[1] E. Alata, M. Dacier, Y. Deswarte, M. Kaaâniche, K. Kortchinsky, V. Nicomette, V. H. Pham, and F. Pouget. Collection and analysis of attack data based on honeypots deployed on the Internet. In *1st Workshop on Quality of Protection*, volume 23 of *Advances in Information Security*, pages 79-91. Springer, September 2005. [bib | DOI]

The CADHo project (Collection and Analysis of Data from Honeypots) is an ongoing research action funded by the French ACI “Securité & Informatique” [1]. It aims at building an environment to better understand threats on the Internet and also at providing models to analyze the observed phenomena. Our approach consists in deploying and sharing with the scientific community a distributed platform based on honeypots that gathers data suitable to analyze the attack processes targeting machines connected to the Internet. This distributed platform, called *Leurré.com* and administrated by Institut Eurecom, offers each partner collaborating to this initiative access to all collected data in order to carry out statistical analyzes and modeling activities. So far, about thirty honeypots have been operational for several months in twenty countries of the five continents. This paper presents a brief overview of this distributed platform and examples of results derived from the data. It also outlines the approach investigated to model observed attack processes and to describe the intruders behaviors once they manage to get access to a target machine.

Keywords: analysis, attacks, collection, data, honeypot, internet, leurrecom

[2] L. F. C. an Simson Garfinkel, editor. *Security and Usability*, chapter Security Administration Tools and Practices, pages 374-393. Theory In Practice. O'Reilly, August 2005. [bib]

[3] A. Årnes, K. Sallhammar, K. Haslum, T. Brekne, M. E. G. Moe, and S. J. Knapskog. Real-time risk assessment with network sensors and intrusion detection systems. In *Computational Intelligence and Security*, volume 3802 of *Lecture Notes in Computer Science*, pages 388-397, 2005. [bib | DOI]

This paper considers a real-time risk assessment method for information systems and networks based on observations from networks sensors such as intrusion detection systems. The system risk is dynamically evaluated using hidden Markov models, providing a mechanism for handling data from sensors with different trustworthiness in terms of false positives and negatives. The method provides a higher level of abstraction for monitoring network security, suitable for risk management and intrusion response applications.

Keywords: Risk management ; Abstraction ; Hidden Markov model ; Confidence ; Sensor array ; Intrusion detection systems ; Risk assessment ; Risk analysis ; Surveillance ; Monitoring ; Measurement sensor ; Data handling ; Hidden Markov models ; Information network ; Information system ; Computer security ; Real time ; Artificial intelligence

[4] A. Atzeni and A. Lioy. Why to adopt a security metric? a brief survey. In *1st Workshop on Quality of Protection*, volume 23 of *Advances in Information Security*, pages 1-12. Springer, September 2005. [bib | DOI]

No doubt that computer security is a hot topic nowadays: given the importance of computer-assisted activities, protection of computer system is of the utmost importance. However we have insofar failed to evaluate the actual security level of a system and thus to justify (either in technical or economical terms) the investments in security. This paper highlights the motivations to improve security measurement techniques, analyses the existing approaches, and discusses whether their are appropriate or some new directions should be explored.

Keywords: security metric, computer system security

[5] S. Axelsson. The base-rate fallacy and its implications for the difficulty of intrusion detection. In *Proceedings of the 6th ACM conference on Computer and communications security (CCS)*, pages 1-7, 1999. [bib | DOI]

Many different demands can be made of intrusion detection systems. An important requirement is that it be effective i.e. that it should detect a substantial percentage of intrusions into the supervised

system, while still keeping the false alarm rate at an acceptable level.This paper aims to demonstrate that, for a reasonable set of assumptions, the false alarm rate is the limiting factor for the performance of an intrusion detection system. This is due to the base-rate fallacy phenomenon, that in order to achieve substantial values of the Bayesian detection rate, P(Intrusion|Alarm), we have to achieve (a perhaps unattainably low) false alarm rate.A selection of reports of intrusion detection performance are reviewed, and the conclusion is reached that there are indications that at least some types of intrusion detection have far to go before they can attain such low false alarm rates.

[6] S. Axelsson. The base-rate fallacy and the difficulty of intrusion detection. *ACM Transactions on Information and System Security*, 3(3):186-205, August 2000. [bib | DOI]

Many different demands can be made of intrusion detection systems. An important requirement is that an intrusion detection system be effective; that is, it should detect a substantial percentage of intrusions into the supervised system, while still keeping the false alarm rate at an acceptable level. This article demonstrates that, for a reasonable set of assumptions, the false alarm rate is the limiting factor for the performance of an intrusion detection system. This is due to the base-rate fallacy phenomenon, that in order to achieve substantial values of the Bayesian detection rate P(Intrusion|Alarm), we have to achieve a (perhaps in some cases unattainably) low false alarm rate. A selection of reports of intrusion detection performance are reviewed, and the conclusion is reached that there are indications that at least some types of intrusion detection have far to go before they can attain such low false alarm rates.

