

On the Features and Challenges of Security & Privacy in Distributed Internet of Things

Rodrigo Roman^a, Jianying Zhou^a, Javier Lopez^b

^a*Institute for Infocomm Research, Singapore*

^b*University of Malaga, Spain*

Abstract

In the Internet of Things, services can be provisioned using centralized architectures, where central entities acquire, process, and provide information. Alternatively, distributed architectures, where entities at the edge of the network exchange information and collaborate with each other in a dynamic way, can also be used. In order to understand the applicability and viability of this distributed approach, it is necessary to know its advantages and disadvantages – not only in terms of features but also in terms of security and privacy challenges. The purpose of this paper is to show that the distributed approach has various challenges that need to be solved, but also various interesting properties and strengths.

Keywords: Internet of Things, Distributed Architectures, Security

1. Introduction

The concept of the Internet of Things (IoT) has evolved over time [1, 2, 3]. Nevertheless, its core idea can be summarized in a sentence: ‘A worldwide network of interconnected entities’. In most cases, these heterogeneous entities, ‘things’ (e.g. Human beings and computers, books and cars, appliances and food) have a locatable, addressable, and readable counterpart on the Internet. They can open a communication channel with any other entity, providing and receiving services at any time, any place, and in any way. Many technologies serve as the building blocks of this new paradigm, such as

Email addresses: rroman@i2r.a-star.edu.sg (Rodrigo Roman),
jyzhou@i2r.a-star.edu.sg (Jianying Zhou), jl@lcc.uma.es (Javier Lopez)

wireless sensor networks (WSN), RFID, cloud services, machine-to-machine interfaces (M2M), and so on. Also, this paradigm has a multitude of application domains, such as automotive, healthcare, logistics, environmental monitoring, and many others.

There is no single strategy for realizing the vision of the IoT, as services can be provisioned in various ways. In a centralized approach, application platforms located in the Internet (e.g. cloud services) acquire information from entities located in data acquisition networks, and provide raw data and services to other entities. These application platforms control the whole information flow, and there is little or no support for accessing the information providers directly. In fact, there are multiple industrial solutions that make use of this approach [4, 5]. On the other hand, in a distributed approach, not only the intelligence and the provisioning of services is located at the edge of the network, but also various application platforms can collaborate with each other dynamically.

In the context of the IoT, the importance of the distributed approach as an element of the Future Internet of Things has been previously mentioned in the literature (cf. [1]). However, there have been no explicit analyses of its features and its challenges. In order to understand the viability and applicability of this distributed approach, it is necessary to explicitly know its actual features and major principles, including the benefits and disadvantages. Also, as security and privacy are important factors that will influence the adoption of the IoT paradigm, it is essential to know what are the security and privacy challenges – and benefits – of the distributed approach, and what are the most promising approaches in this field. If the challenges are too complex and the benefits too small, it might make sense to focus mainly on the centralized approach for IoT deployments. The purpose of this paper is to assess and answer these questions.

The structure of this paper is as follows. Section 2 will focus on the analysis of the centralized and distributed approaches. In this section we will summarize the state of the art, introduce a taxonomy of the different approaches, and provide an analysis of the features of these approaches. Section 3 will focus on the analysis of the different security challenges. In this section we will overview the existing IoT security challenges (3.1), introduce an attacker model that can be applied to both centralized and distributed IoT architectures (3.2), and study the main challenges and promising solutions in the design and deployment of the security mechanisms (3.3). Finally, conclusions are presented in Section 4.

2. A Distributed Internet of Things

2.1. Related Work: Government, Academia and Industry

The concept of a distributed IoT is not novel. In fact, various official documents consider it as one of the possible strategies that can push the dream of the IoT into the real world, and it has been explicitly mentioned that the development of decentralized autonomic architectures and the location of intelligence at the very edge of the networks are issues that need to be addressed [2]. Still, some key questions have to be answered to make the most of this strategy in the real world, such as the specific situations on which the network intelligence should be distributed [1]. In order to answer these questions, it is necessary to study the specific requirements of applications. For example, whether an application needs support for distributed ownership of data [3]. This and other issues that have been raised by these governmental studies are being carefully considered by the research community.

There are various research articles that study different instances of distributed IoT architectures. For example, Gomez-Goiri et al. [6] combine the concept of the web of things (using web protocols to implement the IoT) with the concept of triple spaces (using semantic web techniques to exchange knowledge in a distributed local shared space) to create a distributed environment where devices located in two or more spaces can collaborate with each other through Internet services. In another example, which follows a more holistic point of view, Ning et al. [7] describe a heterogeneous system known as U2IoT that comprises two subsystems: Unit IoTs, which are basic local cells that provide solutions for special applications, and Ubiquitous IoT, which comprises the different Unit IoTs plus other managers and controls the collaboration between all entities.

There are also many research projects funded by various government bodies that, directly or indirectly, are studying as of 2012 the needs of a distributed IoT architecture. Precisely, one of these projects, IoT-A [8], is aiming to provide an architectural reference model for the interoperability of Internet of Things systems. Note that such a reference architecture does not mandate how all entities should collaborate, or who should analyze the data and provide the different services. Still, the communication model provides the foundations for the creation of distributed applications, allowing digital entities to directly connect and interact with other digital entities. Moreover, the location of intelligence at the edge of the network is implicitly considered, as digital entities range from simple devices to abstract entities made up of

various distributed devices. Therefore, its building blocks [9] could be used in the future to create fully distributed IoT applications.

Some concrete building blocks, which can help to build a distributed IoT, have been indirectly studied in other research projects. For example, the HYDRA project [10] developed an open source middleware that allows legacy devices to provide web services over the Internet – directly or indirectly. HYDRA also provides some tools that can be used to enable collaboration, such as a device and service discovery interface. This interface can make use of an ontology to describe the available services, achieving semantic consistency. Another project, SENSEI [11], was more focused on providing a consistent interface to access the services of Wireless Sensor Networks (WSN) islands. But it produced other relevant results, such as semantically-enabled resource directories, and local management systems that benefit of the existence of such directories. Finally, other projects, like CUBIQ [12] and SMARTPRODUCTS [13], studied and developed various P2P-based distributed mechanisms, such as a distributed publish/subscribe system and a distributed storage system.

Beyond theoretical research, there are numerous companies and start-ups that are making use of cloud technologies to provide IoT services. The key idea is that all edge devices and intranet of things will send their information periodically to an application platform located in the cloud. This platform stores all the data and provides specialized API interfaces that can be used by 3rd parties to create their IoT applications. There are various approaches for implementing these types of platforms: from closed environments where even the sensors are controlled by the company [4] to more open platforms that allow the integration of external devices and databases [5]. Most of these solutions are completely centralized: edge systems act mainly as data acquisition networks, and application platforms from different vendors are not prepared to interact with each other. Yet there are some platforms that, pursuing the idea of creating private and hybrid clouds, can be deployed in a local environment [14]. These platforms not only enable the existence of local intelligence but also can exchange information and services with external systems, thus they can easily become instances of the distributed IoT.

2.2. A Taxonomy of the Vision

In the previous section, we have seen that there are two principles that have been applied to most distributed IoT architectures: i) the location of

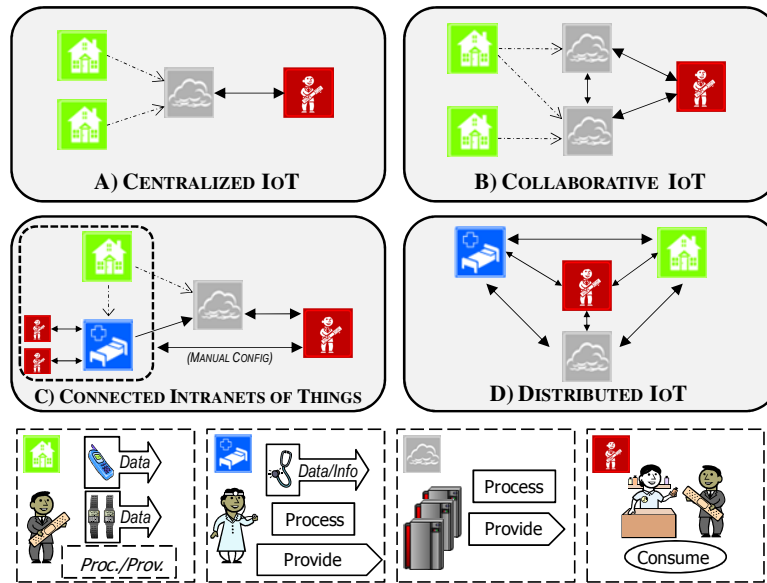


Figure 1: Overview of the centralized and distributed approaches

the intelligence and the provisioning of services at the edge of the network (*edge intelligence*), and ii) the collaboration between diverse entities in order to achieve a common goal (*collaboration*). In fact, these two principles are core elements in the construction of ‘decentralized systems’ and ‘distributed systems’, respectively. In organizational theory, decentralized systems delegate the decision-making authority to entities located in the lower levels. Such delegation can also allow the implementation of any decisions without relying on the approval of high-level entities [15]. On the other hand, a distributed system consists of multiple entities that collaborate with each other and appear to users as a single coherent system [16].

