

INVITED PAPER

Authorization Models for Secure Information Sharing: A Survey and Research Agenda

Farzad Salim^{a,*}, Jason Reid^a, and Ed Dawson^a

^aInformation Security Institute, Queensland University of Technology, Australia

ARTICLE INFO.

Article history:

Received: 30 November 2009

Revised: 14 February 2010

Accepted: 20 January 2010

Published Online: 2 June 2010

Keywords:

Security, Access Control,
Authorisation Model, Multi
Objective Decision Making, Secure
Information Sharing, Dynamic
Environment

ABSTRACT

This article presents a survey of authorisation models and considers their ‘fitness-for-purpose’ in facilitating information sharing. Network-supported information sharing is an important technical capability that underpins collaboration in support of dynamic and unpredictable activities such as emergency response, national security, infrastructure protection, supply chain integration and emerging business models based on the concept of a ‘virtual organisation’. The article argues that present authorisation models are inflexible and poorly scalable in such dynamic environments due to their assumption that the future needs of the system can be predicted, which in turn justifies the use of persistent authorisation policies. The article outlines the motivation and requirement for a new flexible authorisation model that addresses the needs of information sharing. It proposes that a flexible and scalable authorisation model must allow an *explicit* specification of the objectives of the system and access decisions must be made based on a late trade-off analysis between these explicit objectives. A research agenda for the proposed Objective-Based Access Control concept is presented.

© 2010 ISC. All rights reserved.

1 Introduction

Dynamic environments are rapidly emerging as computing systems morph from monolithic and closed entities into globally disaggregated collaborating entities that may need to share sensitive information. There is an emerging need for scalable access control solutions for systems operating in such dynamic and uncertain environments, where changes are frequent

and there are unpredictable threats as well as opportunities [1, 2]. The dynamism and uncertainty that exist in such environments are challenging the most basic foundation of current access control approaches: a security policy as a set of rules, which is the essence of an *already-made trade-off* analysis between a range of system objectives [3–6].

The problem of access control is becoming more contextual as the dynamism of the environment increases. In an environment where the demand for information and the incentives provided for disclosure change, an entity’s posture towards information disclosure also changes. What is and is not acceptable regarding access must be decided on the basis of the

* Corresponding author.

Email addresses: farzad@isi.qut.edu.au (F. Salim),
reid@isi.qut.edu.au (J. Reid), e.dawson@isi.qut.edu.au
(E. Dawson).

ISSN: 2008-2045 © 2010 ISC. All rights reserved.

context of current threats and opportunities present in the environment.

In such an environment the window of *predictability* that old models are based on has narrowed. This challenges the applicability of such models, which at their core assume a relatively static policy that is the result of an already made trade-off analysis between the competing requirements of the system. As Blakley [3] points out, such policies do not scale well and their complexity quickly increases as systems grow and diverge. Similarly, Baker [4] points out that constructing a security policy is a very complex task and a single security policy may not be appropriate in a complex system.

Authorisation underpins the ability to share sensitive information electronically since information must only be disclosed to authorised entities. One area where current approaches are demonstrably inadequate is in critical infrastructure protection [7]. Information sharing among separate communities such as government and private sector infrastructure operators (e.g. telecommunications, energy, finance) has become a priority for many countries, including Australia, Canada, the UK and the US. In Australia specifically, forums such as *Trusted Information Sharing Network (TISN)* [8] have been formed to allow owners and operators of critical infrastructure to work together and share sensitive information. However, information sharing in such networks occurs predominantly in physical meetings where representatives of different organisations meet face-to-face. There are significant advantages to supplementing such physical meetings with a virtual information sharing network. This would allow for example, organisations to share information on cyber attacks in real time. A significant barrier to such information sharing stems from complex technology challenges [9]. A key technology challenge is to overcome the *inflexibility* and poor *scalability* of existing authorisation models in dealing with changes of the environment [3–6].

Current authorisation models are based on a persistent policy (i.e. usually written as a set of *action* rules: if *condition* then *decision*) [10]. These have several characteristics that make them undesirable for dynamic environments. First, they are inherently rigid, situation agnostic and poorly scalable because the policy is assumed to be correct and the essence of an already made trade-off analysis between various organisational objectives [11]. In this context a new trade-off analysis to produce a new policy is possible but has two important drawbacks, inefficiency and instability of the successive policies [12, 13]. Second, policy objectives are *implicit* rather than being *explicitly* specified. Hence, there is no relationship between

the rules in the policy, the decision made, and the objectives. In other words, there is no way to answer *why* the decision has been made [14–16]. Third, such policies are usually *closed* to ensure decidability (i.e. they use a default rule, usually a deny rule, that is returned for those access requests for which there is no explicit rule in the policy). Hence for them, the lack of knowledge for decision making is not explicit; it corresponds to knowing that the request must be denied. Given these issues, neither the concept of *compromise* nor *opportunity* make sense, because the assumption is that the initial trade-off analysis has already predicted and taken into account any important factor for authorisation decisions. Based on this analysis, the required rules are assumed to exist.

There have been several attempts to address the inflexibility of authorisation models through allowing environment conditions to be input parameters for the policy. These approaches have been mostly referred to as context-based policies [17–19]. However, at their core, they still require the environment conditions to be predicted, the value of these conditions to be determined and a priori decisions to be made concerning what to do in each condition.

Other approaches have been recently proposed to address the inflexibility of existing authorisation models. The core of these approaches is to introduce a grey area between granting and denying a request. In this area exceptional and unpredicted access could be granted [5, 13, 20]. These works have rightly identified the need for more flexible authorisation models, however, so far their focus has been on incremental improvements to make the existing access control systems more flexible, rather than identifying why they are inflexible. In other words, what is it that all access control models have in common which makes them inflexible? The existing approaches are ad-hoc attempts to address part of the inflexibility problem and fall short in their generality and systematic specification of the problem or solution to the problem. Based on our analysis of current authorisation models in Section 4, we have two hypotheses:

- (1) The inflexibility and poor scalability of the existing authorisation models is due to their assumption about the *predictability* of the future *needs* of the system, which in turn justifies the use of *persistent access control policies*.
- (2) A flexible and scalable authorisation model must allow an *explicit* specification of the objectives of the system and an access decision must be made based on the trade-off analysis between these explicit objectives performed at or near runtime.

These hypotheses have led us to propose the con-

cept of Objective-Based Access Control. The rest of this paper is organised as follows: Section 2 introduces the terminology used in this article. Section 3 sets the context of the paper by briefly describing the characteristics of an authorisation model for emerging dynamic environments, while Section 4 surveys authorisation models in terms of their suitability for such environments. Section 5 outlines our proposal for the concept of Objective-Based Access Control and Section 6 presents our research agenda.

2 Terminology

Here, we informally define several important terms used in the article to clarify their intended meaning.

- *Objective* is a state that one would like to achieve, maintain or maximize. For example, confidentiality can be considered as an objective that one would like to maintain. Maximising throughput or income, or maintaining reputation are further examples of objectives.
- *Policy* is a set of rules and conditions for achieving an individual objective. For example, the privacy policy is a set of rules which define how to meet the privacy objective.
- *System/agent/player* refer to an entity that has some objectives and can make a decision regarding the release of some information to another entity.
- *Environment* is a set of conditions that are not directly under the control of the system but might affect the access decision. For example, government rules and regulations, an emergency event, the number of entities joining a coalition, etc. are considered as defining elements of the environment.
- *Incentives/forces* are considered as a system's interpretation of the changes in the environment with respect to its objectives. For example, a government regulation to fine those who breach privacy is an incentive to consider the privacy objective importance. The access requester that offers a payment for an access provides an incentive to a financially motivated system to grant the access.

3 Emergence of Dynamic Environments

In the past decades information systems have been revolutionised by low-cost information and communications technology which has led organisations to pursue their mission and derive competitive advantage through strategic partnerships and collaboration. The Internet has been a major enabling factor in this

transition by providing a flexible medium to provide or request *resources*¹, which in turn has made ad-hoc collaboration between these entities not only a possibility but a necessity for their survival and competitiveness. An instance of this is an organisation in a mobile ad-hoc network that forms a coalition to respond to an emergency or disaster. Similarly, government and private agencies that are part of nation's critical infrastructure (e.g. electricity, telecommunication) collaborate and share information to recognise and address threats and system vulnerabilities and to minimise the consequence of adverse events. In supply chains, entities often form dynamic coalitions that require sharing of information and resources [21].

All of these coalitions are highly dependent on each of the involved entities to provide the information required for the coalition to function. However, such a need is highly dynamic, because the information required depends on external uncontrolled factors in the environment; the participants who need the information and the channel through which the information should be shared are also dynamically determined. These are just a short list of factors to indicate the unpredictability involved. Furthermore, every entity in a coalition may have several requirements or needs to satisfy, that determine its posture towards information sharing. For the sake of example, consider the *secrecy* of a piece of information, *reputation* and *monetary profit* as objectives. Note that the balance and the importance of each of these needs may change based on the *incentives* that exist at any point in time for information sharing [13, 14, 20]. In the simple case where an entity has two important objectives: *preserving privacy* and *national security*, the policy for the privacy objective binds the entity to release information to *no one*. Now, in a case of an epidemic outburst, the entity might face a decision to release confidential data for the sake of national security. Hence, the entity can *compromise* its privacy objective to take the *opportunity* of satisfying national security objective and thereby address the threat of the epidemic.

