Optimal Patching in Clustered Epidemics of Malware

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

***Abstract*—Studies on the propagation of malware in mobile networks 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, the clustered nature of the network, etc. In this paper, a general formal framework is proposed for leveraging such information about heterogeneity to derive 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, it is analytically proven that in the mean-field deterministic regime, optimal patch disseminations are simple single-threshold policies that are amenable to implementation in a distributed manner. Through numerical calculations, the behavior of optimal patching policies is investigated in sample topologies and their advantages are demonstrated.**

***Index Terms*—Security, Wireless Networks, Immunization and Healing, Belief Propagation, Technology Adoption**

I. Introduction

Worms, i.e., self-propagating malicious codes, are a decades-old threat in the realm of the Internet. Worms undermine the network by performing various malicious activities: they can eavesdrop on and analyze the traversing data, access privileged information, hijack sessions, disrupt network functions 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 messages (MMS), they can also have ubiquitous access to mobile Internet and peer to peer networks via a telecommunications 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 has, in fact, been 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 fixed by installing security patches that immunize the susceptible and potentially remove the malware from the infected, hence simultaneously healing and immunizing infective nodes. However, the distribution of these patches burdens the limited resources of the network, and hence, if not carefully controlled, this can wreak havoc on the system. 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 where bandwidth is limited and is more sensitive to overload, and nodes have limited 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 that attain desired trade-offs between the efficacy of patching and the extra taxation of network resources. However, as we will explain next, these studies suffer from a common drawback; namely, a strong simplifying assumption.

Malware spreads when an infective node *contacts*, i.e., communicates with, a susceptible node, i.e., 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 significantly 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 spread 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 therefore higher amongst *friendship cliques* in the social network of users. Malware that 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

S. Eshghi, S. Sarkar and S. S. Venkatesh are with the Department of Electrical and Systems Engineering at the University of Pennsylvania, Philadelphia, PA, U.S.A. Their email addresses are *eshghi,swati,venkates@seas.upenn.edu*. MHR. Khouzani is with the Department of Electrical and Computer Engineering at the Ohio State University, Columbus, OH. His e-mail address is *khouzani@ece.osu.edu*.

This paper was presented [in part] at the IEEE Information Theory and Applications Workshop (ITA '12) , San Diego, CA, February, 2012

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, characterize and incorporate such inhomogeneities to more accurately predict the spread of infection [2], [7], [8], [13]–[16], 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 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 of static policies (i.e., those that do not vary patching rates over time) therein. Patching performance can be significantly improved if we allow the patching rates to vary dynamically in accordance with the evolution of the infection. Characterization of the optimal controls in the space of dynamic and clustered policies has, however, so far remained elusive.

We propose a formal framework for deriving *dynamic optimal* patching policies that leverage heterogeneity in the network structure to attain the minimum possible aggregate cost due to the spread of malware and the overhead of patching. We assume arbitrary (potentially non-linear) functions for the cost rates of the infective nodes. 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 that it contacts, by a mechanism similar to the spread of the malware itself. In our model, patching can immunize the susceptible nodes and may or may not heal the infective nodes. The framework in each case relies on optimal control formulations that 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 associated with the controls and a *stratified*[1] mean-field deterministic 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 of different types contact each other at rates particular to that 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, which in turn leads to multidimensional (dynamic) optimal control formulations.

Multidimensional optimal control formulations, particularly those in the solution space of functions rather than variables, are usually associated with the pitfall of amplifying the complexity of the optimization. An important contribution of the paper, therefore, is to prove that for both non-replicative and replicative settings the optimal control associated with each type has a simple structure provided the corresponding patching cost is either concave or convex. Furthermore, our analysis, leveraging Pontryagin's Maximum Principle, reveals that the structure of the optimal control for a specific type depends only on the nature of the corresponding patching cost and not on those of other types. This holds even though the control for each type affects immunization and healing in other types and the spread of the infection in general. Specifically, if the patching cost associated with the control for a given type is concave, irrespective of the nature of the patching costs for other types, the corresponding optimal control turns out to be a bang-bang function with at most one jump: up to a certain threshold time it selects the maximum possible patching rate up and subsequently it stops patching altogether. The thresholds will be different for different types and may now be computed using efficient off-the-shelf numerical techniques. If the patching cost is strictly convex, the decrease from the maximum to the minimum patching rate is continuous rather than abrupt, and monotonous. To the best of our knowledge, such simple structure results have not been established in the context of (static or dynamic) control of heterogeneous epidemics. Our numerical calculations reveal a series of interesting behaviors of optimal patching policies for different sample topologies.

As a final comment, although we focus on malware propagation, stratified or clustered epidemics can capture state evolutions in a diverse set of applications, such as technology adoption, belief propagation over social media, and health care. The framework that we propose may be used to attain desired tradeoffs between the resources consumed in applying the control and the damage induced by the proliferation of undesirable states in these contexts. This is again achieved by exploiting heterogeneities inherent to such systems through multi-dimensional dynamic optimal control. Our work advances the state of the art of dynamic optimal control in these settings by considering a model that captures any degree of inhomogeneity. Existing research in these areas either considers a homogeneous epidemic (e.g. [17], [18]), or investigates the optimal control only numerically for a limited number of distinct clusters and guarantees no structural results (e.g. [19], [20]).

II. System Model and Objective Formulation

In this section we 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, motivate the model (§II-C), formulate the aggregate cost of patching and cast the resource-aware patching as a multi-dimensional optimal control problem (§II-E). This formulation relies on a key property of the state dynamics which we isolate in §II-D. We develop solutions in this model framework and present our main results in §III and §IV.

[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.

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 yet been infected, and **recovered** if it is immune to the malware. An infective node 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* (equivalently, *clusters, segments, populations, categories, classes, compartments, strata*, etc.). A node of type i contacts another of type j at rate $\beta_{ij}^{(N)}$.

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. Therefore, 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. Dispatchers can transmit patches to both susceptible and infective nodes, *immunizing* susceptibles and possibly *healing* infectives; in either case successful transmission converts the target node to the recovered state. In *non-replicative* patching (as opposed to *replicative* patching, which is discussed in §II-B) the recipient nodes of the patch do not propagate it further.[2] Dispatchers of type i contact nodes of type j at rate $\bar{\beta}_{ij}^{(N)}$, which may be different from the contact rate $\beta_{ij}^{(N)}$ of the malware between these two types. Examples where contact rates may be different include settings where the network manager may utilize a higher priority option for the distribution of patches, ones where the malware utilizes legally restricted means of propagation not available to dispatchers, or ones where the patch is not applicable to all types, with the relevant $\bar{\beta}_{ij}^{(N)}$ now being 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 $0 < R_i^0 < 1$.

Place the time origin $t = 0$ at the earliest moment the infection is detected and the appropriate patches generated. Suppose that at $t = 0$, for each i, an initial fraction $0 \leq I_i(0) = I_i^0 < 1$ of nodes of type i are infected; we set $I_i^0 = 0$ if the infection does not initially exist amongst a type i. At the onset of the outbreak of the infection, the dispatchers are the only agents immune to the malware, hence constituting the initial population of recovered nodes. We therefore identify $R_i(0) = R_i^0$ whence the number of dispatchers $n_i^R(0) = N_i R_i^0$ also represents the initial number of recovered nodes of type i. In view of node conservation, it follows that $S_i^0 = 1 - I_i^0 - R_i^0$ represents the initial fraction $S_i(0)$ of susceptible nodes of type i.

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 type j at rate $\beta_{ji}^{(N)}$. We may fold resistance to infection into the contact rates and so we may suppose, from a modeling perspective, that susceptibles contacted by infectious agents are instantaneously infected. Accordingly, susceptibles of type i are transformed to infectives (of the same type) at an aggregate rate of $n_i^S(t) \sum_j \beta_{ji}^{(N)} n_j^I(t)$ by contact with infectives of any type.

The system manager regulates the resources consumed in the distribution of the patches by dynamically controlling the rate at which dispatchers contact susceptible and infective elements. For each j, let the control function $u_j(t)$ represent the rate of transmission attempts of dispatchers of type j at time t. We suppose that the controls are non-negative and bounded,

$$0 \leq u_j(\cdot) \leq u_{j,\max}. \tag{1}$$

We will restrict consideration to control functions $u_j(\cdot)$ that have a finite number of points of discontinuity, and for whom bounded derivatives of up to order M exist in a closed neighborhood of the origin.[3] To save on frequent repetition, we say that a control (vector) $\mathbf{u}(t) = \big(u_1(t), \ldots, u_M(t)\big)$ is *admissible* if it satisfies these conditions.

Given the controls $u_1(\cdot)$, …, $u_M(\cdot)$, susceptibles of type i are transformed to recovered nodes of the same type at an aggregate rate of $n_i^S(t) \sum_j \bar{\beta}_{ji}^{(N)} n_j^R(0) u_j(t)$ by contact with dispatchers of any type. A subtlety in the setting is that the dispatcher may find that the efficacy of the patch is lower when treating infective nodes. This may model situations, for instance, where the malware attempts to prevent the reception or installation of the patch in an infective host, or the patch is designed only to remove the vulnerability that leaves nodes exposed to the malware but does not remove the malware itself if the node is already infected. We capture such possibilities by introducing a (type-dependent) coefficient $0 \leq \pi_{ji} \leq 1$ which represents the efficacy of patching on an infective node: $\pi_{ji} = 0$ represents one extreme where a dispatcher of type j can only immunize susceptibles but can not heal infective elements of type i, while $\pi_{ji} = 1$ represents the other extreme where contact with a dispatcher of type j both immunizes and heals nodes of type i equally well; we also allow π_{ij} to assume intermediate values between the above extremes. An infective node changes state to recovered if a patch heals it; otherwise, it remains an infective. Infective nodes of type i accordingly recover at an aggregate rate of $n_i^I(t) \sum_{j=1}^M \pi_{ji} \bar{\beta}_{ji}^{(N)} n_j^R(0) u_j(t)$ by contact with dispatchers.

In the large (continuum) population regime, suppose that the two limits $\beta_{ij} := \lim_{N\to\infty} N\beta_{ij}^{(N)}$ and $\bar{\beta}_{ij} := \lim_{N\to\infty} N\bar{\beta}_{ij}^{(N)}$ exist. We say that a type j is a *neighbour* of a type i if $\beta_{ij} > 0$ (i.e., infected nodes of type i can contact nodes of type j). There is now a natural notion of a topology that is inherited from these rates with types as vertices and edges between neighboring types. Figure 1 illustrates some simple topologies. For a given topology inherited from the rates $\{\beta_{ij}, 1 \leq i, j \leq M\}$ there is now another natural notion, that of *connectivity*: we say that type j is connected

[2] And this may be preferred if the patches themselves may be contaminated and cannot be reliably authenticated.

