

**UNDERSTANDING
INTRUSION DETECTION
THROUGH VISUALIZATION**

Advances in Information Security

Sushil Jajodia

Consulting Editor

Center for Secure Information Systems

George Mason University

Fairfax, VA 22030-4444

email: jajodia@gmu.edu

The goals of the Springer International Series on ADVANCES IN INFORMATION SECURITY are, one, to establish the state of the art of, and set the course for future research in information security and, two, to serve as a central reference source for advanced and timely topics in information security research and development. The scope of this series includes all aspects of computer and network security and related areas such as fault tolerance and software assurance.

ADVANCES IN INFORMATION SECURITY aims to publish thorough and cohesive overviews of specific topics in information security, as well as works that are larger in scope or that contain more detailed background information than can be accommodated in shorter survey articles. The series also serves as a forum for topics that may not have reached a level of maturity to warrant a comprehensive textbook treatment.

Researchers, as well as developers, are encouraged to contact Professor Sushil Jajodia with ideas for books under this series.

Additional titles in the series:

HOP INTEGRITY IN THE INTERNET by Chin-Tser Huang and Mohamed G. Gouda; ISBN-10: 0-387-24426-3

PRIVACY PRESERVING DATA MINING by Jaideep Vaidya, Chris Clifton and Michael Zhu; ISBN-10: 0-387- 25886-8

BIOMETRIC USER AUTHENTICATION FOR IT SECURITY: From Fundamentals to Handwriting by Claus Vielhauer; ISBN-10: 0-387-26194-X

IMPACTS AND RISK ASSESSMENT OF TECHNOLOGY FOR INTERNET SECURITY: Enabled Information Small-Medium Enterprises (TEISMES) by Charles A. Shoniregun; ISBN-10: 0-387-24343-7

SECURITY IN E-LEARNING by Edgar R. Weippl; ISBN: 0-387-24341-0

IMAGE AND VIDEO ENCRYPTION: From Digital Rights Management to Secured Personal Communication by Andreas Uhl and Andreas Pommer; ISBN: 0-387-23402-0

INTRUSION DETECTION AND CORRELATION: Challenges and Solutions by Christopher Kruegel, Fredrik Valeur and Giovanni Vigna; ISBN: 0-387-23398-9

THE AUSTIN PROTOCOL COMPILER by Tommy M. McGuire and Mohamed G. Gouda; ISBN: 0-387-23227-3

ECONOMICS OF INFORMATION SECURITY by L. Jean Camp and Stephen Lewis; ISBN: 1-4020-8089-1

PRIMALITY TESTING AND INTEGER FACTORIZATION IN PUBLIC KEY CRYPTOGRAPHY by Song Y. Yan; ISBN: 1-4020-7649-5

SYNCHRONIZING E-SECURITY by Godfried B. Williams; ISBN: 1-4020-7646-0

Additional information about this series can be obtained from <http://www.springeronline.com>

UNDERSTANDING INTRUSION DETECTION THROUGH VISUALIZATION

by

Stefan Axelsson

*Chalmers University of Technology
Göteborg, Sweden*

David Sands

*Chalmers University of Technology
Göteborg, Sweden*



Springer

Dr. Stefan Axelsson
Dept. of Computer Science and Engineering
Chalmers University of Technology
412 96 GÖTEBORG
SWEDEN

Prof. David Sands
Dept. of Computer Science and Engineering
Chalmers University of Technology
412 96 GÖTEBORG
SWEDEN

Library of Congress Control Number: 2005933712

UNDERSTANDING INTRUSION DETECTION THROUGH VISUALIZATION

by Stefan Axelsson and David Sands

ISBN-13: 978-0-387-27634-2

ISBN-10: 0-387-27634-3

e-ISBN-13: 978-0-387-27636-6

e-ISBN-10: 0-387-27636-X

Printed on acid-free paper.

© 2006 Springer Science+Business Media, Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, Inc., 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed in the United States of America.