The following important points should be noted; first, the *weight* of objectives in relation to each other can be considerably more subtle and environment-dependent than the above example *privacy* vs. *national security*. For example, it could have been *privacy* vs. *monetary profit*, which could be highly dependent on the entity's financial situation at the time. Further, note that in the above example, the entity that highly values privacy protection, even in the face of a national security threat may want to specify approaches that are less privacy invasive, such as, one

¹ An abstract term for objects, information, files, documents or services.

time information use, or only allowing on-site evaluation of records. Hence, there is no crisp boundary of *satisfying* privacy or not. Second, the entity that compromises an objective (e.g. privacy) actually operates against its policy for that specific objective, but its decision is still aligned with its overall need (combined objectives). Finally, note that the privacy policy (for achieving privacy objective) is correct and the entity must not add an exception to the policy. Ad hoc exceptions are a common and undesirable response to address the unpredicted need for sharing information [13]. A better approach is to handle the exception through the trade-off analysis based on the weight/importance of the objectives.

In the following sections we will survey existing authorisation models and show why each of them fails to address the above mentioned issues.

4 Survey of Authorisation Models for Information Sharing

This section presents a survey of authorisation models, which are analysed in terms of their appropriateness in facilitating information sharing. The survey is structured in three broad categories: traditional models, credential-based models, and risk-based models. The categorisation of these approaches to authorisation is mainly based on their basic assumptions regarding the *predictability* of the *needs* of the system which is closely related to the predictability of the *forces* and *incentives* of the *environment* in which the system operates.

Section 4.1 briefly discusses the building blocks of an access control system and their theoretical boundaries. The aim is to narrow the focus of this paper to the authorisation aspect of access control.

Section 4.2 describes traditional approaches that assume a closed, controlled and predictable environment, where a user and their required access rights are known, further, possible changes in the environment are also assumed to have been foreseen and incorporated into the policy.

Section 4.3 introduces credential-based approaches that are divided into two categories: 1) the trust management approach, which assumes the system knows what it needs to satisfy (i.e. trust) but does not know the identity of those users that can satisfy the need; 2) the Digital Rights Management (DRM) approach, which was originally designed for payment-based systems. The DRM user model is less comprehensive and its focus is mostly on client side enforcement of access rights.

Section 4.4 analyses the recent risk-based ap-

proaches to authorisation. They assume the system has an inherent need and benefits from information sharing. Hence, they blend risk management with authorisation to allow an access within an acceptable level of risk. The assumption is that a more flexible authorisation model would lead to more information sharing that ought to be beneficial. Risk-based approaches are discussed under three sections, based on how they view risk: the survivability approach (Section 4.4.1) and optimistic approach (Section 4.4.2) attempt to account for the risks associated with denying an access, while quantified risk-based approaches (Section 4.4.3) focus on quantifying the risks that are associated with granting access. Hence, all the approaches are commonly attempting to reduce the overall risk.

4.1 Access Control and Authorisation

Historically, *administrators* have controlled access to sensitive data by associating appropriate access rights to long-term local identities that may need to access the resource. The actual granting of access then requires establishing a level of confidence in a user's claimed identity through *authentication* and a determination of access rights for the identity through *authorisation*. The combination of these two tasks plus *audit* is referred to as *access control*.

According to Samarati *et al.* [10], the development of an authorisation system can be theoretically divided into three phases: first, the specification of the rules, on which access is to be constrained. The collection of these rules are referred to as a *security policy*. Second, there is a formal representation of security policies, referred to as *security model*. This allows the properties of the model to be mathematically analysed and proved. Third, there is the development of the necessary software and hardware required to implement the security policy within the constraints of the model, referred to as *security mechanism*. Such a separation in theory allows for security policies to be defined, analysed and compared independent of mechanisms and vice-versa [10].

Here we are only interested in the authorisation aspect of access control that Anderson [22] defines as the process of mediating every request to resources and data maintained by a system and determining whether the request should be granted or denied - the process of making an access decision. Furthermore, this paper presents an analytical study of the characteristics of novel authorisation models in terms of their fitness for dynamic environments. It is not intended to be a complete survey of all the existing authorisation models.

4.2 Traditional Authorisation

Traditionally, the need for access control² came from two major fields: firstly, military, mainly focusing on confidentiality of data; secondly, businesses and civilian governments, primarily demanding flexible models for data integrity [23]. The division between these two needs led to the evolution of two distinct access control models, known as *mandatory* and *discretionary* access control. However, the limitations and rigidity of each of these promoted further research in this area that resulted in alternatives such as role-based access control [24], and task-based access control [17, 25] that will also be described in the following sections.

4.2.1 Mandatory Access Control

Historically, Mandatory Access Control (MAC) is associated with the multi-level security (MLS) model of Bell and LaPadulla [26]. MLS-based MAC is a way of restricting access to objects based on security clearances assigned to users and security labels attached to objects within the system. The model is designed to restrict information flow from more secure classification levels to less secure levels. The controls are mandatory in the sense that they are system-enforced and cannot be modified by users or their programs.

While the MLS model protects the confidentiality of information, it lacks the necessary control for enforcing an *integrity* objective, since subjects with a lower clearance can still make modifications to objects of a higher classification. To address this limitation Biba [27] proposed a model to prevent subjects from indirectly modifying information they cannot read.

MAC models have one important characteristic that could be identified both as a strength and a weakness: they are concerned with one single objective, either confidentiality or integrity. This is very interesting as one can reason as to whether that specific objective is being satisfied or not, every other objective ignored. For example, neither the Bell Lapadulla nor the Biba model care if the systems in which they are implemented have other objectives for which the release of information becomes a necessity. However, this is a shortcoming as well. These policies can reduce productivity by limiting the necessary information flow [13]. Also, the implementation of such models requires a trusted central administrator to assign labels for all the subject and objects within the system. The insensitivity to the other objectives, in parallel with the growing complexity of such systems (with several

objectives) means that MAC models are theoretically interesting but impractical in dynamic contexts [7].

4.2.2 Discretionary Access Control

Discretionary Access Control (DAC) is a way of restricting access to objects based on the identity of the user. Explicit access rules specify the type of access to objects granted to each identity. Access to the resource is only granted when such an association exists. One major difference between DAC and MAC is the discretionary nature of control, in the sense that a subject with a certain access permission is capable of passing that permission onto other subjects [28]. This granting and revocation is however done under the provision of an administrative policy.

The earliest approach to implement a DAC policy involves the use of an *Access Control Matrix* [28]. In the access control matrix model, the triple of (S, O, A) is the representation of the system state, where S is the set of subjects, O is the set of objects and A is the access matrix, where rows correspond to subjects, columns correspond to objects, and entry $A[s, o]$ reports the rights of s on o .

The access control matrix lies at the heart of all the existing discretionary policy models, even though the languages used to express the policy and approaches to implementation differ. They all assume that the author of the policy has already *predicted* the *needs* of the system and *made a decision* on who (identity, role, property, etc.) and under what condition is to be authorised for a resource. As we will show in the following sections, they only differ in the approach and expressiveness to specify these elements.

4.2.3 Chinese Wall Model

Nash and Brewer [29] proposed the Chinese wall security model to address the needs of financial institutions where information flows may cause conflict of interest. The aim of the Chinese wall policy is to prevent users from accessing the information that is in conflict with any other information that they have already accessed. It combines the free choice element of DAC with mandatory controls by initially allowing a user to choose an object they wish to access, however once an object is accessed, the other objects that may trigger a conflict of interest rule may not be accessed any more.

4.2.4 Role-Based Access Control

To address the management complexities of traditional access control models, Ferraiolo *et al.* [30] proposed the Role-Based Access Control (RBAC) model. The main attraction behind the use of RBAC is that it can

² Note that in the rest of this section when we use the term *access control model* instead of *authorisation model*, this is only to be consistent with the terminology used in the literature.

reflect the internal structure of the organisation for which the system is being designed. RBAC restricts access to a resource based on the business function or role the subject is performing. The permissions to access a resource are then assigned to the appropriate role(s), rather than directly being assigned to subjects' identifiers. Because permissions no longer need to be repeatedly assigned to individual users, RBAC scales much better than the identity based DAC models [24, 31].

RBAC has several advantages that make it the primary choice for the implementation of access control within a centralised system. First, RBAC greatly simplifies the management of the security policy. The administrator grants each user the roles corresponding to their job function within an organisation and when their job changes, the administrator simply changes the roles associated with that user. Second, in several variations of RBAC, roles can be structured as hierarchies which greatly simplify the management task. Third, the least privilege concept can be implemented in RBAC as users can log-in using their least privileged roles and change to the higher privileged ones only as required. Fourth, RBAC can simulate the concept of separation of duty by defining roles that are incompatible and cannot be assigned to the same user (static separation of duty) or concurrently activated (dynamic separation of duty) [10, 32].

Whilst RBAC provides great advantages in comparison to traditional MAC or DAC models, in reality the authorisation model is poorly scalable, as a correct set of roles must still be associated to each potential user and a correct set of permissions must be associated with each role. The need to uniquely identify each potential user within a single administrative domain remains. In addition, since such models have a view that resources belong to the domain rather than individuals within the domain, authorisations are always driven from an administrator to the users and the *delegation* of rights is limited to the roles/identities within the boundary of the administrative domain [33].

