in 2007 IEEE Symposium on Security and Privacy

Fuzzy Multi–Level Security : An Experiment on Quantified Risk–Adaptive Access Control[*]

Extended Abstract

Pau–Chen Cheng Pankaj Rohatgi
Claudia Keser Paul A. Karger
IBM Thomas J. Watson Research Center
{*pau,rohatgi,ckeser*}*@us.ibm.com*
karger@watson.ibm.com

Grant M. Wagner Angela Schuett Reninger
US Department of Defense
{*gmw,amschue*}*@tycho.ncsc.mil*

Abstract

This paper presents a new model for, or rather a new way of thinking about adaptive, risk–based access control. Our basic premise is that there is always inherent uncertainty and risk in access control decisions that is best addressed in an explicit way. We illustrate this concept by showing how the rationale of the well–known, Bell–Lapadula model based, Multi–Level Security (MLS) access control model could be used to develop a risk-adaptive access control model. This new model is more like a Fuzzy Logic control system [9] than a traditional access control system and hence the name "Fuzzy MLS". The long version of this paper is published as an IBM Research Report [3].

1. Introduction

Our work is motivated by the fact that many organizations, especially those in the national security and intelligence arena, are unable to rapidly process, share and disseminate large amounts of sensitive information in order to support informed decision making to rapidly respond to external events. A major inhibitor is the inflexibility of current access control models to deal with such dynamic environments and needs. Understanding the significance of isolated events and formulating an effective response may require pooling together of information available within multiple departments and systems (i.e., to *connect the dots*). Clearly, the information that needs to be pooled together would depend on the external events and the analysis approach adopted, and this *cannot be predicted in advance*. Traditional access control policies based on roles aligned with the organizational boundaries are too rigid to allow for such information pooling. This view is reinforced by the JASON Report [8] that investigated barriers to information sharing. The report also paints a disturbing picture of the over-reaction within organizations implicated in failing to "connect the dots". Such organizations have resorted to various ad-hoc means to share information: users have been granted near-blanket access rights or "temporary" authorizations that are never revoked; data have been made more available by understating its sensitivity; a culture has developed along the line of the old saying "it is better to ask for forgiveness rather than for permission" [8]. This has resulted in an unaccountable risk of information leakage.

Our work is geared towards creating and validating a novel, risk based access control model that can revolutionize the way work is conducted in such organizations. The problem is that access control is essentially about balancing risk vs. benefit tradeoffs and existing access control policies specify these tradeoffs *statically* rather than being adaptive to the *dynamic* environments and the needs to which the policies are applied. Thus the ideal case where an organization continually optimizes access control based on risk vs. benefit tradeoffs while capping overall risk cannot be realized. We show how the scenario above can be addressed by making access control much more dynamic and flexible, using a risk management approach based on quantified risk estimates. Essentially, our Fuzzy MLS model quantifies the risk associated with an access; it can even allow risky information flows needed by a user, provided the risk can be accounted for and controlled. The eventual goal is to create a system that encourages information sharing and prudent risk-taking behavior among its users to maximize the benefit to the organization while keeping users accountable for their actions and capping the expected damage an organization could suffer due to sensitive information disclosure. In

[*] This Research is continuing through participation in the International Technology Alliance sponsored by the U.S. Army Research Laboratory and U.K. Ministry of Defense.

0-7695-2848-1/07 $25.00 © 2007 IEEE

addition, the organization will be able to dynamically control risky information flows based on its current operational needs, risk tolerance and environment.

This paper is organized in the following way: section 2 discusses the general idea of quantified risk–adaptive access control, section 3 discusses risk vs. benefit tradeoff, section 4 discusses related work, section 5 presents the Fuzzy MLS model, section 6 discusses on–going and future research.

2. Quantified Risk–Adaptive Access Control

This section discusses *Quantified Risk–Adaptive Access Control* (*QRAAC*). We will first discuss our intuitive interpretation of risk and its relationship to access control, and then expand the discussion into QRAAC.

