J. Lopez, and J. E. Rubio, "Access control for cyber-physical systems interconnected to the cloud", *Computer Networks*, vol. 134, pp. 46 - 54, 2018.
http://doi.org/https://doi.org/10.1016/j.comnet.2018.01.037
NICS Lab. Publications: https://www.nics.uma.es/publications

Access control for Cyber-Physical Systems interconnected to the Cloud

Javier Lopez, Juan E. Rubio
Department of Computer Science, University of Malaga,
Campus de Teatinos s/n, 29071, Malaga, Spain
{jlm,rubio}@lcc.uma.es

Abstract

The continuous advance in manufacturing and information analytics has improved the connectivity between computational and physical elements within the industry, hence increasing the effectiveness and reliability of Cyber-Physical Systems (CPS). This progress has been further enhanced by Cloud computing technologies, by externalizing services and interconnecting different industrial networks. As a consequence, there has been an increase of cyber-security threats in the industrial sector in recent years. Among other security measures, it is of paramount importance to introduce flexible access control mechanisms to avoid unauthorized access to the heterogeneous systems that coexist in this context. In this paper, we identify the requirements for such techniques, and propose a novel industrial architecture where multiple access control models are assessed when cloud technologies are integrated. In particular, we emphasize their adaptability to new heterogeneous scenarios through diverse indicators, achieving a trade-off between security and efficiency. **Keywords:** Cyber, Physical, Systems, Cloud, Security, Access, Control

1 Introduction

A Cyber-Physical System (CPS) refers to a mechanical mechanism that is controlled by computational entities which work collaboratively: namely, sensors and actuators that capture data from the procedure and regulate its parameters according to a set of defined rules, hence achieving an interaction between the physical and computational components [1]. These systems have been deeply integrated in critical infrastructures (i.e., energy sector, transport) and generally in all industrial control systems for years. The initial goal of achieving intelligent, resilient and self-adaptable machines in this context has been eased in recent years by the increasing affordability of sensors and the rapid development of new communication networks and protocols. This has resulted in the continuous generation of high volumes of data and the integration with information technologies (IT). The most evident case is Cloud Computing. It is of

key importance to understand the evolution of the industry towards a model where the product is flexibly manufactured by a network of suppliers accessible via the cloud along the whole production chain, with an extensive integration between customers and business partners.

The counterpart of the modernization of industrial technologies (which we will refer to as "operational technologies", OT) and the interconnection of CPSs with external networks like the Internet brings with it the appearance of new cyber-security threats. Some of them are inherited from the IT paradigm and others arise from the growing integration between IT and OT assets. As a result, there has been an increase of vulnerabilities in the industrial sector in recent years, as some reports show [2][3]. We are talking about attack vectors such as denial of service, presence of malware, exploitation of vulnerabilities in communication protocols to intercept traffic, phishing and social engineering, etc. In terms of authorization and access control, which are the main focus of this work, they imply the misuse of resources and the misappropriation of the identity of nodes, that can even influence the overall behavior of the system. Altogether, these issues make security the main concern for the adoption of these technologies in such a critical scenario.

In this complex environment, where any element could potentially interact and cooperate with any other element, access control is essential to manage permissions of users, peripheral devices or programs when they request to use certain resources within the infrastructure. The integration of IT technologies and especially the cloud hinders the application of conventional access control models in industrial systems, for several reasons. These can be summarized in the sharing of information among heterogeneous entities with different degrees of sensitivity, performance and regulations. Therefore, it becomes mandatory to analyze the full range of requirements that access control presents in the upcoming scenario, in order to accurately tailor the available models and propose new approaches that meet these conditions. In particular, it is useful to consider how these security techniques can affect the physical world by introducing an extra overhead in the control and monitoring procedures.

In this paper, we identify the set of requirements that access control mechanisms must have in the industry as a consequence of the CPS and cloud interconnection, assessing the adaption of particular models. The paper is organized as follows: section 2 exposes the requirements that access control solutions must match, taking the new architecture into consideration. Section 3 presents traditional approaches, whereas Section 4 describes new mechanisms in the literature. They are ultimately analyzed according to the aforementioned requirements in Section 5. Finally, conclusions and future work are presented in Section 6.

2 Access control requirements

In order to identify the requirements of access control in the CPS infrastructure, it is mandatory to firstly review how industrial networks are affected by the integration of IT technologies. A traditional control network follows the

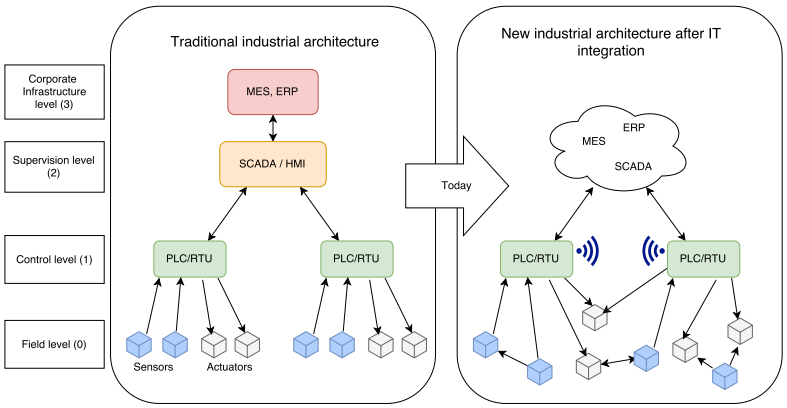

Figure 1: Evolution of the traditional industry architecture

architecture described in the ISA-95 standard [4]. In this way, the productive process itself constitutes the base of the pyramid (level 0), whereas devices that interact with it (i.e., PLCs, RTUs) are set in level 1. Level 2 represents those devices that control the production process (i.e., SCADAs, HMIs), and those that control the workflow (i.e., MES systems) are located in level 3. Lastly, the highest level contains the infrastructure of logistics, inventory, ERP or planning.