4.2.5 Task-Based Authorisation Control

The underlying design of all the above models assumes that all the necessary privileges are available to a subject regardless of the progress of a business function or process. For example, assume there exists a *manager* role within an organisation and *sign-off project* is one of the actions a manager is allowed to perform on an object *contract*. The security administrator creates a policy which authorises the sign-off action, which would be allowed for the manager regardless of the status of the project.

Task-Based Authorisation Control (TBAC) seeks

to model access control from a work-flow perspective rather than the traditional subject-object perspective [34]. TBAC models security and enforcement by considering run-time activities and tasks as they progress from start to completion. To allow such awareness, permissions (P) are constantly monitored and activated/deactivated based on the context of each task. Strictly speaking, in a subject-object access control model, $P \subseteq S \times O \times A$, while TBAC requires information about two additional domains: *usage* (U), and *authorisation-steps* (AS). These permissions are defined by $P \subseteq S \times O \times A \times U \times AS$. These are the additional domains that embed task-based contextual information and draw a distinction between TBAC and other traditional access control models [17, 25]. In TBAC, authorisation steps maintain their own protection state. Each protection state is initialised with a set of valid permissions that become active as a result of the authorisation-step. However, the contents of this set will change as the authorisation-step progresses and the relevant permissions are consumed. There is a limited usage count associated with each permission that will deactivate the permission when the limit is reached and actions are no longer allowed in that state.

The most obvious application of TBAC is in work-flow management, where the granting, usage tracking and revoking of permissions need to be coordinated with the progression of the various tasks. Without such an active model, permissions will, in most cases, be “turned on” too early or too late and will probably remain “on” long after they are needed [17, 25].

Although TBAC proposes an interesting model for access control, it puts a great burden on two components of its authorisation mechanism: First, it requires a very detailed and precise *forecast* on the side of security administrators of the tasks and the necessary permissions as well as conditions and their durations. Second, it demands a detailed monitoring of tasks, which requires very complex and distributed reference monitors that are not widely available currently.

4.2.6 Shortcomings of Traditional Models

While the traditional authorisation models described above address the access control requirements for closed systems, they fall short in several important aspects to provide access control for open distributed systems. The shortcomings were primarily identified by researchers such as Wee *et al.* [35] and Gasser *et al.* [36, 37], and triggered the research on access control approaches for distributed systems.

Some of the weaknesses of these models for providing support for information sharing in distributed systems stem from the fact that such models require that

identities of subjects and objects to be determined before access could be granted [35]. The second problem is the reliance of such models on Access Control Lists (ACL) to express a policy, which is usually stored on a central server under the control of a trusted administrator. However, in distributed systems resources are usually shared between entities spread across multiple administrative domains [38]. The third issue is due to the need of users in distributed systems to delegate some or all of their rights to others in order for tasks to be shared and be completed [36]. The fourth issue is with respect to the questionable assumption of traditional access control models regarding the trustworthiness of the hardware/software of the client machines [39].

These shortcomings have led to research in credential-based authorisation architectures and models that are discussed in the next section.

4.3 Credential-Based Authorisation

4.3.1 Trust Management

With the increasing popularity of Public Key Infrastructure (PKI) and credential based systems the research on authorisations for strangers in open distributed systems has been pursued under the name of *trust management* [40–42]. Trust management in general attempts to address authorisation scenarios where the authoriser and requester do not know each other.

At the core of any trust management system is the *authorisation procedure* that determines whether an access to a resource should be granted or not, based on a number of conditions including a users' capabilities or properties in the form of digital credentials or certificates as well as the authoriser's *local policy* which defines the properties required for an access to be allowed. The semantics of such an authorisation procedure is the main focus of research in trust management as it provides meaning to the features supported for both the authoriser and the access requester [35, 43]. Trust management relies on the formal specification of policies and in this respect, several logics have been used for policy expression and evaluation [44] and several formal trust management frameworks have been introduced [45–47].

The approach of trust management to authorisation is the binding of identities with a set of authorisations referred to as *credentials* which allow the capabilities of identity to be determined and judged based on the relevance of their credentials to the local policy of the resource provider [40, 48]. The trust management model allows every entity to act as an authoriser, a credential issuer, or a requester.

Trust management systems as described by Chapin

et al. [49], are comprised of three major components: *authorisation decision*, *certificate storage and retrieval* and *trust negotiation*. The first component focuses on access control decisions, the second is concerned with the physical location of certificates, credentials and policies as well as the mechanisms involved in acquiring them in order to make authorisation decisions, and the third component attempts to provide the necessary protocols that allow authorisers and requesters to bargain on the required credentials with respect to the access being requested.

The authorisation decision component focuses on providing formalisms for specifying local policy and credentials. Some of the major problems that exist include the expressiveness of the language versus its decidability, and complexity for implementation purposes. The majority of existing formal models are based on three types of formalism: graph theory, logic and relational calculus. Certificate storage and retrieval have been mostly kept away from the context of trust management even though the first and third components strongly depend on this for implementation [49, 50]. Trust negotiation goes beyond the basic model for authorisation in which requesters are assumed to provide all their credentials to the authorisers, trusting them to pick the ones required for access to be permitted. In a more realistic model requesters would like to provide the least number of credentials needed for access to be granted. Further, they may have policies they need the authoriser to abide by, such as privacy constraints over use, disclosure and retention of personal information. The trust negotiation literature aims to provide a framework for negotiating such bargains [51].

There are currently several trust management systems in the literature that mainly focus on the area of authorisation decision or trust negotiation. For example *PolicyMaker* [46], *KeyNote* [52, 53], *REFEREE* [54] and *Binder* [55] are some examples addressing the former problem, while *PeerTrust* [56] and *Portune* [57] focus on the trust negotiation problem.

Although trust management systems provide a comprehensive and interesting authorisation framework for distributed systems, they must still be investigated in terms of their applicability for information sharing environments when information is sent from the realm of the authoriser to the realm of requesters. This is due to the fact that trust management, like other traditional access control models, has focused on protecting digital resources within server systems and does not deal with client-side controls for locally stored digital information. This is the shortcoming that motivated the research on Digital Rights Management.

4.3.2 Digital Rights Management

Digital Rights Management (DRM) is a generic term for a set of technologies and standards for client-side enforcement of access rights [58]. The main goal of DRM is to provide persistent access control, which allows digital content (e.g. music files, video streams, digital books) to be distributed between consumers and to be conditionally accessed using different mediums such as personal computers, mobile devices, etc. [59, 60].

A typical DRM model consists of a *data provider* who holds rights to the content and is the only entity that can create licenses, a *distributor* that is responsible for the distribution of encrypted resources, and the *consumer* who is the user of the resource. The consumer downloads the resource from the distributor and acquires *usage licenses* from data provider. The licenses are then used by the specialised consumer devices to allow controlled access to the resource [60–62].

To build a DRM system three main enabling factors are required: a rights model, a management model and a set of tools [63]. The first deals with modelling the subjects, objects and attributes that need to be considered within a DRM system. The second focuses on introducing the necessary architecture, procedures and protocols to allow resources/licenses to be acquired, distributed, delegated or revoked. The final component attempts to implement trusted clients and management software for a DRM system.

The core concept in DRM relates DRM components to enable the communication of access rights in a *digital license* that bundles the usage rules as well as attributes such as cryptographic keys associated with a digital resource. The rules usually specify a range of restrictions on usage criteria such as no print, no transfer or an expiration date.

Over the past years many rights expression languages have been developed to address the construction of a rights model for DRM systems. The most popular of these languages are the *eXtensible rights Markup Language (XrML)* that was adopted by *MPEG-21* standard [64] and the *Open Digital Rights Language (ODRL)* [65] that was accepted by *Open Mobile Alliance (OMA)*. The major differences between these languages relate to their expressive power to model rights, and the extent to which they are capable of addressing management functions, which in theory is beyond rights modelling but more towards management of rights [66].

As we have mentioned, DRM was introduced to address the access enforcement problem of payment-based systems, which puts more weight on authorisation enforcement rather than the authorisation deci-

sion and requires less expressive authorisation models. DRM was designed for settings where there are some known objects that need to be shared and accessed under some specific conditions rather than the more general problem of whether one should authorise or trust or share the information with a user or not.

Enterprise Rights Management (ERM) is the term used to describe the DRM approach when it is applied to the protection of information in a corporate or enterprise setting. As with DRM, ERM's emphasis is on reliable client enforcement of access policies, not the details of the authorisation model or policy framework [67]. Most ERM proposals deal with client-side enforcement via a client application which is relied on to enforce access policies. Sandhu *et al.* [68] provide an early example. Assurance of the correct behaviour and ongoing integrity of the client application is crucial to the DRM/ERM approach. In the case of a client application running on open computers platform that can also run arbitrary program code, assurance can be provided by a careful combination of a secure operating system, memory isolation features of recently-available processors from Intel and AMD, and a Trusted Platform Module (TPM) based on the Trusted Computing Group specification [68, 69].

DRM and ERM are important for information sharing because of their focus on client-side access enforcement. However, their lack of emphasis on authorisation means the problems with scalability and flexibility in dynamic environments are not addressed.

4.3.3 Usage Based Control

Park and Sandhu [70] proposed a new approach to access control that adopts ideas from traditional access control approaches, trust management and digital rights management. They coined the term *Usage CONtrol* (UCON) to make a clear distinction between the scope of their model and existing ones [43].

