A Quasi-species Approach for Modeling the Dynamics of Polymorphic Worms

Brad Stephenson and Biplab Sikdar
Department of ECSE, Rensselaer Polytechnic Institute
Troy, New York 12180 USA

***Abstract*— Polymorphic worms can change their byte sequence as they replicate and propagate, thwarting the traditional signature analysis techniques used by many intrusion detection systems (IDSes). As the incidence of such worms becomes more frequent, it is important to understand their behavior and interaction with the IDSes in order to develop effective strategies to control their propagation. In this paper, we propose a model based on coevolution of biological quasi-species to characterize the propagation of polymorphic worms and the effects of dynamic IDSes which improve their detection capability with time. The model is used to derive the maximum allowable response time of the IDS in order to contain the worm and the optimal mutation rate the worm should use in order to escape an IDS with a given response time. The observations from the model are validated using simulations with the `ADMmutate` polymorphic engine.**

I. Introduction

The increasing complexity and sophistication of worms in the Internet is evident in recently captured specimens [11]. Polymorphic techniques, which allow a worm to change its appearance with every instance, are increasingly being used to disguise worm payloads and attempt to bypass both signature-based [13] and anomaly-based [12] IDSes. Polymorphic worms try to disguise themselves by changing their byte sequence using a variety of techniques with the most common being encryption with random keys. This disguise may allow aggressively spreading worms to evade discovery and analysis for days, and more stealthy worms may go unnoticed for many months.

In order to develop effective defense strategies and counter measures against polymorphic worms, a key step is to understand their propagation dynamics and the limitations they put on the IDS. The objective of this paper is to develop a model for the propagation of polymorphic worms which also considers the effect of a dynamic IDS which can learn and detect worms in real time. The model is then used to obtain two important factors that impact the propagation of the worm and provide system design guidelines: **(1)** the maximum allowable response time of the IDS in order to contain the spread of the worm and **(2)** the optimal mutation rate that a polymorphic worm may employ in order to evade an IDS with a known response time.

To the best of our knowledge, the few existing papers which directly deal with polymorphic worms either focus on developing mechanisms to detect them [17], [16] or study techniques they may use to evade the IDS [18]. This paper fills an important void in this area by developing a analytic framework for modeling and evaluating the dynamics of polymorphic worms and their interaction with an IDS. Our model is based on biological models for the coevolution of viral quasi-species and their interaction with the immune system. The morphing of code in polymorphic worms is analogous to the modifications in the genetic material of biological organisms in successive generations. The ability of the evolved organisms to survive depends on the ability of immune system of the host to detect the new quasi-species, analogous to the dependence of the survivability of new strains of the polymorphic worm on the effectiveness of the IDS. The evolution of the polymorphic worm's dynamics is modeled by a set of differential equations governing the co-evolutions of the quasi-species represented by the code sequences generated by the polymorphic worms and the fitness landscape representing the capabilities of the IDS. We use these equations to evaluate the conditions under which the polymorphic worm may evade the IDS and provide closed form expressions for the maximum allowable response time at the IDS and the optimal mutation rates of the worm. The observations from the model are validated qualitatively using simulations with the `ADMmutate` polymorphic engine.

The rest of the paper is organized as follows. In Section II related work is discussed. We discuss polymorphism and detail the properties of specific polymorphic engines in Section III while Section IV presents the assumptions made in our model. In Section V, we present our model for the dynamics of polymorphic worm and the results form the model are analyzed in Section VI. Finally, Section VII presents the simulation results and conclusions are presented in Section VIII.

II. Related Work

Understanding polymorphic worms remains a difficult and largely open problem. We are aware of only a handful of papers that directly address polymorphic worms and possible solutions to limit their negative effects. Kolesnikov and Lee analyzed techniques that a polymorphic worm may use to hide itself from both signature and anomaly based IDSes [18]. Particularly, they proposed that stealthy worms could collect and use knowledge about the normal characteristics of traffic on a local network to hide their propagation. Tang and Chen propose a novel double honeypot scheme to capture previously unseen worms and a position-aware signature generation algorithm [17]. Newsome et al. argue that polymorphic worms are not immune to signature-based IDSes and present the

Polygraph system for generating signatures from polymorphic payloads [16]. We will discuss some of these ideas in further detail in IV.

A number of recent papers have focused on the developing propagation models for worms and other malware in the Internet [1], [2], [3], [4], [5]. These models are usually either based on epidemiological models [1], [3], [4] or based on measurement studies [5]. However, none of these models address polymorphic worms. Intrusion detection systems and their requirements have also been extensively studied (see [6], [8] and the references therein). Again, these either specifically focus on single strain worms and viruses or do not consider the coevolution of the worm and the IDS.

With the seminal work of Eigen on replicating molecules [20], the evolution of quasi-species has been extensively studies by biologists and theoretical physicists. A vast majority of these studies consider the asymptotic behavior with static immune systems. However, recent work on time dependent immune systems [9] has allowed the development of co-evolution models of the viral quasi-species and the immune system [10]. These studies have been extended to obtain the optimal mutation rates for B-cells [10] as well as solutions for arbitrary gene networks [15]. Finally, the coevolution of computer viruses and the immune system has been alluded to in [7] in the sense of emergence of more sophisticated worms as IDSes improve in their ability to detect worms.

III. Discussion on Polymorphism

Our objective in this section is to discuss some ideas and principles of polymorphism as they relate to worms. We will present a conceptual explanation of polymorphism followed by a treatment of publicly available polymorphic engines.

As a tool for disguising malware, polymorphism is not new. Polymorphic viruses first surfaced in the early 1990s [19]. The encryption of a byte sequence to alter its appearance has obvious applications to Internet worms. IDS systems handle enormous amounts of data in real-time, and are susceptible to any attack which requires non-trivial computational power to detect. In an effort to evade detection, polymorphic worms use several techniques to remove any static signature which may be obtained from the payload while making it computationally difficult to obtain other meaningful statistics from the byte stream.

The simplest method for implementing a polymorphic worm is to encrypt the worm's payload using a key. Each time the worm replicates, it chooses a key and encrypts the new copy before sending it over the wire. The key is usually sent together with the worm and a decryption routine (or decoding engine). When the worm executes on the target machine, the decryption routine uses the key to obtain the worm code and begin the next round of propagation. This method still allows IDSes to identify the worm because the decoding engine is a constant byte sequence. If the polymorphic engine uses several decryption routines, it will be more successful in evading an IDS. However, the finite number of decryption routines still yields a finite number of static signatures. Thus, the decryption routine is still an easy target for IDSes. If a worm contains a large number of decoding engines, it may lead to a very large signature space. Suppose there are n decryption routines and one is chosen as the routine to be used for this iteration. The remaining $n-1$ routines are appended in random order, resulting in $n!$ possible signatures. But the worm's propagation rate is proportional to the payload size, so large worms travel more slowly. Also, since many exploits rely on overwriting buffers in the target's memory, a large worm may cause unforeseen consequences on the target system such as system failure before the worm is successfully executed.

Another method for implementing a polymorphic worm is to rearrange the order of instructions and use jump instructions to keep the instruction flow intact. This shuffling may result in a high frequency of jump instructions, which will create an anomalous byte distribution capable of being identified by most IDSes.

Alternatively, *non-meaningful* instructions can be inserted to create polymorphic worms. These instructions are garbage code, and usually consist of two or more idempotent instructions. For example, executing $n+2$ followed by $n-1$ and then $n-1$ is equivalent to executing nothing at all, but could serve to vary the appearance of the worm code surrounding it. Also, replacing *meaningful* sequences of instructions with different, but equivalent, sequences can render a worm polymorphic. Using this method without altering the functionality of the code requires a good deal of knowledge regarding machine language and architecture.

Lastly, changing the registers which are used in the worm executable could alter its appearance. This is a weak form of polymorphism which results in less extensive mutation between successive iterations.

By combining two or more of these techniques, more powerful polymorphism results. Consider, for example, a worm which uses the code shuffling method to change its signature and hides the jump instructions by encoding them to match the byte distribution expected by an anomaly-based IDS. The result is a highly flexible payload which attempts to evade signature- and anomaly-based IDSes.

We now consider some specific examples of polymorphic worm generators.

