

## Access control in Internet-of-Things: A survey

Sowmya Ravidas<sup>a,\*</sup>, Alexios Lekidis<sup>a</sup>, Federica Paci<sup>b</sup>, Nicola Zannone<sup>a</sup>

<sup>a</sup> Eindhoven University of Technology, the Netherlands

<sup>b</sup> University of Southampton, UK

### ARTICLE INFO

#### Keywords:

IoT  
Access control  
Literature study

### ABSTRACT

The Internet of Things (IoT) is an emerging technology that is revolutionizing the global economy and society. IoT enables a collaborative environment where different entities – devices, people and applications – exchange information for service provision. Despite the benefits that IoT technology brings to individuals, society and industry, its wide adoption opens new security and privacy challenges. Among them, a vital challenge is the protection of devices and resources produced within IoT ecosystems. This need has attracted growing attention from the research community and industry, and several authorization frameworks have been designed specifically for IoT. In this survey, we investigate the main trends in access control in IoT and perform an extensive analysis of existing authorization frameworks tailored to IoT systems. Driven by the needs of representative IoT applications and key requirements for IoT, we elicit the main requirements that authorization frameworks for IoT should satisfy along with criteria for their assessment. These criteria and requirements form a baseline for our literature study. Based on this study, we identify the main open issues in the field of access control for IoT and draw directions for future research.

### 1. Introduction

A recent technological evolution in the area of pervasive computing is the Internet of Things (IoT). IoT is a “*system of entities (including cyber-physical devices, information resources, and people) that exchange information and interact with the physical world by sensing, processing information, and actuating*” (Standard, 2016). IoT provides advanced applications to industry and citizens that improve individuals’ quality of life and contribute to the world’s digital economy. The adoption of IoT is steeply increasing and several IoT applications are emerging, ranging from smart home (Darianian and Michael, 2008), patient monitoring (Hassanalieragh et al., 2015; Mohammed et al., 2014) and industry automation (Shrouf et al., 2014) to intelligent transportation (Guerero-ibanez et al., 2015), disaster management (Yang et al., 2013) and infrastructure monitoring (Kelly et al., 2013).

IoT combines the current Internet infrastructure and emerging technologies, to ensure the seamless interconnection of hundreds of billions of embedded systems and manage the services they provide while reducing the Internet infrastructure’s cost and making it more scalable, flexible and reactive (van der Meulen, 2015). The adoption of IoT initially relied on the use of web-services to facilitate software reusability

and reduce application development complexity. Although the integration with web-services was an important addition to existing Wireless Sensor Network (WSN) technologies, it also came with new challenges. While web-services are based on long-lived transactions, IoT applications are usually deployed in constrained-resource devices that only wake-up for a short period of time. To overcome the limited capabilities of IoT devices, a recent trend is to shift data storage, communication and computation from resource-constrained devices to the cloud (Mell and Grance, 2011) and edge devices (Stojmenovic and Wen, 2014).

Although IoT brings several benefits for individuals, industry and society (Vermesan et al., 2011), the use of resource-constrained devices along with the adoption of a plethora of technologies enlarges the attack surface and introduces new security vulnerabilities. According to the OWASP IoT project (OWASP, 2018; Miessler, 2018), insecure access to web, backend APIs, cloud and mobile interfaces is one of the top vulnerabilities for IoT applications. Indeed, smart devices are typically configured and controlled via vendor apps, which can have a smartphone-based interface and a Web-based interface through a service running on a cloud infrastructure. Services expose a Web-API that allows to query and control user data and devices from the same vendor and other compliant devices from other vendors. Services from

\* Corresponding author.

E-mail addresses: [s.ravidas@tue.nl](mailto:s.ravidas@tue.nl) (S. Ravidas), [a.lekidis@tue.nl](mailto:a.lekidis@tue.nl) (A. Lekidis), [f.m.paci@soton.ac.uk](mailto:f.m.paci@soton.ac.uk) (F. Paci), [n.zannone@tue.nl](mailto:n.zannone@tue.nl) (N. Zannone).

<https://doi.org/10.1016/j.jnca.2019.06.017>

Received 4 May 2018; Received in revised form 29 March 2019; Accepted 27 June 2019

Available online 3 July 2019

1084-8045/© 2019 Elsevier Ltd. All rights reserved.

**Table 1**  
Comparison with existing surveys.

	Our Survey	Sicari et al. [103]	Roman et al. [90]	Zhang et al. [122]	Ouaddah et al. [82]
<b>Aspects</b>					
Policy Specification	✓	✱	✱	✓	✓
Policy Management	✓	✱	✗	✱	✱
Policy Evaluation & Enforcement	✓	✗	✱	✗	✱
IoT applications	✓	✗	✗	✗	✱

#### Legend

✓: full coverage   ✱: partial coverage   ✗: no coverage

vendors can be composed with third party services e.g Facebook, Instagram using IFTTT Web service.

In this complex IoT ecosystem access control should be enforced at each of these interfaces. However, commercial IoT frameworks fall short in implementing access control to these interfaces. Most of the IoT frameworks enforce coarse-grained access control policies (He et al., 2018; Schuster et al., 2018): for instance, Nest Thermostat<sup>1</sup> grants access to all the capabilities of a smart device or to none, or the Apple Home Kit<sup>2</sup> distinguishes between full control of the device, view only control and local or remote control. Other IoT frameworks enforce slightly richer access control policies based on environmental conditions: for example, Samsung SmartThings<sup>3</sup> grants access whether the user is at home or away. But to track these factors SmartThings gains access to the GPS coordinates of the user smartphone, which allows real-time tracking of users and therefore violates their privacy.

These flaws in implementing access control policies leads to devices and apps to be easily exploited to gain unauthorized access to devices and to users and devices' data that they collect and store (Babun et al., 2009; Celik et al., 2018; Fernandes et al., 2016). A real-world example of the consequences of having permissive or overprivileged interfaces is Internet-enabled baby monitors being remotely hacked and controlled (Stanislav and Beardsley, 2015). The remote hackers could intercept live video feeds from baby monitor's camera and perform different actions including talking to babies and changing camera settings and even permissions to remotely control the baby monitor.

To address these limitations, a fine-grained authorization system that restricts access to IoT device's interfaces and data only to authorized users should be implemented. Such a system should enforce policies based on several factors like the capabilities of the smart devices, the relationships among the users using the devices, and environmental conditions such as time, and location (He et al., 2018).

Driven by this and similar considerations from other studies, recent years have seen an increasing interest in the field of access control for IoT in both academia and industry, which resulted in the emergence of several authorization frameworks for IoT. These frameworks are often based on different IoT technologies and rely on different underlying assumptions. This variety of solutions makes it difficult to evaluate their effectiveness, especially with respect to the target IoT applications.

**Motivation.** While there are a number of surveys that discuss security challenges in IoT (Mahmoud et al., 2015; Sadeghi et al., 2015; Vasilomanolakis et al., 2015; Weber, 2010) such as privacy and network security, only a few address access control (Ouaddah et al., 2017b; Roman et al., 2013; Sicari et al., 2015; Zhang and Wu, 2016). Table 1 provides an overview of existing surveys on access control for IoT. As shown in the table, existing surveys have the following limitations:

- They only discuss some aspects related to access control. Most of the surveys focus on policy specification whereas policy management and evaluation are only partially considered or not investigated at all.
- They do not identify the requirements that access control systems for IoT should satisfy along with evaluation criteria to systematically analyze existing authorization solutions for IoT.
- They do not discuss the demands and requirements of the IoT application for which authorization frameworks are designed.
- They do not discuss the suitability of existing solutions to representative IoT applications.
- They do not analyze most recent papers proposing authorization solutions for IoT.

**Contribution.** In this survey we present a systematic analysis of existing authorization solutions for IoT that addresses the above issues in existing survey papers. Our goal is to identify open challenges in existing authorization solutions to drive the research and development of more effective access control solutions for IoT.

The main contributions of our survey are the following:

- A framework to enable a systematic and comparative analysis of authorization solutions for IoT. The framework consists of a set of requirements that authorization solutions for IoT should meet and a number of criteria for their assessment.
- A review of several recent authorization frameworks for IoT and their evaluation with respect to the requirements and criteria in the framework.
- Guidelines to design an authorization framework tailored to specific needs and constraints of the most common IoT applications.
- Open challenges that need to be addressed when designing access control solutions for IoT.

**Methodology.** To perform our survey, we first provide an overview of relevant characteristics of IoT systems and enabling technologies based on a study of the literature and on current developments of IoT. Our analysis revealed that cloud computing and edge computing are often adopted as a baseline technology in IoT to facilitate the management of devices and resources in IoT ecosystems. To this end, we investigate how these computing paradigms have been adapted to IoT. Driven by real-world scenarios, we identify a set of non-functional requirements for IoT systems.

From these requirements, we elicit the requirements that authorization solutions for IoT should meet and a number of criteria for their assessment. Our requirements cover the main activities of the access control process, ranging from policy specification and management to policy evaluation and enforcement. To assess to what extent existing authorization frameworks meet the elicited requirements, we study the characteristics of the IoT environment in which authorization frameworks have been deployed (*IoT architecture style, communication protocols, data format*) to analyze the assumptions underlying the IoT environment and, in particular, the capabilities of nodes and their interconnections. Moreover, we study the properties of the proposed authorization frameworks (*access control model, policy evaluation strategy, deployment configuration*). We review several recent authorization frameworks for IoT and evaluate them with respect to the identified requirements and criteria. We also discuss the suitability of existing solutions to representative IoT applications.

Our literature review reveals several important insights and help define further research directions in access control for IoT. In partic-

<sup>1</sup> <https://nest.com/thermostats/nest-learning-thermostat>.

<sup>2</sup> <https://www.apple.com/lae/ios/home>.

<sup>3</sup> <https://www.smarthings.com>.

ular, we have observed an increasing interest in the development of authorization frameworks tailored to IoT systems. However, most of the proposed frameworks aim to provide a general solution to address the problem of authorization in IoT. Our analysis, on the other hand, shows that different IoT applications are characterized by different demands and, thus, there is not one solution that fit all IoT applications. Accordingly, the design of an authorization framework tailored to IoT should account for the specific needs and constraints of the target IoT application.

**Organization.** The remainder of the paper is structured as follows. We present an introduction to IoT in Section 2, followed by an overview of IoT enabling technology in Section 3. We discuss typical IoT applications and define the primary functional and non-function requirements for IoT systems in Section 4. We delineate the main requirements that authorization frameworks for IoT should satisfy in Section 5 and review existing frameworks in Section 6. We identify open issues and draw research directions for future research in Section 7 and conclude the paper in Section 8.

## 2. Internet of things

In this section, we provide an overview of the main IoT elements and present how these elements are connected to each other.

### 2.1. IoT elements

Fig. 1 presents the IoT metamodel (in form of UML class diagram) representing the main elements within an IoT system and their relationships. An IoT system consists of devices or smart objects that can interrelate and interconnect among themselves and with the environment to provide services to end-users. Hereafter, we refer to the components of an IoT system as *nodes*. To be able to connect with other nodes, a node should be equipped with a communication interface. In particular, every node is characterized by a Uniform Resource Identifier (URI) that uniquely identifies the node over the network. We distinguish three types of nodes: physical nodes, intermediate nodes and application nodes.

A *physical node* consists of things. A *thing* can be a sensor, an actuator or any other entity that can interact with the environment. A *sensor* is a device that detects events or changes in the environment. An *actuator* is responsible for controlling a mechanism or a system. The output of sensors and actuators is usually referred to as *resource*. An *application node* consumes resources produced by physical nodes to provide services to end-users.

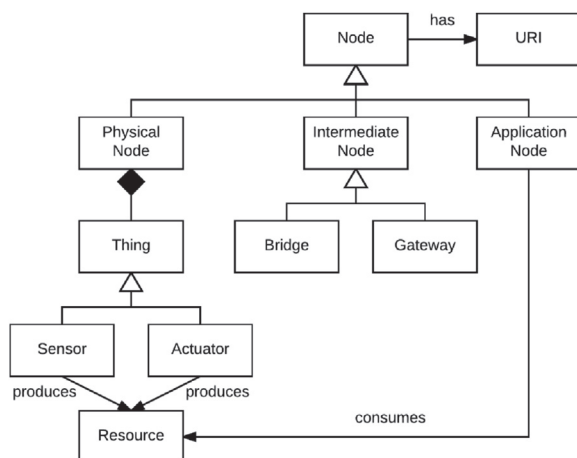


Fig. 1. IoT metamodel.

Physical nodes and application nodes can belong to different networks. *Intermediate nodes* are used to route traffic and connect two or more local area networks. An intermediate node can be a bridge or a gateway. A *bridge* connects local area networks that use the same protocol. In particular, a bridge forwards messages from one network to another based on the MAC address of the destination node. A *gateway* connects networks that use different protocols. Unlike bridges that are only able to forward messages, gateways are also able to perform message conversion to achieve connectivity across different networks.

### 2.2. IoT architecture

An IoT architecture provides a high-level view of the functionalities and connectivity within an IoT ecosystem. In this section, we present an architecture pattern for IoT and discuss the architecture styles typically used in IoT.

#### 2.2.1. Architecture pattern

Several architecture patterns have been proposed for IoT (Aazam et al., 2014; Abdmeziem et al., 2016; Alshehri and Sandhu, 2016; Da Xu et al., 2014; Gubbi et al., 2013; Khan et al., 2012; Wu et al., 2010). The layers in an IoT architecture pattern provide a specific view on the IoT system, and their choice depends on the scope of the study. We observed that authorization solutions tailored to IoT environments are often deployed in the middleware due to the limited capabilities of IoT devices. To this end, differently from many existing architecture patterns for IoT, we separate network communication from the middleware. This distinction allows us to reason on the sharing of resources among nodes and end-users by abstracting from the actual protocols used for their transmission. In particular, we adopt a four-layered IoT architecture pattern (Fig. 2), which consists of physical, network, middleware and application layers. Next, we describe the main functionalities of each layer.

**Application Layer:** The application layer aims to provide services to end-users. This layer comprises application nodes that handle the application logic as well as data semantics and presentation (Alshehri and Sandhu, 2016). These nodes receive data from the middleware and

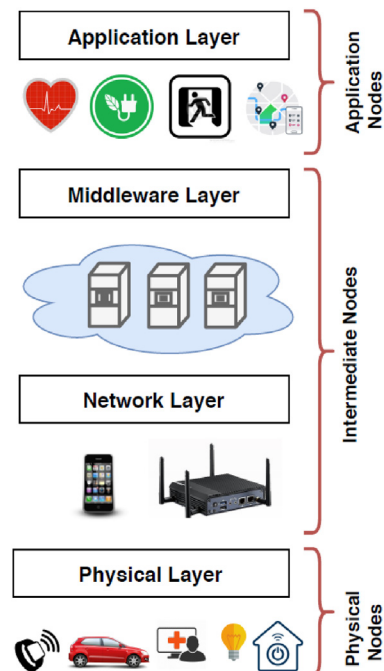


Fig. 2. IoT architecture.

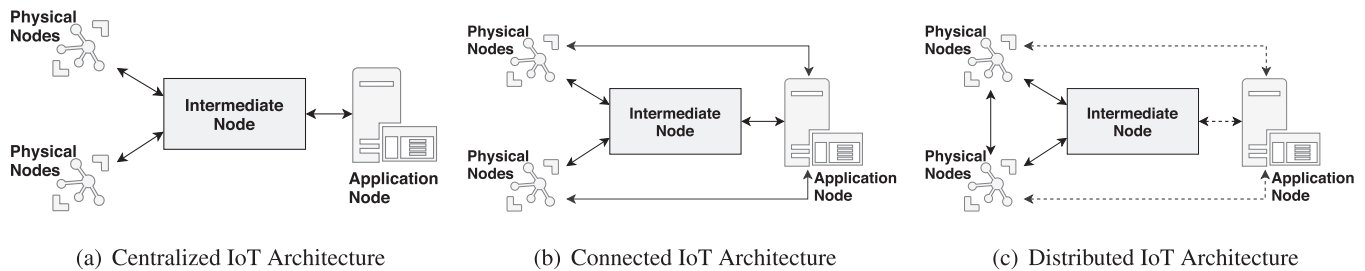


Fig. 3. IoT architecture styles.

process them depending on the end-user requirements and type of service provided. Moreover, the application layer encompasses APIs to facilitate the communication with the middleware and user interfaces through which end-users can access the services.

**Middleware Layer:** The purpose of the middleware is to ensure connectivity and interoperability within the IoT ecosystem. It consists of intermediate nodes<sup>4</sup> that process data received from lower layers and pass them on to the application layer. We will provide an overview of the main types of middleware used in IoT in Section 3.2.

**Network Layer:** The goal of the network layer is to support networking and data transfer between nodes (Da Xu et al., 2014). The network layer implements the communication protocols required for data exchange within an IoT ecosystem. We will discuss the communication protocols typically used in IoT in Section 3.1.

**Physical Layer:** The purpose of the physical layer is to characterize the sensing and control capabilities of an IoT system. This layer comprises physical nodes such as sensors and actuators (Aazam et al., 2014) that sense the environment and interact with it in response to changes or users' requests. These nodes produce resources (e.g., sensing data) that are passed to application nodes through the network and middleware layers.

