

Personal Anomaly-based Intrusion Detection Smart Card Using Behavioural Analysis

A.M. Rossudowski, H.S. Venter, and J.H.P. Eloff

Information and Computer Security Architectures Research Group (ICSA)
Department of Computer Science, University of Pretoria, South Africa
{amrossudowski, hsventer, eloff}@cs.up.ac.za

Abstract. Intrusion Detection Systems play an invaluable role within organisations by detecting attempted attacks on their IT systems. However, Intrusion Detection Systems are complex to set-up and require large quantities of memory and processing power to effectively analyse the large volumes of network traffic involved. Behavioural analysis plays an important role within Intrusion Detection Systems by looking for suspicious behaviour or behaviour out of the ordinary within the network traffic. This paper identifies several problems that decrease the overall performance of Intrusion Detection Systems. It proposes the use of a personal smart card-based Intrusion Detection System to increase the performance and effectiveness of Intrusion Detection Systems as a whole.

Key words: Intrusion Detection System, IDS, smart card, behavioural analysis, personal IDS.

1 Introduction

An Intrusion Detection System (IDS) is just one of the security tools an organisation can use to protect itself from a wide range of attacks designed to disrupt its systems or steal sensitive information. An IDS tries to detect these attacks by monitoring traffic through the organisation's network. A pattern-based IDS looks for a pre-defined pattern of traffic that could constitute an attack [1, 2]; whilst an anomaly-based IDS looks for anomalies within traffic or behaviours that exceed a certain threshold or specified base-line [1, 2]. The base-line represents the typical behaviour of the organisation's network traffic.

The problem facing an anomaly-based IDS, particularly within a large organisation, is that every employee will browse the Internet or communicate across the network in a unique way. Hence it is quite difficult—if not impossible—to determine what network behaviour constitutes the base-line. While a pattern-based IDS requires tremendous amounts of processing power and time to analyse the large quantities of information [3] passing through a network, which results in inefficiencies.

This paper proposes the use of smart card technology, in conjunction with behavioural analysis, to implement an anomaly-based Intrusion Detection System. The smart card-based IDS (SCIDS) implements several approaches, discussed later, that decrease the time taken to discover certain types of attacks. The SCIDS improves the efficiency of intrusion detection, while simultaneously reducing the complexity of anomaly detection.

Section 2 provides relevant background information concerning this study. Section 3 identifies the limitations within current IDSs and proposes solutions to these limitations. Section 4 presents the smart card-based Intrusion Detection System model. The proposed model is compared to existing IDSs within section 5 to illustrate its advantages. The paper is concluded by section 6.

2 Background

This section provides brief information on behavioural analysis, various types of Intrusion Detection Systems, and an overview of smart cards.

2.1 Behavioural Analysis

Alexander [4] states that *“the study of behaviour encompasses all of the movements and sensations by which animals and men mediate their relations with their external environment—physical, biotic and social”*. Behavioural analysis is defined in the context of this paper as *“the study of how an employee behaves under different conditions and environments, with various internal and external stimuli applied to those environments”*.

Many disparate disciplines incorporate the use of behavioural analysis. Biologists use behavioural analysis of chimpanzees as a model of early hominid behaviour [5]. Computer scientists use behavioural analysis to determine the performance and behaviour of complex systems [6]. More frequently, a cross-disciplinary approach to behavioural analysis is being applied, such as the RoboCup Initiative [7]. Researchers are also studying how humans interact with robots so they can design more socially interactive robots in the future [8].

Behavioural analysis can also be applied to the way human beings browse the Internet under changing stimuli, such as night and day and at various times of the day, week, month and/or year. Every individual will browse the Internet in a unique way, for example, at different times on different days and this pattern can be used to build a user’s web browsing profile.

2.2 Intrusion Detection Systems

An Intrusion Detection System (IDS) detects unauthorised access to, or use of, a system or an application [1]. Most IDSs are passive systems using pattern-based (also known as signature-based) detection mechanisms. The most challenging IDS to implement is an active system using anomaly-based detection methods.

A pattern-based IDS detects an attack on a system by looking for a particular series of actions, commands, or events (i.e. a pattern). This pattern is usually created from the records of previous attacks [1, 2]. An anomaly-based IDS, on the other hand, looks for any actions, commands, or events that fall outside the scope of normal user behaviour (the base-line) [1, 2]. Anomaly-based IDSs are usually taught what normal system activity is and generate heuristics or rules according to this behaviour. Any actions that do not comply with these heuristics or rules are flagged as possible intrusions.

