

A Framework for Threat Assessment in Access Control Systems

Hemanth Khambhammettu, Sofiene Boulares, Kamel Adi, and Luigi Logrippo

Laboratoire de Recherche en Sécurité Informatique
Université du Québec en Outaouais, Canada

{hemanth.khambhammettu, bous42, kamel.adi, luigi.logrippo}@uqo.ca

Abstract. We describe a framework for threat assessment specifically within the context of access control systems, where subjects request access to resources for which they may not be pre-authorized. The framework that we describe includes four different approaches for conducting threat assessment: an object sensitivity-based approach, a subject trustworthiness-based approach and two additional approaches which are based on the difference between object sensitivity and subject trustworthiness. We motivate each of the four approaches with a series of examples. We also identify and formally describe the properties that are to be satisfied within each approach. Each of these approaches results in different threat orderings, and can be chosen based on the context of applications or preference of organizations.

Keywords: Security, Access control, Threat assessment

1 Introduction

The “need to share” information in dynamic environments has prompted the development of risk-based access control systems [2–6, 9]. Essentially, in order to facilitate information sharing, risk-based access controls extend traditional access control paradigms to provide support for flexible decision-making by specifying acceptable security risk, operational needs and situational conditions [5]. Risk-based access control mechanisms make access decisions by determining the security risk associated with access requests and weighing such security risk against operational needs together with situational conditions. Specifically, an access request will be *permitted* if the operational benefits outweigh the security risk of granting access to information, and *denied* otherwise.

Clearly, computing the security risk of access requests is an important aspect of risk-based access control systems. However, determining security risk is a complex task, which requires the consideration of a variety of factors, such as the trustworthiness of subjects (or users), sensitivity of data, type of access being requested, access history of subjects and objects, physical or logical location or device from which access to data is being requested as well as protection capabilities and robustness of the system that maintains data [5]. Furthermore,

the interpretation and computation of security risk might differ based on the context of applications or culture of organizations.

The NIST Special Publication (SP) 800-30 [7] is a well-known risk management standard for enterprise systems. In this standard, risk is computed as a product of threat likelihood and impact values.

We adapt the risk assessment function of NIST SP 800-30 to develop a risk assessment function for access control systems. Specifically, the risk of permitting a subject $s \in S$ to perform an action $a \in A$ on an object $o \in O$ is given by the following function:

$$Risk(s, o, a) = Threat(s, o) \times Vulnerability(o) \times Impact(o, a). \quad (1)$$

where $Threat(s, o)$ represents the threat that a subject (threat source) s may present towards an object (threat target) o , $Vulnerability(o)$ represents the weakness within the existing controls for protecting (threat target) o and $Impact(o, a)$ represents the adverse impact on the satisfaction of security objectives that results from successfully performing action a on o .

Note that the NIST SP 800-30 provides no concrete suggestions for estimating threats that subjects present towards data objects. Furthermore, none of the existing work on estimating risk of access requests [1, 3, 6, 9] has explicitly provided approaches to estimate the threats that subjects may present towards data objects, except [2] which will be discussed in Section 3.

In this paper, we focus on estimating the threats that subjects may present towards objects, which are needed to compute risk estimates of access requests. We offer different approaches to estimate such threats which could be applied based on the context of applications.

Consider, for example, a business-to-business scenario that enables organizations to successfully execute their missions, which require subjects (or users) from both intra-business and inter-business to access sensitive data.

- Assume that access requests for sensitive data objects have been initiated by subjects who are employees of the business that owns the requested data objects. In other words, access requests are initiated by subjects who are directly known to (and trusted to some degree by) the system. In such situations, data owners might be more concerned about the sensitivity of data being requested than the trustworthiness of subject. Hence, sensitivity of data objects may be more important than trustworthiness of subjects for estimating the threats posed by subjects towards objects.
- Alternatively, assume that access requests for sensitive objects have been initiated by subjects who are employed by business partners. In other words, subjects who initiate access requests may not be (directly) known to data owners. In such situations, data owners might be greatly concerned about the trustworthiness of subjects for granting access to the requested data objects. Consequently, trustworthiness of subjects may be more important than sensitivity of data objects while estimating the threats posed by subjects towards objects.