### 2.2.2. Architecture styles

Several architecture styles for IoT have been proposed in the last years. Although these architecture styles can vary based on the application domain, they can be classified into three main types based on the connectivity between physical nodes, middleware (intermediate nodes) and application nodes (Bouij-Pasquier et al., 2015b; Roman et al., 2013). An overview of the main IoT architecture styles is presented in Fig. 3.

**Centralized:** Physical and application nodes are required to communicate with each other through intermediate nodes. This means that if an application node wants to retrieve resources from a physical node, a connection has to be established using the interfaces provided by the intermediate node. This architecture style is typically required for physical nodes with limited processing and storage capabilities.

**Connected:** Physical nodes have the ability to process information and forward it to intermediate nodes.

In addition, physical nodes can provide resources directly to application nodes. This means that application nodes can directly connect to physical nodes through the interfaces they provide.

**Distributed:** Every node can communicate with each other. This means that every node has the potentiality to process information and provide services. Note that, differently from the other IoT architecture styles, the distributed IoT architecture does not require an intermediate node, although its use can facilitate the communication between nodes.

<sup>4</sup> Note that we use the term intermediate nodes to indicate any node that is in between the physical layer (comprised by physical nodes) and application layer (comprised by application nodes).

## 3. IoT enabling technology

Recent years have seen the emergence of new technology (and the adaptation of existing technology) to meet the demands of IoT applications and low-power and resource-constraint IoT devices. This section presents an overview of the main technologies enabling IoT, with a particular focus on communication protocols and middleware.

### 3.1. Network layer

Communication protocols and their relations are usually represented using a network stack. A network stack is represented in layers for easier design and evaluation.<sup>5</sup> Each layer represents different functions and offers different methods for data handling. In this work, we use a four layer model that resembles the traditional five layer network model (Socolofsky and Kale, 1991). It comprises the application layer, transport layer, network layer and data link & physical layer.

Fig. 4 presents the network stack describing the protocols commonly used in IoT. Moreover, we relate the standards upon which the physical & datalink layer protocols are defined. In the figure, layers are separated by solid lines. Arrows indicate that a given protocol is built on top of another protocol or built on a given standard. Next, we review the main protocols used in each layer.

**Physical layer & datalink layer:** Several protocols have been used in the physical layer & datalink layer within IoT ecosystems. We classify them based on the network type they support: Local Area Network (LAN), Personal Area Network (PAN) and Wide Area Network (WAN). The PAN protocols commonly used in IoT are Radio-Frequency Identification (RFID) (Information technology, 2013), Bluetooth (Bluetooth SIG Working Group, 2017), ZigBee (ZigBee Specification, 2014) and ZWave. RFID (Information technology, 2013) is largely used within IoT environments to identify devices (Sethi and Sarangi, 2017). Bluetooth and its variant for low energy devices are a short-range wireless technology based on the IEEE 802.15.1 standard (Bluetooth SIG Working Group, 2017).

The IEEE 802.15.4 standard (IEEE Std 802.15.4-2015, 2015) is intended for low-rate wireless personal area networks (LRWPAN). It offers low-cost low-power communication to devices in close proximity. A protocol based on LRWPAN specification is Zigbee (ZigBee Specification, 2014). Although this protocol builds on LRWPAN, it has additional components for the network and application layers. Similar to Zigbee, the Z-Wave<sup>6</sup> protocol also works on low-frequency radio bandwidth. Z-Wave is a proprietary protocol that is not built on any specific standard. It provides the complete network stack from the physical layer to the application layer.

Among LAN protocols, traditional technologies such as Ethernet and Wi-Fi are often used in IoT. Ethernet (IEEE Standard, 2015) is a highly reliable protocol based on IEEE 802.3. WiFi is based on the

<sup>5</sup> Note that the layers in the network stack are different from the architecture layers presented in Section 2.2, although some layers may have the same name.

<sup>6</sup> <http://www.z-wave.com>.

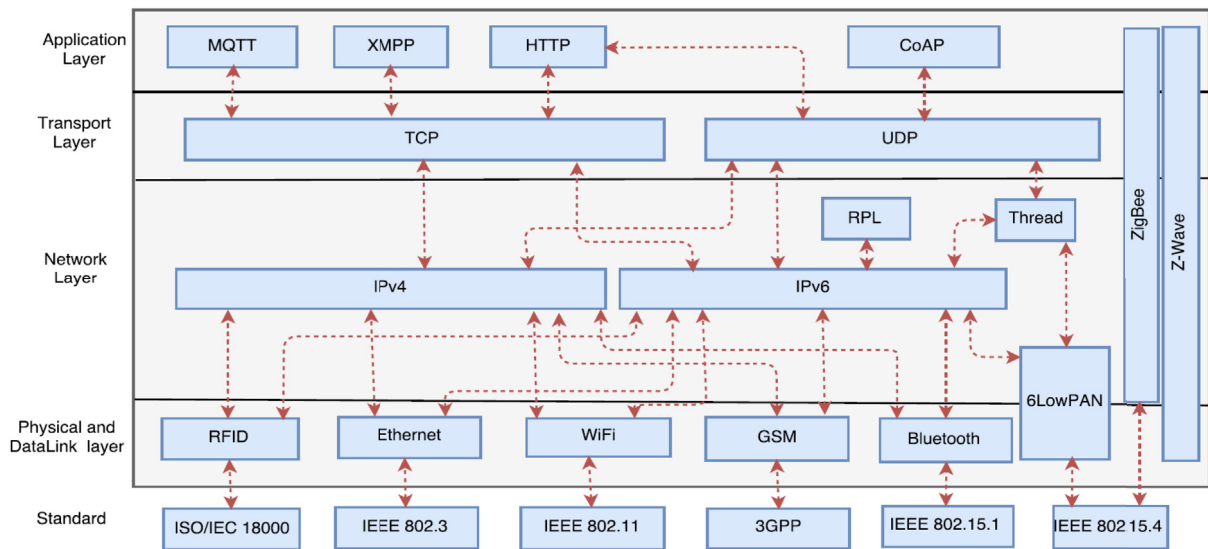


Fig. 4. IoT network stack and standards.

IEEE 802.11 standard (Wireless LAN, 2013) and allows gateway devices to transmit information using radio-waves over high speed Internet connection.

For Wide Area Network (WAN), cellular technologies are often used in IoT. Global System for Mobile Communications (GSM) (GSM, 2017) is the most commonly adopted protocol for long range communication, primarily used for data voice transmission and services based on the 3GPP specification.

**Network layer:** The most commonly used network layer protocols in IoT are IPv4 and IPv6. These protocols are variations of the Internet Protocol (IP) (Internet Protocol, 1981), both used to identify devices on the Internet based on unique addresses. They provide an addressing scheme that is used to identify a group of IoT devices geographically (Gubbi et al., 2013). IPv6 requires a minimum MTU (Maximum Transmission Unit) size of 1280 bytes, whereas the IEEE 802.15.4 link layer allows a maximum frame size of 127 bytes. Hence, there is a need for additional protocols to perform packet compression in order to transmit IPv6 packets over IEEE 802.15.4.

The IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs) (Montenegro et al., 2007) is built on top of the LRWPAN specification. 6LoWPAN uses encapsulation and header compression mechanisms to transmit IPv6 packets over IEEE 802.15.4 networks, thereby creating a mapping between the link and network layer. This protocol aims to support IP for low power IoT devices. Thread<sup>7</sup> is another network layer protocol specifically designed for device-to-device communication in building automation, which is based on IPv6 and 6LoWPAN.

IoT systems also extend the architectures and protocols used in Wireless Sensor Networks (WSN) through the addition of web resources. A particular type of WSN IoT network architecture is low-power and lossy network (LLN). In such networks, devices and routers have memory and processing constraints. Moreover, routers typically support low data rates and are unstable. RPL (Winter et al., 2012) is an IPv6 routing protocol for LLN that efficiently routes multipoint-to-point (from devices to a central point), point-to-multipoint (from a central point to devices) and point-to-point (between the devices) traffic.

**Transport layer:** Transmission Control Protocol (TCP) (Transmission Control Protocol, 1981) and User Datagram Protocol (UDP) (Postel, 1980) are widely used protocols in the transport layer for IoT. TCP

is connection-oriented whereas UDP is connection less. This distinction makes TCP more reliable than UDP as TCP guarantees that all packets are delivered. However, it is not scalable for small data transmissions in IoT devices (Zanella et al., 2014). UDP is more suitable for real-time communication where delay is not tolerated.

**Application layer:** In the traditional Internet protocol stack, the most common protocol in the application layer is HTTP. It can be run over both TCP and UDP. However, HTTP is verbose and complex, and adds a significant parsing overhead. This may not be suitable for constrained devices. Moreover, HTTP inherits the limitations of the protocol at the transport layer on which it runs (Zanella et al., 2014). To overcome these limitations several application layer protocols have been developed. Their goal is to enhance communication performances by minimizing the overhead.

Constrained Application Protocol (CoAP) (Shelby et al., 2014) is one of the most commonly used protocols for IoT devices and is based on the client-server model. It runs over UDP and performs asynchronous message exchanges. CoAP has low header overhead and hence simplifies message parsing (Sheng et al., 2013). HTTP and CoAP can also be used in association with representational state transfer (REST) architecture (REST, 2011) that makes it possible to access the resources of an IoT device through a URI. Another commonly used application layer protocol is Message Queue Telemetry Transport (MQTT) (MQTT, 2014), a lightweight messaging protocol on top of the TCP/IP protocol. It is mainly used for communication with remote locations where network bandwidth can be limited. MQTT is based on the publish-subscribe paradigm, where the sender (i.e., the publisher) transfers the message to a broker that distributes the messages to the interested clients (i.e., subscribers). One of the commonly used MQTT brokers is Mosquitto (2017). Since MQTT runs on top of TCP, it may not be suitable for applications that require real-time processing. Another application layer protocol commonly used in IoT is Extensible Messaging and Presence Protocol (XMPP) (Saint-Andre, 2014). This protocol is used for streaming XML elements and real-time exchange of structured data.

### 3.2. Middleware layer

The middleware is often required to ensure connectivity, interoperability, storage and computation of data within an IoT ecosystem. Different types of middleware have been proposed for IoT (Razzaque et al., 2016; Sethi and Sarangi, 2017; Song et al., 2010). In our study of the literature on access control for IoT, we observed that most of

<sup>7</sup> <https://threadgroup.org>.

the existing solutions rely on cloud computing and edge computing as middleware.

According to the NIST (Mell and Grance, 2011), *cloud computing* is a computing paradigm that enables on-demand network, storage, applications and other services without any management effort. Cloud computing is become a core component of most of IoT platforms because it provides elastic and scalable data storage and processing. The adoption of cloud computing has opened new directions for technological enhancements in several IoT applications. The cloud has different role within the IoT architecture, depending on the application needs and requirements. We refer to Section 4.3 for a discussion.

*Fog computing* (Stojmenovic and Wen, 2014) and *edge computing* (Garcia Lopez et al., 2015) have been recently proposed as new computing paradigms to reduce the communication latency and bandwidth required by the use of a remote cloud platform for data storage and processing. The underlying idea behind fog computing is that data generated by IoT devices are processed at the edge of the network, close to where they are generated. A network of micro data centers process and store critical data locally and then push all received data from IoT devices to a remote cloud platform for long-term storage. The computational processes being done by the micro data centers is usually referred as edge computing, while the network connections between the micro data centers and the cloud platform is referred to as fog computing.

#### 4. IoT applications and requirements

Consumers, industries and governments are starting to realize the numerous benefits that IoT brings to the society, economy and environment. However, in order for these stakeholders to reap off the benefits of IoT, an IoT ecosystem has to address key requirements related to scalability, interoperability, performance, availability, reliability and dynamicity. As shown below, these requirements are not equally relevant for all IoT applications; instead, each application has different characteristics and, thus, different sets of requirements. In this section, we introduce representative IoT applications and the key requirements that they have to satisfy.

##### 4.1. Applications

IoT technology has been applied in a variety of applications domains. In this section, we examine the most representative applications of IoT (Interagency Report, 2018; Gerdes et al., 2014).

**Smart Homes** (🏠) Smart home applications involve the use of smart objects like thermostats, door bells, door locks, smart appliances (e.g., smart tv, smart fridge) and smart light bulbs that are remotely controlled by the home owners via their smart phones or home assistants. A smart home involves both human-to-machine and machine-to-machine interactions. An example of the former is the one in which a home owner has installed smart door locks and an alarm system at the door and windows and she would like to allow other users, e.g. visitors or family members, to unlock the door using their smart phone. An example of machine-to-machine interaction is a thermostat linked to a garage door that automatically increases the temperature in the house when the garage door opens or a smart fridge that automatically notifies the owner's smart phone that the milk is running low.

**Health IoT** (❤️) Health IoT comprises medical devices such as insulin pumps, cochlear implants and pacemakers that include sensors to monitor patients' vital signs and actuators to respond to situations that can potentially threaten patients' life. One typical use case is diabetes monitoring where a patient is equipped with a wireless-enabled glucose monitor and an injection device that allows monitoring her blood glucose level, and receives alerts on her mobile phone for hypoglycemia and insulin dosage updates from her primary physician.

**Smart Buildings** (🏢) Smart building applications focus on the use of smart components such as heating, ventilation, air conditioning

(HVAC), smart lighting systems, and safety & security systems such as fire alarms to enhance the overall living quality of tenants and to save energy. These devices are usually managed centrally by a facility manager using a Building and Lighting Management System (BLMS). However, different areas of these buildings can be leased to different companies. These companies should retain control of the lighting and HVAC in their part of the buildings. Other parts of the buildings automation system such as entrance illumination and fire-alarm systems should be controlled either by all companies together or by a facility management company.

**Connected Vehicles** (🚗) Connected vehicle applications involve vehicles, roadside units and other infrastructure to communicate and share traffic and road information. In this setting, vehicles transmit their location, direction, speed, and other information such as vehicle ID and size to other vehicles in proximity. Vehicle-to-vehicle communication (V2V) enables safety use cases such as forward collision warning, blind-spot detection and traffic congestion. Similarly, vehicles' information can be shared with other components of the road infrastructure such as traffic lights, stop signs, toll booths and road crossing in (vehicle-to-infrastructure (V2I)) to support traffic jam notification, prediction of potential traffic jams and dynamic traffic light control.

**Smart Manufacturing** (🏭) Smart manufacturing applications are typically based on an open and interconnected infrastructure that allows the management and monitoring of industrial and manufacturing processes. The infrastructure integrates different technologies such as microprocessors, cloud services, new generation control systems, software application, sensors and actuators to collect real-time data and process them to take prompt decisions based on reliable data. A popular use case is predictive maintenance where sensors, cameras and data analytics are used to determine when a piece of equipment is going to fail before it actually does. By leveraging streaming data from sensors and devices, the infrastructure can quickly assess the current condition of equipment, recognize warning signs, and deliver alerts to operators, who can trigger appropriate maintenance commands and processes.






##### 4.2. Requirements

Realizing an IoT system is not an easy task due to the many requirements that need to be addressed. In this section, we discuss key functional and non-functional requirements that should be satisfied by any IoT system and for which IoT applications these requirements are more relevant. These requirements have been gathered from a study of the literature and an analysis of the main characteristics of the IoT applications and use cases presented in the previous section. A summary of these characteristics is presented in Table 2 and the requirements along with the IoT applications in which these requirements are particularly relevant are presented in Table 3.






































**Scalability:** Scalability refers to the ability of being extensible in terms of number of users and physical nodes without negatively affecting the quality of the services provided by the IoT system (Al-Fuqaha et al., 2015). Implementing this requirement involves an efficient way to manage physical nodes within the IoT system (GR1). Node management includes aspects such as node registration and identification, and the storage and processing of huge volume of data generated by physical nodes (GR2). This is a key requirement in application scenarios like connected vehicles where millions of cars can join the road infrastructure or in the smart manufacturing sector where new equipment to be monitored is often added to the interconnected infrastructure.

**Interoperability:** Interoperability is a significant requirement for all IoT applications. IoT systems usually consist of heterogeneous devices, services and applications from different vendors and service providers that use different communications technologies and formats for data exchange (Miorandi et al., 2012). Interoperability should be considered by both service providers and device manufactures to make sure that nodes can exchange information and resources with

**Table 2**  
Main characteristics of IoT applications.

					
Scalability	low	medium	medium	high	high
Heterogeneity	high	high	high	high	high
Lightweight	high	high	high	medium	high
Latency sensitive	low	high	medium	high	high
Reliability	low	high	medium	high	high
Dynamicity	medium	high	high	high	high
User involvement	high	high	low	low	low
Automation	medium	high	high	high	high
IoT architecture	connected	connected	centralized	distributed	connected
Administrative domain	single	single/ multi	multi	multi	multi

**Table 3**  
IoT requirements along with the relevant IoT applications.

Category	ID	Requirement	IoT application
Scalability	GR1	An IoT system must be scalable to manage a large number of devices.	 
	GR2	An IoT system must be scalable to handle the resources produced by devices.	 
Interoperability	GR3	All components of an IoT system must be able to communicate with each other.	    
Performance	GR4	The communication overhead should be low on the device side.	    
	GR5	The computation overhead should be low on the device side.	    
	GR6	The latency of information exchange between the nodes must be minimal.	  
Reliability & Availability	GR7	An IoT system must guarantee an adequate level of reliability and availability of its nodes.	  
Dynamicity	GR8	An IoT system must be able to handle the dynamicity of the nodes.	    
	GR9	An IoT system must be able to handle the dynamicity of the environment.	    
Usability	GR10	An IoT system must reduce user effort in system administration and configuration.	 

each other regardless the specific technology or protocols being used (GR3).