A defining feature of UCON is the use of attributes for authorisation. One similarity between traditional access control models and trust management is the use of subject attributes as well as object attributes to produce an authorisation decision. For example, within a MAC model object classification or the clearance level of a subject can be considered an attribute. By the same token, in DAC capability lists could be viewed as subjects' attributes and access control lists (ACL) as object attributes. The second feature of UCON is its consideration of environmental *conditions* for the authorisation decision. For example, employees may be forced to access sensitive resources during business hours at certain locations [70, 71].

Furthermore, UCON allows for a usage decision

to be made conditional on the fulfilment of some prior actions. This characteristic is referred to as an *obligation* and is required in addition to authorisation to enforce a simple form of sequencing of actions in a similar manner to workflow-focused TBAC. For example, consenting to a list of terms and conditions by clicking on a box prior to being given access to a sensitive report is an obligation [43, 70].

In addition to the above three concepts that could be found in the traditional access control literature, UCON introduces two important properties that are referred to as *continuity* and *mutability*.

Continuity requires the ongoing evaluation of usage requirements (e.g. conditions, rights, attributes) while an access is being performed rather than the approach taken by the traditional access control models, in which the act of authorisation is always performed before access.

Mutability allows attributes (e.g. subject/object) to be updated based on the subject's actions. Traditionally, such an update of attributes could only be done by administrators. As Zhang *et al.* [72] describe “in UCON, authorisation decisions are not only checked and made before the access, but may be repeatedly checked during the access and may revoke the access if some policies are not satisfied, according to the changes of the subject or object attributes, or environmental conditions”.

From the architectural point of view UCON supports both traditional access control and trust management models as well as the DRM model to enable server-side access control, and persistent access control on the client-side (after the resources are distributed).

In the past couple of years and particularly in 2008 UCON has gained considerable attention from the access control research community. In terms of languages and formalisms for UCON, Katt *et al.* [73] propose a general obligation model and an enforcement engine, Jamkhedkar *et al.* [66] provide a formalism for expressing rights, Salim *et al.* [74] propose an administrative framework to enable delegation and administration of rights and attributes, and Pretchner *et al.* [75] introduce a formal model for different mechanisms that can enforce usage control policies on the consumer side. From the implementation and architectural perspective, UCON is being used for virtual organisations in data grids [76].

4.4 Risk-Based Authorisation

All the authorisation approaches that we have described so far have one thing in common at their core; they classify actions into two categories: authorised and unauthorised (e.g. good/bad) and try to ensure

this separation is not violated. The problem is that for a complex system with several requirements, the policy that tries to predefine a crisp division is doomed to be either bypassed or be cluttered with too many seemingly ad-hoc exceptions.

As the demand for information sharing has increased, the rigidity of existing approaches has motivated several research proposals which aim to find a more flexible method of authorisation. The work of Hosmer [77, 78] provides an early reference to the need for more flexible approaches in constructing policies in order to bridge the gap between the imprecision that exists in the real world and the precision required by classical logic. Hosmer suggested the use of fuzzy logic [79] to bridge this gap and provided some examples of how the application of fuzzy logic may deliver the required flexibility. More recent proposals that aim to go beyond binary decision making (authorised/unauthorised) focus on the concept of risk. In the rest of this section we will review these recent authorisation approaches which have been directly or indirectly inspired by the work of Hosmer.

4.4.1 Survivability Approach

Survivability research targets systems that operate in highly dynamic environments. The fundamental assumption is that the system has a *mission* that must be completed even if some compromise in system policy enforcement is to be made. Hence, it often involves trade-offs among several functional and non-functional requirements determined by the mission of the system. In other words, the dynamic state of the mission provides the contextual inputs that inform the trade-off decisions.

From a technical perspective the research blends computer security (i.e. as one of the requirements) with business *risk management* [80] and at its core the research departs from being mostly about confidentiality and integrity of information and focuses more on the availability and continuity of service [6, 81].

While the proposal for contextual decision making is one of the primary motivations for our work and will be motivated in more detail in Section 5, the survivability research has so far remained in the realm of software and requirements engineering rather than security and access control. Further, the proposals remain mostly informal and abstract rather than concrete formal models or methodologies.

4.4.2 Optimistic Approach

Another approach that attempts to introduce flexibility into current authorisation models is pursued under the title of *optimistic security*. It is based on a very

important assumption that, regardless of how flexible or expressive authorisation models (policies) are, they will not be able to take into account the dynamic nature of current environments. There will be unforeseen circumstances that are not accounted for in the access policy but which need to be effectively handled [5].

Povey [5] discussed the need and importance of an optimistic authorisation scheme in dynamic environments alongside what he calls existing *pessimistic models*. Povey states that the static nature of current authorisation models can cause unexpected and unjustifiable risks in dynamically changing environments such as disasters, medical emergencies or time-critical events, when unnecessary access restrictions may have catastrophic consequences. He assumes the risk of failure and the cost of recovery is low compared to the cost of not granting access in a given situation. To ensure minimisation of the likelihood and consequences of a user maliciously or inadvertently misusing the system, he proposes that access entries must be constrained; accountability, auditability and recovery must be possible. Povey introduces the concept of a *partially formed transaction* in the Clark Wilson Integrity model [23] that refers to transactions where the integrity of the data is not guaranteed, but where a compensating transaction exists to return the system to a valid state. While he discusses how to introduce such flexibility, the focus is not to discuss or formally justify why and under what situations these otherwise denied requests should be granted.

Ferreira *et al.* [82] have also suggested that traditional authorisation models do not allow overriding. Their domain of interest is healthcare and they motivate an authorisation model that can allow *unanticipated* access to be provided in emergency situations. They have proposed a “Break-The-Glass” policy to allow override whilst providing a non-repudiation mechanism for its usage. Similar to Povey’s proposal, they also assume users have a legitimate need that will actually benefit the whole system. However, their approach is more application-oriented rather than formal, and like Povey, they keep the question of how users would know and decide outside the model.

There are three important assumptions underlying all the above approaches: First there are circumstances where the negative consequences that flow from not granting access may outweigh the potential damage caused by granting access. Second, users are known and dependent on the organisation and can be sanctioned for their unnecessary accesses. Third, users are assumed to be competent to know what is beneficial for the organisation, hence they can judge whether the access must be made. In other words, they take the problem of decision making to the realm of users and

outside the authorisation model. In order to discourage users from making selfish decisions, they assume users can be punished for such behaviour. Moreover, the answer to the question of whether the access was “necessary” assumes monitoring, audit and recovery techniques are in place.

4.4.3 Quantified Risk-Based Approach

More recently, several other approaches have been proposed to address the inflexibility problem of authorisation by defining and estimating the *risk* of granting an access. The MITRE Jason Report [13] studied the requirements of access control for information sharing in government, the intelligence, law enforcement and emergency response community. It notes that organisations deal with inflexible access control systems using various ad-hoc means to share information, such as providing near-blanket access rights or “temporary” authorisations that are never revoked. They suggest there is a need for a parameterisable control that governs a trade-off between security and operational needs. To address this, the report recommends that a new authorisation model must focus on *risk* and they propose a three step procedure for building such a risk-based access control system. The first step is to *measure* risk; as they put it “if you can’t measure it, you can’t manage it”. The second step is to specify the maximum amount of risk for each document. For example, how many copies of a classified document can an organisation afford to lose. The third and the most controversial step is to ensure that information is distributed up to the maximum acceptable risk limit. They also introduce the concept of *risk tokenisation*, where a token is something with exchange value that the holder can trade for access. The tokenisation of risk allows for greater flexibility as it allows limited access to classified information by uncleared users when such access is so important that the holder of a token is willing to pay the price. However, although the research introduces a change in the paradigm of thinking about the authorisation problem, like the previously mentioned approaches, it leaves the decision as to what should be regarded as an important and beneficial access (for the system) to the users. As a result, their approach is still dependent on recovery and audit mechanisms to revert the system if a user’s decision was not aligned to what the system considers beneficial. Further, the idea that pushing information sharing to the highest acceptable risk would imply that maximum sharing of information is an objective for the system, while in reality, information is shared to satisfy an objective, and when there is no known benefit in sharing, there may not be a reason to share even if there are no known risks.

Cheng *et al.* [16, 20] also focus on the inflexibility

of authorisation models for information sharing in dynamic environments, where the set of users with whom the information must be shared depends on external events. They believe that authorisation is a mechanism to manage the risk of leakage of sensitive information by human users (within an organisation), i.e. “to balance the information *needs* of the users in order to perform their jobs with the *need* of the organisation to protect its sensitive information”. Further, they continue that “since the *future needs* and behaviours of users are unpredictable, the authorisation policy is essentially an educated guess that tries to balance *future risks with future needs*”. They blame the core of the problem on the static nature of existing authorisation policies. The educated guesses encoded in the policy will always be imprecise and incomplete in dynamic environments, even if the policy had provisions for pre-specified exceptions, since not all risk vs. benefit trade-offs could have been foreseen by the policy author.

