



HAL
open science

Monitoring the edges of a graph using distances

Florent Foucaud, Shih-Shun Kao, Ralf Klasing, Mirka Miller, Joe Ryan

► **To cite this version:**

Florent Foucaud, Shih-Shun Kao, Ralf Klasing, Mirka Miller, Joe Ryan. Monitoring the edges of a graph using distances. 2020. hal-03043870v1

HAL Id: hal-03043870

<https://hal.science/hal-03043870v1>

Preprint submitted on 7 Dec 2020 (v1), last revised 10 Dec 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Monitoring the edges of a graph using distances^{*†‡}

Florent Foucaud^{§¶} Shih-Shun Kao^{§||} Ralf Klasing[§] Mirka Miller^{**}

Joe Ryan^{††}

October 26, 2020

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 twice its feedback edge set number. Moreover, we characterize connected graphs G with $\text{dem}(G) = 2$.

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.

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

*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.

†A preliminary version of this paper appeared as [12].

‡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).

§Université de Bordeaux, Bordeaux INP, CNRS, LaBRI, UMR 5800, Talence, France. E-mails: florent.foucaud@gmail.com, ralf.klasing@labri.fr

¶Univ. Orléans, INSA Centre Val de Loire, LIFO EA 4022, F-45067 Orléans Cedex 2, France.

|| Department of Computer Science and Information Engineering, National Cheng Kung University, Tainan 701, Taiwan.

**School of Mathematical and Physical Sciences, The University of Newcastle, Australia and Department of Mathematics, University of West Bohemia, Czech Republic.

††School of Electrical Engineering and Computing, The University of Newcastle, Newcastle, Australia. E-mail: joe.ryan@newcastle.edu.au

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* [8, 14]. They are also frequently used for problems concerning *network verification* [2, 4, 6].

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 Section 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.

DISTANCE-EDGE-MONITORING SET

Instance: A graph G , an integer k .
--

Question: Do we have $\text{dem}(G) \leq k$?

MIN DISTANCE-EDGE-MONITORING SET

Instance: A graph G .

Task: Build a smallest possible distance-edge-monitoring set of G .

Related notions. A weaker model is studied in [2, 4] 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 .

A different (also weaker) model is the *Link Monitoring problem* studied in [5], in which one seeks to monitor the edges of a graph network by selecting vertices to act as probes. To each probe is assigned a routing tree (a DFS tree spanning the whole graph), and it is essentially required that each edge of the graph belongs to one of the trees.

Distance-edge-monitoring sets are also related to *resolving sets* and *edge-resolving sets*, that 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 [15, 26]. 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 [19]. A set that distinguishes all vertices and edges is a *mixed resolving set* [18]. Note that there is no relation, in general, between the (edge-)metric dimension of G , and $\text{dem}(G)$, as shown by the examples of trees and grids: trees can have arbitrarily large metric dimension and edge-metric dimension [19], but when G is a tree $\text{dem}(G) = 1$ (Theorem 9). Conversely, $\text{dem}(G)$ can be arbitrarily large for grids (Theorem 10), while their edge-metric dimension and metric dimension is 2 [19].

Another related concept is the one of *strong resolving sets* [24, 25]: 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 .

The concept of an (*strong*) *edge-geodetic set* is also related to ours. 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 [22].

A model related (only by name) is the one of *edge monitoring sets* [3]: a set S is edge monitoring if for every edge xy of G , e belongs to a triangle xyz with $z \in S$. But this is very far from our definition.

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

In Section 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 $\text{dem}(G) \leq 2 \text{fes}(G) - 2$ when $\text{fes}(G) \geq 3$ (when $\text{fes}(G) \leq 2$, $\text{dem}(G) \leq \text{fes}(G) + 1$).¹

In Section 4, we characterize connected graphs G with $\text{dem}(G) = 2$.

In Section 5, 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 Section 6.

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 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

¹This is an improvement over the results from the conference version [12], where we proved $\text{dem}(G) \leq 5 \text{fes}(G) - 5$.

$(x, v) \in P(M, e')$, 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

We next introduce the following terminology from [11]. In a graph, a vertex is a *core vertex* if it has degree at least 3. A path with all internal vertices of degree 2 and whose end-vertices are core vertices is called a *core path* (note that we allow the two end-vertices to be equal, but all other vertices must be distinct). A core path that is a cycle (that is, both end-vertices are equal) is a *core cycle*. 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).

Lemma 3 implies the following.

Observation 4. *Let G be a graph, and G' be its base graph. Then, $\text{dem}(G) = \text{dem}(G')$.*

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 5. *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 5.

Lemma 6. *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 5, 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 7. *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 5. \square

We next give a characterization for the cases where no other edge (other than those incident with x) belongs to $M(x)$.

Lemma 8. *Let G be a connected graph with a vertex x of G . The two following conditions are equivalent.*

- (i) $M(x)$ is the set of edges incident with x .
- (ii) For every vertex y of G with $x \neq y$, there exist two shortest paths from x to y sharing at most one edge (the one incident with x).

Proof. If (ii) holds, then for any $i > 1$, every vertex in level $L_i(x)$ has at least two neighbours in level $L_{i-1}(x)$, so (i) is implied by Lemma 5.

Conversely, if (i) holds, we can prove (ii) by induction on $d_G(x, y)$. The claim is true for all vertices in $L_1(x)$. Assume it is true for vertices at levels $L_1(x), \dots, L_i(x)$ and let y be a vertex in $L_{i+1}(x)$. By Lemma 5, y has at least two distinct neighbours, u and v , in $L_i(x)$. Consider two shortest paths P_u from u to x and P_v from v to x . If (ii) is satisfied, we are done. Otherwise, let st be the edge common to both P_u and P_v that is closest to u and v (with $s \in L_j$ and $t \in L_{j-1}$ for some $j > 1$). By the induction hypothesis, there are two shortest paths P_1 and P_2 from s to x that share at most an edge incident with x . Using these two shortest paths, we can easily combine them with P_u and P_v to obtain the two shortest paths from y to x , as required. \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.

