

Framework for Intrusion Detection Inspired by the Immune System

Melanie Middlemiss

Information Science Department,
University of Otago,
Dunedin, N.Z.
`mmiddlemiss@infoscience.otago.ac.nz`

Abstract. The immune system is a complex and distributed system. It provides a multilevel form of defence, capable of identifying and reacting to harmful pathogens that it does not recognise as being part of its “self”. The framework proposed in this paper incorporates a number of immunological principles, including the multilevel defence and the co-operation between cells in the adaptive immune system. It is proposed that this approach could be used to provide a high level of intrusion detection, while minimising the level of false negative detections.

1 Introduction

As reliance on computers and networks increases, so does the need to provide appropriate security measures. Unfortunately there is no ‘silver bullet’ that is able to provide complete protection which incorporates different levels of protection for different levels of security risk. One component of this multilevel approach is intrusion detection - the process of monitoring a computer system in order to detect intrusive activity.

There are two main approaches to intrusion detection: misuse (also known as signature) detection, and anomaly detection, with misuse detection the approach most commonly used in commercial intrusion detection systems (IDS) [1]. Clearly, if the signatures are tuned correctly it is possible to obtain high levels of detection accuracy. A problem with this approach is that if a new attack occurs that is significantly different to any of the existing signatures, the IDS is unable to detect the intrusion.

Anomaly detection attempts to solve this problem whereby any activity that differs from what is determined to be “normal” system activity is classified as an attack or intrusive behaviour. This has the potential to identify new attacks that have not been seen before, but has the drawback that it can be difficult to determine what is “normal” system activity. The problem then arises that either normal activity can differ slightly and become classified as intrusive, or intrusive activity can be similar to normal activity and is therefore unable to be identified. Whichever approach is taken, the goal is to have a system that is able to identify existing attacks as well as new and unseen attacks in an accurate and timely manner.

An analogy has been made between IDS and the human immune system, particularly in the area of anomaly detection [2–5]. Our immune system, for the most part, successfully protects our body from harmful pathogens. These pathogens are categorised as virus, worm, bacteria, fungi, or protozoa. Each has a different cellular structure, method of replication and mechanism for entering the body. The immune system has evolved complex methods of identifying these pathogens and removing the threat they possess. The widely held view of the immune system is that its main function is to distinguish between “self” and “non-self” (pathogens) [6]. This distinction is analogous to anomaly detection systems where normal activity (self) is distinct from attacks or intrusive activity (non-self). This analogy has been implemented by Forrest et al. [2] in the form of the negative selection algorithm (NSA), which is based on the immune system principle of clonal selection. The limitation of this and many immune inspired systems (artificial immune systems- AIS), is that it only takes into consideration a small part of the immune system. A recent report by Dasgupta and Yu [7] cites 456 journal articles, conference papers and technical reports in the field of AIS, but despite this volume of work there are only a small number of immunological principles that have been used in these models [8, 9]. If we take a broader look at the immune system there are several features that could be used in the development of an intrusion detection system.

In this paper a framework is presented that incorporates principles from both innate immunity (the “built-in” defence of the immune system) and adaptive immunity (the memory and specific defence). The co-operation of cells within the adaptive immune system is also investigated. It is suggested that the framework provided here could be used to develop an adaptive intrusion detection system which, while being tolerant to normal activity, is able to detect specific attacks as well as anomalous activity.

The remainder of this paper is structured as follows. Section 2 contains an overview of the main approaches to intrusion detection providing a basis for the application of this framework. Rather than presenting an overview of the human immune system, Section 3 outlines the main principles that have inspired the framework presented in Section 4. The paper concludes in Section 5 with the direction of future work.

2 Intrusion Detection

An intrusion is defined as “an attempt to compromise the confidentiality, integrity, availability, or to bypass the security mechanisms” [10] of a computer system¹. Intrusion detection is the process of monitoring a computer system in order to detect misuse or anomalous activity. The aim of an intrusion detection system is to provide accurate and timely detection of intrusive activity. The two main approaches to intrusion detection are discussed below.

¹ In this case the term *computer system* could refer to either a single computer (host-based intrusion detection) or a network of computers (network intrusion detection).