[7] I. Balepin, S. Maltsev, J. Rowe, and K. Levitt. Using specification-based intrusion detection for automated response. In *Recent Advances in Intrusion Detection (RAID)*, volume 2820 of *Lecture Notes in Computer Science*, pages 136-154, Sept. 2003. [bib | DOI]

One of the most controversial issues in intrusion detection is automating responses to intrusions, which can provide a more efficient, quicker, and precise way to react to an attack in progress than a human. However, it comes with several disadvantages that can lead to a waste of resources, which has so far prevented wide acceptance of automated response-enabled systems. We feel that a structured approach to the problem is needed that will account for the above mentioned disadvantages. In this work, we briefly describe what has been done in the area before. Then we start addressing the problem by coupling automated response with specification-based, host-based intrusion detection. We describe the system map, and the map-based action cost model that give us the basis for deciding on response strategy. We also show the process of suspending the attack, and designing the optimal response strategy, even in the presence of uncertainty. Finally, we discuss the implementation issues, our experience with the early automated response agent prototype, the Automated Response Broker (ARB), and suggest topics for further research.

[8] D. Balfanz, G. Durfee, D. Smetters, and R. Grinter. In search of usable security: five lessons from the field. *Security & Privacy, IEEE*, 2(5):19-24, Sept.-Oct. 2004. [bib | DOI]

A new system reduces the time to enroll in a secure wireless network by two orders of magnitude, and it also gets high marks for usability and user satisfaction. This article provides a real-world example revealing five general lessons for usable, secure system design.

Keywords: computer network management, human computer interaction, public key cryptography, security of data, telecommunication security, wireless LAN PKI, enrolling time, public key infrastructure, secure system design, secure wireless network, usability, usable security, user satisfaction

[9] D. Balfanz, D. K. Smetters, P. Stewart, and H. C. Wong. Talking to strangers: Authentication in ad-hoc wireless networks. In *Network and Distributed System Security Symposium (NDSS)*, February 2002. [bib | .html]

In this paper we address the problem of secure communication and authentication in ad-hoc wireless networks. This is a difficult problem, as it involves bootstrapping trust between strangers. We present a user-friendly solution, which provides secure authentication using almost any established public-key-based key exchange protocol, as well as inexpensive hash-based alternatives. In our approach, devices exchange a limited amount of public information over a privileged side channel, which will

then allow them to complete an authenticated key exchange protocol over the wireless link. Our solution does not require a public key infrastructure, is secure against passive attacks on the privileged side channel and all attacks on the wireless link, and directly captures users' intuitions that they want to talk to a particular previously unknown device in their physical proximity. We have implemented our system in Java for a variety of different devices, communication media, and key exchange protocols.

[10] M. Cukier, R. Berthier, S. Panjwani, and S. Tan. A statistical analysis of attack data to separate attacks. In *International Conference on Dependable Systems and Networks (DSN)*, pages 383-392, June 2006. [bib | DOI]

This paper analyzes malicious activity collected from a test-bed, consisting of two target computers dedicated solely to the purpose of being attacked, over a 109 day time period. We separated port scans, ICMP scans, and vulnerability scans from the malicious activity. In the remaining attack data, over 78% (i.e., 3,677 attacks) targeted port 445, which was then statistically analyzed. The goal was to find the characteristics that most efficiently separate the attacks. First, we separated the attacks by analyzing their messages. Then we separated the attacks by clustering characteristics using the K-Means algorithm. The comparison between the analysis of the messages and the outcome of the K-Means algorithm showed that 1) the mean of the distributions of packets, bytes and message lengths over time are poor characteristics to separate attacks and 2) the number of bytes, the mean of the distribution of bytes and message lengths as a function of the number packets are the best characteristics for separating attacks

Keywords: computer crime, data analysis, data mining, pattern clustering, statistical analysisICMP scans, K-Means algorithm, attack data statistical analysis, attack separation, data mining, port scans, vulnerability scans

[11] J. E. Gaffney, Jr and J. W. Ulvila. Evaluation of intrusion detectors: A decision theory approach. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, pages 50-61, 2001. [bib | DOI]

This paper presents a method of analysis for evaluating intrusion detection systems. The method can be used to compare the performance of intrusion detectors, to evaluate performance goals for intrusion detectors, and to determine the best configuration of an intrusion detector for a given environment. The method uses a decision analysis that integrates and extends ROC (receiver operating characteristics) and cost analysis methods to provide an expected cost metric. We provide general results and illustrate the method in several numerical examples that cover a range of detectors operating that meet a performance goal and two actual detectors operating in a realistic environment. We demonstrate that, contrary to common advice, the value of an intrusion detection system and the optimal operation of that system depend not only on the system's ROC curve, but also on cost metrics and the hostility of the operating environment as summarized by the probability of intrusion. Extensions of the method are outlined, and conclusions are drawn.

