Department of
Information Engineering
and Computer Science DISI

UNIVERSITY
OF TRENTO - Italy

Compartmental differential equations models of botnets and epidemic malware

Marco Ajelli, Renato Lo Cigno, Alberto Montresor

June 2010

Technical Report # DISI-10-011

Compartmental differential equations models of botnets and epidemic malware

Marco Ajelli
Fondazione Bruno Kessler, Trento, Italy
and DISI, University of Trento, Italy
Email: ajelli@fbk.eu

Renato Lo Cigno, Alberto Montresor
DISI, University of Trento, Italy
Email: {Renato.locigno,alberto.montresor}@disi.unitn.it

Abstract—**Botnets, i.e., large systems of controlled agents, have become the most sophisticated and dangerous way of spreading malware. Their damaging actions can range from massive dispatching of e-mail spam messages, to denial of service attacks, to collection of private and sensitive information. Unlike standard computer viruses or worms, botnets spread silently without actively operating their damaging activity, and then are activated in a coordinated way to maximize the "benefit" of the malware. In this paper we propose two models based on compartmental differential equations derived from "standard" models in biological disease spreading. These models offer insight into the general behavior of botnets, allowing both the optimal tuning of botnets' characteristics, and possible countermeasures to prevent them. We analyze, in closed form, some simple instances of the models whose parameters have non-ambiguous interpretation. We conclude the paper by discussing possible model extensions, which can be used to fine-tune the analysis of specific epidemic malware in the case that some parameters can be obtained from actual measurements of the botnet behavior.**

I. Introduction

A *botnet* is jargon terminology for a collection of infected end-hosts, called *bots*, that are under the control of a human operator known as *botmaster*. While originally the term was coined for legitimate purposes, such as the automated management of IRC channels through distributed scripting, botnets are now inherently linked to malicious activities. They have recently been identified as one of the most important threats to the security of the Internet [6], [18].

Lifecycle of a botnet can be briefly described as follows [23]. First, the botnet creator sends out a virus, or worm, that infects unprotected (or under-protected) machines over the Internet. Several infection strategies exist, mostly borrowed from other classes of malware (e.g., e-mail viruses). The worm payload is the bot itself, a malicious application that logs into a centralized *command-and-control* server (C&C). Bots may remain hidden for a long period, waiting for commands from the C&C. When one of such commands is received, the bot autonomously perform the malicious actions for which it has been instructed. Eventually, the bot code may be discovered and removed, e.g. by an anti-virus software, bringing back the end-host to a clean state. While installed the bot code can switch freely between the hidden and the active state.

Example of attacks performed by botnets include sending of spam and phishing messages [5], click fraud against major search engines [8], blackmail based on denial-of-service [14], and identity theft.

To improve their ability to survive, recent botnets have tried to avoid the centralized point of failure represented by the C&C in favor of modern peer-to-peer architectures [12]. The anti-virus community has responded with more sophisticated mechanisms to detect and disrupt botnets, in an eternal arms race between good and evil.

We believe that an appropriate modeling of the botnet threat is a necessary condition to play this "cops and robber" game. We are interested in modeling two interconnected phases of the lifecycle of botnet: the *creation phase*, which include both the growth and the maintenance of the botnet network, and the *activity phase*, when botnets are actually used by their botmasters for their malicious purposes. We analyze these phenomena from a theoretical, high-level point of view, and we simplify the model by considering only the most common malicious activity: sending spam messages. From now on, the term 'spam' is used as a synonym of maliciousness; in reality, any malicious actions can fit in this model, from illicitly using the CPU, to re–distribute illegal material, to scan the local bulk memory in search of private/sensitive/important information.

Mathematical models have been already proposed to better understand both the propagation of Internet worms and dissemination protocols [2]–[4], [7], [26], [29] and the evaluation of intervention options in response to the propagation itself [20], [30]. The aim of this work is to propose the first mathematical representation of both the botnet growth and the activity phase, in order to understand what are the conditions supporting and what are the ones limiting these processes.

In this paper we show that, by changing the proportion of active bots in the botnet, botmasters are able to build a botnet also if the worm is not very transmissible, maximize the amount of damage performed, and generate multiple waves of malicious activity without the need of creating a new botnet.

From a mathematical point of view, we adopt the classical differential equation approach used to model biological diseases [1], [9], [27], which allows both closed form theoretical results and numerical solutions. In fact, biological diseases and botnets present surprising similarities: both spread using an infection mechanism based on contact (or communication); once infected, both may stay inactive for a long period, still remaining infectious to other members of the population; and eventually, both can be cured returning to a non-susceptible, non-infectious, inactive state.

The paper is structured as follows. Section II introduces the two botnet models discussed in the following. Section III analyzes the theoretical properties of the two models, and Section IV provides some insights on the dynamical behavior, both analytically and experimentally. Section V discusses related work, while Section VI introduces extensions to these models that take into account different and more sophisticated behaviors. Section VII concludes the work.

II. The Models

We take inspiration from years of modeling human diseases and natural epidemic phenomena to build a class of models that we deem help understanding modern worms and botnets in particular. We consider two different cases: the first one assumes that nodes or agents[1], after infection will be recovered, and recovered nodes cannot be infected anymore, i.e., they are immune forever as if vaccinated. This scenario corresponds e.g. to a botnet that exploits a single vulnerability, and this vulnerability is fixed once and for all. The second one assumes instead that agents can be re-infected. This latter case is suitable for worms and malware for which recovery is not definitive. Examples of this behavior include botnets that mutate during their lifecycle, exploiting different vulnerabilities at different times (with some attention in defining the infection rate of different mutations); but it can be even happen in machines that are re-installed and become vulnerable again, maybe only for a finite period of time [24], [25]. Our study is based on compartmental ordinary differential equations models. Moreover, we assume a finite population, which can be arbitrarily large, since our models are scale–invariant with respect to the number of nodes.

A. Botnets Subject to Immunization

We consider four classes of agents.

Class S: *susceptible* nodes with a positive risk of infection;
Class I: *infectious* nodes able to infect the susceptibles;
Class V: *spamming* nodes able to infect the susceptibles and actively executing their specific malware;
Class R: *removed* nodes that have been cured (de–infected) and are immune to the worm.

We call these botnets I-Botnets since the nodes are *immunized* against the worm after recovery.

The epidemic flow among these classes is graphically represented in Figure 1 with the notation we use after the normalization with respect to the parameter μ and N (see Table I for the parameters synopsis). It can be described as follows.

A susceptible can become an infectious node of the botnet I at the rate $b\frac{I+V}{N}$, where b is the transmission rate of the worm ($b > 0$) and N is the total population ($N > 0$).

TABLE I
Model variables and parameters

Notation	Description
N	total population (nodes in the system)
μ	switching rate between hidden and active
s, S	proportion/number of at–risk–of–infection nodes
i, I	proportion/number of (hidden) infectious nodes
v, V	proportion/number of infectious and spamming nodes
R	number of definitely recovered nodes
β, b	normalized, absolute worm transmission rate
γ, g	normalized, absolute (definitive) recovery rate
p	apportioning coefficient of infected hidden nodes
ρ	rate of temporary recovery

Infectious nodes can start or stop spamming. Since worms building botnets are normally entirely quiescent, and consequently very hard to detect, if they do not spam, we assume that only spamming nodes can be detected and removed.

