Modeling Botnets and Epidemic Malware

Marco Ajelli
Fondazione Bruno Kessler, Trento, Italy

Renato Lo Cigno, Alberto Montresor
DISI, University of Trento, Italy

Abstract—**Botnets have become the most sophisticated and dangerous way of spreading malware. Their damaging actions can range from massive dispatching of e-mail 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.**

I. Introduction and Related Work

A *botnet* is a collection of infected end-hosts, called *bots*, that are under the control of a human operator known as *botmaster*. Today, botnets are inherently linked to malicious activities and have been identified as one of the most important threats to the security of the Internet [1].

The lifecycle of a botnet can be briefly described as follows [2]. First, the botnet creator sends out a virus, or worm, that infects unprotected 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, awaiting for commands from the C&C. When one of such commands is received, the bots autonomously perform the malicious actions for which they have been instructed. The bot code can switch freely between the *hidden* and the *active* state. Recent botnets have tried to avoid the centralized point of failure represented by the C&C in favor of modern P2P architectures.

Examples of attacks performed by botnets include sending of spam messages, click fraud against major search engines and blackmail based on denial-of-service.

An appropriate modeling of the botnet threat is a necessary condition to understand the dynamics of the threat they pose. We are interested in modeling two interconnected phases of the lifecycle of a 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 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.

Mathematical models have been already proposed to better understand both the propagation of Internet worms [3], and the evaluation of intervention options in response to the propagation itself [4]. 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 or limiting these processes. We adopt a differential equation approach derived from models of biological diseases [5]. Due to space restrictions, we had to omit several proofs, formula derivations, and additional evaluations. For an extended version of this paper, please refer to [6].

II. The Models

We consider two different models: in the first one, nodes eventually recover after infection, and recovered nodes cannot be infected any more. This scenario corresponds, for instance, to a botnet that exploits a single vulnerability, and this vulnerability is fixed once and for all. The second one assumes instead that nodes can be re-infected. Examples of this behavior include botnets that mutate during their lifecycle, exploiting different vulnerabilities at different times. Our study is based on compartmental ordinary differential equations models; we assume a finite population, which can be arbitrarily large, since our models are scale–invariant w.r.t. the number of nodes.

A. Botnets Subject to Immunization

The *I-Botnet* model ("I" stands for immunization) is composed of four classes of nodes:

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.

The epidemic flow among these classes is represented in Fig. 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 node 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 ($N > 0$).

Infectious nodes can start or stop spamming. Since worms building botnets are normally entirely quiescent, and consequently very hard to detect, we assume that only spamming nodes can be detected and removed.

A fraction p of infected nodes is hidden (class I), and a fraction $(1 - p)$ is actively spamming (class V). Since the

This work is supported by the Italian MIUR – Prin 2008 project, Autonomous Security.

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

behavior of nodes with respect to the spamming/hiding activity is autonomous and dynamic (an node decides on its own whether to spam or be silent), there is a flow $I \rightarrow V$ with normalized rate $1/(1-p)$ and a flow $V \rightarrow I$ with normalized rate $1/p$, rates that at equilibrium gives exactly the p, $(1-p)$ apportioning between the two classes.

The actual rate of transition between I and V depends on how often nodes decide to switch between the two states. Let μ be the average switching rate, so that the the mean time spent in class I is $(1-p)/\mu$ and in class V it is p/μ.

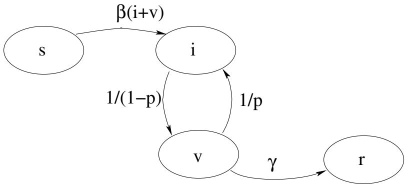

Fig. 1. Normalized epidemic flow among nodes classes for I–botnets

The model above is depicted in Fig. 1. Since we consider a closed population the relation $N = S(\tau)+I(\tau)+V(\tau)+R(\tau)$ $\forall\tau$ holds, so we can formalize the model with the following system of three differential equations:

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

The system is written in adimensional form, by defining $t := \mu\tau$, $s := S/N$, $i := I/N$, $v := V/N$, $\beta := b/\mu$ and $\gamma := g/\mu$. To complete the definition, we assign the 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;
- $v(0) = 0$, i.e. no spamming nodes 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 nodes in the population; therefore epidemics cannot occur.

