Controlling Risk of Data Exfiltration in Cyber Networks Due to Stealthy Propagating Malware

Brian Thompson
U.S. Army Research Lab
Adelphi, MD 20783
Email: bthompso8784@gmail.com

James Morris-King
U.S. Army Research Lab
Adelphi, MD 20783
Email: james.r.morris-king.ctr@mail.mil

Hasan Cam
U.S. Army Research Lab
Adelphi, MD 20783
Email: hasan.cam.civ@mail.mil

***Abstract*—Infamous recent cyber attacks on businesses and governments have demonstrated that even the best contemporary security systems can not prevent well-resourced adversaries from infiltrating their networks and gaining access to sensitive information. Stealthy malware can spread through a network undetected by utilizing zero-day exploits to propagate and hiding malicious behavior in normal activity, potentially doing significant damage before exploited vulnerabilities can be identified or patches developed. In this work, we consider a scenario in which an attacker deploys propagating malware enabling the exfiltration of data from infected devices, and a defender deploys detection and recovery mechanisms designed to control malware spread while obeying network-wide resource constraints. We use a stochastic model to represent changes in the state of the network and analytically derive an upper bound on the total rate at which an optimal attacker can exfiltrate data from the network, expressed in terms of several network parameters, when the detection rate is proportional to the outgoing data rate at each infected device. Our results can help inform cybersecurity decision-makers in judiciously allocating resources to manage risk.**

I. Introduction

A. Motivation

Rapidly advancing technologies and evolving operational requirements increasingly drive both private and public sector organizations toward highly interconnected computer networks [1]. This increases their susceptibility to cyber attack as well as the magnitude of damage an attack could cause. It is therefore essential that organizations develop robust network defense mechanisms to provide security against cyber threats.

Several recent cyber attacks making use of propagating malware to cause large-scale damage have been brought to public attention [2]. In 2010, the Stuxnet worm exploited five vulnerabilities, four of which were zero-day, to sabotage the computer hardware of Iran's nuclear program [3]. Kelihos (2010) and Dorkbot (2014) botnets used the Dropbox service and Amazon's Cloud, respectively, to deliver updates to zombie systems [4], [5]. Cloud Atlas, discovered in 2014, used CloudMe cloud providers located in Sweden to store collected data from infiltrated machines [6]. In 2015, Kaspersky Lab was successfully attacked using a variant of Stuxnet known as Duqu 2.0, which used several zero-day exploits to spy on the new anti-virus technologies being developed the lab [7]. Bitdefender customers' data was leaked after an attack by DetoxRansome in 2015 which hijacked several servers in Amazon's Elastic Compute Cloud (EC2) [8]. These cases are by no means isolated and suggest that traditional security technologies are insufficient to protect today's computer systems from cyber attacks.

A common defensive countermeasure is an *intrusion detection system* (IDS), which monitors activity on a computer or network and sets off an alert when suspected malicious activity occurs, prompting human analysts to investigate and take defensive action as deemed necessary. Alternatively, an *intrusion prevention system* (IPS) takes automated actions to block or purge a potential intrusion when an alert goes off. Regardless of which method is used or how it is implemented, the result is that when the detection mechanism sets off an alert for a device, the device is taken offline for some period of time while a recovery operation is performed and then comes back online clean of the malware. Due to service availability needs, monetary constraints, or other operational requirements, there is often a limit to the number of devices that can be undergoing recovery at any one time.

In this work, we consider a scenario in which an attacker deploys propagating malware enabling the exfiltration of data from infected devices, and a defender deploys detection and recovery mechanisms designed to control malware spread while obeying network-wide resource constraints. We present a compartmental stochastic model to represent changes in the state of the network that integrates models of communication between devices, malware propagation, and a recovery operation. We analytically derive an upper bound on the total rate at which an optimal attacker can exfiltrate data from the network, expressed in terms of several network parameters, when the detection rate is proportional to the outgoing data rate at each infected device. Our approach can help to inform cybersecurity decision-makers in assessing risk of data exfiltration and estimating the benefit of investing additional resources in improving the security of their networks.