These two principles, edge intelligence and collaboration, can be used to define a taxonomy of possible Distributed Internet of Things approaches, which is presented below. Two of these approaches (collaborative IoT and connected intranet of things) comply with only one of the principles, while a “full” distributed IoT complies with both principles. We will also include the definition of a centralized IoT for the sake of completeness.

- (A) *Centralized IoT*. A Centralized Internet of Things (cf. Figure 1A) does not provide any of the previously mentioned principles. In this scenario, the data acquisition networks (i.e. networks of things such as mobile phones, radiation sensors [17], cars, etc) are passive: their only task is to provide data. All this data will be retrieved by a single central entity, which will process it into information, combine it, and provide it to its customers. Consequently, if users want to make use of IoT services, they must connect through the Internet to the interfaces provided by this central entity. Note that there are various strategies to implement this approach. For example, the central entity can be instantiated using a simple server or a cluster of devices forming a cloud (or even located in the cloud itself, cf. solutions such as [4]). Also, its interfaces can provide both raw and preprocessed data, enabling the creation of more complex 3rd party services.
- (B) *Collaborative IoT*. While in this approach the ‘intelligence’ of the network is still located within the central entities (data acquisition networks still behave as passive entities, users access the information through the central entity interfaces), the main difference with a centralized IoT is its compliance with the collaboration principle. As a result, there are various central entities that can exchange data and/or information with each other, generating new services or enriching existing ones (cf. Fig-

ure 1B). For example, IoT service providers that analyze the radiation in the atmosphere of different cities can collaborate in order to provide a snapshot of the radiation levels in the whole country.

- (C) *Connected Intranets of Things*. In this approach, data acquisition networks (Intranets of Things) can actually process local information, and also provide it not only to central entities but also to local and remote users (cf. Figure 1C). However, there are no underlying mechanisms (e.g. discovery services, ontologies) that facilitate the collaboration between entities. As a result, the information mainly flows from the intranets to a central entity, which will be able to provide a holistic point of view of the whole system. For example, IoT-enabled hospitals need to access the services of a central IoT entity to obtain global information (e.g. overall bed occupancy). Note, however, that if the central entities fail, the local services (e.g. the vital signs records of local patients) can still be accessed.
- (D) *Distributed IoT*. In this vision, all entities can have the ability to retrieve, process, combine, and provide information and services to other entities (cf. Figure 1D). Intranet of things (ranging from personal area networks (PAN) [18] to smart city infrastructures [19]) evolve from isolated entities to fully interconnected systems, not only providing services at a local level but also collaborating with each other and with other IoT architectures towards common goals. Observe that it is also possible to integrate higher-level cloud-based services or other centralized entities (e.g. data repositories) within this architecture, but they are not required. Following the e-health example highlighted above, the IoT of a hospital can interact with the IoT located in the household of a patient, or even with the PANs of the personnel located inside the premises. Moreover, all hospitals can easily collaborate so as to obtain the overall bed occupancy.

2.3. Analysis of Distributed IoT Features

After presenting the taxonomy of the different distributed IoT approaches, this section analyzes their features, pointing out their benefits and disadvantages. This is specially important because, as shown in section 2.1, centralized IoT architectures (mostly based on cloud technologies) are not only gaining momentum but also satisfying the requirements of users as of 2012. Therefore, it is necessary to review and understand the benefits of all these

approaches in order to measure their viability, even if the notion of a distributed IoT has been explicitly mentioned as one of the elements of the Future Internet infrastructure [3, 2, 8]. For this analysis, we will use various requirements and properties of IoT deployments that have been gathered from existing reports and research documents. They are enumerated below:

- *Openness.* Beyond presenting raw data and other specialized services, an IoT platform can also be flexible enough to allow 3rd parties to develop complex applications through the provision of an API.
- *Viability.* This property encompasses two concepts: business model (whether it is viable to market this technology) and vendor lock-in (whether a company can take the long-term risk of depending on a particular provider).
- *Reliability.* Not only the IoT architecture must be resilient enough to assure a certain level of availability, but also needs to provide a performance that is tailored to the specific needs of the applications.
- *Scalability.* Within this paradigm, it is expected that the number of devices and the amount of data generated and processed by those devices will grow exponentially (i.e. the concept of “data deluge”). Thus, we have to take scalability and extensibility into account.
- *Interoperability.* Even if the Internet of Things is inherently heterogeneous, all its components must be able to interact with each other. Therefore, it is necessary to achieve service and semantic interoperability, amongst other things.
- *Data Management.* As the different elements of the Internet of Things produce data, either by sensing or by processing, we must take certain design decisions: where the data should be stored? how the data is accessed?
- *Security Issues.* There are various security issues that must be considered in order to achieve a trusted and fault-tolerant IoT: how to protect the communications? how to manage authentication and access control in a world of billions of things? what about the privacy of the users, and the security of the data generated by the things?

PROP./REQ.	CENTRALIZED IoT	PRINCIPLES	
		(Collaboration)	(Edge Intelligence)
<i>Openness</i>	High (Simple)	High (Multiple APIs)	
<i>Viability</i>	<i>Business Model</i>	Already in market	Similar to hybrid clouds
	<i>Vendor Lock-in</i>	Possible	More choices Access to data sources
<i>Reliability</i>	<i>Availability</i>	Zero if failure	Partial if failure Local data if failure
	<i>Performance</i>	Service level + Latency	Distr. bandwidth Limited latency (local)
<i>Scalability</i>	Limited to cloud resources	Scalable	Scalable at edge
<i>Interoperability</i>	Simple	Complex (global)	Complex (raw)
<i>Data Management</i>	Pull, Data at cloud		Pull, Push (partial)

Table 1: Analysis of properties and requirements of the different distributed IoT principles

Table 1 presents an overview of the features (minus the Security issues) of the centralized IoT approach, together with the features of approaches that follow the collaboration and edge intelligence principles. From the results of this table, it is possible to infer why the **centralized approach** was the first to enter the market. In terms of *openness*, a centralized solution usually provides a small set of (mostly) proprietary APIs for acquiring and providing data. This way, application developers can use these APIs to develop rich and complex IoT applications. Regarding *availability*, most companies build their infrastructures through cloud companies, which usually have a very good service uptime: in 2012 [20] it was 99.99% with a standard deviation of 0.00215%. As for *interoperability*, it is easy to achieve: all data sources will interact with the data acquisition API provided by the centralized system, thus it is only necessary to create one adaptor per data source. Finally, the viability of the *business model* has been proved by the existence of profitable ventures and companies.

Although the centralized approach has a great potential to bring the IoT into life, the other distributed approaches also provide interesting advantages. In the **collaborative IoT approach** (which follows the collaboration principle), the risk of *vendor lock-in* becomes smaller, as customers can combine different service providers to obtain a particular service. *Availability* is improved too: if one of the service providers fails, customers can not only try to search another entity that manages a similar data set, but also use the other providers to retrieve a partial view of the information. In addition, the *scalability* feature is greatly improved by the distribution of the computational and data management resources. Finally, it is important to note that the *performance* of this approach can be slightly worse due to the information

exchange between the different service providers, but the risk of bottleneck decreases.

Regarding the advantages of the **connected intranets of things approach** (which follows the edge intelligence principle), while the risk of *vendor lock-in* remains the same (i.e. a customer uses only one service provider for a particular service), customers can still directly access the interfaces provided by the intranets. This is also beneficial for the *availability* of the services, as customers can still retrieve raw and processed data from the intranets in case a problematic situation arises. There are also some minor improvements in *scalability*, since the complexity of the central entities can be simplified by delegating various processing tasks to the intranets. Note that these advantages are mostly related to external customers (i.e. users that access the information produced by various intranets through a central entity), but this approach also provides specific benefits to local customers, that is, users that make use of the specific local services of the intranet. Firstly, users do not need to be connected to the Internet in order to obtain relevant information about its surroundings. Secondly, in case of failure of the central entities or the Internet connection, the local services are still available. Thirdly, the *performance* of local services is also improved, as services are available through the local communication channels.

