Optimal Patching in Clustered Malware Epidemics

MHR. Khouzani, S. Eshghi, S. Sarkar, S. S. Venkatesh

***Abstract*—Studies of propagation of malware in mobile network have revealed that the spread of malware can be highly inhomogeneous across different regions. Heterogeneous rates of contact can also be due to diverse platforms, utilization of contact lists by the malware, clustered nature of the network, etc. In this paper, we propose a general formal framework to leverage such heterogeneity information into devising optimal patching policies that attain the minimum aggregate cost due to the spread of malware and the surcharge of patching. Using Pontryagin's Maximum Principle for a stratified epidemic model, we analytically prove that in the mean-field deterministic regime optimal patch disseminations are simple single-threshold policies which are amenable to implementation in a distributed manner. Through numerical calculations, we investigate the behavior of optimal patching policies in sample topologies and demonstrate their advantages.**

I. Introduction

Worms, i.e. self-propagating malicious codes, are a decades-old threat in the realm of Internet. Worms undermine the network by performing various malicious activities: they can eavesdrop and analyze the traversing data, access privileged information, hijack sessions, disrupt network functionalities such as routing, etc. Although Internet is the traditional arena for malicious codes such as trojans, spyware and viruses, the battle is expanding to new territories: the current boom in mobile devices combined with their spectacular software and hardware capabilities has created a tremendous opportunity for future malware. Mobile devices communicate with each other and with computers through a myriad of means. Not only can they interact using Bluetooth or infrared when they are in each other's proximity or through an exchange of multimedia content messages (MMS), they can have ubiquitous access to mobile Internet and peer to peer networks via a telecom provider. Current smartphones are equipped with operating systems, CPUs and memory powerful enough to execute increasingly more complex codes. Incidents of spread of wireless malware such as *cabir*, *skulls*, *mosquito*, *commwarrior*, *etc.* have already sounded the alarm [1]. It is in fact theoretically predicted [2] that it is only a matter of time before major malware outbreaks are witnessed in the wireless domain.

The spread of malware can be countered through patching [3]: the underlying vulnerability utilized by the worm can be amended by installing security patches that immunize the susceptible, and potentially remove the malware and thus heal and immunize the infective nodes. However, the distribution of these patches burdens the limited resources of the network, and hence, if not carefully controlled, can lead to major havoc. In wired networks, the spread of *Welchia*, a counter-worm to thwart *Blaster*, created substantial traffic which in turn rapidly destabilized important sections of the Internet. Resource constraints are even more pronounced in wireless networks in which bandwidth is constrained and is more sensitive to overload, and nodes are limited in their energy reserves. Recognizing the above, works such as [4], [5] have included the cost of patching in the aggregate damage of the malware and have characterized the optimal dynamic patching policies which attain desired trade-offs between the efficacy of patching and the extra taxation of the network resources. However, as we will explain next, these studies suffered from a drawback: a strong simplifying assumption.

Malware spreads when an infective node *contacts*, i.e. communicates with, a susceptible node, which is a node without a copy of the malware and vulnerable to it. The results in [4], [5] critically rely on the *homogeneous mixing* assumption that all pairs of nodes have identical expected inter-contact times. While this assumption may serve as an approximation in cases where detailed information about the network is not available, a series of studies demonstrate that the spread of malware in mobile networks can be considerably inhomogeneous [2], [6]–[9], owing primarily to the non-uniform distribution of nodes. Wireless nodes in high density areas, sometimes referred to as "popular content" regions or "hot-spots", have more frequent opportunities to contact each other than to contact nodes in distant and less dense areas. Heterogeneity in the contact process can arise for other reasons too. Malware may have a lower rate of contact between devices with differing operating systems or communication protocols [10]–[12]. Mobile malware may also select targets from the address books of the infective hosts [3]: the contact rate is thus higher amongst *friendship cliques* in the social network of users. Malware which spreads using (mobile or wired) Internet can have easier access to the IP-addresses of the subnet to which the infective host belongs compared to the rest of the masked IP addresses [13]. The behavioral pattern of the users can also cause heterogeneous contact rates, e.g. a safe user may avoid unsolicited mass-messages or may install firewalls, hence hindering the spread of malware as compared to one with risky behavior. Moreover, cloud-computing seems to be a natural context for heterogeneous mixing: computers inside the same cluster have a much higher speed of communication amongst themselves than with computers of distant clusters.

Indeed many works have proposed practical methods to identify and characterize and incorporate such inhomogeneities to more accurately predict the spread of infection [2], [7], [8], [13]–[15], etc. Relatively few, e.g., [3], [9], [11], consider the cost of patching and seek to minimize it in the presence of heterogeneous contact processes. The proposed

The authors are with the Department of Electrical and Systems Engineering at University of Pennsylvania, Philadelphia, PA, U.S.A. Their email addresses are *khouzani,eshghi,swati,venkates@seas.upenn.edu*.

policies in [3], [9] are heuristic and apply to specific settings. The only paper we could find that provides *provably optimal* patching policies for heterogeneous networks is [11]. They however focus on SIS models and optimize only in the space static policies (i.e., those that do not vary patching rates over time) therein. Patching performance can significantly improve if we allow the patching rates to vary dynamically in accordance with the evolution of the contagion.

We propose a formal framework for devising dynamic optimal patching policies which leverage heterogeneity in the network structure to attain the minimum aggregate cost due to the spread of malware combined with the overhead of patching. We consider both *non-replicative* and *replicative* patching: in the former, some of the hosts are pre-loaded with the patch which they transmit to the rest. In the latter, each recipient of the patch can also forward the patch to nodes which it contacts, by a mechanism similar to the spread of the malware itself. The framework in each case relies on optimal control formulations which cogently capture the effect of the patching rate controls on the state dynamics and their resulting trade-offs. We accomplish this by using a combination of damage functions and a *stratified*[1] epidemic model in which nodes are divided into different types. Nodes of the same type homogeneously mix with a rate specific to that type, and nodes across different types contact each other at rates particular to each pair of types. If two types do not interact, the corresponding inter-contact rates are set to zero. The model can therefore capture any communication topology between different groups of nodes. Above and beyond, it can exploit the inhomogeneity in the network to enable a better utilization of the resources. Such higher patching efficacy is achieved by allowing the patching controls to depend on the node types. This leads to multidimensional (dynamic) optimal control formulations in the solution space of functions rather than variables, which may significantly add to the complexity of the optimization. However, using Pontryagin's Maximum Principle, we analytically prove that in the mean-field deterministic regime, for both non-replicative and replicative settings, the optimal patching controls are simple single-threshold policies which are amenable to implementation in a distributed manner. Furthermore, the thresholds may now be computed through off-the-shelf efficient numerical techniques. To the best of our knowledge, such structure results have not been established in the context of (either static or dynamic) control of heterogeneous epidemics. Through numerical calculations, we investigate a series of interesting behaviors of optimal patching policies for different sample topologies.

II. System Model

We first describe and develop the model of the state dynamics of the system as a general *stratified* epidemic for both non-replicative (§II-A) and replicative (§II-B) patching. Next in §II-C, we motivate our model in each of the different contexts which we discussed in section (§I). Subsequently in §II-D, we characterize the aggregate cost of patching and cast the multi-type resource-aware patching as a multi-dimensional optimal control problem, whose solutions we develop later in §III and §IV.

A. Dynamics of non-replicative patching

A node is **infective** if it has been contaminated by the malware, **susceptible** if it is vulnerable to the infection but has not been infected yet, and **recovered** if it is immune to the worm. An infective spreads the malware to a susceptible while transmitting data or control messages. The network consists of a total of N nodes which can be stratified into M different *types*.[2] A node of type i contacts another of type j at rate $\hat{\beta}_{ij}$.

There are $N_i = \alpha_i N$ ($\alpha_i > 0$) nodes of type i in the network, among which $n^i_S(t)$, $n^i_I(t)$ and $n^i_R(t)$ are respectively in the susceptible, infective and recovered states at time t. Let the corresponding fractions be $S_i(t) = n^i_S(t)/N_i$, $I_i(t) = n^i_I(t)/N_i$, and $R_i(t) = n^i_R(t)/N_i$. We assume that during the course of the epidemic, the populations of each type, N_i, are stable and do not change with time. Thus, for all t and all i, we have $S_i(t) + I_i(t) + R_i(t) = 1$.