**Performance:** IoT systems often consist of resource constrained devices that have limited storing, networking and processing capabilities. Therefore, IoT solutions and protocols should be lightweight (Nguyen et al., 2015; Seitz et al., 2013) meaning that overhead due to communication (GR4) and computation (GR5) should be as low as possible on the device side. Performance can also be affected by the latency of transferring data between nodes due to the underlying network and middleware infrastructure. Delays in the transmission and processing of data should be minimal (GR6). This is a key requirement for safety-critical IoT applications like connected vehicles and patient monitoring, and for real-time applications like smart manufacturing. For example, in forward collision warning a delay in the transmission of the speed and location of a vehicle to the vehicle traveling on front could cause a collision between the two vehicles.

**Reliability & Availability:** Reliability refers to the proper functioning of an IoT ecosystem (Al-Fuqaha et al., 2015). The data sensing, transmission and processing should be reliable in the sense that even if a failure or a malfunction occurs, the IoT ecosystem should still guarantee service delivery (Razzaque et al., 2016). Reliability entails the availability of data sensing, communication and processing, and of the services and applications that consume the data. If critical data from sensors are not available, the IoT system may actuate the wrong decision. Therefore, an IoT system needs to guarantee reliability and availability of data, applications and services over time (GR7). The need of reliability and availability depends on the type of service delivered by an IoT application. Occasional unavailability and/or failure can be tolerated in smart homes, e.g. the smart fridge failing to notify the

owner that she is running out of milk or a home assistant failing to remotely control the smart lighting system. However, reliability and availability are major requirements for safety related IoT applications, like connected vehicles and health IoT, and time critical applications, like smart manufacturing. For instance, failures in a smart glucose monitoring device or in an injection device could be life threatening for the patient. Similarly, if a warning about a piece of equipment malfunctioning is not delivered to the operator, maintenance processes are not conducted with consequent disruption and delays of the manufacturing process.

**Dynamicity:** An IoT ecosystem is dynamic by design wherein the network topology and connectivity can constantly change (Tönjes et al., 2014). For instance, physical nodes can leave the system or new physical nodes can join the system (Dar et al., 2011). IoT systems should be able to adapt to the dynamicity of nodes (GR8). This is a critical requirement for IoT applications such as connected vehicles and smart manufacturing. Moreover, IoT systems are often employed in cyber-physical systems to monitor and manage IT infrastructures and the surrounding environment. In this setting, it is crucial that an IoT system is able to adapt to changes in the environment (GR9). The ability to adapt to changes is a key requirement for all IoT application scenarios.

**Usability:** Usability is a primary requirement in IoT applications that are characterized by a high user involvement and by the use of wearable devices like in smart homes and health IoT (GR10). Wearable devices often have very small displays, which makes user interaction and determining what information to display a tricky, but important factor. Device interfaces should also be easily customizable by users and facilitate the management and administration of the device itself.

The requirements above strictly influence the design and deployment of security mechanisms employed for the protection of the IoT system itself and resources produced and processed by the system. In Section 5, we discuss how the requirements in Table 3 affect the design of authorization systems tailored to IoT.

#### 4.3. Discussion

IoT ecosystems are required to satisfy the requirements in Table 3, especially the ones of the target IoT application. The use of cloud provides a natural basis for the achievement of some of these requirements. However, the achievement of other requirements might require additional measures, depending on the type of IoT architecture adopted.

As the number of connected devices is increasingly growing, scalability (GR1 and GR2) is of utmost importance for IoT architectures. Cloud computing allows meeting these requirements since it offers a structural way to manage and remotely control the overall IoT ecosystem. The interaction of physical and application nodes with the cloud is enabled through the definition of interfaces. Physical and application nodes can use these interfaces to store or retrieve resources (Fox et al., 2012), thus providing interoperability (GR3). On the other hand, IoT applications like connected vehicles in which physical and application nodes can interact with each other without the presence of the cloud (distributed IoT architecture) are required to provide interfaces for the interaction with other nodes.

Cloud computing can also help in meeting performance requirements by relieving physical nodes from heavy computations (GR5). However, it might bring communication overhead to physical nodes (GR4) and introduce a delay in the communication between nodes (Roman et al., 2013) (GR6). This delay can be alleviated by the use of edge computing that brings cloud capabilities closer to physical nodes.

Cloud-based systems have usually a high uptime, thus ensuring the availability of the infrastructure (GR7) (Patel et al., 2013). Hence, cloud-based IoT systems are usually reliable, although they are not robust against failures in the connectivity. One advantage of the connected and distributed IoT architectures is that, even if the connectivity to the cloud fails, application nodes can still access resources directly from physical nodes (Roman et al., 2013). Cloud and edge computing also support the dynamicity of IoT environments (Botta et al., 2016) (GR8). In particular, the appearing and disappearing of physical nodes are typically handled by the cloud and edge nodes. On the other hand, in a distributed architecture, nodes are self-organizing and, thus, dynamicity is often handled using routing protocols such as RPL. Apart from the dynamicity of nodes, the IoT environment can also change rapidly (GR9). However, the cloud may not have any information about potential environment changes and has to rely on sensors and actuators to gather such information.

### 5. Access control in IoT

While offering attractive opportunities and new business models, IoT opens several security and privacy issues. In this work, we focus on one of the main security issues in IoT, namely how to protect IoT devices and resources (data, applications, services) from being accessed by unauthorized users. A typical solution to address this issue is the adoption of an access control system that guarantees that only authorized entities (users and devices) gain access to IoT devices and resources. In this section, we investigate the requirements that access control systems for IoT should meet and identify design principles and criteria to evaluate the current state of the art on this field.

The design of an access control system typically comprises three main components (Samarati and Capitani de Vimercati, 2000): *policy*, which defines authorization requirements according to which access control is regulated; *model*, which provides a formal representation of access control policies and their evaluation; *mechanism*, which defines the low level implementation of the control imposed by the policy as

formalized in the model. In this work, we identify requirements for the access control model and mechanism and discuss to which extent the access control models, policy evaluation strategies and enforcement architectures supported by existing access control systems for IoT satisfy the requirements. To this end, we first present an overview of the most popular access control models, reference architectures and policy evaluation strategies. Then, we introduce the requirements for an access control system and the criteria to evaluate whether the requirements are satisfied by authorization frameworks proposed for IoT. These requirements and criteria will serve as the baseline for the evaluation of existing authorization frameworks for IoT in Section 7.

#### 5.1. Access control basics concepts

In this section, we first review existing access control models and which concepts they support to specify an access control policy. Then, we introduce the main architectures and policy evaluation strategies that have been proposed to implement an access control system.

##### 5.1.1. Access control models

Access control policies are formally represented according to an access control model. Several models have been proposed in the literature. These models have different characteristics, which can influence the suitability of an authorization mechanism for IoT. Next, we present an overview of the most popular access control models.

**Discretionary Access Control (DAC)** (Graham and Denning, 1972): DAC is based on the notions of ownership where a user has complete control over its own resources and devices, and can determine the permissions other users have on those resources and devices. Although many variations have been proposed, DAC is generally considered an identity-based access control model where access rights are assigned to users based on their identity. Various approaches to implement DAC have been proposed: *access matrix*, *authorization table*, *access control list* (ACL) and *capability list*. We refer to (Samarati and Capitani de Vimercati, 2000) for an overview of these approaches and, later in our analysis (Section 7), we differentiate between these implementations.

**Mandatory Access Control (MAC)** (LaPadula et al., 1973): Differently from DAC, MAC relies on a set of system rules rather than being at the discretion of an object's owner. These rules are typically defined based on security labels associated to subjects and objects. Thus, similarly to DAC, MAC is considered an identity-based access control model.

**Role-based Access Control (RBAC)** (Sandhu et al., 1996): RBAC relies on the notion of role to simplify the specification and management of access rights within an organization. A role comprises the set of permissions needed to carry out a certain job function. Users are assigned to roles and inherit the permissions assigned to the roles they have. Roles are often organized in a role hierarchy, which defines the inheritance of permissions between roles.

**Organization-Based Access Control (OrBAC)** (Kalam et al., 2003): OrBAC is based on three main concepts for the specification of access control policies, namely organization, concrete and abstract levels, and context. Organization is a structured group of active entities. Similarly to other access control models, concrete authorizations are specified in terms of subject, action, and object, defining which action a user can (or cannot) perform on an object. Concrete authorizations are derived from abstract permissions, which are defined in terms of roles, activities and views. As in RBAC, a role represents a job function within the organization. Activities group actions into an abstract set and views represent sets of abstract objects. Subjects in the concrete level are mapped to roles in the abstract level, actions are mapped to activities, and objects are mapped to views. The context represents a specific situation and is used in OrBAC to express dynamic rules.

**Attribute-based Access Control (ABAC)** (Hu et al., 2014; Yuan and Tong, 2005): ABAC is a general-purpose access control model in which access rights are constrained with respect to the attributes of subjects, objects, actions and the environment. Policies and access requests are



defined in terms of attribute names/values pairs. The applicability of a policy to a request is determined by matching the attributes in the request with the attributes in the policy. ABAC models often provides constructs to combine policies authored by different stakeholders and mechanisms to solve conflicts that can arise from these policies.

**Usage Control (UCON)** (Park and Sandhu, 2004): Similarly to ABAC, UCON allows the specification of policies in terms of subjects and objects' attributes. It also uses conditions to express access constraints on the environment, thus providing the same expressiveness of ABAC. Moreover, UCON supports two additional decision properties, namely mutability of attributes and continuity of decision. Mutability of attributes accounts for changes of subjects and objects' attributes as a consequence of the usage. Continuity of decision denotes that permissions are checked not only at access time but also during the entire usage.

### 5.1.2. Reference architectures

An authorization mechanism defines the low-level implementation of the access control model within the system. An authorization mechanism can be logically decomposed into key components that are responsible for the evaluation and enforcement of access control policies specified according to an access control model. Here, we discuss various reference architectures that have been proposed as a foundation for access control systems, namely *policy-based architecture*, *token-based architecture* and *hybrid architecture*, to study how the access control process is spread across the IoT ecosystem.

**Policy-based Architecture:** A widely adopted policy-based architecture is the one proposed by XACML (XACML, 2013), the de facto standard for the specification and enforcement of access control policies. The architecture comprises four main components<sup>8</sup>

- *Policy Enforcement Point (PEP)* provides an interface with the system and is responsible for enforcing access decisions.
- *Policy Decision Point (PDP)* evaluates access requests against access control policies and determines whether access should be granted or denied.
- *Policy Administration Point (PAP)* acts as a policy repository and offers facilities for policy management.
- *Policy Information Point (PIP)* denotes the source of information (e.g., context information) needed for policy evaluation.

Fig. 5 shows the interaction between these components. The PAP makes the policies available to the PDP (1). Upon receiving an access request (2), the PEP forwards the request to the PDP (3), which evaluates the request against the policies fetched from the PAP. If additional information is required for policy evaluation, the PDP queries the PIP (4,5). The PDP evaluates the request against the policies and returns a response specifying the access decision to the PEP (6), which enforces the decision.

**Token-based Architecture:** Solutions adopting a policy-based architecture typically provide a single, centralized point for the evaluation and enforcement of access control policies. This solution may not be suitable when resources are distributed across different nodes, which is a typical situation in many IoT applications. The last years have seen the emergence of token-based architectures as an alternative to policy-based architectures to deal with the needs of open and decentralized systems. Roughly speaking, in a token-based architecture, an authorization service encodes the permissions of users and devices in a token, which is then used to grant them access to resources and services. Various standards have defined reference token-based architectures and authorization protocols. These architectures and protocols vary in the way tokens are generated and in the flow of the authorization process.

<sup>8</sup> The XACML reference architecture includes an additional component, called *Context Handler*. We omit this component here as its main function is to support the authorization process.




Fig. 5. Policy-based architecture.

A widely-used token-based authorization protocol is OAuth (Dennis and Bradley, 2017). OAuth allows client applications (web, mobile or desktop) to access resources hosted on an HTTP server with the authorization of the resource owner. The authorization granted by the resource owner is encoded in an access token. The OAuth architecture encompasses two main components: an Authorization Server, which is responsible to generate access tokens, and a Resource Server, which hosts the resources and is responsible for their disclosure thus acting as the PEP. Fig. 6 presents the OAuth architecture. The client application requests access to the resource owner (1), who provides the client application an authorization grant representing the resource owner's authorization (2). The client requests an access token to the Authorization Server by presenting the authorization grant received by the resource owner (3). The Authorization Server authenticates the client application and validates the authorization grant and, if valid, issues an access token to the client application (4). The client application requests access to the resource to the Resource Server by presenting the access token (5). The Resource Server validates the access token and, if valid, it serves the request (6). Tokens can be reused for subsequent accesses until it is valid. When the validity of the token expires, the token is renewed through a refresh token without user intervention.

**Hybrid Architecture:** Token-based architectures require user intervention during policy evaluation. For instance, OAuth requires the resource owner to authorize a device or an application acting in her behalf the first time that the device/application requires access to a service or a resource, which may be inconvenient in IoT applications characterized by a large number of devices like smart manufacturing. To address this drawback, the Kantara Initiative has proposed User-Managed Access (UMA) (User-Managed Access (UMA), 2017). UMA extends OAuth with the possibility of configuring policies in the Authorization Server to autonomously generate authorization tokens without user involvement. In this respect, UMA adopts a hybrid approach that combines features of policy-based and token-based architectures.

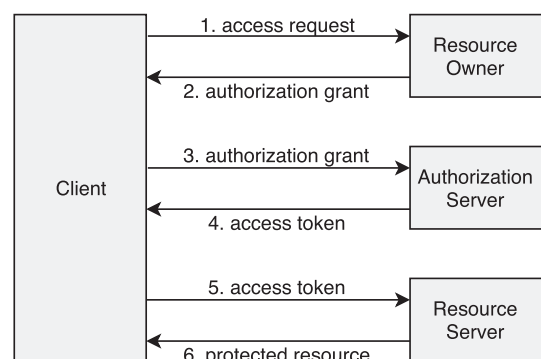


Fig. 6. Token-based architecture (Dennis and Bradley, 2017).

**Table 4**  
Requirements for access control systems tailored to IoT environments.

Category	IoT Req.	ID	Requirement
Policy Specification	GR1	ACR1	Access control models should allow the specification of fine-grained access control policies
	GR8, GR9	ACR2	Access control models should allow the specification of policies able to handle the dynamicity of nodes and IoT environments
Policy Management	GR1,GR2	ACR3	Access control models should be able to handle the complexity of IoT environments
	GR10	ACR4	Access control systems should facilitate users in policy management
	Multi Administrative domain	ACR5	Access control systems should enable policy management across multiple administrative domains
Policy Evaluation & Enforcement	GR2, GR10	ACR6	Access decision making should be automated
	GR4, GR5	ACR7	Access control systems should not significantly impact the computing and communication capabilities of resource-constraint devices
	GR4, GR5, GR6	ACR8	Access control systems should not affect the performance of the IoT system
	GR3	ACR9	The outcome of policy evaluation should be coherent across multiple administrative domains
	GR7	ACR10	Access control systems should be always operational

5.1.3. Policy evaluation strategy

The architecture underlying the access control mechanism determines the strategy used for policy evaluation. In particular, we identified three main classes of policy evaluation strategies:

**Run-time policy evaluation:** policy evaluation is performed at request time. Upon receiving an access request, the PDP evaluates the request against the policies made available by the PAP and returns an access decision to the PEP for enforcement. This strategy is typically supported by frameworks adopting a policy-based architecture.

**Off-line policy evaluation:** policy evaluation is precomputed. An entity obtains assertions on its credentials, access permissions and other attributes from the resource owner or a trusted party. Then, when the entity requests access to a resource, the precomputed assertions are verified by the PEP at run time for enforcement. This strategy is typically supported by frameworks based on OAuth and UMA.<sup>9</sup>

**Hybrid policy evaluation:** Hybrid policy evaluation lies in between run-time and off-line policy evaluation. Specifically, part of policy evaluation is performed off-line, for instance using a token service that asserts the attributes and permissions of an entity within the IoT system. Then, additional verification activities are performed at request time. These activities encompass the verification of the token by the token service or by other external components (hereafter referred to as hybrid<sup>f</sup> strategy) or the verification of context or other constraints in order to make an access decision (hereafter referred to as hybrid<sup>d</sup> strategy).

5.2. Requirements for access control in IoT

In this section, we discuss the requirements that access control systems for IoT should meet. These requirements aim to identify the main concepts and design principles that have to be considered in the design and development of access control systems tailored to IoT applications. In particular, the requirements have been distilled by applying the non-functional requirements for IoT introduced in Section 3 (see Table 3) to the different components of an access control system and from the analysis of key characteristics of IoT applications. Additionally, we have identified the ability to deal with the protection of resources and devices across multiple administrative domains, where different authorities are in control of (different parts of) the IoT ecosystem, as an important requirement in several IoT applications. Our list of requirements is reported in Table 4. The relevance of these requirements to IoT applications is presented in Table 5.

The requirements specified in Table 4 serve as a baseline for the analysis of existing authorization frameworks for IoT and the identification of gaps in the current state of the art. Note that the requirements

**Table 5**  
Relevance of requirements for IoT applications. Symbol ● is used to denote high relevance, ◐ medium relevance, and ◑ low relevance.

