Learning to Suppress SIS Epidemics in Networks

Argyris Kalogeratos **Kevin Scaman** **Nicolas Vayatis**
CMLA, ENS Cachan, CNRS, Université Paris-Saclay, 94235 Cachan, France
{kalogeratos, scaman, vayatis}@cmla.ens-cachan.fr

Abstract

In this paper, we introduce the class of *priority planning* strategies for suppressing SIS epidemics taking place in a network. This is performed via dynamic allocation of treatment resources with limited efficiency to the infected nodes, according to a precomputed *priority-order*. Using recent theoretical results that highlight the role of the *maxcut* of a node ordering and the extinction time of an epidemic, we propose a simple and efficient strategy called *MaxCut Minimization* (MCM) that outperforms competing state-of-the-art strategies in simulated epidemic scenarios.

1 Introduction

Diffusion processes arise in systems involving agents whose behaviors depend on their close environments. Many recent articles have considered the problem of influence maximization [1, 2], which is to maximize the spread of a diffusion process. Although critical in real-life situations, the dynamic suppression of undesired information or social diffusion processes has received relatively less attention. In public health, epidemiologists study scenarios in which the spread of a virus needs to be controlled. Moreover, various analogues emerge in modern information networks in which a diffusion can be engineered to be viral and may cause positive, but sometimes also hugely negative, social and economic effects [3, 4, 5].

The control of diffusion processes has been studied in various fields in the past, including epidemiology and computer networks resilience. The respective literature can generally be divided in three complementary lines of research, the third of which is the line where our work lays:

a) *Static vaccination strategies*. Most of the epidemic literature focuses on static control actions such as permanently removing a set of edges or nodes of the network [6, 7, 8, 9, 10]. In this case, the available budget is considered fixed, and the effect of a control action *permanent* [11, 12].

b) *Budget optimization*. Complementary to resource allocation, the determination of the budget size to be spent at each time step, which aims to fulfill cost and efficiency constraints, is critical for the resulting strategy. Several such studies assume that the network administrator is capable of storing resources for later use [13, 14, 15]. Also, a usual simplifying assumption is the *uniform mixing*, i.e. the infected nodes are uniformly scattered in the network. Therefore, these studies do not address the problem of how exactly to allocate the resources on the nodes of the network, but rather estimate the budget size that can cause a desired macroscopic result.

c) *Dynamic resource allocation*. A few studies consider dynamic strategies for allocating resources for dealing with epidemics. One of the most well-known such strategy is *contact-tracing* [16] that consists in healing the neighbors of infected nodes. In practice, this approach was shown inefficient in containing an epidemic, especially when it is beyond a very initial state. In the definition of efficient strategies, among many graph features, the role of the *cutwidth* has already been underlined in [17, 18] and [19] independently. In this paper, we further explore the power of this particular concept to analyze and control epidemics in a large variety of contexts from the viewpoint of the diffusion and the resource allocation process.

Dynamic resource allocation (DRA) strategies [19, 20] aim to suppress undesired Susceptible-Infected-Susceptible (SIS) diffusion processes. To this end, they consider a *budget of resources* and each of them represents a *targeted* and *temporal* action that can affect the behavior of an individual node of the network. The network administrator is allowed to change the distribution of a set of *treatment resources* as the diffusion takes place. However, reacting in an online fashion to fast spreading phenomena is difficult to achieve.

Our major contribution is the introduction of a particular class of DRA strategies called *priority planning* that rely on a *priority-order* precomputed offline. The strategy heals the infected nodes following the priority-order and, this way, manages to gradually suppress the diffusion. Using recent theoretical results [19] that highlight the connection between the *maxcut* of such a priority-order of the nodes and the extinction time of an SIS epidemic from a network, we propose and implement an efficient strategy called *MaxCut Minimization* (MCM).

The most related method to MCM is the *CURE policy* [17] which was developed independently to our work. Notable differences are that: i) MCM algorithm is more intuitive and applies to a more general setting with important additional aspects such as node self-recovery and the allocation of multiple treatment resources, ii) in contrast with the theoretical work in [17], we also provide a test bed for experimental assessment of healing strategies on benchmark real networks and others artificially generated, and iii) we experimentally show that MCM is more efficient than CURE and other possible approaches to create a priority-order based on the related literature.