Remark Let us note that if $p \rightarrow 1$ or $p \rightarrow 0$, Eq. (1) tends to a classical “SI” model, as those analyzed in [4].

B. Botnets with Re-Infection

In some cases immunity is not possible, for instance because there is no known anti–worm footprint; we use the term *R-Botnet* to define such botnets (“R” stands for immunization).

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 nodes. The simplest model is when nodes in class v can either be removed with rate γ or return susceptible with rate ρ as shown in Fig. 2.

Another straightforward extension is the possibility of a preventive cure, which would introduce a direct transition from class s to class r (dashed arrow in Fig. 2). Additionally, to take into consideration 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 (dashed box in Fig. 2). However, for the sake of simplicity, we will only consider the basic model (without considering dashed lines ‘extensions’).

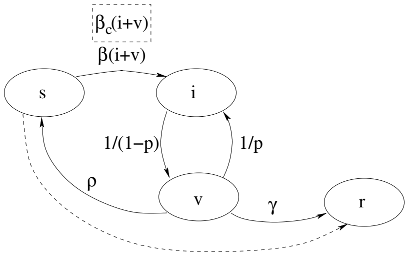

Fig. 2. Normalized epidemic flow among nodes classes for R–Botnets

Using the same normalization approach already used, the model in Fig. 2 is formalized as follows:

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

notice that the special case $\rho = 0$ is the I-Botnet model.

C. Plausible Range for Transmission and Recovery Rates

Before proceeding with the analysis of the models, we restrict the meaningful range of the β and γ parameters. We normalized the system with respect to μ, which describes the transition rate of botnets nodes between the spamming and non-spamming states. With the assumption that nodes in state i cannot be detected, we have $\gamma < 1$. In the following we fix $\gamma = 0.25$, but any value less than 1 can be used as the fundamental properties of the system do not change.

Fig. 3(a) reports the maximum fraction of nodes which are simultaneously infected during the epidemic; Fig. 3(b) reports the fraction of nodes r_∞ that have contracted the worm and recovered by the end of the epidemic itself as a function of the ratio β/γ. A quick inspection of the behavior shows that for $\beta/\gamma > 10$, the fraction of recovered 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

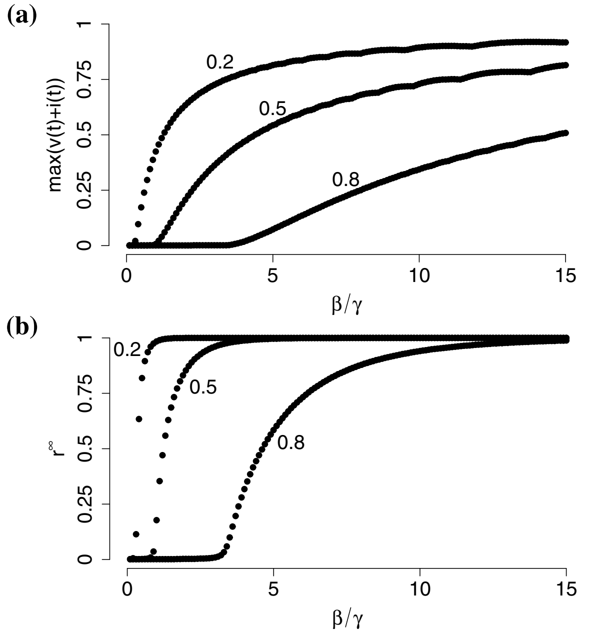

Fig. 3. I–Botnets. **(a)**: Maximum value of $i(t) + v(t)$ as a function of the fraction β/γ for a fixed value of $\gamma = 0.25$. The value of p is reported in the figure. The initial condition is $s(0) = 0.99999$, $i(0) = 0.00001$, $v(0) = r(0) = 0$. **(b)**: Proportion of recovered nodes $r(t)$ at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

also seems to be an unlikely case in modern Internet, where the diversity of hardware and operating systems thereof make it difficult to imagine a single worm that can infect the entire population, even if no known countermeasures are available when it appears. Thus in the following we will restrict the analysis to values $\beta/\gamma < 10$.

