# Optimal and Robust Epidemic Response for Multiple Networks

Michael Bloem [a,1] Tansu Alpcan [b] Tamer Başar [c]

[a] *NASA Ames Research Center, Moffett Field, CA 94035-1000, USA.*

[b] *Deutsche Telekom Laboratories (T-Labs), D-10587 Berlin, Germany.*

[c] *Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA.*

**Abstract**

This paper studies the optimization of malicious software removal or patch deployment processes across multiple networks. The well-known classical epidemic model is adapted to model malware propagation in this multi-network framework. The trade-off between the infection spread and the patching costs is captured in a cost function, leading to an optimal control problem. In the single network case the optimal feedback controller is found by solving an associated Hamilton-Jacobi-Bellman equation. This control law is numerically compared to the proportional response strategy typically assumed by the epidemic model. In the higher dimensional multiple-networks case, the system is linearized to derive feedback controllers using pole-placement, linear quadratic regulator (LQR) optimal control, and $H^\infty$ optimal control, where the measurement errors in the number of infected clients are explicitly modeled. The resulting patching strategies are analyzed numerically and their results are compared.

*Key words:* Epidemic response, optimal control, LQR, H-infinity robust control.

## 1 Introduction

Self-spreading attacks on computer networks, such as worm epidemics, are expensive not only due to the damage they cause but also due to the challenge of prevent-

*Email addresses:* michael.bloem@nasa.gov (Michael Bloem), tansu.alpcan@telekom.de (Tansu Alpcan), tbasar@control.csl.uiuc.edu (Tamer Başar).

[1] M. Bloem was with the University of Illinois at Urbana-Champaign and partly supported by Deutsche Telekom AG Laboratories during this project.

ing and removing them. For example, the cost of the Code Red virus attack (Moore et al., 2002) of 2001 was estimated at $450 million in lost productivity, and an even larger $740 million in cleanup, monitoring, and system checking (Reuters, 2001).

Unsurprisingly, the problems of malware response and removal have caught the interest of the research community. A detailed analysis of the detection of a particular type of worm epidemic has been provided in (Rohloff and Başar, 2005). Specific worm epidemics such as Code Red and Slammer have been studied in (Moore et al., 2002, 2003a). The question of how to contain a worm epidemic on the Internet has been investigated in (Moore et al., 2003b). Recently, quarantine strategies relying on dividing the networks into subnetworks have been proposed in (Chen and Jamil, 2006) and (Cisco, 2006).

The well-known *classical epidemic model* has been applied extensively to medical epidemics (Hethcote, 2000). More recently, this model has been used to model the propagation of worm epidemics in computer networks as well as to evaluate potential responses to worm epidemics (Moore et al., 2003b; Chen and Jamil, 2006; Zou et al., 2003b; Rohloff and Başar, 2005; Zou et al., 2002, 2003a).

Optimization as well as decision and control theory have been applied to network security in several contexts. In (Alpcan and Başar, 2004) and (Alpcan and Başar, 2006), game theory was used for quantitative modeling and to develop decision-making strategies for network intrusion detection and response. Control theory has been applied to the problem of worm propagation by controlling the number of connections made by an infected host (Dantu et al., 2004). Optimization of system administrator time and efforts has been studied in (Bloem et al., 2006). A host-based defense system has been developed by making use of tools such as stochastic control and numerical optimization (Kreidl and Frazier, 2004). However, optimization and control theory have not been applied to the classical epidemic model to determine an optimal response to a worm epidemic.

### *1.1 Summary of Contributions*

The classical epidemic model and the Kermack-Mckendrick epidemic model have been used to model the spread of and response to worm epidemics in a computer network. In (Zou et al., 2002), the Kermack-Mckendrick model was extended to consider two additional factors: a time-varying infection rate and human countermeasures other than fixing infectious hosts, such as patching susceptible hosts. The resulting *two-factor worm model* was shown to have superior worm propagation modeling qualities. Furthermore, the Kermack-Mckendrick model has been used to evaluate a dynamic quarantine worm response (Zou et al., 2003a).

While these models have proven useful in modeling worm propagation and evaluating possible responses to worms, their utility can extend further. An extensive

literature in optimal control theory has developed tools for minimizing a cost subject to system dynamics such as those described by these epidemic models. By applying these optimal control tools to this problem, a provably optimal feedback response to a worm epidemic can be derived. This work demonstrates how optimal control theory can be used to find an optimal "feedback control dynamic quarantine system" by explicitly considering network system dynamics when deriving the worm response rather than only using them for simulation purposes (Zou et al., 2003a).

The optimization of patching response strategies to a worm epidemic is studied within a cost-benefit framework. The behavior of a worm epidemic in several networks is considered to obtain a multi-dimensional version of the classic epidemic model. While hosts infected with a worm are costly for a network, patching costs may also be significant Reuters (2001). The approaches studied here balance these costs against each other when multiple connected networks are threatened by malware.

Tools from optimal control theory are utilized to find appropriate malware response strategies. In order to apply optimal control theory to this model, the costs of infected hosts and of the effort required to patch them must be specified. This leads to an explicit quadratic cost function.

The resulting cost function is also used with the single and multiple network versions of the classical epidemic model differential equations to determine closed-form expressions for the feedback patching strategies in each case. Determining these expressions in the single network case involves the use of the Hamilton-Jacobi-Bellman equation to solve for a value function. In the multiple networks case, first the model differential equations are linearized. Then, controllers are derived using pole placement, LQR optimal control theory, and $H^\infty$ optimal control theory. The advantage of $H^\infty$ optimal control theory in particular is that it considers worst-case system and measurement noises. Indeed, previous research has shown that detecting malware is a significant challenge, justifying the need for a robust solution (Cohen, 1987).

After obtaining the patching response strategies as above, the paper analyzes them numerically and compares them with other heuristic patching strategies. One outcome here is that the proportional patching response rate is not necessarily optimal for the classical epidemic model. This appears to be the first application of optimal control theory to the classical epidemic model.

In the next section, the epidemic model is presented. Section 3 discusses optimal malware removal response for single and multiple networks and provides a stability analysis of the feedback control scheme. Section 4 contains the simulation results. The paper concludes with remarks in Section 5 and an appendix.

## 2 The Epidemic Model

The network analysis in this paper is based on the *classical epidemic model*. It uses a differential equation to model the spread of worm or virus in a computer network. For a single network, this classical model is described by

$$\dot{x}(t) = \beta \left[N - x(t)\right] x(t) - u(t), \tag{1}$$

where $u(t)$ is the number of patches applied at a given time, $x(t)$ is the number of infected hosts, $N$ is the number of hosts in the system, and $\beta$ is a parameter that captures the rate of spread of the epidemic and is referred to as the *pairwise rate of infection*.

