The SOFIE Approach to Address the Security and Privacy of the IoT using Interledger Technologies

Dmitrij LAGUTIN[a], Priit ANTON[b], Francesco BELLESINI[c], Tommaso BRAGATTO[d], Alessio CAVADENTI[d], Vincenzo CROCE[e], Nikos FOTIOU[f], Margus HAAVALA[b], Yki KORTESNIEMI[g], Helen C. LELIGOU[h], Ahsan MANZOOR[i], Yannis OIKONOMIDIS[h], George C. POLYZOS[f], Giuseppe RAVEDUTO[e], Francesca SANTORI[d], Vasilios SIRIS[f], Panagiotis TRAKADAS[h], Matteo VERBER[e]

[a]*Department of Communications and Networking, Aalto University, Espoo, Finland*
[b]*Guardtime, Tallinn, Estonia*
[c]*Emotion s.r.l., Bastia, Italy*
[d]*ASM Terni, Terni, Italy*
[e]*Engineering Ingegneria Informatica S.p.A., Rome, Italy*
[f]*Mobile Multimedia Laboratory, Athens University of Economics and Business, Athens, Greece*
[g]*Department of Computer Science, Aalto University, Espoo, Finland*
[h]*Synelixis Solutions, Athens, Greece*
[i]*Rovio Entertainment, Espoo, Finland*

Abstract. The Internet of Things (IoT) suffer from lack of interoperability as data, devices, and whole sub-systems are locked in 'silos' because of technical, but mostly business reasons. Many new applications would be enabled and existing ones could be implemented in a more cost-efficient way, if the 'silos' could be bridged in a secure and privacy preserving manner. The SOFIE approach provides an effective way of accomplishing this by using interledger technologies that leverage the distributed trust enabled by distributed ledgers. The federated approach of SOFIE facilitates the creation of cross-organisational applications. This chapter presents the SOFIE approach and details the benefits it provides in four real-world pilots.

Keywords. Internet of Things, Distributed Ledger Technologies (DLT), blockchains, privacy, security, smart contracts

1. Introduction

Fragmentation and lack of interoperability among different platforms is a major issue for the Internet of Things (IoT). Currently, IoT platforms and systems are vertically oriented silos unable (or unwilling) to exchange data with, or perform actions across, each other. This leads to multiple problems: reduced competition and vendor lock-ins, as it is difficult for customers to switch IoT providers or combine IoT devices and data from multiple vendors in a single system, worse security as vendors often use proprietary security solutions that have not been properly audited, worse privacy as vendors usually force their customers to move at least some of their data or metadata to the vendor's

cloud, and reduced functionality compared to what better interoperability between platforms would afford. As IoT systems are becoming prevalent in everyday life, lack of interoperability and the resultant reduced use of relevant data is growing into a significant problem for the whole society.

IoT systems face many important security and privacy challenges. Since IoT systems interact with the real world, security is extremely important as breaches can cause significant physical damage and even loss of life. Similarly, as using the IoT is becoming a compulsory part of everyday life and IoT devices are able to collect increasing amounts of personal data, people should be able to carry out their lives without compromising their privacy. IoT systems usually contain large numbers of devices, therefore manually configuring and managing every IoT device is not feasible. Hence, security and privacy solutions for IoT must support high degrees of automation.

Authorisation mechanisms are an integral part of IoT security. The device owner should be able to *authorise* other parties to access the device or its data in a secure, flexible, and decentralised manner. Decentralised authorisation is important as it allows authorisation without a central control point, which may become a bottleneck, an increased failure risk or attack target, or require manual work. As there are numerous IoT devices interacting with each other, people, and the rest of the world, strong *auditability* is also a very important security feature for IoT. This is necessary for the normal operation of the system (e.g., goods have been delivered, therefore the payment should be made), troubleshooting in case of a problem, and dispute resolution between the parties involved.

In addition to the above-mentioned challenges, there are security challenges which will not be covered in this chapter. These include IoT device-level security, including the secure firmware updates, and verifying that IoT data is authentic and correct. The latter problem is very difficult to resolve in practice: it is not enough that the device is properly designed, implemented, calibrated, installed and certified, as it is easy to e.g., manipulate a thermometer by installing a heat source next to it.

From the privacy point of view, it is important to minimise both the *data collection* and *storage*. Especially, long-term storage is dangerous, since encrypted data will be revealed after the used encryption algorithm is broken, which will eventually happen. Protection against *correlation attacks* should also be provided, as in many situations the service should not be able to identify the user or even be aware that the user has used the service previously.

The EU H2020 project SOFIE[1] enables applications to link heterogeneous IoT platforms and autonomous things across technological, organisational, and administrative borders in an open and secure manner, thus simplifying the reuse of existing infrastructure and data, and allowing the creation of open business platforms, which in turn enable new kinds of services. This goal is accomplished by using Distributed Ledger Technologies (DLTs) [1] and interledger techniques – without requiring any modifications to the existing IoT platforms. Decentralised identifiers (DIDs) [2] can be used to manage users' identifiers in a privacy-preserving way. In the long term this will also enable open data markets, where participants can buy and sell IoT data and access to IoT actuation (or more generally: dictate rules for access to data and actuation) in a decentralised and automated manner.

[1] Secure Open Federation for Internet Everywhere (SOFIE), funded by EU's Horizon 2020 Programme under Grant 779984, 1.1.2018 – 31.12.2020, https://www.sofie-iot.eu.

The contributions of this chapter include descriptions of: 1) how DLTs, interledger techniques, and DIDs can be utilized to resolve problems with IoT security and privacy 2) how to realise a secure and open federation among heterogeneous IoT platforms and 3) examples of how these techniques can be leveraged in complex real-world systems, namely: (a) food supply chain provenance and transport conditions tracing, (b) electricity distribution grid balancing through electrical vehicle (EV) charging, (c) mixed reality mobile gaming with interactions between real and virtual worlds through IoT devices, and (d) secure sharing of electricity smart meter data.

This book chapter is organized as follows: Section 2 provides background information about DLTs and DIDs. The SOFIE pilots are presented in Section 3, while Section 4 describes the SOFIE IoT federation approach. Section 5 highlights benefits of SOFIE from the pilot use cases point of view. Related work is discussed in Section 6, while Section 7 provides a discussion of relevant issues and Section 8 concludes the chapter.

2. Background

This section describes key related technologies such as distributed ledgers, interledger technologies, and decentralised identifiers.

