

Computation in the Wild*

Stephanie Forrest
Justin Balthrop
Matthew Glickman
David Ackley

Dept. of Computer Science
University of New Mexico
Albuquerque, NM 87131
forrest@cs.unm.edu

July 18, 2002

1 Introduction

The explosive growth of the Internet has created new opportunities and risks by increasing the number of contacts between separately administered computing resources. Widespread networking and mobility has blurred many traditional computer system distinctions, including those between operating system and application, network and computer, user and administrator, and program and data. An increasing number of executable codes, including applets, agents, viruses, email attachments, and downloadable software, are escaping the confines of their original systems and spreading through communications networks. These programs coexist and coevolve with us in our world, not always to good effect. Our computers are routinely disabled by network-borne

* To appear in K. Park and W. Willinger (Eds.), The Internet as a Large-Scale Complex System. Oxford University Press.

infections, our browsers crash due to unforeseen interactions between an applet and a language implementation, and applications are broken by operating system upgrades. We refer to this situation as *computation in the wild*, by which we mean to convey the fact that software is developed, distributed, stored, and executed in rich and dynamic environments populated by other programs and computers, which collectively form a *software ecosystem*. The thesis of this chapter is that networked computer systems can be better understood, controlled, and developed when viewed from the perspective of living systems.

Taking seriously the analogy between computer systems and living systems requires us to rethink several aspects of the computing infrastructure—developing design strategies from biology, constructing software that can survive in the wild, understanding the current software ecosystem, and recognizing that all nontrivial software must evolve. There are deep connections between computation and life, so much so that in some important ways, “living computer systems” are already around us, and moreover, such systems are spreading rapidly and will have major impact on our lives and society in the future.

In this paper we outline the biological principles we believe to be most relevant to understanding and designing the computational networks of the future. Among the principles of living systems we see as most important to the development of robust software systems are: Modularity, autonomy, redundancy, adaptability, distribution, diversity, and use of disposable components. These are not exhaustive, simply the ones that we have found most useful in our own research. We then describe a prototype network intrusion-detection system, known as LISYS, which illustrates many of these principles. Finally, we present experimental data on LISYS’ performance in a live network environment.

2 Computation in the Wild

In this section we highlight current challenges facing computer science and suggest that they have arisen because existing technological approaches are inadequate. Today's computers have significant and rapidly increasing commonalities with living systems, those commonalities have predictive power, and the computational principles underlying living systems offer a solid basis for secure, robust, and trustworthy operation in digital ecosystems.

2.1 Complexity, modularity, and linearity

Over the past fifty years, the manufactured computer has evolved into a highly complex machine. This complexity is managed by deterministic digital logic which performs extraordinarily complicated operations with high reliability, using modularity to decompose large elements into smaller components. Methods for decomposing functions and data dominate system design methods, ranging from object-oriented languages to parallel computing. An effective decomposition requires that the components have only limited interactions, which is to say that overall system complexity is reduced when the components are nearly independent [24]. Such well-modularized systems are “linear” in the sense that they obey an analog of the superposition principle in physics, which states that for affine differential equations, the sum of any two solutions is itself a solution (see [10] for a detailed statement of the connection). We use the term “linear” (and “nonlinear”) in this sense—linear systems are decomposable into independent modules with minimal interactions and nonlinear systems are not. In the traditional view, nearly independent components are composed to create or execute a program, to construct a model, or to solve a problem. With the emergence of large-scale networks and distributed computing, and the concomitant increase in aggregate com-

plexity, largely the same design principles have been applied. Although most single computers are designed to avoid component failures, nontrivial computer networks must also survive them, and that difference has important consequences for design, as a much greater burden of autonomy and self reliance is placed on individual components within a system.

The Achilles heel of any modularization lies in the interactions between the elements, where the independence of the components, and thus the linearity of the overall system, typically breaks down. Such interactions range from the use of public interfaces in object-oriented systems, to symbol tables in compilers, to synchronization methods in parallel processes. Although traditional computing design practice strives to minimize such interactions, those component interactions that are not eliminated are usually assumed to be deterministic, reliable, and trustworthy: A compiler assumes its symbol table is correct; an object in a traditional object-oriented system doesn't ask where a message came from; a closely-coupled parallel process simply waits, indefinitely, for the required response. Thus, even though a traditionally engineered system may possess a highly decoupled and beautifully linear design, at execution time its modules are critically dependent on each other and the overall computation is effectively a monolithic, highly nonlinear entity.

As a result, although software systems today are much larger, perform more functions, and execute in more complex environments compared to a decade ago, they also tend to be less reliable, less predictable, and less secure, because they critically depend on deterministic, reliable and trustworthy interactions while operating in increasingly complex, dynamic, error prone, and threatening environments.

2.2 Autonomy, living software, and reproductive success

Eliminating such critical dependencies is difficult. If, for example, a component is designed to receive some numbers and add them up, what else can it possibly do but wait until they arrive? A robust software component, however, could have multiple simultaneous goals, of which ‘performing an externally assigned task’ is only one. For example, while waiting for task input a component might perform garbage collection or other internal housekeeping activities. Designing components that routinely handle multiple, and even conflicting, goals leads to robustness.

This view represents a shift in emphasis about what constitutes a practical computation, away from traditional algorithms which have the goal of finding an answer quickly and stopping, and towards processes such as operating systems that are designed to run indefinitely. In client-server computing, for example, although client programs may start and stop, the server program is designed to run forever, handling requests on demand. And increasingly, many clients are designed for open-ended execution as well, as with web browsers that handle an indefinite stream of page display requests from a user. Peer-to-peer architectures such as the ccr system [1] or Napster [18], which eliminate or blur the client/server distinction, move even farther toward the view of computation as interaction among long-lived and relatively independent processes. Traditional algorithmic computations, such as sorting, commonly exist not as stand-alone terminating computations performed for a human user, but as tools used by higher-level nonterminating computational processes. This change, away from individual terminating computations and toward loosely-coupled ecological interactions among open-ended computations, has many consequences for architecture, software design, communication protocols, error recovery, and security.

Error handling provides a clear example of the tension between these two approaches. Viewed

algorithmically, once an error occurs there is no point in continuing, because any resulting output would be unreliable. From a living software perspective, however, process termination is the last response to an error, to be avoided at nearly any cost.

As in natural living systems, successful computational systems have a variety of lifestyles other than just the ‘run once and terminate’ algorithms on the one hand and would be immortal operating systems and processes on the other. The highly successful open-source Apache web server provides an example of such a strategy. Following the pattern of many Unix daemons, Apache employs a form of queen & workers organization, in which a single long-lived process does nothing except spawn and manage a set of server subprocesses, each of which handles a given number of web page requests and then dies. This hybrid strategy allows Apache to amortize the time required to start each new worker process over a productive lifetime serving many page requests, while at the same time ensuring that the colony as a whole can survive unexpected fatal errors in server subprocesses. Further, by including programmed subprocess death in the architecture, analogous to cellular apoptosis (programmed cell death), Apache deals with many nonfatal diseases as well, such as memory leaks. A related idea from the software fault tolerance community is that of software rejuvenation [21], in which running software is periodically stopped and restarted again after “cleaning” of the internal state. This proactive technique is designed to counteract software aging problems arising from memory allocation, fragmentation, or accumulated state.

