Relaxation and Matrix Randomized Rounding for the Maximum Spectral Subgraph Problem

Cristina Bazgan[1], Paul Beaujean[1,2(✉)], and Éric Gourdin[2]

[1] Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016 Paris, France
{cristina.bazgan,paul.beaujean}@lamsade.dauphine.fr
[2] Orange Labs, Châtillon, France
{paul.beaujean,eric.gourdin}@orange.com

Abstract. Modifying the topology of a network to mitigate the spread of an epidemic with epidemiological constant λ amounts to the NP-hard problem of finding a partial subgraph with maximum number of edges and spectral radius bounded above by λ. A software-defined network (SDN) capable of real-time topology reconfiguration can then use an algorithm for finding such subgraph to quickly remove spreading malware threats without deploying specific security countermeasures.

In this paper, we propose a novel randomized approximation algorithm based on the relaxation and rounding framework that achieves a $O(\log n)$ approximation in the case of finding a subgraph with spectral radius bounded by $\lambda \in (\log n, \lambda_1(G))$ where $\lambda_1(G)$ is the spectral radius of the input graph and n its number of nodes. We combine this algorithm with a maximum matching algorithm to obtain a $O(\log^2 n)$ approximation algorithm for all values of λ. We also describe how the mathematical programming formulation we give has several advantages over previous approaches which attempted at finding a subgraph with minimum spectral radius given an edge removal budget.

Keywords: Approximation algorithm · Relaxation and rounding
Semidefinite programming · Spectral graph theory · Random graphs

1 Introduction

In recent years, a sequence of results [2,4,25,28] have established a relationship between the convergence of Markovian models representing an epidemic spreading over a network and the spectral characteristics of the underlying graph. The generalization of these theorems by Prakash et al. [16] states that in the case of a graph G and an epidemic model with epidemiological characteristic λ, fast convergence of the Markovian model to its absorbing state is guaranteed if the spectral radius of the graph $\lambda_1(G) < \lambda$. This has led the mathematical epidemiology community to look for algorithms that modify the topology of a network to ensure that a given epidemic converges rapidly to extinction.

© Springer Nature Switzerland AG 2018
D. Kim et al. (Eds.): COCOA 2018, LNCS 11346, pp. 108–122, 2018.
https://doi.org/10.1007/978-3-030-04651-4_8

At the same time, the software-defined networking (SDN) paradigm has transformed network administration by allowing real-time statistics [20] and topology reconfiguration [27]. This new paradigm has deep consequences for the management of network security as it is now possible for a SDN controller to automatically detect malware spreading over its network via machine learning [10] and react to such threat by deploying adequate security countermeasures. In this work we are following epidemiological practice and propose to use topology modification as a disease-agnostic countermeasure to the spread of malware in networks.

We are looking to preserve as much as possible the existing network topology by keeping the largest number of edges in the graph while guaranteeing that a given epidemic of epidemiological characteristic λ would rapidly disappear. For this purpose, we introduce the maximum spectral subgraph problem (MSSP) defined formally as follows. Denoting by $\lambda_1(G)$ the spectral radius of G, i.e. the largest eigenvalue of its adjacency matrix A, we have:

Definition 1. Maximum spectral subgraph problem (MSSP)

Input: $G = (V, E)$ *an undirected graph and* $1 \leq \lambda < \lambda_1(G)$.
Output: $H = (V, E')$ *with* $E' \subseteq E$ *such that* $|E'|$ *is maximum and* $\lambda_1(H) \leq \lambda$.

1.1 Related Work

Spectral graph theory has often been a decisive tool in the design and analysis of algorithms. However, to the best of our knowledge, surprisingly few computational problems have been defined in terms of finding graphs with appropriate spectrum. The mathematical epidemiology community has proposed and analyzed several problems related to the spectrum of the adjacency matrix [19,26] while systems and control researchers have considered optimization problems related to the spectrum of the Laplacian matrix [5]. In a separate effort, the theoretical computer science community has focused on problems related to the design of expander graphs and graphs with high algebraic connectivity i.e. the second smallest eigenvalue of the Laplacian matrix [7,11]. In this line of research, all problems are NP-hard and the algorithms proposed in the literature are often simple to state. We contrast this with the fact that their analysis can be involved and yet, to the best of our knowledge, only amount to conditional approximation guarantees. Throughout this paper we qualify approximation algorithms by their performance guarantee $r > 1$ which corresponds to returning a solution whose value is at least a fraction $1/r$ of the optimal value for maximization problems or at most a factor r of the optimal value for minimization problems.

A minimization version of MSSP has been studied by Saha et al. [19] where the task is to remove the minimum amount of edges from a graph G such that the resulting subgraph H satisfies $\lambda_1(H) \leq \lambda$. They give a $(1 + \varepsilon, \varepsilon^{-1} \log n)$ bi-criteria approximation algorithm which guarantees that if an optimal solution is to remove k edges to achieve a spectral radius less than or equal to λ then the algorithm will remove $O(\varepsilon^{-1} \log n)$ times more edges (with $n = |V|$

the number of nodes in G) and returns a graph with spectral radius less than or equal to $(1+\varepsilon)\lambda$. Zhang et al. [29] study the problem of maximizing the drop in spectral radius $\lambda_1(G) - \lambda_1(H)$ where H is a subgraph of G obtained by deleting at most k edges. Their randomized algorithm, inspired by the relaxation and rounding framework, has the following conditional guarantees: if the weighted graph obtained from the solution of the relaxed semidefinite programming problem has maximum weighted degree $\Delta^* = \Omega(\log^4 n)$, then the returned subgraph satisfies the constraint on the number of edge deletions in expectation and, with high probability, the remaining graph has a spectral radius within an additive $O(\sqrt{\Delta^*})$ factor of the optimal solution. If the condition on the maximum weighted degree is not satisfied, they do not obtain any performance guarantee.