ADMmutate: In 2001 a hacker known as K2 presented a tool for obfuscating shellcode [13]. Shellcode is machine-executable code that is intended to open a command interpreter, or shell, on a remote machine so that an attacker can type in commands just like an authorized user. In order to launch the command interpreter, the attacker must use an exploit to trick the remote machine into executing the shell command (e.g. `/bin/sh`). Such exploits are usually some type of buffer overflow attack. For a detailed discussion of buffer overflow attacks, see [14]. For purposes of this discussion, it is sufficient to understand Figure 1. When a vulnerable function is called, the attacker's return address, R, is written over the correct return address on the stack of the current function. When the function attempts to return to its

calling function, it instead jumps to the attacker's address, slides through a series of non-operational instructions (nop), and executes the attacker's shellcode S. This "nop sled" is necessary because it is difficult to guess the exact address where the shellcode is located. The sled allows the attacker to establish a range of addresses which are acceptable for the exploit to work.

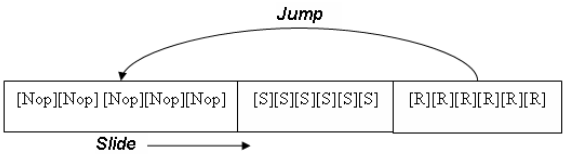

Fig. 1. How a buffer overflow attack executes arbitrary code.

In order to successfully exploit a remote machine, the attacker usually has to create a nop sled of several hundred bytes. For an IA32 machine, the nop code is 0x90, and a string of several hundred of these is highly anomalous, allowing it to be easily identified by current IDSes. The same is true for nop codes on other platforms. Thus, an IDS can stop these obvious attacks before any damage is done.

ADMmutate attempts to disguise the nop sled by taking the exploit shown in Figure 1 and changing it to the form:

[JJJJ][DJDJ][EEEE][RRRR]

where J is a "junk nop," or garbage code which is equivalent to no operation. K2 states that there are 55 possible junk nops on the IA32 architecture. E represents the encoded shellcode. The shellcode is encoded (or encrypted) because current IDS systems also alert on strings that are known to launch command interpreters. For example, if the attacker disguises the nop sled, but leaves the constant string */bin/sh* in the exploit code, the IDS will still prevent the attack. For this reason, ADMmutate encodes the shellcode with a sliding key. The decryption routine D is interlaced with garbage code as well, attempting to hide its location and make signature identification more difficult. Lastly, the return address R points to the junk nop sled. The return address is usually the least variable part of the payload, although ADMmutate attempts to modulate the address by cycling through several values that are close to the desired address.

CLET: In 2003 the CLET Polymorphic Engine was released in an issue of phrack magazine. Included in the article is a uuencoded executable of the polymorphic engine. CLET attempts to further the work done by K2 by adding more garbage instructions and including capability for crafting exploits which evade byte distribution analysis. CLET takes an input file containing the normal traffic spectrum and makes an effort to select encoded bytes which will deviate the least from the normal traffic. CLET also attempts to alter the decryption routine using equivalent instruction substitution and register swaps.

JempiScodes: Also in 2003, Argentinian coder Matias Sedalo released a polymorphic engine he entitled JempiScodes. This engine is quite easy for an untrained attacker to use. The main contribution of JempiScodes is that it provides options for the attacker to select from 4 different encryption algorithms (XOR, AND and Chained XOR in 8- or 16-bit blocks) and also tells the attacker the encryption key used. Although JempiScodes requires more coding to automate the generation of polymorphic shellcode than the other two engines, it can still be developed into an efficient tool with a minimal understanding of the underlying mechanisms.

The polymorphic methods discussed here are not a closed set. The possibilities for other techniques abound. As the polymorphic engines mentioned above demonstrate, tools for creating polymorphic worms are moving toward easier application and more sophisticated methods. It is fortunate that these techniques have not been widely deployed in Internet worms yet, but developing an understanding of polymorphic worms will prepare us to deal more effectively with their spread and effects.

IV. Model Preliminaries and Assumptions

Our purpose here is to discuss the preliminary assumptions which lead us to use the quasi-species model and justify our choice of simulation parameters. We first discuss our assumptions about the IDS followed by the assumptions about the polymorphic worm.

As has been shown in [17], it is possible to isolate worm traffic from normal traffic using a two honeypot scheme. An inbound honeypot presents an image of the server to be protected, but is configured to make no outbound connections. When it does begin to make outbound connections, the machine is assumed to be compromised and its connection requests are forwarded to an outbound honeypot which captures and analyzes the worm packets. This system allows for the complete isolation of worm flows from normal traffic flows. We assume, therefore, in the following analysis that polymorphic worms are recognizable as worms (i.e. not considered benign traffic) and can be classified accurately according to the quasi-species to which each mutation belongs. Note that in this paper we refer to each distinct worm with a different code sequence as a *quasi-species*.

In [16], it was shown that the code generated by a large class of polymorphic worms has some constant content which may be used to generate high quality signatures. In this paper, we assume that a system like that proposed in [16] is available to produce signatures for classifying a polymorphic worm. The time taken to generate successive signatures determines the effectiveness of the IDS and the likelihood that the worm population will reach epidemic proportions before being detected. We denote by τ_{ids} the time required by the IDS to generate a new signature for a new polymorphic worm quasi-species.

In Figure 2, we demonstrate some sample mutation rates from the ADMmutate polymorphic engine. We applied the engine to the same code 2000 times, counting the number of times the byte values in the worm payload did *not* change.

From the figure we note that while a number of bytes change frequently, some stay fairly constant or nearly always constant. The large values shaded in gray, then, represent bytes which could be used as part of the worm's signature.

The procedure above was also used to obtain an estimate of the mutation rate or the probability that a byte in the code sequence changes in successive generations ($\hat{\epsilon}$). To estimate $\hat{\epsilon}$ we identify N bytes from the code sequence corresponding to the signature and apply the mutation engine m times. We calculate each n_i, the number of mutations that occurred in the i^{th} byte of the signature. If m is sufficiently large, the mutation rate for the signature is then said to be

$$\hat{\epsilon} = \frac{\sum_i n_i}{Nm} \tag{1}$$

Using $N = 2000$ and $m = 4$ for the ADMmutate polymorphic engine, $\hat{\epsilon}$ was estimated to be 0.52.

37	111	43	55	28
10	137	32	53	196
172	79	93	4	172
1886	19	1955	1955	1955
11	12	17	19	15
17	10	12	14	14
12	14	9	11	15
15	15	16	7	11
9	11	13	13	11
18	13	14	6	15
13	19	12	15	19
11	10	16	11	15
12	13	15	7	5
27	6	113	1858	195
47	128	1896	199	48
128	1896	199	48	128
1894	201	48	128	1896
199	48	128	1896	199
48	128	1896	199	48

Fig. 2. Mutation rates from ADMmutate. The shaded values represent bytes that mutate infrequently. The non-shaded boxes show bytes that have a high mutation rate.

We assume that the IDS systems tries to identify the constant byte sequences in the worm instances that it comes across in order to generate the signatures. Let the number of strings or sequences constituting the signature be denoted by n. Note that the signature initially may contain sequences that change in many of the worm quasi-species and the signature may be improved with time to capture all quasi-species. Instead of considering the entire worm payload in order to classify worm instances into the corresponding quasi-species which leads to an extremely large number of possible quasi-species, we only consider the n sequences in the worm corresponding to the signature for classification. For the sequences in the signature, we assume that each sequence may mutate or morph in each successive replication independently with probability ϵ.

V. Dynamics of Polymorphic Worms

In this section we develop our model for the evolution and spreading related dynamics of polymorphic worms. We explicitly consider the impact of the IDS on the growth of the worm. The main goal of this section is to obtain the response time required from the IDS in order to contain the worm and to compute the optimal mutation rates for the worm in order to evade the IDS system.