2.1 Misuse detection

Misuse detection identifies intrusive activity by searching for attacks which exploit known vulnerabilities within the system [11]. The assumption is made that all attacks can be described by a pattern or “signature”. A signature is developed for an attack strategy and inserted into the detection system so that next time this attack occurs it can be flagged as misuse of the system. This is a widely used approach in commercial IDS [1]. If an attack is known and can be described by a signature, subsequent detection rates are normally high and the rate of false positive detections is relatively low.

A false positive alert occurs when activity that is not actually intrusive is identified by the IDS as an intrusion. In misuse detection systems this arises if the signature is not as unique as the signature writer assumes. False positives are undesirable and the rate at which they occur is important because if too many incorrect alerts are produced, the system administrator is likely to begin to ignore the alerts or even turn the system off completely [12].

The alternative to false positive alerts is a false negative system response. Misuse detection systems are by design, prone to high levels of false negative system responses. This occurs when the system is unable to detect true intrusive activity and is unable to provide an alert. As misuse detection systems require a signature to be defined in order to be able to identify the attack, slight modifications to the attack can lead to the intrusive activity going undetected by the IDS.

One of the strongest arguments against the use of the misuse detection approach is that the next generation of computer system attacks are likely to come from agents that adapt their behaviour or character over time. Using the analogy of the human immune system, the influenza virus mutates over time with minor changes constantly occurring to virus strains [6]. The immune system is able to adapt to these changes and respond in an appropriate manner to new virus strains that it has not previously encountered. Similarly, an intrusion detection system should be able to adapt its behaviour in order to identify new attacks that it has not encountered before. An approach based purely on misuse detection would be unable to support this adaptive approach. To identify an attack using misuse detection, properties of the attack must already be known so that a signature can be developed and entered into the detection system.

2.2 Anomaly detection

Anomaly detection assumes that intrusive activities are abnormal and searches for anomalous activity that could therefore constitute an attack or intrusion. A profile of “normal system activity” is developed and any activity that differs significantly from this profile is flagged as anomalous and therefore intrusive [13]. This type of intrusion detection cannot specifically identify the attack, rather it indicates that something unusual has occurred.

The advantage of anomaly detection over misuse detection is that it does not need to know about an attack before it can be identified as an intrusion.

However a comprehensive set of “training data” is required in order to form an accurate description of normal system activity.

While misuse detection systems can be tuned for relatively low levels of false positive alerts, anomaly detection systems are prone to high levels of false positives. This arises when some activities, while not actually being intrusive, are not explained by the “normal” profile. In this case they are likely to be incorrectly flagged as intrusive, therefore interfering with normal system activity. According to Axelsson [14] the limiting factor of the performance of IDSs is not the ability to correctly identify intrusions, but rather the ability to suppress false alarms.

3 Principles employed from the immune system

There are many resources available detailing both the human immune system [6, 15–18] and the development of artificial immune systems [8, 19]. For this reason, this section gives an overview of only the particular principles and features of the immune system used in the development of this intrusion detection framework. The principles are summarised in Table 1 along with their application to intrusion detection.

Immune System	Intrusion Detection Framework
Multilevel protection	Multilevel protection
Innate immunity	Signature / misuse detection
Adaptive immunity	Anomaly detection
Lymphocyte	Detector / rule
B cell receptor	Level 1 detector
T cell receptor	Level 2 detector
B cell - T cell co-operation	Error Checking
Memory B and T lymphocytes	Memory detectors (Level 1 and 2)

Table 1. Principles of the immune system and their application to the intrusion detection framework.

3.1 Multilevel defence

The human immune system provides a multilevel defence system. The highest level is the external defence of the skin and other mucosal membranes. This level provides a physical barrier that the pathogen must penetrate in order to enter the body. If the pathogen is able to penetrate the external defences of the body it will then encounter the defences of the innate immune system. If the innate immune system is unable to remove the threat the pathogen presents, the adaptive immune system takes over and removes the pathogen.

