** could be the incorporation of external data management systems (e.g., edge, cloud or fog computing) and the implementation of a hierarchy of nodes in the Blockchain that permits the decoupling of responsibilities. For instance, Hyperledger Fabric allows to deploy nodes that are only in charge of maintaining a copy of the ledger, submitting transactions or creating blocks, thereby alleviating the load of the overall system [38]. Another solution is the introduction of sidechains, which consist in separate blockchains that are attached to its parent blockchain to exchange assets between them at a predetermined rate, increasing the speed of transactions in the main network [68].

More existing optimization approaches as stated in [12]. For example, J. Bruce in [69] defines an account tree as a database containing only simple non-empty addresses without incorporating the entire content of each transaction. This, in our architecture, could mean that each auditor SECAUD is able to validate the blockchain receipts through **P3**, build/update a simple account tree-based database with non-empty addresses, and transfer the entire blockchain receipt with its associated addresses to an external data management system. Other optimization solutions can also be considered such as the patent [70], which focuses on reducing the size of the blockchain by continuously examining the content of the transactions.

To conclude, we stress once again the multiple benefits of blockchain technology for Industry 4.0, but without forgetting the diverse complexities that the technology could bring about the operational performance. It is still necessary to find a good trade-off between operational performance and security as also is underlined in [33]. This also means that we still need to (i) progress in IIoT resources to support new services, and (ii) provide optimized blockchain solutions in terms of computation and storage, and more specifically for **P3**. Our going work revolves around supporting these findings with quantitative tests that assess the performance of these architectures and processes with a real setup. The ultimate aim is to achieve a seamless integration between Blockchain technologies with federated SG scenarios and authorization applications.

Acknowledgments

This work has been mainly supported by the EU H2020 project SealedGRID (H2020-MSCA-RISE-2017 programme under grant agreement No777996) financed by the EU Commission, and partially supported by the EU H2020-SU-ICT-03-2018 Project No. 830929 CyberSec4Europe (cybersec4europe.eu). Likewise, the work of the second author has also been partially financed by the Spanish Ministry of Education under the FPU program (FPU15/03213).

References

- [1] O. Tuttokmagi, A. Kaygusuz, Smart Grids and Industry 4.0, in: 2018 International Conference on Artificial Intelligence and Data Processing (IDAP), 2018, pp. 1–6.
- [2] NIST, NIST Smart Grid Conceptual Model, Update of the NIST Smart Grid Conceptual Model, available at https://www.nist.gov/sites/default/files/documents/2018/09/10/draft_smart_grid_conceptual_model_update.pdf, last access in February 2020 (2018).
- [3] C. Alcaraz, J. Lopez, Secure interoperability in cyber-physical systems, in: Security Solutions and Applied Cryptography in Smart Grid Communications, IGI Global, USA, IGI Global, USA, 2017, Ch. 8, pp. 137–158.
- [4] J. Lopez, J. E. Rubio, Access control for cyber-physical systems interconnected to the cloud, *Computer Networks* 134 (2018) 46 – 54.
- [5] C. Alcaraz, Secure interconnection of it-ot networks in industry 4.0, in: Critical Infrastructure Security and Resilience: Theories, Methods, Tools and Technologies, no. Advanced Sciences and Technologies for Security Applications book series (ASTSA), Springer International Publishing, 2019, pp. 201–217.
- [6] C. Alcaraz, J. Lopez, S. Wolthusen, Policy enforcement system for secure interoperable control in distributed smart grid systems, *Journal of Network and Computer Applications* 59 (2016) 301–314.
- [7] A. Valenzano, Industrial cybersecurity: improving security through access control policy models, *IEEE Industrial Electronics Magazine* 8 (2) (2014) 6–17.
- [8] G. Ryba, M. Jung, W. Kastner, Authorization as a service in smart grids: Evaluating the paas paradigm for xacml policy decision points, in: 2013 IEEE 18th Conference on Emerging Technologies Factory Automation (ETFA), 2013, pp. 1–4.
- [9] J. Vollbrecht, P. Calhoun, S. Farrell, L. Gommans, G. Gross, B. de Bruijn, C. de Laat, M. Holdrege, D. Spence, AAA authorization framework, RFC 2904 (2000).