The implementation of cyber-physical systems within this context means the introduction of advanced connectivity technologies and computational capabilities to ensure a real-time data acquisition from the physical world and an intelligent data management. The goal is to gather information from every connected machine and run specific analytics to extract additional insights, providing feedback from cyber space back to the physical space. In practice, this evolution is fostered by the implementation of several communication protocols due to the standardization of software and hardware: ranging from field bus protocols (i.e., HART; wirelessHART, etherCAP, IO-Link) to protocols working with Ethernet and TCP/IP, such as Ethernet/IP, Ethernet POWERLINK, CANopen, PROFINET, Modbus/TCP or HART/IP. The case of standards devised for the interoperable management of all types of industrial equipment, like CIP, OPC UA and MTConnect are especially interesting. On the whole, this results in the evolution of the traditional architecture towards a distributed and decentralized model, as Figure 1 shows.

According to the new architecture model, devices located in the lower levels of the architecture interoperate with each other to interconnect all the components of the infrastructure, ranging from machines to operators or the product itself, in order to gather data. On the other hand, the cloud is leveraged to provide supervision as a service and interconnect different substations easily. By this means, a collaborative environment can be created by diverse companies whose applications and constraints may differ, making it difficult to reach a global agreement or the adoption of any common specification.

In this complex scenario, access control mechanisms deployed (either in field devices, PLCs or cloud resources) aim to restrict what each entity should be able to access and the connections that can be accepted, having the ability to deal with a diversity of devices [5]. Actual solutions are still in their infancy, due to the need for a dynamic and fine-grained mechanism that deals with several users and constrained resources. We can thereby define the following set of specific requirements, based on an extensive review of the literature with the aim to study which features the models need for this particular context:

- **Dynamicity:** services in modern CPSs are accessed remotely by a large number of technologies and protocols, which are also added or removed on demand. Due to cloud computing, several applications could be integrated in the product life cycle, ranging from monitoring procedures (e.g., inventory, real-time performance) to dynamic manufacturing processes defined on the go, which could change their parameters dynamically. Virtualization elements of cloud computing also offer scalability in terms of resource allocation, which in turns introduces a challenge for access control systems with the control of multiple user accounts.

- **Scalability:** access control must accept the definition of new users and complex policies, while not introducing operational costs. It should be extensible with respect to the number of users and resources controlled, including the adaptation to new technologies (e.g., communication protocols, operating systems) through well defined interfaces. It is important that the access control system has a situational awareness of all factors involved in the authorization decisions at all times: this involves parameters such as the number of connected devices, and their available resources.

- **Flexibility:** proposed mechanisms must provide an easy administration to define which attributes should be used for authorization, which credentials could be transferred or with the definition of trust relationships between entities (e.g., the cloud provider and its users). The access control system can be permanently updated with information about the multiple workflows within the organization, by making use of specification languages that support complex logic rules.

- **Quality of service:** concerns the computational complexity of access control procedures, that may increase the response time for authorization decisions [6], especially in resource constrained devices (i.e., both for computation and autonomy). In this respect, wireless communications must be taken into account, since they can limit bandwidth and experience transmission delays. Therefore, the access control service must manage the connection requests between networks with different demands of service quality, by checking if there are free resources to accept such connections. As a result, that admittance control balances the whole system overload. However, this may be difficult to achieve in practice, as we are dealing with a decentralized architecture. One solution could be the implementation of a hierarchical strategy that implements local control mechanisms

over the entire infrastructure, located in the bridges between networks. Another alternative consists in the deployment of a parallel architecture that retrieves all data coming from the plant.

Not all access control models have these characteristics in the upcoming industrial scenario, since each one is specifically designed with different security requirements. Regardless of the application domain, an effective access control system should provide the security properties of *confidentiality* (avoiding unauthorized disclosure of data), *integrity* (keeping information secure from modification without authorization) and *availability* (assuring the access to information under reliable conditions). Each domain has specific features and particular security requirements to be taken into account when adopting an access control system. Specifically, availability is critical in real-time systems that are involved in automation tasks, such as the SCADA servers. In the next section, we provide the background to traditional access control models that aim to match the scenario we are analyzing.

3 Applicability of traditional access control models for the CPS-Cloud interconnection

In the context of industrial cyber-physical systems, an access control system is in charge of managing the access permissions of all users and entities (e.g., databases, firewalls, or routers) when they require the use of a specific resource within the organization. Firstly, the user or entity authenticates (i.e., using password/login or identity certificates) and then the system approves or refuses the access in order to keep the resources protected against the unauthorized misuse. This decision is computed based on different criteria. For instance, the permission granted to the entity requiring the access, the security of the connection, or the availability of global resources. Regardless of the outcome, information about each request is recorded afterwards for further analysis.

Generally, traditional access control models can be classified into three main categories, that are reviewed here: The discretionary access control (DAC), the mandatory access control (MAC), and the role-based access control model (RBAC). Even though these mechanisms are different, they are not mutually exclusive and can coexist within the plant. What they do have in common is the definition of three main elements:

- Object: the resource whose access is controlled, ranging from files to specific programs, devices, or communication ports, for instance.
- Subject: the entity that potentially requests the access to the object (e.g., operators, groups of workers, or network equipment). It is normally represented by a process, since every user or application has access to a resource through a process executed in the subject environment that represents it.
- Access right: describes the way in which the subject can finally access and use the object (e.g., read, write, execute, delete, create).