A. Polymorphic Worm Propagation Model

Consider a worm composed of n strings or sequences. In each generation the worm mutates some of the sequences to create new quasi-species. We assume that each sequence is mutated independently with probability ϵ and that the mutated sequences are chosen from an alphabet of size S. Each copy of the worm replicates at a rate of β. Let the concentration or the total number of copies of quasi-species k generated up until time t be denoted by $y_k(t)$. The time evolution of $y_k(t)$ is characterized by the following equation

$$\frac{dy_k(t)}{dt} = \sum_l \beta W_{l,k} A(y_l(t)) y_l(t) \tag{2}$$

where $W_{l,k}$ denotes the probability of a worm quasi-species of type l morphing to a quasi-species of type k in the next generation. $A(y_l(k))$ denotes the fitness function of quasi-species l and accounts for the likelihood that the IDS is able detect and prevent the propagation of quasi-species l. The summation above is carried out over all possible worm quasi-species. Assuming each sequence to be of length b bits, we have $(2^b)^n$ possible worm quasi-species and thus the state space required for modeling the worm's dynamics can easily become quite large. For example, if we just consider 10 byte-long sequences characterizing a quasi-species, we have a total of 2^{80} possible quasi-species sequences.

To make the model tractable, we group all quasi-species according to their Hamming distance (HD) (i.e. the number of sequences in which the two quasi-species differ) from an arbitrarily chosen master quasi-species. Denoting the master quasi-species by y_0, the l^{th} group is defined as

$$w_l = \sum_{y_k \in \{y_k | HD(y_k, y_0) = l\}} y_k \tag{3}$$

and its fitness function is defined by $A(l)$. This reduces the problem from $S^n = (2^b)^n$ dimensions to $n+1$ dimensions. We now characterize the equations governing the time evolution of $w_l(t)$.

Claim 1: The time evolution of the worm quasi-species group with Hamming distance l from the master sequence is approximated by

$$\frac{dw_l(t)}{dt} = \sum_{l'=0}^{l} \binom{n-l'}{l-l'} \beta A(l') \epsilon^{l-l'} (1-\epsilon)^{n-(l-l')} w_{l'}(t) \tag{4}$$

Proof: Mutations into group l may occur in two possible ways: (1) up-mutations from groups with lower Hamming distances and (2) down-mutations from groups with larger Hamming distances. Consider the up-mutation case first with mutations from group l' to group l with $l' \leq l$. In each quasi-species of group l' there are l' sequences which have mutated previously and are different from the master quasi-species. There are three possibilities for the up-mutation:

i. $l - l'$ of the $n - l'$ sequence which are identical with the master sequence mutate and all other sequences stay the same. This case, C1, happens with probability

$$P[\text{C1}] = \left(\begin{array}{c} n-l' \\ l-l' \end{array} \right) \epsilon^{l-l'}(1-\epsilon)^{n-(l-l')} \tag{5}$$

where $\left(\begin{array}{c} n-l' \\ l-l' \end{array} \right)$ is the number of ways to choose $l - l'$ sequences to mutate from $n-l'$ non mutated sequences, $\epsilon^{l-l'}$ is the probability that there are $l-l'$ mutations and $(1-\epsilon)^{n-(l-l')}$ is the probability that the remaining sequences do not mutate.

ii. i, $0 < i \leq \min\{l', n - l\}$, of the already mutated sequences mutate back to the master quasi-species and there are mutations in $l-l'+i$ of the $n-l'$ non-mutated sequences. Given that a mutation occurs, the sequence is equally likely to change to any of the other $S - 1$ possibilities. Thus the probability that a sequence mutates and changes back to the corresponding sequence in the master quasi-species is given by

$$\frac{\epsilon}{S-1} \tag{6}$$

Then the probability of this case, C2, is given by

$$P[\text{C2}] = \sum_{i=1}^{\min\{l',n-l\}} \left[\left(\begin{array}{c} l' \\ i \end{array} \right) \left(\begin{array}{c} n-l' \\ l-l'+i \end{array} \right) \left(\frac{\epsilon}{S-1} \right)^i \epsilon^{l-l'+i}(1-\epsilon)^{n+l'-l-2i} \right] \tag{7}$$

where the first two combinatorial terms represent the number of ways of choosing $l - l' + 1$ sequences out of $n - l'$ non-mutated sequences and the number of ways to select i sequences from the l' mutated ones. $(\epsilon/(S-1))^i$ is the probability that i already mutated sequences mutate back to the master quasi-species' sequences, $\epsilon^{l-l'+i}$ is the probability that there are $l-l'+i$ new mutations and $(1-\epsilon)^{n-l+l'-2i}$ is the probability that the remaining sequences do not mutate.

iii. Of the l' already mutated sequences, i mutate back to the corresponding sequences in the master quasi-species, j $(j > 0)$ mutate to other sequences, $l' - i - j$ stay the same with $i + j \leq l'$ and $0 < i \leq \min\{l', n - l\}$ and of the $n - l'$ non-mutated sequences, $l - l' + i$ sequences mutate and the remaining $n - l - i$ sequences stay the same. This case, C3, occurs with probability

$$P[\text{C3}] = \sum_j \sum_i \left[\left(\begin{array}{c} l' \\ i \end{array} \right) \left(\begin{array}{c} l'-i \\ j \end{array} \right) \left(\begin{array}{c} n-l' \\ l-l'+i \end{array} \right) \left(\frac{\epsilon}{S-1} \right)^i \left(\frac{\epsilon(S-2)}{S-1} \right)^j \epsilon^{l-l'+i}(1-\epsilon)^{n+l'-l-2i-j} \right] \tag{8}$$

where the summations over i and j are carried out over the region where $j > 0$, $i + j \leq l'$ and $0 < i \leq \min\{l', n - l\}$. The three combinatorial terms represent the number of ways to choose i sequences from l' mutated sequences, j sequences from the remaining $l - ' i$ mutated sequences and $l - l' + i$ sequences from $n - l'$ non-mutated sequences respectively. The term $(\epsilon(S-2)/(S-1))^j$ represents the probability that j sequences mutate to sequences other than the corresponding ones in the master quasi-species and the remaining terms are along the lines of case C2. Note that the number of sequences that do not mutate is $n-(i)-(j)-(l-l'+i) = n+l'-l-2i-j$.

We now consider the down-mutations where quasi-species with a higher Hamming distance l' back-mutate to generate quasi-species with lower Hamming distance l $(l < l')$ from the master quasi-species. Again, there are three possibilities:

i. $l' - l$ of the already mutated sequences mutate back to the corresponding sequences in the master quasi-species while the remaining sequences stay the same. This case, C4, occurs with probability

$$P[\text{C4}] = \left(\begin{array}{c} l' \\ l'-l \end{array} \right) \left(\frac{\epsilon}{S-1} \right)^{l'-l} (1-\epsilon)^{n-(l'-l)} \tag{9}$$

The explanation for the terms in the expression above follows the explanation for the earlier cases.

ii. i of the already mutated sequences mutate back to the corresponding sequences in the master quasi-species with $l' - l < i \leq \min\{l', n - l\}$ and $l - l' + i$ of the $n - l'$ non-mutated sequences mutate while the remaining sequences stay the same. The probability of the occurrence of this case, C5, is given by

$$P[\text{C5}] = \sum_{i=l'-l+1}^{\min\{l',n-l\}} \left[\left(\begin{array}{c} l' \\ i \end{array} \right) \left(\begin{array}{c} n-l' \\ l-l'+i \end{array} \right) \left(\frac{\epsilon}{S-1} \right)^i \epsilon^{l-l'+i}(1-\epsilon)^{n-l'+l} \right] \tag{10}$$

The explanation for the terms in the expression above follows the explanation for case C2 in up-mutations.

iii. Of the l' already mutated sequences, i mutate back to the corresponding sequences in the master quasi-species (where $l' - l < i \leq \min\{l', n - l\}$), j $(0 < j < i$ and $i + j \leq l')$ mutate to other sequences, $l' - i - j$ stay the same and of the $n - l'$ non-mutated sequences, $l - l' + i$ sequences mutate and the remaining $n - l - i$ sequences stay the same. This case, C6, occurs with probability

$$P[\text{C6}] = \sum_j \sum_i \left[\left(\begin{array}{c} l' \\ i \end{array} \right) \left(\begin{array}{c} l'-i \\ j \end{array} \right) \left(\begin{array}{c} n-l' \\ l-l'+i \end{array} \right) \left(\frac{\epsilon}{S-1} \right)^i \left(\frac{\epsilon(S-2)}{S-1} \right)^j \epsilon^{l-l'+i}(1-\epsilon)^{n+l'-l-2i-j} \right] \tag{11}$$

where the summations over i and j are carried out over the region where $j > 0$, $i + j \leq l'$ and $l' - l < i \leq \min\{l', n - l\}$ and the explanation for the rest of the terms is similar to case C3 in the case of up-mutations.

