



Monitoring the Edges of a Graph Using Distances

Florent Foucaud^{1(✉)}, Ralf Klasing¹, Mirka Miller^{2,3}, and Joe Ryan⁴

¹ Univ. Bordeaux, Bordeaux INP, CNRS, LaBRI, UMR5800, Talence, France
florent.foucaud@gmail.com, ralf.klasing@labri.fr

² School of Mathematical and Physical Sciences, The University of Newcastle,
Callaghan, Australia

³ Department of Mathematics, University of West Bohemia, Pilsen, Czech Republic

⁴ School of Electrical Engineering and Computing, The University of Newcastle,
Callaghan, Australia
joe.ryan@newcastle.edu.au

Abstract. We introduce a new graph-theoretic concept in the area of network monitoring. A set M of vertices of a graph G is a *distance-edge-monitoring set* if for every edge e of G , there is a vertex x of M and a vertex y of G such that e belongs to all shortest paths between x and y . We denote by $\text{dem}(G)$ the smallest size of such a set in G . The vertices of M represent distance probes in a network modeled by G ; when the edge e fails, the distance from x to y increases, and thus we are able to detect the failure. It turns out that not only we can detect it, but we can even correctly locate the failing edge.

In this paper, we initiate the study of this new concept. We show that for a nontrivial connected graph G of order n , $1 \leq \text{dem}(G) \leq n - 1$ with $\text{dem}(G) = 1$ if and only if G is a tree, and $\text{dem}(G) = n - 1$ if and only if it is a complete graph. We compute the exact value of dem for grids, hypercubes, and complete bipartite graphs.

Then, we relate dem to other standard graph parameters. We show that $\text{dem}(G)$ is lower-bounded by the arboricity of the graph, and upper-bounded by its vertex cover number. It is also upper-bounded by five times its feedback edge set number.

Then, we show that determining $\text{dem}(G)$ for an input graph G is an NP-complete problem, even for apex graphs. There exists a polynomial-time logarithmic-factor approximation algorithm, however it is NP-hard to compute an asymptotically better approximation, even for bipartite graphs of small diameter and for bipartite subcubic graphs. For such instances, the problem is also unlikely to be fixed parameter tractable when parameterized by the solution size.

This paper is dedicated to the memory of Mirka Miller, who sadly passed away in January 2016. This research was started when Mirka Miller and Joe Ryan visited Ralf Klasing at the LaBRI, University of Bordeaux, in April 2015.

© Springer Nature Switzerland AG 2020

M. Changat and S. Das (Eds.): CALDAM 2020, LNCS 12016, pp. 28–40, 2020.

https://doi.org/10.1007/978-3-030-39219-2_3

1 Introduction

The aim of this paper is to introduce a new concept of network monitoring using distance probes, called *distance-edge-monitoring*. Our networks are naturally modeled by finite undirected simple connected graphs, whose vertices represent computers and whose edges represent connections between them. We wish to be able to monitor the network in the sense that when a connection (an edge) fails, we can detect this failure. We will select a (hopefully) small set of vertices of the network, that will be called *probes*. At any given moment, a probe of the network can measure its graph distance to any other vertex of the network. Our goal is that, whenever some edge of the network fails, one of the measured distances changes, and thus the probes are able to detect the failure of any edge.

Probes that measure distances in graphs are present in real-life networks, for instance this is useful in the fundamental task of *routing* [6,11]. They are also frequently used for problems concerning *network verification* [2–4].

We will now proceed with the formal definition of our main concept. In this paper, by *graphs* we refer to connected simple graphs (without multiple edges and loops). A graph with loops or multiple edges is called a *multigraph*.

We denote by $d_G(x, y)$ the distance between two vertices x and y in a graph G . When there is no path connecting x and y in G , we let $d_G(x, y) = \infty$. For an edge e of G , we denote by $G - e$ the graph obtained by deleting e from G .

Definition 1. *For a set M of vertices and an edge e of a graph G , let $P(M, e)$ be the set of pairs (x, y) with x a vertex of M and y a vertex of $V(G)$ such that $d_G(x, y) \neq d_{G-e}(x, y)$. In other words, e belongs to all shortest paths between x and y in G .*

For a vertex x , let $M(x)$ be the set of edges e such that there exists a vertex v in G with $(x, v) \in P(\{x\}, e)$. If $e \in M(x)$, we say that e is monitored by x .

A set M of vertices of a graph G is distance-edge-monitoring if every edge e of G is monitored by some vertex of M , that is, the set $P(M, e)$ is nonempty. Equivalently, $\bigcup_{x \in M} M(x) = E(G)$.

We note $\text{dem}(G)$ the smallest size of a distance-edge-monitoring set of G .

Note that $V(G)$ is always a distance-edge-monitoring set of G , so $\text{dem}(G)$ is always well-defined.

Consider a graph G modeling a network, and a set M of vertices of G , on which we place probes that are able to measure their distances to all the other vertices. If M is distance-edge-monitoring, if a failure occurs on any edge of the network (in the sense that the communication between its two endpoints is broken), then this failure is detected by the probes.

