


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>

available at www.sciencedirect.comwww.compseconline.com/publications/prodinf.htmInformation
Security Technical
Report

Immunology, diversity, and homeostasis: The past and future of biologically inspired computer defenses

Anil Somayaji

ABSTRACT

While biological metaphors have a long history of use in computer security, they have been more successful in describing malicious software than in helping to create better defenses. Immune system mechanisms, diversity, and homeostasis have helped inspire the development of promising computer security technologies. To fulfill the promise of biologically inspired security, however, more work is needed in understanding how living systems defend themselves and how those ideas can be brought to computers. This paper presents past successes, limitations, and opportunities for future work in this promising area. It then also addresses why there are significant barriers to such advances.

© 2007 Elsevier Ltd. All rights reserved.

1. Introduction

From the coining of the term “computer virus” (Cohen, 1985), biological metaphors have been used to describe computer security problems. Computer security defenses, however, have generally been described using other terminologies, e.g., firewalls, virus scanners, patches, and intrusion detection systems. This difference in terminology is telling: while self-propagating malware has lifelike characteristics (Spafford, 1994), current computer security defenses have few similarities with those used by living systems.

This is not to say, however, that there have not been efforts to change this. The ideas of computer immune systems (Kephart, 1994) and autonomic “self-healing” software (Kephart and Chess, 2003) have been heavily promoted both in the literature and in the press. One possible explanation for this discrepancy is that biological defense strategies may be inappropriate for computers. After all, living systems are not good at keeping secrets—we give away our genomes every time we lose a hair, for example. Further, parts of computer systems are not nearly as “disposable” as those of living systems—we don’t like individual programs or computers to die.

While biological defense strategies may not always be appropriate, it is more accurate to say that we don’t know

whether they are appropriate or not. As explained in the rest of this paper, biologically inspired defenses have barely tapped the richness of biological systems. Successes or failures, then, are better seen as evidence of immaturity, not inapplicability.

To explore both the potential and the limitations of past work in biologically inspired security, this paper first reviews the structure of natural immune systems in Section 2 and then selectively reviews past work on applying the biological concepts of immune systems, diversity, and homeostasis to building computer security defenses in Sections 3–5. Section 6 discusses future opportunities for biologically inspired defenses, while Section 7 discusses potential barriers to those opportunities. Section 8 then concludes.

2. Immunology

Perhaps the most important sources of biological inspiration for computer security researchers have been natural immune systems. Like most aspects of living systems, natural immune systems are remarkable in their robustness, flexibility, and complexity. While the broad outlines of how the immune systems of mammals function are known (particularly those of mice and humans), there are numerous unknowns that

E-mail address: soma@scs.carleton.ca

1363-4127/\$ – see front matter © 2007 Elsevier Ltd. All rights reserved.

doi:10.1016/j.istr.2007.10.004

greatly complicate any effort to use these systems as design templates for engineered systems. What follows, then, is a high-level description of the threats facing natural immune systems and the general approaches that are taken to address them. For a more detailed (but still high-level) introduction to the human immune system, see Hofmeyr (2000); for more on the design principles underlying the immune system, see Somayaji et al. (1998). (For readers with some background in immunology, please skip to Section 3.)

The human immune system consists of numerous cell types: T-cells, B-cells, macrophages, mast cells, natural killer (NK) cells, neutrophils, and others—and each of these types is made up of many subtypes. These cells are constantly in motion, patrolling independently and self-organizing to classify and respond to problems. To keep these many cells coordinated and controlled, dozens of distinct chemicals (some of which are identical to those used by nerve cells) are used to exchange messages. The details of this extremely complex system, and many of the basic operating principles, are still being uncovered by immunologists everyday, meaning that textbooks are out of date the moment they are published. Because of these factors, it can be extremely difficult to understand what the immune system does and how it does it; however, by understanding the “threat model” of immune systems, though, we can begin to see some patterns.