[12] C. Gates and C. Taylor. Challenging the anomaly detection paradigm: a provocative discussion. In *Proceedings of the 2006 workshop on New security paradigms (NSPW)*, pages 21-29, 2006. [bib | DOI]

In 1987, Dorothy Denning published the seminal paper on anomaly detection as applied to intrusion detection on a single system. Her paper sparked a new paradigm in intrusion detection research with the notion that malicious behavior could be distinguished from normal system use. Since that time, a great deal of anomaly detection research based on Denning's original premise has occurred. However, Denning's assumptions about anomalies that originate on a single host have been applied essentially unaltered to networks. In this paper we question the application of Denning's work to network based anomaly detection, along with other assumptions commonly made in network-based detection research. We examine the assumptions underlying selected studies of network anomaly detection and discuss these assumptions in the context of the results from studies of network traffic patterns. The purpose of questioning the old paradigm of anomaly detection as a strategy for network intrusion detection is to reconfirm the paradigm as sound or begin the process of replacing it with a new paradigm in light of changes in the operating environment.

[13] C. E. Gates. A case study in testing a network security algorithm. In *Proceedings of the 4th International Conference on Testbeds and research infrastructures for the development of networks & communities (TridentCom)*, pages 1-6, 2008. [bib]

Several difficulties arise when testing network security algorithms. First, using network data captured at a router does not guarantee that any instances of the security event of interest will be captured. Similarly, if the event of interest is not detected, this does not guarantee that it does not exist in the captured data. Further, such network data is often not publicly available, making comparisons with other detectors difficult. On the other extreme, purely simulated data can be made publicly available and can provide guarantees that the event of interest exists in the data set. However, simulated data often has unintended artifacts and may also incorporate the biases of the particular simulator. In this paper I describe an emulation approach that takes advantage of captured data while using the DETER network to generate realistic traffic for the event of interest. The problem domain was described in terms of seven variables, where the DETER network provided a flexible medium for examining the complete problem domain. The results of a set of experiments using this approach are provided, along with regression equations that describe the expected true and false positive rates.

[14] G. Gu, A. A. Cárdenas, and W. Lee. Principled reasoning and practical applications of alert fusion in intrusion detection systems. In *Proceedings of the 2008 ACM symposium on Information, computer and communications security (ASIACCS)*, pages 136-147, 2008. [bib | DOI]

It is generally believed that by combining several diverse intrusion detectors (i.e., forming an IDS ensemble), we may achieve better performance. However, there has been very little work on analyzing the effectiveness of an IDS ensemble. In this paper, we study the following problem: how to make a good fusion decision on the alerts from multiple detectors in order to improve the final performance. We propose a decision-theoretic alert fusion technique based on the likelihood ratio test (LRT). We report our experience from empirical studies, and formally analyze its practical interpretation based on ROC curve analysis. Through theoretical reasoning and experiments using multiple IDSs on several data sets, we show that our technique is more flexible and also outperforms other existing fusion techniques such as AND, OR, majority voting, and weighted voting.

[15] G. Gu, P. Fogla, D. Dagon, W. Lee, and B. Skorić. Measuring intrusion detection capability: an information-theoretic approach. In *Proceedings of the 2006 ACM Symposium on Information, computer and communications security (ASIACCS)*, pages 90-101, 2006. [bib | DOI]

A fundamental problem in intrusion detection is what metric(s) can be used to objectively evaluate an intrusion detection system (IDS) in terms of its ability to correctly classify events as normal or intrusive. Traditional metrics (e.g., true positive rate and false positive rate) measure different aspects, but no single metric seems sufficient to measure the capability of intrusion detection systems. The lack of a single unified metric makes it difficult to fine-tune and evaluate an IDS. In this paper, we provide an in-depth analysis of existing metrics. Specifically, we analyze a typical cost-based scheme [11], and demonstrate that this approach is very confusing and ineffective when the cost factor is not carefully selected. In addition, we provide a novel information-theoretic analysis of IDS and propose a new metric that highly complements cost-based analysis. When examining the intrusion detection process from an information-theoretic point of view, intuitively, we should have less uncertainty about the input (event data) given the IDS output (alarm data). Thus, our new metric, C_{ID} (*Intrusion Detection Capability*), is defined as the ratio of the mutual information between the IDS input and output to the entropy of the input. C_{ID} has the desired property that: (1) It takes into account all the important aspects of detection capability naturally, i.e., true positive rate, false positive rate, positive predictive value, negative predictive value, and base rate; (2) it objectively provides an intrinsic measure of intrusion detection capability; and (3) it is sensitive to IDS operation parameters such as true positive rate and false positive rate, which can demonstrate the effect of the subtle changes of intrusion detection systems. We propose C_{ID} as an appropriate performance measure to maximize when fine-tuning an IDS. The obtained operation point is the best that can be achieved by the IDS in terms of its intrinsic ability to classify input data. We use numerical examples as well as experiments of actual IDSs on various data sets to show that by using C_{ID}, we can choose the best (optimal) operating point for an IDS and objectively compare different IDSs.