2.1. Distributed Ledger and Interledger Technologies

Distributed Ledger Technologies (DLTs), such as blockchains, offer decentralised solutions for collaboration and interoperability. One of the main features of DLTs is the *immutability* of data: ledgers are append-only databases where existing data cannot be modified and only new data can be added. Another major feature of DLTs is a distributed *consensus mechanism* [3], which controls what and how data is added to the ledger. Finally, DLTs also *replicate* data to participating nodes thus improving availability. Because of these three properties, DLTs avoid a single point of failure and offer resilience against many attacks. It is relatively easy to determine if any of the participating nodes in the DLT are misbehaving and even in an extreme case where an attacker manages to control the majority of the DLT's resources, the attacker can only control the addition of new data and in some extreme cases modify the very latest previously added data (but not the older data).

DLTs can be implemented with different levels of openness. They can be fully *open* (permissionless), which means that anyone can join the DLT and propose transactions; most well-known DLTs such as Bitcoin[2] and Ethereum[3] are based on this principle. However, DLTs can also be permissioned, either *semi-open*, in which case read access is open to everyone but write access is restricted, or *closed*, in which case both read and write access are restricted.

Overall, the main practical innovation of DLTs is the enablement of distributed trust. While there have been multiple proposals for distributed databases in the past, they have mostly concentrated on the distributed implementation, while the trust model has remained firmly centralized. In contrast, DLTs allow various entities, such as individuals, organizations, and companies, which may not fully trust each other, to collaborate in a

[2] https://bitcoin.org/en/

[3] https://www.ethereum.org/

safe and transparent manner, with only a low risk of being cheated by others. This makes DLTs a natural approach for solving the (business rather than technical) interoperability problem among IoT platforms.

Smart contracts [4] are another important feature provided by several DLTs: they are distributed applications that are executed on the ledger. Whenever an entity interacts with a smart contract, all operations are executed by all (full) nodes in the DLT network in a deterministic and reliable way; one of these nodes is selected to store the contract's execution outcome (if any) in the ledger. Smart contracts can verify DLT identities and digital signatures, perform general purpose computations, and invoke other smart contracts. The code of the smart contract is immutable and cannot be modified even by its owner. Moreover, since all transactions sent to a contract are recorded in the DLT, it is possible to obtain all historical values of the contract. Smart contracts typically refer to code running on Ethereum (in which case they are Turing-complete), but similar functionality is available in other DLTs. In particular, in the permissioned Hyperledger Fabric [5], similar functionality is named *chaincode*, while simpler, more constrained scripts can be run on Bitcoin. Smart contracts or similar functionality is critical for automating processes and will be exploited in the techniques described later.

There exists a large number of DLTs, each offering different trade-offs in terms of latency, throughput, consensus algorithm, functionality, etc., thus rendering them suitable for different types of applications. For example, a DLT can focus on cryptocurrency payments, recording of IoT events, or access authorisation. In complex systems it is therefore often not feasible to use only a single DLT for everything, hence the *interledger* approach that allows different DLTs to exchange data with each other is required in many situations. Using multiple ledgers is also beneficial for privacy reasons: participants within a DLT need to be able to access all data stored in that DLT to independently verify its integrity, which encourages the participants to use private ledgers, and store only a subset of the data to the main ledger used for collaboration with others. Multiple ledgers are also necessary for crypto-agility, as cryptographic algorithms used by DLTs (such as SHA-256) will not stay safe forever, thus it is necessary to have a mechanism to transfer data from one ledger to another. Siris et al. [6] present a review of interledger approaches, which differ in their support for transferring and/or trading value between ledgers, whether they support the transfer of information in addition to payments across ledgers, the balance between decentralised trust and cost (which can include both transaction cost and delay), the level of privacy, and their overall scalability and functionality that can facilitate the innovation of the DLT ecosystem.

A concrete example of the use of interledger is the following: Some parties decide to use the Hyperledger Fabric blockchain, which provides low-cost transactions and chaincode for transaction automation, for recording IoT and authorisation-related events. Parties also decide to use the Ethereum blockchain in order to make payments and fully automate the whole process with smart contracts. An interledger mechanism can be used to interconnect these two ledgers in a way that ensures atomic transactions, i.e., either both the authorisation and payment related transactions succeed, or both fail.

2.2. Decentralised Identifiers (DIDs)

Currently, an identity technology receiving much attention are decentralised identifiers (DIDs). A key aspect of DIDs is that they are designed not to be dependent on a central issuing party (Identity Provider or IdP) that creates and controls the identity. Instead,

DIDs are managed by the identity owner (or a guardian on the owner's behalf), an approach known as *self-sovereign identity* [7].

There are several different DID technologies in development [8], some of the most prominent being Sovrin[4], uPort[5], and Veres One[6]. These technologies started with similar but distinct goals in mind, but lately many of them have adopted the approach and format of the W3C DID specification [2], thus rendering them more and more interoperable. The specification defines a DID as a random string, often derived from the public key used with the identity. If a new DID is allocated for every party one operates/communicates with, correlating one's activities with different parties would be significantly harder to achieve. This property can be further enhanced by replacing existing DIDs with new ones at suitable intervals, e.g., even after just a single use.

Yet DIDs alone do not suffice, as some means of distributing the related public keys, any later changes to the keys, or other identity-related information is required. To this end, many of the DID solutions rely on a DLT for public DIDs (used by parties that want to be known publicly), whereas for private DIDs (e.g., used by individuals) application specific channels are used to distribute the information. Some DID technologies, e.g., Sovrin and Veres One, are launching their own permissioned DLTs, while others rely on existing blockchains (e.g., uPort is built on top of Ethereum). All three example technologies originally intended to use DLTs/blockchains for distributing information about DIDs belonging to individuals and IoT devices in addition to organisations, but the emergence of the General Data Protection Regulation (GDPR) [9] in the EU and other similar requirements have made storing personally identifiable information on a non-mutable platform such as a DLT/blockchain problematic. For this reason, Sovrin and Veres One have already excluded individuals DIDs from the ledger – and similar treatment may face the DIDs of IoT devices if they reveal personal information.

In many cases, there is also a need to associate machine verifiable properties to the identifier of an entity. This is accomplished with Verifiable Credentials (VCs) [10], which are analogous to traditional authorisation certificates. In a VC, the party issuing the credential (i.e. the *issuer*) states that according to them, the party about which the credential is made, known as the *prover*, has the stated properties. These could be e.g., the person's name, date of birth etc. in the case of driver's license issued by the police. To rely on a credential to prove something, the prover also has to demonstrate that the credential was issued to them. This can be done e.g., by proving the possession of the private key corresponding to the public key used in the credential (if the credential format supports such information), or with a separate proof built onto the credential. With a suitably created credential, a proof can also be used to only reveal *some of the attributes* of the credential (known as *selective disclosure*) or even prove that e.g., one is over a certain age without revealing the actual age attribute (a property known as *zero-knowledge proof*).