[3] The assumption that the bounded derivatives of up to order M exist for the control functions in a closed neighborhood of the origin is the most stringent of our technical assumptions on the control functions. We need this assumption only to prove Lemma 1; specifically Lemma 6 which the proof for Lemma 1 uses. Note that if $\mathbf{I}(0) \succ 0$, Lemma 6 follows from the continuity of $\mathbf{I}(\cdot)$ even in the absence of the above assumption on differentiability.

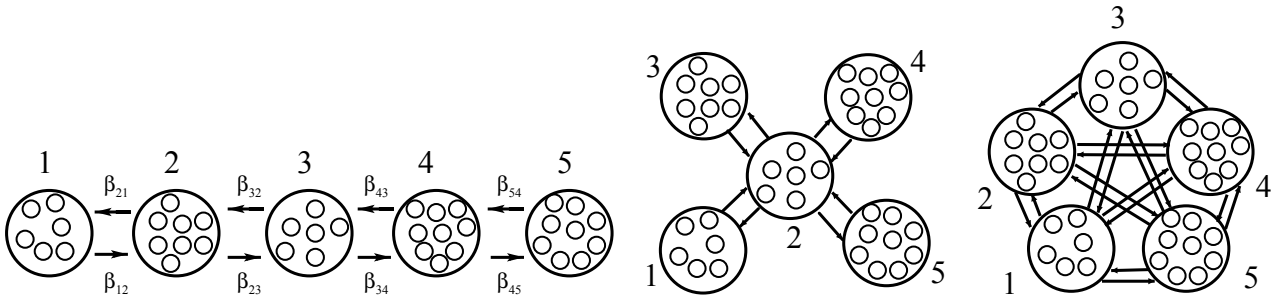

Fig. 1. Three sample topologies of 5 hotspot regions: linear, star and complete. For instance, nodes of hotspot 1 in the linear topology can only communicate with nodes of hotspots 1 and 2: they contact nodes of hotspot 1 at rate $\beta_{11}^{(N)}$ and nodes of hotspot 2 at rate $\beta_{12}^{(N)}$.

to type i if, for some k, there exists a sequence of types $i = s_1 \mapsto s_2 \mapsto \ldots \mapsto s_{k-1} \mapsto s_k = j$ where type s_{l+1} is a neighbour of type s_l for $1 \leq l < k$. We assume that for every type i, there exists a type j for which $\bar{\beta}_{ij} > 0$, i.e., type i can immunize nodes of at least one other type, and there exist types k and l for which $\beta_{ki} > 0$ and $\beta_{il} > 0$, i.e., the infection can spread to and from that type. (In most settings we may expect, naturally enough, that $\beta_{ii} > 0$ and $\bar{\beta}_{ii} > 0$.)

If the total number of nodes, N, is large and $\alpha_i > 0$ for all i,[4] then in the mean-field limit (c.f. [21, p.1] or [22]), $S_i(t)$ and $I_i(t)$ converge pathwise to the solution of the system of differential equations[5]

$$\dot{S}_i = -\sum_{j=1}^{M} \beta_{ji} I_j S_i - S_i \sum_{j=1}^{M} \bar{\beta}_{ji} R_j^0 u_j, \tag{2a}$$

$$\dot{I}_i = \sum_{j=1}^{M} \beta_{ji} I_j S_i - I_i \sum_{j=1}^{M} \pi_{ji} \bar{\beta}_{ji} R_j^0 u_j, \tag{2b}$$

where, by writing $\mathbf{S}(t) = \big(S_1(t), \ldots, S_M(t)\big)$, $\mathbf{I}(t) = \big(I_1(t), \ldots, I_M(t)\big)$, and $\mathbf{R}(t) = \big(R_1(t), \ldots, R_M(t)\big)$ in a compact vector notation, the initial conditions and state constraints are given by

$$\mathbf{S}(0) = \mathbf{S}^0 \succ \mathbf{0}, \quad \mathbf{I}(0) = \mathbf{I}^0 \succeq \mathbf{0}, \tag{3}$$

$$\mathbf{S}(t) \succeq \mathbf{0}, \quad \mathbf{I}(t) \succeq \mathbf{0}, \quad \mathbf{S}(t) + \mathbf{I}(t) \preceq \mathbf{1} - \mathbf{R}^0. \tag{4}$$

In these expressions $\mathbf{0}$ and $\mathbf{1}$ represent vectors all of whose components are 0 and 1, respectively, and the vector inequalities are to be interpreted as component-wise inequalities. Note that the evolution of $\mathbf{R}(t)$ need not be explicitly considered since at any given time, node conservation gives $R_i(t) = 1 - S_i(t) - I_i(t)$. We henceforth drop the dependence on t and make it implicit whenever we can do so without ambiguity. The initial conditions $\mathbf{S}^0 \succ \mathbf{0}$ and $\mathbf{I}^0 \succeq \mathbf{0}$ given in (3) require that each type has a non-vanishing fraction of susceptible agents initially, though we permit some of the types to be initially infection-free. This allows us to examine the propagation of infection and the evolution of the optimal controls in initially uninfected regions.

Let $U := \{\, i : I_i^0 > 0 \,\}$ be the set of the initially infected types. To obviate trivialities, we assume that each type not in U is connected to some type in U (i.e., the topology allows the infection to be propagated to all regions).

[4]So that each type comprises a non-vanishing fraction of the total population of the nodes.

[5]We use dots to denote *time* derivatives throughout the paper, $\dot{S}_i(t) = dS_i(t)/dt$, etc.

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, whence the fraction of dispatchers of type i grows from $R_i(0) = R_i^0$ initially to $R_i(t)$ at time t. This should be contrasted with the non-replicative model in which the fraction of dispatchers of type i is fixed at R_i^0 for all time. The system dynamics equations given in (2) for the non-replicative setting now need to be modified to take into account the growing pool of dispatchers. While in the non-replicative case we chose the pair $\big(\mathbf{S}(t), \mathbf{I}(t)\big)$ to represent the system state, in the replicative case it is slightly more convenient to represent the system state by the explicit triple $\big(\mathbf{S}(t), \mathbf{I}(t), \mathbf{R}(t)\big)$. The new system dynamics are now governed by

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

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

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

with corresponding initial conditions and state constraints given by

$$\mathbf{S}(0) = \mathbf{S}^0 \succ \mathbf{0}, \quad \mathbf{I}(0) = \mathbf{I}^0 \succeq \mathbf{0}, \quad \mathbf{R}(0) = \mathbf{R}^0 \succ \mathbf{0}, \tag{6}$$

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

We make the same assumptions on controls and connectivity as in (§II-A).

C. Motivation of the models and instantiations

In the introduction (§I) we described the motivations for the stratified epidemic models (2) and (5) through different examples. In this section, we provide more detail and specify the different types in each context.

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, i.e., each type corresponds to a region. IP eavesdropping techniques (using software such as `AirJack`, `Ethereal`, `FakeAP`, `Kismet`, etc.) allow 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. Accordingly, malware (and also patch) propagation rates β_{ij} (respectively $\bar{\beta}_{ij}$) 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: A network that relies on a homogeneous software/protocol is vulnerable to an attack that exploits a common weakness (e.g., a buffer overflow vulnerability). Accordingly, inspired by the natural observation that *the chances of survival are improved by heterogeneity*, increasing the network's heterogeneity without sacrificing interoperability has been proposed as a defense mechanism [10]. In practice, mobile nodes use different operating systems and communication protocols, e.g., `Symbian`, `Android`, `IOS`, `RIM`, `webOS`, etc. 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 response should take such inhomogeneities into account in order to optimally utilize network resources, since the rate of patching can also be dissimilar among 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* (ASs) with each type representing an AS [13]. A worm either scans IP addresses uniformly randomly 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 differing clique sizes: Malware that specifically spreads in social networks has been recorded in the past few years [23]. Examples include `Samy` in MySpace in 2005 and `Koobface` in MySpace and Facebook in 2008. `Koobface`, for instance, spread by delivering (contaminated) messages to the "friends" of an infective user. MMS based malware such as `commwarrior` can also utilize the contact list of an infective host to access new handsets. In such cases the social network graph can be approximated by a collection of friendship *cliques*.[6] Users of the same clique can be regarded as the same type with the rate of contact within cliques and across cliques differing depending on the relative sizes of the cliques.

5) Cloud-computing—heterogeneity through cluster sizes: In cluster (or grid, or volunteer) computing [25], 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.

6) Clustered epidemics in technology adoption, belief-formation over social media and health care: We now elaborate on the application of our clustered epidemics model in these diverse set of contexts. First consider a rivalry between two technologies or companies for adoption in a given population, e.g., Android and iPhone, or cable and satellite television. Individuals who are yet to choose either may be considered as susceptibles and those who have chosen one or the other technology would be classified as either infective or recovered depending upon their choice. Dispatchers constitute the promoters of a given technology (the one whose subscribers are denoted as recovered). Awareness about the technology and subsequent subscription to either may spread through social contact between infectives and susceptibles (infection propagation in our terminology), and dispatchers and the rest (patching in our terminology). Immunization of a susceptible corresponds to her adoption of the corresponding technology, while healing of an infective corresponds to an alteration in her original choice. The stratifications may be based on location or social cliques, and the control $\mathbf{u}$ would represent promotion efforts, which would be judiciously selected by the proponent of the corresponding technology. Patching may either be replicative or non-replicative depending on whether the newly subscribed users are enticed to attract more subscribers by referral rewards. Similarly, clustered epidemics may be used to model belief management over social media, where infective and recovered nodes represent individuals who have conflicting persuasions and susceptibles represent those who are yet to subscribe to either doctrine. Last, but not least, the susceptible-infective-recovered classification and immunization/healing/infection have natural connotations in the context of a biological epidemic. Here, the dispatchers correspond to health-workers who administer vaccines and/or hospitalization and the stratification is based on location. Note that in this context, patching can only be non-replicative.

D. Key observations

A natural but important observation is that if the initial conditions are non-negative, then the system dynamics (2) and (5) yield states satisfying the positivity and normalisation constraints (4) and (7), respectively. The proof is elementary but technical and not needed elsewhere in the paper; we relegate it accordingly to the appendix.

Lemma 1. *Suppose* $\mathbf{u}(\cdot)$ *is any admissible control and consider the evolution of the dynamical system* (2) *with initial conditions* (3) *or the dynamical system* (5) *with initial conditions* (6). *Then the state constraints* (4) *and* (7) *hold in the non-replicative and replicative cases, respectively. In both cases,*

[6] A clique is a maximal complete subgraph of a graph [24, p. 112].

moreover, the strict positivity conditions $\mathbf{S}(t) \succ \mathbf{0}$, $\mathbf{I}(t) \succ \mathbf{0}$, *and* $\mathbf{R}(t) \succ \mathbf{0}$ *are satisfied for all* $t > 0$.