Combining the six cases above, the probability of mutations from group l' to group l, $P[w_{l' \to l}]$, is given by

$$P[w_{l' \to l}] = P[\text{C1}] + P[\text{C2}] + P[\text{C3}] + P[\text{C4}] + P[\text{C5}] + P[\text{C6}] \tag{12}$$

Note that the expressions for all cases except $P[\text{C1}]$ have the term $(S-1)^i$ in the denominator. For even moderate alphabet size S, (for example $S = 256$ for sequences of byte length) these probabilities thus become very small compared to $P[\text{C1}]$ and can be neglected to a good degree of approximation. Thus

$$\begin{aligned} P[w_{l' \to l}] &\approx P[\text{C1}] \\ &= \begin{pmatrix} n-l' \\ l-l' \end{pmatrix} \epsilon^{l-l'} (1-\epsilon)^{n-(l-l')} \end{aligned} \quad (13)$$

The time evolution of the l^{th} quasi-species group is thus governed by the the rate of up-mutations from quasi-species groups with lower Hamming distances. In each time unit, the concentration of the quasi-species group l' increases by $\beta A(l') w_{l'}(t)$ and a fraction $P[w_{l' \to l}]$ of these mutate to quasi-species group l. Summing up these contributions, the time evolution of $w_l(t)$ is then given by

$$\begin{aligned} \frac{dw_l(t)}{dt} &= \sum_{l'=0}^{l} P[w_{l' \to l}] \beta A(l') w_{l'}(t) \\ &= \sum_{l'=0}^{l} \begin{pmatrix} n-l' \\ l-l' \end{pmatrix} \beta A(l') \epsilon^{l-l'} (1-\epsilon)^{n-(l-l')} w_{l'}(t) \end{aligned}$$

which completes the proof of Claim 1. ∎

In the equation governing the growth and dynamics of each group of worm quasi-species as given in Claim 1, a key parameter is the fitness landscape $A(l')$ corresponding to each group. In the absence of any defense mechanism or known signatures, the fitness landscape for each group will be the same since all of them are equally likely to propagate without detection. We thus consider a scenario where initially all the quasi-species have an identical fitness landscape (we introduce the effect of the IDS later in this section)

$$A(y_l) = \eta \qquad \forall l \quad (14)$$

which implies

$$A(w_l) = \eta \qquad \forall l \quad (15)$$

In order to determine the optimal worm mutation rates and the maximum allowable response time of the IDS, we now obtain the concentrations of arbitrary members of different quasi-species groups. In the following, we assume that each quasi-species has the same initial concentrations, i.e.

$$y_l(0) = y_0(0) \qquad \forall l \quad (16)$$

Claim 2: The concentration of an arbitrary member of the 0^{th}, 1^{th} and n^{th} quasi-species group, represented by w_0^a, w_1^a and w_n^a respectively is given by

$$w_0^a(t) = w_0(0) e^{(1-\epsilon)^n \beta \eta t} \quad (17)$$

$$w_1^a(t) = w_0(0) \frac{e^{(1-\epsilon)^n \beta \eta t}}{n(S-1)} \left[1 + n\beta\eta\epsilon(1-\epsilon)^{n-1} t\right] \quad (18)$$

$$w_n^a(t) = w_0(0) e^{(1-\epsilon)^n \beta \eta t} \quad (19)$$

where $w_0(0) = y_0(0)$ is the initial concentration of the master quasi-species $w_0(t)$.

Proof: We consider each quasi-species group separately and solve Equation (4) to obtain the expressions above.

Quasi-species group 0: Note that there is only one quasi-species (the master quasi-species) which belongs to the 0^{th} quasi-species group and thus $w_0^a(t) = w_0(t)$. From Equation (4), substituting $l = 0$, we have

$$\begin{aligned} \frac{dw_0^a(t)}{dt} &= \frac{dw_0(t)}{dt} \\ &= A(w_0)(1-\epsilon)^n \beta w_0(t) \\ &= \eta(1-\epsilon)^n \beta w_0(t) \end{aligned} \quad (20)$$

Solving the ordinary differential equation above, we obtain

$$w_0^a(t) = w_0(0) e^{(1-\epsilon)^n \beta \eta t} \quad (21)$$

where $w_0(0) = w_0^a(0) = y_0(0)$, the initial concentration of the master quasi-species.

Quasi-species group 1: Each member of quasi-species group 1 has a Hamming distance of 1 from the master quasi-species. With each worm quasi-species consisting of n sequences and an alphabet of size S, there are $n(S-1)$ possible quasi-species which have a Hamming distance of 1 from the master quasi-species and thus form group 1. Thus the dynamics of the group ($w_1(t)$) is $n(S-1)$ times faster than an arbitrary quasi-species ($w_1^a(t)$) in the group. Substituting $l = 1$ in Equation (4) we then have

$$\begin{aligned} \frac{dw_1^a(t)}{dt} &= \frac{1}{n(S-1)} \frac{dw_1(t)}{dt} \\ &= \sum_{l'=0}^{1} \begin{pmatrix} n-l' \\ l-l' \end{pmatrix} \frac{\beta A(l') \epsilon^{l-l'} (1-\epsilon)^{n-(l-l')}}{n(S-1)} w_{l'}(t) \\ &= \frac{\eta\beta\epsilon(1-\epsilon)^{n-1}}{S-1} w_0(t) + \frac{\beta\eta(1-\epsilon)^n}{n(S-1)} w_1(t) \\ &= \frac{\eta\beta\epsilon(1-\epsilon)^{n-1}}{S-1} e^{(1-\epsilon)^n \beta\eta t} w_0(0) + \frac{\beta\eta(1-\epsilon)^n}{n(S-1)} w_1(t) \end{aligned}$$

Solving the differential equation above and using the fact that the initial concentrations of all the quasi-species are the same, i.e. $w_1^a(0) = w_0^a(0) = w_0(0)$, we obtain

$$w_1^a(t) = w_0(0) \frac{e^{(1-\epsilon)^n \beta\eta t}}{n(S-1)} \left[1 + n\beta\eta\epsilon(1-\epsilon)^{n-1} t\right] \quad (22)$$

Quasi-species group n: Consider an arbitrary member of quasi-species group n, w_n^a. Increase in its concentration results from the mutations from members of all other groups, other members of its own group as well as its own growth. These contributions can be written as

$$\begin{aligned} \frac{dw_n^a(t)}{dt} = &\sum_{i=0}^{n-1} A(w_i) \beta \left(\frac{\epsilon}{S-1}\right)^{n-i} (1-\epsilon)^i \frac{w_i(t)}{(S-1)^i} + \\ &\sum_{w_n^j \in \{w_n | w_n^j \neq w_n^a\}} A(w_n) \beta \left(\frac{\epsilon}{S-1}\right)^n w_n^j(t) + A(w_n)\beta(1-\epsilon)^n w_n^a(t) \end{aligned}$$

In the expression above we note that a mutation from a quasi-species of group i to the quasi-species w_n^a occurs only when: (1) the $n-i$ sequences that were not mutated so far mutate

to the corresponding sequences in w_n^a which happens with probability $(\epsilon/(S-1))^{n-i}$ and (2) the i sequences that had already mutated had mutated to the corresponding sequences in w_n^a, the probability of which is $1/(S-1)^i$. Similarly, other members of the quasi-species group n mutate to w_n^a only if all of their sequences mutate to the corresponding sequence in w_n^a, an event which occurs with probability $(\epsilon/(S-1))^n$. Finally, members of the quasi-species w_n^a continue to replicate their own if no mutation occurs, i.e. with probability $(1-\epsilon)^n$. Note that in the expression above, all terms except for the last have $(S-1)^n$ in the denominator. For even small alphabet size S and number of sequences n, the contribution of these terms becomes quite small and we can thus write

$$\begin{aligned}\frac{dw_n^a(t)}{dt} &\approx A(w_n)\beta(1-\epsilon)^n w_n^a(t) \\ &= \beta\eta(1-\epsilon)^n w_n^a(t) \end{aligned} \quad (23)$$