2.1. Risk in Access Control Decisions

The Merriam–Webster dictionary defines the word *risk* as "*the possibility of loss or injury*". Essentially risk is about some incident that may occur in the future and cause damage. One such risk is the leakage of sensitive information by human users. Access control is one of the mechanisms used to manage this risk, i.e., to balance the information needs of the users in order to perform their jobs with the need of the organization to protect its sensitive information. Since, fundamentally, *the future needs and behaviors of users are unpredictable*, the access control policy is essentially an *educated guess* that tries to balance *future risks with future needs*. For example, in national defense settings, access rights are commensurate with the level of background investigation undergone by a user, yet these investigations are no guarantee for future behaviors; in fact most leaks of classified information are done by people with high security clearance. 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. need tradeoffs could have been foreseen by the policy author. It would be important to bring these *unforseen tradeoffs* into the access control model so that exceptions can be granted in a timely manner and their associated risk is accounted for. This would require a computer system to know when to bend the policy to grant an exception, and the system has to know *how much* the rule would be bent. QRAAC is meant to answer this "how much bending" question; it goes even further and enables the system to take proportional risk mitigation measures [3] to account for and reduce the risk.

2.2. Risk–Adaptive Access Control

In this section we discuss adaptive access control using quantified risk, defined as the *expected value of damage*.

$$quantified\,risk = (probability\;of\;damage) \times (value\;of\;damage) \qquad (1)$$

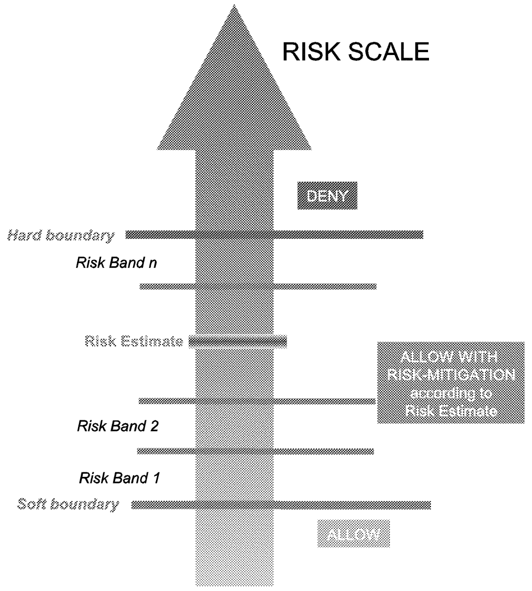

Figure 1. QRAAC on a Risk Scale

The *probability of damage* is the chance that an event happens to incur the damage. The *value of damage* is a quantified measurement of the damage. We do not define the unit of "value" but consider it is the job of a policy writer to determine the proper unit for his/her particular context. Quantifying risk means determining the probability and the value. Due to the *unpredictability of the future*, the probability and the value can at best be *good enough estimates* to compute reasonable *quantified risk estimates*. Using such estimates may seem dangerous, but as we argued above, traditional access control policy are also guesses and imprecise as well. Also, with usage experience, such risk estimates can be fine tuned over time, just as is done in other fields such as the insurance business.

We propose that the existing static access control models with binary "allow/deny" decisions be replaced by a dynamic, multi–decision access control model based on quantified risk estimates and risk tolerance. This model is shown in Figure 1, where the *risk scale* represents the range of quantified risk estimates that is further divided into multiple *bands of risk*. The quantified risk estimate for any access falls into one of these risk bands. Each band is associated with *a decision and an action*; *the decision, the action and band boundaries are all determined according to risk tolerance* and *can be changed when risk tolerance changes*. The top band would be associated with the decision "deny" because the risk is too high; we call the lower bound of the top band the *hard boundary*. The bottom band would be associated with the decision "allow" because the risk is low enough; we call the upper bound of the bottom band the *soft boundary*. A band between the hard and the soft boundaries

0-7695-2848-1/07 $25.00 © 2007 IEEE

can be associated with the decision *allow with risk mitigation measures*, which are actions such as increased auditing, application sandboxing, charging the risk to the user, etc.

Section 5 will first discuss the Fuzzy MLS model as an example of computing quantified risk estimates in an MLS context and then discuss how the model's ideas may be generalized and applied in other contexts.

3. Risk vs. Benefit Trade–off

A primary cause of an access control policy being subverted is that the policy conflicts with individual users' legitimate needs. The QRAAC model addresses this issue by allowing some risk taking when the risk of an access is between the hard and the soft boundaries. An organization's optimal goal should be encouraging prudent, calculated risk taking by users to achieve better results while still keeping the overall risk within the organization's risk tolerance, without micro-managing the human users. Once risk has been quantified, this optimal goal can be achieved in different ways, based on how an organization chooses to influence its user behaviors.

