Intrusion Detection via Artificial Immune System: a Performance-based Approach

Andrea Visconti, Nicoló Fusi, Hooman Tahayori

Abstract In this paper, we discuss the design and engineering of a biologically-inspired, host-based intrusion detection system to protect computer networks. To this end, we have implemented an Artificial Immune System (AIS) that mimics the behavior of the biological *adaptive immune system*. The proposed AIS, consists of a number of running artificial white blood cells, which search, recognize, store and deny anomalous requests on individual hosts. The model monitors the system through analysing the set of parameters to provide a general information on its state — ill or not. When some parameters are discovered to have anomalous values, then the artificial immune system takes a proper action. To prove the effectiveness of the suggested model, an exhaustive test on the AIS is conducted, using a server running Apache, Mysql and OpenSSH, and results are reported. Four types of attacks were tested: remote buffer overflow, Distributed Denial of Service (DDOS), port scanning, and dictionary-attack. The test proved that our definition of self/non-self system components is quite effective in protecting host-based systems.

1 Introduction

Artificial Immune Systems (AISs) are inspired by the workings of the biological immune systems [1, 2, 3], and focus their capability to recognize elementary *self*

Andrea Visconti
Universitá degli Studi di Milano, Dipartimento di Informatica e Comunicazione, Via Comelico 39/41 Milano 20135 Italy, e-mail: andrea.visconti@unimi.it

Nicoló Fusi
Universitá degli Studi di Milano, Dipartimento di Informatica e Comunicazione, Via Comelico 39/41 Milano 20135 Italy, e-mail: nicolo.fusi@studenti.unimi.it

Hooman Tahayori
Universitá degli Studi di Milano, Dipartimento di Scienze dell'Informazione, Via Comelico 39/41 Milano 20135 Italy, e-mail: hooman.tahayori@unimi.it

components of the body — endogenous or innocuous — and elementary *non-self* components of the body — exogenous or potentially pathogenic. Artificial immune systems proposed in the last decade were mainly studied and implemented for solving real-world problems (spam filtering, intrusion detection, pattern recognition, etc.). In particular, the most representative applications of artificial immune systems are from the area of computer security and fault detection [4].

Artificial immune systems, based on the recognition of self/non-self behavior, require unambiguous definitions of all permitted and/or non-permitted actions in a system. A number of ground-breaking solutions to this problem are proposed [5, 6, 7, 8, 2]. In particular, interesting approaches suggested in [5, 9] introduce the possibility of using a sequences of system calls, executed by running UNIX processes, as discriminator between normal and abnormal behavior. Moreover, references [10, 11, 12] discuss the design and testings of Lisys, a LAN traffic anomaly detector that monitors TCP SYN packets for detecting unusual requests, alerting administrator when abnormal connections are detected. In [13], authors suggested collecting data on normal and abnormal behaviors in host-based and network-based systems. Such data, collected in a realistic context, contains information that may be used for automatically detecting, analyzing and controlling future anomaly behaviors due to new and unpredictable network attacks. Tarakanov et al. in [7] described an artificial immune system based on a rigorous mathematical approach, that applied the singular value decomposition to the matrix of connection logs and mapped the users' requests into a real two-dimensional vector space. Authors argued that similar self/non-self requests lump together. In [14, 15] authors suggested interesting approaches based on the Danger Theory [16]. This new theory has shifted control of immunity to the tissues that need protection. Inspired by the behavior of innate immune system, Pagnoni and Visconti in [13] illustrated an artificial immune system based on the working of the macrophages. The authors argued that the main idea of an intrusion detection system is not to recognize and kill a specific intruder in the most effective way, but rather to find and kill any intruder as soon as possible. In addition to the artificial immune systems previous mentioned, authors in [5, 17, 8, 18] suggested applying the negative selection algorithm to the problem of network intrusion detection.

Unfortunately, none of these solutions has achieved one hundred percent precision. Nevertheless, real-world applications have the necessity of (a) providing a strong, reliable discrimination between normal and abnormal behavior and (b) maintaining a complete database of "good or bad behavior" to be used by the self/non-self recognition algorithms. The choice of self/non-self behaviors is crucial because some bad behavior not stored in the database may not be recognized as network attacks; moreover a large database of self/non-self behaviors entails a substantial degree of slowness that is not acceptable in real-time applications.