Amongst each type, a pre-determined set of nodes, referred to as *dispatchers*, are loaded with the appropriate patch. The dispatchers can transmit the patches to the susceptible and infective nodes and *immunize* the susceptibles and possibly *heal* the infectives to the recovered state. In *non-replicative* patching, as opposed to *replicative* patching in §II-B, the recipient nodes of the patch do not propagate it further.[3] Dispatchers of type i can contact nodes of type j at the rate of $\tilde{\beta}_{ij}$ which may be different from the contact rates $\hat{\beta}_{ij}$ of the malware (fig. 1). Indeed the network manager may utilize a higher priority option for distribution of the patches, or the malware may utilize legally restricted means of propagation not practicable for dispatchers. Moreover, the patch might not be applicable to all types, forcing relevant cross $\tilde{\beta}_{ij}$ to be zero. The number of dispatchers of type i, which is fixed over time in the non-replicative setting, is $N_i R^i_0$ where $0 < R^i_0 < 1$. We assume that the dispatchers are themselves immune to the malware and hence they are recovered from the beginning. If the network manager can patch only parts of the network, then R^i_0 for other parts can be assumed zero. A node does not know which other nodes are infected; it can however identify the dispatchers.

Let $t = 0$ designate the earliest moment that the infection is detected by the network and the appropriate patches are generated. At $t = 0$, a fraction of $0 \leq I_i(0) = I^i_0 \leq 1$ of nodes of type i are infected, and $S_i(0) = 1 - I^i_0 - R^i_0$. If the infection does not initially exist amongst a type i, then $I^i_0 = 0$.

At any given t, any one of the $n^i_S(t)$ susceptibles of type i may be contacted by any of the $n^j_I(t)$ infectives of the type j

[1] known by other terms such as structured, clustered, multi-class, multi-type, multi-population, compartmental epidemic models, and sometimes loosely as heterogeneous, inhomogeneous or spatial epidemic models.

[2] equivalently, *clusters, segments, populations, categories, classes, compartments, strata*, etc.

[3] This may be preferred when the patches may themselves be contaminated and can not be reliably authenticated.

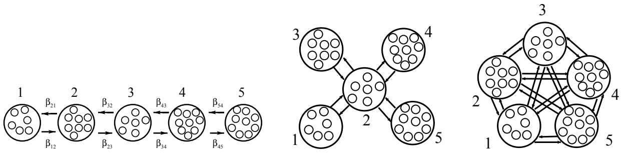

Fig. 1. Three sample topologies of 5 hotspot regions: linear, star and complete. For instance, nodes of hotspot 1 in the linear topology can only communicate with nodes of hotspots 1 and 2: they contact nodes of hotspot 1 at rate $\hat{\beta}_{11}$ and nodes of hotspot 2 at rate $\hat{\beta}_{12}$.

at rate $\hat{\beta}_{ji}$. Susceptibles of type i are therefore transformed to infectives (of the same type) at rate $n_S^i(t) \sum_j \hat{\beta}_{ji} n_I^j(t)$. The system manager controls the resources consumed in distribution of the *patches* by dynamically determining the distribution rates of the dispatchers. Let the rate of transmission attempts of dispatchers of type i at time t be $u_i(t)$. Hence, the patches immunize the susceptibles of type i, transforming them to recovered nodes of type i, at rate $n_S^i(t) \sum_j \tilde{\beta}_{ji} u_j(t) N_j R_0^j$ at each t. The attainable transmission rates of dispatchers of each type must be non-negative and bounded: [4]

$$0 \leq u_i(t) \leq u_{i,\max} \quad \text{at each } t. \tag{1}$$

The efficacy of the patch may be lower for treating infective nodes than susceptible nodes. For instance, the malware may try to prevent the reception or installation of the patch in an infective host, or the patch may only be able to remove the vulnerability that leaves the nodes exposed to the malware, but fail to remove the malware itself. We capture the above possibility by introducing a (type-dependent) coefficient $0 \leq \pi_i \leq 1$ as follows: $\pi_i = 0$ occurs when the patch is completely unable to heal the infectives of type i and only immunizes the susceptibles of type i, whereas $\pi_i = 1$ represents the other extreme scenario where a patch can equally well immunize and heal susceptibles and infectives of type i.[5] If the patch heals an infective node, its state changes to recovered, otherwise it remains an infective. Thus, the infectives of type i recover at rate $\pi_i n_I^i(t) \sum_{j=1}^M \tilde{\beta}_{ji} u_j(t) N_j R_0^j$ at each t. Define $\beta_{ij} := \lim_{N\to\infty} N\hat{\beta}_{ij}$ and $\bar{\beta}_{ij} := \lim_{N\to\infty} N\tilde{\beta}_{ij}$. If the total number of nodes, N, is large and $\alpha_i > 0$ for all i[6], then according to the mean field convergence results (e.g. in [16, p.1] or recently in [17]), $S_i(t)$, $I_i(t)$ converge pathwise to the solution of the following system of differential equations:[7]

$$\dot{S}_i = -\sum_{j=1}^{M} \beta_{ji} I_j S_i - S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_0^j u_j \tag{2a}$$

$$\dot{I}_i = \sum_{j=1}^{M} \beta_{ji} I_j S_i - I_i \sum_{j=1}^{M} \pi_i \bar{\beta}_{ji} R_0^j u_j \tag{2b}$$

where the initial conditions and the the state constraints are:

$$\mathbf{S}(0) = \mathbf{S}_0 := (S_0^1, \ldots, S_0^M), \mathbf{I}(0) = \mathbf{I}_0 := (I_0^1, \ldots, I_0^M) \tag{3}$$

$$\mathbf{S} \succeq 0, \quad \mathbf{I} \succeq 0, \quad \mathbf{S} + \mathbf{I} \preceq \mathbf{1} - \mathbf{R}_0 \tag{4}$$

in which, $\mathbf{1}_{M\times 1} := (1, \ldots, 1)^T$, and the inequality signs represent element-wise constraints. Note that the evolution of $\mathbf{R}(t)$ need not be explicitly considered since at any given time, $R_i(t) = 1 - S_i(t) - I_i(t)$. Henceforth wherever not ambiguous, we drop the dependence on t and make it implicit.

B. Dynamics of replicative patching

In the replicative setting, a recipient of the patch can forward it to other nodes upon subsequent contact. Thus, recovered nodes of type i add to the pool of dispatchers of type i, and hence the fraction of dispatchers of type i grows to $R_i(t)$ at time t, whereas in the non-replicative model, the fraction of dispatchers of type i continues to be R_0^i at all times. Thus the system dynamics equations in (2) need to be modified. In the non-replicative case we chose $(\mathbf{S}(t), \mathbf{I}(t))$ to represent the system, whereas In the replicative case we retain $(\mathbf{S}(t), \mathbf{I}(t), \mathbf{R}(t))$ instead. As it turns out, the specific choices make the analyses more convenient in each case. The new system dynamics equations are hence as follows:

$$\dot{S}_i = -\sum_{j=1}^{M} \beta_{ji} I_j S_i - S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_j u_j \tag{5a}$$

$$\dot{I}_i = \sum_{j=1}^{M} \beta_{ji} I_j S_i - I_i \sum_{j=1}^{M} \pi_i \bar{\beta}_{ji} R_j u_j \tag{5b}$$

$$\dot{R}_i = S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_j u_j + I_i \sum_{j=1}^{M} \pi_i \bar{\beta}_{ji} R_j u_j \tag{5c}$$

with the initial conditions and the state constraints as:

$$\begin{aligned} &\mathbf{S}(0) = \mathbf{S}_0 := (S_0^1, \ldots, S_0^M), \\ &\mathbf{I}(0) = \mathbf{I}_0 := (I_0^1, \ldots, I_0^M) \\ &\mathbf{R}(0) = \mathbf{R}_0 := (R_0^1, \ldots, R_0^M) \end{aligned} \tag{6}$$

$$\mathbf{S} \succeq 0, \quad \mathbf{I} \succeq 0, \quad \mathbf{R} \succeq 0, \quad \mathbf{S} + \mathbf{I} + \mathbf{R} = \mathbf{1}. \tag{7}$$

The following lemma shows that the state constraints in both non-replicative and replicative models hold as long as the control functions satisfy (1) - thus, these constraints can be ignored henceforth.

Lem. 1. *For any* $\mathbf{u}(.)$ *that satisfies* (1)*, the state constraints* (4)*,* (7) *hold in the non-replicative and replicative cases respectively. In both cases, for all* i, t, $S_i(t) > 0$ *and* $I_i(t) > 0$ *if* $I_0^i > 0$.

C. Motivation of the models and instantiations

In the introduction section (§I), we described the motivations for the our stratified epidemic model ((2) and (5)) through different examples. Here, we provide more detail and specify the different types in each context. Note that in general, different types can refer to different combination of the following cases as well.

[4] In general, if the network manager cannot control a certain type, the corresponding $u_{i,\max}$ will be zero.

[5] Also, irregular success chances of healing are representable by intermediate values of π.

[6] so that each type comprises a non-vanishing fraction of the total population of the nodes

[7] Throughout the paper, variables with dot marks (e.g., $\dot{S}_i(t)$) will represent their <u>time</u> derivatives (e.g., time derivative of $S_i(t)$ at t).