Regardless of the way of defining and managing these three elements by the proposed mechanisms, all of them should follow the principle of minimum privilege to achieve the security of the entire architecture. Users should only possess the minimum access permission that is enough to meet their responsibilities.

Discretionary Access Control (DAC)

The Discretionary Access Control (DAC) model [7] restricts the access to the regulated objects based on the identity of the subjects and/or the groups they belong to. It is discretionary because there are no rigid rules to assign permissions but in turn it does allow some legitimate users to grant access to other subjects, once the system has verified his/her identity and checked it against access control list permissions. There are two ways to implement the DAC model [8]:

- Access Control Lists: each list corresponds to an object and shows the set of subjects assigned to it and their access rights.
- Access Control Capability Lists: each list corresponds to a subject and shows the set of objects that he/she has access to and his/her rights.

DAC constitutes a model that features flexibility when it is used in local environments, since it provides flexibility for the users to assign access permissions to their own resources; this is the reason why it is extensively used in Windows and UNIX-based platforms to provide access control to the file system. However, it has several issues when it is used in the CPS-cloud scheme depicted in Figure 1. In this sense, we can stress two main concerns: firstly, the low control over the information flow in a dynamic network, since a user can, for e.g., deliberately read a file from one company and copy its content in another company. Equally, that user can grant permissions or transfer his/her rights to another malicious user, through Trojan horses [9]. Secondly, it is not scalable or efficient enough for the industry: the potential interconnection with devices of all kinds makes it burdensome to hold the complete data structure that this model proposes.

Mandatory Access Control (MAC)

The Mandatory Access Control (MAC) model [10] defines a central authority in charge of deciding whether a subject can access a given object or its information. It is inspired by military and commercial security policies, and assigns a security class to the authorized objects. Namely, these security labels are classified according to the resource criticality or the sensitivity of the information involved (e.g., *top secret* $\rightarrow$ *confidential* $\rightarrow$ *unclassified*). Accordingly, subjects of the organization are given security clearances that indicate the permission or level of trustworthiness associated with that user. This way, the clearance of a subject is compared with the object class in order to check if he/she has sufficient privileges to access that resource. Contrary to the DAC model, MAC centralizes

the control to define policies that are enforced for all users, not letting them override or modify permissions, either accidentally or intentionally.

Contrary to DAC, MAC guarantees that the information flow is protected at all times by giving the control to a central authority. However, this turns out to be the main drawback because a unique party has to be in charge of determining what information is accessible and by whom. This model only concerns controlling the data flow and achieving confidentiality. For this reason, *Biba* [11] proposed an enhanced model that also defines integrity levels to prevent against unauthorized changes. For our purpose, it also becomes mandatory to deploy a distributed model to put into practice a more accurate control over the resources beyond the clearances declared in the MAC approach, which do not offer flexibility. With the inclusion of the cloud computing infrastructure and the integration of complex databases (commonly accessed through web interfaces), a fine-grained control is needed. For example, in a network of suppliers accessed via the cloud, a manager might have permission to access a client's file in order to retrieve his/her address, but not the bank account details, which would be restricted to the financial staff. This could be solved with a dynamic activation of access rights for certain tasks, which is not supported by MAC. In addition, the separation of duties, least privilege, delegation or inheritance of rights are not supported either.

Role Based Access Control (RBAC)

The role-based access control (RBAC) model [12] was originally designed to solve the issues of the previous proposals, under the principle that "a subject's responsibility is more important than whom the subject is"[13]. For this reason, this model bases the access control on the roles assigned to users instead of their identities. This simplifies the management of rights over DAC and MAC, while still supporting three key security rules: apart from the minimum privilege principle mentioned before, the separation of duties and the data abstraction principle.

In the RBAC model, a role represents a specific profile within the organization that has a set of responsibilities and actions associated with it. Subsequently, roles are assigned to users (subjects), thereby acquiring the corresponding role's access permissions. Conversely, users can have different roles, which are granted with the authority needed to perform his/her tasks in the infrastructure. This eases the authorization management, since it is no longer necessary to assign individual authorizations to subjects; they are simply assigned with a role that already reflects this authority, which incidentally prevents that a user is granted more permissions than needed. In addition, this makes it easier to change the functions of a user inside the organization, since he/she only needs to revoke or replace the corresponding role. This principle is fundamental in the industrial context analyzed in the present work, where the set of roles involved in the production life-cycle may change dynamically (as a consequence of the multiple interactions achieved by cloud), as well as the access rights associated with a particular role.

At a glance, this model represents a natural way to provide access control within organizations with heterogeneous resources and security requirements. The RBAC core is conceived to be extended with three components: the *role hierarchy*, that permits the inheritance of rights; *the static separation of duties* (that allows the definition of role constraints) and the *dynamic separation of duties* (that supports time-dependent role constraints). Altogether, they allow RBAC to support important principles such as the least privilege, separation of administrative functions and separation of duties. However, despite its advantages over DAC and MAC models, it has some problems when it is implemented in the complex industrial scenario considered [14]. In particular, we highlight the following issues:

1. It does not make a differentiation of sensitivity in the data (e.g., location of workers in the plant); every access is only regulated through the roles possessed, without taking other factors (e.g., risk, current workload or time constraints) into consideration, as it does not separate tasks from roles.