In order to overcome these problems, we have designed a biologically-inspired intrusion detection system based on the paradigms of the *acquired immune system*. Being more slow than the innate one, the acquired immune system is the only one that remembers the previously encountered attacks, recognizes new attacks of unwanted intruders entering the system, and provides a proper response to the attack of

the enemies. Diagnosing an abnormal behavior of a specific type requires knowing which, if any, set of parameters characterizes the anomaly. This set of parameters is called *antigen signature*. Some such signatures are well-known, and can be easily recognized automatically, others are just less well-defined, and can be more difficult to recognize, whilst others are completely unknown. To this end, we analyze and improve existing solutions in computer security through the design, engineering and testing an intrusion detection system that recognizes anomalous values of parameters of a given system. These anomalous values can be interpreted as a sign of an attack to the system comparable with fever in the case of presence of infection in body.

In this paper, we suggest an artificial immune system based on several agents that mimics the behavior of white blood cells in the acquired immune system. Such white blood cells — or lymphocytes, — cooperate using a specific communication protocol in order to protect the system against exogenous or endogenous attacks. Every white blood cell is a separate process that monitors the parameters of the system and checks the presence of non-self antigen signatures.

In the sequel, in section 2 we discuss the principles of immune systems, in particular, we focus our attention on the adaptive immune systems. In Section 3, we describe the design and engineering of our artificial immune system; while an intensive testing is presented in section 4. Finally in Section 5, we provide pros and cons of our model and draw conclusions.

2 Biological Immune System

Biological immune systems draw up several lines of defense to protect the organism. This defense systems include *chemical* and *physical barriers*, *innate immune system* and *adaptive immune system*.

Chemical and physical barriers provide the first line of defense in the fight against invaders. Examples of the chemical and physical barriers are skin, gastric acid in the stomach, eyelashes, tears, and so on. These barriers try to protect the body against pathogens that enter an organism, and consequently reduce the probability that the pathogens will lead to an illness. Unfortunately, some foreign invaders that are present on the skin surface pass through injuries on the skin.

When the chemical and physical barriers fail to stop unwanted intruders, invaders are attacked by the cells of the second line of defense: the *innate immune system*. The innate immune system recognizes and attacks invaders in a generic way, with no necessity of previous exposure to them. The cells involved in the innate reaction are *leukocytes* such as macrophages, natural killer cells, mast cells, basophils, and so on. These leucocytes (a) release chemical factors that cause inflammation, swelling and local blood vessel dilation; (b) recruit immune cells to sites of infection; (c) attack everything of a foreign nature, engulfing pathogens and dead cells in a process called phagocytosis; (d) and finally, activate the adaptive immune system.

The *adaptive* or *acquired immune system* provides the third line of defense in the fight against intruders. It does not replace the innate immune system, but rather improves it. The ability of the adaptive immune system to kill invaders is based on the capacity of recognizing several kinds of pathogens and remembering specific antigen signatures after the resolution of the infection. Comparing to the previous two lines of defense, the adaptive immune system works in a more complex way because its responds to an attack is antigen-specific. Being exposed to different pathogens, the adaptive immune system learns to identify enemies and as a result its specific response will be more effective than a generic response of the innate immune system.

The cells involved in the acquired reaction are *leukocytes* such as memory B cell, killer T cell, helper T cells, and so on. When activated, these leucocytes are able to (a) distinguish the cells of the body from unwanted invaders; (b) recognize specific signature for each non-self antigen; (c) generate a specific immune response against invaders; (d) remember specific signatures for each non-self antigen; (e) and eventually, quickly remove the previously encountered non-self antigen.

Unfortunately, these lines of defense, that generate a powerful barrier against intruders, are not perfect and sometimes fail. Failures occur when the ability to fight invaders of one or more components is reduced — immunodeficiency — or the ability to recognize self and non-self cells is compromised — autoimmunity. In both cases, an organism will be vulnerable to infections.

3 Design and Engineering of Artificial Immune System

To achieve the ultimate goal of designing and engineering an intrusion detection system (IDS) based on the workings of the acquired immune system, clear understanding of the characteristics of the acquired immunity is of great importance. In particular, we concentrated on (a) the acquisition of a clear discrimination between self — regular — and non-self — unwanted — system behaviors, (b) the elimination of recognized infections — recognized attacks —, (c) the take care of system injuries — bugs, — (d) the detection and elimination of new infections — new attacks, — and (e) the absence of autoimmune reactions.