1) Proximity-based spread – heterogeneity through locality: The overall roaming area of the nodes can be divided into regions (e.g., hotspots, office/residential areas, central/peripheral areas, etc.) of different densities (fig. 1). One can therefore stratify the nodes based on their locality, that is, each type corresponds to a region. IP eavesdropping techniques (using softwares such as `AirJack`, `Ethereal`, `FakeAP`, `Kismet`, etc.) allows a malware to detect new victims in the vicinity of the host. Distant nodes have more attenuated signal strength (i.e., lower SINR) and are therefore less likely to be detected. Thus, malware (and also patch) propagation rates $\beta_{ij}, \bar{\beta}_{ij}$ are related to the local densities of the nodes in each region and decay with an increase in the distance between regions i and j: typically β_{ii} exceeds β_{ij} for $i \neq j$, likewise for $\bar{\beta}_{ij}$. The same phenomenon was observed for malware such as `cabir` and `lasco` that use Bluetooth and Infrared to propagate.

2) heterogeneity through software/protocol diversity: Mobile nodes use different operating systems and communication protocols, e.g., `Symbian`, `Android`, `IOS`, `RIM`, `webOS`, etc. In fact, a network which relies on a homogeneous software/protocol is vulnerable to an attack that exploits a common weakness (e.g. a buffer flow vulnerability). Thus, inspired by the natural observation of *survivability through heterogeneity*, increasing the network's heterogeneity is proposed as a defense mechanism without sacrificing interoperability [10]. Such heterogeneities lead to dissimilar rates of propagation of malware amongst different types, where each type represents a specific OS, platform, software, protocol, etc. In the extreme case, the malware may not be able to contaminate nodes of certain types. The patching should also be wary of such inhomogeneities in order to optimally utilize the network resources since the rate of patching might also be dissimilar amongst different types.

3) heterogeneity through available IP space: Smartphone trojans like `skulls` and `mosquito` spread using Internet or P2P networks. In such cases the network can be decomposed into *autonomous systems* (ASes) that each represent an AS [13]. A worm either uniformly randomly scans the IP addresses or uses the IP masks of ASs to restrict its search domain and increase its rate of finding new susceptible nodes. In each of these cases the contact rates differ between different AS's depending on the actual number of assigned IPs in each IP sub-domain and the maximum size of that IP sub-domain.

4) heterogeneity through size of the cliques: Malware which specifically spread in social networks has been recorded in the past few years [18]. Examples include `Samy` in MySpace in 2005 and `Koobface` in MySpace and Facebook in 2008. Specifically, `Koobface` spread by delivering (contaminated) messages to the "friends" of an infective user. Moreover, MMS based malware such as `commwarrior` may also utilize the contact list of an infective host to access new handsets. The social network graph can be approximated by a collection of friendship *cliques*. A clique is a (close to) complete subgraph (a collection of nodes and all (most of) the possible links between them) within a larger graph. Users of the same clique can be regarded as the same type. Indeed, the rate of contact within cliques and across cliques differ depending on the relative sizes of the cliques

5) Cloud-computing: heterogeneity through cluster sizes: In cluster (or grid, or volunteer) computing [19], the types are simply each cluster of CPUs in the cloud. Any two computers in the same cluster can communicate at faster rates than those in different clusters. These contact rates depend on the communication capacity of connecting lines as well as the relative number of computers in each cluster.

D. The Objective Function

The network seeks to minimize the overall combined cost of infection and the resource overhead of patching in a operation time window of $[0, T]$, where T is a parameter of choice. At any given time t, the system incurs costs at rates of $f(\mathbf{I}(t))$ due to the malicious activities of the malware. For instance, the malware may use the infected hosts to (i) eavesdrop and analyze and/or (ii) misroute, alter or destroy the traffic that the hosts generate or relay. The scalar function $f(.)$ is non-decreasing and differentiable with respect to each I_i. Also, WLoG $f(\mathbf{0}) = 0$. The most natural candidate for $f(\mathbf{I})$ is $\sum_{i=1}^{M} f_i(I_i)$, that is, a weighted summation of the costs of infection of each type. The weights reflect the number of nodes as well as the relative importance of the nodes in each type, which in turn depend on the criticality of their contained data and/or the nodes' functionality.[8] The network also benefits at the rate of $L(\mathbf{R}(t))$, i.e. incurs a cost at the rate of $-L(\mathbf{R}(t))$ owing to the removal of incertitude about the state of the nodes brought about by patching. Moreover, this addition allows us to recover the results of optimal epidemic forwarding policies in multi-type DTN networks [20], [21] as a special case of our model.[9]

In addition to the cost of infection, each dispatcher incurs a cost on the network by consuming the available bandwidth and the energy reserves of the nodes to disseminate the patches. We will capture this portion of the cost with a fairly general combination of cost terms for two patching scenarios. In the first simple scenario, the dispatchers *broadcast* the patch. The overhead bandwidth/energy at time t is thus proportional to

[8] Such differences themselves may be a source of stratification. In general, different types need not exclusively reflect disparate mixing rates.

[9] In a Delay Tolerant Network, a server may seek to broadcast a message to as many nodes as possible before a deadline, by employing minimal resources such as energy and bandwidth. In this case, susceptibles are the nodes that are yet to receive the message, and the recovered are those that have received it. Dissemination of the message may either be performed in non-replicative or replicative manner. Infectives are absent in this problem. The system reward increases with an increase in the number of nodes which have received the disseminated message. Also, the sooner the message is disseminated, the better, hence the integration of $-L(\mathbf{R}(t))$ over time (the negative sign in there because the optimization is cast as a minimization problem). [22, appendix-A] directly relates the integral over time of the fraction of the recipient nodes to the probability that a message is delivered to sink nodes before deadline T. Hence the minimum delay problem is transferred to the maximization of $\int_0^T \sum_{i=1}^{M} N_i R_i(t)\, dt$, which corresponds to the special case of linear $L(\mathbf{x}) = \sum_{i=1}^{M} c_i x_i$ in our setting for appropriate scaling. This observation is used in papers such as [20], [21] to propose resource-optimal epidemic forwarding policies in multi-type DTNs. Our model therefore captures the systems discussed in [20], [21] as special cases by assuming $I_0 = 0$ and $f(\mathbf{I}) \equiv \mathbf{0}$. This connection is also recognized in [5].

the weighted summation $\sum_{i=1}^{M} R_0^i h_1^i(u_i)$, as each dispatcher of type i incurs a cost at the rate of $h_1^i(u_i(t))$. In the alternative scenario, dispatchers may transmit only to the nodes that have not yet received the patch.[10] Hence, the cost of patching in this case can in general be represented by a combination of terms of the following structure: $\sum_{i=1}^{M}\sum_{j=1}^{M} R_0^i \bar{\beta}_{ij}(S_j + I_j) h_2^i(u_i)$. The scalar functions $h_1^i(.)$ and $h_2^i(.)$ for each i are selected to represent how much resource is consumed for transmission of the patch to nodes of each type and how significant such extra taxation of the resources is for each type. These functions are naturally assumed to be non-decreasing and differentiable with respect to their argument. Along the same lines of reasoning, the corresponding terms for the cost of replicative patching can be expressed as $\sum_{i=1}^{M} R_i h_1^i(u_i)$ and $\sum_{i=1}^{M}\sum_{j=1}^{M} \bar{\beta}_{ij} R_i (S_j + I_j) h_2^i(u_i)$ respectively.

Thus, the aggregate cost for the non-replicative patching is:

$$J_N = \int_0^T \Bigg(f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i=1}^{M} R_0^i h_1^i(u_i) + \sum_{i=1}^{M}\sum_{j=1}^{M} \bar{\beta}_{ij} R_0^i (S_j + I_j) h_2^i(u_i) \Bigg)\, dt \tag{8}$$

and for the replicative patching is:

$$J_R = \int_0^T \Bigg(f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i=1}^{M} R_i h_1^i(u_i) + \sum_{i=1}^{M}\sum_{j=1}^{M} \bar{\beta}_{ij} R_i (S_j + I_j) h_2^i(u_i) \Bigg)\, dt \tag{9}$$

Problem Statement: The system seeks to minimize the aggregate cost (**A**) in (8) for non-replicative patching and (**B**) in (9) for replicative patching by appropriately regulating $\mathbf{u}(t)$ at all t subject to (1) and $\mathbf{u}(\cdot)$ having only a finite number of points of discontinuity, allowing the states to evolve (**A**) as per (2) for non-replicative, and (**B**) as per (5) for replicative patching, and satisfy the respective initial state conditions in (3) and (6).

