Optimal and Robust Epidemic Response for Multiple Networks

Michael Bloem[1,2]
Coordinated Science Lab
University of Illinois
1308 West Main Street
Urbana, IL 61801, USA
mbloem2@control.csl.uiuc.edu

Tansu Alpcan[1]
Deutsche Telekom Laboratories
Technische Universität Berlin
Ernst-Reuter-Platz 7, 10587, Germany
tansu.alpcan@telekom.de

Tamer Başar[2]
Coordinated Science Lab
University of Illinois
1308 West Main Street
Urbana, IL 61801, USA
tbasar@control.csl.uiuc.edu

***Abstract*— We study 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. We capture the trade-off between the infection spread and the patching costs in a cost function, leading to an optimal control problem. We linearize the system to derive feedback controllers using pole-placement, linear quadratic regulator (LQR) optimal control, and H^∞ optimal control, where we explicitly model measurement errors in the number of infected clients. The resulting patching strategies are analyzed numerically and their results are compared.**

I. 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 preventing and removing them. For example, the cost of the Code Red virus attack [1] of 2001 was estimated at $450 million in lost productivity, and an even larger $740 million in cleanup, monitoring, and system checking [2].

Unsurprisingly, the problems of malware response and removal have caught the interest of the research community. The well-known *classical epidemic model* has been applied extensively to medical epidemics [3]. More recently, this model has been used to study the propagation of and response to worm epidemics in computer networks [4]–[7].

A detailed analysis of the detection of a particular type of worm epidemic has been provided in [7]. Specific worm epidemics such as Code Red and Slammer have been studied in [1], [8]. The question of how to contain a worm epidemic on the Internet has been investigated in [4]. Recently, quarantine strategies relying on dividing the networks into subnetworks have been proposed in [5] and [9].

Optimization as well as decision and control theory have been applied to network security in several contexts. In [10] and [11], 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 [12]. Optimization of system administrator time and efforts has been studied in [13].

[1]Research supported in part by Deutsche Telekom AG.

[2]Research supported in part by a grant from the Boeing company, through the Information Trust Institute at the University of Illinois at Urbana-Champaign.

A. Summary of Contributions

In this paper, we study the optimization of patching response strategies to a worm epidemic. While hosts infected with a worm are costly for a network, patching costs are also significant [2]. We study how to balance these costs against each other when multiple connected networks are threatened by malware.

The classical epidemic model has been successfully applied to the spread of and response to worm epidemics in a single computer network. We consider the behavior of a worm epidemic in several networks to obtain a multi-dimensional version of the classic epidemic model.

To find appropriate malware response strategies, we utilize tools from optimal control theory. In order to apply optimal control theory to this model, we must first explicitly specify the costs of infected hosts and of the effort required to patch them. The resulting cost function is used in conjunction with the multiple network version of the classical epidemic model differential equations to determine closed-form expressions for the feedback patching strategies. To derive these strategies we first linearize the model differential equations. Then we can derive controllers using pole placement, LQR optimal control theory, and H^∞ optimal control theory. The advantage of H^∞ optimal control theory in particular is that it considers worst-case system and measurement noise, which will capture the model inaccuracies and noisy measurements that impact the application of this theory.

We numerically analyze these patching response strategies and compare them with other heuristic patching strategies. We show that the proportional patching response rate is not necessarily optimal for the classical epidemic model. To the best of the knowledge of the authors, this is the first application of optimal control theory to the worm epidemic response problem and to the classical epidemic model.

In the next section, we introduce the epidemic model. Then, in Section III, we derive malware removal control laws using stability theory, LQR optimal control theory, and H^∞ optimal control theory. Simulations of these controllers follow in Section IV. Finally, Section V contains concluding remarks and suggestions for future research.

II. The Epidemic Model

We base our multiple network model on the *classical epidemic model*. This model uses a differential equation to model the spread of worm or virus in a computer network. For a single network, the classical model is

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

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 β is a parameter that captures the rate of spread of the epidemic and is referred to as the *pairwise rate of infection*.