B. Related Work

1) Epidemic models of malware spread: Mathematical models of biological viral epidemics such as the well-known Susceptible-Infected-Recovered (SIR) model [9] and its numerous variants have also been used to describe malware spread in cyber networks. Kephart and White apply such models to study the dynamics of computer virus spreading in a variety of contexts [10], [11]. They consider several network topologies, such as Erdos-Renyi random graphs, connected regular graphs, and sparse graphs with a high clustering coefficient. In all of these topological models they study the number of infected nodes in the population over time and how

various factors affect convergence to a steady state, finding that in many cases there exists a sharp epidemic threshold. Others have extended such models to additional network topologies and contexts. Boguna et al. [12] and Dezső et al. [13] focus on epidemic models for power-law networks. Mickens et al. examine device-to-device spreading of malicious software in mobile ad-hoc networks and present a queue-based technique [14] and a differential equations-based technique [15] to overcome some of the limitations of the earlier homogeneous models. Valler et al. develop a framework for analyzing epidemic spreading processes on mobile ad-hoc networks and find that mobility does not have a strong impact on virus spreading [16].

There are many similarities between our model and other epidemiologically-inspired models used in the cybersecurity literature. The main difference that distinguishes our work from existing work is that in the context of stealthy malware, there is a degree of uncertainty about which nodes are infected, so countermeasures can not be applied solely to infected nodes as is typically done in the epidemiological models. Due to constraints on network resources, we must address the additional challenge of decision-making under uncertainty.

2) Control of malware spread: Okhravi and Nicol evaluate the tradeoff between the time spent on pre-deployment testing and the timely deployment of patches for software vulnerabilities [17]. Khouzani et al. explore how to allocate resources to prevent the spread of an aggressive malware infection controlled by an attacker seeking to do maximum damage to a mobile wireless network [18]. They show that it is optimal for the attacker to have a propagation phase during which killing of any infected nodes is deferred, followed by an exploit phase during which all infected nodes are killed simultaneously. From the defender's perspective, they show that higher immunization and healing rates reduce the damage of the attack. Eshghi et al. propose patching strategies for countering propagating malware in both a replicative (patches can be transmitted by other patched devices) and a non-replicative (patches are only disseminated by designated sources) context [19].

These approaches rely on the existence of patches for known vulnerabilities, whereas our approach is also effective in combating stealthy attacks that may be exploiting unknown vulnerabilities.

3) Proactive cyber defense: Evans et al. evaluate the effectiveness of moving target defense against a variety of real-world attacks [20]. Colbaugh and Glass take a game-theoretic approach and propose moving target defense strategies against an adaptive attacker [21]. Ben-Asher et al. study the effectiveness of migration-based moving target defense against attackers with a range of skills and resources [22]. Shan et al. propose proactive restart of smartphone apps to reduce side channel time series predictability [23].

These works all consider defense mechanisms for a single system rather than a coordinated effort over networked devices, and are therefore not sensitive to the needs of the network as a whole. We propose an approach to security in cyber networks that aims to achieve cybersecurity goals while obeying resource constraints on a network-wide scale.

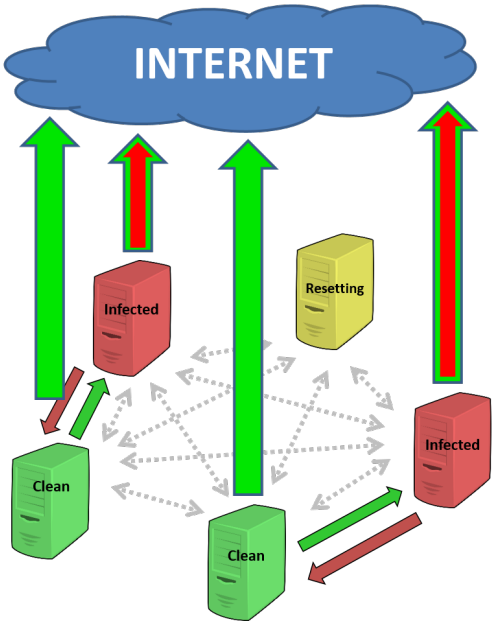

Fig. 1: A sample network illustrating our model. Green nodes are clean. Red nodes are infected. The yellow node is resetting. When a clean node and an infected node interact, malware spreads from the infected node to the clean node. Infected nodes exfiltrate sensitive data to an external location accessible to the attacker.

