On the Inapproximability of Finding Minimum Monitoring Edge-Geodetic Sets (short paper)

Davide Bilò[1], Giordano Colli[1,*], Luca Forlizzi[1,*] and Stefano Leucci[1]

[1]Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, Italy

Abstract

Given an undirected connected graph $G = (V(G), E(G))$ on n vertices, the minimum Monitoring Edge-Geodetic Set (MEG-set) problem asks to find a subset $M \subseteq V(G)$ of minimum cardinality such that, for every edge $e \in E(G)$, there exist $x, y \in M$ for which all shortest paths between x and y in G traverse e.

We show that, for any constant $c < \frac{1}{2}$, no polynomial-time $(c \log n)$-approximation algorithm for the minimum MEG-set problem exists, unless $\mathsf{P} = \mathsf{NP}$.

Keywords

Geodetic Set, Edge Monitoring, Inapproximability

1. Introduction

We study the minimum Monitoring Edge-Geodetic Set (MEG-set) problem. Given an undirected and connected graph $G = (V(G), E(G))$ on n vertices, we say that an edge $e \in E(G)$ is monitored by a pair $\{x, y\}$ if e lies on *every* shortest path between x and y in G. Moreover, we say that a set $M \subseteq V(G)$ monitors e if there exist $x, y \in M$ such that $\{x, y\}$ monitors e. A MEG-set of G is a subset $M \subseteq V(G)$ that monitors all edges in $E(G)$. The goal of the minimum MEG-set problem is that of finding a MEG-set of G of minimum cardinality. This problem naturally arises in a scenario where network failures are detected by probes that can detect the current distance that separates them, placed in some network nodes. Then, finding a MEG-set of minimum cardinality means determining the minimum number of probes needed to monitor the entire network. Furthermore, the importance of the problem also derives from its connections with other concepts and scenarios related to network monitoring, such as distance edge-monitoring and the computation of edge-geodetic sets.

The minimum MEG-set problem has been introduced in [1], where the authors focus on providing upper and lower bounds to the size of minimum MEG-sets for both general graphs and special classes of graphs. Further bounds on the size of MEG-sets have been given in [2, 3, 4, 5, 6]. The problem was proven to be NP-hard by [2] on general graphs, and by [3] on 3-degenerate 2-apex graphs. However, to the best of our knowledge, no inapproximability result is currently known.

ICTCS'24: Italian Conference on Theoretical Computer Science, September 11–13, 2024, Torino, Italy
*Corresponding author.
davide.bilo@univaq.it (D. Bilò); giordano.colli@student.univaq.it (G. Colli); luca.forlizzi@univaq.it (L. Forlizzi); stefano.leucci@univaq.it (S. Leucci)
0000-0003-3169-4300 (D. Bilò); 0000-0002-3923-7668 (L. Forlizzi)

In this paper, we show that the problem of finding a MEG-set of minimum size is not approximable within a factor of $c \ln n$, where $c < \frac{1}{2}$ is a constant of choice, unless $\mathsf{P} = \mathsf{NP}$.

2. Preliminaries

Lemma 1. *Let v be a vertex of degree 1 in G. Vertex v belongs to all MEG-sets of G.*

Proof. Assume towards a contradiction that some MEG-set $M \subseteq V(G) \setminus \{v\}$ of G exists. Let $\{x, y\}$ with $x, y \in M$ be a pair of vertices that monitors the unique edge incident to v. Since $\{x, y\} \cap \{v\} = \emptyset$, all shortest path from x to y in G contain v as an internal vertex, thus v must have at least two incident edges, contradicting the hypothesis of the lemma. □

Lemma 2. *Let u be a vertex of degree 1 in G and let v be its sole neighbor. If $|V(G)| \geq 3$ and M is a MEG-set of G then $M \setminus \{v\}$ is a MEG-set of G.*

Proof. From Lemma 1 we know that $u \in M$. We start by observing that $\{u, v\}$ only monitors edge (u, v) and, since $|E(G)| \geq |V(G)| - 1 \geq 2$, there must exist some vertex $x \in M \setminus \{u, v\}$.

Since (u, v) is a bridge of G, it is traversed by all paths between u and any vertex in $V(G) \setminus \{u\}$ and, in particular, it is monitored by $\{u, x\}$.