3.1 Trees and grids

Theorem 9. *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 6, 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 9, we have $\text{dem}(G_{1,a}) = 1$. We can compute all other values.

Theorem 10. *For any integers $a, b \geq 2$, we have $\text{dem}(G_{a,b}) = \max\{a, b\}$.*

Proof. Without loss of generality we assume that $\max\{a, b\} = a$, and $G_{a,b}$ has a rows and b columns.

Clearly, for any vertex x , by Lemma 5, $M(x)$ consists of the edges that are in the same row and column as x . This shows that any distance-edge-monitoring set M must contain a vertex of each row and column, and $\text{dem}(G_{a,b}) \geq a$.

To see that $\text{dem}(G_{a,b}) \leq a$, we can choose any set M of a vertices containing exactly one vertex of every row and at least one vertex of every column of $G_{a,b}$: M is a distance-edge-monitoring set of $G_{a,b}$ of size a . \square

3.2 Connection to arboricity and clique number

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 11. *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 6, 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

3.3 Connection to vertex covers and consequences for hypercubes and complete bipartite graphs

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 12. *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 7, 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 13. *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 13 does not provide a characterization of graphs with $\text{dem}(G) = \text{vc}(G)$. For example, as seen in Theorem 10, 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 . Both $K_{a,b}$ and H_d satisfy Condition (ii) of Lemma 8. In a connected bipartite graph, the smallest vertex cover consists of the smallest of the two parts. Thus, by Observation 13, we obtain the following.

Corollary 14. *We have $\text{dem}(K_{a,b}) = \text{vc}(K_{a,b}) = \min\{a, b\}$, and $\text{dem}(H_d) = \text{vc}(H_d) = 2^{d-1}$.*

3.4 Connection to feedback edge set number

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).

The following folklore lemma uses the terminology defined in Section 2 (a proof can be found for example in Section 5.3.1 of [11] or Section 4.1 of [20]).

Lemma 15 ([11, 20]). *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 core paths.*

In other words, Lemma 15 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 9 states that for any tree G (that is, a graph with feedback edge set number 0), we have $\text{dem}(G) \leq \text{fes}(G) + 1$. We now show the same bound for graphs G with $\text{fes}(G) \leq 2$.

Theorem 16. *If $\text{fes}(G) \leq 2$, then $\text{dem}(G) \leq \text{fes}(G) + 1$. Moreover, if $\text{fes}(G) \leq 1$, then equality holds.*

Proof. Let $k = \text{fes}(G)$. If $k = 0$, G is a tree and we are done by Lemma 6.

Assume next that $k = 1$, that is, G is a connected unicyclic graph. Since G is not a tree, we have $\text{dem}(G) \geq 2$.

Let C be the unique cycle of G and let x be a vertex of C . If C has even length, $M(x) \cap E(C)$ consists of all edges of C except the two that are incident to the antipodal vertex of x in C . Thus, for any two non-adjacent vertices x and y of C , $(M(x) \cup M(y)) \cap E(C) = E(C)$. By Lemma 3, all other edges of G are monitored and thus $\{x, y\}$ is distance-edge-monitoring.

If C has odd length, $M(x) \cap E(C)$ consists of all edges of C except the one that connects the two vertices that are farthest from x . Thus, for any two vertices x and y of C , $(M(x) \cup M(y)) \cap E(C) = E(C)$ and again $\{x, y\}$ is distance-edge-monitoring.

Assume finally that $k = 2$ and let G_b be the base graph of G . By Lemma 15, G_b contains at most two core vertices joined by at most three core 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.

If one of the core paths of G_b is a core cycle, then in fact G_b must have two core cycles, with either one or two core vertices. Then, select one vertex from each core cycle that is farthest from the core vertex on this cycle. One can check that these two vertices form a distance-edge-monitoring set.

Otherwise, G_b consists of two core vertices x, y joined by three parallel core paths. If x, y are not adjacent, then $\{x, y\}$ forms a distance-edge-monitoring set of G_b . Otherwise, the edge xy is one of the three core paths. Then, one may select x and a middle vertices of each of the two core paths of length at least 2. This forms a distance-edge-monitoring set and completes the proof. \square

We will now give a weaker (but similar) bound for any value of $\text{fes}(G)$. This improves on our previous bound from [12].

Theorem 17. *Let G be a graph with $\text{fes}(G) = k$. If $k \geq 3$, then $\text{dem}(G) \leq 2k - 2$.*

Proof. Let G_b be the base graph of G . By Lemma 15, G_b contains at most $2k - 2$ core vertices, that are joined by at most $3k - 3$ core 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 $2k - 2$.

A first candidate would be to select the set C of all core vertices of G_b . However, this might not be a distance-edge-monitoring set, indeed for any odd-length core cycle of G_b , its middle edge is not monitored. Similarly, if G_b contains two distinct core vertices c_1, c_2 that are joined by a core path P of odd length at least 3 and by another core path of length 1 (an edge), then the middle edge of P is not monitored. (Note that every other edge e is monitored, indeed there is one endpoint x of e such that the unique shortest path from x to its closest core vertex goes through e .) We call these core paths, *problematic* core paths. Let \mathcal{P} be the set of problematic core paths. In what follows, we will start from C as a solution set and adjust it to handle problematic core paths.