One such approach is similar to a credit card system. Each human user will be given a risk budget as a *line of risk credit* in some units of risk. If a user makes an access whose risk is between the soft and the hard boundaries, then the difference between the risk and the soft boundary (in units of risk) will be charged against the user's risk credit. This charge can be considered *the price paid for "purchasing" exceptional access to information and the necessary risk mitigation measures*. Periodically, the user's *return on investment* (ROI) will be evaluated; the return is the evaluation of the results delivered by the user, and the investment is the amount of risk charged. Greater reward will be given to those users with higher ROI. This process could be part of performance evaluations that an organization already conducts for its employees. A user's line of risk credit could be adjusted based on his/her ROI. The total risk for the organization is always below the sum of all lines of risk credit. Also, each "purchase" will be logged so the users' behaviors can be reviewed and the overall risk–based security policy can be regularly fine-tuned to be more aligned with the actual needs. The lines of credit also provide a means for users to tide over minor conflicts between their needs and the current policy in real–time, i.e., provides flexibility in the short term whereas the fine–tuning process which is to be done off–line adjusts the policy for long–term trends.

Another option would be to create a *market-based mechanism* for users to "trade" quantified risk as a commodity. There will be a finite number of risk units in the market, based on the cap on risk that the organization is willing to accept. As before, exceptional accesses will need to be paid for by the users based on the difference between the risk of access and the soft boundary in risk units. A market could be set up so that users who see an opportunity that requires risk taking could purchase risk units from the market with the hope of realizing tangible rewards from the organization if the opportunity is realized. Users who don't foresee any opportunities could sell their risk units into the market for tangible benefits. Provided the market is set up properly, this mechanism could potentially pool the information and knowledge distributed over the users to optimally allocate the risk to the most promising opportunities [12, 6].

One problem with this approach is that risks associated with multiple actions may not add up linearly. The same *aggregation* problem also exists with most existing access control models, where the combined effect of multiple allowed accesses could create a high risk situation. Chinese–Wall [2] and other history dependent policies are one way to address this problem in traditional models and this may be the way to address the problem even in our model. If risk calculations are made dependent on the history of accesses by the user, such risks could be managed. For example, Chinese–Wall–like constraints could be easily added to the quantified risk model: a user's request to access information in a category would show up as having much higher risk if this user has accessed information in another category that is mutually exclusive with the requested category under the Chinese–Wall policy.

The JASON report [8] also presents some ideas on market mechanism; it discusses the notion of an *access token* which grants access rights to certain kinds of access. The report gives the following example: "*1 token = risk associated with one–day, soft–copy–only access to one document by the average Secret–cleared individual.* " A token associated with a specific kind of access is assigned a value using some common denomination. This allows different tokens, and therefore different access rights, to be traded. So it is more like a barter system; and the report neither presents an uniform way nor a mathematical model to quantify the risk associated with information access or to compute the value of a token.

4. Related Work

Research on risk in access control models, flexible access control models, and risk management in general has been done for many years. We highlight a few recent ones that are related to our work. The JASON report [8] discusses the importance of a risk–based access control system in which the risk is measurable. McDaniel [10] discusses how the context of an access control decision can affect the decision. Nissanke and Khayat [11] analyze the risk associated with permissions assigned to a role in a RBAC system where the risk is assessed by an independent assessment process. None of these works present a way to quantify risk. Dimmoc et al. [4] discuss a computational approach to estimate risk and uses the estimate to make optimal deci-

0-7695-2848-1/07 $25.00 © 2007 IEEE

sions. However, the subjects in their model are autonomous agents, not humans; and it seems that the model requires a prior knowledge of outcomes of all possible combinations of states and actions when a decision is being made. We doubt that such knowledge is obtainable in general.

5. Fuzzy MLS Model

We now discuss the Fuzzy MLS model in detail. Fuzzy MLS was developed as the access control model for humans to access information from IBM System S, an exploratory, very high performance data analysis system, designed and built to continuously analyze a huge amount of input data flow. In particular, the Fuzzy MLS model was developed for the "Brokerage of the Future" scenario, where System S could be used to analyze information for use by stock analysts in a brokerage. The brokerage needs to constantly access, analyze and protect a large amount of sensitive and privileged information and provide the best possible information to its traders and fund managers while managing the risk of sensitive information disclosure. More details of the brokerage scenario and IBM System S are given in [3] and [7].