We now turn our attention to the edges in $E \setminus \{(v, u)\}$ and show that any such edge e that is monitored by $\{v, y\}$, with $y \in M$, is also monitored by $\{u, y\}$. Notice, since $e \neq (u, v)$, we have $y \neq u$. Consider any shortest path P from u to y in G and observe that P consists of the edge (u, v) followed by a path P' from v to y. By suboptimality of shortest paths, P' is a shortest path from v to y in G. Since $\{v, y\}$ monitors e, P' contains e and so does P. Thus, $\{u, y\}$ monitors e. □

3. Our Inapproximability Result

We reduce from the Set Cover problem. A Set Cover instance $\mathcal{I} = \langle X, \mathcal{S} \rangle$ is described as a set of η items $X = \{x_1, \ldots, x_\eta\}$, and a collection $\mathcal{S} = \{S_1, \ldots, S_h\}$ of $h \geq 2$ distinct subsets of X, such that each subset contains at least two items and each item appears in at least two subsets.[1] The goal is that of computing a collection $\mathcal{S}^* \subseteq \mathcal{S}$ of minimum size such that $\cup_{S_i \in \mathcal{S}^*} S_i = X$.[2]

It is known that, unless $\mathsf{P} = \mathsf{NP}$, the Set Cover problem is not approximable within a factor of $(1 - \varepsilon) \ln |\mathcal{I}|$, where $\varepsilon > 0$ is a constant and $|\mathcal{I}|$ is the size of the Set Cover instance [7].

Given an instance $\mathcal{I} = \langle X, \mathcal{S} \rangle$ of Set Cover, we can build an associated bipartite graph H whose vertex set $V(H)$ is $X \cup \mathcal{S}$ and such that H contains edge (x_i, S_j) if and only if $x_i \in S_j$. We define $N = h + \eta$. Observe that $|\mathcal{I}| \geq N$.

[1] This can be guaranteed w.l.o.g. by repeatedly reducing the instance by applying the first applicable of the following two reduction rules. Rule 1: if there exists an item x_i that is contained by a single subset S_j, then S_j belongs to all feasible solutions, and we reduce to the instance in which both S_j and x_i have been removed. Rule 2: if there exists a subset S_j that contains a single element, then (due to Rule 1) there is an optimal solution that does not contain S_j, and we reduce to the instance in which S_j has been removed. Notice that this process can only decrease the values of η and h.

[2] We assume w.l.o.g. that $\cup_{i=1}^{h} S_i = X$, i.e., that a solution exists.

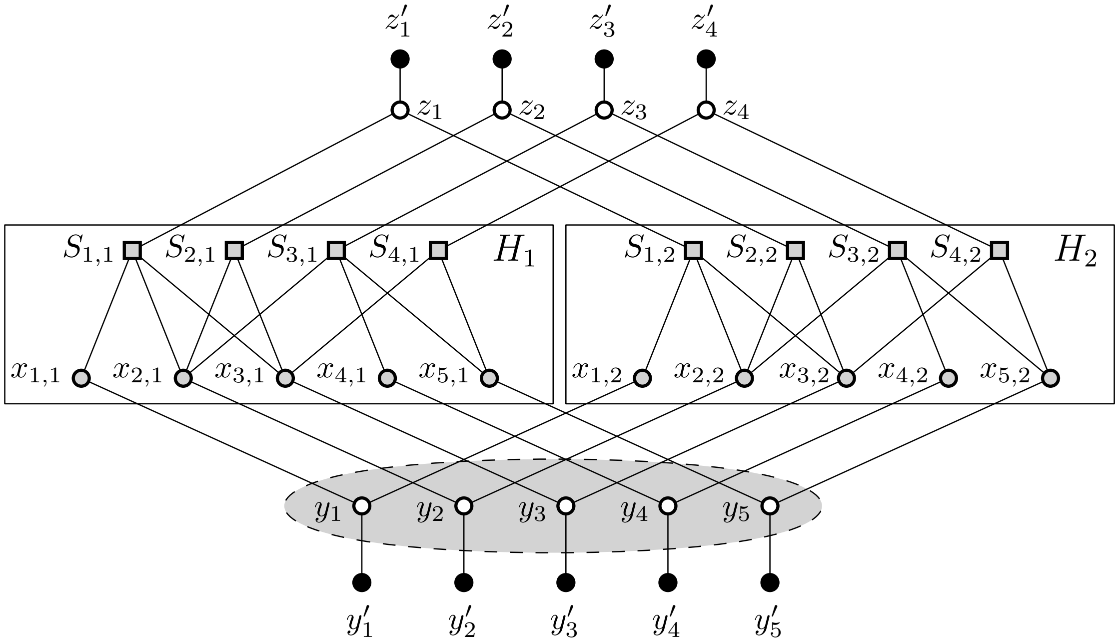

Figure 1: The graph G obtained by applying our reduction with $k = 2$ to the Set Cover instance $\mathcal{I} = \langle X, \mathcal{S} \rangle$ with $\eta = 5$, $h = 4$, $S_1 = \{x_1, x_2, x_3\}$, $S_2 = \{x_2, x_3\}$, $S_3 = \{x_2, x_4, x_5\}$, and $S_4 = \{x_3, x_5\}$. To reduce clutter, the edges of the clique induced by the vertices y_i (in the gray area) are not shown.