In this article we introduce the maximum spectral subgraph problem (MSSP) and our main contribution is the design of a $O(\log^2 n)$-approximation algorithm for MSSP obtained by combining a randomized algorithm based on the relaxation and rounding framework with a maximum matching algorithm. We also describe some shortcomings of existing mathematical programming formulation for variants of MSSP that attempt at minimizing the spectral radius of a given graph within a prescribed edge deletion budget.

The rest of this paper is organized as follows. In Sect. 2 we recall some simple facts from spectral graph theory and introduce appropriate notations and known results. In Sect. 3 we describe our relaxation and rounding algorithm and illustrate its usage on star graphs. Then, in Sect. 4, we prove its approximation ratio for the range $\lambda \in (\log n, \lambda_1(G))$ in general graphs. In Sect. 5 we show that a maximum matching is a $O(\lambda^2)$-approximation algorithm for MSSP. Finally, perspectives and concluding remarks are provided in Sect. 6.

2 Preliminaries

We review here useful facts about the spectrum of adjacency matrices of graphs. Unless specified, all graphs are assumed to be undirected. Recall that the adjacency matrix A of a graph $G = (V, E)$ is a symmetric matrix defined as follows:

$$A_{ij} = \begin{cases} 1 & \text{if } ij \in E \\ 0 & \text{otherwise} \end{cases}$$

Property 1. [24] (General bounds) Given a graph $G = (V, E)$, we denote by $\Delta(G)$ its largest degree. The spectral radius of the graph, defined as the largest eigenvalue of its adjacency matrix, lies between the following quantities:

$$\max\left(\sqrt{\Delta(G)}, \frac{2|E|}{|V|}\right) \leq \lambda_1(G) \leq \Delta(G) \tag{1}$$

2.1 Computational Complexity

The problem of deciding whether there exists a subgraph with at least k edges and spectral radius at most λ was studied by van Mieghem et al. [26]. We can

see that it is the decision problem associated with both MSSP and the problem of minimum edge removal introduced by Saha et al. [19] that was mentioned in Sect. 1. Van Mieghem et al. proved that the decision problem is NP-complete by reduction from the Hamiltonian path problem. It follows from this result that MSSP is NP-hard.

The reduction uses a fact from extremal spectral graph theory: the path graph on $|V|$ nodes is the graph with minimum spectral radius among all connected graphs with $|V|$ nodes and $|V|-1$ edges. Setting $\lambda = \lambda_1(P_{|V|}) = 2\cos(\pi/(|V|+1))$ and $k = |V|-1$ completes the reduction. Recall that while the spectral radius of a graph might be a real number, verifying a candidate solution amounts to checking whether the eigenvalues of a given adjacency matrix are bounded above by a given value which can be done in polynomial time to any precision [14].

Note that if the bound on the spectral radius $\lambda = 1$, then MSSP becomes the maximum matching problem which can be solved in polynomial time. Indeed, from Property 1, it is easy to see that the problem consists in finding a subgraph of degree at most 1 with maximum number of edges. Furthermore, note that all undirected graphs that are not matchings have a spectral radius larger than or equal to $\sqrt{2}$ which is the spectral radius of the path graph on 3 nodes. From this consideration, we will study the range where the bound on the spectral radius is meaningful, that is $\sqrt{2} \leq \lambda < \lambda_1(G)$.

We now present our algorithm based on the relaxation and randomized rounding framework.

3 Relaxation and Matrix Randomized Rounding

The relaxation and randomized rounding framework [17] is a general algorithmic technique composed of two steps: first, solving a continuous relaxation of the original combinatorial programming and then, sampling a discrete solution based on an optimal solution of the relaxed problem. This technique has resulted in the design of a large number of approximation algorithms for a broad range of combinatorial problems and has been the cornerstone of the application of the sum of squares hierarchy developed by Lasserre [8] and Parrilo [15] in combinatorial optimization. There are often two steps in the analysis of a relaxation and randomized rounding algorithm: finding a tight relaxation of the original problem that is solvable in polynomial time and proving that the random discrete solution is feasible with high probability.

Here we propose a mathematical programming formulation of MSSP that uses semidefinite programming (SDP) to model the constraint on the spectral radius. While linear programming allows to define optimization problems with non-negative vector variables written $x \geq 0$, SDP extends to the larger class of problems with positive semidefinite matrix variables written $X \succeq 0$ i.e. all eigenvalues of X are non-negative: $\forall\, i \in [1,n]$, $\lambda_i(X) \geq 0$. Given an input graph $G = (V, E)$ and a bound on the spectral radius λ, we write the following semidefinite programming problem with binary variables:

$$\begin{aligned} \max \quad & \sum_{ij \in E} y_{ij} \\ \text{s.t.} \quad & \sum_{ij \in E} y_{ij} A_{ij} \preceq \lambda I \\ & \sum_{j \in \Gamma(i)} y_{ij} \leq \lambda^2, \ \forall i \in V \\ & y_{ij} \in \{0,1\}, \ \forall \ ij \in E \end{aligned} \tag{SDP$_{0,1}$}$$

where A_{ij} is the adjacency matrix of the graph $G_{ij} = (V, \{ij\})$ with a single edge ij and I is the identity matrix of size $|V|$. The decision variables y_{ij} represent whether an edge ij belongs to the subgraph when $y_{ij} = 1$ or not when $y_{ij} = 0$. Recall that for a n by n square matrix, $M \preceq tI \iff \forall\, i \in [1, n]$, $\lambda_i(M) \leq t$. This means that the semidefinite constraint ensures that the adjacency matrix of the subgraph defined by y_{ij} has its spectral radius bounded above by λ. The linear constraint ensures that the degree of each node $i \in V$ in the subgraph is bounded above by λ^2. Note that this constraint is redundant given that the general bounds (1) state that the maximum degree of a graph is bounded above by the square of its spectral radius i.e. $\Delta \leq \lambda_1^2$. However this is in general not the case with weighted graphs, which will be discussed in Sect. 3.3.