III. Theoretical Properties

In order to understand the dynamics of systems (1)–(2), the key parameters are the basic and the effective reproductive numbers. In epidemiological models, the *basic reproductive number* R_0 is defined as the threshold parameter which determines if the introduction of an infected node can lead to an epidemic ($R_0 > 1$) or not ($R_0 < 1$). R_0 can essentially be interpreted as the number of new infections generated by an infectious node during its entire infectivity period, assuming a completely susceptible population [5]. Similarly, the *effective reproductive number* $R_e(t)$ can be defined as the number of new infections generated by an infectious node 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 a decreasing phase, i.e., $(\dot{i}(t) + \dot{v}(t)) < 0$, or in other words, the botnet is decreasing in size, and the worm is being eliminated from the system.

To compute the reproductive numbers we adopt the next–generation matrix technique with the method described in [7].

A. Threshold Condition for I-Botnets

Eq. (1) admits a continuum of equilibria (called disease–free equilibrium) given by $(s^\star, 0, 0)$. Let us consider the Jacobian J of Eq. (1) restricted to the infectious classes i and v, computed at the disease–free equilibrium. It can be written as $J = T + \Sigma - D$, where

$$T = s^\star \begin{pmatrix} \beta & \beta \\ 0 & 0 \end{pmatrix},\ \Sigma = \begin{pmatrix} -\frac{1}{1-p} & \frac{1}{p} \\ \frac{1}{1-p} & -\frac{1}{p} \end{pmatrix},\ D = \begin{pmatrix} 0 & 0 \\ 0 & \gamma \end{pmatrix}$$

T is a real matrix whose elements are non–negative numbers corresponding to the transmission rates; Σ is a real matrix with positive off–diagonal elements corresponding to transition between the infectivity classes; D is a real non-negative diagonal matrix whose strictly positive element represents the recovery rate. Since $\Sigma - D$ is invertible we can compute

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

Noticing that all its elements are real and positive, it is possible to define R_0 as the dominant eigenvalue of the next–generation matrix K [8], where $K = -T(\Sigma - D)^{-1}$ is equal to:

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

Since $\det(K) = 0$, the dominant eigenvalue of K (i.e. R_0) is

$$R_0 = \frac{\beta}{\gamma} \frac{1 + \gamma p(1-p)}{p} s^\star \qquad (3)$$

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:

$$R_e(t) = \frac{\beta}{\gamma} \frac{1 + \gamma p(1-p)}{p} s(t)\ . \qquad (4)$$

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 Eq. (2) for R-Botnets:

$$R_0 = \frac{\beta}{\gamma + \rho} \frac{1 + (\gamma + \rho)p(1-p)}{p} s^\star \qquad (5)$$

$$R_e(t) = \frac{\beta}{\gamma + \rho} \frac{1 + (\gamma + \rho)p(1-p)}{p} s(t)\ . \qquad (6)$$

IV. System Dynamics

I–Botnets show a threshold behavior, with the threshold parameter given in (3). If $R_0 > 1$, the shape of the epidemic is classical: susceptibles (s) and removed (r) tend to some real positive value, while the two infectious classes (i, v) tend to zero. 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 Fig. 4(b), intermediate values of p favor the proliferation