In computer security there is no single component or application that can be employed to keep a computer system completely secure. For this reason it is

recommended that a multilevel defence approach be taken to computer security. Northcutt *et. al.* [20] define this multilevel approach as ‘*defence in depth*’ and describe how the layering of multiple security components can provide a more secure computing environment. This would see an intrusion detection system used in conjunction with other security components and procedures, such as encryption, firewall and antivirus systems.

The analogy of the human immune systems multilevel defence could be extended further to the intrusion detection system itself. Once an attack had penetrated the external defences of the computer security system (for example the firewall or other security components), the first level of the intrusion detection system would attempt to identify the attack. If this level was unable to identify the attack the second, more complex, level of detection within the IDS would be enabled. This multilevel approach could provide more specific levels of defence and response to attacks or intrusions.

3.2 Innate immunity

The innate immune system provides the non-specific defence of the immune system. It consists of specialised cells and molecules that are responsible for the initial response to any pathogens that have entered the body. The responses are “built-in” and do not change with age or experience of infection. These cells recognise surface molecules on invading pathogens that have been conserved through evolution and are common to many pathogens.

In intrusion detection there are a number of large databases and public resources that provide expert knowledge of existing computer attacks, for example the Common Vulnerabilities and Exposures List [21]. This information relates to common attacks and the intrusion detection system should be able to recognise these attacks in the way that the innate immune system recognises patterns that are common to many pathogens.

While the innate immune system is non-specific in its defence, the IDS implementation of this would actually be specific. The innate immune system is unable to determine which specific pathogen it is responding to, but the IDS responds to a signature for a specific attack or intrusion. However, as addressed in the previous section, this type of approach (misuse detection) is not adaptive. It is unable to modify its recognition abilities in order to identify new attacks or intrusions, just as the innate immune system does not change with age or experience of infection.

3.3 Adaptive immunity

As opposed to the innate immune system, which is non-specific in its defence against harmful pathogens, the adaptive immune system provides a level of defence that initiates a response specific to the pathogen that has entered the body. Specialist cells (called lymphocytes) have receptors on their cell surface that recognise pathogens (antigen). While each lymphocyte has multiple receptors on its surface, they are all identical and recognise the same specific pathogen.

The adaptive immune system also offers memory capabilities to the immune system. During the adaptive immune system response memory lymphocytes are formed which are specific to the antigen that is being targeted. Subsequent infections by this pathogen will cause the memory cells to be activated, speeding up the response time.

When designing an intrusion detection system it is desirable to have an adaptive system. The system should be able to recognise attacks it has not seen before and then respond appropriately. This kind of adaptive approach is used in anomaly detection, although where the adaptive immune system is specific in its defence, anomaly detection is non-specific. Anomaly detection identifies behaviour that differs from “normal” but is unable to the specific type of behaviour, or the specific attack. However, the adaptive nature of the adaptive immune system and its memory capabilities make it a useful inspiration for an intrusion detection system.

3.4 Lymphocytes

In the adaptive immune system there are two types of lymphocytes involved in the recognition of antigen: B cells and T cells. These lymphocytes begin as a stem cell found in the bone marrow and go through a maturation process. B cells complete their maturation in the bone marrow, while T cells move to the thymus to complete maturation. This process involves clonal selection during which the lymphocytes encounter “self” antigens and are destroyed if they recognise, or match, any of the self antigens. Only lymphocytes that are able to recognise “non-self” antigens will remain.

There have been a number of papers in this area of artificial immune systems that have used the principle of clonal selection, implemented in the negative selection algorithm (NSA), for anomaly detection systems [2]. Of particular interest here is the fact that while B cells and T cells both go through negative selection, T cells also go through positive selection [8, 22]. The reason for this is the difference between B cells and T cells - a distinction that is not made in the NSA.

The B cell receptor (BCR) is the cell surface-bound form of the antibody that is secreted during the immune response. This receptor recognises and binds directly to antigen that is circulating within the body. The point at which the receptor binds to the antigen is called the epitope. During the maturation process of negative selection, B cells are removed if they match any self antigens. This means mature B cells should not recognise self antigens and initiate an immune response to self. Obviously there may be some B cells that manage to circumvent this selection process and end up as mature cells that recognise self. The immune system has mechanisms in place to counter this (e.g. refer Section 3.5).