[16] K. Haslum, A. Abraham, and S. Knapskog. DIPS: A framework for distributed intrusion prediction and prevention using hidden markov models and online fuzzy risk assessment. In *Third International Symposium on Information Assurance and Security (IAS)*, pages 183-190, August 2007. [bib | DOI]

This paper proposes a Distributed Intrusion Prevention System (DIPS), which consists of several IPS over a large network (s), all of which communicate with each other or with a central server, that facilitates advanced network monitoring. A Hidden Markov Model is proposed for sensing intrusions in a distributed environment and to make a one step ahead prediction against possible serious intrusions. DIPS is activated based on the predicted threat level and risk assessment of the protected assets. Intrusions attempts are blocked based on (1) a serious attack that has already occurred (2) rate of packet flow (3) prediction of possible serious intrusions and (4) online risk assessment of the assets possibly available to the intruder. The focus of this paper is on the distributed monitoring of intrusion attempts, the one step ahead prediction of such attempts and online risk assessment using fuzzy inference systems. Preliminary experiment results indicate that the proposed framework is efficient for real time distributed intrusion monitoring and prevention.

[17] M. Jahnke, C. Thul, and P. Martini. Graph based metrics for intrusion response measures in computer networks. In *Proceedings of the 32nd IEEE Conference on Local Computer Networks (LCN)*, pages 1035-1042, October 2007. [bib | DOI]

This contribution presents a graph based approach for modelling the effects of both attacks against computer networks and response measures as reactions against the attacks. Certain properties of the model graphs are utilized to quantify different response metrics which are well-kown from the pragmatic view of network security officers. Using these metrics, it is possible to (1) quantify practically relevant properties of a response measure after its application, and (2) estimate these properties for all available response measures prior to their application. The latter case is the basis for the selection of an appropriate reaction to a given attack. Our graph-based model is similar to those used in software reliability analysis and was designed for a scalable granularity in representing properties of the network and its components to be protected. Different examples show the applicability of the model and the resulting metric values.

[18] A. D. Keromytis. Characterizing self-healing software systems. In *Proceedings of the 4th International Conference on Mathematical Methods, Models and Architectures for Computer Networks Security (MMM-ACNS)*, pages 22-33, September 2007. [bib | DOI | .pdf]

The introduction of self-healing capabilities to software systems could offer a way to alter the current, unfavorable imbalance in the software security arms race. Consequently, self-healing software systems have emerged as a research area of particular interest in recent years. Motivated by the inability of traditional techniques to guarantee software integrity and availability, especially against motivated human adversaries, self-healing approaches are meant to complement existing approaches to security.

In this paper, we provide a first attempt to characterize self-healing software systems by surveying some of the existing work in the field. We focus on systems that effect structural changes to the software under protection, as opposed to block-level system reconfiguration. Our goal is to begin mapping the space of software self-healing capabilities. We believe this to be a necessary first step in exploring the boundaries of the research space and understanding the possibilities that such systems enable, as well as determining the risks and limitations inherent in automatic-reaction schemes.

Keywords: Self-healing, reliability, availability, software security

[19] C. Leita and M. Dacier. SGNET: a distributed infrastructure to handle zero-day exploits. Technical Report RR-07-187, Institut Eurecom, France, February 2007. [bib | http]

This work builds upon the Leurre.com infrastructure and the Scriptgen technology. Leurré.com is a worldwide distributed setup of low interaction honeypots whereas Scriptgen is a new class of honeypot: a medium interaction one. In this paper, we see how Scriptgen can be enriched thanks to the Argos and Nepenthes open source software in order to build a distributed system able to collect rich information about ongoing attacks and to collect malware, even for zero-day attacks, without

facing the same liability and complexity issues encountered by classical high interaction honeypots. The design is precisely exposed as well as its implementation. Experimental results are offered that highlight the validity of the proposed solution.

[20] C. Leita and M. Dacier. SGNET: A worldwide deployable framework to support the analysis of malware threat models. In *Seventh European Dependable Computing Conference (EDCC)*, pages 99-109, May 2008. [bib | DOI]

The dependability community has expressed a growing interest in the recent years for the effects of malicious, external, operational faults in computing systems, ie. intrusions. The term intrusion tolerance has been introduced to emphasize the need to go beyond what classical fault tolerant systems were able to offer. Unfortunately, as opposed to well understood accidental faults, the domain is still lacking sound data sets and models to offer rationales in the design of intrusion tolerant solutions. In this paper, we describe a framework similar in its spirit to so called honey-farms but built in a way that makes its large-scale deployment easily feasible. Furthermore, it offers a very rich level of interaction with the attackers without suffering from the drawbacks of expensive high interaction systems. The system is described, a prototype is presented as well as some preliminary results that highlight the feasibility as well as the usefulness of the approach.

Keywords: fault tolerant computing, security of dataaccidental faults, computing systems, fault tolerant systems, intrusion tolerance, malware threat models, operational faults

[21] C. Leita, V. Pham, O. Thonnard, E. Ramirez-Silva, F. Pouget, E. Kirda, and M. Dacier. The leurre.com project: collecting internet threats information using a worldwide distributed honeynet. In *Proceedings of the 1st WOMBAT workshop*, April 2008. [bib]

[22] R. Lippmann, D. Fried, I. Graf, J. Haines, K. Kendall, D. McClung, D. Weber, S. Webster, D. Wyschogrod, R. Cunningham, and M. Zissman. Evaluating intrusion detection systems: The 1998 DARPA off-line intrusion detection evaluation. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX)*, volume 2, pages 1012-1026, January 2000. [bib | DOI]

An intrusion detection evaluation test bed was developed which generated normal traffic similar to that on a government site containing 100's of users on 1000's of hosts. More than 300 instances of 38 different automated attacks were launched against victim UNIX hosts in seven weeks of training data and two weeks of test data. Six research groups participated in a blind evaluation and results were analyzed for probe, denial-of-service (DoS) remote-to-local (R2L), and user to root (U2R) attacks. The best systems detected old attacks included in the training data, at moderate detection rates ranging from 63% to 93% at a false alarm rate of 10 false alarms per day. Detection rates were much worse for new and novel R2L and DoS attacks included only in the test data. The best systems failed to detect roughly half these new attacks which included damaging access to root-level privileges by remote users. These results suggest that further research should focus on developing techniques to find new attacks instead of extending existing rule-based approaches.