An infectious node can become a spamming node V with rate $\mu/(1-p)$, where $\mu > 0$ is the speed of the interchange between classes V and I; a spamming node stops spamming and go back to the infectious status I with rate μ/p or can be removed at rate $g > 0$; p (0¡p¡1) is an apportioning parameter describing the aggressiveness of the botnet (smaller p means more aggressive botnets that tend to spam continuously as $p \to 0$), p is also the fraction of infected nodes in class I.

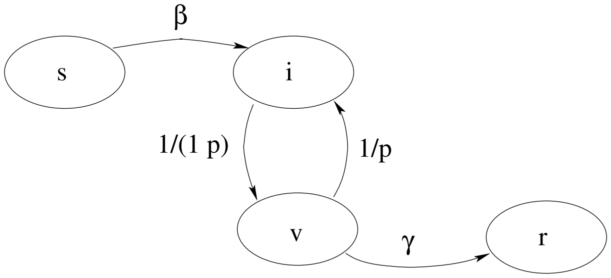

Fig. 1. Normalized epidemic flow among agents classes for non re-infecting botnets

The model described above can be formalized with the following system of differential equations:

$$\begin{cases} \dot{S}(\tau) &= -b\frac{I(\tau)+V(\tau)}{N}S(\tau) \\ \dot{I}(\tau) &= b\frac{I(\tau)+V(\tau)}{N}S(\tau) - \frac{\mu}{1-p}I(\tau) + \frac{\mu}{p}V(\tau) \\ \dot{V}(\tau) &= \frac{\mu}{1-p}I(\tau) - \left(\frac{\mu}{p} + g\right)V(\tau) \\ \dot{R}(\tau) &= gV(\tau) \end{cases} \quad (1)$$

a) Normalization: Since we consider a closed population the relation $N = S(\tau)+I(\tau)+V(\tau)+R(\tau)\ \forall\tau$ holds. Thus we can consider only three differential equations. Moreover, we can rewrite Equation (1) in an adimensional form, by defining $t := \mu\tau$, $s := S/N$, $i := I/N$, $v := V/N$, $\beta := b/\mu$ and $\gamma := g/\mu$. Therefore, we have:

$$\begin{cases} \dot{s}(t) &= -\beta\left[i(t)+v(t)\right]s(t) \\ \dot{i}(t) &= \beta\left[i(t)+v(t)\right]s(t) - \frac{1}{1-p}i(t) + \frac{1}{p}v(t) \\ \dot{v}(t) &= \frac{1}{1-p}i(t) - \left(\frac{1}{p} + \gamma\right)v(t) \end{cases} \quad (2)$$

[1] We use the terms “node” and “agent” interchangeably, indicating anything that may be subject to the infection and become part of the botnet. For instance an agent can be a PC or a single virtual machine. We instead reserve the terms ‘bot’ and ‘bot code’ to the actual infectious malware when we discuss examples.

as graphically described in Figure 1. For a complete definition of Equation (2) we have to assign the initial conditions. Here we assume the following initial conditions:

- $s(0) = s_0$, where $0 < s_0 \leq 1$; note that $s_0 = 0$ means that the population cannot be infected by the worms and epidemics cannot occur, while $s_0 = 1$ means that the population is completely susceptible to the worm;
- $v(0) = 0$, meaning that there are no active spamming node at the beginning;
- $i(0) = \epsilon_0$, where ϵ_0 is typically of the order of 10^{-N}, and it is anyway included in the range $]0, 1[$; note that $i(0) = 0$ and $v(0) = 0$ means that are no infectious agents in the population; therefore epidemics cannot occur.

b) Remark: Let us note that, if $p \to 1$ or $p \to 0$, Equation (2) tends to a classical "SI" model, as the one analyzed in [20], [29]. Moreover we used μ as normalization factor because the actual rate of transition between i and v states does not influence the system, since it is only the proportion time spent in v that makes the botnet 'apparent' and generate spamming. Finally, notice that using an apportioning coefficient (p) controlling the transition rates between i and v ensures that the sum of agent in I and V states remains equal to the total number of infected nodes, which is coherent with the fact that switching between spam and no–spam states is controlled by the botnet itself.

B. Botnets with Re-Infection

In some cases the vaccination of nodes is not feasible, either because there is no known anti–worm footprint or simply because infected nodes must be re–installed from scratch from a CD-ROM and then upgraded to be vaccinated: the on-line time between installation and vaccination leave them susceptible [24], [25]. These are R-Botnets since a node can be *Re-infected* after recovery.

Indeed, if re-infection is introduced, many interesting additional parameters can be considered based on the re-infection model, even without changing the classes of agents. The simplest model is when agents in class v can either be removed with rate γ or return susceptible with rate ρ as shown in Figure 2 only considering the solid transition rates.

A straightforward extension is considering also the possibility of a preventive cure, which would introduce a direct transition from class s to class r as indicated by the dashed arrow in Figure 2. Additionally, if one wants to consider worms which can mutate to make complete recovery more difficult, one can introduce different infectious rates parameterized on the class as indicated by the parameter β_c in the dashed box in Figure 2. However, to maintain discussion simple, we will only consider the fundamental case (without considering dashed lines 'extensions').

Using the same normalization approach we used in Section II-A, the model in Figure 2, considering only the solid-line transitions and and not mutations is formalized by the following set of normalized differential equations:

$$\begin{cases} \dot{s}(t) &= -\beta\left[i(t) + v(t)\right]s(t) + \rho v(t) \\ \dot{i}(t) &= \beta\left[i(t) + v(t)\right]s(t) - \frac{1}{1-p}i(t) + \frac{1}{p}v(t) \\ \dot{v}(t) &= \frac{1}{1-p}i(t) - \left(\frac{1}{p} + \rho + \gamma\right)v(t) \end{cases} \tag{3}$$

notice that the special case $\rho = 0$ is the model we discussed in Section II-A.

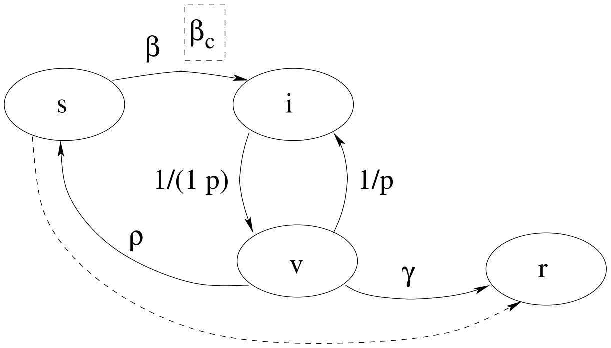

Fig. 2. Normalized epidemic flow among agents classes for re-infecting botnets

C. Plausible Ranges for Transmission and Recovery Rates

Before proceeding with the analysis of the models, we would like to restrict the meaningful range of the β and γ parameters. Let us start with the recovery rate. We normalized the system with respect to μ, which describes the transition rate of botnets agents between the spamming and the non-spamming states. With the assumption that agents in state i cannot be detected, we must have $\gamma < 1$. In the following we will always use $\gamma = 0.25$, but any other values < 1 can be used and the fundamental properties of the system will not change.

Figure 3 reports the maximum fraction of nodes which are simultaneously infected during the epidemic (plot (A)), and the fraction of nodes r_∞ that have contracted the worm and recovered at the end of the epidemic itself (plot (B)) as a function of the ratio β/γ. A quick inspection of the behavior shows that for $\beta/\gamma > 10$, the fraction of infected nodes tends to 1 and, for small p, also the fraction of nodes infected at the same time tends to one. This is a behavior that, to the best of our knowledge, has not yet been observed in botnets. Indeed, this also seems to be an unlikely case in modern Internet, where the diversity of hardware, operating systems and versions thereof make it difficult to imagine a single worm that can infect the entire population, even if no known countermeasures are yet available when it appears. Thus in the following we will restrict the analysis to values $\beta/\gamma < 16$.