The T cell receptor (TCR) is similar to the BCR, but with some significant differences. The TCR is not antibody and it is unable to bind directly to antigen. The TCR recognises antigen that has been consumed by another cell and is presented on the surface of that cell [16]. This can be performed by specialist cells (e.g. Antigen Presenting Cells - APC) whose job it is to “eat” antigen and

present it to T cells, or by cells that are infected by the pathogen (e.g. a cell infected by a virus). When the cell consumes the antigen it is broken down into molecules called peptides. When a peptide is presented on the cell surface, it is presented in a complex that includes a special molecule called MHC (Major Histocompatibility Complex) that is unique to the individual.

MHC is a “self” molecule and for T cells to recognise the MHC:peptide complex, they must have a certain level of tolerance towards self [17]. During lymphocyte maturation T cells are subjected to negative selection in the same way that B cells mature. However for T cells, this negative selection is performed in combination with positive selection. During positive selection any T cells that bind with the self MHC are given survival signals by the immune system [23]. If no signal is received the cell dies. T cells that remain after positive selection then undergo negative selection where they are removed if they match the self MHC:peptide complex with a high affinity. This combination of positive and negative selection ensures that the T cells will be tolerant to the self MHC:peptide complex, while also being able to recognise non-self antigens.

In terms of the intrusion detection framework presented in this paper, the distinction between the two types of lymphocyte and the generation of their receptors can be used as an analogy for detectors in the IDS. In the immune system B cells and T cells can be distinguished by their receptors and the level at which they recognise antigen: B cells at the epitope level, and T cells at the peptide level. In an intrusion detection system this distinction can be represented by two types of detector that recognise patterns in data at different levels.

As a potential example, we can think of an intrusion detection system that analyses packets transported across a network. In this situation we can describe packet header information as data at the epitope level. This is the raw data that can be collected from each packet and presented to the Level 1 detectors (analogous to B cells). We can then go a step further and describe the connection statistics, such as volume analysis, as data at the peptide level. This is data relating to the packets that has been internalised by the system and the connection statistics presented to the Level 2 detectors (analogous to T cells).

This is a simplified version of the approach taken by Dasgupta et. al. [24] in which they describe the use of four types of detector analogous to T helper cells, T suppressor cells, B cells and antigen presenting cells.

3.5 B cell - T cell co-operation

One feature of the immune system that is not widely used in artificial immune systems is the co-operation between cells of the adaptive immune system. This arises because the systems deal with lymphocytes as generic detectors and do not make a distinction between T lymphocytes and B lymphocytes.

In the adaptive immune system there are two parts - the humoral immune response and the cell-mediated immune response. The humoral response is driven by direct B cell interaction with antigen and defends against extracellular pathogens. The cell-mediated response is driven by indirect activation of T cells through the recognition of the self MHC:peptide complex, and defends against intracellular

pathogens and cancer [25]. Attacks on a computer system can be thought of as analogous to extracellular pathogens and so, in terms of the intrusion detection framework presented here, we are interested in the humoral immune response.

A B cell directly interacting with antigen initiates the humoral immune response. However the B cell does not become activated and respond to the antigen immediately because it is possible that it may not actually be non-self antigen that has been recognised. The B cell has to wait for a co-stimulation signal to confirm the need to respond to the foreign antigen. This co-stimulation signal comes from a T cell (this particular T cell is called a helper T cell). In order for the T cell to provide the signal it must also have recognised the antigen. Once the B cell has received the co-stimulation signal it is able to become activated and gives rise to cells that can defend against the pathogen.

The problem with anomaly detection systems is that often normal activity is classified as intrusive activity and so the system is continuously raising alarms. The co-operation and co-stimulation between cells in the immune system ensures that an immune response is not initiated unnecessarily, thus providing some regulation to the immune response. Implementing an error-checking process provided by co-operation between two levels of detectors could reduce the level of false positive alerts in an intrusion detection system.

3.6 Memory

The innate immune system does not provide any memory to the immune system. It does not remember patterns or features of new pathogens it encounters. The memory capability of the immune system comes from the adaptive immunity.