	ACR1	ACR2	ACR3	ACR4	ACR5	ACR6	ACR7	ACR8	ACR9	ACR10
Smart Home (🏠)	●	●	◐	●	◐	◐	◐	◐	◐	●
Health IoT (❤️)	●	●	◐	●	◐	●	●	●	◐	●
Smart Buildings (🏢)	●	●	●	◐	●	●	◐	●	●	●
Connected Vehicles (🚗)	●	●	●	◐	●	●	◐	●	●	●
Smart Manufacturing (🏭)	●	●	●	◐	●	●	●	●	●	●

in Table 4 are to be considered complementary to other conventional requirements typical of access control systems. We omit these conventional requirements here, as they are not specific to our discussion.

**Policy Specification:** IoT systems are open systems that are continuously growing with more and more entities (physical nodes and users) connected. As discussed in Section 4, IoT systems should be scalable to handle this complexity (GR1). Moreover, IoT systems should be able to handle the dynamicity of devices (GR8) and of the environment (GR9). These requirements require the access control system to support the specification of fine-grained access control policies (ACR1). In particular, the system should allow the specification of different access constraints for different users and physical nodes, which are tight to the devices’ functionalities rather than to the devices themselves (Lee et al., 2017). This requirement is key for all IoT applications. For example, in a smart home, the home owner may want to define a different access control policy on the smart door lock for their parents and for the cleaner (Gerdes et al., 2014). Moreover, every IoT application requires the access control system to adapt to the dynamicity of the IoT ecosystem (ACR2), and support the specification of context-aware access control policies that impose conditions on the IoT ecosystem such as access time, location and status of the entities requesting access (Seitz et al., 2013; Tian et al., 2017). For instance, a smart home’s owner may want to grant her parents access only when they are in front of the door or specify a policy that allows the cleaner to unlock the smart door lock only on a specific day of the week and time of the day. Failing to meet these requirements results in assigning users and applications more permissions than what needed (the so-called *overprivilege* (Jia et al., 2017; Tian et al., 2017)), which can be exploited, e.g., to compromise the system or to leak sensitive information.

**Policy Management:** The scalability and dynamicity of the IoT ecosystem can also lead to challenges in policy management. An access control system should be human-centric (Tian et al., 2017) and able to

<sup>9</sup> Recall that user intervention is only requested for the first access.

effectively manage the policies of multiple entities within the IoT system (Ouaddah et al., 2017b; Salonikias et al., 2015). In particular, an access control system should keep at the minimum the effort required from users to administer access control policies for multiple entities (ACR3) and facilitate users in policy management (ACR4). Usability is particularly important in scenarios such as smart homes and health IoT where end-users are in charge to define access control policies for their IoT devices and related resources but have little or no knowledge of security (Kim et al., 2010). Moreover, in an IoT system, not all nodes might be under the control of a single authority; instead, nodes can belong to or be managed by different administrative domains interacting together (Miorandi et al., 2012). Therefore, an access control system for IoT should be able to support the management of access control policies for devices and resources across multiple domains (ACR5). This requirement is relevant in smart building applications where different parts of a building are rent out to different companies for office space (Gerdes et al., 2014). Each company is given access to the building automation components such as the HVAC, lighting and fire alarm systems. For instance, the building owner may impose a policy stating that the lighting in the corridors is not adjustable, and it is automatically switch off when no occupancy is detected. The tenant companies instead would like to override the policy to be able to adjust the light brightness and color.

**Policy Evaluation & Enforcement:** Several IoT applications like health IoT, smart buildings, connected vehicles and smart manufacturing, are characterized by a strong presence of machine-to-machine interactions. These interactions require a high level of automation for the activities performed within the IoT ecosystem (e.g., data processing, communication) to guarantee the scalability (GR2) and usability (GR10) of the IoT ecosystem. This need for automation is also reflected in the access decision making process (ACR6). In particular, access decision making should ideally require no user involvement. To determine whether an entity is allowed to access a certain device or a resource, its access requests have to be evaluated against the employed policies. However, IoT devices can have resource constraints (Salonikias et al., 2015) and, thus, computation should be minimal on the device side (GR6) and the employed network protocols lightweight (GR4,GR5). Therefore, it is desirable to limit the involvement of physical nodes in the authorization process (ACR7) and minimize the latency introduced by the authorization mechanism (ACR8), especially in IoT applications characterized by the use of resource constrained devices or latency critical applications like health IoT, connected vehicles and smart manufacturing. These requirements impose constraints also on the storage, retrieval and processing of context information for access decision making. When the IoT system is governed by multiple authorities like in connected vehicles, smart buildings and smart manufacturing applications, each administrative domain can employ an authorization mechanism based on a different access control model and/or use different data semantics. This can lead to interoperability issues during policy evaluation, which can result in an unauthorized disclosure of sensitive data and resources (Alam et al., 2011; Salonikias et al., 2015). Hence, the evaluation of an access control policy should be consistent across multiple administrative domains (ACR9). IoT systems should also guarantee the availability of nodes and, in particular, the ones involved in the authorization process. If one of these nodes fails, the access control system must still be operational (ACR10) to meet the reliability requirements of IoT (GR7). Reliability requirements should be satisfied by any access control system regardless the IoT application where they are deployed.

### 5.3. Evaluation criteria

The requirements in Table 4 define basic and desirable characteristics that an access control system for IoT should satisfy. To determine whether existing authorization solutions meet such desiderata, we have identified a number of evaluation criteria. These criteria aim to provide

the basis for an assessment of the similarities and differences amongst existing authorization frameworks for IoT and their evaluation against the requirements in Table 4. The identified criteria can be grouped into two main categories.

The first category encompasses criteria concerning the properties of the authorization system. In particular, we identify the *access control model*, *policy evaluation strategy* and the *deployment configuration* of the access control mechanism within the IoT system as the main criteria to assess whether existing authorization frameworks meet the requirements in Table 4. The second category is used to assess the purpose of the proposed framework and the assumptions underlying the IoT ecosystem. It includes *IoT architecture style*, *communication protocol* and *application domain*. The IoT architecture style (Section 2.2.2) provides insights on the capabilities of nodes and their interconnections. The communication protocol used in the framework determines the communication and computing burden on physical nodes, providing additional insights on the capabilities required from physical nodes. The application domain provides additional constraints and assumptions for the proposed framework. In our study, we have also observed that existing solutions differ significantly for *maturity level*. The extent to which access control mechanisms for IoT are actually applicable (and therefore tested) to real-world systems is important, and constitutes an additional key criteria for the evaluation of existing solutions.

In the remainder of the section, for each of the requirements in Table 4, we introduce the main criteria used to determine whether existing authorization frameworks for IoT meet the identified requirements.

**Policy Specification:** Authorization frameworks for IoT should support the specification of fine-grained (ACR1) and context-aware (ACR2) access control policies to meet the scalable and dynamic nature of IoT applications. The satisfaction of ACR1 mainly depends on the underlying access control model and its support for specifying fine-grained access control policies. In particular, we evaluate the ability of a framework to selectively control access to devices and their resources based on the level of granularity in which policies can be expressed. On the other hand, we consider ACR2 *fully satisfied* when the underlying access control model allows the specification of conditions on the context in a policy and *not satisfied* otherwise.

**Policy Management:** A key requirement of authorization frameworks for IoT is the ability to minimize the efforts required from users to administer access control policies of multiple entities (ACR3). We evaluate the ease of policy administration by considering whether an authorization framework offers a single point for policy administration and facilitates the administration of policies for a large number of entities (e.g. it does not require defining a new policy every time a new entity is added to the IoT ecosystem) (Ahmad et al., 2018). Accordingly, we consider the requirement *fully satisfied* when the access control system provides a single administrative point and adopts an access control model that minimizes the number of policies to be defined, *partially satisfied* when one of the two is supported, and *not satisfied* otherwise.

Usability (ACR4) is another important requirement especially in those IoT applications where the users in charge of defining the access control policies have no security knowledge. This requirement is considered *fully satisfied* by access control systems that support both approaches for semi-automatically or automatically generating access control policies and interfaces for policy configuration. The requirement is *partially satisfied* by approaches that support only one of these features and *not satisfied* if none of them is provided.

Management of policies across multiple domains (ACR5) is an important requirement in various IoT applications like smart manufacturing. Policy management is particularly challenging when devices are controlled by different authorities under different context conditions (Alshehri and Sandhu, 2017). We consider this requirement *satisfied* when the authorization framework supports administrative policies or provides functions that allow entities to manage and delegate the control over devices and resources across multiple domains.

**Policy Evaluation & Enforcement:** Several IoT applications require a high level of automation to guarantee scalability and usability of the IoT ecosystem. This demand also reflects on the access decision making process (ACR6). To assess this requirement, we evaluate the degree of user involvement in the access decision process required by a given authorization framework. This involvement mainly depends on the architecture style adopted by the framework. For instance, in policy-based architecture, access decisions are made autonomously based on predefined policies without user intervention, thus *fully satisfying* the requirement. On the other hand, the satisfaction of the requirement by frameworks based on a token-based architecture depends on the standard adopted. For instance, OAuth requires users to give his/her consent the first time an application requests access to a resource or a service, thus *partially satisfying* ACR6. As discussed in Section 5.1.2, this issue has been addressed by hybrid architectures, e.g. based on UMA, in which tokens are generated by means of policies. Therefore, frameworks adopting a hybrid architecture *satisfy* the requirement.

The type of architecture also affects other requirements related to the performance of the access control system and to the overall performance of the IoT system in general. The performance of an access control system depends on a number of factors: (i) the capabilities of the components involved in making and enforcing an access decision; (ii) the time taken at request time to make the decision; and (iii) the communication among the components. Authorization frameworks for IoT should not introduce communication and computation overhead on resource-constrained devices (ACR7). This clearly depends on the architecture adopted to evaluate and enforce access control policies along with the deployment of its components and the communication protocol employed by the IoT system. The architecture and its deployment determine the impact on the processing capabilities of IoT devices. Two opposite deployment solutions can be conceived for an access control system: one solution in which policy evaluation and enforcement are performed on constrained devices, and one solution in which the whole authorization process is externalized to other components. Clearly, the former does *not satisfy* the requirement while the latter *fully satisfies* it. Between these two extremes, we can find a large variety of solutions that satisfy ACR7 to a certain degree, depending on the deployment of the access control components and the architectural style adopted. In addition, the communication protocol and data exchange format have a significant impact on the communication and computation overhead on devices.

The overhead in the computation and communication should not only be minimized for constrained devices but also for the whole access control architecture (ACR8). The overhead in this case is influenced not only by the location of the components involved in policy evaluation and by the adopted communication protocol, but also by the policy evaluation strategy. In particular, the policy evaluation strategy determines when access decisions are computed. On the other hand, we analyze the deployment of the components involved in policy evaluation together with the adopted communication protocol to assess the communication overhead in access decision making. For instance, the location of the PIP can help assess the overhead required to retrieval of context information. If context information is stored in a different node from where policy evaluation is performed, there can be a delay due to its transfer.

Another critical requirement for policy evaluation and enforcement is interoperability across different administrative domains, which could potentially use different authorization frameworks (ACR9). To achieve interoperability among domains, policies should be interpreted and evaluated in the same way across different domains, e.g. using the same semantics across domains (Gusmeroli et al., 2013; Trivellato et al., 2013). We assume that this requirement is *satisfied* when the administrative domains use a standard format and data (semantic) model for communicating policies, tokens and authorization decisions, or when a solution for aligning policy semantics is provided.

Reliability and availability are other important requirements for an authorization framework for IoT (ACR10). Two main types of failures can affect the availability of the access control mechanism and the reliability of access decision decisions: failures occurring in the nodes involved in the authorization process and the lack of connectivity. In our analysis, we consider ACR10 *satisfied* by authorization frameworks that adopt measures to deal with both types of failures. On the other hand, we consider the requirements *partially satisfied* if only one type of failure is addressed and *not satisfied* if none of these failures is addressed.

## 6. Analysis of authorization frameworks for IoT

Several frameworks and architectures have been proposed in the literature to enable authorization in IoT. In this section, we review existing proposals and analyze them with respect to the requirements and evaluation criteria presented in the previous section. Then, we discuss their suitability to the IoT applications presented in Section 4.1.

### 6.1. Overview

Our analysis of the literature shows that a variety of approaches have been designed and developed to enable access control in IoT. These approaches can be broadly classified in two main categories based on the policy evaluation strategy and architecture. On one side, we have authorization frameworks (Alshehri and Sandhu, 2016, 2017; Barka et al., 2015; Bouij-Pasquier et al., 2015a; Dorri et al., 2016, 2017; Fernández et al., 2017; Garcia-Morchon and Wehrle, 2010; Guoping and Wentao, 2011; Jindou et al., 2012; Kim et al., 2011, 2012; Lee et al., 2017; Mahalle et al., 2013b; Neisse et al., 2014; Ouaddah et al., 2017a; Pinno et al., 2017; Ray et al., 2017; Salonikias et al., 2015; Sciancalepore et al., 2018; Tian et al., 2017; Ye et al., 2014; Zhang and Tian, 2010) that adopt a policy-based architecture and a runtime policy evaluation strategy. Most of these frameworks are inspired to the XACML standard. On the other side, we have frameworks (Cirani et al., 2015; Gusmeroli et al., 2013; Hernandez-Ramos et al., 2013; Hussein et al., 2017; Islam et al., 2018; Mahalle et al., 2013b; Rivera et al., 2015; Seitz et al., 2013) that adopt a hybrid-based architecture and policy evaluation strategy. A number of these frameworks build on top of OAuth by extending this standard to enable the generation of tokens based on the evaluation of access control policies like in (Cirani et al., 2015; Fernández et al., 2017) whereas Rivera et al. (2015) adopt UMA.

Regardless the type of access control architecture, different deployments and technologies are used to implement architecture. For example, PDP, PEP, PAP, and PIP could all be deployed in the cloud like in (Alshehri and Sandhu, 2016, 2017; Neisse et al., 2014) or they could all be implemented on edge devices (Kim et al., 2012; Tian et al., 2017) or a combination of both (Salonikias et al., 2015). Some works (Dorri et al., 2016, 2017; Ouaddah et al., 2017a) have also proposed authorization mechanisms based on blockchain technology.

Existing frameworks also vary significantly for maturity level. While a few (Bouij-Pasquier et al., 2015a; Cirani et al., 2015; Fremantle et al., 2014; Garcia-Morchon and Wehrle, 2010; Hernandez-Ramos et al., 2013; Hussein et al., 2017; Jindou et al., 2012; Kim et al., 2012; Lee et al., 2017; Mahalle et al., 2013a, 2013b; Neisse et al., 2014; Seitz et al., 2013) provide a prototype implementation, many (Alshehri and Sandhu, 2016, 2017; Barka et al., 2015; Dorri et al., 2016, 2017; Fernández et al., 2017; Guoping and Wentao, 2011; Gusmeroli et al., 2013; Islam et al., 2018; Kim et al., 2011; Ouaddah et al., 2017a; Pinno et al., 2017; Ray et al., 2017; Rivera et al., 2015; Salonikias et al., 2015; Sciancalepore et al., 2018; Ye et al., 2014) only remain at a conceptual level. In particular, Ray et al. (2017) and Zhang and Tian (2010) only propose an access control model tailored to IoT ecosystems and do

**Table 6**  
Analysis of existing authorization frameworks for IoT with respect to requirements.

	Policy Specification		Policy Administration			Policy Evaluation & Enforcement				
	ACR1	ACR2	ACR3	ACR4	ACR5	ACR6	ACR7	ACR8	ACR9	ACR10
Neisse et al. [78]	✓	✓	✓	✗	✗	✓	✓	✗	✗	✓
Alshehri et al. [22]	✓	✓	✓	✗	✗	✓	✓	✗	✗	✳
Alshehri et al. [23]	✓	✓	✓	✗	✓	✓	✓	✗	✗	✳
Fremantle et al. [44]	—	✗	✗	✗	✗	✳	✓	✳	✗	✗
Fernandez et al. [41]	✳	✗	✓	✗	✗	✓	✓	✗	✗	✗
Cirani et al. [32]	?	✗	✳	✗	✗	✓	✓	✗	✗	✗
Rivera et al. [89]	?	✗	✳	✗	✗	✓	✓	✳	✗	✗
Seitz et al. [98]	✓	✓	✓	✗	✗	✓	✓	✓	✗	✗
Salonikias et al. [93]	✓	✓	✓	✗	✓	✓	✓	✗	✗	✳
Ye et al. [118]	✓	✓	✗	✗	✗	✓	✗	✓	✗	✗
Hussein et al. [57]	✓	✓	✓	✗	✗	✓	✓	✳	✗	✗
Hernandez-Ramos et al. [55]	?	✓	✳	✗	✗	✓	✓	✓	✗	✗
Gusmeroli et al. [52]	?	✗	✳	✗	✗	✓	✓	✗	✗	✗
Garcia et al. [46]	✳	✓	✗	✗	✗	✓	✗	✓	✗	✗
Dorri et al. [37, 38]	✗	✗	✗	✗	✗	✓	✓	✗	✗	✳
Ouaddah et al. [81]	✳	✗	✗	✗	✗	✓	✓	✗	✗	✳
Kim et al. [66]	✗	✳	✗	✓	✗	✓	✓	✳	✗	✗
Tian et al. [108]	✗	✳	✗	✓	✗	✓	✳	✳	✗	✗
Zhang & Tian [121]	✳	✓	?	✗	✗	✓	?	?	✗	?
Jindou et al. [60]	✳	✗	✓	✗	✗	✓	✳	✗	✗	✗
Guoping & Wentao [51]	✓	✓	✓	✗	✗	✓	✓	✗	✗	✗
Islam et al. [58]	✓	✓	✓	✗	✗	✓	✓	✗	✗	✳
Barka et al. [25]	✳	✗	✓	✗	✗	✓	✓	✗	✗	✗
Cirani & Picone [31]	—	✗	✳	✗	✗	✳	✓	✗	✗	✗
Bouij-Pasquier et al. [28]	✳	✓	✓	✗	✓	✓	✓	✳	✗	✗
Lee et al. [68]	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗
Ray et al. [87]	✓	✓	?	✗	✓	✓	?	?	✗	?
Mahalle et al. [69]	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
Pinno et al. [85]	✓	✓	✳	✗	✗	✓	✗	✗	✗	✳
Kim et al. [64]	✓	✓	✓	✗	✗	✓	✓	✓	✗	✗
Sciancalepore et al. [97]	✳	✗	✗	✗	✗	✓	✗	✗	✗	✗
Mahalle et al. [70]	✗	✗	✳	✗	✗	✓	✗	✗	✗	✗
Schuster et al. [96]	?	✓	✓	✗	✓	✓	✓	✗	✓	✓

**Legend**

- ✓: satisfied
- ✳: partially satisfied
- ✗: not satisfied
- : not applicable
- ?: information missing

not provide detail of the underlying IoT architecture and access control mechanism.