Both of the previous approaches also have some disadvantages that are worth mentioning. Due to the interactions between various heterogeneous devices and infrastructures (e.g. intranets providing interfaces to external users, different service providers exchanging information), the underlying mechanisms that are needed to achieve *openness* and *interoperability* (e.g. ontologies, search and discovery, interfaces) are more complex and probably need to be standardized.

Finally, although the **distributed IoT approach** combines most of the advantages and disadvantages of the previous two approaches (superior scalability, limited vendor lock-in, infrastructure complexity), there are some new aspects that need to be explicitly mentioned. For example, in terms of *data management*, the provisioning of data can follow the ‘push’ model (provide only when it is needed), as it is not necessary to provide all data to a central system. In terms of *availability*, the service uptime is more dependant on how many resources are invested in maintaining the underlying IoT infrastructures, but a failure in one element of the infrastructure will not affect the whole system. As for the *business model*, it might be less well-defined in comparison to the model of a centralized IoT, but there are some approaches

that can be taken, such as maintenance fees or management of open source (OSS) services.

All centralized and distributed approaches have their own advantages and disadvantages; and in case of the distributed approaches, the challenges that need to be solved are more numerous (e.g. locate and manage different APIs from multiple and heterogeneous service providers, achieve semantic interoperability). Still, the partially distributed approaches, which comply either with the collaboration or with the edge intelligence principle, provide several interesting advantages such as better availability and higher scalability. Moreover, the distributed IoT approach also allows diverse IoT entities to cooperate even if no central systems are available, amongst other benefits. Therefore, they can be seriously considered as an enabler of future IoT deployments, such as remote healthcare management. Note that all approaches are not mutually exclusive: central entities (e.g. data storage systems) can coexist with distributed IoT entities (e.g. service providers) within the distributed IoT ecosystem.

There is one open question that must be answered, though. Security has not been studied in this section, due to the need of analyzing its impact on all approaches more thoroughly.

3. Security issues in Distributed IoT systems

Although academic research on the topic of security in the Internet of Things is still in its infancy, there is a substantial body of work that analyzes the existing challenges and possible protection mechanisms (cf. Section 3.1). However, existing research mainly provides an overview of the generic problems – without considering the impact of specific features such as the ones studied in this paper. In order to understand the specific security issues of a distributed IoT, it is necessary not only to analyze the impact of the distributed IoT principles (collaboration, edge intelligence) over the existing threats and attacker models (Section 3.2), but also to study the influence of these principles in the integration of the different security mechanisms (Section 3.3)

3.1. IoT Security: An Overview

One of the major challenges that must be overcome in order to push the Internet of Things into the real world is security. IoT architectures are supposed to deal with an estimated population of billions of objects, which

will interact with each other and with other entities, such as human beings or virtual entities. And all these interactions must be secured somehow, protecting the information and service provisioning of all relevant actors and limiting the number of incidents that will affect the entire IoT.

However, protecting the Internet of Things is a complex and difficult task. The number of attack vectors available to malicious attackers might become staggering, as global connectivity (“access anyone”) and accessibility (“access anyhow, anytime”) are key tenets of the IoT. The threats that can affect the IoT entities are numerous, such as attacks that target diverse communication channels, physical threats, denial of service, identity fabrication, and others [21]. Finally, the inherent complexity of the IoT, where multiple heterogeneous entities located in different contexts can exchange information with each other, further complicates the design and deployment of efficient, interoperable and scalable security mechanisms.

Some of the previously mentioned challenges, alongside with the security mechanisms that should be integrated into the Internet of Things, have been already enumerated by the research community [3, 22, 23]. They are as follows:

- Heterogeneity has a great influence over the **protocol and network security** services that must be implemented in the IoT. Constrained devices will interact with various heterogeneous devices (e.g. other constrained devices, full-fledged web servers) either directly or through gateways. In this scenario, not only it is essential to implement efficient *cryptographic algorithms* that can provide a high throughput even in 8-bit or 16-bit devices, but also to adapt or create lightweight *security protocols* that offer an end-to-end secure communication channel. These protocols require credentials, thus optimal *key management systems* must be implemented to distribute these credentials and to help in establishing the necessary session keys between peers.
- The existence of billions of heterogeneous objects also affects **identity management**. Beyond defining the actual scope of ‘identity’ in this context (e.g. underlying identity vs real identity, core identity vs temporary identity), we also need to provide some mechanisms for achieving universal *authentication*. Without authentication, it will not be possible to assure that the data flow produced by a certain entity contains what it is supposed to contain. Another important aspect related to authentication is *authorization*. If there is no access control

whatsoever, everything will be accessed by everyone, which is neither viable nor realistic.

- In fact, the data deluge caused by billions of entities creating information is a big threat to **privacy**. Users must have tools that allow them to retain their *anonymity* in this superconnected world. Other tools must provide a snapshot of the information and policies surrounding a particular user, enabling *transparency* and preventing the notion that the IoT is silently controlling our lives. In fact, the IoT itself must seriously consider the implementation of the *privacy by design* principles [24], providing user-centric support for security and privacy from its very own foundations.
- The size and heterogeneity of the IoT also affects its **trust and governance**. There are actually two dimensions of *trust*: a) trust in the interaction between entities, where we have to deal with uncertainty about the future actions of all collaborating entities, and b) trust in the system from the point of view of the user, as users must be able to manage their things so as to not feel under some unknown external control. Regarding *governance*, it is a double-edged sword that must be wielded with care. On the one hand, it offers stability, support for political decisions, and the possibility to define common frameworks and interoperability mechanisms. On the other hand, governance can easily become excessive, fostering an over-controlled environment.
- The number of vulnerable systems and attacks vectors will surely increase in the context of the IoT, thus **fault tolerance** becomes essential. Not only we must strive for *security by default* (robust implementations, usable systems, etc) in the IoT, but also we need to develop *awareness* mechanisms that can be used to create the foundations of *intrusion detection and prevention mechanisms*, which will help IoT entities to protect or even gracefully degrade their services. Finally, *recovery services* must be able to locate unsafe zones (i.e. zones affected by attacks) and redirect the functionality of the systems to other trusted zones.

3.2. Analysis of Attacker Models and Threats

As aforementioned, in order to understand how the different approaches presented in section 2.2 should be secured in the future, it is firstly necessary

to enumerate and analyze the attacker models. These models have been defined in a way that they can be applied to both centralized and distributed IoT approaches. Note, however, that the concept of ‘perimeter’ in the Internet of Things is a bit fuzzy: an attacker can control part of the network, but due to the inherent distributed nature of the IoT, it is nearly impossible for an attacker to fully control the whole system. As a result, an attacker can be both ‘internal’ and ‘external’ at the same time. These attacker models, categorized by threats, are introduced in the following paragraph.

- *Denial of service (DoS)*. There are a wide number of DoS attacks that can be launched against the IoT. Beyond traditional Internet DoS attacks that exhaust service provider resources and network bandwidth, the actual wireless communication infrastructure of most data acquisition networks can also be targeted (e.g. jamming the channels). Malicious internal attackers that take control of part of the infrastructure can create even more mayhem.
- *Physical damage*. This threat can be seen as a subset of the *DoS* threat. In this attacker model, active attackers usually lack technical knowledge, and can only hinder the provisioning of IoT services by destroying the actual ‘things’. This is a realistic attack in the IoT context, because things might be easily accessible to anyone (e.g. a street light). If that is not possible, the attacker can simply target the hardware module in charge of creating the ‘virtual persona’ of the thing.
- *Eavesdropping*. Passive attackers can target various communication channels (e.g. wireless networks, local wired networks, Internet) in order to extract data from the information flow. Obviously, an internal attacker that gains access to a particular infrastructure will be able to extract the information that circulates within that infrastructure.
- *Node Capture*. As aforementioned, things (e.g. household appliances, street lights) are physically located in a certain environment. Instead of destroying them, an active attacker can try to extract the information they contain. Note also that, instead of things, active attackers can also target other infrastructures that store information, such as data processing or data storage entities.

- *Controlling.* As long as there is an attack path, active attackers can try to gain partial or full control over an IoT entity. The scope of the damage caused by these attackers depends mainly on a) the importance of the data managed by that particular entity, b) the services that are provided by that particular entity.