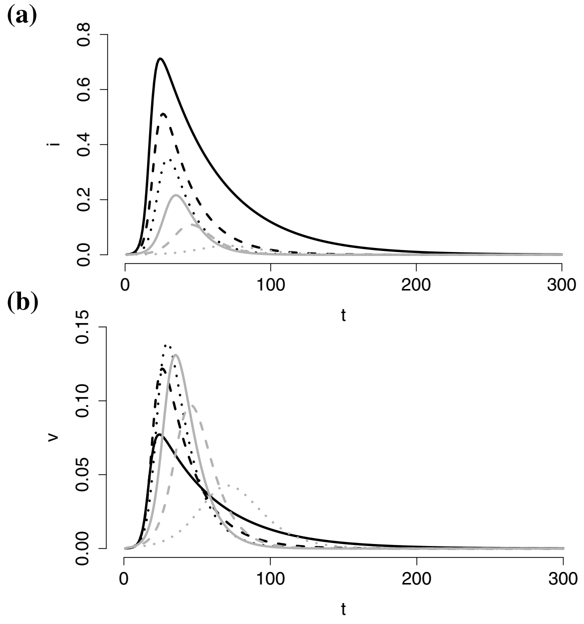

Fig. 4. Temporal evolution of class i **(a)** and v **(b)** in I-Botnets. Parameters: $\beta = 0.5$; $\gamma = 0.25$. $p = 0.1$ (black line), $p = 0.2$ (dashed black line), $p = 0.3$ (dotted black line), $p = 0.4$ (grey line), $p = 0.5$ (dashed grey line), $p = 0.6$ (dotted grey line). The initial condition is $s(0) = 0.999$, $i(0) = 0.001$, $v(0) = r(0) = 0$.

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 (2) at the same value of γ in (1), the dynamics of both i and v in R–Botnets are higher than the corresponding for I–Botnets [6].

A. Building Botnets

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 Eq. (3) 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

$$p \;<\; \frac{1}{2}\left(1 - \frac{1}{\beta s^\star} + \sqrt{\left(1 - \frac{1}{\beta s^\star}\right)^2 + \frac{4}{\gamma}}\right) \;. \qquad (7)$$

Therefore, by varying p in $(0, 1)$, Eq. 7 can always be verified and thus the threshold condition can always be satisfied. In conclusion, independently of the countermeasures (represented by γ), the botnet can always be built simply by acting on the aggressivity p of the botnet itself.

Similarly, 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

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

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 contradictory. Fig. 5 shows quantitative estimation on the value of ρ needed to bring the state of the system in an under-threshold condition.

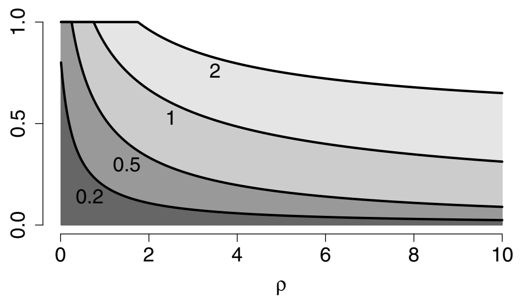

Fig. 5. R–Botnets. Black lines represent the parameters subsets where $R_0 = 1$. In the areas below each line $R_0 > 1$. Parameters: $s^\star = 1$, $\gamma = 0.25$; the value of β is reported in the figure.

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 $v^\infty := k\int_0^\infty v(t)dt$ 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 $0 \leq v(t) < 1 \; \forall t$); therefore, the aim of the botmaster is to maximize v^∞.

The value of v^∞ can be controlled by varying p. 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 nodes spend in states infective or susceptible. On the other hand, 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, Fig. 6 shows that optimal values of p are around 0.1, corresponding to botnets not very aggressive; 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 Fig. 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.

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

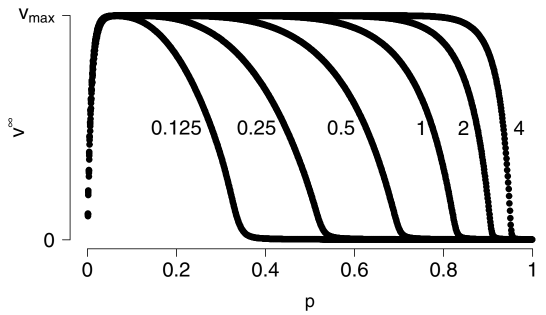

Fig. 6. I–Botnets. Value of v^{∞} as a function of p. Parameters: $\gamma = 0.25$; the value of β is reported in the figure.