Organisms must always contend with the fact that others wish to take advantage of their success. Being reproductively successful means that one becomes a potential resource for others. In the case of animals, predators are best addressed through changes in high-level behavior (e.g., running away, fighting, and hiding). For threats where the attack takes place inside the body, however, other strategies are needed. The skin protects most of the perimeter of an animal. Organisms are more vulnerable at entry and exit points, but those are also protected. For example, the stomach uses enzymes and powerful acids to destroy foreign organisms, while the lungs use mucous and cilia to trap and expel invaders.

It is when foreign organisms get passed these barriers the immune system gets involved. In architecture, the immune system is a distributed, decentralized network of autonomous agents (cells) that collectively can recognize and respond to attacking entities, or pathogens. Immune systems such as the human immune system face four basic challenges:

- **Scale:** Parasites can be single or multicellular eukaryotic organisms, so they are often much, much larger than the individual immune system cells that must identify and destroy them. In contrast, bacteria and viruses are much smaller than the cells of the immune system, meaning that they can often hide *inside* the cells of the body—even immune system cells. This range of scales means the immune system needs a variety of detection and response strategies.
- **Communication:** Immune system cells coordinate their actions through the use of chemical signals; those same chemicals, however, can also be produced by pathogens. Compromised communications can result in responses being dulled or even prevented.
- **Errors:** Most responses produce some amount of “friendly fire”—the body’s own cells are killed in the process of

defending the body. Some deaths are unavoidable; however, others result from classification errors, some of which are exacerbated by the actions of the attackers. In extreme cases the actions of the immune system can cause more damage than the pathogen itself, even resulting in the death of the organism.

- **Evolution:** As pathogens are killed, those that survive reproduce. This evolutionary process means that they are constantly getting better at bypassing immune systems, and doing so on relatively short timescales: bacteria reproduce every 20 min, but humans reproduce every 20 years.

To address these challenges, the human immune system has to be able to respond quickly to well-known threats yet retain the capability to deal with the unknown. To address this dichotomy, the cells of the immune system are divided into two classes, the innate and adaptive immune systems. The innate immune system functions in a way analogous to most modern virus scanners: its cells can recognize specific molecular markers and behaviors that are “known” signatures of pathogens. Because such signals can be easily distinguished from those produced by the normal cells of the body, the responses of the innate immune system are both swift and accurate. Unfortunately, like virus scanners, their responses are insufficient for novel threats.

To bridge this gap, the adaptive immune system does not attempt to develop detectors that match specific new pathogens a priori; instead, it creates a large variety of detectors that can match virtually anything and then destroys or carefully regulates those that match the body’s own cells, or *self*. By being so selected, the remaining detectors will only match against *nonself*, which in general will be pathogens. Note that the primary advantage of this *negative selection* process is that individual cells decide on their own whether something doesn’t belong without each of them requiring a full inventory of what is allowed. Coordination only needs to take place to decide on the type and magnitude of response, rather than on the detection side.

When deployed in the form of antibodies or B-cells, the detectors primarily match exposed components of pathogens (e.g., proteins on the surface of a parasite). To detect cells that have been compromised internally (e.g., a virus-infected cell), additional mechanisms are required because under normal circumstances one cell cannot look inside another. One way this problem is circumvented is through the interaction between T-cells and MHC-peptide complexes.

Proteins are the workhorse molecules of the cell: they are the ones that make chemical reactions happen inside a cell, and they provide much of its structural support as well. As each cell recycles its proteins, they are chopped up into fragments known as *peptides* that are then recycled into new proteins. The MHC molecule intervenes in this process by snatching up peptides and displaying them on the surface of the cell. These displayed peptides are queried by passing T-cells in the negative selection style of the immune system: if the T-cell matches the peptide, it was part of an abnormal protein, meaning that the cell has been somehow compromised. Thus, the T-cell signals that the identified cell should be killed.