Keywords: safety systems, security of dataDARPA off-line intrusion detection evaluation, automated attacks, denial-of-service, intrusion detection systems, remote-to-local, rule-based approaches, user to root

[23] M. E. Locasto, A. Stavrou, G. F. Cretu, and A. D. Keromytis. From STEM to SEAD: speculative execution for automated defense. In *Proceedings of the USENIX Annual Technical Conference*, pages 219-232, 2007. [bib | http]

Most computer defense systems crash the process that they protect as part of their response to an attack. Although recent research explores the feasibility of self-healing to automatically recover from an attack, self-healing faces some obstacles before it can protect legacy applications and COTS (Commercial Off-The-Shelf) software. Besides the practical issue of not modifying source code, self-healing must know both when to engage and how to guide a repair.

Previous work on a self-healing system, STEM, left these challenges as future work. This paper improves STEM's capabilities along three lines to provide practical speculative execution for

automated defense (SEAD). First, STEM is now applicable to COTS software: it does not require source code, and it imposes a roughly 73% performance penalty on Apache's normal operation. Second, we introduce repair policy to assist the healing process and improve the semantic correctness of the repair. Finally, STEM can create behavior profiles based on aspects of data and control flow.

[24] D. Nicol, W. Sanders, and K. Trivedi. Model-based evaluation: from dependability to security. *IEEE Transactions on Dependable and Secure Computing*, 1(1):48-65, 2004. [bib | DOI]

The development of techniques for quantitative, model-based evaluation of computer system dependability has a long and rich history. A wide array of model-based evaluation techniques is now available, ranging from combinatorial methods, which are useful for quick, rough-cut analyses, to state-based methods, such as Markov reward models, and detailed, discrete-event simulation. The use of quantitative techniques for security evaluation is much less common, and has typically taken the form of formal analysis of small parts of an overall design, or experimental red team-based approaches. Alone, neither of these approaches is fully satisfactory, and we argue that there is much to be gained through the development of a sound model-based methodology for quantifying the security one can expect from a particular design. In this work, we survey existing model-based techniques for evaluating system dependability, and summarize how they are now being extended to evaluate system security. We find that many techniques from dependability evaluation can be applied in the security domain, but that significant challenges remain, largely due to fundamental differences between the accidental nature of the faults commonly assumed in dependability evaluation, and the intentional, human nature of cyber attacks.

Keywords: Markov processes, fault tolerant computing, security of data, software reliability Markov reward models, dependability evaluation, discrete-event simulation, model-based evaluation, performability evaluation, security evaluation, stochastic modeling, system dependability, system security

[25] Y. U. Ryu and H.-S. Rhee. Evaluation of intrusion detection systems under a resource constraint. *ACM Transactions on Information and System Security*, 11(4):1-24, 2008. [bib | DOI]

An intrusion detection system plays an important role in a firm's overall security protection. Its main purpose is to identify potentially intrusive events and alert the security personnel to the danger. A typical intrusion detection system, however, is known to be imperfect in detection of intrusive events, resulting in high false-alarm rates. Nevertheless, current intrusion detection models unreasonably assume that upon alerts raised by a system, an information security officer responds to all alarms without any delay and avoids damages of hostile activities. This assumption of responding to all alarms with no time lag is often impracticable. As a result, the benefit of an intrusion detection system can be overestimated by current intrusion detection models. In this article, we extend previous models by including an information security officer's alarm inspection under a constraint as a part of the process in determining the optimal intrusion detection policy. Given a potentially hostile environment for a firm, in which the intrusion rates and costs associated with intrusion and security officers' inspection can be estimated, we outline a framework to establish the optimal operating points for intrusion detection systems under security officers' inspection constraint. The optimal solution to the model will provide not only a basis of better evaluation of intrusion detection systems but also useful insights into operations of intrusion detection systems. The firm can estimate expected benefits for running intrusion detection systems and establish a basis for increase in security personnel to relax security officers' inspection constraint.