They proposed a Quantified Risk-Adaptive Access Control (QRAAC) for Bell-LaPadulla MLS [26] that attempts to bring these trade-offs into the authorisation model so that exceptions can be granted where the associated risk can be accounted for. Consistent with the MITRE [13] report, their goal is to encourage prudent, calculated risk taking by users to achieve better results while still keeping the overall risk within the organisation’s risk tolerance. In their model, there is a flexible gap between allow and deny. Within this area, transactions could be allowed by using some risk mitigation mechanisms to avoid unaccepted overall risk while increasing information sharing. In their approach a risk is defined as the expected value of loss due to unauthorised disclosure:

$$risk = v \times p$$

where v is the value of the information and p is the probability of unauthorised disclosure. The value of information is defined to be the potential maximum damage sustained if the information is disclosed in an unauthorised manner, the unit of damage being system specific. Determining the probability of unauthorised disclosure is more difficult as it requires predicting future user behaviour. For example, within an MLS system it is intuitive to assume that the probability would be higher when a person without security clearance is given access to top secret information and lower if the person has a top secret clearance. Given this, such probability can be estimated based on two independent components: *temptation*, which is a function of both the subject’s clearance level, that indicates the subject’s trustworthiness, and the object sensitivity level, which indicates the value of the object. Temptation increases as the subject’s trust-

worthiness decreases or the object’s value increases. The second component is *inadvertent disclosure*. This value is represented by the difference in compartment membership between the subject and object. More specifically, subjects are given a fuzzy membership for a category, which indicates the subject’s need for information in that category. They also give a fuzzy membership to objects for that category that determines the relevance of this object to the category. Hence, the *willingness* to share increases as the subject and object membership increase.

While they believe in a trade-off analysis between the risk and benefit of information sharing, they also fall short in providing a comprehensive mechanism to allow such a trade-off to be made. Their concept of benefit is simplistically incorporated in the function that calculates the probability of inadvertent disclosure, within the willingness index. It only assumes one factor and that is the degree of membership of a category. Further, their approach is not general and only focuses on the Bell-LaPadulla model. Notwithstanding these limitations they have proposed a novel approach incorporating the preliminary step of quantifying risk.

Inspired by the MITRE proposal [13], Zhang et.al [15] also introduced a new authorisation model, Benefit And Risk Access Control (BARAC), based on balancing the risk of information disclosure and the benefits of information sharing. One major distinction between this approach and other risk based approaches is the explicit treatment of benefits in BARAC. The authors strongly believe that measuring the benefit of access along with the associated risk is of crucial importance for making an access control decision. The model is composed of subjects, objects, read and update transactions. Each BARAC model has a configuration on which transactions are associated with risk and benefit vectors and some subjects are associated with the risk of being compromised. Further, the configuration defines an allowed transaction graph (AT), that captures allowed transactions and their flow path, as well as an accessibility graph (AC) which describes the accessibility of objects by subjects in terms of the underlying communication system. Finally, they introduce two properties that must be satisfied by the AT graph: *risk cover* that ensures the total system benefit outweighs the total system risk and *weak optimality* which ensures that the AT graph cannot be improved (in terms of benefit vs. risk) by adding or deleting a transaction.

In BARAC benefit and risk are defined based on a multidimensional resource space, where each dimension represents a different component of risk and/or benefit. For example, some of the components of risk

for a risk space may include monetary damage, risk to national strategic interest, or human life. For each of these components they assume an underlying discrete probability distribution, with a finite set of outcomes, each associated with probability and damage. So risk becomes a mathematical expectation of damage for that probability distribution. The benefit is measured in terms of how much one would “pay” for the benefit, in terms of the risk that one is ready to accept. However, since risk and benefit are multi-dimensional, not every two vectors are comparable, hence they consider one vector outweighing another in all dimensions.

One important question that needs to be answered in their model arises from the fact that users are to decide what they are willing to pay for the risk, by actually taking the risk (i.e. paying for it). However, in many cases what a user assesses as beneficial may not actually realise a benefit for the system. This approach also takes the decision making about the risk taking outside the model and into the realm of users of the system. While this could be acceptable for models proposed by Povey [5] and Ferreira *et al.* [82] that assumed several recovery mechanisms, it is not clear how BARAC would deal with this problem.

Molloy *et al.* [14] suggested the field of information security should be viewed as a problem of *risk management*. They define risk as the *expected value of damages* and treat it as a finite resource. Then, damages are the possible outcomes of security decisions and actions. They argue that an access control system is an attempt to model the organisation’s notion of risk and the central issue is where and how much risk to take, which they refer to as the *risk allocation problem*. By taking this approach, benefit and risk based authorisation would directly address the goal of an access control system: to manage the risk of access to sensitive data. They focus on how to cap the aggregated organisational damage while maximising information flow within an organisation. To achieve this they suggest that the organisation must set up a risk token market where it releases a fixed number of risk tokens that can be traded by users amongst themselves using the internal currency issued to them. For a given access request by a user for an object, the access control system determines a risk value quantified in terms of risk tokens. Further, they assume the information objects are accessed to produce benefits which are enumerated in terms of the internal currency. One very important observation is that they assume these benefits to be *context dependent*, evolving over time. Therefore, they cannot be determined *a priori*. Such a dynamic situation implies that the same information object becomes less or more beneficial in different contexts and may cause more or less damage. As a proof of concept they have built a simulation to

show the risk-based authorisation out-performs the existing traditional approaches by increasing information sharing as well as security in Bell-LaPadulla [26] (without compartments).

There have been other proposals in the literature that attempt to incorporate risk with access control. For example, Agrawal [12] mentions that for systems operating in a dynamic environment, the traditional static policy is insufficient and they suggest the need for mechanisms to monitor the overall environment and feed the observations back to the access control system. They assume the authorisation model in the system is already capable of making risk vs. benefit trade-offs, such as one of the above mentioned models. Nissanke and Khayat [83] proposed “risk graphs” to be used to analyse the risk associated with permissions of roles in an RBAC model.

5 Towards an Objective-Based Access Control

In this section we analyse the weakness of current risk based models and outline a new authorisation approach whose defining feature is the inclusion of the concept of multiple competing objectives.

When an access control system or policy interferes with the attainment of legitimate organisational objectives, ad-hoc exceptions and ‘work-arounds’ are commonly introduced to circumvent the restrictions. The MITRE report [13] highlights the problems associated with this response, the chief of which is an uncapped and unknown risk exposure. We believe, based on our analysis of the authorisation literature, that when exceptions to an access control system are rife, they must be regarded as a symptom of the inadequate specification of authorisation policy which itself is a product of a poor *prediction* of the dynamic *incentives* and *forces* that govern the needs of the system.

The research on risk-based authorisation approaches demonstrates the need to focus on another aspect of authorisation theory that has been taken for granted, and that is *a priori* to any policy: that is, how the *system decides* whether to share information or not. Here, we say ‘system decides’ rather than authorisation system to emphasize that the decision is based on several system objectives which evolve according to system needs within a changing environment (context). The primary step for the shift in thinking about authorisation began with the risk-based approaches to authorisation that we have surveyed, where they all attempt to allow *unpredicted exceptions* to the policy to be accounted for. However, these approaches are still entangled with the traditional view, that a comprehensive policy exists that

governs actions. These approaches attacked the inflexibility problem by introducing a grey area within which the access is viable if it stays within a defined risk limit. The underlying motivation of all the above approaches to flexible authorisation is to address the systems' *needs*. This need is abstracted in terms of a *benefit* to be gained from sharing information.

These approaches divide the problem into two aspects: first they attempt to make the authorisation policy flexible by incorporating risk; second, they introduce the concept of *benefit* to justify the risk. However we observe that in practice, the benefits cannot be determined a priori. Therefore, these models must pass the decision making about what the benefits are together with the risk of attempting to secure such benefits, to the users through, for example, providing risk tokens (risk based approaches) [1, 14] or enabling users to ask for exceptional permissions (i.e. optimistic approaches) [5, 82]. However, the limitation of these proposals is that the concept of benefit is not clear. What is considered to be beneficial is assumed to be decided by the users, without reference to what the system recognises as a benefit within a context. Furthermore, the complexity arises when the benefit for the users differs from the system's benefit. We speculate that this gap can be bridged when there is a framework through which the system can define its *needs*. This would essentially provide a reference point to where the compromises are viable and what a benefit is from system's perspective (i.e. given the context).

It is our belief that the shift towards an analytic approach for authorisation is an important one. Like any other decision where there are alternatives, there must be a way to justify the decision. It is our hypothesis that an authorisation framework must allow for *explicit* specification of underlying factors for an authorisation decision rather than just an expressive language for expressing already made decisions as a policy. Policies can perform well if considered as a definition of a single objective. The problem arises when several objectives (requirements) are analysed and distilled into a set of rules of a policy for *future decision making* in a *dynamic environment*.

5.1 Introducing Objectives to Authorisation

The sharing of information and protection of its 'security' are two inherently conflicting objectives of today's interrelated collaborating systems. One fundamental problem is how and based on what factors these needs must be balanced such that the overall objective of the systems is best satisfied. To date, the research in authorisation has worked around this complex problem by assuming the trade-off between these objectives can be made a priori. Figure 1 depicts the

current boundary of the authorisation aspect of access control research.




Figure 1. The boundary of an existing authorisation system

While the above approach reduces the complexity of authorisation significantly and directs its focus to languages and mechanisms for expressing and enforcing policies, it also makes the authorisation system rigid and inflexible. Figure 1 shows that since the trade-off analysis is considered a priori and external, the changing needs (objectives, labelled O in the figure) of the system cannot directly reflect on the authorisation function. The scale of this problem is directly proportional to the dynamism of the environment and the frequency of the changing needs.