Another way to interpret the expression above is that for quasi-species far removed from the master sequence, mutations in and out of the quasi-species cancel and its growth is determined by its own environmental or uninhibited growth rate of $\beta\eta(1-\epsilon)^n$. Solving the differential equation defined in Equation (23) we obtain

$$w_n^a(t) = w_n^a(0)e^{\beta\eta(1-\epsilon)^n t} \quad (24)$$

This completes the proof of Claim 2. ■

B. Maximum Allowable IDS Response Time

With the expressions developed for the evolution of the concentration of the various worm quasi-species, we now derive the maximum allowable IDS response time in order to contain the worm. In the following analysis, we assume that the IDS is first detects the quasi-species with the highest concentration and we call the quasi-species which currently has the highest concentration the *dominant* quasi-species. Without loss of generality, we can assume that the quasi-species detected first contains the master sequence, and thus we call it the master quasi-species. Based on the detected quasi-species' code sequence, the IDS now tries to generate signatures for other morphed versions of the worm. If the IDS cannot detect other worm quasi-species fast enough, it will not be able to contain their spread. Our interest in this section is to obtain the maximum allowable difference between the time of detection of the worm quasi-species which currently has the highest concentration and the detection time of the quasi-species with the next highest concentration. We show that if this detection time is below a certain threshold, successive peaks of the currently dominant quasi-species will be smaller, eventually leading to the worm's elimination.

We now consider the interaction between the polymorphic worm and the IDS in greater detail. Once the IDS detects the quasi-species with the highest concentration, it applies a large kill rate to that specific worm quasi-species. Thus the value of its fitness function decreases and its effective growth rate drops. Until the IDS is capable of detecting the other strains of the polymorphic worm, other worm quasi-species will now have a higher growth rate as compared to the earlier dominant strain. In other words, the fitness landscape for the worm moves. From Equations (17), (18) and (19) we note that a member of the quasi-species group 1 will now have the next highest growth rate. The polymorphic worm adapts to this change in the fitness function on a timescale τ_w, the time required for the population of the new dominant strain to overtake that of the old. The IDS system now tries to adapt to this new dominant quasi-species and we denote the time required for it to detect the new dominant quasi-species by τ_{ids}. Through the iteration of the steps above, the worm scours through the sample space of quasi-species and the IDS system follows on its heels. Note that the iteration repeats every $\tau = \tau_w + \tau_{ids}$ seconds.

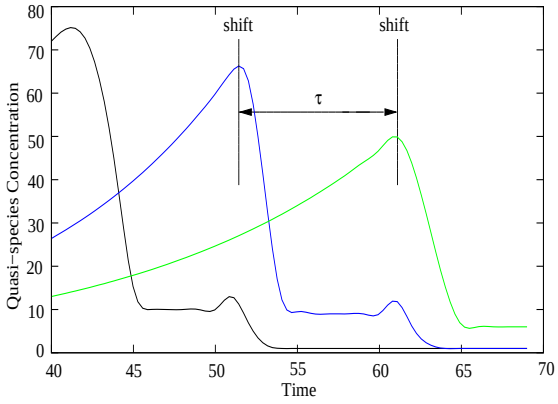

Fig. 3. An example of the time evolution of three worm quasi-species for a case where the IDS is capable of containing the worm's spread. At each shift a new quasi-species becomes dominant but is unable to attain the concentration of the previous dominant quasi-species before it is detected.

If τ_{ids} is large, the IDS moves slowly and the new dominant strain reaches or exceeds the concentration of the previous dominant strain and the worm will survive for all time. If however the new dominant strain cannot regenerate fast enough (or equivalently, the IDS is fast enough) the new dominant quasi-species will be detected and consequently a shift in the worm's fitness landscape will occur before it reaches the population level of the previous dominant strain. The next dominant quasi-species will not be able to reach the concentration level of the previous dominant quasi-species and so on. This continues till there is no discernible worm presence in the network. This pattern of chaining quasi-species concentrations is shown in Figure 3.

We now determine the required growth rate of the new dominant quasi-species over a full cycle τ as a criterion for the quasi-species' survival. Consider time $t = 0$ at the instant when the fitness landscape shifts and an arbitrary quasi-species from the 1^{st} quasi-species group becomes dominant. Using Equation (18) the normalized growth of this quasi-species over a period τ is then given by

$$\frac{w_1^a(\tau)}{w_1^a(0)} = \frac{e^{(1-\epsilon)^n\beta\eta\tau}}{n(S-1)}\left[1 + n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau\right] \quad (25)$$

The quasi-species w_1^a is currently the fittest. But if another

quasi-species far away from the fitness peak is able to surpass its population in the interval τ, then the currently dominant quasi-species will die out. With the current dominant quasi-species being from the quasi-species group 1 and the quasi-species group 0 having already been detected, we turn to an arbitrary member of the quasi-species group n, w_n^a, since it has the largest Hamming distance from the master sequence. Using Equation (19), its normalized growth rate is then given by

$$\frac{w_n^a(\tau)}{w_n^a(0)} = e^{\beta\eta(1-\epsilon)^n\tau} \tag{26}$$

The ratio of these growth rates is then

$$k = \frac{\frac{w_1^a(\tau)}{w_1^a(0)}}{\frac{w_n^a(\tau)}{w_n^a(0)}} = \frac{1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}{n(S-1)} \tag{27}$$

The quasi-species will die out if $k < 1$ and survives only when $k \geq 1$.

1) Obtaining τ_w and τ_{ids}: In order to obtain the maximum allowable limit on τ_{ids}, we first obtain τ_w and substitute it in Equation (27). Consider the situation when the IDS system detects the original dominant quasi-species and exerts a kill rate of δ on it. To estimate the timescale for the shift in the worm's fitness landscape, τ_w, we first iterate the quasi-species propagation model for a full cycle of length τ starting at $t = 0$. The switch in the dominant quasi-species is made at $t = \tau$ when the IDS starts applying the decay rate of δ on the previous dominant quasi-species. The relative sizes of the new and old dominant quasi-species are then determined for another interval τ_w and the timescale τ_w is given by the waiting time until the new dominant quasi-species population exceeds the old one. Note that the populations of the old and new dominant quasi-species as the end of τ are given by $w_0^a(\tau)$ and $w_1^a(\tau)$ in Equations (17) and (18) respectively. In the subsequent interval τ_w, the growth rates of the old and new dominant quasi-species are $e^{(1-\epsilon)^n\eta\beta-\delta}$ and $e^{(1-\epsilon)^n\beta\eta}$ respectively. Equating the populations of the old and new dominant quasi-species at τ_w and using $w_0^a(0) = w_1^a(0)$ i.e. equal initial conditions, we obtain

$$\begin{aligned}
& e^{(1-\epsilon)^n\beta\eta\tau_w} w_1^a(\tau) = e^{(1-\epsilon)^n(\beta\eta-\delta)\tau_w} w_0^a(\tau) \\
\Rightarrow \quad & e^{(1-\epsilon)^n\beta\eta\tau_w} e^{(1-\epsilon)^n\beta\eta\tau} \frac{1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}{n(S-1)} w_0(0) \\
& \qquad = e^{(1-\epsilon)^n(\beta\eta-\delta)\tau_w} e^{(1-\epsilon)^n\beta\sigma\tau} w_0(0) \\
\Rightarrow \quad & e^{-\delta\tau_w} = \frac{1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}{n(S-1)} \\
\Rightarrow \quad & \tau_w = \frac{1}{\delta}\ln\left(\frac{n(S-1)}{1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}\right)
\end{aligned} \tag{28}$$

where we need $n(S-1) > 1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau$ for τ_w to be positive.

