Survey on Access Control for Community-Centered Collaborative Systems

FEDERICA PACI, University of Southampton
ANNA SQUICCIARINI, Pennsylvania State University
NICOLA ZANNONE, Eindhoven University of Technology

The last decades have seen a growing interest and demand for community-centered collaborative systems and platforms. These systems and platforms aim to provide an environment in which users can collaboratively create, share and manage resources. While offering attractive opportunities for online collaboration and information sharing, they also open several security and privacy issues. This has attracted several research efforts towards the design and implementation of novel access control solutions that can handle the complexity introduced by collaboration. Despite these efforts, transition to practice has been hindered by the lack of maturity of the proposed solutions. The access control mechanisms typically adopted by commercial collaborative systems like online social network websites and collaborative editing platforms, are still rather rudimentary and do not provide users with a sufficient control over their resources. This survey examines the growing literature on access control for collaborative systems centered on communities, and identifies the main challenges to be addressed in order to facilitate the adoption of collaborative access control solutions in real-life settings. Based on the literature study, we delineate a roadmap for future research in the area of access control for community-centered collaborative systems.

CCS Concepts: •**Security and privacy** → **Access control;** •**Human-centered computing** → **Collaborative and social computing systems and tools;**

Additional Key Words and Phrases: collaborative access control, policy specification, data governance, usability, literature study.

ACM Reference Format:
Federica Paci, Anna Squicciarini, Nicola Zannone, 2016. Survey on Access Control for Community-Centered Collaborative Systems. *ACM Comput. Surv.* 0, 0, Article 0 (2017), 36 pages.
DOI: 0000001.0000001

1. INTRODUCTION

The advent of new technologies such as cloud computing and social network sites has enabled new types of collaborative systems. Traditional collaborative systems (e.g., workflow management systems) focus on task sharing among multiple users [Tolone et al. 2005]. In contrast, modern collaborative systems enable both users and organizations to build communities to promote common interests, and share content with each other. Therefore, collaborative systems have evolved from being *task-centered* – wherein users gather together in "groups" to perform a common "task" – to *community-centered*, wherein "communities" of online users share information and resources of common interests.

Community-centered systems have a number of unique characteristics, in addition to the ones of task-centered collaborative systems. First, they are *complex* and *dynamic* systems, where interpersonal relations regulate the interplay among users. Users build a community by establishing online relationships that resemble their real-life interpersonal relationships. Moreover, community-

Authors are listed in alphabetical order. This work is partially supported by the ITEA2 project M2MGrids (13011) and the ITEA3 project APPSTACLE (15017). Portion of Squiciarini's work is supported by National Science Foundation under grant 1453080.

Author's addresses: F. Paci, Department of Electronics and Computer Science, University of Southampton; A. Squicciarini, College of Information Sciences and Technology, Pennsylvania State University; N. Zannone, Department of Mathematics and Computer Science, Eindhoven University of Technology.

DOI: 0000001.0000001

centered systems are *multi-party* systems. Based on established interpersonal relationships, users cooperate to create, manage and protect resources within the community. Last but not least, communities can bring together users with different background and expertise. Similar to traditional information systems and task-centered collaborative systems, the interplay of lay users calls for *easy-to-use* and *transparent* way for managing shared resources. However, this demand becomes even more crucial within community-centered systems, compared to other types of information systems, due to their complex and multi-party nature.

As community-centered collaborative systems have emerged and gained increasing popularity, the need of proper mechanisms for protecting sensitive resources shared in these systems is becoming a primary concern. Unfortunately, traditional access control models are not well suited for community-centered systems, as they are either too rigid (e.g., Mandatory Access Control), too structured (e.g., Role-Based Access Control), or not sufficiently powerful (e.g., Discretionary Access Control) to support the dynamic sharing and access control needs of communities. As a consequence, several research efforts have focused on the design of novel models and mechanisms able to provide fine-grained control over the access and usage of sensitive resources shared in these environments. In particular, much attention has been placed on access control issues in the social computing domain, wherein a large amount of personal content is often shared among users (e.g., [Fong 2011; Ahn et al. 2012; Carminati et al. 2009; Damen and Zannone 2013; Hu et al. 2013; Squicciarini et al. 2010; Xiao and Tan 2012; Rajtmajer et al. 2016; Vishwamitra et al. 2017]).

In this survey, we present a thorough analysis of the current state-of-the-art in access control for community-centered collaborative systems. First, we have identified relevant characteristics of community-centered collaborative systems based on a study of the literature, on the analysis of real-world scenarios wherein collaboration demands proper access control support and on the current state of affair of community-centered systems. Driven by these characteristics, we have identified a set of non-functional requirements that access control models and mechanisms for these systems should meet. For each requirement, we have then identified the features that access control systems should support to satisfy the requirement.

Requirements and related features used as a baseline for our analysis of the literature encompass various aspects of an access control system, and range from policy specification and governance to usability, transparency and evaluation. Some requirements are used to assess the ability of existing access control models to capture and reason on the users' interplay and dynamics of communities, e.g. interpersonal relationships and context, for access decision making. We have marked these requirements as related to 'policy specification'. Governance requirements aim to study the management of shared resources among multiple stakeholders, e.g. multiple ownership, policy combination and conflict resolution. Usability requirements are used to assess how simple an access control system is from the end user point of view in terms of its usage. Transparency requirements focus on the support provided to users for understanding the consequences of enforcing access control policies on shared resources. In addition to those requirements, we also investigate the type of evaluation used to assess the effectiveness of an access control system. This last feature aims to provide an indication of the level of maturity of existing access control systems.

We review the relevant literature according to these requirements, and provide an overview of many recent access control models and mechanisms for community-centered systems. In particular, we analyze existing proposals with respect to the identified features and, for each feature, we identify the main research trends that have emerged in the last decade. Finally, we assess to what extent existing access control models and mechanisms satisfy the elicited requirements.

Our analysis leads to several important insights. First, it unveils a growing interest in this sub-field of access control, and highlights some of the most important contributions and accomplishments in the last decades. Further, it identifies gaps in the literature and provides pointers for future research in this space. In particular, we noticed that existing works only deal with one of the characteristics of community-centered systems at the time, namely either relationship-based access control or multi-party data governance. Another major gap is the lack of concrete mechanism deployments and of their evaluation.

	Our Survey	Tolone et al. [2005]	Carminati and Ferrari [2008]	Carminati and Ferrari [2010]	Kayes and Iamnitchi [2015]	Asim and Malik [2016]
Aspects						
Policy Specification	✓	✓	✓	✓	✓	✓
Governance	✓					
Usability	✓					
Transparency	✓					
Evaluation	✓					
Collaborative Systems						
Task-centered		✓				
Community-centered	✓		✓	✓	✓	✓

Table I: Comparison of our Survey to Other Related Survey Articles

1.1. Related Work

Access control for collaborative systems has been the subject of a large body of research work. Much of this literature is summarized in a number of literature reviews [Asim and Malik 2016; Carminati and Ferrari 2008; Ferrari 2010; Samarati and De Capitani di Vimercati 2000; Suhendra 2011; Tolone et al. 2005], which we briefly discuss here. Tolone et al. [2005] discuss the merits and shortcomings of early access control models like Access Matrix [Lampson 1974], Role-Based Access Control [Sandhu et al. 1996], Task-Based Access Control [Thomas and Sandhu 1997] and Team-Based Access Control [Thomas 1997] when applied to collaborative systems. Carminati and Ferrari [2008] analyze the main research proposals on security and privacy for social networks with particular emphasis on access control. The authors list a set of requirements for access control models and systems with respect to policy specification and enforcement. Accordingly, they review existing proposals (up to the year 2008) on access control for social network sites in light of these requirements. Similarly, Carminati and Ferrari [2010] identify the requirements for the specification of access control policies tailored to online social networks (i.e., relationship-based) and for their privacy-preserving enforcement (i.e., not centralized, efficiency). Based on these requirements, the authors review the relationships-based access control models proposed at that time and cryptographic techniques for preventing the disclosure of sensitive information concerning interpersonal relationships during policy enforcement. Kayes and Iamnitchi [2015] discuss security and privacy issues in social networks. Authors focus on privacy and security risks in online social networks, and briefly discuss access control models as solutions to protect users from attacks from other social network users. More recently, Asim and Malik [2016] reviewed access control models in the context of social networks by categorizing them in relationship-based, attribute-based, community-structure-based and user activity centric models.

Table I provides a comparative analysis of our work with the literature reviews mentioned above. The comparison is based on the type of collaborative systems considered as well as on the access control aspects considered in this survey, namely policy specification, policy governance, usability, transparency and evaluation. Compared to our survey, Tolone et al. [2005] review only access control models for traditional task-centered collaborative systems rather than community-centric collaborative ones. In contrast, Carminati and Ferrari [2008], Carminati and Ferrari [2010], Kayes and Iamnitchi [2015], Asim and Malik [2016] focus on community-centered collaborative systems, but only review some of the aspects related to access control that are relevant for this type of systems. They mainly address issues related to the specification of policies, and do not consider other issues related to collaboration such as collaborative authoring and administration of policies as well as the usability, transparency and evaluation of existing solutions. In contrast, our survey provides a much more comprehensive overview of existing research efforts on access control for community-centered collaborative systems.

1.2. Organization

The paper is organized as follows. The next section provides an overview of the research challenges in designing community-centric collaborative systems and delineates the main requirements that access control models for those systems should meet. Section 3 reviews existing literature on policy specification and, in particular, analyzes existing access control models in the context of collaboration and communities. Section 4 reviews recent developments in the area of administration of shared resources within community-centered systems. Section 5 reviews approaches aiming to enhance the usability and transparency of access control systems. Section 6 discusses the evaluation methods currently used to evaluate access control systems for community-centered systems. Section 7 identifies open issues and draws a roadmap for future research. Finally, Section 8 concludes the paper.

2. SECURITY CHALLENGES IN COMMUNITY-CENTERED COLLABORATIVE SYSTEMS

This section provides an overview of the research challenges in community-centered collaborative systems and identifies the main requirements for the design of access control solutions tailored to such systems.

2.1. Community-Centered Collaborative Systems

Collaboration is becoming a fundamental part of modern information systems. The term “collaborative system” in this paper is used to refer to systems that bring geographically distributed users together, supporting communication, coordination and cooperation [Bafoutsou and Mentzas 2002]. Examples of collaborative systems range from virtual organizations, supply chains and international military coalitions [den Hartog and Zannone 2016; Trivellato et al. 2013] to remote collaborative editing and programming environments [Shen and Dewan 1992] and collaborative design [Kim et al. 2006]. These collaborative systems are task-centered in the sense that a group of users or organizations cooperate to accomplish a common task.

In this survey, we focus our analysis on an increasingly popular type of collaborative systems, namely “community-centered” (collaborative) systems. A community is “a persistent, sustained [social-technical] network of individuals who share and develop an overlapping knowledge base, set of beliefs, values, history, and experiences” [Barab et al. 2004]. Community-centered systems aim to support the interplay among the members of a community by providing them with platforms and tools that promote and facilitate the sharing of information and resources of interest.

Community-centered systems have specific characteristics that are more prominent compared to other types of collaborative systems. First, interpersonal relationships among users plays a key role in the dynamics of community-centered systems [Gates 2007; Carminati et al. 2009; Carminati and Ferrari 2010; Fong 2011], wherein relationships mirror offline social connections (e.g., friends, colleagues) among users. These relationships are typically ad-hoc and can evolve over time; moreover, they are typically not governed by a central administrative unit or a business agreement. In particular, communities are characterized by the longevity of interpersonal relationships and collaboration among their members in contrast to the task-oriented nature of other types of collaboration, e.g. virtual organizations and supply chains, which typically dissolves once the task is completed [Barab et al. 2004].

Another key characteristic of community-centered systems is given by their multi-party nature. In community-centered systems, multiple members can contribute to the creation and management of resources, being data or any other object [Damen and Zannone 2013; Hu et al. 2013]. The collaborative management of shared resources has a significant impact on their protection. In fact, it has been largely recognized that individuals can have different privacy concerns and privacy attitudes [Kokolakis 2017], which can result in the specification of possibly conflicting protection requirements for shared resources.

Last but not least, communities can bring together people with different technical background and expertise. Psychological acceptability of protection mechanisms has been largely recognized as a main challenge in the security research community since the seminal work of Saltzer and Schroeder [1975]. For instance, policy specification has been recognized as a difficult and costly task for end-users [Fang and LeFevre 2010; Klemperer et al. 2012; Squicciarini et al. 2017]. This is even more challenging in community-centered systems mainly due to the multi-party nature of these systems, wherein the members of the community should collaboratively specify protection requirements for shared resources. In addition, users can find challenges in understanding the effects of the specified policies (e.g., in terms of authorized users). This is because of the (possibly conflicting) protection requirements provided by multiple stakeholders [Hu et al. 2013; Mahmudlu et al. 2016] and because of the continued evolution of interpersonal relationships between users [Carminati and Ferrari 2010; Fong 2011; Fong et al. 2013]. This issue is further aggravated by the functionality for information sharing and community building provided by many existing community-centered systems (e.g., tagging). While promoting information sharing, these functionalities can lead to a disclosure of information that is unforeseeable and unexpected by users [Damen and Zannone 2013; McLaughlin and Vitak 2012], who can sense it as a lack of control over their information.

The most popular and largely studied form of community-centered systems is represented by social network websites. Social network websites provide an environment that allows ease of connecting people to one another and share information based on common interests. We note however, that, in addition to social networks, many other applications are emerging in the context of community-centered collaborative systems. A concrete example of non-social networking application is the emerging practice in healthcare of sharing sensitive health-revealing information in bio repositories for the purpose of genetic personalized medicine. Bio repositories collect sensitive data that reveal health-related information not only about patients but also about their families [Egea et al. 2013]. Conflicts may arise among the policies generated from patients' and other family members. For example, a patient may want to take a genetic test to verify his chances of having a cardiac arrest, which has affected most of the members of his family, and he wants to make the results available to his insurance company. However, other members of his family might be against that because they are afraid that the cost of their health insurance will increase if test results are disclosed. In this setting, it is desirable to capture the interpersonal relationships among the members of the community (the patient and his family in the example) as well as to support ways to reconcile possibly inconsistent or conflicting access control requirements.

Currently, however, most (if not all) real-world systems mandate that access decisions are manually entered by a single user. This approach does not support ways for patients to transparently control how their data will be accessed by others. Even more difficult is the question of how to relieve lay users (e.g., the patient and the members of his family) from the burden of managing the disclosure of the genetic test results. They should be able to easily express consent and exercise control over the data. Even though it is critical for patients and their families to gain better control over their genetic data, there is currently no mechanism to support these complex processes in an automated fashion. Therefore, given the highly sensitive nature of this data, it is important to support a fine-grained specification of access rights and their human-mediated enforcement. In case of inconsistencies, rather than applying blanket policies, access control mechanisms should support easy-to-use semi-automated conflict resolution methods and enable transparent data management in a coherent fashion.

Similarly, beyond this example, even social computing sites and other modern collaborative systems only support minimal features for truly collaborative access control. Most of the real-world existing access control solutions manage resources and data through simple owner-specified policies, which may later be reviewed and edited by co-managers or other stakeholders, leaving room

Req	Type	Description
R1	Policy Specification	Access control models should be able to manage the increased complexity that collaboration and communities introduce.
R2	Policy Specification	Access control models should be able to capture the dynamics and context of collaboration and communities.
R3	Governance	Access control models should allow the collaborative administration of shared resources.
R4	Governance	Access control models should support conflict resolution methods.
R5	Usability	Access control systems should be unobtrusive and should not impose extra overhead on users.
R6	Usability	Access control systems should support users in the inspection and configuration of their access preferences.
R7	Transparency	Access control systems should be transparent to users and should allow users to understand collaborative decisions and their effect.

Table II: Requirements for access control models and systems tailored to community-centered systems

for user errors, security issues, inconsistencies and other security and functional concerns. For instance, in many social network websites like Facebook the user hosting the information in their profile, has full control of the information regardless of the privacy preferences given by the user who posted the information or the user(s) to whom the information refers [Hu et al. 2011]. Besides that, existing social network websites offer other capabilities to promote information sharing and community building, possibly introducing additional complications and undesired effects on the protection of users' privacy. A typical example of such capabilities is the tagging feature provided by Facebook. It has been shown that tagging could increase the visibility of tagged objects beyond users' expectations [Damen and Zannone 2013; Pesce et al. 2012; Wisniewski et al. 2015]. In order to deal with these concerns new solutions that allow the understanding of the consequences of enforcing (possible conflicting) protection requirements on shared resources are needed.