3. Artificial immune systems

Given the complexity and scope of natural immune systems, it has not made sense to translate them “wholesale” into computer mechanisms; instead, what has happened is that aspects of immune systems have become templates for specific security mechanisms or architectures. Broadly speaking, applications of immune system ideas in computer security can be divided into three categories: negative selection-based systems, MHC-based systems, and other systems. What follows is a brief overview of the first two; we discuss other approaches in Section 6.

3.1. Negative selection-based systems

As discussed previously, negative selection is the process whereby the immune system selects detectors (e.g., T-cells): it randomly generates detectors that have the potential to match everything, and then it throws away those that match normally behaving cells. Forrest et al. (1994) first applied this process to a computer security problem, specifically the detection of virus-infected files. Later, Hofmeyr (1999) proposed LISYS, a network intrusion detection system based upon the negative selection algorithm, along with various immunologically inspired mechanisms for managing false positives.

While these systems were both major advances, they both had certain basic limitations. When applied to individual computer systems, the virus detection mechanism didn't have any clear advantages to storing cryptographic checksums of files (e.g., Kim and Spafford, 1994). LISYS had the limitation that it detected unusual network connections on the basis of IP addresses and ports, a choice that restricts its ability to detect many kinds of network attacks. Efforts by other researchers to increase the number of features observed ran into issues regarding accuracy and detector generation time (Kim and Bentley, 2001); a response by Balthrop et al. (2002) argued that such problems arose not from limitations in the negative selection algorithm, but from the algorithm being applied in inappropriate ways.

In the field of artificial immune systems (AISs) (Dasgupta, 1999), many researchers have taken up the task of expanding the bounds of negative selection, applying it to new domains, studying its properties, and also developing algorithms based upon other immune system mechanisms. Such work has now largely moved from the area of computer security to that of machine learning. However, the more general idea of representing a set (self) by using detectors that match its complement (nonself) may have significant applications in the area of privacy-preserving databases (Esponda et al., 2006), among others.

While the negative selection algorithm has inspired significant amounts of research, its legacy, for the moment, is found in work that has very little connection with the basic problems for which immune systems evolved. While disappointing, this should probably be expected: current computer systems are not composed of billions of autonomous, disposable, self-replicating computing elements. Architectures and algorithms that work well in such circumstances, then, should not be expected to automatically apply to other systems.

3.2. MHC-based systems

The MHC mechanism, by allowing T-cells to observe the internal state of other cells, acts as a kind of cell-level anomaly detector that allows the immune system to uncover cells running “malicious code” (viruses). By equating cells with running programs, MHC suggests that intrusions could be reliably detected through the use of program-level anomaly detection. The question then becomes, what is an appropriate analog for a peptide?

In 1996 we first proposed that a possible peptide analog could be short, contiguous sequences of system calls (Forrest et al., 1996). Because most modern operating systems are still relatively centralized (even on multiprocessor systems), we did not need the distributed properties of the negative selection algorithm; instead, we used simple arrays (Forrest et al., 1996) and later, trees (Hofmeyr et al., 1998) to represent sequences seen when training the system to recognize normal program behavior. The result of this work was a system that could detect a wide variety of security violations with few false positives.

These results stimulated significant amount of work in the area of program-level intrusion detection. Some researchers studied whether sequence-based analysis could be applied to other datastreams (Stillerman et al., 1999); many more, though, chose to see whether other algorithms could do better than the ones we proposed (e.g., Lee et al., 1997). Because program behavior is extremely regular at the level of invoked system calls, it turns out that complex learning algorithms are not needed and provide little benefit (Warrender et al., 1999). However, by adding additional state such as program counter information, it is possible to build more precise models (Sekar et al., 2001).

The direction of research in program-level intrusion detection changed, however, with the development of *mimicry attacks* (Wagner and Soto, 2002). The idea here was that attacks could be crafted such that they looked “normal” (i.e., they produced no novel system-call sequences). Mimicry attacks have been perceived as a significant weakness of the system-call sequence method. Many researchers have now proposed “mimicry-resistant” intrusion detection methods (e.g., Feng et al., 2003); unfortunately, these approaches generally can either detect fewer attacks overall and/or have significantly higher false alarm rates, making them less suitable for production use.