9 8 7 6 5 4 3 2 1

SPIN 11425250, 11524885

springeronline.com

Contents

List of Figures	ix
List of Tables	xi
Foreword	xiii
Preface	xvii
Acknowledgments	xix
1. INTRODUCTION	1
1 Context	1
2 Computer Security	2
3 Rationale and Problem Statement	3
4 Information Visualization	4
5 Overview of the book	5
2. AN INTRODUCTION TO INTRUSION DETECTION	15
1 Intrusion Prevention	15
2 Intrusion Detection	17
3. THE BASE-RATE FALLACY AND THE DIFFICULTY OF INTRUSION DETECTION	31
1 Problems in Intrusion Detection	32
2 The Base-Rate Fallacy	32
3 The Base-Rate Fallacy in Intrusion Detection	35
4 Impact on Intrusion Detection Systems	40
5 Future Directions	46
6 Further Reading	46
7 Conclusions	47

4.	VISUALIZING INTRUSIONS: WATCHING THE WEBSERVER	49
1	The Experimental System	50
2	The Log Reduction Scheme	51
3	Visualizing the Lowest Scoring Requests	52
4	Detailed Analysis of the Features Found	56
5	Effectiveness of the Log Reduction Scheme	59
6	Discussion	63
7	Future Work	66
8	Conclusions	66
9	Further Reading	67
5.	COMBINING A BAYESIAN CLASSIFIER WITH VISUALIZATION	69
1	Automated Learning for Intrusion Detection	69
2	Naive Bayesian Detection	70
3	The Experimental Data	71
4	Visualizing a Bayesian IDS	73
5	The Training Data	80
6	The Experiment	80
7	Conclusions	86
6.	VISUALIZING THE INNER WORKINGS OF A SELF LEARNING CLASSIFIER	89
1	Introduction	89
2	Markovian Matching with Chi Square Testing	90
3	Visualizing the Detector	92
4	The Experimental Data	98
5	Experimental Results	99
6	Conclusions	109
7	Future Work	109
7.	VISUALIZATION FOR INTRUSION DETECTION —HOOKING THE WORM	111
1	Introduction	111
2	The Monitored System	112
3	Scientific Visualization	114

<i>Contents</i>	vii
4 Visual Analysis of the Log File	119
5 Results of the Investigation	121
6 Discussion	125
7 Conclusion	126
8 Future Work	126
8. EPILOGUE	129
1 Results in Perspective	129
2 Further Reading	130
3 Conclusions and Future Work	132
References	133
Author Index	141
Index	143

List of Figures

2.1	Anti-intrusion techniques	15
2.2	Organisation of a generalised intrusion detection system	21
2.3	Classical detection theory model	23
2.4	One dimensional detection model	25
3.1	Venn diagram of medical diagnostic example	34
3.2	Plot of Bayesian detection rate versus false alarm rate	39
3.3	ROC-curves for the “low performers”	41
3.4	ROC-curve for the “high performers”	41
4.1	Frequencies of component frequencies	51
4.2	Requests sorted by lowest score	51
4.3	Graph of the lowest scoring requests	53
4.4	Zoom on feature (Spam attack in this case)	54
4.5	Zoom on feature (benign accesses forming a subgraph, isolated)	55
4.6	Zoom on feature (benign accesses forming a subgraph, in vivo)	56
4.7	Zoom on feature (benign accesses forming a not very clear subgraph)	57
4.8	The remaining accesses deemed to be intrusion attempts, 2D graph	60
4.9	The remaining accesses deemed to be intrusion attempts, 3D graph	61
5.1	The <i>Bayesvis</i> tool	75
5.2	The <i>Bayesvis</i> tool after retraining on false alarms	78
5.3	The <i>Bayesvis</i> tool after having corrected under training	79
5.4	False positives during the training phase	83
5.5	Examples of false alarms in February log	84