E. The optimality objective

The network seeks to minimize the overall cost of infection and the resource overhead of patching in a given operation time window $[0, T]$. At any given time t, the system incurs costs at a rate $f\big(\mathbf{I}(t)\big)$ due to the malicious activities of the malware. For instance, the malware may use infected hosts to eavesdrop, analyze, or misroute, alter or destroy the traffic that the hosts generate or relay. We suppose that $f(\mathbf{0}) = 0$ and make the natural assumption that the scalar function $f(\mathbf{I})$ is increasing and differentiable with respect to each I_i. The simplest natural candidate for $f(\mathbf{I})$ is of the form $\sum_{i=1}^{M} f_i(I_i)$; in this setting each f_i is a non-decreasing function of its argument representing the cost of infection of type i which is in turn determined by the criticality of the data of that type and the nature of its functions.[7] The network also benefits at the rate of $L\big(\mathbf{R}(t)\big)$, i.e., incurs a cost at the rate of $-L\big(\mathbf{R}(t)\big)$, owing to the removal of uncertainty about the state of the nodes brought about by patching. We suppose that the scalar function $L(\mathbf{R})$ is non-decreasing and differentiable with respect to each R_i.

In addition to the cost of infection, each dispatcher burdens the network with a cost by consuming either available bandwidth or the energy reserves of the nodes (e.g., in communication and computing networks) or money (e.g., in technology adaptation, propaganda, health-care) to disseminate the patches. Suppose each dispatcher of type i incurs cost at a rate of $h_i\big(u_i(t)\big)$. We suppose that the overhead of extra resource (bandwidth or energy or money) consumption at time t is then given by a sum of the form $\sum_{i=1}^{M} R_i^0 h_i(u_i)$. The scalar functions $h_i(\cdot)$ represent how much resource is consumed for transmission of the patch to nodes of each type and how significant this extra taxation of resources is for each type. Naturally enough, we assume these functions are non-decreasing and satisfy $h_i(0) = 0$ and $h_i(\gamma) > 0$ for $\gamma > 0$. We assume, additionally, that each h_i is twice differentiable. Following the same lines of reasoning, the corresponding expression for the cost of replicative patching is of the form $\sum_{i=1}^{M} R_i h_i(u_i)$.

With the arguments stated above for motivation, the aggregate cost for non-replicative patching is given by an expression of the form

$$J_{non-rep} = \int_0^T \left(f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i=1}^{M} R_i^0 h_i(u_i) \right) dt, \quad (8)$$

while for replicative patching, the aggregate cost is of the form

$$J_{rep} = \int_0^T \left(f(\mathbf{I}) - L(\mathbf{R}) + \sum_{i=1}^{M} R_i h_i(u_i) \right) dt. \quad (9)$$

Problem Statement: The system seeks to minimize the aggregate cost **(A)** in (8) for non-replicative patching (2, 3) and **(B)** in (9) for replicative patching (5, 6) by an appropriate selection of an optimal admissible control $\mathbf{u}(t)$.

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

In this cost setting, it is clear that by a scaling of $\bar{\beta}_{ji}^{(N)}$ and $h_j(\cdot)$ we may simplify our admissibility conditions on the optimal controls by supposing that $u_{j,\max} = 1$. We do so without further ado.

It is worth noting that any control in the non-replicative case can be emulated in the replicative setting: this is because the fraction of the dispatchers in the replicative setting is non-decreasing, hence at any time instance, a feasible $u_{rep}(t)$ can be selected such that $R(t)u_{rep}(t)$ is equal to $R^0 u_{non-rep}(t)$. This means that the minimum possible cost of replicative patching is always less than the minimum cost of its non-replicative counterpart. Our numerical results will show that this improvement is substantial. However, replicative patches increase the risk of patch contamination: the security of a smaller set of known dispatchers is easier to manage than that of a growing set, many of whose identity may be ambiguous. Hence, in a nutshell, if there is a dependable mechanism of authenticating the patches, replicative patching is the preferred method, otherwise one needs to evaluate the trade-off between the risk of compromised patches and the efficiency of the patching.

III. Optimal Non-Replicative Patching

A. Numerical framework for computing the optimal controls

The main challenge in computing the optimal state and control functions $\big((\mathbf{S}, \mathbf{I}), \mathbf{u}\big)$ is that while the differential equations (2) can be solved once the optimal controls $\mathbf{u}(\cdot)$ are known, an exhaustive search for an optimal control is infeasible as there are an uncountably infinite number of control functions. *Pontryagin's Maximum Principle (PMP)* provides an elegant technique for solving this seemingly intractable problem (c.f. [26], for example). PMP derives from the classical calculus of variations. It bears a close analogy to the primal-dual non-linear optimization framework and indeed generalizes it in the sense that 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* $\mathcal{H}$ will play that of the objective function of the relaxed optimization in the primal-dual framework. Referring to the integrand in (8) as $\xi_{non-rep}$ and the RHS of (2a) and (2b) as ν_i and μ_i, we define the Hamiltonian to be

$$\mathcal{H} = \mathcal{H}(\mathbf{u}) := \xi_{non-rep} + \sum_{i=1}^{M} (\lambda_i^S \nu_i + \lambda_i^I \mu_i) \quad (10)$$

where, the *adjoint* (or *costate*) functions λ_i^S and λ_i^I are continuous functions that for each $i = 1 \dots M$ and at each point of continuity of $\mathbf{u}(\cdot)$, satisfy

$$\dot{\lambda}_i^S = -\frac{\partial \mathcal{H}}{\partial S_i}, \quad \dot{\lambda}_i^I = -\frac{\partial \mathcal{H}}{\partial I_i} \quad (11)$$

along with the final (i.e., transversality) 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 must satisfy

$$\mathbf{u} \in \arg\min_{\mathbf{v}} \mathcal{H}(\mathbf{v}) \tag{13}$$

where the minimization is over the space of admissible controls (i.e., $H(u) = \min_v H(v)$).

In economic terms, the adjoint functions represent a shadow price (or imputed value); they measure the marginal worth of an increment in the state at time t when moving along an optimal trajectory. Intuitively, in these terms, λ_i^I ought to be positive as it represents the additional cost that the system incurs per unit time with an increase in the fraction of infective nodes. Furthermore, as an increase in the fraction of the infective nodes has worse long-term implications for the system than an increase in the fraction of the susceptibles, we anticipate that $\lambda_i^I - \lambda_i^S > 0$. The following result confirms this intuition. It is of value in its own right but as its utility for our purposes is in the proof of our main theorem of the following section, we will defer its proof (to §III-C) to avoid breaking up the flow of the narrative at this point.

Lemma 2. *The positivity constraints $\lambda_i^I(t) > 0$ and $\lambda_i^I(t) - \lambda_i^S(t) > 0$ hold for all $i = 1, \ldots, M$ and all $t \in [0, T)$.*

The abstract maximum principle takes on a very simple form in our context. Using the expression for $\xi_{non-rep}$ from (8) and the expressions for ν_i and μ_i from (2), trite manipulations show that the minimization (13) may be expressed in the much simpler scalar formulation

$$u_i(t) \in \arg\min_{0 \le x \le 1} \psi_i(x, t) \qquad (1 \le i \le M) \tag{14}$$

where

$$\psi_i(x, t) := R_i^0(h_i(x) - \phi_i(t)x) \tag{15}$$

and

$$\phi_i := \sum_{j=1}^{M} \bar{\beta}_{ij} \lambda_j^S S_j + \sum_{j=1}^{M} \bar{\beta}_{ij} \pi_{ij} \lambda_j^I I_j. \tag{16}$$

Equation (14) allows us to characterise u_i as a function of the state and adjoint functions at each time instant. Plugging these controls into (2) and (11), we obtain a system of (nonlinear) differential equations that involves only the state and adjoint functions (and not the control $\mathbf{u}(\cdot)$), and where the initial values of the states (3) and the final values of the adjoint functions (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 (14).

We conclude this section by proving an important property of $\phi_i(\cdot)$, which we will use in subsequent sections.

Lemma 3. *For each i, $\phi_i(t)$ is a decreasing function of t.*

Proof: We examine the derivative of $\phi_i(t)$; we need expressions for the derivatives of the adjoint functions towards that end. From (10), (11), at any t at which $\mathbf{u}$ is continuous, we have:

$$\dot{\lambda}_i^S = -\frac{\partial L(\mathbf{R})}{\partial R_i} - (\lambda_i^I - \lambda_i^S) \sum_{j=1}^{M} \beta_{ji} I_j + \lambda_i^S \sum_{j=1}^{M} \bar{\beta}_{ji} R_j^0 u_j$$

$$\dot{\lambda}_i^I = -\frac{\partial L(\mathbf{R})}{\partial R_i} - \frac{\partial f(\mathbf{I})}{\partial I_i} - \sum_{j=1}^{M} \left((\lambda_j^I - \lambda_j^S) \beta_{ij} S_j \right) + \lambda_i^I \sum_{j=1}^{M} \pi_{ji} \bar{\beta}_{ji} R_j^0 u_j \tag{17}$$

Using (16), (17) and some reassembly of terms, at any t at which $\mathbf{u}$ is continuous,

$$\dot{\phi}_i(t) = -\sum_{j=1}^{M} \bar{\beta}_{ij} \Bigg[S_j \frac{\partial L(\mathbf{R})}{\partial R_j} + \pi_{ij} I_j \left(\frac{\partial L(\mathbf{R})}{\partial R_j} + \frac{\partial f(\mathbf{I})}{\partial I_j} \right) + \sum_{k=1}^{M} (1 + \pi_{ij}) \lambda_j^I \beta_{kj} S_j I_k + \sum_{k=1}^{M} \pi_{ij} I_j S_k \beta_{jk} (\lambda_k^I - \lambda_k^S) \Bigg].$$

The assumptions on $L(\cdot)$ and $f(\cdot)$ (together with Lemma 1) show that the first two terms inside the square brackets on the right are always non-negative. Lemmas 1 and 2 (together with our assumptions on π_{ij}, β_{ij} and $\bar{\beta}_{ij}$) show that the penultimate term is positive and the final term on the right is non-negative. It follows that $\dot{\phi}_i(t) < 0$ for every $t \in (0, T)$ at which $\mathbf{u}(t)$ is continuous. As $\phi_i(t)$ is a continuous function of time and its derivative is negative except at a finite number of points (where $\mathbf{u}$ may be discontinuous), it follows indeed that, as advertised, $\phi_i(t)$ is a decreasing function of time. ∎

B. Structure of optimal non-replicative patching

We are now ready to identify the structures of the optimal controls $(u_1(t), \ldots, u_M(t))$, which constitute an important result of this paper.