To summarize, the initial idea of MHC-style attack detection has given rise to the area of program-level intrusion detection. In its approaches and mechanisms, however, there is little left of the original biological inspiration.

4. Diversity

Computer security researchers often complain about the dangers of overly popular software: a single vulnerability can potentially lead to millions of compromised machines. Such criticisms have more recently been framed in terms of a “software monoculture”, in reference to the disease susceptibility of farms that raise single varieties of plants on large numbers of fields (Geer et al., 2003).

Uniformity in software, however, is something that is generally advantageous: it reduces the effort required for users to learn systems and for administrators to maintain them. However, it has long been recognized in the fault-tolerance literature that having diversity in implementations can increase robustness (Chen and Avizienis, 1995). Creating multiple implementations by hand, however, is an expensive proposition, one that is impractical for all but the most safety-critical systems.

To achieve the benefits of diversity without the typical costs, we proposed that software implementations be “diversified” (Forrest et al., 1997). The idea is that some attacks, particularly memory corruption-based attacks such as buffer overflows, require the attacker to understand how code and data are organized in a targeted program. By changing this organization, attacks can be made to fail without disrupting normal program behavior. We implemented a simple diversification based on padding stack variables; other researchers have diversified (or randomized) memory layouts (Bhatkar et al., 2003), instruction sets (Barrantes et al., 2005; Gaurav et al., 2003), and other aspects of low-level code behavior. Some, such as ASLR (PaX Team), have been so successful that such mechanisms are now part of Microsoft Windows Vista and many Linux distributions.

While these advances are significant, it should be noted that while the problem identified has been framed biologically, the solutions are not biological at all. Diversity in living systems comes from very different sources: sexual reproduction, death, and speciation. Sex allows for the production of unique but similar individuals, while death makes sure new variants (i.e., young individuals) have an opportunity to survive and prosper. Speciation allows groups of closely related individuals to split so that they may become more different (genetically diverge) over time.

While sexual reproduction is the inspiration for genetic algorithms and related work, these methods cannot be directly applied to conventional code bases. Thus, it would appear that the scope for truly biologically inspired diversity techniques is quite limited. We will return to this issue in Section 6.

5. Homeostasis

All biological systems maintain a stable internal state by monitoring and responding to internal and external changes. This self-monitoring is one of the defining properties of life and is known as “homeostasis”. Homeostatic mechanisms are typically autonomic, i.e., they operate below the level of consciousness. But homeostatic mechanisms have a very specific purpose, to minimize variations in the internal state of an organism. In effect, homeostasis refers to mechanisms that ensure that the “assumptions” made by an organism’s internal systems stay true.

Two classic examples of homeostasis are the temperature and fluid balance mechanisms of our own bodies. Such mechanisms are essential because significant changes in these can have catastrophic consequences. The human body detects internal changes through sensors (specialized nerve cells) in the skin and inside the body. As we become too cold, our peripheral blood vessels constrict and we shiver; when we become

too hot, peripheral blood vessels dilate and we sweat. Similarly, our kidneys secrete more or less water in our urine depending upon how hydrated we are.

Because electronics, like living systems, are sensitive to temperature and voltage extremes, computer systems have sensors (thermistors and silicon-embedded diodes) and regulatory mechanisms (fans and voltage regulators) analogous to biological homeostatic mechanisms. At the software level, though, there is little correspondence.

This discrepancy is significant because natural immune systems can be thought of more as homeostatic mechanisms than strict defense mechanisms. Any immune response must be proportional to the change detected, much as the violence of our shivering depends on how cold we are. Further, the type and severity of an immune response must be balanced against other factors rather than just being relative to the maliciousness of the threat. Because immune responses kill healthy cells, the immune system must always trade-off destroying the pathogen enemy with destroying the body itself. Thus, rather than just trying to destroy pathogens, the immune system controls pathogens such that the cells of the body can continue to live.