In systems such as the Internet, which are large, open, and spontaneously evolving, individual components must be increasingly self-reliant and autonomous, able to function without depending on a consistent global design to guarantee safe interactions. Individual components must act autonomously, protect themselves, repair themselves, and increasingly, as in the example of Apache,

decide when to replicate and when to die. Natural biological systems exhibit organizational hierarchies (e.g., cells, multi-cellular organisms, ecosystems) similar to computers, but each individual component takes much more responsibility for its own interactions than what we hitherto have required of our computations.

The properties of autonomy, robustness, and security are often grouped together under the term “survivability,” the ability to keep functioning, at least to some degree, in the face of extreme systemic insult. To date, improving application robustness is largely the incremental task of growing the set of events the program is expecting—typically in response to failures as they occur. For the future of large-scale networked computation, we need to take a more basic lesson from natural living systems, that survival is more important than the conventional strict definitions of correctness. As in Apache’s strategy, there may be leaks, or even outright terminal programming errors, but they won’t cause the entire edifice to crash.

Moving beyond basic survival strategies, there is a longer term survivability achieved through the reproductive success of high-fitness individuals in a biological ecosystem. Successful individuals reproduce more frequently, passing on their genes to future generations. Similarly, in our software ecosystem, reproductive success is achieved when software is copied, and evolutionary dead-ends occur when a piece of software fails to be copied—as in components that are replaced (or “patched”) to correct errors or extend functionality.

2.3 Disposability, adaptation, and homeostasis

As computers become more autonomous, what happens when a computation makes a mistake, or is attacked in a manner that it cannot handle? We advocate building autonomous systems from

disposable components, analogous to cells of a body that are replicating and dying continually. This is the engineering goal of avoiding single points of failure, taken to an extreme. Of course, it is difficult to imagine that individual computers in a network represent disposable components; after all, nobody is happy when it's their computer that 'dies', for whatever reason, especially in the case of a false alarm. However, if the disposability is at a fine enough grain-size, an occasional inappropriate response is less likely to be lethal. Such lower-level disposability is at the heart of Apache's survival strategy, and it is much like building reliable network protocols on top of unreliable packet transfer, as in TCP/IP. We will also see an example of disposability in the following section, when we discuss the continual turnover of detectors in LISYS, known as *rolling coverage*.

Computation today takes place in highly dynamic environments. Nearly every aspect of a software system is likely to change during its normal life cycle, including: Who uses the system and for what purpose, the specification of what the system is supposed to do, certain parts of its implementation, the implementors of the system, the system software and libraries on which the computation depends, and the hardware on which the system is deployed. These changes are routine and continual.

Adaptation and homeostasis are two important components of a robust computing system for dynamic environments. Adaptation involves a component modifying its internal rules (often called "learning") to better exploit the current environment. Homeostasis, by contrast, involves a component dynamically adjusting its interactions with the environment to preserve a constant internal state. As a textbook says, homeostasis is "the maintenance of a relatively stable internal physiological environment or internal equilibrium in an organism" [5, p. G-11].

Given the inherent difficulty of predicting the behavior of even static nonlinear systems, and

adding the complication of continuously evolving computational environments, the once plausible notion of finding *a priori* proofs of program correctness is increasingly problematic. Adaptive methods, particularly those used in biological systems, appear to be the most promising near-term approach for modeling inherent nonlinearity and tracking environmental change. The phenomenon of “bit rot” is widespread. When a small piece of a program’s context changes, all too often either the user or a system administrator is required to change pathnames, apply patches, install new libraries (or reinstall old ones), etc. Emerging software packaging and installation systems such as RPM [22] or Debian’s `apt-get` utility, which provide common formats for ‘sporifying’ software systems to facilitate mobility and reproduction, and provide corresponding developmental pathways for adapting the software to new environments; we believe this sort of adaptability should take place not just at software installation but as an ongoing process of accommodation, occurring automatically and autonomously at many levels, and that programs should have the ability to evolve their functionality over time.

Turning to homeostasis, we see that current computers already have many mechanisms that can be thought of as homeostatic, for example, temperature and power regulation at the hardware level, and virtual memory management and process scheduling at the operating systems level. Although far from perfect, they are essential to the proper functioning of almost any application program. However, mechanisms such as virtual memory management and process scheduling cannot help a machine survive when it is truly stressed. This limitation stems in part from the policy that a kernel should be fair. When processor time is shared among all processes, and memory requests are immediately granted if they can be satisfied, stability cannot be maintained when resource demands become extreme. Fairness necessitates poor performance for all processes, either

through “thrashing,” when the virtual memory system becomes stressed, or through unresponsiveness, failed program invocations, and entire system crashes, when there are too many processes to be scheduled. One fair response to extreme resource contention, used by AIX [4], is random killing of processes. This strategy, however, is too likely to kill processes that are critical to the functioning of the system. Self-stabilizing algorithms [23] are similar in spirit, leading to systems in which an execution can begin in any arbitrary initial state and be guaranteed to converge to a “legitimate” state in a finite number of steps.

A more direct use of homeostasis in computing is Anil Somayaji’s pH system (for process homeostasis) [25, 26]. pH is a Linux kernel extension which monitors every executing process on a computer and responds to anomalies in process behavior, either by delaying or aborting subsequent system calls. Process-level monitoring at the system-call level can detect a wide range of attacks aimed at individual processes, especially server processes such as sendmail and apache [11]. Such attacks often leave a system functional, but give the intruder privileged access to the system and its protected data. In the human body, the immune system can kill thousands of cells without causing any outward sign of illness; similarly, pH uses execution delays for individual system calls as a way to restore balance. By having the delays be proportional to the number of recent anomalous system calls, a single process with gross abnormalities will effectively be killed (e.g., delayed to the point that native time-out mechanisms are triggered), without interfering with other normally behaving processes. Likewise, small deviations from normal behavior are tolerated in the sense that short delays are transient phenomena which are imperceptible at the user level. pH can successfully detect and stop many intrusions before the target system is compromised. In addition to coping with malicious attacks, pH detects a wide range of internal system-level problems, slowing down

the offending program to prevent damage to the rest of the system.

2.4 Redundancy, diversity, and the evolutionary mess

Biological systems use a wide variety of redundant mechanisms to improve reliability and robustness. For example, important cellular processes are governed by molecular cascades with large amounts of redundancy, such that if one pathway is disrupted an alternative one is available to preserve function. Similarly, redundant encodings occur throughout genetics, the triplet coding for amino acids at the nucleotide level providing one example. These kinds of redundancy involve more than the simple strategy of making identical replicas that we typically see in redundant computer systems. Identical replicas can protect against single component failure but not against design flaws. The variations among molecular pathways or genetic encodings thus provide additional protection through the use of diverse solutions.