In the following sections we discuss the key theoretical properties of our models, that will be used to show how a botnet can be efficiently built and exploited, and, at the same time, how countermeasures can be taken against botnets, even in face of very aggressive behaviors.

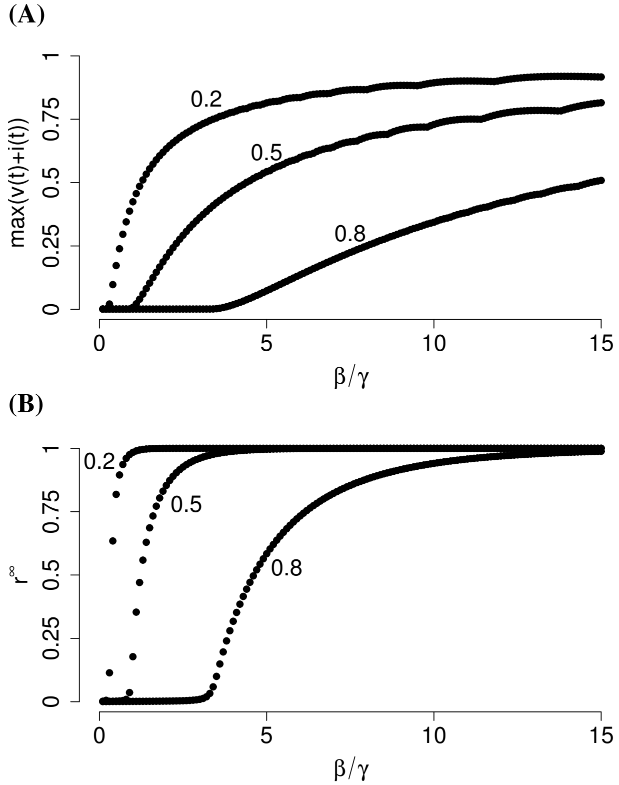

Fig. 3. I–Botnets. **Plot (A):** Maximum value of $i(t)+v(t)$ as a function of the fraction β/γ for a fixed value of $\gamma = 0.25$. The value of p is reported in the figure. The initial condition is $s(0) = 0.99999$, $i(0) = 0.00001$, $v(0) = r(0) = 0$. **Plot (B):** Proportion of recovered agents $r(t)$ at the end of the epidemic as a function of the fraction β/γ for a fixed value of $\gamma = 0.25$. In both **a** and **b** the value of p are reported in the figure and the initial condition is $s(0) = 0.99999$, $i(0) = 0.00001$, $v(0) = r(0) = 0$.

III. Theoretical Properties

In order to understand the dynamics of Equations (2)–(3) the key parameters are the basic and the effective reproductive numbers.

In mathematical models for the epidemiology of infectious diseases (included computer worms, as the cases of system defined by Equations (2)–(3)), the basic reproductive number R_0 is defined as the threshold parameter which determines if the introduction of an infected agent can lead to an epidemic ($R_0 > 1$) or not ($R_0 < 1$). Therefore, if $R_0 < 1$, the introduction of infectious individuals in the population can not lead to a major epidemic but only to small clusters of cases; whereas, if $R_0 > 1$, an epidemic can start with a positive probability.

The basic reproductive number R_0 can essentially be interpreted as the number of new infections generated by an infectious agent during its entire infectivity period, assuming a completely susceptible population [1]. Similarly, the effective reproductive number $R_e(t)$ can be defined as the number of new infections generated by an infectious agent at time t, and hence defines the sustainability of the infection after time t.

Considering botnets, if $R_0 < 1$, then we can assume that the spam generated is very limited and the game is not worth the effort, thus a worm whose parameters leads to $R_0 < 1$ can be considered non dangerous. Similarly, when $R_e(t) < 1$, even if $R_0 > 1$, the epidemic is in decreasing phase, i.e., $(\dot{i}(t) + \dot{v}(t)) < 0$, or in other words, the nodes in the botnet are decreasing, and the worm is being eliminated from the system.

To compute the reproductive numbers we adopt the next–generation matrix technique introduced by Diekmann and Heesterbeek in [9]. In particular, we exploit the same method described in [19], which is a specific adaptation for models with multiple infectivity classes with mutual interchanges.

A. Threshold Condition for I-Botnets

Note that Equation (2) admits a continuum of equilibria (called disease–free equilibrium) given by $(s^\star, 0, 0)$ (see Appendix A). Therefore, let us consider the Jacobian J of Equation (2) restricted to the infectious classes i and v, computed at the disease–free equilibrium. It can be written as

$$J = T + \Sigma - D ,$$

where

$$T = s^\star \begin{pmatrix} \beta & \beta \\ 0 & 0 \end{pmatrix}$$

is a matrix whose elements are all real non–negative numbers, and they correspond to the transmission rates;

$$\Sigma = \begin{pmatrix} -\frac{1}{1-p} & \frac{1}{p} \\ \frac{1}{1-p} & -\frac{1}{p} \end{pmatrix}$$

is a real matrix with positive off–diagonal elements corresponding to transition between the infectivity classes;

$$D = \begin{pmatrix} 0 & 0 \\ 0 & \gamma \end{pmatrix}$$

is a real non-negative diagonal matrix whose strictly positive element represents the recovery rate.
Since $\Sigma - D$ is invertible we can compute

$$-(\Sigma - D)^{-1} = \frac{1-p}{\gamma} \begin{pmatrix} \frac{1}{p} + \gamma & \frac{1}{p} \\ \frac{1}{1-p} & \frac{1}{1-p} \end{pmatrix}$$

and, noticing that all its elements are real and positive, it is possible to estimate R_0 as the dominant eigenvalue of the next–generation matrix K defined as

$$\begin{aligned} K &= -T(\Sigma - D)^{-1} \\ &= s^\star \frac{1-p}{\gamma} \begin{pmatrix} \beta\left(\frac{1}{p} + \gamma + \frac{1}{1-p}\right) & \beta\left(\frac{1}{p} + \frac{1}{1-p}\right) \\ 0 & 0 \end{pmatrix} . \end{aligned}$$

Since $\det(K) = 0$, it follows that the dominant eigenvalue of K (namely R_0) is

$$R_0 = \frac{\beta}{\gamma} \frac{1 + \gamma p(1-p)}{p} s^\star \qquad (4)$$

R_0 represents the main parameter of the system, in fact if $R_0 > 1$ epidemic will occur. Otherwise if $R_0 < 1$ epidemic can not occur.

Using similar arguments we obtain a formula for the effective reproductive number:

$$R_e(t) \quad = \quad \frac{\beta}{\gamma}\frac{1+\gamma p(1-p)}{p}s(t) \ . \qquad (5)$$

B. Threshold Condition for R-Botnets

Using the same approach, not repeated here for the sake of brevity, we compute the basic and effective reproductive number of Equation (3) for R-Botnets:

$$R_0 \quad = \quad \frac{\beta}{\gamma+\rho}\frac{1+(\gamma+\rho)p(1-p)}{p}s^\star \qquad (6)$$

and

$$R_e(t) \quad = \quad \frac{\beta}{\gamma+\rho}\frac{1+(\gamma+\rho)p(1-p)}{p}s(t) \ . \qquad (7)$$

IV. System Dynamics

I–Botnets (Equation (2)) show the classical threshold behavior, with the threshold parameter given in Equation (4). If $R_0 > 1$, the shape of the epidemic is the classical one: susceptibles (s) and removed (r) tend to some real positive value, while the two infectious classes (i, v) tend to zero (see Appendix A).