[26] K. Sallhammar. *Stochastic Models for Combined Security and Dependability Evaluation*. PhD thesis, Norwegian University of Science and Technology, Faculty of Information Technology, Mathematics and Electrical Engineering, Department of Telematics, Trondheim, Norway, June 2007. [bib | http]

Security is a topic of ever increasing interest. Today it is widely accepted that, due to the unavoidable presence of vulnerabilities, design faults and administrative errors, an ICT system will never be totally secure. Connecting a system to a network will necessarily introduce a risk of inappropriate access resulting in disclosure, corruption and/or loss of information. Therefore, the security of a system

should ideally be interpreted in a probabilistic manner. More specifically, there is an urgent need for modelling methods that provide operational measures of the security. Dependability, on the other hand, is the ability of a computer system to deliver service that can justifiably be trusted. In a dependability context one distinguishes between accidental faults, which are modelled as random processes, and intentional faults, i.e., attacks, which in most cases are not considered at all. A major drawback of this approach is that attacks may in many cases be the dominating failure source for today's networked systems. The classical way of dependability evaluation can therefore be very deceptive: highly dependable systems may in reality fail much more frequently than expected, due to the exploitation from attackers. To be considered trustworthy, a system must be both dependable and secure. However, these two aspects have so far tended to be treated separately. A unified modelling framework for security and dependability evaluation would be advantageous from both points of view. The security community can benefit from the mature dependability modelling techniques, which can provide the operational measures that are so desirable today. On the other hand, by adding hostile actions to the set of possible fault sources, the dependability community will be able to make more realistic models than the ones that are currently in use. This thesis proposes a stochastic modeling approach, which can be used to predict a system's security and dependability behavior. As will be seen, the basic model has a number of possible applications. For example, it can be used as a tool for trade-off analysis of security countermeasures, or it can be used as a basis for real-time assessment of the system trustworthiness.

[27] K. Sallhammar, B. E. Helvik, and S. J. Knapskog. A game-theoretic approach to stochastic security and dependability evaluation. In *2nd IEEE International Symposium on Dependable, Autonomic and Secure Computing (DASC)*, pages 61-68, September 2006. [bib | DOI]

The complex networked computer systems of today are often vulnerable to a large number of failures, accidental as well as intentional. To be able to assess to what degree one can rely on such a system, new methods for quantitative evaluation is needed. This paper presents a stochastic model for integrated security and dependability evaluation, which models malicious attacker behavior as transitions between system states. To predict the probabilities of attack actions a game theoretic approach is applied. We demonstrate the method by computing security and dependability measures in two different case studies.

Keywords: game theory, security of data, stochastic processesdependability evaluation, game theory, malicious attacker behavior model, stochastic security

[28] K. Sallhammar, B. E. Helvik, and S. J. Knapskog. On stochastic modeling for integrated security and dependability evaluation. *Journal of Networks (JNW)*, 1(5):31-42, September/October 2006. [bib | .html]

This paper presents a new approach to integrated security and dependability evaluation, which is based on stochastic modeling techniques. Our proposal aims to provide operational measures of the trustworthiness of a system, regardless if the underlying failure cause is intentional or not. By viewing system states as elements in a stochastic game, we can compute the probabilities of expected attacker behavior, and thereby be able to model attacks as transitions between system states. The proposed game model is based on a reward- and cost concept. A section of the paper is devoted to the demonstration of how the expected attacker behavior is affected by the parameters of the game. Our model opens up for use of traditional Markov analysis to make new types of probabilistic predictions for a system, such as its expected time to security failure.

Keywords: stochastic models, integrating security and dependability, security measures, game theory

[29] K. Sallhammar, B. E. Helvik, and S. J. Knapskog. Towards a stochastic model for integrated security and dependability evaluation. In *First International Conference on Availability, Reliability and Security (ARES)*, pages 156-165, April 2006. [bib | DOI]

We present a new approach to integrated security and dependability evaluation, which is based on stochastic modelling techniques. Our proposal aims to provide operational measures of the trustworthiness of a system, regardless if the underlying failure cause is intentional or not. By viewing system states as elements in a stochastic game, we can compute the probabilities of expected attacker behavior, and thereby be able to model attacks as transitions between system states. The

proposed game model is based on a rewardand cost concept. A section of the paper is devoted to the demonstration of how the expected attacker behavior is affected by the parameters of the game. Our model opens up for use traditional Markov analysis to make new types of probabilistic predictions for a system, such as its expected time to security failure.

Keywords: Markov processes, probability, security of data, stochastic games, system recovery Markov analysis, dependable computing, probabilistic prediction, reward and cost concept, security failure, stochastic game theory, stochastic model

[30] K. Sallhammar, B. E. Helvik, and S. J. Knapskog. A framework for predicting security and dependability measures in real-time. *International Journal of Computer Science and Network Security (IJCSNS)*, 7(3):169-183, March 2007. [bib | .html | .pdf]

The complex networked systems of today that our technological and social society relies upon are vulnerable to a large number of failures, accidental as well as intentional. Ideally, the service delivered by such a system should be both dependable and secure. This paper presents a framework for integrated security and dependability assessment. The proposed model is based on traditional stochastic analysis techniques, supported by live data from network sensors, which is used to estimate the current state and predict the future behavior of a system in real-time. The method is demonstrated by a small case study.