C. Contributions and Outline

The main contributions of this work are:

- A model of distributed cyber networks that incorporates elements of network communication, malware propagation, detection of malicious activity, and a defensive recovery operation
- Theoretical analysis of the equilibrium state of the network under a compartmental stochastic model
- An upper bound on the total rate of data exfiltration by an optimal attacker, expressed in terms of network parameters

In Section II we describe our model of cyber networks and present a stochastic model to capture network dynamics. In Section III we derive analytical results for the maximal total rate of data exfiltration by an optimal attacker. We conclude with some discussion and directions for future work in Section IV.

II. Model

A. Network Model

We model a cyber network as a set of interacting *nodes* corresponding to networked devices. In practice, nodes could be computers in a distributed computing cluster, components of a SCADA system, or terminal machines on a local area network. When operational, pairs of nodes periodically communicate with one another, providing a channel for data transfer between the two nodes, and each node additionally transmits some

information externally, e.g. to respond to user queries, report sensor readings, or send an email. For the purpose of this paper, we model each node as interacting with other operational nodes at *communication rate* λ and transmitting data externally at *upload rate* υ. Our model can equally be applied to represent communication on the network or application layer, and could be extended to consider different communication patterns or heterogeneous behavior of the nodes.

We consider an attacker trying to stealthily establish and maintain a presence on nodes in the network, which it uses to exfiltrate data. The attacker begins with a presence on an initial set of *seed* nodes and spreads stealthily to other nodes via self-propagating malware, which leverages existing communication by embedding itself in transmitted data. If the attacker has established a presence on a node, we say that the node is *infected*, otherwise we say it is *clean*. Each infected node, in addition to its normal activity, exfiltrates sensitive data to an external location accessible to the attacker at *exfiltration rate* ξ, determined by the attacker. By only spreading across existing communication channels and by limiting its exfiltration rate, the attacker hides evidence of its behavior in normal network activity. While in practice an attacker may not be able to spread between every pair of networked devices (e.g. due to protocol constraints, incompatibilities in system configurations, or defense mechanisms such as firewalls), we consider a worst-case attacker who can utilize every possible channel to spread to additional nodes. Our network model is illustrated in Figure 1.

We consider a defender using a detection mechanism to monitor network activity and flag the nodes that are most likely to be exfiltrating data. We model the detection rate at a node as being directly proportional to the average rate of outgoing traffic from that node, which includes both normal and malicious activity, and scaled by the *detector sensitivity* σ, determined by the defender. When a node is flagged by the detector, a *reset* operation is performed on the node, which restores it to a known malware-free state but renders it inoperable in the meantime. In practice, resetting could entail reimaging the disk or switching to a new virtual machine created from a clean disk image; the implementation is domain-specific and outside the scope of this paper. It should be understood that our approach will only be effective against malware that can be removed from a device by performing the designated reset operation.

Let r denote the *reset time*, the time required to reset a node and return it to operational status. When the reset is complete, we say that the node is *activated*. In practice, there may be costs associated with performing a reset, for example due to the loss of cached data or software that needs to be reinstalled. We capture such costs by including the time to recompute previously cached data or reinstall software in the reset time.

There is an obvious tradeoff between security and functionality of the network. Too high a threshold will allow infected nodes to exfiltrate data at higher rates without being detected. Too low a sensitivity threshold will result in a large number of resetting nodes. If the network is providing critical services, this may result in a self-imposed denial-of-service, which could be just as damaging as a successful cyber attack. Let θ denote the *operational threshold*, the minimum fraction of nodes required to be operational in order to sustain the desired level of network functionality.

Parameter	Description
λ	communication rate for each node
υ	upload rate for each node
ξ	exfiltration rate for each infected node
σ	detector sensitivity
r	reset time for each node
θ	operational threshold for the network
α	activation rate for each node
$\beta(t)$	infection rate at time t for each node
ρ_C	reset rate for clean nodes
ρ_I	reset rate for infected nodes
$\overline{\pi_C}$	limiting probability of clean nodes
$\overline{\pi_I}$	limiting probability of infected nodes
$\overline{\pi_R}$	limiting probability of resetting nodes

TABLE I: Notation. λ, υ, r, and θ are the given network parameters. ξ is a tunable parameter for the attacker. σ is a tunable parameter for the defender. α, $\beta(t)$, ρ_C, and ρ_I correspond to transition rates in our Markov model. $(\overline{\pi_C}, \overline{\pi_I}, \overline{\pi_R})$ is the limiting distribution over node states.