6.2. Requirements assessment

A summary of our analysis with respect to the requirements is presented in Table 6. Tables 7 and 8 report a detailed analysis against the evaluation criteria concerning the access control system whereas Table 9 presents the analysis with respect to the criteria concerning the

general characteristics of the underlying IoT ecosystem. This detailed analysis provides a rationale for the analysis in Table 6. In the tables, we use symbol “—” when a requirement/criterion is not applicable to a given framework, and symbol “?” when the information is not available.

**Policy Specification:** Two key requirements that have to be satisfied by an authorization framework for IoT are the specification of fine-grained (ACR1) and context-aware (ACR2) access control policies. The satisfaction of these two requirements mainly depends on

**Table 7**

Analysis of existing authorization frameworks for IoT with respect to the evaluation criteria concerning the access control system (1).

	Access Control Model	Context Awareness	Policy Generation	Policy Configuration	Multi Domain Administration
Neisse et al. (2014)	ABAC	Yes	No	No	No
Alshehri and Sandhu (2016)	ABAC	Yes	No	No	No
Alshehri and Sandhu (2017)	ACL RBAC ABAC	Yes	No	No	Yes
Fremantle et al. (2014)	–	No	No	No	No
Fernández et al. (2017)	RBAC	No	No	No	No
Cirani et al. (2015)	?	No	No	No	No
Rivera et al. (2015)	?	No	No	No	No
Seitz et al. (2013)	ABAC	Yes	No	No	No
Salonikias et al. (2015)	ABAC	Yes	No	No	Yes
Ye et al. (2014)	ABAC	Yes	No	No	No
Hussein et al. (2017)	ABAC	Yes	No	No	No
Hernandez-Ramos et al. (2013)	?	Yes	No	No	No
Gusmeroli et al. (2013)	?	No	No	No	No
Garcia-Morchon and Wehrle (2010)	RBAC	Yes	No	No	No
Dorri et al., 2016, 2017	ACL	No	No	No	No
Ouaddah et al. (2017a)	RBAC	No	No	No	No
Kim et al. (2011)	ACL	Yes	Yes	Yes	No
Tian et al. (2017)	ACL	Yes	Yes	Yes	No
Zhang and Tian (2010)	RBAC	Yes	No	No	No
Jindou et al. (2012)	RBAC	No	No	No	No
Guoping and Wentao (2011)	UCON	Yes	No	No	No
Islam et al. (2018)	ABAC	Yes	No	No	No
Barka et al. (2015)	RBAC	No	No	No	No
Cirani and Picone (2015)	–	No	No	No	No
Bouij-Pasquier et al. (2015a)	OrBAC	Yes	No	No	Yes
Lee et al. (2017)	ACL	No	No	No	No
Ray et al. (2017)	ABAC	Yes	No	No	Yes
Mahalle et al. (2013a)	DAC	No	No	No	No
Pinno et al. (2017)	ACL + Capability + RBAC + OrBAC + ABAC + UCON	Yes	No	No	No
Kim et al. (2012)	ABAC	Yes	No	No	No
Sciancalepore et al. (2018)	ABAC	No	No	No	No
Mahalle et al. (2013b)	Trust-based	No	No	No	No
Schuster et al. (2018)	?	Yes	No	No	No

**Table 8**

Analysis of existing authorization frameworks for IoT with respect to the evaluation criteria concerning the access control system (2).

	Evaluation Strategy	Architecture Style	Deployment			
			PAP	PDP	PEP	PIP
Neisse et al. (2014)	run-time	policy-based	cloud	cloud	cloud	cloud
Alshehri and Sandhu (2016)	run-time	policy-based	cloud	cloud	cloud	cloud
Alshehri and Sandhu (2017)	run-time	policy-based	cloud	cloud	cloud	cloud
Fremantle et al. (2014)	off-line	token-based	–	external service	cloud	–
Fernández et al. (2017)	run-time	policy-based	external service	external service	local service	–
Cirani et al. (2015)	hybrid <sup>f</sup>	hybrid	external service	external service	local service	–
Rivera et al. (2015)	off-line	hybrid	external service	external service	local service	–
Seitz et al. (2013)	hybrid <sup>c</sup>	hybrid	cloud	cloud + physical node	physical node	physical node
Salonikias et al. (2015)	run-time	policy-based	cloud	edge	edge	cloud
Ye et al. (2014)	run-time	policy-based	physical node	physical node	physical node	physical node
Hussein et al. (2017)	hybrid <sup>c</sup>	hybrid	external service	external service + edge	edge	–
Hernandez-Ramos et al. (2013)	hybrid <sup>c</sup>	hybrid	external service	external service + physical node	physical node	physical node
Gusmeroli et al. (2013)	hybrid <sup>f</sup>	hybrid	external service	external service	physical node	–
Garcia-Morchon and Wehrle (2010)	run-time	policy-based	physical node	physical node	physical node	physical node
Dorri et al., 2016, 2017)	run-time	policy-based	edge	edge	edge	–
Ouaddah et al. (2017a))	run-time	policy-based	physical node	edge + physical node	physical node	physical node
Kim et al. (2011)	run-time	policy-based	edge	edge	edge	external service
Tian et al. (2017)	run-time	policy-based	edge	edge	edge	edge
Zhang and Tian (2010)	run-time	policy-based	?	?	?	?
Jindou et al. (2012)	run-time	policy-based	external service 1	external service 1	physical node	external service 2
Guoping and Wentao (2011)	run-time	policy-based	external service 1	external service 1	external service 1	external service 2
Islam et al. (2018)	hybrid <sup>c</sup>	hybrid	cloud	cloud (+physical node)	cloud (physical node)	cloud
Barka et al. (2015)	run-time	policy-based	external service 1	external service 1	external service 2	–
Cirani and Picone (2015)	hybrid <sup>f</sup>	token-based	external service	external service	local service	–
Bouij-Pasquier et al. (2015a)	run-time	policy-based	external service	edge	edge	physical node
Lee et al. (2017)	run-time	policy-based	external service	external service	external service	–
Ray et al. (2017)	run-time	policy-based	?	?	?	?
Mahalle et al. (2013a)	run-time	hybrid	physical node	physical node	physical node	–
Pinno et al. (2017)	run-time	policy-based	physical node	physical node	physical node	physical node + external service
Kim et al. (2012)	run-time	policy-based	edge	edge	edge	edge
Sciancalepore et al. (2018)	run-time	policy-based	physical node	application node + physical node	physical node	external service
Mahalle et al. (2013b)	run-time	policy-based	physical nodes	physical node	physical node	peers
Schuster et al. (2018)	run-time	policy-based	?	?	?	physical node + external service

**Table 9**  
 Analysis of existing authorization frameworks for IoT. Symbol † indicates that only a toy example is provided for demonstration purposes, but the framework is not designed specifically for that application domain.

	IoT Architecture style	Communication Protocol	Data Format	Cross-domain Data Semantics	Node Failure Robustness	Application Domain	Maturity Level
Neisse et al. (2014)	centralized	MQTT	?	No	Yes	–	prototype
Alshehri and Sandhu (2016)	centralized	–	–	No	Yes	Lighting†	design
Alshehri and Sandhu (2017)	centralized	–	–	No	Yes	Connected Vehicles†	design
Fremantle et al. (2014)	centralized	MQTT	JSON	No	No	–	prototype
Fernández et al. (2017)	centralized	–	–	No	No	–	design
Cirani et al. (2015)	distributed	6LoWPAN CoAP	?	No	No	–	prototype
Rivera et al. (2015)	centralized	–	–	No	No	Traffic Lights†	design
Seitz et al. (2013)	distributed	CoAP	XACML JSON	No	Yes	–	prototype
Salonikias et al. (2015)	connected	–	–	No	Yes	Connected Vehicles	design
Ye et al. (2014)	distributed	–	–	No	No	Wireless Sensor Network†	design
Hussein et al. (2017)	connected	?	JSON	No	No	Smart Home†	prototype
Hernandez-Ramos et al. (2013)	connected	6LoWPAN CoAP	JSON	No	Yes	–	prototype
Gusmeroli et al. (2013)	connected	–	–	No	Yes	–	design
Garcia-Morchon and Wehrle (2010)	distributed	6LoWPAN	?	No	No	Medical Sensor Networks	prototype
Dorri et al., 2016, 2017	distributed	–	–	No	Yes	Smart Home†	design
Ouaddah et al. (2017a)	distributed	–	–	No	No	–	design
Kim et al. (2011)	connected	–	–	No	No	Smart Home	design
Tian et al. (2017)	connected	HTTP	?	No	Yes	Smart Home	product
Zhang and Tian (2010)	?	–	–	No	No	–	design
Jindou et al. (2012)	connected	HTTP	JSON	No	No	Smart Home†	prototype
Guoping and Wentao (2011)	centralized	–	–	No	No	–	design
Islam et al. (2018)	connected	–	JSON	No	Yes	Health Prescription Assistant	design
Barka et al. (2015)	centralized	–	–	No	No	–	design
Cirani and Picone (2015)	distributed	CoAP	?	No	No	–	prototype
Bouij-Pasquier et al. (2015a)	connected	CoAP	JSON	No	No	Health IoT†	prototype
Lee et al. (2017)	distributed	Wi-Fi	?	No	No	Smart Lock†	prototype
Ray et al. (2017)	?	–	–	No	No	Remote Healthcare Monitoring	design
Mahalle et al. (2013a)	distributed	Wi-Fi	?	No	No	–	prototype
Pinno et al. (2017)	distributed	–	–	No	Yes	–	design
Kim et al. (2012)	centralized	ZigBee	?	No	No	Smart Home	prototype
Sciancalepore et al. (2018)	connected	–	–	No	No	–	design
Mahalle et al. (2013b)	distributed	?	?	No	No	–	prototype
Schuster et al. (2018)	distributed	HTTP	?	No	Yes	Smart Home	prototype



the access control model adopted by the authorization framework to express access control policies.<sup>10</sup>

A number of existing frameworks adopt ABAC (Alshehri and Sandhu, 2016, 2017; Hussein et al., 2017; Islam et al., 2018; Kim et al., 2012; Neisse et al., 2014; Ray et al., 2017; Salonikias et al., 2015; Seitz et al., 2013; Ye et al., 2014) or UCON (Guoping and Wentao, 2011) as the underlying access control model and, therefore, satisfy both the requirements. In fact, both ABAC and UCON provide a flexible means to specify access control policies as conditions on attributes related to the entity requesting access and the resource being accessed. Moreover, ABAC intrinsically supports the specification of environment attributes that can be used to denote the context in which permissions hold, whereas UCON allows the specification of context-sensitive access constraints through conditions. Of particular interest is the work by Neisse et al. (2014) that uses event-condition rules to express fine-grained authorizations: the action in the rule is performed when the event is observed and the condition is satisfied. Besides being able to specify conditions on the context using data values acquired in a given moment in time (e.g., GPS location), these rules make it possible to specify context situations, which are composed data types modeling specific, complex conditions (e.g., the entity requesting access has to be not more than 100 m from other entities nearby).

An exception is the framework proposed in (Sciancalepore et al., 2018). Although based on ABAC, this framework only supports the verification of subject attributes, thus not satisfying ACR2 and only partially satisfying ACR1. Similarly, frameworks based on RBAC or OrBAC (Barka et al., 2015; Bouij-Pasquier et al., 2015a; Fernández et al., 2017; Garcia-Morchon and Wehrle, 2010; Guoping and Wentao, 2011; Jindou et al., 2012; Ouaddah et al., 2017a; Zhang and Tian, 2010) only partially satisfy ACR1 since these models only allow abstracting permissions at the level of role. Moreover, core RBAC does not support the notion of context and, thus, frameworks based on this model do not satisfy ACR2. On the other hand, OrBAC allows the specification of context-sensitive access constraints by explicitly representing the context in which permissions holds and, thus, frameworks based on this model satisfy ACR2.

We also found a few frameworks (Dorri et al., 2016; Kim et al., 2011; Lee et al., 2017; Mahalle et al., 2013a; Tian et al., 2017) based on identity-based access control models and, in particular, ACLs. These models, in general, do not satisfy requirements ACR1 and ACR2 because they only allow a direct assignment of access rights to users and do not support the notion of context. It is worth noting that some of the frameworks based on RBAC (i.e. (Garcia-Morchon and Wehrle, 2010; Zhang and Tian, 2010),) and DAC (i.e. (Kim et al., 2011; Tian et al., 2017),) have been extended, often in an ad-hoc fashion, to account for contextual information during policy evaluation. For instance, in (Garcia-Morchon and Wehrle, 2010) access is granted based on the user's health condition and other context information: if the user's health condition is critical, access is given to any doctor or medical staff to face emergency. However, some of them are limited in the type of context information that can be specified in policies, thus only partially satisfying ACR2. In particular, the frameworks in (Kim et al., 2011; Tian et al., 2017) only support the specification concerning the location of the requester.

Some authorization frameworks for IoT are not based on standard access control models. For instance, Mahalle et al. (2013b) propose a trust-based access control model in which access is granted based on the trustworthiness of the requester. This model, however, only allows the specification of permissions at device level, resulting in coarse grained access control policies. This leads to the problem of overprivilege as

<sup>10</sup> Note that we consider the access control model in which users have to specify their policies. Therefore, even if a framework use capabilities tokens for authorization purposes, here we identify the access control model used to generate the tokens.

users and applications have more capabilities than needed (Lee et al., 2017; Tian et al., 2017). Moreover, it does not account for context information for access decision making, thus not meeting ACR2.

Finally, we observed that the framework in (Fremantle et al., 2014) does not evaluate access requests against access control policies represented according to a specific model. This framework is based on the OAuth 2.0 authorization protocol where access to a resource is granted explicitly by the resource owner. On the other hand, the frameworks in (Cirani et al., 2015; Gusmeroli et al., 2013; Hernandez-Ramos et al., 2013; Rivera et al., 2015) use policies to generate capability tokens but they do not discuss the access control model employed to generate those tokens. Accordingly, ACR1 and the access control model for these frameworks are marked with symbol “?” in Tables 6 and 7, respectively.

The analysis of existing frameworks in light of requirements ACR1 and ACR2 shows that existing frameworks have adopted a variety of access control models (i.e., ABAC, UCON, OrBAC, RBAC, ACL or capabilities) to express access control policies or rely on user intervention (i.e., frameworks based on OAuth 2.0). However, only frameworks based on ABAC and UCON allow the specification of fine-grained and context-aware access control policies and therefore fully satisfy both requirements. It is interesting to observe that we did not find any framework that uses MAC as the underlying access control model. We speculate that this is due to the fact that MAC is a very static and rigid access control model and, thus, not suitable to cope with the dynamicity characterizing most IoT applications.