Theorem 1. *If $h_i(\cdot)$ is concave for type i, then the optimal control for type i has the following structure: $u_i(t) = 1$ for $0 < t < t_i$, and $u_i(t) = 0$ $t_i < t \le T$, where $t_i \in [0, T)$. If $h_i(\cdot)$ is strictly convex then the optimal control for type i, $u_i(t)$ is continuous and has the following structure: $u_i(t) = 1$ for $0 < t < t_i^1$, $u_i(t) = 0$ for $t_i^2 < t \le T$, and $u_i(t)$ strictly decreases in the interval $[t_i^1, t_i^2]$, where $0 \le t_i^1 < t_i^2 \le T$.*

Intuitively, at the onset of the epidemic, a large fraction of nodes are susceptible to the malware ("potential victims"). Bandwidth and power resources should hence be used maximally in the beginning (in all types), rendering as many infective and susceptible nodes robust against the malware as possible. In particular, there is no gain in deferring patching since the efficacy of healing infective nodes is less than that of immunizing susceptible nodes (recall that $\pi_{ij} \le 1$). While the non-increasing nature of the optimal control is intuitive, what is less apparent is the nature of the decrease, which we establish in this theorem. For concave $h_i(\cdot)$, nodes are patched at the maximum possible rate until a time instant when patching stops abruptly, while for strictly convex $h_i(\cdot)$,

this decrease is continuous. It is instructive to note that the structure of the optimal action taken by a type only depends on its own patching cost and not on that of its neighbours. This is somewhat counter-intuitive as the controls for one type affect the spread of infection and recovery of other types. The timing of the decrease in each type differs and depends on the location of the initial infection as well as the topology of the network, communication rates, etc.

Proof: For strictly concave $h_i(\cdot)$, (14) requires the minimization of the (strictly concave) difference between a strictly concave function of a scalar variable x and a linear function of x at all time instants; hence the minimum can only occur at the end-points of the interval over which x can vary. Thus all that needs to be done is to compare the values of $\psi_i(x,t)$ for the following two candidates: $x=0$ and $x=1$. Note that $\psi_i(0,t)=0$ at all time instants and $\psi_i(1,t)$ is a function of time t. Let

$$\gamma_i(t) := \psi_i(1,t) = R_i^0 h_i(1) - R_i^0 \phi_i(t). \tag{18}$$

Then the optimal u_i satisfies the following condition:

$$u_i(t) = \begin{cases} 1 & \gamma_i(t) < 0 \\ 0 & \gamma_i(t) > 0 \end{cases} \tag{19}$$

From the transversality conditions in (12) and the definition of $\phi_i(t)$ in (16), for all i, it follows that $\phi_i(T)=0$. From the definition of the cost term, $h_i(1)>0$, hence, since $R_i^0>0$, therefore $\gamma_i(T)>0$. Thus the structure of the optimal control for the strictly concave case will follow if we can show that $\gamma_i(t)$ is an increasing function of time t, as that implies that it can be zero at most at one point t_i, with $\gamma_i(t)<0$ for $t<t_i$ and $\gamma_i(t)>0$ for $t>t_i$. Hence the theorem will follow from (19). From (18), it follows that γ_i will be an increasing function of time if ϕ_i is a decreasing function of time, which we showed in Lemma 3.

If $h_i(\cdot)$ is linear (i.e., $h_i(x)=K_i x$, K_i is positive,[8]) $\psi_i(x,t) = R_i^0 x(K_i - \phi_i(t))$ and from (14), the condition for an optimal u_i is:

$$u_i(t) = \begin{cases} 1 & \phi_i(t) > K_i \\ 0 & \phi_i(t) < K_i \end{cases} \tag{20}$$

But from (12), $\phi_i(T) = 0 < K_i$ and as by Lemma 3, $\phi_i(t)$ is decreasing, it follows that $\phi_i(t)$ will be equal to K_i at most at one time instant $t=t_i$, with $\phi_i(t)>0$ for $t<t_i$ and $\phi_i(t)<0$ for $t>t_i$. This, along with (20), concludes the proof of the theorem for the concave case.

We now consider the case where $h_i(\cdot)$ is strictly convex. In this case, the minimization in (14) may also be attained at an interior point of $[0,1]$ (besides 0 and 1) at which the partial derivative of the right hand side with respect to x is zero. Hence,

$$u_i(t) = \begin{cases} 1 & 1 < \eta(t) \\ \eta(t) & 0 < \eta(t) \le 1 \\ 0 & \eta(t) \le 0. \end{cases} \tag{21}$$

[8]since $h_i(\gamma) > 0$ for $\gamma > 0$.

where $\eta(t)$ is such that

$$\left.\frac{dh_i(x)}{dx}\right|_{(x=\eta(t))} = \phi_i(t) \tag{22}$$

Note that $\phi_i(t)$ is a continuous function due to the continuity of the states and adjoint functions. We showed that it is also a decreasing function of time (Lemma 3). Since $h_i(\cdot)$ is double differentiable, its first derivative is continuous, and since it is strictly convex, its derivative is a strictly increasing function of its argument. Therefore, $\eta(t)$ must be a continuous and decreasing function of time. ∎

C. Proof of Lemma 2

Proof: From (17), at time T we have:

$$\begin{aligned} \lambda_i^I|_{t=T} &= (\lambda_i^I - \lambda_i^S)|_{t=T} = 0, \\ \dot{\lambda}_i^I|_{t\uparrow T} &= -\frac{\partial L(\mathbf{R})}{\partial R_i}(T) - \frac{\partial f(\mathbf{I})}{\partial I_i}(T) < 0 \\ (\dot{\lambda}_i^I - \dot{\lambda}_i^S)|_{t\uparrow T} &= -\frac{\partial f(\mathbf{I})}{\partial I_i}(T) < 0 \end{aligned}$$

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, going backward in time from $t=T$, (at least) one of the inequalities is first violated at $t=t^*$ for i^*, i.e., for all i, $\lambda_i^I(t)>0$ and $(\lambda_i^I(t) - \lambda_i^S(t)) > 0$ for all $t > t^*$ and either (A) $(\lambda_{i^*}^I(t^*) - \lambda_{i^*}^S(t^*)) = 0$ or (B) $\lambda_{i^*}^I(t^*) = 0$ for some $i = i^*$. Note that from continuity of the adjoint functions $\lambda_i^I(t^*) \ge 0$ and $(\lambda_i^I(t^*) - \lambda_i^S(t^*)) \ge 0$ for all i.

We investigate case (A) first. We have:[9] [10]

$$\begin{aligned} (\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)(t^{*+}) = &-\frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \sum_{j=1}^{M}[(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j] \\ &- \lambda_{i^*}^I \left(\sum_{j=1}^{M} \bar{\beta}_{ji^*}(1-\pi_{ji^*})R_j^0 u_j \right) \end{aligned}$$

First of all, $-\partial f(\mathbf{I})/\partial I_{i^*} < 0$. Also, the terms $-\sum_{j=1}^{M}[(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j]$ and $-(1-\pi_{ji^*})\lambda_{i^*}^I u_{i^*}$ are non-positive, according to the definition of t^*. Hence, $(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)(t^{*+}) < 0$, which is in contradiction with Property 1 of real-valued functions, stated below and proved in the appendix.

Property 1. *Let $g(t)$ be a continuous and piecewise differentiable function of t. If $g(t_0) = L$ and $g(t) < L$ ($g(t) > L$) for all $t \in (t_0, t_1]$. Then $\dot{g}(t_0^+) \le 0$ (respectively $\dot{g}(t_0^+) \ge 0$).*

On the other hand, for case (B) we have:[10]

$$\dot{\lambda}_{i^*}^I(t^{*+}) = -\frac{\partial L(\mathbf{R})}{\partial R_{i^*}} - \frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \sum_{j=1}^{M}[(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j]$$

which is again negative according to the definition of t^* and since $f(\cdot)$ and $L(\cdot)$ are increasing and non-decreasing in their arguments respectively. This contradicts Property 1, and the claim follows. ∎

[9]We define $g(t_0^+)$ and $g(t_0^-)$ to be respectively equal to $\lim_{t\downarrow t_0} g(t)$ and $\lim_{t\uparrow t_0} g(t)$ for a general function $g(.)$.

[10]The RHS of the equation is evaluated at $t=t^*$ due to continuity.

IV. Optimal Replicative Patching

A. Numerical framework for computing the optimal controls

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. Referring to the integrand of (9) as ξ_{rep}, and the RHS of equations (5-a,b,c) as ν_i, μ_i and ρ_i respectively, the new Hamiltonian becomes:

$$\mathcal{H} = \mathcal{H}(\mathbf{u}) := \xi_{rep} + \sum_{i=1}^{M}[(\lambda_i^S \nu_i + \lambda_i^I \mu_i + \lambda_i^R \rho_i) \tag{23}$$

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

$$\dot{\lambda}_i^S = -\frac{\partial \mathcal{H}}{\partial S_i},\ \dot{\lambda}_i^I = -\frac{\partial \mathcal{H}}{\partial I_i},\ \dot{\lambda}_i^R = -\frac{\partial \mathcal{H}}{\partial R_i} \tag{24}$$

with the final constraints:

$$\lambda_i^S(T) = \lambda_i^I(T) = \lambda_i^R(T) = 0. \tag{25}$$

According to PMP, any optimal controller must satisfy:

$$\mathbf{u} \in \arg\min_{\mathbf{v}} \mathcal{H}(\mathbf{v}) \tag{26}$$

where the minimization is over the set of admissible controls as before.

Using the expressions for ξ_{rep} from (9) and the expressions for ν_i, μ_i and ρ_i from (5), it can be shown that the vector minimization (26) can be expressed in scalar minimization form

$$u_i(t) \in \arg\min_{0 \le x \le 1} \psi_i(x,t) \qquad (1 \le i \le M) \tag{27}$$

where

$$\psi_i(x,t) := R_i(t)(h_i(x) - \phi_i(t)x) \tag{28}$$

and

$$\phi_i := \sum_{j=1}^{M} \bar{\beta}_{ij}(\lambda_j^S - \lambda_j^R)S_j + \sum_{j=1}^{M} \pi_{ij}\bar{\beta}_{ij}(\lambda_j^I - \lambda_j^R)I_j. \tag{29}$$

Equation (27) characterizes the optimal control u_i as a function of the state and adjoint functions at each instant. Plugging the optimal u_i into the state and adjoint function equations (respectively (5) and (24)) will again leave us with a system of (non-linear) differential equations that involves only the state and adjoint functions (and not the control $\mathbf{u}(\cdot)$), the initial values of the states (6) and the final values of the adjoint functions (25). Similar to the non-replicative case, the optimal controls may now be obtained (via (27)) by solving the above system of differential equations.