- [10] K. Wüst, A. Gervais, Do you need a blockchain?, in: 2018 Crypto Valley Conference on Blockchain Technology (CVCBT), IEEE, 2018, pp. 45–54.
- [11] J. Wu, N. Tran, Application of blockchain technology in sustainable energy systems: An overview, *Sustainability* 10 (9) (2018) 3067.
- [12] Z. Zheng, S. Xie, H. Wang, Blockchain challenges and opportunities: A survey, in: *Work Pap.*, no. 1-25, 2018.
- [13] ID. Yaga, P. Mell, N. Roby and K. Scarfone, Blockchain Technology Overview, NISTIR 8202, available at <https://nvlpubs.nist.gov/nistpubs/ir/2018/NIST.IR.8202.pdf>, last access in February 2020 (2018).
- [14] ENISA, Distributed Ledger Technology & Cybersecurity - Improving information security in the financial sector, NISTIR 8202, available at <https://www.enisa.europa.eu/publications/blockchain-security>, last access in February 2020 (2016).
- [15] M. Andoni, V. Robu, D. Flynn, S. Abram, D. Geach, D. Jenkins, P. McCallum, A. Peacock, Blockchain technology in the energy sector: A systematic review of challenges and opportunities, *Renewable and Sustainable Energy Reviews* 100 (2019) 143 – 174.
- [16] M. Mylrea, S. N. G. Gourisetti, Blockchain for smart grid resilience: Exchanging distributed energy at speed, scale and security, in: *Resilience Week (RWS)*, 2017, IEEE, 2017, pp. 18–23.
- [17] Indigo Advisory Group, Blockchain in energy and utilities use cases, vendor activity, market analysis, available at <https://www.indigoadvisorygroup.com/blockchain>, last access in February 2020 (2019).
- [18] Dal Canto D. Enel, Blockchain: which use cases in the energy industry, *CIRE* 2017, Glasgow, Round table discussion (2017).
- [19] Z. Li, J. Kang, R. Yu, D. Ye, Q. Deng, Y. Zhang, Consortium blockchain for secure energy trading in industrial internet of things, *IEEE Transactions on Industrial Informatics* 14 (8) (2018) 3690–3700.
- [20] C. Liu, K. K. Chai, X. Zhang, E. T. Lau, Y. Chen, Adaptive blockchain-based electric vehicle participation scheme in smart grid platform, *IEEE Access* 6 (2018) 25657–25665.
- [21] R. Alcarria, B. Bordel, T. Robles, D. Martín, M.-Á. Manso-Callejo, A blockchain-based authorization system for trustworthy resource monitoring and trading in smart communities, *Sensors* 18 (10).
- [22] Y. Zhou, Y. Guan, Z. Zhang, F. Li, A blockchain-based access control scheme for smart grids, *IACR Cryptology ePrint Archive* 2019 (2019) 880.