2 Setup and main definitions

2.1 Epidemic and control model

We consider the standard Susceptible-Infected-Susceptible (SIS) diffusion model [21] which is widely used and allows for theoretical analysis. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a network of $N = |\mathcal{V}|$ nodes with adjacency matrix A, where $A_{ij} = 1$ if $i \neq j$ and edge $(i, j) \in \mathcal{E}$, else $A_{ij} = 0$. Let also $\mathbf{0}$ and $\mathbf{1}$ be vectors of size N that are, respectively, all-zeros and all-ones, and $\mathbb{1}\{\cdot\}$ be the indicator function. A *state vector* $X(t) \in \mathbb{R}^N$ represents the state of the diffusion process via the nodes' infection states: for each node $i \in \{1, ..., N\}$, $X_i(t) = 1$ if node i is infected at time t, else $X_i(t) = 0$. We assume no *incubation period*, therefore, a node becomes contagious upon infection. Let the control action be represented as a *resource allocation vector* $\rho(t)$, where $\rho_i(t) > 0$ iff a resource is being given to node i at time t. In such a case, we say that node i is being *healed* by the resource. Following the formalism of [22], we model an epidemic under a control action as a *stochastic process* with the following transition rates:

$$\begin{aligned} &X_i(t) : 0 \to 1 \text{ at rate } \beta \textstyle\sum_j A_{ji} X_j(t); \\ &X_i(t) : 1 \to 0 \text{ at rate } \delta + \rho_i(t), \end{aligned} \tag{1}$$

where β is the transmission rate over an edge, and δ is the self-recovery rate, both being essential characteristics of the infection.

We consider *dynamic resource allocation* (DRA) strategies that only depend on the current and past network states $X(t' \leq t)$. In other words, a DRA strategy $\rho(t)$ is a stochastic process that is adapted to the natural filtration associated to $X(t)$. In addition to this constraint, we also limit the amount of resources in the network by r and on each node by ρ. Hence, for any $t > 0$, $\|\rho(t)\|_1 = \sum_{i \in V} \rho_i(t) \leq r$ and $\|\rho(t)\|_\infty = \max_{i \in V} \rho_i(t) \leq \rho$.

Extinction time. In order to account for the criticality of the contagion, we will consider the standard quality measure of *expected extinction time*:

$$\tau_x = \min\{t \in \mathbb{R}_+ | X(0) = x,\ X(t) = \mathbf{0}\}\ .$$

The extinction time depends on the chosen DRA strategy, and the main quality measure that we will consider in this work is its expectation $\mathbb{E}[\tau_x]$. Note that this expectation is never infinite, and may present sub-critical (respectively super-critical) behavior [22], in the sense that $\mathbb{E}[\tau_x]$ may be upper bounded by a polynomial function (respectively lower bounded by an exponential function) of the network size N. In the super-critical regime, the extremely long extinction time implies that the epidemic will reach a stable behavior, a case where we consider that the DRA strategy is not efficient enough to remove the epidemic.

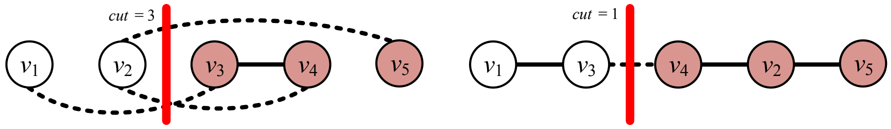

(a) Priority-order $\ell \colon \mathcal{V} \to \{1, 2, 3, 4, 5\}$ (b) Priority-order $\ell' \colon \mathcal{V} \to \{1, 3, 4, 2, 5\}$

Figure 1: Two priority-orders (from left to right) leading to different maxcuts: $\mathcal{C}^*(\ell) = 3$ for (a) and $\mathcal{C}^*(\ell') = 1$ for (b). The cut (vertical red line) separates the nodes in two sets (white are healthy, and red are infected). The second priority-order ℓ' is optimal and the network has a cutwidth $\mathcal{W} = 1$.

2.2 Maxcut and cutwidth

The strategy we propose in the Sec. 3 relies on creating a fixed linear arrangement (i.e. an ordering) of the network nodes, that we call *priority-order*. Formally, a priority-order is a bijective mapping $\ell : \mathcal{V} \to \{1, ..., N\}$ of the N nodes of the network s.t. $\ell(v)$ is the position of node v in the order. The importance of the graph-theoretic concept of *cutwidth* [23] for containing epidemics has been recently pointed out [17, 19]. Here we provide a definition adapted to our priority-order notations.