This model can readily be extended to the multiple networks case. Given $M$ networks, let $x_i(t)$ denote the number of infected hosts in network $i$, where $i = 1, 2, \ldots, M$. Likewise, let $u_i(t)$ be the malware removal rate for network $i$. Let $\alpha$ be the *cross-network pairwise rate of infection*. Note that the more security measures are used between various networks, the smaller is $\alpha$ relative to $\beta$. Further, let $N_i$ denote the number of hosts on a particular network $i$. In general, because computers on a network are more likely to communicate with each other than those on different networks, and because individual networks typically have independent security measures, malware will be assumed to spread more rapidly within a network than between networks. Therefore, $\beta > \alpha$. Overall, one arrives at the model

$$\dot{x}_i(t) = \beta \left[N_i - x_i(t)\right] x_i(t) + \sum_{j=1, j\neq i}^{M} \alpha \left[N_i - x_i(t)\right] x_j(t) - u_i(t),$$

for $i = 1, \ldots, M$.

Another epidemic model considers the case where hosts that have had malware removed are no longer susceptible to malware infection. This model is referred to as the epidemic model *with removals*. When only one network is considered, this model becomes

$$\begin{aligned} \dot{x}_1(t) &= \beta \left[N - x_1(t) - x_2(t)\right] x_1(t) - u(t) \\ \dot{x}_2(t) &= u(t). \end{aligned} \tag{2}$$

Note that for each network, there are two state variables. The first is the number of infected hosts in the network. Its dynamics are very similar to those of the regular epidemic model. The second one keeps track of the number of hosts that have been patched and thus are no longer vulnerable to attack.

The epidemic model with removals can also be extended to the case where there

are multiple networks, as described above. This leads to the set of $2M$ coupled differential equations

$$\begin{aligned}\dot{x}_i(t) &= \beta\left[N_i - x_i(t) - x_{M+i}(t)\right] x_i(t) \\ &\quad + \sum_{j=1, j\neq i}^{M} \alpha\left[N_i - x_i(t) - x_{M+i}(t)\right] x_j(t) - u_i(t) \\ \dot{x}_{M+i}(t) &= u_i(t),\end{aligned} \tag{3}$$

for $i = 1, \ldots, M$. Here $x_1, \ldots, x_M$ are the number of infected hosts in networks $1, \ldots, M$, and $x_{M+1}, \ldots, x_{2M}$ are the number of patched hosts in networks $1$ through $M$, respectively.

Traditionally, when patching infected hosts, it is assumed that a particular proportion of them are patched at each time instance, i.e., $u_i(t) = \kappa x_i(t)$, for some $\kappa \in (0, 1)$, and for all $i = 1, \ldots, M$. The coefficient $\kappa$ is known as the *removal rate* of infectious hosts. Here, this will be referred to as a *proportional patching* controller.

In order to find an optimal control strategy, a cost must be chosen. Traditionally, quadratic costs are implemented on both the state (number of infected hosts) and control (patching rate). This structure is reasonable theoretically and mathematically tractable. Consider the cost function

$$J(\mathbf{x}(t), \mathbf{u}(t)) = \int_0^\infty \left[\mathbf{x}^T(t) Q \mathbf{x}(t) + \mathbf{u}^T(t) R \mathbf{u}(t)\right] dt, \tag{4}$$

where $\mathbf{x}$ and $\mathbf{u}$ are vectors of the state and control variables. In the classical epidemic model the $Q$ and $R$ matrices are chosen as diagonal matrices, with the $(i, i)$ entry designating the cost of an infected host in network $i$ (for $Q$) and a particular patching response rate in network $i$ (for $R$). In the epidemic model with removal, the $Q$ matrix is similarly structured but with no cost placed on states $x_{M+1}$ to $x_{2M}$, as these merely keep track of the number of patched hosts. The $R$ matrix is unchanged in this case.

## 3 Feedback Malware Removal Response

### *3.1 Single Network*

The optimal malware removal controllers for the single network epidemic models are presented in this subsection. Details of the derivation can be found in the

Appendix. The optimal feedback controller for the classical model is

$$\mu(x(t)) = \left(a(x) + \sqrt{a^2(x) + \frac{q}{r}}\right) x, \tag{5}$$

where $a(x) = \beta(N - x(t))$.

On the other hand, finding an explicit analytical solution of the optimal controller for the epidemic model with removal (2) proves to be difficult. Therefore, two approximations are considered. The first one is

$$\mu_{r1}(x(t)) \approx \left(\beta N + \sqrt{\beta^2 N^2 + \frac{q}{r}}\right) x_1. \tag{6}$$

The second one can be inferred by investigating and slightly altering the optimal solution derived for the classic epidemic model (5):

$$\mu_{r2}(x(t)) \approx \left(a(x_1 + x_2) + \sqrt{a^2(x_1 + x_2) + \frac{q}{r}}\right) x_1, \tag{7}$$

where $(x_1 + x_2)$ is simply substituted into $a(x)$ instead of $x$.

The resulting strategies (5), (6), and (7) each describe the (approximately) optimal patching or malware removal rate in the form of a feedback controller for a given set of cost parameters. All of these strategies differ significantly in form from a proportional patching controller.

### *3.2 Multiple Networks*

#### *3.2.1 Stabilizing Response*

One possible approach to this problem is to stabilize the system at a point where there are no infected machines. However, deriving a feedback controller to stabilize the nonlinear equations in the models (2) and (3) is not straightforward.

On the other hand, by studying the particular properties of these models it is possible to devise a strategy that results in a reasonable stabilizing feedback controller. One crucial observation is that each $x_i(t)$ has to be nonnegative. This leads to the insight that all of the cross terms and squared terms in the models (2) and (3) decrease the magnitude of the infection rates ($\dot{x}_i(t)$). Therefore, if these helpful squared and cross terms are disregarded, then one would be working with systems of equations that are actually more difficult to stabilize than the original models. Moreover, when these terms are disregarded, the models reduce to the same linear model

$$\dot{\mathbf{x}}(t) = A\mathbf{x}(t) + B\mathbf{u}(t), \tag{8}$$

where

$$A = \begin{bmatrix} \beta N_1 & \alpha N_1 & \alpha N_1 & \cdots & \alpha N_1 \\ \alpha N_2 & \beta N_2 & \alpha N_2 & \cdots & \alpha N_2 \\ \vdots & & \ddots & & \vdots \\ \vdots & & & \ddots & \vdots \\ \alpha N_M & \alpha N_M & \cdots & \alpha N_M & \beta N_M \end{bmatrix} \tag{9}$$

and $B$ is simply the negative identity matrix of dimension $M \times M$.