Towards this end, we have focused our efforts on developing a suite of threat assessment techniques by considering the sensitivity of objects and trustworthiness of subjects. Of course, as mentioned earlier, we may have to consider a variety of additional factors to compute threat metrics in a comprehensive manner. Nevertheless, even by only considering trustworthiness of subjects and sensitivity of objects, our framework provides significant insights into various ways for assessing the threats posed by subjects towards objects.

We believe that the threat assessment techniques presented in this paper can be used to compute risk metrics by using Formula 1 that is given above. Due to page limitations, we restrict ourselves in this paper to focus on describing approaches which assess threats posed by subjects towards objects and defer the risk assessment of access requests as a subject for our future work.

The following are the main contributions of this paper.

- We present a family of approaches for assessing the threats posed by subjects towards objects: an object sensitivity-based approach, a subject trustworthiness-based approach and two additional approaches which are based on the difference between object sensitivity and subject trustworthiness. We use a series of examples as a basis for developing and identifying the properties of our threat assessment approaches, which provide support for *qualitative* threat assessment of subject-object accesses. Each of these approaches results in different threat orderings, and can be chosen based on the context of applications or preference of organizations.
- We demonstrate that the order in which “general threat principles” (described in Section 2.2) are applied makes a difference in the resulting threat vectors.

The rest of the paper is organized as follows. Section 2 describes our threat assessment framework. In Section 3, we compare our work with notable works of the literature. We draw conclusions for this paper and outline opportunities for future work in Section 4.

2 Threat Assessment Approaches

In this section, we develop our framework for estimating the threats posed by subjects towards objects within the context of access control systems by considering object sensitivity and subject trustworthiness. Throughout this section, we present a series of examples for developing the conceptual underpinnings of our threat assessment approaches.

2.1 Assumptions

We assume the existence of the following entities within an access control system: a set of subjects S and a set of objects O . Furthermore, we assume that every object is associated with a sensitivity score that reflects the protection needs of the data it holds. Typically, sensitivity scores are assigned to data objects by

data owners. A function $ol : O \rightarrow [0, 100]$ formally represents the assignment of sensitivity scores to objects. We also assume that every subject is associated with a trustworthiness score that reflects the trust bestowed upon the subject by the organization that owns the data. Such trustworthiness scores of subjects may be computed either statically by referring to attributes of subjects or dynamically by referring to access histories of subjects or system. A function $sl : S \rightarrow [0, 100]$ represents the assignment of trustworthiness scores to subjects.

We assume a “risk-based” system where a subject labeled to a certain trustworthiness score is *always* permitted to access objects whose sensitivity score is up to the subject’s trustworthiness score. In other words, all accesses initiated by a subject $s \in S$ to an object $o \in O$, such that $sl(s) \geq ol(o)$, will be *permitted*. However, should $sl(s) < ol(o)$ then access decisions are made by the system by computing the risk of granting access to o for s and referring to the risk acceptance level specified within the access control policy. In particular, an access request initiated by a subject $s' \in S$ to object o , such that $sl(s') < ol(o)$, will be *permitted* if the risk of granting access to o for s' is lower than the specified risk acceptance level, and *denied* otherwise.

2.2 Defining “threat”

As discussed earlier in Section 1 and also shown above in Formula 1, threat metrics are a pre-requisite to compute risk metrics. We take the point of view that permitting a subject s to access an object o , such that $sl(s) < ol(o)$, presents by itself a “measurable threat”, independently of what might happen to the information that is accessed. In this section, we define the notion of “threat” in the context of the right to access objects by subjects. In particular, threat of subject-object accesses is defined as follows.

Definition 1 *We say that there exists a threat if a subject $s \in S$ is able to access an object $o \in O$, such that $sl(s) < ol(o)$.*

In other words, any attempt by a subject s to access an object o , such that $sl(s) \geq ol(o)$ does not present a threat.