2. The model does not contemplate the delegation principle present in DAC, which could be leveraged in critical environments when a given system is down or the respective staff member is absent. Also related to relationships between roles, it does not support the control of operations in sequence: for example, a manager in charge of purchasing new material needs to examine the inventory file and then contact the corresponding suppliers. Each action needs different permissions, which can not all be controlled by the model. In general, RBAC should be enforced with policies to cope with anomalous behaviors of users.

3. The scalability can be affected by the thousands of dynamic users and permissions involved in a conglomerate of industrial networks, which hardens the task of accurately and securely creating and assigning roles, especially in the first phases of integration. This requires a complex role engineering process for the security administrators [15], where roles are identified and assigned permissions in two ways [16]: either via *top-down* approach, which takes a job function and associates a right to it; or with the *bottom-up* approach, which takes needed permissions and aggregates them into roles. The latter allows the automation of the process, which is also known as *role mining* in the literature [17].

4. Concerning the quality of service, the model has to ensure an access decision in a timely manner, irrespective of the number of devices or the environmental conditions. This becomes difficult in distributed networks of critical infrastructures composed by resource-constrained systems, especially when dealing with diverse users and roles. It is also essential to carry out tests prior to deploying these decision systems, in order to report conflicts between roles and privileges.

4 New access control models for the CPS-Cloud interconnection

From Section 2 we can deduce that the success of any access control solution in the modern industry depends on accurately identifying a complete set of requirements that may vary with the type of organization and business network. In the previous section, we have determined why traditional models cannot be fully applied to this scenario. This is because they are not able to address the requirements of complexity and heterogeneity. There have also been multiple efforts to adapt RBAC to these environments as well as designing new techniques that extend its principle, where there are diverse sets of policy decision points and users who may not be known throughout the network. In the following, we analyze the most representative models that have been explored and taken as a basis in the literature for the development of custom access control approaches. We begin by reviewing those that extend RBAC through the concept of attribute, to later present those that address access control from the organizational perspective. Then, we describe models that consider other parameters such as the privacy or the risk implied by the operation whose access is requested.

Attribute-Based Access Control (ABAC) model

The Attribute-Based Access Control (ABAC) model [18] addresses the shortcomings of the traditional models. It proposes an approach where access is granted according to a set of attributes associated with the subject and the object, which are presented and compared once the requester has been authenticated in the system. These attributes can refer to roles possessed, location of the requester, the computational resources available, etc. or a combination of them. The access decision is finally made by a policy decision point, which checks whether all security conditions are met: namely, the subject-object attributes, together with environment attributes and the type of operation performed. This kind of model is also known as a Policy- Based Access Control (PBAC).

As a result, this enhances the granularity of access controls in comparison with traditional models, as it is able to cope with different risk levels and context-related characteristics. However, the main drawback arises with the accurate selection of the attributes used for access decisions, which could be challenging in a cloud computing scenario.[19]

Access Control based on Usage Control (UCON)

As ABAC, the usage control (UCON) access control model [20] extends the functionality of conventional approaches by proposing a more fine-grained control of the users permissions. It bases the access decision on the same concept of attributes defined in ABAC, but it additionally introduces the mutability of their values over time, which ultimately influences the security policies while the

access request is still in progress. This way, it performs a continuous authorization (before, during and after the access itself). One example of this on-going authorization is a procedure (e.g., acquisition of materials) that involves a set of steps with fluctuating factors and multiple users (e.g., weather conditions, demand peaks).

On the one hand, UCON is characterized for being suitable for open dynamic scenarios like the one addressed in this paper. Here, a distributed authority is usually deployed instead of a central authority, with an administrator or the resource owner acting as the authority root to check the attributes asserted by users. On the other hand, we must stress the complexity of this model, which needs a continuous database support to maintain the status of the attribute values.

Capability-Based Access Control (CapBAC) model

Access control models like RBAC or ABAC base their decision making on the definition of roles and attributes, which are checked in complex access policies. They have been successfully applied in multiple security scenarios, but they need further enhancement to adapt to fully distributed environments where thousands of heterogeneous end-devices may coexist. In recent years, this problem has often been overcome with access control models based on capabilities (CapBAC), which have been widely used in the Internet of Things (IoT) field [21]. The capability concept refers to a token that contains rights granted to the subject that holds it, and must be tamper-proof and unequivocally identified. This token is issued by the provider of the resource whose access is required, and is commonly presented as an authorization certificate. Consequently, the user needs to show this certificate to the provider prior to request an operation. This way, the entity that receives the request already knows the permissions granted to the requestor, in contrast with the MAC models, where the security levels and the corresponding clearances are decoupled. On the whole, this makes the authorization simpler for scenarios with resource-constrained devices, since complex policies are no longer required.

CapBAC can therefore be considered an adequate model since it satisfies the principle of least privilege and ensures the validity, revocation and updatability of the authorization certificates. The main disadvantage is the need to create and maintain all certificates, together with the need to enforce the delegation of privileges. The authorization certificates should also be improved with a standard structure, in order to be compatible with cross-domain or cross-enterprise applications.

Risk-Based Access Control model

The Risk-Based Access Control [22] was designed for highly dynamic environments involving multiple organizations, where there are diverse security policies being applied and it is not possible to predict the number of existing users and resources whose access is required. It is based on real-time decision making

that firstly assesses the risk that the required access involves, creating a scale of different levels that depend on various environmental conditions following the principle of operational need. This risk is computed as $risk = VxP$, where V is the information value (that reflects the resource sensitivity) and P represents the probability of unauthorized disclosure (and reflects the user's trustworthiness). The information value can be expressed as costs from loss of availability (if the access turns into a Denial of Service attack), loss of confidentiality (if the access causes information to be disclosed), and loss of integrity (if the resource accessed is illegitimately modified). On the other hand, the probability P can be estimated by considering various threatening scenarios. Nevertheless, such values can be dynamically changed based on the security policy and the set of security levels.