While both centralized and distributed approaches share the same attacker models, there are subtle differences caused by the distributed IoT features and principles. They change various aspects of the underlying infrastructures, such as the deployment strategies of the different IoT entities, the actual information flow, and the availability of certain interfaces and services. Such changes can create new threats and facilitate the work of attackers, but also can reduce the effectiveness of certain attack vectors. In the following paragraphs we will discuss the different aspects that are influenced by the distributed IoT features and principles, and how they impact the threats and attacker models.

One aspect is the *centralization of resources*. Most adversaries will aim to target systems that provide the biggest payoff, and central entities fall under this category – they store, manage, and process a huge amount of information. Theoretically, these central entities will have better protection mechanisms, but any vulnerability can make the whole system fall apart. On the other hand, if the actual intelligence of the Internet of Things is distributed, the information will be created and processed in different entities, thus adversaries need to redouble their efforts in order to control the same amount of resources. However, the distribution of resources is a double-edged sword. If the adversary is only interested in a specific piece of information, it can target the system that manages that particular information – which might not be as protected as a central entity. Besides, node capture attacks become more dangerous, as more logic is placed within the things themselves. In fact, an adversary can use a guerrilla warfare strategy and gradually take control of small parts of the network, so as to affect the whole system in a covert way.

Another aspect, related to the centralization of resources, is the nature of the *information flow*. In centralized IoT deployments, the information flow will follow a hierarchical pattern, as a central entity will receive information from every ‘thing’. On the other hand, in more distributed approaches, the information flow will resemble a peer-to-peer system, where information is only exchanged when needed. In this particular case, an adversary that

eavesdrops on a section of the network will not be able to obtain a holistic point of view of the whole system. There is a caveat here: if an adversary targets an intranet of things (e.g. an IoT-enabled hospital) in a distributed scenario, he might be able to retrieve processed information instead of raw data.

Regarding the *overall connectivity of the network*, in approaches that follow the edge intelligence principle, constrained entities are expected to be directly locatable and addressable via the Internet. Therefore, they must be able to accept connections from external entities. This situation allows malicious adversaries to launch attacks that can easily exhaust their resources. Observe that this situation can also arise in networks (either centralized or distributed) with actuators (e.g. electric motors, industrial machinery), as the behavior of actuators can change when receiving orders from remote administrators. Note also that it is possible to implement additional protection mechanisms to control these incoming connections, such as firewalls and additional middleware layers.

Finally, we also have to consider the *user involvement* in the configuration of the security mechanisms. User-centric networks, such as personal area networks [18], are one of the elements of the Internet of Things. By pushing the intelligence onto the edge of the network, it is possible for the owners of these networks to create and manage their own policies (cf. section 3.3.4). However, as most users are not experts, mistakes will happen if the security mechanisms are not usable enough. Such misconfigurations can be exploited by malicious adversaries to access personal data or even take control of that particular user-centric network. Note that in centralized entities the configuration of the security mechanisms will be made by experts, but any misconfiguration will create a very rewarding window of opportunity that can be exploited by any knowledgeable adversary.

By reviewing these attacker models, we can conclude that no approach is better in terms of threats and attacker models – all of them have various advantages and disadvantages. In a centralized IoT the central entity becomes a single point of failure; and although the number of attack vectors are smaller (and the protection mechanisms might be better), a single vulnerability or a misconfiguration can cause extreme damage to the whole network. If the resources of the network are distributed, the impact caused by a successful attack will be smaller, but the number of attack vectors will increase. Note that in all approaches there will be a huge number of data providers, the things, that can be highly constrained and physically accessible – in other

words, easy targets. Therefore, it is clear that any IoT application will have to deal with a certain amount of bogus data.

3.3. Specific Challenges and Promising Solutions

Once the analysis of the threats and attacker models is finished, we can study what are the main challenges in the design and deployment of the security mechanisms. Such study, which will be performed in the next sections, will help to point out specific problems that must be considered if we want to bring the distributed IoT architectures to the real world. Moreover, within this study, we will explore not only existing IoT security mechanisms, but also promising approaches that could be used to provide security in a distributed IoT environment.

3.3.1. Identity and Authentication

It is essential to consider how to manage identity and authentication in the Internet of Things, as multiple entities (e.g. data sources, service providers, information processing systems) need to authenticate each other in order to create trustable services [25]. When defining these security mechanisms, we also have to consider some of the inherent features of the Internet of Things. As interactions can be quite dynamic, the entities of the network might not even know in advance which partners can be used to create a certain service. Vehicular networks (VANETs [26]) are an example of this: cars are expected to provide data not only to devices located on the roadside but also to other cars. Besides, if billions of things are going to be interconnected, it is necessary to manage their identities in a scalable way.

In a *centralized IoT architecture*, some of these challenges are inherently more simple. In this particular approach, the application logic is mainly located in one central entity (e.g. a cloud-based IoT application platform) that provides a limited set of well-known entry points (e.g. APIs). Both data providers, such as sensors, and information consumers, such as user applications and other customers, connect to this central entity. As a consequence, all the authentication logic can be centralized in this entity or in an identity provider associated with it. In case there are data providers that have their own identity provider, there are no scalability problems, as such identity providers can establish a relationship of trust with the central entity (a N-to-1 scenario). Note that if an IoT complies with the collaboration principle (Collaborative IoT), it might be possible to make use of a federated

identity management system, where all the service providers belong to the same circle of trust.

This simplification cannot be found in purely *distributed IoT architectures*, which fulfill both the collaboration and edge intelligence principles. In this context we find a dynamic N-to-N scenario, where data providers are no longer passive and are able to acquire and process information from other sources. Moreover, due to the edge intelligence principle, local users can query local information providers directly, without intervention from external entities. As a result, some kind of authentication logic must be present in every service provider – including the tiniest of objects. Note, however, that things do not exist in a vacuum: they usually belong to a specific group, are located in a particular context, and are owned by certain entities. These aspects must be taken into account.

Promising approaches. As aforementioned, it is essential to manage the identities of the things in a scalable way. However, as of 2012, there are various mechanisms that can be used to identify things uniquely, such as the tag code standards EPC and ucode [27]. Therefore, it is expected that in the future various systems will coexist – not only at the universal level but also at a local scale [28]. Note, however, that in many scenarios the ‘who’ is less important than the ‘where’ and the ‘what’. As a consequence, things should be able to identify themselves using their attributes and their context (e.g. radiation sensor #2044A can simply state that is a radiation sensor located in Shibuya, Tokyo).

Regarding things authentication, we have to consider that in many scenarios things belong to a certain group (e.g. intranets of things, personal area networks) located in the same spatial area (e.g. IoT-enabled hospitals). In such environments, local identity providers can manage the identities of those things, and also can create a circle of trust with relevant external resource providers (e.g. the household of a chronic patient, other hospitals). Consequently, local entities are not only able to authenticate to each other within the group, but also can provide a proof of identity when interacting with external entities. Also, external entities can receive a temporary persona (e.g. long-term patient) from the local identity provider if necessary. This group-based strategy has been, in fact, partially considered in the interactions between WSN islands, where interdomain collaboration is possible through federated identity management and access tokens translation [29]. Traditional Web 2.0 SSO such as OpenID and Shibboleth could also be used

in this situation, although it should be noted that they were not designed to fulfill certain IoT requirements such as identity disclosure (i.e. support for privacy) [30], thus more analyses are needed.

If the thing is actually a human being, it can also be possible to use existing authentication mechanisms (e.g. web credentials, electronic identity cards) if the resource provider understands them. For example, Guinard et al. [31] proposed a smart gateway infrastructure (Social Access Controller, or SAC) that allows users to retrieve data from local sensors using their social network (e.g. Facebook) credentials. Note that this approach might not work in case the human being does not directly interact with the IoT entities. In such cases, it is necessary to develop surrogate mechanisms that can act on behalf of the human users. One existing example is the concept of the Minimal Entity (ME) [32] – a device that stores the digital identity of the user and acts as his representative in the virtual world. Not only it does provide end-to-end secure communication and collaboration with anonymous receivers, but also allows the implementation of pseudonyms. Another example is the concept of a digital shadow [33], where users can delegate their credentials (including access control credentials) to multiple objects or virtual entities.