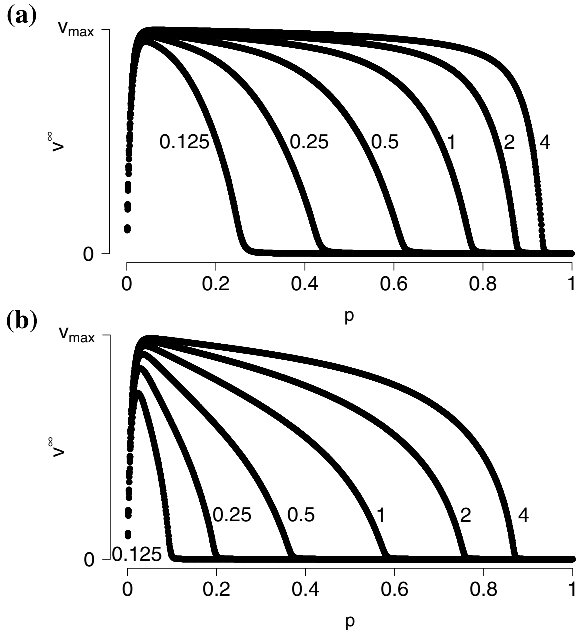

Fig. 7. R–Botnets. Value of v^{∞} as a function of p. Parameters: $\gamma = 0.25$, $\rho = 0.125$ in **(a)**, and $\rho = 1$ in **(d)**; the value of β is reported in the figure.

temporal evolution of I–Botnets under the hypothesis of a non–constant value of p. In this context, an interesting behavior is the possibility of having multiple waves of spam.

As stated in Sect. 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. Eq. (4) shows the dependence of $R_e(t)$ on p. Using the same argument of Sect. 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 can increase p for generating a new wave of spam.

Let p be defined by the following piecewise function:

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

Fig. 8 reports the temporal evolution of i (v is complementary) and shows a multiple waves pattern. In general, bots can generate a new wave at any time, simply by increasing p.

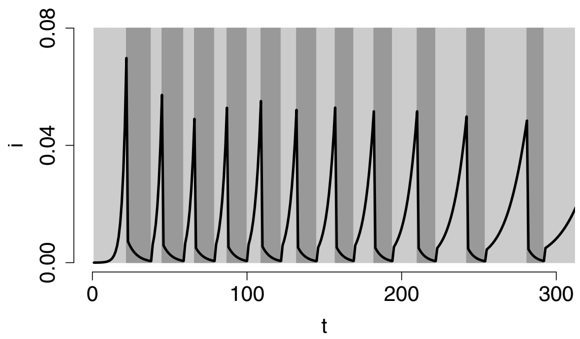

Fig. 8. I–Botnets. Temporal evolution of i obtained modifying p. In light grey colored areas $p = 0.1$, while in dark grey colored areas $p = 0.9$. Parameters: $\beta = 0.5$, $\gamma = 0.25$, $\bar{v} = 0.005$.

V. Conclusions

This work introduced the use of compartmental differential equations to model the spreading of botnets and the amount of damage 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 dormant states.

We analyzed extremely simple models, 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. The contribution of this work is the definition of an easy, yet powerful modeling framework that can be used for testing the spreading capabilities of epidemic malware and proof-of-concept countermeasures. The modeling framework can be specialized and tuned to specific cases; additionally, stochastic versions of the techniques can be developed if data is available to test the stochastic hypotheses.

References

[1] E. Cooke, F. Jahanian, and D. McPherson, "The zombie roundup: Understanding, detecting, and disrupting botnets," in *Proc. of the USENIX SRUTI Workshop*, 2005.

[2] M. Rajab, J. Zarfoss, F. Monrose, and A. Terzis, "A multifaceted approach to understanding the botnet phenomenon," in *Proc. of the 6th ACM SIGCOMM Conf. on Internet Measurement*. ACM, 2006.

[3] D. Dagon, C. Zou, and W. Lee, "Modeling botnet propagation using time zones," in *Proc. of NDSS'06*, Feb. 2006.

[4] D. Moore, C. Shannon, G. M. Voelker, and S. Savage, "Internet quarantine: Requirements for containing self-propagating code," in *Proc. of INFOCOM'03*, 2003, pp. 1901– 1910.

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

[6] Ajelli, M. and Lo Cigno, R. and Montresor, A., "Compartmental differential equations models of botnets and epidemic malware (extended version)," University of Trento, T.R. DISI-10-011, 2010, available at http://disi.unitn.it/locigno/preprints/TR-DISI-10-011.pdf

[7] S. Merler, P. Poletti, M. Ajelli, B. Caprile, and P. Manfredi, "Coinfection can trigger multiple pandemic waves," *Journal of Theoretical Biology*, vol. 254, no. 2, pp. 499–507, 2008.

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