Fuzzy MLS:
An Experiment on Quantified Risk–Adaptive Access Control

Pau–Chen Cheng
pau@us.ibm.com

Pankaj Rohatgi
rohatgi@us.ibm.com

Claudia Keser
ckeser@us.ibm.com

IBM Thomas J. Watson Research Center

January 3, 2007

Abstract

The goal of this paper is to present a new model for, or rather a new way of thinking of adaptive, risk–based access control. Our basic premise is that there is always inherent uncertainty in access control decisions and such uncertainty leads to unpredictable risk that should be quantified and addressed in an explicit way. The ability to quantify risk makes it possible to treat risk as countable resource. This enables the use of economic principles to manage this resource with the goal of achieving the optimal utilization of risk, i.e, allocate risk in a manner that optimizes the risk vs. benefit tradeoff. We choose to expand the well known and practiced Bell–Lapadula multi–level security (MLS) access control model as a proof–of–concept case study for our basic premise. The resulting access control model is more like a Fuzzy Logic control system [Jyh97] than a traditional access control system and hence the name "Fuzzy MLS".

1 Introduction

In today's information and knowledge driven business environment, there is an increasing need to share information across traditional organizational boundaries and with partners to support informed decision making and to rapidly respond to external events, yet sensitive business information must be protected from unauthorized disclosure. Traditional approaches to access control and information security that are aligned with organization charts and roles are not flexible enough to accommodate this new paradigm. Organizations essentially have a choice to either set up a rigid policy that may inhibit necessary sharing or set up ad-hoc controls or provide some users near-blanket access rights, which can result in unaccountable risk of information leakage. Studies such as the JASON Report [JPO04] were explicitly commissioned to investigate barriers to information sharing and have reached a similar conclusion. The problem is due to the fact that existing access control policies specify access decisions *statically* whereas the environments in which the policy is applied are *dynamic*. Thus the ideal case where an organization continually optimizes access control based on risk vs. benefit tradeoffs while capping overall risk cannot be realized.

In this paper, we introduce *Fuzzy MLS*, a new access control model, which in a limited context can be used to quantify risk associated with information access. The ability to quantify risk makes it possible to treat risk that an organization is willing to take as limited and countable resource. This enables the use of a variety of economic principles to manage this resource with the goal of achieving the optimal utilization of risk, i.e, allocate risk in a manner that optimizes the risk vs. benefit tradeoff.

This paper is structured as follows: section 2 discusses the problem with traditional access control models, section 3 presents Fuzzy MLS as a solution in a limited context, section 4 presents the scenario under which the model is developed, section 5 presents the mathematics

of the model, section 6 discusses the prototype implementation of the model and future work, section 7 discusses related work.

2 The Problem

Controlling access to resources is a fundamental security concept through which an organization tries to minimize its exposure to potential damage from mishaps and attacks by limiting *illegitimate* access while optimizing its operations by allowing *legitimate* access. With the advent of computing, access control and access control policy models became a fundamental, well studied and practiced area in computer security and several models such as Lattice Based Access Control (LBAC) [Den76], Role Based Access Control (RBAC)[FKC03], Domain Type Enforcement (DTE)[WSB+96], MLS (multi–level security, a.k.a the Bell Lapadula Model [BL76]), ACLs and Unix file permissions have been invented and deployed.

Given the multitude of policy models one would expect that an organization should be able to select one (or more) of these models to achieve their access control goals. After all, any security model can be used to write a security policy that specifies who gets access to what resources; the different models mostly differ in terms of granularity, expressibility, confinement and manageability properties. Unfortunately, our experience and the experience of other security practitioners [JPO04] suggests otherwise: in many cases, especially for dynamic organizations that have a lot of sensitive data that needs to be shared, the organization's basic need to discriminate between legitimate vs illegitimate access is not met by adopting any of these models.