5.6	Generalized detection of Unicode attacks	86
6.1	The <i>Chi2vis</i> tool after training one bad and one good	93
6.2	The <i>Chi2vis</i> tool after training one bad and two good	96
6.3	The <i>Chi2vis</i> tool after training two bad and two good	97
6.4	Generalising the Unicode training to detect new instances	102
6.5	False alarms: Example of the <i>HEAD</i> -pattern	103
6.6	Results from training on syscall data	104
6.7	All the false alarms of Bayesvis	107
6.8	The “cgi-bin” pattern false alarms of Bayesvis	108
7.1	Sample records from the webserver log file	114
7.2	A simple parallel coordinate plot	116
7.3	A trellis of parallel coordinate plots	120
7.4	A plot of the “Code-red” worm access pattern	122
7.5	The six different requests made by pattern 1 from Figure 7.3	123

List of Tables

4.1	Detection rates of the log reduction mechanism	62
4.2	Summary of the true and false alarms of the log reduction mechanism	63
4.3	Number of previously unseen (new) accesses for the following months	64
5.1	Summary of the types of accesses in the training data	81
5.2	Summary of the results of the experiment (approximate values)	87
6.1	Summary of the types of accesses in the training data	99
6.2	False negatives (<i>misses</i>) in testing data	101
6.3	False positives (<i>false alarms</i>) in testing data	102
6.4	False negatives (<i>misses</i>) in testing data for Bayesvis	105
6.5	False positives (<i>false alarms</i>) in testing data for Bayesvis	106

Foreword

This monograph is the outgrowth of Stefan Axelson's PhD Dissertation at Chalmers University in Göteborg, Sweden. The dissertation, in turn collects a number of research efforts performed over a period of six years or so into a coherent whole. It was my honor to serve as the "opponent" at Dr. Axelsson's examination. In the Swedish system, it is the job of the opponent to place the candidate's work into a broader perspective, demonstrating its significance and contributions to the field and then to introduce the work to the attendees at the examination. This done, the candidate presents the technical details of the work and the opponent critiques the work giving the candidate the opportunity to defend it¹. This forward is adapted from the introduction that I gave at the examination and should serve to acquaint the reader, not only with the work at hand, but also with the field to which it applies. The title of the work, "Understanding Intrusion Detection Through Visualization," is particularly telling. As is the case with any good piece of research, we hope to gain an understanding of a problem, not just a recipe or simple solution of immediate, but limited utility.

For much of its formative period, computer security concentrated on developing systems that, in effect, embodied a fortress model of protection. These systems were intended to be immune to most of the attacks that we see today and were supposed to be capable of processing classified material at multiple levels of security (MLS). The problem of building highly secure systems was harder than thought, but, by the early 1990s, a number of promising systems were beginning to emerge.

In the mid 1980s commodity personal computers emerged. These were initially produced without any regard for security – not even protecting the user

¹It is interesting to note that Swedish technical universities received the ability to award PhDs rather late (1940 in the case of Chalmers), as it was felt that the work of the master engineer had to stand above *any* criticism and it was thus inappropriate to subject it to a form of examination which in its very form relied on the work being subjected to critique.

from himself. The military adopted these platforms wholesale in spite of their insecurity and stopped substantial MLS research efforts in the mid 1990s.

By the late 1980s, broad band networks were available to most corporations and to many educational institutions. Increasingly, these were using PC based platforms as network nodes. Node level security was minimal and difficult. Managing large numbers of machines securely was difficult or impossible. Firewalls were introduced to provide a single point of protection for an organization. Intrusion Detection Systems (IDS) were introduced to detect attacks either from the outside or from the inside, providing another line of defense for the increasingly difficult to manage firewalls.

In 1980, James Anderson produced a report entitled “Computer Security Threat Monitoring and Surveillance” that sets up the framework for what we now know as intrusion detection. Anderson (and later Denning) assumed that user behavior was regular enough to permit statistical models that would equate unusual (or anomalous) with malicious. In general, this is not true, but anomaly based systems are still the focus of much research. The other primary area of activity is signature based systems in which patterns of activity that match previously known intrusions are sought. Finding the right pattern at an appropriate level of abstraction is not easy and most truly new attacks are undetectable using signatures.