This diversity is an important source of robustness in biology. As one example, a stable ecosystem contains many different species which occur in highly conserved frequency distributions. If this diversity is lost and a few species become dominant, the ecosystem becomes unstable and susceptible to perturbations such as catastrophic fires, infestations, and disease. Other examples include the variations among individual immune systems in a population and other forms of genetic diversity within a single species. Computers and other engineered artifacts, by contrast, are notable for their lack of diversity. Software and hardware components are replicated for several reasons: Economic leverage, consistency of behavior, portability, simplified debugging, and cost of distribution and maintenance. All these advantages of uniformity, however, become potential weaknesses when they replicate errors or can be exploited by an attacker. Once a method is created

for penetrating the security of one computer, all computers with the same configuration become similarly vulnerable [13]. Unfortunately, lack of diversity also pervades the software that is intended to defend against attacks, be it a firewall, an encryption algorithm, or a computer virus detector. The potential danger grows with the population of interconnected and homogeneous computers.

Turning to evolution, we see that the history of manufactured computers is a truly evolutionary history, and evolution does not anticipate, it reacts. To the degree that a system is large enough and distributed enough that there is no effective single point of control for the whole system, we must expect evolutionary forces—or ‘market forces’—to be significant. In the case of computing, this happens both at the technical level through unanticipated uses and interactions of components as technology develops and at the social level from cultural and economic pressures. Having humans in the loop of an evolutionary process, with all their marvelous cognitive and predictive abilities, with all their philosophical ability to frame intentions, does not necessarily change the nature of the evolutionary process. There is much to be gained by recognizing and accepting that our computational systems resemble naturally evolving systems much more closely than they resemble engineered artifacts such as bridges or buildings. Specifically, the strategies that we adopt to understand, control, interact with, and influence the design of computational systems will be different once we understand them as ongoing evolutionary processes.

In this section we described several biological design principles and how they are relevant to computer and software systems. These include modularity, redundancy, distribution, autonomy, adaptability, diversity, and use of disposable components. As computers have become more complex and interconnected, these principles have become more relevant, due to an increasing variety

of interactions between software components and a rising number of degrees of freedom for variation in computational environments. Software components are now confronted with challenges increasingly similar to those faced by organic entities situated in a biological ecosystem. An example of how biological organisms have evolved to cope with their environments is the immune system. In the next section we explore a software framework that is specifically patterned after the immune system. In addition to exhibiting many of the general biological design principles advocated above, this framework illustrates a specific set of mechanisms derived from a particular biological system and applied to a real computational problem.

3 An Illustration: LISYS

In the previous section, we outlined an approach to building and designing computer systems that is quite different from that used today. In this section we illustrate how some of these ideas could be implemented using an artificial immune system framework known as ARTIS and its application to the problem of network intrusion detection in a system known as LISYS. ARTIS and LISYS were originally developed by Steven Hofmeyr in his dissertation [16, 15].

The immune system processes peptide patterns using mechanisms that in some cases correspond closely to existing algorithms for processing information (e.g., the genetic algorithm), and it is capable of exquisitely selective and well-coordinated responses, in some cases responding to fewer than ten molecules. Some of the techniques used by the immune system include learning (affinity maturation of B-cells, negative selection of B- and T-cells, and evolved biases in the germline), memory (cross-reactivity and the secondary response), massively parallel and distributed computations with highly dynamic components (on the order of 10^8 different varieties of

receptors [27] and 10^7 new lymphocytes produced each day [19]), and the use of combinatorics to address the problem of scarce genetic resources (V-region libraries). Not all of these features are included in ARTIS, although they could be (affinity maturation and V-region libraries are the most notable lacunae).

ARTIS is intended to be a somewhat general artificial immune system architecture, which can be applied to multiple application domains. In the interest of brevity and concreteness, we will describe the instantiation of ARTIS in LISYS, focusing primarily on how it illustrates our ideas about computation in the wild. For more details about LISYS as a network intrusion-detection system, the reader is referred to [16, 3, 2].

3.1 The Network Intrusion Problem

LISYS is situated in a local-area broadcast network (LAN) and used to protect the LAN from network-based attacks. In contrast with switched networks, broadcast LANs have the convenient property that every location (computer) sees every packet passing through the LAN.¹ In this domain, we are interested in building a model of normal behavior (known as *self* through analogy to the normally occurring peptides in living animals) and detecting deviations from normal (known as *nonself*). Self is defined to be the set of normal pairwise TCP/IP connections between computers, and non-self is the set of connections, potentially an enormous number, which are not normally observed on the LAN. Such connections may represent either illicit attacks or significant changes in network configuration or use. A connection can occur between any two computers in the LAN

¹There are several ways in which the architecture could be trivially modified to run in switched networks. For example, some processing could be performed on the switch or router, SYN packets could be distributed from the switch to the individual nodes, or a combination could be tried.

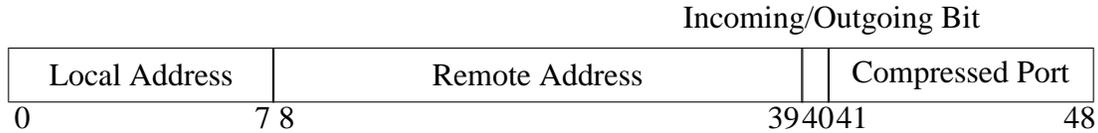


Figure 1: The 49-bit compression scheme used by LISYS to represent TCP SYN packets.

as well as between a computer in the LAN and an external computer. It is defined in terms of its “data-path triple”—the source IP address, the destination IP address, and the port by which the computers communicate [17, 14].

LISYS examines only the headers of IP packets. Moreover, at this point we restrict LISYS to examining only a restricted set of features contained in the headers of TCP SYN packets. The connection information is compressed to a single 49-bit string that unambiguously defines the connection. This string is compressed in two ways [16]. First, it is assumed that one of the IP addresses is always internal, so only the final byte of this address needs to be stored. Secondly, the port number is also compressed from 16 bits to 8 bits. This is done by re-mapping the ports into several different classes. Figure 1 shows the 49-bit representation.

3.2 LISYS Architecture

LISYS consists of a distributed set of *detectors*, where each detector is a 49-bit string (with the above interpretation) and a small amount of local state. The detectors are distributed in the sense that they reside on multiple hosts; there exists one *detector set* for each computer in the network. Detector sets perform their function independently, with virtually no communication. Distributing detectors makes the system more robust by eliminating the single point of failure associated with centralized monitoring systems. It also makes the system more scalable as the size of the network being protected increases.

A perfect *match* between a detector and a compressed SYN packet means that at each location in the 49-bit string, the symbols are identical. However, in the immune system matching is implemented as binding affinities between molecules, where there are only stronger and weaker affinities, and the concept of perfect matching is ill defined. To capture this LISYS uses a partial matching rule known as r -contiguous bits matching [20]. Under this rule, two strings match if they are identical in at least r contiguous locations. This means there are $l - r + 1$ windows where a match can occur, where l is the string length. The value r is a threshold which determines the specificity of detectors, in that it controls how many strings can potentially be matched by a single detector. For example, if $r = l$ (49 in the case of LISYS' compressed representation), the match is maximally specific, and a detector can match only a single string—itsself. As shown in [9], the number of strings a detector matches increases exponentially as the value of r decreases.