On the other hand, the transition phase strongly depends on p. In particular, by increasing the value of p, class i is favored and the epidemic peak increases in height and width (see Fig. 4 (A)). The dynamics of v is instead more complicated. As shown in Figure 4(B), intermediate values of p favor the proliferation of v. This means that being excessively aggressive or not sufficiently aggressive are both bad strategies.

R–Botnets show the same qualitative behavior of I–Botnets, but quantitative results are different. Not surprisingly, by fixing the vale of $\gamma+\rho$ in Equation (3) at the same value of γ in Equation (2), the dynamics of both i and v in R–Botnets are higher than the corresponding for I–Botnets (see Fig. 4).

A. Building the Botnet

I–Botnets can always be built, even if the worms are not much transmissible (i.e. if β is small). As stated before, an epidemic occurs (i.e., a botnet can be built) only if $R_0 > 1$. From Equation (4) we can deduce the conditions under which $R_0 > 1$ holds:

C1: $\frac{\beta}{\gamma}s^\star > 1$;

or

C2: $\frac{\beta}{\gamma}s^\star < 1$ and

$$p \quad < \quad \frac{1}{2}\left(1-\frac{1}{\beta s^\star}+\sqrt{\left(1-\frac{1}{\beta s^\star}\right)^2+\frac{4}{\gamma}}\right) \ . \ (8)$$

Therefore, by varying p in $(0, 1)$, Equation 8 can always be verified and thus the threshold condition can always be satisfied. In conclusion, independently for the countermeasures (represented by γ), the botnet can always be built simply by acting on the aggressivity p of the botnet itself.

In the same way, we can derive the conditions for having $R_0 > 1$ for R-botnets:

C1: $\frac{\beta}{\gamma+\rho}s^\star > 1$;

or

C2: $\frac{\beta}{\gamma+\rho}s^\star < 1$ and

$$p \quad < \quad \frac{1}{2}\left(1-\frac{1}{\beta s^\star}+\sqrt{\left(1-\frac{1}{\beta s^\star}\right)^2+\frac{4}{\gamma+\rho}}\right) \ .$$

Differently from I–Botnets, for each value of p it is possible to find a value of ρ which is able to interrupt the creation of a botnet. Therefore, creation of R–Botnets can be interrupted by a sufficiently large rate of temporary countermeasures ρ, a results apparently contraddictory. Fig. 5 shows quantitative estimation on the value of ρ needed to bring the state of the system in an under-threshold condition.

B. Maximizing the Damage of Botnets

For both I–Botnets and R–Botnets, the total number of spam sent by the botnet can be defined as

$$v^\infty \quad := \quad k\int_0^\infty v(t)dt$$

for any $k > 0$ representing the number of spam sent by a bot in the unit of time. Obviously, v^∞ increases as v increases (because $0 \le v(t) < 1 \ \forall t$); therefore, the aim of the botmaster is to maximize v^∞.

The value of v^∞ can be controlled by varying p, as suggested by Figure 4(B),(D). Given the difficulty in analytically obtaining the dependence between v^∞ and p, a numerical evaluation on $\frac{d}{dp}v^\infty$has been performed.

Given γ, I–Botnets have the potential to reach a maximum amount of damage $v_{\max}$. Figure 6 shows how, in order to reach $v_{\max}$ a value of p can always be found, for each (reasonable) transmission rate β. Surprisingly, β does not influence the potential damage a botnet can do, which is instead controlled only by the countermeasures rate γ, which turns out to be the only parameter that controls the overall time (integral value on the botnet existence) that agents spend in states infective of susceptible. On the ohter hand, not surprisingly, the higher is the transmission potential of the worm (β), the larger is the set of p such that the damage is maximum. In particular, Figure 6 shows that optimal values of p are around 0.1, corresponding to not very aggressive botnets; i.e. hidden botnets are much more dangerous than very aggressive ones. Acting in an extremely covert way ($p \approx 0$) is however not efficient as well.

For R–Botnets, the behavior is slightly more complex than for I–Botnets. The chance of reaching the maximum damage $v_{\max}$ depends on the temporary recovery rate (ρ) (as shown in Figure 7). This proves that the temporary countermeasures have positive effects on limiting the damage of the botnets. Obviously, the higher the transmission rate, the less effective are the countermeasures. Interestingly from Figure 7 we can

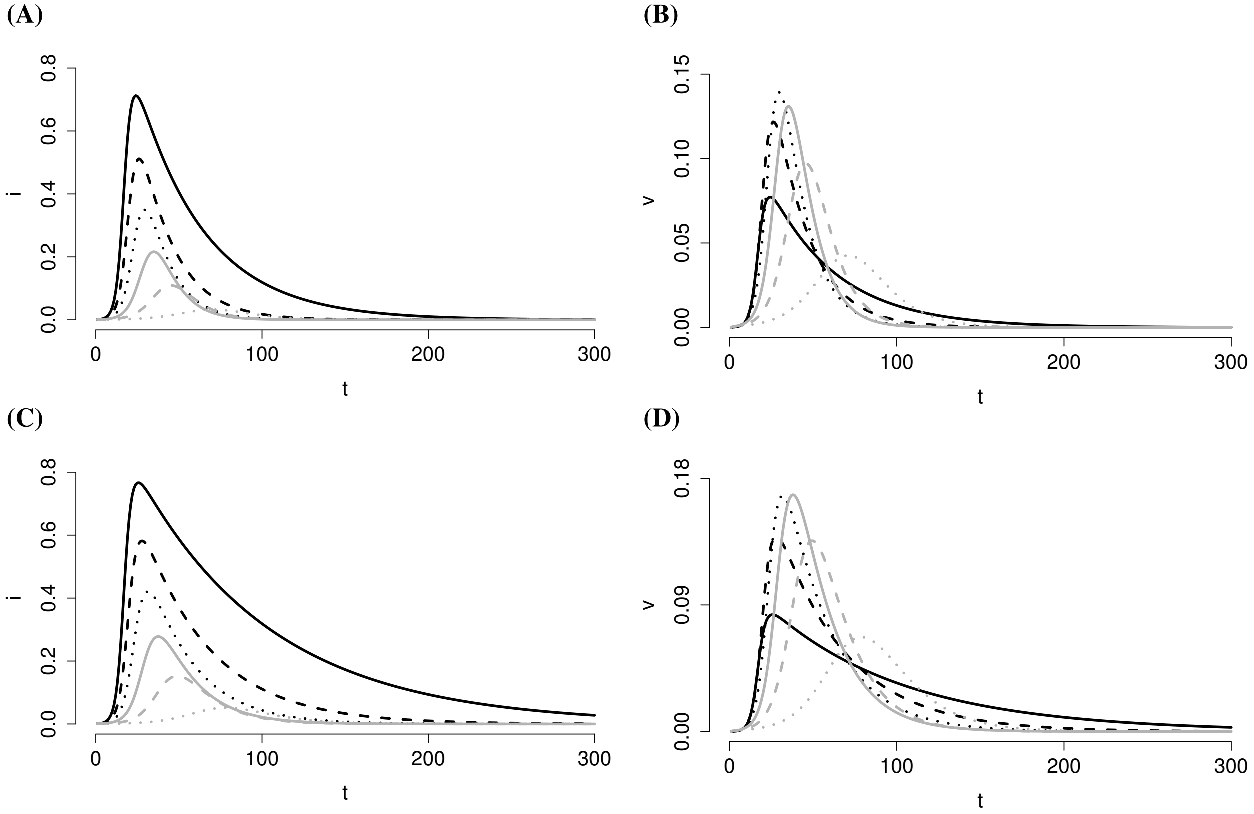

Fig. 4. *I–Botnets*: Temporal evolution of class i **(A)** and v **(B)**. Parameter choices: $\beta = 0.5$, $\gamma = 0.25$; $p = 0.1$ (black line), $p = 0.2$ (dashed black line), $p = 0.3$ (pointed black line), $p = 0.4$ (grey line), $p = 0.5$ (dashed grey line), $p = 0.6$ (pointed grey line). The initial condition for the entire figure is $s(0) = 0.999$, $i(0) = 0.001$, $v(0) = r(0) = 0$. *R–Botnets*: Temporal evolution of class i **(C)** and v **(D)**. Parameter choices and initial condition are the same of figure **(A)** and **(B)** but for $\gamma = 0.125$ and $\rho = 0.125$.