In the last years, several research efforts have been devoted to the study and design of access control models and mechanisms able to deal with the characteristics and challenges specific to community-centered systems (e.g., [Fong 2011; Ahn et al. 2012; Carminati et al. 2009; Damen et al. 2014; Hu et al. 2013; Squicciarini et al. 2010; Xiao and Tan 2012; Rajtmajer et al. 2016]). In the next section, we elicit the requirements that access control systems for community-centered systems should meet based on the key characteristics of these systems discussed above. Then, we review the relevant literature and identify prominent trends in the area of access control for community-centered systems with respect to the elicited requirements.

2.2. Requirements

In order to support access control in community-centered systems, in addition to classic security and usability requirements considered in many access control systems (e.g., correctness, safety, reachability, feasibility etc. [Crampton and Sellwood 2014; Fong 2011; Stoller et al. 2011]), ad hoc non-functional requirements must be met. We derive a list of such requirements based on the key characteristics of community-centered systems that we have discussed in the previous section (see Table II). The identified requirements can be organized in three main classes. The first class includes requirements related to policy specification in the context of community and its dynamic nature. The second class of requirements is related to the governance of shared resources. Finally, the last class of requirements promotes the usability of access control systems and the transparency of access decision making for the members of community-centered systems (who are usually lay users). We discuss each class of requirements separately and link them with the specific characteristics of community-centered collaborative systems.

Requirements for Policy Specification. Access control models should be able to handle the complexity and dynamics of community-centered systems (**R1** and **R2**). To this end, an access control model should provide elements that facilitate the specification of access control policies in the context of community-centered systems (i.e., support for interpersonal relationships and related constraints, context). As shown in the bio data example of Section 2.1, high-level specification of access rights would simplify patients' understanding and doctors' data practices. Lay users should be able to specify whom they want their data (e.g., health record, blood work, etc.) to be shared with, without entering into complex notions of access roles or security requirements. For instance, interpersonal relationships (e.g., family members) can prove intuitive and effective to use in order to capture complex connections among users in a given domain. In fact, as first noted by Carminati et al. [2006], and later discussed by Gates [2007] and Carminati et al. [2009], one effective way to improve policy specification is to constrain access rights with respect to the interpersonal relationships or level of trust between the resource owner and resource requester. As subsequently suggested by other researchers, these relationships can also capture more complex scenarios, i.e. multi-party access control, when a binary relationship between the resource owner and requester is not sufficient [Squicciarini et al. 2010; Hu et al. 2011]. In our bio data example, this requirement translates to supporting access control policies against the different interpersonal relationships between the patient and his family and the – possibly complex – dynamics of communities. We will discuss proposals aiming to achieve those requirements in Section 3.

Requirements for Governance. Multi-party systems like community-centered systems require mechanisms for the collaborative administration of shared resources (**R3**). Of particular interest is the ability to reconcile or solve possible conflicts due to multiple administrators managing the same resource (**R4**). It is worth noting that generic policy conflicts are a long-standing issue in access control (e.g., [Bertino et al. 2003; McDaniel and Prakash 2006]), particularly in the context of database systems and for XACML policies. Yet, in the context of community-centered systems, this issue arises due to the presence of multiple stakeholders at play, and their potentially conflicting security goals on a same resource, and the lack of a central security authority or administrator. These conflicts may be solved using predefined conflict resolution strategies [Jajodia et al. 2001; Li et al. 2009; Shen and Dewan 1992] or automated conflict resolution solutions that strive for mutual agreement [Squicciarini et al. 2014; Such and Criado 2016; Xiao and Tan 2012]. Solutions addressing those requirements are reviewed in Section 4.

Requirements for Usability and Transparency. The need of ease-of-use access control systems is not new and has been largely recognized since the seminal work of Saltzer and Schroeder [1975] as a main design principle for psychological acceptability of protection mechanisms. Usability of access control systems is even more critical and challenging to achieve in community-centered systems. The complex nature of community-centered environments, where resources can be managed by several users, can make the specification and configuration of access preferences even more challenging and error-prone than in generic access control systems. Policy specification should be assisted by supporting interfaces that can help offset the burden from lay users, and address potential dependencies and inconsistencies (**R5** and **R6**). Moreover, access control systems usually make decisions in a blackbox manner [Ghai et al. 2010; Mahmudlu et al. 2016] and do not inform users about the privacy risks arising from collaborative decisions. In the context of collaboration, where data can be administrated by several users with possibly conflicting privacy and sharing preferences, it should be possible for users to understand why a certain decision on their data has been taken and what the effect of such a decision is (**R7**). We will discuss solutions aiming to enhance usability of access control mechanisms and user awareness about access decision making in Section 5.

Last but not least, we have observed a large gap between the solutions proposed by the research community and the solutions currently employed by commercial community-centered systems like online social networks. Although this can be due to the lack of financial incentive from the online collaborative platform providers, it may also be an indication of the limited maturity of this research field. In particular, the extent to which mechanisms for community-based access control are actually *applicable* (and therefore tested) to real-world systems is important, and constitutes an additional, ad-hoc requirement. We provide an overview of the methods currently used for the evaluation of access control solutions for community-centered systems in Section 6.

3. POLICY SPECIFICATION

A main component of access control systems is represented by the *policies* used to regulate the access to data. An access control policy defines high-level rules specifying who can access a protected object and under which conditions [Samarati and De Capitani di Vimercati 2000]. Ideally, policies should be able to express conditions based on the needs and characteristics of the target application domain. In the context of community-centered systems, we have identified two main desiderata for the specification of access control policies, namely the ability to specify conditions under which access is granted based on interpersonal relationships among users (req. **R1**) and context information related to collaboration and community (req. **R2**).

Several access control models have been proposed for the specification of access control policies. An *access control model* provides a formal representation of access control policies and their evaluation [Samarati and De Capitani di Vimercati 2000], including the concepts used to represent access conditions tailored to the application domain. In this section, we review existing access control models in the context of collaboration and community. In particular, we first present an overview of access control models that have not been designed specifically for community-centered systems, and discuss their limitations; then, we review access control models that are used to handle authorization in community-centered environments and, in particular, a novel access control paradigm called relationship-based access control.

3.1. Overview on Conventional Access Control Models

Many early access control models rely on the notion of access matrix [Lampson 1974]. In an access matrix A, rows represent users, columns represent objects, and each entry $A[s, o]$ indicates the access rights that user s has on object o. Access matrix has several limitations, especially when considering community-centered systems and collaborative environments in general. First of all, access matrix is an identity-based access control model, in which access rights have to be specified for each subject and object individually. This makes the specification and management of access rights impractical for open and complex systems like community-centered systems. Moreover, this model does not account for context information and, thus, is not able to capture the dynamic nature of collaboration and community.

To overcome these limitations, several other access control models have been proposed over the years. Examples of such models are Role-Based Access Control (RBAC) [Sandhu et al. 1996] and Task-Based Access Control (TBAC) [Thomas and Sandhu 1997]. RBAC introduces the notions of role and role hierarchy to simplify the specification and management of access control policies within an organization. Roles are used to model job functions within an organization; specifically, a role represents the set of permissions needed to carry out a certain job function [Sandhu 1996]. A role hierarchy defines the inheritance of permissions between roles based on the organization structure. In RBAC, permissions are not assigned to users directly; instead, permissions are assigned to roles and users inherit the permissions assigned to the role(s) they have (directly or through role hierarchy). TBAC proposes to organize and manage user permissions with respects to the tasks to be executed. Although the aforementioned access control models provide some basic features to facilitate the specification and management of access control policies for collaborative environments, they are not able to fully handle the complexity of collaboration and communities. Both RBAC and TBAC provide very limited support for modeling and reasoning on the context; in these models context information is usually limited to tasks and workflow progress, which is insufficient to capture the complex nature of collaboration [Tolone et al. 2005]. Moreover, they do not use interpersonal relationships among users in access decision making.

Other researchers have introduced the notion of team in access control to capture the effect of group dynamics on the protection of sensitive resources. For instance, TeaM-based Access Control (TMAC) [Thomas 1997] extends RBAC with the notion of team to group users for access control purposes. Teams are structured in terms of roles, thus providing a means for defining the collaboration context for the activities to be executed. However, TMAC does not account for the interpersonal relationships among team members, limiting its applicability to community-centered environments. Similarly to TMAC, Bullock and Benford [1999] present an access control mechanism that aims to support people working together in collaborative teams. The authors investigate the group access functionality within the SPACE access control model [Bullock and Benford 1997] and, in particular, define a number of policies that constrain the access rights of a group of users based on the properties of the team as a whole.

A prolific research stream has explored how to balance the ability to collaborate and the security of shared resources across organizational boundaries. This stream has resulted in a number of coalition-based access control models [Cohen et al. 2002; Phillips et al. 2002; Atluri and Warner 2004; Trivellato et al. 2013] and workflow-based access control models [Kang et al. 2001; Wainer et al. 2003]. These models usually extend earlier access control models like RBAC, TBAC and TMAC to address requirements typical of inter-organizational resource sharing. However, these models as their precursory models do not presume any connection among users; they are based on well-defined coalitions and organizations, with predefined roles and hierarchies among users. On the contrary, in community-centered systems, access decisions are often made according to ad-hoc relationships, which can evolve over time and are not governed by a central administrative unit.

The need of flexible and dynamic access control systems has also led to the emergence of new access control models like Attribute-Based Access Control (ABAC) [Hu et al. 2014b] and Usage Control (UCON) [Park and Sandhu 2004]. ABAC is a general-purpose access control model in which permissions are constrained with respect to the attributes of entities (subjects and objects), operations and the environment. Similarly, UCON expresses authorizations and obligations in terms of subject and object attributes; moreover, it employs conditions to check the environmental status. The use of attributes makes policy specification extremely flexible and expressive, and many earlier access control models can be expressed in ABAC and UCON. Among the others, environment attributes (and conditions in UCON) allow these models to be context-aware, thus making them suitable for a variety of applications including community-centered systems. Moreover, interpersonal relationships can also be encoded as attributes. However, it has been observed in [Crampton and Sellwood 2014] that it may be difficult to capture and manage the complex dynamics of communities (e.g., chains of interpersonal relationships) in ABAC.

3.2. Access Control Models for Community-centered Systems

Community-centered systems like social networks provide an environment where users can establish interpersonal relationships and share information with other users. These relationships play an important role for users to determine to whom their data can be disclosed. For instance, a user may want to share some information with his/her colleagues and other information with his/her friends and family. This has prompted several efforts to accommodate the use of interpersonal relationships between users in access decision making.

Most of the existing real-world and commercial community-centered systems represent interpersonal relationships among their users by employing the notion of group. A group is a collection of users [Sandhu 1996].[1] The notion of group has led to the emergence and adoption of Group-based Access Control (GBAC) models, especially in online social networks. GBAC shares some ideas of RBAC: permissions are assigned to groups and users inherit all permissions assigned to the group(s) they belong. Membership to a group can be established differently depending on the type of community-centered systems. This difference has led to the definition of different types of GBAC models. For instance, in some community-centered systems, group membership is based on

[1] Despite the similarity between the concepts of ‘role’ and ‘group’, in this work we make a clear distinction between these two concepts. In particular, roles are defined in terms of permissions, whereas groups are defined in terms of users. We refer to [Sandhu 1996] for a discussion on the differences between these two concepts.

subscription; users can subscribe to a group, for instance, based on common interests.[2] An access control model targeting this type of communities is Group-Centric Secure Information Sharing (g-SIS) [Krishnan et al. 2011]. g-SIS proposes to bring users and information together into groups in order to facilitate information sharing and collaboration by allowing group members to access information belonging to the group.

In other types of community-centered systems like social networks, users should assign their contacts to groups and specify which of these groups can access to a certain object. Initially, social networks only allowed users to group their contacts based on a general notion of 'friendship'. This, however, is not sufficient to capture the complexity and diversity of interpersonal relationships between users, thus failing to provide users with fine-grained control over their information. To offer users more control over their information, online social networks have started providing them with the ability to manage multiple social groups as well as to define customized groups. Although this enables a more fine-grained control on the disclosure of information, some studies [Lederer et al. 2004; Jones and O'Neill 2010] have shown that managing groups can be a significant burden for users, especially when the number of contacts and relationships is large. To alleviate this burden, researchers have investigated automated approaches to grouping based on network clustering techniques [Jones and O'Neill 2010; Jones and O'Neill 2011; Xu et al. 2007]. Achieving full automation may however not be possible due to the variability of criteria that different people use to group their contacts and the context of the collaboration [Olson et al. 2005]. Moreover, the GBAC models usually employed by existing social networks do not allow users to specify groups arbitrarily, for instance, based on chains of interpersonal relationships; in these systems, groups are typically used to denote direct interpersonal relationships among users.

To overcome these limitations, the last years have witnessed the emergence of a new trend aiming to support users in the specification of fine-grained access control policies for community-centered systems. As first noted by Carminati et al. [2006], the sole use of direct interpersonal relationships is not flexible enough in denoting authorized users in community-centered systems. Based on this observation, Gates [2007] has introduced a novel access control paradigm based on interpersonal relationships, called Relationship-Based Access Control (ReBAC), and several ReBAC models [Aktoudianakis et al. 2013; Bruns et al. 2012; Carminati et al. 2009; Crampton and Sellwood 2014; Fong et al. 2009; Fong 2011; Fong et al. 2013; Squicciarini et al. 2014] have been proposed since her seminal paper. These models allow the specification of access control policies that employ social relationships as the key factor in access decision making. In particular, access decisions are made on the basis of primitive (e.g., friend) and composite (e.g., friend-of-friend) relationships between the resource requester and the resource owner. It is worth noting that composite relationships resemble chains of trust and delegation of authority, which have been largely studied in the area of trust management [Ahn et al. 2012; Li et al. 2002; Trivellato et al. 2014].

Interpersonal relationships among users are often represented using a *social graph*. In a social graph, nodes denote the entities within the system and edges denote the interpersonal relationships between those entities. Many ReBAC models like the ones proposed by Bruns et al. [2012], Carminati et al. [2009], Fong [2011] and Crampton and Sellwood [2014], rely on poly-relational social graphs. In these models, social graphs are extended by associating edges with a label indicating the *type* of relationship between two entities (e.g., friend, colleague, family). ReBAC models usually express access constraints in terms of paths within the social graph. The main difference between these models lies in the way constraints are represented and the security properties that can be verified. For instance, Fong [2011] represents access constraints as formulas in modal logic, thus providing a robust mathematical foundation for relationship-based access control and formal analysis capabilities to assess policy correctness. However, it has been shown in [Fong and Siahaan 2011] that modal logic is not representationally complete and several relationship-based policies typical of Facebook-style social networks cannot be expressed in the language proposed by Fong [2011].

[2] Note that in this type of community-centered systems interpersonal relationships are built upon affiliation to a group.

Bruns et al. [2012] address this issue by proposing the use of (a fragment of) hybrid logic for the specification and enforcement of relationship-based policies. Differently from modal logic, hybrid logic allows binding a node to a principal in the social graph. This makes it possible to express graded modalities that are necessary for the specification of relationship-based policies such as "access is granted to requesters who have at least k friends in common with the owner" [Bruns et al. 2012]. Crampton and Sellwood [2014] propose a generic ReBAC model based on path conditions. Path conditions, represented as sequences of relationships, are used to map the resource requester to a (set of) authorization principal(s). Intuitively, authorization principals are the ReBAC counterpart of roles in RBAC. In this model, the subject and resource specified in the access request are first used to find the set of applicable principals. These principals are then used to determine whether the action specified in the request should be authorized or denied based on a given authorization policy.

Although accounting for the (type of) interpersonal relationships with other users provides users with a powerful control over who can access their resources, a more fine-grained control may be desirable. In order to increase users' empowerment in the control of their information, Carminati et al. [2009] propose to use the *depth* and *strength* of relationships. The depth of relationships represents the length of the shortest path between two nodes in the social graph, and it is used to control the radius of the social circle to whom access should be granted, along the line of previous work in the areas of trust management [Li et al. 2003] and security requirements engineering [Giorgini et al. 2006]. Moreover, Carminati and colleagues annotate interpersonal relationships with a trust level representing the strength of relationships and specify access control policies in terms of access conditions. Intuitively, an access condition is a statement specifying the type of relationship that the resource requester should have with the resource owner possibly along with constraints on the depth and trust level of the relationship. Similarly to Carminati and colleagues, Hu et al. [2011], Such and Criado [2016] and Ilia et al. [2017] explicitly model the strength of relationships between the resource owner and requester and use this information to resolve multi-party conflicts (see Section 4).