Intuitively, any measure of threat is affected by one or more of the following three general principles:

- **Principle 1:** Threat increases as object sensitivity score increases.
- **Principle 2:** Threat increases as subject trustworthiness score decreases.
- **Principle 3:** Threat increases as the difference between the object sensitivity score and subject trustworthiness score increases.

We define a function **Threat** : $S \times O \rightarrow [0, 1]$ that represents the threat value of a subject $s \in S$ accessing an object $o \in O$. We use relation \preceq^T to denote an ordering that represents threat on a set of subject-object accesses. In particular, \preceq^T can be defined in terms of subject-object accesses in the following way: $(s, o) \preceq^T (s', o')$ iff **Threat**(s, o) \leq **Threat**(s', o').

The relation \preceq^T allows threats to be compared, and “greater” and “lesser” threats assessed. We define $(s, o) \simeq^T (s', o')$ iff $(s, o) \preceq^T (s', o')$ and $(s, o) \succeq^T (s', o')$.




Fig. 1. Configuration of the running example

(s', o') , and say (s', o') is a greater threat than (s, o) and write $(s, o) \prec^T (s', o')$ if $(s, o) \preceq^T (s', o')$ and $(s, o) \not\approx^T (s', o')$. We may write $(s', o') \succ^T (s, o)$ whenever $(s, o) \prec^T (s', o')$.

2.3 Running example

In this section, we describe the setting of a scenario that is used in the rest of the paper for motivating our threat assessment approaches.

We assume the existence of the following four subjects: Alice, Bob, Carol and Dave. Figure 1(a) illustrates the trustworthiness scores of these four subjects.

Let us consider two objects o and o' within a military context, such that o maintains information regarding the location for nuclear weapons and o' maintains launch codes for nuclear weapons. It is reasonable to assume here that information regarding the launch codes for nuclear weapons is *more sensitive* than the location of nuclear weapons. Hence, we assume that object o is assigned a sensitivity score 90 and object o' is assigned a sensitivity score 100. Figure 1(b) shows the sensitivity scores of objects o and o' .

Recall that the objective of our work is to assess the threat of access requests, where a subject $s \in S$ who initiated the request may not be pre-authorized for the requested data object $o \in O$. Hence, throughout this paper, we only cite examples where $sl(s) < ol(o)$.

2.4 Object-based threat assessment

It is possible that certain applications which maintain “highly” sensitive data, such as government or military systems, may understand or interpret threat in terms of access to such data. We now give examples that motivate our technique for threat assessment that primarily is based on the sensitivity score of objects.

Example 2 Suppose that Alice requests access to object o' , and Bob requests access to object o . We have the following from Figure 1: $sl(\text{Alice}) = 90$, $sl(\text{Bob}) = 80$, $ol(o) = 90$ and $ol(o') = 100$.

If we were to consider object sensitivity score to be the basic criteria for determining threat measures, then according to Principle 1 stated in Section 2.2

6 Lecture Notes in Computer Science: Authors' Instructions

allowing **Alice** to access object o' is a greater threat than allowing **Bob** to access object o . This is simply because the sensitivity score of object o' is higher than the sensitivity score of object o .

In the above example, we were able to understand which access poses a greater threat by simply comparing the sensitivity scores of those two objects. However, as we show below, such a technique is no longer sufficient when object sensitivity scores are the same.

Example 3 Let us extend Example 2 by considering an additional subject **Carol** whose trustworthiness score is given in Figure 1 as follows: $sl(\text{Carol}) = 70$. Suppose that **Carol** requests access to object o' , where $ol(o') = 100$. In other words, both **Alice** and **Carol** request access to object o' .

Now, if we were to determine which of these two accesses is a greater threat, then according to Principle 2 (see Section 2.2) one is likely to conclude that allowing **Carol** to access object o' is a greater threat than allowing **Alice** to access o' . This is because **Carol** who has a trustworthiness score of 70 is less trusted than **Alice** who has a trustworthiness score of 90.