Definition 1 (**Maxcut of a priority-order**)**.** *For a network with N nodes and adjacency matrix A, and for a given priority-order ℓ, the maxcut of ℓ is defined as:*

$$\mathcal{C}^*(\ell) = \max_{c=1,\ldots,N} \sum\nolimits_{i,j} A_{ij} \mathbb{1}\{\ell(v_i) < c \leq \ell(v_j)\}. \tag{2}$$

Furthermore, the minimal value of the maxcut over all possible priority-orders is an inherent property of the network structure known as *cutwidth*, which we denote as $\mathcal{W} = \min_\ell \mathcal{C}^*(\ell)$. Fig. 1 illustrates two priority-orders for a small network and their respective maxcut (see Definition 1). The priority-order of Fig. 1(b) is better, indeed optimal, with maxcut equal to the cutwidth of the network.

3 The *MaxCut Minimization* strategy

3.1 Priority planning: a healing plan to gradually remove a contagion

The control action $\rho(t)$ of a dynamic resource allocation (DRA) strategy depends on the history of variations of the diffusion process $X(t)$. In this work, we introduce *priority planning*: a novel class of DRA strategies that involve a priority-order)see Sec. 2.2) accounting for the criticality of each node w.r.t. the overall diffusion process. This concept is formally defined below.

Definition 2 (**Priority planning**)**.** *Priority planning is a DRA strategy under limited budget r and resource threshold ρ that distributes resources to the top-q infected nodes according to a fixed priority-order ℓ of the network nodes, where q is the number of nodes such that the allocated amount of resources matches the available resource budget r. More specifically, the strategy heals the first $q(t) = \min\{\lceil \frac{r}{\rho} \rceil, \sum_i X_i(t)\}$ infected nodes according to the respective mapping ℓ, and allocates the resource budget as follows:*

$$\rho_i(t) = \begin{cases} \frac{r}{\lceil \frac{r}{\rho} \rceil} & \text{if } X_i(t) = 1 \text{ and } \ell(v_i) \leq \theta(t); \\ 0 & \text{otherwise}, \end{cases} \tag{3}$$

where $\theta(t)$ is a threshold adjusted s.t. $\sum_i \mathbb{1}_{\{\rho_i(t) > 0\}} = q(t)$.

This definition may be regarded as a description of a class of simple planning strategies for the removal of a contagion: a *healing plan*, i.e. a priority-order for healing the nodes, is determined prior to the beginning of the diffusion and is followed no matter how the diffusion process evolves. The plan proceeds from the first to the last node in the priority-order, hence aims to remove gradually the contagion from the network. In what follows, we refer interchangeably to a priority-order and its corresponding mapping ℓ.

Recent theoretical results [17, 19] highlighted the role of the *maxcut* of a priority-order for assessing the quality of a priority-planning strategy. More specifically, under technical assumptions on the structure of the network, the extinction time behaves linearly w.r.t. the number of nodes if and only if the resource budget r is lower than $\beta\mathcal{C}^*(\ell)$, i.e. the spreading rate multiplied by the maxcut of the priority-order. This means that a priority-order with lower maxcut would in essence lower the size of the resource budget needed to accomplish the worst step of the plan, and thus remove the infection from the network. Following this idea, we propose the novel DRA srategy for arbitrary networks, called *MaxCut Minimization* (MCM), that distributes resources to infected nodes in the priority-order that minimizes $\mathcal{C}^*(\ell)$. Given a network $\mathcal{G}$, we compute, prior to the diffusion process, a *priority-order* $\ell_{_{MCM}}(\mathcal{G})$ with minimum maxcut $\mathcal{C}^*(\ell)$:

$$\ell_{_{MCM}}(\mathcal{G}) = \underset{\ell}{\operatorname{argmin}}\, \mathcal{C}^*(\ell). \tag{4}$$

This is done using any available optimization algorithm for this problem. Then, during the diffusion, the strategy distributes the resource budget to the infected nodes according to the order $\ell_{_{MCM}}(\mathcal{G})$. Alg. 1 presents the pseudocode of our strategy.