Accounting for interpersonal relationships between users in access decision making has led to the definition of topology-based policies [Carminati et al. 2009; Fong et al. 2009; Squicciarini et al. 2014]. Intuitively, topology-based policies rely on topological properties of the social graph to determine the access rights of users. A variety of topology-based policies has been proposed by Fong et al. [2009]: degree of separation, in which access is permitted if the resource requester is within a social circle of a given radius from the resource owner; known quantity, in which access is permitted if the resource owner and resource requester share at least a given number of friends; clique, which strengthens the known quantity policy by requiring that the resource owner and resource requester are in a k-clique (i.e., in a fully connected graph of order k); trusted referral, in which access is permitted if the requester has a friendship relation with a given number of users in a given set; and stranger, in which the resource requester is granted access if it has at least a certain social 'distance' from the resource owner. Carminati et al. [2009] extend the 'degree of separation' policy by associating a trust level to relationships and requiring that the strength (i.e., the trust level) of the path between the resource owner and requester is greater that a certain threshold. Squicciarini et al. [2014] propose to base access decisions on the closeness between the resource owner and requester in terms of common neighbor where access is allowed if the resource owner and requester share k common neighbor.

Community-centered systems are usually dynamic, wherein the interpersonal relationships between users can change and evolve over time. The dynamics of IT systems is often captured and supported in access control models by accommodating the notion of *context* [Bhatti et al. 2005; XACML v2.0 2005; XACML v3.0 2013]. These models, however, only focus on context information related to the environment. Fong [2011] observes that interpersonal relationships are also contextual and may only exist in certain contexts. Fong has dealt with this issue by mapping contexts to social graphs. Intuitively, a social graph represents the interpersonal relationships that hold in a given context. On the other hand, Fong et al. [2013] account for the evolution of communities by extending the language proposed in [Fong 2011] with temporal operators. This extended lan-

	Mazzoleni et al. [2008]	Wishart et al. [2010]	Squicciarini et al. [2010]	Hu et al. [2011]	Carminati and Ferrari [2011]	Xiao and Tan [2012]	Hu et al. [2013]	Damen et al. [2014]	Ilia et al. [2015]	Mahmudlu et al. [2016]	Mehregan and Fong [2016]	Rajtmajer et al. [2016]	Such and Criado [2016]	Ilia et al. [2017]	Vishwamitra et al. [2017]
Data Governance Model	MO	MO	MO	As	MO	MO	As(SO)*	As	SO	As	MO	MO(SO)**	MO	MO	MO
Policy Combination	MA	MA	–	–	–	–	–	Au	–	Au	MA	–	–	MA	–
Conflict Resolution	Pr	CF	MP	MP	Pr	MP	Pr	Pr	–	Pr	CF	MP	MP	Pr	MP

Data Governance Model:		**Policy Combination:**		**Conflict Resolution:**	
SO:	Single-Ownership	Au:	Authoritative	Pr:	Predefined
MO:	Multi-Ownership	MA:	Mutual-Agreement	MP:	Multi-Party Decision
As:	Asymmetric			CF:	Conflict-Free

* This work assumes a single owner of data to possess the highest priority in the control of sharing.
** This work assumes a leader that can set an upper- and a lower-bound of sharing.

Table III: Governance for collaborative systems

guage allows the specification of policies encompassing both relationship-based and history-based elements, thus accounting for past user interactions in access decision making.

4. GOVERNANCE

Most of the existing access control models are based on the assumption that every resource is administered by a single entity, usually called resource *owner*. Intuitively, the owner of a resource is the entity that has full control on the resource and defines the access control policy to regulate how and with whom the resource can be shared. This view, however, has to be revised in order to deal with collaborative systems and, in particular, with community-centered systems. In such systems, resources can be simultaneously administered by several entities, each of which might have a different degree of "ownership" or stake in the resource protection. Further, resource owners may not be security experts, or have a complete view of the domain within which the resources are to be shared. This complex scenario raises a number of new challenges. For instance, each entity should be able to specify its own authorization requirements autonomously, which should be eventually merged or replaced by a single access control policy in order to determine to whom a shared resource can be disclosed (req. **R3**). This autonomy in policy specification can lead to obvious conflicts, i.e. an entity may want to grant access to a certain user and another entity denies the access to the very same user. Ideally, an access control system should provide a means to resolve such conflicts (req. **R4**).

In this section, we review recent developments in the area of policy administration within collaborative systems. By administration, here we refer to models and mechanisms supporting the authoring and enforcement of policies, based on some underlying administrative protocols. More specifically, we have identified three main features concerning the collaborative specification and administration of access control policies: *data governance model*, which defines the authority that entities have over a resource; *policy composition*, which describes how the authorization requirements authored by multiple entities are combined or reconciled to regulate the access to a resource; *conflict resolution*, which indicates the method used to resolve policy conflicts in order to obtain a conclusive decision. A summary of our review is presented in Table III.

Data Governance Model. Within collaborative environments, several entities can be involved in the creation, management and protection of sensitive resources [Damen and Zannone 2013; Hu et al. 2013; Mazurek et al. 2014]. For instance, in social network websites data can be posted by one user, can be hosted in the profile of some other user, and refer to yet other users. Another illustrative example is the protection of information derived by the fusion of data collected from different sources under the control of different entities [den Hartog and Zannone 2016]. In these scenarios, each entity within the collaboration retains a share of authority on the data. Specifically, each entity related to a data object can define its own authorization requirements indicating which users should be able to access the object and which should not [Hu et al. 2013]. The requirements of every entity should be then combined in order to determine the actual permissions on the object. To this end, various data governance models have been proposed. Intuitively, the data governance model specifies the level of authority or stake in a resource that users have.

A recent trend proposes to handle the governance of objects involving multiple stakeholders by decomposing an object into subparts and assigning each subpart to the corresponding stakeholder. In this respect, Ra et al. [2013] present the P3 algorithm that allows the selective encryption of a portion of an image. Ilia et al. [2015] exploit this type of techniques to shift the level of control from photos to the faces appearing in a photo. In particular, the authors use face recognition techniques to identify the users appearing in a photo. Each of these users can specify his/her own permission on the photo, indicating which subjects can see his/her face in the photo. When a subject attempts to view the photo, a new photo is created based on the permissions specified by each user appearing in the original photo. The newly created photo only shows the faces that the subject is allowed to see whereas the other faces in the photo are blurred. These solutions, however, are mainly tailored to control the sharing of photos/images and might not be suitable for other types of objects and resources. Moreover, they do not account for users' interplay to deal with multi-party conflicts, which is a main desideratum for community-centered environments. Thus, we do not consider these approaches further in our analysis.

A number of access control models for community-centered systems assume a *multi-ownership* data governance model. These models enable the collaborative specification of policies aiming at the protection of resources co-owned by multiple users [Bahri et al. 2015; Carminati and Ferrari 2011; Ilia et al. 2017; Mehregan and Fong 2016; Rajtmajer et al. 2016; Squicciarini et al. 2010; Wishart et al. 2010; Xiao and Tan 2012; Vishwamitra et al. 2017]. Data governance models based on multi-ownership implicitly assume that all entities related to a data object have the same level of authority over the object. More recently, researchers have acknowledged the possible differences among various stakeholders involved in the protection of sensitive resources. For instance, Hu and colleagues identify four main controllers for shared data in a social network environment based on the relation that users have with a given data object [Hu et al. 2011; Hu et al. 2013]: owner, contributor, stakeholder and disseminator. The owner is the user hosting the data; the contributor is the user who published the data; a stakeholder corresponds to a tagged user; and the disseminator is a user who shares a data object from another user's profile to its profile. Each of these entities can have a different level of authority on the resource to be protected.

This has resulted in the definition of *asymmetric* data governance models in which the degree of authority an entity has over a resource to be protected depends on its relation with the resource. In particular, different (types of) controllers may have a different level of influence on the final decision. For example, Hu et al. [2013] assume that the owner of the data (i.e., the data host) has the highest priority in the control of the data and is responsible to decide how controllers' authorization requirements should be combined. This, however, results in an access control model that is single owner-centric. Moreover, from a privacy perspective, it may not be desirable to grant full authority to the data host. According to privacy and data protection regulations, at least for selected applications, data subjects should be able to influence how their data are processed [Guarda and Zannone 2009; Xiao and Tan 2012].

A new line of work is now focusing on asymmetric data governance models that go beyond the notion of ownership. In this respect, Damen et al. [2014] have introduced the notion of archetype to characterize the relations of a user with a resource. The archetypes for a shared resource are used to determine the extent of control that users have over the object. Mahmudlu et al. [2016] have extended the work in [Damen et al. 2014] by proposing a general data governance framework for collaborative environments. This governance model relies on an archetype hierarchy to reason on and prioritize the degree of authority that users have over a shared resource based on their relation with the given resource. Based on the archetype hierarchy, the framework provide a means to combine the authorization requirements specified by multiple controllers in order to form the access control policy used to regulate the access to the resource.

Policy Composition. In collaborative systems, each entity related to an object can independently define authorization requirements specifying who is authorized to access the object. These requirements have to be combined in order to build the access control policy defining the permissions on objects managed by multiple entities. Most of the existing solutions (marked with a dash in Table III) evaluate the policies of each entity individually and then apply strategies to combine the (possibly conflicting) decisions obtained by such evaluations (we will discuss these strategies below). However, this approach limits the strategies that can be applied, and does not allow the use of combinations of strategies where the decisions of some entities are combined using one strategy, the decisions of other entities are combined using another strategy and the obtained decisions are combined using yet another strategy.

To deal with this issue, a few researchers have proposed to combine the authorization requirements of each entity into a single (possibly structured) policy. Approaches for policy combination can be grouped in two main classes. Solutions like the ones in [Damen et al. 2014; Mahmudlu et al. 2016] use an *authoritative* approach in which the authorization requirements provided by the various entities are combined in a predefined manner. These solutions usually adopt or extend languages that allow the specification of federated access control policies such as OrBAC [Kalam et al. 2003] and XACML [XACML v2.0 2005; XACML v3.0 2013], to support the collaborative authoring of access control policies.

Another line of research aims to reconcile the authorization requirements of the entities related to the object to be protected through a *mutual agreement* between these entities. Among these solutions, Mazzoleni et al. [2008] propose a policy similarity process and various policy integration algorithms to combine policies authored by multiple authorities. The policy similarity process aims to determine which policy is more restrictive. Possible access requests are matched against the controllers' policies and the outcome of this analysis drives the selection of the best-fitting policy integration algorithm for the co-owned resource.

Some researchers advocate the need of interactive negotiation of access control policies to determine how co-owned objects should be protected. Wishart et al. [2010] use weak conditions and strong conditions to represent the negotiable and non-negotiable constraints of each co-owner respectively. However, the authors do not provide any automatic or semi-automated mechanism to support the collaborative authoring of privacy policies and, in particular, the negotiation of privacy restrictions between co-owners. Entities have to reach an agreement through a multi-round negotiation process in which, at each round, entities have to manually review or accept the weak conditions defined by other entities. Mehregan and Fong [2016] propose an interactive policy negotiation framework to support the authoring of access control policies in collaborative manner based on ReBAC. The framework comprises a multi-ownership administrative model in which access control policies are defined as graph patterns, and an interactive policy negotiation protocol, which allows co-owners to reach an agreement on the access control policy to be enforced. The protocol supports co-owners in the iterative revision of the policy of co-owned objects in which co-owners can express their authorization requirements by counteracting the authorization requirements of other co-owners. At each iteration, the obtained policy is verified against three availability criteria, namely satisfiability, feasibility and resilience, thus providing a balance between privacy and sharing. Ilia et al. [2017] propose a multi-party access control model based on a threshold-based secret sharing scheme. The basic idea underlying this model is to associate a set of shares to a secret and distribute these shares among co-owners. Shares can be redistributed by co-owners to trusted users (so called shareholders) based on selection rules specified in ReBAC. Along with a share, co-owners provide shareholders with a share provision rule that defines to whom a delegated shared can be disclosed. Similarly to selection rules, shared provision rules are specified in ReBAC. In order to access an object, a requester has to collect from the shareholders a number of shares sufficient to reconstruct the decryption key (i.e., the secret), where the needed number of shares depends on the sensitivity of the object.

Although solutions based on mutual agreement between the parties involved can potentially result in conflict-free policies, this approach requires renegotiating the policy every time a co-owner changes its authorization requirements. Thus, a large body of research proposes to deal with policy conflicts when they occur by providing methods for conflict resolution. In the remainder of the section, we present an overview of these approaches.

Conflict Resolution. Every user demands that his/her authorization requirements are enforced. However, users can define conflicting requirements for the same resource and, thus, it is not possible to satisfy all users' requirements. Most access control mechanisms employ policy conflict resolution strategies to determine how conflicting requirements should be reconciled. Existing conflict resolution strategies can be broadly classified into two main classes: *predefined strategies* and *multi-party decisions*.

Predefined conflict resolution strategies consist of predefined rules that are used to determine how (possibly conflicting) access decisions should be combined [Shen and Dewan 1992]. These strategies are usually based on a (partial) ordering of access decisions or on the policy structure. An overview of the predefined conflict resolution strategies proposed in the literature is given in Table IV. Typical examples of strategies based on priority between decisions are permit-overrides and deny-overrides. According to permit-overrides, a policy evaluates Permit if at least one of its sub-policies evaluates Permit; deny-overrides is the dual of permit-overrides: a policy evaluates Deny if at least one of its subpolicies evaluates Deny. An example of strategy based on the policy structure is first-applicable. This strategy requires that subpolicies are evaluated in the order they are specified and the access decision is the one of the first applicable policy. The first-applicable strategy is typically used to capture that an entity predominates another entity.

In the literature, we can also find a number of combining algorithms based on consensus and majority voting. Among the others, we mention the work of Hu et al. [2013] and Ilia et al. [2017] in which voting schemes are combined with weights to capture the fact that the policies of some users weighs more than the policies of other users. To deal with policy conflicts, Carminati and Ferrari [2011] introduce a new type of policies, called collaborative security policies. The goal of these policies is to involve collaborative users in decision making by gathering feedback from the collaborative users on whether a certain access should be granted. In particular, these policies specify how this feedback should be combined: all indicates that access is granted if all collaborative users agree on releasing the resource (i.e., consensus); one indicates that access is granted if at least one user agrees on releasing the resource; majority indicates that access is granted if the majority of collaborative users agree on releasing the resource. These collaborative policies resemble strong-consensus, permit-overrides and weak-majority strategies respectively. The aforementioned studies provide strategies to implement a variety of specific and ad-hoc policy integration requirements. A few works (e.g., [Bruns and Huth 2011; Rao et al. 2011]) have proposed policy algebras for the specification of policy combination strategies. It is worth mentioning that these algebras are able to encode most of the conflict resolution strategies presented in Table IV.

Some recent work has focused on conflict resolution in the context of multi-party decisions. In particular, researchers have analyzed conflict resolution from a game-theoretic point of view and proposed negotiation protocols based on game-theoretic concepts such as the Nash equilibrium [Hu et al. 2014a; Rajtmajer et al. 2016; Squicciarini et al. 2010; Such and Rovatsos 2016; Xiao and Tan 2012]. For instance, Squicciarini et al. [2010] use the Clarke-Tax model and game theory in the context of an auction to resolve policy conflicts by selecting the privacy policy that maximizes social utility. CAPE [Xiao and Tan 2012] addresses the problem of multi-party conflicts by taking into account the social interactions among the co-owners of the data in conflict resolution. In par-

Conflict resolution strategy	Shen and Dewan [1992]	Jajodia et al. [2001]	Ashley et al. [2003]	Li et al. [2009]	Reeder et al. [2009]	Carminati and Ferrari [2011]	Matteucci et al. [2012]	Hu et al. [2013]	Ilia et al. [2017]	XACML v2.0 [2005]	XACML v3.0 [2013]
permit-overrides		✓				✓				✓	✓
deny-overrides		✓			✓					✓	✓
first-applicable			✓		✓					✓	✓
only-one-applicable										✓	✓
permit-unless-deny											✓
deny-unless-permit											✓
specificy-precedence	✓	✓			✓		✓				
weak-consensus				✓							
strong-consensus				✓		✓		✓			
weak-majority				✓		✓					
strong-majority				✓				✓	✓		
super-majority-permit				✓				✓	✓		
owner-overrides								✓			

Table IV: Overview of existing conflict resolution strategies

ticular, CAPE uses a graph theoretic model to reach consensus between the co-owners of shared resources by considering the intensity with which the co-owners are willing to share their co-owned object and the extent to which they want their decisions to be affected by the action of their acquaintances. These proposals, however, are not able to capture users' behavior accurately [Such and Criado 2016]. In fact, they assume that users' behavior is perfectly rational and do not account for the social facets that are typically considered by users when they have to deal with multi-party conflicts. Rajtmajer and colleagues have dealt with this issue in a recent effort [Rajtmajer et al. 2016]. The authors followed research suggesting that realistic models of human behavior should relax perfect-rationality assumptions in favor of *bounded rationality* models accounting for limited time and information, as well as cognitive limitations [Kahneman 2003; Simon 1957]. Accordingly, they propose a bounded rationality model building on the quantal response and recency models [Rajtmajer et al. 2016], which accounts for users' limitations and yet provides guarantees of success and well models users' actions.