Remark 4 (from Examples 2 and 3) A threat assessment technique that primarily is based on object-sensitivity scores should support the following:

1. always apply Principle 1 (that is, threat always increases as object sensitivity score increases),
2. whenever object sensitivity scores are the same, apply Principle 2 (that is, threat increases as subject trustworthiness score decreases).

Based on Remark 4, we obtain the following ordering of threat for the accesses which were considered in Examples 2 and 3: $(\text{Bob}, o) \prec^T (\text{Alice}, o') \prec^T (\text{Carol}, o')$.

It can easily be seen that we can construct a “priority order” of excessive accesses by subjects for objects, when sensitivity scores are higher than trustworthiness scores, in terms of their threat by adhering to the properties of Remark 4. Essentially, the properties of Remark 4 can be generalized as follows: $(s, o) \prec^T (s', o')$ if either

1. $ol(o) < ol(o')$ or
2. $ol(o) = ol(o')$ and $sl(s') < sl(s)$.

Within an object-based threat assessment approach, whenever object sensitivity scores are the same, unlike Remark 4 we may wish to apply Principle 3 as a secondary criterion. In other words, we may use “the difference of object sensitivity and subject trustworthiness scores” as a secondary parameter, rather than subject trustworthiness scores. Note however that whenever the object sensitivity score is fixed, the difference between object sensitivity and subject trustworthiness scores increases only if subject trustworthiness scores decrease. This means that the threat priority order remains the same irrespective of whether we apply Principle 2 or Principle 3 as a secondary criterion. Hence, we do not describe the subcase that applies Principle 3 as a secondary criterion.

2.5 Subject-based threat assessment

As discussed earlier in Section 1, in certain scenarios (such as business-to-business environments), access requests could be initiated by subjects who may not be (directly) known to data owners. In such situations, trustworthiness of subjects may take higher preference than sensitivity of data objects while estimating access threats.

We now give examples that motivate our technique for threat assessment that, primarily, is based on trustworthiness scores of subjects.

Example 5 *Let us reuse the setting of Example 2 here. That is, we consider subjects Alice and Bob, and suppose that Alice requests access to object o' , and Bob requests access to object o .*

Now, should subject trustworthiness score be the basic criteria for conducting threat assessment, then according to Principle 2 (see Section 2.2) one is likely to conclude that granting access to Bob for object o is a greater threat than granting access to Alice for o' . This is because Bob, who has a subject trustworthiness score of 80, is less trusted than Alice, who has a subject trustworthiness score of 90.

Example 6 *Let us extend Example 5 by considering an additional user Dave where $sl(\text{Dave}) = 80$ (see Figure 1(a)). Now both Bob and Dave have the same trustworthiness score. Suppose that Dave is requesting access to object o' .*

Now, if we were to determine which one of the above two accesses of Bob and Dave poses a greater threat, then according to Principle 1 (see Section 2.2) we may reasonably say that granting access to Dave for o' is a greater threat than granting access to Bob for o . This is because—although both Bob and Dave have the same trustworthiness scores—Dave is requesting access to object o' which has a higher sensitivity score than object o that Bob is requesting access to.

Remark 7 (from Examples 5 and 6) *A threat assessment technique that primarily is based on subject-trustworthiness scores should support the following properties:*

1. always apply Principle 2 (that is, threat increases as subject trustworthiness score decreases),
2. whenever subject trustworthiness scores are the same, apply Principle 1 (that is, threat increases as object sensitivity score increases).

Based on Remark 7, we obtain the following ordering of threat for the subject-object accesses which were considered in Examples 5 and 6: $(\text{Alice}, o') \prec^T (\text{Bob}, o) \prec^T (\text{Dave}, o')$.

Essentially, the properties of Remark 7 can be generalized as follows: $(s, o) \prec^T (s', o')$ if either

1. $sl(s') < sl(s)$ or
2. $sl(s') = sl(s)$ and $ol(o') > ol(o)$.