The inadequacy of these models in this scenario is not a reflection of their lack of expressibility, but rather the fact that when a security administrator creates the policy, she is guessing and codifying what risk-benefit tradeoffs will be acceptable for information accesses that will happen *in the future*. Clearly, for an organization with dynamic needs the future risk-benefit tradeoffs are not predictable and the guesses made about future risk-benefit tradeoffs, encoded in the security policy are likely to be in conflict with the real risk-benefit tradeoffs at the time of access. For a traditional access control policy, these unforseen tradeoffs often result in the *creation of exceptions* to the policy in order to meet practical needs [JPO04]. The creation of these exceptions often needs human approval and is usually time–consuming. Furthermore, exceptions are outside the access control policy and therefore the risk carried by an exception is not accounted for by the policy. This *unaccounted risk defeats the purpose of having an access control policy.*

Thus current access control models that are not adaptive to changing needs are usually successful only in static environments and new, adaptive models are needed for highly dynamic environments. Such models have to be designed so that they can meet the real time needs of the users and of the organization, while bounding the potential damage, even as the needs of the users, the organization and its tolerance for damage varies.

3 Fuzzy MLS: A Solution by Quantifying Risk

While building a general purpose, risk–adaptive access control model appears quite daunting, the Fuzzy MLS model works in at least some settings where the traditional MLS Bell Lapadula model can no longer meet an organization's need for adaptive access control.

The basic premise of the traditional MLS Bell Lapadula model is to determine if a subject is *trustworthy* enough and has the legitimate *need–to–know* to access an object. A subject is usually a person or an application running on behalf of a person. An object is usually a piece of information such as a file. Each subject or object is tagged with a *security label* which is a $<$*sensitivity level, categories set*$>$ tuple. A subject's sensitivity level reflects the degree of trust placed on the subject; a subject's categories set specifies the categories of objects to which the subject has a legitimate need–to–know. An object's sensitivity level indicates how sensitive the object is or the magnitude of the damage incurred by an unauthorized disclosure of the object; an object's categories set specifies the categories to which the object belong. All tuples in a system form a partial–order relation set where $< SL_1, CS_1 > \geq < SL_2, CS_2 >$ if and only if

$SL_1 \geq SL_2$ and $CS_1 \supseteq CS_2$. This relation is called "*dominate*". A subject can read an object only if the subject's label dominated the object's label; this means that the subject is trustworthy enough and has the legitimate need–to–know to read the object. The trustworthiness means that the chance that the subject intentionally leak the information is low. The need–to–know means that there is no unnecessary exposure of the information; such an exposure may lead to an unintentional leakage the risk of which the organization is not willing to accept.

The main feature of Fuzzy MLS is that it considers access control as an exercise in risk management where access control decisions are made on the basis of risk, risk tolerance, and risk mitigation, where risk has the usual connotation of *expected damage*. Viewed in terms of risk, the process of setting a traditional MLS policy is actually determining a fixed tradeoff between the risk of leakage of sensitive information versus the need of the organization to provide such information to its employees for them to perform their job. This fixed tradeoff sets up a non-adaptive, *binary* access control decision model where accesses have been pre-classified as having either acceptable risk or non acceptable risk and only the accesses with acceptable risk are allowed.

Fuzzy MLS devises a way based on the rationale and experience behind MLS to compute *a quantified estimate of risk* associated with a *human subject* reading an object by quantifying the "gap" between the subject's and the object's labels. With these quantified estimates of risk, a *risk scale*, depicted in Figure 1, can be built such that each access is associated with a point on the scale. With such an scale, the access control model can be made risk–adaptive by adjusting