- [23] G. Suci, C.-I. Istrate, A. Vulpe, M.-A. Sachian, M. Vochin, A. Farao, C. Xenakis, Attribute-based access control for secure and resilient smart grids, in: 6th International Symposium for ICS & SCADA Cyber Security Research 2019 6, 2019, pp. 67–73.
- [24] X. Liang, S. Shetty, D. Tosh, C. Kamhoua, K. Kwiat, L. Njilla, Prochain: A blockchain-based data provenance architecture in cloud environment with enhanced privacy and availability, in: 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), 2017, pp. 468–477.
- [25] T.-y. Feng, A survey of interconnection networks, *Computer* 14 (12) (1981) 12–27.
- [26] Y. Lu, Industry 4.0: A survey on technologies, applications and open research issues, *Journal of Industrial Information Integration* 6 (2017) 1–10.
- [27] N. J. Edwards, J. Rouault, Multi-domain authorization and authentication, Patent US 7,444,666B2, available at <https://patentimages.storage.googleapis.com/30/f9/cd/d0a8524d06b12d/EP1280317A1.pdf>, last access in February 2020 (2003).
- [28] H. F. Atlam, R. J. Walters, G. B. Wills, Fog computing and the internet of things: A review, *Big Data and Cognitive Computing*, MDPI 2 (2) (2018) 1–18.
- [29] C. Alcaraz, L. Cazorla, G. Fernandez, Context-awareness using anomaly-based detectors for smart grid domains, in: 9th International Conference on Risks and Security of Internet and Systems, Vol. 8924, Springer International Publishing, Springer International Publishing, Trento, 2015, pp. 17–34.
- [30] J. E. Rubio, R. Roman, C. Alcaraz, Y. Zhang, Tracking apts in industrial ecosystems: A proof of concept, *Journal of Computer Security* 27 (2019) 521–546.
- [31] C. Alcaraz, J. Lopez, K.-K. R. Choo, Resilient interconnection in cyber-physical control systems, *Computers & Security* 71 (2017) 2–14.
- [32] C. Alcaraz, R. Roman, P. Najera, J. Lopez, Security of industrial sensor network-based remote substations in the context of the internet of things, *Ad Hoc Networks* 11 (2013) 1091–1104.
- [33] C. Alcaraz, J. Lopez, Analysis of requirements for critical control systems, *International Journal of Critical Infrastructure Protection (IJCIP)* 5 (2012) 137–145.
- [34] BMG, Put the power of local energy choice in your hands, <https://www.brooklyn.energy>, last access in April 2020 (2017).
- [35] C. Burger, A. Kuhlmann, P. Richard, J. Weinmann, Blockchain in the energy transition. A survey among decision-makers in the german energy industry, DENA German Energy Agency 60.

- [36] C. Mohan, State of public and private blockchains: Myths and reality, in: Proceedings of the 2019 International Conference on Management of Data, 2019, pp. 404–411.
- [37] Y. Xiao, N. Zhang, W. Lou, Y. T. Hou, A survey of distributed consensus protocols for blockchain networks, *IEEE Communications Surveys & Tutorials* (2020) 1–34.
- [38] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, et al., Hyperledger fabric: a distributed operating system for permissioned blockchains, in: Proceedings of the Thirteenth EuroSys Conference, 2018, pp. 1–15.
- [39] A. Baliga, I. Subhod, P. Kamat, S. Chatterjee, Performance evaluation of the quorum blockchain platform, *arXiv preprint arXiv:1809.03421* (2018) 1–8.
- [40] R. G. Brown, J. Carlyle, I. Grigg, M. Hearn, Corda: an introduction, *R3 CEV*, August 1 (2016) 15.
- [41] R. B. Uriarte, R. De Nicola, Blockchain-based decentralized cloud/fog solutions: challenges, opportunities, and standards, *IEEE Communications Standards Magazine* 2 (3) (2018) 22–28.
- [42] C. Alcaraz, Cloud-assisted dynamic resilience for cyber-physical control systems, *IEEE Wireless Communications* 25 (1) (2018) 76–82.
- [43] F. ul Hassan, A. Ali, S. Latif, J. Qadir, S. Kanhere, J. Singh, J. Crowcroft, Blockchain and the future of the internet: a comprehensive review (2019). *arXiv:1904.00733*.
- [44] P. Buneman, S. Khanna, W.-C. Tan, Data provenance: Some basic issues, in: *International Conference on Foundations of Software Technology and Theoretical Computer Science*, Springer, 2000, pp. 87–93.
- [45] F. Zafar, A. Khan, S. Suhail, I. Ahmed, K. Hameed, H. M. Khan, F. Jabeen, A. Anjum, Trustworthy data: A survey, taxonomy and future trends of secure provenance schemes, *Journal of Network and Computer Applications* 94 (2017) 50 – 68.
- [46] M. Herschel, R. Diestelkämper, H. Ben Lahmar, A survey on provenance: What for? what form? what from?, *The VLDB Journal* 26 (6) (2017) 881–906.
- [47] M. R. Asghar, M. Ion, G. Russello, B. Crispo, Securing data provenance in the cloud, in: J. Camenisch, D. Kesdogan (Eds.), *Open Problems in Network Security*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 145–160.
- [48] S. Sultana, G. Ghinita, E. Bertino, M. Shehab, A lightweight secure scheme for detecting provenance forgery and packet dropattacks in wireless sensor networks, *IEEE Transactions on Dependable and Secure Computing* 12 (3) (2015) 256–269.