The continuous relaxation of Problem (SDP$_{0,1}$) is obtained by relaxing integer constraints into box constraints. We underline that the semidefinite constraint does not originate from the relaxation as is the case for some problems which relax vector variables with quadratic constraints into a SDP problem e.g. the one used in the algorithm given by Goemans and Williamson for the maximum cut problem [6]. Our relaxation is limited to the binary variables.

$$\begin{aligned} \max \quad & \sum_{ij \in E} y_{ij} \\ \text{s.t.} \quad & \sum_{ij \in E} y_{ij} A_{ij} \preceq \lambda I \\ & \sum_{j \in \Gamma(i)} y_{ij} \leq \lambda^2, \ \forall i \in V \\ & y_{ij} \in [0,1], \ \forall \ ij \in E \end{aligned} \tag{SDP$_{\lambda\Delta}$}$$

As semidefinite programming is in P, we can solve Problem (SDP$_{\lambda\Delta}$) in polynomial time. This allows us to state our relaxation and randomized rounding algorithm. In the rest of this article, we denote scalar random variables by lowercase bold letters e.g. $\boldsymbol{x}$ and matrix random variables by uppercase bold letters e.g. $\boldsymbol{X}$. Furthermore, we denote by $\boldsymbol{x} \sim \text{Ber}(\mu)$ the fact that $\boldsymbol{x}$ is a random variable following a Bernoulli distribution of mean μ.

Algorithm 1. Relaxation & Randomized Rounding

Input: $G = (V, E)$, $\sqrt{2} \leq \lambda < \lambda_1(G)$, and $r > 1$.
Output: $H = (V, E')$ such that $\lambda_1(H) \leq \lambda$ with probability p_r.
$y^* \leftarrow \arg$ Problem $(\mathrm{SDP}_{\lambda\Delta})$
Sample $\forall ij \in E$, $\boldsymbol{x}_{ij} \sim \mathrm{Ber}(y^*_{ij}/r)$
return $H = (V, \{\, ij \in E : \boldsymbol{x}_{ij} = 1 \,\})$

We will now turn to a simple application of Algorithm 1 to the case of star graphs and determine the adequate sampling factor r that results in a feasible solution with high probability i.e. $p_r = 1 - 1/n$ where $n = |V|$.

3.1 The Case of Star Graphs

Before giving the complete analysis of our relaxation and randomized rounding algorithm we focus on a specific class of input graphs to illustrate the methodology of relaxation and randomized rounding but also to highlight the importance of the degree constraint in our proposed mathematical formulation.

Recall that a star graph $S_n = K_{1,n}$ is a graph with $V = \{0, \ldots, n\}$ and $E = \{(0,1), \ldots, (0,n)\}$. It is a well-known fact from spectral graph theory that the spectral radius of a star equals the square root of its number of edges i.e. $\lambda_1(S_n) = \sqrt{n}$. More generally, it is easy to see that a weighted star graph S_w, where each edge ij is associated with weight w_{ij}, has spectral radius $\lambda_1(S_w) = ||w||_2 = \sqrt{\sum_{ij \in E} w_{ij}^2}$. Notice that we recover the non-weighted case by setting every weight to be 1. Using this property, we determine that the number of edges in an optimal solution of Problem $(\mathrm{SDP}_{0,1})$ is exactly $\lfloor \lambda^2 \rfloor$ edges. We denote the optimal value of MSSP on a star graph S_n and parameter λ by $\mathrm{opt}(S_n, \lambda) = \lfloor \lambda^2 \rfloor$.

To analyze the gap between the combinatorial problem and our relaxation, we now compute the value of an optimal solution of Problem $(\mathrm{SDP}_{\lambda\Delta})$. First, we can use the above definition of the spectral radius of a star graph to replace the semidefinite constraint by $||y||_2 \leq \lambda$. Second, we interpret the degree constraint as a constraint on the ℓ_1-norm of y. This means that we can compute an optimal solution of Problem $(\mathrm{SDP}_{\lambda\Delta})$ by solving the following second-order cone programming problem:

$$\begin{aligned} \max_{y \in [0,1]^{|E|}} \quad & ||y||_1 \\ \text{s.t.} \quad & ||y||_2 \leq \lambda \\ & ||y||_1 \leq \lambda^2 \end{aligned} \qquad (\mathrm{SOCP}_{\lambda\Delta})$$

It is now easy to see that an optimal solution of this problem has value at most λ^2 and that can be achieved by any y such that $||y||_2 \leq \lambda$ e.g. the uniform solution where $\forall ij \in E, y^*_{ij} = \lambda^2/n$ has ℓ_2-norm $||y^*||_2 = \lambda/\sqrt{n}$. We denote by $\mathrm{opt}_{\mathrm{rel}}(S_n, \lambda) = \lambda^2$ the optimal value of the relaxation.