3.3.2. Access Control

In the Internet of Things, the challenges related to Access Control are closely related to those found in any distributed system. A particular service is constructed by aggregating several services and data sources from different locations and contexts (e.g. a hospital retrieving information from home patients and ambulances). All these information providers will have their own access control policies and permissions whose life cycle (creation, enforcement, maintenance, translation) needs to be managed.

There are also some specific issues that must be taken into account in the context of the IoT. Granularity (i.e. providing more information to people with the right credentials) and location (i.e. checking whether users are accessing the services of a thing locally or remotely) become important elements of the access control policies in certain scenarios. For example, in case of an accident, everyone at the crash site can access my blood group, but only certified doctors and nurses can access my vital signs. Also, whenever access control mechanisms are implemented at the thing level, it is necessary to consider the amount of computational resources that are available, as constrained devices might not have enough space to implement a complex access control mechanism. Finally, as many things are owned by their users (either

permanently or temporarily) and may belong to a group (e.g. personal area network), it is necessary to consider the design of delegation mechanisms, as these things may act in the name of the user / group.

As with authentication, access control policies are easier to manage in *centralized IoT architectures*: all access control policies are stored and managed within a single central entity. Therefore, data providers do not need to implement any kind of access control logic: they will send all their data to those whom they trust (i.e. the central entity). As a side effect of this configuration, both data providers and information consumers must completely trust the central entity, as it will store the information generated by all network entities. On the other hand, purely *distributed IoT architectures* have to deal with all previously mentioned challenges: management of heterogeneous policies, multiple enforcement points, etc. Nevertheless, as will be mentioned in Section 3.3.4, the overall privacy of the network improves once the things can control directly who accesses their own data.

Observe that additional mechanisms must be implemented whenever the collaboration principle is applied to centralized IoT architectures (e.g. tools for maintaining consistency between access control lists, resource delegation mechanisms). Note also that we need to manually configure the direct links between the intranets and the external entities in networks that only comply with the edge intelligence principle.

Promising approaches. There have been very few advances in the management of access control policies for distributed IoTs. In fact, it is not trivial to apply existing access control approaches to completely distributed environments. For example, there are scalability and consistency issues when storing the list of users and their associated access rights in access control lists (ACLs). Role-based access control (RBAC) mechanisms need to define the different roles that users can take, which might be different in various contexts even if they refer to the same type of entity (e.g. custodian vs. janitor). Finally, RBAC policies that use attribute certificates [34] need of an infrastructure that allows validating such certificates in a cross-domain environment. Note, however, that due to the specific features of the Internet of Things, it is possible to consider certain factors such as context as part of the access control model [35]. As a consequence, with adequate technological support, certain policies (e.g. only authenticated users located within my vicinity during working hours can access today's reports) can be easily implemented.

Besides, there are various simple strategies that could be used whenever the things belong to a certain group (cf. Section 3.3.1). For example, the access control logic could be pushed to specific trusted entities, which will act as token-granting services à la Kerberos (i.e. a thing will grant access to anyone that has a valid signature created by a trusted entity). In another approach, the access control logic can be implemented within the things themselves, but relying only on locally-defined roles (e.g. a doctor from another hospital must retrieve his locally-issued role before interacting with the local things). A drawback of all these strategies is that users must first access the trusted entity before requesting information from the things.

3.3.3. Protocol and network security

A secure communications channel is, in most cases, a byproduct of a successful authentication (e.g. server authentication or mutual authentication using protocols such as TLS/DTLS). This process will make use of certain user credentials, such as shared keys or X.509 certificates. If there is a limited set of well-known *centralized application providers* (i.e. central entities), the distribution and management of these credentials becomes easier, as it is possible to preload information in the devices. However, in *distributed IoT architectures*, extra challenges arise: any entity can connect with any other entity at any time, these entities might not know each other in advance, and also limited devices can exchange information with other limited devices. Therefore, in this scenario key management becomes a significant problem.

There are some additional challenges related to the computational resources available to things. When opening a secure channel, devices should be able to negotiate the actual parameters of that channel, such as algorithms (e.g. RSA vs. ECC), strength (AES-128 vs. AES-256), and protection mechanisms (only integrity vs. confidentiality and integrity). The first reason is obvious: constrained devices might not be able to implement certain configurations. There is another reason, though: adaptability. Depending on various factors such as the level of criticality of the data, it might not be necessary to apply strong protection mechanisms to a particular information flow (e.g. confidentiality and the on/off status of a street light). Another challenge is the need to analyze the number of security protocols that can be implemented within a constrained device. In fact, it is necessary to carefully study whether existing Internet protocols should be adapted to this context or not. Finally, things that can be accessed directly (e.g. in the distributed IoT approach) need to be careful about the overhead caused by incoming con-

nections (e.g. multiple incoming connections that require the use of public key cryptography).

Promising approaches. As the Internet of Things inhabits the Internet ecosystem, it is important to provide support for existing security protocols. In fact, the security of IoT-designed web transfer protocols, such as CoAP (Constrained Application Protocol), is largely dependant on the implementation of these security protocols [36]. Some protocols can be implemented without any major changes. For example, there are commercial implementations available of DTLS for constrained devices [37]. However, other protocols need to be adapted due to the complexity of their design. Such protocols must achieve a tradeoff between simplicity and compatibility. For example, one approach seeks to apply IPsec to constrained environments by balancing link-layer security and IPsec security (cf. Raza et al. [38]).

As for the distribution of the credentials, there are various strategies that could be used to tackle this problem. As aforementioned (cf. Sections 3.3.1 and 3.3.2), whenever things belong to a particular local group, it is possible to have one or various entities in charge of managing and distributing the credentials. Also, in scenarios where clients and servers know each other in advance, it is also possible to use certain symmetric key-based protocols, which can provide good properties such as high resilience to attacks [39]. Finally, beyond the optimization of these security protocols, there are various researchers that are pursuing the implementation of fast and compact cryptographic algorithms. There are various research areas, which are not mutually exclusive: from the design of novel hash functions and symmetric algorithms [40] to the optimization of existing primitives [41].

3.3.4. *Privacy*

Up to this point we have seen that a distributed IoT architecture requires more complex security mechanisms. There is, however, one area where *distributed IoTs* provide immediate benefits: Data management and privacy. The core idea is that, due to the edge intelligence principle, every entity has more control over the data it generates and processes. There are several consequences of this approach. Firstly, entities can control the granularity of the data they produce. For example, a portable radiation sensor can announce that it is located in a certain area without providing its exact coordinates. Secondly, entities can define their own access policies. The previously mentioned object can provide the city where it is located (Tokyo) to anonymous

entities, the area where it is deployed (Shibuya, Tokyo) to entities with adequate permissions, or even detailed GPS location information to local entities in case of emergencies. Thirdly, entities do not need to provide all the data they produce, only the data that is needed by the external entities for a particular service. This is closely related to privacy, as it will be more difficult to create a profile of a certain entity if not all information is available.

As for *centralized IoT architectures* (including those who comply with the collaboration principle), a data provider can also decide whether to share or not a particular data stream. Still, as the intelligence is located on the central entity, the type of services it provides will be limited to the amount of data it receives. Another approach can be used if the centralized architecture complies with the edge intelligence principle: as data providers and information consumers are able to communicate directly, they might negotiate a set of secret keys in order to protect their information. However, in this case the central entity cannot process the data, thus it becomes a simple storage system unless it implements advanced cryptographic mechanisms that can manipulate encrypted data, such as homomorphic encryption.

As a final note, we have to point out that the previous paragraphs focus mainly on the protection of personal information, but there is another dimension of privacy that is especially relevant in the context of the IoT: the existence of entities that profile and track users without their consent. Here, the benefits of a distributed IoT might turn into nightmares when misused. By following the edge intelligence principle, entities can adapt their behavior and track users more effectively. Moreover, thanks to the collaboration principle, these entities can share up-to-date information about the target.

Promising approaches. The distributed IoT approach facilitates the implementation of the privacy-by-design principles [24], as all entities can directly manage their own data. However, it is necessary to go beyond the implementation of user-centric access policies and mechanisms to control the granularity of the provided data. Whenever human beings are involved, aspects such as the usability of the user interface (e.g. what can be accessed and to what extent [42]) should be taken into account. As data will be distributed amongst various entities, it also is necessary to study the applicability of existing privacy-preserving distributed data mining algorithms [43]. For example, certain privacy enhancing technologies (PETs) [44] such as multi-party computations [45] can be used to provide protection to some cooperative protocols (e.g. cooperative benchmarking and forecasting). For

especially sensitive data, advanced concepts such as active bundles (i.e. a container with a payload of sensitive data, metadata, and a virtual machine (VM) [46]) might be used. Finally, the legal privacy regulations should be revised to fully consider the intricacies of an always connected Internet of Things [47].