Next, we present a compartmental stochastic model to represent changes in the state of the network.

B. State Model

The state of the network can be conceptualized using a *Petri net*, which is defined by a set of *places*, a set of *transitions* between places, and a set of possible *configurations*. In our context, places correspond to the node states (Clean, Infected, and Resetting); transitions correspond to actions on the nodes (activate, infect, and reset); and a configuration specifies how many nodes are in each node state.

The Petri net diagrams in Figure 2 illustrate state changes for a sample network of 10 nodes. Figure 2 (a) shows the initial configuration, consisting of 6 clean nodes, 2 infected nodes, and 2 resetting nodes. Figure 2 (b) shows the configuration after two more nodes have been infected. At this point the network is highly vulnerable — if the attacker were to exploit the infected nodes, for example by deploying malware that causes all infected nodes to fail simultaneously, only the four clean nodes would remain operational. In Figure 2 (c), a node completes resetting and is activated. In Figure 2 (d), the defender begins resetting three nodes, all of which were infected. Although the number of infected nodes has decreased, so has the total number of operational nodes.

C. Stochastic Model

We model the behavior of each node with a continuous-time Markov chain corresponding to the states and transitions of the Petri net model described in Section II-B, illustrated in Figure 3. The states of the Markov chain are Clean, Infected, and Resetting, and transitions occur at *activation rate* α, *infection rate* β, and *reset rates* ρ_C and ρ_I for clean and infected nodes, respectively.

The activation rate α determines how long a node spends in the Resetting state before activating. We set $\alpha = \frac{1}{r}$ so that the expected time in the Resetting state is equal to the time required to perform the reset operation.

Transitions from Clean to Infected occur at the infection rate $\beta = \beta(t)$, which is a function of the communication rate λ and the current states of the other nodes. In particular, if

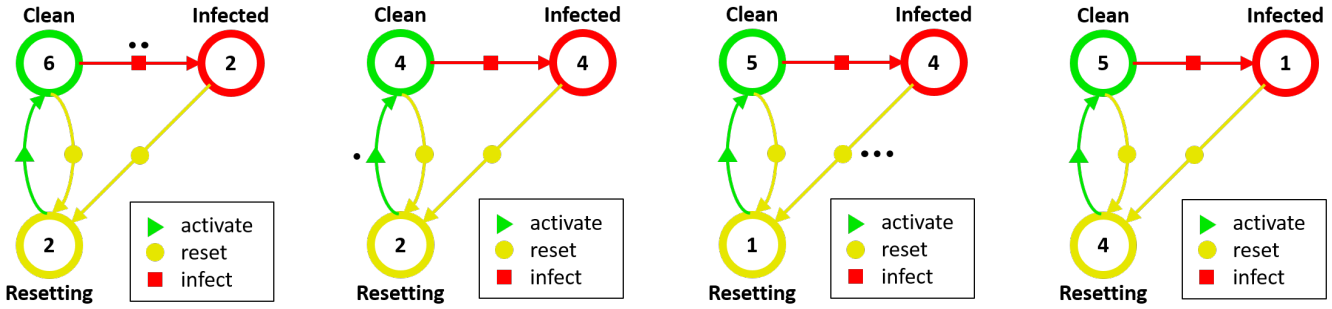

(a) (b) (c) (d)

Fig. 2: Petri net diagrams for a sample network of 10 nodes. Figure (a) shows the initial configuration. Figure (b) shows the configuration after two more nodes have been infected. In Figure (c), a node completes resetting and is activated. In Figure (d), the defender begins resetting three nodes, all of which were infected.

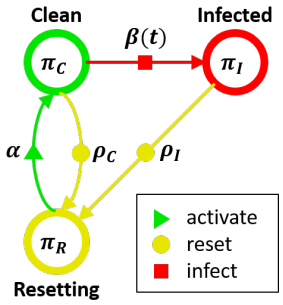

Fig. 3: Markov diagram for our compartmental model, indicating transition rates between the states. The values (π_C, π_I, π_R) correspond to the fractions of nodes in the Clean, Infected, and Resetting states, respectively.

there are k clean nodes and l infected nodes at time t, then the probability that a clean node gets infected upon interacting with another node is $\frac{l}{k+l}$. Therefore, the infection rate at time t is $\beta(t) = \lambda\frac{l}{k+l}$.

