

Understanding Mutable Internet Pathogens, or How I Learned to Stop Worrying and Love Parasitic Behavior

Kevin R. B. Butler and Patrick D. McDaniel

Systems and Internet Infrastructure Security Laboratory
Pennsylvania State University
University Park PA 16802, USA

Abstract. Worms are becoming increasingly hostile. The exponential growth of infection rates allows small outbreaks to have worldwide consequences within minutes. Moreover, the collateral damage caused by infections can cripple the entire Internet. While harmful, such behaviors have historically been short-lived. We assert the future holds much more caustic malware. Attacks based on mutation and covert propagation are likely to be ultimately more damaging and long lasting. This assertion is supported by observations of natural systems, where similarly behaving *parasites* represent by far the most successful class of living creatures. This talk considers a parasite for the Internet, providing biological metaphors for its behavior and demonstrating the structure of pathogens. Through simulation, we show that even with low infection rates, a mutating pathogen will eventually infect an entire community. We posit the inevitability of such parasites and consider ways that they can be mitigated.

1 Introduction

Internet worms are possibly the most intimidating of all the malicious entities that can attack systems and users. They are representative of the most volatile attacks currently available. From as early as 1988, security researchers have been cognizant of the ease at which worms can rapidly propagate across a network [1]; the rates of propagation have only increased in the period of time since then. Consider the Slammer worm, which was able to strike 90% of its intended targets within ten minutes of being released [2]. Accordingly, so-called *flash worms* have been considered the most dangerous of all the worm variants, and attempting to contain them is a very active area of research.

In this work, however, we posit that as troubling as these fast-moving worms might be, with their ability to infect large portions of the population in a short amount of time (i.e., individual hosts potentially infected in under one second [3]), they still do not represent the worst possible scenario. We consider a new form of malevolent digital organism, the Internet parasite. The parasite exhibits worst-case behavior as follows:

1. It operates silently within the host, which remains unaware as to its presence.
2. It transmits itself to other hosts with the same frequency and behavior as other traffic, making it non-anomalous and undetectable to intrusion detection systems.
3. It acts autonomously and evolves new methods of learning behavior and attack patterns against new systems.

We simulate how such a parasite would propagate in a network, using parameters for infection and recovery determined from historical epidemiological research. We find that small changes in the rates of infection, mutation, and inoculation can have dramatic changes on whether a parasite will die out or eventually fully propagate to every host in the network. For a sufficiently high rate of mutation within a parasite, even a well-defended network will eventually succumb. The effectiveness of a parasite's infection vector and its resistance to host inoculation also play major roles in determining whether the network will fall.

In order to understand these new digital organisms, we begin by examining biological species that provide the metaphor for parasitic behavior.

2 Physical Parasites

Unbeknownst to many, parasites are the most abundant lifeform on the planet, with as many as three parasite species existing for every one “free-living” species [4]. Their success has been predicated on many factors. Parasites can foist degrees of unwanted behavior on their hosts, spawning across multiple generations before reaching their intended target. This can be manifested through children not necessarily resembling their parents, a characteristic rarely found in other species. The reasons for these differences in resemblance are purely functional, in order to continue the multi-generational life cycle that some of these parasites exhibit. A vitally important trait is that the methods a parasite uses to attack a host can change depending on the species of the intermediate or final victim.

Toxoplasma gondii is a parasite that lives in cats as their ultimate hosts. While many species can be infected by the parasite, cats are the only mammalian species that *T. gondii* will sexually reproduce within. The parasite multiplies within a cat's gut, and is shed in its feces [5]. Rats will eat cat feces that carry the parasite, which propagates back to the cat in an ingenious manner. Rats are naturally afraid of cats as a predator species. However, when *T. gondii* is ingested by the rat, it works through several organs to the rat's brain, creating cysts that alter the rat's behavior. Infected rats lose many of their environmental fears, making them act “bolder” and therefore more susceptible to being eaten by cats, returning the parasite back to its preferred host [6]. This is an example of exploiting vulnerabilities in a straightforward way—the parasite is able to control the rat's behavior by affecting its brain. In a similar manner, computer pathogens can take over a compromised system, making it work in a manner not in its best interest.