Regarding the problem of user tracking and profiling, there are some ongoing efforts in the research community that aim to provide solutions for this particular threat. For example, there is an interesting perspective that considers a local environment as an operative system [48]. In short, incoming and outgoing items need to be scanned for rogue devices and malicious software that can threaten the privacy of the user. This can be achieved by using mechanisms such as the privacy coach [49]. However, as users could be tracked anywhere and anytime, these concepts should be extended in order to help users to become more aware of how their surroundings capture and use their information. Frameworks like uTRUSTit (cf. Section 3.3.5) might help in this area. Besides, existing studies on surveillance systems such as CCTVs [50] might also provide a clue on the specific legal challenges that our society will face once the Internet of Things becomes a reality.

3.3.5. Trust and Governance

There are other areas where both centralized and distributed IoT approaches have their own specific advantages and disadvantages. One of those areas is Trust Management. As aforementioned, in the IoT we can consider two dimensions of trust: trust in the interactions between entities, and trust in the system from the users' perspective. In a *centralized IoT*, uncertainty comes from the interactions with the data providers ('Which data is more reliable and fresh?'). The holistic point of view of a central entity can help in calculating the reputation of other entities (e.g. a radiation sensor cannot give a warning if all sensors in the vicinity provide a low value). However, if different central entities collaborate with each other, they must be able to exchange trust information in order to fix inconsistencies in the reputation values. In a *distributed IoT*, there is uncertainty in both the interactions with the data providers and the interactions with the service providers ('Who can give me a robust and timely service?'). The distributed infrastructure makes the management of trust more complicated: how can reputation and trust be calculated and shared? Which ontology should I use? Can I trust the reports from other systems? Still, these trust management systems can make better use of second-hand information sources: when a certain entity is given a low

reputation, this reputation can be propagated to other entities that might interact with such an outlier in the future.

As for the trust in the system, it is largely dependant on knowing the internal state of the Internet of Things that surround us. In a *centralized IoT* not all information will be available: in order to provide services, a central entity is more interested in retrieving physical and entity data instead of status and network data. Still, if a centralized system provides an additional ‘internal status’ service, it can be able to supply this kind of information very quickly, as a) it stores internally most of the information from the data providers and b) if fresh data is not available, it can send immediate queries to the specific data providers. As for a *distributed IoT*, this kind of service is more complex and needs more time to be completed, as relevant data providers must be discovered and queried. Nevertheless, the more intelligence at the edge of the network, the more relevant information (e.g. network status, existing connections between entities) that can be retrieved. This way, it can be possible to have a more accurate picture of the status of the whole system.

Regarding the issues associated with Governance, it is not clear how this problem will be solved in the context of the IoT [51], although the *distributed IoT approach* can provide some solutions. As policies in a distributed IoT can be defined at the edge of the network, it could be possible to implement and enforce certain rules such as limiting the countries that can access to our data. This is not possible in a centralized system, whose data servers will be located in most cases in foreign countries. Still, as with many other security mechanisms, a distributed IoT needs to implement various distributed mechanisms to control and enforce these policies, which is not trivial. This same problem affects the management of accountability in the IoT. As logging subsystems will be distributed throughout the network, it will be more difficult to retrieve all the relevant information that might be needed for forensic analysis. There is one clear benefit, though: if a balance between accountability and privacy is achieved, it will be possible to pinpoint the source of a particular problem thanks to the detailed information about the behavior of the system.

Promising approaches. There are some theoretical studies that analyze the suitability of trust management systems for the IoT. For example, Køien [52] points out that subjective logic systems such as TNA-SL [53] can capture dynamic environments where beliefs and uncertainties change over time. There

are other open issues that the state of the art needs to address, such as the management of trust without central authorities. Still, it might be possible to develop preliminary solutions for such problems by analyzing how they are solved in the building blocks of the Internet of Things (e.g. sensor networks, ad hoc networks). The reason is simple: these building blocks have several features in common with the distributed IoT approach. For example, ad-hoc networks are dynamic environments where the network is created, operated and managed by the nodes themselves. In such networks, the decentralization of trust, which is essential in a distributed IoT context, has been extensively studied [54]. Moreover, there are other holistic paradigms closely related to the Internet of Things, such as ambient intelligence and pervasive computing, whose existing works in the area might also provide additional information on how to deal with multidisciplinary challenges [55].

All the previous approaches do not consider the interactions between human beings and IoT entities. In fact, as the Internet of Things can (and will) contain user-generated content, we also have to ask ourselves how to model this type of trust. One promising approach is the existence of user-managed circles of trust, as described in the shoppingLense system by Robinson et al. [56]. This system increases users' trust in the IoT by including trusted metadata in the information flow. In particular, patterns (e.g. QR codes) located in the environment (e.g. shopping mall) are digitally signed and owned by a user-defined group. Members of that group can also add ratings to a particular pattern. This way, if a user trusts a particular group, it can acquire both information from the pattern and trusted ratings from other users. Finally, regarding the trust in the system from the users' perspective, one particular research project (uTRUSTit [57]) has already produced promising results in this area. In particular, the framework developed in this project not only provides an inventory of the local devices that are connected to the Internet of Things, but also enables users to know their status, allowing the creation of a mental model of the virtual world.

3.3.6. Fault Tolerance

Regardless of the approach, centralized or distributed, there is an expected population of billions of things that will act as data and information providers. Such things can become faulty and stop working, but they also can send bogus or even manipulated data. As mentioned in Section 3.2, it is unrealistic to assume that a data processing entity will never have to deal with such problematic data. Therefore, in the IoT context, it is essential

to consider fault tolerance. We must not only aim to provide a ‘best-effort’ service in case parts of the network are not accessible, but also assume that every entity can receive bogus information from other entities.

In case one of the things fails and stops sending data, it is necessary to discover another thing that can provide a similar set of data. In *centralized IoT architectures* this task is more simple, as the central entity will have access to all data flows. As for *distributed IoT architectures*, they need to develop a discovery mechanism that is able to pinpoint related data flows. Note that additional mechanisms need to be implemented in order to assure the survivability of the network in case of a failure of part of the infrastructure: not only data providers need to be located, but also service providers and data processing entities as well.

As for the existence of bogus data, it is possible to develop holistic (centralized) and detailed (distributed) mechanisms that deal with this problem. A *centralized system* can analyze the consistency of the data, pinpointing data providers who seem to behave erratically. A *distributed system* can make use of the additional information (e.g. network information) retrieved at a local level or in the interactions with other entities to apply advanced intrusion detection systems. Both approaches have their own challenges, but they are not mutually exclusive (e.g. in a distributed environment there can be certain entities that provide high-level services and behave like cloud-based IoT infrastructures), thus it is advisable to take full advantage of both of them if possible.

Promising approaches. As of 2012, there are almost no explicit analyses on the mechanisms that could be used to provide service survivability in the IoT. Still, there are various research approaches that can be used as a foundation to enable such fault tolerance. For example, the tools that allow human users to create a mental model of their surroundings (highlighted in Section 3.3.5) can also be used by the network entities to discover devices that are faulty. There are also various theoretical platforms whose aim is to provide service look-up, discovery and composition mechanisms for the Internet of Things [58]. However, it is necessary to study their applicability in an heterogeneous distributed environment. The use of local clusters can help with this task: if entities are clustered in local groups, that cluster can incorporate mechanisms that not only provide up-to-date information about local things, but also enable the interaction of different service discovery protocols through specialized middleware [58]. Besides, all these services can make use

of the functionality provided by existing security mechanisms such as trust management (e.g. only reports from trusted entities will be considered, zones with high reputation will take care of the extra workload).

Regarding the detection of bogus data and malicious entities, most existing intrusion detection mechanisms and rules focus on internal adversaries that try to attack the specific protocols of data acquisition networks (e.g. sensor networks) [59], but do not consider attacks that target the interactions between different IoT domains (e.g. a DoS attack or a malformed packet attack targeting a smart door service [60]). In fact, the state of the art on this specific area is very limited and only few works are available [61, 62]. It is then necessary to implement new detection mechanisms that take into account the distributed IoT specific attacker models. Note that it is also possible, in certain scenarios, to adapt existing mechanisms. For example, centralized entities can make use of clustering-based mechanisms and other data mining techniques to detect outliers and intrusions [63]. Moreover, lessons might be taken from existing distributed intrusion detection systems implemented in similar environments such as smart grids [64].