Recent work also recognized that users can be influenced by their acquaintances and, thus, can adjust their decision accordingly [Hu et al. 2011; Such and Criado 2016; Vishwamitra et al. 2017]. For instance, Hu et al. [2011] propose a mechanism for collaborative data management in online social networks. In particular, authors propose to balance the users' need of privacy protection and their desire for information sharing in the detection and resolution of multi-party privacy conflicts. The detection of multi-party conflicts is based on the segmentation of the space of possible access requesters with respect to the settings of all controllers of the shared resource. The identified conflicts are resolved by finding an optimal balance between the potential privacy threats caused by the disclosure of the information (the so called privacy risk) and the loss caused by the non-disclosure of the information (the so called sharing loss). Privacy risk and sharing loss are assessed for each single controller based on the sensitivity of the shared resource and the strength of the relationship of the controller with the resource requester. Such and Criado [2016] propose to resolve multi-party conflicts based on the willingness of each co-owner to change its privacy settings. Similar to the notion of privacy risk and sharing loss proposed by Hu and colleagues, a co-owner's willingness is estimated based on the sensitivity of the object and the importance of the resource requester for the co-owner. To alleviate users from the burden of providing such measures, Such and colleagues show how they can be computed on the basis of the strength of the relationships between the co-owners and the resource requesters. Co-owners' concession to relax their privacy constraints is modeled through predefined fuzzy rules (i.e., "I Do Not Mind", "I Understand", "No Concession") that define how co-owners would actually negotiate their privacy constraints. During negotiation, if every co-owner is willing to change its decision, the decision that is preferable by the majority of the co-owners is selected. Otherwise, if a co-owner has high willingness and another low willingness, the one with high willingness will concede and the one with low willingness will determine the access. If at least two co-owners with low willingness have conflicting policies, access will be denied, implicitly assuming a deny-overrides strategy.

5. USABILITY & TRANSPARENCY

A cornerstone of community-centered collaborative systems is that multiple users are in charge of specifying access control policies to be applied to shared resources. Often, these users have no technical expertise and, therefore, policies readability and usability of the policy specification mechanisms plays a central role in the effectiveness of these systems [Cao and Iverson 2006]. Easy-to-use access control systems should reduce the effort spent by users in authoring and configuring conflict-free access control policies (req. **R5**). Moreover, user-friendly interfaces should help users in specifying and configuring policies for shared resources (req. **R6**). Ideally, such interfaces should facilitate users' understanding of the interaction of their policies with the policies specified by other users and help resolve possible conflicts. Another crucial aspect of access control systems tailored to community-centered systems is transparency. Transparency pertains to the information available to users with respect to a (possibly collaborative) access decision (req. **R7**). In particular, users should be able to understand the effect of the specified policies and be notified when and why the decision enforced on shared resources differs from their own.

These requirements are largely understudied compared to other more conventional requirements presented in Section 2. Nonetheless, as the research community continues to focus on community-centered access control, we note that usability-related issues are of increasing importance for adoption and wide-spread use of access control mechanisms. To date, the problem of supporting end-users in policy authoring and configuration is attracting growing attention. Yet, only a few works that aim to promote transparency in collaborative access control exist thus far. In the reminder of this section, we review the main research efforts on policy authoring, comprehension and configuration as well as transparency in access control, with a particular emphasis on the ones specific to community-centered access control. An overview of these contributions is shown in Table V.

Policy Generation. A number of researchers have focused on the automated or semi-automated generation of meaningful access control policies to assist non-experts users in policy authoring, especially for image sharing in social networks. Existing approaches vary significantly with respect to the features used to generate privacy preferences; see Table VI for an overview. Some approaches derive privacy settings from image metadata (e.g., tags) [Klemperer et al. 2012; Squicciarini et al. 2010], other from personal traits (e.g., location, age) [Squicciarini et al. 2014; Squicciarini et al. 2014] and/or social context (e.g., friend list) [Fang and LeFevre 2010]. Squicciarini et al. [2010] discuss two alternative approaches to assist users in the automated generation of privacy settings, thus, freeing them from the burden of specifying a privacy policy for co-owned images. One approach is based on similarity analysis among pictures based on the tags associated with them: if two pictures are similar, the same privacy policy will be applied to both. The other approach leverages collaborative filtering: privacy policies are determined based on the ones specified by similar users on similar images. Fang and LeFevre [2010] propose a social networking privacy wizard to configure users' privacy setting automatically. The wizard requires a user to assign a privacy label

Category	**Approach**	**Source**
Policy Generation	Similarity Analysis Collaborative Filtering	[Squicciarini et al. 2010]
	Rule-based	[Klemperer et al. 2012; Squicciarini et al. 2014; Squicciarini et al. 2014]
	Association Rule Mining	[Squicciarini et al. 2015]
	Machine Learning-based	[Fang and LeFevre 2010; Spyromitros-Xioufis et al. 2016; Zerr et al. 2012; Squicciarini et al. 2017; Fogues et al. 2017]
	Deep Learning-based	[Tonge and Caragea 2016; Yu et al. 2017]
Interfaces for Policy Comprehension	Circle Pie Chart	[Rode et al. 2006]
	Expandable Grid	[Reeder et al. 2008]
	Eyes Metaphor	[Schlegel et al. 2011]
	Mirror-Looking Metaphor	[Anwar and Fong 2012], Facebook, LinkedIn
Interfaces for Policy Configuration	Bubble chart	[Mazzia et al. 2012]
	Genome bar sequence	[Wang et al. 2015]
	Circle management	[Hu et al. 2012]
Feedback Generation	Justification of denied requests to resource requester	[Bonatti et al. 2001; Kapadia et al. 2004; Ghai et al. 2010]
	Notification of policy conflicts to resource controller(s)	[Damen et al. 2014; Mahmudlu et al. 2016]
	Justification of policy conflicts to resource controller(s)	[den Hartog and Zannone 2016]
	Effect of policy conflicts to resource controller(s)	[Hu et al. 2013]

Table V: Mechanisms and interfaces for usability and transparency in access control systems

to selected friends. Machine-learning techniques are then employed to create a model of the user's privacy preferences based on these labels, which is then used to configure the privacy settings concerning unlabeled friends automatically. Squicciarini et al. [2014] adopt a rule-based approach to automatically generate access rules for users' profile information using multiple dimensions like users' privacy preferences, strength of relationship, sensitivity of information and visibility of the information.

Other approaches (e.g., [Spyromitros-Xioufis et al. 2016]) look at the images themselves to automatically generate privacy settings. In particular, a recent trend leverages deep learning techniques for image privacy prediction and privacy-aware image classification [Tonge and Caragea 2016; Yu et al. 2017]. For instance, Yu et al. [2017] use deep multi-task learning to automatically detect privacy-sensitive objects in images and recommend privacy settings for their protection. A number of approaches (e.g., [Zerr et al. 2012; Squicciarini et al. 2015; Squicciarini et al. 2017]) leverage both visual features of images and other information for privacy-aware image classification and the generation of privacy settings. For instance, Squicciarini et al. [2015] propose an adaptive policy prediction system based on association rule mining to generate personalized privacy settings based on social context and personal traits as well as on image content and metadata.

The aforementioned approaches, however, have only been applied to the automated generation of privacy settings from the perspective of individual users. Recently, some steps towards the generation of privacy settings for multi-party scenarios have been made by Fogues et al. [2017]. A main characteristic of this approach is the use of arguments about privacy settings for the generation of privacy settings besides other factors like social context, the sentiment associated with the information and the privacy settings themselves.

Feature	Squicciarini et al. [2010]	Fang and LeFevre [2010]	Klemperer et al. [2012]	Zerr et al. [2012]	Squicciarini et al. [2014]	Squicciarini et al. [2014]	Squicciarini et al. [2015]	Tonge and Caragea [2016]	Spyromitros-Xioufis et al. [2016]	Yu et al. [2017]	Squicciarini et al. [2017]	Fogues et al. [2017]
Social Context		✓			✓	✓	✓					✓
Personal Traits							✓					
Image Metadata	✓		✓	✓			✓	✓			✓	
Object Sensitivity					✓							✓
Sentiment											✓	✓
Image Content				✓			✓	✓	✓	✓	✓	
Privacy Preferences					✓							✓
Arguments												✓

Table VI: Overview of features used for policy generation

Most of these approaches, especially the ones based on machine learning, are only able to determine whether an image should be treated as a private information based on privacy labels assigned to similar images. However, this provides a too coarse-grained control for sharing in community-centered environments. Our analysis revealed that only a few works exploit the privacy preferences provided by users on previously updated objects. Recent studies have showed that deep learning provides a promising approach to generate fine-grained policies. However, existing solutions only consider a limited set of features, typically related to the image content. Other features, for instance concerning the social context, are needed to capture the complex user interplay in community-centered environments. In this respect, the arguments about the privacy preferences provided by users provide a promising feature to assess users' willingness to change their settings based on the privacy needs of their peers.

Interactive Interfaces for Policy Comprehension and Configuration. Other research efforts provide users with tools that support them in comprehending and configuring access control policies [Rode et al. 2006; Reeder et al. 2008; Schlegel et al. 2011; Anwar and Fong 2012]. The main difference among these tools is the type of visualization used to represent the policies. Rode et al. [2006] use circled pie charts (Figure 1a) to represent a shared user workspace where each slice represent a user, labeled dots represent the shared files and multiple concentric regions represent the permissions applied to each file (e.g., files in the center of the pie are readable, writable and available persistently). Similarly, Reeder et al. [2008] propose to use an Expandable Grids interface (Figure 1b) to set file permissions in Windows XP. The interface provides a matrix-based visualization of a policy, resembling an access matrix. In particular, the trees along the vertical axis represent the resources, the trees along the horizontal axis denote the principals and the entries at the intersection of the two trees indicate the access rights that each principal has on the resources. In particular, each entry is partitioned into five boxes, each of them representing an access right (i.e., read, write, execute, delete, administrate). Colors are used to indicate whether the access is permitted (green), denied (red) or some access is allowed (yellow). Schlegel et al. [2011] use the eyes metaphor (Figure 1c) to let users understand who has access to their personal information: the presence of a pair of eyes denotes that another user has gained access to the information, whereas the size of those eyes denotes the number of accesses to the information by that particular user. However, only the

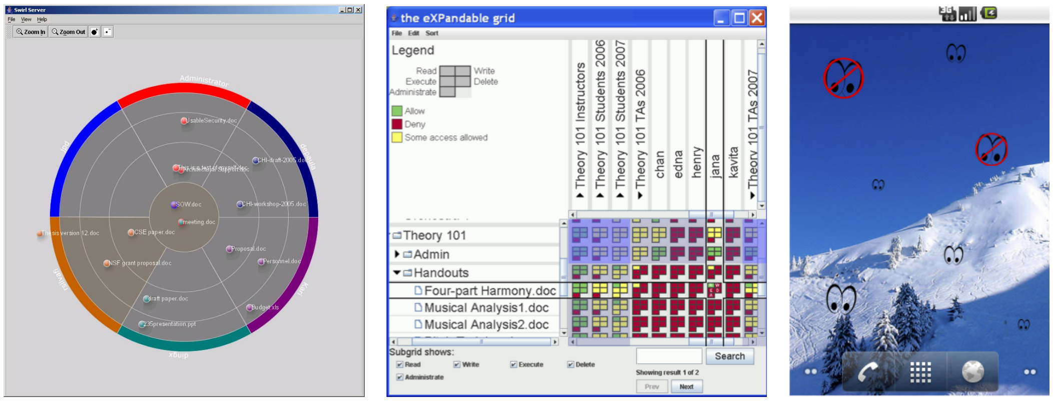

(a) Circled Pie Chart [Rode et al. 2006]

(b) Expandable Grids [Reeder et al. 2008]

(c) Eyes Metaphor [Schlegel et al. 2011]

Fig. 1: Interactive interfaces for Policy Comprehension

category (e.g., colleague, friend) to which the user belongs is revealed, while the identity of the user is kept private. Anwar and Fong [2012] adopt a mirror-looking metaphor where a profile owner is given a visual representation of its neighborhood. In particular, it can select a user and examine its profile from the perspective of the selected user. It is worth noting that solutions based on the mirror-looking metaphor have also been adopted by popular social network websites like Facebook and LinkedIn. For example, Facebook provides the “View As” tool that allows a user to see how its profile appears to the public, friends or a specific social network user.[3]

Other tools have been proposed to support users in configuring their social network privacy settings like PViz [Mazzia et al. 2012] and VeilMe [Wang et al. 2015]. PViz provides a graphical interface based on a bubble chart to visualize sharing settings in Facebook and provides support to generate social groups from public profile information automatically in order to reduce configuration efforts (Figure 2a). VeilMe allows users to explore their personality traits using a genome like bar sequence where each bar represents a trait and different trait types are denoted with different colors (Figure 2b). It also supports users in setting their privacy preferences based on their personality traits and uses a social distance visual metaphor to show how much personality traits would be disclosed with a specific social group, e.g. close colleague, distant colleague or public.

Among the tools above, the Expandable Grid and solutions based on the mirror-looking metaphor have potential in the context of community-centered collaborative access control systems. The Expandable Grid could be adapted and extended to give a global view of the policies specified by each stakeholder on a shared object and the interaction among them. The adoption of the mirror-looking metaphor can help users understand the effect of the privacy policies specified by other stakeholders on shared objects.

However, none of these tools has been deployed in the context of community-centered collaborative access control systems. The only notable initiative towards the development of a comprehensive framework for assisting users in the comprehension and configuration of privacy policies for shared objects is the SNGuard framework.[4] This framework comprises three tools aiming to enhance user comprehension and assist users in the configuration of their policies, namely MController, Retinue and Sigma. MController [Hu et al. 2013] is a voting-based tool for the collaborative management of shared resources (Figure 3a). MController allows the owner of a shared photo to select and configure the conflict resolution mechanism, possibly accounting for the sharing preferences of the other con-

[3] https://www.facebook.com/about/basics/what-others-see-about-you/profile/

[4] http://honeynet.asu.edu/snguard

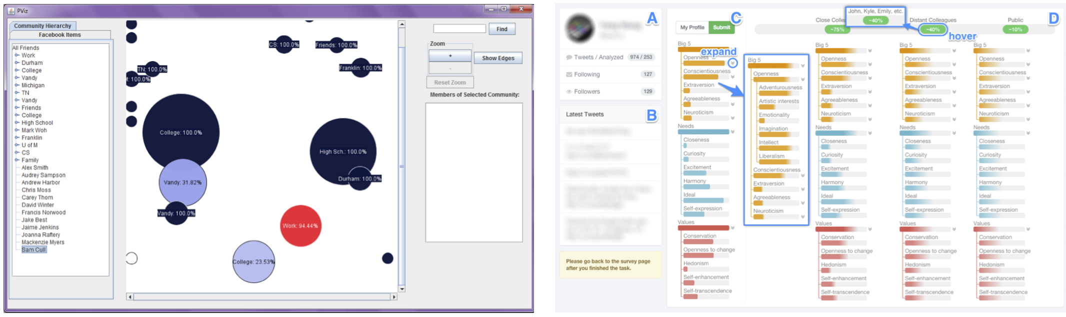

(a) PViz [Mazzia et al. 2012]
(b) VeilMe [Wang et al. 2015]

Fig. 2: Interactive interfaces for Privacy Policy Configuration

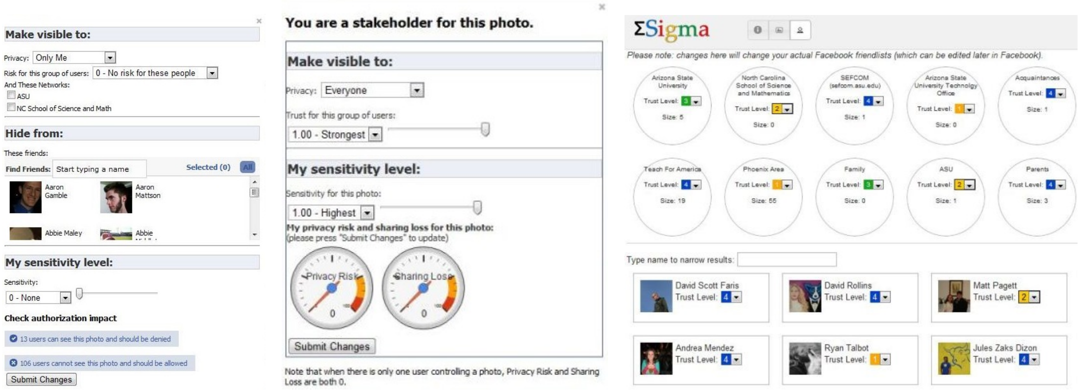

Make visible to:
Privacy: Only Me
Risk for this group of users: 0 - No risk for these people
And These Networks:
□ ASU
□ NC School of Science and Math
Hide from:
These friends:
Find Friends: Start typing a name Selected (0) All
Aaron Gamble
Aaron Mattson
Abbie Maley
Abbie
My sensitivity level:
Sensitivity: 0 - None
Check authorization impact
13 users can see this photo and should be denied
106 users cannot see this photo and should be allowed
Submit Changes

You are a stakeholder for this photo.
Make visible to:
Privacy: Everyone
Trust for this group of users: 1.00 - Strongest
My sensitivity level:
Sensitivity for this photo: 1.00 - Highest
My privacy risk and sharing loss for this photo:
(please press "Submit Changes" to update)
Privacy Risk 0
Sharing Loss 0
Submit Changes
Note that when there is only one user controlling a photo, Privacy Risk and Sharing Loss are both 0.