III. Optimal Non-Replicative Patching

A. Numerical framework for computing the optimal controls

The main challenge in computing the optimal state and control functions $((\mathbf{S},\mathbf{I}),\mathbf{u})$ is that the differential equations (2) can be solved once the optimal controls $\mathbf{u}(\cdot)$ are known. Thus, the only approach seems to be that of an exhaustive search, which is ruled out since there are an uncountably infinite number of control functions. *Pontryagins Maximum Principle (PMP)*, however, provides an elegant technique for solving this seemingly intractable problem. PMP bears a close analogy to the primal-dual non-linear optimization framework, except that it allows optimization in the function (as opposed to the variable) space. (8) and (2) are analogous to the objective function and constraints of a primal optimization formulation; *adjoint* functions, to be defined shortly, will play the role of Lagrange multipliers and the *Hamiltonian* $\mathcal{H}$ is similar to the objective function of the relaxed optimization in the primal-dual framework:

$$\begin{aligned}\mathcal{H} := {} & f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i=1}^{M} R_0^i h_1^i(u_i) \\ & + \sum_{i=1}^{M}\sum_{j=1}^{M} \bar{\beta}_{ij} R_0^i (S_j + I_j) h_2^i(u_i) \\ & + \sum_{i=1}^{M} \Bigg((\lambda_i^I - \lambda_i^S) S_i \sum_{j=1}^{M} \beta_{ji} I_j \Bigg) \\ - \sum_{i=1}^{M} \lambda_i^S S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_0^j u_j & - \sum_{i=1}^{M} \lambda_i^I I_i \sum_{j=1}^{M} \pi_i \bar{\beta}_{ji} R_0^j u_j \end{aligned} \tag{10}$$

where at each point of continuity of $\mathbf{u}(\cdot)$ for all $i = 1 \ldots M$ the *adjoint* functions λ_i^S and λ_i^I satisfy

$$\begin{aligned}\dot{\lambda}_i^S = -\frac{\partial \mathcal{H}}{\partial S_i} = & - L_i(\mathbf{R}) - \sum_{j=1}^{M} \bar{\beta}_{ji} R_0^j h_2^j(u_j) \\ & - (\lambda_i^I - \lambda_i^S) \sum_{j=1}^{M} \beta_{ji} I_j + \lambda_i^S \sum_{j=1}^{M} \bar{\beta}_{ji} R_0^j u_j \\ \dot{\lambda}_i^I = -\frac{\partial \mathcal{H}}{\partial I_i} = & - L_i(\mathbf{R}) - f_i(\mathbf{I}) - \sum_{j=1}^{M} \bar{\beta}_{ji} R_0^j h_2^j(u_j) \\ & - \sum_{j=1}^{M} \big((\lambda_j^I - \lambda_j^S) \beta_{ij} S_j \big) + \lambda_i^I \sum_{j=1}^{M} \pi_i \bar{\beta}_{ji} R_0^j u_j \end{aligned} \tag{11}$$

along with the final conditions:

$$\lambda_i^S(T) = 0, \quad \lambda_i^I(T) = 0. \tag{12}$$

Then according to PMP, the optimal control at any time t is derived as:

$$u_i \in \arg \min_{u_{i,\min} \le \underline{u}_i \le u_{i,\max}} \mathcal{H}.$$

This yields:

$$u_i \in \arg\min \Bigg(h_1^i(u_i) + R_i^0 h_2^i(u_i) \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j) - \phi_i u_i \Bigg) \tag{13}$$

$$\text{where} \quad \phi_i := R_0^i \sum_{j=1}^{M} \bar{\beta}_{ij} \lambda_j^S S_j + R_0^i \pi_j \sum_{j=1}^{M} \bar{\beta}_{ij} \lambda_j^I I_j \tag{14}$$

[10] This can be achieved by keeping a common database of nodes that have successfully received the patch, or a turn-taking algorithm preventing double targeting. This choice of policy can remove some unnecessary transmissions of the patches and hence save on the patching overhead, but it should be immediately clear that its implementation involves some extra effort. Note that we naturally assume that the network does not know a priori with certainty which nodes are infective, and hence it cannot differentiate between susceptible and infective nodes. Thus, even when $\pi = 0$, i.e. the system manager knows that the patch cannot remove the infection and only immunizes the susceptible, still the best it may be able to do is to forward the message to any node which has not previously received it.

Combining (2), (11), (13), we obtain a system of (nonlinear) differential equations that involve only the state and adjoint functions (and not the control $\mathbf{u}(\cdot)$), and the initial values of the states (eq.(3)) and the final values of the adjoint functions (eq.(12)) are known. Numerical methods for solving *boundary value* nonlinear differential equation problems may now be used to solve for the state and adjoint functions corresponding to the optimal control, which will provide the optimal controls using (13).

B. Structure of Optimal Non-Replicative Patching

We now analytically prove that the optimal controls $(u_1(t), \ldots, u_M(t))$ follow simple structures. We separately consider the cases of concave and convex $h^i(\cdot)$.

Theorem 1. *When for all i, $h_1^i(\cdot)$ and $h_2^i(\cdot)$ are concave, then for each i, there exists a $0 \leq t_i \leq T$, such that an optimal $u_i(t)$ can be expressed as $u_i(t) = u_{i,\max}$ for $0 < t < t_i$, and $u_i(t) = 0$ for $t_i \leq t \leq T$.*

Intuitively, at the onset of the epidemic, a large fraction of the nodes are susceptible to the malware, each of which is a potential victim. Bandwidth and power resources hence should be used maximally in the beginning (in all regions), rendering as many infective and susceptible nodes robust against the malware. Specifically, there is no gain in deferral of patching since the efficacy of healing infectives is less than that of the immunization of the susceptible nodes (recall that $\pi_i \leq 1$). The fact that the process of curbing the patching in this case is abrupt rather than gradual is however less apparent. The drop time in each region differs and depends on the location of the initial infection as well as the topology of the network, communication rates, etc.

Proof: When all $h_1^i(\cdot)$ and $h_2^i(\cdot)$ are concave functions then the minimization in (13) is achieved by comparing the values for the following two candidates: 0, $u_{i,\max}$. Let

$$\varphi_i(u_i) := h_1^i(u_i) + R_i^0 h_2^i(u_i) \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j) - \phi_i u_i.$$

Thus, the condition for an optimal u_i is:

$$u_i = \begin{cases} u_{i,\max} & \varphi_i(u_{i,\max}) < 0 \\ 0 & \varphi_i(u_{i,\max}) > 0 \end{cases} \tag{15}$$

Note that for all i, $\phi_i(T) = 0$. From lemma 1 $S_j(T) > 0$ for all j. Thus, $\varphi_i(u_{i,\max})(T) > 0$. $\varphi_i(u_{i,\max})(\cdot)$ is a continuous function of time. We will next show that $\varphi_i(u_{i,\max})(\cdot)$ is a strictly increasing function of time as well. Therefore, either (i) $\varphi_i(u_{i,\max})(t) > 0$ at all $t \in (0, T)$, or (ii) $\varphi_i(u_{i,\max})(t) < 0$ at all $t \in (0, t_i)$, and $\varphi_i(u_{i,\max})(t) > 0$ at all $t \in (t_i, T)$ for some $t_i \in (0, T)$. The Theorem follows from (15) in both cases.

$\varphi_i(u_{i,\max})(\cdot)$ is differentiable whenever $\mathbf{u}(\cdot)$ is continuous. At any such t, we have:

$$\frac{1}{R_0^i u_{i,\max}} \dot{\varphi}(u_{i,\max}) = \frac{1}{R_0^i u_{i,\max}} \frac{d}{dt}\Big(h_1^i(u_{i,\max})$$
$$+ R_0^i h_2^i(u_{i,\max}) \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j) - \phi_i u_{i,\max}\Big)$$

Using the definition of the Hamiltonian in Section III-A, the derivatives of $\mathbf{S}(\cdot), \mathbf{I}(\cdot)$ from (2), we have:

$$\begin{aligned}
\frac{1}{R_0^i u_{i,\max}} \dot{\varphi}(u_{i,\max}) &= \frac{h_2^i(u_{i,\max})}{u_{i,\max}} \sum_{j=1}^{M} \bar{\beta}_{ij}(\dot{S}_j + \dot{I}_j) \\
&\quad - \sum_{j=1}^{M} \bar{\beta}_{ij} \dot{\lambda}_j^S S_j - \sum_{j=1}^{M} \bar{\beta}_{ij} \lambda_j^S \dot{S}_j \\
&\quad - \sum_{j=1}^{M} \pi_j \bar{\beta}_{ij} \dot{\lambda}_j^I I_j - \sum_{j=1}^{M} \pi_j \bar{\beta}_{ij} \lambda_j^I \dot{I}_j \\
= \frac{h_2^i(u_{i,m})}{u_{i,m}} \sum_{j=1}^{M} \bar{\beta}_{ij} &\left(- \sum_{k=1}^{M} S_j \bar{\beta}_{kj} R_0^k u_k - \sum_{k=1}^{M} I_j \pi_j \bar{\beta}_{kj} R_0^k u_k \right) \\
&\quad - \sum_{j=1}^{M} \bar{\beta}_{ij} S_j \Bigg(- L_j(\mathbf{R}) - \sum_{k=1}^{M} \bar{\beta}_{kj} R_0^k h_2^k(u_k) \\
&\quad - (\lambda_j^I - \lambda_j^S) \sum_{k=1}^{M} \beta_{kj} I_k + \lambda_j^S \sum_{k=1}^{M} \bar{\beta}_{kj} R_k^0 u_k \Bigg) \\
&\quad - \sum_{j=1}^{M} \bar{\beta}_{ij} \lambda_j^S \left(- \sum_{k=1}^{M} \beta_{kj} I_k S_j - \sum_{k=1}^{M} S_j \bar{\beta}_{kj} R_0^k u_k \right) \\
&\quad - \sum_{j=1}^{M} \pi_j \bar{\beta}_{ij} I_j \Bigg(- L_j(\mathbf{R}) - f_j(\mathbf{I}) - \sum_{k=1}^{M} \bar{\beta}_{kj} R_0^k h_2^k(u_k) \\
&\quad - \sum_{k=1}^{M} \left(\lambda_k^I - \lambda_k^S\right) \beta_{jk} S_k + \lambda_j^I \sum_{k=1}^{M} \pi_j \bar{\beta}_{kj} u_k \Bigg) \\
&\quad - \sum_{j=1}^{M} \pi_j \bar{\beta}_{ij} \lambda_j^I \left(- \sum_{k=1}^{M} \beta_{kj} I_k S_j - \sum_{k=1}^{M} I_j \pi_j \bar{\beta}_{kj} R_0^k u_k \right) \\
&= \sum_{j=1}^{M} \bar{\beta}_{ij} S_j L_j(\mathbf{R}) + \sum_{j=1}^{M} \pi_j \bar{\beta}_{ij} I_j \left(L_j(\mathbf{R}) + f_j(\mathbf{I})\right) \\
&\quad + \sum_{j,k=1}^{M} \bar{\beta}_{ij}(S_j + \pi_j I_j) \bar{\beta}_{kj} R_0^k \left(h_2^k(u_k) - \frac{u_k h_2^k(u_{i,m})}{u_{i,m}} \right) \\
+ \sum_{j,k=1}^{M} \bar{\beta}_{ij}(1 + \pi_j) \lambda_k^I \beta_{jk} S_j I_k &+ \sum_{j,k=1}^{M} \pi_j \bar{\beta}_{ij} I_j S_k \beta_{kj} (\lambda_k^I - \lambda_k^S)
\end{aligned}$$

and thus, at the points of continuity of $\mathbf{u}(\cdot)$, we have:

$$\frac{1}{R_0^i u_{i,\max}} \dot{\varphi}(u_{i,\max}) \geq \sum_{j,k=1}^{M} \bar{\beta}_{ij}(1 + \pi_j) \lambda_k^I \beta_{jk} S_j I_k$$
$$+ \sum_{j,k=1}^{M} \pi_j \bar{\beta}_{ij} I_j S_k \beta_{kj} (\lambda_k^I - \lambda_k^S)$$

The inequality follows from the non-negativity of the states and the following trivial property of $h(.)$ function:

Property: For any continuous, increasing and concave function $f(\cdot)$ such that $f(0) = 0$ we have $af(b) \geq bf(a)$ if $a \geq b \geq 0$.

The right-hand-side of the last inequality is strictly positive for any $t \in (0, T)$ because:

Lem. 2. *For all $i = 1, \ldots, M$ and all $t \in (0, T)$, we have $\lambda_i^I > 0$ and $(\lambda_i^I - \lambda_i^S) > 0$.*

Thus, since $\mathbf{u}(\cdot)$ has finite number of points of discontinuity and also because $\varphi(u_{i,\max})$ is continuous, $\varphi(u_{i,\max})$ is strictly increasing in $(0, T)$. The theorem therefore follows once we prove Lem. 2. Intuitively, λ_i^I ought to be positive as it represents the additional cost the system incurs per unit time with increase in the fraction of the infectives. Also, an increase in the fraction of the infectives has worse long term implications for the system than that of the susceptibles, hence, $(\lambda_i^I - \lambda_i^S) > 0$. The formal proof confirms this intuition. ■

Proof: $\lambda_i^I|_{t=T} = (\lambda_i^I - \lambda_i^S)|_{t=T} = 0$ and $\dot{\lambda}_i^I|_{t=T} = -L_i(\mathbf{R}(T)) - f_i(\mathbf{I}(T)) - \sum_{j=1}^M \bar{\beta}_{ji} R_0^j h_2^j(u_j(T)) < 0$ and $(\dot{\lambda}_i^I - \dot{\lambda}_i^S)|_{t=T} = -f_i(\mathbf{I}(T)) < 0$, for all i. Hence, $\exists \epsilon > 0$ s.t. $\lambda_i^I > 0$ and $(\lambda_i^I - \lambda_i^S) > 0$ over $(T - \epsilon, T)$. Now suppose that (at least) one of the inequalities is first (going backward in time from $t = T$) violated at $t = t^*$ for i^*. At such a point $(\lambda_{i^*}^I - \lambda_{i^*}^S) = 0$ and $(\lambda_i^I - \lambda_i^S) \geq 0$ for all $i \neq i^*$ and $\lambda_i^I \geq 0$ for all i. Now,

$$(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)|_{t^{*+}} = -f_i(\mathbf{I}) - \sum_{j=1}^{M} [(\lambda_j^I - \lambda_j^S)\beta_{i^*j} S_j] \\ - (1 - \pi_i)\lambda_{i^*}^I \left(\sum_{j=1}^{M} \bar{\beta}_{ji} R_0^j u_j \right)$$

First of all, $-f_i(\mathbf{I}) < 0$. Also, the terms $-\sum_{j=1}^M [(\lambda_j^I - \lambda_j^S)\beta_{i^*j} S_j]$ and $-(1 - \pi)\lambda_{i^*}^I u_{i^*}$ are non-positive, according to the definition of t^*. Thus $(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)|_{t^{*+}} < 0$ which is in contradiction with our supposition. Let $\lambda_{i^*}^I = 0$ for an i^* and hence $\lambda_i^I \geq 0$ and $(\lambda_i^I - \lambda_i^S) \geq 0$ for all i. But then:

$$\dot{\lambda}_{i^*}^I|_{t^{*+}} = -L_i(\mathbf{R}) - f_i(\mathbf{I}) - \sum_{j=1}^{M} \bar{\beta}_{ji} R_0^j h_2^j(u_j) \\ - \sum_{j=1}^{M} [(\lambda_j^I - \lambda_j^S)\beta_{i^*j} S_j]$$

which is again negative according to the definition of t^*. Hence the supposition could not be true. The claim follows. ■

We now consider the case that $h_1^i(\cdot)$ is convex for all i and strictly convex for some i. For simplicity, we assume that $h_2^i(\cdot)$ is a zero function for all i. In this case, the minimization in (13) may also be attained at an interior point of $(0, u_{i,\max})$ (besides 0 and $u_{i,\max}$) at which the partial derivative of the right hand side with respect to u_i is zero. Hence,

$$u_i = \begin{cases} u_{i,\max} & u_{i,\max} < \eta \\ \eta & 0 < \eta \leq u_{i,\max} \\ 0 & \eta \leq 0. \end{cases} \tag{16}$$

where η is such that $h_1^{i'}(\eta) = \phi_i$ and ϕ_i has been defined in (14). In this case, the structure of optimal u_i for each i is similar to the concave case, except that the transition between extreme values is continuous rather than abrupt:

Theorem 2. *When for each i, $h_1^i(\cdot)$ is convex (and strictly convex at some i), and $h_2^i(\cdot)$s are zero functions, there exist for each i, t_i^1 and t_i^2, $0 \leq t_i^1, t_i^2 \leq T$, such that $u_i(t) = u_{i,\max}$ for $0 < t < t_i^1$, and $u_i(t)$ is continuously and strictly decreasing for $t_i^1 \leq t \leq t_i^2$, and $u_i(t) = 0$ for $t_i^2 \leq t \leq T$.*

Proof: From (16), it is clear that $u_i(t)$ is a continuous function of time. For the interval over which $0 < \eta \leq u_{i,\max}$, from the implicit equation $h_1^{i'}(\eta) = \phi_i$, we have:

$$\dot{u}_i h_1^{i''}(u_i) = \dot{\phi}_i \Rightarrow \dot{u}_i h_1^{i''}(u_i) = -(-\dot{\phi})$$

Following similar steps as for the concave case, we will show that the right hand side is strictly negative. Also note that $h_1^{i''}(u_i)$is strictly positive as h^is are strictly convex here. Thence the theorem follows.