Even though this model constitutes an acceptable solution to address the network heterogeneity with dynamic behavior, it is not very applicable in practice due to the burden of analyzing the criticality and probability of abuse of every system involved in the organization, which is necessary to compute the risk tolerance levels. This task commonly requires expertise and is subjective, according to the organization's security objectives.

Organizational-Based Access Control (OrBAC) model

The Organizational-Based Access Control (OrBAC) is a model that also addresses the RBAC flaws, concerning the lack of control over predefined tasks comprising various operations that imply information about the context to make the access decision. For this goal, it proposes the organizational dimension on top of the original model, defining a new level of abstract entities separated from the concrete level, that comprises the subject, the object and the action whose permission is requested. Namely, it is based on roles, activities and views:

- Subjects are assigned with roles, that are given with a set of permissions.
- An activity is a set of actions with the same security policy (i.e., associated permissions).
- Equivalently, views encompass objects with the same security.

This way, OrBAC extends RBAC from a high-level perspective to flexibly group roles that are given with permissions to realize activities on determined views. Concrete privileges are derived from theses abstract privileges. Altogether, this provides a common framework to describe security policies between many organizations at the same time. In addition, this model is not only restricted to permissions, but also contemplates the possibility of establishing prohibitions and obligations.

In summary, OrBAC can be considered a context-sensitive model that takes dynamism into account, which also supports the concepts of hierarchy and role constraints. However, despite these advantages, it is only conceived for centralized structures and does not focus on the distributed nature and interoperability

of modern industry control networks that interconnect different infrastructure resources through Internet technologies. Security policies deployed in these environments must be implemented taking into consideration all the different entities of each organization and hence the mutual trust, in order to ensure a secure collaboration between them.

Task-Role-Based Access Control (T-RBAC)

The Task-Role-Based Access Control model [23] is another scheme that improves RBAC that, like the OrBAC, tries to reflect the organization workflows in the authorization mechanism. It was originally proposed for cloud computing applied to health care systems [24] to synchronize the workflow with the authorization flow. It is based on the activation of permissions depending on the process taking place in the system, which is known as a task. Tasks are assigned to users by their job positions or roles, and they finally perform read and write operations over objects when executing them. The decision of which access rights should be assigned to the user depends on the tasks he/she requires to execute, so they become the basis of the access control. Consequently, a workflow is a business process composed by some tasks that are connected to achieve a common goal. All tasks are therefore classified into two classes, according to whether they belong to a certain workflow or not. In the case the execution of a task depends on the output of another previous task, the user can activate his/her access rights in accordance with the workflow, which is called *active* access control. To the contrary, if we are dealing with static tasks that are always executed under the same conditions or do not belong to a particular workflow, we talk about *passive* access control. Therefore tasks are leveraged to support active access control and roles support passive access control.

There are many advantages associated with T-RBAC, which is successfully implemented in various domains like the enterprise or the military environment [25]. One example is the case of Amazon Elastic Compute Cloud (Amazon EC2). The fact that it uses the task as the unit to separate duties (instead of the roles) means it supports more elaborate policies than RBAC, since in real world access rights are given to roles according to their tasks. However, it also has some cons, such as lacking the ability to cope with heterogeneous technologies and data. This model does not provide any sensitive levels for information, and it is intended to be deployed in a centralized manner, which also hinders the scalability of resources in the upcoming industrial scenario.

Privacy-Aware Role-Based Access Control (P-RBAC)

So far, we have reviewed the literature, searching for approaches that basically extend the principles of RBAC to consider other parameters involved in the production cycle. These include workflows, their risks or the environmental parameters associated with the users and objects whose access is required. However, they are not designed taking privacy into consideration and consequently they hardly meet any privacy protection requirement [26]. This concern is of

particular importance with the integration of cloud technologies within the industry: on the one hand, the trust in external cloud providers may lead to an exposure of sensitive data about the organization and its technical procedures. On the other hand, the growing interconnection of devices to exchange data between machines and operators might put privacy of workers at risk, since sensitive data may be exposed, like their location or their throughput. The Privacy-Aware Role-Based (P-RBAC) Access Control [27] was developed to extend the classical RBAC model to support privacy policies. This protection of privacy is achieved by analyzing the purpose of the access request and the fulfillment of obligations within time intervals. These obligations may differ from users performing the access, and make reference to privacy laws or conditions to be guaranteed, which are represented with high-level context variables. For example, a tuple of subject, action, object, time window and accountability of the obligation fulfillment. In some ways, P-RBAC behaves similarly to UCON, regarding the continuous access control; however, UCON only focuses on updating certain mutable attributes and does not provide any temporal constraints in obligations. Privacy enforcement demands a more concrete, temporal constraint model, and P-RBAC is capable of detecting conflicts and redundancies in UCON authorizations, obligations and conditions.

Nevertheless, there are open issues concerning P-RBAC as well. The translation of high-level privacy protection policies to low-level obligations is a burdensome task that requires automation to achieve the flexibility that we are trying to address. It would be desirable to introduce an obligation model with more flexible flow control, that also copes in cases where such translation is not expressed well. This can cause indeterminism when multiple policies can apply to the same access request.