ΣSigma
Please note: changes here will change your actual Facebook friendlists (which can be edited later in Facebook).
Arizona State University Trust Level: 3 Size: 5
North Carolina School of Science and Mathematics Trust Level: 2 Size: 0
SEFCOM (sefcom.asu.edu) Trust Level: 4 Size: 1
Arizona State University Technology Office Trust Level: 1 Size: 0
Acquaintances Trust Level: 4 Size: 1
Teach For America Trust Level: 4 Size: 19
Phoenix Area Trust Level: 1 Size: 55
Family Trust Level: 3 Size: 0
ASU Trust Level: 2 Size: 1
Parents Trust Level: 4 Size: 3
Type name to narrow results:
David Scott Faris Trust Level: 4
David Rollins Trust Level: 4
Matt Pagett Trust Level: 2
Andrea Mendez Trust Level: 4
Ryan Talbot Trust Level: 1
Jules Zaks Dizon Trust Level: 4

(a) MControl
(b) Retinue
(c) Sigma

Fig. 3: SNGuard Framework for Privacy Policy Comprehension and Configuration (http://honeynet.asu.edu/snguard)

trollers of the photo. Retinue [Hu et al. 2011] is a risk-based collaborative data sharing mechanism that allows a systematic detection and resolution of multi-party conflicts in online social networks (Figure 3b). Retinue allows the controllers of a shared photo to set the sensitivity of the photo and assess privacy risk and sharing loss based on the defined sensitivity and the preferences given by other controllers of the photo. Sigma [Hu et al. 2012] is a circle-based approach to enable collaborative control of photos based on Google+ circles (Figure 3c). Sigma enables users to group their friends in circles and to specify the trust level for single friends or for circles.

Feedback Generation. A growing body of research focuses on promoting user awareness by providing users with feedback about access decision making. Bonatti et al. [2001] propose an access control system for web services that gives feedback about a policy protecting a resource: when access to a resource cannot be granted to a user, further actions, e.g. signing an agreement that would result in the required access, are communicated to the user. However, they do not consider the problem of protecting access control policies that, if disclosed during the feedback process, can provide insights on the policies themselves both to intruders and legitimate users. Intruders can use this information to find out how and where to direct their attacks; also, legitimate users can deduce what other users can access, thus violating the privacy of those other users.

Other works focused on providing feedback to access requesters while limiting the disclosure of access control policies [Kapadia et al. 2004; Ghai et al. 2010]. Kapadia et al. [2004] propose KNOW, a system for providing feedback to access requesters when access is denied while limiting the amount of information revealed about the applied access control policies. This is achieved by protecting the disclosure of access control policies with meta-policies and providing the requester possible alternative ways of fulfilling the access control policies that do not violate the meta-policies. Similarly to Kapadia et al., Ghai et al. [2010] propose a system called Cue to generate feedback explaining why an access request is denied while protecting the disclosure of access control policies. However, there are two main differences between these approaches: KNOW generates feedback for RBAC models, which mainly suggests a user to change role or wait for an event to occur. On the other hand, Cue is based on an ABAC model and exploits the attributes provided in an access request to provide feedback on which conditions were not met by the access requester. Another difference is that Cue provides feedback not only when an access request is denied but also when no policy applies to the request.

Some steps in the direction of providing feedback on access decision making in collaborative systems have been done in [Damen et al. 2014; Hu et al. 2013; Mahmudlu et al. 2016; den Hartog and Zannone 2016]. Hu et al. [2013] provide the controllers of shared objects with feedback information to assess the effect of multi-party conflicts and, thus, evaluate the impact of collaborative authorizations within a social network site. In particular, a controller can visualize which users are allowed to see a photo but access should be denied and which users are not allowed to see a photo but access should be granted according to the controller's privacy settings. Damen et al. [2014] propose a transparency service for collaborative access control that notifies a user (at policy evaluation time) when the decision enforced by the system differs from the authorization requirements specified by the user for the requested shared object. Mahmudlu et al. [2016] extend the work in [Damen et al. 2014] by showing how XACML-based mechanisms can be augmented with transparency. The authors investigate two possible deployments of the transparency service within an XACML implementation, namely either as part of the Policy Enforcement Point (PEP) or as part of the Policy Decision Point (PDP). A recent work [den Hartog and Zannone 2016] makes a step further by providing an approach for generating justifications of why a controller's decision was overridden. The feedback is provided at a level of detail that takes into account the relationship of the controller with the resource as well as the visibility preferences of the other controllers. This approach is the most promising to promote transparency in collaborative systems because not only notifies users when the decision enforced on a shared resource is different from their own, but it also explains why a certain decision has been taken on the basis of the (visible) policies specified by other users on the shared resource.

6. EVALUATION METHODS

Evaluation should be an essential component of the design and development of an access control system. In general, evaluation can be *formative*, to drive the design process of the access control system, e.g. it could be used during pre-design to understand users' privacy requirements and specify the access control policy accordingly; or it can be *summative* to check the access control system's effectiveness and correctness or to identify usability problems. The evaluation of access control systems is usually summative: it is carried out after the design of the access control system has been completed. The main evaluation methods to evaluate access control systems are listed below:

- *Controlled Experiment*. This type of study is carried out with representative users in a laboratory environment where users conduct access control tasks defined by the designers of the experiment. The main goal of this type of study is to assess the effect of design elements or features of the proposed access control model or system.
- *Usability Study*. Similarly to a controlled experiment, this type of study involves observing users while they are performing specific access control tasks. However, the main difference lies in the

Evaluation Method	**Feature studied**	**Source**
Controlled Experiment	Usefulness, Understanding, Fairness of Collaborative Privacy Settings	[Squicciarini et al. 2010]
	Balance between privacy protection and data sharing	[Hu et al. 2012]
	Likability, Simplicity	[Hu et al. 2011; Hu et al. 2013]
	Control	[Hu et al. 2011; Hu et al. 2013; Vishwamitra et al. 2017]
	Adoption Willingness	[Vishwamitra et al. 2017]
	% of Matched Concession Behavior	[Such and Criado 2016]
	Peer Pressure Effect, Consensus-building Approximation	[Rajtmajer et al. 2016]
Performance Evaluation	Collective Privacy Setting Execution Time	[Squicciarini et al. 2010]
	Policy Generation Time	[Squicciarini et al. 2014]
	Policy Integration Time	[Mazzoleni et al. 2008]
	Policy Evaluation Time	[Damen et al. 2014; Hu et al. 2013; Mahmudlu et al. 2016; Mazzoleni et al. 2008; Vishwamitra et al. 2017]
	Certificate Path Discovery Time, Assertion Generation Time, Proof Generation Time	[Carminati et al. 2009]
	Encryption and Decryption Time, Secret Share Creation Time, Secret Reconstruction Time	[Ilia et al. 2017]
	Policy Satisfiability Execution Time	[Mehregan and Fong 2016]
	Privacy Risk, Sharing Loss	[Hu et al. 2011]
	Effectiveness	[Hu et al. 2012; Vishwamitra et al. 2017]
	Consensus-building Convergence, Consensus-building Approximation	[Rajtmajer et al. 2016]
	Amount of Feedback	[Damen et al. 2014]

Table VII: Evaluation methods applied to access control systems tailored to community-centered environments

effect that is assessed: in controlled experiments it can be the cost, effectiveness, correctness of policies, etc., while in this type of study the focus is placed only on usability.

- *Performance Evaluation*. It is an automated or computer-generated analysis of an access control system with respect to load or response times under particular use conditions, e.g. number of policies, number of users.

Most of the access control systems we have studied have been demonstrated using one or more usage scenarios (e.g., [Carminati and Ferrari 2011; Ahn et al. 2012]), especially in the social network setting. Usage cases are, however, hard to generalize and often do not provide insights on the effectiveness of the approach in realistic settings. Only controlled experiments, usability studies and performance evaluation can provide such insights. Nonetheless, only a few access control systems tailored to community-centered environments have been evaluated using controlled experiments and/or performance evaluation [Hu et al. 2013; Squicciarini et al. 2010; Such and Criado 2016]; we are not aware of any usability study for these systems. Table VII summarizes how proposed access control systems tailored to community-centered environments have been evaluated using controlled experiments and performance evaluation. The analyzed studies are classified based on two main criteria: *feature studied*, i.e. the properties of the access control system investigated in the evaluation (e.g., effectiveness, usefulness), and the *evaluation method*, i.e. performance evaluation or controlled experiment.

Controlled Experiment. To guarantee that a controlled experiment is correctly executed and allows drawing valid conclusions, it is important to follow a formal process. Typically, this process consists of the following steps as suggested by Wohlin et al. [2000]: *scoping*, *planning*, *operation*, *analysis and interpretation*, and *presentation and package*. During the scoping activity, the goal of the experiment is defined in terms of *object of study* (what is studied), the *purpose* (what is the intention of the study), which is the *quality focus* (what is the effect being studied), the *perspective* of who and the *context* (where the study is going to be conducted). The planning phase is the core step in conducting the experiment. First, the type of participants and the environment where the experiment is going to be run are determined. Then, the hypothesis of the experiment is stated formally and the variables to be measured are selected. The experiment design is also selected, e.g. randomization of subjects. The measurement procedures, the objects and possible guidelines are defined. The threats to internal, external, construct and conclusion validity of the experiment are also considered in this phase. Internal validity is related to the reliability of the results; external validity is related to the ability of generalizing the findings beyond the context of the experiment; construct validity is related to the generalization of the result to the theory behind the experiment; conclusion validity is concerned with coming to the correct conclusions about the relations between the treatment (the object of study) and the outcome of the experiment. The operation phase is where the experiment is conducted and the data are collected. The analysis and interpretation is where the data are analyzed to decide whether the hypothesis should be rejected or accepted. The presentation and package phase concerns the documentation of how the experiment was designed and conducted along with the presentation of the findings of the experiment.

Most of the controlled experiments reported in Table VII followed only partially the steps of the above process and therefore their findings may not be valid. They only focused on some of the activities required by the planning step: the selection of the subjects, environment and objects and the identification of the variables to be measured. However, no discussion of the experimental design and of the threat to validity is typically reported. For example, Squicciarini et al. [2010] conducted a controlled experiment to study users' understanding of co-ownership, usefulness and fairness of collaborative privacy settings provided by their approach for collective privacy management on shared content. Participants had to watch a video illustrating the concept of co-ownership, the collective privacy management approach and three scenarios illustrating the features provided by the approach. Participants were then requested to fill in a questionnaire to measure their level of understanding of co-ownership, usefulness and fairness. The design of this experiment, however, has one major threat to internal validity because participants did not have any direct experience with the approach.

A different experimental design that does not suffer of internal validity has been adopted in [Hu et al. 2011; Hu et al. 2013; Such and Criado 2016; Rajtmajer et al. 2016; Vishwamitra et al. 2017]. In these experiments, the subjects performed a task using the proposed access control system. For instance, Hu and colleagues evaluated the likeability, simplicity and control of two of their prototypes, *Retinue* [Hu et al. 2011] and *MController* [Hu et al. 2013]. Likeability measures a user's satisfaction with the system, simplicity measures the intuitiveness and usefulness of the system, and control measures the perceived control users have over their data. The participants of the experiment were asked to complete a set of tasks using the proposed systems, and later complete a post-session questionnaire on likeability, simplicity and control. Such and Criado [2016] took a more thorough approach to evaluate the accuracy of their conflict resolution method. Specifically, authors asked participants to first specify their privacy preferences on shared photos and then automatically generated preferences conflicting with the ones specified by the users. Based on these conflicts, Such and Criado asked the participants if they were willing to concede and, thus, change their setting to solve conflicts with the other people depicted in the photo. Accuracy was measured as the percentage by which their approach matched participants' concession behavior. In similar spirit, Rajtmajer et al. [2016] measured users' response to peers' privacy behavior, as compared to their theoretical model. The experiment was carried out in a controlled setting and based on a simulated network, wherein participants participate in group social network activities and make privacy setting decisions accordingly. The above experiments may not suffer of internal validity but they are affected by construct validity since the hypotheses of the experiment have not been specified.

Performance Evaluation. To date, there does not exist a standard evaluation framework for access control; every proposal has been evaluated against different criteria and with respect different parameters, thus making it difficult to compare existing systems. Performance measures range from the time needed for policy specification/integration and conflict resolution to the success rate in meeting users' expectations. Mazzoleni et al. [2008] evaluated the time required by their XACML-based mechanism to integrate access control policies for a varying number of policies to be integrated and size. More recently, Squicciarini et al. [2010] evaluated the scalability of their collective policy conflict resolution mechanism based on auctions with respect to the number of co-owners that appear in a photo. Similarly, Hu et al. [2013] assessed the performance of MController by varying the number of controllers of a shared photo. Ilia et al. [2017] assessed the overhead due to the cryptographic operations introduced in their multi-party access control model by varying the number of shares and sensitivity of objects.

Damen et al. [2014] evaluated their transparency service by measuring how policy complexity and social network density impact the amount of feedback sent to users. A similar approach has been adopted by Mahmudlu et al. [2016] to assess the performance of their transparency service. In this case, authors assessed the overhead introduced by the transparency service in the policy evaluation process. Squicciarini et al. [2014] assessed the time needed to semi-automatically generate access rules varying two parameters: the number of traits in the user profile and the number of friends. To this end, they have developed a prototype in a content-management system (i.e., Drupal) and executed it focusing one user's profile. In a first experiment, they increased the number of the traits in the user profile according to a uniform distribution over the possible types of traits, e.g. interpersonal relationships and comments. In a second experiment, they increased the number of user's friends while leaving the other traits equal. Other works focused on evaluating performance with respect to other system parameters. For example, Mehregan and Fong [2016] evaluated the execution time of their algorithm for verifying policy satisfiability in multi-ownership settings. Hu et al. [Hu et al. 2011; Hu et al. 2012] assessed their strategy for privacy conflict resolution with respect to two metrics, privacy risk and sharing loss.

Although effective in measuring specific features and algorithms of their own proposals, the performance evaluations lack generalizability. All of them suffer from a threat to external validity because they are conducted using synthetic ad-hoc generated policies rather than on real-world access control policies. As a consequence, the set of policies on which they are tested is always different and this hinders the comparison of the performance of the various solutions.

7. RESEARCH CHALLENGES

In this section, we first give a summary of the literature review with respect to the requirements presented in Section 2; then, based on this analysis, we identify open challenges for access control systems tailored to community-centered systems and delineate a roadmap for future work.