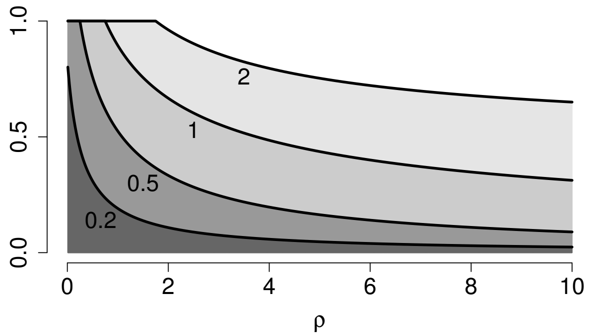

Fig. 5. R–Botnets. Black lines represent the parameters subsets where $R_0 = 1$. In the areas below each line $R_0 > 1$. Parameter choices are: $s^\star = 1$, $\gamma = 0.25$; the value of β is reported in the figure.

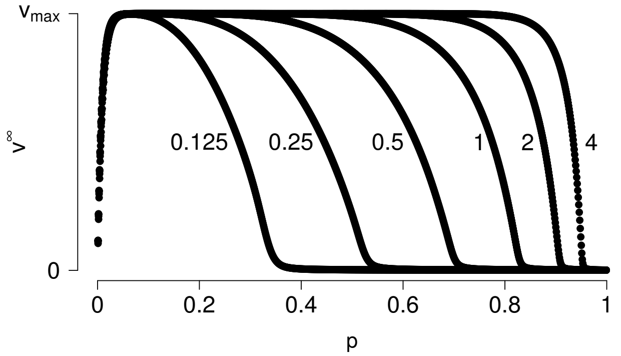

Fig. 6. I–Botnets. Value of v^∞ as a function of p. Parameter choices are: $\gamma = 0.25$; the value of β is reported in the figure.

deduce that very aggressive botnets are not much efficient, as observed for the case of I–Botnets.

The maximum value of v^∞ that can be produced by the models depends on γ as suggested by Figure 8. This happens because γ is related to the time that the agents can potentially spend in the infectivity classes (v and i) as we already noted.

Moreover, Figure 8 shows how v_{max} depends on γ and not on the sum $\gamma + \rho$. This proves that temporary countermeasures have measurable effects on the reduction of the optimal set of p and can reduce the actual damage done by the botnet as discussed in Figure 7, but only permanent countermeasures (γ) are able to considerably reduce the maximum potential damage of the botnet. Once more, we note that the transmittability β of the worm does not influence at all the maximum potential damage of the botnet, but only the range of parameters that can achieve this maximum.

C. Multiple Waves of Spam Messages

Since the parameter p can be remotely changed by the botmaster controlling the botnet, we want to investigate the temporal evolution of I–Botnets under the hypothesis of a non–constant value of p. In this context, an interesting behavior is

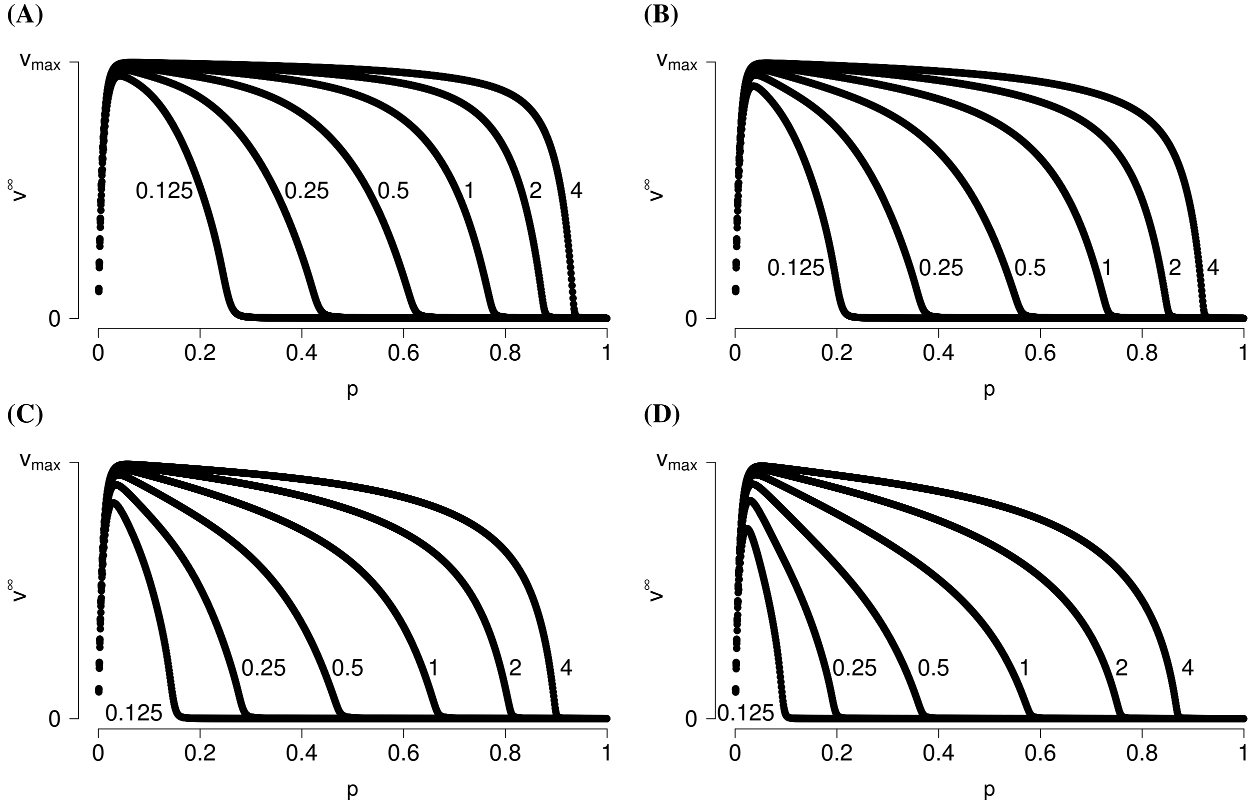

Fig. 7. R–Botnets. Value of v^{∞} as a function of p. Parameter choices are: $\gamma = 0.25$, $\rho = 0.125$ in **(A)**, $\rho = 0.25$ in **(B)**, $\rho = 0.5$ in **(C)**, $\rho = 1$ in **(D)**; the value of β is reported in the figure.