To find the maximum allowable τ_{ids}, we note that the worm population dies out and its spread is contained if $k < 1$ in Equation (27). Thus by substituting $\tau = \tau_w + \tau_{ids}$ in Equation (27) and solving for τ_{ids} we can obtain the maximum allowable IDS response time. Using Equation (28), we can write τ as

$$\begin{aligned}
\tau &= \tau_w + \tau_{ids} \\
&= \frac{1}{\delta}\ln\left(\frac{n(S-1)}{1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}\right) + \tau_{ids} \\
&= \frac{1}{\delta}\text{LamW}(x) - \frac{1}{n\beta\eta\epsilon(1-\epsilon)^{n-1}}
\end{aligned} \tag{29}$$

where x is given by

$$x = \frac{\delta(S-1)}{\beta\eta\epsilon(1-\epsilon)^{n-1}} e^{\frac{\delta(1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau_{ids})}{n\beta\eta\epsilon(1-\epsilon)^{n-1}}} \tag{30}$$

and $\text{LamW}(\cdot)$ is the Lambert W function, i.e., a function which satisfies

$$\text{LamW}(y)e^{\text{LamW}(y)} = y \tag{31}$$

Now, solving $k < 1$ using Equation (27) for τ, we obtain

$$\begin{aligned}
& k = \frac{1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}{n(S-1)} < 1 \\
\Rightarrow \quad & \tau < \frac{n(S-1)-1}{n\beta\eta\epsilon(1-\epsilon)^{n-1}}
\end{aligned}$$

Substituting the value of τ in terms of τ_{ids} from Equations (29) and (30) in the expression above, we have

$$\frac{1}{\delta}\text{LamW}(x) < \frac{1}{n\beta\eta\epsilon(1-\epsilon)^{n-1}} + \frac{n(S-1)-1}{n\beta\eta\epsilon(1-\epsilon)^{n-1}} \tag{32}$$

which can be solved for τ_{ids} to obtain the maximum allowable IDS response time in order to contain the worm and this solution is given by

$$\tau_{ids} < \frac{n(S-1)-1}{n\beta\eta\epsilon(1-\epsilon)^{n-1}} \tag{33}$$

C. Optimal Worm Mutation Rates

We now further analyze the quasi-species model to determine the optimal mutation rate of the polymorphic worm. Given an estimate of the response time τ_{ids} of the IDS, mutations at a rate greater than the optimal mutation rate guarantees that the IDS will not be able to contain the worm's proliferation and the spread will reach epidemic proportions. With the ratio k in Equation (27) governing the survivability of the worm, optimizing the worm's mutation rate requires solving for

$$\frac{\partial k}{\partial \epsilon} = 0 \tag{34}$$

which can be written as

$$\begin{aligned}
\frac{\partial k}{\partial \epsilon} &= \frac{\partial}{\partial \epsilon}\frac{1+n\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}{n(S-1)} \\
&= \frac{\partial}{\partial \epsilon}\frac{\beta\eta\epsilon(1-\epsilon)^{n-1}\tau}{S-1}
\end{aligned}$$

Substituting the value of τ from Equations (28) and (30) into the equation above, we have

$$0 = \frac{\partial}{\partial \epsilon}\frac{\beta\eta\epsilon(1-\epsilon)^{n-1}}{S-1}\left[\frac{1}{\delta}\text{LamW}(x) - \frac{1}{n\beta\eta\epsilon(1-\epsilon)^{n-1}}\right] \tag{35}$$

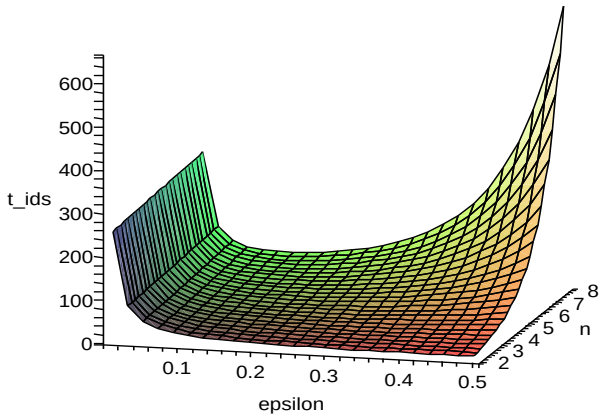

Fig. 4. τ_{ids} versus ϵ and n. The parameters used are $\eta = 1.0$, $\beta = 100$, $\delta = 100$ and $S = 256$.

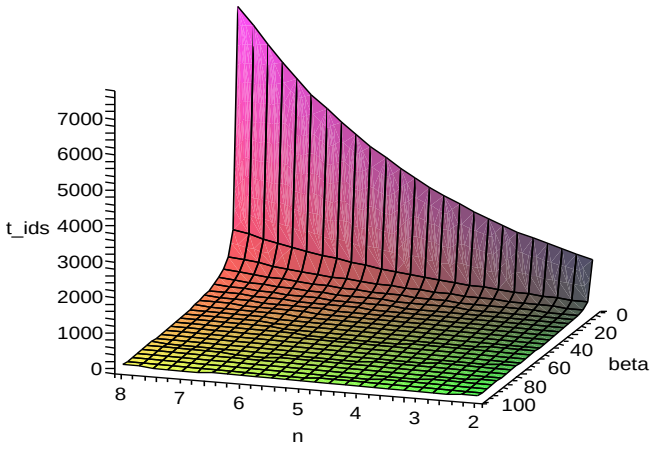

Fig. 5. τ_{ids} versus β and n. The parameters used are $\eta = 1.0$, $\epsilon = 0.25$, $\delta = 100$ and $S = 256$.

After some algebraic manipulations of the equation above we have the optimal mutation rate ϵ_{opt} as the solution of the following equation for ϵ

$$\big[1 + \text{LamW}(x)\big]x = \frac{\epsilon(1-\epsilon)}{n\epsilon - 1} \tag{36}$$

VI. Model Evaluation

In this section we evaluate the results of the previous section to draw insights into the behavior of the complex system characterized by the jointly evolving polymorphic worm and the IDS. We first consider the limit on the response time of the IDS and the effect of various parameters on its value and then consider the optimal worm mutation rates.

A. Maximum Allowable IDS Response Time

In Figure 4 we show τ_{ids} as a function of the mutation rate ϵ and n for $\eta = 1.0$, $\beta = 100$, $\delta = 100$ and $S = 256$. We note that as expected, when the mutation rate decreases, the maximum allowable IDS response time increases. This is because with lower ϵ the IDS has more time to generate the signatures since now the rate of appearance of new quasi-species is lower. Also, τ_{ids} increases as n increases since larger signatures allow for more accurate signatures and thus better chances of detecting the worm. Also, note that for low values of ϵ, n does not have an appreciable impact on the value of τ_{ids}.

In Figure 5 we show τ_{ids} as a function of the worm growth rate β and n for $\eta = 1.0$, $\epsilon = 0.25$, $\delta = 100$ and $S = 256$. We note that as the worm growth rate increases, the the maximum allowable IDS response time decreases exponentially. As in the previous figure, a larger signature or worm length n leads to higher allowable values for τ_{ids}. We also note that as the worm growth rate increases, the signature length has very little impact on the maximum allowable value of τ_{ids}.

In Figure 6 we show τ_{ids} as a function of ϵ and β, $\eta = 1.0$, $n = 6$, $\delta = 100$ and $S = 256$. We observe that at high values

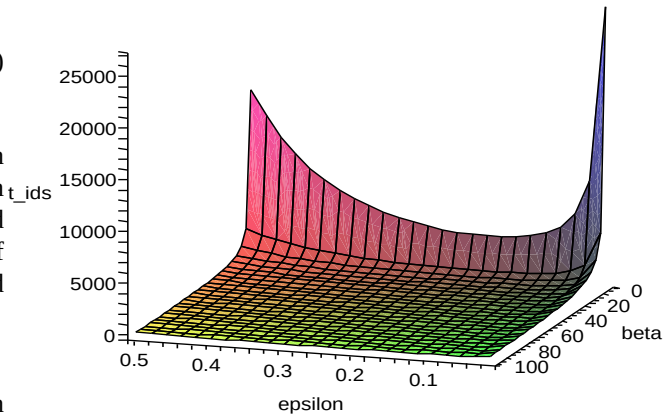

Fig. 6. τ_{ids} versus ϵ and β. The parameters used are $\eta = 1.0$, $\beta = 100$, $\delta = 100$ and $S = 256$.

of the growth rate β the mutation rate does not affect the IDS response time appreciably. However, at small growth rates the allowable IDS response time increases and ϵ has a considerable impact on the value of τ_{ids}.