3. SOFIE Pilots

This section describes security and privacy challenges of four SOFIE real-life pilots that rely heavily on IoT: 1) agricultural/food supply chain, where produce growth and

[4] https://sovrin.org/

[5] https://www.uport.me

[6] https://veres.one/

transportation conditions are tracked from field to fork, 2) power balancing of the electrical grid by offering incentives to EV owners to charge their cars at certain times and locations, 3) mixed reality mobile gaming, where gamers can interact with the real world through IoT devices, and 4) utilizing data from electricity smart meters to develop various applications, e.g., to suggest the best electricity provider for a given user profile.

3.1. Tracking Food from Farm to Fork

The farm-to-fork pilot demonstrates a community-supported heterogeneous end-to-end agricultural food chain scenario. The main goal is to provide accountability, immutability, and auditability for both IoT data and transactions between different parties. As a result, consumers can reliably verify the provenance of a specific product from farm to fork. This gives consumers the ability to make decisions about their food based on e.g., health and ethical concerns, including environmental sustainability, fair labour practices, the use of fertilizers and pesticides, and other similar issues. The producers will be able to launch new products with a description, pricing, quantity and photos, while customers may interact with the marketplace, looking for products that fulfil certain requirements or preferences. Enabling immutable transactions also helps with dispute resolution between all the parties involved, reducing the chances of fraud and cutting out corresponding mediation expenses and transaction costs.

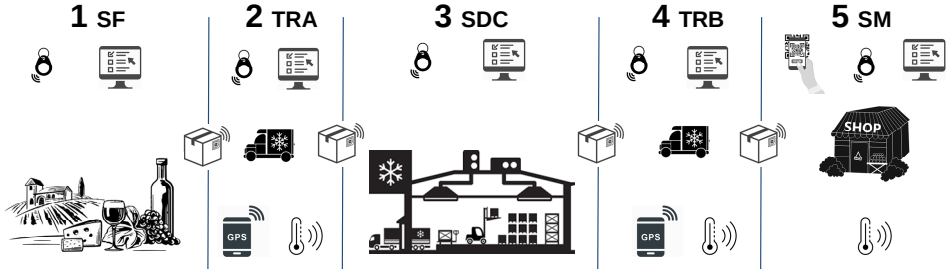

Figure 1. An overview of the SOFIE food-chain pilot, describing how agricultural produce moves from the farm to the supermarket through transporters and distributers

The path from farm to fork is split into 5 segments as depicted in Figure 1, and between segments the produce is handed over to the party responsible of the next segment. Each IoT platform uses its own data management and storage infrastructure (which can be either a database or a DLT).

Smart Farm (SF): In the farm, there are multiple sensor nodes capable of measuring e.g. temperature, humidity, wind speed/direction, rainfall, and soil moisture.

Transportation Routes A (TRA) and B (TRB): These segments cover the paths from the SF to the Storage & Distribution Centre (SDC), and from the SDC to the Supermarket (SM). The vehicles are equipped with GPS and temperature sensors.

Storage & Distribution Centre (SDC): SDC is where the smart boxes with farm crops will be stored until they are transported to the Supermarket. In SDC, a number of sensors monitor, among other parameters, temperature and presence of the boxes.

Supermarket (SM): SM contains the storage area, where the boxes are kept until they are placed in the customer area, and the customer area, where the products are available for the customers. Before the products are removed from the smart boxes to be placed to the customers' area, QR labels are created and applied on the crop packages, enabling the retrieval of the relevant information by customers.

The security and privacy challenges of this pilot include: how to accurately record the IoT data and handover events of the produce between different parties, how to provide sufficient audit trail for dispute resolutions, and how to minimize the leakage of private data (e.g., real identity of the workers should not be revealed to other party during the handovers).

3.2. Grid Balancing with Scheduled Electrical Vehicles Charging

In a second pilot, the goal is to balance a load on a real electricity network, namely the distribution grid of the city of Terni located in central Italy. There, a notable amount of energy is produced locally by distributed photovoltaic plants [11], which on occasion can cause Reverse Power Flow, when unbalances between locally produced and consumed electricity occur. To avoid this abnormal operation [12][13], electrical vehicles (EVs) will be offered incentives to match their EV charging needs with the distribution network's requirements.

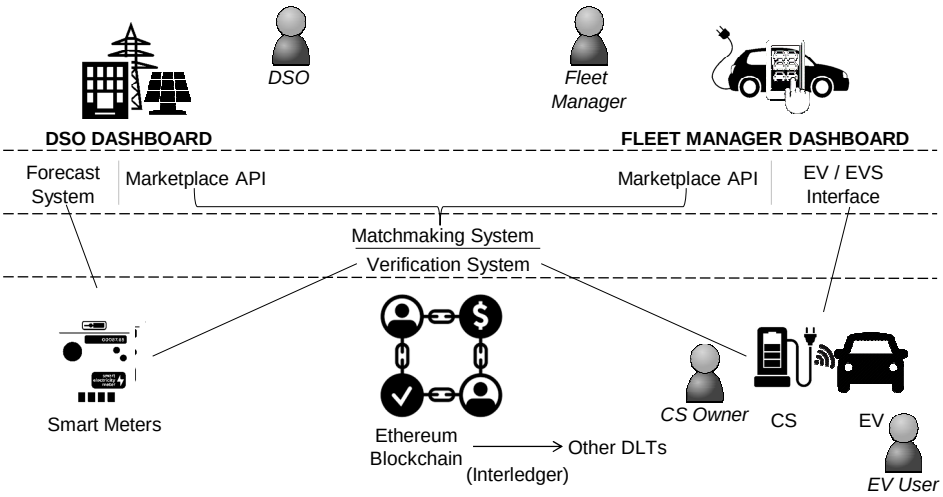

Figure 2. An overview of the SOFIE energy pilot, describing how DSO, EV fleet manager, and EV users utilise a decentralised marketplace to optimize the load on electrical grid

The actors in the pilot, as depicted in Figure 2, are the Distribution System Operator (DSO), who is responsible for grid management, the charging station (CS) Owner, who owns and manages the EV charging stations, the Fleet Manager, who represents EVs in energy price negotiations, and the EV users, who receive information and requests about the optimal scheduling of the charging of their vehicle. The main part of the pilot is a decentralised marketplace enabling DSO and fleet manager to negotiate on scheduled electricity consumption (using EV charging) and associated incentives, thus forming an end-to-end scenario from production via distribution to storage and consumption.