To introduce a scalable authorisation framework we propose casting the authorisation problem as a multi-objective decision problem [84], based on the core objectives of the system (O), rules specifying their satisfaction definition (P), and a set of functions that determine the relevance and importance of the objectives for a given context.




Figure 2. The boundary of an objective-based access control system

Figure 2 provides a high level representation of the components of the proposed framework. We believe such an authorisation framework has several interesting properties (e.g. transparency, intuitiveness, scalability, and systematic handling of exceptions) that are missing from existing authorisation models, simply due to their limited boundary.

5.2 Properties of an Objective-Based Access Control

A typical system usually has several objectives that often conflict (e.g. confidentiality and availability). Further, given the context, the importance of these objectives may change. This makes a simple authorisation based on a predefined static policy infeasible for

dynamic environments. As was mentioned in Section 5.1, we propose the explicit inclusion of objectives in the authorisation model to address this problem. We hypothesize that our proposal will have the following desired properties:

- *Transparency*: the objectives (goals/reasons) informing an access decision can be *explicitly* specified.
- *Dynamism*: given changes in the weight and importance of objectives, the posture of the authorisation changes. Note that this is without changing the definition of objectives.
- *Scalability/Incrementally Evolving*: to allow new objectives to be added incrementally and take effect in decision making, without the need to re-evaluate all the existing objectives and procedures for addressing them.
- *Compromising*: to allow some of the explicit objectives (of less importance) be sacrificed for the satisfaction of the objectives of higher importance. This is based on how the trade-off machinery is specified.
- *Opportunistic*: opportunism arises as a consequence of the scalability and compromising capabilities. Opportunism implies that the system can take advantage of a new and unpredictable circumstance which can positively contribute to the satisfaction (increase) of some of the objectives.
- *Justifiability and systematic handling of exceptions*: to be able to explain *why* a specific decision was made based on the objectives. In existing authorisation models, exceptions are usually due to unpredicted circumstances and their existence is to satisfy specific objectives. Hence, their occurrence may be explained with respect to the conflicting objectives and the trade-off analysis of the system. This property may address increasing demands for accountable governance.

6 Research Agenda for Objective-Based Access Control

In this section we outline a research agenda for the exploration of the concept of Objective-Based Access Control. First, we raise some difficult questions that must be answered to fully operationalise the concept. Their difficulty flows from the motivation, inherent in the proposal, to systematically formalise processes of analysis and decision making that are currently carried out by humans based on imprecise and incomplete information. The concept does not require that such complex trade-offs be fully automated though the necessary extent of human involvement is presently unclear. We then propose a set of simplifying assumptions that permit a foothold on a simpler, core set of

problems, which are enumerated.

The very concept of an Objective-Based Access Control system for dynamic environments raises several important and challenging questions. For example:

- How can changes in the environment that affect the information sharing attitude of a system be registered or recognised in or by the system? Furthermore, how can such changes translate into a modification of the relative importance of objectives?
- How can the consequences of a decision be detected and interpreted? The idea here is that every decision has a consequence and since the consequences can be only seen in the future, how should these changes be monitored?
- How can a system learn from the consequences and make more “desirable” decisions in the future?
- How can a system deal with uncertainty about consequences? In other words, how can it make an authorisation decision when it is unclear what the consequences of the access will be?
- How can the relative importance of objectives be negotiated and agreed by multiple, interested but independent decision makers in an information sharing context?

While these are interesting research questions it is clear that they do not have simple answers. Fortunately, we do not believe that it is necessary to address them as a *precondition* of being able to investigate the basic concept of an Objective-Based Access Control. We propose a number of simplifying assumptions to permit a focus on the basic problem of *how* objectives can be explicitly included in authorisation decisions, namely that:

- changes in the environment *do* affect the objectives within the system. However, at least initially, it is not necessary to consider *how* these external changes actually translate to the internal objective changes.
- there is a unitary decision maker who specifies the relative importance of objectives.
- the decision maker is capable of predicting the consequences of modifying the importance of objectives or introducing the new objectives. This simplifies the problem as it permits a focus on the internal changes of objectives rather than the question of whether the authorisation decision turns out to bring about the desired results.

With these simplifying assumptions in place, it should be possible to focus the core concept of an Objective-Based Access Control. The following concepts, languages and mechanisms need to be developed

and are the focus of our work.

- a theory and language for the succinct classification and specification of objectives,
- a related specification for defining the achievement of objectives and the quantification of this achievement,
- theories and mechanisms to allow simple trade-off analysis between the objectives in the system.

We hypothesize that techniques from the discipline of economics related to game theory may provide a useful starting point. Autonomic computing may also provide useful insights.

Building on this fundamental theory, the following opportunities for further investigation are suggested:

- How can the notion of objectives be used within the existing risk-based approaches to authorisation? We conjecture, the explicit specification of objectives and their importance given the context can be used to provide a guideline for the market based approaches in determining what benefits are with respect to the systems' needs. In other words, the proposed framework is to be considered as a central definition of what objectives the system is willing to take the most risk on (highest importance) and given these, users in a market based model can use their risk tokens on the tasks which they believe satisfy the system's need.
- Based on (presently unknown) strengths and limitations, what are the existing environments that the proposed model is most suitable for? One criteria is the frequency of the changes in incentives and forces that prompt the relative changes in objectives of the system. Another is the complexity of objectives.
- What is the assurance level required? Note that most of our focus has been on the flexibility and scalability of authorisation framework. Here, we need to ask the question, whether the framework can be used for applications that need a specific level of assurance with respect to information security goals such as confidentiality, integrity and availability.
- What are the complexity issues with respect to the number of objectives?

The concept of Objective-Based Access Control represents a seismic shift in approach, raising a raft of challenging questions. However, we believe that independent developments in fields relevant to the understanding and modelling of how humans make complex decisions, and successful approaches to formalise some of these decisions, give cause for some optimism that the idea is worth exploring further.

7 Conclusion

The application context of information sharing has revealed an important internal limitation in the current theory of access control: namely the required existence of a static policy which is based on a priori trade-off analysis between competing objectives. Information sharing in pursuit of activities such as emergency response, national security and critical infrastructure protection, occur in an inherently dynamic environment where the opportunities and threats are impossible to predict. Thus, the outcome of a priori trade-off is likely to be sub-optimal and demand a broad rewriting of policy or ad-hoc exceptions in response to actual circumstances. We have proposed the concept of Objective-Based Access Control to address this fundamental problem. Based on our survey of the authorisation literature, we have argued that it is necessary to stretch the boundary of authorisation from a *decision expression* problem that focuses on enriching policy expression languages, to a *decision making* problem. Our proposed Objective-Based Access Control concept aims to provide a *modular* multi-objective authorisation framework that explicitly specifies and considers objectives of a decision problem. The trade-offs between decisions can then be defined based on *needs* to determine an authorisation decision. To do this, we have proposed casting the authorisation problem as a decision problem to answer “why” an access should be granted based on a late trade-off analysis between the explicitly defined objectives. We believe that this will enable the framework to realize two important concepts, *compromise* and *opportunity* necessary to act in unpredicted circumstances where information sharing may become *necessary* or become *desirable* as the changes in the environment introduce new forces or incentives. We hypothesize that this can deliver the necessary flexibility and scalability without compromising increasing important principles of governance where decisions are required to be justified.

Acknowledgements

This research has been partially funded by the Australian Research Council - Project DP0773706.

References

- [1] Xia Zhao and M. Eric Johnson. Information Governance: Flexibility and Control through Escalation and Incentives. In *Proceedings of Seventh Workshop on the Economics of Information Security (WEIS'08)*, Hanover, NH (USA), June 2008.
- [2] Yow Tzu Lim, Pau-Chen Cheng, John Andrew