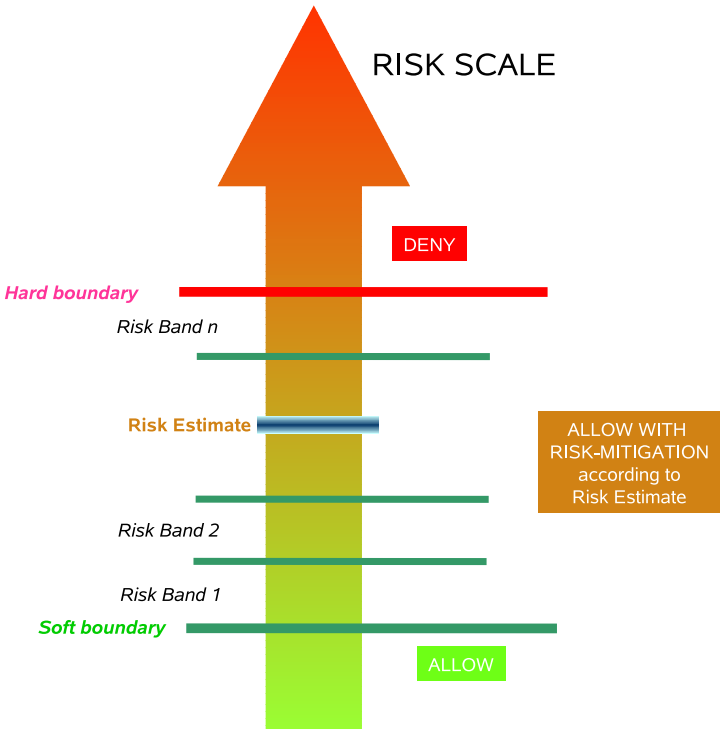

Figure 1: Risk–Adaptive Access Control on a Risk Scale

the point of trade-off on the scale as the needs and environment change. Fuzzy MLS goes one step further by expanding this point of trade–off into a region on the scale. An access associated with a point below the lower-bound of the region (also called the *soft boundary*) is allowed, an access associated with a point above the upper-bound of the region (the *hard boundary*) is denied. The region is further divided into *bands of risk* such that each band is associated with a *risk*

mitigation measure. An access located in a band is allowed only if the risk mitigation measure associated with that band can be applied to the access[1]. Thus, the Fuzzy MLS model depicts a risk management system that resembles a Fuzzy control system [Jyh97] and thus the name "Fuzzy MLS".

A primary cause of a access control policy being subverted is that the policy conflicts with individual users' legitimate needs. The Fuzzy MLS model addresses this issue by allowing some risk taking when the risk of an access is within the region. 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. Fuzzy MLS model has been designed so that this optimal goal can be achieved in different ways based on how an organization chooses to influence its user's behavior.

One such system we propose is similar to a credit card system. Each human user will be given a *line of risk credit* in some units of risk. If a user makes an access whose risk is within the region, 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 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 organizations anyway conduct for their employees. A user's line of risk credit could be adjusted based on his/her ROI. The total risk for the organization's 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 security policy, including the hard and soft boundaries, lines of risk credit, and users' security labels can be regularly fine-tuned to be more aligned with the actual needs. The line of credit also provides 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 system which extends the credit card approach above would be to create a *market-based mechanism* for users to "purchase risk" using a pseudo currency. 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. Each user may be allocated some amount of risk units and pseudo currency initially to get her started, but there would be a market for users to buy and sell risk units for pseudo currency. To motivate prudent-risk taking in such a market setting, a user's contribution to the organization will be evaluated periodically and a score will be given in an amount of pseudo–currency. It is important that pseudo-currency should have direct value to a user, possibly by linking it directly to actual monetary benefit. This way, a user who has knowledge and reason to believe that a particular risky access has a disproportionate chance of yielding benefit compared to current market rate for risk, would be motivated to purchase risk units from the market and pursue the opportunity; whereas users who do not see any good opportunity to use their risk units would be encouraged to sell their risk units in the market to acquire pseudo–currency. This way, it would be possible to aggregate the collective knowledge of the users to optimally allocate the risk towards maximizing the benefit to the organization.

The exact market mechanism to be used and how it should be run (e.g., time periods for risk unit distribution and evaluations etc) would be subjects of further research.