$$-\dot{\phi}_i = (\lambda_i^I - \lambda_i^S) \sum_{j=1}^{M} \beta_{ji} I_j S_i - \lambda_i^S u_i S_i \\ -\lambda_i^S(-\sum_{j=1}^{M} \beta_{ji} I_j S_i - u_i S_i) \\ -\pi(-f_i(\mathbf{I})I_i - \sum_{j=1}^{M} (\lambda_j^I - \lambda_j^S)\beta_{ij} S_j I_i + \lambda_i^I \pi u_i I_i) \\ -\pi\lambda_i^I(\sum_{j=1}^{M} \beta_{ji} I_j S_i - \pi u_i I_i) \\ = (1 - \pi)\lambda_i^I \sum_{j=1}^{M} \beta_{ji} I_j S_i + \pi f_i(\mathbf{I}) I_i + \pi \sum_{j=1}^{M} (\lambda_j^I - \lambda_j^S)\beta_{ij} S_j I_i$$

The right hand side is positive according to lemma 2. Note that in the proof of the lemma 2, we did not use any assumption on the concavity/convexity of the function $h(.)$. ■

IV. Optimal Replicative Patching

As in the non-replicative setting, we first develop a numerical framework for calculation of the optimal solutions using PMP, and subsequently we establish the structure of the optimal controls. For simplicity, we assume both $h_1^i(\cdot)$ and $h_2^i(\cdot)$ to be linear functions. Thus,

$$J_R = \int_0^T f(\mathbf{I}) + \sum_{i=1}^{M} K_{1i} R_i u_i + \sum_{i=1}^{M} K_{2i} R_i u_i \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j)$$

The new Hamiltonian is:

$$
\begin{aligned}
\mathcal{H} &:= f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i}^{M} K_{1i} u_i R_i \\
&+ \sum_{i=1}^{M} K_{2i} R_i u_i \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j) + \sum_{i=1}^{M} [(\lambda_i^I - \lambda_i^S) S_i \sum_{j=1}^{M} \beta_{ji} I_j] \\
&+ \sum_{i=1}^{M} (\lambda_i^R - \lambda_i^S) S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_j u_j + \sum_{i=1}^{M} (\lambda_i^R - \lambda_i^I) I_i \sum_{j=1}^{M} \pi_i \bar{\beta}_{ji} R_j u_j
\end{aligned}
$$

where at each point of continuity of $\mathbf{u}(\cdot)$ for all $i = 1 \ldots M$ the *adjoint* functions $\lambda_i^S, \lambda_i^I, \lambda_i^R$ satisfy

$$
\dot{\lambda}_i^S = -\frac{\partial \mathcal{H}}{\partial S},\ \dot{\lambda}_i^I = -\frac{\partial \mathcal{H}}{\partial I},\ \dot{\lambda}_i^R = -\frac{\partial \mathcal{H}}{\partial R} \qquad (17)
$$

with the final constraints:

$$
\lambda_i^S(T) = \lambda_i^I(T) = \lambda_i^R(T) = 0. \qquad (18)
$$

According to PMP, any optimal controller must satisfy:

$$
u_i \in \arg \min_{[u_{\min}, u_{\max}]} \mathcal{H}
$$

Hence,

$$
u_i \in \arg\min \left(K_{1i} R_i u_i + K_{2i} R_i u_i \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j) + \phi_i u_i \right) \qquad (19)
$$

where

$$
\phi_i := R_i \sum_{j=1}^{M} \bar{\beta}_{ij} (\lambda_j^R - \lambda_j^S) S_j + R_i \sum_{j=1}^{M} \pi_j \bar{\beta}_{ij} (\lambda_j^R - \lambda_j^I) I_j.
$$

Let $\varphi_i := K_{1i} R_i + K_{2i} R_i \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j) + \phi_i$. Then the optimality condition is reduced to

$$
u_i = \begin{cases} u_{i,\max} & \varphi_i < 0, \\ 0 & \varphi_i > 0. \end{cases}
$$

Combining the above with (5), (17) we will again have a system of (non-linear) differential equations that involve only the state and adjoint functions (and not the control $\mathbf{u}(\cdot)$), and the initial values of the states and the final values of the adjoint functions (eq.(6) and eq.(18) respectively). Similarly to the non-optimal case, the optimal controls may now be obtained by solving the above system of differential equations.

The optimal controls for replicative patching exhibit similar structure as that in the non-replicative setting:

Theorem 3. *When for all i, $h_1^i(\cdot)$ and $h_2^i(\cdot)$ are linear, then for each i, there exists a $0 \le t_i \le T$, such that an optimal $u_i(t)$ can be expressed as $u_i(t) = u_{\max}$ for $0 < t < t_i$, and $u_i(t) = 0$ for $t_i \le t \le T$.*

Proof: We have:

$$
\begin{aligned}
\dot{\varphi}_i|_{\varphi_i = 0} &= K_{1i} \dot{R}_i + K_{2i} \Bigg(\dot{R}_i \sum_{j=1}^{M} \bar{\beta}_{ij}(S_j + I_j) \\
&\qquad + R_i \sum_{j=1}^{M} \bar{\beta}_{ij}(\dot{S}_j + \dot{I}_j) \Bigg) + \dot{\phi}_i \\
&= \dot{R}_i \frac{\varphi_i}{R_i} + K_{2i} R_i \sum_{j=1}^{M} \bar{\beta}_{ij}(\dot{S}_j + \dot{I}_j) + R_i \sum_{j=1}^{M} \bar{\beta}_{ij} [(\dot{\lambda}_j^R - \dot{\lambda}_j^S) S_j \\
&\qquad + (\lambda_j^R - \lambda_j^S) \dot{S}_j + \pi_j (\dot{\lambda}_j^R - \dot{\lambda}_j^I) I_j + \pi_j (\lambda_j^R - \lambda_j^I) \dot{I}_j] \\
&= K_{2i} R_i \Bigg(- \sum_{j=1}^{M} \bar{\beta}_{ij} S_j \sum_{k=1}^{M} \bar{\beta}_{kj} R_k u_k \\
&\qquad - \sum_{j=1}^{M} \bar{\beta}_{ij} I_j \sum_{k=1}^{M} \pi_j \bar{\beta}_{kj} R_k u_k \Bigg) \\
&+ R_i \sum_{j=1}^{M} \bar{\beta}_{ij} [L_j(\mathbf{R}) S_j + \sum_{k=1}^{M} K_{2k} \bar{\beta}_{kj} R_k u_k S_j \\
&\qquad + \sum_{k=1}^{M} (\lambda_j^I - \lambda_j^S) \beta_{kj} I_k S_j + (\lambda_j^R - \lambda_j^S) \sum_{k=1}^{M} \bar{\beta}_{kj} R_k S_j u_k \\
&\qquad + (\lambda_j^R - \lambda_j^S)(- \sum_{k=1}^{M} \beta_{kj} I_k S_j - S_j \sum_{k=1}^{M} \bar{\beta}_{kj} R_k u_k) \\
&\qquad + \pi_j (L_j(\mathbf{R}) I_j + f_j(\mathbf{I}) I_j + \sum_{k=1}^{M} K_{2k} \bar{\beta}_{kj} R_k u_k I_j \\
&+ I_j \sum_{k=1}^{M} (\lambda_k^I - \lambda_k^S) \beta_{jk} S_k + (\lambda_j^R - \lambda_j^I) I_j \sum_{k=1}^{M} \pi_j \bar{\beta}_{kj} R_k u_k \\
&\qquad + (\lambda_j^R - \lambda_j^I)(\sum_{k=1}^{M} \beta_{kj} I_k S_j - I_j \sum_{k=1}^{M} \pi_j \bar{\beta}_{kj} R_k u_k))] \\
&= R_i \sum_{j=1}^{M} \bar{\beta}_{ij} [(1 - \pi_j) \sum_{k=1}^{M} (\lambda_j^I - \lambda_j^R) \beta_{kj} I_k S_j + \pi_j f_j(\mathbf{I}) I_j \\
&\qquad (S_j + \pi I_j) L_j(\mathbf{R}) + \pi_j I_j \sum_{k=1}^{M} (\lambda_k^I - \lambda_k^S) \beta_{jk} S_k]
\end{aligned}
$$

Now, the right hand side of the last equation is strictly positive according in part to the following lemma.

Lem. 3. *For all $t \in [0, T)$, for all i we have $(\lambda_i^I - \lambda_i^S) > 0$ and $(\lambda_i^I - \lambda_i^R) > 0$.*

Proof: The proof is similar to that of lemma 2. $(\lambda_i^I - \lambda_i^S)|_{t=T} = 0$ and $(\dot{\lambda}_i^I - \dot{\lambda}_i^S)|_{t=T} = -f_i(\mathbf{I}) < 0$, for all i. Also, $(\lambda_i^I - \lambda_i^R)|_{t=T} = 0$ and $(\lambda_i^I - \dot{\lambda}_i^R)|_{t=T} = -f_i(\mathbf{I}) - L_i(\mathbf{R}) < 0$. Hence, $\exists \epsilon > 0$ s.t. $(\lambda_i^I - \lambda_i^S) > 0$ and $(\lambda_i^I - \lambda_i^R) > 0$ over $(T - \epsilon', T)$.