The design and engineering of the proposed AIS is based on the previous points and on the observation that getting into a system without leaving any track is virtually impossible. We search for these tracks by considering the parameters of the system when our server is under attack. In order to identify such tracks i.e. the fingerprint of attacks, the values of different system parameters were surveyed. The gathered data was analyzed and more than 150 graphs were generated. Some of them did not show any significant change during an attack, while some did. We conceived the artificial immune system with these ideas in mind.

The proposed AIS is a host-based system that consists of a set of processes — helper T-cells, killer T-cells, and memory B-cells — running on a server. These processes act and cooperate as digital lymphocytes in order to discover suspicious

values of the parameter of the system and face external attacks. The artificial immune system must be initialized through a training phase, in which the AIS defines the number of running lymphocytes. In fact, this number is not constant, but is optimized experimentally because it depends on the hardware features and the workload of the server. This number can vary between a lower and an upper bound, improving the performance of the AIS when the system is under attack. Furthermore, if such number falls under the optimized threshold, new lymphocytes are automatically created. In addition to the task of defining the number of running lymphocytes, the training phase is also responsible for setting up the parameters of all processes. Indeed, the digital lymphocytes learn to identify self/non-self behaviors, analyzing the data of system parameters under the supervision of an expert. After the training phase, the artificial immune system is ready to be activated.

Helper T-cells Are processes, or agents, of the artificial immune system. Their main objectives are identifying an anomalous behavior in the monitored parameters and promoting the activation of adaptive immune response. In order to do so, helper T-cells collect a sample data of actual system parameters, compute the mean and the standard deviation of such data, and compare the current values to the previously stored values. If the mean of each parameter monitored exceeds a given threshold — mean stored in memory plus or minus the standard deviation stored in memory, — the current interval is defined non-self. Moreover, helper T-cells try to identify the type of attack using type-1 fuzzy rules and promoting a quick immune response. Indeed, they stimulate other cells of immune system, controlling and inhibiting immune attacks against self-antigens.

Every helper T-cell has a lifespan at the end of which the cell dies. This means that the cell will be regenerated, will undergo the negative selection phase, and hence the ability to recognize possible unwanted attacks will be improved.

Memory B-Cells Are processes that remember attacks previously encountered. When they recognize a set of parameters that show an anomalous behavior previously identified, the artificial immune system stops the recognition phase and memory B-cells inform killer T-cells that the system is under attack, specifying what kind of attacks is.

Killer T-cells Are processes that take proper actions for denying and eliminating unwanted behaviors. If the attack type is known, a predefined action will take place. For example, in presence of a denial of service attack, killer T-cells deny unwanted requests, banning the IP addresses of the senders, while in presence of a reverse shell techniques, killer T-cells eliminate such shells. On the other end, if the attack type is unknown, a notification is sent to the system administrator, alerting him of the security threat.

All digital cells are separate Java processes, so if one process crashes it will not affect the working of the artificial immune system. Designing the system as a set of different processes offered greater security and stability, at the price of a more difficult communication between processes. In order to solve this problem, we implemented a communication protocol. When authenticated, processes can commu-

nicate with each other by calling specific functions. This communication protocol is necessary to stimulate groups of lymphocytes, or the entire artificial immune system, when the system is under attack.

4 Testing

All tests were made on a dual Intel® PIII® server with 1.5Gb of RAM on which Gentoo linux [19] was installed; running Apache 2.0.59 with PHP 5.2.5, Mysql 5.0.40 and OpenSSH 4.7. Several types of attacks were tested: port scanning (NMAP), remote buffer overflow, Distributed Denial of Service (DDOS) and dictionary-attack against SSH authentication. These types of attack were chosen because they have very different performance fingerprints and are the most common.

In order to simulate a real-world situation, a fake institutional website was implemented and the system was tested with different amounts of traffic:

None	No one, or a small number of users, surf the website. This profile has been used in order to provide baseline data.
Moderate	An average number of users surf the website.
Intense	A really high amount of users surf the website, making a huge number of self requests. This profile has been used in order to evaluate the behavior of the system under severe stress condition.