The JASON report [JPO04] also presents some ideas on market mechanism; it discusses the notion of an *access token* which grants access right 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

[1] More discussion on this can be found in the Appendix

it is more like a "my eggs for your milk" barter system; and the report does not present an uniform way nor a mathematical model to quantify the risk associated with information access or to compute the value of a token.

4 Fuzzy MLS : Example Scenario

Consider a (futuristic) brokerage that has set up an information processing system for monitoring and analyzing different data sources such as news reports, trading data as well other data from non-public, sensitive sources. This system is available to its traders, fund managers and brokers through a query interface that produces discrete and streaming results to the user's inquiries. Examples of such inquiries could be "what is the short term pricing trend for security X". Each query can dynamically result in a chain of analysis being performed within the system using public as well as sensitive data that the brokerage has purchased at great cost or whose usage it doesn't want other business rivals to know about, and some of the analysis components themselves could utilize the brokerage's internal secrets such as models of markets. Different queries could dynamically result in different chain of analysis processing being performed on the input data to produce the results. Such a dynamic composition of analysis to respond to different inquiries is possible using planning techniques from AI [RL05, RL06].

For this system, data centric access control model such as MLS with provision for data downgrading is more appropriate than user centric models since results are generated from a dynamic combination of analysis algorithms and data sources and a fundamental requirement is that access restrictions on any result should be easy to determine. With this approach data sources are labeled with sensitivity levels commensurate with the monetary loss if the information is disclosed. For important stocks, categories are used to protect stock-specific sensitive sources and algorithms. E.g., the brokerage estimates production figures for company X (a packaged food supplier) using revenue estimates its packaging supplier which are available from an expensive market intelligence newsletter. The brokerage also has a custom mathematical model for X's stock. Both the data source (newsletter) and the model are protected by the category X. Brokers and traders specializing in the packaged foods industry are cleared for category X and can receive detailed reports about X that include packaging revenues and stock model parameters. Other users not cleared for category X only get limited trend prediction for the stock X by means of downgraders. Valuable sources that can predict multiple stocks are given their own categories and downgraders are used to indicate their influence on each important stock. E.g., a sensitive source may be the daily sales figures for different food items from a major grocery chain. It is assigned category Y and a downgrader can utilize only a part of this data to produce competitive analysis of company X with respect to its peers.

However, any traditional access control model is not suitable in this dynamic environment. When there is a market anomaly, the brokerage would be willing to accept more risk rather than suffer huge losses and would temporarily want to allow wider access to sensitive information. But, when times are good it would want to exercise tighter control on sensitive information to avoid the risk of disclosure. Also, with this setup many traders will have their needs unmet and no satisfactory way to meet them. Consider a hedge fund manager. From time to time the manager needs to make huge short term bets on particular stocks, but requires detailed information before making the bet. In the MLS model, either the hedge fund manager needs to be given access to all major stock categories, giving him unfettered access to most of the brokerage's secrets or given no categories which mean he he gets only sanitized information about stocks and that doesn't serve his needs. Fuzzy MLS can solve both problem by having adjustable hard and soft boundaries that can globally (or just for the traders) be adjusted by security officer based on business conditions. The hedge fund manager will be given partial membership in major stock categories and a risk budget so that, as needed, he can use his budget to get detailed information about any particular stock but that access consumes his budget and gets audited. This way, the hedge fund manager is able to perform his job, but the company can control its risks with respect to how much information the fund manager gets over time and audit records of what information he has accessed.

5 Fuzzy MLS Details : Computing Risk

The rationale for the MLS model was essentially risk based [JPO04] but it suffers from a binary decision model based on risk avoidance. 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:

$$risk \;=\; (value\ of\ information) \;\times\; (probability\ of\ unauthorized\ disclosure) \qquad (1)$$

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 [DoD97]. Typically, sensitivity levels correspond to 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 a precise prediction of future actions of the user. Instead, the Fuzzy MLS model strives to develop a way to assign such probabilities that is commensurate with common sense and intuition which largely comes from prior research done on 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 [BL76] can be viewed as estimating such a probability P from two probabilities P_1 and P_2 and combining them.