5.1. Fuzzy MLS: Computing Risk

The rationale for the MLS model is essentially risk based [5] but it suffers from a binary decision model based on risk avoidance [8]. Fuzzy MLS utilizes and extends the underlying risk–based rationale of MLS but changes the access model to be based on risk management. For a *human user's read access*, the risk is defined as the expected value of loss due to unauthorized disclosure:

$$risk \;=\; (value\,of\,information) \times (probability\,of\; unauthorized\,disclosure) \qquad (2)$$

The "value" of information is defined to be the damage sustained if this information is disclosed in an unauthorized manner, where units of damage would be organization specific. Estimating value may appear difficult, but any organization already practicing MLS is expected to assign sensitivity levels to information based on a rough estimate of its value, as prescribed by the principles in [5]. *Typically, sensitivity levels correspond to the order of magnitude of loss* and thus *approximate "value" can be derived from a traditional sensitivity level by an exponential function.*

Determining the probability of unauthorized disclosure requires more work. A precise determination is generally impossible since that would require predicting future user behaviors. Instead, the Fuzzy MLS model strives to develop a way to assign such probabilities that is commensurate with common sense and intuition coming from prior experience with the traditional MLS model. For example, the probability should be very high when a person without security clearance is given access to top secret information but relatively low if the access is given to a person with top secret clearance. The Bell–Lapadula MLS model [1] can be viewed as estimating such a probability P from two *independent* probabilities, P_1 and P_2, and combining them.

$$P_1 = \begin{cases} 0 & \text{human subject clearance level } \geq \\ & \text{object sensitivity level} \\ 1 & \text{otherwise} \end{cases}$$

$$P_2 = \begin{cases} 0 & \text{human subject category set } \supseteq \\ & \text{object category set} \\ 1 & \text{otherwise} \end{cases}$$

$$P \;=\; P_1 + P_2 - P_1 P_2 \qquad (3)$$

The Fuzzy MLS model considers P_1 to be the probability that a human subject (a user) leaks the information by succumbing to *temptation* and P_2 to be the probability of *inadvertent disclosure*, regardless of a subject's intention. The model estimates P_1 and P_2 but they are no longer binary. Thus, the Fuzzy MLS model quantifies the risk that is the concern of the *simple security property* of the Bell–LaPadula model; this property states that a subject can not *read up* and is meant to prevent unauthorized disclosure of information to human subjects. How IBM System S addresses the concern of the *$*$–property*[1] of the Bell–LaPadula Model is discussed in [3].

5.1.1 Computing P_1

A human subject's temptation would be a function of the subject's clearance level (*sl*), which indicates the subject's trustworthiness, and object sensitivity level (*ol*), which indicates the value of the object. Temptation should monotonically increase with respect to *ol* and monotonically decrease with respect to *sl*. Traditional MLS takes a binary view of temptation: no temptation when $ol \leq sl$ and full temptation otherwise. MLS also uses a step function to relate temptation to the probability of disclosure P_1: no disclosure when there is no temptation and disclosure with probability 1 when there is temptation. We take a more nuanced view that all accesses result in temptation, which we quantify by a temptation index TI that varies over a scale. TI is then converted to P_1. There could be many ways to derive TI, which is a function of *sl* and *ol*, but we suggest that any such function should have the following properties:

- Temptation increases as object sensitivity increases or subject trustworthiness decreases.
  $ol_1 > ol_2 \Rightarrow TI(sl, ol_1) > TI(sl, ol_2)$
  $sl_1 > sl_2 \Rightarrow TI(sl_1, ol) < TI(sl_2, ol)$
- TI is always greater than 0.
- TI is biased toward more sensitive objects.

[1] No "write–down".

