Panel: The Future of Biologically-Inspired Security: Is There Anything Left to Learn?

Anil Somayaji
School of Computer Science
Carleton University
Ottawa, Ontario, Canada
soma@scs.carleton.ca

Michael Locasto
Dept. of Computer Science
Columbia University
New York, New York, USA
locasto@cs.columbia.edu

Jan Feyereisl
School of Computer Science
University of Nottingham
Nottingham, England, UK
jqf@cs.nott.ac.uk

ABSTRACT

While biology has inspired much of the vocabulary in computer security, biologically-inspired security remains a controversial research strategy. This panel was convened to address the issue of biologically-inspired security by raising the question of whether there is anything left to learn. The discussion at NSPW touched on many issues, ranging from the nature of evolved and intelligent systems to whether anything in security works. The final consensus, however, was that while there may be promise in biologically-inspired defenses, we need to clarify our goals and develop better evaluation methodologies if we are to see further successes in such approaches.

Categories and Subject Descriptors

D.4.6 [**Operating Systems**]: Security and Protection

General Terms

Security, Design

Keywords

biologically-inspired security, computer immune systems, diversity, autonomic computing

1. INTRODUCTION

Although the language of computer security is dominated by the metaphors of war, biology has had a significant role in shaping the language of the field. From viruses and worms to immune systems and self-healing software, living systems have been held up as examples of what to fear and what to imitate. The question for this panel discussion was what is the future of biology and computer security. Specifically, as computer security experts, what more should we learn from biology, if anything at all?

This question is difficult to address because of the broad nature of the topic; however, there are three basic positions to be taken:

1. We've learned all we should from biology.
2. We haven't learned enough, our methods and technology should be further informed by biology.
3. We've borrowed too much, biology has mislead us as to the proper solutions to computer security problems.

In the rest of this document we present these positions as they were presented at the start of the panel discussion at NSPW 2007. These simplified, somewhat extreme viewpoints then served as the basis for a lively discussion that illuminated several facets of biologically-inspired security.

The rest of this paper proceeds as follows. To provide context for the panel discussions, Section 2 presents background material on the interactions between biology and computer security. Sections 3–5 present the three initial positions taken by the panelists. Section 6 summarizes the subsequent discussion, and Section 7 concludes.

NSPW'07, September 18–21, 2007, North Conway, NH, USA.
ACM 978-1-60558-080-7/07/09.

2. BACKGROUND

Biological systems have long inspired computer scientists. Biological inspiration in computer security dates, at least, to the coining of the term "computer virus" in the early 1980's [5]. While self-propagating malware has clear lifelike properties [24], in contrast, currently used defenses are not very biological in flavor. That is not to say that the ideas and mechanisms of biology have not influenced the development of computer defenses. We classify past efforts into three broad categories: computer immune systems, diversity, and autonomic computing. In the rest of this section, we describe past work in these areas and how they have been informed by biology.

2.1 Computer Immune Systems

The most potent metaphor has been that of a "computer immune system." Natural immune systems are able to respond to novel threats autonomously and effectively, something that no current computer defense mechanism can do reliably. Natural immune systems, particularly those of large mammals such as humans, are remarkably complex systems that are only incompletely understood. Thus, "computer immune systems" are not complete copies of whole natural immune systems; instead, they are excerpts that have been simplified and translated.

Much of this metaphorical translation has taken place at the level of architecture: many researchers (e.g., [6, 8, 19]) have proposed agent-based systems in which individual agents function analogously to the various cell types of the human immune system. What has been more influential, though, has been work on translating specific mechanisms and strategies. For example, the negative selection algorithm [11], an abstraction of how detectors are generated

and used in the human immune system, has been applied to the problem of detecting viruses [11], detecting anomalous network connections [15], and in large part has given rise to the field of Artificial Immune Systems (AISs) [7]. It is important to note, though, that work in AISs has largely shifted away from computer security. This shift is primarily because the key properties of the negative selection algorithm, redundant detectors that can be used in a distributed, decentralized manner, are not properties that are useful for computer defenses, at least as commonly envisioned.

Another immune system metaphor that has been influential has been MHC virus detection mechanism [23]. MHC molecules display fragments of "running code" (protein fragments, or peptides) on the surface of each cell for passing T-cells to inspect. The detectors on T-cells are negatively selected such that they do not match most normal peptides. Thus, when the detectors on a T-cell match the presented peptide, the T-cell concludes the cell is running code that it shouldn't and kills it.