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

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

$$P \;=\; P_1 + P_2 - P_1 P_2 \qquad (2)$$

The Fuzzy MLS model also estimates P_1 and P_2 but they are no longer binary.

Computing P_1: We consider P_1 to be the probability that a human subject leaks the information by succumbing to *temptation*. For a human user, the temptation would be a function of user's sensitivity level (sl) which indicates the user'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. MLS takes a binary view of temptation: no temptation when $ol \leq sl$ and temptation otherwise. MLS also uses a step function to relate temptation to 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 which varies over a scale. The probability of leakage due to temptation should monotonically increase with TI. While one could choose different ways to relate TI to probability of disclosure, in order to closely parallel the MLS step function approach we choose a *sigmoid* function to relate P_1 to TI, i.e., P_1 is defined as

$$P_1 \;=\; \frac{1}{1 \;+\; exp(\;(-k) \times (TI - mid)\;)} \qquad (3)$$

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 with regard to TI. There could be countless many ways to derive TI which is a function of sl and ol but we submit that any such function should have the following properties that are consistent with our intuition:

- The more sensitive an object is, the higher the temptation,
  $ol_1 > ol_2 \Rightarrow TI(sl, ol_1) > TI(sl, ol_2)$.

- The more trustworthy a subject is, the lower the temptation,
  $sl_1 > sl_2 \Rightarrow TI(sl_1, ol) < TI(sl_1, ol)$
- TI is always greater than 0. This implies our belief that no human subject is above temptation nor completely trustworthy.
- TI is biased toward more sensitive objects.
  - 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 and some other nice properties as well[2]. 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

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

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 and such access control decisions should not be made by machines. Formula 4 can also be easily related to the Bell–LaPadula model based MLS policy since TI is greater than $1/(m - ol)$ if $sl < ol$, less than $1/(m - ol)$ if $sl > ol$ and equal to $1/(m - ol)$ if $sl = ol$. Thus, with this formula we have that the the Bell–LaPadula model is violated iff TI is greater than $1/(m - ol)$.

Computing P_2: Our intuition for P_2 comes from the probability of *inadvertent disclosure*. This is the probability that a human subject discloses the information *unintentionally*; this kind of "slip of tongue" is always possible once the information is in a human's mind. When a human subject has a very strong, legitimate need–to–know of information in a category, the organization is more willing to accept this probability as the usual risk associated with conducting its business. When the subject only has marginal or no need–to–know, 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 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 difference as P_2.

More research is needed to determine the probability of inadvertent disclosure for a category. We are currently experimenting with the following formulation to compute P_2. For a category c, a subject is given a fuzzy membership in $[0, 1]$ that indicates the subject's need–to–know for information in the category; an object is also given a fuzzy membership that indicates the relevance of this object to the category. The subject and object membership can be used as ol and sl in formula in 4 to compute a "*willingness index*". This index can be used in place of TI in formula 3 to compute w_c ="willingness to accept for c" which is a number in $[0, 1]$. If P_c denotes the probability of inadvertent disclosure for category c,

$$P_2 \;=\; Maximum\{\; P_c(1 - w_c) \mid c \; is \; a \; category \;\} \tag{5}$$

Computing Risk: For a given subject and object label, the value of the object is computed from its sensitivity level, the probabilities P_1 and P_2 are computed as above, the probability of disclosure is computed using formula 2 and finally the risk is computed using formula 1.

6 Implementation, Extentions and Future Work

We have experimented with Fuzzy MLS on a prototype with 10 levels and a few categories with fuzzy subject membership to gain experience in setting the risk calculation parameters. Fuzzy

[2] more information provided in the Appendix

MLS is now being implemented in a larger prototype system and will undergo adjustments as further experience is gained from its usage, especially with how risk estimates and risk budgets should be managed. In this prototype, the core data processing is done using the MLS model whereas requests for exceptional access by users are processed using Fuzzy MLS. We use a lattice alteration strategy to bridge between the two models. I.e., the results are moved to a new temporary lattice point introduced within the MLS lattice where they can be picked up by the user's process operating at this lattice point. Even at this stage, some advantages of Fuzzy MLS are apparent which may help address issues that arise in current MLS systems. These will be the subject of further research, we briefly discuss these ideas.