We now have a complete description of the integrality gap g_{S_n} of our relaxation for star graphs. The gap is the largest ratio between the optimal value of the relaxation and the optimal value of the original problem:

$$g_{S_n} \stackrel{\text{def}}{=} \max_{S_n,\lambda} \frac{\text{opt}_{\text{rel}}(S_n,\lambda)}{\text{opt}(S_n,\lambda)} = \frac{\lambda^2}{\lfloor \lambda^2 \rfloor} \leq \frac{4}{3} \tag{2}$$

where the last inequality comes from the fact that $\lambda \geq \sqrt{2}$.

3.2 Erdös-Rényi Stars

Now that we have solved our relaxation of MSSP, we will use the computed optimal solution y^* to sample a discrete solution, here a random subgraph $S_{\boldsymbol{x}}$ of the original star graph S_n. For this purpose we introduce for each edge ij an independent random variable $\boldsymbol{x}_{ij} \sim \text{Ber}(y^*_{ij}/r)$. By definition the random number of edges $\boldsymbol{x}$ of the random subgraph $S_{\boldsymbol{x}}$ is a sum of independent Bernoulli random variables with mean $\mathbb{E}\,\boldsymbol{x} = \sum_{ij} y^*_{ij}/r = \lambda^2/r$.

If for some $r > 1$ the random subgraph $S_{\boldsymbol{x}}$ satisfies the spectral radius constraint with high probability, i.e. $p_r = 1 - O(1/n)$, then we would have a polynomial time randomized r-approximation algorithm. We obtain the following approximation algorithm in the case of star graphs:

Theorem 1. *(Feasible with constant probability) Given a star graph S_n, a bound on the spectral radius $\lambda \geq \sqrt{2}$, and an optimal solution y^* of Problem ($SOCP_{\lambda\Delta}$), the random partial subgraph $S_{\boldsymbol{x}}$ obtained by keeping edges according to independent random variables $\boldsymbol{x}_{ij} \sim Ber(y^*_{ij}/r)$ is a feasible solution of* MSSP *with probability $p_r \geq 2/3$ for $r = 4$.*

Proof. As is common practice in the analysis of randomized algorithms [12], we use the Chernoff bound to get an estimate of the probability that our sampled solution is feasible. Recall that the Chernoff bound gives an upper bound on the probability that a sum of independent random variables exceeds a certain value.

Theorem. [12] *(Chernoff bound) Let $\boldsymbol{x} = \sum_{i=1}^{n} \boldsymbol{x}_i$ where each $\boldsymbol{x}_i$ is an independent random variable. We have for a given value $a > 0$ the following estimate:*

$$\Pr(\boldsymbol{x} \geq a) \leq \min_{t>0} e^{-ta} \prod_{i=1}^{n} \mathbb{E}\, e^{t\boldsymbol{x}_i}.$$

We directly apply the Chernoff bound on the random number of edges $\boldsymbol{x} = \sum_{ij\in E} \boldsymbol{x}_{ij}$ for the value $a = \lambda^2$. After using the fact that $\mathbb{E}\, e^{t\boldsymbol{x}_{ij}} \leq e^{\frac{y^*_{ij}}{r}(e^t-1)}$, we obtain the following:

$$\Pr(\boldsymbol{x} \geq \lambda^2) \leq \min_{t>0} e^{-t\lambda^2} \exp\left(\sum_{ij\in E} \frac{y^*_{ij}}{r}(e^t - 1)\right)$$

$$\Pr(\boldsymbol{x} \geq \lambda^2) \leq \min_{t>0} \exp\left(\frac{\lambda^2}{r}(e^t - 1) - t\lambda^2\right).$$

The minimum of the r.h.s. is attained at $t = \log r$ under the condition that $t > 0$ from which we deduce that $r = 1 + h$ for some $h > 0$. The bound then simplifies into:

$$\Pr(\boldsymbol{x} \geq \lambda^2) \leq \exp\left(\lambda^2\left(\frac{h}{1+h} - \log(1+h)\right)\right).$$

We choose $r = 1 + h = 4$ to produce the following readable bound which remains valid for any $\lambda \geq \sqrt{2}$:

$$\Pr(\boldsymbol{x} \geq \lambda^2) \leq \exp(-0.64\lambda^2) \leq \frac{1}{3}$$

which concludes the proof of Theorem 1. Recall indeed that the event $\lambda_1(S_{\boldsymbol{x}}) \geq \lambda$ is equivalent to the event $\boldsymbol{x} \geq \lambda^2$. □

Since our success probability for a single sample $p_r \geq 2/3$ we can amplify it by repetition in polynomial time to obtain a solution of expected value $\lambda^2/4$ and such that the solution is feasible with high probability $p_r = 1 - O(1/n)$.

To summarize the case of star graphs, our relaxation and randomized rounding algorithm is a polynomial time algorithm which returns with high probability a feasible star graph of expected size $\lambda^2/4$.

3.3 Without the Degree Constraint

It is important to notice that the degree constraint played a significant role in the tightness of the relaxation in the case of star graphs. Reusing the same analysis as in Sect. 3.1 we can see that Problem ($\text{SDP}_{\lambda\Delta}$) without the degree constrained is equivalent to the following problem:

$$\begin{aligned} \max_{y \in [0,1]^{|E|}} \quad & ||y||_1 \\ \text{s.t.} \quad & ||y||_2 \leq \lambda \end{aligned} \tag{SOCP_λ}$$

By a geometrical argument, we notice that the uniform solution $\forall ij \in E, y_{ij} = \lambda/\sqrt{n}$ is the unique optimal solution of Problem (SOCP_λ). It follows that the associated optimal value $\text{opt}_{\text{rel}'}(S_n, \lambda) = \lambda\sqrt{n}$.