Keywords: Stochastic modeling, integrating security and dependability, security measures, real-time prediction, hidden Markov models

[31] N. Stakhanova. *A framework for adaptive, cost-sensitive intrusion detection and response system*. PhD thesis, Iowa State University, 2007. [bib]

Intrusion detection has been at the center of intense research in the last decade owing to the rapid increase of sophisticated attacks on computer systems. Typically, intrusion detection refers to a variety of techniques for detecting attacks in the form of malicious and unauthorized activities. There are three broad categories of detection approaches: (a) misuse-based technique that relies on pre-specified attack signatures, (b) anomaly-based approach, that typically depends on normal patterns classifying any deviation from normal as malicious; and (c) specification-based technique that although operates in a similar fashion to anomaly-based approach, employs a model of valid program behavior in a form of specifications requiring user expertise.

When intrusive behavior is detected, it is desirable to take (evasive and/or corrective) actions to thwart attacks and ensure safety of the computing environment. Such countermeasures are referred to as intrusion response. Although the intrusion response component is often integrated with the Intrusion Detection System (IDS), it receives considerably less attention than IDS research owing to the inherent complexity in developing and deploying response in an automated fashion. As such, traditionally, triggering an intrusion response is left as part of the administrators responsibility, requiring a high-degree of expertise.

In this work we present an integrated approach to intrusion detection and response based on the technique for monitoring abnormal patterns in the program behavior. The proposed model effectively combines the advantages of anomaly-based and specification-based approaches recognizing a known behavior through the specifications of normal and abnormal patterns and classifying unknown patterns using a machine-learning algorithm. Such combination not only allows adaptation of the specification-based detection to the new patterns, but also provides a method for automatic development of specifications.

In addition to detection, our framework incorporates preemptive response. By preemption, we imply deploying response before a monitored pattern is classified completely as an intrusion. Such response deployment is likely to stop an intrusion before it can affect the system. However, preemption also inherently suffers from false positives; i.e., responses are deployed to deter correct execution which may look intrusive in its initial phase. To reduce false positives, we have developed a multi-phase response selection and deployment mechanism based on the evaluation of the cost information of the system damage caused by potential intrusion and candidate responses.

[32] N. Stakhanova, S. Basu, and J. Wong. A taxonomy of intrusion response systems. *Int. J. Inf.*

Comput. Secur., 1(1/2):169-184, 2007. [bib | DOI]

Recent advances in the field of intrusion detection brought new requirements to intrusion prevention and response. Traditionally, the response to an attack is manually triggered by an administrator. However, increased complexity and speed of the attack-spread during recent years show acute necessity for complex dynamic response mechanisms. Although intrusion detection systems are being actively developed, research efforts in intrusion response are still isolated. In this work we present a taxonomy of intrusion response systems, together with a review of current trends in intrusion response research. We also provide a set of essential features as a requirement for an ideal intrusion response system.

[33] F. Stevens, T. Courtney, S. Singh, A. Agbaria, J. Meyer, W. Sanders, and P. Pal. Model-based validation of an intrusion-tolerant information system. In *Proceedings of the 23rd IEEE International Symposium on Reliable Distributed Systems (SRDS)*, pages 184-194, October 2004. [bib | DOI]

An increasing number of computer systems are designed to be distributed across both local and wide-area networks, performing a multitude of critical information-sharing and computational tasks. Malicious attacks on such systems are a growing concern, where attackers typically seek to degrade quality of service by intrusions that exploit vulnerabilities in networks, operating systems, and application software. Accordingly, designers are seeking improved techniques for validating such systems with respect to specified survivability requirements. In this regard, we describe a model-based validation effort that was undertaken as part of a unified approach to validating a networked intrusion-tolerant information system. Model-based results were used to guide the system's design as well as to determine whether a given survivability requirement was satisfied.

Keywords: computer networks, fault tolerant computing, formal verification, information systems, security of data application software, computer systems, critical computational task, critical information-sharing, intrusion tolerant information system, local area network, malicious attacks, model-based validation, networked intrusion-tolerant information system, operating system, quality of service, survivability requirement, wide area network

[34] C. Strasburg, N. Stakhanova, S. Basu, and J. Wong. The methodology for evaluating response cost for intrusion response systems. Technical Report 08-12, Dept. of Computer Science, Iowa State University, Dec. 2008. [bib | http]

Recent advances in the field of intrusion detection brought new requirements to intrusion prevention and response. Traditionally, the response to the detected attack was selected and deployed manually, in the recent years the focus has shifted towards developing automated and semi-automated methodologies for responding to intrusions. In this context, the cost-sensitive intrusion response models have gained the most interest mainly due to their emphasis on the balance between potential damage incurred by the intrusion and cost of the response. However, one of the challenges in applying this approach is defining consistent and adaptable measurement of these cost factors on the basis of requirements and policy of the system being protected against intrusions. In this paper we present a structured methodology for evaluating cost of responses based on three factors: the response operational cost associated with the daily maintenance of the response, the response goodness that measures the applicability of the selected response for a detected intrusion and the response impact on the system that refers to the possible response effect on the system functionality. The proposed approach provides consistent basis for response evaluation across different systems while incorporating security policy and properties of specific system environment. We demonstrate the advantages of the proposed cost model and evaluate it on the example of three systems.