3.2 Maxcut optimization and implementation

Minimum maxcut linear arrangement. Finding a priority-order of a network with minimum $\mathcal{C}^*(\ell)$ can be performed under the framework of *linear arrangement* (LA) problems [24, 25, 26], and specifically the *minimum maxcut linear arrangement* (MMLA) problem[1]. MMLA is an NP-hard combinatorial problem since all node permutations should be examined in order to find the optimal $\mathcal{C}^*(\ell)$. However, approximation heuristics do exist in literature [25, 26]. One of the major difficulties of this problem is that the cost function to optimize is extremely *flat* in the search space, i.e. slight changes in the arrangement will most probably not change $\mathcal{C}^*(\ell)$.

Relaxation of the MMLA problem. For the latter reasons, we propose to relax the MMLA problem and instead of their *maximum* to optimize the *sum of the cuts*. This is known as the *minimum linear arrangement* (MLA) and is part of the larger class of *minimum p-sum linear arrangement* problems (MpLA) [27, 28] that minimize the following functional:

$$\text{MpLA:}\;\; \phi(\mathcal{G},\ell) = \Big(\sum\nolimits_{i,j} A_{ij}\,|\ell(v_i) - \ell(v_j)|^p\Big)^{1/p}. \tag{5}$$

For $p=1$, a simple calculation shows that MLA minimizes the average cut in the linear arrangement, instead of its maximum for MMLA (see Definitions 1). MLA is easier than MMLA and more suited to gradient descent or simulated annealing methods, and it produces a smoother priority-order w.r.t. the cuts at each position of the ordering.

Practical implementation. MLA and MMLA are very challenging problems and, interestingly, most related works on their optimization conduct experiments on relatively small benchmark networks for which the optimal cost is not known. Designing a procedure that can be applied on large social networks with tens of thousands of nodes is by itself a remarkable contribution. We should also note that MCM strategy seeks for a priority-order with as low as possible maxcut, but not necessarily the optimal one. The MLA solver we developed for our simulations follows the steps below and uses a hierarchical approach to take advantage of the group structure of social and contact networks:

- **s1)** first, we identify dense clusters by applying *spectral clustering* and we order those clusters (considered as high-level nodes) using *spectral sequencing* [29],
- **s2)** then, we compute a good ordering of the nodes inside each cluster independently using spectral sequencing followed by an iterative approach which is based on random node swaps (swap heuristics inspired by [30]),
- **s3)** finally, the swap-based approach is reapplied to optimize the overall ordering.

Scalability. The scalability of the MCM strategy is highly dependent on the employed offline algorithm for finding the optimal node order. The whole process described above achieves fairly good results in reasonable time. Since spectral clustering and spectral sequencing depend on the computation of eigenvectors for the highest eigenvalues of an $N \times N$ sparse matrix with $|\mathcal{E}|$ non-zero entries, the overall complexity of the algorithm is $O(|\mathcal{V}| + |\mathcal{E}|)$ [31]. Hence, MCM is generally scalable to the size of real social and contact networks. It is worth noticing that, for networks that are close to

[1] In the related literature, a linear arrangement is formally defined exactly as the priority-order (see Sec. 2.2).

Algorithm 1 MCM strategy

▷ ***Prior to the diffusion process*:**
Compute the priority-order $\ell = \ell_{MCM}(\mathcal{G})$ by minimizing the maxcut $\mathcal{C}^*(\ell)$
Order the nodes of $\mathcal{G}$ according to ℓ, i.e. compute the node list $(v_1, ..., v_N)$ s.t. $\forall i \in \{1, ..., N\}, \ell(v_i) = i$

▷ ***During the diffusion process*:**
Input: network $\mathcal{G}$, state vector $X(t)$, resource budget r, resource threshold ρ
Output: the resource allocation vector $\rho(t)$

$q \leftarrow \lceil \frac{r}{\rho} \rceil$
if $\sum_i X_i(t) < q$ **then**
  return $\frac{r}{q} X(t)$
end if
$\rho(t) \leftarrow \mathbf{0}$ // a zero vector in $\mathbb{R}^N$
$budget \leftarrow q$
$i \leftarrow 1$
while $budget > 0$ **do**
  if $X_{v_i}(t) = 1$ **then**
    $\rho_{v_i}(t) \leftarrow \frac{r}{q}$
    $budget \leftarrow budget - 1$
  end if
  $i \leftarrow i + 1$
end while
return $\rho(t)$