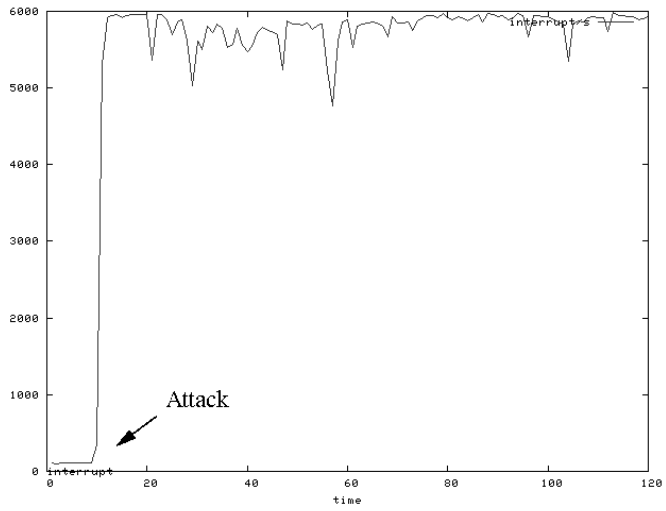

Figure 1 Interrupts per second during a DDOS attack without legitimate traffic

The population of digital cells used for the testing is as follows: 3 to 5 killer T-cells, 5 to 10 helper T-cells and 3 to 5 memory B cells.

The traffic has been simulated with JMeter [20], a stress testing tool for web applications provided by the Apache software foundation.

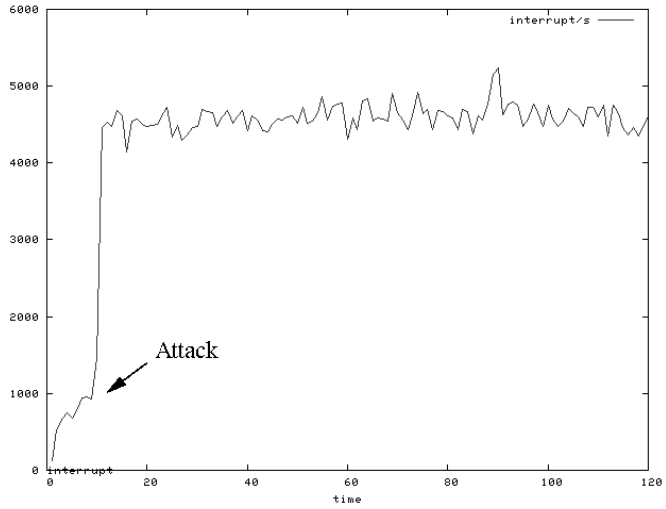

Figure 2 Interrupts per second during a DDOS attack under moderate traffic

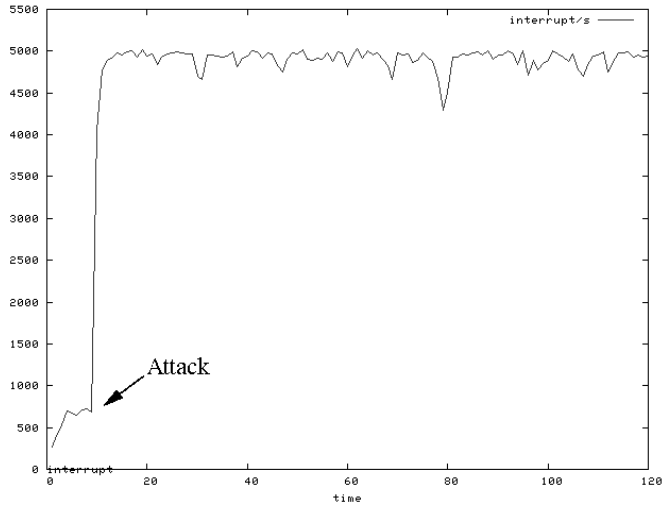

Figure 3 Interrupts per second during a DDOS attack under intense traffic

As mentioned in Section 3, our AIS monitors the system parameters in order to identify unusual patterns that may be related to unwanted behaviors. For example,

figures 1, 2 and 3 show the behavior of a system parameter during a DDOS attack. Analyzing the values of this parameter, an unusual pattern in the number of interrupt requests can be recognized.