B. Worm Survivability and Optimal Mutation Rates

In this section we observe the effect of various parameters on the optimal worm mutation rates and the worm's survivability. To show the existence of an optimum and the region around it, in Figure 7 we plot k given by Equation (27) as a function of of the mutation rate ϵ for $\eta = 1.0$, $\beta = 150$, $\delta = 150$, $n = 6$, $S = 256$ and $\tau_{ids} = 40.0$. Note that values of $k \leq 1$ implies that the worm is able to evade the IDS and able to survive at all times. From Figure 7 it is clear that there is an optimal mutation rate for a given set of parameters.

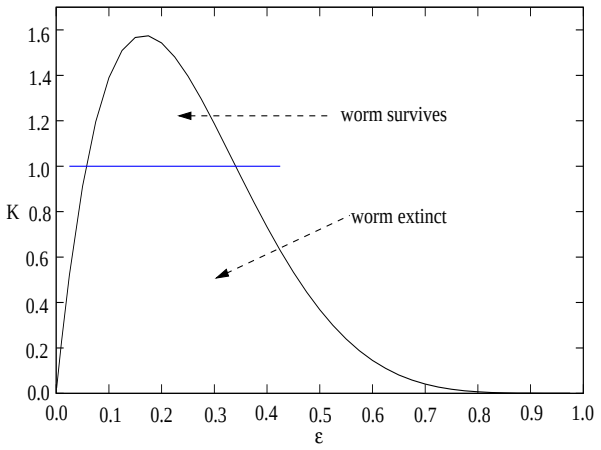

Fig. 7. K versus ϵ showing the region where the worm quasi-species survive and the optimal mutation rate. The parameters used are $\eta = 1.0$, $\beta = 150$, $\delta = 150$, $n = 6$, $S = 256$ and $\tau_{ids} = 40.0$.

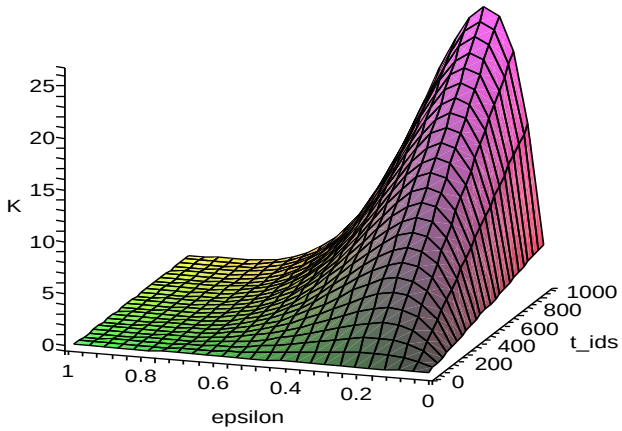

Fig. 9. k versus ϵ and τ_{ids}. The parameters used are $\eta = 1.0$, $\beta = 100$, $\delta = 100$, $n = 6$ and $S = 256$.

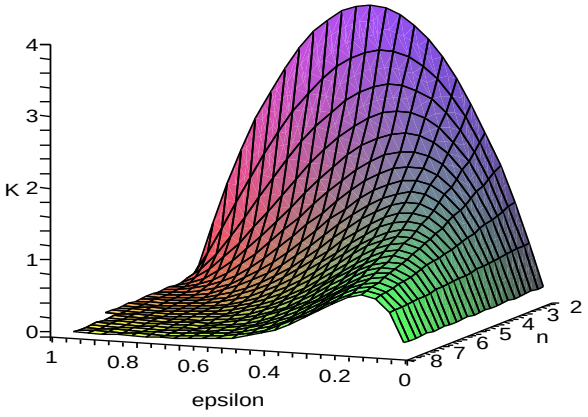

Fig. 8. k versus ϵ and n. The parameters used are $\eta = 1.0$, $\beta = 100$, $\delta = 100$, $\tau_{ids} = 40$ and $S = 256$.

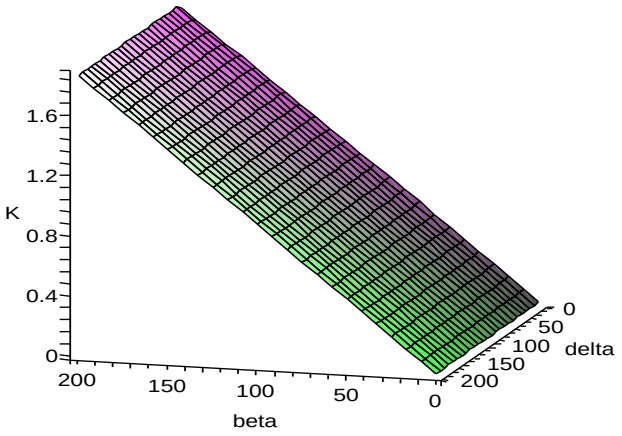

Fig. 10. k versus β and δ. The parameters used are $\eta = 1.0$, $\epsilon = 0.25$, $\tau_{ids} = 40$, $n = 6$ and $S = 256$.

In Figure 8 we plot k as a function of ϵ and n for $\eta = 1.0$, $\beta = 100$, $\delta = 100$, $\tau_{ids} = 40$ and $S = 256$. We note that as n increases, the value of k decreases. This is because as the signature length increases, the IDS becomes better at detecting the worm quasi-species. We note that a higher mutation rate does not necessarily mean greater survivability for the quasi-species.

In Figure 9 we plot k as a function of ϵ and the IDS response time τ_{ids} for $\eta = 1.0$, $\beta = 100$, $\delta = 100$, $n = 6$ and $S = 256$. We note that as τ_{ids} increases, the values of k increases since now the worm quasi-species have more time to grow before the IDS acts on them. This is because as the signature length increases, the IDS becomes better at detecting the worm quasi-species. Also, the optimal ϵ changes very little with changes in τ_{ids}.

In Figure 10 we plot k as a function of β and the IDS kill rate δ for $\eta = 1.0$, $\epsilon = 0.25$, $\tau_{ids} = 40$ and $S = 256$. We note that as the worm growth rate β increases, the worm quasi-species is more likely to survive and k increases linearly with β. On the other hand k stays almost constant with δ.

VII. Simulation Results

In this section we present the simulation results to qualitatively validate the proposed model for polymorphic worms. We first present the details of the simulation environment before presenting the actual simulation results.

A. Worm Generator

To simulate polymorphic worm propagation and IDS response, we developed a simulator which used the ADMmutate polymorphic engine to produce variants of exploit code with a total size of 640 bytes. We used the same growth model as assumed in Section V, with the number of active worms

doubling with each time step, i.e., with $\beta = 1.0$. We begin by generating a worm whose signature is used as the master sequence, and then begin replicating with growth rate β = 1.0. Each generated worm is then sent to the IDS simulator for processing.

B. IDS Response

The IDS simulator takes one parameter, $\hat{\tau}_{ids}$, which determines the time required by the IDS to generate the signature for the currently dominant quasi-species. As a worm is received from the worm generator, the worm is classified based on its Hamming distance from the master sequence's code sequence. The Hamming distance indicates which quasi-species group the worm belongs to and the corresponding concentration is incremented. After all polymorphic worms from the current time step have been processed, the IDS checks to see if $\hat{\tau}_{ids}$ time units have elapsed. If so, a kill rate of $\delta = 1.0$ is applied to the quasi-species *group* with the highest concentration. That is, no more worms of that group are allowed to pass the IDS and propagate in the network.

C. Simulation Results

The objective of the simulations is to verify our assumptions about applying the quasi-species model to polymorphic worm propagation and containment. Specifically, we wish to demonstrate the effect of $\hat{\tau}_{ids}$ and β on the concentration of each worm quasi-species where $\hat{\tau}_{ids}$ is the time required by the IDS to generate the signature for the currently dominant quasi-species and n_0 is the initial size of the worm pool. For convenience, we will refer to these parameters using the form $\{\hat{\tau}_{ids}, n_0\}$.

Our simulations used the following values: S = 256 (we look at byte-long sequences); n = 4 (signature bytes); ϵ = 0.52 (determined experimentally for the 4 signature bytes in Section IV); β = 1.0 (exponential growth rate).