Notice that the epidemic models have the inherent physical constraints

$$0 \leq x_i \leq N_i, \quad i = 1, \ldots, M.$$

However, if $x_i = N_i$ for any $i$, then $\dot{x}_i < 0$ for the original nonlinear system (2) under the condition $\mathbf{u} < 0$ which is discussed in detail in the next section. In other words, the trajectory leaves the boundary $[N_1, \ldots, N_M]$ immediately. A similar argument can also be made for (3). Therefore, it is ignored for the simplified linear system (8) and focus on the constraint set

$$x_i \geq 0, \quad i = 1, \ldots, M. \tag{10}$$

Under this set of constraints (10), the constrained linear model becomes

$$\dot{x}_i = \begin{cases} [A\mathbf{x} + B\mathbf{u}]_i & \begin{cases} \text{if } x_i > 0 \text{ or} \\ \text{if } [A\mathbf{x} + B\mathbf{u}]_i \geq 0 \text{ and } x_i = 0 \end{cases} \\ 0 & \text{else} \end{cases} \tag{11}$$

for all $i = 1, \ldots, M$.

While it is known that a linear feedback controller can stabilize the linear model (8), whether such a controller also stabilizes the nonlinear models (2), (3), and (11) is a question which is investigated in the next section.

### *3.3 Stability Analysis*

The stability of the system (11) is now shown under the set of boundary constraints (10) when controlled by the linear feedback controller

$$\mathbf{u}_s = -K\mathbf{x}, \tag{12}$$

where $K$ is the feedback matrix. Obviously, the origin constitutes the unique equilibrium for this system.

A sufficient condition for stability can be found by considering the special structure of this problem. Since it is known that the components of $\mathbf{x}$ will never become negative, the closed-loop system matrix does not even need to be Hurwitz. A sufficient condition for stability is that the diagonal elements are negative and that the non-diagonal elements are nonpositive. This condition is easy to verify upon inspection of the closed-loop matrix.

**Theorem 1** *Given a nonlinear system of the form in (11) under a control of the form* $\mathbf{u} = -K\mathbf{x}$*, the system is stable if the closed-loop matrix* $(A - BK)$ *has negative diagonal entries and nonpositive off-diagonal entries.*

**Proof** In the context of the system (11), components of $\mathbf{x}$ can only be positive or zero. In this case the diagonal entries of the closed-loop matrix $(A - BK)$ are assumed to be negative and the off-diagonal entries are nonpositive. Clearly, this implies that each component of $\dot{\mathbf{x}}$ is nonpositive.

However, this is not enough to guarantee stability. It must also be known that any positive component of $\mathbf{x}$ will decrease to zero. This is ensured because the diagonal elements of the closed-loop system matrix are assumed to be negative. Therefore all positive components of $\mathbf{x}$ will decrease to zero at a rate faster than or equal to that specified by the corresponding diagonal entry in $(A - BK)$. $\square$

The feedback controllers derived from LQR and $H^\infty$ optimal control may not have this property. Nevertheless, in many situations controllers derived with optimal control theory will meet this condition and therefore can be used to stabilize the system under consideration. The closed-loop matrices for three dimensional versions of all of the simulations in this paper also met this requirement. Cases where the numbers of hosts differ between different networks or the costs on control differ between networks sometimes led to closed-loop matrices with slightly positive off-diagonal elements. However, even in cases where the difference in the numbers of hosts and costs on control were large ($10^5$ times larger), the positivity of the off-diagonal terms was small and setting these terms to zero would stabilize the system without significantly affecting its performance.

It is next shown that the stabilizing controller (12) which meets the conditions of Theorem 1 will stabilize the actual nonlinear systems (2) and (3). Given that $\mathbf{x} \geq 0$, the nonlinear terms in these equations will only decrease the magnitude of the components of $\dot{\mathbf{x}}$. If $\mathbf{x}$ could become negative this may destabilize the system. However, in this case it only adds additional negative drift, leading to faster stabilization. In conclusion, the stabilizing condition in Theorem 1 ensures stability over the entire state space, even in the nonlinear case.

### 3.3.1 Linear Quadratic Regulator Optimal Response

Determining the optimal malware removal strategy relative to the cost (4) for the epidemic models (2) and (3) is nontrivial. When multiple networks are considered, even if only two networks are studied and a very simple form for the value function is assumed, the approach presented in the Appendix leads to an over-determined set of non-linear equations, which is not tractable. Therefore a more tractable but sub-optimal approach to this problem was developed based on the linear model (8).

By making use of this linear model (8) and the quadratic cost function (4), one arrives at the well-studied linear quadratic regulator (LQR) optimal control problem, whose optimal solution can readily be obtained.

$$\mathbf{u}_o(t) = -R^{-1}B^T P\mathbf{x}(t), \tag{13}$$

where $P$ is the positive definite solution to the algebraic Riccati equation (ARE)

$$A^T P + PA - PBR^{-1}B^T P + Q = 0. \tag{14}$$

Here the fact that $(A, B)$ is controllable is used because $B$ is the negative identity matrix. Also, note that because $Q > 0$, the solution to (14), $P > 0$, exists and is unique.

### 3.3.2 $H^\infty$-Optimal Response

While the model (8) developed in Sub-section 3.3.1 is useful, it involves some assumptions. First, the nonlinear terms in the more precise models (2) and (3) have been ignored. Moreover, availability of perfect measurement of the number of infected hosts in each network has been assumed. Finally, the original epidemic models (2) and (3) themselves only approximate malware propagation.