- Clark, and Pankaj Rohatgi. Policy Evolution with Genetic Programming: A Comparison of Three Approaches. In *IEEE Congress on Evolutionary Computation*, pages 1792–1800, 2008.
- [3] Bob Blakley. The Emperor's Old Armor. In *Proceedings of the 1996 Workshop on New Security Paradigms (NSPW'96)*, pages 2–16, New York, NY, USA, 1996. ACM.
- [4] Dixie B. Baker. Fortresses Built Upon Sand. In *Proceedings of the 1996 Workshop on New Security Paradigms (NSPW'96)*, pages 148–153, New York, NY, USA, 1996. ACM.
- [5] Dean Povey. Optimistic Security: A New Access Control Paradigm. In *Proceedings of the 1999 Workshop on New Security Paradigms (NSPW'99)*, pages 40–45, New York, NY, USA, 2000. ACM.
- [6] Howard F. Lipson and David A. Fisher. Survivability - A New Technical and Business Perspective on Security. In *Proceedings of the 1999 Workshop on New Security Paradigms (NSPW'99)*, pages 33–39, New York, NY, USA, 2000. ACM.
- [7] John and Mary R. Markle Foundation. Mobilizing Information to Prevent Terrorism: Accelerating Development of a Trusted Information Sharing Environment: Third Report of the Markle Foundation Task Force. Technical report, John and Mary R. Markle Foundation, 2006.
- [8] TISN. Trusted Information Sharing Network. [online: <http://www.tisn.gov.au/>].
- [9] J. F. Reid, S. Coronos, E. Dawson, A. McCullagh, and E. Foo. High Assurance Communication Technologies Supporting Critical Infrastructure Protection Information Sharing Networks. In *Proceedings of RNSA Security Technology Conference 2007*, pages 156–167, Melbourne, Australia, September 2007.
- [10] Pierangela Samarati and Sabrina De Capitani di Vimercati. Access Control: Policies, Models, and Mechanisms. In *International School on Foundations of Security Analysis and Design (FOSAD)*, pages 137–196, London, UK, 2001. Springer-Verlag.
- [11] Achille Fokoue, Mudhakar Srivatsa, Pankaj Rohatgi, Peter Wrobel, and John Yesberg. A Decision Support System for Secure Information Sharing. In *Proceedings of the 14th ACM Symposium on Access Control Models and Technologies (SACMAT'09)*, pages 105–114, New York, NY, USA, 2009. ACM.
- [12] Dakshi Agrawal. A New Schema for Security in Dynamic Uncertain Environments. Technical Report RC-24759 A(W0903-025), IBM Research Division, Thomas J. Watson Research Centre, NY 10598, March 2009.
- [13] MITRE Corporation Jason Program Office. Horizontal Integration: Broader Access Models for Realizing Information Dominance. Technical Report JSR-04-132, MITRE Corporation, 2004.
- [14] Ian Molloy, Pau-Chen Cheng, and Pankaj Rohatgi. Trading in Risk: Using Markets to Improve Access Control. In *New Security Paradigms Workshop (NSPW)*, California, USA, 2008.
- [15] Lei Zhang, Alexander Brodsky, and Sushil Jajodia. Toward Information Sharing: Benefit And Risk Access Control (BARAC). In *POLICY*, pages 45–53, 2006.
- [16] Pau-Chen Cheng and Paul A. Karger. Risk Modulating Factors in Risk-Based Access Control for Information in a MANET. Technical Report RC24494, IBM Research Division, Thomas J. Watson Research Center, February 2008.
- [17] Roshan K. Thomas and Ravi S. Sandhu. Task-Based Authorization Controls (TBAC): A Family of Models for Active and Enterprise-Oriented Authorization Management. In *Proceedings of the IFIP TC11 WG11.3 Eleventh International Conference on Database Security XI: Status and Prospects*, pages 166–181, London, UK, UK, 1998. Chapman & Hall, Ltd.
- [18] Elisa Bertino, Piero Andera Bonatti, and Elena Ferrari. TRBAC: A Temporal Role-Based Access Control Model. *ACM Transactions on Information and System Security*, 4(3):191–233, 2001.
- [19] Arun Kumar, Neeran Karnik, and Girish Chafle. Context Sensitivity in Role-Based Access Control. *SIGOPS Operating System Review*, 36(3):53–66, 2002.
- [20] Pau-Chen Cheng, Pankaj Rohatgi, Claudia Keser, Paul A. Karger, Grant M. Wagner, and Angela Schuett Reninger. Fuzzy Multi-Level Security: An Experiment on Quantified Risk-Adaptive Access Control. In *IEEE Symposium on Security and Privacy*, pages 222–230, 2007.
- [21] Keah Choon Tan. A Framework of Supply Chain Management Literature. *European Journal of Purchasing & Supply Management*, 7(1):39 – 48, 2001.
- [22] James P. Anderson. Computer Security Technology Planning Study. Technical Report ESD-TR-73-51, Electronic Systems Division, Hanscom Air Force Base, Mass., 1972.
- [23] David D. Clark and David R. Wilson. A Comparison of Commercial and Military Security Policies. *IEEE Symposium on Security and Privacy*, pages 184–193, April 1987.
- [24] David F. Ferraiolo, Ravi Sandhu, Serban Gavrila, D. Richard Kuhn, and Ramaswamy Chandramouli. Proposed NIST Standard for Role-Based Access Control. *ACM Transactions on Information and System Security*, 4(3):224–274, 2001.

- [25] Roshan K. Thomas and Ravi S. Sandhu. Conceptual Foundations for a Model of Task-Based Authorizations. In *Proceedings of the 7th IEEE Computer Security Foundations Workshop*, Franconia, NH.
- [26] D. Elliott Bell and Leonard J. La Padula. Secure Computer Systems: Mathematical Foundations. Technical report, March 1973.
- [27] Kenneth J. Biba. Integrity Considerations for Secure Computer Systems. Technical Report TR-3153, MITRE Co., technical report, Bedford MA, 1977.
- [28] Butler Lampson. Protection. In *Proceedings of the 5th Annual Princeton Conference on Information Sciences and Systems*, pages 437–443, Princeton University, 1971.
- [29] David F.C. Brewer and Michael J. Nash. The Chinese Wall Security Policy. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 206–214, May 1989.
- [30] David F. Ferraiolo and D.R. Kuhn. Role Based Access Control. *15th National Computer Security Conference*, pages 554–563, Oct 13-16 1992.
- [31] Ravi S. Sandhu. Rationale for the RBAC96 Family of Access Control Models. In *Proceedings of the first ACM Workshop on Role-Based Access Control*, 1995.
- [32] Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-Based Access Control Models. *IEEE Computer*, 29(2):38–47, 1996.
- [33] Ezedin Barka and Ravi Sandhu. Framework for Role-Based Delegation Models. In *Proceedings of the 16th Annual Computer Security Applications Conference (ACSAC'00)*, page 168, Washington, DC, USA, 2000. IEEE Computer Society.
- [34] Myong H. Kang, Joon S. Park, and Judith N. Froscher. Access Control Mechanisms for Inter-organizational Workflow. In *Proceedings of the Sixth ACM Symposium on Access Control Models and Technologies (SACMAT'01)*, pages 66–74, New York, NY, USA, 2001. ACM.
- [35] Thomas Y. C. Woo and Simon S. Lam. A Framework for Distributed Authorization. In *Proceedings of the 1st ACM Conference on Computer and Communications Security (CCS'93)*, pages 112–118, New York, NY, USA, 1993. ACM.
- [36] Morrie Gasser and E. McDermott. An Architecture for Practical Delegation in a Distributed System. In *IEEE Symposium on Security and Privacy*, pages 20–30, 1990.
- [37] Morrie Gasser, Charles Kaufman, J. Linn, Y. Le Roux, and Joseph Tardo. Distributed Authentication Security Service (DASS). In *IFIP Congress (2)*, pages 447–456, 1992.
- [38] Morrie Gasser, Andy Goldstein, Charlie Kaufman, and Butler Lampson. The Digital Distributed System Security Architecture. In *Proceedings of the 12th National Computer Security Conference*, volume NIST/NCSC, 1989.
- [39] Ram Krishnan, Ravi S. Sandhu, and Kumar Ranganathan. PEI Models Towards Scalable, Usable and High-Assurance Information Sharing. In *Proceedings of the 12th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 145–150, 2007.
- [40] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized Trust Management. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, pages 164–173, 1996.
- [41] Amir Herzberg, Yosi Mass, Joris Michaeli, Yiftach Ravid, and Dalit Naor. Access Control Meets Public Key Infrastructure, or: Assigning Roles to Strangers. In *Proceedings of the 2000 IEEE Symposium on Security and Privacy (SP'00)*, page 2, Washington, DC, USA, 2000. IEEE Computer Society.
- [42] Stephen Weeks. Understanding Trust Management Systems. *Security and Privacy, IEEE Symposium on*, 0:0094, 2001.
- [43] Jaehong Park and Ravi S. Sandhu. The UCON_{ABC} Usage Control Model. *ACM Transactions on Information and System Security*, 7(1):128–174, 2004.
- [44] Martín Abadi. Logic in access control. In *Proceedings of the 18th Annual IEEE Symposium on Logic in Computer Science (LICS'03)*, pages 228–233. IEEE Computer Society, June 2003.
- [45] Martín Abadi, Michael Burrows, Butler W. Lampson, and Gordon D. Plotkin. A Calculus for Access Control in Distributed Systems. *ACM Transactions on Programming Languages and Systems*, 15(4):706–734, 1993.
- [46] Matt Blaze, Joan Feigenbaum, and Martin Strauss. Compliance Checking in the Policy Maker Trust Management System. In *Proceedings of the Second International Conference on Financial Cryptography (FC'98)*, pages 254–274, 1998.
- [47] Ninghui Li, J.C. Mitchell, and W.H. Winsborough. Design of A Role-Based Trust-Management Framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 114–130. IEEE Computer Society Press, 2002.
- [48] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. The Role of Trust Management in Distributed Systems Security. In *Secure Internet Programming*, pages 185–210, 1999.
- [49] Peter C. Chapin, Christian Skalka, and X. Sean Wang. Authorization in Trust Management: Features and Foundations. *ACM Computing Surveys*, 40(3):1–48, 2008.