To this end, we build a subgraph G'_b of G_b , by iteratively selecting a new core path P of \mathcal{P} , and by removing the inner-vertices and edges of P from G_b , unless we decrease the number of core vertices of the graph by more than 1. When there is no such candidate core path left in \mathcal{P} , we stop the process and obtain the graph G'_b . Let p be the number of core paths of \mathcal{P} that we have removed from G_b in this process, and let $k' = \text{fes}(G'_b)$.

First, we show that $k' = k - p$. First, to see that $k' \geq k - p$, we show that $k \leq k' + p$. Observe that one can obtain a feedback edge set of G_b from one of G'_b by additionally selecting one edge of each of the p core paths that were deleted from G_b to obtain G'_b , which shows that $k \leq k' + p$, as claimed. Second, to see that $k' \leq k - p$, consider any feedback edge set F of G_b , and note that it necessarily contains one edge of each core cycle of G_b . Moreover, for every pair c_1, c_2 of distinct core vertices of G_b joined by $p(c_1, c_2) \geq 2$ core paths, it contains an edge of all of them, except possibly one. In fact we may assume that F is chosen so that it contains an edge of all of them that have length at least 2. Thus, every problematic core path of G_b has an edge in F . Removing these p edges from F , we obtain the desired feedback edge set of G'_b of size $k - p$.

Let C' be the set of core vertices of G'_b . Since at each step of the procedure used for building G'_b , we decrease the number of core vertices by at most 1, G'_b has at least $|C| - p$ core vertices. Thus, since $k' = k - p$ and applying Lemma 15 to G'_b , we have:

$$\begin{aligned}
|C| &\leq |C'| + p \\
&\leq 2k' - 2 + p \\
&= 2(k - p) - 2 + p \\
&= 2k - p - 2
\end{aligned}$$

We now build our distance-edge-monitoring set M as follows. First, we let $M = C$. Moreover, for each problematic core path P that was deleted to obtain G'_b , we add one arbitrary inner-vertex of P to M . At this point, M monitors all edges, except those of the problematic paths of G_b that were not deleted when constructing G'_b . Observe that in G'_b , each such core path P joins two distinct adjacent core vertices c_1, c_2 of G'_b , both having degree exactly 3 in G'_b . (Moreover, since $k \geq 3$, c_1 and c_2 are joined only by two core paths in G'_b .) We select an arbitrary core vertex among c_1 and c_2 (say c_1), remove it from M , and replace it with the neighbour of c_1 on P . We do this for every such remaining problematic core path. Again, since $k \geq 3$, now every problematic core path contains a vertex of M .

It is clear that M has size at most $2k - 2$. We now show that M is distance-edge-monitoring. To do so, consider an edge e of G_b , which lies on some core path P . If P is not a problematic path or e is not the middle edge of P , then similarly as for the set C , there is an endpoint x of e whose unique shortest path to one of the core vertices c of P (or the neighbour of c , if c was removed from M) goes through e , and so e is monitored by c or its neighbour. For the remainder, assume that P is problematic and that e is the middle edge of P . If P was among the p paths from which we added an extra vertex x to M , then x monitors e . Otherwise, we have removed one of the core vertices of P , and replaced it with its neighbour on P . Then, this neighbour monitors e . This completes the proof. \square

We do not believe that the bound of Theorem 17 is tight. 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 10, 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 10, we have $\text{dem}(G_{a,b}) = \max\{a, b\}$, but $\text{fes}(G_{a,b})$ is linear in the order ab .

4 Graphs G with $\text{dem}(G) = 2$

In this section, we characterize connected graphs G with $\text{dem}(G) = 2$.

For two vertices u, v of a graph G , we denote by $B_{i,j}$ the set of vertices at distance i of u and distance j of v in G .

Observation 18. *Let G be a graph, and let u, v be two vertices in G . Let $B_{i,j}$ be the set of vertices at distance i of u and distance j of v in G . Then, $|i - j| \leq d(u, v)$.*

Observation 19. *Let G be a graph, and let u, v be two vertices in G . Let $x \in B_{i,j}$, and let $y \in B_{i',j'}$ be a neighbor of x . Then, $i' \in \{i - 1, i, i + 1\}$ and $j' \in \{j - 1, j, j + 1\}$.*

We are now ready to state our characterization in the next theorem. An example of a graph G with $\text{dem}(G) = 2$ following the characterization is given in Figure 1.

Theorem 20. *Let G be a connected graph with at least one cycle, and let G' be the base graph of G . Then, $\text{dem}(G) = 2$ if and only if there are two vertices u, v in G' such that all of the following conditions (I) – (IV) hold in G' .*

- (I) *Every $B_{i,j}$ is an independent set.*
- (II) *Every vertex x in $B_{i,j}$ has at most one neighbor in each of the four sets $B_{i-1,j} \cup B_{i-1,j-1}$, $B_{i-1,j} \cup B_{i-1,j+1}$, $B_{i,j-1} \cup B_{i-1,j-1}$ and $B_{i,j-1} \cup B_{i+1,j-1}$.*
- (III) *There is no 4-vertex path $xyz'z'$ with $z \in B_{i-1,a}$, $z' \in B_{a',j}$, $x \in B_{i,j}$, $y \in B_{i-1,j+1}$, $a \in \{j - 1, j + 1\}$, $a' \in \{i - 2, i\}$.*
- (IV) *$x \in B_{i,j}$ has neighbors in at most two sets among $B_{i-1,j+1}$, $B_{i-1,j-1}$, $B_{i+1,j-1}$.*

Proof. \Rightarrow Let us assume $\text{dem}(G) = 2$, hence $\text{dem}(G') = 2$ according to Observation 4. Let $\{u, v\}$ be a distance-edge-monitoring set of G' . We will show that the properties (I) - (IV) hold.