In fact, it turns out that not only the probes can *detect* a failing edge, but they can also precisely *locate* it (the proof of the following is delayed to Sect. 2).

Proposition 2. *Let M be a distance-edge-monitoring set of a graph G . Then, for any two distinct edges e and e' in G , we have $P(M, e) \neq P(M, e')$.*

Thus, assume that we have placed probes on a distance-edge-monitoring set M of a network G and initially computed all the sets $P(M, e)$. In the case a unique edge of the network has failed, Proposition 2 shows that by measuring the set of pairs (x, y) with $x \in M$ and $y \in V(G)$ whose distance has changed, we know exactly which is the edge that has failed.

We define the decision and optimization problem associated to distance-edge-monitoring sets.

<p>DISTANCE-EDGE-MONITORING SET</p>

<p>Instance: A graph G, an integer k.</p>

<p>Question: Do we have $\text{dem}(G) \leq k$?</p>
--

<p>MIN DISTANCE-EDGE-MONITORING SET</p>

<p>Instance: A graph G.</p>
--

<p>Task: Build a smallest possible distance-edge-monitoring set of G.</p>
--

Related Notions. A weaker model is studied in [2, 3] as a *network discovery* problem, where we seek a set S of vertices such that for each edge e , there exists a vertex x of S and a vertex y of G such that e belongs to *some* shortest path from x to y .

Distance-edge-monitoring sets are related to *resolving sets* and *edge-resolving sets*, that also model sets of sensors that can measure the distance to all other vertices in a graph. A resolving set is a set R of vertices such that for any two distinct vertices x and y in G , there is a vertex r in R such that $d_G(r, x) \neq d_G(r, y)$. The smallest size of a resolving set in G is the *metric dimension* of G [12, 20]. If instead, the set R distinguishes the edges of G (that is, for any pair e, e' of edges of G , there is a vertex $r \in R$ with $d_G(x, e) \neq d_G(x, e')$, where for $e = uv$, $d_G(x, e) = \min\{d_G(x, u), d_G(x, v)\}$), we have an edge-resolving set [15].

Another related concept is the one of *strong resolving sets* [18, 19]: a set R of vertices is strongly resolving if for any pair x, y of vertices, there exists a vertex z of R such that either x is on a shortest path from z to y , or y is on a shortest path from z to x . It is related to distance-edge-monitoring sets in the following sense. Given a distance-edge-monitoring set M , for every pair x, y of *adjacent* vertices, there is a vertex z of M such that either x is on *every* shortest path from z to y , or y is on *every* shortest path from z to x .

Another related concept is the one of an (*strong*) *edge-geodetic set*. A set S of vertices is an edge-geodetic set if for every edge e of G , there are two vertices x, y of S such that e is on some shortest path from x to y . It is a strong edge-geodetic set if to every pair x, y of S , we can assign a shortest $x - y$ path to $\{x, y\}$ such that every edge of G belongs to one of these $\binom{|S|}{2}$ assigned shortest paths [16].

Our Results. We first derive a number of basic results about distance-edge-monitoring sets in Sect. 2 (where we also give some useful definitions).

In Sect. 3, we study dem for some basic graph families (like trees and grids) and relate this parameter to other standard graph parameters such as arboricity arb , vertex cover number vc and feedback edge set number fes . We show that

$\text{dem}(G) = 1$ if and only if G is a tree. We show that for any graph G of order n , $\text{dem}(G) \geq \text{arb}(G)$. Moreover, $\text{dem}(G) \leq \text{vc}(G) \leq n - 1$ (with equality if and only if G is complete). We show that for some families of graphs G , $\text{dem}(G) = \text{vc}(G)$, for instance this is the case for complete bipartite graphs and hypercubes. Then we show that $\text{dem}(G) \leq 5\text{fes}(G) - 5$ when $\text{fes}(G) \geq 2$ (when $\text{fes}(G) \leq 1$, $\text{dem}(G) = \text{fes}(G) + 1$).

In Sect. 4, we show that DISTANCE-EDGE-MONITORING SET is NP-complete, even for apex graphs (graphs obtained from a planar graph by adding an extra vertex). Then, we show that MIN DISTANCE-EDGE-MONITORING SET can be approximated in polynomial time within a factor of $\ln(|E(G)| + 1)$ by a reduction to the set cover problem. Finally, we show that no essentially better ratio can be obtained (unless $P = NP$), even for graphs that are of diameter 4, bipartite and of diameter 6, or bipartite and of maximum degree 3. For the same restrictions, the problem is unlikely to be fixed parameter tractable when parameterized by the solution size. These hardness results are obtained by reductions from the set cover problem.

We conclude our paper in Sect. 5. Due to the given space constraints, some proofs are not included.

2 Preliminaries

We now give some useful lemmas about basic properties of distance-edge-monitoring sets. We start with the proof of Proposition 2.