3.4. Summary

Security Challenges	Centralized IoT	Distributed IoT
<i>Identity and Authentication</i>	N-to-1	N-to-N
<i>Access control</i>	Homogeneous policies	Heterogeneous policies
<i>Protocol and Network Security</i>	Known centralized provider	Unknown peers
<i>Privacy</i>	Less flexible	More flexible
<i>Trust management</i>	Holistic point of view	More detailed information
<i>Governance</i>	Less flexible, more simple	More flexible, more complex
<i>Fault tolerance</i>	Holistic point of view	Detailed point of view

Table 2: Analysis of security challenges in different IoT strategies

A summary of the challenges studied in the previous sections is shown in Table 2. We can conclude that the decentralized and heterogeneous nature of the distributed approach increases the complexity of most security mechanisms (*Identity and Authentication*, *Access Control*, *Protocol and Network Security*, *Trust management* and *Fault Tolerance*). Still, there are some security mechanisms (*Privacy*, *Trust management* and *Governance*, *Fault Tolerance*) where i) the distributed approach provides interesting features, ii) both approaches (centralized and distributed) can complement each other.

In the previous sections we have also highlighted various strategies that could be used in the near future to design and deploy IoT-specific security mechanisms. One such strategy assumes that things belong to a certain group (intranet of things, personal area network) located in a certain spatial area (IoT-enabled hospital, household). These groups comply with the edge intelligence and collaboration principles, thus they are part of the distributed IoT. Once the things are grouped, the implementation of certain security mechanisms becomes easier: local identity providers can be defined, the access control logic can be pushed onto specific entities, a mental model of the virtual world can be created, and so on. Note that this strategy might be partially applicable to highly dynamic environments such as VANETs [26] if we consider the existence of logical groups (“all cars that have been registered in Singapore”), although more research is needed to validate this point of view.

Other strategies focus on the interactions of human users with the Internet of Things. For example, as digital social infrastructures have been already deployed, they can be used in the implementation of specific security mechanisms such as user-defined access control and circles of trust. Finally, another strategy consists of adapting the security mechanisms that have been developed in i) the building blocks of the Internet of Things (e.g. sensor networks, ad hoc networks) and ii) other paradigms closely related to the Internet of Things (ambient intelligence, pervasive computing). Note that while the building blocks lack the complexity of the distributed IoT approach, they share certain similarities such as the decentralization of resources. In fact, some security mechanisms, such as Key Management, have been successfully adapted to certain IoT scenarios [39].

4. Conclusions

The main goal of this paper was to provide an explicit analysis of the features and security challenges of the distributed approach of the Internet of Things, in order to understand what is its place in the Future Internet. There are numerous challenges that must be solved, such as assuring interoperability, reaching a business model, and managing the authentication and authorization of entities. Still, there are multiple benefits as well. Since intelligence is not concentrated on a limited set of centralized application platforms – although these platforms can also exist in order to provide additional support – scalability is improved. Data is managed by the distributed

entities, thus it is possible not only to push/pull data only when needed, but also to implement specific privacy policies. Besides, additional trust and fault tolerance mechanisms can be specifically created for this approach. These and other benefits show that this approach is actually useful and applicable to the real world. As a final note, we would like to stress that both centralized and distributed approaches can coexist with each other, providing the foundations of a full-fledged Internet of Things.

Acknowledgement

This work was partially supported by the Spanish Ministry of Science and Innovation through the ARES (CSD2007-00004) and SPRINT (TIN2009-09237) projects. The latter is cofinanced by the European Regional Development Fund (FEDER).

References

- [1] INFISO D.4 Networked Enterprise & RFID INFISO G.2 Micro & Nanosystems, in co-operation with the Working group RFID of the ETP EPOSS, Internet of things in 2020: Roadmap for the future, 27 May 2008.
- [2] CERP-IoT Cluster, Visions and Challenges for Realising the Internet of Things, European Commission, 2010.
- [3] O. Vermesan, P. Friess, P. Guillemin, S. Gusmeroli, H. Sundmaeker, A. Bassi, I. S. Jubert, M. Mazura, M. Harrison, M. Eisenhauer, P. Doody, Internet of Things Strategic Research Roadmap, Cluster of European Research Projects on the Internet of Things, CERP-IoT, 2011.
- [4] ThingWorx, Accessed on November, 2012. [link].
URL <http://www.thingworx.com/>
- [5] Cosm, Accessed on November, 2012. [link].
URL <https://cosm.com/>
- [6] A. Gómez-Goiri, D. López-de-Ipiña, On the complementarity of Triple Spaces and the Web of Things, in: 2nd International Workshop on Web of Things (WoT'11), San Francisco, USA, 2011.

- [7] H. Ning, H. Liu, Cyber-Physical-Social Based Security Architecture for Future Internet of Things, *Advances in Internet of Things 2 (1)* (2012) 1–7.
- [8] IoT-A project - Internet of Things Architecture, Accessed on November, 2012. [link].
URL <http://www.iot-a.eu>
- [9] Joachim W. Walewski (ed.), D1.2 - Initial Architectural Reference Model for IoT, IoT-A project, 2011.
URL <http://www.iot-a.eu/public/public-documents>
- [10] Hydra project - Heterogeneous Physical Devices in a Distributed Architecture, Accessed on November, 2012. [link].
URL <http://www.hydramiddleware.eu>
- [11] Sensei project - Integrating the Physical with the Digital World of the Network of the Future, Accessed on November, 2012. [link].
URL <http://www.sensei-project.eu>
- [12] M. Ohashi, Introduction of Ubiquitous Service Platform Project CUBIQ, in: *10th International Symposium on Autonomous Decentralized Systems (ISADS'11)*, Kobe, Japan, 2011, pp. 456–460.
- [13] Smartproducts project - Proactive Knowledge for Smart Products, Accessed on November, 2012. [link].
URL <http://www.smartproducts-project.eu>
- [14] Sensinode NanoService, Accessed on November, 2012. [link].
URL <http://www.sensinode.com>
- [15] G. Jones, *Organizational Theory, Design, and Change* (7th ed.), Prentice Hall, 2012.
- [16] A. Tanenbaum, M. van Steen, *Distributed Systems: Principles and Paradigms*, Prentice Hall, 2002.
- [17] A. Williams. How the Internet of Things Helps us Understand Radiation Levels, Accessed on November, 2012. [link].
URL <https://cosm.com/press>

- [18] D. Viehland, F. Zhao, The Future of Personal Area Networks in a Ubiquitous Computing World, *International Journal of Advanced Pervasive and Ubiquitous Computing* 2 (2) (2010) 30–44.
- [19] H. Schaffers, N. Komninos, M. Pallot, B. Trousse, M. Nilsson, A. Oliveira, Smart Cities and the Future Internet: Towards Cooperation Frameworks for Open Innovation, in: *The Future Internet*, Vol. 6656 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, 2011, pp. 431–446.
- [20] Status of Cloud Services. Cloud Harmony, Accessed on November, 2012. [link].
URL <http://cloudharmony.com/status>
- [21] S. Babar, P. Mahalle, A. Stango, N. Prasad, R. Prasad, Proposed Security Model and Threat Taxonomy for the Internet of Things (IoT), in: *3rd International Conference on Recent Trends in Network Security and Applications (CNSA'10)*, Chennai, India, 2010, pp. 420–429.
- [22] S. Turner, T. Polk, Security Challenges For the Internet of Things, in: *IAB Interconnecting Smart Objects with the Internet Workshop*, Prague, Czech Republic, 2011.
- [23] R. Roman, P. Najera, J. Lopez, Securing the Internet of Things, *IEEE Computer* 44 (9) (2011) 51–58.
- [24] A. Cavoukian, Privacy by Design... Take the Challenge, Information and Privacy Commissioner of Ontario, Canada, 2009.
- [25] P. Mahalle, S. Babar, N. R. Prasad, R. Prasad, Identity Management Framework towards Internet of Things (IoT): Roadmap and Key Challenges, in: N. Meghanathan, S. Boumerdassi, N. Chaki, D. Nagamalai (Eds.), *Recent Trends in Network Security and Applications*, Vol. 89 of *Communications in Computer and Information Science*, Springer Berlin Heidelberg, 2010, pp. 430–439.
- [26] H. Hartenstein, K. Laberteaux, A Tutorial Survey on Vehicular ad hoc Networks, *IEEE Communications Magazine* 46 (6) (2008) 164–171.