For multiple networks this model can be expanded and generalized. If we consider M networks, let $x_i(t)$ denote the number of infected hosts in network i, where $i = 1, 2, \ldots, M$. Likewise, let $u_i(t)$ denote the malware removal rate for network i. Let α stand for the *cross-network pairwise rate of infection*. Note that the more security measures are used between various networks, the smaller α is relative to β. Also, 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, we arrive at the model

$$\begin{aligned}\dot{x}_i(t) =& \beta\left[N_i - x_i(t)\right]x_i(t) \\ &+ \sum_{j=1, j\neq i}^{M} \alpha\left[N_i - x_i(t)\right]x_j(t) - u_i(t),\end{aligned} \tag{2}$$

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 restricted to one network, this model is

$$\begin{aligned}\dot{x}_1(t) =& \beta\left[N - x_1(t) - x_2(t)\right]x_1(t) - u(t) \\ \dot{x}_2(t) =& u(t).\end{aligned} \tag{3}$$

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 keeps track of the number of hosts that have been patched and are thus not vulnerable to attack.

The epidemic model with removals can also be extended to the case where we have multiple networks. This leads to the set of $2M$ coupled differential equations

$$\begin{aligned}\dot{x}_i(t) =& \beta\left[N_i - x_i(t) - x_{2i}(t)\right]x_i(t) \\ &+ \sum_{j=1, j\neq i}^{M} \alpha\left[N_i - x_i(t) - x_{2i}(t)\right]x_j(t) - u_i(t) \\ \dot{x}_{2i}(t) =& u_i(t),\end{aligned} \tag{4}$$

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.

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 all $i = 1 \ldots M$. The coefficient κ is known as the *removal rate*. Here, we will refer to this scheme 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 makes sense theoretically and is tractable mathematically. Consider the cost

$$J(\mathbf{x}(t), \mathbf{u}(t)) = \int_0^\infty \left[\mathbf{x}^T(t)Q\mathbf{x}(t) + \mathbf{u}^T(t)R\mathbf{u}(t)\right]dt, \tag{5}$$

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.

III. Feedback Malware Removal Response

We now use the multiple dimension versions of the classical epidemic model and the epidemic model with removals to find the feedback responses to malware epidemics in multiple networks. We can choose such responses to be stabilizing or optimal in some sense. Deriving such responses is facilitated by first linearizing the models.

A. Stabilizing Response

One approach to this problem is to simply attempt to stabilize the system, which corresponds to having no infected machines. Finding a feedback controller to stabilize the nonlinear equations in the models (2) and (4) is not simple.

However, by studying these models in light of some particular properties of this context we were able to devise a strategy that results in a stabilizing feedback controller. One crucial observation is that each $x_i(t)$ is nonnegative because of the definition of this variable. This leads to the insight that all of the cross terms and squared terms in the models (2) and (4) decrease the magnitude of the infection rates ($\dot{x}_i(t)$). Therefore, if these helpful squared and cross terms are disregarded, then we will 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 linear models. In fact, both models reduce to the same linear model

$$\dot{\mathbf{x}}(t) = A\mathbf{x}(t) + B\mathbf{u}(t), \tag{6}$$

where

$$A = \begin{bmatrix} \beta N_1 & \alpha N_1 & \alpha N_1 & \cdots & \alpha N_1 \\ \alpha N_2 & \beta N_2 & \alpha N_2 & \cdots & \alpha N_2 \\ \vdots & & \ddots & & \vdots \\ \vdots & & & \ddots & \vdots \\ \alpha N_M & \alpha N_M & \cdots & \alpha N_M & \beta N_M \end{bmatrix} \tag{7}$$

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

Notice that the epidemic models have the inherent physical constraints

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

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 we discuss in detail in the next section. In other words, the trajectory leaves the boundary $[N_1, \dots, N_M]$ immediately. A similar argument can also be made for (4). Therefore, we ignore it for the simplified linear system (6) and focus on the constraint set

$$x_i \geq 0, \quad i = 1 \dots M. \tag{8}$$

Then, under the set of constraints (8), the constrained linear model becomes

$$\dot{x}_i = \begin{cases} [A\mathbf{x} + B\mathbf{u}]_i & \begin{cases} \text{if } x_i > 0 \text{ or} \\ \text{if } [A\mathbf{x} + \mathbf{u}]_i \geq 0 \text{ and } x_i = 0 \end{cases} \\ 0 & \text{else} \end{cases} \tag{9}$$

for all $i = 1 \dots M$.