- (I) Assume that $B_{i,j}$ is not an independent set. Let $e = xy$ be an edge such that $x, y \in B_{i,j}$. Then, the distance from x to u is i and the distance from y to u is i , hence e is not monitored by u , according to Lemma 5. Likewise, e is not monitored by v , a contradiction.
- (II) Let $x \in B_{i,j}$.
- (i) Assume x has two neighbors $y, y' \in B_{i-1,j} \cup B_{i-1,j-1}$. Assume first that x has two neighbors y, y' such that either both $y, y' \in B_{i-1,j}$, both $y, y' \in B_{i-1,j-1}$, or $y \in B_{i-1,j}$, $y' \in B_{i-1,j-1}$. Let $e = xy$. Then, the distance from y to u and the distance from y' to u are both $i-1$, hence e is not monitored by u , according to Lemma 5. The distance from y to v and the distance from x to v are also the same, hence e is not monitored by v , according to Lemma 5. This is a contradiction.
- (ii) Assume x has two neighbors $y, y' \in B_{i-1,j} \cup B_{i-1,j+1}$. If both $y, y' \in B_{i-1,j}$ we are done by Case (i). Thus, assume that $y \in B_{i-1,j} \cup B_{i-1,j+1}$ and $y' \in B_{i-1,j+1}$, and let $e = xy$. If $y, y' \in B_{i-1,j+1}$, as in Case (i), the ordered pair of distances from y to u and v and the ordered pair of distances from y' to u and v are the same, and by Lemma 5 e is not monitored. If $y \in B_{i-1,j}$, $y' \in B_{i-1,j+1}$, the same argument works to show that u does not monitor e . Moreover, the distances from x to v and from y to v are both j , and by Lemma 5 v also does not monitor e . This is a contradiction.
- (iii) The cases where x has two neighbors in $B_{i,j-1} \cup B_{i-1,j-1}$ or in $B_{i,j-1} \cup B_{i+1,j-1}$ are completely symmetric to Cases (i) and (ii).
- (III) Assume there is a path $zxy'z'$ with $z \in B_{i-1,a}$, $z' \in B_{a',j}$, $x \in B_{i,j}$, $y \in B_{i-1,j+1}$, $a \in \{j-1, j+1\}$, $a' \in \{i-2, i\}$, and let $e = xy$, then the distance from u to z is $i-1$, and the distance from u to y is $i-1$, hence by Lemma 5 e is not monitored by u . The distance from v to z' is j , and the distance from v to x is j , hence e is also not monitored by v according to Lemma 5, a contradiction.
- (IV) Let us assume by contradiction that $x \in B_{i,j}$ has three neighbors y, y', z such that $y \in B_{i-1,j-1}$, $y' \in B_{i+1,j-1}$, and $z \in B_{i-1,j+1}$, and let $e = xy$. The distance from z to u is $i-1$, and the distance from y to u is $i-1$, hence e is not monitored by u . The distance from y to v is $j-1$, and the distance from y' to v is $j-1$, hence e is not monitored by v , according to Lemma 5, a contradiction.

\Leftarrow Let us assume that there are two vertices u, v in G' such that all of the conditions (I) - (IV) hold in G' . We will show that $\{u, v\}$ is a distance-edge-monitoring set in G' , and hence $\text{dem}(G) = 2$ according to Observation 4. Let $e = xy$ be an edge in G such that $x \in B_{i,j}$. According to condition (I), $y \notin B_{i,j}$. We will distinguish three cases with respect to y . All the other cases are symmetric to these three main cases.

Case(a): $y \in B_{i-1,j-1}$. (The case $y \in B_{i+1,j+1}$ is symmetric.)

By contradiction, assume e is not monitored. Then, there is a path P_i of length i from x to u not using e , and there is a path P_j of length j from x to v not using e . Let z be the neighbor of x in P_i . Then, z belongs to $B_{i-1,a}$. Let z' be the neighbor of x in P_j . Then, z' belongs to $B_{a',j-1}$. According to Observation 19 and condition (II), $a = j+1$ and $a' = i+1$. Hence, x has neighbors in all three sets $B_{i-1,j+1}$, $B_{i-1,j-1}$, $B_{i+1,j-1}$, in contradiction to condition (IV).

Case(b): $y \in B_{i,j-1}$. (The cases $y \in B_{i-1,j}$, $B_{i+1,j}$ or $B_{i,j+1}$ are symmetric.)

By contradiction, assume e is not monitored. Then, there is a path P_j of length j from x to v not using e . Let z' be the neighbor of x in P_j . Then, z' belongs to $B_{a',j-1}$. As $a' \in \{i-1, i, i+1\}$, this contradicts condition (II).

Case(c): $y \in B_{i-1,j+1}$. (The case $y \in B_{i+1,j-1}$ is symmetric.)

By contradiction, assume e is not monitored. Then, there is a path of length i from x to u avoiding y (let z be the neighbor of x on this path). Similarly, there is a path of length $j+1$ from y to v avoiding x (let z' be the neighbor of y on this path). Thus, there is a 4-vertex path $zxyz'$ with $z \in B_{i-1,a}$, $z' \in B_{a',j}$, $a \in \{j-1, j, j+1\}$, $a' \in \{i-2, i-1, i\}$. According to condition (II), $a \neq j$ and $a' \neq i-1$. However, the other cases ($a \in \{j-1, j+1\}$ and $a' \in \{i-2, i\}$) contradict condition (III).