To capture these approximations and imperfections in the linear model (11), one can alter this model to include a noise term, with the problem then cast in an $H^\infty$ optimal control framework. Let $\delta_i = [A\mathbf{x} + B\mathbf{u} + D\mathbf{w}_a]_i$. Then

$$\dot{x}_i = \begin{cases} \delta_i & \begin{cases} \text{if } x_i > 0 \text{ or} \\ \text{if } \delta_i \geq 0 \text{ and } x_i = 0 \end{cases} \\ 0 & \text{else.} \end{cases} \tag{15}$$

Here, $\mathbf{w}_a(t)$ is a noise term that accounts for model assumptions and approximations. The $D$ matrix describes how this noise term impacts the dynamics of $\mathbf{x}(t)$ and will be set to the identity matrix. In addition, a measurement error can be introduced: if $\mathbf{y}(t)$ is the measurement of the number of infected hosts, then

$$\mathbf{y}(t) = \mathbf{x}(t) + \mathbf{w}_n(t). \tag{16}$$

which says that the noise vector $\mathbf{w}_n(t)$ impacts the measurement of the number of infected hosts on each network (each element of $\mathbf{y}(t)$). In order to develop the $H^\infty$ optimal controller, also the *controlled output* has to be defined, which is

$$\mathbf{z}(t) = H\mathbf{x}(t) + G\mathbf{u}(t). \tag{17}$$

It is assumed here that $G^TG$ and $H^TH$ are positive definite and that $H^TG = 0$. This says that there is no cost placed on the product of patching response and infected hosts, although each of those quantities individually contributes to the cost. So that the cost $\|\mathbf{z}\|^2$ (defined below) corresponds to the cost (4) for the LQR controller, let $Q = H^TH$ and $R = G^TG$. A few other constraints that must be met for this $H^\infty$ optimal control theory to apply are that $(A,\ B)$ and $(A,\ D)$ be stabilizable, and $(A,\ H)$ and $(A,\ I)$ be detectable. Also define $\mathbf{w} := [\mathbf{w}_a^T\ \mathbf{w}_n^T]^T$ as the total disturbance to the system. Let the cost ratio used in the $H^\infty$ analysis be

$$L(\mathbf{x}, \mathbf{u}, \mathbf{w}) = \frac{\|\mathbf{z}\|}{\|\mathbf{w}\|}, \tag{18}$$

where $\|\mathbf{z}\|^2 := \int_{-\infty}^{\infty} |\mathbf{z}(t)|^2 dt$ and a similar definition applies to $\|\mathbf{w}\|^2$. This captures the proportional changes in $\mathbf{z}$ due to changes in $\mathbf{w}$.

$H^\infty$ optimal control theory also produces a performance factor (the $H^\infty$ norm) that one can guarantee will be met, as described in Section 1. This norm can be thought of as the worst possible value for the cost $L$. The lowest possible value of $\gamma$ is

$$\gamma^* := \inf_{\mathbf{u}} \sup_{\mathbf{w}} L(\mathbf{u}, \mathbf{w}) \tag{19}$$

which can also be viewed as the optimal performance level in the $H^\infty$-control context.

In order to actually solve for the optimal controller $\mu(\mathbf{y})$, a corresponding differential game is defined, which is parameterized by $\gamma$. The optimal worst case controller $\mathbf{u}_w = \mu_\gamma(\mathbf{y})$ can be determined from this differential game for any $\gamma > \gamma^*$. It is given by (Başar and Bernhard, 1995) as

$$\mu_\gamma(\mathbf{y}) = -(G^TG)^{-1}B^T\bar{Z}_\gamma\hat{\mathbf{x}}, \tag{20}$$

where $\bar{Z}_\gamma$ is solved from

$$A^TZ + ZA - Z(B(G^TG)^{-1}B^T - \gamma^{-2}DD^T)Z + Q = 0, \tag{21}$$

as its unique minimal positive definite solution, and $\hat{\mathbf{x}}$ is given by

$$\begin{aligned}\dot{\hat{\mathbf{x}}} = &\left[A - (B(G^TG)^{-1}B^T - \gamma^{-2}DD^T)\bar{Z}_\gamma\right]\hat{\mathbf{x}} \\ &+ \left[I - \gamma^{-2}\bar{\Sigma}_\gamma\bar{Z}_\gamma\right]^{-1}\bar{\Sigma}_\gamma(\mathbf{y} - C\hat{\mathbf{x}}),\end{aligned} \tag{22}$$

where $\bar{\Sigma}_\gamma$ is the unique minimal positive definite solution of

$$A\Sigma + \Sigma A^T - \Sigma(I - \gamma^{-2}H^T H)\Sigma + DD^T = 0. \tag{23}$$

Note that $\gamma^*$ is the smallest $\gamma$ such that the spectral radius $\rho(\bar{\Sigma}_\gamma \bar{Z}_\gamma) < \gamma^2$.

The linear $H^\infty$-optimal feedback controller (20) provides a robust malware response or epidemic removal strategy based on the estimate of the number of infected hosts. It can be calculated off-line using only the linear quadratic system model.

## 4 Simulations and Results

### *4.1 Single Network*

First for the single network classical epidemic model, the optimal controller is compared with the proportional controller $u(x(t)) = \kappa x(t)$. The value for $\beta$ is assumed to be $5.6 \times 10^{-5}$. This is the estimated value of $\beta$ for the SQL Slammer worm (Liljenstam et al., 2003). The cost parameter ratio $q : r$ is varied between 100:1 and 1:1. The network in question is assumed to have 500 hosts, 100 of which are initially infected. The $\kappa$ value for this particular proportional controller is optimistically set to 1. This means that, initially, this proportional controller will be patching at a rate of 100 patches per time unit. The optimal controller (5) patches at a significantly higher rate than this when there is a relatively high cost put on the state. In order to compare these controllers more fairly, a bounded controller inspired by the optimal controller (5) is considered, on which a maximum patching rate of $u_{\max} = 100$ is imposed. The bounded controller is identical to (5) except in cases where the control specified by (5) exceeds $u_{\max}$, in which case the bounded controller applies $u_{\max}$ as the input.

The cost results for the scenario described are shown in Table 1. Note that this proportional controller performs well when the infected hosts and patching have equal cost weights. However, it incurs a fivefold larger cost when the cost ratio is set to 100:1. Of course this behavior would change with a different $\kappa$ parameter value. The bounded optimal response performs better than this proportional response but not as well as the optimal response. When the cost ratio is 100:1, the optimal patching strategy starts patching at a very high rate (around 1000 hosts per time unit). The bounded optimal controller performs significantly worse than the optimal controller, but significantly better than this proportional controller.