While it is known that a linear feedback controller can stabilize the linear model (6), whether such a controller also stabilizes the nonlinear models (2), (4), and (9) is a question we investigate in the next section.

1) Stability Analysis: We now show the stability of the system (9) under the set of boundary constraints (8) when controlled by the linear feedback controller

$$\mathbf{u}_s = -K\mathbf{x}, \tag{10}$$

where the matrix K denotes the feedback coefficients. Obviously, the origin constitutes the unique equilibrium for this system. Our analysis focuses on the epidemic model (2) because the extension to the case of (4) is straightforward.

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

Theorem 1. *Given a nonlinear system of the form in (9) 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 (9) we know that components of $\mathbf{x}$ can only be positive or zero. In this case we have assumed that the diagonal entries of the closed loop matrix $(A - BK)$ are 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. We must also know 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 $(A - BK)$. □

We note that the feedback controllers derived from LQR and H^∞ optimal control are not guaranteed to 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 all of the simulations in this paper met this conditions. The closed loop matrices for three dimensional versions of all of the simulations in this paper also met this requirement. Cases where the number of hosts differs between different networks or the cost on control differs 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.

We next show that the stabilizing controller (10) which meets the criteria of Theorem 1 will stabilize the nonlinear systems (2) and (4). 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. The stabilizing condition in Theorem 1 ensures stability over the entire state space, even in the nonlinear case.

B. Linear Quadratic Regulator Optimal Response

Determining the optimal malware removal strategy relative to the cost (5) for the epidemic models (2) and (4) 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, attempting to solve explicitly for the optimal feedback solution using the Hamilton-Jacobi-Bellman equation and the value function leads to an over-determined set of non-linear equations that is not tractable. Therefore a more tractable but sub-optimal approach to this problem was developed based on the linear model (6).

By making use of this linear model (6) and the quadratic cost function (5), we arrive at the well-studied linear quadratic regulator (LQR) optimal control problem. Finding the optimal feedback controller is straightforward once we have established this LQR problem. The optimal feedback controller is

$$\mathbf{u}_o(t) = -R^{-1}B^T P\mathbf{x}(t), \tag{11}$$

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

$$A^T P + PA - PBR^{-1}B^T P + Q = 0. \tag{12}$$

Here we use the fact that (A, B) is controllable because B is the negative identity matrix. Also, note that because $Q > 0$ we know that $P > 0$ exists and is unique.

C. H^∞ Optimal Response

While the model (6) developed in section III-B is useful, it involves some assumptions. First, we ignore the nonlinear terms in the more precise models (2) and (4). Moreover, we assumed that we have a perfect measurement of the number of infected hosts in each network. Finally, the original epidemic models (2) and (4) themselves only approximate malware propagation.

To handle some of these imperfections, we cast this problem in an H^∞ optimal control framework. To capture the assumptions and imperfections in the linear model (9), we alter this model to include a noise term. Let $\delta_i := [A\mathbf{x} + B\mathbf{u} + D\mathbf{w}_a]_i$. Then

$$\dot{x}_i = \begin{cases} \delta_i & \begin{cases} \text{if } x_i > 0 \text{ or} \\ \text{if } \delta_i \geq 0 \text{ and } x_i = 0 \end{cases} \\ 0 & \text{else.} \end{cases} \tag{13}$$