We conclude this subsection by stating and proving some important properties of the function $\psi_i(\cdot,\cdot)$, the adjoint functions (Lemma 4 below) and $\phi_i(\cdot)$ (Lemma 5 subsequently), which we use later.

First, from (25), $\psi_i(0,t) = 0$, hence (27) results in

$$\psi_i(u_i,t) \le 0. \tag{30}$$

Lemma 4. *For all $t \in [0,T)$ and for all i, we have $(\lambda_i^I - \lambda_i^S) > 0$ and $(\lambda_i^I - \lambda_i^R) > 0$.*

Using our previous intuitive analogy, Lemma 4 implies that infective nodes are always worse for the evolution of the system than either susceptible or healed nodes, and thus the marginal price of infectives is greater than that of susceptible and healed nodes at all times before T. As before, we defer the proof of this lemma (to §IV-C) to avoid breaking up the flow of the narrative. We now state and prove Lemma 5.

Lemma 5. *For each i, $\phi_i(t)$ is a decreasing function of t, and $\dot{\phi}_i(t^+) < 0$ and $\dot{\phi}_i(t^-) < 0$ for all t.*

Proof: $\phi_i(t)$ is continuous everywhere (due to the continuity of the states and adjoint functions) and differentiable whenever $\mathbf{u}(\cdot)$ is continuous. At any t at which $\mathbf{u}(\cdot)$ is continuous, we have:

$$\begin{aligned}\dot{\phi}_i(t) =& \sum_{j=1}^{M} \bar{\beta}_{ij}[(\dot{\lambda}_j^S - \dot{\lambda}_j^R)S_j + (\lambda_j^S - \lambda_j^R)\dot{S}_j \\ &+ \pi_{ij}(\dot{\lambda}_j^I - \dot{\lambda}_j^R)I_j + \pi_{ij}(\lambda_j^I - \lambda_j^R)\dot{I}_j]\end{aligned}$$

From (23) and the adjoint equations (24), at points of continuity of the control, we have:

$$\begin{aligned}\dot{\lambda}_i^S =& -(\lambda_i^I - \lambda_i^S)\sum_{j=1}^{M}\beta_{ji}I_j - (\lambda_i^R - \lambda_i^S)\sum_{j=1}^{M}\bar{\beta}_{ji}R_j u_j \\ \dot{\lambda}_i^I =& -\frac{\partial f(\mathbf{I})}{\partial I_i} - \sum_{j=1}^{M}(\lambda_j^I - \lambda_j^S)\beta_{ij}S_j - (\lambda_i^R - \lambda_i^I)\sum_{j=1}^{M}\pi_{ji}\bar{\beta}_{ji}R_j u_j \\ \dot{\lambda}_i^R =& \frac{\partial L(\mathbf{R})}{\partial R_i} - h_i(u_i) + u_i\sum_{j=1}^{M}\bar{\beta}_{ij}(\lambda_j^S - \lambda_j^R)S_j \\ &+ u_i\sum_{j=1}^{M}\pi_{ij}\bar{\beta}_{ij}(\lambda_j^I - \lambda_j^R)I_j = \frac{\partial L(\mathbf{R})}{\partial R_i} - \frac{\psi_i(u_i,t)}{R_i}\end{aligned} \tag{31}$$

Therefore, after some regrouping and cancellation of terms, at any t, we have

$$\begin{aligned}-\dot{\phi}_i(t^+) =& \sum_{j=1}^{M}\bar{\beta}_{ij}[(1-\pi_{ij})\sum_{k=1}^{M}(\lambda_j^I - \lambda_j^R)\beta_{kj}I_k S_j \\ &+\pi_{ij}\frac{\partial f(\mathbf{I})}{\partial I_j}I_j + (S_j + \pi_{ij}I_j)\frac{\partial L(\mathbf{R})}{\partial R_j} \\ &+\pi_{ij}I_j\sum_{k=1}^{M}(\lambda_k^I - \lambda_k^S)\beta_{jk}S_k - (S_j + \pi_{ij}I_j)\frac{\psi_j(u_j,t)}{R_j}].\end{aligned}$$

Now, since $0 \le \pi_{ij} \le 1$, the assumptions on $\bar{\beta}_{ij}$, β_{ki} and β_{il}, Lemma 1, and Lemma 4 all together imply that the sum of the first and fourth term of the RHS will be positive. The second and third terms will be non-negative due to the definitions of $f(\cdot)$ and $L(\cdot)$. The last term will also be non-negative due to (30). So $\dot{\phi}_i(t^+) < 0$ for all t. The proof for $\dot{\phi}_i(t^-) < 0$ is exactly as above. In a very similar fashion, it can be proved that $\dot{\phi}_i(t) < 0$ at all points of continuity of $\mathbf{u}(\cdot)$, which coupled with the continuity of $\phi_i(t)$ shows that it is a decreasing function of time. ■

B. Structure of optimal replicative dispatch

We now prove that the structures indicated in Theorem 1 hold

for replicative dispatch.

Proof: First consider an i such that $h_i(\cdot)$ is strictly concave. Note that hence $\psi_i(x,t)$ is a strictly concave function of x. Thus, the minimum can only occur at extremal values of x, i.e., $x=0$ and $x=1$. Now $\psi_i(0,t)=0$ at all times t, so to obtain the structure of the control, we need to examine $\psi_i(1,t)$. Let $\gamma_i(t) := \psi_i(1,t) = R_i(t)(h_i(1)-\phi_i(t))$ be a function of time t. From (27), the optimal u_i satisfies:

$$u_i(t) = \begin{cases} 1 & \gamma_i(t) < 0, \\ 0 & \gamma_i(t) > 0. \end{cases} \tag{32}$$

We now show that $\gamma_i(t) > 0$ for an interval $(t_i, T]$ for some t_i, and $\gamma_i(t) < 0$ for $[0, t_i)$ if $t_i > 0$. From (25) and (29), $\gamma_i(T) = h_i(1)R_i(T) > 0$. Since $\gamma_i(t)$ is a continuous function of its variable (due to the continuity of the states and adjoint functions), it will be positive for a non-zero interval leading up to $t = T$. If $\gamma_i(t) > 0$ for all $t \in [0,T]$, the theorem follows. Otherwise, from continuity, there must exist a $t = t_i$ such that $\gamma_i(t_i) = 0$. We show that for $t > t_i$, $\gamma_i(t) > 0$, from which it follows that $\gamma_i(t) < 0$ for $t < t_i$ (by a contradiction argument). The theorem will then follow from (32).

Towards establishing the above claim, we show that $\dot{\gamma}_i(t^+) > 0$ and $\dot{\gamma}_i(t^-) > 0$ for any t such that $\gamma_i(t) = 0$. If we show this, there will exist an interval $(t_i, t_i+\epsilon)$ over which $\gamma_i(t) > 0$. If $t_i + \epsilon \geq T$, then the claim holds, otherwise there must exist a $t = t_i^{'} > t_i$ such that $\gamma_i(t_i^{'}) = 0$ and $\gamma_i(t) \neq 0$ for $t_i < t < t_i^{'}$ (from the continuity of $\gamma_i(t)$). Note that $\dot{\gamma}_i(t_i^{'-}) > 0$, contradicting a property of real-valued functions stated below and proved in the appendix, thus establishing the claim.

Property 2. *If $g(x)$ is a continuous and piecewise differentiable function over $[a,b]$ such that $g(a) = g(b)$ while $g(x) \neq g(a)$ for all x in (a,b), $\frac{dg}{dx}(a^+)$ and $\frac{dg}{dx}(b^-)$ cannot be positive simultaneously.*

We now show that $\dot{\gamma}_i(t^+) > 0$ and $\dot{\gamma}_i(t^-) > 0$ for any t such that $\gamma_i(t) = 0$. Due to the continuity of $\gamma_i(t)$ and the states, and the finite number of points of discontinuity of the controls, for any t we have:

$$\dot{\gamma}_i(t^+) = (\dot{R}_i(t^+)\frac{\gamma_i(t)}{R_i(t)} - R_i(t)\dot{\phi}_i(t^+)) \tag{33}$$

and

$$\dot{\gamma}_i(t^-) = (\dot{R}_i(t^-)\frac{\gamma_i(t)}{R_i(t)} - R_i(t)\dot{\phi}_i(t^-)). \tag{34}$$

If $\gamma_i(t) = 0$, then $\dot{\gamma}_i(t^+) = -R_i(t)\dot{\phi}_i(t^+)$ and $\dot{\gamma}_i(t^-) = -R_i(t)\dot{\phi}_i(t^-)$, and the theorem will follow from Lemmas 5 and 1.

The proofs for linear and strictly convex $h_i(\cdot)$'s are virtually identical to the corresponding parts of the proof of the non-replicative case and are omitted for brevity; the only difference is that in the linear case we need to replace R_i^0 with $R_i(t)$. ■

C. Proof of Lemma 4

Proof: First, from (29) and (25), we have $\phi_i(T) = 0$, which, combined with (28) results in $\psi_i(x,T) = R_i(T)h_i(x)$. Since Lemma 1 dictates that $R_i(T) > 0$, (27) results in $u_i(T) = 0$, since all other values of x would produce a positive $\psi_i(x,T)$. Therefore, $h_i(u_i(T)) = 0$, a result which we will use shortly.

The rest of the proof has a similar structure to that of Lemma 2. $(\lambda_i^I - \lambda_i^S)|_{t=T} = 0$ and $(\dot{\lambda}_i^I - \dot{\lambda}_i^S)|_{t\uparrow T} = -\partial f(\mathbf{I})/\partial I_i < 0$, for all i. Also, for all i, $(\lambda_i^I - \lambda_i^R)|_{t=T} = 0$ and $(\dot{\lambda}_i^I - \dot{\lambda}_i^R)|_{t\uparrow T} = -\partial f(\mathbf{I})/\partial I_i - \partial L(\mathbf{R})/\partial R_i + h_i(u_i(T)) < 0$, since $h_i(u_i(T)) = 0$.