LISYS uses *negative detection* in the sense that valid detectors are those that fail to match the normally occurring behavior patterns in the network. Detectors are generated randomly and are initially *immature*. Detectors that match connections observed in the network during the *tolerization period* are eliminated. This procedure is known as *negative selection* [12]. The tolerization period lasts for a few days, and detectors that survive this period become *mature*. For the r -contiguous bits matching rule and fixed self sets which don't change over time, the random detector generation algorithm is inefficient—the number of random strings that must be generated and tested is approximately exponential in the size of self. More efficient algorithms based on dynamic programming methods allow us to generate detectors in linear time [6, 7, 29, 28, 30]. However, when generating detectors asynchronously for a dynamic self set, such as the current setting, we have found that random generation works sufficiently well. Negative detection allows LISYS to be distributed,

because detectors can make local decisions independently with no communication to other nodes. Negative selection of randomly generated detectors allows LISYS to be adaptive to the current network condition and configuration.

LISYS also uses *activation thresholds*. Each detector must match multiple connections before it is activated. Each detector records the number of times it matches (the *match count*) and raises an alarm only when the match count exceeds the activation threshold. Once a detector has raised an alarm, it returns its match count to zero. This mechanism has a time horizon: Over time the match count slowly returns to zero. Thus, only repeated occurrences of structurally similar and temporally clumped strings will trigger the detection system. The activation threshold mechanism contributes to LISYS's adaptability and autonomy, because it provides a way for LISYS to tolerate small errors that are likely to be false positives.

LISYS uses a "second signal," analogous to costimulatory mechanisms in the immune system. Once a detector is activated, it must be costimulated by a human operator or it will die after a certain time period. Detectors which are costimulated become *memory detectors*. These are long-lived detectors which are more easily activated than non-memory detectors. This secondary response improves detection of true positives, while costimulation helps control false positives.

Detectors in LISYS have a finite lifetime. Detectors can die in several ways. As mentioned before, immature detectors die when they match self, and activated detectors die if they do not receive a costimulatory signal. In addition, detectors have a fixed probability of dying randomly on each time step, with the average lifespan of a single detector being roughly one week. Memory detectors are not subject to random death and may thus survive indefinitely. However, once the number of memory detectors exceeds a specified fraction of the total detector population, some are

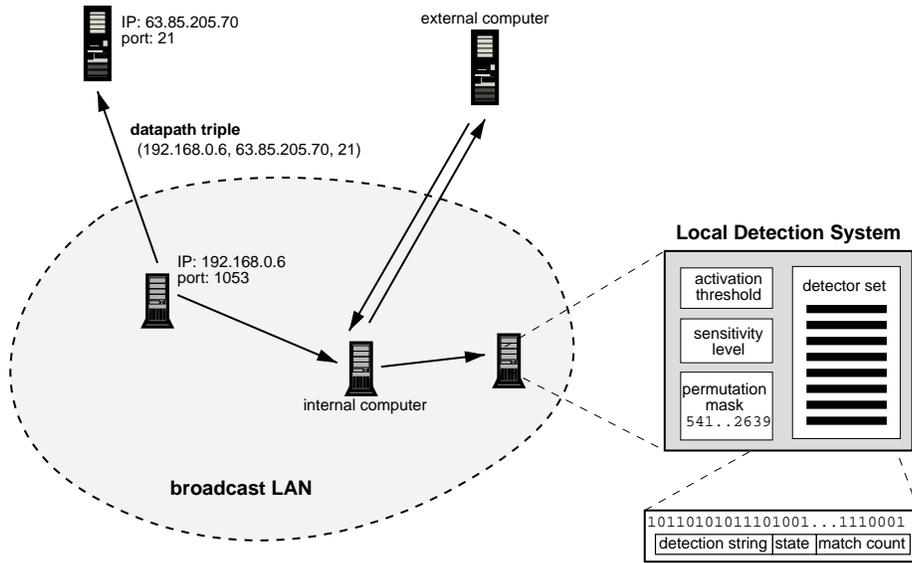


Figure 2: The architecture of LISYS.

removed to make room for new ones. The finite lifetime of detectors, when combined with detector regeneration and tolerization, results in *rolling coverage* of the self set. This rolling coverage clearly illustrates the principles of disposable components and adaptability.

Each independent detector set (one per host) has a *sensitivity level*, which modifies the local activation threshold. Whenever the match count of a detector in a given set goes from 0 to 1, the effective activation threshold for all the other detectors in the same set is reduced by one. Hence, each different detector that matches for the first time “sensitizes” the detection system, so that all detectors on that machine are more easily activated in future. This mechanism has a time horizon as well; over time, the effective activation threshold gradually returns to its default value. This mechanism corresponds very roughly to the role that inflammation, cytokines, and other molecules play in increasing or decreasing the sensitivity of individual immune system lymphocytes within a physically local region. Sensitivity levels help the system detect true positives, especially dis-

tributed coordinated attacks, and it is another simple example of an adaptive mechanism.

LISYS uses *permutation masks* to achieve diversity of representation². A permutation mask defines a permutation of the bits in the string representation of the network packets. Each detector set (network host) has a different, randomly generated, permutation mask. One feature of the negative-selection algorithm as originally implemented is that it can result in undetectable patterns called *holes*, or put more positively *generalizations* [6, 7]. Holes can exist for any symmetric, fixed-probability matching rule, but by using permutation masks we effectively change the match rule on each host, which gives us the ability to control the form and extent of generalization in the vicinity of self [8]. Thus, the permutation mask controls how the network packet is presented to the detection system, which is analogous to the way different MHC types present different sets of peptides on the cell surface.

Although LISYS incorporates many ideas from natural immune systems, it also leaves out many features. One important omission is that LISYS has only a single kind of detector “cell,” which represents an amalgamation of T cells, B cells, and antibodies. In natural immune systems, these cells and molecules have distinct functions. Other cell types missing from LISYS include effector cells, such as natural killer cells and cytotoxic T cells. The mechanisms modeled in LISYS are almost entirely part of what is known as “adaptive immunity.” Also important is innate immunity. Moreover, an important aspect of adaptive immunity is clonal selection and somatic hypermutation, processes which are absent from LISYS. In the future, it will be interesting to see which of these additional features turn out to have useful computational analogs.

²Permutation masks are one possible means of generating *secondary representations*. A variety of alternative schemes are explored in Hofmeyr [16].

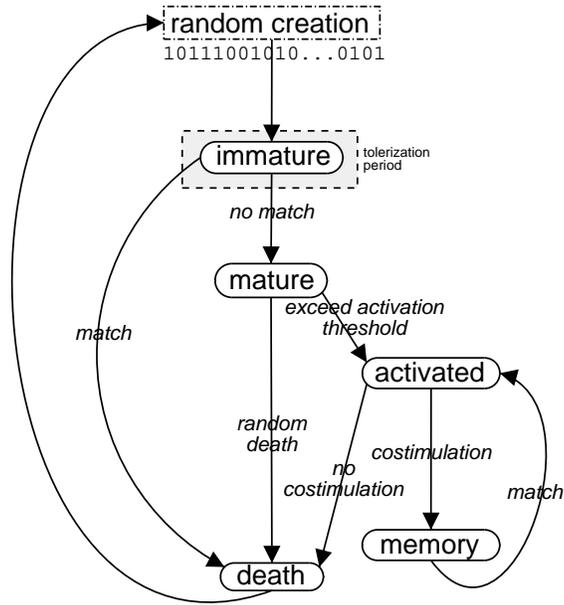


Figure 3: The lifecycle of a detector in LISYS.