- The more sensitive an object is, the faster TI increases as sl decreases.
  $ol_1 > ol_2 \Rightarrow$
  $0 > \partial TI(sl, ol_2)/\partial sl > \partial TI(sl, ol_1)/\partial sl$
- For a constant difference $(sl - ol)$, TI increases as ol increases.
  $TI(sl_1, ol_1) > TI(sl_2, ol_2)\ if$
  $ol_1 > ol_2\ and\ (sl_1 - ol_1) = (sl_2 - ol_2)$

As an example formulation for TI, we choose formula 4 below since it is simple, analytic and has all the above properties. Let a be a real number that is greater than 1 and m be a real number that is greater than the maximum allowed value of ol. We further assume that sl and ol are non–negative, then

$$TI(sl, ol) = (a^{-(sl-ol)})/(m - ol) \qquad (4)$$

Here a^{ol} *corresponds to the estimate value of loss as explained in section 5.1*; and a^{sl} *corresponds to the trustworthiness of a human subject* [1, 5] such as "John can be trusted with information worthy of at most $10M". In this formulation TI approaches infinity as ol approaches m. The intuition behind m is that the temptation for a human subject is considered to be too great if an object is as sensitive as m or more sensitive than m, implying that such access control decisions should not be made by machines.

P_1 should monotonically increase with TI. While there could be many different ways to relate TI to P_1, we choose a *sigmoid* function [9] in order to closely parallel the MLS step function approach. P_1 is defined as

$$P_1 = \frac{1}{1 + exp(\,(-k) \times (TI - mid)\,)} \qquad (5)$$

where the parameter mid is the value of TI when P_1 is 0.5 and k determines the slope of the P_1 curve.

5.1.2 Computing P_2

When a human subject has a very strong, legitimate need for information in a category, the organization is more willing to accept the probability of inadvertent disclosure as the usual risk associated with conducting its business. When the subject only has marginal or no need, the organization is less willing to accept the probability. If a subject accesses an object belonging to only one category, P_2 is the difference between the probability of inadvertent disclosure and the probability that the organization is willing to accept for that subject; P_2 is zero if the difference is negative. If the object belongs to multiple categories, we make the *simplifying assumption that the object is a monolithic entity* and compute a difference for each category and use the maximum as P_2.

More research is needed to determine the probability of inadvertent disclosures for a category and an organization's willingness to accept such disclosures. We expect different categories to have different considerations for specifying the probability and willingness [3]. Many formulations for P_2 or even explicit table listings are possible. We are currently experimenting with the following. For a category c, a subject is given a *fuzzy membership* in $[0, 1]$ that indicates the subject's need for information in the category; an object is also given a fuzzy membership that indicates the relevance of this object to the category. Thus the willingness decreases as the subject membership decreases and the object membership increases. The subject and object memberships can be used to compute a *willingness index* using formula 6 where $b > 1$, sm and om are subject and object memberships, and m_{max} is the maximum category membership.

$$wi_c(sm, om) = (b^{-(om-sm)})/(m_{max} - sm), \qquad (6)$$

Formula 6 is similar to 4, but its bias is on the subject membership so that the willingness decreases rapidly as the subject membership decreases. This index can be used in place of TI in formula 5 to compute w_c =*willingness to accept for c*, which is a number in $[0, 1]$.

$$w_c = \frac{1}{1 + exp(\,(-k') \times (wc_i - mid')\,)} \qquad (7)$$

If P_c denotes the probability of inadvertent disclosure for category c,

$$P_2 = Maximum\{\ P_c(1 - w_c) \mid c\ is\ a\ category\ \} \quad (8)$$

5.2. Computing Risk In General

This section discusses how the Fuzzy MLS ideas can be generalized and applied in contexts other than MLS. Our research has been focused on deriving the probability in formula 1. We believe that any organization practicing risk management or MLS has a procedure in place to assess, or at least to classify the values of information [5].

While the true probabilities cannot be accurately derived, it is feasible to estimate the probabilities using the following two–step process, if qualitative comparison between two accesses can be made to determine which access is more likely to result in misuse of the accessed resource.

1. Encoding the comparison using a formula that computes *indices*, such that a larger index implies a higher likelihood. Since the range of indices is not constrained, they could offer a high *resolution* to encode the intuition behind the qualitative comparison. The indices can be put on a scale, and the scale can be *calibrated* such that some points on the scale correspond to real access scenarios. The calibrated scale provides a frame of reference for step 2.

2. Assigning probabilities of misuse to the indices in a way that is commensurate with experience, intuition and threat assessment. These probabilities should increase with the indices. These assignments should be fine tuned over time; but the indices can be kept fixed to make the fine tuning easier. Such probability assignments are guesses, but *all access control policies and decisions are guesses*, as discussed in section 2.1.