Proof (Proof of Proposition 2). Suppose by contradiction that there are two distinct edges e and e' with $P(M, e) = P(M, e')$. Since by definition, $P(M, e) \neq \emptyset$, we have $(x, v) \in P(M, e)$ (and thus $(x, v) \in P(M, e')$), where $x \in M$ and $v \in V(G)$. Thus, all shortest paths between x and v contain both e and e' . Assume that on these shortest paths, e is closer to x than e' , and let w be the endpoint of e that is closest to v . Then, we have $(x, w) \in P(M, e)$ but $(x, w) \notin P(M, e')$: a contradiction. \square

An edge e in a graph G is a *bridge* if $G - e$ has more connected components than G . We will now show that bridges are very easy to monitor.

Lemma 3. *Let G be a connected graph and let e be a bridge of G . For any vertex x of G , we have $e \in M(x)$.*

Proof. Assume that $e = uv$. Since e is a bridge, we have $d_G(x, u) \neq d_G(x, v)$. If $d_G(x, u) < d_G(x, v)$, then $(x, v) \in P(\{x\}, e)$. Otherwise, we have $(x, u) \in P(\{x\}, e)$. In both cases, $e \in M(x)$. \square

Given a vertex x of a graph G and an integer i , we let $L_i(x)$ denote the set of vertices at distance i of x in G .

Lemma 4. *Let x be a vertex of a connected graph G . Then, an edge uv belongs to $M(x)$ if and only if $u \in L_i(x)$ and v is the only neighbour of u in $L_{i-1}(x)$.*

Proof. Let $uv \in M(x)$. Then, there exists a vertex y such that all shortest paths from y to x go through uv . Thus, one of u and v (say, u) is in a set $L_i(x)$ and the other (v) is in a set $L_{i-1}(x)$, for some positive integer i . Moreover, v must be the only neighbour of u in $L_{i-1}(x)$, since otherwise there would be a shortest path from y to x going through u but avoiding uv .

Conversely, if u is a vertex in $L_i(x)$ with a unique neighbour v in $L_{i-1}(x)$, then the edge uv belongs to $M(x)$ since all shortest paths from u to x use it. \square

We obtain some immediate consequences of Lemma 4.

Lemma 5. *For a vertex x of a graph G , the set of edges $M(x)$ induces a forest.*

Proof. Let G_x be the subgraph of G induced by the edges in $M(x)$. By Lemma 4, an edge e belongs to $M(x)$ if and only if $e = uv$, $u \in L_i(x)$ and v is the only neighbour of u in $L_{i-1}(x)$. In this case, we let v be the *parent* of u . Each vertex in G_x has at most one parent in G_x , and each edge of G_x is an edge between a vertex and its parent. Thus, G_x is a forest. \square

Lemma 6. *Let G be a graph and x a vertex of G . Then, for any edge e incident with x , we have $e \in M(x)$.*

Proof. For every vertex y of $L_1(x)$ (that is, every neighbour of x), x is the unique neighbour of y in $L_0(x) = \{x\}$. Thus, the claim follows from Lemma 4. \square

3 Basic Graph Families and Bounds

In this section, we study dem for standard graph classes, and its relation with other standard graph parameters.

Theorem 7. *Let G be a connected graph with at least one edge. We have $\text{dem}(G) = 1$ if and only if G is a tree.*

Proof. In a tree, every edge is a bridge. Thus, by Lemma 3, if G is a tree, any vertex x of G is a distance-edge-monitoring set and we have $\text{dem}(G) \leq 1$ (and of course as long as there is an edge in G , $\text{dem}(G) \geq 1$).

For the converse, suppose that $\text{dem}(G) = 1$. Then, clearly, G must have at least one edge. Moreover, since all edges of G must belong to $M(x)$, by Lemma 5, G must be a forest. Since G is connected, G is a tree. \square

Let $G_{a,b}$ denote the grid of dimension $a \times b$. By Theorem 7, we have $\text{dem}(G_{1,a}) = 1$. We can compute all other values.

Theorem 8. *For any integers $a, b \geq 2$, we have $\text{dem}(G_{a,b}) = \max\{a, b\}$.*

The *arboricity* $\text{arb}(G)$ of a graph G is the smallest number of sets into which $E(G)$ can be partitioned and such that each set induces a forest. The *clique number* $\omega(G)$ of G is the size of a largest clique in G .

Theorem 9. *For any graph G of order n and size m , we have $\text{dem}(G) \geq \text{arb}(G)$, and thus $\text{dem}(G) \geq \frac{m}{n-1}$ and $\text{dem}(G) \geq \frac{\omega(G)}{2}$.*

Proof. By Lemma 5, for each vertex x of a distance-edge-monitoring set M , $M(x)$ induces a forest. Thus, to each edge e of G we can assign one of the forests $M(x)$ such that $e \in M(x)$. This is a partition of G into $|M|$ forests, and thus $\text{arb}(G) \leq \text{dem}(G)$.