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^*, i.e., for all i, $(\lambda_i^I(t) - \lambda_i^S(t)) > 0$ and $(\lambda_i^I(t) - \lambda_i^R(t)) > 0$ for all $t > t^*$, and either (A) $(\lambda_{i^*}^I(t^*) - \lambda_{i^*}^S(t^*)) = 0$ or (B) $(\lambda_{i^*}^I(t^*) - \lambda_{i^*}^R(t^*)) = 0$ for some i^*. Note that from continuity of the adjoint functions, $(\lambda_i^I(t^*) - \lambda_i^S(t^*)) \geq 0$ and $(\lambda_i^I(t^*) - \lambda_i^R(t^*)) \geq 0$ for all i.

Case (A): . Here, we have:[12]

$$\begin{aligned}(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)(t^{*+}) = &-\frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \sum_{j=1}^{M}(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j \\ &- (\lambda_{i^*}^I - \lambda_{i^*}^R)\sum_{j=1}^{M}\bar{\beta}_{ji^*}(1-\pi_{ji^*})R_j u_j\end{aligned}$$

First of all, $-\partial f(\mathbf{I})/\partial I_{i^*} < 0$. Also, the terms $-\sum_{j=1}^{M}[(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j]$ and $-(\lambda_{i^*}^I - \lambda_{i^*}^R)\sum_{j=1}^{M}\bar{\beta}_{ji^*}(1-\pi_{ji^*})R_j u_j$ are non-positive, according to the definition of t^*. Hence, $(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^S)(t^{*+}) < 0$, which contradicts Property 1, therefore case (A) does not arise.

Case (B):. In this case, we have:[12]

$$\begin{aligned}(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^R)(t^{*+}) = &-\frac{\partial f(\mathbf{I})}{\partial I_{i^*}} - \frac{\partial L(\mathbf{R})}{\partial R_{i^*}} \\ &- \sum_{j=1}^{M}(\lambda_j^I - \lambda_j^S)\beta_{i^*j}S_j + \frac{\psi_{i^*}(u_{i^*},t)}{R_{i^*}}\end{aligned}$$

The terms $-\partial f(\mathbf{I})/\partial I_{i^*}$ and $-\partial L(\mathbf{R})/\partial R_{i^*}$ are negative and non-positive respectively. The term $-(\lambda_{i^*}^I - \lambda_{i^*}^S)\sum_{j=1}^{M}\beta_{ji^*}S_j$ is non-positive, according to the definition of t^*. The last term will be non-negative due to (30). This shows that $(\dot{\lambda}_{i^*}^I - \dot{\lambda}_{i^*}^R)(t^{*+}) < 0$, which contradicts Property 1, and hence case (B) does not arise either. This completes the proof. ■

V. An Alternative Cost Functional

Recall that in our objective function, the cost of non-replicative patching was defined as $\sum_{i=1}^{M} R_i^0 h_i(u_i)$ (respectively $\sum_{i=1}^{M} R_i h_i(u_i)$ for the replicative case), which corresponds to a scenario in which the dispatchers are charged for every instant they are immunizing/healing (distributing the patch), irrespective of the number of nodes they are delivering patches to. This represents a *broadcast* cost model where each

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

[12] The RHS of the equation is evaluated at $t = t^*$ due to continuity.

transmission can reach all nodes of the neighbouring types. In an alternative *unicast* scenario, different transmissions may be required to deliver the patches to different nodes. This model is particularly useful if the dispatchers may only transmit to the nodes that have not yet received the patch.[13] Hence, the cost of patching in this case can in general be represented by : $\sum_{i=1}^{M}\sum_{j=1}^{M} R_i^0\bar{\beta}_{ij}(S_j + I_j)p(u_i)$ (respectively $\sum_{i=1}^{M}\sum_{j=1}^{M} R_i\bar{\beta}_{ij}(S_j + I_j)p(u_i)$ for the replicative case), where $p(.)$ is an increasing function. More generally, the patching cost can be represented as a sum of the previously seen cost (§II-E) and this term.

For non-replicative patching, if all $h_i(\cdot)$ and $p(\cdot)$ are concave, then Theorem 1 will hold if for all pairs (i,j), $\pi_{ij} = \pi_j$ (i.e., healing efficacy only depends on the type of an infected node, not that of the immunizer). The analysis will change in the following ways: A term of $R_i^0 p(u_i)\sum_{j=1}^{M}\bar{\beta}_{ij}(S_j + I_j)$ is added to (15), and subsequently to (18) (with $u_i = 1$ in the latter case). Also, (17) is modified by the subtraction of $\sum_{j=1}^{M}\bar{\beta}_{ji}R_j^0 p(u_j)$ from the RHS of both equations. This leaves $\dot{\lambda}_i^I - \dot{\lambda}_i^S$ untouched, while subtracting a positive amount from $\dot{\lambda}_i^I$, meaning that Lemma 2 still holds. As $\phi_i(t)$ was untouched, this means that Lemma 3 will also go through. Thus the RHS of $\dot{\gamma}_i$ is only modified by the subtraction of $\sum_{j,k=1}^{M}\bar{\beta}_{ij}(S_j + \pi_j I_j)\bar{\beta}_{kj}R_k^0\left(p(u_k) - u_k p(1)\right)$ which is a positive term, as for any continuous, increasing, concave function $p(\cdot)$ such that $p(0) = 0$, we have $ap(b) \geq bp(a)$ if $a \geq b \geq 0$.[14]. This yields: $(p(u_k) - u_k p(1) \geq 0)$. Therefore the conclusion of Theorem 1 holds. Furthermore, using similar techniques, it may be shown that Theorem 1 also holds for strictly convex $h_i(\cdot)$ provided $p(\cdot)$ is linear.

For the replicative case, if $p(\cdot)$ is linear in its argument $(p(x) = Cx)$ and again $\pi_{ij} = \pi_j$ for all (i,j), Theorem 1 will hold. The modifications of the integrand and ψ_i are as above. The adjoint equations (31) are modified by the subtraction of $\sum_{j=1}^{M} C\bar{\beta}_{ji}R_j u_j$ from $\dot{\lambda}_i^I$ and $\dot{\lambda}_i^I$, and the subtraction of $Cu_i\sum_{j=1}^{M}\bar{\beta}_{ij}(S_j + I_j)$ from $\dot{\lambda}_i^R$. Due to the simultaneous change in ψ_i, however, we still have $\lambda_i^R = \partial L(\mathbf{R})/\partial R_i - \psi_i(u_i,t)/R_i$. Therefore, Lemma 4 still holds, as $\dot{\lambda}_i^I - \dot{\lambda}_i^S$ is unchanged, and a positive amount is subtracted from $\dot{\lambda}_i^I - \dot{\lambda}_i^R$. We absorb $\sum_{j=1}^{M} C\bar{\beta}_{ij}(S_j + I_j)$ into $\phi_i(t)$, where all the $p(\cdot)$ terms in $\dot{\phi}_i$ will cancel out, leaving the rest of the analysis, including for Lemma 5, to be the same. The theorem follows.

VI. Numerical Investigations

In this section, we numerically investigate the optimal control policies for a range of malware and network parameters.[15] Recalling the notion of topologies presented in §II-A (in the paragraph before (2)), we consider three topologies: *linear*, *star* and *complete*, as was illustrated in fig. 1. At time zero, we assume that only one of the regions is infected, i.e., $I_i^0 > 0$ only for $i = 1$. Also, $R_i^0 = 0.2$, $\beta_{ii} = \beta = 0.223$ for all i.[16] 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. It should be noted that $\beta_{ij} * T$ denotes the average number of contacts between nodes of regions i and j within the time period, and thus β and T are dependent variables. For simplicity, we use equal values for $\beta_{ji}, \beta_{ij}, \bar{\beta}_{ij}$, and $\bar{\beta}_{ji}$ for every i and j (i.e., $\beta_{ji} = \beta_{ij} = \bar{\beta}_{ij} = \bar{\beta}_{ji}$). 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): $\int_0^T \left(K_I\sum_{i=1}^{M} I_i(t) + K_u\sum_{i=1}^{M} R_i^0 u_i(t)(S(t) + I(t))\right)dt$ (discussed in §V), where in both cases we select $T = 35$, $K_I = 1$, $K_u = 0.5$ unless stated otherwise. (To aid understanding, we assume $\pi_{ij} = \pi$ to be equal for all regions) For replicative patching, R_i^0 in both cost models is replaced with $R_i(t)$.

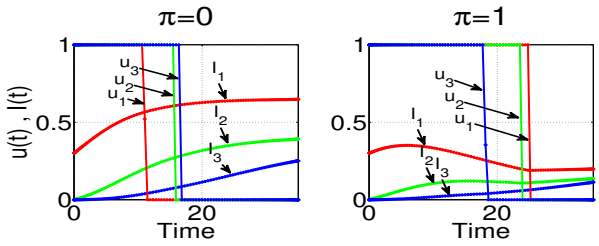

Fig. 2. Illustration of optimal patching policies and the corresponding levels of infection in each of the three regions for a simple linear topology. Note how the infection that initially only exists in region 1 spreads in region 1 and then to region 2, and from there to region 3.

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 (due to healing).

We now investigate the effect of topology on the optimal patching policy. We study the *drop-off* times (the time thresholds at which the bang-bang optimal 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 $\pi = 0$, a middle region is patched for the longest time, whereas for $\pi = 1$,

[13]This can be achieved by keeping a common database of nodes that have successfully received the patch, or by implementing 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 with a priori certainty which nodes are infective, and hence it cannot differentiate between susceptible and infective nodes. Consequently, even when $\pi_{ij} = 0$, i.e., the system manager knows that the patch cannot remove the infection and can only immunize the susceptible, still the best it may be able to do is to forward the message to any node that has not previously received it.

[14]Due to the increasing nature of a finite $\frac{p(x)}{x}$ when $p(0) = 0$ in these conditions

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

[16]This specific value of β is chosen to match the average inter-meeting times from the numerical experiment reported in [27].

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 susceptibles. In regions closer to the origin of the infection, the fraction of 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 for 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 susceptible and infective nodes 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 given intuition.

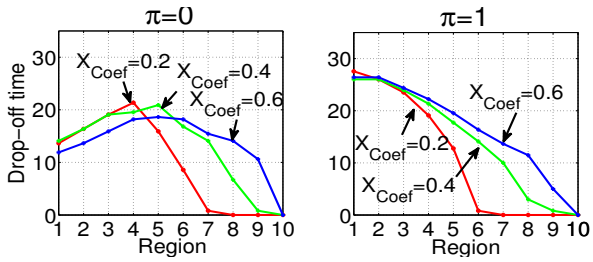

Fig. 3. The drop-off times of different regions in a linear topology with $M = 10$ regions. We consider a type-A cost function and non-replicative patching with $X_{Coef} = 0.2, 0.4, 0.6$.