The reset rates are determined by the upload rate υ, the exfiltration rate ξ, and the detector sensitivity σ. Since nodes are reset when they are flagged by the detector, the reset rate at a node is the same as its detection rate, which we model as being directly proportional to the average rate of outgoing traffic from that node, including both normal and malicious activity (see Section II-A). Specifically, we have that $\rho_C = \sigma\upsilon$ and $\rho_I = \sigma(\upsilon + \xi)$.

Table I summarizes the notation used in our analysis.

In the following section, we use our stochastic model to analytically derive an upper bound on the rate at which an attacker can exfiltrate data without arousing suspicion.

III. Theoretical Analysis

Our goal is to determine the maximal total rate of data exfiltration by an optimal attacker, which corresponds to the individual exfiltration rate at each infected node times the number of infected nodes.

To solve this problem, we represent the network as a dynamical system, letting $\pi_C = \pi_C(t)$, $\pi_I = \pi_I(t)$, and $\pi_R = \pi_R(t)$ denote the fractions of nodes in the Clean, Infected, and Resetting states at time t, respectively, as depicted in Figure 3. Although for a network of n nodes the π values must be one of $\left\{\frac{1}{n}, \ldots, \frac{n}{n}\right\}$, approximating them as being real-valued and continuous allows us to apply tools from statistical mechanics. This can be considered a compartmental model, meaning that all nodes in the same state at a given time are assumed to behave homogeneously. We provide the master equations expressing the instantaneous rate of change for each of the compartments:

$$\begin{aligned}
\frac{d\pi_C}{dt} &= \alpha\,\pi_R - (\beta + \rho_C)\,\pi_C \\
\frac{d\pi_I}{dt} &= \beta\,\pi_C - \rho_I\,\pi_I \\
\frac{d\pi_R}{dt} &= \rho_C\,\pi_C + \rho_I\,\pi_I - \alpha\,\pi_R
\end{aligned}$$

Next, we derive the limiting distribution over node states by solving a nonlinear system of equations. At equilibrium, there is no change in the distribution, i.e. $\frac{d\pi}{dt} = 0$ for all states. Additionally, since the π values form a probability distribution over node states, they must sum to 1. Substituting the expressions for the transition rates α, β, ρ_C, and ρ_I from Section II-C yields the following system of equations, where $\overline{\pi_C}$, $\overline{\pi_I}$, and $\overline{\pi_R}$ are the limiting probabilities:

$$\begin{aligned}
\frac{1}{r}\,\overline{\pi_R} &= \left(\lambda\frac{\overline{\pi_I}}{\overline{\pi_C} + \overline{\pi_I}} + \sigma\upsilon\right)\,\overline{\pi_C} \\
\lambda\frac{\overline{\pi_I}}{\overline{\pi_C} + \overline{\pi_I}}\,\overline{\pi_C} &= \sigma(\upsilon + \xi)\,\overline{\pi_I} \\
\sigma\upsilon\,\overline{\pi_C} + \sigma(\upsilon + \xi)\,\overline{\pi_I} &= \frac{1}{r}\,\overline{\pi_R} \\
\overline{\pi_C} + \overline{\pi_I} + \overline{\pi_R} &= 1
\end{aligned}$$

Treating λ, r, υ, ξ, and σ as constants, and assuming that $\overline{\pi_C}, \overline{\pi_I}, \overline{\pi_R} > 0$, this system of equations of dimension three (the first three equations are linearly dependent) with three unknowns ($\overline{\pi_C}$, $\overline{\pi_I}$, and $\overline{\pi_R}$) can be solved analytically, yielding the following solution expressed in terms of the

network parameters:

$$\overline{\pi_C} = \frac{\sigma(\upsilon+\xi)}{\lambda + r\sigma(\upsilon+\xi)(\lambda-\sigma\xi)} \quad (1)$$

$$\overline{\pi_I} = \frac{\lambda-\sigma(\upsilon+\xi)}{\lambda + r\sigma(\upsilon+\xi)(\lambda-\sigma\xi)} \quad (2)$$

$$\overline{\pi_R} = 1 - \frac{\lambda}{\lambda + r\sigma(\upsilon+\xi)(\lambda-\sigma\xi)} \quad (3)$$