In that case, the integrality gap of the relaxation given by Problem (SOCP_λ) is

$$g'_{S_n} \stackrel{\text{def}}{=} \max_{S_n, \lambda} \frac{\text{opt}_{\text{rel}'}(S_n, \lambda)}{\text{opt}(S_n, \lambda)} = \frac{\lambda\sqrt{n}}{\lfloor \lambda^2 \rfloor} = O\left(\frac{\sqrt{n}}{\lambda}\right) \tag{3}$$

which translates into a much higher $r = O(g'_{S_n})$ than the constant obtained in Sect. 3.1. Problem formulations focusing on minimizing the spectral radius given an edge deletion budget cannot a priori bound the maximum degree of the resulting weighted graph. This additional information is a key advantage over problems that optimize the spectral parameter.

We are now ready to describe our matrix randomized rounding whose analysis follows a similar structure to the one for star graphs. However we need to use more powerful concentration inequalities than the Chernoff bound to obtain bounds on the spectral radius of the random matrix we sample. This sampling can be seen as a special case of inhomogeneous Erdös-Rényi random graphs.

4 Spectral Subgraphs in General Graphs

In order to extend the analysis of Algorithm 1 to arbitrary graphs we turn to more advanced concentration inequalities that describe the behavior of random matrices and in particular their spectrum. Fortunately, recent results in the analysis of random matrices (cf. the survey by Tropp [22]) provide tail bounds for the largest eigenvalue of random matrices. These results are directly applicable to the analysis of Algorithm 1 for finding the sampling factor r that guarantees that the returned solution is feasible with high probability $p_r = 1 - 1/n$.

We start by presenting the generic matrix Bernstein bound and its application to adjacency matrices following the work of Radcliffe and Chung [3]. Finally we give the proof that Algorithm 1 is a randomized $O(\log n)$-approximation algorithm with the following property:

Theorem 2. *(Feasible with constant probability) Given a graph $G = (V, E)$ with $|V| = n$, a bound on the spectral radius $\lambda \geq \log n$, and an optimal solution y^* of Problem ($SDP_{\lambda\Delta}$), the random subgraph H obtained by keeping edges $ij \in E$ according to independent random variables $\boldsymbol{x}_{ij} \sim Ber(y^*_{ij}/r)$ is a feasible solution of MSSP with probability $p_r \geq 2/3$ for $r = O(\log n)$.*

4.1 Following the Matrix Bernstein Bound

The matrix Bernstein bound is a generalization of the classical Bernstein bound to the setting of independent random matrices. The theorem states the following:

Theorem. [3] *(Matrix Bernstein) Let $\boldsymbol{X} = \sum_i \boldsymbol{X}_i$ where each $\boldsymbol{X}_i$ is an independent symmetric random matrix of size n which is centered $\mathbb{E}\,\boldsymbol{X}_i = 0$ and bounded in spectral norm $\lambda_1(\boldsymbol{X}_i) \leq L$. We define the matrix variance of $\boldsymbol{X}$ by $v(\boldsymbol{X}) = \lambda_1(\sum_i \mathbb{E}\,\boldsymbol{X}_i^2)$. The following tail inequality holds:*

$$\Pr(\lambda_1(\boldsymbol{X}) \geq a) \leq n \exp\left(-\frac{a^2}{2v(\boldsymbol{X}) + 2La/3}\right). \tag{4}$$

The output of Algorithm 1 corresponds to a random adjacency matrix $\boldsymbol{A}$ which is the sum of independent random adjacency matrices each corresponding to an edge in the random graph. Let $A_{ij} = (E_{ij} + E_{ji})$ where the E_{ij} form the canonical basis for $M_{n,n}$ and denote by $\boldsymbol{x}_{ij}$ a Bernoulli random variable of mean y^*_{ij}/r. We have the following:

$$\boldsymbol{A} = \sum_{ij \in E} \boldsymbol{x}_{ij} A_{ij} \tag{5}$$

Note that our random adjacency edges have non-zero mean $\mathbb{E}\,\boldsymbol{x}_{ij}A_{ij} = (y^*_{ij}/r)A_{ij}$. Fortunately, applying Weyl's inequalities on $\boldsymbol{A}$ and $\mathbb{E}\,\boldsymbol{A}$ will give us control over the spectral radius of $\boldsymbol{A}$ by proxy.

Theorem. [1] *(Weyl's inequalities) Let X and Y be two symmetric matrices,*

$$\lambda_1(X - Y) \leq \varepsilon \implies |\lambda_1(X) - \lambda_1(Y)| \leq \varepsilon \tag{6}$$

This theorem implies that bounding the spectral radius of our centered random adjacency matrix by $(1-1/r)\lambda$ will give us the adequate bound on the spectral radius of $\boldsymbol{A}$. Since we only consider the event where $\boldsymbol{A}$ has greater spectral radius than $\mathbb{E}\,\boldsymbol{A}$, we drop the absolute value:

$$\lambda_1(\boldsymbol{A}-\mathbb{E}\,\boldsymbol{A}) < \left(1-\frac{1}{r}\right)\lambda \implies \lambda_1(\boldsymbol{A}) - \lambda_1(\mathbb{E}\,\boldsymbol{A}) < \left(1-\frac{1}{r}\right)\lambda$$

and by feasibility of an optimal solution of the relaxed SDP, i.e. Problem $(\mathrm{SDP}_{\lambda\Delta})$, we have $\lambda_1(\mathbb{E}\,\boldsymbol{A}) \leq \lambda/r$ which gives:

$$\lambda_1(\boldsymbol{A}-\mathbb{E}\,\boldsymbol{A}) < \left(1-\frac{1}{r}\right)\lambda \implies \lambda_1(\boldsymbol{A}) < \lambda.$$