Both the DSO and the fleet manager interact with the system through their dedicated dashboards that show near real-time data collected from the two IoT subsystems (i.e. smart meters for the DSO and EV sensors for the fleet manager). The actors create requests and offers accordingly on the decentralised marketplace.

From the security point of view, it is important that agreements made on the marketplace cannot be tampered with, there is a secure way to verify that the terms of the agreement have been carried out, and parties will be compensated accordingly after

the agreement has been fulfilled. From the privacy point of view, it is important to protect privacy of the electric vehicle users, therefore the DSO or the CS Owner should not be able to determine EV users' real identities or correlate their charging activities.

3.3. Context-aware Mobile Gaming

In-game assets are a large market, with rare assets costing thousands of euros. However, currently in-game asset market poses significant risks to the players, since the gaming company can create unlimited instances of the assets, which in turn devaluates them, or the assets can disappear completely if the gaming company closes down. The goals of this pilot are to 1) provide a mechanism for recording asset ownership and trades in a secure and transparent manner, and 2) allow interactions between the mobile games and physical world through IoT devices.

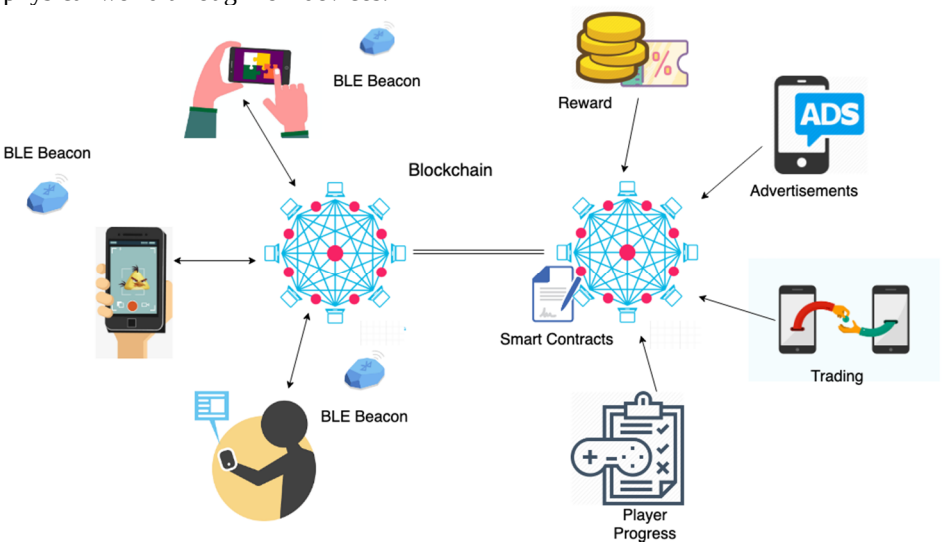

Figure 3. Overview of the SOFIE context-aware gaming pilot

An overview of the pilot is shown in Figure 3. The main actors of pilot include Game player, who can play any challenge, manage their profiles and assets, and claim reward data, through a mobile application, Game company, which is responsible for developing and maintaining the game servers, Challenge designer, who can create new challenges, assets, tasks, and puzzles using the existing game infrastructure, and the Asset designer, who can also list their creation for the trade on the SOFIE platform.

Multiple use-cases will be studied throughout the pilot. In the first use case players can collect and trade in-game content (e.g., characters, weapons, equipment, parts, etc.).

The second use case will utilize a scavenger hunt location-based game using IoT beacons. The player needs to solve the riddles using the received clues to reveal the location of the IoT beacon, which needs to be visited by the player. The player must perform some tasks (such as viewing an advertisement) to collect the points, which can be later redeemed for rewards.

The third use case allows generic trading between IoT resources and gaming assets. For example, as an extension of the IoT beacon use case, gamer who would perform certain real world activity (physical exercise, solving puzzles, etc.) with IoT devices could receive a gaming asset as a reward. The possession or sale of gaming assets could

in turn enable the gamer to, for example, receive a discount from the vending machine, temporary control of a robot in a mall, or some other IoT resource access.

There are several security and privacy issues in this pilot. Securing access control (for both IoT data and actuation) is very important as gamers are interacting with third-party IoT devices. The system should also offer an audit trail to help with dispute resolution in case something goes wrong. Furthermore, the owner of IoT beacons or other IoT devices should not be able to track players or determine their real identity. Finally, when IoT resources are exchanged with gaming assets, parties managing abovementioned resources should not be able to determine the other side of the trade. For example, if a player uses the gaming asset to gain access to an IoT resource, the owner of the IoT resource should only see that the player receives the access to the resource, without being able to determine whether access has been granted with the help of gaming assets or by other means (e.g., a monetary payment). In a similar way, if a player receives a gaming asset as a reward for solving physical challenge, other parties should not be able to determine how the gaming asset was received.

3.4. Decentralised Energy Data Exchange

The core idea of this pilot is to provide secure data exchange of smart meter data between end users, infrastructure owners, and energy service providers (intermediaries, distributors, and brokers). This in turn enables novel services such as fine-grained energy trading and energy flexibility marketplace. The overview of the pilot is shown in Figure 4 where participants, the SOFIE approach and the added value are presented.

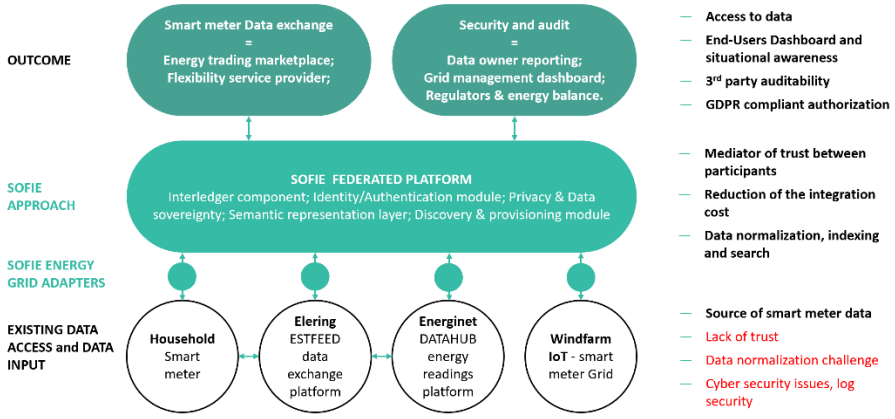

Figure 4. An overview of the SOFIE decentralised energy data exchange pilot

The key input for the pilot is the Estfeed open software platform (connecting 700 000 smart meters in Estonia). In order to demonstrate the cross-border data exchange and transfer of trust between network grid participants, the Danish Datahub (Energinet) will be the secondary input for the pilot. Besides national hubs integration, the pilot will also develop adapters and connection to two other instances: a local IoT network (windfarm) and a household metering point.