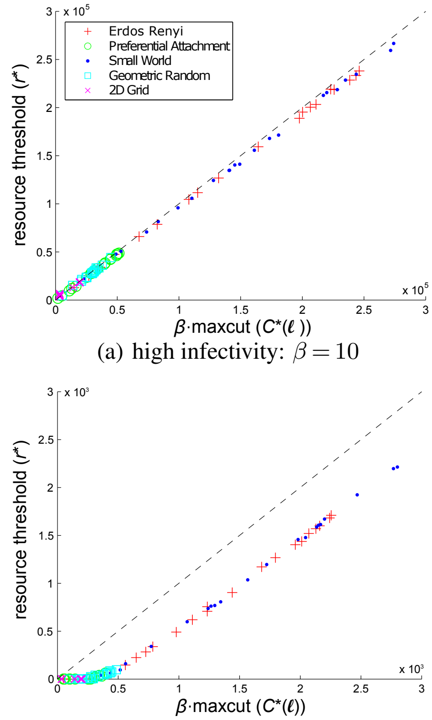

(a) high infectivity: $\beta = 10$

(b) low infectivity: $\beta = 0.1$

Figure 2: Epidemic threshold w.r.t. the maxcut for random networks of 1,000 nodes ($\delta = 1$, $\beta = 10$).

planar and can be embedded in the 2D plane without many edge intersections (e.g. contact networks), the clustering step could be skipped since the spectral sequencing method is a good initial approximation that can be refined with a node swapping process.

4 Experimental results

4.1 Setup and competitors

In the experimental study, we compare MCM against five DRA strategies, which are grouped into three types: *static vaccination* by comparing to centrality-based priority-orders, *uniform mixing* by considering a random allocation of resources, and a *state-of-the-art direct competitor*:

- *Most neighbors* (MN): gives priority to high degree nodes, hence aims to first remove the contagion from the network's core before dealing with the periphery.
- *Least neighbors* (LN): gives priority to low degree nodes and works conversely to MN.
- *Largest reduction in spectral radius* (LRSR): gives priority to nodes whose removal will lead to the maximum decrease of the *spectral radius* of the adjacency matrix of the resulting network, and is a state-of-the-art method from the vaccination literature [7].
- *Random baseline* (RAND): the resources is assigned to $\lceil \frac{r}{\rho} \rceil$ nodes at random at each time.
- *CURE policy* (CURE): this is a state-of-the-art method developed in parallel to our work [17]. CURE follows a healing plan with minimal maxcut, called *crusade*. A crusade can be considered to be a priority-order similar to that of MCM, however in [17] this was formulated as a sequence of nested bags which differ by one node each time.

A brief summary of the most significant differences between our work and the work in [17] is given in Sec. 1. An additional practical difference is that, when many reinfections occur, CURE enters a *waiting phase* in order to return to a previous step of the removal plan. This waiting phase is triggered when the number of infected nodes before the position of the front exceeds $\frac{r}{8d}$, where d is

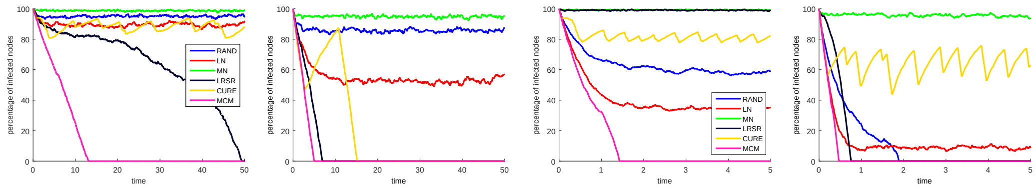

(a) low budget: $r = 100$ (b) high budget: $r = 250$ (c) low budget: $r = 3000$ (d) high budget: $r = 7000$

Figure 3: Simulation of an SIS epidemic in the GermanSpeedway network ((a) and (b), $N = 1168$) and the OpenFlights airport network ((c) and (d), $N = 2939$), under various DRA strategies ($\delta = 0$, $\rho = r$, and $\beta = 1$).

the maximum degree of the network. In practice, this threshold value can be very small, and we will see in our experiments that this may lead to substantial delay, or even failure of the healing plan.

For all the experiments, we consider a fixed number of treatments $q \in \{1, ..., 100\}$ leading to $\rho = \frac{r}{q}$.