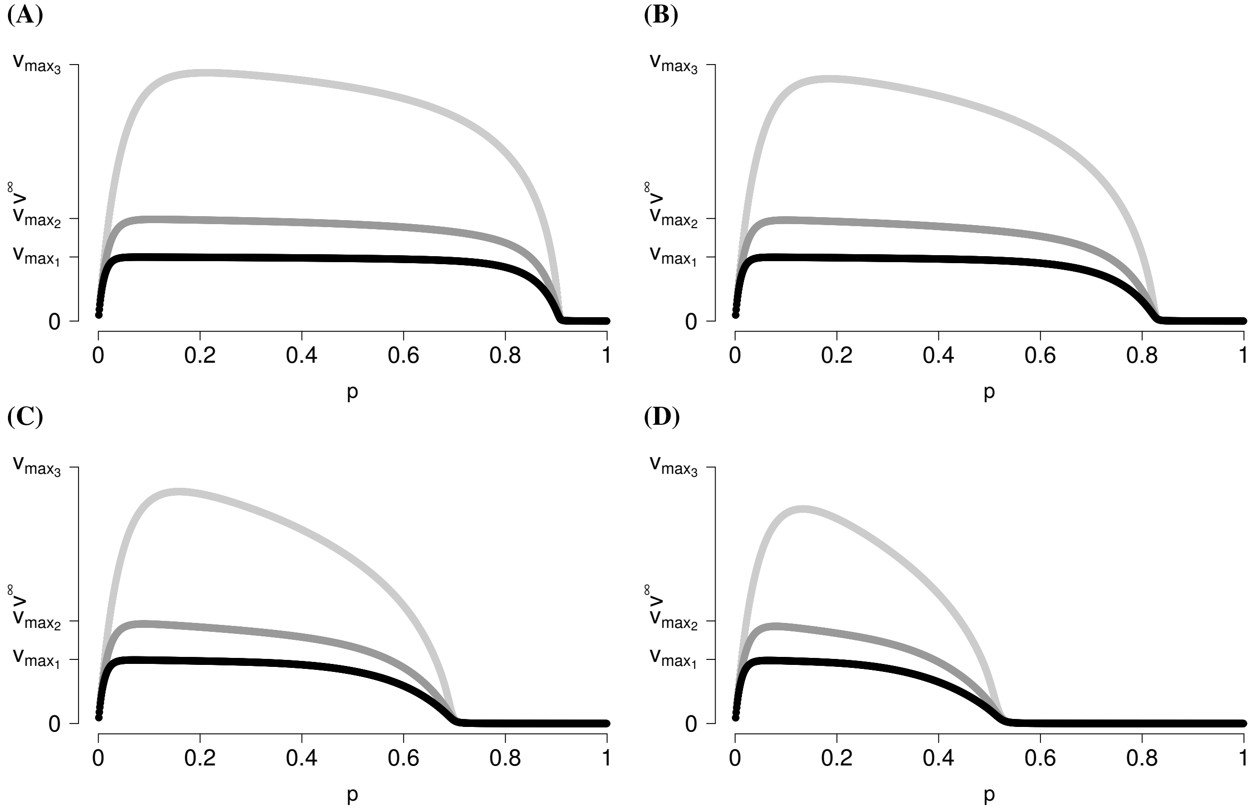

Fig. 8. R–Botnets. Value of v^{∞} as a function of p obtained by keeping fixed the quantity $\gamma + \rho$. Parameters choices are: $\gamma = 0.2$, $\rho = 0.05$ for black dots; $\gamma = 0.125$, $\rho = 0.125$ for dark grey dots; $\gamma = 0.05$, $\rho = 0.2$ for light grey dots; $\beta = 2$ in **(A)**, $\beta = 1$ in **(B)**, $\beta = 0.5$ in **(C)**, $\beta = 0.25$ in **(D)**.

the possibility of having multiple waves of spam messages.

As stated in Sec. III, the number of new infections over time depends on the value of $R_e(t)$: in particular if $R_e(t) > 1$ the number of new infections increases, while if $R_e(t) < 1$ it decreases. Equation (5) shows the dependence of $R_e(t)$ on p. Using the same argument of Section IV-A, it follows that, for every value of $s(t)$, it is possible to choose p in such a way to have $R_e(t) > 1$. Therefore, at any time the botmaster could increase p for generating a new wave of messages, without the need of increasing the transmission rate. In general, increasing the transmission rate β is very hard, because it depends on the structure of the population where the worm is spreading and on the infecting capability of the worm. Therefore, to generate a new wave of spam in botnets, it is an attractive option to vary the aggressivity parameter (e.g., by using remote controls) instead of varying the transmissibility parameter.

Here we show an example of a possible alternative definition of p which implies a multiple waves pattern of i and v. We assume that p can be defined by the following piecewise function:

$$p \quad := \quad \begin{cases} p_1 & \text{if } v(t) \leq \bar{v} \\ p_2 & \text{otherwise} \end{cases} \quad . \tag{9}$$

Evaluating the number of nodes in a botnet is very simple, if the botnet is controlled by a C&C centralized server. But even modern peer-to-peer botnet can easily estimate the size of an overlay network [13].

In Fig. 9 the temporal evolution of i and v show a multiple waves pattern. Even if, in general, the botmaster could increase p at any time for generating a new wave of messages, by using Equation (9) to define p this is not true. In fact, it can be proved that $v(t) \to 0$ repeating the same reasoning of Appendix A by taking into account that the solutions of Equation (2) are continuous even if p is defined as in Equation (9).

The same arguments hold also for R–Botnets, not reported here for the sake of brevity.

V. Related Work

The idea that mathematical epidemiology could be adapted to computer viruses is 20-year old [21]; few years later, the first serious attempt in this sense was published [15]. Since then, the number of papers on this subject has grown steadily, presenting a large collection of analytical models to better understand the propagation of Internet worms [4], [26], [29] and the evaluation of intervention options in response to the propagation itself [20], [30].

The idea of viruses and worms that adapt their behavior in order to improve their possibility of staying stealthy is not new. Self-stopping worms [17], for example, may reduce the infection pace after having detected a substantial saturation of the system, becoming quiescent "sleeper agents". This reduce the possibility of detection, as it is the prodigious infection activity that is typically used to identify compromised nodes.

Varying Scan Rate (VSR) worms deliberately vary their scan rate to avoid being effectively detected by existing worm detection techniques. Yu et al. [28] provide an epidemiological model of such worms, and discuss a new detection scheme capable to act as a countermeasure; however, their work does not consider the effects of VSR on the malicious activity, however.

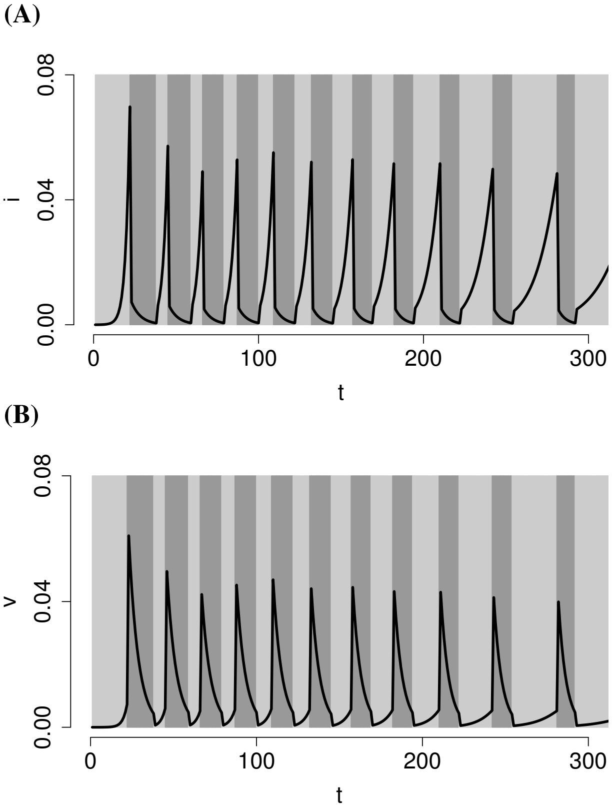

Fig. 9. I–Botnets. Temporal evolution of i (**A**) and v (**B**) obtained by considering p evolving according to Equation (9). In light grey colored areas $p = p_1$, while in dark grey colored areas $p = p_2$. Parameter choices for both are: $\beta = 0.5$, $\gamma = 0.25$, $p_1 = 0.1$, $p_2 = 0.9$, $\bar{v} = 0.005$. The initial condition is $s(0) = 0.99999$, $i(0) = 0.00001$, $v(0) = r(0) = 0$.