Moreover, it is not difficult to see that $\text{arb}(G) \geq \frac{m}{n-1}$ (since a forest has at most $n - 1$ edges) and $\text{arb}(G) \geq \frac{\omega(G)}{2}$ (since a clique of size $k = \omega(G)$ has $k(k - 1)/2$ edges but a forest that is a subgraph of G can contain at most $k - 1$ of these edges). \square

We next see that distance-edge-monitoring sets are relaxations of vertex covers. A set C of vertices is a *vertex cover* of G if every edge of G has one of its endpoints in C . The smallest size of a vertex cover of G is denoted by $\text{vc}(G)$.

Theorem 10. *In any graph G of order n , any vertex cover of G is a distance-edge-monitoring set, and thus $\text{dem}(G) \leq \text{vc}(G) \leq n - 1$. Moreover, we have $\text{dem}(G) = n - 1$ if and only if G is the complete graph of order n .*

Proof. Let C be a vertex cover of G . By Lemma 6, for every edge e , there is a vertex in x with $e \in M(x)$, thus C is distance-edge-monitoring.

Moreover, any graph G of order n has a vertex cover of size $n - 1$: for any vertex x , the set $V(G) \setminus \{x\}$ is a vertex cover of G .

Finally, suppose that $\text{dem}(G) = n - 1$: then also $\text{vc}(G) = n - 1$. If G is not connected, we have $\text{vc}(G) \leq n - 2$ (starting with $V(G)$ and removing any vertex from each connected component of G yields a vertex cover), thus G is connected. Suppose by contradiction that G is not a complete graph. Then, we have three vertices x , y and z in G such that xy and yz are edges of G , but xz is not. Then, $V(G) \setminus \{x, z\}$ is a vertex cover of G , a contradiction.

This completes the proof. \square

In some graphs, any distance-edge-monitoring set is a vertex cover.

Observation 11. *If, for every vertex x of a graph G , $M(x)$ consists exactly of the sets of edges incident with x , then a set M is a distance-edge-monitoring set of G if and only if it is a vertex cover of G .*

Note that Observation 11 does not provide a characterization of graphs with $\text{dem}(G) = \text{vc}(G)$. For example, as seen in Theorem 8, for the grid $G_{a,2}$, we have $\text{dem}(G_{a,2}) = \text{vc}(G_{a,2}) = a$, but for any vertex x of $G_{a,2}$, $M(x)$ consists of the whole row and column of $G_{a,2}$.

Let $K_{a,b}$ be the complete bipartite graph with parts of sizes a and b , and let H_d denote the hypercube of dimension d . We can deduce the following from Observation 11.

Corollary 12. *We have $\text{dem}(K_{a,b}) = \text{vc}(K_{a,b}) = \min\{a, b\}$, and $\text{dem}(H_d) = \text{vc}(H_d) = 2^{d-1}$.*

A *feedback edge set* of a graph G is a set of edges such that removing them from G leaves a forest. The smallest size of a feedback edge set of G is denoted by $\text{fes}(G)$ (it is sometimes called the *cyclomatic number* of G).

Theorem 7 states that for any tree G (that is, a graph with feedback edge set number 0), we have $\text{dem}(G) = \text{fes}(G) + 1$. We now show the same bound for graphs G with $\text{fes}(G) = 1$, that is, unicyclic graphs (graphs with a unique cycle).

Proposition 13. *If G is a unicyclic graph, then $\text{dem}(G) = 2$.*

We will now give a weaker (but similar) bound for any value of $\text{fes}(G)$. But first, we need the following terminology and lemma from [9] about the structure of graphs G with given feedback edge set number.

In a graph, a vertex is a *core vertex* if it has degree at least 3. The *base graph* of a graph G is the graph obtained from G by iteratively removing vertices of degree 1 (thus, the base graph of a forest is the empty graph).

The following lemma can be found in Sect. 5.3.1 of [9].

Lemma 14 ([9]). *Let G be a graph with $\text{fes}(G) = k \geq 2$. The base graph of G has at most $2k - 2$ core vertices, that are joined by at most $3k - 3$ edge-disjoint paths with internal vertices of degree 2 and whose endpoints are (not necessarily distinct) core vertices.*

In other words, Lemma 14 says that the base graph of a graph G with $\text{fes}(G) = k \geq 2$ can be obtained from a multigraph H of order at most $2k - 2$ and size $3k - 3$ by subdividing its edges an arbitrary number of times.

Theorem 15. *Let G be a graph with $\text{fes}(G) \geq 2$. Then, $\text{dem}(G) \leq 5\text{fes}(G) - 5$.*

Proof. Let $\text{fes}(G) = k$ and G_b , the base graph of G . By Lemma 14, G_b contains at most $2k - 2$ core vertices, joined by at most $3k - 3$ edge-disjoint paths with internal vertices of degree 2 and whose endpoints are core vertices (both endpoints can be the same vertex). Let $\{P_1, \dots, P_t\}$ be the set of these paths.