In an attempt to make a more homeostatic computer defense mechanism, I developed pH (“process homeostasis”) (Somayaji, 2002), a Linux kernel extension that delays the execution of unusually behaving processes. pH is, in effect, an unusual CPU scheduler that favors normally behaving processes. By greatly reducing the CPU cycles received by compromised programs, attacks are effectively stopped while legitimate computation proceeds.

The delay mechanism of pH was a forerunner to the network-level virus throttling of Twycross and Williamson (2003). In their system, the rate of network connections was throttled in order to limit the spread of malicious programs. While their approach is less adaptive than pH, it is also easier to predict and understand. No current commercial intrusion detection or prevention system delays unusually behaving processes (as measured by an adaptive anomaly detector); the virus throttle technology, however, has been commercialized by HP and is available in some of their router products.

6. Future opportunities

From the preceding overview of work inspired by biological immune systems, diversity, and homeostasis, it should be clear that existing work has borrowed very few ideas from nature. There has been other work at biologically inspired security mechanisms—innate immunity (Greensmith et al., 2005), the “danger” model of immune system function (Aickelin et al., 2003), and self-healing systems (Sidiroglou et al., 2005), for example. However, negative selection, MHC, and the general ideas of diversity and (to a lesser extent) homeostasis are the biological ideas that have most influenced the development of existing computer security technology.

But even the influence of these ideas has been more inspirational than practical; while the respective phenomena in nature are all immensely complicated and are far from being fully understood, the ideas at the heart of the computer science translations can be summarized in a few sentences.

Translating negative selection does not mean that we understand the dynamics of T-cells, let alone B-cells or macrophages; translating MHC does not mean that we understand exactly how the immune system detects and responds to virus-infected cells. There are many, many more deep and inspiring ideas in living systems waiting to be translated.

There is also very good reason to continue with this translation process: many of the properties that we admire in living systems are not properties of our computers—we do not have computers that can autonomously detect and respond effectively to novel attacks, for example, even in the research literature. Software systems need constant monitoring and maintenance by people, and single flaw can still result in millions of computers being compromised completely silently. Process-level anomaly detection, implementation diversity, and delay-based automated response—these are steps forward, but only small steps.

For example, consider the lackluster adoption of program-level anomaly detection. One fundamental reason why anomaly detection has not been widely adopted is that when anomalies are detected, it is not at all clear what they mean. Right now an administrator has to examine system logs, network traces, and whatever other information they can obtain in order to decide whether an anomaly was or was not an attack.

It turns out that the “anomaly classification” problem is significant for living systems as well. A T-cell cannot just kill a matching cell on its own; it must first be activated in some way by a secondary signal. The processes controlling such secondary signals are very complex and are still being unravelled. Indeed, the complexity of the immune system seems to lie not so much in the sensors but in how information from them is combined with other inputs in order to make correct decisions reliably in the face of attacker interference. Thus, in order to address the limitations of biologically inspired security mechanisms, perhaps the solution is to see how biology dealt with those limitations (Hofmeyr, 1999).

Along such a path of adopting ever more biology, however, there are many barriers to be faced. We consider these issues in the next section.

7. Barriers

After having now long worked in this divide between computer security and biology, I've come to see that there are significant technical, cultural, and conceptual barriers that have hampered, and likely will continue to hamper, such efforts. There is reason to think, however, that we will have to overcome these barriers because of current trends in the computer systems. In what follows I discuss the technical, cultural, and conceptual barriers to biologically inspired security, followed by a discussion of future trends.

7.1. Technical barriers

As explained earlier, biological systems are extremely complex, interconnected systems. As such, biological mechanisms cannot be imported wholesale to computer systems—translations must be done piecemeal. But even if one chooses

the “right part” to translate, how do you do the translation? How do you make the metaphor concrete? Simply giving parts of your system biological names isn't enough—the mapping has to somehow capture and use relevant patterns from biology.