4.2 Quality of the theoretical bound

Fig. 2 shows the relationship between the *maxcut* $\mathcal{C}^*(\ell)$ and the resource threshold r^* under a specific priority planning and budget. The resource threshold is computed by running simulations with a fixed number of treatments q, and finding the resource budget above which the strategy is able to remove the epidemic. Each plotted point is a simulation with fixed network, number of treatments, and epidemic parameters β and δ. To cover a wide range of scenarios, each of the 100 points plotted in each subfigure of Fig. 2 involves: i) the priority-order of a DRA control strategy, randomly chosen among MCM, RAND, MN, LN, and LRSR, ii) a fixed q value, set at random in $\{1, ..., 100\}$, and iii) a random network of 1,000 nodes, constructed by employing at random a generator for: *Erdös-Rényi*, *preferential attachment*, *small-world*, *geometric random* [32], and *2D regular grids* (see details on these networks in [33]).

According to the results illustrated in Fig. 2, the resource threshold is always below, but very close, to $\beta\mathcal{C}^*(\ell)$ which seems to be a very good approximation of the former. The very stable, nearly linear, behavior holds even for low infectivity where the random self-recoveries of nodes become more significant (Fig. 2(b)). Overall, this result justifies the minimization of $\mathcal{C}^*(\ell)$ as a proxy for removing a contagion with less resources.

4.3 Empirical evaluation of simulated contagion on real networks

In this section, we perform simulations on two real networks matching different use cases of DRA strategies: the GermanSpeedway network [34] for analyzing the growth of an epidemic through road network, and the OpenFlights airport network [35] for epidemics spreading through air routes. In order to compare to the CURE policy, we consider a simplified setting (matching the limitations set in [17]) for the experiments on the GermanSpeedway network and OpenFlights aiport network: we use only one treatment ($q = 1$) and let no self-healing ($\delta = 0$).

GermanSpeedway network. This is the German Autobahn network from [34]. Due to the spatial properties of road networks, the respective graph is symmetric, with a single connected component, and close to being *planar* (i.e. a graph embedding on the plane would create only very few edge intersections other than the endpoint connections). It contains 1,168 nodes and 1,243 edges, while the degree distribution is particularly flat: 101 nodes are leaves, 971 nodes have degree 2, and 96 nodes have degree 3. Finally, the maximum degree is $d = 12$.

Two scenarios of SIS epidemics with different resource budget r are shown in Fig. 3(a) and Fig. 3(b). In both of them, MN and RAND are the worst strategies. MCM is the best performing strategy, and all the other strategies are strongly affected by the resource budget. LRSR is the second best performing strategy, although the low budget causes an increase in the extinction by a factor of 7. CURE policy presents a behavior with characteristic ups-and-downs, which is due to its waiting phase. Even for a high resource budget, entering the waiting phase can happen with non negligible probability (see Fig. 3(b) for an example of such a scenario) and largely degrade the performances of the CURE policy. The capacity of the strategies to remove the epidemics is correctly predicted by their maxcut values, with CURE to be an exception since its behavior largely depends on whether or not the strategy enters its waiting phase. Specifically, the maxcut values for the strategies are: 650 ± 50 for RAND, 379 for MN and LN, 104 for LRSR, and 29 for CURE and MCM.

OpenFlights airport network. This network represents the US air traffic for the year 2010 [35]. The nodes are the US airports, plus those non-US airports connected through flights with the former. We used a symmetric, undirected, and unweighted version of this network containing $2{,}939$ nodes in a single connected component with $30{,}501$ edges. For this network, $d = 242$.

Fig. 3(c) and Fig. 3(d) presents two epidemic scenarios similar to Fig. 3(a) and Fig. 3(b). MCM is the best performing strategy and the least affected by the variation in the resource efficiency. LRSR, on the other hand, fails completely with a low resource budget, which is due to the fact that its maxcut is located at the beginning of the considered node ordering. The CURE strategy presents again an unstable behavior as an effect of its waiting phase, which also shows that the conditions under which this policy gets into the waiting phase are not rare at all and can be catastrophic. Indeed, CURE is outperformed even by RAND in these simulations. Again, the capacity of the strategies to remove the epidemics is correctly predicted by their maxcuts (except CURE): $7{,}800 \pm 100$ for RAND, $7{,}504$ for MN and LN, $6{,}223$ for LRSR, and $2{,}231$ for CURE and MCM.