Next, we determine the optimal detection sensitivity σ in terms of the other parameters. Because increasing the detection sensitivity simultaneously increases the number of resetting nodes and decreases the number of infected nodes, the optimal value is simply the maximum value of σ that respects the operational threshold θ. This can be derived by solving the equation $\overline{\pi_R}(\sigma) = 1-\theta$ with $\overline{\pi_R}$ as derived above, subject to the constraints $\overline{\pi_C}, \overline{\pi_I}, \overline{\pi_R} > 0$, which yields

$$\sigma = \frac{\lambda}{2\xi} \cdot \left(1 \pm \sqrt{1 - \frac{4\xi(1-\theta)}{r\lambda\theta(\xi+\upsilon)}}\right).$$

It can be verified that only the "−" sign yields a feasible solution over the domain on which σ is defined.

Finally, we consider an optimal attacker, who wants to maximize the total rate of data exfiltration. This corresponds to the optimization problem

$$\max_{\xi} f(\xi) \quad \text{where} \quad f(\xi) = \xi \cdot \overline{\pi_I}.$$

Substituting the solutions derived above for $\overline{\pi_I}$ and σ yields

$$f(\xi) = \frac{1}{2}\left(\theta(\xi-\upsilon) + X\right),$$

where $X = \sqrt{1 - \frac{4(1-\theta)\xi}{\lambda r\theta(\xi+\upsilon)}}$, which has a derivative of

$$f'(\xi) = \frac{\theta}{2X}\left(X + 1 - \frac{2(1-\theta)(2\xi+\upsilon)}{\lambda r\theta(\xi+\upsilon)}\right).$$

The critical points occur when $f'(\xi) = 0$ and when $X = 0$.

Setting $f'(\xi) = 0$ yields

$$\xi = \frac{\upsilon}{2} \cdot \frac{2 \pm Y - Y^2}{Y^2 - 1},$$

where $Y = \sqrt{\frac{4(1-\theta)}{\lambda r\theta}}$. It can be verified that only the "+" sign yields a feasible solution over the defined parameter space, and that it is a local maximum. Simplifying yields a locally optimal value of ξ for the attacker in terms of the other network parameters:

$$\xi^* = \operatorname*{argmax}_{\xi} f(\xi) = \frac{\upsilon}{2} \cdot \frac{2-Y}{Y-1} = \upsilon \cdot \frac{1 - \sqrt{\frac{1-\theta}{\lambda r\theta}}}{2\sqrt{\frac{1-\theta}{\lambda r\theta}} - 1}. \quad (4)$$

Substituting ξ^* back into $f(\xi)$ yields the corresponding total rate of data exfiltration by the attacker:

$$\max_{\xi} f(\xi) = f(\xi^*) = \frac{1}{4}\left(\theta\upsilon \cdot \frac{4-3Y}{Y-1} - 2(Y-1)\right)$$
$$= \frac{1}{2}\left(\theta\upsilon \cdot \frac{2 - 3\sqrt{\frac{1-\theta}{\lambda r\theta}}}{2\sqrt{\frac{1-\theta}{\lambda r\theta}} - 1} - \left(2\sqrt{\frac{1-\theta}{\lambda r\theta}} - 1\right)\right). \quad (5)$$

However, note that ξ is only defined on the domain $[0, \infty)$. This puts a constraint on the feasibility of Equation 4. In particular, letting $Z = \frac{1-\theta}{\lambda r\theta}$, we see that if $Z > 1$ or $Z < \frac{1}{4}$, then $\xi^* < 0$. Next we examine the boundary cases: When $Z = 1$, $\xi^* = 0$ and therefore $f(\xi^*) = 0$, indicating that the attacker can not exfiltrate any data without getting purged from the network. Using l'Hôpital's Rule, we find that $\lim_{\xi\to 0}\sigma = \frac{1-\theta}{r\theta\upsilon}$, which substituted into Equation 2 along with $Z = 1$ and $\xi = 0$ yields $\overline{\pi_I} = 0$, implying that even when no data exfiltration occurs, the attacker is purged from the network. This represents the ultimate security guarantee: if the defender can reduce the time it takes to reset a device or reduce the operational requirement (e.g. by adding more devices) such that $\frac{1-\theta}{\lambda r\theta} \geq 1$, our proactive reset strategy is guaranteed to purge even the stealthiest malware from the network.