As a step towards this, we previously identified several organizing principles that appear to be at work in natural immune systems (Somayaji et al., 1998). Some principles, such as “defense in depth”, are straightforward to apply and are already widely used; others, such as “no secure layer” (or, more properly, “no trusted layer”), are much more problematic: currently, we simply don't have any design methodologies that produce systems that do not assume that some part of the system is always secure. The real problem, however, comes with concepts like “disposability”—while our bodies barely notice the destruction of hundreds or even thousands of cells, the corruption of a single bit can bring down computers and even networks. Thus, solutions based upon disposability may have requirement that current computers cannot satisfy.

Of course, this all assumes that we understand what is going on in living systems—an assumption that is all too often false. For example, how exactly does the immune system detect rogue cells (e.g., precancerous cells) without aggressively attacking normal cells? How are specific responses chosen? How does the immune system reliably sample “self”, and how does it account for the self-peptides that it must overlook? These basic questions are not completely understood by immunologists. With such incomplete knowledge, it is hard to look to biology to create systems that solve analogous problems.

7.2. Cultural barriers

Initially, biological terminology was welcomed in the computer security community as an interesting new approach. Now, though, it is generally recognized that biological terms in new papers are signs of poor quality research. While I would prefer this not to be the case, this rule of thumb is generally true.

Why should this be the case? One factor is that it is hard to master even one field; to translate from one field to another, you need to have a solid understanding of *both*. Current training programs are not very good at producing computer scientists who understand biology, or vice versa. Thus, most work in biologically inspired security tends to be done by researchers with inadequate background, resulting in work that either draws little from biology or, more commonly, is poor quality security work.

Ultimately, good security mechanisms should be good security mechanisms, no matter their source of inspiration or design. Work in biologically inspired security, then, should stand or fall without the crutch of biological rhetoric.

7.3. Conceptual barriers

While technical and cultural challenges are significant, the biggest barrier to making advances in biologically inspired security is conceptual. It comes down to this simple problem: how biological do we really want our computers, and if they do acquire biological properties, how will we deal with it?

While living systems may be robust, flexible, and relatively secure, they are also very hard to understand or “debug” and they can be extremely unpredictable at times. Really borrowing from biology may mean creating artifacts that we don’t understand and cannot completely predict. Can we accept this?

The robustness and flexibility of biological control patterns arise, in significant part, from intertwined nonlinear feedback loops. By the fundamentally chaotic nature of such systems, they cannot be designed in a conventional sense; instead, they must be developed through an empirical process, one in which potential solutions are developed, tested, and adjusted.

These are major changes in how computer security, and indeed computer science, is currently practiced. With such difficulties, is it any wonder that when borrowing from biology, computer scientists throw the biology away as soon as it is possible?

7.4. A necessary future

While borrowing from biology may seem to lead computer scientists down a dark and windy path, perhaps we have already chosen this path on our own. Modern computer systems consist of millions upon millions of lines of code—much more than can be understood in full detail by any single person. While we still engage in designing work, the actual development and deployment of systems is a highly empirical affair: problems regularly arise that were not anticipated and, by design, should never occur—yet they do, and so we find ways to manage such problems on an ad hoc basis.

Software systems have already become so complex that techniques used to study genetic data have been applied to diagnosing computer problems (Wang et al., 2003); as software systems become ever bigger and more complex, we should expect such exchanges to become more common.

Further, the rapidly evolving and changing nature of security threats are making manual, vulnerability-targeted responses less and less tenable. Keeping up with patches and virus signature updates is hard today; what will happen when the rate and severity of attacks increase to 10- or even 100-fold?

As the sheer complexity of computers begins to approach that of living systems, and as the threats become more dynamic and numerous, the world of biology will inevitably start looking less unfamiliar. Perhaps only then we may be able to better understand what living systems have to teach us about computer security.