There are a number of problems that beset both production and research intrusion detection systems. These provide a context for the monograph and include: 1) Lack of a fundamental theoretical basis for intrusion detection and 2) Poor understanding of environments in which intrusion detection systems function. These lead to excessive false alarms, inappropriate training for machine learning systems, poorly formed signatures for abuse detection and many other problems. The monograph directly addresses several of these problems. It is the result of a series of investigations that began late in the last century. Although individual results have appeared in a variety of forums, they represent a coherent body of work and a significant contribution to the field.

In the next few paragraphs, we will introduce each of these works and place them in perspective. The technical details of each form a chapter in the monograph.

The Base-Rate Fallacy and the Difficulty of Intrusion Detection

Originally presented at RAID 99, this was my first introduction to the Dr. Axelsson and his work. It deals with the problem of excessive false alarm rates, a problem that plagues many intrusion detection systems.

The problem of false alarms is troubling. Every alarm requires investigation and uses (typically human) resources. Alarms are often described in terms of percentages of cases examined. If there are a lot of cases, even a low alarm rate can require excessive resources to examine every alarm. While this is well

known in epidemiology (where it is called the Base-Rate Fallacy), its impact was not understood in the IDS community. As a result of this work, the IDS community is now aware that very low intrusion rates require even lower false alarm rates to prevent operator overload. The consequences of this observation inform much of the subsequent work.

Visualizing Intrusions

Watching the Webserver represents a *tour de force* in primary data analysis as well as providing a beautiful example of an observational study. In many cases, the quantity of data available defies individual analyses. Only by clustering and abstraction can the data be reduced to manageable size.

Most researchers in this area are more interested in their algorithms than in the data. In this work the analysis is properly viewed as a means to understanding the processes that produced the data. While the way the log reduction and visualization were performed are significant contributions, some of the observations in the discussion have the potential to be even more significant as they provide a possible basis for defining a necessary property of certain intrusions.

Combining a Bayesian Classifier with Visualization

Understanding the IDS is an often overlooked aspect of research in this field. Much of the current work in intrusion detection involves machine learning. Even using carefully labeled data, classifiers often learn the right thing for the wrong reasons. As far as I know, the approach here of using visualization with interactive classification during the learning phase as an aid to understanding both the data and the detector, is unique.

While the simple Bayesian detector used in the study is not particularly strong as an IDS, the training approach can be extended to other detectors and the results are impressive for the detector involved. This work is significant in its own right, however, it also sets forth a significant agenda of future work.

Visualizing the Inner Workings of a Self Learning Classifier

Following the previous work with a more complex learning system is logical next step. The detector used in this study is much more complex and its operation, as originally defined, opaque. Not knowing why a classifier made a particular classification impedes training and hampers use.

The work performed here demonstrates, for this more complex case, that it is possible to develop a visualization that gives insight into both the classifier and the data allowing the “why” to be understood. As in the previous case, the insights into the reasons why the detectors function as they do on the data provides insight into the intrusive behavior.

Visualization for Intrusion Detection

Hooking the Worm is an interesting study of attempts to attack a small web server. This work takes a neutral view of the dataset involved, developing visual techniques for clustering and displaying web accesses. As we noted earlier, clustering and abstracting allow us to reduce many individual records to a manageable set of classes.

In this case, reducing the records to a few essential characteristics still allow the production of useful patterns. The primary contribution of the work is a simple mechanism for providing insight into system activity in a way that supports classification into malicious and benign activity.

Beyond the Monograph

In addition to providing specific insights in a number of specific areas of intrusion detection, a number of less tangible contributions are made. All of the studies serve as exemplars of the utility of observational studies in computer security. The astute reader will see that the work has benefited from deep thought into the activities manifest in the data and tools studied. The resulting insights are carefully and clearly set forth.

The works also show that there is no easy substitute for primary data collection and analysis. Researchers who expect to have data sets handed to them, should take note that significant results require hard and tedious work. In many other fields, primary data collection and data management may consume as much as 90% of a project's budget. There is no reason to expect observational studies in computer security to be different.