In the adaptive immune system, after an immune response is initiated memory cells are created. In the humoral response these are memory B cells and helper memory T cells. In the first instance of exposure to a pathogen the adaptive immune system takes several days before it takes over the immune response. However on subsequent exposure to the same pathogen, memory cells are already present and are ready to be activated and defend the body.

It is important for an intrusion detection system to be adaptive. There are always new attacks being generated and so an IDS should be able to recognise these attacks. It should also then be able to use the information gathered through the recognition process so that it can quickly identify the attack in the future.

4 Framework

The previous section has outlined principles from the human immune system that are seen as important and applicable to intrusion detection systems. This section will describe the intrusion detection framework that has been inspired by these principles. There are basically three parts to the framework (Figure 1): the pre-processing engines and then the two levels of detection - innate detection and adaptive detection. The description provided in this section is general and no attempt has been made to describe specific implementation issues. It is assumed

that this framework could be implemented on a single computer (host-based IDS) or on a network of computers (network-based IDS) and therefore the data being analysed will be dependent on the implementation.

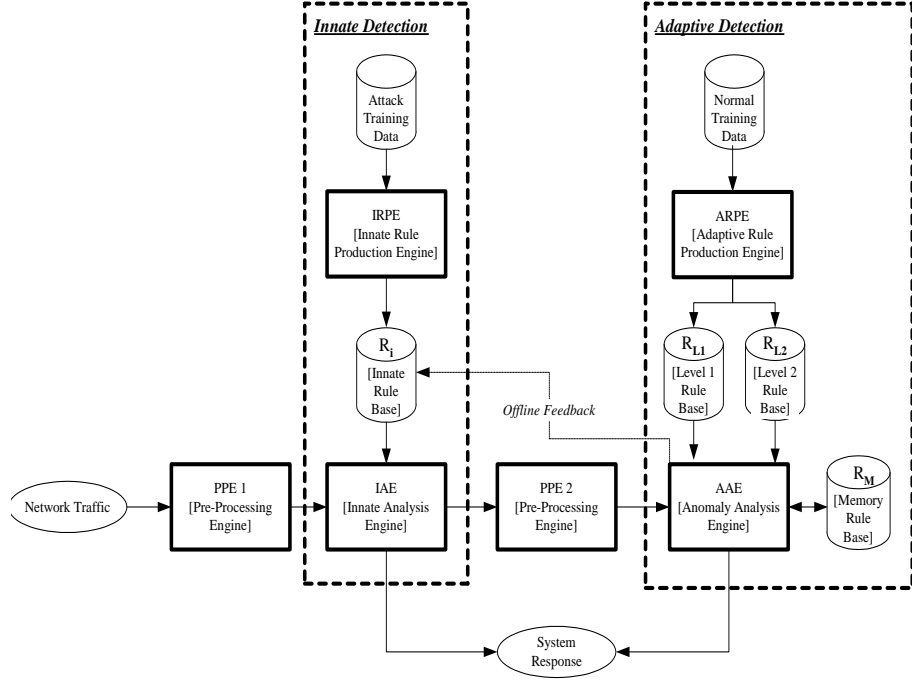


Fig. 1. Immune inspired intrusion detection framework.

4.1 Pre-processing engines (PPE1 and 2)

There are two pre-processing engines in the framework. Data to be analysed is processed into a format that is compatible with the analysis engines for each part of the framework - PPE1 for the Innate Analysis Engine and PPE2 for the Adaptive Analysis Engine. These formats will depend on how the rules have been generated by the rule production engines. Other standard pre-processing techniques, such as filtering and noise reduction, can also be included here.

4.2 Innate detection

This part of the framework implements standard misuse detection and is analogous to the innate immune system. The innate rule base (R_i) is used to detect known attacks using the Innate Analysis Engine. This could be an existing

database of known attack signatures, or could be generated using the Innate Rule Production Engine.

Innate rule production engine (IRPE) The Innate Rule Production Engine is used to generate rules that describe known attacks or intrusions. Any rule production system could be used here to generate the rules using data containing labelled attacks. The rules are stored in the innate rule base (R_i) and used by the analysis engine.