Here $\mathbf{w}_a(t)$ is a noise term that accounts for model assumptions or imperfections. The D matrix shows how this noise term impacts the dynamics of $\mathbf{x}(t)$ and will be set to the identity matrix. Perhaps more importantly, we introduce a measurement error. Let $\mathbf{y}(t)$ be the measurement of the number of infected hosts. Then, we have

$$\mathbf{y}(t) = C\mathbf{x}(t) + E\mathbf{w}_n(t). \tag{14}$$

The matrices C and E will be assumed to be the identity matrix for simplicity here. We define $N := EE^T$ and note that N is also the identity matrix. This implies 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^∞ optimal controller, we also must define the *controlled output*. This quantity is

$$\mathbf{z}(t) = H\mathbf{x}(t) + G\mathbf{u}(t). \tag{15}$$

Here we assume that G^TG is positive definite and that $H^TG = 0$. This implies 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 (5) for the LQR controller we will set $Q = H^TH$ and $R = G^TG$. A few other constraints that must be met for this H^∞ optimal control theory to apply are that $(A,\, B)$ and $(A,\, D)$ be stabilizable, and $(A,\, H)$ and $(A,\, C)$ be detectable, all of which hold under the assumptions made heretofore. We also define $\mathbf{w} := [\mathbf{w}_a\, \mathbf{w}_n]$, the total disturbance to the system. Let the cost ratio used in the H^∞ analysis be

$$L(\mathbf{x}, \mathbf{u}, \mathbf{w}) = \frac{\|\mathbf{z}\|}{\|\mathbf{w}\|}, \tag{16}$$

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^∞ optimal control theory also produces a performance factor (the H^∞ norm) that we can guarantee will be met, as described in Section I. This norm can be thought of as the worst possible value for the cost L. The lowest possible γ is defined by $\gamma^* := \inf_{\mathbf{u}} \sup_{\mathbf{w}} L(\mathbf{u}, \mathbf{w})$ and could also be viewed as the optimal performance level in this H^∞ context.

In order to actually solve for the optimal controller $\mu(\mathbf{y})$, a corresponding differential game is defined, which is parameterized by γ. The optimal worst case controller $\mathbf{u}_w = \mu_\gamma(\mathbf{y})$ can be determined from this differential game formulation for any $\gamma > \gamma^*$. It is given by [14] as

$$\mu_\gamma(\mathbf{y}) = -(G^TG)^{-1}B^T\bar{Z}_\gamma\hat{\mathbf{x}}, \tag{17}$$

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

$$A^TZ + ZA - Z(B(G^TG)^{-1}B^T - \gamma^{-2}DD^T)Z + Q = 0, \tag{18}$$

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

$$\begin{aligned} \dot{\hat{\mathbf{x}}} = &\left[A - (B(G^TG)^{-1}B^T - \gamma^{-2}DD^T)\bar{Z}_\gamma\right]\hat{\mathbf{x}} \\ &+ \left[I - \gamma^{-2}\bar{\Sigma}_\gamma\bar{Z}_\gamma\right]^{-1}\bar{\Sigma}_\gamma C^TN^{-1}(\mathbf{y} - C\hat{\mathbf{x}}), \end{aligned} \tag{19}$$

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

$$A\Sigma + \Sigma A^T - \Sigma(C^TN^{-1}C - \gamma^{-2}H^TH)\Sigma + DD^T = 0. \tag{20}$$

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

The linear H^∞-optimal feedback controller (17) provides a robust malware response or epidemic removal strategy based on an estimate of the number of infected hosts. It can be calculated offline using only the linear quadratic model.

IV. Simulations and Results

Simulations of the models (2) and (4) with the various controllers were performed in Matlab. We simulate two networks, each containing 500 hosts. Initially network 1 has 250 infected hosts while network 2 only has 100 infected hosts. The β parameter is set to 5.6×10^{-5}, the estimated value of β for the SQL Slammer worm [15]. The α parameter is set to $\beta/4$, reflecting an assumption that worms will spread more slowly between networks if proper security measures are taken for each network. We set the Q and H^TH matrices at 0.01 times the identity matrix and set R and G^TG equal to the identity matrix. This reflects a situation where patching is more expensive than an infected host.

Moreover, we added a noise term to the system dynamics when simulating each of the models, identical to the noise term in the H^∞ model (13). This normally distributed noise term captures some of the imperfections in the models.