**Policy Management:** Several IoT applications require the access control system to deal with the management and protection of several entities (ACR3). As discussed in Section 5.3, this requirement is satisfied by authorization frameworks that offer users a single point for policy administration and make use of a flexible access control model. These criteria are usually met by authorization frameworks based on ABAC and UCON (and on RBAC and OrBAC to a certain extent) that are used in combination with a message broker to regulate the subscription to and the publishing of resources (e.g. (Alshehri and Sandhu, 2016; Alshehri and Sandhu, 2017; Neisse et al., 2014),). In particular, ABAC and UCON provide a flexible approach for the specification of access control policies. RBAC and OrBAC have been introduced to simplify the specification and management of access control policies compared to identity-based models (e.g., DAC and MAC), but they may require defining new roles with specific permissions leading to role explosion (Elliott and Knight, 2010). A single administrative point is also offered by authorization frameworks deployed as external services (Barka et al., 2015; Cirani et al., 2015; Fernández et al., 2017; Guoping and Wentao, 2011; Hussein et al., 2017; Jindou et al., 2012; Rivera et al., 2015) or in the cloud (Islam et al., 2018; Seitz et al., 2013). This is also the case of the frameworks proposed by Salonikias et al. (2015) and by Bouij-Pasquier et al. (2015a) in which the PAP is deployed in the cloud or in an external service respectively, whereas the other components of the access control system are deployed in edge nodes. On the other hand, the requirement is not satisfied by frameworks in which policies have to be deployed in physical nodes and/or adopt an identity-based access control model (e.g. (Dorri et al., 2016; Dorri et al., 2017; Garcia-Morchon and Wehrle, 2010; Kim et al., 2011; Lee et al., 2017; Mahalle et al., 2013a; Ouaddah et al., 2017a; Tian et al., 2017; Ye et al., 2014),). In fact, both approaches require to set the permissions for each new device. In particular, some frameworks use ACLs, motivated by the simplicity of policy specification, but they have not been proven at a large scale. An exception is the framework by Mahalle et al. (2013b) that, although policies are deployed in physical nodes, they are predefined and do not need to be deployed when new devices or services are added to the system. However, the evaluation of these policies requires users to define and retrieve information about experience, knowledge and recommendations about other devices. Thus, we mark ACR3 partially satisfied by this framework. On the other hand, ACR3 is not satisfied by frameworks of Fremantle et al. (2014). This framework is based on OAuth and, thus, a user has to manually grant

permission to each application and device that request access to his/her resources (see Section 5.1.2). The frameworks proposed in (Cirani et al., 2015; Gusmeroli et al., 2013; Hernandez-Ramos et al., 2013; Rivera et al., 2015) could potentially satisfy the requirement depending on the access control model adopted to issue capability tokens. Since those frameworks offer a single point for policy administration but the access control model is not discussed, we mark the requirement partially satisfied for these frameworks.

An aspect that is neglected by most of the existing authorization frameworks for IoT is usability (ACR4). This is a key requirement in scenarios like smart homes and health IoT where the users that are in charge of protecting devices and resources, often lack the security expertise necessary to specify access control policies (Kim et al., 2010; Mazurek et al., 2010). Therefore, access control systems for these IoT applications should provide users with an interface that suggests the access control policies to be enforced, displays the current policies, and allows to modify the policies as needed. Only two frameworks (Kim et al., 2011; Tian et al., 2017) provide a full solution to usability. Kim et al. (2011) proposed an access control system called CARA, which automatically suggests the access control policies to be assigned to the visitors of a smart home. These policies are defined based on three main access control constraints: *presence*, which requires the visitor to be in the house in order to access a device/resource, *logs*, which requires the device to maintain logs, and *ask for permission*, which requires the visitor to explicitly ask the permission to access a device/resource to the home owner. These three constraints are used to define four basic policy configurations – full, restricted, partial and minimal control – on resources/devices that reflect the level of trust the home owner places into the visitor. For instance, if the visitor is highly trusted by the home owner, e.g. a family member, he will be assigned to the full policy configuration that gives him full access and control to all devices and resources in the house when he is physically in the house. Policy configurations are automatically assigned to users based on the social relationship between the home owner and the user, which is inferred based on the social network graph information or phone usage information, e.g. users called more frequently. The automatic assignment of a policy configuration to a visitor, on one side, simplifies the owner's task but, on the other side, can lead to assign a wrong policy configuration to visitors. The authors suggest that policy configurations are preloaded into the devices by the manufacturer and changed manually by the home owner if needed.

Existing authorization frameworks typically assume that policies are predefined by users (possibly with automated aid as in the case of (Kim et al., 2011)). However, this permission model is not suitable when users have to confirm the permissions asked by IoT applications. In this setting, applications can require more permissions than what actually needed, thus resulting in *overprivilege*. To address this issue while minimizing user burden, Tian et al. (2017) propose an approach to automatically generate access control policies to grant access to devices/resources to the smart phone app that the home owner uses to control the resources/devices. The approach derives the policies by identifying any possible discrepancy between the functionality exhibited by the mobile app's code, e.g. unlocking the door, and the app description, e.g. switching on the coffee machine. If the functionality of the app code matches the one in the app description, the app is automatically authorized to perform it. Otherwise, if there is a mismatch, the app is automatically blocked. If the analysis of the app code reveals that the app not only switches on the coffee machine but also unlock the door, the unlock door functionality is flagged as a mismatched functionality. The generated policies are displayed to users through the mobile app interface: the verified functionalities are labeled in green while mismatches are labeled in red so that the home owner can understand that they represent a potential risky behavior of the app.

Another aspect that has been marginally investigated is the management of access control policies across multiple administrative domains (ACR5) that is particularly relevant in IoT applications like connected

vehicles, smart buildings and smart manufacturing. Our analysis of the literature shows that most of the existing authorization frameworks fail to meet this requirement as they implicitly assume that resources and devices are under the control of a single authority. Notable exceptions are the works in (Alshehri and Sandhu, 2017; Ray et al., 2017), which propose an approach for policy administration tailored to IoT ecosystems, and the work by Salonikias et al. (2015). In particular, the latter introduces the notion of propagation policy and proposes a policy propagation method to update all PDPs (deployed in edge nodes and under the control of possible different authorities) when policies are modified in a centralized PAP. On the other hand, Bouij-Pasquier et al. (2015a) introduce a collaboration layer to handle multi-party collaborative interactions. In particular, the authors propose a negotiation of access rules for cross-domain sharing of resources and information.

**Policy Evaluation & Enforcement:** The ability to automate the evaluation of an access request (ACR6) is an important requirement in all IoT applications where a multitude of devices and users share information. Assuming that a user is always available to evaluate if access to certain resource should be granted is not realistic. To automate the evaluation of an access request, existing frameworks adopt either a policy-based (Alshehri and Sandhu, 2016, 2017; Barka et al., 2015; Bouij-Pasquier et al., 2015a; Dorri et al., 2016, 2017; Fernández et al., 2017; Garcia-Morchon and Wehrle, 2010; Kim et al., 2011, 2012; Lee et al., 2017; Mahalle et al., 2013b; Neisse et al., 2014; Ouaddah et al., 2017a; Pinno et al., 2017; Ray et al., 2017; Salonikias et al., 2015; Sciancalepore et al., 2018; Tian et al., 2017; Ye et al., 2014; Zhang and Tian, 2010) or a hybrid architecture (Cirani et al., 2015; Gusmeroli et al., 2013; Hernandez-Ramos et al., 2013; Hussein et al., 2017; Islam et al., 2018; Rivera et al., 2015; Seitz et al., 2013). In a policy-based architecture access requests are evaluated against a predefined set of access control policies; while in hybrid architectures authorization tokens are issued based on the evaluation of access control policies. The only frameworks that do not fully satisfy the requirement is the one by Fremantle et al. (2014) and Cirani and Picone (2015). The framework proposed in (Fremantle et al., 2014) is based on the OAuth protocol, which requires the resource owner to grant access to the application the first time an authorization token is issued. Similarly, Cirani and Picone (2015) require the resource owner's involvement in the issuing of tokens. In particular, they account for three operational modes to obtain the tokens: *owner-to-owner*, in which a user registers his/her own device and obtains a token with all permissions on the device; *reactive owner-to-any*, in which the owner grants permission upon a user's request; and *proactive owner-to-any*, in which the owner proactively grants permission to a user.

Another key requirement for authorization frameworks designed for IoT applications is that they should not introduce communication and computation overhead on resource-constrained devices (ACR7). This requirement is typically addressed by outsourcing the most computationally expensive operation, namely policy evaluation, to an external service while performing only the enforcement of the access decision on constrained devices. Our analysis shows that most of the frameworks that adopt a policy-based architecture (Alshehri and Sandhu, 2016, 2017; Barka et al., 2015; Bouij-Pasquier et al., 2015a; Dorri et al., 2016, 2017; Fernández et al., 2017; Kim et al., 2011; Lee et al., 2017; Neisse et al., 2014; Ouaddah et al., 2017a; Salonikias et al., 2015; Tian et al., 2017) externalize the PDP and the PAP (i.e., these components are not deployed in the physical node), thus fully satisfying the requirement. The only exceptions are the frameworks proposed in (Garcia-Morchon and Wehrle, 2010; Ye et al., 2014) in which the PDP and the PAP run on the constrained device. Similarly, frameworks that rely upon a token-based or a hybrid architecture (Cirani et al., 2015; Gusmeroli et al., 2013; Hernandez-Ramos et al., 2013; Hussein et al., 2017; Islam et al., 2018; Rivera et al., 2015; Seitz et al., 2013) fully satisfy the requirement. These frameworks employ dedicated services for the generation and issue of authorization tokens, and only the validation of the

token is performed on the device. However, an aspect that should be considered is the size and format of the authorization token that could introduce a computation overhead on a constrained device. Lightweight standards to represent tokens like JSON should be preferred over XML-based formats like the one supported by SAML.

On the other hand, the performance of an access control system (ACR8) not only depends on the location of the components involved in the policy evaluation and communication protocol, but also on the policy evaluation strategy. Frameworks that use an off-line evaluation strategy (i.e. (Fremantle et al., 2014; Rivera et al., 2015),) or a hybrid strategy in which only context constraints are verified at run-time (hybrid<sup>c</sup>) and their verification does not require retrieving information from other components (e.g. (Garcia-Morchon and Wehrle, 2010; Hernandez-Ramos et al., 2013; Seitz et al., 2013; Ye et al., 2014),), do not introduce latency in the access decision making process. Similarly, latency is limited if policy evaluation is performed on the edge like in (Bouij-Pasquier et al., 2015a; Hussein et al., 2017; Kim et al., 2012; Tian et al., 2017). On the other hand, policy-based frameworks in which the access control mechanism is deployed in the cloud or provided as an external service (e.g. (Alshehri and Sandhu, 2016; Alshehri and Sandhu, 2017; Barka et al., 2015; Fernández et al., 2017; Jindou et al., 2012; Neisse et al., 2014; Zhang and Tian, 2010),) might introduce delay due to additional communication. This is also the case of frameworks that require validating tokens at run-time (hybrid<sup>d</sup>) like in (Cirani and Picone, 2015; Cirani et al., 2015; Gusmeroli et al., 2013), or that require retrieving contextual information from external sources or from the cloud like in (Jindou et al., 2012; Kim et al., 2010; Salonikias et al., 2015; Schuster et al., 2018; Zhang and Tian, 2010). On top of this, the communication protocol has a significant impact on the overall performance of the IoT ecosystem where frameworks based on lightweight protocols like MQTT and CoAP provide better performance compared to the ones based on HTTP. Frameworks based on blockchain technology (e.g. (Dorri et al., 2016; Dorri et al., 2017; Ouaddah et al., 2017a; Pinno et al., 2017),) also do not satisfy the requirement due to time required to confirm a transaction. Every time an access control policy has to be added to or retrieved from the blockchain, a new transaction has to be created and added to the blockchain. Before a transaction can be added to the blockchain, special nodes called miners run a consensus protocol that requires them to verify each transaction. The time to complete the validation process is typically in the order of minutes (Ouaddah et al., 2016), which is clearly unsuitable for most IoT applications, especially for the ones that are latency sensitive.

Other key requirements for policy evaluation are interoperability (ACR9) and reliability/availability of components (ACR10) involved in the evaluation of the policies. However, despite their importance these two requirements are only marginally considered by existing authorization frameworks for IoT. Some of the frameworks only scratch the surface of the interoperability problem because they use a standard like XACML to specify the access control policies (Seitz et al., 2013) or they encode the capability token in JSON (Fremantle et al., 2014; Hernandez-Ramos et al., 2013; Hussein et al., 2017; Jindou et al., 2012; Seitz et al., 2013). Interestingly, Seitz et al. (2013) provide an encoding of SAML assertions in JSON, while the others propose an ad-hoc format to encode tokens. However, using a standard only facilitates the exchange of policies or tokens across multiple domains but not their interpretation. If different authorities define their policies based on different semantic models, the collaborative evaluation of these policies can result in granting access to users for which access should be denied.

Reliability and availability (ACR10) is *fully satisfied* by those frameworks that can tolerate the failure of an architectural component and of the communication among them. Most of the frameworks partially satisfy the requirement because they only address the reliability/availability of the components but not of the communication among them. To address the failure of an architectural component, three main solutions have been adopted by existing authorization frameworks. Some frameworks have deployed the components in the cloud

(Alshehri and Sandhu, 2016, 2017; Islam et al., 2018; Neisse et al., 2014; Salonikias et al., 2015; Seitz et al., 2013), which guarantees that the components are evenly distributed across different servers, which are connected to work as one. Therefore, if one server fails, downtime is avoided. Salonikias et al. (2015) instead ensure reliability and availability by replicating the PDP and the PEP and by defining propagation policies that specify how access control policies should be exchanged between PDPs. Frameworks based on blockchain (Dorri et al., 2016, 2017; Ouaddah et al., 2017a; Pinno et al., 2017) propose to deploy and maintain a copy of the components of the authorization framework in all nodes forming the blockchain, thus ensuring resilience against failures of architecture components. The only framework proposed that fully satisfies ACR10 is the one proposed by Neisse et al. (2014), which adopts a reliable communication protocol like MQTT besides addressing the reliability of the architectural components.

**Implementation and Evaluation:** An important aspect is the *applicability* of an access control framework to real IoT applications. In this respect, most of the proposed frameworks (Alshehri and Sandhu, 2016, 2017; Barka et al., 2015; Dorri et al., 2016, 2017; Fernández et al., 2017; Guoping and Wentao, 2011; Gusmeroli et al., 2013; Islam et al., 2018; Kim et al., 2011; Ouaddah et al., 2017a; Rivera et al., 2015; Salonikias et al., 2015; Sciancalepore et al., 2018; Ye et al., 2014) only present the architecture of the access control mechanism and demonstrate the authorization flow among the components based on a realistic IoT use cases. For example, Dorri et al. (Rivera et al., 2015) have illustrated their access control framework based on a smart home scenario. However, use cases do not provide insights on the effectiveness of the framework in realistic IoT settings. Only implementing the framework on a real IoT system and evaluating its performance and usability can provide such insights. Nonetheless, only few of the proposed frameworks have been implemented and evaluated (Cirani et al., 2015; Garcia-Morchon and Wehrle, 2010; Lee et al., 2017; Mahalle et al., 2013b; Neisse et al., 2014; Seitz et al., 2013), while other works only report a prototype implementation (Bouij-Pasquier et al., 2015a; Fremantle et al., 2014; Hernandez-Ramos et al., 2013; Hussein et al., 2017; Jindou et al., 2012; Kim et al., 2012). For instance, Neisse et al. (2014) have proposed an authorization framework for MQTT brokers. The enforcement of access control policies is performed by a PEP that is integrated into the browser, while policy evaluation is done by an external PDP and Context Manager. The MQTT broker has been implemented using the Mosquitto library and its performance evaluated in terms of overhead introduced in the communication by implementing the PEP in the MQTT broker. Cirani et al. (2015) have instead focused on evaluating the performance of their access control framework on constrained devices. In particular, they evaluated the energy and memory consumption of policy evaluation on a Contiki-based devices. To run the evaluation, they used the Cooja simulator and considered Zolertia Z1 nodes with 92 KB ROM and 8 kb RAM. Similarly, Garcia et al. (Garcia-Morchon and Wehrle, 2010) have evaluated the performance of their framework on constrained devices but using a real testbed rather than a simulation environment like Cooja. The testbed consisted of Arduino Mega 2560 board3 with 16 MHz processor, 256 kB of Flash Memory, 8 kB of SRAM, and 4 kB of EEPROM.

### 6.3. Discussion

Our analysis of the literature shows that there is no *one-size-fits-all* authorization framework for all IoT applications. Each IoT application has its own set of requirements that should be satisfied when designing an authorization framework specific to that application. The main difference lies in the requirements imposed by each application on policy management and evaluation, while the requirements on policy specification are the same for all IoT applications. Regardless the specific application, an authorization framework for IoT should support the specification of fine-grained (ACR1) and context-aware (ACR2) poli-

cies. Both requirements are satisfied when the authorization framework adopts either ABAC or UCON as the underlying access control model (see Section 6.2 for a detailed discussion).

With respect to the requirements on policy management and evaluation, we can divide IoT applications in three main groups: the first group is formed by smart homes, the second one by health IoT and the third group is composed by smart buildings, connected vehicles and smart manufacturing. Requirements imposed by each group of applications on policy management and evaluation are quite different as shown in Table 5. The only exception is requirement ACR10 on the reliability and availability of the architectural components involved in the authorization process, which is required by each group of applications.

In smart homes, home owners are in charge of specifying the access control policies to protect a relatively small number of IoT devices but they typically do not have the necessary security knowledge. Therefore, usability (ACR4) is a key requirement in smart homes and should be addressed by minimizing the efforts of home owners in specifying access control policies. The ideal authorization framework for smart homes should be based on a centralized and policy-based architecture where access decision are made based on access control policies that are not defined by the home owners but automatically generated taking into account the context of access. Moreover, the PAP should support home owners in the configuration and modification of their policies. Since latency can be tolerated in smart homes applications, a run-time policy evaluation strategy can be adopted. The PDP could be deployed on edge devices like IoT gateways or on a local cloud. The frameworks presented in (Kim et al., 2011, 2012; Tian et al., 2017) are specifically designed for smart homes but they do not address all relevant requirements for this IoT application. Kim et al. (2011) and Tian et al. (2017) provide a mechanism to suggest access control policies to home owners, but the policies generated are coarse-grained and only impose simple conditions on the environment. In contrast, Kim et al. (2012) does not consider the usability issues related to the specification of access control policies by lay users, but allow the specification of fine-grained and context-aware access control policies. None of the frameworks ensure reliability and availability of the components involved in the policies evaluation and enforcement.