The epidemic model with removals is next evaluated. Here a less aggressive attack with $\beta = 2.8 \times 10^{-5}$ and the proportional response, the two approximations of the

Table 1
Costs of Patching Strategies for the Classical Model ($\times 10^3$)

| Cost Ratio | Proportional | Optimal | Bounded Optimal |
|---|---|---|---|
| 100:1 | 518.2 | 101.7 | 349.8 |
| 10:1 | 56.44 | 32.01 | 43.08 |
| 3:1 | 20.52 | 17.61 | 18.50 |
| 1:1 | 10.26 | 10.26 | 10.26 |

Table 2
Costs of Patching Strategies for the Removal Model ($\times 10^3$)

| Cost Ratio | Proportional | $\mu_{r1}$ | $\mu_{r2}$ | Bounded $\mu_{r1}$ |
|---|---|---|---|---|
| 100:1 | 511.4 | 101.6 | 101.6 | 346.6 |
| 10:1 | 55.70 | 31.88 | 31.88 | 42.69 |
| 3:1 | 20.26 | 17.48 | 17.48 | 17.99 |
| 1:1 | 10.13 | 10.13 | 10.13 | 10.13 |

optimal controller (6) and (7), and a bounded version of the linear approximation (6) are simulated. The cost results for this scenario can be seen in Table 2.

The two approximations to the optimal feedback controller incur nearly identical costs. Once again, when patching costs are relatively high (e.g. a $q : r$ cost parameter ratio of 1:1) this proportional controller behaves almost as well as the optimal controller.

The response of the proportional, linear optimal, and bounded linear optimal controllers are shown in Fig. 1, Fig. 2, and Fig. 3, respectively. These graphs are for the cost ratio 3:1. The proportional controller does not patch at a high enough rate for long enough at the start of the simulation, leading to its inferior performance. The approximate optimal solution patches at a very high rate initially, thereby incurring very high patching costs for some time. This indicates the intuitive result that system administrators should spare no expense at the initial phase of an epidemic. Then, when the worm attack is under control, administrators can allocate fewer resources in order to balance the costs of removal against the now lower costs of infections.

In conclusion, while the heuristic approach of patching a fixed proportion of infected hosts may perform well in certain circumstances, the optimal controller always achieves the lowest cost. At the same time it provides a framework for quantifying the system trade-offs and expressing preferences.

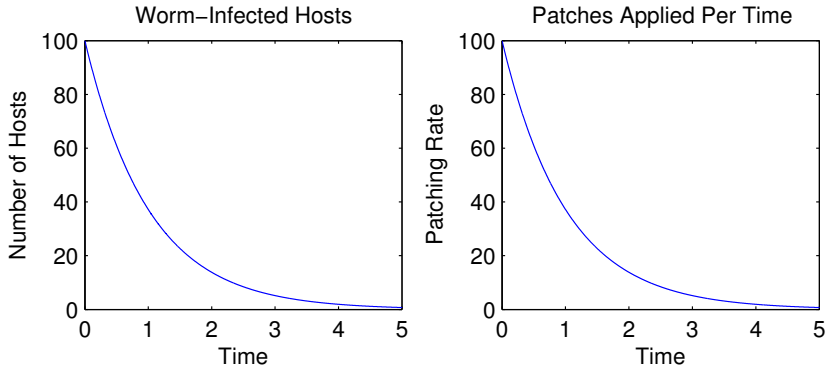


Fig. 1. Behavior of proportional patching response to worm infections with removal of patched hosts (3:1 cost structure).

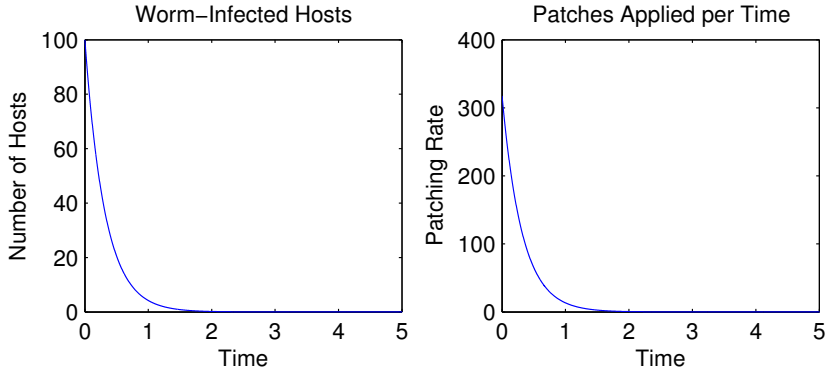


Fig. 2. Behavior of approximate optimal feedback patching response $\mu_{r1}$ (6) to worm infections with removal of patched hosts (3:1 cost structure).

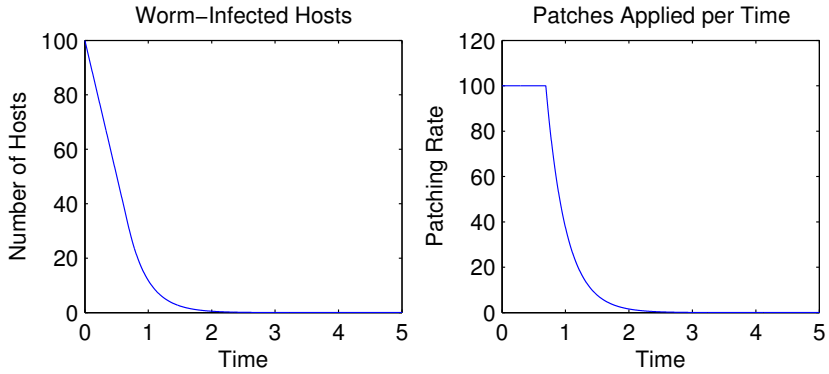


Fig. 3. Behavior of bounded approximate optimal feedback patching response $\mu_{r1}$ (6) to worm infections with removal of patched hosts (3:1 cost structure).

### *4.2 Multiple Networks*

Simulations of the models (2) and (3) with the various controllers were performed in Matlab. Two networks are simulated, each containing 500 hosts. Initially network 1 has 250 infected hosts while network 2 only has 100 infected hosts. The $\beta$

parameter is set again to $5.6 \times 10^{-5}$, the estimated value of $\beta$ for the SQL Slammer worm. The $\alpha$ parameter was set to $\beta/4$, reflecting an assumption that worms will spread more slowly between networks if proper security measures are taken for each network individually. For this case $Q$ and $H^T H$ are set at 0.01 times the identity matrix, and the $R$ and $G^T G$ are taken as the identity matrix. This reflects a situation where patching is more expensive than an infected host.

Moreover, a noise term was added to the system dynamics when simulating each of the models, identical to the noise term in the $H^\infty$ model (15). This normally distributed noise term is meant to capture some of the imperfections in the models.