Forrest [10] first proposed that short sequences of system calls could be a peptide analog and be used to detect security violations in the form of unusually-behaving processes. This work has inspired a large amount of follow-up work, primarily in the form of applying "sequence-based analysis" to other data streams [25] and in developing alternative techniques for modeling system calls [21]. The development of mimicry attacks [28], in particular, has inspired researchers to look into more precise ways of characterizing program behavior [9]. One noticeable pattern in the literature is that while program-level anomaly detection was first inspired by the immune system, current work in the area does not draw upon biology, except as a metaphor for styles of attack.

In addition to these, there has also been work in applying immunological "danger theory" to computer security [3, 14] and concepts from the innate immune system [26]. To date, however, such work has had limited influence on the computer security community.

2.2 Diversity

It has long been recognized that diversity can improve the robustness of computer systems; methodologies such as n-version programming [4] traditionally have been of interest more to the fault-tolerance community than computer security experts. Over the past five years, though, many researchers and security practitioners have recognized the dangers of a "software monoculture" [13], one where an overabundance of a single software platform is susceptible to attacks much as fields of a single variety of plant (e.g. wheat, potatoes) can be destroyed by a single virus or parasite.

In living systems, diversity is largely maintained through two processes: sexual reproduction and speciation. Sexual reproduction allows a population of similar individuals to combinatorially recombine different characteristics so as to create unique, but similar, individuals. Speciation, on the other hand, allows populations to split and diverge, thereby increasing the overall diversity in an ecosystem.

While the problem of diversity in computer systems has been framed in biological terms, the solutions have not. Since first proposed by Forrest [12], efforts to secure systems through automated diversity mechanisms have focused more on randomized implementations (of memory layout [2] and instruction sets [1, 16]) than on functional diversity—systems that have different purposes and specifications, not just implementations. However, diversity-introducing mechanisms such as address space layout randomization (ASLR) [20] have proven to be useful mitigation mechanisms and are now part of Windows Vista and many Linux distributions.

2.3 Autonomic Computing

Autonomic computing [17] refers more to a dream than a specific technical proposal. The dream is for computers to maintain themselves with minimal human assistance, much as living systems are able to do. A part of such self-maintenance would be the ability to automatically detect and respond to security violations.

While the goals of autonomic computing are biologically inspired, most proposed solutions have their origins more in standard computer science than in biology. For example, efforts at self-healing software (a necessary component of an autonomic system) are based on fault tolerance expertise as well as technologies such as thread-level speculation, checkpoints, and error virtualization [18]—technologies that have no clear biological analogy.

One hallmark of autonomic mechanisms in living systems is that they are based on multiple, interconnected regulatory feedback loops. The complexity and potential unpredictability of highly-coupled, nonlinear systems makes engineering analogous systems a process of trial and error, at best. One early effort to use low-level, incremental feedback for security purposes was pH [22], a Linux kernel extension that delays unusually behaving processes, as determined by their system call-level behavior. Twycross and Williamson [27], with their virus throttle, showed that delays can be an effective security response without adaptive feedback. It remains an open question whether multiple interconnected feedback loops, in the style of biology, can form the foundation of practical, effective security mechanisms.

3. POSITION: WE'VE LEARNED WHAT WE NEED

One view to take on biologically-inspired security is that we've already learned what we need. As discussed above, we've already borrowed at least three concepts from biology: anomaly detection and response, diversity, and autonomic computing. These ideas, now divorced from their biological origins, continue to guide the work of many researchers and developers. Over the past decade, though, actual biological concepts have had little impact on computer security research and practice.

One potential reason for this lack of inspiration is that the implementation details of living systems are just too different from computer systems. Even if we want to imitate biological properties, we are better off ignoring how biology achieves these goals as they will not provide adequate solutions to computer security problems. To this end, consider the known weaknesses of both negative selection (limited anomaly detection performance) and sequence-based detection (mimicry attacks), arguably two of the most successful biologically-inspired algorithms. Further work on borrowing detailed mechanisms from biology are likely to lead to similarly flawed defense mechanisms.

In contrast, consider the relative success of the higher-level ideas: the field of program-level intrusion detection remains quite active, and program randomization mechanisms are now part of mainstream operating systems. Indeed, the

entire field of autonomic computing is built upon this philosophy of being inspired by biological capabilities but not be limited by biological implementations.