5 Analysis and Discussion

Having reviewed the main access control models surveyed in the literature that concern the environment of the industry and cloud computing, we can conclude that there is no single solution to address authorization for all the scenarios that involve the integration of CPS with modern IT technologies. In turn, the correct choice and tailoring of the model will depend on two main factors. From a logical perspective, the intrinsic characteristics of the business model, that impose restrictions on data due to the information sensitivity and the network of stakeholders (e.g., workers, clients, providers) with whom the company collaborates.

On the other hand, from a physical perspective the infrastructure architecture also plays an important role when deciding which model to integrate within the organization. The complexity and distributed nature of the network topology can force security administrators to integrate various solutions or combine them to achieve an optimal solution depending on the set of elements to be protected in the control network. Namely, we can establish a division of the components belonging to the industrial control system into multiple sections,

Table 1: Comparison of models

MODEL	Traditional Access Control Requirements			Modern Access Control Requirements			
	Least privilege principle	Separation of duties	Delegation of capabilities	Dynamicity	Scalability	Flexibility	Quality of Service
DAC	l	l	h	l	l	l	l
MAC	l	l	l	l	l	l	l
RBAC	h	h	l	m	m	l	m
ABAC	h	h	l	m	m	m	m
UCON	h	h	l	h	h	h	m
CapBAC	h	h	m	m	m	l	h
Risk-BAC	h	h	l	h	m	l	m
OrBAC	h	h	l	h	m	h	m
T-RBAC	h	h	l	h	m	h	m
P-RBAC	h	h	l	h	m	h	m

l=low, m=medium, h=high

ranging from the corporative network, the intermediate network between the SCADA and control devices (i.e., PLCs or RTUs) and the lower layers, comprising field devices. These networks have diverse access control requirements, mainly concerning flexibility to manage heterogeneous data and quality of service to deal with systems with different computational resources.

In accordance with such logical and physical constraints, multiple access control solutions are available, as explained in this paper. We now offer a comparison of the suitability of these mechanisms for a set of scenarios that take the requirements listed in Section 2 as a baseline, thereby offering guidance for a security administrator on which technique to introduce. This comparison is shown in Table 1, and it is based upon basic security principles and the additional requirements explored in Section 2. Symbols in its cells represent how well the model matches each requirement, having three different levels: *low*, *medium* and *high*. In the following paragraphs, we give a discussion about the compliance of these aspects in every access control model.

Beginning with the traditional principles of access control, it is strongly recommended for an open environment like the one proposed in this paper that the model can support well-known principles such as the least privilege, separation of duties (both static and dynamic), and the delegation of capabilities. As for the **least privilege principle**, it is achieved when the entity can only access the resources that are necessary for its legitimate purpose; this ensures, for example, that a maintainer cannot change the behavior of other component within the production system apart from the one whose rights he owns. Related to this, the **separation of duties** is enabled when the access rights can be grouped by the different actors involved in the system (which can change over time due to the dynamicity of the proposed scenario). Both principles are satisfied by the role-based schemes, which allow to model an organizational hierarchy of privileges. This is the reason why DAC and MAC fail to ensure these two principles, since they focus on individual users (not positions) and resources when regulating permissions. However, DAC ensures that a subject can grant access rights to another user, hence complying with the **delegation of capabilities** (resulting in an h in the table). Among the rest of models, CapBAC is the only one that also supports this principle, that still needs to be improved when applied to the CPS-cloud context, with respect to a secure

management of certificates (consequently resulting in m).

As for **dynamicity** and the ability to deal with heterogeneous devices, it is ensured when the model can keep up with changes on the authorizations due to environmental conditions and the attributes presented by the subject over time (thus achieving a *high* level in the table). Pure RBAC weakly addresses this requirement by introducing constraints to roles, but lacks the definition of dynamic permissions, which is overcome in ABAC by means of attributes, thereby defining more detailed policies. UCON fully extends this functionality by including mutable attributes that represent temporary constraints. The rest of the techniques presented also achieve dynamicity over time (with an h in the table) through the definition of context-related parameters taken into account when making an access decision. Risk-BAC opts to focus on assessing the risk implied by the action performed, whose value can fluctuate; OrBAC and T-RBAC use the perspective of activities and tasks as units to group actions with the same set of access rights, that can dynamically change depending on the organizational workflows; and finally P-RBAC leverages the concept of obligations when allowing a subject to access a resource taking a fixed time. In the case of CapBAC, the inclusion of context information beyond the authorization certificates (which must be constantly maintained) is not considered.

As far as **scalability** is concerned, it will be achieved in accordance with the topology deployed by the model and its ability to cope with new users and objects in the system. in a distributed security approach. Each architecture has its own advantages. On the one hand, a centralized access control can support complex and resource-consuming computations, which makes it easier to manage the policies. However, it represents a single point of failure and the person responsible or the owners of resources do not have full control over their data. On the other hand, a distributed approach ensures a better control and granularity over data control at its source, as well as a enhanced resilience against global failures (since it can still work offline regardless of central decision points). Despite this, it is a challenge to manage and update access control policies remotely, especially when these systems are placed in a little-attended location. In addition, the integration of a complex model in certain regions of the industrial network could be complicated in presence of constrained devices. Related to this, the authors in [28] discuss the advantages and drawbacks of each approach to implement access control in a heterogeneous scenario like the Internet of Things. For the purposes of this paper, we have considered a model to be scalable when it leverages a distributed architecture (represented with h in the table). This is the case of UCON, where there exist multiple authorities distributed over the network. In some cases the model is mildly scalable (m in the table), whereas others like DAC or MAC use burdensome structures to conduct the access decision (represented with l).