- *Label Uncertainty*: In an MLS system labels are assumed to be correct. Also most MLS systems include the notion of *perfect* secrecy downgraders that sanitize data [STH85, SRS+00]. This situation makes data difficult to share because human labelers and downgraders err on side of higher secrecy. If labels were uncertain or probabilistic to account for the uncertainty in ascertaining the right secrecy level, then the situation of over-classification could be addressed. But traditional MLS model cannot make access decisions with labels having uncertainty. However this is not a problem with Fuzzy MLS since it can still compute the risk associated with the access and make the decision based on risk.
- *Loss variance based access decisions*: Fuzzy MLS can be easily extended so that both the expected loss and the variance of loss can be used to make access decisions; users may have variance in their trustworthiness and data may have variance with respect to their secrecy and these can be combined to compute risk and variance of loss for making access decisions.
- *Aggregation Problem*: This problem has not been satisfactorily resolved in MLS systems, even in the simplest form, where a sequence of allowed accesses to less sensitive data results in a collection of data that is more sensitive than the individual items. With Fuzzy MLS, the aggregation problem gets exposed, as each user access to data incurs a risk and multiple accesses should accumulate risk. We are exploring ways in which label uncertainty can be used to address the aggregation problem. If individual data items have label X but collectively have higher secrecy label Y, one idea is to assign the individual data item an uncertain label which indicates that it has a small probability of having the label Y. Repeated access to such objects would then incur an aggregate risk comparable to an access to Y.

7 Related Work

Research on risk in access control models, flexible access control models, and risk management in general have been done for many years. We highlight a few recent ones that are more related to our work. The JASON report [JPO04] discusses the importance for a risk–based access control system in which the risk is measurable. McDaniel [McD03] discussed how the context of an access control decision can affect the decision. Nissanke and Khayat [NK04] analyzed the risk associated with permissions assigned to a role in an RBAC system but the risk is assessed by an independent assessment process. None of the works presented a way to quantify risk. Dimmoc et al. [DBIM05] discussed a computational approach to estimate risk and uses the estimate to make optimal decisions. 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 and we doubt if such knowledge is obtainable in general.

References

[BL76] David E. Bell and Leonard 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, 6

[CC03] Suresh N. Chari and Pau-Chen Cheng. BlueBox: A Policy–Driven, Host–Based Intrusion Detection System. *ACM Transactions on Information and System Security*, 6(2), May 2003. 11

[DBIM05] Nathan Dimmock, Jean Bacon, David Ingram, and Ken Moody. Risk Models for Trust–Based Access Control (TBAC). In *the Third Annual Conference on Trust Management (iTrust 2005)*. Springer–Verlag, May 2005. 8

[Den76] Dorothy E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236 – 243, May 1976. ISSN 0001–0782. 2

[DoD97] INFORMATION SECURITY PROGRAM, DOD 5200.1–R, US Department of Defense, January 1997. http://www.fas.org/irp/doddir/dod/5200-1r/. 6

[FKC03] David F. Ferraiolo, D. Richard Kuhn, and Ramaswamy Chandramouli. *Role–Based Access Control*. Artech House Publishers, April 2003. ISBN 1580533701. 2

[JPO04] MITRE Corporation Jason Programrm 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. 1, 2, 4, 6, 8

[Jyh97] 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. 1, 4

[McD03] Patrick McDaniel. On Context in Autorization Policy. In *SACMAT*, Como, Italy, June 2003. 8

[NK04] Nimal Nissanke and Etienne J. Khayat. Risk Based Security Analysis of Permissions in RBAC. In *2nd International Workshop on Security in Information Systems*, Porto, Portugal, April 2004. 8