### *4.3 Stabilizing Response*

The stabilizing controller is derived simply by placing in the left half plane the poles of the closed-loop system that results from the application of a linear feedback controller to the linearized system model (8). The resulting closed-loop system matrix is also verified to meet the conditions in Theorem 1. Therefore, one can choose exactly where to place the poles and achieve varying degrees of stability. Note the proportional response for the single dimensional version of this problem is actually a pole placement strategy. For these simulations the poles are placed at a few locations: -1, -0.5, and -0.1. Fig. 4 shows the results of a simulation of the epidemic model with removals when the poles are placed both at -0.5.

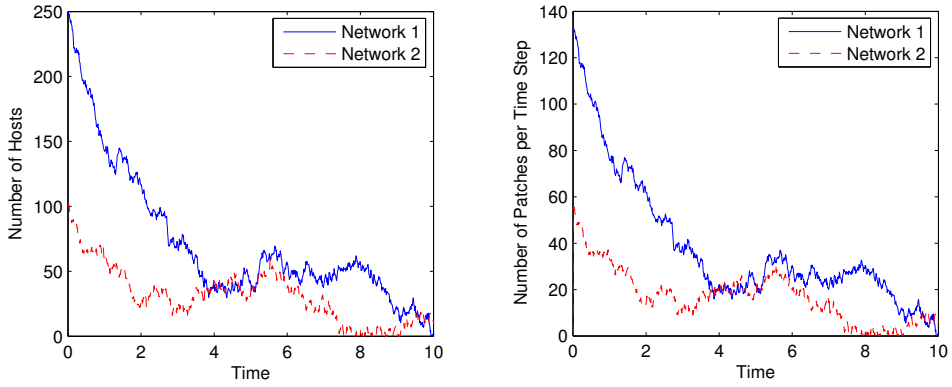


(a) (b)

Fig. 4. (a) The number of infected hosts and (b) the patching rate over time when the feedback controller that places the closed-loop system poles at -0.5 is applied to epidemic model with removals (3).

When this controller is applied to the various models under consideration the resulting cost values are found to be very similar. The cost results of these simulations are given in Table 3. The graphs of the simulations of each of the systems are also nearly identical, implying that disregarding the nonlinear terms in order to derive the controller was reasonable. Note that the linear model has a higher cost than the

Table 3
Cost of Models Under Stabilizing Controllers

| Model | Cost (-0.1) | Cost (-0.5) | Cost (-1) |
|---|---|---|---|
| Linear | 12,470 | 30,730 | 47,520 |
| Classical Epidemic | 11,310 | 30,170 | 47,120 |
| Epidemic with Removals | 11,080 | 29,890 | 46,870 |

other two models because it does not consider the nonlinear terms, which actually contribute to stabilizing the system. The epidemic model with removals contains the most helpful nonlinear terms, and thus it achieves the lowest cost with this controller.

### *4.4 LQR Optimal Response*

The results of a simulation of this scenario for the epidemic model with removal are shown in Fig. 5. The number of infected hosts in each network are shown in Fig. 5(a) while the patching rates for each network are shown in (b). This controller is somewhat less aggressive than the stabilizing controller simulated in Fig. 4.

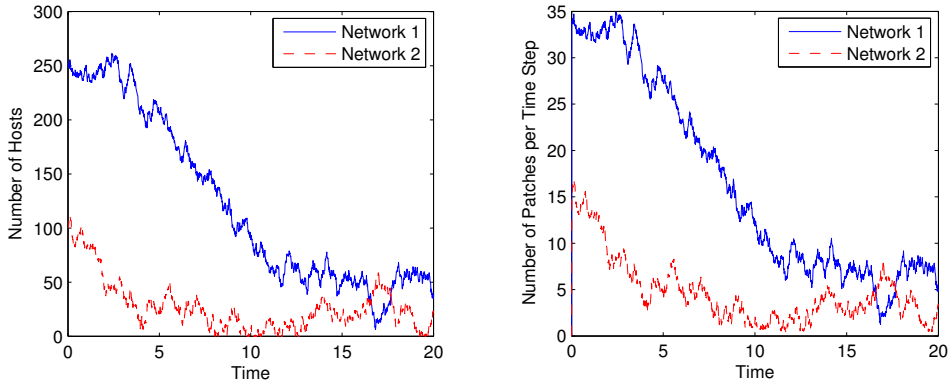


(a) (b)

Fig. 5. (a) The number of infected hosts and (b) the patching rate over time when the LQR optimal controller is applied to epidemic model with removals (3).

When this controller is applied to the various models under consideration, the resulting cost values are found to be very similar. These costs are shown in Table 4. The graphs of the simulations of each of the systems are also nearly identical, implying that disregarding the nonlinear terms in order to derive the controller was reasonable. The linear model has a higher cost than the other two models because it ignores the helpful nonlinear terms. The epidemic model with removals contains the most nonlinear terms, and thus it achieves the lowest cost with this controller.

Interestingly, these costs are slightly higher than the costs of the stabilizing con-

Table 4
Cost of Models Under LQR-Optimal Controller

| Model | Cost | Cost with Measure Noise |
|---|---|---|
| Linear | 14,390 | 33,080 |
| Classical Epidemic | 12,980 | 31,340 |
| Epidemic with Removals | 12,700 | 31,120 |

troller with poles placed at $-0.1$. This occurs because the LQR optimal controller is designed assuming a linear model when actually the model has nonlinearities that help stabilize the system. Thus the LQR controller is more aggressive than is optimal (poles are $-0.1059$ and $-0.1022$), leading to sub-optimal performance. Other contributing factors to this behavior are errors introduced in the discretization of continuous time theory for simulation purposes and variations in the noise faced by each system. This additional cost of about $15\%$ is the magnitude of the cost increase resulting from using a linearized version of the system model to derive an optimal controller. In practice, the best approach would be to tune the controller resulting from LQR optimal control theory, as it will be close to the actual optimal controller for the nonlinear system.

This system is also simulated under measurement noise. Note that the system cost more than doubles for each model (see Table 4). Moreover, the feedback controller becomes highly oscillatory, as seen in Fig. 6.