From the general bounds of (1) we know that the spectral radius of the centered adjacency matrix of a random edge ij is either y^*_{ij}/r (no edge) or $1-y^*_{ij}/r$ (one edge) which lets us bound the spectrum of each summand. In the worst case we have, for each edge ij:

$$\lambda_1\left(\left(\boldsymbol{x}_{ij}-\frac{y^*_{ij}}{r}\right)A_{ij}\right) \leq 1 \tag{7}$$

4.2 Proof of Theorem 2

We start by computing the matrix variance:

$$v(\boldsymbol{A}-\mathbb{E}\,\boldsymbol{A}) = \lambda_1\left(\sum_{ij\in E}\mathrm{Var}(\boldsymbol{x}_{ij}A_{ij})\right).$$

Since $\mathrm{Var}(\boldsymbol{x}_{ij}A_{ij}) = \mathrm{Var}(\boldsymbol{x}_{ij})A^2_{ij}$ and $A^2_{ij} = D_i + D_j$ where $D_v = E_{vv}$, we obtain a clean expression for the variance of the centered adjacency matrix as the spectral radius of the matrix of degree variances:

$$\begin{aligned} v(\boldsymbol{A}-\mathbb{E}\,\boldsymbol{A}) &= \lambda_1\left(\sum_{ij\in E}\mathrm{Var}(\boldsymbol{x}_{ij})(D_i+D_j)\right) \\ &= \max_{i\in V}\sum_{j\in\Gamma(i)}\frac{y^*_{ij}}{r}\left(1-\frac{y^*_{ij}}{r}\right) \\ &\leq \max_{i\in V}\sum_{j\in\Gamma(i)}\frac{y^*_{ij}}{r} \end{aligned}$$

and by feasibility of an optimal solution of the relaxation, the degree constraint holds which means that $\max_{i\in V}\sum_{j\in\Gamma(i)} y^*_{ij} \leq \lambda^2$ and gives:

$$v(\boldsymbol{A}-\mathbb{E}\,\boldsymbol{A}) \leq \frac{\lambda^2}{r}.$$

We now fulfill all the prerequisites to apply the matrix Bernstein bound on $\boldsymbol{A} - \mathbb{E}\,\boldsymbol{A}$ and $L = 1$. To explicitly describe the fact that the approximation ratio $r > 1$ we introduce as earlier $h > 0$ such that $r = 1 + h$. We apply the Bernstein bound for the value $a = (h/(1+h))\lambda$:

$$\begin{aligned}\Pr\left(\lambda_1(\boldsymbol{A} - \mathbb{E}\,\boldsymbol{A}) \geq \frac{h}{1+h}\lambda\right) &\leq n \exp\left(-\frac{1}{2}\frac{a^2}{v(\boldsymbol{A} - \mathbb{E}\,\boldsymbol{A}) + \frac{a}{3}}\right) \\ &= n \exp\left(-\frac{1}{2}\frac{a^2}{\frac{\lambda^2}{1+h} + \frac{1}{3}\frac{h}{1+h}\lambda}\right) \\ &\leq n \exp\left(-\frac{1}{2}\frac{h^2}{(1+h)^2}\,\frac{\lambda^2}{\frac{\lambda^2}{1+h} + \frac{1}{3}\frac{h}{1+h}\lambda}\right).\end{aligned}$$

We simplify the above expression to obtain:

$$\begin{aligned}\Pr\left(\lambda_1(\boldsymbol{A} - \mathbb{E}\,\boldsymbol{A}) \geq \frac{h}{1+h}\lambda\right) &\leq n \exp\left(-\frac{1}{2}\frac{h^2}{(1+h)^2}\,\frac{\lambda^2}{\frac{\lambda}{1+h}\,(\lambda + h/3)}\right) \\ &= n \exp\left(-\frac{1}{2}\,\frac{h^2}{1+h}\,\frac{\lambda}{\lambda + h/3}\right).\end{aligned}$$

As in the case of star graphs, we will derive possible values for r (resp. for h) such that the probability of our subgraph H being infeasible is less than $1/3$. For this, we attempt to derive an upper bound for the argument of the exponential as $n\exp(-x) \leq \frac{1}{3}$ implies that $x \geq \log 3n$.

We are looking for values of h and λ such that the following inequality holds:

$$\frac{1}{2}\,\frac{h^2}{1+h}\,\frac{\lambda}{\lambda + h/3} \geq \log 3n$$

We start by deriving a lower bound on λ function of h. In the above inequality, $\lambda/(\lambda + h/3)$ can be arbitrarily small if h is unbounded. To prevent this, we impose that, for a certain constant $c > 0$:

$$\frac{1}{2}\,\frac{\lambda}{\lambda + h/3} \geq c$$

which implies that

$$\lambda \geq \frac{2c}{3 - 6c}h.$$

Choosing $c = 1/4$ gives us the condition that $\lambda \geq h/3$.