Many factors contribute to risk and it may be difficult to design one index formula covering all factors. Such a formula will contain many tunable parameters and be difficult to maintain. We could first design the index and probability formulas for each factor and then divide these factors into smaller groups, such that the relationship among members of a group can be understood or at least conjectured. This will allow a group's joint probability to be computed. Then treat the groups as independent and compute their joint probability. For example, Fuzzy MLS computes the probability through P_1 and P_2 using formula 3.

6. On–Going and Future Research

Besides the "risk market" discussed in section 3, and the fine tuning of the Fuzzy MLS parameters, we believe that quantified risk–adaptive access control (QRAAC) may help address issues that exist in current access control systems. These will be the subjects of further research :

- *Uncertainty in Security Labels*: Security labels, such as MLS labels are assumed to be exact and correct. For example, most MLS systems include the notion of *perfect* secrecy downgraders that sanitize data [13]. In reality, label assignments are not exact and tend to either err on the side of security and be too restrictive or err on the side of convenience and be too loose. If the uncertainty in label assignments could be expressed explicitly, such as probability distribution functions or fuzzy set memberships, labels could be more accurate and QRAAC may be used to make better access control decisions as long as risk can be computed from the "uncertain" labels. This approach also has the potential to address some issues related to the classical *Aggregation Problem* in MLS. The uncertain label of a piece of aggregated information could be skewed toward higher sensitivity than its components.
- *Loss variance based access decisions*: Using probabilistic security labels, both the expected loss and the variance of loss may be computed and used to make access decisions.
- *Risk Modulating Factors*: Many factors other than security labels, such as usages of risk mitigation measures, security of the physical environments, history, properties of information delivery channels can affect risk estimates. Taking these factors into account will result in a more holistic and realistic model for dynamic environments found in mobile settings.

7. Acknowledgment

We thank Shai Halevi, Trent Jaeger, Ronald Perez, Michael Steiner, Douglas L. Schales, Jeff Kravitz and Josyula R. Rao for valuable suggestions.

References

[1] D. E. Bell and L. J. LaPadula. Computer Security Model: Unified Exposition and Multics Interpretation. Technical Report ESD–TR–75–306, The MITRE Corporation, Bedford, MA. HQ Electronic Systems Division, Hanscom AFB, MA, March 1976. http://csrc.nist.gov/publications/history/bell76.pdf.

[2] D. Brewer and M. Nash. The Chinese Wall Security Policy. In *IEEE Symposium on Security and Privacy*, pages 206–214, 1989. Oakland, California.

[3] 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*, 2007. IBM Research Report RC24190.

[4] N. Dimmock, J. Bacon, D. Ingram, and K. Moody. Risk Models for Trust–Based Access Control (TBAC). In *the Third Annual Conference on Trust Management (iTrust 2005)*. Springer–Verlag, May 2005.

[5] INFORMATION SECURITY PROGRAM, DOD 5200.1–R, US Department of Defense, January 1997. http://www.fas.org/irp/doddir/dod/5200-1r/.

[6] F. A. Hayek. The Use of Knowledge in Society. *American Economic Review*, 30:519–530, 1945.

[7] Security for IBM System S, 2006. http://domino.research.ibm.com/comm/research_projects.nsf/pages/system_s_security.index.html.

[8] M. C. Jason Prograrm Office. HORIZONTAL INTEGRATION: Broader Access Models for Realizing Information Dominance, JSR–04–132, December 2004. http://www.fas.org/irp/agency/dod/jason/classpol.pdf.

[9] Jyh–Shing Roger Jang and Chuen–Tsai Sun and Eiji Mizutani. *Neuro–Fuzzy AND Soft Computing: A Computational Approach to Learning and Machine Intelligence*. Prentice Hall, 1997. ISBN 0–13–261066–3.

[10] P. McDaniel. On Context in Autorization Policy. In *SACMAT*, Como, Italy, June 2003.

[11] N. Nissanke and E. J. Khayat. Risk Based Security Analysis of Permissions in RBAC. In *2nd International Workshop on Security in Information Systems*, Porto, Portugal, April 2004.

[12] V. L. Smith. Markets as Economizers of Information: Experimental Examination of the "Hayek Hypothesis". *Economic Inquiry*, 20(2):165–175, 1982.

[13] D. Stork. Downgrading in a Secure Multilevel Computer System: The Formulary Concept. Technical Report MTR 2924, ESD–TR–75–62, The MITRE Corporation: Bedford, MA. HQ Electronic Systems Division, Hanscom AFB, MA, May 1975.