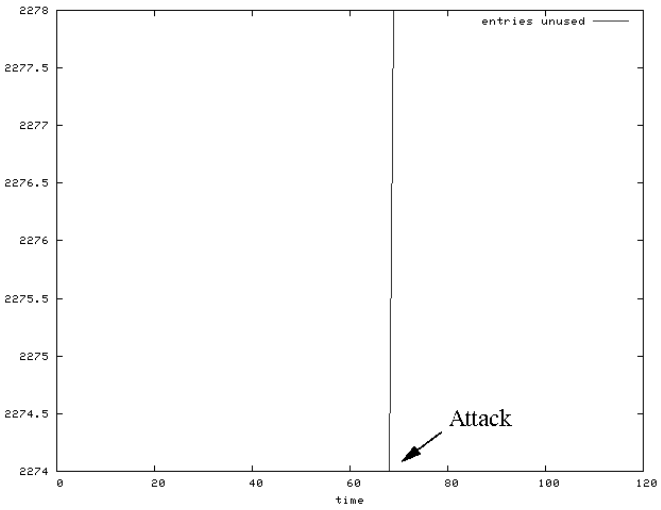

Figure 4 Unused cache entries during a Buffer Overflow attack under moderate traffic

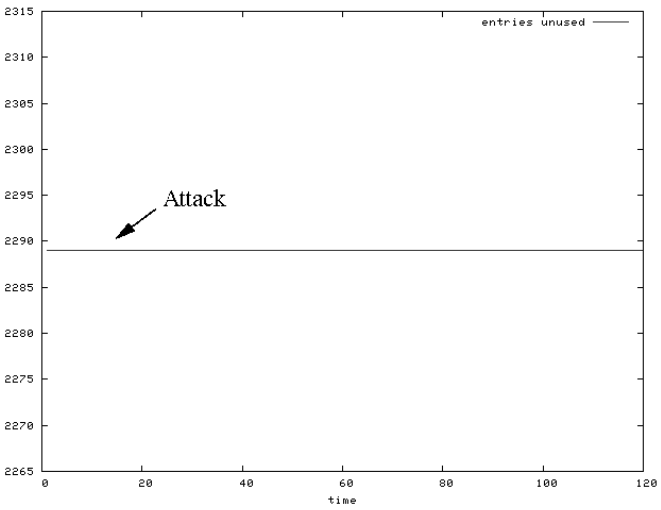

Figure 5 Unused cache entries during a Buffer Overflow attack under intense traffic

Figures 1, 2, and 3 show clearly that DDOS attack is rather easy to spot, because it largely affects the system performance. Unfortunately, as can be seen in the figures 4 and 5, not all kinds of attack are always so easy to recognize. Such figures represent the unused cache entries during a buffer overflow attack while an average and large number of users, respectively, surf the website. In the first two cases, — figures 4 and 5, low and moderate workload — the artificial immune system is able to recognize the anomaly. In the last case — figure 6, high workload — the artificial immune system fails. It is easy to see that a high workload situation may introduce an excessive level of background noise, decreasing the ability of the system to recognize anomalous behaviors. Such situations affect negatively the performance of the artificial immune system, enhancing the risk of false positives.

Tables 1, 2, and 3 summarize the results of the testing activity. For each of the four attacks tested, we mark the set of system parameters that may indicate the presence of an attack.

Table 1 No traffic

Attacks	UsedIH	UnusedCE	PGfault	SysCPU	Interr	TRate
Buf Overfl.	X	X	X			X
DDOS				X	X	
Scan NMAP	X	X	X	X		X
Bruteforce			X	X		

Table 2 Moderate traffic

Attack	UsedIH	UnusedCE	PGfault	SysCPU	Interr	TRate
Buf Overfl.	X	X	X			
DDOS				X	X	
Scan NMAP						
Bruteforce			X	X		

Table 3 Intense traffic

Attack	UsedIH	UnusedCE	PGfault	SysCPU	Interr	TRate
Buf Overfl.	X		X			
DDOS				X	X	
Scan NMAP	X	X				X
Bruteforce			X	X		

As tables 1, 2, and 3 illustrate, recognizing an attack in the presence of many net-surfing users is increasingly difficult, and in some cases is rather impossible.

For example, recognizing a port scanning attack is almost impossible, given the anomalous behavior of the monitored parameters under different traffic profiles (see tables 1, 2 and 3).

The results of an exhaustive testing are summarized in table 4.

Table 4 Test results

	No	Moderate	Intense
Buf Overfl.	95%-100%	75%-90%	65%-85%
DDOS	100%	95%-100%	90%-100%
Scan NMAP	0%-20%	0%	0%-5%
Bruteforce	100%	90%-100%	85%-100%

5 Conclusions and Future Works

The suggested artificial immune system is an IDS based on the idea of equipping servers with the technological equivalence of an acquired immune system. To this end, our AIS monitors and analyzes a set of system parameters to check for anomalous behaviors. Although still at a preliminary stage, the exhaustive testing revealed that AIS is able to quickly detect anomalous behaviors previously encountered, deny proliferation of foreign processes by killing dangerous processes before they will widely used, and recognize attacks with a strong fingerprint such as denial of service and dictionary attack. On the other hand, the intensive testing has proved that in order to avoid a large number of false positives, we have to lower the sensibility of the system, affecting the recognition of some kinds of attack. Indeed, under these circumstances the system cannot recognize an NMAP scan; moreover, in presence of many net-surfing users recognition is increasingly difficult, and sometimes is impossible.