In a IDS that uses purely misuse detection, attention must be paid to the sensitivity of the rules or signatures that describe known attacks. They must be designed specific enough to identify an attack without detecting normal activity, while also being generic enough to identify a slightly modified attack. As the proposed framework combines a misuse detection approach with anomaly detection, the rules generated from the Innate Rule Production Engine can be weighted towards specificity in order to be certain of identifying the attacks that they are intended for.

Innate analysis engine (IAE) The Innate Analysis Engine is an implementation of misuse detection, which is a standard technique used in intrusion detection systems. The data to be analysed is matched against the rule base (R_i). If there is a match the system provides an alarm response. This technique can provide a high level of detection accuracy, and a low level of false positive responses, if the rules are tuned correctly.

If the system does not provide an alarm response at this point it could either be that the data being analysed is “normal”, or it could be an attack that does not have a similar or matching rule in the rule base. To determine which situation it is, if no alarm is raised the data is passed to the next level of the system for analysis - the Adaptive Detection.

4.3 Adaptive detection

This part of the framework is analogous to standard anomaly detection in intrusion detection systems, but also includes several immune system principles. It consists of a rule production engine (ARPE) and an analysis engine (AAE).

Adaptive rule production engine (ARPE) The adaptive rule production engine is inspired by the formation and maturation of lymphocytes in the immune system. There are two levels of rules generated - level 1 is analogous to the B cells which interact with antigen at the epitope level, and level 2 is analogous to the T cells which interact with antigen at the peptide level.

The details of the relationship between these levels and the data they relate to is left as an implementation issue. However it can be assumed that network traffic will be analysed (collected using a tool such as *TCPdump*²), with data

² TCPdump is a network tool that can be used to collect and display TCP/IP packets that are transmitted and received on a network. <http://www.tcpdump.org>

at level 1 relating to packet header information and data at level 2 relating to connection analysis statistics.

The algorithm used for generating each level of rule is described below.

- Level 1:
 1. Randomly generate a set of rules (R_{init}^1).
 2. Match these rules to normal data (N_{data}).
 3. Remove all rules that match normal above a selected affinity threshold (α_1) using an affinity function $\nu_1(r_i, d_j)$; where $r_i \in R_{init}^1$, and $d_j \in N_{data}$.
 4. Add the remaining rules to the level 1 rule base (R_{L1}).
 5. Repeat until there are a specified number of rules in the rule base (R_{L1}).

- Level 2:
 1. Randomly generate a set of rules (R_{init}^2).
 2. Match these rules to normal data (N_{data}).
 3. Remove all rules that do not match normal above a selected affinity threshold (β_1) using an affinity function $\nu_2(r_i, d_j)$; where $r_i \in R_{init}^2$, and $d_j \in N_{data}$.
 4. Of the rules remaining from step 3, remove all rules that match normal above a selected affinity threshold (β_2), where $\beta_2 \gg \beta_1$, using the affinity function ν_2 .
 5. Add the remaining rules to the level 2 rule base (R_{L2}).
 6. Repeat until there are a specified number of rules in the rule base (R_{L2}).

Adaptive analysis engine (AAE) The recognition and co-operation processes that occur in the adaptive immune system inspire this part of the framework (Figure 2). It is similar to the anomaly detection presented by Forrest et al. [2], except that it uses two levels of detectors and requires co-operation or co stimulation between the two. It also models the memory capabilities of the adaptive immune system. If an alarm is raised, the rules that were involved are inserted into the memory rule base (R_M). It is this memory rule base that is searched initially when the data is being analysed. This will mean that with subsequent occurrences of the same attack, the system will be quicker to respond.

The Adaptive Analysis Engine can also provide feedback to the innate rule base. If there is an attack that occurs multiple times the memory rule base will contain rules relating to the attack. These rules could be modified and added to the innate rule base so that in the future the system would be able to detect the attack at the first level of analysis.

The processes involved in this part of the framework are shown in Figure 2. Level 1 and level 2 data are provided to the analysis engine from the pre-processing engine (PPE2). Each level of data is matched with the rules in the memory rule base (R_M) and if there is a match on both levels an alarm response is raised. If there is no match at one or other level the data continues to be analysed. It is now matched with the rules in the level 1 (R_{L1}) and level 2 (R_{L2}) rule bases. If there is no match at one or other of these levels there is no system response. If there is a match at both levels the rules are added to the memory rule base and an alarm is raised.