The main objective of the pilot is to enable trust between parties who exchange energy meter readings, which in turn creates several security and privacy challenges. From the data owner side, it is critical to guarantee control of the data (including the ability to grant and revoke access to/from third parties), as well as to have access to audit

logs for transparent overview of to whom the data access rights are given and how private data are handled. From the smart meter system operator side (transmission or distribution system operator), there is a need for mechanisms to agree on and prove the responsibility of the smart meter data after the data exchange. Data consumers (brokers, aggregators, and energy traders) need to access authentic smart meter data and be able to reliably verify the whole data provenance chain. The auditors require access to audit logs and tamper-proof evidence of the activities that have taken place in data exchange process. The pilot can be divided into the following two scenarios:

1. Data exchange - covering the full chain from identification, authorisation to requesting and granting access and exchanging the smart meter data.
2. Data exchange verification - including audit logs, tamper-proof evidence in case of disputes, and verification of the integrity of smart meter data.

4. The SOFIE Federation Approach

The main goal of SOFIE is to federate existing IoT platforms in an open and secure manner, in order to enable interoperability and without making any internal changes to the platforms themselves. Here, openness refers both to technical aspects (interfaces, implementation, etc.) and to flexible and (at least partially) open business models. A key benefit of SOFIE is that it allows the creation of solutions that connect many individual systems to a whole that provides significant new functionality. The approach also preserves users' privacy and is compliant with the EU General Data Protection Regulation (GDPR), which requires the minimisation of personal data collection.

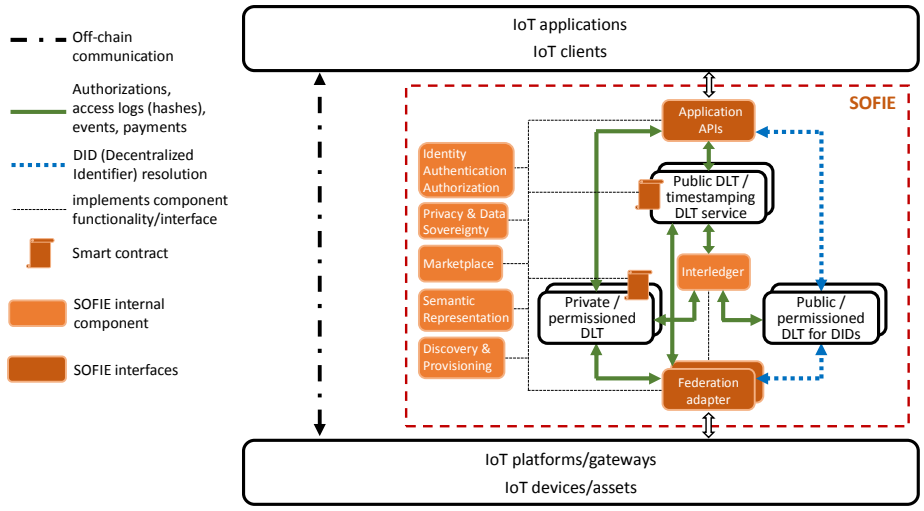

Figure 5. SOFIE framework architecture

The SOFIE framework architecture is depicted in Figure 5. The lowest level of the architecture contains IoT assets (or resources), that include e.g., IoT sensors for sensing the physical environment, actuators for acting on the physical environment, and boxes with RFID tags that are used to transport products. The IoT assets can be either connected to or integrated in actual devices. IoT platforms include platforms with data stores, where the measurements from sensors are collected and made available to third parties, in addition to servers providing IoT services. The federation adapters are used to interface

the IoT platforms with the SOFIE framework. This allows the IoT platforms to interact with the SOFIE framework without requiring any changes to the IoT platforms themselves. Moreover, different scenarios and pilots can utilize different types of federation adapters, which implement only the required parts of the SOFIE functionality.

The architecture emphasises the interledger functionality responsible for interconnecting different types of DLTs, which can have quite diverse features and functionality. The architecture also illustrates the separation of data transfer and control message exchanges. Some IoT data can be transferred directly between the IoT platforms and IoT clients. Control messages related to authorisation logs, events, payments, etc. go through the SOFIE framework.

The other SOFIE framework components [14] are: *Identity, Authentication, and Authorization* (IAA), which provides identity management and supports multiple authentication and authorisation techniques; *Privacy and data sovereignty*, which provides mechanisms that enable data sharing in a controlled and privacy preserving way; *Semantic representation*, which provides tools for describing services, devices, and data in an interoperable way; *Marketplace*, which allows participants to trade resources by placing bids and offers in a secure, auditable, and decentralised way; and *Discovery and provisioning*, which provides functionality for the discovery and bootstrapping of services. Finally, in the upper part of the figure are the application APIs, which provide the interfaces for IoT clients and applications to interact with the SOFIE framework. The rest of this section describes the most important components from privacy and security perspective in more detail: interledger, IAA, and Privacy and data sovereignty.

SOFIE results are open source and the source code for the SOFIE framework is available from https://github.com/SOFIE-project.

4.1. Interledger

The main purpose of the SOFIE interledger component is to enable transactions between actors and devices belonging to different (isolated) IoT platforms or silos. Each IoT silo either utilizes or is connected to one or more DLTs.

SOFIE's pilots and evaluation scenarios will utilize the following ledgers: Ethereum (both private deployments of the Ethereum client code and public test networks such as Rinkeby and Ropsten), Hyperledger Fabric, Guardtime's KSI blockchain, and Hyperledger Indy. If the federated IoT silo relies upon a ledger, such a ledger can be connected via SOFIE to the degree allowed by both the silo owner and the connected ledger governance or owner, provided that the silo has been enabled to support SOFIE federation.