With respect to **flexibility**, it is accomplished when the model allows to express granular access control policies, introducing more conditions to subjects and objects that ultimately lead to a more accurate access decision. In this regard, we must firstly mention the UCON model, that ensures a high level of expressiveness through their mutable attributes. In general, attribute-based

access control models offer a high level of granularity, which also includes OrBAC, taking an organizational approach that regulates different procedures with obligations and prohibitions. T-RBAC is also based on this managerial perspective with the definition of tasks, which achieve both active and passive access control (represented with h in Table 1). Despite such ability to declare high-level policies, this solution does not make a differentiation of sensitivity levels on data, which is the main contribution of P-RBAC. However, P-RBAC only restricts the obligations to the subjects, contrary to UCON approach, which defines attributes for both subjects and objects. In the last range of the flexibility ranking comes the Risk-BAC solution, that although contemplates a granular decision based on the current resource risk, it lacks the conception of more precise context-related attributes. Finally, CapBAC is classified in a low level of granularity since it does not feature context-awareness parameters beyond the rigid management of certificates. Nevertheless, this is precisely the reason why CapBAC turns out to be one of the most usable and user-driven models, as it overcomes the complexity of the decision by centering the access rights on these certificates, that can be issued and propagated by the own resource authority or one of its users.

Lastly, **Quality of Service** is concerned by the complexity of the access model, that affects the response time [29] in two ways: firstly, the setup time required to model the relationships among all the entities involved in the organization and their access rights, as well as creating the respective authorities. Secondly, the response time to take an access decision, that can be excessive in presence of complex access rules and subjects. As for the latter, RBAC and attribute-based access control solutions present a lower level of lightness, along with solutions that concentrate all the computation in a centralized authority, as is the case of OrBAC or T-RBAC. On the other hand, multiple of the solutions surveyed here need a first phase to design and tailor the model to the business constraints prior to implement it, which can delay the addition and setup new devices. It is the case, for instance, of Risk-BAC, which leverages a subjective process of analyzing the risks associated to the resources, or P-RBAC with the translation of high-level policies into low-level obligations. In general, all attribute-based solutions require an accurate selection of attributes, which is hard to automate (hence having a *medium* level in Table 1). Taking this twofold implication of the complexity on the responsiveness of the solution, CapBAC represents the most efficient solution, since it leverages a distributed topology and the light access decision based on capabilities ensure a lower impact on the general throughput.

Having analyzed all the requirements on the mechanisms explained, and in light of table 1, UCON poses an adequate solution for the scenario of interconnecting the CPS networks to the Cloud. Despite its complexity, it is the most suitable model to tackle the dynamic functionalities of these systems. However, it is important to stress the constant need to support these mechanisms with further security measures. There are various standards that help organizations comply with security requirements. Specially, it is worth mentioning the IEC 62351 [30], a reference in the industry that concerns security in control systems

and the protection of communication channels. In terms of access control, it suggests the customization of the RBAC model with contextual attributes to enhance its suitability for large control distributions. Similarly, NIST SP800-82 [31] proposes guidelines about the inclusion of security measures in control systems, giving recommendations on access control also when using wireless protocols and information technologies. In this regard, NIST-IR-7316 [32] introduces the main concepts of access control in industrial control systems and perform the assessment of multiple models. Nevertheless, aside from achieving an accurate and adaptive access control, it is vital to support it with features like identity management, mutual authentication, session key agreement between the users, etc. together with monitoring abilities to provide a full AAA service (authentication, authorization, and accounting). This is crucial to keep track of eventual failures or unauthorized accesses taking place in different devices of the infrastructure, that may not count on extensive memory to be able to save this evidence.

6 Conclusions and future work

In recent years, there have been a growing interconnection of traditional CPS to external networks and the integration of IT technologies, such as cloud computing. This evolution has brought with it several cyber-security issues. In terms of access control, it is mandatory to introduce exhaustive data control and permissions management among all the entities that collaborate along the production life cycle, because of the multiplicity of points of attack and the heterogeneity of technologies. In this work, we have extracted a set of requirements for access control solutions in this context, and we have assessed various solutions of the literature according to these principles. Future work will involve the creation of a richer analysis that also takes into consideration other technologies being integrated in the industrial network besides Cloud computing, as well as an accurate tailoring and creation of specific solutions adapted to this context.

Acknowledgements

The second author is supported by the Spanish Ministry of Education through the National F.P.U. Program under Grant Agreement No. FPU15/03213. In addition, this work has been partially supported by the Spanish Ministry of Economy and Competitiveness through the research project SADCIP (RTC-2016-4847-8) and the research project SMOG (TIN2016-79095-C2-1-R).

References

[1] L. Monostori, B. Kdr, T. Bauernhansl, S. Kondoh, S. Kumara, G. Reinhart, O. Sauer, G. Schuh, W. Sihn, and K. Ueda. Cyber-physical systems in

manufacturing. *CIRP Annals - Manufacturing Technology*, 65(2):621 – 641, 2016.

[2] ICS-CERT. Overview of cyber vulnerabilities. `https://ics-cert.us-cert.gov/content/overview-cyber-vulnerabilities`, last retrieved in May 2017, 2017.

[3] Federal Office for information Security. Industrial control system security: Top 10 threats and countermeasures 2016. `https://www.allianz-fuer-cybersicherheit.de/ACS/DE/\_/downloads/BSI-CS\_005E.pdf?\_\_blob=publicationFile\&v=3`, last retrieved in May 2017, 2017.