The other boundary case is when $Z = \frac{1}{4}$, in which case $f(\xi^*) \to \infty$, indicating that the attacker can exfiltrate data at an arbitrarily high rate without being purged from the network. This represents the worst case in security: if $\frac{1-\theta}{\lambda r\theta} \leq \frac{1}{4}$, the rate at which the defender is able to reset devices cannot keep up with the spread of the malware, and thus the attacker can exfiltrate data without fear of being stopped.

Recall that there was also a critical point when $X = 0$. Setting $X = 0$ yields

$$\xi = \frac{\upsilon}{Y^2 - 1} = \frac{\upsilon}{\frac{4(1-\theta)}{\lambda r\theta} - 1}.$$

Substituting back into $f(\xi)$ yields the corresponding total rate of data exfiltration by the attacker:

$$f(\xi) = \frac{1}{2}\theta\upsilon \cdot \frac{2-Y^2}{Y^2-1} = \frac{1}{2}\theta\upsilon \cdot \frac{2 - \frac{4(1-\theta)}{\lambda r\theta}}{\frac{4(1-\theta)}{\lambda r\theta} - 1}.$$

It can be verified that the solution is infeasible for $Z < \frac{1}{4}$, and that for $Z \geq \frac{1}{4}$, this solution is less than $f(\xi^*)$ as derived in Equation 5. Thus the locally optimal solution found in Equation 5 is also the global optimum.

Finally, we verify our results computationally by plotting $f(\xi)$ for a sample scenario with $\lambda = 1$, $r = 1$, $\theta = 0.7$, and $\upsilon = 1$ in Figure 4. The maximum point matches with the analytical solution $\xi^* \approx 1.1165$ and $f(\xi^*) \approx 0.2699$.

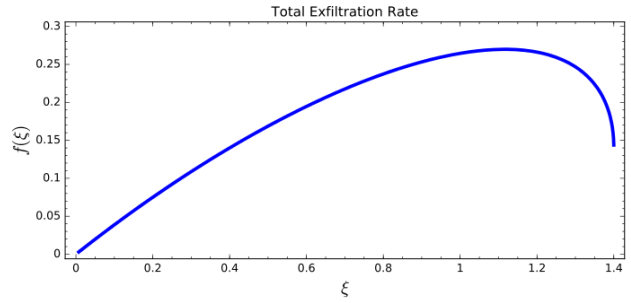

Fig. 4: Plot of the total exfiltration rate $f(\xi^*)$ as ξ varies across the feasible domain, with $\lambda = 1$, $r = 1$, $\theta = 0.7$, and $\upsilon = 1$.

IV. Conclusions

In this work we explored the relationship between technological capabilities, operational requirements, and the effectiveness of cyber defense mechanisms in limiting the rate at which an attacker can exfiltrate data from a cyber network. We first presented a model of cyber networks that captures communication between devices, silent propagation of malware over existing network communication, and a defensive reset operation that removes malware from a device at the cost of reduced availability to perform normal network functions. We used a continuous-time stochastic model to derive an upper bound on the total rate at which an optimal attacker can exfiltrate data from the network, expressed in terms of several network parameters, when the detection rate is proportional to the outgoing data rate at each infected device.

This is a significant result because it allows cybersecurity decision-makers to assess the maximal risk to their network in the case of an optimal attacker, and to estimate the benefit of investing additional resources in improving the robustness of their network by acquiring additional devices.

Directions for future work include considering arbitrary network structures instead of assuming all-pairs reachability, incorporating a mobility model, modeling multiple simultaneous attacks or the propagation of different types of malware, and developing additional policies such as one that estimates nodes' risk and factors that into the strategic decision-making process.

References

[1] R. Colbaugh, "Does coevolution in malware adaptation enable predictive analysis?" in *IFA Workshop: Exploring Malware Adaptation Patterns*, 2010.

[2] D. Lobo and S. Lertputtarak, "Botnets: From irc to android," in *The Fourth International Conference on Digital Information Processing and Communications (ICDIPC2014)*. The Society of Digital Information and Wireless Communication, 2014, pp. 23–27.

[3] R. Naraine, E. Protalinski, and D. Danchev, "Stuxnet attackers used 4 windows zero-day exploits," *ZDnet Blog*, 2010.