A. Stabilizing Response

Recall that 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 (6). Therefore we can choose exactly where to place the poles and achieve varying degrees of stability. For these simulations we will investigate placing the poles in a few locations: -1, -0.5, and

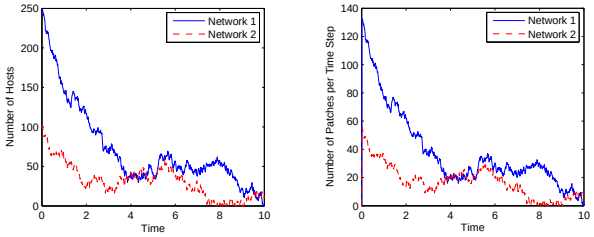

(a) (b)

Fig. 1. (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 (4).

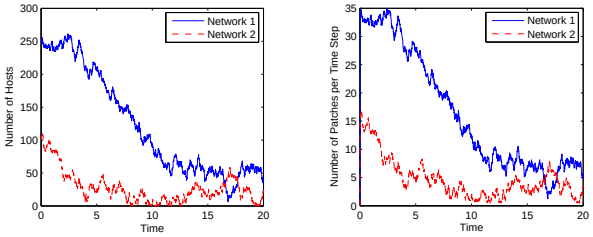

(a) (b)

Fig. 2. (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 (4).

TABLE I

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

TABLE II

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

-0.1. Fig. 1 shows the results of a simulation of the epidemic model with removals when the poles are placed both at -0.5.

When this controller is applied to the various models under consideration, we find that the resulting cost values are very similar. The cost results of these simulations are expressed in Table I. The graphs of the simulations of each of the systems are also nearly identical, implying that disregarding the nonlinear terms in order to derive our controller was reasonable. The linear model has a higher cost than the other two models because it does not consider the nonlinear terms, which contribute to stabilizing the system. The epidemic model with removals contains the most helpful nonlinear terms, so it achieves the lowest cost with this controller.

B. LQR Optimal Response

The results of a simulation of this scenario for the epidemic model with removal are shown in Fig. 2. The number of infected hosts in each network are shown in Fig. 2 (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. 1. The costs incurred when this controller is applied to the various models are shown in Table II. Again the similarity of costs across different models implies that the linearization was reasonable.

When this controller is applied to the various models under consideration we find that the resulting cost values are very similar. These costs are shown in Table II. The graphs of the simulations of each of the systems are also nearly identical, implying that disregarding the nonlinear terms in order to derive our controller was not unreasonable. The linear model has a higher cost than the 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.

Interestingly, these costs are slightly higher than the costs of the stabilizing controller 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 for the nonlinear system (poles are -0.1059 and -0.1022). 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% provides one instance of the magnitude of the cost increase resulting from deriving an optimal controller with a linearized version of the model. 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.

We also simulate this system under measurement noise. The system cost more than doubles for each model (see Table II). Moreover, the feedback controller becomes highly oscillatory, as seen in Fig. 3.

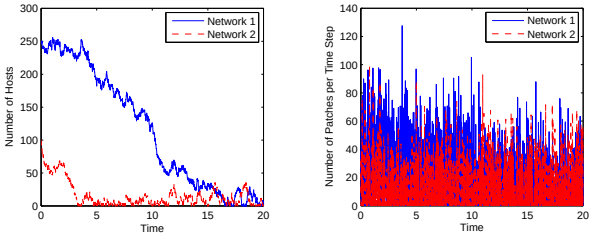

(a) (b)

Fig. 3. (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 (4) and when there are noisy measurements.

TABLE III
COST OF MODELS UNDER H^∞ OPTIMAL CONTROLLER

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

C. H^∞ Optimal Response

The simulations of the H^∞ optimal controller are unique because we design the controller to operate when there is the worst-case possible noise on the system measurements and dynamics. Thus in these simulations we incorporate noisy state measurements and also the H^∞ state estimate (19).