8. Conclusion

Biologically inspired computer security has given us ways to describe the problems facing computer systems and, in areas such as immune systems, diversity, and homeostasis, it has helped researchers to develop interesting and sometimes practical technologies. Nevertheless, there still remains much to learn from biology: we have borrowed relatively few mechanisms and concepts, and work to date has not

been sufficient to make computers as robust or self-defending as living systems.

Future work in biologically inspired security faces significant barriers of a technical, cultural, and conceptual nature. As our systems become ever more complex and as they become more threatened by rapidly adapting adversaries, however, current defensive strategies are proving to be increasingly inadequate. In such circumstances, the concepts of biology and the techniques of biologists may come to play an increasingly important role in computer security.

REFERENCES

- Aickelin U, Bentley P, Cayzer S, Kim J, McLeod J. Danger theory: the link between AIS and IDS? In: Proceedings of the second international conference on artificial immune systems (ICARIS 2003). LNCS, vol. 2787; 2003. p. 147–55.
- Balthrop J, Forrest S, Glickman MR. Revisiting LISYS: parameters and normal behavior. In: CEC '02. Proceedings of the 2002 congress on evolutionary computation; 2002. p. 1045–50.
- Barrantes Elena Gabriela, Ackley David H, Forrest Stephanie, Stefanović Darko. Randomized instruction set emulation. *ACM Transactions on Information and System Security* 2005; 8(1):3–40.
- Bhatkar Sandeep, DuVarney Daniel C, Sekar R. Address obfuscation: an efficient approach to combat a broad range of memory error exploits. In: Proceedings of the 12th USENIX security symposium; 2003. p. 105–20.
- Chen Liming, Avizienis Algirdas. 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. IEEE Computer Society; 1995. p. 113–9.
- Cohen Fred. Computer viruses. PhD thesis. University of Southern California; 1985.
- Dasgupta Dipankar, editor. Artificial immune systems and their applications. Berlin: Springer-Verlag, Inc.; 1999.
- Esponda Fernando, Ackley Elena S, Helman Paul, Jia Haixia, Forrest Stephanie. Protecting data privacy through hard-to-reverse negative databases. In: Proceedings of the 9th information security conference (ISC'06). LNCS. Springer; 2006. p. 72–84.
- Feng H, Kolesnikov O, Fogla P, Lee W, Gong W. Anomaly detection using call stack information. In: Proceedings of the 2003 IEEE symposium on security and privacy; May 2003.
- Forrest S, Perelson AS, Allen L, Cherukuri R. Self-nonsel discrimination in a computer. In: Proceedings of the 1994 IEEE symposium on research in security and privacy. IEEE Computer Society Press; 1994. p. 202–12.
- Forrest S, Hofmeyr S, Somayaji A, Longstaff T. A sense of self for Unix processes. In: Proceedings of the 1996 IEEE symposium on computer security and privacy. IEEE Press; 1996.
- Forrest Stephanie, Ackley David, Somayaji Anil. Building diverse computer systems. In: Sixth workshop on hot topics in operating systems (HotOS-VI); 1997.
- Gaurav SKc, Keromytis Angelos D, Prevelakis Vassilis. Countering code-injection attacks with instruction-set randomization. In: CCS '03: Proceedings of the 10th ACM conference on computer and communications security. New York, NY, USA: ACM Press; 2003. p. 272–80.
- Geer Daniel. Cyberinsecurity: The cost of monopoly – How the dominance of Microsoft’s products poses a risk to security, <<http://cryptome.org/cyberinsecurity.htm>>; September 2003.
- Greensmith Julie, Aickelin Uwe, Cayzer Steve. Introducing dendritic cells as a novel immune-inspired algorithm for