□

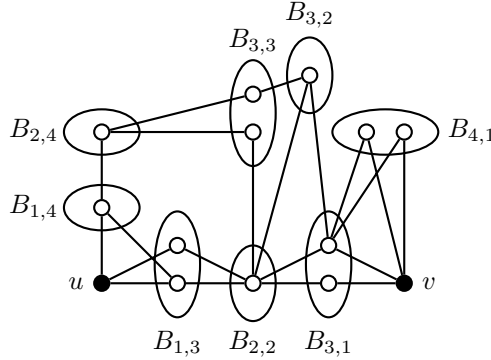


Figure 1: A (base) graph G with $\text{dem}(G) = 2$, where $\{u, v\}$ is a distance-edge-monitoring set.

5 Complexity

We now study the algorithmic complexity of finding an optimal distance-edge-monitoring set.

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, 16] 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 $W[i]$ (with $i \geq 1$) denote complexity classes with parameterized decision problems that are believed not to be fpt. We refer to the books [10, 23] for more details.

5.1 NP-hardness

We will now use the connection to vertex covers hinted in Section 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 21. *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 $\text{vc}(G) = \text{dem}(G \bowtie K_1)$.*

Proof. We denote by u the vertex from K_1 . By Theorem 12, we have $\text{dem}(G \bowtie K_1) \leq \text{vc}(G \bowtie K_1) \leq \text{vc}(G) + 1$ since u together with any vertex cover of G is a vertex cover of $G \bowtie K_1$.

To see that $\text{vc}(G) \leq \text{dem}(G \bowtie K_1)$, assume that we have a distance-edge-monitoring set M of $G \bowtie K_1$, and let vw be any edge of G . There exists a vertex x in M such that $vw \in M(x)$. Then, we must have $x \in \{v, w\}$, indeed, if not, then vw would be part of a unique shortest path of length 2 from x to some vertex y , which is not possible since the path from x to y going through u is another path of length 2.

Now, assume that G has radius at least 4: it remains to prove that $\text{dem}(G \bowtie K_1) \leq \text{vc}(G)$. Let M be a vertex cover of G . By Lemma 7, all edges incident with a vertex of M are monitored

by M , which includes all (original) edges of G . Let $e = uv$ be an edge incident with the vertex u of K_1 such that $v \notin M$. We claim that there is a vertex x of M with $d_G(v, x) \geq 3$. Indeed, since G has radius at least 4, there is a vertex at distance 4 of v in G . Consider a neighbour of that vertex that is itself at distance 3 of v in G . Since M is a vertex cover, one of these two vertices is in M , which proves the existence of x . Now, we know that $d_{G \setminus K_1}(v, x) = 2$, and the only path of length 2 from v to x goes through e . Thus $e \in M(x)$. This shows that all edges are monitored by M , and completes the proof. \square

Corollary 22. DISTANCE-EDGE-MONITORING SET is NP-complete, even for graphs obtained from a planar subcubic graph by attaching a universal vertex.

Proof. It is known (see [13]) that VERTEX COVER is NP-complete for planar subcubic graphs with radius at least 4 (for the latter condition, note that a subcubic graph of radius at most 3 has constant order). Theorem 21 provides a polynomial reduction from such instances, and this completes the proof. \square

5.2 Approximation algorithm

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 $\ln(|X| + 1)$ [17], but, unless $P=NP$, not within a factor of $(1 - \epsilon) \ln |X|$ (for every positive ϵ) [9]. Moreover, SET COVER is $W[2]$ -hard (when parameterized by k) and not solvable in time $|X|^{o(k)} |H|^{O(1)}$ unless $FPT=W[1]$ [7].

Theorem 23. 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 [17]. \square

5.3 Approximation and parameterized hardness

Theorem 24. 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 [9], the parameterized hardness of SET COVER (parameterized by solution size), and the lower bound on its running time [7].

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.

(1) 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 $V = V_1 \cup V_2 \cup \dots \cup V_5 \cup V'_1 \cup V'_2 \cup V'_3$, $E = E_1 \cup E_2 \cup E_3 \cup E_4 \cup E'_1 \cup E'_2 \cup E'_3 \cup E'_4$ and

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

An example is given in Figure 2. Notice that the purpose of the vertices in V'_1, V'_2, V'_3 is only to reduce the diameter. A simpler reduction without these vertices, and with an additional edge joining u_2 and u_3 would also function.

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 5, 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 $E'_2 \cup E'_3 \cup E'_4$. 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

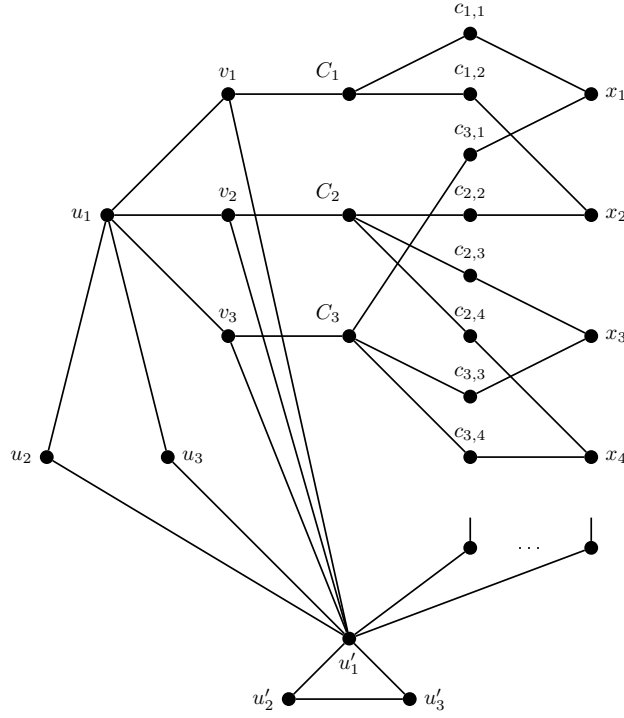


Figure 2: First reduction from SET COVER to DISTANCE-EDGE-MONITORING SET from the proof of Theorem 24 applied to the hypergraph $(\{x_1, x_2, x_3, x_4\}, \{C_1 = \{x_1, x_2\}, C_2 = \{x_2, x_3, x_4\}, C_3 = \{x_1, x_3, x_4\}\})$. Vertices and edges of V'_i and E'_i for $i = 2, 3$ are only suggested.