In summary, this is work to be emulated by researchers as well as students. It has been a great pleasure to correspond with Stefan Axelsson as he performed the studies leading to the thesis this monograph is based on, and it is a pleasure to be able to introduce the work to the readers of this monograph.

John MCHugh
Canada Research Chair
Director, Privacy and Security Laboratory
Dalhousie University
Halifax, Nova Scotia, Canada
July 2005

Preface

With the ever increasing use of computers for critical systems, computer security, the protection of data and computer systems from intentional, malicious intervention, is attracting much attention. Among the methods for defense, *intrusion detection*, i.e. the application of a tool to help the operator identify ongoing or already perpetrated attacks, has been the subject of considerable research in the past ten years. A key problem with current intrusion detection systems is the high number of false alarms they produce. This book presents research into why false alarms are and will remain a problem, and proposes to apply results from the field of *information visualization* to the problem of intrusion detection. This approach promises to enable the operator to correctly identify false (and true) alarms, and also aid the operator in identifying other operational characteristics of intrusion detection systems. Four different visualization approaches are presented, mainly applied to data from web server access logs. The four approaches studied can be divided into *direct* and *indirect* methods. In the direct approaches the system puts the onus of identifying the malicious access requests on the operator by way of the visualization. For the indirect approaches the state of two self learning automated intrusion detection systems are visualized to enable the operator to examine their inner workings. The aim here being to provide the operator with an understanding of how the intrusion detections systems operate and whether that level of operation, and the quality of the output, is satisfactory. Several experiments were performed and many different attacks in web access data from publicly available web servers were found. The visualization helped the operator either detect the attacks herself and more importantly the false alarms.

Website

A website for the book can be found at “www.cs.chalmers.se/~dave/VisBook”. Most importantly the website contains the more detailed figures from the book, in full size and color.

Acknowledgments

This book is based on the PhD thesis of the first author, under the supervision of the second. Even though writing the thesis on which this book was based was at times a lonely task, the research was not done in isolation. Far from it; I owe many more people my thanks than I can mention here. That said I would still like to take the opportunity to mention a few people who have been instrumental in helping to bring this work to completion.

While all my colleagues are too numerous to mention, I would especially like to thank (in no particular order) Daniel Hedin, Ulf Norell, Nils Anders Danielsson, Tobias Gedell, Claes Nyberg and Thorbjörn Axelsson. I would be less knowledgeable without having worked with you and I would certainly have had a much drearier time doing it. My erstwhile climbing partner, now turned colleague Dr. Rogardt Heldal deserves special mention, as he is put up with my comings and goings and still managed to provide valuable insights over the past few years. My erstwhile supervisor Prof. Erland Jonsson also deserves special mention, as he was the one that put me on to the idea of applying visualisation to the area of intrusion detection in the first place, many years ago now. I would also like to thank previous and present colleagues at the department of Computer Engineering here at Chalmers and at Ericsson where I have been employed for the past few years.

Outside of Chalmers we would like to thank Prof. John McHugh for his helpful comments and support on a number of aspects of the work presented here, and to *Spotfire inc* for letting us use “Spotfire Decision Suite” for some of our visualization experiments.

Last but not least are the two people without whose support this work would not have got far. I am talking of course of my wife Hanna Tornevall who has had to bear the brunt of the work keeping the family going this autumn, and our son Oskar. In fact, Oskar’s first proper two syllable word was “dat-oo” (clearly legible Swedish for *dator*, i.e. *computer*), as in: “Oskar, where’s daddy?”, “Dat-oo!” I know I have been an absent father at times when preparing the thesis

this book was based on, even when present in the flesh. Thank you Oskar for not holding that against me.

Stefan Axelsson

Since Stefan did all the hard work, there are considerably fewer acknowledgments needed from my side. However, this work would not have been possible without the support of the Department of Computer Science and Engineering at Chalmers, and research grants from *SSF* (the Swedish foundation for Strategic Research) and *Vinnova* (The Swedish Agency for Innovation Systems).

David Sands

Göteborg, August 2005