In humans, the blood fluke (*Schistosoma mansoni*) has existed for hundreds of years and causes schistosomiasis, a malady affecting over 200 million people worldwide [7]. The parasite exists as larva in freshwater snails. When the larvae are mature, they burst out of the snail into water in a free-swimming form, where they can penetrate into the skin of humans venturing into the contaminated water. They enter the bloodstream, and although they grow to between 9 and 12 mm in length, they evade detection from the body's immune system by sloughing off their own proteins and covering themselves with proteins from the host (i.e., human antigens) until they are ready to reproduce [8]. The parasites seek out the human liver as a spawning ground, and eggs enter the large intestine or the bladder, where they are passed through urine or feces into fresh water, to hatch into larvae and attack the snails that serve as their intermediate hosts [9].

In this example, we see the multiple attack vectors a parasite can employ in order to target different hosts, and how it transforms its shape and behavior in quest of a goal state that can be several stages away. *S. mansoni* exploits particular vulnerabilities specific to the host in question and transfers itself in an innocuous manner. It also makes itself indistinguishable in the human host, acting like part of the host and obfuscating its signature. Combined with its ability to change shape and behavior, the parasite presents a compelling analogy to a particularly malicious computer virus that possesses the ability to morph into multiple forms.

The final key to a fully-realized parasite is the specialized manner in which it manages to attack and flourish within its host. As we have seen, these behaviors are complex enough to appear as if they had been thought out in advance. However, they are the result of countless generations of evolutionary behavior. As Darwin wrote about the human eye,

Although the belief that an organ so perfect as the eye could have been formed by natural selection, is more than enough to stagger any one; yet in the case of any organ, if we know of a long series of gradations in complexity, each good for its possessor, then, under changing conditions of life, there is no logical impossibility in the acquirement of any conceivable degree of perfection through natural selection. [10]

This provides a powerful metaphor for artificial organisms, in that the most successful digital organisms should be able to morph and mutate to dynamically propagate.

3 A Model for a Computer Parasite

Let us consider a piece of malcode that exhibits parasitic properties. We propose that such an entity would have the ability to lie undetected in its host and transmit in a manner undifferentiated from other traffic. To that end, we suggest a parasite that exhibits the following behavior:

- It listens to incoming and outgoing traffic on its host and determines which ports are open based on this information.

- It infers the protocol in use by constructing a finite state machine based on the traffic flows observed.
- Using automated methods, it dynamically discovers new vulnerabilities, saving successful exploits as part of its attack arsenal.
- It exploits the found vulnerabilities and propagates to the victim host in an undetectable manner.

While these elements have been studied individually, it is their combination in this manner that makes them particularly dangerous. We examine the different elements of this behavior in greater detail below.

3.1 Traffic Observation and Protocol Inference

The first stage of the parasite’s life cycle is to eavesdrop on the series of messages and infer flow relationships based on this information. Such a task is not onerous, as TCP sequence numbers and other methods can be used to determine this information. The parasite will listen over all of a host’s interfaces to determine potential relationships. In the taxonomy of worms presented by [11], *passive worms* exhibit this form of behavior, waiting for host machines to contact or be contacted by other machines.

The parasite will then attempt to infer the protocol represented by a given flow through construction of a finite state machine. For example, it may be able to infer the existence of the FTP protocol in use by noting that an outgoing connection is made to port 21 of a remote server, with a USER message sent, and the subsequent message containing a PASS message followed by a string, followed by file transfer activity. Observing the USER and PASS messages at the beginning of every transaction to a given port can provide the basis for reconstructing the protocol.

Inferring a protocol from network flows has been studied at the network level. Such inference engines already exist for measuring TCP connection characteristics by observing traffic [12] and as model checking tools for probabilistic systems [13] and communication protocols [14].

3.2 Generating Attack Vectors

Based on the information inferred, the parasite constructs messages to send to peers it has already communicated with over previously used protocols. It can try fault-injection methods [15] to craft messages that exploit potential weaknesses in the protocol; for example, if the largest message seen is 250 bytes, it can try sending a 500-byte message to a peer and seeing if the connection terminates, or what the failure mode is.

One important quality of the parasite is that it can be capable of learning methods of exploitation and employing what works against future potential hosts. It can also continue probing and attempting random attacks against other hosts, making it usable against any potential platform and network protocol. The AGENT architecture [16] is a key piece allowing the generation of new attacks.