In other words, the high-level metaphor of biology is helpful, not the details. But if this is the case, then we've already learned what we need to know: we know the metaphors, and we are actively developing systems that exploit such metaphors. Looking for further inspiration to biology, then, is just a distraction from the main problem of building better computer defenses.

4. POSITION: WE'VE JUST SCRATCHED THE SURFACE

The view that "we've learned what we need" from biology misses the larger picture. Living systems are like an alien technology: they can do many things we don't understand, and they do so in almost magical ways. Some of the things living systems do, such as reproduce on their own, are generally not what we want our computers to do. Biological systems, however, have amazing capabilities that we haven't begun to replicate: self-defending, self-healing, self-aware systems that are massively parallel and fault-tolerant. Biology should not be dismissed so long as the gap between the natural and artificial remains so large.

Another reason to learn from biology is that many of the solutions used by living systems are "evolutionarily stable," meaning that they are robust to attacker innovation. Species adapt to new pathogens, and even when individual species get wiped out ecosystems survive all but the most severe disruptions. Biology has been engaged in "arms races" for millions of years: given the adaptive and creative power of evolution, solutions that come of of this process should be robust, even in the face of intelligent adversaries. After all, it is evolution that created intelligent adversaries in the first place!

If we wish to take advantage of what biology has to offer, we need to translate biological concepts and mechanisms to computational ones. There are three steps to this process:

1. Identify a property of living systems that we wish to replicate.
2. Understand how living systems achieve this property.
3. Translate the biology into computational mechanisms (i.e., choose the right metaphor).

Many properties of living systems that preserve their "security" have not been studied in the context of computer security: camouflage, risk assessment (fear), context-dependent trust (friendship)—these are just a few possibilities. Further, for virtually all significant biological properties, we don't understand how they are achieved: pathogen defense, homeostasis, self-repair—these are all still being actively studied by biologists. (Consider how different the world of medicine would be if we understood how these mechanisms worked!) And, for the translations, only a few have been tried: are we really sure that we've found the "correct" translation of MHC, for example?

We have a very long way to go before we've achieved parity with living systems, and there remain many avenues to explore on the path to replicating biological capabilities. That's why we've just scratched the surface of biologically-inspired security.

5. POSITION: WE'VE BORROWED TOO MUCH

It is true that there have been positive developments in security that owe some debt to biological inspiration. It is also true that there remain significant gaps between what biological systems and computer systems can do from a security perspective. However, it is also true that past work in biologically-inspired security has, in turn, inspired a lot of poor quality security research. Further, by looking to biology we are studying speculative approaches that are highly unlikely to lead to better defenses, while overlooking feasible solutions that address the real threats that computers face.

If one looks at the evolutionary computation, machine learning, and AIS literatures, one can find a number of papers on security that would never be published in a respectable computer security venue. Typically these papers report work on moderately interesting algorithms; however, the evaluation of these methods is spotty at best—results based on insufficient data, comparisons against inappropriate benchmarks. Biologically-inspired security opened the door for experts in other domains to try their hand at security. As the authors are not experts in security, they do not know what good security research looks like. The result is many papers purporting to address security problems but making negligible contributions to the field.

The problem with biologically-inspired security, though, extends far beyond this exposure to the "unwashed masses" of non-security researchers. Biology does not provide the guarantees that we want and need for computer systems. Even worse, for security guarantees that are maintained, they are provided using means that we cannot (or do not) want to adopt.

Living systems, at a low level, do not keep secrets, yet we value privacy and secrecy in most security-critical applications. Living systems treat cells and individuals as disposable; for many security-critical systems, such as military computers and financial systems, routine breaches of individual machines is not acceptable. Further, living systems just need to survive; we, however, want our computers to perform specific tasks. In these and other ways, we *want* our computers to be different from living systems.

As for desirable properties such as self repair, self-organization, and self-defense: these are properties of systems, not of individual mechanisms. In order to get the guarantees of living systems, we will likely have to implement most of their mechanisms. But living systems are very different in substrate and are staggering in complexity. Even if we successfully copied biology's mechanisms, they wouldn't necessarily work properly when translated, and that is assuming they were appropriately translated. And because of the highly nonlinear, feedback-driven, "spaghetti-code"-like structure of living systems, debugging the resulting system would be a nightmare.