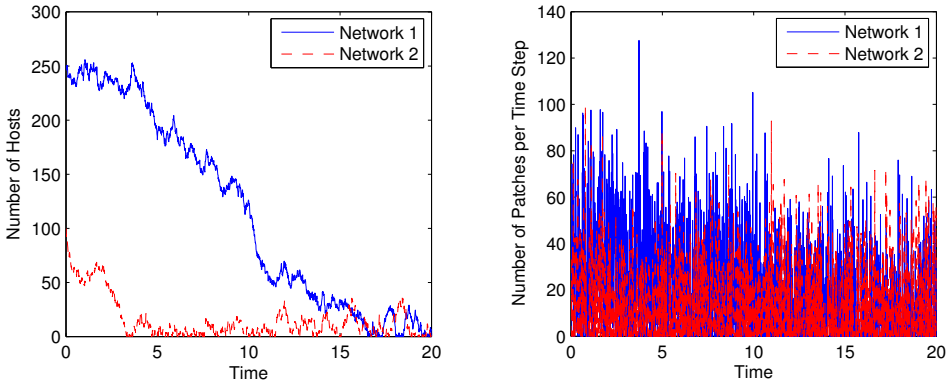


(a) (b)

Fig. 6. (a) The number of infected hosts and (b) the patching rate over time when the LQR optimal controller is applied to the epidemic model with removals and when there are noisy measurements (3).

### *4.5 $H^\infty$ Optimal Response*

The simulations of the $H^\infty$ optimal controller are unique because the controller is designed to operate when there is the worst-case possible noise on the system

Table 5
Cost of Models Under $H^\infty$ Optimal Controller

| Model | Cost |
|---|---|
| Linear | 59,530 |
| Classical Epidemic | 59,270 |
| Epidemic with Removals | 59,150 |

measurements and dynamics. Thus, in these simulations noisy state measurements and also the $H^\infty$ state estimate (22) are incorporated.

This controller will lead to unnecessarily high costs due to over-aggressive responses. This aggressive response can be seen in Fig. 7, where the application of the $H^\infty$ optimal controller to the epidemic model with removals is simulated. There are very few infected hosts remaining even after just 5 time units. Relatively high patching rates were called for in order to generate this aggressive response. However, the control applied in this case is relatively stable and robust in the face of noise, as compared with the control applied by the LQR optimal controller shown in Fig. 6.

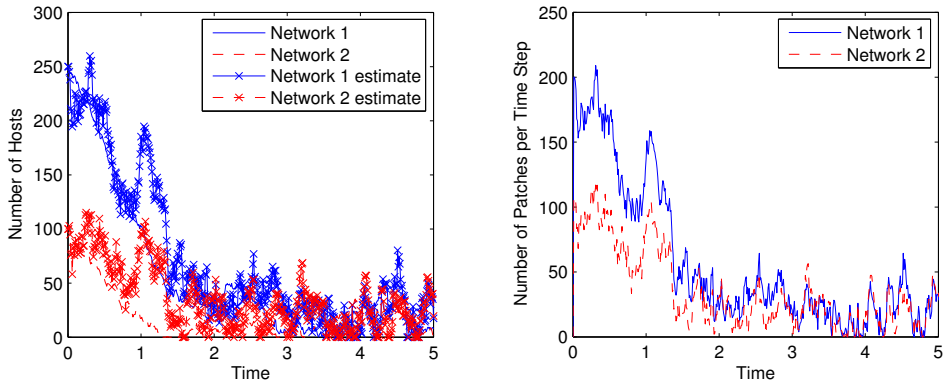


(a) (b)

Fig. 7. (a) The number of infected hosts and (b) the patching rate over time when the $H^\infty$ optimal controller is applied to epidemic model with removals (3).

The costs under this controller for each of the three models are shown in Table 5. The aggressive nature of the $H^\infty$ optimal controller and the noisy measurements lead to costs that are significantly higher than those resulting from the use of the LQR optimal controller. While these costs are high, the $H^\infty$ optimal controller offers a cost ratio guarantee that the other controllers cannot (here $\gamma^* = 1.05$). Again it can be seen that the nonlinear terms in the system do not significantly change the system dynamics.

Table 6
Cost Benefit of Secure Subnetworks for Epidemic Model with Removals

| $(x_1(0), x_2(0))$ | Insecure Subnets | Secure Subnets | % Cost Gain |
|---|---|---|---|
| (250, 0) | 7,835 | 7,492 | 4.38% |
| (400, 0) | 19,180 | 18,320 | 4.48% |
| (250, 250) | 16,480 | 14,630 | 11.2% |
| (400, 400) | 37,483 | 34,324 | 8.43% |

### *4.6 Cost Implications of Secure Networks*

This sub-section briefly discusses the impact on system costs that structuring a network so as to have multiple subnetworks that are securely isolated from each other. The LQR optimal controller is simulated but this time $\alpha$ is set equal to $\beta$, which would correspond to a situation where a network is divided into subnetworks but where a worm or virus is just as likely to spread outside of a subnetwork as it is to spread inside. The costs, and those from the case where $\alpha = \beta/4$, are shown in Table 6.

Clearly, there is a cost incentive to implementing secure subnetworks, but the cost savings are only about 4-11% in this scenario. This may not be significant enough to justify the additional expense of building securely divided subnetworks, depending on the magnitude of security-related expenditures and the cost of securing subnetworks.

## 5 Conclusions

Some solutions to the malware removal problem for a single infected network have been presented. The classical epidemic models assume a response strategy that is not optimal. By constructing a cost function and utilizing basic optimal control methods, an optimal feedback controller has been derived for malware removal and patching. This controller not only allows for different weights to be assigned to infected hosts and patching efforts, but also performs optimally in relation to these costs, and hence is more efficient than the traditional proportional response strategy.

This theory was also extended to the case where multiple networks are designed to inhibit the flow of malware between them. By linearizing the nonlinear system of differential equations describing this situation it was possible to derive and numerically evaluate several feedback malware removal controllers. While fine-tuned stabilizing controllers can out-perform optimal controllers in some cases, optimal controllers are more flexible to changes in the cost structure. While designing secure subnetworks does lead to costs savings, these savings may not be significant

in some scenarios.

Several extensions to this theory exist. In some situations parameters like $\alpha$, which capture the degree to which networks are quarantined from each other, can be control variables. Optimal control theory could be applied to this problem to determine to what extent networks should be quarantined in a given situation. Many variations of the epidemic models considered in this paper have been developed to more realistically describe the spread of infectious diseases and computer viruses as well as to capture other possible response actions (Hethcote, 2000; Zou et al., 2002). Optimal control theory could be applied to these models.

## Acknowledgment

The authors would like to thank the Boeing Corporation and Deutsche Telekom, AG for their support of this research, for the former through the Information Trust Institute at the University of Illinois at Urbana-Champaign. An earlier, more concise version of this paper was presented at the 46th IEEE Conference on Decision and Control, New Orleans, December 12-14, 2007, with the same title.