- anomaly detection. In: Proceedings of the fourth international conference on artificial immune systems (ICARIS 2005). LNCS, vol. 3627; 2005. p. 153-67.
- Hofmeyr Steven A. An immunological model of distributed detection and its application to network security. PhD thesis. University of New Mexico; 1999.
- Hofmeyr Steven A. An interpretative introduction to the immune system. In: Segel Lee A, Cohen Irun R, editors. Design principles for the immune system and other distributed autonomous systems. Oxford University Press; 2000.
- Hofmeyr S, Somayaji A, Forrest S. Intrusion detection using sequences of system calls. *Journal of Computer Security* 1998; 6:151-80.
- Kephart Jeffrey O. A biologically inspired immune system for computers. In: Artificial life IV: Proceedings of the fourth international workshop on the synthesis and simulation of living systems; 1994. p. 130-39.
- Kephart Jeffrey O, Chess David M. The vision of autonomic computing. *Computer* 2003;36(1):41-50.
- Kim Gene H, Spafford Eugene H. The design and implementation of tripwire: a file system integrity checker. In: CCS '94: Proceedings of the second ACM conference on computer and communications security. ACM Press; 1994. p. 18-29.
- Kim Jungwon, Bentley Peter J. An evaluation of negative selection in an artificial immune system for network intrusion detection. In: Proceedings of the genetic and evolutionary computation conference (GECCO-2001), 7-11. Morgan Kaufmann; 2001. p. 1330-7.
- Lee Wenke, Stolfo Salvatore, Chan Patrick. Learning patterns from Unix process execution traces for intrusion detection. In: Proceedings of the AAAI97 workshop on AI methods in fraud and risk management; 1997.
- PaX Team. PaX address space layout randomization (ASLR). <<http://pax.grsecurity.net/docs/aslr.txt>>.
- Sekar R, Bendre M, Bollineni P, Dhurjati D. A fast automaton-based method for detecting anomalous program behaviors. In: IEEE symposium on security and privacy; 2001.
- Sidiroglou Stelios, Locasto Michael E, Boyd Stephen W, Keromytis Angelos D. Building a reactive immune system for software services. In: Proceedings of the USENIX 2005 annual technical conference; 2005.
- Somayaji Anil. Operating system stability and security through process homeostasis. PhD thesis. University of New Mexico; 2002.
- Somayaji A, Hofmeyr S, Forrest S. Principles of a computer immune system. In: New security paradigms workshop. New York: Association for Computing Machinery; 1998.
- Spafford Eugene H. Computer viruses as artificial life. *Journal of Artificial Life* 1994;1(3):249-65.
- Stillerman Matthew, Marceau Carla, Stillman Maureen. Intrusion detection for distributed applications. *Communications of the ACM* July 1999;42(7):62-9.
- Twycross Jamie, Williamson Matthew M. Implementing and testing a virus throttle. In: Proceedings of the 12th USENIX security symposium; 2003. p. 285-94.
- Wagner David, Soto Paolo. Mimicry attacks on host-based intrusion detection systems. In: CCS '02: Proceedings of the ninth ACM conference on computer and communications security. New York, NY, USA: ACM Press; 2002. p. 255-64.
- Wang Yi-Min, Verbowski Chad, Dunagan John, Chen Yu, Wang Helen J, Yuan Chun, et al. STRIDER: A black-box, state-based approach to change and configuration management and support. In: Proceedings of the 17th large installation systems administration conference (LISA '03); 2003. p. 159-72.
- Warrender Christina, Forrest Stephanie, Pearlmutter Barak. Detecting intrusions using system calls: alternative data models. In: Proceedings of the 1999 IEEE symposium on security and privacy; 1999.

Professor Anil Somayaji is an Assistant Professor in the School of Computer Science at Carleton University. He received a B.S. (1994) degree in mathematics from the Massachusetts Institute of Technology and the Ph.D. (2002) degree in computer science from the University of New Mexico. He has served on the program committees of the USENIX Security Symposium and the New Security Paradigms Workshop, among others. His research interests include computer security, operating systems, complex adaptive systems, and artificial life.