In contrast, as any survey of the current security literature will reveal, there are a number of promising security solutions that have no connection to biology. These solutions address the requirements and constraints of real systems, and they are solutions *that we understand*.

While biology may be useful for inspiring metaphors, it is a siren song that distracts researchers from more productive approaches. Enough damage has already been done—we should move on.

6. DISCUSSION

During the panel discussion each of the three panelists presented one of the aforementioned viewpoints, followed by an interactive exchange of ideas and opinions. One surprising pattern in the subsequent discussion was that most NSPW attendees looked favorably upon biologically-inspired security. Even the most negative did not dismiss the approach out of hand. Rather than debate the issue as presented, the focus of the discussion instead was on problems in current approaches to security, the extent to which biologically-inspired approaches could hope to address these problems, and the barriers that prevent biologically-inspired mechanisms from being developed and deployed. These views are discussed in further detail below.

6.1 The State of Security

While some laypersons may believe their systems are sufficiently secure, security experts realize that even the best maintained systems are still remarkably vulnerable to attack. As was noted during the discussion, "nothing in security works"—even combinations of all available technologies cannot provide robust protection at an acceptable cost (in terms of administrative effort and end-user usability). Some attendees argued that given this situation, unconventional security proposals should be given the benefit of the doubt.

Computer security is a very young academic discipline and is even younger as an area of commercial interest. In contrast, biological systems started simple and grew complex over a very long period of time. Various branches of engineering proceeded through a sequence of failures before becoming increasingly reliable and useful. Thus we should not be surprised that current security solutions are inadequate.

6.2 The Knowledge Gap

Specialization is an essential part of modern research—there is simply too much for any one individual to learn. Too much specialization, however, can also impede progress. Achieving the right balance between "holistic" and specialist research strategies is particularly difficult when bringing together fields as different as biology and computer science. The gap between them is so large that specialists in either field often lack the necessary background to see the utility of collaboration. Such knowledge gaps are not necessarily symmetric, however. Consider the training required to debug and fix the state of a human (medical school) with that required to fix computers (vocational computer training). Thus, there may in fact be a bigger burden on computer scientists working with biological concepts than those going the other way.

A natural strategy for addressing such gaps is interdisciplinary collaboration. Even so, scientific communication still must be based upon a foundation of shared concepts and shared language. Differences in vocabulary are always an issues in interdisciplinary work; with computer security, however, we have the additional challenge that we don't necessarily agree on the basic concepts within our own discipline. While there is much uncertainty in some parts of biology, biologists do have the foundation concepts such as evolution, cells, and the genetic code. What are the equivalent fundamental concepts of security? Given the failures of security, are we even sure that those concepts are correct? And, do we know how to communicate them to people outside our discipline?

6.3 Unlike Adversaries?

One argument that has been repeatedly used against the idea of biologically-inspired security research is the difference between biological and digital adversaries. Computer systems are threatened by intelligent, self-aware adversaries who often intend deliberate harm and can consciously adapt to any defense. In contrast, in biology the adversaries—other organisms—are not intelligent and are motivated by a need to survive, not malevolence.

While evolution can refine the quality of their attacks, some attendees argued the lack of guiding intelligences with dangerous agendas makes their attacks fundamentally easier to repel. To capture the essence of intelligent adaptation, some argued that we should look to fields such as sociology and political science, ones that study conflicts between intelligent systems (humans). Other attendees, however, noted that evolved human diseases can appear to be pretty malevolent and hard to defeat. While several points were made on the issue, no consensus was reached.

There was agreement, however, that current computer defenses are very bad at adapting on their own to new attacks. The question, then, is whether biologically-inspired defense mechanisms will help or hinder humans when adapting to new threats.

6.4 CIA or AIC?

A larger debate followed on the differences between designed and evolved systems. Some panelists argued that designed systems are categorically different from evolved systems, while others argued that they were both fundamentally the same. A less controversial point was that designed systems are easier to understand, while evolved systems are generally more robust. Given this trade-off, the question then becomes to what extent it is appropriate to borrow "designs" from evolved systems.

Key to this question are the "design goals" of biological and computer security systems. Traditionally, computer security has focused on methods that guarantee confidentiality first, integrity second, and availability last. Attacks that transform from a breach in confidentiality or integrity to a reduction in availability are typically seen as being a good trade-off. In contrast, living systems maximize availability (survival) even if it compromises integrity—organisms will happily live with a virus's DNA intermingled with their own, for example. Further, living systems do not seem to rigorously preserve confidentiality.