Cross-chain transactions can take different forms depending on the specific scenario and its requirements. For example, interactions between a public and a permissioned ledger can use hashed time-lock contracts (HTLCs) to cryptographically link transactions and events on the two ledgers. In that scenario, the public ledger can record payments while the permissioned ledger can record authorisation message exchanges and IoT events. Alternatively, hashes of records stored on the permissioned ledger can be periodically recorded on the public ledger in order to provide a timestamped anchoring point. This approach exploits the wide-scale decentralised trust provided by the public ledger, while keeping the actual records accessible only by a permissioned set of nodes. Finally, interactions between a public or permissioned ledger and a ledger storing DID documents can focus on the resolution of DIDs to DID documents. This allows the interoperability between different DID implementations and different trust, privacy, and

cost tradeoffs with different selections of the ledgers for storing the transactions and the DID documents. The interledger functionality can be implemented in different entities, which include the entities that are interacting, a third party, or multiple third parties. In the case of third parties offering interledger services, such services can be provided for some fee. Moreover, in the case where multiple third parties offer interledger services, some coordination between the different parties is necessary.

4.2. Identity, authentication, authorization (IAA)

The goal of the Identity, Authentication, Authorization (IAA) component is to provide mechanisms that can be used for entities' and services' identification and authentication, and consumers' authorisation. To this end, it supports the following Identification and Authentication mechanisms: URIs (e.g., Web of Things URIs) for identification coupled with digital certificates for authentication, usernames for identifications bounded to secret passwords used for authentication, and decentralised identifiers (DIDs) associated with a DID documents used for authentication. A popular DID implementation, also considered by our component, is Hyperledger Indy. Consumers' authorisation is primarily implemented with the widely-used OAuth 2.0 protocol. The IAA component supports plain OAuth 2.0, OAuth 2.0 tailored for constrained devices as defined by the IETF's Authentication and Authorization for Constrained Environments (ACE) working group, and OAuth 2.0 combined with DIDs. Furthermore, it supports various token types and encodings. In addition to OAuth 2.0, the IAA component supports the UMA (User-Managed Access) protocol. An example of utilising DIDs together with OAuth 2.0 is presented in [15], while general authorisation solutions for IoT utilising DLTs and smart contracts are presented in [16] and [17].

The IAA component can use smart contracts in order to link authorisation decisions with payments, as well as for logging transaction-specific information that can be later used for auditing and dispute resolution. Moreover, authorisation decisions can be linked to IoT events that are recorded on the blockchain.

4.3. Privacy and Data sovereignty

The goal of the Privacy and Data sovereignty component is to enable data sharing in a controlled and privacy preserving way. This component considers privacy preservation as a two-dimensional problem. The first dimension concerns the privacy of the data provider, whereas the second dimension concerns the privacy of the data consumer. Data provider privacy is related to the amount and the accuracy of information a third party (including the consumer) can deduce about the provider from all the available data. This can be achieved by reducing or obfuscating the data stored in ledgers. A mechanism to reduce the data is to store only hashes on a more accessible platform. Depending on the use case it could mean storing data in private databases and storing hashes on permissioned ledger, or storing data in permissioned ledger and storing hashes on public ledger. Mechanisms to obfuscate data include differential privacy mechanisms. In particular, data obfuscation can be provided by selecting a special purpose node that acts as a data accumulator and also adds noise to the (encrypted) collected data. An alternative can be adding noise directly at the sources; however, in order to achieve the required degree of privacy and accuracy of the results, this approach requires a large number of sources. The coordination among the entities, namely the data provider, data consumer, and data accumulator, can be achieved through a smart contract. Consumer privacy is

related to the amount and the accuracy of information a third party (including the provider) can deduce about the consumer during the authentication, authorisation, and payment processes. To this end, this component supports attribute-based access control, where consumers can attest some of their attributes using verifiable credentials and zero-knowledge proofs. The underlying mechanisms support the minimum disclosure of information necessary to obtain a service. Additionally, multiple identifiers can be used to further improve privacy.

Data sovereignty is achieved by supporting two access control mechanisms: access control through delegation to an authorisation server and cryptotoken-based access control imposed by smart contracts. The first scheme enables data owners to define an authorisation server (AS), i.e., a special type of mediator that vouches about the eligibility and/or handles payments made by a consumer to access an IoT resource. The second scheme leverages blockchain-backed cryptotokens and enables owners to define access control policies based on these tokens. Cryptotokens can be granted only through a blockchain transaction and blockchain-specific functions, such as transfer, aggregation, etc. can be applied on these tokens.

5. SOFIE Benefits

This section describes how the SOFIE approach provides benefits for real-world pilot use cases in terms of interoperability, security, and privacy.

From the interoperability and security point of view, smart contracts, immutability, decentralisation, and other properties of DLTs allow high-level of automation, low risk of fraud, and efficient dispute resolution between participants. This enables interoperability between multiple parties in a secure and transparent manner, without requiring changes to the underlying IoT platforms. DLTs also allow maintaining non-repudiation and transparency without compromising privacy and business secrets by keeping the critical data in private data stores, while storing hashes of that data to DLTs. Only in case of a dispute the actual data will be revealed, and hashes stored in DLT guarantee that the data has not been tampered with in the meantime.

DIDs are used to enhance the privacy, since they allow the user to be in charge of their digital identifiers and solely be in possession of the associated private key (in contrast to some schemes that rely on centralised key management and distribution). DIDs also allow identifiers to be changed frequently, which offers protection against correlation attacks. In most of use cases, it is not necessary for third parties to know the real identity of the user, or even be aware that the user is the same who used the system previously, it is enough to determine that the user has a right to perform some action (such as to deliver the package on behalf of the company or charge the electrical vehicle).

5.1. Tracking Food from Farm to Fork

This pilot utilises a Consortium Ledger (CL) (private Ethereum) with smart contracts to record all the relevant data and meta-data related to the whole provenance chain from the farm to the supermarket. The members of the Consortium Ledger are the participants of the provenance chain (this if, for example, some of the produce is transported by other companies, another CL is formed) and a Legal Entity on a national or European level (association or public authority). Most of the measurements (temperature, soil conditions, humidity, etc.) are stored in private databases of IoT

platforms with hashes being frequently stored on CL. Aggregated data, such as average, maximum, and minimum temperature during the stage, is also stored on CL, along with all handover events between the participants (e.g., when a package is delivered by the transportation company to the warehouse). Finally, the hashes of CL transactions and data are periodically stored in public DLTs through interledger operations for extra accountability and transparency.

DIDs are used to protect the identities of the involved employees as, for example, the warehouse accepting a package does not need to know the real identity of the truck driver, or even whether the driver is the same as yesterday, it is enough to know that the truck driver is authorised by the transportation company to deliver that package.