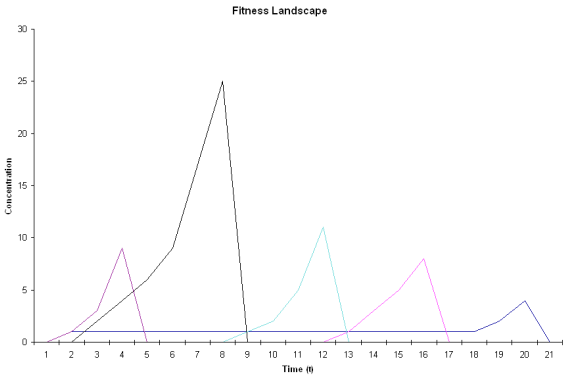

Fig. 11. Simulation with parameters {4, 1}. The choice of $\hat{\tau}_{ids}$ is sufficiently small to contain the worm.

The graph in Figure 11 shows the result of a simulation with {4, 1} and plots the concentration of the 5 quasi-species groups as a function of time. This plot matches the expected result from the quasi-species model. For $\hat{\tau}_{ids} = 4$, the IDS is sufficiently effective to generate the signature for the dominant quasi-species and thus control their propagation. The peak concentrations of the successive dominant quasi-species die out until all 5 quasi species are eliminated from the network. It is interesting to note that the highest peak concentration occurs during the second interval, not the first. This is due to the fact that the IDS misidentified the most dominant quasi-species in the first interval because of the small worm population and the initial transients in the systems. However, it can be seen that it does identify the dominant strain and apply the kill rate after $2\hat{\tau}_{ids}$. This was common among many of our simulations with small n_0. After the initial transient, we see that the concentration of each successive dominant strain never exceeds the previous dominant strain's peak. Note that even though we plot and apply the kill rates on a quasi-species group rather than individual quasi-species for illustrative purposes, the validity of the results is not affected since this is a special case of our model.

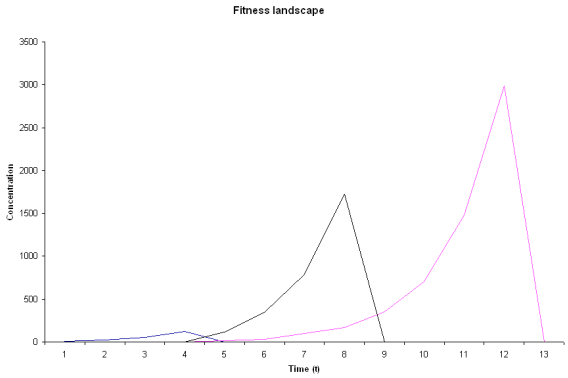

Fig. 12. Simulation with parameters {4, 8}. We can see how the size of the initial population has made the IDS ineffective for the same $\hat{\tau}_{ids}$.

In Figure 12 we can see the effect of a larger initial worm population. Choosing the parameters {4, 8}, we cannot contain the worm. The exponential growth of each quasi-species group, which begins with a population of 8 worms, overwhelms the IDS during each interval $\hat{\tau}_{ids}$ and the worm continues to spread. In practical terms, even though the IDS will apply the kill rate δ after each interval, the propagation rate is high enough to ensure the survival of the worm. For example, suppose that the signature of the worm consists of n bytes. It will take the IDS $(n+1)\hat{\tau}_{ids}$ time steps to eliminate all quasi-species groups. If the peak concentration of the new dominant quasi-species group grows at a rate γ, the last quasi-species group will attain a peak concentration of

$$n_0 e^{\gamma \hat{\tau}_{ids}(n+1)} \tag{37}$$

For the values $n_0 = 8$, $\gamma = 1.0$, $\hat{\tau}_{ids} = 4$, $n = 4$, the peak concentration of the final quasi species is 3.88 billion. This would be enough to infect the entire Internet with just the last iteration. The worm would achieve a total concentration of over 7.5 billion, clearly indicating successful spread under these conditions.

Other simulations were performed with varying values for n, n_0, $\hat{\tau}$, and ϵ, all supporting the claims made in Section VI.

VIII. Conclusions

With their ability to change appearance at every replication step, polymorphic worms have the ability to evade existing IDSes and have the potential to cause large scale damage. In this paper we presented a framework based on coevolution of quasi-species to model the dynamics of polymorphic worms and their interaction with an IDS that has the capability to evolve with time. The model provides a theoretical basis to explain polymorphic worm outbreaks and the limitations they put on the IDS and thereby aid the development of defense strategies which can respond effectively to such outbreaks.

The model is used to obtain the conditions that govern the survivability of the worm and we obtain closed form expressions for the maximum allowable response time of the IDS in order to contain the spread of the worm. We also obtain the optimal mutation rate that a polymorphic worm may employ in order to evade an IDS with a known response time. The analytic results are qualitatively validated using simulation results. The results of the paper provide valuable insights and design guidelines for the development of realistic IDSes capable of containing polymorphic worms.

References

[1] Z. Chen, L. Gao and K. Kwait, "Modeling the spread of active worms," *Proceedings of IEEE INFOCOM*, San Francisco, CA, March 2003.

[2] C. Zou, W. Gong and D. Towsley, "Code red worm propagation modeling and analysis," *Proceedings of ACM Conference on Computer and Communication security*, Washington, DC, October 2002.

[3] C. Zou, W. Gong and D. Towsley, "Worm propagation modeling and analysis under dynamic quarantine defense," *Proceedings of ACM Workshop on Rapid Malcode*, pp. 51-60, Washington, DC, November 2003.

[4] J. Kephart and W. White, "Direct-graph epidemiological models of computer viruses," *Proceedings of IEEE Symposium on Security and Privacy*, pp. 343-361, Oakland, CA, May 1991.

[5] D. Moore, C. Shannon and K. Claffy, "Code-Red: A case study on the spread and victims of an Internet worm," *Proceedings of Internet Measurement Workshop*, pp. 273-284, Marseille, France, November 2002.

[6] D. Moore, C. Shannon, G. Voelker and S. Savage, "Internet quarantine: Requirements for containing self-propagating code," *Proceedings of IEEE INFOCOM*, San Francisco, CA, March 2003.

[7] C. Nachenberg, "Computer virus-antivirus coevolution," *Communications of the ACM*, vol. 40. no. 1, pp. 46-51, January 1997.

[8] G. Vigna, W. Robertson and D. Balzarotti, "Testing network based intrusion detection signatures using mutant exploits," *Proceedings of ACM Conference on Computer and Communication Security*, Washington, DC, October 2004.

[9] M. Nilsson and N. Snoad, "Error thresholds for quasispecies on dynamic fitness landscapes," *Physical Review Letters*, vol. 84, no. 1, January 2000.

[10] C. Kamp and S. Bornholdt, "Coevolution of quasispecies: B-cell mutation rates maximize viral error catastrophes," *Physical Review Letters*, vol. 88, no. 6, February 2002.

[11] Lurhq Threat Intelligent Group, "Phatbot trojan analysis." `http://www.lurhq.com/phatbot/html`.

[12] T. Detristan, T. Ulenspiegel, Y. Malcom, and M. V. Underduk "Polymorphic shellcode engine using spectrum analysis," `http://www.phrack.org/show.php?p=61&a=9`.

[13] K2, ADMmutate. `http://www.ktwo.ca/c/ADMmutate-0.8.4.tar.gz`.

[14] Aleph One, Smashing The Stack For Fun And Profit. `http://www.phrack.org/phrack/49/P49-14`.

[15] E. Tannenbaum and E. Shakhnovich, "Solution of the quasispecies model for an arbitrary gene network," *Physical Review E*, vol. 70, 2004.

[16] J. Newsome, B. Karp and D. Song, "Polygraph: Automatically Generating Signatures for Polymorphic Worms," In *Proc. IEEE Symposium on Security and Privacy*, Oakland, CA, May 2005.

[17] Y. Tang and S. Chen. "Defending Against Internet Worms: A Signature-Based Approach," In *Proc. of IEEE INFOCOM'05*, Miami, Florida, USA, May 2005.

[18] O. Kolesnikov and W. Lee. "Advanced Polymorphic Worms: Evading IDS by Blending in with Normal Traffic," `http://www.cc.gatech.edu/ ok/w/ok_pw.pdf`.

[19] 1990, `http://www.viruslist.com/en/viruses/encyclopedia?chapter=153311162`.

[20] M. Eigen and P. Schuster, *the Hypercycle - A Principle of Natural Self-Organization*, Springer-Verlag, Berlin, 1979.