[4] D. Dittrich, "So you want to take over a botnet," in *Proceedings of the 5th USENIX conference on Large-Scale Exploits and Emergent Threats*. USENIX Association, 2012, pp. 6–6.

[5] A. Adamov, "Lavasoft security bulletin," *Lavasoft*, September 2013.

[6] Kaspersky, "Cloud atlas: Redoctober apt is back in style," https://securelist.com/blog/research/68083/cloud-atlas-redoctober-apt-is-back-in-style/, 2014, kaspersky Lab Global Research and Analysis Team.

[7] E. Kaspersky, "Kaspersky lab ivestigates hacker attack on its own network," *Kaspersky Lab*, 2015.

[8] T. Fox-Brewster, "Anti-virus firm bitdefender admits breach, hacker claims stolen passwords are unencrypted," *Forbes*, 2015.

[9] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," in *Proceedings of the Royal Society of London A: mathematical, physical and engineering sciences*, vol. 115. The Royal Society, 1927, pp. 700–721.

[10] J. Kephart and S. White, "Directed-graph epidemiological models of computer viruses," in *Research in Security and Privacy, 1991. Proceedings., 1991 IEEE Computer Society Symposium on*, May 1991, pp. 343–359.

[11] J. O. Kephart and S. R. White, "Measuring and modeling computer virus prevalence," in *Research in Security and Privacy, 1993. Proceedings., 1993 IEEE Computer Society Symposium on*. IEEE, 1993, pp. 2–15.

[12] M. Boguná, R. Pastor-Satorras, A. Vespignani *et al.*, "Epidemic spreading in complex networks with degree correlations," in *Proceedings of the XVIII Sitges Conference on Statistical Mechanics, Lecture Notes in Physics, Springer, Berlin*, 2003.

[13] Z. Dezső and A.-L. Barabási, "Halting viruses in scale-free networks," *Physical Review E*, vol. 65, no. 5, p. 055103, 2002.

[14] J. W. Mickens and B. D. Noble, "Modeling epidemic spreading in mobile environments," in *Proceedings of the 4th ACM Workshop on Wireless Security*, ser. WiSe '05. New York, NY, USA: ACM, 2005, pp. 77–86. [Online]. Available: http://doi.acm.org/10.1145/1080793.1080806

[15] ——, "Analytical models for epidemics in mobile networks," in *Third IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob 2007)*. IEEE, 2007, pp. 77–77.

[16] N. C. Valler, B. A. Prakash, H. Tong, M. Faloutsos, and C. Faloutsos, "Epidemic spread in mobile ad hoc networks: Determining the tipping point," in *Proceedings of the 10th International IFIP TC 6 Conference on Networking - Volume Part I*, ser. NETWORKING'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 266–280. [Online]. Available: http://dl.acm.org/citation.cfm?id=2008780.2008807

[17] H. Okhravi and D. Nicol, "Evaluation of patch management strategies," *International Journal of Computational Intelligence: Theory and Practice*, vol. 3, no. 2, pp. 109–117, 2008.

[18] M. Khouzani, S. Sarkar, and E. Altman, "Maximum damage malware attack in mobile wireless networks," *Networking, IEEE/ACM Transactions on*, vol. 20, no. 5, pp. 1347–1360, 2012.

[19] S. Eshghi, M. Khouzani, S. Sarkar, and S. Venkatesh, "Optimal patching in clustered malware epidemics," *Networking, IEEE/ACM Transactions on*, vol. 24, no. 1, pp. 283–298, Feb 2016.

[20] D. Evans, A. Nguyen-Tuong, and J. Knight, "Effectiveness of moving target defenses," in *Moving Target Defense*. Springer, 2011, pp. 29–48.

[21] R. Colbaugh and K. Glass, "Predictability-oriented defense against adaptive adversaries," in *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*, Oct 2012, pp. 2721–2727.

[22] N. Ben-Asher, J. Morris-King, B. Thompson, and W. J. Glodek, "Attacker skill, defender strategies, and the effectiveness of migration-based moving target defense in cyber systems," in *Proceedings of the International Conference on Cyber Warfare and Security*. ACPI, 2016.

[23] Z. Shan, I. Neamtiu, Z. Qian, and D. Torrieri, "Proactive restart as cyber maneuver for android," in *Proceedings of the Conference on Military Communications*, 2015.