An IDS is a tool used to monitor, identify and respond to attacks on a given system and/or network. There are several different types of IDSs:

- **Host-based IDSs** are designed to run in the background on systems presumed to be critical and/or sensitive, such as web servers, mail servers and DNS servers[2].
- **Network-based IDSs** sit on the network and monitor traffic at the packet level. The system's network interface is set to *stealth* mode and *promiscuous* mode and has no IP address [9] to help hide it from the network and protect it from attacks [2].
- **Pattern-based IDSs** (also known as misuse, signature or knowledge-based IDSs) monitor the log files (host-based) or network traffic (network-based) looking for specific patterns that could indicate suspicious behaviour [1, 2].
- **Anomaly-based IDSs** (also known as behaviour-based IDSs) use statistical techniques or a trained neural-net to detect penetrations or attacks on the system. This is achieved by determining a statistical base-line of behaviour on the system. Actual behaviour on the system is then analysed and compared to the base-line and an alert issued if a certain threshold is exceeded [1, 2].

Most IDSs are *passive* systems, determining or discovering an intrusion *post factum* through the examination of log files. System and/or Network Administrators then need to determine which vulnerabilities the attack exploited and correct the problem. Unlike a passive IDS, an *active* IDS is able to detect or discover intrusions while they are occurring by monitoring the network traffic in real-time. However, the active IDS can do little to prevent the attack from proceeding.

Intrusion Prevention Systems (IPS) [10, 11] try to either prevent or mitigate the damage caused by an attack. Intrusion Prevention [10, 11] is a relatively new term and is, essentially, a combination of access control (firewall/router) and Intrusion Detection [9]. Hence, an IPS can be defined as a product that focuses on identifying and blocking malicious network activity using preventative measures in real time [9, 11].

2.3 Smart Cards

A smart card is a token that contains an Integrated Circuit Chip (ICC) and is available in a variety of shapes and sizes [12]. The ICC—stored either on the card's surface or within its structure—contains a Central Processing Unit

(CPU), non-volatile memory (RAM, ROM, and/or EEPROM) and an Operating System (OS), usually stored in the EEPROM memory. Information can be stored and retrieved from a smart card in a similar fashion to a magnetic strip card, but a smart card has certain advantages over magnetic strip cards [12, 13], such as:

- The ability to process information stored on the card or passed to it.
- The ability to encrypt the information stored on the card [14].

Moreover, certain information and functionality stored on the smart card can only be accessed if the user (owner of the card) enters an authorised Personal Identification Number (PIN). If the user enters an incorrect PIN several times, either the smart card permanently destroys itself i.e. over loads the internal circuitry, or locks itself and requires the user to enter a longer PIN to regain access [12].

The limitations of pattern-based and anomaly-based IDSs are identified in the next section, followed by a discussion of the possible ways of improving their performance.

3 IDS Limitations and Proposed Solutions

As mentioned previously, there are various problems associated with the effective implementation of different types of Intrusion Detection Systems (IDSs), especially within large organisations. These problems are elaborated on below.

A Pattern-based IDS needs to analyse all the network traffic (packets) looking for specific patterns that suggest an attack is occurring or has occurred. This can be a time consuming process if the IDS has to analyse a large quantity of network information. It is also a computationally expensive process, especially if the attacks are of a more sophisticated and complex nature. In addition, attack patterns need to be pre-recorded within the system to be discovered, especially within a real-time IDS, thus any new type of attack that has not been previously recorded will go undetected.

An Anomaly-based IDS analyses network traffic for any activity that does not fit within the “norm” i.e. base-line. The complexity and difficulty in determining what constitutes “normal” network traffic makes implementing an anomaly-based IDS particularly challenging. This is especially true in large organisations, where different departments and users are likely to generate different types of traffic, such as web traffic, SMTP traffic and telnet sessions.

As a result of the above-mentioned problems, an IDS is likely to require excessive processing power to analyse the network information in a timely manner and is rather complex to implement [3]. The authors, therefore, propose incorporating the following techniques into an IDS implementation to improve its performance: Distributed Analysis, Attack-time Isolation, and Base-line Reduction.

Distributed Analysis would accelerate the detection of attacks by distributing the computational load of analysing the network information over multiple computers, as used in a distributed processing environment and a Distributed Intrusion Detection System (DIDS) [2, 15].