Now suppose that (at least) one of the inequalities is first[11] violated at $t = t^*$ for i^*. Suppose it is that $(\lambda_{i^*}^I - \lambda_{i^*}^S) = 0$ and $(\lambda_i^I - \lambda_i^S) \ge 0$ for all $i \ne i^*$ and $(\lambda_i^I - \lambda_i^R) \ge 0$ for all

[11] going backward in time from $t = T$

i. Now, let us investigate $(\dot\lambda^I_{i^*} - \dot\lambda^S_{i^*})|_{t^*+}$:

$$(\dot\lambda^I_{i^*} - \dot\lambda^S_{i^*})|_{t^*+} = -f_i(\mathbf{I}) - \sum_{j=1}^{M}(\lambda^I_j - \lambda^S_j)\beta_{i^*j}S_j - (1-\pi_{i^*})(\lambda^I_{i^*} - \lambda^R_{i^*})\sum_{j=1}^{M}\bar\beta_{ji^*}R_j u_j$$

First of all, $-f_i(\mathbf{I}) < 0$. Also, the terms $-\sum_{j=1}^{M}[(\lambda^I_j - \lambda^S_j)\beta_{i^*j}S_j]$ and $-(1-\pi_{i^*})(\lambda^I_{i^*} - \lambda^R_{i^*})\sum_{j=1}^{M}\bar\beta_{ji^*}R_j u_j$ are non-positive, according to the definition of t^*. Hence, $(\dot\lambda^I_{i^*} - \dot\lambda^S_{i^*})|_{t^*+} < 0$ which is in contradiction with our supposition, and hence $(\lambda^I_i - \lambda^S_i) > 0$ for all i, since i^* was arbitrary.

Now suppose that the first violated inequality is $(\lambda^I_{i^*} - \lambda^R_{i^*}) = 0$ and $(\lambda^I_i - \lambda^S_i) \geq 0$ for all i and $(\lambda^I_i - \lambda^R_i) \geq 0$ for all $i \neq i^*$. Now, let us investigate $(\dot\lambda^I_{i^*} - \dot\lambda^R_{i^*})|_{t^*+}$:

$$(\dot\lambda^I_{i^*} - \dot\lambda^R_{i^*})|_{t^*+} = -f_i(\mathbf{I}) - L_i(\mathbf{R}) - \sum_{j=1}^{M}(\lambda^I_j - \lambda^S_j)\beta_{i^*j}S_j + \frac{\varphi_i u_i}{R_i}$$

The first terms $-f_i(\mathbf{I})$ and $-L_i(\mathbf{R})$ are trivially negative. The next term $-(\lambda^I_i - \lambda^S_i)\sum_{j=1}^{M}\beta_{ji^*}S_j$ is non-positive, according to the definition of t^*. The last term is non-positive according to PMP, from (19). This shows that $(\dot\lambda^I_{i^*} - \dot\lambda^R_{i^*})|_{t^*+} < 0$ which is in contradiction with existence of t^*, and hence $(\lambda^I_i - \lambda^R_i) > 0$ for all i, since i^* was arbitrary. This completes the proof. ■

■

V. Numerical Investigations

In this section, we numerically investigate the optimal control policies for a range of malware and network parameters.[12] We consider three topologies: *linear*, *star* and *complete*, as was illustrated in fig. 1. Specifically, there is a link between two regions i, j $i \neq j$ if and only if the communication rate between them $\beta_{ij} = \bar\beta_{ij} = \beta_{ji} = \bar\beta_{ji} \neq 0$. At time zero, we assume that only one of the regions (types) is infected, i.e. $I^i_0 > 0$ for only $i = 1$. Also, $R^0_i = 0.2$, $\beta_{ii} = \beta = 0.223$ for all i.[13] The value of β_{ij}, $i \neq j$ is equal to $X_{Coef} \cdot \beta$ if link ij is part of the regional topology graph, and is zero otherwise. We examine two different aggregate cost structures, for non-replicative patching: (cost-A) $\int_0^T \left(K_I \sum_{i=1}^{M} I_i(t) + K_u \sum_{i=1}^{M} R^i_0 u_i(t)\right) dt$, and (cost-B): $\int_0^T \left(K_I \sum_{i=1}^{M} I_i(t) + K_u \sum_{i=1}^{M} R^i_0 u_i(t)(S(t) + I(t))\right) dt$, where in both cases we take $T = 35$, $K_I = 1$, $K_u = 0.5$ unless stated otherwise. For replicative patching, R^i_0 in both cost models is replaced with $R_i(t)$.

First, with the intention of illustrating our analytical results, in fig. 2 we have depicted an example of the optimal dynamic patching policy along with the corresponding evolution of the levels of infection as a function of time. In this example, we are looking at a simple 3-region linear topology where the infection starts in region 1 with $I^1_0 = 0.3$, and $I^2_0, I^3_0 = 0$. X_{Coef} is taken to be 0.1, i.e. the internal mixing rate in each region is ten times the cross-region mixing rates. The cost model is of type-A and patching is non-replicative. Note that for $\pi = 0$ the levels of infection are non-decreasing, whereas for $\pi = 1$ they may go down as well as up.

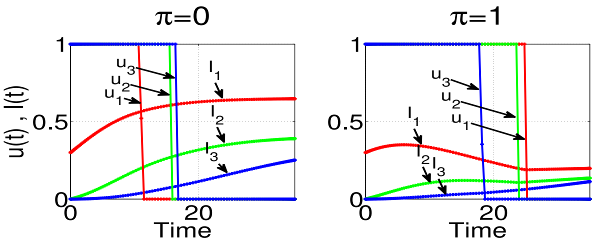

Fig. 2. Illustration of optimal patching policies and the corresponding levels of infection in each of the three regions for a simple linear topology. Note how the infection that initially only exists in region 1 spreads in region 1 and then to region 2, and from there to region 3.

We now investigate the effect of topology on the optimal patching policy. We study the *drop-off* times (the time thresholds at which the patching halts) in different regions for linear and star topologies. Fig. 3 reveals two different patterns for $\pi = 0$ and $\pi = 1$ in a linear topology with 10 regions. For $\pi = 0$, the middle region is patched for the longest time, whereas for $\pi = 1$, as we go farther from the origin of the infection (region 1), the drop-off point decreases. The reason for this is that for $\pi = 0$, patching can only benefit the network by recovering the susceptibles. In regions closer to the origin of the infection, the fraction of the susceptibles decreases quickly as a result of the spread of the infection, making continuation of the patching comparatively less beneficial. In the middle regions, where there are more susceptibles at risk of contamination, patching should be continued longer. For regions far from the origin, patching can be stopped earlier, as even when the susceptibles are not immunized, infection barely reaches them within the time horizon of consideration. For $\pi = 1$, the patching is able to recover both susceptible and infective nodes. Hence, the drop-off times depend only on the exposure to the infection which decreases as the distance from the origin of the infection increases. An interesting phenomenon is that as X_{Coef} is increased, the value of the drop-off points in the $\pi = 1$ case get closer together. Intuitively, this is because higher cross-mixing rates have a homogenizing effect, as the levels of susceptibles and infectives in different region rapidly become comparable. Also, fig. 3 reveals that as X_{Coef} increases and more infection reaches the farther regions, they are patched for longer durations, which agrees with the intuition.

We next investigate a star configuration where the infection starts from a peripheral region (region 1). Fig. 4 reveals the following interesting phenomenon: although the central region is the only one that is connected to all the regions, for $\pi = 0$ it is patched for shorter times compared to the peripherals.

[12] For our calculations, we use a combination of *C* programming and *PROPT*®, by *Tomlab Optimization Inc* for *MATLAB*®.

[13] Specific value of β is chosen to match the average inter-meeting times from the numerical experiment reported in [23].

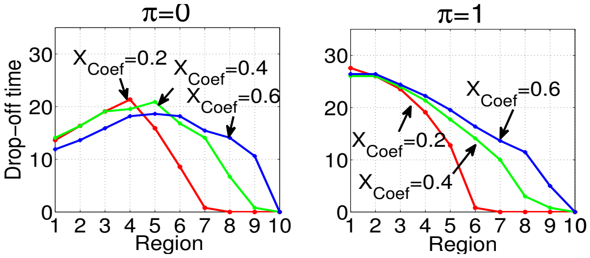

Fig. 3. The drop-off times of different regions in a linear topology with $M = 10$ regions. We consider type-A cost function and non-replicative patching with $X_{Coef} = 0.2, 0.4, 0.6$.

In retrospect, this is again because only the susceptible nodes can be patched and their number at the central region drops quickly due to its interactions with all the peripheral regions, rendering the patching inefficient relatively swiftly. Following this explanation, as expected this effect is amplified with higher number of peripheral regions. For $\pi = 1$, on the other hand, the central node is patched for the longest time. This is because the infective nodes there can infect the susceptible nodes in all of the regions, and hence the patching, which can now heal the infectives as well, does not stop until it heals almost all of infectives in this region.