To complete our examination of contact rates, we fix X_{Coef} and vary β, the base value of the contact rates for a simple linear topology with $M = 3$ in a replicative setting with $\pi = 0$ (fig. 4). The insights given will extend to any $M \geq 3$. For low values of β, the infection will not propagate to the farthest regions in the time-frame beyond a negligible amount, and thus they will not need immunizing. As β increases, the importance of the central region(s) increases, as they can both propagate the infection and spread the immunization to multiple regions, and, therefore, negligence in immunizing them could have a significant toll on system performance. Hence, for intermediate values of β, the center region(s) are immunized for longer. As β increases even further, both the infection and the cure spread very fast, and as for $\pi = 0$, I_i is an unrecoverable state, immunization becomes pointless very quickly (due to the dwindling pool of susceptibles) and thus it is ceased sooner in all regions.

We next investigate a star configuration where the infection starts from a peripheral region (region 1). Fig. 5 reveals the following interesting phenomenon: although the central region is the only one that is connected to all the regions,

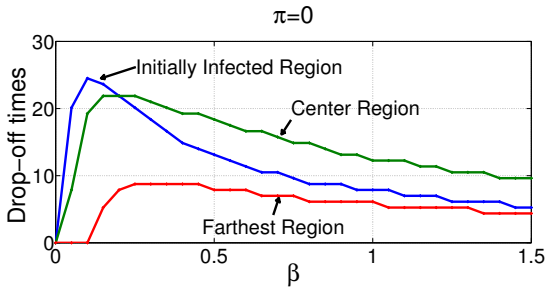

Fig. 4. The drop-off times of different regions in a linear topology with $M = 3$ regions as a function of the base-line β. We consider a type-A cost function with replicative patching without healing ($\pi = 0$) with $X_{Coef} = 0.1$. Here, $K_u = 0.04$, with all other parameters being as previously described.

for $\pi = 0$ it is patched for shorter lengths of time compared to the peripherals. 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 numbers of peripheral regions. For $\pi = 1$, on the other hand, the central region 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 infective nodes in this region.

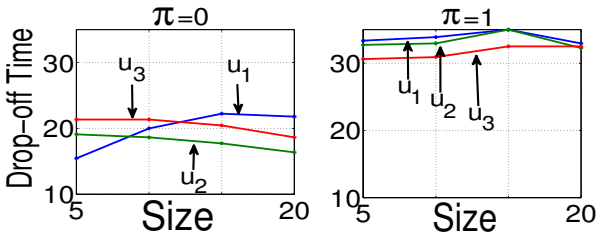

Fig. 5. Trends in the drop-off points in the star topology. The cost is type-B and the patching is non-replicative, and $I_1^0 = 0.6$.

Now we consider a star topology in which there are multiple initially infected regions, and we examine the relationship between drop-off times and the number of initially infected (non-central) regions for a fixed size star topology (here, $M = 12$) [fig. 6]. For the case where there is no healing ($\pi = 0$), it can be seen that increasing the number of initially infected regions increases the cost of immunizing all of them, while decreasing the effect of immunizing each one (since there are others that can take up the mantle of infecting the susceptible), and thus immunization in the initially infected regions stops progressively sooner as their number grows. Conversely, this results in less initially susceptible regions, making the collective decision to immunize less costly, while putting them at greater risk of infection. Thus, the initially susceptible regions are immunized for a longer period of time as the number of initially infected regions grows. The center feels the brunt of both these conflicting effects, and their effect more-or-less cancels out, leaving the drop-off time of the center region practically unchanged. It is also interesting

to note that as the number of initially infected regions grows, infection spreads much faster and thus the system becomes somewhat homogeneous, a fact borne out by the closeness of the drop-off times of the controls. For the case of $\pi = 1$ (immunization and healing) these effects are much more mild, as in this case the target population of the immunizer only changes with a change in R_i, and not both R_i and I_i as in the case of $\pi = 0$. This ensures smoother changes for both the states and the control drop-off times. as can be seen from a cursory comparison between the two cases.

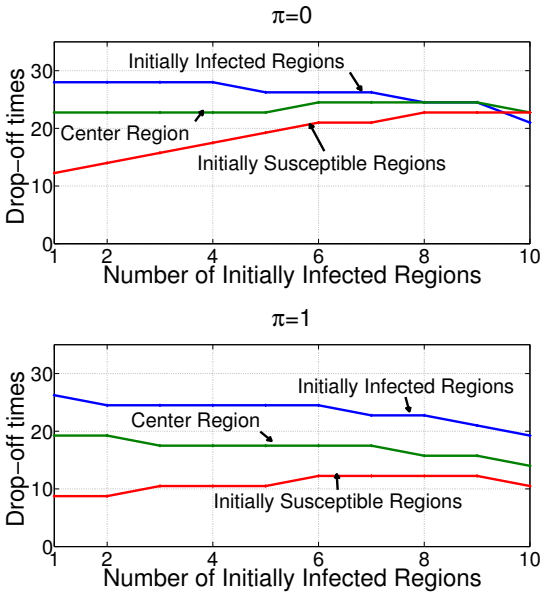

Fig. 6. Trends in the drop-off points in the star topology of total size $M = 12$ when the number of initially infected regions is increased for $\pi = 0, 1$. The cost is type-A and the patching is replicative. We have$I_1^0 = 0.3$ in all initially infected regions, with cost coefficients $K_I = 5$ and $K_u = 0.1$.

Next, in order to evaluate the efficacy of our dynamic heterogeneous patching policy, we compare our aggregate cost against those of four alternative patching policies. We label our policy *Stratified Dynamic* (S.D.). In the simplest alternative policy, all regions use identical patching intensities that 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 derived based on this model and applied across all regions to calculate the aggregate cost. We call this policy *Simplified Homogeneous* (H. in short). The simplified homogeneous policy is a special case of Spatially Static (Sp. St.) policies, where one one-jump bang-bang control is applied to all regions to find the optimum uniform control.

Fig. 7 depicts the aggregate costs of all four policies for a linear topology with $M = 2, 3, 4$ and 5 regions. The cost is type-A and patching is replicative, with $\pi = 1$. Here, $I_1^0 = 0.2$, $K_u = 0.2$ and the rest of parameters are as before. As we can clearly observe, our stratified policy achieves the least cost, outperforming the rest. Also of note is that when the number of regions is small, H. and Sp. St. perform better than S. St., all of which obviously outperform St. However, as the number of regions increases and the network becomes more spatially heterogeneous, the homogeneous approximation, and all uniform controls in general, worsen and the S. St. policy quickly overtakes them as the best approximation to the optimal. For example for $M = 5$ regions, our policy outperforms the best static policies by 40% and the homogeneous approximation by 100%, which points to the importance of the optimality of controls in such systems and shows that our results about the structure of the optimal control can result in large cost improvements. For $\pi = 0$, a similar performance gap is observed. This underscores the significance of considering heterogeneity in the controls. Specifically, as we discussed before, optimal drop-off times for this problem should vary based on the distance from the originating region, a factor that the Sp. St., H., and St. policies ignore.

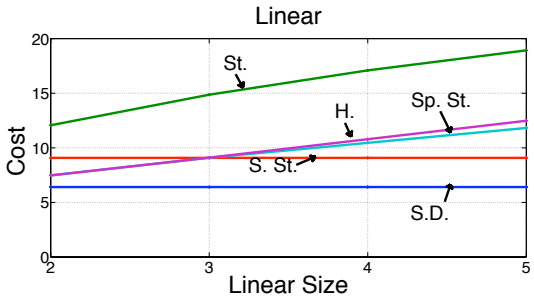

Fig. 7. Linear Topology, trends in the cost.

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. 8. 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.

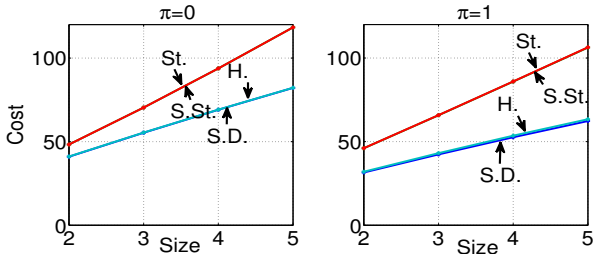

Fig. 8. Complete topology, trends in the cost.

Finally, we compare replicative to non-replicative patching (fig. 9). As previously stated, any solution to the non-replicative patching problem can be emulated in the replicative patching case, making this scenario worth investigation, even with the security vulnerabilities that the system has to contend with in these settings. Here, we see the aggregate cost of optimal replicative and non-replicative patching in a complete topology as a function of the size of the network for $M \leq 11$. Even for such modest sizes, it can be seen that replicative patching can be 60% more efficient than non-replicative patching, a very significant amount which verifies the above assertion. This is especially true for the complete topology and other edge-dense topologies, as in replicative patching, the patch can spread in ways akin to the malware.

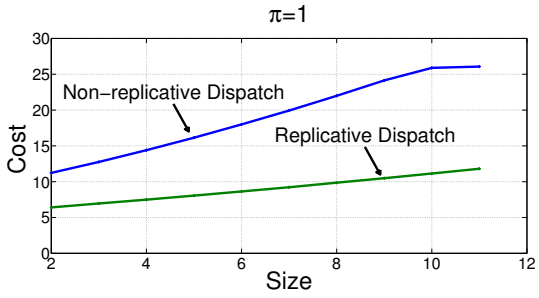

Fig. 9. Complete topology, comparison of replicative and non-replicative patching cost for $\pi = 1$, $K_I = 1$, $K_u = 0.2$ and every other parameter as outlined in the introduction of this section. It can be seen that as size grows, the replicative patch spreads farther throughout the graph, widening the performance gap to the non-replicative case.

VII. 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, pp. 1071–1076, 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, L. Su, D. Jin, and L. Zeng, "An optimal distributed malware defense system for mobile networks with heterogeneous devices," in *IEEE SECON*, pp. 314–322, 2011.

[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)*, vol. 52, pp. 73–104, 1990.

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

[16] T. Antunović, Y. Dekel, E. Mossel, and Y. Peres, "Competing first passage percolation on random regular graphs," *ArXiv e-prints*, 2011.

[17] S. Sethi and G. Thompson, *Optimal control theory: applications to management science and economics*. Springer Netherlands, 2000.

[18] H. Behncke, "Optimal control of deterministic epidemics," *Optimal control applications and methods*, vol. 21, no. 6, pp. 269–285, 2000.

[19] M. Ndeffo Mbah and C. Gilligan, "Optimization of control strategies for epidemics in heterogeneous populations with symmetric and asymmetric transmission," *Journal of theoretical biology*, vol. 262, no. 4, pp. 757–763, 2010.