## Appendix

The optimal feedback controller for the single network classical model was derived with standard optimal control methods. This derivation is standard, and utilizes the Hamilton-Jacobi-Bellman (HJB) equation and continuous time dynamic programming, as described, for example, in Chapter 4, Section 5 of Sage and White III (1977). The entirety of the derivation is not shown for simplicity and to conserve space.

In order to apply the HJB equation first the optimal controller $u^*(t)$ is derived and substituted back into the equation. Differentiating the HJB equation with respect to $u$, setting the result equal to zero, and solving for $u^*(t)$ yields

$$u^*(t) = \frac{p(t)}{2r}, \tag{24}$$

where $p(t)$ is the co-state function. Replacing $p$ with $V_x$ and substituting the result into the HJB equation leads to a quadratic equation in $V_x$, which can be solved for $V_x$:

$$V_x(x(t), u(t)) = 2ra(x)x(t) + 2r\sqrt{(a(x)x(t))^2 + \frac{q}{r}x^2(t)}, \tag{25}$$

where $a(x(t)) = \beta(N - x(t))$. Finally $p$ is replaced in (24) with $V_x$ to obtain the optimal feedback controller (5).

The derivation of the optimal controller for the single network epidemic model with removal (2) is similar. In this case the HJB equation becomes

$$\begin{aligned} 0 = \min_u \{ & q x_1^2(t) + r u^2(t) \\ & + V_{x_1} \left[ \beta \left[ N - x_2(t) - x_1(t) \right] x_1(t) - u(t) \right] \\ & + V_{x_2} \left[ u(t) \right] \}. \end{aligned} \tag{26}$$

The minimizing $u^*(t)$ becomes

$$u^*(t) = \frac{V_{x_1} - V_{x_2}}{2r}, \tag{27}$$

Substituting this back into (26) yields

$$0 = q x_1^2(t) - \frac{1}{4r}(V_{x_1} - V_{x_2})^2 + V_{x_1} \left[ \beta (N - x_1(t) - x_2(t)) x_1(t) \right].$$

Solving for $V_{x_1}$ and $V_{x_2}$ explicitly proves difficult in this case. Instead, an approximate solution for the optimal controller with removal, $\mu_r(x(t))$, is obtained. The approach is to approximate $V$ as

$$V = k_0 + k_1 x_1 + k_2 x_2 + k_3 x_1^2 + k_4 x_1 x_2 + k_5 x_2^2. \tag{28}$$

The partial derivatives of this assumed form with respect to $x_1$ and $x_2$ are substituted back into (28). When the coefficients of the various terms in this equation ($x_1$, $x_2$, $x_1^2$, etc.) equal to zero, a system of 9 equations in 5 variables is produced. Investigating these equations and the solutions they provide yields two possible solutions. Numerical analysis of these two solutions (as in section 4.1) confirms that one is the approximate optimal feedback controller (6).

## References


Alpcan, T., Başar, T., December 2004. A game theoretic analysis of intrusion detection in access control systems. In: Proc. 43rd IEEE Conf. Decision and Control. Paradise Island, Bahamas, pp. 1568–1573.

Alpcan, T., Başar, T., July 2006. An intrusion detection game with limited observations. In: Proc. of 12th International Symposium on Dynamic Games and Applications. Sophia Antipolis, France.

Başar, T., Bernhard, P., 1995. H$^\infty$-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach, 2nd Edition. Birkhäuser, Boston, MA.

Bloem, M., Alpcan, T., Başar, T., December 2006. Intrusion response as a resource allocation problem. In: Proc. 45th IEEE Conf. Decision and Control. San Diego, CA, USA.

Chen, T. M., Jamil, N., June 2006. Effectiveness of quarantine in worm epidemics. In: Proc. of IEEE ICC 2006. Istanbul, Turkey, pp. 2142–2147.

Cisco, 2006. NAT and stateful inspection in Cisco IOS firewall. white paper, http://www.cisco.com/en/US/tech/tk648/tk361/technologies_white_paper 09186a0080194af8.shtml.

Cohen, F., 1987. Computer viruses: Theory and experiments. Computers and Security.

Dantu, R., Cangussu, J., Yelimeli, A., April 2004. Dynamic control of worm propagation. In: Proc. of IEEE Conference on Information Technology: Coding and Computing. Vol. 1. pp. 419–423.

Hethcote, H. W., 2000. The mathematics of infectious diseases. SIAM Review 42 (4), 599–653.

Kreidl, O., Frazier, T., 2004. Feedback control applied to survivability: A host-based autonomic defense system. IEEE Transactions on Reliability 53 (1).

Liljenstam, M., Nicol, D., Berk, V., Gray, R., 2003. Simulating realistic network worm traffic for worm warning system design and testing. In: Proc. of ACM Workshop on Rapid Malcode. Washington, DC, pp. 24–33.

Moore, D., Paxson, V., Savage, S., Shannon, C., Staniford, S., Weaver, N., July-Aug. 2003a. Inside the slammer worm. IEEE Security & Privacy Magazine 1, 33–39.

Moore, D., Shannon, C., Claffy, K., 2002. Code-red: A case study on the spread and victims of an internet worm. In: Proc. of ACM SIGCOMM Workshop on Internet Measurement. Marseille, France, pp. 273–284.

Moore, D., Shannon, C., Voelker, G., Savage, S., April 2003b. Internet quarantine: Requirements for containing self-propagating code. In: Proc. 22nd IEEE Infocom. Vol. 3. pp. 1901 – 1910.

Reuters, Aug. 1, 2001. The cost of 'code red': $1.2 billion. USA Today, http://usatoday.com/tech/news/2001-08-01-code-red-costs.htm.

Rohloff, K., Başar, T., 2005. The detection of RCS worm epidemics. In: Proc. of ACM Workshop on Rapid Malcode. Fairfax, VA, pp. 81–86.

Sage, A., White III, C., 1977. Optimum Systems Control, 2nd Edition. Prentice Hall, Englewood Cliffs, NJ.

Zou, C., Gong, W., Towsley, D., October 2003a. Worm propagation modeling and analysis under dynamic quarantine defense. In: Proceedings of ACM Workshop on Rapid Malcode (WORM). Washington, DC.

Zou, C., Gong, W., Towsley, D., 2003b. Worm propogation modeling and analysis under dynamic quarantine defense. In: Proc. of ACM Workshop on Rapid Malcode. Washington, DC, pp. 51–60.

Zou, C. C., Gong, W., Towsley, D., November 2002. Code red worm propagation modeling and analysis. In: Proceedings of ACM Conference on Computer and Communications Security. Washington, DC.