It is important to note the effect of the basic criterion on threat orderings or metrics when subjects request access to objects, where sensitivity scores are higher than trustworthiness scores. In particular, should the sensitivity score of objects be the basic criterion for assessing threat, then $(\mathbf{Bob}, o) \prec^T (\mathbf{Alice}, o')$ (see Example 2). Whereas, if the trustworthiness score of subjects is considered as the basic criterion for assessing threat, then $(\mathbf{Alice}, o') \prec^T (\mathbf{Bob}, o)$ (see Example 5).

Note that, whenever subject trustworthiness scores are the same, unlike Remark 7 we may wish to apply Principle 3 as a secondary criterion. That is, we may wish to use “the difference of object sensitivity and subject trustworthiness scores” as a secondary parameter, rather than object sensitivity scores. However, note that whenever the subject trustworthiness score is fixed, the difference between object sensitivity and subject trustworthiness scores increases only if object sensitivity scores increase. This means that, within a subject-based threat assessment approach, the threat priority order remains the same irrespective of whether we apply Principle 1 or Principle 3 as a secondary criterion. Hence, we do not describe the subcase that applies Principle 3 as a secondary criterion.

2.6 Difference of scores-based threat assessment

In certain scenarios, we may not be directly concerned with either the object sensitivity scores or subject trustworthiness scores; however, our objective could be to understand threat simply in terms of the difference between object sensitivity and subject trustworthiness scores. Essentially, in such an approach, the degree of threat proportionally increases with the difference between object sensitivity and subject trustworthiness scores.

In this section, we adopt such a notion of threat (as described above) and develop two different techniques for threat assessment which, primarily, are based on the difference between the sensitivity scores of objects and subjects. We first give examples below for motivating our threat assessment techniques and then formalize their properties.

Example 8 *Let us reuse the setting from Examples 2 and 3 here. In particular, we consider user **Bob** from Example 2 and **Carol** from Example 3. As before, we suppose that **Bob** requests access for object o and **Carol** requests access for object o' .*

*Now, should the basic criteria for determining threat measures be the difference between object sensitivity and subject trustworthiness scores, then according to Principle 3 (see Section 2.2) granting access to **Carol** for object o' is a greater threat than granting access to **Bob** for object o . This is because the difference between the sensitivity score of object o and trustworthiness score of **Carol** (which is $100 - 70 = 30$) is greater than the difference between the sensitivity score of object o' and trustworthiness score of **Bob** (which is $90 - 80 = 10$).*

Note, in the above example, that the differences between object sensitivity and subject trustworthiness scores for the two accesses under consideration are

not the same. Hence, we were able to compare the threat of granting accesses by simply computing and comparing the difference between object sensitivity and subject trustworthiness scores of those two subject-object accesses.

We show in the following example that such a technique is no longer sufficient when the difference between object sensitivity and subject trustworthiness scores of subject-object accesses is the same.

Example 9 *Let us extend Example 8 by also considering subject Alice. As before, we suppose that Alice requests access for object o' , where $sl(\text{Alice}) = 90$ and $ol(o') = 100$ (see Figure 1).*

Note that the difference between the sensitivity score of object o' and trustworthiness score of Alice (which is $100 - 90 = 10$) – is the same as – the difference between the sensitivity score of object o and trustworthiness score of Bob (which is $90 - 80 = 10$). Hence, it is not immediately obvious which of the above two subject-object accesses poses a greater threat.

If we were to determine which of the two subject-object accesses considered in Example 9 is a greater threat, then we may choose between applying either Principle 1 or Principle 2 (see Section 2.2) yielding two different approaches, which consider different secondary parameters, for resolving the parity observed in Example 9. These two approaches are described below.

Difference weighted by object sensitivity score In this approach, we consider object sensitivity scores as a secondary criterion and apply Principle 1 which says that threat increases with an increase in object sensitivity scores (see Section 2.2) for resolving the parity observed in Example 9.

This means that, in Example 9, granting access to object o' for Alice is a *greater threat* than granting access to object o for Bob, because $ol(o') > ol(o)$.