AGENT will systematically generate real attacks based on the information afforded it. Using rule-sets, it will exhaustively generate all possible attacks from a known attack instance, and can prove, based on a sequence of packets, that a sequence comprises an attack. Evolutionary and genetic algorithms are the foundation for this approach [17], and other systems such as THOR [18] add injection attacks to determine intrusion detection. Additionally, tools such as GARD [19] will generate a signature for complex attacks; hence, it is possible for a permuted set of functions to be compared against the signature for determination of whether an attack will be successful and potentially increasing the effectiveness of mutated instances. The generation of new attacks is indicative of the polymorphic behavior we seek to exploit.

3.3 Covert Transmission

If an exploit is discovered that the parasite can take advantage of, it transfers itself to the host using the communication channel discovered in a “low and slow” manner. That is, it divides itself into small blocks that will fit into a packet of typical size for the protocol it is exploiting, and reassembles itself through when it has finished transferring. Current protocols such as BitTorrent [20] are capable of subdividing files into small blocks to be transferred; the parasite would employ similar behavior.

The host should be unable to distinguish parasitic traffic because of the strict use of already-established relationships. Additionally, the only clues to its existence would be the occasional dropping of connections or potential system crashes, depending on the nature of the potential exploit attempted by the parasite. However, given the varied reasons for connection failure and the computational cycles consumed by spyware and adware on many users’ machines, attempting to diagnose such random activity could prove extremely difficult. Worms that behave in this manner (i.e., only propagating across currently existing communication channels) are defined as *contagion* worms [21], capable of stealthily spreading across networks but limited by the reliance on pre-programmed vulnerabilities to exploit in the client and server hosts. While these worms infect only hosts that are connected to legitimately, they derive much of their power from the *small-world* nature of the network topology [22], where an infection of highly connected machines allows the potential for many more hosts to be infected. In addition, peer-to-peer communication is another manner in which these worms can spread between hosts.

4 Simulating Parasitic Propagation

We have considered the methods by which propagation of parasites can occur. In this section, we provide empirical results based on simulation of a system that models real propagation. Our simple model is instructive for showing how parasitic behavior differs in important ways from regular worm propagation.




Fig. 1. Percentage of trials (out of 100) where entire network was infected.


Our *parasim* simulator models a network of 500 nodes. For simplicity, we assume that all nodes have the ability to directly connect with each other. We model the probability of infection P_i by the pathogen from an infected host to a victim, the probability of inoculation P_n that cures the host of the infecting pathogen, and the probability of mutation P_m , which considers the likelihood of a particular pathogen changing into a new attack strain.

We examined literature in epidemiology and parasitology to determine a numerical basis for the values of P_i , P_n , and P_m . Previous work on the introduction of the parasite *Plagiorchis muris* in mice—itsself based on the pioneering epidemiological studies of Greenwood et al. [23]—found that within the colony of mice, there was a transmission coefficient of 0.0056 per day, with a corresponding recovery coefficient of 0.04 per day [24, 25]. We assume that the probabilities of infection, mutation and inoculation are exponentially distributed. Recall that an exponential distribution has the form


$$f(x) = \lambda e^{-\lambda x} \tag{1}$$

with mean $E(X) = \frac{1}{\lambda}$. Accordingly, we select $P_i = \frac{1}{0.0056}$ and $P_n = \frac{1}{0.04}$, equivalent to the transmission and recovery coefficients discovered from the forementioned experiments. Newly generated parasites receive values randomly selected from the distributions.

Figure 1 shows the results of 100 trials for a range of mutation probabilities. For each trial, we assume an initial infection of 25 hosts, representing 5% of the total population. We simulated 5000 rounds in each trial, and found that in the vast majority of cases, either the number of infected hosts converged on zero or the network became fully saturated. We observe that with the chosen



(a) Average number of attempted infections per time period.



(b) Median number of attempted infections per time period.

Fig. 2. Average and median attempted infections for varying values of mutation rates.




Fig. 3. Distributions of infected hosts from the first 100 virus mutations of a parasite trial with 5% mutation rate. Note that the majority of mutations fail, but a very small number infect large numbers of hosts.

parameters, the percentage of fully infected hosts increases almost linearly from a mutation probability of 0.002 until an equilibrium point is reached at $P_m = 0.03$. Variations in the graph and the lack of full saturation are attributable to randomness within the distribution.