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 $E_1, E_2, E_3, E'_1, E'_2, E'_3$ and E'_4 are monitored by $\{u_1, u'_2\}$. However, no edge of E_4 is 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 $c'_{i,j} \in M$, 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 $|M' \cap V_1| = 1, |M' \cap V'_1| \geq 1, |M' \cap V_3| \leq k, M' \cap (V_2 \cup V_4 \cup V_5 \cup V'_2 \cup V'_3) = \emptyset$. 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 5, 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 .

(2) We now describe the second reduction, in which the obtained instance is bipartite and has diameter 6. 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 $V = V_1 \cup V_2 \cup \dots \cup V_5 \cup V'_1 \cup V'_2 \cup V'_3 \cup V''$, $E = E_1 \cup E_2 \cup E_3 \cup E_4 \cup E'_1 \cup E'_2 \cup E'_3 \cup E''$ and

$$\begin{aligned} V_1 &= \{u_1, u_2, u_3, u_4\}, \\ V_2 &= \{v_i \mid 1 \leq i \leq |S|\}, \\ V_3 &= S, \\ V_4 &= \{c_{i,j} \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}, \end{aligned}$$

$$\begin{aligned}
V_5 &= X, \\
V'_1 &= \{u'_1, u'_2, u'_3, u'_4\}, \\
V'_2 &= \{v'_i \mid 1 \leq i \leq |S|\}, \\
V'_3 &= \{w'_j \mid 1 \leq j \leq |X|\}, \\
V'' &= \{u''_1\}, \\
E_1 &= \{(u_1, u_2), (u_2, u_3), (u_3, u_4), (u_4, u_1)\}, \\
E_2 &= \{(u_1, v_i), (v_i, C_i) \mid 1 \leq i \leq |S|\}, \\
E_3 &= \{(C_i, c_{i,j}) \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}, \\
E_4 &= \{(c_{i,j}, x_j) \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}, \\
E'_1 &= \{(u'_1, u'_2), (u'_2, u'_3), (u'_3, u'_4), (u'_4, u'_1)\}, \\
E'_2 &= \{(u'_1, v'_i), (v'_i, C_i) \mid 1 \leq i \leq |S|\}, \\
E'_3 &= \{(u'_1, w'_j), (w'_j, x_j) \mid 1 \leq j \leq |X|\}, \\
E'' &= \{(u_1, u''_1), (u''_1, u'_1)\}.
\end{aligned}$$

The reduction is depicted in Figure 3.

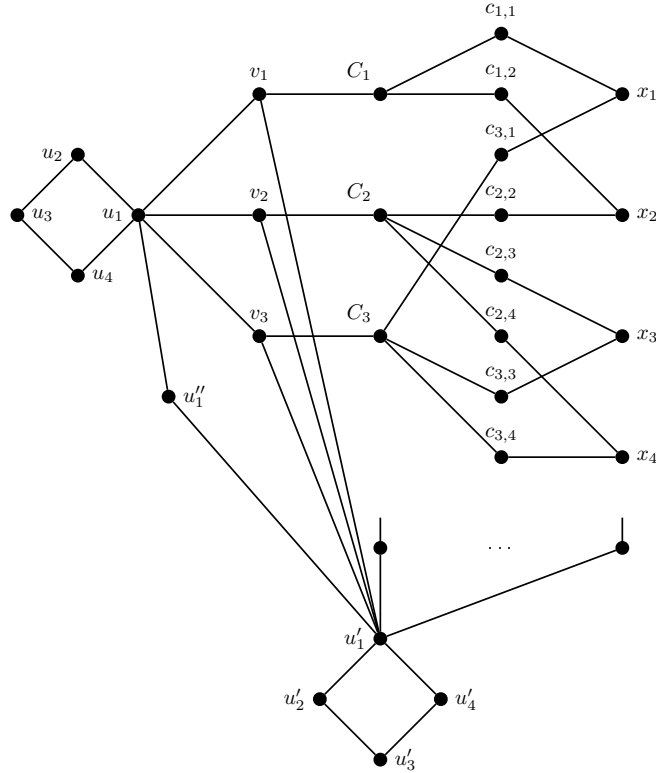


Figure 3: Second reduction from SET COVER to DISTANCE-EDGE-MONITORING SET from the proof of Theorem 24 applied to the hypergraph $(\{x_1, x_2, x_3, x_4\}, \{C_1 = \{x_1, x_2\}, C_2 = \{x_2, x_3, x_4\}, C_3 = \{x_1, x_3, x_4\}\})$. Vertices and edges of V'_i and E'_i for $i = 2, 3$ are only suggested.

Let C be a set cover of H of size k . Define $M = C \cup \{u_3, u'_3\}$. Then, by Lemma 5, u_3 monitors the edges (u_2, u_3) , (u_3, u_4) , (u_1, u''_1) , (u'_2, u'_1) , (u'_4, u'_1) and all the edges in $E_2 \cup E_3$, and u'_3 monitors the edges (u'_2, u'_3) , (u'_3, u'_4) , (u''_1, u'_1) , (u_2, u_1) , (u_4, u_1) and all the edges in $E'_2 \cup E'_3 \cup E'_4$. 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 $k + 2$.