- [27] E. Ilie-Zudor, Z. Kemeny, F. van Blommestein, L. Monostori, A. van der Meulen, A Survey of Applications and Requirements of Unique Identification Systems and RFID Techniques, *Computers in Industry* 62 (3) (2011) 227–252.
- [28] J. Takalo-Mattila, J. Kiljander, M. Etelaperä, J.-P. Soininen, Ubiquitous Computing by Utilizing Semantic Interoperability with Item-Level Object Identification, in: *Second International ICST Conference on Mobile Networks and Management (MONAMI'10)*, Santander, Spain, 2010, pp. 198–209.
- [29] T. Bauge (ed.), D3.5 - Global and Pluggable Sensor and Actuator Networking Framework, SENSEI project, 2011.
URL <http://www.sensei-project.eu/>
- [30] H. Akram, M. Hoffmann, Supports for Identity Management in Ambient Environments - The Hydra Approach, in: *3rd International Conference on Systems and Networks Communications (ICSNC'08)*, Sliema, Malta, 2008, pp. 371–377.
- [31] D. Guinard, M. Fischer, V. Trifa, Sharing Using Social Networks in a Composable Web of Things, in: *1st International Workshop on the Web of Things (WoT'10)*, Mannheim, Germany, 2010, pp. 702–707.
- [32] S. G. Weber, L. A. Martucci, S. Ries, M. Mühlhäuser, Towards Trustworthy Identity and Access Management for the Future Internet, in: *4th International Workshop on Trustworthy Internet of People, Things & Services (Trustworthy IoPTS'10)*, 2010.
- [33] A. Sarma, J. a. Girão, Identities in the Future Internet of Things, *Wireless Personal Communications* 49 (3) (2009) 353–363.
- [34] Z. Wei, C. Meinel, Implement Role Based Access Control with Attribute Certificates, in: *6th International Conference on Advanced Communication Technology (ICACT'04)*, Phoenix Park, Korea, 2004, pp. 536–540.
- [35] G. Bai, L. Yan, L. Gu, Y. Guo, X. Chen, Context-aware Usage Control for Web of Things, *Security and Communication Networks*, In Press.

- [36] M. Brachmann, S. L. Keoh, O. G. Morchon, S. S. Kumar, End-to-End Transport Security in the IP-Based Internet of Things, in: 21st International Conference on Computer Communications and Networks (ICCCN'12), Munich, Germany, 2012, pp. 1–5.
- [37] Mocana - NanoDTLS, Accessed on November, 2012. [link].
URL <https://mocana.com/products.html>
- [38] S. Raza, S. Duquennoy, J. Hglund, U. Roedig, T. Voigt, Secure Communication for the Internet of Things - A Comparison of link-layer Security and IPsec for 6LoWPAN, Security and Communication Networks, In Press.
- [39] R. Roman, C. Alcaraz, J. Lopez, N. Sklavos, Key management systems for sensor networks in the context of the internet of things, Computers & Electrical Engineering 37 (2011) 147–159.
- [40] European Network of Excellence in Cryptology II, Accessed on November, 2012. [link].
URL <http://www.ecrypt.eu.org/>
- [41] I. Verbauwhede, J. Fan, Light-weight Public Key Implementations for Constrained Devices, in: Workshop on Cryptography for the Internet of Things, Antwerp, Belgium, 2012.
- [42] K. Beznosov, P. Inglesant, J. Lobo, R. Reeder, M. Zurko, Panel: Usability Meets Access Control: Challenges and Research Opportunities, in: 14th ACM Symposium on Access Control Models and Technologies (SACMAT'09), Stresa, Italy, 2009.
- [43] C. C. Aggarwal, P. S. Yu, A General Survey of Privacy-Preserving Data Mining Models and Algorithms, in: Privacy-Preserving Data Mining, Vol. 34 of Advances in Database Systems, Springer US, 2008, pp. 11–52.
- [44] Y. Shen, S. Pearson, Privacy Enhancing Technologies: A Review, Tech. rep., HP Laboratories (2011).
- [45] V. Oleshchuk, Internet of Things and Privacy Preserving Technologies, in: 1st International Conference on Wireless Communication, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology (WirelessVITAE'09), Aalborg, Denmark, 2009, pp. 336–340.

- [46] P. Angin, B. Bhargava, R. Ranchal, N. Singh, M. Linderman, L. B. Othmane, L. Lilien, An Entity-Centric Approach for Privacy and Identity Management in Cloud Computing, in: 29th IEEE Symposium on Reliable Distributed Systems (SRDS'10), New Delhi, India, 2010, pp. 177–183.
- [47] R. H. Weber, Internet of Things New Security and Privacy Challenges, *Computer Law & Security Review* 26 (1) (2010) 23–30.
- [48] S. Radomirovic, Towards a Model for Security and Privacy in the Internet of Things, in: 1st International Workshop on the Security of the Internet of Things (SecIoT'10), Tokyo, Japan, 2010.
- [49] G. Broenink, J.-H. Hoepman, C. van't Hof, R. van Kranenburg, D. Smits, T. Wisman, The Privacy Coach: Supporting customer privacy in the Internet of Things, in: Pervasive 2010 Conference Workshop on What can the Internet of Things do for the citizen? (CIoT'10), Helsinki, Finland, 2010, pp. 72–81.
- [50] M. Button, Setting the Watch Privacy and Ethics of CCTV Surveillance, *International Journal of Law, Crime and Justice* 39 (4) (2011) 215–217.
- [51] J.-H. Hoepman, In Things We Trust? Towards trustability in the Internet of Things, CoRR abs/1109.2637.
- [52] G. Køien, Reflections on Trust in Devices: An Informal Survey of Human Trust in an Internet-of-Things Context, *Wireless Personal Communications* 61 (3) (2011) 495–510.
- [53] A. Jøsang, R. Hayward, S. Pope, Trust Network Analysis with Subjective Logic, in: 29th Australasian Computer Science Conference (ACSC'06), Hobart, Australia, 2006.
- [54] J.-H. Cho, A. Swami, I.-R. Chen, A Survey on Trust Management for Mobile Ad Hoc Networks, *IEEE Communications Surveys & Tutorials* 13 (4) (2011) 562–583.
- [55] D. Trcek, Trust Management in the Pervasive Computing Era, *IEEE Security & Privacy* 9 (4) (2011) 52–55.

- [56] J. Robinson, I. Wakeman, D. Chalmers, B. Horsfall, Trust and the Internet of Things, in: Joint International Workshop on Trust in Location and Communications in Decentralised Computing (TruLoco'10), Morioka, Japan, 2010.
- [57] uTRUSTit project - Usable Trust in the Internet of Things, Accessed on November, 2012. [link].
URL <http://www.utrustit.eu>
- [58] T. Teixeira, S. Hachem, V. Issarny, N. Georgantas, Service Oriented Middleware for the Internet of Things: A Perspective, in: Proceedings of the 4th European Conference on Towards a Service-based Internet (ServiceWave'11), Poznan, Poland, 2011, pp. 220–229.
- [59] T. Giannetsos, I. Krontiris, T. Dimitriou, F. Freiling, Intrusion Detection in Wireless Sensor Networks, in: Security in RFID and Sensor Networks, Auerbach Publications, CRC Press, 2009.
- [60] D. Petro, G. Vesztergombi, L. Fritsch, D3.2 - Threat Analysis, uTRUSTit project, 2011.
URL <http://www.utrustit.eu/>
- [61] S. Amin, Y. jig Young, M. Siddiqui, C. S. Hong, A Novel Intrusion Detection Framework for IP-based Sensor Networks, in: International Conference on Information Networking (ICOIN'09), 2009, pp. 1–3.
- [62] R. Chen, C. M. Liu, L. X. Xiao, A Security Situation Sense Model Based on Artificial Immune System in the Internet of Things, *Advanced Materials Research* 403–408 (2011) 2457–2460.
- [63] G. Singh, F. Masegla, C. Fiot, A. Marascu, P. Poncelet, Data Mining for Intrusion Detection: from Outliers to True Intrusions, in: 13th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD-09), Bangkok, Thailand, 2009.
- [64] Y. Zhang, L. Wang, W. Sun, R. C. G. II, M. Alam, Distributed Intrusion Detection System in a Multi-Layer Network Architecture of Smart Grids, *IEEE Transactions on Smart Grid* 2 (4) (2011) 796–808.