As the probability of mutations rises, the rate at which hosts are infected increases dramatically. Figure 2(a) shows the average number of infected hosts per time period for mutation probabilities between 0.01 and 0.05, with the number of time periods limited to 600 rounds for clarity. Some interesting trends emerge in this graph: note that for each data series, the average number of infected hosts decreases slightly in the first few rounds of the infection, because of the inoculation rate being higher than the attack rate. For $P_m = 0.01$, the average number of infected hosts does not increase appreciably in this time period. For $P_m = 0.03$ and above, however, note that there are points in the curve where the rate of infection increases dramatically. These *points of criticality* are dependent on the parameters, but it is clear that past this point, infections that previously had been contained to a small number of hosts suddenly become epidemic in nature. Figure 2(b), which shows the median number of infected hosts at each time period, illustrates these points of criticality even more starkly, as they approximately appear at times $t = 280$, 210 and 120 for $P_m = 0.03$, 0.04 and 0.05, respectively. As these graphs show, the rate at which the parasite mutates can have dramatic implications on whether it will saturate the network and how quickly this will occur.

To see the effect that individual mutation strains have on the infection rate, we consider the distribution from a single trial with $P_m = 0.05$, shown in Fig-




Fig. 4. Number of infected hosts for differing values of the initially infected set.

ure 3. The first 100 mutant strains are shown in the graph. Notice that in this trial, the initial strain that begins the infection does not go on to infect many hosts; the number of infected hosts steadily declines to zero as hosts inoculate themselves. The majority of the mutated variants, in fact, infect no more than a handful of hosts before dying out. However, a strain begins spreading at $t = 200$ and begins to infect hosts, rapidly increasing the rate of infection at approximately $t = 350$ and eventually infecting over 100 hosts before hosts are inoculated. This curve, which displays a peak after rapid growth, followed by dwindling to zero, is commonly seen in epidemiological studies of pathogens dating back to some of the first quantitative studies [26]. The graph also shows that while the vast majority of mutations are failures, some mutations will result in spectacularly successful growth, a key observation in the evolutionary process of any organism.

We now consider the effect of the other variables considered on rates of infection. Unless stated otherwise, the tests keep the same parameters for P_i and P_n as previously described (0.0056 and 0.04, respectively), with 25 hosts initially infected. For clarity in the graphs, we assume a mutation rate P_m of 0.01. The first variable considered is the number of initial hosts. Figure 4 shows how the number of infections varies depending on the initial number of infected hosts. As shown, increasing the number of hosts does not dramatically change the characteristics of the infection, although with a sufficiently small number of initial infections, the infection will end (note that for an initial set of 10 infected hosts, the total number of infected hosts quickly diminishes to zero). By contrast, Figure 5 displays that saturation will occur for $P_i \geq 0.008$. This graph bears similarities with the median infections found by varying the mutation rate




Fig. 5. Number of infected hosts for differing rates of infection.

in Figure 2(b). Past the saturation parameters, increasing the infection rate merely causes saturation to occur more quickly. This is another case where a small increase in the effectiveness of an infection vector can cause a network to quickly be overcome.

In a similar fashion, varying the rates of infection, as displayed in Figure 6, shows a saturating point at 0.03 (the number of hosts saturates over a long time period at 0.03). As P_n decreases, the amount of time required for network saturation decreases.

While changing the parameters will change the slopes of these lines, the lessons are clear: for a sufficiently high rate of mutation within the parasite, tantamount to it learning new avenues for infection, even a well-defended network will eventually succumb. The effectiveness of a parasite's infection vector and its resistance to host inoculation also play major roles in determining whether the network will fall. We defer more detailed analysis and simulation, and consideration of network topologies that mimic real-world operation, for future work. In particular, research in modern parasitology has considered the virulence of parasites and its effect on host mortality [27–31] and coevolution between parasites and their hosts [32, 33], including the possibility of hosts losing immunity to infection after being inoculated. We will revisit these issues in detail in future work.




Fig. 6. Number of infected hosts for differing rates of inoculation.

5 Resisting Countermeasures

Detection and containment systems to resist attackers have been extensively examined. In this section, we consider how parasites act in the face of these systems.

Traditional intrusion detection systems (IDS) attempt to detect attackers based on their signature. Because parasites hide themselves as normal files and conceal themselves by transmission through innocuous protocols, they will not trigger alerts from an IDS. Similarly, transmitted parasites will not trigger traditional filtering mechanisms unless they are set to be aggressive enough that they attempt to quarantine or contain every incoming file. Additionally, the random protocol failures that would likely precede a successful exploit and corresponding parasite transmission could leave the defending IDS in a more vigilant state. As with a biological parasite, transmission can be detected in some cases, but methods of transmission and attack vectors can be difficult and non-intuitive to determine.