4 Experiments

In this section we summarize some experiments that were performed in order to study LISYS's performance in a network intrusion-detection setting and to explore how LISYS' various mechanisms contribute to its effectiveness. For the experiments described in this section, we used a simplified form of LISYS in order to study some features more carefully. Specifically, we used a version of LISYS that does not have sensitivity levels, memory detectors, or costimulation by a human operator. Although we collected the data on-line in a production distributed environment, we performed our analysis off-line on a single computer. This made it possible to compare performance across many different parameter values. The programs used to generate the results reported in this paper are available from <http://www.cs.unm.edu/~immsec>. The programs are part of LISYS and are found in the LisysSim directory of that package.

4.1 Data Set

Our data were collected on a small but realistic network. The normal data were collected for two weeks in Nov, 2001 on an an internal restricted network of computers in our research group at UNM. The six internal computers in this network connected to the Internet through a single Linux machine acting as a firewall, router and masquerading server. In this environment we were able understand all of the connections, and we could limit attacks. Although small, the network provides a plausible model of a corporate intranet where activity is somewhat restricted and external connections must pass through a firewall. And, it resembles the increasingly common home network that connects to the Internet through a cable or DSL modem and has a single external IP address. After collecting the normal data set, we used a package known as “Nessus” to generate attacks against the network.

Three groups of attacks were performed. The first attack group included a denial-of-service (DOS) attack from an internal computer to an external computer, attempted exploitations of weaknesses in the configuration of the firewall machine, an attack against FTP (File Transfer Protocol), probes for vulnerabilities in SSH (Secure Shell), and probes for services such as chargen and telnet. The second attack group consisted of two TCP port scans, including a stealth scan (difficult to detect) and a noisy scan (easy to detect). The third attack group consisted of a full nmap³ port scan against several internal machines.

Most of these attacks are technically classified as probes because they did not actually execute an attack, simply checking to see if the attack could be performed. However, in order to succeed, a probe would typically occur first; actually executing the attack would create additional TCP traffic,

³Nmap is a separate software package used by Nessus.

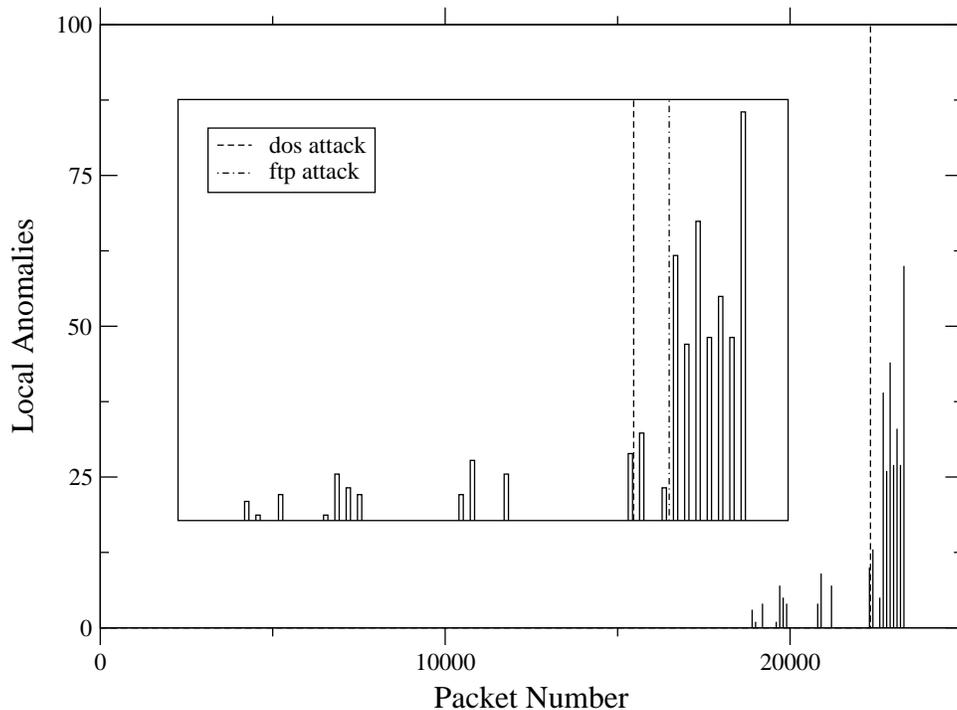


Figure 4: Local anomalies for attack group one. The inset is a zoomed view of the same data.

making it easier for LISYS to detect the attack. More details about this data set are given in [3].

4.2 Detecting Abnormal Behavior

Here we report how a minimal version of LISYS performs. In addition to the simplifications mentioned earlier, we also eliminated the use of permutation masks for this set of experiments and used a single large detector set running on a single host. Because we are operating in a broadcast network environment and not using permutation masks, this configuration is almost equivalent to distributing the detectors across different nodes (there could be some minor differences arising from the tolerization scheme and parameter settings for different detector sets).

We used the following parameter values [3]: $r = 10$, number of detectors = 6,000,

tolerization period = 15,000 connections, and activation threshold = 10. The experiments were performed by running LISYS with the normal network data followed by the attack data from one of the attack groups.

Figure 4 shows how LISYS performed during the first group of attacks. The x-axis shows time (in units of packets) and the y-axis shows the number of anomalies per time period. An anomaly occurs whenever the match count exceeds the activation threshold. The vertical line indicates where the normal data ended and the attack data began. Anomalies are displayed using windows of 100. This means that each bar is an indication of the number of anomalies signalled in the last 100 packets. There are a few anomalies flagged in the normal data, but there are more anomalies during the group of attacks.

The graph inset shows that LISYS was able to detect the denial-of-service and FTP attacks. The vertical lines indicate the beginning of these two attacks, and there are spikes shortly after these attacks began. The spikes for the FTP attack are significantly higher than those for the DOS attack, but both attacks have spikes that are higher than the spikes in the normal data, indicating a clear separation between true and false positives. This view of the data is interesting because the height of the spikes indicates the system's confidence that there is an anomaly occurring at that point in time.

Figure 5 shows the LISYS anomaly data for attack group two. By looking at this figure, we can see that there is something anomalous in the last half of the attack data, but LISYS was unable to detect anomalies during the first half of the attack.⁴ Although the spikes are roughly the same

⁴In LISYS, detectors are continually being generated, undergoing negative selection, and being added to the repertoire. Some new detectors were added during the first half of the attack, but these detectors turned out not to be crucial for the detection in the second half.