Remark 10 *A threat assessment technique that primarily is based on the difference between object sensitivity and subject trustworthiness scores, and that uses object sensitivity scores as a secondary criterion should support the following properties:*

1. *always apply Principle 3 (that is, threat increases as the difference between object sensitivity and subject trustworthiness scores increases),*
2. *whenever parity is observed on the difference between object sensitivity and subject trustworthiness scores, apply Principle 1 (that is, threat increases as object sensitivity score increases).*

Based on Remark 10, we obtain the following ordering of threat for the subject-object accesses which were considered in Examples 8 and 9: $(\text{Bob}, o) \prec^T (\text{Alice}, o') \prec^T (\text{Carol}, o')$.

The properties of Remark 10 can be generalized as follows: $(s, o) \prec^T (s', o')$ if either

1. $(ol(o) - sl(s)) < (ol(o') - sl(s'))$ or
2. $(ol(o) - sl(s)) = (ol(o') - sl(s'))$ and $ol(o') > ol(o)$.

Difference weighted by subject trustworthiness score In this approach, we consider subject trustworthiness scores as a secondary criterion and apply Principle 2 which says that threat increases with a decrease in subject trustworthiness scores (see Section 2.2) for resolving the parity observed in Example 9.

Recall from Example 9 that the difference between the sensitivity score of object o' and trustworthiness score of Alice (which is $100 - 90 = 10$)—*is the same as*—the difference between the sensitivity score of object o and trustworthiness score of Bob (which is $90 - 80 = 10$). In this approach, granting access to object o Bob poses a *greater threat* than granting access to object o' Alice, because $sl(\text{Bob}) < sl(\text{Alice})$.

Remark 11 *A threat assessment technique that primarily is based on the difference between object sensitivity and subject trustworthiness scores, and that uses the subject trustworthiness scores as a secondary criterion should support the following properties:*

1. always apply Principle 3 (that is, threat increases as the difference between object sensitivity and subject trustworthiness scores increases),
2. whenever parity is observed on the difference between object sensitivity and subject trustworthiness scores, apply Principle 2 (that is, threat increases as subject trustworthiness score decreases).

Based on Remark 11, we obtain the following ordering of threat for the subject-object accesses which were considered in Examples 8 and 9: $(\text{Alice}, o') \prec^T (\text{Bob}, o) \prec^T (\text{Carol}, o')$.

The properties of Remark 11 can be generalized as follows: $(s, o) \prec (s', o')$ if either

1. $(ol(o) - sl(s) < ol(o') - sl(s'))$ or
2. $(ol(o) - sl(s) = ol(o') - sl(s'))$ and $sl(s') < sl(s)$.

3 Discussion and Related work

Cheng *et al* have proposed Fuzzy Multi-Level Security (Fuzzy MLS), which quantifies the risk of an access request in multi-level security systems as a product of the value of information and probability of unauthorized disclosure [2]. Fuzzy MLS considers that all subject-object accesses include a temptation to leak information and aims to quantify the risk of “unauthorized disclosure” of information by subjects.

In comparison with Fuzzy MLS, the aim of our framework is to assess the threat posed by subjects towards objects by referring to object sensitivity and subject trustworthiness scores. Although our framework considers simpler requirements than Fuzzy MLS, we have described four different approaches for assessing threat where each approach is biased towards a different set of criteria (unlike Fuzzy MLS whose temptation index is biased only towards object sensitivity scores).

Ni *et al* used fuzzy inference techniques as an approach for estimating access risks and developed an enforcement mechanism for risk based access control [6]. In comparison, we exclusively focus on threat assessment which is a pre-requisite for estimating access risks. Bartsch proposed a policy override calculus for qualitative risk assessment in the context of role-based access control systems [1]. In comparison with the work of Bartsch, our framework is developed in the context of generic access control systems by referring to the sensitivity of objects and trustworthiness of subjects. Diep *et al* described an access control model with context-based decisions that includes quantitative risk assessment [3]. However, there is no description for computing threat measures in [3].