[20] R. Rowthorn, R. Laxminarayan, and C. Gilligan, "Optimal control of epidemics in metapopulations," *Journal of the Royal Society Interface*, vol. 6, no. 41, pp. 1135–1144, 2009.

[21] T. Kurtz, "Solutions of ordinary differential equations as limits of pure jump markov processes," *Journal of Applied Probability*, vol. 7, pp. 49–58, 1970.

[22] 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.

[23] 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, 2009.

[24] B. Bollobás, *Modern graph theory*, vol. 184. Springer Verlag, 1998.

[25] M. Altunay, S. Leyffer, J. Linderoth, and Z. Xie, "Optimal response to attacks on the open science grid," *Computer Networks*, vol. 55, pp. 61–73, 2010.

[26] R. F. Stengel, *Optimal control and estimation*. Dover, 1994.

[27] 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.

Appendix

A. Proof of Lemma 1

We will need a few definitions in the course of this proof: let $d(i,j)$ be the *distance* from type j to type i (i.e., for all i, $d(i,i) = 0$ and for all pairs (i,j), $d(i,j) = 1+$ minimum number of types in a path from type j to type i). Now, define $d(i,U) := \min_{j \in U} d(i,j)$. Since we assumed that every type i is either in U or is connected to a type in U, $d(i,U) < M$ for all types i. We also define $\theta^{(n)} := d^n\theta/dt^n$.

Proof: Non-replicative patching: If we show that $S_i(t)$ and $I_i(t)$ are positive for all i and $t \in (0,T]$, then $\dot{S}_i + \dot{I}_i \leq 0$,

and given that $S_i^0 + I_i^0 = 1 - R_i^0$, the inequality $\mathbf{S}+\mathbf{I} \leq 1-\mathbf{R_0}$ will be true throughout the interval, so we focus on the first two inequalities. We prove this argument in two steps: we first show that the inequalities hold strictly over an interval $(0, \epsilon)$ for some $\epsilon > 0$, and then we show that $\epsilon \geq T$.

First Step: Since $\mathbf{S}^0 > 0$, $\mathbf{S} > 0$ for an open interval $(0, \epsilon)$ for some $\epsilon > 0$ (due to the continuity of the states). A similar conclusion for $\mathbf{I}$ is not direct because I_i^0 may equal 0, and therefore it requires validation, which we provide in the following lemma.

Lemma 6. *There exists an $\epsilon > 0$ such that $\mathbf{I} > 0$ in the interval $(0, \epsilon)$.*

Our differentiability assumption on each u_i for a closed neighborhood of 0, the finite number of points of discontinuity of $\mathbf{u}$, and the continuity of the state functions mean that the first M derivatives of I_i and S_i exist and are bounded in $[0, \delta)$ for some $\delta > 0$. We first show that the first derivative of $I_i(t)$ that is not equal to zero at $t = 0$ will be positive (Lemma 7). We use this fact to conclude that our claim will hold by appealing to a property of continuous functions (Property 3), completing our proof for Lemma 6.

Lemma 7. *For all i and for all integers $r \geq 0$, if $d(i, U) > r$, then $I_i^{(r)}(0) = 0$. Else if $d(i, U) = r$, then $I_i^{(r)}(0) > 0$.*

Proof: By induction on r.

Base case: $r = 0$. If $d(i, U) = 0$, this means that the type is initially infected, and thus $I_i^{(0)}(0) = I_i^0 > 0$. On the other hand, if $d(i, U) > 0$, then $i \notin U$ and $I_i^{(0)}(0) = I_i^0 = 0$ by definition. Therefore the base case holds.

Induction step: We assume that the statement holds for $r = 0 \ldots, k$ and consider the case of $r = k+1$. Therefore, we need to examine types i such that $d(i, U) \geq k + 1$. As $I_i^{(k+1)}(0) = dI_i^{(k)}(0)/dt$, we focus on equation (2b). The k-th derivative of the first term on the right involves terms like $I_j^{(s)}(0) S_i^{(k-s)}(0)$ where j is a neighbor of i, while the contribution of the second term involves terms like $I_i^{(s)}(0) u_j^{(k-s)}(0)$ for $s = 0, \ldots, k$. Since $d(i, U) \geq k+1$, we have $I_i^{(s)}(0) = 0$ for all $s = 0, \ldots, k$ (by the induction hypothesis), and as the M derivatives of u_i are bounded at $t = 0$, the contribution of the second term of equation (2b) to the $(k + 1)$-th derivative of I_i at $t = 0$ is always zero. So our $(k + 1)$-th derivative is governed by the first term at $t = 0$.

Now if $d(i, U) > k + 1$, then $d(j, U) > k$ for all nodes j that neighbour i, and thus from the induction hypothesis, $I_i^{(0)}(0) = I_i^{(1)}(0) = \ldots = I_i^{(k)}(0) = 0$, and there is no contribution from the first term either. Thus $I_i^{(k+1)}(0) = 0$.

Finally, if $d(i, U) = k + 1$, i will have some neighbours j for whom $d(j, U) = k$, and it might have neighbours for whom $d(j, U) \geq k + 1$. Due to the induction hypothesis, all the $I_j^{(s)}(0)$ terms resulting from the second kind of neighbours will be zero. For the types for which $d(j, U) = k$, the only non-zero derivative (at $t = 0$) in the first k will be $I_j^{(k)}(0)$, which will be positive by the induction hypothesis. This term will be multiplied by S_i^0, which is positive, and β_{ji}, which is positive due to i and j being neighbours, and so it will contribute a positive amount to $I_i^{(k+1)}(0)$, making it positive. This completes the proof of this lemma. ∎

Since $d(i, U) < M$, there exists an l such that $l < M$ for which $I_i^{(l)}(0) > 0$ and $I_i^{(h)}(0) = 0$ for all $h < l$ (from Lemma 7). Lemma 6 now follows from the following property of real-valued functions.

Property 3. *Consider a function $g(\cdot)$ for which the first $k+1$ derivatives exist in a closed neighborhood of 0. If $g^{(r)}(0) = 0$ for $r = 1, \ldots, k$ and $g^{(k+1)}(0) > 0$, there exists an $\epsilon > 0$ such that $g(t) > 0$ for $t \in (0, \epsilon)$.*

Proof: By the definition of a one-sided limit, $g^{(k+1)}(0) > 0$ means that there exists an $\epsilon > 0$ such that $g^{(k+1)}(t) > 0$ for all $t \in (0, \epsilon)$. Taking the integral of this function from $(0, t)$ for $t \in (0, \epsilon)$ results in $g^{(k)}(t) - g^{(k)}(0) = g^{(k)}(t) > 0$. Now the same argument can be recursively applied for all non-negative $r < k$, leading to the fact that $g(t) > 0$ for $t \in (0, \epsilon)$. ∎

So we have proved that there exists an $\epsilon > 0$ such that all S_i and I_i's are positive for $t \in (0, \epsilon)$, and *a fortiori* $S_i(\epsilon/2) > 0$ and $I_i(\epsilon/2) > 0$.

Second Step: Now we show that $\epsilon \geq T$. If the conditions $S_i(t) > 0$ and $I_i(t) > 0$ do not hold for all i and all $t \in (0, T]$, there must be a time in $t \in [\epsilon, T)$ at which these strict inequalities becomes an equality (due to the continuity of S_i and I_i) and the inequalities are strictly satisfied beforehand. Define $t^* \in [\epsilon, T)$ such that for all $t \in (0, t^*)$, $S_i(t) > 0$, $I_i(t) > 0$, and either $S_{i^*}(t^*) = 0$ or $I_{i^*}(t^*) = 0$ for some $i = i^*$, with the other inequalities holding (not necessarily strictly) at $t = t^*$. But $\dot{S}_{i^*} = S_{i^*}(-\sum_{j=1}^M \beta_{ji^*} I_j - \sum_{j=1}^M \bar{\beta}_{ji^*} R_j^0 u_j)$, and the coefficient of S_{i^*} in the RHS is lower bounded by some $-K < 0$ in $[0, t^*]$ (because continuous functions are bounded on a closed and bounded interval). Thus $S_{i^*}(t^*) \geq S_{i^*}(0) e^{-Kt^*} > 0$. Next, from (2b), for $t \in [\epsilon/2, t^*]$, we have $\dot{I}_{i^*}(t) \geq -I_{i^*}(t) \sum_{j=1}^M \pi_{ji^*} \bar{\beta}_{ji^*} R_j^0 u_j(t) > -K I_{i^*}(t)$ for some $K > 0$, and thus $I_{i^*}(t^*) \geq I_i(\epsilon/2) e^{-K(t^* - \epsilon/2)} > 0$ (from Property 3). Therefore we have arrived at a contradiction, and the lemma follows.

Replicative patching: The proof is very similar to the above, with the following differences: R_i^0 is universally replaced with R_i, $S_i(t) + I_i(t) + R_i(t) = 1$ is satisfied at the original state and holds for all t as $\dot{S}_i(t) + \dot{I}_i(t) + \dot{R}_i(t) = 0$ for all t, and in addition to $S_i(t) > 0$ and $I_i(t) > 0$, we prove $R_i(t) > 0$ (using $\dot{R}_i(t) \geq 0$). ∎

B. Proof of Property 1

Proof: By contradiction. We assume $g(t_0) = L$ and $\dot{g}(t_0^+) > 0$. Then there exists a $\delta \in (0, t_1 - t_0)$ such that $\dot{g}(t) > 0$ for $t \in (t_0, t_0 + \delta)$. But $g(t_0 + \delta) = g(t_0) + \int_{t_0}^{t_0+\delta} \dot{g}(x)\, \mathrm{d}x > L$, which is a contradiction, and thus the property holds. The proof for the $g(t) > L$ is exactly as above, with the signs interchanged. ∎

C. Proof of Property 2

Proof: We denote the value of $g(a)$ and $g(b)$ by L. If $\frac{dg}{dx}(a^+) > 0$, there exists an $\epsilon > 0$ such that $g(x) > L$ for all $x \in (a, a + \epsilon)$ and if $\frac{dg}{dx}(b^-) > 0$ then there exists an $\alpha > 0$

such that $g(x) < L$ for all $x \in (b-\alpha, b)$. Now $g(a+\frac{\epsilon}{2}) > L$ and $g(b-\frac{\alpha}{2}) < L$; thus, due to the continuity of $g(t)$, the intermediate value theorem states that there must exist a $y \in (a+\frac{\epsilon}{2}, b-\frac{\alpha}{2})$ such that $g(y) = L$, which is in contradiction with the assumption that $g(x) \neq L$ for $x \in (a, b)$. The property follows. ■