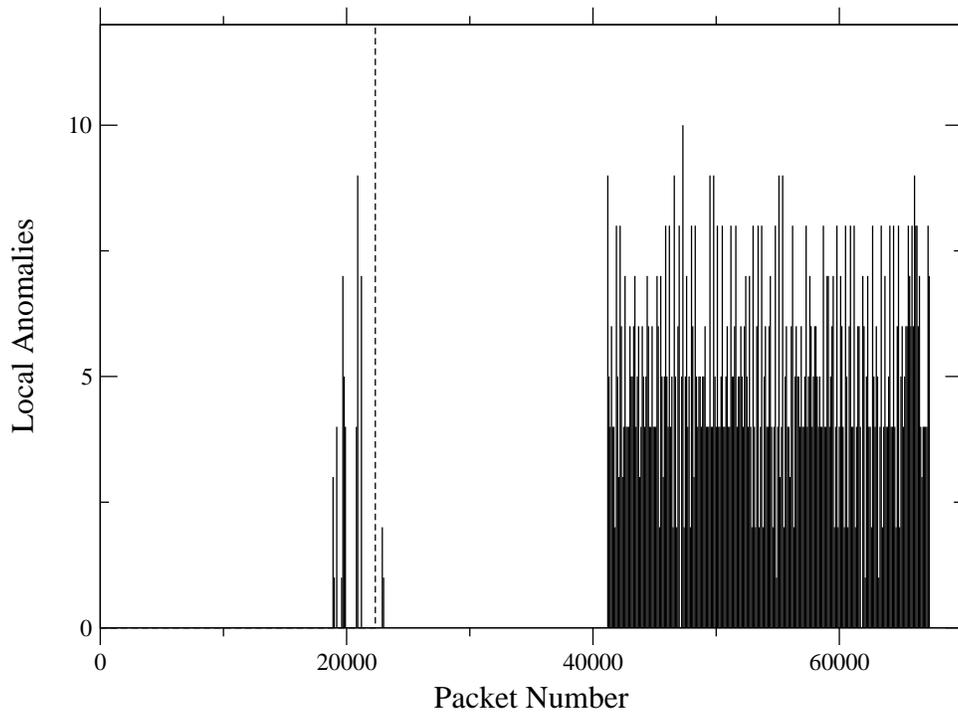


Figure 5: Local anomalies for attack group two.

height as the spikes in the normal data, the temporal clustering of the spikes is markedly different. Figure 6 shows the LISYS anomaly data for attack group three. The figure shows that LISYS overwhelmingly found the nmap scan to be anomalous. Not only are the majority of the spikes significantly higher than the normal data spikes, but there is a huge number of temporally clustered spikes.

These experiments support the results reported in [16], suggesting that the LISYS architecture is effective at detecting certain classes of network intrusions. However, as we will see in the next section, LISYS performs much better under slightly different circumstances.

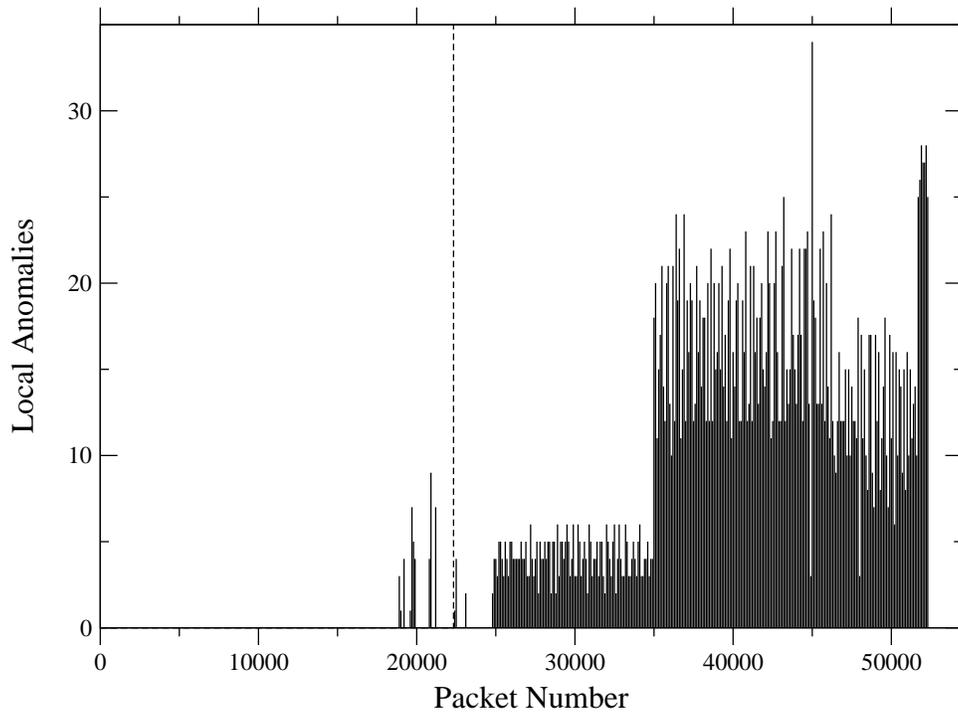


Figure 6: Local anomalies for attack group three.

5 The Effect of Diversity

The goal of the experiments in this section was to assess the effect of diversity in our artificial immune system. Recall that diversity of representation (loosely analogous to MHC diversity in the real immune system) is implemented in LISYS by permutation masks. We were interested to see how LISYS' performance would be affected by adding permutation masks. Because performance is measured in terms of true and false positives, this experiment tested the effect of permutations on the system's ability to generalize (because low false-positive rates correspond to good generalization, and high false positives correspond to poor generalization). For this experiment, 100 sets of detectors were tolerized (generated randomly and compared against self using negative selection) using the first 15,000 packets in the data set (known as the *training set*), and each detector set was

assigned a random permutation mask. Each detector set had exactly 5,000 mature detectors at the end of the tolerization period [3], and for consistency we set the random death probability to zero. 5,000 detectors provides maximal possible coverage (i.e. adding more detectors does not improve subsequent matching) for this data set and r threshold. Each set of detectors was then run against the remaining 7,329 normal packets, as well as against the simulated attack data. In these data (the *test sets*), there are a total of 475 unique 49-bit strings. Of these 475, 53 also occur in the training set and are thus undetectable (because any detectors which would match them are eliminated during negative selection). This leaves 422 potentially detectable strings, of which 26 come from the normal set and 396 are from the attack data, making the maximal possible coverage by a detector set 422 unique matches.

An ideal detector set would achieve zero false positives on the normal test data and a high number of true positives on the abnormal data. Because network attacks rarely produce a single anomalous SYN packet, we don't need to achieve perfect true-positive rates at the packet level in order to detect all attacks against the system. For convenience, however, we measure false positives in terms of single packets. Thus, a perfect detector set would match the 396 unique attack strings, and fail to match the 26 new normal strings in the test set.