Attack-time Isolation is a method for identifying the general time period during which an attack occurred, reducing the network information that needs to be analysed by allowing the IDS to “zoom-in” on network information that occurred during that specific time period. Thus, in turn, reduces the time required to isolate the actual attack information and, in the case of a real-time IDS, allows countermeasures to be deployed much quicker. Administrators can also quickly fix the security “hole” that is being exploited and mitigate the damage caused by the attack.

Base-line Reduction reduces the complexity inherent in determining a base-line for an anomaly-based IDS by creating an individual base-line for each employee (hereafter called a user) within the organisation. The following example demonstrates the advantages of individual base-lines over an organisational base-line.

Assume an organisation implements an anomaly-based IDS on an SMTP (email) server. The base-line it sets for the SMTP server is that *“no more than 50 emails are sent to the SMTP server per second, with a threshold of 5 emails per second”*. In other words, should more than 55 emails per second be sent, it should be considered a possible attack on the system.

If more than 55 legitimate emails are sent to the SMTP server by employees within the organisation—due to an internal organisational poll or survey, for example—the IDS on the SMTP server will register these events as an attack, in this instance a “false-positive”. If, however, the individual base-line is that *“no more than one email is sent to the SMTP server per second”*, even if the whole organisation sends an email at exactly the same instance no attack would be registered because the event would still be within the threshold of the individual base-lines.

To facilitate the use of individual base-lines, a personal IDS must monitor a specific user’s network requests. Therefore, instead of a single IDS monitoring all network requests from all users and comparing that to a single base-line, a single IDS monitors network requests from a single user and compares them to a base-line specific to that user.

The solutions mentioned above—Distributed Analysis, Attack-time Isolation and Base-line Reduction—may sound simple, but are not feasible because of the overhead involved in implementing individual IDSs and creating individual base-lines for every employee within the organisation and analysing all the data within the log files to determine the time period during which an attack occurred. Consequently, the following section proposes a smart card-based IDS implementation model that overcomes these issues.

4 A Smart Card-based IDS Model

The following sections outline the proposed model: Section 4.1 illustrates how smart cards can be used to introduce IDSs at an individual level; Section 4.2 details how user behaviour can be logged to implement effective intrusion detection; Section 4.3 outlines how the smart card detects anomalous behaviour through a request service, and in section 4.4 a feedback process that can be used to analyse the discovered anomaly in detail is discussed.

4.1 Introducing the Smart Card IDS

The model proposes an IDS environment based on the principle of Distributed Analysis and Base-line Reduction to detect anomalies within individual users' network requests. The IDS environment is implemented on a smart card and is referred to as a smart card IDS (SCIDS). Smart card technology is used within this model due to its inherent security, mobility, scalability and low manufacturing/implementation costs. However, the technical specifications of the smart card are beyond the scope of this paper. The IDS stored on the smart card is a host-based IDS. All users within an organisation are issued with a smart card. The SCIDS then monitors all network requests originating and terminating at the computer from which the user is currently working on. Figure 1 shows a UML component diagram of the SCIDS, together with the organisation's network and conventional IDS.

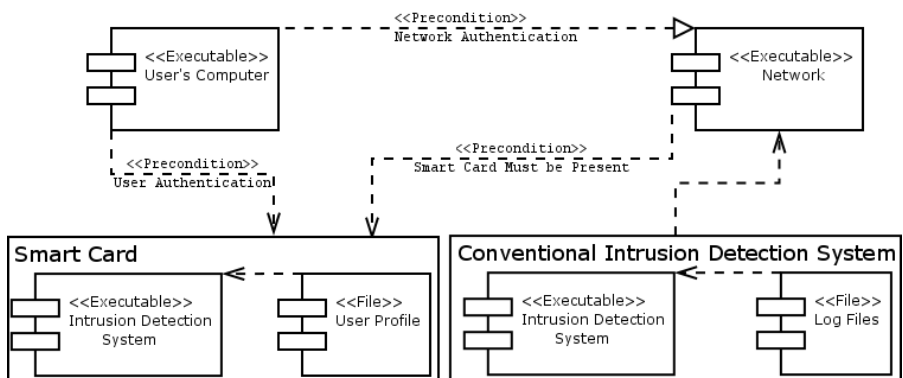


Fig. 1. UML component diagram of the Smart Card-based Intrusion Detection System.

In order to unlock the computer's network interface, the user inserts his/her smart card and authenticates his/her identity to the smart card, as shown in figure 2. While the smart card can be viewed as an authentication token, in this case its primary role is to act as an IDS. Therefore, the user still needs to

authenticate him/herself to the computer network. To make certain that the SCIDS can monitor the user’s network requests at all times, the smart card has to be present in the computer for the network to work, a precondition illustrated in figure 1.