5.2. Grid Balancing with Electrical Vehicles

In this pilot, the marketplace matching energy flexibility bids and offers operates on a private Ethereum blockchain, ensuring privacy (i.e., data cannot be read by external parties) and reducing transactions costs and times (i.e., mining is not required). Using SOFIE's interledger capabilities, this "first layer" will be paired with a public DLT acting as a "second layer", where the status of the private blockchain will be periodically synchronized, granting security and auditability, thus protecting the data stored in the first layer DLT from any alterations. The business logic for the requests and offers collection, and for the winning offer selection algorithm is coded in smart contracts, ensuring transparency and auditability of the whole process.

In the current version of the marketplace, a smart contract implements an auction mechanism, in which the best offer is selected following the "lowest bidder" rule. In the upcoming versions of the marketplace, an upgraded version of the smart contract will consider a different matchmaking algorithm, based on the clearing price algorithm used in commodity trades. In addition, the smart meter readings are stored on blockchain to ensure transparency, and the blockchain will also contain data of electric vehicles, charging stations, and charging events. Such data will be used for payments by the DSO to the fleet manager and for rewarding the users (through tokens or discounts) in an automated manner.

The pilot can be easily extended to include a retailer actor in charge of accounting, providing benefits to the two main actors involved: the DSO benefits of the grid stability provided and the fleet manager can reduce the overall charging costs to be paid to the retailer thanks to the incentives awarded by the DSO.

The privacy of the electrical vehicle users can be further enhanced with DIDs to protect their privacy against the DSO or the CS Owner, who do not need to know the real identity of the user charging the vehicle [18][19].

5.3. Context-aware Mobile Gaming

The mobile gaming pilot utilises a permissioned DLT (Hyperledger Fabric) to store ownership information and 'DNA' of in-game assets, enabling transparency and consistency of asset attributes and ownership changes. This also enables verification of the asset's rarity, since new assets cannot be created in secret. For the actual asset trading, this ledger would be interconnected with either cryptocurrencies (such as public Ethereum), or other payment methods. The pilot will also use other DLTs to store information and relevant transactions (such as authorisations for accessing IoT

resources) related to advertisement views, IoT beacons, and other IoT devices that interact with the player.

DIDs are used to protect players' privacy. For example, when the player needs to perform some task near the IoT beacon to collect the reward, the player registers with the entity running the challenge (which can be a gaming company or other party) with pseudonymous or anonymous identifiers *X* and *Y*. After the user completes the task, this event is recorded to the "IoT beacon ledger" using player's identifier *X*. An interledger function written by the challenge designer is monitoring this ledger and when such event occurs, it triggers (perhaps after some random delay to prevent correlation attacks) the ownership change in the "Gaming asset ledger" granting asset ownership to the identifier *Y*. In this way parties monitoring the first ledger will not be able to know what kind of reward the player has received for the completion of the task, while parties monitoring the asset ledger wil not know which event triggered the ownership change of the asset.

5.4. Decentralised Energy Data Exchange

As with the previous pilots, SOFIE enables strong auditability while preserving user's privacy through usage of DLTs and DIDs. All authorisation related messages concerning smart meter data are signed with the KSI blockchain. No smart meter data is handled by the SOFIE framework, the data is stored by the data owner or in the data hub. For the actual data exchange, a secure communication channel is created between the participants. SOFIE's semantic representation functionality is used to describe available datasets.

6. Related work

Some existing approaches for solving the IoT interoperability problem rely on creating a new interoperability layer, which is not feasible in most cases, since it requires making changes to the existing IoT platforms. Other approaches, including BIG IoT [20], aim to allow interoperability between IoT systems through an API and Marketplace; however the proposed marketplace is designed to be centralized, limiting its applicability and flexibility. WAVE [21] provides a decentralised authorisation solution for IoT devices using a private Ethereum blockchain and smart contracts, however it assumes that all IoT devices are able to interact with the blockchain, which is not a feasible assumption for many constrained devices.

There are also application-specific approaches utilizing DLTs for, e.g., energy trading [22][23][24]. They often utilise tokens issued by a single party as currency, which can lead to speculation and harm the actual users of the system. While cryptocurrency was the original use case of blockchains, it is important to use separate DLTs for payments and for other use cases, such as asset tracking, logging, etc. In this way, price fluctuations of the cryptocurrency will not affect the cost of, e.g. recording asset ownership changes. Furthermore, performance limitations and transaction cost issues associated with public, permissionless DLTs typically used for cryptocurrencies will not limit other uses of DLTs in IoT systems, which need to be responsive and highly scalable.

Therefore, the existing work does not fully address the need for an open, secure, and decentralised solution for the IoT interoperability problem, which supports existing IoT platforms, enables new open business models, while taking security and privacy into account.

7. Discussion

In order to enable real, efficient, and secure IoT interoperability, the SOFIE approach relies on multiple, appropriate, distributed ledgers glued together using interledger technologies, to provide openness, decentralization, trust, security, privacy, automation, and auditability.

Openness may be undesirable for (myopic) individuals or businesses, but when it comes to the whole society, openness becomes beneficial (or even critical), especially from an economic and business perspective. Openness enables inclusion, preventing powerful players from excluding new entrants (either directly or by creating entrance barriers that are hard to overcome). Two key ingredients of openness are decentralization and trust. DLTs are a successful example of a decentralized system that is robust and does not have a single (or few) point(s) of failure. Some DLTs even rely on distributed governance, thus allowing the evolution of the rules from within the system. Similarly, DIDs facilitate decentralized trust management, rendering overlay applications more secure and also more usable.

Security is of paramount importance for SOFIE and the IoT. For this reason, SOFIE does not focus only on security issues at the level of each individual IoT system (which by itself is critical, since the IoT is bridging the cyber with the physical world, therefore security breaches can lead to major safety issues), but it also considers end-to-end security at the level of the whole system, including the interfacing mechanisms and components, which may be even more susceptible to attacks. SOFIE defence mechanisms consider both internal and external threats, as well as threats originating from interconnected systems, which need to be provided controlled access.

A federation system that includes various actors and even expands across the borders of a single country creates challenging privacy issues, since the privacy policies that govern user data depend on many entities and possibly many jurisdictions, legal systems, and rules. Privacy preservation becomes even more challenging when public DLTs are involved, since not only do all parties have access to all information, but replication makes information access easier, and immutability facilitates correlation. Moreover, since data stored in public DLTs never disappears, future advances in de-anonymisation techniques could compromise currently anonymised data. On other hand, open systems facilitate verifiability and auditability, hence there is a critical trade-off. SOFIE tries to get the best of the two worlds by leveraging pseudonymous, self-sovereign identifiers, such as DIDs, which can be frequently changed.