Figure 7 shows the results of this experiment. The performance of each detector set is shown as a separate point on the graph. Each detector set has its own randomly generated permutation of the 49 bits, so each point shows the performance of a different permutation. The numbers on the x-axis correspond to the number of unique self-strings in the test set which are matched by the detector set, i.e. the number of false positives (up to a maximum of 26). The y-axis plots a similar value with respect to the attack data, i.e. the number of unique true positive matches (up to

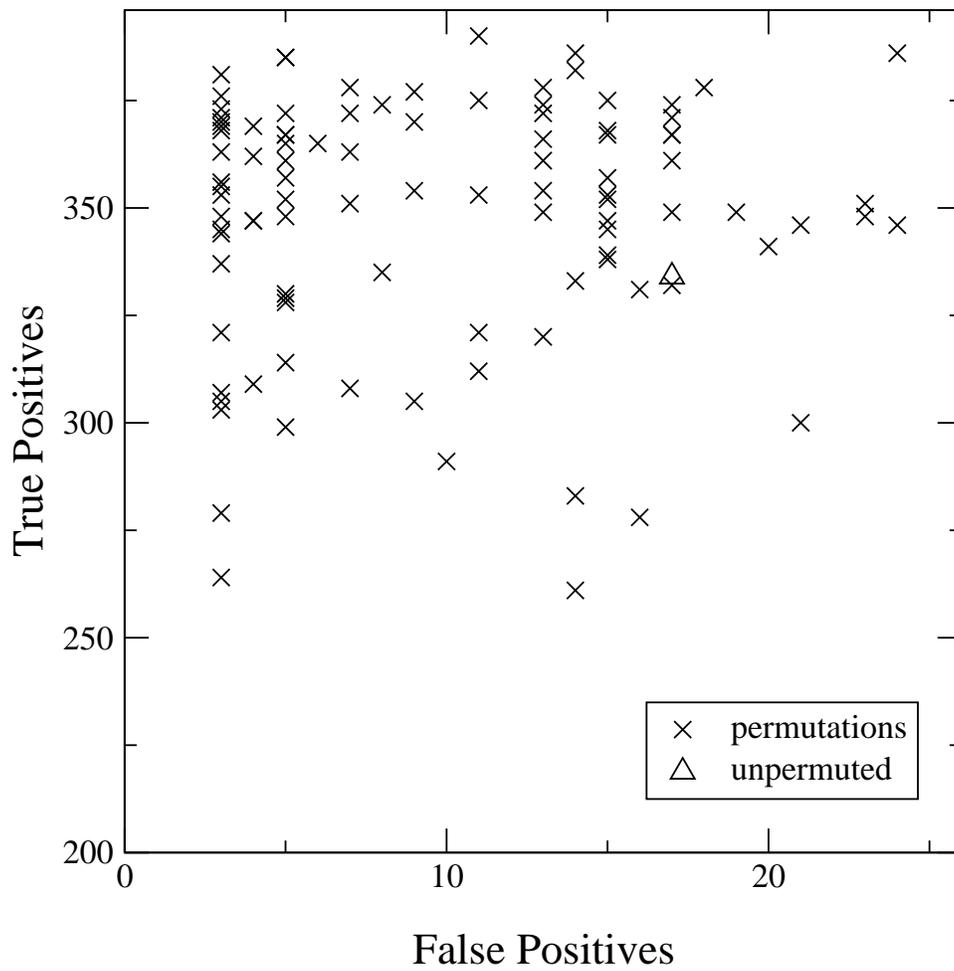


Figure 7: LISYS performance under different permutations. Each plotted point corresponds to a different permutation, showing false positives (x-axis) and true positives (y-axis).

a maximum of 396). The graph shows that there is a large difference in the discrimination ability of different permutations. Points in the upper left of the graph are the most desirable, because they correspond to permutations which minimize the number of false positives and maximize the number of true positives; points toward the lower right corner of the graph indicate higher false positives and/or lower true positives. Surprisingly, the performance of the original (unpermuted) mapping is among the worst we found, suggesting that the results reported in the previous section are a worst case in terms of true vs. false positives. Almost any other random permutation we tried outperformed the original mapping. Although we don't yet have a complete explanation for this behavior, we believe that it arises in the following way.

The LISYS design implicitly assumes that there are certain predictive bit patterns that exhibit regularity in self, and that these can be the basis of distinguishing self from non-self. As it turns out, there are also deceptive bit patterns which exhibit regularity in the training set (observed self), but the regularity does not generalize to the rest of self (the normal part of the test set). These patterns tend to cause false positives when self strings that do not fit the predicted regularity occur. We believe that the identity permutation is bad because the predictive bits are at the ends of the string, while the deceptive region is in the middle. Under such an arrangement, it is difficult to find a window that covers many predictive bit positions without also including deceptive ones. It is highly likely that a random permutation will break up the deceptive region, and bring the predictive bits closer to the middle, where they will appear in more windows.

The preceding analysis is based on the choice of a single permutation. In the original system developed by Hofmeyr, however, each host on the network used a different (randomly chosen) permutation of the 49 bits. Using a diverse set of random permutation masks reduces the overall

number of undetectable patterns, thus increasing population-level robustness, but there is a risk of increasing false positives. In the experiments described here, using all permutations (with each permutation consisting of 5000 detectors) would raise the number of true positives to 396 out of 396, and it would also raise the number of false positives to 26 out of 26 (compared with 334/396 true positives and 17/26 false positives when the original representation is used). However, if we chose only those permutations with low false-positive rates, we can do significantly better. The relative tradeoffs between diversity of representation and the impact on discrimination ability is an important area of future investigation.

6 Discussion

In the previous section we described a prototype network intrusion-detection system based on the architecture and mechanisms of natural immune systems. Although LISYS is not a complete model of the immune system nor a production quality intrusion-detection system, it illustrates many of the principles elucidated in Section 2. We identified the principles of redundancy, distribution, autonomy, adaptability, diversity, use of disposable components, and modularity as key features required for successful “computation in the wild.”

Of these, LISYS is weakest in the autonomy component. There are a number of ways in which LISYS could be made more autonomous. One way would be to model the so-called constant region of antibody molecules. Detectors in LISYS represent generic immune cells, combining properties of T cells, B cells, and antibodies. However, as described to date they model only the recognition capabilities of immune cells (receptors) and a small amount of state. A simple extension would be to concatenate a few bits to each detector to specify a response (analogous to different antibody

isotypes which comprise an antibody’s “constant region”). This feature would constitute a natural extension to LISYS, although it leaves open the important question of how to interpret the response bits—that is, what responses do we want LISYS to make?

A second approach to adding autonomy to LISYS involves the use of a second signal. Costimulation in LISYS is an example of a second signal, in which a second independent source must confirm an alarm before a response can be initiated. There are several examples of second signals in immunology, including helper T-cells (which confirm a B-cell’s match before the B-cell can be activated) and the complement system used in the early immune response. The complement cascade provides a general and early indication of trouble, and operates in combination with highly specific T-cell responses. One possibility for LISYS would be to incorporate the pH system described earlier to provide the second signal. pH can detect anomalous behavior on a host (e.g., during intrusion attempts) and could serve as a generic indicator of “damage” in conjunction with the specific detectors of LISYS. Thus, when a LISYS detector became activated it would respond only if pH confirmed that the system was in an anomalous state. Either of these two approaches would represent an important next step towards autonomy.