By Lemma 3, any non-empty distance-edge-monitoring set of G_b is also one of G , since all edges of G not present in G_b are bridges. Thus, it is sufficient to construct a distance-edge-monitoring set M of G_b of size at most $5k - 5$. We build M as follows: first of all, M contains all core vertices of G_b . Moreover, for each path P_i of length at least 2 connecting two core vertices x and y of G_b , we add to M an internal vertex z of P_i . If $x \neq y$ or P_i has length at most 3, we let z be a neighbour of, say, x . Otherwise, $x = y$ and P_i has length at least 4; we choose z as a vertex at distance 2 of x on P_i . Thus, $|M| \leq 5k - 5$.

Let us prove that M is distance-edge-monitoring in G_b . Let e be an edge of G_b ; necessarily, e belongs to some path P_i connecting two core vertices x and y . If e is incident with x or y , by Lemma 6 e is monitored by M . Thus, suppose this is not the case. So, P_i has length at least 3 and so there is an internal vertex z of P_i in M . Again if e is incident with z we are done.

If $x = y$, P_i forms a cycle and for any vertex v , $M(v) \cap E(P_i)$ consists of the two edges (if P_i has even length) or the unique edge (if P_i has odd length)

that are opposite to v on the cycle. Since x and z are non-adjacent, we have $e \in M(x) \cup M(z)$ and so z is monitored.

If $x \neq y$, assume that $e = uv$ with u the vertex of e closest to x . Let d_x be the distance from x to u on P_i , and d_y the distance from y to v on P_i . If $e \in M(x) \cup M(y)$, we are done. Otherwise, there must exist a path from x to y that is edge-disjoint with P_i (otherwise e would be a bridge and would be monitored by x or y). Let p be the length of a shortest such path. Since $e \notin M(x)$, there exists a shortest path from v to x avoiding e . This shortest path must go through y , and then use a path of length p . Thus, we have $d_y + p \leq d_x + 1$. By a similar argument with u and x , we have $d_x + p \leq d_y + 1$. We obtain from this that $d_x = d_y$, $|P_i|$ is odd, and $p = 1$. But then, the unique shortest path from v to z goes through e , and we have $e \in M(z)$. This completes the proof. \square

We do not believe that the bound of Theorem 15 is tight. For the sake of simplicity, we have not tried to optimize the bound of Theorem 15, as we do not believe that this method will provide a tight bound. There are examples of graphs G where $\text{dem}(G) = \text{fes}(G) + 1$, for example this is the case for the grid $G_{a,2}$: by Theorem 8, when $a \geq 2$ we have $\text{dem}(G_{a,2}) = a$, and $\text{fes}(G_{a,2}) = a - 1$.

Note that $G_{a,b}$ for $a, b \geq 2$ provides examples of a family of graphs where for increasing a, b the difference between $\text{fes}(G_{a,b})$ and $\text{dem}(G_{a,b})$ is unbounded by a constant. Indeed by Theorem 8 we have $\text{dem}(G_{a,b}) = \max\{a, b\}$ but $\text{fes}(G_{a,b})$ is linear in the order ab .

4 Complexity

A c -approximation algorithm for a given optimization problem is an algorithm that returns a solution whose size is always at most c times the optimum. We refer to the books [1, 13] for more details. For a decision problem Π and for some parameter p of the instance, an algorithm for Π is said to be *fixed parameter tractable* (fpt for short) if it runs in time $f(p)n^c$, where f is a computable function, n is the input size, and c is a constant. In this paper, we will always consider the solution size k as the parameter. The class FPT contains the parameterized decision problems solvable by an fpt algorithm. The classes $\text{W}[i]$ (with $i \geq 1$) denote complexity classes with parameterized decision problems that are believed not to be fpt. We refer to the books [8, 17] for more details.

We will now use the connection to vertex covers hinted in Sect. 3 to derive an NP-hardness result. For two graphs G, H , $G \bowtie H$ denotes the graph obtained from disjoint copies of G and H with all possible edges between $V(G)$ and $V(H)$. We denote by K_1 the graph on one single vertex.

Theorem 16. *For any graph G , we have $\text{vc}(G) \leq \text{dem}(G \bowtie K_1) \leq \text{vc}(G) + 1$. Moreover, if G has radius at least 4, then $\text{vc}(G) = \text{dem}(G \bowtie K_1)$.*

We deduce the following from the second part of the statement of Theorem 16 and the NP-hardness of VERTEX COVER for planar subcubic graphs with radius at least 4 [10].

Corollary 17. DISTANCE-EDGE-MONITORING SET is NP-complete, even for graphs obtained from a planar subcubic graph by attaching a universal vertex.

Given a hypergraph $H = (X, S)$ with vertex set X and edge set S , a set cover of H is a subset $C \subseteq S$ of edges such that each vertex of X belongs to at least one edge of C . We will provide reductions for MIN DISTANCE-EDGE-MONITORING SET to and from SET COVER and MIN SET COVER.

SET COVER

Instance: A hypergraph $H = (X, S)$, an integer k .

Question: Is there a set cover $C \subseteq S$ of size at most k ?