Some attendees argued that biology has its priorities correct: availability should be the first, not last, priority when securing systems. Others argued that with a true loss of confidentiality and integrity, availability isn't worthwhile and might even be dangerous for some computer systems. But maybe perfect confidentiality is impossible. Is it better to focus on what is achievable and necessary, namely availability, and accept imperfect confidentiality and integrity?

What this discussion highlighted is that to make the appropriate trade-off between these goals, security issues cannot be considered in isolation; rather, they are but a few of the many important properties that computer systems should possess. One attendee suggested that perhaps what we need is a more "ecological" view of security, one that encompasses users, administrators, developers, and attackers. This idea resonated with many attendees and lead to further discussion.

6.5 New Approaches

Several ideas were mentioned for new sources of biological inspiration in addition to looking to ecology (in a way deeper than the monoculture metaphor). One attendee also suggested that behavioral defenses such as fear might be a useful metaphor. A broader view was taken by another, who pointed out that living systems use many mechanisms to ensure survival, and many of these mechanisms have security aspects. For example, systems for reproduction, from the level of determining genome compatibility to that of mate choice, can be thought of as involving issues of confidentiality, integrity, and availability. More work is needed to explore the potential utility of such mechanisms in computer security.

6.6 Evaluating Biologically-Inspired Security

Because biologically-inspired methods are quite fashionable in security, at least with certain funding agencies, many researchers have attempted to bridge the gap. The results, however, have overall been very disappointing, with many poor-quality papers being produced. Why is this the case?

The consensus view was that poor papers are produced because of improper evaluation methodologies—a problem that is largely the responsibility of the computer security community. For example, in the area of intrusion detection, there are no public datasets that are suitable for evaluating a new approach, and there are no clear guidelines as to how one could generate a suitable dataset. By opening the door to biology we are inviting new approaches. If we cannot tell others, particularly outsiders to security, how to evaluate their work, should be surprised that "bad papers" are produced?

While this problem is particularly acute in the area of intrusion detection, it also applies to any heuristic or adaptive security mechanism—in other words, to anything that cannot be formally verified. While medicine has a tradition of experimentally evaluating candidate therapies, computer security does not. But how do we get beyond the limited views of academic evaluations and commercial benchmarks? While ideas for new commercial organizations were discussed, no conclusion was reached.

6.7 The Future of Biology

One significant closing point in the discussion was that with the development of widespread and ever more ambitious genetic engineering, the gap between computer security and biology may be narrowing in a new way: soon, we may be coping with the effects of malicious, intelligent attackers exploiting flaws in engineered biological systems. What are the security implications of engineering our bodies? In such a world, maybe concepts from computer security will have applicability to biological defenses. While there is something intellectually satisfying to biologically-inspired security coming "full circle," the results are not so pleasant to contemplate.

7. CONCLUSION

This NSPW panel discussion showed that there was much interest still in biologically-inspired security. To move beyond past work in anomaly detection, diversity, and autonomic computing, however, several issues need to be addressed. Security researchers need to learn more about biology and learn how to communicate the fundamentals of their field to others. The differences between evolution and intelligence need to be better understood, both in their effects on attacker capabilities and on the organization of systems. We need to clarify our goals: how do we value confidentiality, integrity, and availability relative to each other? And while we should look to other biological mechanisms, we need to keep in mind that new approaches will only be adopted if they are evaluated appropriately. Developing better evaluation methodologies is something that will benefit the entire field of computer security.

8. ACKNOWLEDGMENTS

We would like to thank all the NSPW participants for their invaluable views and comments on the issue in question and especially Bob Blakely who acted as a scribe and without whom it would be difficult to reproduce in this paper what really went on during the panel.

9. REFERENCES

[1] E. G. Barrantes, D. H. Ackley, S. Forrest, and D. Stefanovic. Randomized instruction set emulation. *ACM Transactions on Information and Systems Security (TISSEC)*, 8(1):3–40, 2005.

[2] S. Bhatkar, D. C. DuVarney, and R. Sekar. Address obfuscation: An efficient approach to combat a broad range of memory error exploits. In *Proceedings of the 12th USENIX Security Symposium*, 2003.

[3] M. Burgess. Computer immunology. In *Proceedings of the 12th system administration conference (LISA '98)*. USENIX Association, 1998.