7 Conclusions

The implementation we described is based on similarities between computation and immunology. Yet, there are also major differences, which it is wise to respect. The success of all analogies between computing and living systems ultimately rests on our ability to identify the correct level of abstraction—preserving what is essential from an information-processing perspective and discarding what is not. In the case of immunology, this task is complicated by the fact that real

immune systems handle data that are very different from that handled by computers. In principle, a computer vision system or a speech-recognition system would take as input the same data as a human does (e.g., photons or sound waves). In contrast, regardless of how successful we are at constructing a computer immune system, we would never expect or want it to handle pathogens in the form of actual living cells, viruses, or parasites. Thus, the level of abstraction for computational immunology is necessarily higher than that for computer vision or speech, and there are more degrees of freedom in selecting a modeling strategy.

Our study of how the architecture of the immune system could be applied to network security problems illustrates many of the important design principles introduced in the first part of this chapter. Many of the principles discussed here are familiar, and in some cases have long been recognized as desirable for computing systems, but in our view, there has been little appreciation of the common underpinnings of these principles, little or no rigorous theory, and few working implementations that take them seriously. Hofmeyr's artificial immune system is, however, an initial step in this direction. As our computers and the software they run become more complex and interconnected, properties such as robustness, flexibility and adaptability, diversity, self reliance and autonomy, and scalability can only become more important to computer design.

8 Acknowledgments

The authors gratefully acknowledge the support of the National Science Foundation (grants CDA-9503064, and ANIR-9986555), the Office of Naval Research (grant N00014-99-1-0417), Defense Advanced Projects Agency (grant AGR F30602-00-2-0584), the Santa Fe Institute, and the Intel Corporation.

We thank the many members of Adaptive Computation group at the University of New Mexico for their help over the past ten years in developing these ideas.

References

- [1] D. H. Ackley. ccr: A network of worlds for research. In C.G. Langton & K. Shimohara, editor, *Artificial Life V. (Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems)*, pages 116–123, Cambridge, MA, 1996. The MIT Press.
- [2] J. Balthrop, F. Esponda, S. Forrest, and M. Glickman. Coverage and generalization in an artificial immune system. In *GECCO-2002: Proceedings of the Genetic and Evolutionary Computation Conference*, 2002.
- [3] J. Balthrop, S. Forrest, and M. Glickman. Revisiting lisy: Parameters and normal behavior. In *CEC-2002: Proceedings of the Congress on Evolutionary Computing*, 2002.
- [4] International Business Machines Corporation. *AIX Version 4.3 System Management Guide: Operating System and Devices*, chapter Understanding Paging Space Allocation Policies. IBM, 1997. http://www.rs6000.ibm.com/doc_link/en_US/a_doc.lib/aixbman/baseadm/pag_space_under.htm.
- [5] Helena Curtis and N. Sue Barnes. *Biology*. Worth Publishers, Inc., New York, 5th edition, 1989.
- [6] P. D’haeseleer, S. Forrest, and P. Helman. An immunological approach to change detection: algorithms, analysis and implications. In *Proceedings of the 1996 IEEE Symposium on Computer Security and Privacy*. IEEE Press, 1996.
- [7] Patrik D’haeseleer. An immunological approach to change detection: theoretical results. In *Proceedings of the 9th IEEE Computer Security Foundations Workshop*. IEEE Computer Society Press, 1996.
- [8] Carlos Fernando Esponda and Stephanie Forrest. Defining self: Positive and negative detection. In *Proceedings of the IEEE Symposium on Security and Privacy*, Submitted Nov. 2001.
- [9] Fernando Esponda. Detector coverage with the r-contiguous bits matching rule. Technical Report TR-CS-2002-03, University of New Mexico, 2002.
- [10] S. Forrest. *Emergent Computation*. MIT Press, Cambridge, MA, (1991).
- [11] S. Forrest, S. Hofmeyr, A. Somayaji, and T. Longstaff. A sense of self for unix processes. In *Proceedings of the 1996 IEEE Symposium on Computer Security and Privacy*. IEEE Press, 1996.

- [12] S. Forrest, A. S. Perelson, L. Allen, and R. Cheru kuri. Self-nonsel self discrimination in a computer. In *Proceedings of the 1994 IEEE Symposium on Research in Security and Privacy*, Los Alamitos, CA, 1994. IEEE Computer Society Press.
- [13] S. Forrest, A. Somayaji, and D. H. Ackley. Building diverse computer systems. In *Sixth Workshop on Hot Topics in Operating Systems*, 1998.
- [14] L. T. Heberlein, G. V. Dias, K. N. Levitte, B. Mukherjee, J. Wood, and D. Wolber. A network security monitor. In *Proceedings of the IEEE Symposium on Security and Privacy*. IEE Press, 1990.
- [15] S. A. Hofmeyr and S. Forrest. Architecture for an artificial immune system. *Evolutionary Computation Journal*, 8(4):443–473, 2000.
- [16] Steven A. Hofmeyr. *An immunological model of distributed detection and its application to computer security*. PhD thesis, University of New Mexico, Albuquerque, NM, 1999.
- [17] B. Mukherjee, L. T. Heberlein, and K. N. Levitt. Network intrusion detection. *IEEE Network*, pages 26–41, 1994.
- [18] Napster. <http://www.napster.org/>, March 2002.
- [19] D. G. Osmond. The turn-over of B-cell populations. *Immunology Today*, 14(1):34–37, 1993.
- [20] J. K. Percus, O. Percus, and A. S. Perelson. Predicting the size of the antibody combining region from consideration of efficient self/non-self discrimination. *Proceedings of the National Academy of Science*, 90:1691–1695, 1993.
- [21] A Pfenning, S. Garg, A. Puliafito, M. Telek, and K. Trivedi. Optimal software rejuvenation for tolerating software failures. *Performance Evaluation*, 1996.
- [22] Rpm documentation. <http://www.rpm.org/>, March 2002.
- [23] M. Schneider. Self-stabilization. *ACM Computing Surveys*, 25(1):45–67, 1993.
- [24] Herbert A. Simon. *The Sciences of the Artificial*. M.I.T. Press, Boston, 1969.
- [25] A. Somayaji and S. Forrest. Automated response using system-call delays. In *Usenix Security Symposium*, 2000.
- [26] Anil Somayaji. *Operating System Stability and Security through Process Homeostasis*. PhD thesis, University of New Mexico, Albuquerque, NM, 2002.
- [27] S. Tonegawa. Somatic generation of antibody diversity. *Nature*, 302:575–581, 1983.
- [28] S. T. Wierzhon. Discriminative power of the receptors activated by k-contiguous bits rule. *Journal of Computer Science and Technology*, 1(3):1–13, 2000.

- [29] S. T. Wierzchon. Generating optimal repertoire of antibody strings in an artificial immune system. In M. A. Kłopotek, M. Michalewicz, and S. T. Wierzchon, editors, *Intelligent Information Systems*, pages 119–133, Heidelberg New York, 2000. Physica-Verlag.
- [30] S. T. Wierzchon. Deriving concise description of non-self patterns in an artificial immune system. In S. T. Wierzchon, L. C. Jain, and J. Kacprzyk, editors, *New Learning Paradigm in Soft Computing*, pages 438–458, Heidelberg New York, 2001. Physica-Verlag.