7.1. Summary

Table VIII presents a summary of our literature review with respect to the requirements presented in Section 2 and the evaluation methods discussed in Section 6. In the table, 'full support' (✔) is used to indicate that an access control system (or the underlying access control model for requirements **R1** to **R4**) fully satisfies a given requirement or a comprehensive evaluation has been performed using a given evaluation method; 'partial support' (❇) is used to indicate that an access control system (or the underlying access control model) addresses a given requirement but the requirement is not fully satisfied by the system or that only certain features of the access control system have been evaluated using a given evaluation method; 'minimal support' (†) is used to indicate that an access control system (or the underlying access control model) only provides some basic features to address a given requirement; and 'no support' (✘) is used to indicate that a requirement is not

	Domain	**Policy Specification**		**Governance**		**Usability & Transparency**			**Evaluation Method**	
		R1	**R2**	**R3**	**R4**	**R5**	**R6**	**R7**	**CE**	**PE**
Mazzoleni et al. [2008]	Gen	✘	❇	❇	†	❇	✘	✘	✘	❇
Fong et al. [2009]	Gen	✔	✘	✘	✘	✘	✘	✘	✘	✘
Carminati et al. [2009]	Gen	✔	✘	✘	✘	✘	✘	✘	✘	❇
Wishart et al. [2010]	Gen	✘	✘	❇	†	✘	✘	✘	✘	✘
Squicciarini et al. [2010]	Gen	❇	✘	❇	†	✔	✘	✘	❇	❇
Hu et al. [2011]	SN	❇	✘	❇	†	❇	✔	❇	❇	❇
Carminati and Ferrari [2011]	Gen	✔	✘	❇	†	❇	✘	✘	✘	✘
Fong [2011]	SN	❇	❇	✘	✘	✘	✘	✘	✘	✘
Krishnan et al. [2011]	Gen	❇	❇	✘	✘	✘	✘	✘	✘	✘
Hu et al. [2012]	SN	❇	✘	❇	†	❇	❇	✘	❇	❇
Xiao and Tan [2012]	SN	❇	✘	❇	†	❇	✘	✘	✘	✘
Bruns et al. [2012]	Gen	✔	✘	✘	✘	✘	✘	✘	✘	✘
Hu et al. [2013]	SN	❇	✘	†	†	❇	✔	❇	❇	❇
Fong et al. [2013]	Gen	❇	❇	✘	✘	✘	✘	✘	✘	✘
Damen et al. [2014]	Gen	✘	✘	✔	❇	❇	✘	†	✘	❇
Squicciarini et al. [2014]	Gen	❇	✘	✘	✘	❇	✘	✘	✘	❇
Crampton and Sellwood [2014]	Gen	❇	✘	✘	✘	✘	✘	✘	✘	✘
Mahmudlu et al. [2016]	Gen	✘	❇	✔	❇	❇	✘	†	✘	❇
Mehregan and Fong [2016]	Gen	❇	✘	❇	†	✘	✘	✘	✘	❇
Rajtmajer et al. [2016]	Gen	†	✘	†	†	❇	✘	✘	❇	❇
Such and Criado [2016]	SN	✘	✘	❇	†	❇	✘	✘	❇	✘
den Hartog and Zannone [2016]	Gen	✘	❇	✔	❇	❇	✘	❇	✘	✘
Ilia et al. [2017]	Gen	❇	✘	❇	❇	❇	✘	✘	✘	❇
Vishwamitra et al. [2017]	SN	❇	✘	❇	❇	❇	✘	✘	❇	❇

Legend

✔:	full support	❇:	partial support	†:	minimal support	✘:	no support
Gen:	General	SN:	Social Network				

Table VIII: Evaluation of existing access control models and systems against requirements highlighted in Section 2.2

addressed by the system (or the underlying access control model) or that the system has not been evaluated using a given evaluation method. In the table, we also indicate whether an access control system is general ('*Gen*') or is specific for a certain application domain. Here, '*SN*' is selected only if the proposed system was casted in social networks *and* suitable for a few (or even a single) specific social network. We note that this is particularly the case for access control systems that are casted on Facebook-like social networks, and build on the specific privacy offerings proposed by the Facebook platform. We mark an access control system as '*Gen*' if it can be applied outside the social network domain, regardless of where its main examples are.

Requirements **R1** and **R2** concern policy specification and, in particular, the ability of an access control model to capture users' interplay and dynamics characterizing community-centered systems. For the evaluation of requirement **R1**, we marked 'full support' for those models that allow the specification of (complex) topology-based policies. Access control models that only consider some aspects of interpersonal relationships like relationship type, depth and strength, are marked with 'partial support' whereas models that only account for the existence of interpersonal relationships between users are marked with 'minimal support'. Models that do not consider interpersonal relationships are marked with 'no support'. Requirement **R2** is considered fully satisfied ('full support') by access control models that are able to capture the dynamics and context of collaboration and of the environment in which access decisions have to be made. Models that allow the specification and reasoning over either general context information like XACML [XACML v2.0 2005; XACML v3.0 2013] (and approaches based on those languages) or the context of interpersonal relationships are marked with ‘partial support’. Access control models that do not consider context information in access decision making are marked with ‘no support’. Looking at Table VIII we can observe that several efforts have been devoted to accommodate interpersonal relationships into access decision making and a number of proposals already allow the authoring of complex topology-based access control policies. On the contrary, very little work has attempted to account for context information, especially concerning community-centered collaborative systems. This lack is mainly due to the fact that most proposals only provide a simplified and ad-hoc language to demonstrate a limited set of features and concepts, which however is often not able to capture the full complexity and subtleties of community-centered systems.

Requirements **R3** and **R4** relate to the governance of shared resources. We consider requirement **R3** fully satisfied by access control models that are able to capture the full complexity of the governance of shared objects, i.e. by approaches that allow accounting for the relation between users and objects (asymmetric governance models). Access control models based on a multi-ownership governance model are marked with ‘partial support’. Models that support features for the governance of shared resources but ultimately rely on a single-ownership governance model (e.g., [Hu et al. 2013]) are marked with ‘minimal support’. Models based on a purely single-ownership governance model are marked with ‘no support’. Regarding conflict resolution (**R4**), it is desirable for an access control model to provide a flexible and expressive way to reconcile conflicting authorization requirements. Thus, we mark models that allow the arbitrary combination of predefined and multi-party decision strategies with ‘full support’. Models that allow the arbitrary combination of strategies but support only either predefined strategies or strategies based on mutual agreement, are marked with ‘partial support’. ‘Limited support’ is used to indicate models that allow the use of only one strategy to solve all possible conflicts arising from stakeholders’ requirements (either predefined or strategies based on mutual agreement). Models that do not address conflict resolution are marked with ‘no support’. We note that most of the existing access control models tailored to community-centered environments are based on a multi-ownership model, and only recently it has been recognized that the level of authority that users have over a shared resource depends on their relation with the resource. Moreover, our literature review has shown that the problem of policy conflict is largely recognized and several conflict resolution methods have been proposed over the years. Nonetheless, these methods are often ad-hoc and based on a single conflict resolution strategy. Only a few proposals attempt to provide a flexible framework that allows the customization and combination of various policy conflict resolution strategies.

The last set of requirements (**R5**, **R6** and **R7**) concerns the usability and the level of transparency offered by access control systems. Requirement **R5** states that access control systems should be unobtrusive and should not impose additional overhead on users. In the evaluation of this requirement, we consider the efforts that users have to spend for conflict resolution as well as the efforts needed for policy authoring and configuration. Access control systems in which all these aspects have been addressed are marked with ‘full support’. Access control systems that address only some of them are marked with ‘partial support’. Requirement **R6** is considered fully satisfied by access control systems that provide interfaces for both policy comprehension and configuration. If only one aspect is supported, a system is marked with ‘partial support’ and, if no interfaces are provided, the system is marked with ‘no support’. To enhance user awareness for conscious decision making (**R7**), several types of feedback can be provided to users, for instance to explain the effect of a certain decision [Hu et al. 2011; Hu et al. 2013] or to justify the decision made [den Hartog and Zannone 2016]. We mark an access control system with ‘full support’ if it provides comprehensive feedback about decision making. If only some type of feedback is provided, a system is marked with ‘partial support’. We mark with ‘minimal support’ systems that notify users about policy mismatches (i.e., the situation in which the final access decision differ from a user’s preferences [Damen et al. 2014]) and with ‘no support’ if transparency is not addressed by the system. It is worth noting that we did not identify any access control system that fully meets all usability and transparency requirements. In particular, none of the reviewed contributions fully supports **R7**, and only few of them offer even partial support. We attribute this gap to the level of maturity of the theoretical models and the underlying algorithms and protocols. Clearly, transparency and support are important for fully deployed access control systems that are implemented in practice. Most of the systems we reviewed aims to demonstrate foundational and ad-hoc features, and not yet brought to deployment stage.

We conclude our study with some observations concerning the methods used for the evaluation of existing access control systems. The last two columns of Table VIII show the evaluation method(s) used for each studied access control system, namely controlled experiments (**CE**) and performance evaluation (**PE**). For both evaluation methods, we mark an access control system with 'full support' if a comprehensive evaluation has been performed. We mark a system with 'partial support' if only some aspects have been evaluated and with 'no support' if no evaluation has been carried out. First, we note that empirical studies are scarce (and often ad-hoc) among the systems being studied here, showing that usability of access control systems (at least for community-centered collaborative systems) is still understudied. In particular, we observed that the lack of empirical studies is more prominent for work proposing models for the specification of ReBAC policies (e.g., [Crampton and Sellwood 2014; Fong et al. 2009; Fong 2011; Fong and Siahaan 2011]). These models are typically demonstrated through usage scenarios, sometimes coupled with formal proofs of the tractability and expressiveness of the proposed model (e.g., [Fong and Siahaan 2011]); however, no studies were carried out to establish the usability of those models. Our study has also revealed that, despite the number of well-established methodologies for the design and evaluation of experiments available in the literature (e.g., [Wohlin et al. 2000]), these methodologies have not been properly applied to the design of controlled experiments and, thus, the obtained findings are not generalizable. Possibly explaining the limitations of above, we note the lack of available datasets of access control policies used in real-life community-centered collaborative systems. As a consequence, many researchers created their own datasets of policies to validate their system. Moreover, most systems have been evaluated against various criteria. We believe that, without a comprehensive validation, it is not possible to demonstrate the applicability of research results in real settings, thus hindering their transition to practice.

7.2. Roadmap

The field of access control has a long history of exciting research findings and development efforts. Many access control models have been proposed to date, some of which have gained traction in real-world systems (e.g., RBAC) and others have remained primarily as theoretical models. Whether access control solutions for community-centered systems can be effectively deployed in real-world environments is still an open question. Despite a growing interest in access control for those systems, to this date, only very simple forms of group-based access control have been deployed in existing community-centered systems like online social networks (e.g., Facebook, LinkedIn). Transition to practice has been hindered by the lack of maturity of these models. To date, many important issues must be addressed before this transition can occur successfully. Below we summarize some of the main open research issues.

— *Lack of comprehensive solutions.* Our study has revealed two main streams of research on access control for community-centered environments. On the one side, we identified several relationship-based access control models that study how to incorporate interpersonal relationships in access decision making. These approaches, however, often assume that resources are owned by a single entity and do not account for governance of resources shared between multiple users. On the other side, we can find studies that investigate the governance of shared resources. These studies, however, only minimally consider existing relationships among users. Only a very few proposals (e.g., [Ilia et al. 2017; Mehregan and Fong 2016]) attempt to reconcile these two views, for instance by studying how interpersonal relationships influence policy composition and conflict resolution. More efforts are required to reconcile these two streams and design access control models able to handle the full complexity of community-centered systems.

— *Lack of standardized solutions.* Many proposals provide ad-hoc languages for the authoring of access control policies for collaborative systems. These languages, however, mainly aim to demonstrate the suitability of (a limited set of) features/concepts to enable access control in community-centered systems. Only a few proposals attempt to integrate those concepts into standardized languages like XACML (e.g., [Mazzoleni et al. 2008; Mahmudlu et al. 2016]). We argue that a successful adoption of access control models in real-life community-centered systems can only be achieved if research efforts are complemented with standardization efforts in order to increase the level of maturity of existing solutions.

- *User-centric.* Ultimately, access control systems for community-centered systems should strike a balance between human-supported work and automation. While users' involvement is necessary to specify access control preferences, daunting, multi-round negotiations are likely to fail and result in low adoption. Moreover, one may argue that lay users often do not have the necessary information and knowledge-set to assess what best policies and concessions apply for a given resource. Accordingly, researchers are still trying to identify the best interaction protocol between the users and the system. This should be efficient and effective: users' input should be consistent with their preferences and easy to specify, but most of the remaining effort (policy reconciliation, conflict resolution and enforcement) should be driven by the system with limited intervention (ideally no intervention) required from the users.
- *Generalizability* beyond online social network domain. To this date, most of the work related to multi-party access control has been designed in the context of web-based social networks. While this is an important application domain, a generic model should be adaptable to a number of application domains, outside the "Facebook-like" model. We believe that most of the existing systems are general and can be potentially applied to other application domains. However, so far there is little evidence of their generalizability as they are manly demonstrated through usage scenarios related to social network. Some important challenges arise when one wants to apply them in a different domain. First, how to specify social graphs underlying multi-party models? And how could these models be applicable when ad-hoc access control languages are not available? Most of the work in the field of collaborative access control so far has focused on discretionary policies, with varying syntax. It is desirable to design access control models for community-centered systems that offer security guarantees (e.g., safety, non-interference) provided from earlier access control models. Extending these models would offer guarantees already provided naturally by traditional access control models, and possibly offer more generic solutions yet supporting the desired multi-user capabilities.
- *Enforcement of multi-party settings.* Existing models have put most of their attention to policy specification and data governance. The enforcement mechanisms have often been assumed easy to deploy and consistent with conventional access models (a single policy enforcement point and centralized access database). We argue that as multi-party models adapt to various domains and infrastructures, related enforcement mechanisms should be revisited and adjusted to support possible asynchronous multi-user input.
- *Interactive User Interface.* Most of the proposed systems to date, attempt to assist users in making possibly complex access control decisions. Yet, usability of these systems is still mostly unexplored. To date, very few works provide interactive interfaces that facilitate users in setting their access control policies on shared content and allow users to understand the policies specified by other stakeholders. We believe that such interfaces are a desirable feature for lay users, especially in domains that require a thorough understanding of access control policies and their effect. In particular, their design and development can enhance user awareness of the risks of data sharing and empower users with the control over their data, thus facilitating the adoption in practice of collaborative access control systems.
- *Transparency.* As noted in the previous section, only a very few systems support transparency in multi-party access control (and in general, in access control for collaborative systems). Existing systems provide either limited insights into the decisions made [Mahmudlu et al. 2016] or a partial view of access decision making, e.g. the effect of decisions [Hu et al. 2013] or why a decision was made [den Hartog and Zannone 2016]. We believe that more research efforts are needed to provide users with an understanding of access decision making in collaborative systems

and of the effects of multi-party access decisions. Without this understanding, users may feel their resources are not properly protected and, thus, be reluctant to contribute to a community.

— *Empirical evaluation.* Empirical evaluation of access control systems for community-centered environment and access control systems in general, is often an underestimated aspect of access control research and development. This is mainly due to the vastness and diversity of evaluation methodologies, which makes it difficult for researchers and practitioners to find the most appropriate methods to achieve their evaluation goals. Another aspect is the lack of literature guidelines on how to evaluate access control systems. Therefore, there is the need of an evaluation framework that allows researchers to answer several questions: Which qualitative and quantitative aspects should be assessed? Which evaluation method should applied to assess these aspects? Which data should be collected? Last but not least, there is the need of available datasets of access control policies used in real-life community-centered systems, which will serve as a benchmark for the evaluation and comparison of the proposed access control models and mechanisms.

8. CONCLUSION

This paper provides a thorough analysis of existing access control models and systems tailored to community-centered collaborative systems. Our goal was to drive research and development of novel access control solutions suitable for community-centered environments. We found that despite some growing interest in the area, much work is left to be done.

We identified several important requirements deemed desirable to support access control in these emerging systems. These span from policy specification (e.g., the ability to capture domain needs and characteristics and support for contextual policies) to governance (e.g., collaborative administration of shared resources). By matching these requirements with the current state-of-the-art, we found a plethora of models and systems that meet some (but never all) such requirements. In particular, our analysis of the literature has identified key gaps in existing research in access control for community-centered systems and helped in the development of a roadmap for future research. Such gaps are particularly prevalent in the usability realm, as access control systems for community-centered systems have not yet been widely deployed and tested. We believe this is one of the main obstacles to the adoption of existing solutions in real-world and commercial community-centered collaborative systems and platforms.

ACKNOWLEDGMENTS

The authors would like to thank Barbara Carminati and the reviewers for helpful comments and suggestions on earlier versions of this paper.

REFERENCES

Gail-Joon Ahn, Jing Jin, and Mohamed Shehab. 2012. Policy-driven Role-based Access Management for Ad-hoc Collaboration. *J. Comput. Secur.* 20, 2-3 (2012), 223–257. http://dl.acm.org/citation.cfm?id=2590720.2590724

Evangelos Aktoudianakis, Jason Crampton, Steve Schneider, Helen Treharne, and Adrian Waller. 2013. Policy templates for relationship-based access control. In *Proceedings of Annual International Conference on Privacy, Security and Trust*. IEEE, 221–228. DOI: http://dx.doi.org/10.1109/PST.2013.6596057

Mohd Anwar and Philip W. L. Fong. 2012. A Visualization Tool for Evaluating Access Control Policies in Facebook-style Social Network Systems. In *Proceedings of the 27th Annual ACM Symposium on Applied Computing (SAC '12)*. ACM, New York, NY, USA, 1443–1450. DOI: http://dx.doi.org/10.1145/2245276.2232007

Paul Ashley, Satoshi Hada, Günter Karjoth, Calvin Powers, and Matthias Schunter. 2003. Enterprise Privacy Authorization Language (EPAL 1.2). (2003).