Let k be an integer parameter, whose exact value will be chosen later, that satisfies $2 \leq k = O(\text{poly}(N))$. We construct a graph G that contains k copies $H_1, \ldots, H_k$ of H as induced subgraphs. In the following, for any $\ell = 1, \ldots, k$, we denote by $x_{i,\ell}$ and $S_{j,\ell}$ the vertices of H_ℓ corresponding to the vertices x_i and S_j of H, respectively. More precisely, we build G by starting with a graph that contains exactly the k copies $H_1, \ldots, H_k$ of H and augmenting it as follows (see Figure 1):

- For each item $x_i \in X$, we add two new vertices y_i, y'_i along with the edge (y_i, y'_i) and all the edges in $\{(x_{i,\ell}, y_i) \mid \ell = 1, \ldots, k\}$;
- We add all the edges $(y_i, y_{i'})$ for all $1 \leq i \leq i' \leq \eta$, so that the subgraph induced by $y_1, \ldots, y_\eta$ is complete.
- For each set $S_j \in \mathcal{S}$, we add two new vertices z_j, z'_j along with the edge (z_j, z'_j), and all the edges in $\{(S_{j,\ell}, z_j) \mid \ell = 1, \ldots, k\}$.

Observe that the number n of vertices of G satisfies $n = 2h + 2\eta + k(\eta + h) = (k+2)(\eta + h) = (k+2)N$.

Let $Y = \{y_i \mid i = 1, \ldots, \eta\}$, $Y' = \{y'_i \mid i = 1, \ldots, \eta\}$, $Z = \{z_j \mid i = 1, \ldots, h\}$, and $Z' = \{z_j \mid i = 1, \ldots, h\}$. Moreover, define L as the set of all vertices of degree 1 in G, i.e., $L = Y' \cup Z'$. By Lemma 1, the vertices in L belong to all MEG-sets of G.

Lemma 3. *L monitors all edges having both endvertices $Y \cup Y' \cup Z \cup Z'$.*

Proof. Observe that all shortest paths from $y'_i \in L$ (resp. $z'_j \in L$) to any other vertex $v \in L \setminus \{y'_i\}$ (resp. $v \in L \setminus \{z'_j\}$) must traverse the sole edge incident to y'_i (resp. z'_i), namely (y'_i, y_i) (resp. z'_j, z_j). Since $|L| \geq 2$, such a v always exists, and all edges incident to L are monitored by L.

The only remaining edges are those with both endpoints in Y. Let $(y_i, y_{i'})$ be such an edge. Since the only shortest path between y'_i and $y'_{i'}$ in G is $\langle y'_i, y_i, y_{i'}, y'_{i'} \rangle$, the pair $\{y'_i, y'_{i'}\}$ monitors $(y_i, y_{i'})$. ☐

Lemma 4. *Let $\mathcal{S}_1, \ldots, \mathcal{S}_k$ be k (not necessarily distinct) set covers of $\mathcal{I}$. The set $M = L \cup \{S_{j,\ell} \mid S_j \in \mathcal{S}_\ell, 1 \leq \ell \leq k\}$ is a MEG-set of G.*

Proof. Since $L \subseteq M$, by Lemma 3, we only need to argue about edges with at least one endvertex in some H_ℓ, with $1 \leq \ell \leq k$. Let $S_j \in \mathcal{S}_\ell$, and consider any $x_i \in S_j$. Edge $(S_{j,\ell}, z_j)$ is monitored by $\{z'_j, S_{j,\ell}\}$. Edges $(S_{j,\ell}, x_{i,\ell})$ and $(x_{i,\ell}, y_i)$ are monitored by $\{S_{j,\ell}, y'_i\}$.