Similarly to smart homes, health IoT applications require lay users to be responsible for the specification of access control policies and, thus, have to satisfy the same requirements with respect to usability. In contrast, health IoT applications involve a large number of IoT devices, e.g. insulin pumps and pacemakers, that feed sensitive medical data directly in patients' electronic healthcare records. Therefore, it is important that the authorization framework minimizes the effort of lay users (e.g., patients and medical staff) in administering policies for multiple devices (ACR3), takes into account the constrained capabilities of medical devices (ACR7) and reduces the latency in decision making (ACR8), which can be potentially life threatening. In order to meet all these requirements, an authorization framework for health IoT applications should adopt a hybrid policy evaluation strategy and architecture. Similarly to smart homes, the PDP should be configured with access control policies that are automatically generated rather than being defined by patients. Moreover, to take into account the constrained capabilities of medical devices and minimize the latency to evaluate and enforce the policies, the PDP should be deployed on an edge device while the PEP could be located on the devices. Garcia et al. (Garcia-Morchon and Wehrle, 2010) and Ray et al. (2017) have proposed an authorization framework that enables remote patient monitoring. While Garcia's framework supports the specification of context-aware access control policies and a lightweight mechanism that efficiently runs on constrained sensor nodes, Ray and colleagues only propose a fine-grained access control model inspired to the XACML and NIST NGAC (Ferraiolo et al., 2016) standards. Both frameworks do not provide a solution to address key requirements like usability of policy specification and adopt a runtime policy evaluation strategy and a

policy-based architecture that increase the latency of access decision making.

Unlike smart homes and health IoT, which are characterized by a high user involvement, smart buildings, connected vehicles and smart manufacturing applications mainly involve a large number of IoT devices that directly communicate with each other. Often these devices are not managed by a single authority but they belong to different administration domains. Therefore, unlike smart homes and health IoT applications, usability is not a key requirement. On the other hand, the ability of supporting the management of policies across different domains (ACR5), ensuring interoperability among domains (ACR9), and automating an access control decision (ACR6) are fundamental requirements. Similarly to health IoT applications, smart buildings, connected vehicles and smart manufacturing applications are time-critical applications and, therefore, the authorization framework should adopt a policy evaluation strategy and an architecture that reduce the latency of the access decision making process (ACR8). Therefore, the authorization framework should be similar to the one discussed for health IoT applications but the PDP should also provide functionalities to take access decisions based on policies from different administrative domains and guarantee the correct interpretation of these policies. While there are no authorization frameworks specific to smart buildings and smart manufacturing applications, Salonikias et al. (2015) proposed an authorization framework for connected vehicles that satisfies most of the above requirements except the one related to latency. The authors proposed an XACML-like architecture that consists of multiple PDPs and PEPs located at the edge, while a single PAP deployed in the cloud is responsible to maintain and propagate access control policies to the PDPs. However, the communication among the PAP and the PDPs increases the time needed to take an access decision, thus not satisfying ACR8.

## 7. Lessons learned and open challenges

This section summarizes the lessons learned that should be taken into account when designing an authorization mechanism for IoT and discusses open challenges.

*There is no need of new access control models.* Several of the analyzed works indicate that access constraints for IoT systems should account for the context (e.g., location, time) and often propose ad-hoc (extensions of) policy languages to represent such access constraints. As discussed in the previous sections, both ABAC and UCON have proven capable to express a large range of access control policies and allow the specification of fine-grained and context-aware access control policies. In particular, these models allow the specification of permissions at the level of devices' functionalities, which is necessary to avoid application overprivilege. Therefore, regardless the IoT application, any authorization framework should adopt one of these models as the underlying access control model.

*There is no one-size-fit-all authorization framework for all IoT applications.* Although many authorization frameworks for IoT have been proposed, only few of them have been designed for a specific IoT application and even fewer address all the unique requirements of the application that they are meant to protect. The main research challenge is thus to design an authorization framework that satisfies the requirements related to policy management and evaluation specific to a target IoT application.

*OAuth is not suitable for most IoT applications.* Many initiatives from standardization bodies and industry aim to adapt the OAuth authorization protocol to IoT. OAuth could be potentially applied to smart home applications because it partially solves the usability issues related to lay users being in charge of specifying access control policies. In particular, OAuth does not require home owners to specify access control policies but they have to grant access to a smart device or appliance when it is requested. However, the permissions granted by home owners are coarse-grained because they give access to the whole IoT device rather

than only on specific resources and services provided by the device itself. Moreover, OAuth requires human intervention to take an authorization decision, which makes it unsuitable for IoT applications involving a high number of devices that directly communicate with each other like connected vehicles and smart manufacturing and for IoT applications involving machine-to-machine interaction. UMA improves over OAuth with respect to human involvement in access decision making because it does not require the resource owner to be online at the time of access request but it handles requests based on access control policies predefined by the owner. However, this introduces usability issues because home owners might not be security experts. UMA has also been proposed to achieve privacy and compliance with data protection principles like informed consent because users explicitly grant access to their personal data and determine who access their data, for how long, and under what circumstances. While UMA certainly empowers users with control over their personal data, it does not address the requirements on collecting users' consent imposed by the GDPR,<sup>11</sup> the new EU regulation on data protection. The GDPR requires individuals to explicitly give their consent to collect their personal data and that when the consent is given they are informed of data collection purposes and the nature of the data processing in a clear, easy to understand and concise language. The mechanism adopted by UMA to obtain users' consent is explicit but not informed because consent is collected without any clear and concise explanation of the purpose for which the data are accessed and how they are going to be processed.

Besides distilling the above lessons learned, we have identified a number of aspects that have been neglected by most of existing frameworks and for which solutions are yet to be provided.

**Usability:** Usability is a largely unexplored aspect for IoT applications like smart home and health IoT, which are characterized by a high user involvement. The few efforts that have aimed to address usability issues do not take into account that these applications involve both machine-to-machine and user-to-machine interactions. They either generate access control policies for different users that could interact with the IoT devices or to restrict a device's access to another device.

**Multi-domain policy administration:** Most of the proposed frameworks assume that access control policies for IoT devices and resources are managed by a single authority. However, this is only a realistic assumption for smart home applications where only home owners are in charge of protecting the smart devices and appliances in their home. Other applications like smart buildings, connected vehicles and smart manufacturing involve users, IoT devices, services managed by authorities that belong to different domains. An authorization framework for these applications should support a policy governance model able to reconcile policies from different domains. The framework should also provide a solution to resolve interoperability issues due to the use of a different semantic to define the policies across different domains.

**Reliability and Availability:** Reliability and availability of the components involved in the evaluation and enforcement of the access control policies is as important as the reliability and availability of the IoT devices deployed in a particular IoT application. If the PDP fails to take an authorization decision, this affects the performance of the whole IoT ecosystem. Despite the importance of this requirement, only one of the analyzed authorization frameworks for IoT has addressed it.

**Lack of validation:** Our analysis reveals that most frameworks (e.g. (Alshehri and Sandhu, 2016; Alshehri and Sandhu, 2017; Fernández et al., 2017; Salonikias et al., 2015; Ye et al., 2014),) are still at a conceptual level and lack a proof-of-concept implementation. Although a few frameworks have been implemented (e.g. (Cirani et al., 2015; Neisse et al., 2014; Seitz et al., 2013),), they often lack a validation within large scale IoT systems. This makes it difficult to evaluate whether they meet requirements such as scalability and performance, and thus to assess their suitability to cope with realistic IoT scenarios. We believe that

this is a step necessary for the transfer of research efforts into real IoT applications.

## 8. Conclusion

This paper has provided an analysis of existing authorization frameworks for IoT. Our goal was to identify the main research trends and developments in this area.

We have identified several important requirements to support access control in IoT driven by the non-functional requirements to be met by IoT systems and the demands of IoT applications. By analyzing the current state-of-the-art against these requirements, we observed that there is no one-size-fits-all access control system for all IoT applications. The main research challenge in the design of an authorization framework for IoT lies in devising an architecture that meets the requirements specific to the target IoT application.

## Acknowledgments

This work is partially funded by the ITEA3 projects APPSTACLE (15017).

## References

- Aazam, M., Khan, I., Alsaffar, A.A., Huh, E.-N., 2014. Cloud of things: integrating internet of things and cloud computing and the issues involved. In: Proceedings of International Bhurban Conference on Applied Sciences & Technology. IEEE, pp. 414–419.
- Abdmeziem, M.R., Tandjaoui, D., Romdhani, I., 2016. Architecting the internet of things: state of the art. In: Robots and Sensor Clouds. Springer, pp. 55–75.
- Ahmad, T., Morelli, U., Ranise, S., Zannone, N., 2018. A lazy approach to access control as a service (ACaaS) for IoT: an AWS case study. In: Proceedings of the 23rd ACM Symposium on Access Control Models and Technologies. ACM, pp. 235–246.
- Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., Ayyash, M., 2015. Internet of things: a survey on enabling technologies, protocols, and applications. IEEE Commun. Surv. Tutorials 17 (4), 2347–2376.
- Alam, S., Chowdhury, M.M., Noll, J., 2011. Interoperability of security-enabled internet of things. Wireless Pers. Commun. 61 (3), 567–586.
- Alshehri, A., Sandhu, R., 2016. Access control models for cloud-enabled internet of things: a proposed architecture and research agenda. In: Proceedings of International Conference on Collaboration and Internet Computing. IEEE, pp. 530–538.
- Alshehri, A., Sandhu, R., 2017. Access control models for virtual object communication in cloud-enabled IoT. In: Proceedings of International Conference on Information Reuse and Integration. IEEE, pp. 16–25.
- L. Babun, A. K. Sikder, A. Acar, and A. S. Uluagac. Iotdots: A Digital Forensics Framework for Smart Environments. CoRR, abs/1809.00745, 2018.
- Barka, E., Mathew, S.S., Atif, Y., 2015. Securing the web of things with role-based access control. In: Codes, Cryptology, and Information Security. Springer, pp. 14–26.
- Bluetooth SIG Working Group, 2017. The Building Blocks of Your Bluetooth Device. Bluetooth Core Specification.
- Botta, A., De Donato, W., Persico, V., Pescap, A., 2016. Integration of cloud computing and internet of things: a survey. Future Gener. Comput. Syst. 56, 684–700.
- Bouij-Pasquier, I., El Kalam, A.A., Ouahman, A.A., De Montfort, M., 2015. A security framework for internet of things. In: Cryptology and Network Security. Springer, pp. 19–31.
- Bouij-Pasquier, I., Ouahman, A.A., El Kalam, A.A., de Montfort, M.O., 2015. SmartOrBAC security and privacy in the internet of things. In: Proceedings of International Conference on Computer Systems and Applications. IEEE, pp. 1–8.
- Celik, Z.B., Babun, L., Sikder, A.K., Aksu, H., Tan, G., McDaniel, P., Uluagac, A.S., 2018. Sensitive information tracking in commodity iot. In: Proceedings of USENIX Security Symposium. USENIX Association.
- Cirani, S., Picone, M., 2015. Effective authorization for the web of things. In: Proceedings of World Forum on Internet of Things. IEEE, pp. 316–320.
- Cirani, S., Picone, M., Gonizzi, P., Veltri, L., Ferrari, G., IoT-OAS, 2015. An OAuth-based authorization service architecture for secure services in IoT scenarios. IEEE Sens. J. 15 (2), 1224–1234.
- Da Xu, L., He, W., Li, S., 2014. Internet of Things in industries: a survey. IEEE Trans. Inf. Inf. 10 (4), 2233–2243.
- Dar, K., Taherkordi, A., Rouvov, R., Eliassen, F., 2011. Adaptable service composition for very-large-scale internet of things systems. In: Proceedings of Middleware Doctoral Symposium. ACM, pp. 11:1–11:2.
- Darianian, M., Michael, M.P., 2008. Smart home mobile RFID-based Internet-of-Things systems and services. In: Proceedings of International Conference on Advanced Computer Theory and Engineering. IEEE, pp. 116–120.
- Denniss, W., Bradley, J., 2017. OAuth 2.0 for Native Apps. RFC 8252. Internet Engineering Task Force (IETF).
- Dorri, A., Kanhere, S.S., Jurdak, R., 2016. Blockchain in Internet of Things: Challenges and Solutions. arXiv: 1608.05187, arXiv.org. .

<sup>11</sup> <https://gdpr-info.eu/>.