Yousra Asim and Ahmad Kamran Malik. 2016. A Survey on Access Control Techniques for Social Networks. In *Innovative Solutions for Access Control Management*. IGI Global, Hershey, PA, 1–32. DOI: http://dx.doi.org/10.4018/978-1-5225-0448-1.ch001

Vijayalakshmi Atluri and Janice Warner. 2004. Automatic Enforcement of Access Control Policies Among Dynamic Coalitions. In *Proceedings of the 1st International Conference on Distributed Computing and Internet Technology (ICDCIT'04)*. Springer-Verlag, Berlin, Heidelberg, 369–378. DOI: http://dx.doi.org/10.1007/978-3-540-30555-2_43

Georgia Bafoutsou and Gregoris Mentzas. 2002. Review and functional classification of collaborative systems. *International Journal of Information Management* 22, 4 (2002), 281–305. DOI: http://dx.doi.org/10.1016/S0268-4012(02)00013-0

Leila Bahri, Barbara Carminati, and Elena Ferrari. 2015. CARDS-Collaborative Audit and Report Data Sharing for A-Posteriori Access Control in DOSNs. In *Proceedings of Conference on Collaboration and Internet Computing*. IEEE, 36–45. DOI: http://dx.doi.org/10.1109/CIC.2015.18

Sasha A. Barab, Rob Kling, and James H. Gray. 2004. *Designing for Virtual Communities in the Service of Learning*. Cambridge University Press.

Elisa Bertino, Barbara Catania, Elena Ferrari, and Paolo Perlasca. 2003. A Logical Framework for Reasoning About Access Control Models. *ACM Trans. Inf. Syst. Secur.* 6, 1 (2003), 71–127. DOI: http://dx.doi.org/10.1145/605434.605437

Rafae Bhatti, Elisa Bertino, and Arif Ghafoor. 2005. A Trust-Based Context-Aware Access Control Model for Web-Services. *Distributed and Parallel Databases* 18, 1 (2005), 83–105. DOI: http://dx.doi.org/10.1007/s10619-005-1075-7

Piero A. Bonatti, Ernesto Damiani, Sabrina de Capitani, and Pierangela Samarati. 2001. A Component-Based Architecture for Secure Data Publication. In *Proceedings of the 17th Annual Computer Security Applications Conference (ACSAC '01)*. IEEE Computer Society, Washington, DC, USA, 309–318. DOI: http://dx.doi.org/10.1109/ACSAC.2001.991546

Glenn Bruns, Philip W.L. Fong, Ida Siahaan, and Michael Huth. 2012. Relationship-based Access Control: Its Expression and Enforcement Through Hybrid Logic. In *Proceedings of the 2nd ACM Conference on Data and Application Security and Privacy*. ACM, New York, NY, USA, 117–124. DOI: http://dx.doi.org/10.1145/2133601.2133616

Glenn Bruns and Michael Huth. 2011. Access Control via Belnap Logic: Intuitive, Expressive, and Analyzable Policy Composition. *ACM Trans. Inf. Syst. Secur.* 14, 1, Article 9 (2011), 27 pages. DOI: http://dx.doi.org/10.1145/1952982.1952991

Adrian Bullock and Steve Benford. 1997. Access Control in Virtual Environments. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '97)*. ACM, New York, NY, USA, 29–35. DOI: http://dx.doi.org/10.1145/261135.261142

Adrian Bullock and Steve Benford. 1999. An Access Control Framework for Multi-user Collaborative Environments. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work (GROUP '99)*. ACM, New York, NY, USA, 140–149. DOI: http://dx.doi.org/10.1145/320297.320313

Xiang Cao and Lee Iverson. 2006. Intentional Access Management: Making Access Control Usable for End-users. In *Proceedings of the 2nd Symposium on Usable Privacy and Security (SOUPS '06)*. ACM, New York, NY, USA, 20–31. DOI: http://dx.doi.org/10.1145/1143120.1143124

Barbara Carminati and Elena Ferrari. 2008. Access control and privacy in web-based social networks. *International Journal of Web Information Systems* 4, 4 (2008), 395–415. DOI: http://dx.doi.org/10.1108/17440080810919468

Barbara Carminati and Elena Ferrari. 2010. Privacy-Aware Access Control in Social Networks: Issues and Solutions. In *Privacy and Anonymity in Information Management Systems: New Techniques for New Practical Problems*. Springer London, London, 181–195. DOI: http://dx.doi.org/10.1007/978-1-84996-238-4_9

Barbara Carminati and Elena Ferrari. 2011. Collaborative access control in on-line social networks. In *Proceedings of International Conference on Collaborative Computing: Networking, Applications and Worksharing*. IEEE, 231–240. DOI: http://dx.doi.org/10.4108/icst.collaboratecom.2011.247109

Barbara Carminati, Elena Ferrari, and Andrea Perego. 2006. Rule-Based Access Control for Social Networks. In *On the Move to Meaningful Internet Systems (LNCS)*, Vol. 4278. Springer-Verlag, Berlin, Heidelberg, 1734–1744. DOI: http://dx.doi.org/10.1007/11915072_80

Barbara Carminati, Elena Ferrari, and Andrea Perego. 2009. Enforcing Access Control in Web-based Social Networks. *ACM Trans. Inf. Syst. Secur.* 13, 1, Article 6 (2009), 38 pages. DOI: http://dx.doi.org/10.1145/1609956.1609962

Eve Cohen, Roshan K. Thomas, William Winsborough, and Deborah Shands. 2002. Models for Coalition-based Access Control (CBAC). In *Proceedings of the 7th ACM Symposium on Access Control Models and Technologies (SACMAT '02)*. ACM, New York, NY, USA, 97–106. DOI: http://dx.doi.org/10.1145/507711.507727

Jason Crampton and James Sellwood. 2014. Path Conditions and Principal Matching: A New Approach to Access Control. In *Proceedings of the 19th ACM Symposium on Access Control Models and Technologies (SACMAT '14)*. ACM, New York, NY, USA, 187–198. DOI: http://dx.doi.org/10.1145/2613087.2613094

Stan Damen, Jerry den Hartog, and Nicola Zannone. 2014. CollAC: Collaborative access control. In *Proceedings of International Conference on Collaboration Technologies and Systems*. IEEE, 142–149. DOI: http://dx.doi.org/10.1109/CTS.2014.6867557

Stan Damen and Nicola Zannone. 2013. Privacy Implications of Privacy Settings and Tagging in Facebook. In *Secure Data Management (LNCS)*, Vol. 8425. Springer International Publishing, Cham, Switzerland, 121–138. DOI: http://dx.doi.org/10.1007/978-3-319-06811-4_16

Jerry den Hartog and Nicola Zannone. 2016. Collaborative Access Decisions: Why has my decision not been enforced?. In *Proceedings of the 12th International Conference on Information Systems Security (LNCS)*, Vol. 10063. Springer International Publishing, Cham, Switzerland, 109–130. DOI: http://dx.doi.org/10.1007/978-3-319-49806-5_6

Jerry den Hartog and Nicola Zannone. 2016. A Policy Framework for Data Fusion and Derived Data Control. In *Proceedings of ACM International Workshop on Attribute Based Access Control*. ACM, New York, NY, USA, 47–57. DOI:http://dx.doi.org/10.1145/2875491.2875492

Marina Egea, Federica Paci, Marinella Petrocchi, and Nicola Zannone. 2013. PERSONA - A Personalized Data Protection Framework. In *Trust Management VII (IFIP Advances in Information and Communication Technology)*, Vol. 401. Springer Berlin Heidelberg, Berlin, Heidelberg, 272–280. DOI:http://dx.doi.org/10.1007/978-3-642-38323-6

Lujun Fang and Kristen LeFevre. 2010. Privacy Wizards for Social Networking Sites. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. ACM, New York, NY, USA, 351–360. DOI:http://dx.doi.org/10.1145/1772690.1772727

Elena Ferrari. 2010. Access Control in Data Management Systems. *Synthesis Lectures on Data Management* 2, 1 (2010), 1–117. DOI:http://dx.doi.org/10.2200/S00281ED1V01Y201005DTM004

Ricard L. Fogues, Pradeep K. Murukannaiah, Jose M. Such, and Munindar P. Singh. 2017. Sharing Policies in Multiuser Privacy Scenarios: Incorporating Context, Preferences, and Arguments in Decision Making. *ACM Trans. Comput.-Hum. Interact.* 24, 1 (2017), 5:1–5:29. DOI:http://dx.doi.org/10.1145/3038920

Philip W.L. Fong. 2011. Relationship-based Access Control: Protection Model and Policy Language. In *Proceedings of the 1st ACM Conference on Data and Application Security and Privacy*. ACM, New York, NY, USA, 191–202. DOI:http://dx.doi.org/10.1145/1943513.1943539

Philip W.L. Fong, Pooya Mehregan, and Ram Krishnan. 2013. Relational Abstraction in Community-based Secure Collaboration. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (CCS '13)*. ACM, New York, NY, USA, 585–598. DOI:http://dx.doi.org/10.1145/2508859.2516720

Philip W.L. Fong and Ida Siahaan. 2011. Relationship-based Access Control Policies and Their Policy Languages. In *Proceedings of the 16th ACM Symposium on Access Control Models and Technologies*. ACM, New York, NY, USA, 51–60. DOI:http://dx.doi.org/10.1145/1998441.1998450

Philip W. L. Fong, Mohd M. Anwar, and Zhen Zhao. 2009. A Privacy Preservation Model for Facebook-Style Social Network Systems. In *Proceedings of the 14th European Symposium on Research in Computer Security (LNCS)*, Vol. 5789. Springer-Verlag, Berlin, Heidelberg, 303–320. DOI:http://dx.doi.org/10.1007/978-3-642-04444-1_19

Carrie E. Gates. 2007. Access control requirements for Web 2.0 security and privacy. In *Proceedings of IEEE Web 2.0 Privacy and Security Workshop*.

Sunil Kumar Ghai, Prateek Nigam, and Ponnurangam Kumaraguru. 2010. Cue: A Framework for Generating Meaningful Feedback in XACML. In *Proceedings of the 3rd ACM Workshop on Assurable and Usable Security Configuration (SafeConfig '10)*. ACM, New York, NY, USA, 9–16. DOI:http://dx.doi.org/10.1145/1866898.1866901

Paolo Giorgini, Fabio Massacci, John Mylopoulos, and Nicola Zannone. 2006. Requirements engineering for trust management: model, methodology, and reasoning. *International Journal of Information Security* 5, 4 (2006), 257–274. DOI:http://dx.doi.org/10.1007/s10207-006-0005-7

Paolo Guarda and Nicola Zannone. 2009. Towards the development of privacy-aware systems. *Information & Software Technology* 51, 2 (2009), 337–350. DOI:http://dx.doi.org/10.1016/j.infsof.2008.04.004

Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2012. Enabling Collaborative data sharing in Google+. In *Proceedings of IEEE Global Communications Conference*. IEEE, 720–725. DOI:http://dx.doi.org/10.1109/GLOCOM.2012.6503198

Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2011. Detecting and Resolving Privacy Conflicts for Collaborative Data Sharing in Online Social Networks. In *Proceedings of the 27th Annual Computer Security Applications Conference (ACSAC '11)*. ACM, New York, NY, USA, 103–112. DOI:http://dx.doi.org/10.1145/2076732.2076747

Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2013. Multiparty Access Control for Online Social Networks: Model and Mechanisms. *IEEE Transactions on Knowledge & Data Engineering* 25, 7 (2013), 1614–1627. DOI:http://dx.doi.org/10.1109/TKDE.2012.97

Hongxin Hu, Gail-Joon Ahn, Ziming Zhao, and Dejun Yang. 2014a. Game Theoretic Analysis of Multiparty Access Control in Online Social Networks. In *Proceedings of the 19th ACM Symposium on Access Control Models and Technologies (SACMAT '14)*. ACM, New York, NY, USA, 93–102. DOI:http://dx.doi.org/10.1145/2613087.2613097

Vincent C. Hu, David Ferraiolo, Rick Kuhn, Adam Schnitzer, Kenneth Sandlin, Robert Miller, and Karen Scarfone. 2014b. *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*. NIST Special Publication 800-162. NIST National Institute of Standards and Technology. DOI:http://dx.doi.org/10.6028/NIST.SP.800-162

Panagiotis Ilia, Barbara Carminati, Elena Ferrari, Paraskevi Fragopoulou, and Sotiris Ioannidis. 2017. SAMPAC: Socially-Aware Collaborative Multi-Party Access Control. In *Proceedings of the 7th ACM on Conference on Data and Application Security and Privacy (CODASPY '17)*. ACM, New York, NY, USA, 71–82. DOI:http://dx.doi.org/10.1145/3029806.3029834

Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. 2015. Face/Off: Preventing Privacy Leakage From Photos in Social Networks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. ACM, New York, NY, USA, 781–792. DOI:http://dx.doi.org/10.1145/2810103.2813603

Sushil Jajodia, Pierangela Samarati, Maria Luisa Sapino, and V. S. Subrahmanian. 2001. Flexible Support for Multiple Access Control Policies. *ACM Trans. Database Syst.* 26, 2 (2001), 214–260. DOI:http://dx.doi.org/10.1145/383891.383894

Simon Jones and Eamonn O'Neill. 2010. Feasibility of Structural Network Clustering for Group-based Privacy Control in Social Networks. In *Proceedings of the 6th Symposium on Usable Privacy and Security (SOUPS '10)*. ACM, New York, NY, USA, Article 9, 13 pages. DOI:http://dx.doi.org/10.1145/1837110.1837122

Simon Jones and Eamonn O'Neill. 2011. Contextual Dynamics of Group-based Sharing Decisions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 1777–1786. DOI:http://dx.doi.org/10.1145/1978942.1979200

Daniel Kahneman. 2003. Maps of Bounded Rationality: Psychology for Behavioral Economics. *American Economic Review* 93, 5 (2003), 1449–1475. DOI:http://dx.doi.org/10.1257/000282803322655392

Anas Abou El Kalam, Salem Benferhat, Alexandre Miège, Rania El Baida, Frédéric Cuppens, Claire Saurel, Philippe Balbiani, Yves Deswarte, and Gilles Trouessin. 2003. Organization based access control. In *Proceedings of the 4th International Workshop on Policies for Distributed Systems and Networks*. IEEE, Washington, DC, USA, 120–131. DOI:http://dx.doi.org/10.1109/POLICY.2003.1206966

Myong H. Kang, Joon S. Park, and Judith N. Froscher. 2001. Access Control Mechanisms for Inter-organizational Workflow. In *Proceedings of the 6th ACM Symposium on Access Control Models and Technologies (SACMAT '01)*. ACM, New York, NY, USA, 66–74. DOI:http://dx.doi.org/10.1145/373256.373266

Apu Kapadia, Geetanjali Sampemane, and Roy H. Campbell. 2004. KNOW Why Your Access Was Denied: Regulating Feedback for Usable Security. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS '04)*. ACM, New York, NY, USA, 52–61. DOI:http://dx.doi.org/10.1145/1030083.1030092

Imrul Kayes and Adriana Iamnitchi. 2015. A Survey on Privacy and Security in Online Social Networks. *CoRR* abs/1504.03342 (2015).

Taeseong Kim, Christopher D. Cera, William C. Regli, Hyunseung Choo, and JungHyun Han. 2006. Multi-Level modeling and access control for data sharing in collaborative design. *Advanced Engineering Informatics* 20, 1 (2006), 47–57.

Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Lorrie Faith Cranor, Nitin Gupta, and Michael Reiter. 2012. Tag, You Can See It!: Using Tags for Access Control in Photo Sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 377–386. DOI:http://dx.doi.org/10.1145/2207676.2207728

Spyros Kokolakis. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security* 64 (2017), 122–134. DOI:http://dx.doi.org/10.1016/j.cose.2015.07.002

Ram Krishnan, Jianwei Niu, Ravi Sandhu, and William H. Winsborough. 2011. Group-Centric Secure Information-Sharing Models for Isolated Groups. *ACM Trans. Inf. Syst. Secur.* 14, 3, Article 23 (2011), 29 pages. DOI:http://dx.doi.org/10.1145/2043621.2043623

Butler W. Lampson. 1974. Protection. *SIGOPS Oper. Syst. Rev.* 8, 1 (1974), 18–24. DOI:http://dx.doi.org/10.1145/775265.775268

Scott Lederer, Jason I. Hong, Anind K. Dey, and James A. Landay. 2004. Personal Privacy Through Understanding and Action: Five Pitfalls for Designers. *Personal Ubiquitous Comput.* 8, 6 (2004), 440–454. DOI:http://dx.doi.org/10.1007/s00779-004-0304-9

Ninghui Li, Benjamin N. Grosof, and Joan Feigenbaum. 2003. Delegation Logic: A Logic-based Approach to Distributed Authorization. *ACM Trans. Inf. Syst. Secur.* 6, 1 (2003), 128–171. DOI:http://dx.doi.org/10.1145/605434.605438

Ninghui Li, John C. Mitchell, and William H. Winsborough. 2002. Design of a Role-Based Trust-Management Framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy (SP '02)*. IEEE Computer Society, Washington, DC, USA, 114–. DOI:http://dx.doi.org/10.1109/SECPRI.2002.1004366

Ninghui Li, Qihua Wang, Wahbeh Qardaji, Elisa Bertino, Prathima Rao, Jorge Lobo, and Dan Lin. 2009. Access Control Policy Combining: Theory Meets Practice. In *Proceedings of the 14th ACM Symposium on Access Control Models and Technologies (SACMAT '09)*. ACM, New York, NY, USA, 135–144. DOI:http://dx.doi.org/10.1145/1542207.1542229

Rauf Mahmudlu, Jerry den Hartog, and Nicola Zannone. 2016. Data Governance & Transparency for Collaborative Systems. In *Data and Applications Security and Privacy (LNCS)*. Springer International Publishing, Cham, Switzerland, 199–216. DOI:http://dx.doi.org/10.1007/978-3-319-41483-6_15

Ilaria Matteucci, Paolo Mori, and Marinella Petrocchi. 2012. Prioritized Execution of Privacy Policies. In *Data Privacy Management and Autonomous Spontaneous Security (LNCS)*, Vol. 7731. Springer, Berlin, Heidelberg, 133–145. DOI:http://dx.doi.org/10.1007/978-3-642-35890-6_10

Michelle L. Mazurek, Yuan Liang, William Melicher, Manya Sleeper, Lujo Bauer, Gregory R. Ganger, Nitin Gupta, and Michael K. Reiter. 2014. Toward strong, usable access control for shared distributed data. In *Proceedings of the 12th USENIX Conference on File and Storage Technologies (FAST '14)*. USENIX, Santa Clara, CA, 89–103. https://www.usenix.org/system/files/conference/fast14/fast14-paper_mazurek.pdf

Alessandra Mazzia, Kristen LeFevre, and Eytan Adar. 2012. The PViz Comprehension Tool for Social Network Privacy Settings. In *Proceedings of the 8th Symposium on Usable Privacy and Security (SOUPS '12)*. ACM, New York, NY, USA, Article 13, 12 pages. DOI: http://dx.doi.org/10.1145/2335356.2335374

Pietro Mazzoleni, Bruno Crispo, Swaminathan Sivasubramanian, and Elisa Bertino. 2008. XACML Policy Integration Algorithms. *ACM Trans. Inf. Syst. Secur.* 11, 1, Article 4 (2008), 29 pages. DOI: http://dx.doi.org/10.1145/1330295.1330299

Patrick McDaniel and Atul Prakash. 2006. Methods and Limitations of Security Policy Reconciliation. *ACM Trans. Inf. Syst. Secur.* 9, 3 (2006), 259–291. DOI: http://dx.doi.org/10.1145/1178618.1178620

Caitlin McLaughlin and Jessica Vitak. 2012. Norm evolution and violation on Facebook. *New Media & Society* 14, 2 (2012), 299–315. DOI: http://dx.doi.org/10.1177/1461444811412712

Pooya Mehregan and Philip W.L. Fong. 2016. Policy Negotiation for Co-owned Resources in Relationship-Based Access Control. In *Proceedings of the 21st ACM on Symposium on Access Control Models and Technologies (SACMAT '16)*. ACM, New York, NY, USA, 125–136. DOI: http://dx.doi.org/10.1145/2914642.2914652

Judith S. Olson, Jonathan Grudin, and Eric Horvitz. 2005. A Study of Preferences for Sharing and Privacy. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI EA '05)*. ACM, New York, NY, USA, 1985–1988. DOI: http://dx.doi.org/10.1145/1056808.1057073

Jaehong Park and Ravi Sandhu. 2004. The UCONABC Usage Control Model. *ACM Trans. Inf. Syst. Secur.* 7, 1 (2004), 128–174. DOI: http://dx.doi.org/10.1145/984334.984339

João Paulo Pesce, Diego Las Casas, Gustavo Rauber, and Virgílio Almeida. 2012. Privacy Attacks in Social Media Using Photo Tagging Networks: A Case Study with Facebook. In *Proceedings of the 1st Workshop on Privacy and Security in Online Social Media (PSOSM '12)*. ACM, New York, NY, USA, Article 4, 8 pages. DOI: http://dx.doi.org/10.1145/2185354.2185358

Charles E. Phillips, Jr., T.C. Ting, and Steven A. Demurjian. 2002. Information Sharing and Security in Dynamic Coalitions. In *Proceedings of the 7th ACM Symposium on Access Control Models and Technologies (SACMAT '02)*. ACM, New York, NY, USA, 87–96. DOI: http://dx.doi.org/10.1145/507711.507726

Moo-Ryong Ra, Ramesh Govindan, and Antonio Ortega. 2013. P3: Toward Privacy-preserving Photo Sharing. In *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation*. USENIX Association, Berkeley, CA, USA, 515–528.

Sarah Rajtmajer, Anna Squicciarini, Christopher Griffin, Sushama Karumanchi, and Alpana Tyagi. 2016. Constrained Social-Energy Minimization for Multi-Party Sharing in Online Social Networks. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 680–688.

Prathima Rao, Dan Lin, Elisa Bertino, Ninghui Li, and Jorge Lobo. 2011. Fine-grained Integration of Access Control Policies. *Comput. Secur.* 30, 2-3 (2011), 91–107. DOI: http://dx.doi.org/10.1016/j.cose.2010.10.006

Robert W. Reeder, Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, Kelli Bacon, Keisha How, and Heather Strong. 2008. Expandable Grids for Visualizing and Authoring Computer Security Policies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 1473–1482. DOI: http://dx.doi.org/10.1145/1357054.1357285

Robert W Reeder, Lujo Bauer, Lorrie Faith Cranor, Michael K Reiter, and Kami Vaniea. 2009. *Effects of access-control policy conflict-resolution methods on policy-authoring usability*. Technical Report CMU-CyLab-09-006. CyLab. 12 pages.

Jennifer Rode, Carolina Johansson, Paul DiGioia, Roberto Silva Filho, Kari Nies, David H. Nguyen, Jie Ren, Paul Dourish, and David Redmiles. 2006. Seeing Further: Extending Visualization As a Basis for Usable Security. In *Proceedings of the 2nd Symposium on Usable Privacy and Security (SOUPS '06)*. ACM, New York, NY, USA, 145–155. DOI: http://dx.doi.org/10.1145/1143120.1143138

Jerome H. Saltzer and Michael D. Schroeder. 1975. The protection of information in computer systems. *Proc. IEEE* 63, 9 (1975), 1278–1308. DOI: http://dx.doi.org/10.1109/PROC.1975.9939

Pierangela Samarati and Sabrina De Capitani di Vimercati. 2000. Access Control: Policies, Models, and Mechanisms. In *Foundations of Security Analysis and Design (LNCS)*, Vol. 2171. Springer Berlin Heidelberg, Berlin, Heidelberg, 137–196. DOI: http://dx.doi.org/10.1007/3-540-45608-2_3

Ravi S. Sandhu. 1996. Roles Versus Groups. In *Proceedings of the 1st ACM Workshop on Role-based Access Control (RBAC '95)*. ACM, New York, NY, USA, 25–26. DOI: http://dx.doi.org/10.1145/270152.270163

Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. 1996. Role-Based Access Control Models. *Computer* 29, 2 (1996), 38–47. DOI: http://dx.doi.org/10.1109/2.485845

Roman Schlegel, Apu Kapadia, and Adam J. Lee. 2011. Eyeing Your Exposure: Quantifying and Controlling Information Sharing for Improved Privacy. In *Proceedings of the 7th Symposium on Usable Privacy and Security (SOUPS '11)*. ACM, New York, NY, USA, Article 14, 14 pages. DOI: http://dx.doi.org/10.1145/2078827.2078846

HongHai Shen and Prasun Dewan. 1992. Access Control for Collaborative Environments. In *Proceedings of the 1992 ACM Conference on Computer-supported Cooperative Work (CSCW '92)*. ACM, New York, NY, USA, 51–58. DOI: http://dx.doi.org/10.1145/143457.143461

Herbert Alexander Simon. 1957. A behavioural model of Rational Choice. In *Models of man: social and rational; mathematical essays on rational human behavior in a social setting*. J. Wiley, New York, 241–260.

Eleftherios Spyromitros-Xioufis, Symeon Papadopoulos, Adrian Popescu, and Yiannis Kompatsiaris. 2016. Personalized Privacy-aware Image Classification. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval (ICMR '16)*. ACM, New York, NY, USA, 71–78. DOI:http://dx.doi.org/10.1145/2911996.2912018

Anna Squicciarini, Cornelia Caragea, and Rahul Balakavi. 2017. Toward Automated Online Photo Privacy. *ACM Trans. Web* 11, 1, Article 2 (2017), 29 pages. DOI:http://dx.doi.org/10.1145/2983644

Anna Squicciarini, Sushama Karumanchi, Dan Lin, and Nicole Desisto. 2014. Identifying Hidden Social Circles for Advanced Privacy Configuration. *Comput. Secur.* 41 (2014), 40–51. DOI:http://dx.doi.org/10.1016/j.cose.2013.07.007

Anna Squicciarini, Dan Lin, Smitha Sundareswaran, and Joshua Wede. 2015. Privacy Policy Inference of User-Uploaded Images on Content Sharing Sites. *IEEE Transactions on Knowledge and Data Engineering* 27, 1 (2015), 193–206. DOI:http://dx.doi.org/10.1109/TKDE.2014.2320729

Anna Squicciarini, Federica Paci, and Smitha Sundareswaran. 2014. PriMa: a comprehensive approach to privacy protection in social network sites. *Annales des Télécommunications* 69, 1-2 (2014), 21–36. DOI:http://dx.doi.org/10.1007/s12243-013-0371-x

Anna Squicciarini, Mohamed Shehab, and Joshua Wede. 2010. Privacy Policies for Shared Content in Social Network Sites. *The VLDB Journal* 19, 6 (2010), 777–796. DOI:http://dx.doi.org/10.1007/s00778-010-0193-7

Scott D. Stoller, Ping Yang, Mikhail I. Gofman, and C.R. Ramakrishnan. 2011. Symbolic reachability analysis for parameterized administrative role-based access control. *Computers & Security* 30, 2–3 (2011), 148–164. DOI:http://dx.doi.org/10.1016/j.cose.2010.08.002

Jose M. Such and Natalia Criado. 2016. Resolving Multi-Party Privacy Conflicts in Social Media. *IEEE Trans. Knowl. Data Eng.* 28, 7 (2016), 1851–1863. DOI:http://dx.doi.org/10.1109/TKDE.2016.2539165

Jose M. Such and Michael Rovatsos. 2016. Privacy Policy Negotiation in Social Media. *ACM Trans. Auton. Adapt. Syst.* 11, 1, Article 4 (2016), 29 pages. DOI:http://dx.doi.org/10.1145/2821512

Vivy Suhendra. 2011. *A Survey on Access Control Deployment*. Communications in Computer and Information Science, Vol. 259. Springer Berlin Heidelberg, Berlin, Heidelberg, 11–20. DOI:http://dx.doi.org/10.1007/978-3-642-27189-2_2

Roshan K. Thomas. 1997. Team-based Access Control (TMAC): A Primitive for Applying Role-based Access Controls in Collaborative Environments. In *Proceedings of the Second ACM Workshop on Role-based Access Control (RBAC '97)*. ACM, New York, NY, USA, 13–19. DOI:http://dx.doi.org/10.1145/266741.266748

Roshan K. Thomas and Ravi S. Sandhu. 1997. Task-based authorization controls (TBAC): A family of models for active and enterprise-oriented authorization management. In *DBSec*. Springer US, Boston, MA, 166–181. DOI:http://dx.doi.org/10.1007/978-0-387-35285-5_10

William Tolone, Gail-Joon Ahn, Tanusree Pai, and Seng-Phil Hong. 2005. Access Control in Collaborative Systems. *ACM Comput. Surv.* 37, 1 (2005), 29–41. DOI:http://dx.doi.org/10.1145/1057977.1057979

Ashwini Kishore Tonge and Cornelia Caragea. 2016. Image Privacy Prediction Using Deep Features. In *Proceedings of the 13th AAAI Conference on Artificial Intelligence*. AAAI Press, 4266–4267.

Daniel Trivellato, Nicola Zannone, and Sandro Etalle. 2014. GEM: A distributed goal evaluation algorithm for trust management. *Theory and Practice of Logic Programming* 14, 3 (2014), 293–337. DOI:http://dx.doi.org/10.1017/S1471068412000397

Daniel Trivellato, Nicola Zannone, Maurice Glaundrup, Jacek Skowronek, and Sandro Etalle. 2013. A Semantic Security Framework for Systems of Systems. *International Journal of Cooperative Information Systems* 22, 1 (2013), 35. DOI:http://dx.doi.org/10.1142/S0218843013500044

Nishant Vishwamitra, Yifang Li, Kevin Wang, Hongxin Hu, Kelly Caine, and Gail-Joon Ahn. 2017. Towards PII-based Multiparty Access Control for Photo Sharing in Online Social Networks. In *Proceedings of the 22Nd ACM on Symposium on Access Control Models and Technologies (SACMAT '17)*. ACM, New York, NY, USA, 155–166. DOI:http://dx.doi.org/10.1145/3078861.3078875

Jacques Wainer, Paulo Barthelmess, and Akhil Kumar. 2003. W-RBAC – A Workflow Security Model Incorporating Controlled Overriding of Constraints. *International Journal of Cooperative Information Systems* 12, 04 (2003), 455–485. DOI:http://dx.doi.org/10.1142/S0218843003000814

Yang Wang, Liang Gou, Anbang Xu, Michelle X. Zhou, Huahai Yang, and Hernan Badenes. 2015. VeilMe: An Interactive Visualization Tool for Privacy Configuration of Using Personality Traits. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 817–826. DOI:http://dx.doi.org/10.1145/2702123.2702293

Ryan Wishart, Domenico Corapi, Srdjan Marinovic, and Morris Sloman. 2010. Collaborative Privacy Policy Authoring in a Social Networking Context. In *Proceedings of the 2010 IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY '10)*. IEEE Computer Society, Washington, DC, USA, 1–8. DOI:http://dx.doi.org/10.1109/POLICY.2010.13

Pamela Wisniewski, Heng Xu, Heather Lipford, and Emmanuel Bello-Ogunu. 2015. Facebook apps and tagging: The trade-off between personal privacy and engaging with friends. *Journal of the Association for Information Science and Technology* 66, 9 (2015), 1883–1896. DOI:http://dx.doi.org/10.1002/asi.23299

Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Bjöorn Regnell, and Anders Wesslén. 2000. *Experimentation in Software Engineering: An Introduction*. Kluwer Academic Publishers, Norwell, MA, USA.

XACML v2.0. 2005. eXtensible Access Control Markup Language (XACML) Version 2.0. OASIS. (2005).

XACML v3.0. 2013. eXtensible Access Control Markup Language (XACML) Version 3.0. OASIS. (2013).

Qian Xiao and Kian-Lee Tan. 2012. Peer-aware collaborative access control in social networks. In *Proceedings of International Conference on Collaborative Computing: Networking, Applications and Worksharing*. IEEE, 30–39. DOI:http://dx.doi.org/10.4108/icst.collaboratecom.2012.250524

Xiaowei Xu, Nurcan Yuruk, Zhidan Feng, and Thomas A. J. Schweiger. 2007. SCAN: A Structural Clustering Algorithm for Networks. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '07)*. ACM, New York, NY, USA, 824–833. DOI:http://dx.doi.org/10.1145/1281192.1281280

Jun Yu, Baopeng Zhang, Zhengzhong Kuang, Dan Lin, and Jianping Fan. 2017. iPrivacy: Image Privacy Protection by Identifying Sensitive Objects via Deep Multi-Task Learning. *IEEE Trans. Information Forensics and Security* 12, 5 (2017), 1005–1016. DOI:http://dx.doi.org/10.1109/TIFS.2016.2636090

Sergej Zerr, Stefan Siersdorfer, Jonathon Hare, and Elena Demidova. 2012. Privacy-aware Image Classification and Search. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '12)*. ACM, New York, NY, USA, 35–44. DOI:http://dx.doi.org/10.1145/2348283.2348292