[RL05] Anton Riabov and Zhen Liu. Planning for Stream Processing Systems. In *The Twentieth National Conference on Artificial Intelligence*, AAAI 2005. 5

[RL06] Anton Riabov and Zhen Liu. Scalable Planning for Distributed Stream Processing Systems. In *The International Conference on Automated Planning and Scheduling (ICAPS)*, 2006. 5

[SRS+00] Gerhard Schellhorn, Wolfgang Reif, Axel Schairer, Paul Karger, Vernon Austel, and David Toll. Verification of a Formal Security Model for Multiapplicative Smart Cards. In *6th European Symposium on Research in Computer Security (ESORICS 2000)*, 2000. 8

[STH85] R. R. Schell, T. F. Tao, and M. Heckman. Designing the GEMSOS security kernel for security and performance. In *the 8th National Computer Security Conference*, pages 108–119. DoD Computer Security Center and National Bureau of Standards, 1985. 8

[WSB+96] Kenneth M. Walker, Daniel F. Sterne, M. Lee Badger, Michael J. Petkac, David L. Shermann, and Karen A. Oostendrop. Confining Root Programs with Domain and Type Enforcement. In *the 6th USENIX Security Symposium*, July 1996. 2

A Temptation Index Formula: Motivation and Example

In Section 5, we proposed the following formula to compute the Temptation Index TI.

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

The basic motivation for this formula follows from experience with MLS. We expect the value of an object to increase exponentially with the sensitivity level. Also, the way the user clearance process works today, subject labeled to a certain sensitivity level are allowed accesses to objects upto that level. This means that a measure of subject trustworthiness also increases exponentially with the subject level and at a rate which is probably commensurate with the rate at which object value increases. So a basic formula TI' where

$$TI'(sl, ol) \;=\; a^{-(sl-ol)}$$

captures the basis intuition in MLS that the temptation from an exponentially rising object value can be balanced by exponentially rising subject sensitivity level. TI' has the property that when its above 1 access is denied. But it does not conform to our intuition that given equal subject and object levels, temptation should increase with higher object level. Therefore we tweaked TI' to create TI which creates a temptation bias towards higher value objects.

Here we show a table of temptation indices (Table 1) and the corresponding probabilities (Table 2). The main point to make here is that temptation are indices usually fairly large or fairly low except the cases where the subject and object levels are close. This is the place where calculated risk taking should be allowed.

$ol \backslash sl$	1	2	3	4	5
1	$1.000e-01$	$1.000e-02$	$1.000e-03$	$1.000e-04$	$1.000e-05$
2	$1.111e+00$	$1.111e-01$	$1.111e-02$	$1.111e-03$	$1.111e-04$
3	$1.250e+01$	$1.250e+00$	$1.250e-01$	$1.250e-02$	$1.250e-03$
4	$1.429e+02$	$1.429e+01$	$1.429e+00$	$1.429e-01$	$1.429e-02$
5	$1.667e+03$	$1.667e+02$	$1.667e+01$	$1.667e+00$	$1.667e-01$
6	$2.000e+04$	$2.000e+03$	$2.000e+02$	$2.000e+01$	$2.000e+00$
7	$2.500e+05$	$2.500e+04$	$2.500e+03$	$2.500e+02$	$2.500e+01$
8	$3.333e+06$	$3.333e+05$	$3.333e+04$	$3.333e+03$	$3.333e+02$
9	$5.000e+07$	$5.000e+06$	$5.000e+05$	$5.000e+04$	$5.000e+03$
10	$1.000e+09$	$1.000e+08$	$1.000e+07$	$1.000e+06$	$1.000e+05$