MIN SET COVER

Instance: A hypergraph $H = (X, S)$.

Task: Build a smallest possible set cover of H .

It is known that MIN SET COVER is approximable in polynomial time within a factor of $\ln(|X|+1)$ [14], but, unless $P = NP$, not within a factor of $(1-\epsilon) \ln |X|$ (for every positive ϵ) [7]. Moreover, SET COVER is $W[2]$ -hard (when parameterized by k) and not solvable in time $|X|^{o(k)}|H|^{O(1)}$ unless $FPT = W[1]$ [5].

Theorem 18. MIN DISTANCE-EDGE-MONITORING SET is approximable within a factor of $\ln(|E(G)|+1)$ in polynomial time.

Proof. Let G be a graph and consider the following hypergraph $H = (X, S)$ with $X = E(G)$ and such that S contains, for each vertex x of G , the set $S_x = \{e \in X \mid e \in M(x)\}$. Now, it is not difficult to see that there is a one-to-one correspondence between set covers of H and distance-edge-monitoring sets of G , where we associate to each vertex x of a distance-edge-monitoring set of G , the set S_x in a set cover of H . Thus, the result follows from the $\ln(|X|+1)$ -approximation algorithm for MIN SET COVER from [14]. \square

Theorem 19. Even for graphs G that are (a) of diameter 4, (b) bipartite and of diameter 6, or (c) bipartite and of maximum degree 3, MIN DISTANCE-EDGE-MONITORING SET is not approximable within a factor of $(1-\epsilon) \ln |E(G)|$ in polynomial time, unless $P = NP$. Moreover, for such instances, DISTANCE-EDGE-MONITORING SET cannot be solved in time $|G|^{o(k)}$, unless $FPT = W[1]$, and it is $W[2]$ -hard for parameter k .

Proof. For an instance (H, k) of SET COVER, we will construct in polynomial time instances $(G, k+2)$ or $(G, k+1)$ of DISTANCE-EDGE-MONITORING SET so that H has a set cover of size k if and only if G has a distance-edge-monitoring set of size at most $k+2$ or $k+1$.

In our first reduction, the obtained instance has diameter 4, while in our second reduction, the obtained instance is bipartite and has diameter 6, and in our third reduction, the graph G is bipartite and has maximum degree 3. The three constructions are similar.

The statement will follow from the hardness of approximating MIN SET COVER proved in [7], the parameterized hardness of SET COVER (parameterized by solution size), and the lower bound on its running time [5].

First of all, we point out that we may assume that in an instance $(H = (X, S), k)$ of SET COVER, there is no vertex of X that belongs to a unique set of S . Indeed, otherwise, we are forced to take S in any set cover of H ; thus, by removing S and all vertices in S , we obtain an equivalent instance $(H', k - 1)$. We can iterate until the instance satisfies this property.

We now describe the first reduction, in which the obtained instance has diameter 4. Let $(H, k) = ((X, S), k)$ be an instance of SET COVER, where $X = \{x_1, x_2, \dots, x_{|X|}\}$, $S = \{C_1, C_2, \dots, C_{|S|}\}$ and $C_i = \{c_{i,j} \mid x_j \in C_i\}$. Construct the following instance $(G, k + 2) = ((V, E), k + 2)$ of DISTANCE-EDGE-MONITORING SET, where $V = V_1 \cup V_2 \cup \dots \cup V_5 \cup V'_1 \cup V'_2 \cup V'_3 \cup V'_4$, $E = E_1 \cup E_2 \cup E_3 \cup E_4 \cup E'_1 \cup E'_2 \cup E'_3 \cup E'_4 \cup E'_5$ and

$$\begin{aligned} V_1 &= \{u_1, u_2, u_3\}, & V_2 &= \{v_i \mid 1 \leq i \leq |S|\}, & V_3 &= S, & V_4 &= \{c_{i,j} \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}, & V_5 &= X \\ V'_1 &= \{u'_1, u'_2, u'_3\}, & V'_2 &= \{v'_i \mid 1 \leq i \leq |S|\}, & V'_3 &= \{c'_{i,j} \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}, & V'_4 &= \{w'_j \mid 1 \leq j \leq |X|\} \\ E_1 &= \{(u_1, u_2), (u_1, u_3), (u_2, u'_1), (u_3, u'_1)\}, & E_2 &= \{(u_1, v_i), (v_i, C_i) \mid 1 \leq i \leq |S|\}, & E_3 &= \{(C_i, c_{i,j}) \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}, & E_4 &= \{(c_{i,j}, x_j) \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\} \\ E'_1 &= \{(u'_1, u'_2), (u'_2, u'_3), (u'_3, u'_1)\}, & E'_2 &= \{(u'_1, v'_i), (v'_i, C_i) \mid 1 \leq i \leq |S|\}, & E'_3 &= \{(u'_1, c'_{i,j}), (c'_{i,j}, c_{i,j}) \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}, & E'_4 &= \{(u'_1, w'_j), (w'_j, x_j) \mid 1 \leq j \leq |X|\}, & E'_5 &= \{(u'_1, v_i) \mid 1 \leq i \leq |S|\}. \end{aligned}$$

An example is given in Fig. 1.