- [50] Elisa Bertino, Elena Ferrari, and Anna Squicciarini. Trust Negotiations: Concepts, Systems, and Languages. *Computing in Science and Engg.*, 6(4):27–34, 2004.
- [51] Kent E. Seamons, Marianne Winslett, Ting Yu, Bryan Smith, Evan Child, Jared Jacobson, Hyrum Mills, and Lina Yu. Requirements for Policy Languages for Trust Negotiation. In *Proceedings of the 3rd International Workshop on Policies for Distributed Systems and Networks (POLICY'02)*, page 68, Washington, DC, USA, 2002. IEEE Computer Society.
- [52] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. KeyNote: Trust Management for Public-Key. In *Infrastructures (Position Paper). Lecture Notes in Computer Science 1550*, pages 59–63, 1999.
- [53] A. Keromytis. The KeyNote Trust-Management System, version 2. *IETF RFC*, 2704:164–173, 1999.
- [54] Yang-Hua Chu, Joan Feigenbaum, Brian A. LaMacchia, Paul Resnick, and Martin Strauss. REFEREE: Trust Management for Web Applications. *Computer Networks*, 29(8-13):953–964, 1997.
- [55] John DeTreville. Binder, a Logic-Based Security Language. *Security and Privacy, IEEE Symposium on*, page 105, 2002.
- [56] Li Xiong and Ling Liu. PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):843–857, 2004.
- [57] Piero A. Bonatti and Daniel Olmedilla. Driving and Monitoring Provisional Trust Negotiation with Metapolicies. In *Proceedings of the Sixth IEEE International Workshop on Policies for Distributed Systems and Networks*, pages 14–23, 2005.
- [58] Qiong Liu, Reihaneh Safavi-Naini, and Nicholas Paul Sheppard. Digital Rights Management for Content Distribution. In *Proceedings of the Australasian Information Security Workshop Conference on ACSW Frontiers 2003 (ACSW Frontiers'03)*, pages 49–58, Darlinghurst, Australia, Australia, 2003. Australian Computer Society, Inc.
- [59] Olin Sibert, David Bernstein, and David Van Wie. DigiBox: A Self-protecting Container for Information Commerce. In *Proceedings of the 1st Conference on USENIX Workshop on Electronic Commerce (WOEC'95)*, pages 15–15, Berkeley, CA, USA, 1995. USENIX Association.
- [60] P.B. Schneck. Persistent Access Control to Prevent Piracy of Digital Information. In *Proceedings of the IEEE*, volume 87 of 7, pages 1239–1250, MRJ Technol. Solutions, Fairfax, VA,; July 1999. IEEE.
- [61] Nicholas Paul Sheppard and Reihaneh Safavi-Naini. Protecting Privacy with the MPEG-21 IPMP Framework. In *Proceedings of 6th Workshop on Privacy Enhancing Technologies*, pages 152–171, 2006.
- [62] Farzad Salim, Nicholas Paul Sheppard, and Reihaneh Safavi-Naini. Enforcing P3P Policies Using a Digital Rights Management System. In *Privacy Enhancing Technologies*, pages 200–217, 2007.
- [63] Xin Wang, Guillermo Lao, Thomas DeMartini, Hari Reddy, Mai Nguyen, and Edgar Valenzuela. XrML – eXtensible rights Markup Language. In *Proceedings of the 2002 ACM Workshop on XML Security (XMLSEC'02)*, pages 71–79, New York, NY, USA, 2002. ACM.
- [64] International Standards Organization. Information Technology - Multimedia Framework (MPEG-21) - part 5: Rights Expression Language. Technical report, ISO/IEC21000-5:2004, 2004.
- [65] Renato Iannella. Open digital rights language (ODRL). Technical report, ISO/IEC21000-5:2004, August 2002.
- [66] Pramod Arvind Jamkhedkar and Gregory L. Heileman. A Formal Conceptual Model for Rights. In *Proceedings of the 8th ACM Workshop on Digital Rights Management (DRM'08)*, pages 29–38, New York, NY, USA, 2008. ACM.
- [67] Muntaha Alawneh and Imad M. Abbadi. Preventing Information Leakage Between Collaborating Organisations. In *Proceedings of the 10th International Conference on Electronic Commerce (ICEC'08)*, pages 1–10, New York, NY, USA, 2008. ACM.
- [68] Ravi Sandhu, Kumar Ranganathan, and Xinwen Zhang. Secure Information Sharing Enabled by Trusted Computing and PEI Models. In *Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security (ASIACCS'06)*, pages 2–12, New York, NY, USA, 2006. ACM.
- [69] Jason F. Reid and William J. Caelli. DRM, Trusted Computing and Operating System Architecture. In *Proceedings of the 2005 Australasian Workshop on Grid Computing and e-Research*, pages 127–136, Darlinghurst, Australia, Australia, 2005. Australian Computer Society, Inc.
- [70] Jaehong Park and Ravi Sandhu. Towards Usage Control Models: Beyond Traditional Access Control. In *Proceedings of the Seventh ACM Symposium on Access Control Models and Technologies (SACMAT'02)*, pages 57–64, New York, NY, USA, 2002. ACM.
- [71] Ravi Sandhu and Jaehong Park. Usage Control:

- A Vision for Next Generation Access Control. In *Lecture Notes in Computer Science*, volume 2776/2003, pages 17–31. Springer Berlin / Heidelberg, 2003.
- [72] Xinwen Zhang, Jaehong Park, Francesco Parisi-Presicce, and Ravi Sandhu. A Logical Specification for Usage Control. In *Proceedings of the Ninth ACM Symposium on Access Control Models and Technologies (SACMAT'04)*, pages 1–10, New York, NY, USA, 2004. ACM.
- [73] Basel Katt, Xinwen Zhang, Ruth Breu, Michael Hafner, and Jean-Pierre Seifert. A General Obligation Model and Continuity: Enhanced Policy Enforcement Engine for Usage Control. In *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT'08)*, pages 123–132, New York, NY, USA, 2008. ACM.
- [74] Farzad Salim, Jason Reid, and Ed Dawson. An Administrative Model for $UCON_{ABC}$. In *Proceedings of the Eight Australasian Information Security Conference (AISC)*, volume 105 of *Conferences in Research and Practice in Information Technology (CRISP)*, pages 32–38, Brisbane, Australia, January 2010. Australian Computer Society (ACS).
- [75] A. Pretschner, M. Hilty, D. Basin, C. Schaefer, and T. Walter. Mechanisms for Usage Control. In *Proceedings of the 2008 ACM Symposium on Information, Computer and Communications Security (ASIACCS'08)*, pages 240–244, New York, NY, USA, 2008. ACM.
- [76] Xinwen Zhang, Masayuki Nakae, Michael J. Covington, and Ravi S. Sandhu. Toward a Usage-Based Security Framework for Collaborative Computing Systems. *ACM Transactions on Information and System Security*, 11(1), 2008.
- [77] Hilary H. Hosmer. Using Fuzzy Logic to Represent Security Policies in the Multipolicy Paradigm. *ACM SIGSAC Review*, 10(4):12–21, 1992.
- [78] Hilary H. Hosmer. Security is Fuzzy!: Applying the Fuzzy Logic Paradigm to the Multipolicy Paradigm. In *Proceedings on the 1992-1993 Workshop on New Security Paradigms (NSPW'92-93)*, pages 175–184, New York, NY, USA, 1993. ACM.
- [79] Lotfi A. Zadeh. Fuzzy Sets. *Information Control*, 8:338–353, 1965.
- [80] Heather M. Hinton. Under-Specification, Composition and Emergent Properties. In *Proceedings of the 1997 Workshop on New Security Paradigms (NSPW'97)*, pages 83–93, New York, NY, USA, 1997. ACM.
- [81] Kevin Sullivan, John C. Knight, Xing Du, and Steve Geist. Information Survivability Control Systems. In *Proceedings of the 21st International Conference on Software Engineering (ICSE'99)*, pages 184–192, New York, NY, USA, 1999. ACM.
- [82] Anna Ferreira, Ricardo Joao Cruz Correia, Luis Antunes, Pedro Farinha, E. Oliveira-Palhares, David W. Chadwick, and Altamiro da Costa Pereira. How to Break Access Control in a Controlled Manner. In *Proceedings of the 19th IEEE Symposium on Computer-Based Medical Systems (CBMS'06)*, pages 847–854, Washington, DC, USA, 2006. IEEE Computer Society.
- [83] Nimal Nissanke and Etienne J. Khayat. Risk Based Security Analysis of Permissions in RBAC. In *Proceedings of 2nd International Workshop on Information Systems*, pages 332–341, 2004.
- [84] Ralph L. Keeney, Howard. Raiffa, and David W. Rajala. Decisions with Multiple Objectives: Preferences and Value Trade-Offs. *IEEE Transactions on Systems, Man and Cybernetics*, 9(7):403–403, July 1979.



Farzad Salim is a PhD candidate at the Information Security Institute - Queensland University of Technology. Prior to this he was a Research Fellow at the School of Computer Science and Software Engineering, University of Wollongong, where he also obtained his MSc in Computer Science. His research interests include access control, DRM, privacy and especially the application of ideas and techniques from the field of economics to information security.



Jason Reid holds a PhD in the area of trusted computing and distributed systems security. Since 1999, he has worked for the Information Security Institute - Queensland University of Technology where he holds the position of Senior Research Fellow. His research interests and areas of expertise include trusted systems, trusted computing and smart card security as well as the privacy implications of these technologies. He has extensive research and consulting experience in such fields as access control, authentication and identity management and general network security.



Ed Dawson is a Professor Emeritus in the Information Security Institute (ISI) at Queensland University of Technology (QUT). Prior to this he was the Research Director of the ISI and has been a member of the academic staff at QUT since 1974. Professor Dawson has extensive research experience in many aspects of cryptology and its applications. He has published over 250 referred research papers, supervised 25 PhD students to completion and received numerous research grants. He is on the editorial board of several journals and has chaired several conferences on various aspects of information security. Currently Professor Dawson is Vice President of the International Association of Cryptologic Research (IACR).