Now we are left with finding the value of h such that:

$$\frac{1}{4}\,\frac{h^2}{1+h} \geq \log 3n.$$

For all values of n, it is sufficient to take $h = 3 \log n$ which completes the proof. □

Algorithm 1 is a randomized algorithm which returns a feasible solution with probability greater than $2/3$ and of expected value within $1 + 3 \log n$ of the value of an optimal solution whenever $\lambda \geq \log n$. Recall that the success probability of such an algorithm can be amplified to high probability in polynomial time. We now turn to a different algorithm to handle the range $\lambda \in [\sqrt{2}, \log n)$.

5 Maximum Matching

After designing an approximation algorithm for MSSP for the range of the spectral bound $\lambda \in (\log n, \lambda_1(G))$, we turn to the well-studied maximum matching problem: finding a subgraph M consisting of the maximum number of non-adjacent edges in a given graph G. The number of edges in M is often called the matching number $\nu(G)$ of the graph. We use a spectral generalization of a classical lower bound on the matching number due to Stevanović [21] which states the following:

Theorem. [21] *(Spectral lower bound on the matching number) Given a graph $G = (V, E)$ we have the following lower bound:*

$$\nu(G) \geq \frac{|E|}{\lambda_1^2(G) - 1}.$$

This static lower bound can be immediately turned into an approximation algorithm since computing a maximum matching can be done in polynomial time.

Algorithm 2. Maximum matching

Input: $G = (V, E)$, $\sqrt{2} \leq \lambda \leq \lambda_1(G)$
Output: $H = (V, E')$ such that $\lambda_1(H) \leq \lambda$
return $H = \arg \nu(G)$

Theorem 3. *(Approximation by maximum matching) Given $G = (V, E)$ and a spectral bound $\lambda > 0$, a maximum matching of G is a $(\lambda^2 - 1)$-approximation for MSSP.*

Proof. Denoting by H^* an optimal solution of MSSP for a graph G and spectral bound λ, we know that H^* is a partial subgraph of G which implies $\nu(G) \geq \nu(H^*)$. We also know that H^* is feasible i.e. $\lambda_1(H^*) \leq \lambda$. Combining these two statements together with the lower bound of Stevanović, we obtain the following inequality:

$$\nu(G) \geq \nu(H^*) \geq \frac{\mathrm{opt}(G, \lambda)}{\lambda^2 - 1}$$

which shows that the size of a maximum matching is within a factor of $\lambda^2 - 1$ of an optimal solution of MSSP. Furthermore any matching has spectral radius equal to 1 i.e. is trivially feasible. □

Used in the range $\lambda \in [\sqrt{2}, \log n)$ a maximum matching is a $O(\log^2 n)$-approximation algorithm in the worst-case. We then combine Algorithm 1 with Algorithm 2 to obtain a $O(\log^2 n)$-approximation algorithm for all values of λ.

6 Conclusion and Perspectives

We have introduced the maximum spectral subgraph problem and designed a randomized $O(\log^2 n)$-approximation algorithm based on the relaxation and rounding framework to solve it.

In terms of lower bounds, we currently do not have any result regarding hardness of approximation, but we are actively exploring this direction. To the best of our knowledge, no inapproximability results have been established for problems related to the spectrum of a graph. Indeed, NP-hardness results found in the literature [11,26] are based on reductions which relate extremal values in spectral graph theory to classical computational problems. These reductions cannot be directly extended to obtain an approximation gap.

Without a better lower bound than NP-hardness, we are compelled to find new techniques to improve our current upper bound. First, the continuous relaxation used in Algorithm 1 is rather natural aside from the redundant degree constraints. It would be interesting to see if stronger relaxations could be used to obtain more information about the random graph e.g. strong bounds on the variance of the random degrees. For this purpose we would like to consider a sum-of-squares relaxation for the binary semidefinite programming problem. Indeed, Nie [13] has given an extension of the classical sum-of-squares hierarchy to include positivity certificates for matrix variables. This relates to the question of generalizing the results of Raghavendra [18] on maximum constraint satisfaction problems where constraints apply to at most k variables to maximum constraint satisfaction problems with spectral constraints which, by definition, involve all variables at once. Aside from strengthening the relaxation, there is opportunity for improvement in developing more precise tail bounds on the spectrum of random adjacency matrices following recent results by van Handel [23] as well as by Le, Levina, and Vershynin [9]. On a separate note, we are currently working on applying the method of conditional probabilities to derandomize Algorithm 1 in order to obtain a deterministic approximation algorithm. The analysis of Sect. 5 focuses on the maximum matching problem as a way of computing a feasible solution for the range $\lambda \in [\sqrt{2}, \log n)$. It is natural to wonder whether the degree constrained subgraph problem with $\Delta \leq \lambda$ (also known as the simple λ-matching problem) could be proven to return a better solution, and possibly match the $O(\log n)$ ratio obtained by Algorithm 1.

Finally, we are also interested in applying a similar strategy to the problem of adding the smallest number of edges to reach a given algebraic connectivity i.e. a lower bound on the second smallest eigenvalue of the Laplacian matrix of the graph. This problem, proven NP-hard by Mosk-Aoyama [11], is a variant of the problem of finding the maximum algebraic connectivity given an edge addition budget proposed by Ghosh and Boyd [5]. While Kolla et al. have designed an

approximation algorithm with conditional guarantees [7] for the original problem, we hope that our methodology could apply to the variant and lead to an unconditional approximation ratio.