To see that G has diameter 4, observe that every vertex of G has a path of length at most 2 to the vertex u'_1 .

Let C be a set cover of H of size k . Define $M = C \cup \{u_1, u'_2\}$. Then, by Lemma 4, u_1 monitors (in particular) the edges (u_1, u_2) , (u_1, u_3) , (u'_1, u'_3) , and all the edges in $E_2 \cup E_3$. Similarly, u'_2 monitors the edges (u'_1, u'_2) , (u'_2, u'_3) , (u_3, u'_1) , (u_2, u'_1) and all the edges in $E'_2 \cup E'_3 \cup E'_4 \cup E'_5$. It thus remains to show that all edges of E_4 are monitored. Notice that among those edges, vertex C_i of S monitors exactly all edges c_{j_1, j_2} with $x_{j_2} \in C_i$. Thus, if $e = (c_{i,j}, x_j)$, there is i' such that $x_j \in C_{i'}$ and $C_{i'} \in C$, and e is monitored by $C_{i'}$ (either $i = i'$ and the only shortest path from x_j to C_i contains e , or $i \neq i'$ and the only shortest path from $c_{i,j}$ to $C_{i'}$ contains e). Hence, M is a distance-edge-monitoring set of G of size at most $k + 2$.

Conversely, let M be a distance-edge-monitoring set of G of size at most $k + 2$. In order to monitor the edge (u'_2, u'_3) , either $u'_2 \in M$ or $u'_3 \in M$. In order to monitor the edges (u_1, u_3) and (u_1, u_2) , there must be a vertex of M in $\{u_1, u_2, u_3\}$. We may replace $M \cap \{u_1, u_2, u_3\}$ by u_1 and $M \cap \{u'_2, u'_3\}$ by u'_2 , as u_1 monitors the same edges as $\{u_1, u_2, u_3\}$ (among those not already monitored by u'_2) and u'_2 monitors the same edges as $\{u'_2, u'_3\}$ (among those not already monitored by u_1). As seen in the previous paragraph, all edges of $E_1, E'_1, E_2, E_3, E'_2, E'_3, E'_4$ and E'_5 are monitored by $\{u_1, u'_2\}$. However, no edge of E_4 is

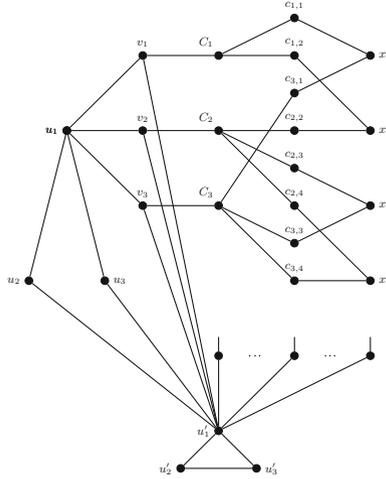


Fig. 1. First reduction from SET COVER to DISTANCE-EDGE-MONITORING SET from the proof of Theorem 19 applied to the hypergraph $(\{x_1, x_2, x_3, x_4\}, \{C_1 = \{x_1, x_2\}, C_2 = \{x_2, x_3, x_4\}, C_3 = \{x_1, x_3, x_4\}\})$. Vertices and edges of V'_i and E'_i for $i = 2, 3, 4$ are only suggested.

monitored by any vertex of $V_1 \cup V'_1$. Thus, all remaining vertices of M are needed precisely to monitor the edges of E_4 .

If $v_i \in M$, let $M = M \setminus \{v_i\} \cup \{C_i\}$, and the set of monitored edges does not decrease. If $c_{i,j} \in M$, let $M = M \setminus \{c_{i,j}\} \cup \{C_i\}$, and the set of monitored edges does not decrease. If $x_j \in M$, let $M = M \setminus \{x_j\} \cup \{C_{i_0}\}$ where $i_0 = \min\{i \mid x_j \in C_i\}$, and the set of monitored edges does not decrease. If $v'_i \in M$, let $M = M \setminus \{v'_i\} \cup \{C_i\}$, and the set of monitored edges does not decrease. If $c'_{i,j} \in M$, let $M = M \setminus \{c'_{i,j}\} \cup \{C_i\}$, and the set of monitored edges does not decrease. If $w'_j \in M$, let $M = M \setminus \{w'_j\} \cup \{C_{i_0}\}$ where $i_0 = \min\{i \mid x_j \in C_i\}$, and the set of monitored edges does not decrease. Iterating this process, we finally obtain a distance-edge-monitoring set M' of G with $|M' \cap V_1| = 1$, $|M' \cap V'_1| \geq 1$, $|M' \cap V_3| \leq k$, $M' \cap (V_2 \cup V_4 \cup V_5 \cup V'_2 \cup V'_3 \cup V'_4) = \emptyset$. Let $C = M' \cap V_3$. If C is not a set cover of H , then there is $x_j \in X$ that is not covered by C . Recall that we assumed that each vertex of X belongs to at least two edges of S . Then, according to Lemma 4, any edge $(c_{i,j}, x_j)$ is not monitored by any of the vertices in $M' = (M' \cap V_1) \cup (M' \cap V'_1) \cup C$, a contradiction. Hence, C is a set cover of H of size at most k .