[35] C. Strasburg, N. Stakhanova, S. Basu, and J. S. Wong. A framework for cost sensitive assessment of intrusion response selection. In *IEEE Computer Software and Applications Conference (COMPSAC)*, July 2009. [bib | .pdf]

[36] R. Werlinger, K. Hawkey, K. Muldner, and P. Jaferian. The challenges of using an intrusion detection system: Is it worth the effort? In *Proceedings of the 4th International Symposium On Usable Privacy and Security (SOUPS)*, pages 107-118??, July 2008. [bib | .pdf]

An intrusion detection system (IDS) can be a key component of security incident response within organizations. Traditionally, intrusion detection research has focused on improving the accuracy of IDSs, but recent work has recognized the need to support the security practitioners who receive the IDS alarms and investigate suspected incidents. To examine the challenges associated with deploying and maintaining an IDS, we analyzed 9 interviews with IT security practitioners who have worked with IDSs and performed participatory observations in an organization deploying a network IDS. We had three main research questions: (1) What do security practitioners expect from an IDS?; (2) What difficulties do they encounter when installing and configuring an IDS?; and (3) How can the usability of an IDS be improved? Our analysis reveals both positive and negative perceptions that security practitioners have for IDSs, as well as several issues encountered during the initial stages of IDS deployment. In particular, practitioners found it difficult to decide where to place the IDS and how to best configure it for use within a distributed environment with multiple stakeholders. We provide recommendations for tool support to help mitigate these challenges and reduce the effort of introducing an IDS within an organization.

[37] Proceedings of the Workshop on Information-Security-System Rating and Ranking (WISSRR), May 2001. [bib | http | .pdf]

[38] Y. Wu and S. Liu. A cost-sensitive method for distributed intrusion response. In *12th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, pages 760-764, April 2008. [bib | DOI]

A method for the evaluation of response cost is proposed. It is based on the principle that one should achieve the maximum security goal through a minimal response cost. On this basis a method for judging the causal relationship between an intrusion and a cooperative intrusion is further suggested. The intrusion response system designed according to the above response strategy can be applied to the distributed network environment. Through the cooperation of more than one management domain and a large scale study of relationships among various intrusion response costs a superior response strategy can be deduced.

Keywords: distributed processing, security of data, software cost estimationcooperative intrusion, cost-sensitive method, distributed intrusion response, distributed network environment, maximum security goal, response cost

[39] Y.-S. Wu, B. Foo, Y.-C. Mao, S. Bagchi, and E. H. Spafford. Automated adaptive intrusion containment in systems of interacting services. *Computer Networks*, 51(5):1334-1360, 2007. [bib | DOI]

Large scale distributed systems typically have interactions among different services that create an avenue for propagation of a failure from one service to another. The failures being considered may be the result of natural failures or malicious activity, collectively called disruptions. To make these systems tolerant to failures it is necessary to contain the spread of the occurrence automatically once it is detected. The objective is to allow certain parts of the system to continue to provide partial functionality in the system in the face of failures. Real world situations impose several constraints on the design of such a disruption tolerant system of which we consider the following - the alarms may have type I or type II errors; it may not be possible to change the service itself even though the interaction may be changed; attacks may use steps that are not anticipated a priori; and there may be bursts of concurrent alarms. We present the design and implementation of a system named Adepts as the realization of such a disruption tolerant system. Adepts uses a directed graph representation to model the spread of the failure through the system, presents algorithms for determining appropriate responses and monitoring their effectiveness, and quantifies the effect of disruptions through a high level survivability metric. Adepts is demonstrated on a real e-commerce testbed with actual attack patterns injected into it.

Keywords: Automated adaptive intrusion response, Intrusion containment, E-commerce system, Survivability, Attack graphs

[40] D. Yu and D. Frincke. Improving the quality of alerts and predicting intruder's next goal with hidden colored petri-net. *Computer Networks*, 51(3):632-654, 2007. [bib | DOI]

Intrusion detection systems (IDS) often provide poor quality alerts, which are insufficient to support rapid identification of ongoing attacks or predict an intruder's next likely goal. In this paper, we propose a novel approach to alert postprocessing and correlation, the Hidden Colored Petri-Net (HCPN). Different from most other alert correlation methods, our approach treats the alert correlation problem as an inference problem rather than a filter problem. Our approach assumes that the intruder's actions are unknown to the IDS and can be inferred only from the alerts generated by the IDS sensors. HCPN can describe the relationship between different steps carried out by intruders, model observations (alerts) and transitions (actions) separately, and associate each token element (system state) with a probability (or confidence). The model is an extension to Colored Petri-Net (CPN). It is so called “hidden” because the transitions (actions) are not directly observable but can be inferred by looking through the observations (alerts). These features make HCPN especially suitable for discovering intruders' actions from their partial observations (alerts) and predicting intruders' next goal. Our experiments on DARPA evaluation datasets and the attack scenarios from the Grand Challenge Problem (GCP) show that HCPN has promise as a way to reducing false positives and negatives, predicting intruder's next possible action, uncovering intruders' intrusion strategies after the attack scenario has happened, and providing confidence scores.

Keywords: Intrusion detection, Alert correlation, Hidden Colored Petri-Net

---