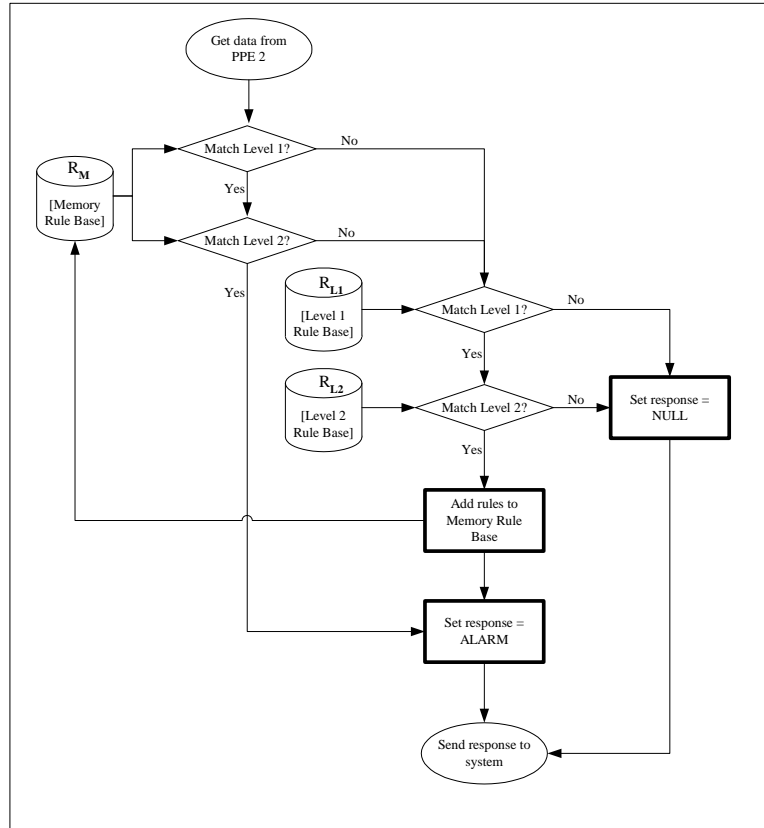


Fig. 2. Algorithm for the Adaptive Analysis Engine.

5 Conclusion

The goal of an effective intrusion detection system is to provide accurate and timely detection of intrusive activity. Problems have been identified with the performance of each of the main intrusion detection approaches. Misuse detection can suffer from high levels of false negative responses (intrusions going undetected) and anomaly detection is prone to high false positive responses (normal activity identified as intrusive). In order to achieve the goal of an IDS a balance between these levels must be reached. The system needs to be tolerant to normal network traffic or system activity, while also being able to detect unusual intrusive activity. This ability to successfully distinguish self from non-self is provided by the human immune system. This has led to the analogy with intrusion detection system and inspired the framework presented in this paper.

The framework presented here has combined a number of immunological principles. The multilevel defence of the immune system has been used to provide the

basis for the framework. Two levels of the immune system, innate immunity and adaptive immunity, have been described as analogous to misuse and anomaly detection respectively and used in the multilevel defence of the framework. Detectors are generated in a manner inspired by the maturation of B and T cells, and their cooperation in the humoral response of the adaptive immune system has been used to develop the algorithm for adaptive anomaly detection.

The focus of this research will now be on implementing the framework. An investigation also needs to be made into the representation of the detectors used in the framework. This will depend on the type of computer system the IDS is being designed to protect and the data that is derived from the system. Finally the performance of the framework should be measured by evaluating its ability to provide a high level of intrusion detection, while minimising the level of false negative detections.