While self-stopping worms and VSR worms are techniques explicitly designed to improve the probability of being undetected, variations in the activity of viruses and worms can be caused by the diurnal variations in the number of active machines. Dagon et al. [7] have analytically modeled such periodic behavior, and consider the existence of distinct time zones using a differential equation approach similar to the one we use. As the previous paper, their work is limited to the propagation phase, without considering the actual damage caused by malicious activities.

The importance of staying inactive for some time after a period of activity (called dormancy) is well–known in biological models. E.g., thanks to this dormancy phenomena, Mycobacterium tuberculosis are able to survive inside human hosts for years and be responsible for latent tuberculosis [22].

Finally we note how compartmental differential equations are closely related to stochastic Markovian models used with success to model dynamic Internet protocols [10], [11], [16].

VI. Extending the Model

So far we have only considered the simplest possible models that capture the property of modern botnet agents of swapping between active and dormient states without the need of a centralized control and with apparently random behavior of single agents.

The technique of compartimental differential equations can be used of model many other aspects and facets of the botnets and malware behaviors. We briefly comment on some possible extensions here.

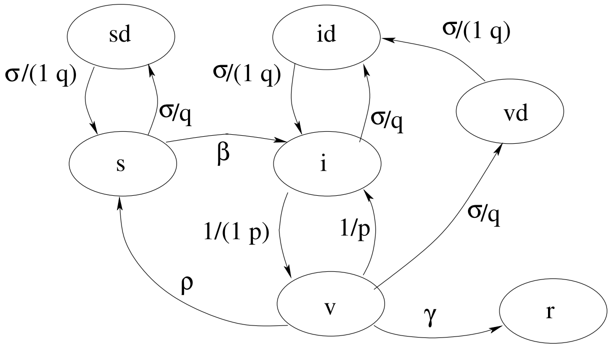

Fig. 10. A model for nodes that switch on and off with rate σ

For instance a model like the one represented in Fig. 10 can represent a system where nodes are not always on line (only susceptibles and infected dormient are meaningful with this respect), but they switch on and off with a rate σ and are on-line with proportion q. Nodes that are switched off while actually spamming will go into the infective but not spamming status when switched on again, thus potentially giving an illusion of problem solution. We set the on-off ratio and rate equal for all states, but this can be easily modifyied, for instance to represent users behavior that aggressively switch node off if they recognize some malicious behavior.

VII. Conclusions

This work introduced the use of compartmental differential equations to model the spreading of botnets-creating worms, as well as, for the first time to the best of our knowledge, to model the amount of damage (e.g., the number of spam messages) the botnet can cause during its lifetime.

Results underline how a botnet can be easily built using worms that are not very transmissible, and how a botnet can have a great impact in terms of spamming. Moreover, it proves how a botnet built by a single worm can cause multiple waves of spamming just by changing the probability with which single agents swap between active and dormient states.

Thus, this study highlights the necessity to identify all the infected agents in order to limit the proliferation of the botnets. In fact, it is not sufficient to recognize only the spamming agents and removing only them, independently of how well it is done (namely the value of the recovery rate).

The models we analyzed are on purpose the simplest ones, that are prone to non-ambiguous interpretation even in the absence of quantitative data to tune the parameters. Many other models can be derived form these ones, and one has been shortly introduced at the end of the paper as example.

Indeed, many other hypotheses should be checked and other containment/mitigation strategies, more sophisticated than temporary or definitive recovery, can be introduced. For instance these models only capture the pool of susceptible computers, while it might be interesting to model the entire population, including nodes that are protected by antivirus software before being infected. All these control options could be quite easily tested by adapting this mathematical model in order to take them into account.

Further work includes considering non-homogeneous systems, for instance considering differences in terms of bandwidth among agents. Being this work a purely theoretical study, it will be interesting to extend this paper by testing its results on real data set.

References

[1] R. M. Anderson and R. M. May. *Infectious diseases of humans: dynamics and control*. Oxford, UK: Oxford University Press, 1992.

[2] D. Carra, R. Lo Cigno, and E. W. Biersack. Graph Based Analysis of Mesh Overlay Streaming Systems. *IEEE Journal on Selected Areas in Communications*, 25(9):1667–1677, Dec. 2007.

[3] D. Carra, R. Lo Cigno, and E. W. Biersack. Stochastic Graph Processes for Performance Evaluation of Content Delivery Applications in Overlay Networks. *IEEE Trans. on Parallel and Distributed Systems*, 19(2):247 –261, feb. 2008.

[4] Z. Chen, L. Gao, and K. Kwiat. Modeling the spread of active worms. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'03)*, volume 3, pages 1890–1900, 2003.

[5] K. Chiang and L. Lloyd. A Case Study of the Rustock Rootkit and Spam Bot. In *Proceedings of First USENIX Workshop on Hot Topics in Understanding Botnets (HotBots'07)*, 2007.

[6] E. Cooke, F. Jahanian, and D. McPherson. The Zombie Roundup: Understanding, Detecting, and Disrupting Botnets. In *Proceedings of the USENIX SRUTI Workshop*, 2005.

[7] D. Dagon, C. Zou, and W. Lee. Modeling botnet propagation using time zones. In *Proceedings of the 13th Annual Network and Distributed System Security Symposium (NDSS'06)*, Feb. 2006.

[8] N. Daswani and M. Stoppelman. The anatomy of Clickbot.A. In *Proceedings of 1st USENIX Workshop on Hot Topics in Understanding Botnets (HotBots'07)*, 2007.

[9] O. Diekmann and J. A. P. Heesterbeek. *Mathematical epidemiology of infectious diseases: model building, analysis and interpretation.* John Wiley & Son, 2000.

[10] M. Garetto, R. Lo Cigno, M. Meo, and M. Ajmone Marsan. Closed queueing network models of interacting long-lived TCP flows. *IEEE/ACM Trans. on Networking*, 12(2):300–311, April 2004.

[11] M. Garetto, R. Lo Cigno, M. Meo, and M. Ajmone Marsan. Modeling short-lived TCP connections with open multiclass queuing networks. *Computer Networks*, 44(2):153 – 176, 2004.

[12] J. B. Grizzard, V. Sharma, C. Nunnery, B. B. Kang, and D. Dagon. Peer-to-peer botnets: overview and case study. In *Proceedings of First USENIX Workshop on Hot Topics in Understanding Botnets (HotBots'07)*, 2007.

[13] M. Jelasity, A. Montresor, and O. Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Trans. Comput. Syst.*, 23(1):219–252, Aug. 2005.