From the host's perspective, the parasite would be similarly lacking in a signature, as we assume it maintains the ability to employ polymorphic code and behavior. The less polymorphism displayed by the parasite, and the less it changes, the easier it is to defend against. Optimally, it carries no demonstrable signature. Because it acts autonomously and learns new behaviors as it evolves and traverses the network, a parasite on one system may well appear considerably different from one on another system, making static detection very difficult. Additionally, if the parasite can exploit a system, it can be forced to act in ways that prevent the parasite from being discovered. Like the rat, whose brain processes are altered by the influence of *T. gondii*, the host system can be subverted

so that the parasite's existence remains undetected through mechanisms such as buffer overflows.

Countermeasures against polymorphic worms have been suggested [34], and the AGENT architecture itself (which forms the basis of the protocol inference engine for our parasite) can be calibrated to work in either white-hat or black-hat mode, making it a potentially valuable defender against parasitic behavior. The parasite's potential for random behavior, however, could stymie efforts to ensure a full 100% success rate against any attacks it could generate. As an analog to biological parasites, although the methods of transmission and the full life cycles of many parasites are known, effectively immunizing and defending against them can still be very difficult.

Methods of detecting behavioral patterns in worms could potentially discover parasites that repeatedly employ the same methods of exploiting hosts [35], however, by searching and randomly testing for new vulnerabilities, the parasite's behavior itself can be seen as polymorphic. Similarly, methods of generating a content-based signature usable by intrusion detection systems exist, and are based on analyzing network flows without understanding the protocol behavior above TCP [36]. These methods could similarly detect parasites using similar behavioral models.

6 Conclusions

In this work, we have put forth the idea of employing parasitic behavior to create a new form of Internet pathogen, merging disparate threads of research to create a new understanding and classification. Because of the undetectable nature of network parasites and their ability to learn and evolve as they move through successive hosts, they have the potential to mimic their biological counterparts and spreading throughout the virtual world. They will form an unwelcome relationship with machines and their users.

References

1. Spafford, E.H.: The Internet worm program: An analysis. *ACM Computer Communication Review* **19** (1989) 17–57
2. Moore, D., Paxson, V., Savage, S., Shannon, C., Staniford, S., Weaver, N.: Inside the Slammer worm. *IEEE Security and Privacy Magazine* (2003) 33 – 39
3. Staniford, S., Moore, D., Paxson, V., Weaver, N.: The top speed of flash worms. In: *Proceedings of the 2nd Workshop on Rapid Malcode (WORM 2004)*, Fairfax, VA, USA (2004)
4. Zimmer, C.: *Parasite Rex: Inside the Bizarre World of Nature's Most Dangerous Creatures*. Free Press (2001)
5. Denkers, E.Y., Gazzinelli, R.T.: Regulation and function of T-cell-mediated immunity during *Toxoplasma gondii* infection. *Clinical Microbiology Reviews* **11** (1998) 569–588
6. Berdoy, M., Webster, J., Macdonald, D.W.: Fatal attraction in rats infected with *Toxoplasma gondii*. In: *Proceedings of the Royal Society of London: Biological Sciences*, London, UK (2000) 1591–1594