$ol \backslash sl$	6	7	8	9	10
ol	6	7	8	9	10
1	$1.000e-06$	$1.000e-07$	$1.000e-08$	$1.000e-09$	$1.000e-10$
2	$1.111e-05$	$1.111e-06$	$1.111e-07$	$1.111e-08$	$1.111e-09$
3	$1.250e-04$	$1.250e-05$	$1.250e-06$	$1.250e-07$	$1.250e-08$
4	$1.429e-03$	$1.429e-04$	$1.429e-05$	$1.429e-06$	$1.429e-07$
5	$1.667e-02$	$1.667e-03$	$1.667e-04$	$1.667e-05$	$1.667e-06$
6	$2.000e-01$	$2.000e-02$	$2.000e-03$	$2.000e-04$	$2.000e-05$
7	$2.500e+00$	$2.500e-01$	$2.500e-02$	$2.500e-03$	$2.500e-04$
8	$3.333e+01$	$3.333e+00$	$3.333e-01$	$3.333e-02$	$3.333e-03$
9	$5.000e+02$	$5.000e+01$	$5.000e+00$	$5.000e-01$	$5.000e-02$
10	$1.000e+04$	$1.000e+03$	$1.000e+02$	$1.000e+01$	$1.000e+00$

Table 1: TI values, $m = 11.0,\ a = 10.0$

$ol \backslash sl$	1	2	3	4	5
1	$5.215e-02$	$4.788e-02$	$4.747e-02$	$4.743e-02$	$4.743e-02$
2	$1.314e-01$	$5.271e-02$	$4.793e-02$	$4.748e-02$	$4.743e-02$
3	$9.999e-01$	$1.480e-01$	$5.340e-02$	$4.799e-02$	$4.748e-02$
4	$1.000e+00$	$1.000e-00$	$1.720e-01$	$5.431e-02$	$4.808e-02$
5	$1.000e+00$	$1.000e+00$	$1.000e-00$	$2.086e-01$	$5.555e-02$
6	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e-00$	$2.689e-01$
7	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e-00$
8	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e+00$
9	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e+00$
10	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e+00$	$1.000e+00$

$ol \backslash sl$	6	7	8	9	10
1	$4.743e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$
2	$4.743e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$
3	$4.743e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$
4	$4.749e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$
5	$4.818e-02$	$4.750e-02$	$4.743e-02$	$4.743e-02$	$4.743e-02$
6	$5.732e-02$	$4.834e-02$	$4.752e-02$	$4.743e-02$	$4.743e-02$
7	$3.775e-01$	$6.009e-02$	$4.857e-02$	$4.754e-02$	$4.744e-02$
8	$1.000e-00$	$5.826e-01$	$6.497e-02$	$4.895e-02$	$4.758e-02$
9	$1.000e+00$	$1.000e+00$	$8.808e-01$	$7.586e-02$	$4.974e-02$
10	$1.000e+00$	$1.000e+00$	$1.000e+00$	$9.991e-01$	$1.192e-01$

Table 2: Probability for TI values in Table 1, $k = 1.0$, $mid = 3.0$

B Risk Mitigation

In Section 3, we introduced the notion of risk credit as currency for purchasing risk mitigation measures for risky access. We describe possible risk mitigation measures in more detail here. Since a subject cannot be made more trustworthy instantly, risk mitigation measures are geared towards making the subject less likely to disclose information. Such measures usually fall into the following categories : *deterrence*, *prevention* and *limiting damage* which are discussed below.

- **Deterrence**: provide (strong) disincentives for wrong doings. For example, detailed auditing and mandatory computer or human review may be used to ensure that risky accesses are made for the right reasons. This could also set the stage for administrative or legal actions.
- **Prevention**: To prevent a user process (not the user) from inadvertently disclosing information, one can insist on that the user process runs in a special environment which has extra physical and/or logical security. For example, exceptional accesses may require the process to run within a specific secure location or on a specific trustworthy system, or the process could be sandboxed [CC03].
- **Limiting Damage**: to assume that bad things will happen and take precaution measures to limit the potential damage. Examples are limiting the output rate of information flow to a user/user process, reduced scheduling priority, etc. Another measure is to further restrict the user's future access to resources based on the already granted access.