It should be stressed that yet, no single method, biological or artificial, can achieve one hundred percent precision. For these reasons, the suggested system is not meant to replace firewalls, login policies, or antivirus because it cannot blocks every kind of attack. The proposed artificial immune system should be used in conjunction with other complementing technologies either biologically inspired or not.

Our future works will be devoted to improve the AIS, extending actual acquired immunity with specific components of a second line of defense: the innate immunity.

6 Acknowledgements

This research was funded by the State of Italy via FIRST (Fondo per gli Investimenti nella Ricerca Scientifica e Tecnologica).

References

1. D'haeseleer, P., Forrest, S., Helman, P.: An immunological approach to change detection: algorithms, analysis and implication. In: Proceedings of the 1996 IEEE Symposium on Computer Security and Privacy, (1996)
2. Forrest, S., Hofmeyr, S., Somayaji, A., Longstaff T.: A sense of self for UNIX processes. In: Proceedings of the 1996 IEEE Symposium on Research in Security and Privacy, (1996)
3. Forrest, S., Hofmeyr, S., Somayaji, A.: Computer immunology. In: Communication of ACM **40**(10), 88-96 (1997)
4. Dasgupta, D.: Advances in Artificial Immune Systems. In: IEEE Computational Intelligence Magazine, (November 2006)
5. Hofmeyr, S., Somayaji, A., Forrest, S.: Intrusion Detection using Sequences of System Calls. In: Journal of Computer Security **6**(3), 151-180 (1998)
6. Dasgupta, D.: Immune-based intrusion detection system: A general framework. In: Proceedings of the 22nd National Information Systems Security Conference, (1999)
7. Tarakanov, A.O., Skormin, V.A., Sokolova, S.P.: Immunocomputing: Principles and Applications. Springer-Verlag, New York (2003)
8. Forrest, S., Glickman, M. R.: Revisiting LISYS: Parameters and Normal behavior. In: Proceedings of the 2002 Congress on Evolutionary Computation, (2002)
9. Warrender, C., Forrest, S., Pearlmutter, B.: Detecting intrusions using system calls: Alternative data models 1999. In: IEEE Symposium on security and Privacy, (1999)
10. Hofmeyr, S., Forrest, S.: Architecture for an artificial immune system. In: Evolutionary Computation, **8**(4), 443-473 (2000)
11. Hofmeyr, S.: An immunological model of distributed detection and its application to computer security. In: PhD thesis, University of New Mexico, (1999)
12. Balthrop, J., Forrest, S., Glickman, M.: Revisiting lisys: Parameters and normal behavior. In: Proceedings of the Congress on Evolutionary Computation, (2002)
13. Pagnoni, A., Visconti, A.: An Innate Immune System for the Protection of Computer Networks. In: Proceedings of the 4th International Symposium on Information and Communication Technologies, (2005)
14. Aickelin, U., Cayzer, S.: The Danger Theory and Its Application to Artificial Immune Systems. In: Proceedings of 1st International Conference on Artificial Immune Systems, (2002)
15. Aickelin, U., Bentley, P., Cayzer, S., Kim, J., McLeod, J.: Danger Theory: The Link between AIS and IDS? LCNS 2787, (2003).
16. Anderson, C., Matzinger, P.: Danger: the view from the bottom of the cliff. In: Seminars in Immunology, **12**(3), 231-238 (2000)
17. Kim, J., Bentley, P.: The human Immune system and Network Intrusion Detection. In: Proceedings of 7th European Congress on Intelligent techniquesSoft Computing, (1999)
18. Gonzalez, F., Dasgupta, D.: An Immunogenetic Technique to Detect Anomalies in Network Traffic. In: Proceedings of the International Conference Genetic and Evolutionary Computation (GECCO), (2002)
19. Gentoo linux, available at *http://www.gentoo.org/*
20. Apache JMeter, available at *http://jakarta.apache.org/jmeter/*