The only remaining edges with at least one endvertex in H_ℓ are those incident to vertices $S_{j,\ell}$ with $S_j \in \mathcal{S} \setminus \mathcal{S}_\ell$. Consider any such S_j, let x_i be an item in S_j, and let $S_k \in \mathcal{S}_\ell$ be any set such that $x_i \in S_k$ (notice that both x_i and S_k exist since sets are non-empty and $\mathcal{S}_\ell$ is a set cover). Edge $(S_{j,\ell}, x_{i,\ell})$ is monitored by $\{S_{k,\ell}, z'_j\}$, which also monitors $(S_{j,\ell}, z_j)$. ☐

We say that a MEG-set M is *minimal* if, for every $v \in M$, $M \setminus \{v\}$ is not a MEG-set. Lemma 2 ensures that any minimal MEG-set M does not contain any of the vertices y_i, for $i = 1, \ldots, \eta$, or z_j for $j = 1, \ldots, h$. Hence, $M \setminus L$ contains only vertices in $\bigcup_{\ell=1}^{k} V(H_\ell)$.

Lemma 5. *Let M be a minimal MEG-set of G. For every $i = 1, \ldots, \eta$ and every $\ell = 1, \ldots, k$, M contains at least one among $x_{i,\ell}$ and all $S_{j,\ell}$ such that $x_i \in S_j$.*

The proof of Lemma 5 is given in the full version of this work [8].

Lemma 6. *Given a MEG-set M' of G, we can compute in polynomial time a MEG-set M of G such that $|M| \leq |M'|$ and, for every $\ell = 1, \ldots, k$, the set $\mathcal{S}_\ell = \{S_j \in \mathcal{S} \mid S_{j,\ell} \in M\}$ is a set cover of $\mathcal{I}$.*

Proof. Let M'' be a minimal MEG-set of G that is obtained from M' by possibly discarding some of the vertices. Clearly $|M''| \leq |M'|$ and M'' can be computed in polynomial time. Moreover, by Lemma 5, for every $i = 1, \ldots, \eta$ and every $\ell = 1, \ldots, k$, M'' contains $x_{i,\ell}$ or some $S_{j,\ell}$ such that S_j covers x_i. We compute M from M'' by replacing each $x_{i,\ell} \in M''$ with $S_{j,\ell}$, where $S_j \in \mathcal{S}$ is any set that covers x_i. As a consequence, for every $\ell = 1, \ldots, k$, the set $\mathcal{S}_\ell = \{S_j \in \mathcal{S} \mid S_{j,\ell} \in M\}$ is a set cover of $\mathcal{I}$. Moreover, since M'' contains all vertices in L by Lemma 1, so does M. Then, Lemma 4 implies that M is a MEG-set of G. ☐

Lemma 7. *Let $\varepsilon > 0$ be a constant of choice. Any polynomial-time $(\alpha \ln n)$-approximation algorithm for the minimum MEG-set problem, where $\alpha > 0$ is a constant, implies the existence of a polynomial-time $((2\alpha + \varepsilon) \ln N)$-approximation algorithm for Set Cover.*

Proof. Given an instance $\mathcal{I} = \langle X, \mathcal{S} \rangle$ of Set Cover and let h^* be the size of an optimal set cover of $\mathcal{I}$. In the rest of the proof we assume w.l.o.g. that $N \geq 4$ and $h^* \geq \frac{4\alpha}{\varepsilon}$. Indeed, if any of the above two conditions does not hold, we can solve $\mathcal{I}$ in constant time.

We now construct the graph G with $n = (k+2)N \leq N^2$ vertices by making $k = N-2$ copies of H. Next, we run the $(\alpha \ln n)$-approximation algorithm to compute a MEG-set M' of G, and we use Lemma 6 to find a MEG-set M with $|M| \leq |M'|$ that contains k set covers $\mathcal{S}_1, \ldots, \mathcal{S}_k$ in polynomial time. Among these k set covers, we output one $\mathcal{S}'$ of minimum size.

To analyze the approximation ratio of the above algorithm, let M^* be an optimal MEG-set of G. Lemma 4 ensures that $|M^*| \leq |L| + kh^* = N + kh^*$, and hence

$$|M| \leq |M'| \leq \alpha(N + kh^*) \ln n = \alpha(N + kh^*) \ln N^2 = 2\alpha k h^* \ln N + 2\alpha N \ln N.$$

Therefore we have:

$$\begin{aligned} |\mathcal{S}'| &\leq \frac{|M|}{k} \leq 2\alpha h^* \ln N + \frac{2\alpha N \ln N}{k} \\ &\leq 2\alpha h^* \ln N + 4\alpha \ln N = \left(2\alpha + \frac{4\alpha}{h^*}\right) h^* \ln N \leq (2\alpha + \varepsilon) h^* \ln N. \end{aligned} \qquad \square$$