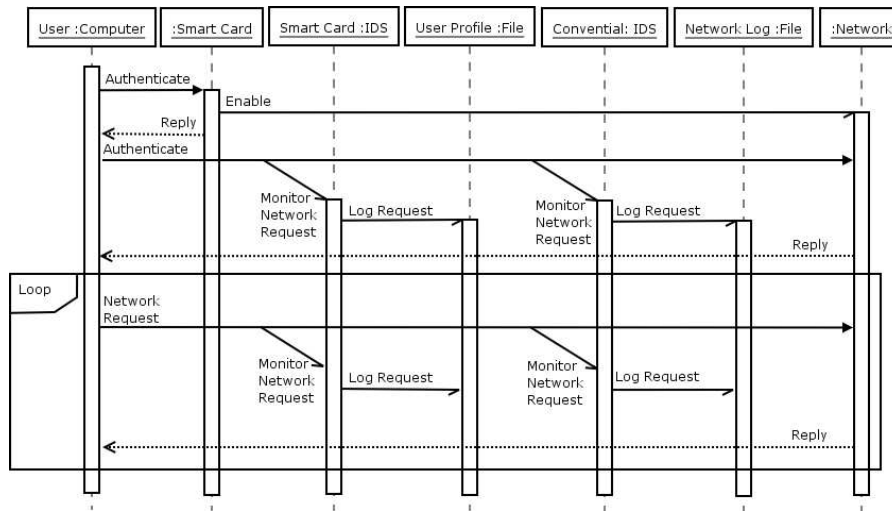


Fig. 2. UML sequence diagram illustrating the user authenticating process (both smart card and network), while the SCIDS and conventional IDS monitor the network requests.

The next section outlines how the SCIDS monitors and records user behaviour in further detail.

4.2 Tracking User Behaviour

The SCIDS tracks the network behaviour of the user, specifically, how many network requests the user makes on the network. For example, web requests, web-search requests, TCP/IP requests and UDP requests (hereafter collectively called network requests). The SCIDS creates a user network behaviour profile by monitoring the network requests made between the user’s computer and the organisation’s network. Due to resource limitations on the smart card, only the number of network requests a user makes per minute will be recorded by the SCIDS. Therefore, the number of network requests a user makes within a given time period, is consider a user’s network behaviour profile. In addition, the organisation’s network is monitored by a conventional IDS which maintains the integrity of the organisation’s systems and protects the network from inside and outside attack as usual. The conventional IDS also creates a user network behaviour profile for each user on the network. Therefore, a user profile exists

both locally, on the user's smart card, and centrally, on the conventional IDS. The user network behaviour profile is represented graphically as a graph in figure 3, showing the number of network requests issued by the user per quarter hour over a particular time period. This user network behaviour profile acts as the base-line that the SCIDS can then use to detect anomalies that might indicate an attack on the system. The next section discusses how this user profile can be used to discover an attack in further detail.




Fig. 3. Visual representation of the user's network behaviour.

4.3 Anomaly Detection

At every network log-on, the SCIDS sends a request to the organisation's conventional IDS for the user's network behaviour profile from the conventional IDS's perspective, as shown in figure 4. While the user's network behaviour profile is recorded by both the SCIDS and the organisation's conventional IDS, the network behaviour profile recorded by the SCIDS is regarded as the base-line for the user. Once the request sent by the SCIDS has been fulfilled, the SCIDS compares the network behaviour profile received from the conventional IDS to that of the user's base-line, looking for any anomalies that could suggest that an attack has occurred, as shown in figure 4. It is important to note, that the smart card does not detect intrusion on the system, but whether an anomaly between the user behaviour profile recorded by the SCIDS and the conventional IDS exists. Anomalies are any points within the network behaviour graph that do not correspond to the graph of the user's base-line, as shown in figure 5 (note the anomaly between time period 14 and 15).