References

1. Jackson, K.: Intrusion detection system product survey. Research report LA-UR-99-3883, Los Alamos National Laboratory (1999)
2. Forrest, S., Perelson, A.S., Allen, L., Cherukuri, R.: Self-nonsel self discrimination in a computer. In: IEEE Symposium on Research in Security and Privacy, IEEE Computer Society Press (1994) 202–212
3. Forrest, S., Hofmeyr, S., Somayaji, A.: Computer immunology. *Communications of the ACM* **40** (1997) 88–96
4. Kim, J., Bentley, P.: Negative selection and niching by an artificial immune system for network intrusion detection. In Brave, S., Wu, A.S., eds.: Late Breaking Papers at the 1999 Genetic and Evolutionary Computation Conference, Orlando, Florida, USA (1999) 149–158
5. Kephart, J.O.: A biologically inspired immune system for computers. In: *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, Cambridge, MA, US, MIT Press (1994) 130–139
6. Playfair, J.H.L.: *Infection and immunity*. Oxford University Press, Oxford ; New York (1995)
7. Dasgupta, D., Yu, S.: Artificial immune systems: A bibliography. CS Technical Report CS-03-002, Computer Science Division, The University of Memphis, USA (2003)
8. de Castro, L.N., Timmis, J.: *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer-Verlag (2002)
9. Dasgupta, D., Ji, Z., Gonzalez, F.: Artificial immune system (AIS) research in the last five years. In Sarker, R., Reynolds, R., Abbass, H., Tan, K.C., McKay, R., Essam, D., Gedeon, T., eds.: *Congress on Evolutionary Computation*, Canberra, Australia, IEEE Press (8-12 Dec 2003) 123–130
10. Bace, R., Mell, P.: Intrusion detection systems. Special Publication SP 800-31, Computer Security Division of the Information Technology Laboratory, National Institute of Standards and Technology (2001)
11. Ryan, J., Lin, M.J., Miikkulainen, R.: Intrusion detection with neural networks. In Jordan, M., Kearns, M., Solla, S., eds.: *Advances in Neural Information Processing Systems*. Volume 10. The MIT Press (1998)

12. Ranum, M.: Intrusion detection: Ideals, expectations and realities. *Computer Security Journal* **XV** (1999)
13. Sundaram, A.: An introduction to intrusion detection. *Crossroads, The ACM Student Magazine*, <http://www.acm.org/crossroads/xrds2-4/intrus.html> (1996)
14. Axelsson, S.: On a difficulty of intrusion detection. In: RAID'99: Second International Workshop on the Recent Advances in Intrusion Detection, Purdue University, West Lafayette, Indiana, USA (1999)
15. Janeway, C., Travers, P.: *Immunobiology : the immune system in health and disease*. 5th edn. *Current Biology ; Garland Pub., London ; San Francisco New York* (2001)
16. Sompayrac, L.: *How the immune system works*. 2nd edn. *Blackwell Science, Malden, Mass.* (2003)
17. Hannigan, B.: *Immunology. Biomedical Sciences Explained*. Arnold, London (2000)
18. Weir, D.M., Stewart, J.: *Immunology*. 8 edn. Churchill Livingstone, Medical Division of Longman Group Limited, New York (1997)
19. Dasgupta, D.: *Artificial Immune Systems and Their Applications*. Springer Verlag (1998)
20. Northcutt, S., Zeltser, L., Winters, S., Frederick, K.K., Ritchey, R.W.: *Inside Network Perimeter Security*. New Riders (2003)
21. The MITRE Corporation: Common vulnerabilities and exposures list, <http://cve.mitre.org/> (2002)
22. Kim, J., Ong, A., Overill, R.E.: Design of an artificial immune system as a novel anomaly detector for combating financial fraud in the retail sector. In Sarker, R., et. al., eds.: *Congress on Evolutionary Computation*, Canberra, Australia, IEEE Press (2003) 405–412
23. Decker, J.M.: *Introduction to Immunology*. 11th Hour Series. Blackwell Science (2000)
24. Dasgupta, D., Yu, S., Majumdar, N.S.: MILA - multilevel immune learning algorithm. In Cantu-Paz, E., et. al., eds.: *Genetic and Evolutionary Computation Conference*, Chicago, USA, Springer-Verlag (2003) 183–194
25. Campbell, N.A., Reece, J.B., Krebs, C.J.: *Biology*. 6th edn. Addison world student series. Benjamin Cummings, San Francisco [Calif.] (2002)