Automation in SOFIE is provided through smart contracts, which enable automation in a reliable, available, secure, and decentralised manner. For instance, in order to support openness and privacy in access to data and actuation, an automatic process is required to control the access, perhaps complemented by an associated payment mechanism (which can be provided by cryptocurrencies).

Using SOFIE's Federation Adapters, diverse IoT platforms can be integrated into the SOFIE framework, without any modifications to the platforms. All integrated systems can then benefit from the intriguing features of the federation approach, including increased functionality and privacy. SOFIE's federation approach enables interesting extensions to its four real-life pilots. For example, in-game assets could be provided as a reward for providing energy consumption measurements and could be used for paying for electrical vehicle charging. Similarly, "ethical" producers could be rewarded with cheaper and cleaner energy. All pilots are currently being implemented and tested and the results will be presented in future publications.

8. Conclusions

This chapter described how SOFIE utilises interledger and Distributed Ledger Technologies (DLTs) for providing interoperability between IoT platforms, while providing strong security, auditability, and privacy. This work has shown that using DLTs and interledger approaches allows more flexible co-operation among various parties in multiple use cases, such as a food supply chain, electricity grid load balancing, context-aware mobile gaming, and secure smart meter data exchange. The SOFIE solution is tested in four real-life pilots, which also raise interesting cross-pilot interactions. In the longer term, this approach will also enable open data markets and allow the creation of new business models around IoT data.

References

[1] M. Rauchs. Distributed Ledger Technology Systems - A Conceptual Framework. University of Cambridge Report, 2018, available at: https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2018-10-26-conceptualising-dlt-systems.pdf (accessed 4.2.2019).

[2] D. Reed et al. Decentralized Identifiers (DIDs) v0.13 - Data Model and Syntaxes. Draft Community Group Report, July 2019, available at: https://w3c-ccg.github.io/did-spec/ (accessed 31.7.2019).

[3] N. Stifter, A. Judmayer, P. Schindler, A. Zamyatin, and E. Weippl. Agreement with Satoshi - On the Formalization of Nakamoto Consensus. Cryptology ePrint Archive, Report 2018/400, 2018.

[4] N. Fotiou and G.C. Polyzos. Smart Contracts for the Internet of Things: Opportunities and Challenges. European Conference on Networks and Communications (EuCNC), Ljubljana, Slovenia, June 2018.

[5] E. Androulaki et al. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. Eurosys 2018, Porto, Portugal, April 2018.

[6] V.A. Siris, P. Nikander, S. Voulgaris, N. Fotiou, D. Lagutin, and G.C. Polyzos. Interledger Approaches. *IEEE Access*, Vol. 7, pp. 89948–89966, July 2019.

[7] C. Allen. The Path to Self-Sovereign Identity. April 2016. Available at: http://www.lifewithalacrity.com/2016/04/the-path-to-self-soverereignidentity.html (Accessed 18.12.2018).

[8] Blockchain and Identity: Projects/companies working on blockchain and identity. Available at: https://github.com/peacekeeper/blockchainidentity (Accessed 7.11.2018).

[9] 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/EC (General Data Protection Regulation).

[10] M. Sporny, D.C. Burnett, D. Longley, and G. Kellogg. Verifiable Credentials Data Model 1.0 - Expressing verifiable information on the Web. W3C Proposed Recommendation, August 2019. Available at: https://w3c.github.io/vc-data-model (Accessed 31.7.2019).

[11] T. Bragatto et al. Statistical Analysis of Prosumer Behaviour in a Real Distribution Network Over Two Years. IEEE EEEIC/I&CPS 2018.

[12] T.O. Olowu, A. Sundararajan, M. Moghaddami, and A.I. Sarwat. Future challenges and mitigation methods for high photovoltaic penetration: A survey. Energies, vol. 11, no. 7, 2018.

[13] S. Rahman, H. Aburub, M. Moghaddami, and A.I. Sarwat. Reverse Power Flow Protection in Grid Connected PV Systems. IEEE SoutheastCon, 2018.

[14] Y. Kortesniemi, et al. SOFIE Deliverable D2.5 - Federation Framework, 2nd version. August 2019, available at: https://media.voog.com/0000/0042/0957/files/SOFIE_D2.5-Federation_Framework%2C_2nd_version.pdf

[15] D. Lagutin, Y. Kortesniemi, N. Fotiou, and V.A. Siris. Enabling Decentralised Identifiers and Verifiable Credentials for Constrained Internet-of-Things Devices using OAuth-based Delegation. Workshop on Decentralized IoT Systems and Security (DISS 2019), in conjunction with the NDSS Symposium 2019, San Diego, USA, February 2019.

[16] N. Fotiou, V.A. Siris, S. Voulgaris, and G.C. Polyzos. Interacting with the Internet of Things using Smart Contracts and Blockchain Technologies. 11th SpaCCS, Melbourne, Australia, December 2018.

[17] N. Fotiou, V.A. Siris, G.C. Polyzos, and D. Lagutin. Bridging the Cyber and Physical Worlds using Blockchains and Smart Contracts. Workshop on Decentralized IoT Systems and Security (DISS 2019), in conjunction with the NDSS Symposium 2019, San Diego, USA, February 2019.
[18] Y. Kortesniemi, D. Lagutin, T. Elo, and N. Fotiou. Improving the Privacy of Internet of Things with Decentralised Identifiers (DIDs). *Journal of Computer Networks and Communications*, 2019.
[19] A. Antonino. A Privacy-preserving approach to grid balancing via scheduled electric vehicle charging. Master's thesis, Aalto University, September 2019.
[20] BIG IoT - Bridging the Interoperability Gap of the Internet of Things, available at: http://big-iot.eu/ (accessed 6.2.2019).
[21] M.P. Andersen et al. WAVE: A Decentralized Authorization System for IoT via Blockchain Smart Contracts. University of California at Berkeley Technical Report UCB/EECS-2017-234, 2017, available at: http://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-234.pdf (accessed 4.2.2019).
[22] LO3 Energy, available at: https://lo3energy.com/ (accessed 5.2.2019).
[23] Grid Singularity, available at: http://gridsingularity.com/ (accessed 5.2.2019).
[24] SolarCoin, available at: https://solarcoin.org/ (accessed 5.2.2019).