[4] International Society of Automation. ISA-95 standard. `https://www.isa.org/isa95/`, last retrieved in December 2017, 2017.

[5] Younis A Younis, Kashif Kifayat, and Madjid Merabti. An access control model for cloud computing. *Journal of Information Security and Applications*, 19(1):45–60, 2014.

[6] Vincent C Hu and Karen Ann Kent. *Guidelines for access control system evaluation metrics*. US Department of Commerce, National Institute of Standards and Technology, 2012.

[7] Ravi S Sandhu and Pierangela Samarati. Access control: principle and practice. *IEEE communications magazine*, 32(9):40–48, 1994.

[8] Butler W Lampson. Protection. *ACM SIGOPS Operating Systems Review*, 8(1):18–24, 1974.

[9] Pierangela Samarati and Sabrina De Capitani Di Vimercati. Access control: Policies, models, and mechanisms. *Lecture notes in computer science*, (2171):137–196, 2001.

[10] Hakan Lindqvist. Mandatory access control. *Master's Thesis in Computing Science, Umea University, Department of Computing Science, SE-901*, 87, 2006.

[11] Kenneth J Biba. Integrity considerations for secure computer systems. Technical report, MITRE CORP BEDFORD MA, 1977.

[12] David Ferraiolo, Janet Cugini, and D Richard Kuhn. Role-based access control (rbac): Features and motivations. In *Proceedings of 11th annual computer security application conference*, pages 241–48, 1995.

[13] Laurie B. Access control (v0.1). `http://www.links.org/files/capabilities.pdf`, last retrieved in May 2017, 2017.

[14] Vivy Suhendra. A survey on access control deployment. *Security Technology*, pages 11–20, 2011.

[15] Qamar Munawer. *Administrative models for role-based access control*. George Mason University, 2000.

[16] Jaideep Vaidya, Vijayalakshmi Atluri, Janice Warner, and Qi Guo. Role engineering via prioritized subset enumeration. *IEEE Transactions on Dependable and Secure Computing*, 7(3):300–314, 2010.

[17] Mario Frank, Joachim M Buhmann, and David Basin. On the definition of role mining. In *Proceedings of the 15th ACM symposium on Access control models and technologies*, pages 35–44. ACM, 2010.

[18] Mohammad A Al-Kahtani and Ravi Sandhu. A model for attribute-based user-role assignment. In *Computer Security Applications Conference, 2002. Proceedings. 18th Annual*, pages 353–362. IEEE, 2002.

[19] Xin Jin, Ram Krishnan, and Ravi S Sandhu. A unified attribute-based access control model covering dac, mac and rbac. *DBSec*, 12:41–55, 2012.

[20] Xinwen Zhang, Francesco Parisi-Presicce, Ravi Sandhu, and Jaehong Park. Formal model and policy specification of usage control. *ACM Transactions on Information and System Security (TISSEC)*, 8(4):351–387, 2005.

[21] Sergio Gusmeroli, Salvatore Piccione, and Domenico Rotondi. A capability-based security approach to manage access control in the internet of things. *Mathematical and Computer Modelling*, 58(5):1189–1205, 2013.

[22] Pau-Chen Cheng, Pankaj Rohatgi, Claudia Keser, Paul A Karger, Grant M Wagner, and Angela Schuett Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In *Security and Privacy, 2007. SP'07. IEEE Symposium on*, pages 222–230. IEEE, 2007.

[23] Sejong Oh and Seog Park. Task–role-based access control model. *Information systems*, 28(6):533–562, 2003.

[24] Hema Andal Jayaprakash Narayanan and Mehmet Hadi Güneş. Ensuring access control in cloud provisioned healthcare systems. In *Consumer Communications and Networking Conference (CCNC), 2011 IEEE*, pages 247–251. IEEE, 2011.

[25] Sejong Oh and Seog Park. Task-role based access control (t-rbac): An improved access control model for enterprise environment. In *International Conference on Database and Expert Systems Applications*, pages 264–273. Springer, 2000.

[26] Simone Fischer-Hübner. *IT-security and privacy: design and use of privacy-enhancing security mechanisms*. Springer-Verlag, 2001.

[27] Qun Ni, Elisa Bertino, Jorge Lobo, Carolyn Brodie, Clare-Marie Karat, John Karat, and Alberto Trombeta. Privacy-aware role-based access control. *ACM Transactions on Information and System Security (TISSEC)*, 13(3):24, 2010.

[28] Jose L Hernandez-Ramos, Antonio J Jara, Leandro Marın, and Antonio F Skarmeta. Distributed capability-based access control for the internet of things. *Journal of Internet Services and Information Security (JISIS)*, 3(3/4):1–16, 2013.

[29] Vincent C Hu, D Richard Kuhn, and David F Ferraiolo. The computational complexity of enforceability validation for generic access control rules. In *Sensor Networks, Ubiquitous, and Trustworthy Computing, 2006. IEEE International Conference on*, volume 1, pages 7–pp. IEEE, 2006.

[30] WG15 of IEC TC57. IEC 62351. `http://www.iec.ch/smartgrid/standards/`, last retrieved in May 2017, 2017.

[31] Keith A Stouffer, Joseph A Falco, and Karen A Scarfone. Guide to industrial control systems (ics) security. *Special Publication (NIST SP)-800-82 Rev 1*, 2013.

[32] Vincent C Hu, David Ferraiolo, and D Richard Kuhn. *Assessment of access control systems*. US Department of Commerce, National Institute of Standards and Technology, 2006.