[4] L. Chen and A. Avizienis. N-version programming: A fault-tolerant approach to reliability of software operation. In *The Twenty-Fifth International Symposium on Fault Tolerant Computing: Highlights from Twenty-Five Years*, pages 113–119. IEEE Computer Society, 1995.

[5] F. Cohen. *Computer Viruses*. PhD thesis, University of Southern California, 1985.

[6] M. Crosbie and G. Spafford. Defending a computer system using autonomous agents. Technical Report CSD–TR–95–008, Department of Computer Sciences, Purdue University, 1995.

[7] D. Dasgupta, editor. *Artificial Immune Systems and Their Applications*. Springer-Verlag, Inc., Berlin, 1999.

[8] D. Dasgupta. Immunity-based intrusion detection system: A general framework. In *Proceedings of the 22nd National Information Systems Security Conference*, 1999.

[9] H. Feng, O. Kolesnikov, P. Fogla, W. Lee, and W. Gong. Anomaly detection using call stack information. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy*, 2003.

[10] 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*, 1996.

[11] S. Forrest, A. Perelson, L. Allen, and R. Cherukuri. Self-nonself discrimination in a computer. In *Proceedings of the 1994 IEEE Symposium on Research in Security and Privacy*, 1994.

[12] S. Forrest, A. Somayaji, and D. Ackley. Building diverse computer systems. In *6th Workshop on Hot*

Topics in Operating Systems (HotOS-VI). IEEE Computer Society, 1997.

[13] D. Geer et al. Cyberinsecurity: The cost of monopoly how the dominance of microsoft's products poses a risk to security. http://cryptome.org/cyberinsecurity.htm, September 2003.

[14] J. Greensmith, U. Aickelin, and S. Cayzer. Introducing dendritic cells as a novel immune-inspired algorithm for anomaly detection. In *Artificial Immune Systems: Proceedings of the 4th International Conference (ICARIS 2005)*, volume LNCS 3627, 2005.

[15] S. A. Hofmeyr. *An Immunological Model of Distributed Detection and its Application to Network Security*. PhD thesis, University of New Mexico, 1999.

[16] G. S. Kc, A. D. Keromytis, and V. Prevelakis. Countering code-injection attacks with instruction-set randomization. In *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*, 2003.

[17] J. O. Kephart and D. M. Chess. The vision of autonomic computing. *Computer*, 36(1):41–50, 2003.

[18] A. D. Keromytis. Characterizing self-healing software systems. In *Proceedings of the 4th International Conference on Mathematical Methods, Models and Architectures for Computer Networks Security (MMM-ACNS)*, St. Petersburg, Russia, September 2007.

[19] J. Kim and P. Bentley. The artificial immune model for network intrusion detection. In *Proceedings of the 7th European Conference on Intelligent Techniques and Soft Computing (EUFIT '99)*, Aachen, Germany, 1999.

[20] PaX Team. PaX address space layout randomization (ASLR). http://pax.grsecurity.net/docs/aslr.txt.

[21] R. Sekar, M. Bendre, P. Bollineni, and D. Dhurjati. A fast automaton-based method for detecting anomalous program behaviors. In *IEEE Symposium on Security and Privacy*, 2001.

[22] A. Somayaji. *Operating System Stability and Security through Process Homeostasis*. PhD thesis, University of New Mexico, 2002.

[23] A. Somayaji, S. Hofmeyr, and S. Forrest. Principles of a computer immune system. In *Proceedings of the 1997 Workshop on New Security Paradigms*. ACM Press, 1998.

[24] E. H. Spafford. Computer viruses as artificial life. *Journal of Artificial Life*, 1(3):249–265, 1994.

[25] M. Stillerman, C. Marceau, and M. Stillman. Intrusion detection for distributed applications. *Communications of the ACM*, 42(7):62–69, July 1999.

[26] J. Twycross and U. Aickelin. Towards a conceptual framework for innate immunity. In *Artificial Immune Systems: Proceedings of the 4th International Conference (ICARIS 2005)*, volume LNCS 3627, 2005.

[27] J. Twycross and M. M. Williamson. Implementing and testing a virus throttle. In *Proceedings of the 12th USENIX Security Symposium*, pages 285–294, 2003.

[28] D. Wagner and P. Soto. Mimicry attacks on host-based intrusion detection systems. In *CCS '02: Proceedings of the 9th ACM conference on Computer and communications security*, 2002.