This controller will lead to unnecessarily high costs due to over-aggressive responses. This aggressive response can be seen in Fig. 4, where we simulate the application of the H^∞ optimal controller to the epidemic model with removals. 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. 3.

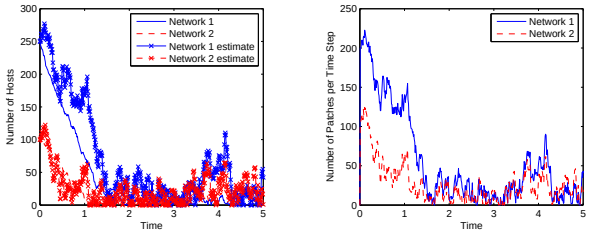

(a) (b)

Fig. 4. (a) The number of infected hosts and (b) the patching rate over time when the H^∞ optimal controller is applied to the epidemic model with removals (4).

The costs under this controller for each of the three models are shown in Table III. The aggression of the H^∞ 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^∞ optimal controller offers a cost ratio guarantee that the other controllers cannot. Again the nonlinear terms in the system do not significantly change the system dynamics.

V. CONCLUSIONS AND FUTURE WORK

We have presented some solutions to the malware removal problem for multiple connected networks. We first extended the classical epidemic response model and the epidemic response model with removals to multiple dimensions. Next we constructed a cost function and linearized the coupled differential equations that make up the model. Then we utilized basic optimal control methods to derive and numerically evaluate several stabilizing and optimal (for the linearized model) feedback malware removal controllers.

We found that while fine-tuned stabilizing controllers can out-perform the optimal controllers in some cases, optimal controllers are more flexible to changes in the cost structure. We also showed that classical epidemic models assume a proportional response strategy that is not necessarily optimal.

Several extensions to this theory exist. In some situations parameters like α, 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 describe the spread of infectious diseases [3]. Optimal control theory could be applied to these models to determine the optimal response to the spread of a disease.

REFERENCES

[1] D. Moore, C. Shannon, and K. Claffy, "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, 2002, pp. 273–284.

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

[3] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM Review*, vol. 42, no. 4, pp. 599–653, 2000.

[4] D. Moore, C. Shannon, G. Voelker, and S. Savage, "Internet quarantine: Requirements for containing self-propagating code," in *Proc. 22nd IEEE Infocom*, vol. 3, April 2003, pp. 1901 – 1910.

[5] T. M. Chen and N. Jamil, "Effectiveness of quarantine in worm epidemics," in *Proc. of IEEE ICC 2006*, Istanbul, Turkey, June 2006, pp. 2142–2147.

[6] C. Zou, W. Gong, and D. Towsley, "Worm propogation modeling and analysis under dynamic quarantine defense," in *Proc. of ACM Workshop on Rapid Malcode*, Washington, DC, 2003, pp. 51–60.

[7] K. Rohloff and T. Başar, "The detection of RCS worm epidemics," in *Proc. of ACM Workshop on Rapid Malcode*, Fairfax, VA, 2005, pp. 81–86.

[8] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver, "Inside the slammer worm," *IEEE Security & Privacy Magazine*, vol. 1, pp. 33–39, July-Aug. 2003.

[9] Cisco, "NAT and stateful inspection in Cisco IOS firewall," white paper, 2006. [Online]. Available: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_white_paper09186a0080194af8.shtml

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

[11] ——, "An intrusion detection game with limited observations," in *Proc. of 12th International Symposium on Dynamic Games and Applications*, Sophia Antipolis, France, July 2006.

[12] R. Dantu, J. Cangussu, and A. Yelimeli, "Dynamic control of worm propagation," in *Proc. of IEEE Conference on Information Technology: Coding and Computing*, vol. 1, April 2004, pp. 419–423.

[13] M. Bloem, T. Alpcan, and T. Başar, "Intrusion response as a resource allocation problem," in *Proc. 45th IEEE Conf. Decision and Control*, San Diego, CA, USA, December 2006.

[14] T. Başar and P. Bernhard, *H^∞-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, 2nd ed. Boston, MA: Birkhäuser, 1995.

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