These anomalous points correspond to an attack on the system using the user's network credentials. The network credentials, i.e. user ID and password,

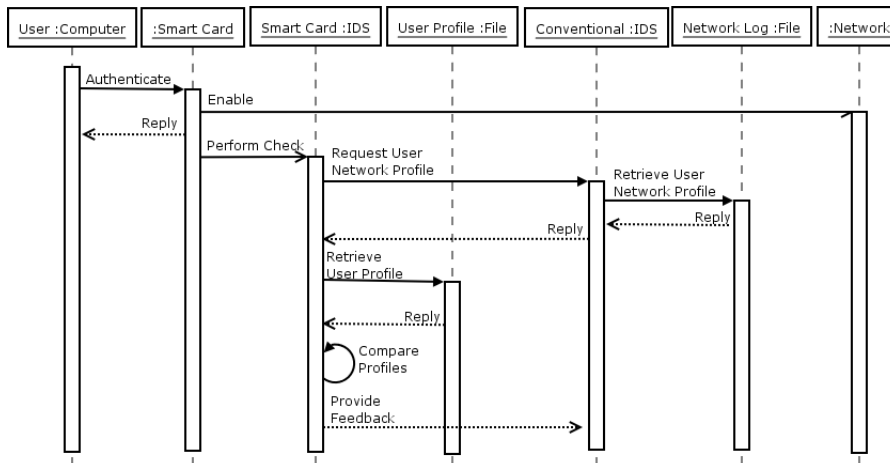


Fig. 4. UML sequence diagram depicting how the SCIDS requests the user behaviour profile from the conventional IDS, analyses the received user profile for anomalies and provides feedback.




Fig. 5. Visual representation of detecting anomalous behaviour.

were stolen and used by another user (hereafter called an attacker) to commit the attack on the system. The following section details how the SCIDS handles the discovery of the attack.

4.4 Anomaly Feedback

A smart card has limited processing power, so the SCIDS cannot completely analyse the extent of the attack, the damage caused by the attack or its origin. The SCIDS has performed an important step though, it has isolated the general time period during which the attack occurred by noting the time the anomaly occurred. Therefore, the SCIDS has been able to achieve Attack-time Isolation. The SCIDS also provides feedback to the organisation’s conventional

IDS regarding the anomaly that has been discovered, as shown in figure 4. This feedback informs the conventional IDS of the general time period that the attack occurred within and the network credentials that were used to perform the attack.

The conventional IDS has detailed log files of the network requests at its disposal to analyse and uncover the exact details of the attack and the attacker. The system administrator will also be able to instruct the conventional IDS to “zoom-in” on a specific time period within these log files (as specified by the information from the SCIDS). The search can be further refined by filtering out only the network requests associated with the specific user ID (again, as obtained from the SCIDS feedback information). Depending on the size of the network, log files typically have large amounts of entries, in the order of millions for large networks. Depending on the duration of the attack, the refined search of the log files could narrow down the enquiry to less than a few hundred entries that need to be analysed. This significantly smaller quantity of log file entries would make the job of applying vulnerability assessment tools and analysing the data to determine the damage that the attack has caused much easier for the conventional IDS.

At every log-on the SCIDS examines the user network behaviour profile to detect whether or not an attack has occurred, the same way a passive IDS would. However, the fact that the SCIDS can effectively isolate the time period during which an attack might have occurred, simultaneously reduces the time required to discover an attack and accelerates the reaction time required to respond. Hence, the SCIDS is defined as a semi-passive, anomaly-based IDS. Semi-passive because it does not detect attacks in real-time, but it actively decreases the time between discovery of the attack and response to the attack. Anomaly-based because, as explained in the previous section, the SCIDS discovers attacks in the same manner as an anomaly-based IDS: by looking for anomalies between the user network behaviour profile created by the SCIDS and the conventional IDS.

The following section compares the advantages and disadvantages of the SCIDS with a conventional IDSs.

5 Smart Card-based IDS versus Conventional IDSs

The smart card-based IDS (SCIDS) proposed in this paper is ideally suited to monitor network requests within an organisation and whether an employee’s network credentials—such as user ID and password—have been stolen by discovering anomalies within the network requests. It is difficult for conventional IDSs to discover an attack disguised with a stolen network credential, as generally, network requests from users within the organisation should be trusted.

However, the SCIDS is not a complete IDS solution in itself, but rather compliments the entire IDS environment. For example, the SCIDS is unable to discover attacks that originate from either, other internal users or from outside

the organisation. Nor is it able to discover an attack perpetrated without the use of network credentials. For example, if attacker *A* initiates a Denial-of-Service (DoS) attack on the SMTP server by flooding it with email messages—an attack that can be performed without the use of network credentials—no SCIDS would have detected it. Nor can a SCIDS discover an attack on the system if the attacker is actually a legitimate user with the organisation.