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 $u_3 \in M$. 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. In order to monitor the edge (u'_2, u'_3) , either $u'_2 \in M$ or $u'_3 \in M$. If $v'_i \in M$, let $M = M \setminus \{v'_i\} \cup \{C_i\}$, 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. If $u''_1 \in M$, let $M = M \setminus \{u''_1\} \cup \{u_3\}$, and the set of monitored edges does not decrease. Iterating this process, we finally obtain a distance-edge-monitoring set M' of G with $|M' \cap V_1| \geq 1$, $|M' \cap V'_1| \geq 1$, $|M' \cap V_3| \leq k$, $M' \cap (V_2 \cup V_4 \cup V_5 \cup V'_2 \cup V'_3 \cup V'') = \emptyset$. 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 . Without loss of generality, assume that each vertex of X belongs to at least two edges of C . Then, according to Lemma 5, 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 .

(3) We now describe the third reduction, in which the obtained instance is bipartite and has maximum degree 3. 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 $C_i = \{c_{i,j} \mid x_j \in C_i\}$. Let $D = \lceil \log_2 \max\{|X|, |S|\} \rceil$. Let $B(r, \{r_1, r_2, \dots, r_\ell\}, d)$ denote a binary tree with root r and ℓ leaves r_1, r_2, \dots, r_ℓ at distance $d \geq \lceil \log_2 \ell \rceil$ from r . Let $P(r, s, \ell)$ denote the path of length ℓ with endpoints r and s . Let $\mathcal{C}_i = \{c_{i,j} \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}$. Let $\mathcal{X}_j = \{c_{i,j} \mid 1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i\}$. Let $B = B(u, \{v_1, v_2, \dots, v_{|S|}\}, D)$. For $1 \leq i \leq |S|, 1 \leq j \leq |X|$, let $P_i = P(v_i, C_i, D)$, $B(C_i) = B(C_i, \mathcal{C}_i, D)$, and $B(x_j) = B(x_j, \mathcal{X}_j, D)$. Construct the following instance $(G, k+1) = ((V, E), k+1)$ of DISTANCE-EDGE-MONITORING SET, where $V = V_1 \cup V_2 \cup \dots \cup V_8$, $E = E_1 \cup E_2 \cup \dots \cup E_5$ and

$$\begin{aligned}
V_1 &= \{u_1, u_2, u_3, u_4\}, \\
V_2 &= V(B), \\
V_3 &= \cup_{1 \leq i \leq |S|} V(P_i), \\
V_4 &= S, \\
V_5 &= \cup_{1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i} V(B(C_i)), \\
V_6 &= \cup_{1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i} C_i, \\
V_7 &= \cup_{1 \leq j \leq |X|} V(B(x_j)), \\
V_8 &= X, \\
E_1 &= \{(u_1, u_2), (u_2, u_3), (u_3, u_4), (u_4, u_1), (u_1, u)\}, \\
E_2 &= E(B), \\
E_3 &= \cup_{1 \leq i \leq |S|} E(P_i), \\
E_4 &= \cup_{1 \leq i \leq |S|, 1 \leq j \leq |X|, x_j \in C_i} E(B(C_i)), \\
E_5 &= \cup_{1 \leq j \leq |X|} E(B(x_j)).
\end{aligned}$$

The reduction is depicted in Figure 4.

Let C be a set cover of H of size k . Define $M = C \cup \{u_2\}$. Then, by Lemma 5, u_2 monitors the edges (u_1, u_2) , (u_2, u_3) , (u_1, u) and all the edges in $E_2 \cup E_3 \cup E_4$, and any vertex in C monitors the edges (u_1, u_4) , (u_1, u_2) . If $e \in E(B(x_j))$, 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 $k + 1$.

Let M be a distance-edge-monitoring set of G of size at most $k + 1$. In order to monitor the edge (u_2, u_3) , either $u_2 \in M$ or $u_3 \in M$. If $v \in (M \cap V(B))$, descend to the leftmost leaf v_i in B with v as its ancestor, and let $M = M \setminus \{v_i\} \cup \{C_i\}$. Then, the set of monitored edges does not decrease. If $v \in (M \cap V(P_i))$, let $M = M \setminus \{v_i\} \cup \{C_i\}$, and the set of monitored edges does not decrease. If $v \in (M \cap V(B(C_i)))$, let $M = M \setminus \{v\} \cup \{C_i\}$, and the set of monitored edges does not decrease. If $v \in (M \cap V(B(x_j)))$, 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. Iterating this process, we finally obtain a distance-edge-

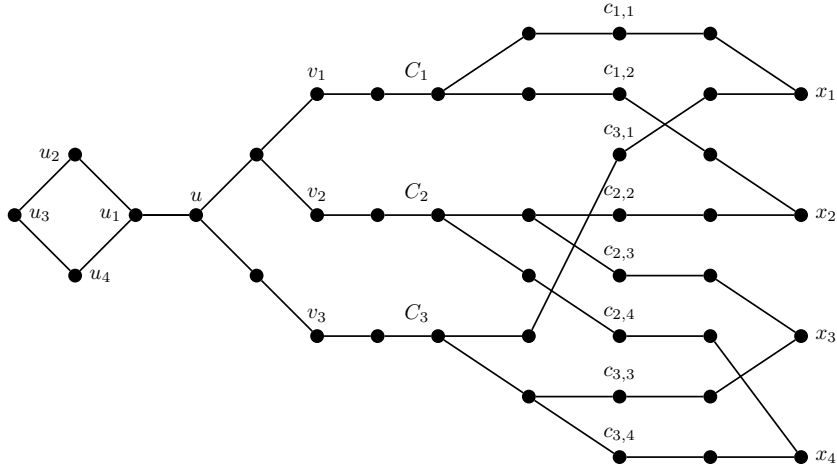


Figure 4: Third reduction from SET COVER to DISTANCE-EDGE-MONITORING SET from the proof of Theorem 24 applied to the hypergraph $(\{x_1, x_2, x_3, x_4\}, \{C_1 = \{x_1, x_2\}, C_2 = \{x_2, x_3, x_4\}, C_3 = \{x_1, x_3, x_4\}\})$, thus $D = \lceil \log_2 \max\{3, 4\} \rceil = 2$.