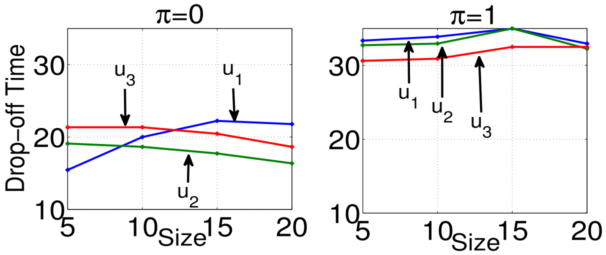

Fig. 4. Trends in the drop-off points in the star topology. The cost is type-B and the patching is non-replicative, and $I_0^1 = 0.6$.

Next, in order to evaluate the efficacy of our dynamic heterogeneous patching policy, we compare our inflicted aggregate cost against those of three alternative patching policies. We label our policy as *Stratified Dynamic* (S.D.). In the simplest alternative policy, all regions use identical patching intensities which do not change with time. We then select this fixed and static level of patching so as to minimize the aggregate cost among all possible choices. We refer to this policy as *Static* (St. in short). The aggregate cost may be reduced if the static value of the patching is allowed to be distinct for different regions. These values (still fixed over time) are then independently varied and the best combination is selected. We refer to this policy as *Stratified Static* (S.St. in short). The third policy we implement is a *homogeneous* approximation to the heterogeneous network. Specifically, the whole network is approximated by a single region model with an equivalent inter-contact rate. This value is selected such that the *average* pairwise contact rates are equal in both systems. The optimal control is hence derived based on this model and applied across all regions to calculate the aggregate cost. We call this policy *Simplified Homogeneous* (H. in short).

Fig. 5 depicts the aggregate costs of all four policies for the linear topology with $M = 2, 3, 4$ and 5 number of regions. The cost is type-A and patching is replicative. Here, $I_0^1 = 0.2$, $K_u = 2.5$ and the rest of parameters are as before. As we can clearly observe, our stratified policy achieves the least costs, outperforming the rest. Also of note is that H. is doing better than S. St., which outperforms St. Notably, as the number of the regions increases and the network becomes more spatially heterogeneous, the homogeneous approximation worsens. For example for $M = 5$ regions, our policy outperforms the static policies by 40% and the homogeneous approximation by 20% for $\pi = 0$. For $\pi = 1$, the performance gap widens to 40% and 65% respectively. This underscores the significance of considering the heterogeneity in the controls. Specifically, as we discussed before, optimal drop-off times should vary based on the distance from the originating region, a factor which the S.H. (and St.) policies ignore.

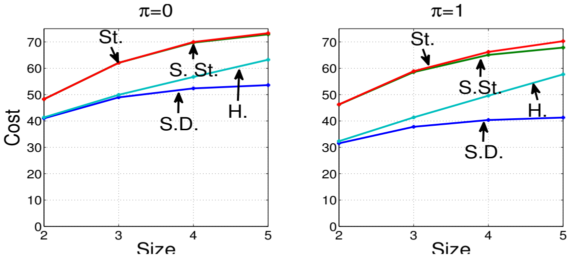

Fig. 5. Linear Topology, trends in the cost.

We repeat the same experiment in a complete topology, again varying the number of regions between 2, 3, 4, and 5 and report the results in fig. 6. As before, our stratified dynamic policy incurs less aggregate cost compared to the rest. In the complete topology, as one could expect, we observe that the homogeneous approximation performs close to the optimal. The contrast in the relative performance of the homogeneous approximation between the linear and complete cases is a testament to the significance of the effect of topology on optimal patching.

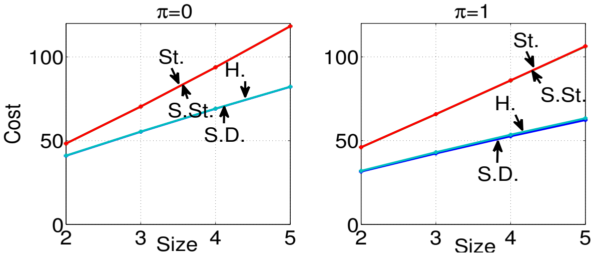

Fig. 6. Complete topology, trends in the cost.

VI. Conclusion

We considered the problem of disseminating security patches in a large resource-constrained heterogeneous network in the mean-field regime. Using tools from optimal control theory, we analytically proved that optimal dynamic policies for each type of node follow simple threshold-based structures, making them amenable to distributed implementation. We numerically demonstrated the advantage of our heterogeneous policies over homogeneous approximations, as well as over static policies. For future research, we would like to further investigate the effects of heterogeneities in the structure of networks on both defense and attack strategies.

REFERENCES

[1] K. Ramachandran and B. Sikdar, "On the stability of the malware free equilibrium in cell phones networks with spatial dynamics," in *ICC'07*, pp. 6169–6174, 2007.

[2] P. Wang, M. González, C. Hidalgo, and A. Barabási, "Understanding the spreading patterns of mobile phone viruses," *Science*, vol. 324, no. 5930, p. 1071, 2009.

[3] Z. Zhu, G. Cao, S. Zhu, S. Ranjan, and A. Nucci, "A social network based patching scheme for worm containment in cellular networks," in *INFOCOM'09*, pp. 1476–1484, IEEE, 2009.

[4] M. Khouzani, S. Sarkar, and E. Altman, "Dispatch then stop: Optimal dissemination of security patches in mobile wireless networks," in *IEEE CDC'10*, pp. 2354–2359, 2010.

[5] M. Khouzani, S. Sarkar, and E. Altman, "Optimal control of epidemic evolution," in *IEEE INFOCOM*, 2011.

[6] J. Mickens and B. Noble, "Modeling epidemic spreading in mobile environments," in *Proceedings of the 4th ACM Workshop on Wireless Security*, pp. 77–86, ACM, 2005.

[7] K. Ramachandran and B. Sikdar, "Modeling malware propagation in networks of smart cell phones with spatial dynamics," in *IEEE INFOCOM'07*, pp. 2516–2520, 2007.

[8] Z. Chen and C. Ji, "Spatial-temporal modeling of malware propagation in networks," *IEEE Transactions on Neural Networks*, vol. 16, no. 5, pp. 1291–1303, 2005.

[9] F. Li, Y. Yang, and J. Wu, "CPMC: an efficient proximity malware coping scheme in smartphone-based mobile networks," in *IEEE INFOCOM'10*, pp. 1–9, 2010.

[10] Y. Yang, S. Zhu, and G. Cao, "Improving sensor network immunity under worm attacks: a software diversity approach," in *Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*, pp. 149–158, ACM, 2008.

[11] Y. Li, P. Hui, D. Jin, L. Su, and L. Zeng, "An optimal distributed malware defense system for mobile networks with heterogeneous devices,"

[12] H. Nguyen and Y. Shinoda, "A macro view of viral propagation and its persistence in heterogeneous wireless networks," in *Fifth International Conference on Networking and Services*, pp. 359–365, IEEE, 2009.

[13] M. Liljenstam, Y. Yuan, B. Premore, and D. Nicol, "A mixed abstraction level simulation model of large-scale internet worm infestations," in *10th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems, MASCOTS 2002*, pp. 109–116, IEEE, 2002.

[14] J. Cuzick and R. Edwards, "Spatial clustering for inhomogeneous populations," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 73–104, 1990.

[15] W. Hsu and A. Helmy, "Capturing user friendship in WLAN traces," *IEEE INFOCOM poster*, 2006.

[16] T. Kurtz, "Solutions of ordinary differential equations as limits of pure jump markov processes," *Journal of Applied Probability*, pp. 49–58, 1970.

[17] N. Gast, B. Gaujal, and J. Le Boudec, "Mean field for Markov decision processes: from discrete to continuous optimization," *Arxiv preprint arXiv:1004.2342*, 2010.

[18] M. Faghani and H. Saidi, "Malware propagation in online social networks," in *4th IEEE International Conference on Malicious and Unwanted Software (MALWARE)*, pp. 8–14.

[19] M. Altunay, S. Leyffer, J. Linderoth, and Z. Xie, "Optimal response to attacks on the open science grid," *Computer Networks*, 2010.

[20] F. De Pellegrini, E. Altman, and T. Başar, "Optimal monotone forwarding policies in delay tolerant mobile ad hoc networks with multiple classes of nodes," in *IEEE WiOpt'10*, pp. 497–504, 2010.

[21] E. Altman, T. Başar, and F. De Pellegrini, "Optimal control in two-hop relay routing," *IEEE TAC*, no. 99, pp. 670 – 675, 2011.

[22] T. Small and Z. Haas, "The shared wireless infostation model: a new ad hoc networking paradigm (or where there is a whale, there is a way)," in *MobiHoc'03*, ACM, 2003.

[23] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *ACM SIGCOMM Workshop on Delay-tolerant Networking*, p. 251, ACM, 2005.