- [49] S. Shetty, V. Red, C. Kamhoua, K. Kwiat, L. Njilla, Data provenance assurance in the cloud using blockchain, *SPIE (Disruptive Technologies in Sensors and Sensor Systems)* 10206 (2017) 10206 – 10206 – 11.
- [50] A. Ramachandran, M. Kantarcioglu, Using blockchain and smart contracts for secure data provenance management, *CoRR* abs/1709.10000 (2017) 1–11. arXiv:1709.10000.
- [51] T. M. Fernández-Caramés, P. Fraga-Lamas, A review on the use of blockchain for the internet of things, *IEEE Access* 6 (2018) 32979–33001.
- [52] M. Pilkington, *Blockchain technology: principles and applications*, Research handbook on digital transformations (2016) 225.
- [53] K. Souali, O. Rahmaoui, M. Ouzzif, An overview of traceability: Definitions and techniques, in: *2016 4th IEEE International Colloquium on Information Science and Technology (CiSt)*, 2016, pp. 789–793.
- [54] G. D. P. Regulation, Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46, *Official Journal of the European Union (OJ)* 59 (1-88) (2016) 294.
- [55] The Cambridge Dictionary, Accountability, available at: <https://dictionary.cambridge.org/dictionary/english/accountability>, last access February 2020 (2019).
- [56] IEC-62351, IEC-62351-(1-8): Information security for power system control operations, international electrotechnical commission, <http://www.iec.ch/smartgrid/standards/>, last access in February 2020 (2007-2011).
- [57] A. P. Joshi, M. Han, Y. Wang, A survey on security and privacy issues of blockchain technology, *Mathematical Foundations of Computing* 1 (2) (2018) 121–147.
- [58] J. Kim, E. Deelman, Y. Gil, G. Mehta, V. Ratnakar, Provenance trails in the wings/pegasus system, *Concurrency and Computation: Practice and Experience* 20 (5) (2008) 587–597.
- [59] M. Imran, H. Hlavacs, I. U. Haq, B. Jan, F. A. Khan, A. Ahmad, Provenance based data integrity checking and verification in cloud environments, *PloS one* 12 (5) (2017) e0177576.
- [60] Y. L. Simmhan, B. Plale, D. Gannon, A survey of data provenance in e-science, *SIGMOD Rec.* 34 (3) (2005) 31–36.
- [61] C. Alcaraz, S. Zeadally, Critical control system protection in the 21st century: Threats and solutions, *IEEE Computer* 46 (10) (2013) 74 – 83. doi:10.1109/MC.2013.69.

- [62] A. Dorri, S. S. Kanhere, R. Jurdak, Towards an optimized blockchain for iot, in: *Proceedings of the Second International Conference on Internet-of-Things Design and Implementation*, ACM, 2017, pp. 173–178.
- [63] J. Lopez, C. Alcaraz, R. Roman, Smart control of operational threats in control substations, *Computers & Security* 38 (2013) 14–27.
- [64] M. Keeney, E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, S. Rogers, Insider threat study: Computer system sabotage in critical infrastructure sectors, Tech. rep., National Threat Assessment Ctr Washington Dc (2005).
- [65] J. Crampton, M. Huth, Towards an access-control framework for countering insider threats, in: *Insider Threats in Cyber Security*, Springer, 2010, pp. 173–195.
- [66] T. Alharbi, Deployment of blockchain technology in software defined networks: A survey, *IEEE Access* 8 (2020) 9146–9156.
- [67] M. P. Andersen, J. Kolb, K. Chen, G. Fierro, D. E. Culler, R. A. Popa, Wave: A decentralized authorization system for iot via blockchain smart contracts, EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2017-234.
- [68] A. Back, M. Corallo, L. Dashjr, M. Friedenbach, G. Maxwell, A. Miller, A. Poelstra, J. Timón, P. Wuille, Enabling blockchain innovations with pegged sidechains, URL: <http://www.opensciencereview.com/papers/123/enablingblockchain-innovations-with-pegged-sidechains> 72.
- [69] J. D. Bruce, The mini-blockchain scheme, White paper (2014).
- [70] R. Childress, R. Gupta, D. B. Kumhyr, M. Mukherjee, Limiting blockchain size to optimize performance, U.S. Patent Application 15396,960 (2018).