7. Centers for Disease Control: Parasitic disease information: Schistosomiasis fact sheet (2005) http://www.cdc.gov/ncidod/dpd/parasites/schistosomiasis/factsht_schistosomiasis.htm.
8. McKerrow, J.H.: Cytokine induction and exploitation in schistosome infections. *Parasitology* **115** (1997) S107–S112
9. Anderson, R., Mercer, J., Wilson, R., Carter, N.: Transmission of *Schistosoma mansoni* from man to snail: experimental studies of miracidial survival and infectivity in relation to larval age, water temperature, host size and host age. *Parasitology* **85** (1982) 339–360
10. Darwin, C.: *The Origin of Species*. 6th edn. John Murray (1872)
11. Weaver, N., Paxson, V., Staniford, S., Cunningham, R.: A taxonomy of computer worms. In: *Proceedings of the 1st Workshop on Rapid Malcode (WORM 2003)*, Washington, DC, USA (2003)
12. Jaiswal, S., Iannaccone, G., Diot, C., Kurose, J., Towsley, D.: Inferring TCP connection characteristics through passive measurements. In: *Proceedings of IEEE INFOCOM 2004*, Hong Kong (2004)
13. de Alfaro, L., Kwiatkowska, M., Norman, G., Parker, D., Segala, R.: Symbolic model checking of probabilistic processes using MTBDDs and the Kroenecker representation. In: *Proceedings of the 6th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'2000)*, Berlin, Germany (2000)
14. Edelkamp, S., Leue, S., Lluch-Lafuente, A.: Directed explicit-state model checking in the validation of communication protocols. *International Journal on Software Tools for Technology Transfer (STTT)* **5** (2004) 247 – 267
15. Voas, J., McGraw, G., Kassab, L., Voas, L.: A “crystal ball” for software liability. *IEEE Computer* **30** (1997) 29–36
16. Rubin, S., Jha, S., Miller, B.P.: Automatic generation and analysis of NIDS attacks. In: *Proceedings of the 20th Annual Computer Security Applications Conference (ACSAC 2004)*, Tuscon, AZ, USA (2004)
17. Spears, W., DeJong, K., Baeck, T., Fogel, D., de Garis, H.: An overview of evolutionary computation. In: *Proceedings of the 4th European Conference on Machine Learning (ECML'93)*, Vienna, Austria (2003)
18. Marty, R.: THOR: A tool to test intrusion detection systems by variations of attacks. Master’s thesis, Swiss Federal Institute of Technology, Zurich, Switzerland (2002)
19. Rubin, S., Jha, S., Miller, B.P.: Language-based generation and evaluation of NIDS signatures. In: *Proceedings of the 2005 IEEE Symposium on Security and Privacy*, Oakland, CA, USA (2005)
20. Qiu, D., Srikant, R.: Modeling and performance analysis of Bit Torrent-like peer-to-peer networks. In: *Proceedings of ACM SIGCOMM 2004*, Portland, OR, USA (2004)
21. Staniford, S., , Paxson, V., Weaver, N.: How to Own the Internet in your spare time. In: *Proceedings of the 11th USENIX Security Symposium*, San Francisco, CA, USA (2002)
22. Bu, T., Towsley, D.: On distinguishing between Internet power law topology generators. In: *Proceedings of IEEE INFOCOM 2002*, New York, NY, USA (2002)
23. Greenwood, M., Bradford Hill, A., Topley, W.W.C., Wilson, J.: *Experimental Epidemiology*. His Majesty’s Stationary Office, London, UK (1936) Privy Council, Medical Research Council.
24. Anderson, R., May, R.: Population biology of infectious diseases, part 1. *Nature* **280** (1979) 361–367

25. May, R., Anderson, R.: Population biology of infectious diseases, part 2. *Nature* **280** (1979) 455–461
26. MacDonald, G.: *The Epidemiology and Control of Malaria*. Oxford University Press, New York, NY (1957)
27. Davies, C., Webster, J., Woolhouse, M.: Trade-offs in the evolution of virulence in an indirectly transmitted macroparasite. In: *Proceedings of the Royal Society of London: Biological Sciences*, London, UK (2001) 251–257
28. Levin, B., Svanborg-Eden, C.: Selection and evolution of virulence in bacteria: an ecumenical and modest suggestion. *Parasitology* **100** (1990) S103–S115
29. Dwyer, G., Levin, S., Buttell, L.: A simulation model of the population dynamics and evolution of myxomatosis. *Ecological Monographs* **60** (1990) 423–447
30. Anderson, R.: Parasite pathogenicity and the depression of host population equilibria. *Nature* **279** (1979) 150–152
31. Lewontin, R.: The units of selection. *Annual Review of Ecology and Systematics* **1** (1970) 1–18
32. May, R., Anderson, R.: Parasite-host coevolution. *Parasitology* **100** (1990) S89–S101
33. Levin, S., Pimentel, D.: Selection of intermediate rates of increase in parasite-host systems. *The American Naturalist* **117** (1981) 308–315
34. Chistodorescu, M., Jha, S.: Static analysis of executables to detect malicious patterns. In: *Proceedings of the 13th USENIX Security Symposium*, Washington, DC, USA (2003)
35. Ellis, D.R., Aiken, J.G., Attwood, K.S., Tenaglia, S.D.: A behavioral approach to worm detection. In: *Proceedings of the 2nd Workshop on Rapid Malcode (WORM 2004)*, Fairfax, VA, USA (2004)
36. Kim, H.A., Karp, B.: Autograph: Toward automated, distributed worm signature detection. In: *Proceedings of the 14th USENIX Security Symposium*, San Diego, CA, USA (2004)