Due to space constraints, we omit the two other reductions. They are very similar to the above one. To obtain a bipartite graph, we remove the odd cycles of the construction at the expense of a higher diameter. To obtain maximum degree 3, we replace the set V_2 by some suitably defined binary trees. \square

5 Conclusion

We have introduced a new graph parameter useful in the area of network monitoring. We have related it to other standard graph parameters by the means of lower and upper bounds. It would be interesting to improve them. In particular, is it true that $\text{dem}(G) \leq \text{fes}(G) + 1$? As we have seen, this bound would be tight.

It would also be interesting to determine graph classes where DISTANCE-EDGE-MONITORING SET has a polynomial-time (or parameterized) exact or constant-factor approximation algorithm.

Acknowledgements. The authors acknowledge the financial support from the ANR project HOSIGRA (ANR-17-CE40-0022), the IFCAM project “Applications of graph homomorphisms” (MA/IFCAM/18/39), and the Programme IdEx Bordeaux – SysNum (ANR-10-IDEX-03-02).

References

1. Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Proietti, M.: Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties. Springer, Heidelberg (1999). <https://doi.org/10.1007/978-3-642-58412-1>
2. Bampas, E., Bilò, D., Drovandi, G., Gualà, L., Klasing, R., Proietti, G.: Network verification via routing table queries. *J. Comput. Syst. Sci.* **81**(1), 234–248 (2015)
3. Beerliova, Z., et al.: Network discovery and verification. *IEEE J. Sel. Areas Commun.* **24**(12), 2168–2181 (2006)
4. Bilò, D., Erlebach, T., Mihalák, M., Widmayer, P.: Discovery of network properties with all-shortest-paths queries. *Theoret. Comput. Sci.* **411**(14–15), 1626–1637 (2010)
5. Chen, J., et al.: Tight lower bounds for certain parameterized NP-hard problems. *Inf. Comput.* **201**(2), 216–231 (2005)
6. Dall’Asta, L., Alvarez-Hamelin, J.I., Barrat, A., Vázquez, A., Vespignani, A.: Exploring networks with traceroute-like probes: theory and simulations. *Theoret. Comput. Sci.* **355**(1), 6–24 (2006)
7. Dinur, I., Steurer, D.: Analytical approach to parallel repetition. In: Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing, STOC 2014, pp. 624–633 (2014)
8. Downey, R.G., Fellows, M.R.: Fundamentals of Parameterized Complexity. TCS. Springer, London (2013). <https://doi.org/10.1007/978-1-4471-5559-1>
9. Epstein, L., Levin, A., Woeginger, G.J.: The (weighted) metric dimension of graphs: hard and easy cases. *Algorithmica* **72**(4), 1130–1171 (2015)
10. Garey, M.R., Johnson, D.S.: The rectilinear Steiner tree problem is NP-complete. *SIAM J. Appl. Math.* **32**(4), 826–834 (1977)
11. Govindan, R., Tangmunarunkit, H.: Heuristics for Internet map discovery. In: Proceedings of the 19th IEEE International Conference on Computer Communications, INFOCOM 2000, pp. 1371–1380 (2000)
12. Harary, F., Melter, R.A.: On the metric dimension of a graph. *Ars Comb.* **2**, 191–195 (1976)

13. Hromkovič, J.: *Algorithmics for Hard Problems: Introduction to Combinatorial Optimization, Randomization, Approximation, and Heuristics*. Texts in Theoretical Computer Science. An EATCS Series, 2nd edn. Springer, Heidelberg (2004). <https://doi.org/10.1007/978-3-662-05269-3>
14. Johnson, D.S.: Approximation algorithms for combinatorial problems. *J. Comput. Syst. Sci.* **9**, 256–278 (1974)
15. Kelenc, A., Tratnik, N., Yero, I.G.: Uniquely identifying the edges of a graph: the edge metric dimension. *Discrete Appl. Math.* **251**, 204–220 (2018)
16. Manuel, P., Klavžar, S., Xavier, A., Arokiaraj, A., Thomas, E.: Strong edge geodesic problem in networks. *Open Math.* **15**, 1225–1235 (2017)
17. Niedermeier, R.: *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, Oxford (2006)
18. Oellermann, O.R., Peters-Fransen, J.: The strong metric dimension of graphs and digraphs. *Discrete Appl. Math.* **155**, 356–364 (2007)
19. Sebő, A., Tannier, E.: On metric generators of graphs. *Math. Oper. Res.* **29**(2), 383–393 (2004)
20. Slater, P.J.: Leaves of trees. *Congr. Numerantium* **14**, 549–559 (1975)