References

1. Bhatia, R.: Matrix Analysis, vol. 169. Springer, Heidelberg (2013). https://doi.org/10.1007/978-1-4612-0653-8
2. Chakrabarti, D., Wang, Y., Wang, C., Leskovec, J., Faloutsos, C.: Epidemic thresholds in real networks. ACM Trans. Inf. Syst. Secur. (TISSEC) **10**(4), 1 (2008)
3. Chung, F., Radcliffe, M.: On the spectra of general random graphs. Electron. J. Comb. **18**(1), 215 (2011)
4. Ganesh, A., Massoulié, L., Towsley, D.: The effect of network topology on the spread of epidemics. In: Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2005), vol. 2, pp. 1455–1466 (2005)
5. Ghosh, A., Boyd, S.: Growing well-connected graphs. In: Proceedings of the IEEE Conference on Decision and Control (CDC 2006), pp. 6605–6611. IEEE (2006)
6. Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. J. ACM (JACM) **42**(6), 1115–1145 (1995)
7. Kolla, A., Makarychev, Y., Saberi, A., Teng, S-H.: Subgraph sparsification and nearly optimal ultrasparsifiers. In: Proceedings of the ACM Symposium on Theory of Computing (STOC 2010), pp. 57–66. ACM (2010)
8. Lasserre, J.B.: Global optimization with polynomials and the problem of moments. SIAM J. Optim. **11**(3), 796–817 (2001)
9. Le, C.M., Levina, E., Vershynin, R.: Concentration and regularization of random graphs. Random Struct. Algorithms **51**(3), 538–561 (2017)
10. Mehdi, S.A., Khalid, J., Khayam, S.A.: Revisiting traffic anomaly detection using software defined networking. In: Sommer, R., Balzarotti, D., Maier, G. (eds.) RAID 2011. LNCS, vol. 6961, pp. 161–180. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23644-0_9
11. Mosk-Aoyama, D.: Maximum algebraic connectivity augmentation is NP-hard. Oper. Res. Lett. **36**(6), 677–679 (2008)
12. Motwani, R., Raghavan, P.: Randomized Algorithms. Chapman & Hall/CRC, Boca Raton (2010)
13. Nie, J.: Polynomial matrix inequality and semidefinite representation. Math. Oper. Res. **36**(3), 398–415 (2011)
14. Pan, V.Y., Chen, Z.Q.: The complexity of the matrix eigenproblem. In: Proceedings of the ACM Symposium on Theory of Computing (STOC 1999), pp. 507–516 (1999)
15. Parrilo, P.A.: Semidefinite programming relaxations for semialgebraic problems. Math. Program. **96**(2), 293–320 (2003)
16. Aditya Prakash, B., Chakrabarti, D., Faloutsos, M., Valler, N., Faloutsos, C.: Threshold conditions for arbitrary cascade models on arbitrary networks. In: Proceedings of the IEEE International Conference on Data Mining (ICDM 2011), pp. 537–546 (2011)
17. Raghavan, P., Tompson, C.D.: Randomized rounding: a technique for provably good algorithms and algorithmic proofs. Combinatorica **7**(4), 365–374 (1987)

18. Raghavendra, P.: Optimal algorithms and inapproximability results for every CSP? In: Proceedings of the ACM Symposium on Theory of Computing (STOC 2008), pp. 245–254. ACM (2008)
19. Saha, S., Adiga, A., Aditya Prakash, B., Vullikanti, A.K.S.: Approximation algorithms for reducing the spectral radius to control epidemic spread. In: Proceedings of the SIAM International Conference on Data Mining (SDM 2015), pp. 568–576 (2015)
20. Shin, S., Gu, G.: CloudWatcher: network security monitoring using OpenFlow in dynamic cloud networks (or: how to provide security monitoring as a service in clouds?). In: Proceedings of the IEEE International Conference on Network Protocols (ICNP 2012), pp. 1–6. IEEE (2012)
21. Stevanović, D.: Resolution of AutoGraphiX conjectures relating the index and matching number of graphs. Linear Algebra Appl. **8**(433), 1674–1677 (2010)
22. Tropp, J.A., et al.: An introduction to matrix concentration inequalities. Found. Trends® Mach. Learn. **8**(1–2), 1–230 (2015)
23. van Handel, R.: Structured random matrices. In: Carlen, E., Madiman, M., Werner, E.M. (eds.) Convexity and Concentration. TIVMA, vol. 161, pp. 107–156. Springer, New York (2017). https://doi.org/10.1007/978-1-4939-7005-6_4
24. Van Mieghem, P.: Graph Spectra for Complex Networks. Cambridge University Press, Cambridge (2010)
25. Van Mieghem, P., Omic, J., Kooij, R.: Virus spread in networks. IEEE/ACM Trans. Netw. (TON) **17**(1), 1–14 (2009)
26. Van Mieghem, P., et al.: Decreasing the spectral radius of a graph by link removals. Phys. Rev. E **84**(1), 016101 (2011)
27. Wang, G., Ng, T.S., Shaikh, A.: Programming your network at run-time for big data applications. In: Proceedings of the Workshop on Hot Topics in Software Defined Networks (HotSDN 2012), pp. 103–108. ACM (2012)
28. Wang, Y., Chakrabarti, D., Wang, C., Faloutsos, C.: Epidemic spreading in real networks: an eigenvalue viewpoint. In: Proceedings of the International Symposium on Reliable Distributed Systems (SRDS 2003), pp. 25–34. IEEE (2003)
29. Zhang, Y., Adiga, A., Vullikanti, A., Aditya Prakash, B.: Controlling propagation at group scale on networks. In: Proceedings of the International Conference on Data Mining (ICDM 2015), pp. 619–628 (2015)