In agreement with our previous analysis, our results show that: i) the uniform mixing hypothesis leads to a massive drop in efficiency, since MCM substantially outperforms the RAND strategy, ii) despite being efficient in the static vaccination problem, centrality-based priority-orders are sub-optimal for the DRA problem, iii) the $\beta\mathcal{C}^*(\ell)$ value can be used as a good criterion for assessing the quality of a priority-order, iv) the CURE policy fails to be effective in practical applications due to its waiting phase, and v) MCM outperforms all its competitors in all our experiments.

5 Conclusion

In this paper, we presented *priority planning*, a novel class of dynamic strategies for allocating resources on the nodes of an arbitrary network in order to suppress an undesired SIS diffusion process. We presented the *MaxCut Minimization* (MCM) strategy, that distributes resources to nodes following a priority-order with minimum maxcut and provided a practical MCM algorithm that includes an efficient solver for the minimum linear arrangement problem suitable for large networks. Finally, we provided a test bed for experimental assessment of healing strategies on benchmark real networks and others artificially generated.

Our experimental results verified that, for a wide range of network types, the maxcut is indeed a good approximation of the epidemic threshold under a given strategy. Moreover, the MCM strategy outperformed other competing strategies in simulations on real-world networks that are interesting for the practical application of epidemic control policies.

As part of our future work, we plan to study some theoretical aspects of the MCM strategy, as well as to investigate the robustness of the method to noise or uncertainty regarding the network structure.

Acknowledgments

This research is part of the SODATECH project funded by the French Government within the program of "*Investments for the Future – Big Data*".

References

[1] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD, pages 137–146, New York, NY, USA, 2003. ACM.

[2] Nan Du, Le Song, Manuel Gomez-Rodriguez, and Hongyuan Zha. Scalable influence estimation in continuous-time diffusion networks. In *NIPS*, pages 3147–3155, 2013.

[3] Rob Procter, Farida Vis, and Alex Voss. Reading the riots on twitter: methodological innovation for the analysis of big data. *Intern. J. Social Research Methodology*, 16(3):197–214, 2013.

[4] Sarah Vieweg, Amanda L. Hughes, Kate Starbird, and Leysia Palen. Microblogging during two natural hazards events: What twitter may contribute to situational awareness. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, pages 1079–1088, New York, NY, USA, 2010. ACM.

[5] Amanda L. Hughes, Lise A. A. St. Denis, Leysia Palen, and Kenneth M. Anderson. Online public communications by police & fire services during the 2012 hurricane sandy. In *Proc. of the 32nd ACM Conf. on Human Factors in Computing Systems*, pages 1505–1514, New York, NY, USA, 2014. ACM.

[6] Reuven Cohen, Shlomo Havlin, and Daniel Ben-Avraham. Efficient immunization strategies for computer networks and populations. *Physical Review Letters*, 91(24):247901, 2003.

[7] Hanghang Tong, B Aditya Prakash, Tina Eliassi-Rad, Michalis Faloutsos, and Christos Faloutsos. Gelling, and melting, large graphs by edge manipulation. In *Proc. of the 21st ACM Intl. Conf. on Information and Knowledge Management*, pages 245–254. ACM, 2012.

[8] Yang Wang, Deepayan Chakrabarti, Chenxi Wang, and Christos Faloutsos. Epidemic spreading in real networks: An eigenvalue viewpoint. In *Proc. of the 22nd Intl. Sym. on Reliable Distributed Systems*, pages 25–34. IEEE, 2003.

[9] Christian M Schneider, Tamara Mihaljev, Shlomo Havlin, and Hans J Herrmann. Suppressing epidemics with a limited amount of immunization units. *Physical Review E*, 84(6):061911, 2011.

[10] Victor M Preciado, Michael Zargham, Chinwendu Enyioha, Ali Jadbabaie, and George Pappas. Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks. *arXiv preprint arXiv:1303.3984*, 2013.

[11] Victor M. Preciado, Michael Zargham, Chinwendu Enyioha, Ali Jadbabaie, and George J. Pappas. Optimal resource allocation for network protection: A geometric programming approach. *CoRR*, abs/1309.6270, 2013.

[12] F. Chung, P. Horn, and A. Tsiatas. Distributing antidote using pagerank vectors. *Internet Mathematics*, 6(2):237–254, 2009.