- Dorri, A., Steger, M., Kanhere, S.S., Jurdak, R., 2017. Blockchain: a dis. arXiv: 1704.00073. arXiv.org. Tributed Solution to Automotive Security and Privacy.
- Elliott, A., Knight, S., 2010. Role explosion: acknowledging the problem. In: *Software Engineering Research and Practice*, pp. 349–355.
- Fernandes, E., Jung, J., Prakash, A., 2016. Security analysis of emerging smart home applications. In: *Proceedings of Symposium on Security and Privacy*. IEEE, pp. 636–654.
- Fernndez, F., Alonso, ., Marco, L., Salvacha, J., 2017. A model to enable application-scoped access control as a service for IoT using OAuth 2.0. In: *Proceedings of Conference on Innovations in Clouds, Internet and Networks*. IEEE, pp. 322–324.
- Ferraiolo, D., Chandramouli, R., Kuhn, R., Hu, V., 2016. Extensible access control markup language (XACML) and next generation access control (NGAC). In: *Proceedings of International Workshop on Attribute Based Access Control*. ACM, pp. 13–24.
- Fox, G.C., Kamburugamuve, S., Hartman, R.D., 2012. Architecture and measured characteristics of a cloud based internet of things. In: *Proceedings of International Conference on Collaboration Technologies and Systems*. IEEE, pp. 6–12.
- Fremantle, P., Aziz, B., Kopeck, J., Scott, P., 2014. Federated identity and access management for the internet of things. In: *Proceedings of International Workshop on Secure Internet of Things*. IEEE, pp. 10–17.
- Garcia Lopez, P., Montresor, A., Epema, D., Datta, A., Higashino, T., Iamnitshi, A., Barcellos, M., Felber, P., Riviere, E., 2015. Edge-centric computing: vision and challenges. *SIGCOMM Comput. Commun. Rev.* 45 (5), 37–42.
- Garcia-Morchon, O., Wehrle, K., 2010. Modular context-aware access control for medical sensor networks. In: *Proceedings of Symposium on Access Control Models and Technologies*. ACM, pp. 129–138.
- Gerdes, L.S.S., Selander, G., Mani, M., Kumar, S., 2014. Use Cases for Authentication and Authorization in Constrained Environments. RFC 7744. Internet Engineering Task Force (IETF).
- Graham, G.S., Denning, P.J., 1972. Protection: principles and practice. In: *Proceedings of Spring Joint Computer Conference*. ACM, pp. 417–429.
- GSM/EDGE Radio Transmission and Reception. 3GPP TS 45.005, 3GPP, 2017.**
- Gubbi, J., Buyya, R., Marusic, S., Palaniswami, M., 2013. Internet of Things (IoT): a vision, architectural elements, and future directions. *Future Gener. Comput. Syst.* 29, 1645–1660.
- Guerrero-ibanez, J.A., Zeadally, S., Contreras-Castillo, J., 2015. Integration challenges of intelligent transportation systems with connected vehicle, cloud computing, and Internet of Things technologies. *IEEE Wireless Commun.* 22 (6), 122–128.
- Guoping, Z., Wentao, G., 2011. The research of access control based on UCON in the internet of things. *J. Softw.* 6 (4), 724–731.
- Gusmeroli, S., Piccione, S., Rotondi, D., 2013. A capability-based security approach to manage access control in the Internet of Things. *Math. Comput. Model.* 58 (5), 1189–1205.
- Hassanlieryagh, M., Page, A., Soyata, T., Sharma, G., Aktas, M., Mateos, G., Kantarci, B., Andrescu, S., 2015. Health monitoring and management using Internet-of-Things (IoT) sensing with cloud-based processing: opportunities and challenges. In: *Proceedings of International Conference on Services Computing*. IEEE, pp. 285–292.
- He, W., Golla, M., Padhi, R., Ofek, J., Drmuth, M., Fernandes, E., Ur, B., 2018. Rethinking access control and authentication for the home internet of things (IoT). In: *Proceedings of USENIX Security Symposium*. USENIX Association, pp. 255–272.
- Hernandez-Ramos, J.L., Jara, A.J., Marn, L., Skarmeta, A.F., 2013. Distributed capability-based access control for the Internet of Things. *J. Int. Serv. Inf. Sec.* 3 (3/4), 1–16.
- Hu, V.C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., Scarfone, K., 2014. Guide to Attribute Based Access Control (ABAC) Definition and Considerations. NIST Special Publication 800. NIST.
- Hussein, D., Bertin, E., Frey, V., 2017. A community-driven access control approach in distributed IoT environments. *IEEE Commun. Mag.* 55 (3), 146–153.
- IEEE Standard for Ethernet. IEEE Std 802.3-2015, IEEE Standard Association, 2015.**
- IEEE Standard for Low-Rate Wireless Networks. IEEE Std 802.15.4-2015, IEEE Standard Association, 2015.**
- Information technology, 2013. Radio Frequency Identification for Item Management Part 6: Parameters for Air Interface Communications at 860 MHz to 960 MHz General. ISO/IEC 18000-6:2013. International Organization for Standardization.
- Interagency Report On Status of International Cybersecurity Standardization for the Internet of Things (IoT).** <https://csrc.nist.gov/CSRC/media/Publications/nistir/8200/draft/documents/nistir8200-draft.pdf>. Accessed: 2018-02-21, 2018.
- Internet Protocol, 1981. RFC 791.
- Islam, S.M.R., Hossain, M., Hasan, R., Duong, T.Q., 2018. A conceptual framework for an IoT-based health assistant and its authorization model. In: *Proceedings of Annual Computing and Communication Workshop and Conference*. IEEE, pp. 616–621.
- Jia, Y.J., Chen, Q.A., Wang, S., Rahmati, A., Fernandes, E., Mao, Z.M., Prakash, A., 2017. ContextIoT: towards providing contextual integrity to appified IoT platforms. In: *Proceedings of Network and Distributed System Security Symposium*.
- Jindou, J., Xiaofeng, Q., Cheng, C., 2012. Access control method for web of things based on role and SNS. In: *Proceedings of International Conference on Computer and Information Technology*. IEEE, pp. 316–321.
- Kalam, A.A.E., Baida, R.E., Balbiani, P., Benferhat, S., Cuppens, F., Deswarte, Y., Mieke, A., Saurel, C., Trouessin, G., 2003. Organization based access control. In: *Proceedings of International Workshop on Policies for Distributed Systems and Networks*. IEEE, pp. 120–131.
- Kelly, S.D.T., Suryadevara, N.K., Mukhopadhyay, S.C., 2013. Towards the implementation of IoT for environmental condition monitoring in homes. *IEEE Sens. J.* 13 (10), 3846–3853.
- Khan, R., Khan, S.U., Zaheer, R., Khan, S., 2012. Future internet: the internet of things architecture, possible applications and key challenges. In: *Proceedings of International Conference on Frontiers of Information Technology*. IEEE, pp. 257–260.
- Kim, T.H.-J., Bauer, L., Newsome, J., Perrig, A., Walker, J., 2010. Challenges in access right assignment for secure home networks. In: *Proceedings of USENIX Conference on Hot Topics in Security*. USENIX Association, pp. 1–6.
- Kim, T.H.-J., Bauer, L., Newsome, J., Perrig, A., Walker, J., 2011. Access right assignment mechanisms for secure home networks. *J. Commun. Netw.* 13 (2), 175–186.
- Kim, J.E., Boulos, G., Yackovich, J., Barth, T., Beckel, C., Mosse, D., 2012. Seamless integration of heterogeneous devices and access control in smart homes. In: *Proceedings of International Conference on Intelligent Environments*. IEEE, pp. 206–213.
- LaPadula, L., Bell, D.E., LaPadula, L.J., 1973. *Secure Computer Systems: Mathematical Foundations*. Draft MTR. The MITRE Corporation, p. 2.
- Lee, S., Choi, J., Kim, J., Cho, B., Lee, S., Kim, H., Kim, J., 2017. FACT: functionality-centric access control system for IoT programming frameworks. In: *Proceedings of Symposium on Access Control Models and Technologies*. ACM, pp. 43–54.
- Mahalle, P.N., Anggorojati, B., Prasad, N.R., Prasad, R., 2013. Identity authentication and capability based access control (IACAC) for the internet of things. *J. Cyber Sec. Mobil.* 1 (4), 309–348.
- Mahalle, P.N., Thakre, P.A., Prasad, N.R., Prasad, R., 2013. A fuzzy approach to trust based access control in internet of things. In: *Proceedings of International Conference on Wireless Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems*. IEEE, pp. 1–5.
- Mahmoud, R., Yousef, T., Aloul, F., Zualkernan, I., 2015. Internet of things (IoT) security: current status, challenges and prospective measures. In: *Proceedings of International Conference for Internet Technology and Secured Transactions*. IEEE, pp. 336–341.
- Mazurek, M.L., Arsenault, J.P., Bresee, J., Gupta, N., Ion, I., Johns, C., Lee, D., Liang, Y., Olsen, J., Salmon, B., Shay, R., Vaniea, K., Bauer, L., Cranor, L.F., Ganger, G.R., Reiter, M.K., 2010. Access control for home data sharing: attitudes, needs and practices. In: *Proceedings of SIGCHI Conference on Human Factors in Computing Systems*. ACM, pp. 645–654.
- Mell, P., Grance, T., 2011. The NIST Definition of Cloud Computing. SP 800-145. NIST.
- Miessler, D., 2018. Preparing to Release the OWASP IoT Top 10. <https://goo.gl/kyeXkf>. Accessed:2018-11-15.
- Miorandi, D., Sicari, S., De Pellegrini, F., Chlamtac, I., 2012. Internet of Things: vision, applications and research challenges. *Ad Hoc Netw.* 10 (7), 1497–1516.
- Mohammed, J., Lung, C.-H., Ocneanu, A., Thakral, A., Jones, C., Adler, A., 2014. Internet of Things: remote patient monitoring using web services and cloud computing. In: *Proceedings of IEEE International Conference on Internet of Things, IEEE Green Computing and Communications, and IEEE Cyber, Physical and Social Computing*. IEEE, pp. 256–263.
- Montenegro, G., Kushalnagar, N., Hui, J., Culler, D., 2007. Transmission of IPv6 Packets over IEEE 802.15.4 Networks. RFC 4944. Network Working Group.
- Mosquitto: An Open Source MQTT v3.1/v3.1.1 Broker.** <https://mosquitto.org>, 2017.
- MQTT Version 3.1.1.** OASIS Standard, OASIS, 2014.
- Neisse, R., Steri, G., Baldini, G., 2014. Enforcement of security policy rules for the internet of things. In: *Proceedings of International Conference on Wireless and Mobile Computing, Networking and Communications*. IEEE, pp. 165–172.
- Nguyen, K.T., Laurent, M., Oualha, N., 2015. Survey on secure communication protocols for the Internet of Things. *Ad Hoc Netw.* 32, 17–31.
- Ouaddah, A., Elkalam, A.A., Ouahman, A.A., 2016. FairAccess: a new Blockchain-based access control framework for the Internet of Things. *Secur. Commun. Network.* 9 (18), 5943–5964.
- Ouaddah, A., Elkalam, A.A., Ouahman, A.A., 2017. Towards a novel privacy-preserving access control model based on blockchain technology in IoT. In *Europe and MENA Cooperation Advances in Information and Communication Technologies*. Adv. Intell. Syst. Comput. 520, 523–533 Springer.
- Ouaddah, A., Mousannif, H., Elkalam, A.A., Ouahman, A.A., 2017. Access control in the Internet of Things: big challenges and new opportunities. *Comput. Netw.* 112, 237–262.
- OWASP Top 10 IoT Vulnerabilities.** [https://www.owasp.org/index.php/Top\\_IoT\\_Vulnerabilities](https://www.owasp.org/index.php/Top_IoT_Vulnerabilities). Accessed:2018-11-15.
- Park, J., Sandhu, R., 2004. The UCON<sub>ABC</sub> usage control model. *ACM Trans. Inf. Syst. Secur.* 7 (1), 128–174.
- Patel, P., Bansal, D., Yuan, L., Murthy, A., Greenberg, A., Maltz, D.A., Kern, R., Kumar, H., Zikos, M., Wu, H., et al., 2013. Ananta: cloud scale load balancing. *SIGCOMM Comput. Commun. Rev.* 43 (4), 207–218.
- Pinno, O.J.A., Gregio, A.R.A., Bona, L.C.E.D., 2017. ControlChain: blockchain as a central enabler for access control authorizations in the IoT. In: *Proceedings of Global Communications Conference*. IEEE, pp. 1–6.
- Postel, J., 1980. User Datagram Protocol. RFC 768. .
- Ray, I., Alangot, B., Nair, S., Achuthan, K., 2017. Using attribute-based access control for remote healthcare monitoring. In: *Proceedings of International Conference on Software Defined Systems*. IEEE, pp. 137–142.
- Razzaque, M.A., Mилоjevic-Jevric, M., Palade, A., Clarke, S., 2016. Middleware for internet of things: a survey. *IEEE Int. Things J.* 3 (1), 70–95.
- REST, 2011. Semantic Web Standard. W3C Semantic Web.
- Rivera, D., Cruz-Piris, L., Lopez-Civera, G., de la Hoz, E., Marsa-Maestre, I., 2015. Applying a unified access control for iot-based intelligent agent systems. In: *Service-Oriented Computing and Applications (SOCA), 2015 IEEE 8th International Conference on*. IEEE, pp. 247–251.

- Roman, R., Zhou, J., Lopez, J., 2013. On the features and challenges of security and privacy in distributed internet of things. *Comput. Netw.* 57 (10), 2266–2279.
- Sadeghi, A.-R., Wachsmann, C., Waidner, M., 2015. Security and privacy challenges in industrial Internet of Things. In: *Proceedings of Design Automation Conference*. IEEE, pp. 1–6.
- Saint-Andre, P., 2014. Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence. RFC 3921. Network Working Group.
- Salonikias, S., Mavridis, I., Gritzalis, D., 2015. Access control issues in utilizing fog computing for transport infrastructure. In: *Critical Information Infrastructures Security*, LNCS 9578. Springer, pp. 15–26.
- Samarati, P., Capitani de Vimercati, S., 2000. Access control: policies, models, and mechanisms. In: *Foundations of Security Analysis and Design*, LNCS 2171. Springer, pp. 137–196.
- Sandhu, R.S., Coyne, E.J., Feinstein, H.L., Youman, C.E., 1996. Role-based access control models. *Computer* 29 (2), 38–47.
- Schuster, R., Shmatikov, V., Tromer, E., 2018. Situational access control in the internet of things. In: *Proceedings of Conference on Computer and Communications Security*. ACM, pp. 1056–1073.
- Sciancalepore, S., Piro, G., Tedeschi, P., Boggia, G., Bianchi, G., 2018. Multi-domain access rights composition in federated IoT platforms. In: *Proceedings of Workshop on Recent Advances in Secure Management of Data and Resources in the IoT*.
- Seitz, L., Selander, G., Gehrmann, C., 2013. Authorization framework for the internet-of-things. In: *Proceedings of International Symposium on A World of Wireless, Mobile and Multimedia Networks*. IEEE, pp. 1–6.
- Sethi, P., Sarangi, S.R., 2017. Internet of things: architectures, protocols, and applications. *J. Electr. Comput. Eng.* 2017.
- Shelby, Z., Hartke, K., Bormann, C., 2014. The Constrained Application Protocol (CoAP). RFC 7252. Internet Engineering Task Force (IETF).
- Sheng, Z., Yang, S., Yu, Y., Vasilakos, A., Mccann, J., Leung, K., 2013. A survey on the IETF protocol suite for the Internet of Things: standards, challenges, and opportunities. *IEEE Wireless Commun.* 20 (6), 91–98.
- Shrouf, F., Ordieres, J., Miragliotta, G., 2014. Smart factories in Industry 4.0: a review of the concept and of energy management approached in production based on the Internet of Things paradigm. In: *Proceedings of International Conference on Industrial Engineering and Engineering Management*. IEEE, pp. 697–701.
- Sicari, S., Rizzardi, A., Grieco, L.A., Coen-Porisini, A., 2015. Security, privacy and trust in Internet of Things: the road ahead. *Comput. Netw.* 76, 146–164.
- Socolofsky, T., Kale, C., 1991. A TCP/IP Tutorial. RFC 1180. Network Working Group.
- Song, Z., Crdenas, A.A., Masuoka, R., 2010. Semantic middleware for the internet of things. In: *Proceedings of International Conference on the Internet of Things*. IEEE, pp. 1–8.
- Standard For an Architectural Framework for the Internet of Things (IoT). <https://standards.ieee.org/develop/project/2413.html/>. Accessed: 2018-02-20, 2016.
- Stanislav, M., Beardsley, T., 2015. HACKING IoT: A Case Study on Baby Monitor Exposures and Vulnerabilities. Technical report, Rapid7.
- Stojmenovic, I., Wen, S., 2014. The fog computing paradigm: scenarios and security issues. In: *Proceedings of Federated Conference on Computer Science and Information Systems*. IEEE, pp. 1–8.
- Tian, Y., Zhang, N., Lin, Y.-H., Wang, X., Ur, B., Guo, X., Tague, P., 2017. SmartAuth: user-centered authorization for the internet of things. In: *Proceedings of USENIX Security Symposium*, pp. 361–378.
- Tnjes, R., Barnaghi, P., Ali, M., Mileo, A., Hauswirth, M., Ganz, F., Ganea, S., Kjrgaard, B., Kuemper, D., Nechifor, S., Puiu, D., Sheth, A., Tsiatsis, V., Vestergaard, L., 2014. Real time IoT stream processing and large-scale data analytics for smart city applications. In: *Poster Presented at European Conference on Networks and Communications*.
- Transmission Control Protocol, 1981. RFC 793.
- Trivellato, D., Zannone, N., Glaundrup, M., Skowronek, J., Etalle, S., 2013. A semantic security framework for systems of systems. *Int. J. Coop. Inf. Syst.* 22 (1).
- User-Managed Access (UMA) 2.0 Grant for OAuth 2.0 Authorization. <https://kantarinitiative.org/file-downloads/rec-oauth-uma-grant-2-0-pdf/>. Accessed: 2018-02-26, 2017.
- van der Meulen, R., 2015. Gartner Says 6.4 Billion Connected Things Will Be in Use in 2016, up 30 Percent from 2015. Press release, Gartner.
- Vasilomanolakis, E., Daubert, J., Luthra, M., Gazis, V., Wiesmaier, A., Kikiras, P., 2015. On the security and privacy of internet of things architectures and systems. In: *Proceedings of International Workshop on Secure Internet of Things*. IEEE, pp. 49–57.
- Vermesan, O., Friess, P., Guillemin, P., Gusmeroli, S., Sundmaeker, H., Bassi, A., Jubert, I.S., Mazura, M., Harrison, M., Eisenhauer, M., Doody, P., 2011. Internet of things strategic research roadmap. In: *Internet of Things Global Technological and Societal Trends*. River Publishers, pp. 9–51.
- Weber, R.H., 2010. Internet of thingsnew security and privacy challenges. *Comput. Law Secur. Rep.* 26 (1), 23–30.
- Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, J., Alexander, R., 2012. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. RFC 6550. Internet Engineering Task Force (IETF).
- Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. IEEE Std 802.11ac-2013, IEEE Standards Association, 2013.
- Wu, M., Lu, T.-J., Ling, F.-Y., Sun, J., Du, H.-Y., 2010. Research on the architecture of internet of things. In: *Proceedings of International Conference on Advanced Computer Theory and Engineering*, vol. 5. IEEE, pp. 484–487.
- eXtensible Access Control Markup Language (XACML) Version 3.0. OASIS Standard, OASIS, 2013.
- Yang, L., Yang, S.-H., Plotnick, L., 2013. How the internet of things technology enhances emergency response operations. *Technol. Forecast. Soc. Change* 80 (9), 1854–1867.
- Ye, N., Zhu, Y., Wang, R.-c., Malekian, R., Qiao-min, L., 2014. An efficient authentication and access control scheme for perception layer of internet of things. *Appl. Math.* 8 (4), 1617–1624.
- Yuan, E., Tong, J., 2005. Attributed based access control (ABAC) for web services. In: *Proceedings of International Conference on Web Services*. IEEE, pp. 561–569.
- Zanella, A., Bui, N., Castellani, A., Vangelista, L., Zorzi, M., 2014. Internet of things for smart cities. *IEEE Int. Things J.* 1 (1), 22–32.
- Zhang, G., Tian, J., 2010. An extended role based access control model for the Internet of Things. In: *Proceedings of International Conference on Information, Networking and Automation*, vol. 1. IEEE, pp. 319–323.
- Zhang, Yungpeng, Wu, Xuqing, 2016. Access Control in Internet of Things: A Survey. *CoRR abs/1610.01065*.
- ZigBee Specification, 2014. Standard. ZigBee Alliance.

**Sowmya Ravidas** is a doctoral candidate in the Department of Mathematics and Computer Science at Eindhoven University of Technology. Her research interests are in the areas of access control, cloud computing and IoT. Previously she obtained a master's degree in Mobile Computing from Aalto University in Finland and has briefly worked at Nokia Bell Labs.

**Alexios Lekidis** is a Research Associate (post-doc) at the Security and Embedded Networked Systems Group of the Eindhoven University of Technology working in security mechanisms for automotive and IoT systems. He holds a PhD in Applied Mathematics / Computer Science, which he obtained from the University of Grenoble in the topic of model-based design in networked embedded systems. Dr. Lekidis research interests lie on system design for the IoT with a special focus on performance evaluation, network monitoring, security and model verification techniques.

**Federica Paci** is a Lecturer in Cyber Security in the Electronics and Computer Science Department at the University of Southampton. She received a PhD in Information Technology from the University of Milan in February 2008. She was postdoctoral fellow at the Computer Science Department of Purdue University and at the Department of Information Engineering and Computer Science of the University of Trento. She published more than 60 contributions as papers in international conferences and journals on data security and privacy.

**Nicola Zannone** received his Ph.D. degree in Computer Science at the University of Trento, Italy, in 2007. He is an associate professor in the Security group at the Eindhoven University of Technology, the Netherlands. His research interests include computer security, data protection, access control and formal methods.