For example, assume user *B*, using his own network credentials, is able to attack the system and access information that he does not have privileged rights to access. His SCIDS is not able to discover this type of an attack because it will not discover any anomalies between the network behaviour recorded by the SCIDS and the conventional IDS, as would be the case if user *B*'s network credential was stolen.

Even though a conventional IDS would be required to discover and handle such attacks, in user *B*'s case, the SCIDS could prove that he perpetrated the attack. Once the conventional IDS has discovered the attack and traced it back to user *B*, he will not be able to use the theft of his network credentials as a defence. This is because the SCIDS has recorded the same network behaviour as the conventional IDS and did not detect any anomalies—verifying the conventional IDS's suspicions that user *B* committed the attacks.

6 Conclusion

The smart card-based IDS model proposed in this paper addressed the capacity problems and inefficiencies IDSs face due to the volume of network requests that need to be analysed and the complexities of implementing an anomaly-based IDS. The SCIDS is able to achieve Distributed Analysis by having a single IDS monitor the network requests of a single user. The SCIDS achieves Base-line Reduction by creating a single base-line for each user and, hence, reduces the complexity of anomaly detections. Finally, the SCIDS is able to achieve Attack-time Isolation by determining the general time period during which an attack occurred and, as a result, increase the efficiency with which attacks can be discovered and handled. Overall, the SCIDS has been shown to complement the conventional IDS environment rather than being a full-blown solution.

There are clear privacy issues that arise once a conventional IDS can generate individual user profiles from data stored in the organisation's log files. These issues need to be addressed by future work and a balance between acceptable levels of privacy and the necessary levels of IDS efficiency achieved.

7 Acknowledgement

This material is based upon work supported by the National Research Foundation under grant number 2054024. Any opinion, findings, conclusions or recommendations expressed in this material are those of the authors and the NRF does not accept any liability thereto.

References

1. Dorothy E. Denning. An intrusion-detection model. *IEEE Transactions on Software Engineering*, SE-13(2):222–232, February 1987.
2. Biswanath Mukherjee, L. Todd Herlein, and Karl N. Levitt. Network intrusion detection. *Network, IEEE*, 8(3):26–41, May 1994.
3. Wenbao Jiang, Hua Song, and Yiqi Dai. Real-time intrusion detection for high-speed networks. *Computers & security*, 24:287–294, 2005.
4. R.D. Alexander. The search for a general theory of behaviour. *Behavioural Science*, 20(2):77–100, 1975.
5. Craig B. Stanford. The social behaviour of chimpanzees and bonobos: Empirical evidence and shifting assumptions. *Current Anthropology*, 39:399–420, August 1998.
6. S. C. Cheung and J. Kramer. An integrated method for effective behaviour analysis of distributed systems. In *ICSE '94: Proceedings of the 16th International Conference on Software Engineering*, pages 309–320, Los Alamitos, CA, USA, 1994. IEEE Computer Society Press.
7. Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, and Eiichi Osawa. Robocup: The robot world cup initiative. In *AGENTS '97: Proceedings of the first international conference on Autonomous agents*, pages 340–347, New York, NY, USA, 1997. ACM Press.
8. A. Kerepesi, E. Kubinyi, G.K. Jonsson, M.S. Magnussin, and Á. Miklósi. Behavioural comparison of human-animal (dog) and human-robot (aibo) interactions. *Behavioural Science*, 20(2):77–100, 1975.
9. Andreas Fuchsberger. Intrusion detection systems and intrusion prevention systems. *Information Security Technical Report*, 10:134–139, 2005.
10. John Wilander and Mariam Kamkar. A comparison of publicly available tools for static intrusion prevention. In *Proceedings of the 7th Nordic Workshop on Secure IT Systems*, pages 68–84, Karlstad, Sweden, November 2002.
11. Roberto Battistoni, Emanuele Gabrielli, and Luigi V. Mancini. A host intrusion prevention system for windows operating systems. *Lecture Notes in Computer Science*, 3193:352–368, 2004.
12. Mike Hendry. *Smart Card Security and Applications*. Artech House, April 2001.
13. Efraim Turban and Debbie McElroy. Using smart cards in electronic commerce. *International Journal of Information Management*, 18(1):61–72, February 1998.
14. Sebastian Münscher. Smartcard security. Technical report, NamITech, Giesecke & Devrient, November 2004.
15. Mauro Cesar Bernardes and Edson dos Santos Moreira. Implementation of an intrusion detection system based on mobile agents. In *Software Engineering for Parallel and Distributed Systems, 2000. Proceedings*, pages 158–164, 2000.