[13] Petra Klepac, Ottar N. Bjørnstad, C. Jessica E. Metcalf, and Bryan T. Grenfell. Optimizing reactive responses to outbreaks of immunizing infections: Balancing case management and vaccination. *PLoS ONE*, 7(8):e41428, 2012.

[14] Graeme A Forster and Christopher A Gilligan. Optimizing the control of disease infestations at the landscape scale. *Proc. of the National Academy of Sciences*, 104(12):4984–4989, 2007.

[15] MHR Khouzani, Saswati Sarkar, and Eitan Altman. Optimal control of epidemic evolution. In *Proc. of the 30th Conf. of the IEEE Computer and Communications Societies*, pages 1683–1691. IEEE, 2011.

[16] Christian Borgs, Jennifer Chayes, Ayalvadi Ganesh, and Amin Saberi. How to distribute antidote to control epidemics. *Random Structures & Algorithms*, 37(2):204–222, 2010.

[17] K. Drakopoulos, A. Ozdaglar, and John N. Tsitsiklis. An efficient curing policy for epidemics on graphs. *IEEE Transactions on Network Science and Engineering*, 1(2):67–75, 2014.

[18] K. Drakopoulos, A. Ozdaglar, and John N. Tsitsiklis. An efficient curing policy for epidemics on graphs. *CoRR*, abs/1407.2241v1, 2014.

[19] K. Scaman, A. Kalogeratos, and N. Vayatis. What Makes a Good Plan? An Efficient Planning Approach to Control Diffusion Processes in Networks. *CoRR*, abs/1407.4760, 2014.

[20] Kevin Scaman, Argyris Kalogeratos, and Vayatis Nicolas. A greedy approach for dynamic control of diffusion processes in networks. In *to appear in IEEE International Conference on Tools with Artificial Intelligence*, 2015.

[21] Piet Van Mieghem, Jasmina Omic, and Robert Kooij. Virus spread in networks. *IEEE/ACM Trans. on Networking*, 17(1):1–14, 2009.

[22] Ayalvadi Ganesh, Laurent Massoulié, and Don Towsley. The effect of network topology on the spread of epidemics. In *Proc. of the 24th Conf. of the IEEE Comp. and Comm. Societies*, volume 2, pages 1455–1466. IEEE, 2005.

[23] F. R. Chung and P. D. Seymour. Graphs with small bandwidth and cutwidth. *Discrete Mathematics*, 75(1-3):113–119, 1989.

[24] Josep Díaz, Jordi Petit, and Maria Serna. A survey of graph layout problems. *ACM Computing Surveys*, 34(3):313–356, September 2002.

[25] JuanJ. Pantrigo, Rafael Mart, Abraham Duarte, and EduardoG. Pardo. Scatter search for the cutwidth minimization problem. *Annals of Operations Research*, 199(1):285–304, 2012.

[26] Eduardo G. Pardo, Nenad Mladenović, Juan J. Pantrigo, and Abraham Duarte. Variable formulation search for the cutwidth minimization problem. *Applied Soft Computing*, 13(5):2242–2252, 2013.

[27] L. H. Harper. Optimal assignments of numbers to vertices. *J. of the Society for Industrial and Applied Mathematics*, 12(1):131–135, 1964.

[28] D. Adolphson and T. C. Hu. Optimal linear ordering. *SIAM J. on Applied Mathematics*, 25(3):403–423, 1973.

[29] Martin Juvan and Bojan Mohar. Optimal linear labelings and eigenvalues of graphs. *Discrete Applied Mathematics*, 36(2):153–168, April 1992.

[30] Eduardo Rodriguez-Tello, Jin-Kao Hao, and Jose Torres-Jimenez. An effective two-stage simulated annealing algorithm for the minimum linear arrangement problem. *Computers & Operations Research*, 35(10):3331–3346, 2008.

[31] S. Arora, E. Hazan, and S. Kale. Fast algorithms for approximate semidefinite programming using the multiplicative weights update method. In *Proc. of the 46th IEEE Symp. on Foundations of Computer Science*, pages 339–348, Oct 2005.

[32] Mathew Penrose. *Random geometric graphs*, volume 5. Oxford University Press, 2003.

[33] Mark Newman. *Networks: An Introduction*. Oxford University Press, New York, NY, USA, 2010.

[34] Marcus Kaiser and Claus C. Hilgetag. Spatial growth of real-world networks. *Phys. Rev. E*, 69:036103, 2004.

[35] Openflights.org.