[14] S. Kandula, D. Katabi, M. Jacob, and A. W. Berger. Botz-4-Sale: Surviving organized DDoS attacks that mimic flash crowds. In *Proceedings of the 2nd Symposium on Networked Systems Design and Implementation (NSDI'05)*, May 2005.

[15] J. O. Kephart and S. R. White. Directed-graph epidemiological models of computer viruses. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 343–359, Los Alamitos, CA, USA, 1991. IEEE Computer Society.

[16] R. Lo Cigno and M. Gerla. Modeling window based congestion control protocols with many flows. *Performance Evaluation*, 3637:289 – 306, 1999.

[17] J. Ma, G. M. Voelker, and S. Savage. Self-stopping worms. In *Proceedings of the 2005 ACM Workshop on Rapid Malcode (WORM'05)*, pages 12–21, New York, NY, USA, 2005. ACM.

[18] B. McCarty. Botnets: Big and Bigger. *IEEE Security and Privacy*, 1:87–90, 2003.

[19] S. Merler, P. Poletti, M. Ajelli, B. Caprile, and P. Manfredi. Coinfection can trigger multiple pandemic waves. *Journal of Theoretical Biology*, 254(2):499–507, 2008.

[20] D. Moore, C. Shannon, G. M. Voelker, and S. Savage. Internet Quarantine: Requirements for Containing Self-Propagating Code. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'03)*, pages 1901– 1910, 2003.

[21] W. H. Murray. The application of epidemiology to computer viruses. *Comput. Secur.*, 7(2):139–145, 1988.

[22] N. M. Parrisha, J. D. Dickb, and W. R. Bishaic. Mechanisms of latency in Mycobacterium tuberculosis. *Trends in Microbiology*, 6:107–12, 1998.

[23] M. Rajab, J. Zarfoss, F. Monrose, and A. Terzis. A multifaceted approach to understanding the botnet phenomenon. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, pages 41–52. ACM New York, NY, USA, 2006.

[24] SANS Istitute Internet Storm Center. Windows XP: Surviving the first day, 2003. White paper.

[25] SANS Istitute Internet Storm Center. Windows Vista: First steps, 2007. White paper.

[26] S. Shakkottai and R. Srikant. Peer to Peer Networks for Defense Against Internet Worms. In *Proceedings of the 2006 Workshop on Interdisciplinary Systems Approach in Performance Evaluation and Design of Computer & Communications Sytems*, 2006.

[27] H. R. Thieme. *Mathematics in population biology*. Princeton: Princeton University Press, 2003.

[28] W. Yu, X. Wang, D. Xuan, and D. Lee. Effective detection of active worms with varying scan rate. In *Proceedings of IEEE International Conference on Security and Privacy in Communication Networks (SecureComm), Baltimore, MD, August*, 2006.

[29] C. C. Zou, W. Gong, and D. Towsley. Code Red Worm Propagation Modeling and Analysis. In *Proceedings of the 9th ACM Conference on Computer and communications security*, pages 138 – 147, 2002.

[30] C. C. Zou, W. Gong, and D. Towsley. Worm propagation modeling and analysis under dynamic quarantine defense. In *Proceedings of the 2003 ACM workshop on Rapid malcode*, pages 51–60, 2003.

APPENDIX

In this appendix we prove the basic properties of Equation (2).

A. Proof of the Stability of Disease Free Equilibrium

Equation (2) admits a continuum of equilibria given by $(s^\star, 0, 0)$, for all $0 < s^\star \leq 1$[2]. This continuum of equilibria, commonly called *disease free equilibrium (DFE)*, is unstable if $R_0 s^\star > 1$; therefore in this case an epidemic occurs. Otherwise, if $R_0 s^\star < 1$, the DFE is stable, but not asymptotically stable, and major epidemics can not occur.

Now, let us prove that the DFE is unstable if $R_0 s^\star > 1$. First of all, we compute the Jacobian matrix computed at the DFE:

$$J_{dfe} \;=\; \begin{pmatrix} 0 & -\beta s^\star & -\beta s^\star \\ 0 & \beta s^\star - \frac{1}{1-p} & \beta s^\star + \frac{1}{p} \\ 0 & \frac{1}{1-p} & -\frac{1}{p} - \gamma \end{pmatrix} .$$

The eigenvalues are the roots of the characteristic polynomial

$$-\lambda \left[\lambda^2 - \lambda \left(\beta s^\star - \frac{1}{1-p} - \frac{1}{p} - \gamma \right) - \left(\beta s^\star \left(\frac{1}{p} + \frac{1}{1-p} + \gamma \right) - \frac{\gamma}{1-p} \right) \right] .$$

If $R_0 s^\star > 1$ it follows that $\beta s^\star \left(\frac{1}{p} + \frac{1}{1-p} + \gamma \right) - \frac{\gamma}{1-p} > 0$ and it proves that the equilibrium is unstable.

B. Proof of the asymptotic behaviour

In this section we will show that the solution of Equation (2) asymptotically tends to $(s_\infty, 0, 0, r_\infty)$, where $0 < s_\infty < 1$ and $0 < r_\infty < 1$.

For simplicity we add the equation related to the removed class to Equation (2). Therefore we have:

$$\begin{cases} \dot{s}(t) &= -\beta \left[i(t) + v(t) \right] s(t) \\ \dot{i}(t) &= \beta \left[i(t) + v(t) \right] s(t) - \frac{1}{1-p} i(t) + \frac{1}{p} v(t) \\ \dot{v}(t) &= \frac{1}{1-p} i(t) - \left(\frac{1}{p} + \gamma \right) v(t) \\ \dot{r}(t) &= \gamma v(t) \end{cases} \tag{10}$$

From the first equation of (10) we have that $\dot{s}(t) < 0$ for all $t \geq 0$, therefore s is always (strictly) decreasing. Moreover, s is bounded ($0 < s \leq 1$). Thus, $\lim_{t\to\infty} s(t) = s_\infty$ and $0 < s_\infty < 1$. From the fourth equation of (10) we have

$$r(t) \;=\; \gamma \int_0^t v(\sigma) d\sigma + r(0)$$

Moreover, for all t, $r(t) = \gamma \int_0^t v(\sigma) d\sigma + r(0) < 1$ because $s(t) + i(t) + v(t) + r(t) = 1$. Thus,

$$\int_0^\infty v(\sigma) d\sigma \;<\; \frac{1}{\gamma} - \frac{r(0)}{\gamma} \;<\; \infty$$

Integrating the third equation of (10) we obtain

$$v(t) \;=\; \frac{1}{1-p} \int_0^t i(\sigma) d\sigma - (\frac{1}{p} + \gamma) \int_0^t v(\sigma) d\sigma \;<\; 1 \; \forall t$$

Therefore,

$$\begin{aligned} \int_0^\infty i(\sigma) d\sigma \;&<\; (1-p) + (1-p) \left(\frac{1}{p} + \gamma \right) \int_0^t v(\sigma) d\sigma \\ &<\; \infty \end{aligned}$$

Since $1 = s + i + v + r$, it follows that

$$\begin{aligned} \lim_{t\to\infty} i(t) + \lim_{t\to\infty} v(t) \;&=\; 1 - s_\infty - \gamma \int_0^\infty v(\sigma) d\sigma \\ &<\; \infty \end{aligned}$$

In conclusion, $\lim_{t\to\infty} i(t) = 0$ and $\lim_{t\to\infty} v(t) = 0$.

[2] If $s^\star = 0$, trivially, no epidemic can occur.