Wang and Jin proposed a method to quantify access risk by considering need-to-know requirements for privacy protection within the context of health information systems [9]. This work exploited the concept of entropy from information theory to compute risk scores of access requests. We believe that our framework could be extended to also consider need-to-know requirements while assessing threats of subject-object accesses. Kandala *et al* developed a framework that captures various components and their interactions in order to develop “abstract models” for RAdAC [4]. However, this work does not consider concrete details of assessing threat or risk.

4 Conclusions

The main contribution of this paper is a framework that includes a family of threat assessment approaches for subject-object accesses, which can be selected based on the context of applications or on the preference of organizations. Specifically, our framework includes four different ways of assessing the threat of subject-object accesses. Our first threat assessment approach, described in Section 2.4, primarily considers the sensitivity scores of objects, and thus gives more priority to the sensitivity of data. We have described another threat assessment approach in Section 2.5 that mainly considers trustworthiness scores of subjects, and thus gives more priority to subject trustworthiness than object sensitivity. A third approach that is based on the idea that threat can be calculated as the difference between object sensitivity and subject trustworthiness scores has been described in Section 2.6, and for this approach we have identified two different subcases by considering the object sensitivity scores and subject trustworthiness scores as secondary parameters.

We have demonstrated that the order in which the “three general threat principles” are applied makes a difference in the resulting threat vectors. This result is important because the approach adopted to assess the threat posed by subjects towards objects will subsequently affect the computation of risk metrics.

To the best of our knowledge, our work represents the first attempt in the literature to conduct a comprehensive study of several alternative approaches for threat assessment by considering object sensitivity and subject trustworthiness scores. We have presented several examples which justify our approaches in intuitive terms.

As mentioned in Section 1, our ultimate goal is to develop a framework for estimating risk of access requests by adopting the well accepted risk assessment function of the NIST SP 800-30. This requires us to extend the work reported in this paper in the following two directions. Firstly, in order to compute impact values of object-action pairs, we intend to exploit “data classification policies” [8]. Subsequently, we will use such impact values together with threat values for quantifying risk of $(subject, object, action)$ triples based on Formula 1. We also aim to use real-world data sets in order to evaluate the efficacy of the threat assessment approaches described in this paper.

Acknowledgements This research was funded in part by grants of the Natural Sciences and Engineering Research Council of Canada and CA Technologies. We thank Serge Mankovski of CA Technologies for having motivated our work.

References

1. S. Bartsch. A calculus for the qualitative risk assessment of policy override authorization. In *Proceedings of the 3rd International Conference on Security of Information and Networks (SIN'10)*, pages 62–70, 2010.
2. P.-C. Cheng, P. Rohatgi, C. Keser, P.A. Karger, G.M. Wagner, and A.S. Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In *Proceedings of IEEE Symposium on Security and Privacy, (SP '07)*, pages 222–230, 2007.
3. N. N. Diep, L. X. Hung, Y. Zhung, S. Lee, Y.-K. Lee, and H. Lee. Enforcing access control using risk assessment. In *Proceedings of the 4th European Conference on Universal Multiservice Networks (ECUMN'07)*, pages 419–424, 2007.
4. S. Kandala, R. Sandhu, and V. Bhamidipati. An attribute based framework for risk-adaptive access control models. In *Proceedings the 6th International Conference on Availability, Reliability and Security (ARES'11)*, 2011.
5. R. McGraw. Risk adaptive access control (RAAdAC). In *Proceedings of NIST & NSA Privilege Management Workshop*, 2009.
6. Q. Ni, E. Bertino, and J. Lobo. Risk-based access control systems built on fuzzy inferences. In *Proceedings of 5th ACM Symposium on Information, Computer and Communications Security (ASIACCS'10)*, pages 250–260, 2010.
7. NIST. Risk management guide for information technology systems. National Institute of Standards and Technology, Special Publication (SP) 800-30, 2002.
8. NIST. Guide for mapping types of information and information systems to security categories. National Institute of Standards and Technology, Special Publication (SP) 800-60, volumes I & II, 2008.
9. Q. Wang and H. Jin. Quantified risk-adaptive access control for patient privacy protection in health information systems. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security (ASIACCS'11)*, pages 406–410, 2011.