Let γ be any positive constant. Since Set Cover cannot be approximated in polynomial time within a factor of $(1-\gamma)\ln|\mathcal{I}|$, unless P = NP [7], and since an invocation of Lemma 7 with $\alpha = \frac{1}{2} - \gamma$ and $\varepsilon = \gamma$ shows that any polynomial-time $((\frac{1}{2}-\gamma)\ln n)$-approximation algorithm for the minimum MEG-set problem can be turned into a polynomial-time approximation algorithm for Set Cover with an approximation ratio of $(1-\gamma)\ln N \leq (1-\gamma)\ln|\mathcal{I}|$, we have:

Theorem 1. *The minimum MEG-set problem cannot be approximated in polynomial time within a factor of $c \ln n$, for any constant $c < \frac{1}{2}$, unless* P = NP.

4. Conclusion

In this work we present inapproximability results on the minimum MEG-set problem, proving that the problem is APX-hard and not approximable within a factor of $c \ln n$, for any constant $c < \frac{1}{2}$. Observe that it is not hard to devise an efficient approximation algorithm for the minimum MEG-set achieving an approximation ratio of $O(\sqrt{n \cdot \ln n})$ based on the well-known approximation algorithm for the Set Cover [9]. It is an open problem to narrow the gap between the lower and upper bounds on approximability of minimum MEG-set. Moreover, it could be interesting to study the approximability of the problem on specific classes of graphs.

References

[1] F. Foucaud, K. Narayanan, L. R. Sulochana, Monitoring edge-geodetic sets in graphs, in: A. Bagchi, R. Muthu (Eds.), Algorithms and Discrete Applied Mathematics - 9th International Conference, CALDAM 2023, Gandhinagar, India, February 9-11, 2023, Proceedings, volume 13947 of *Lecture Notes in Computer Science*, Springer, 2023, pp. 245–256. URL: https://doi.org/10.1007/978-3-031-25211-2_19. doi:10.1007/978-3-031-25211-2_19.

[2] J. Haslegrave, Monitoring edge-geodetic sets: Hardness and graph products, Discret. Appl. Math. 340 (2023) 79–84. URL: https://doi.org/10.1016/j.dam.2023.06.033. doi:10.1016/j.dam.2023.06.033.

[3] F. Foucaud, P. Marcille, Z. M. Myint, R. B. Sandeep, S. Sen, S. Taruni, Monitoring edge-geodetic sets in graphs: Extremal graphs, bounds, complexity, in: S. Kalyanasundaram, A. Maheshwari (Eds.), Algorithms and Discrete Applied Mathematics - 10th International Conference, CALDAM 2024, Bhilai, India, February 15-17, 2024, Proceedings, volume 14508 of *Lecture Notes in Computer Science*, Springer, 2024, pp. 29–43. URL: https://doi.org/10.1007/978-3-031-52213-0_3. doi:10.1007/978-3-031-52213-0_3.

[4] A. Tan, W. Li, X. Wang, X. Li, Monitoring edge-geodetic numbers of convex polytopes and four networks, Int. J. Parallel Emergent Distributed Syst. 38 (2023) 301–312. URL: https://doi.org/10.1080/17445760.2023.2220143. doi:10.1080/17445760.2023.2220143.

[5] R. Ma, Z. Ji, Y. Yao, Y. Lei, Monitoring-edge-geodetic numbers of radix triangular mesh and sierpiński graphs, Int. J. Parallel Emergent Distributed Syst. 39 (2024) 353–361. URL: https://doi.org/10.1080/17445760.2023.2294369. doi:10.1080/17445760.2023.2294369.

[6] X. Xu, C. Yang, G. Bao, A. Zhang, X. Shao, Monitoring-edge-geodetic sets in product networks, Int. J. Parallel Emergent Distributed Syst. 39 (2024) 264–277. URL: https://doi.org/10.1080/17445760.2024.2301929. doi:10.1080/17445760.2024.2301929.

[7] I. Dinur, D. Steurer, Analytical approach to parallel repetition, in: D. B. Shmoys (Ed.), Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014, ACM, 2014, pp. 624–633. URL: https://doi.org/10.1145/2591796.2591884. doi:10.1145/2591796.2591884.

[8] D. Bilò, G. Colli, L. Forlizzi, S. Leucci, On the inapproximability of finding minimum monitoring edge-geodetic sets, 2024. arXiv:2405.13875.

[9] G. Colli, Monitoring graph edges via shortest paths: computational complexity and approximation algorithms, Tesi di laurea, University of L'Aquila, 2023.