monitoring set M' of G with $|M' \cap V_1| \geq 1$, $|M' \cap V_4| \leq k$, $M' \cap (V_2 \cup V_3 \cup V_5 \cup V_6 \cup V_7 \cup V_8) = \emptyset$. Let $C = M' \cap V_4$. If C is not a set cover of H , then there is $x_j \in X$ that is not covered by C . Without loss of generality, assume that each vertex of X belongs to at least two edges of C . Then, according to Lemma 5, any edge (v, x_j) is not monitored by any of the vertices in $M' = (M' \cap V_1) \cup C$, a contradiction. Hence, C is a set cover of H of size at most k . \square

6 Concluding remarks and questions

We have introduced a new graph parameter useful in the area of network monitoring. We conclude the papers with some remarks for future research directions.

6.1 Structural graph parameters

We have related the parameter dem to other standard graph parameters by giving lower and upper bounds. The diagram in Figure 5 shows the relations between these parameters. 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. Moreover, is the bound $\text{dem}(G) \geq \omega(G)/2$ from Theorem 11 tight?

Other structural graph parameters that have not been considered here. A *feedback vertex set* of a graph G is a set of vertices such that removing them from G leaves a forest. The smallest size of a feedback vertex set of G is denoted by $\text{fvs}(G)$. It is not difficult to show that for any graph G , $\text{fvs}(G) \leq \text{vc}(G)$ and $\text{fvs}(G) \leq \text{fes}(G)$. However, we cannot hope to bound $\text{dem}(G)$ by a function of $\text{fvs}(G)$, as we did in Theorems 12 and 17 for $\text{vc}(G)$ and $\text{fes}(G)$. Indeed, consider the example of a path P_a of order $a \geq 9$ to which we add a universal vertex. This graph $P_a \bowtie K_1$ has radius at least 4, thus, by Theorem 21, we have $\text{dem}(P_a \bowtie K_1) = \text{vc}(P_a) = \lfloor a/2 \rfloor$, while we have $\text{fvs}(P_a \bowtie K_1) = 1$. Similarly, the pathwidth of $P_a \bowtie K_1$ is 2. Nevertheless, we do not know whether we can upper-bound $\text{dem}(G)$ by a function of the bandwidth or the treedepth of G .

6.2 Further complexity questions

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. For example, what happens for planar graphs, chordal graphs or interval graphs?

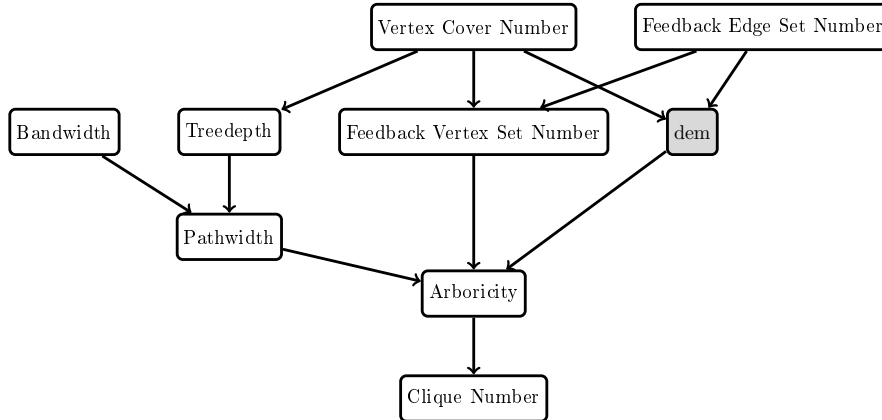


Figure 5: Relations between some structural parameters and dem. Arcs between parameters indicate that the bottom parameter is upper-bounded by a function of the top parameter.

6.3 Strengthening the definition

It is also reasonable to require to be able to simultaneously detect multiple edge failures in a network. In that case, we could propose the following strengthening of distance-edge-monitoring sets, based on extending Proposition 2. For a graph G and a set E of edges of G , we denote by $G - E$ the graph obtained by removing all edges of E from G .

Definition 25. For a set E of edges and a set M of vertices of a graph G , we 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 $d_G(x, y) \neq d_{G-E}(x, y)$.

For a graph G , a distance-edge-monitoring set M of G is called a strong distance-edge-monitoring set of G if for any two distinct subsets E_1 and E_2 of edges of G , we have $P(M, E_1) \neq P(M, E_2)$. We denote by $\text{sdem}(G)$ the smallest size of a strong distance-edge-monitoring set of G .

However, it turns out that the concept of strong distance edge-monitoring sets is in fact equivalent to the one of a vertex cover!

Proposition 26. A set M of vertices of a graph G is strong distance-edge-monitoring if and only if it is a vertex cover of G .

Proof. First, assume that M is a vertex cover. For an edge uv (assume without loss of generality that $u \in M$), we have $(u, v) \in P(M, E)$ if and only if $uv \in E$. For two distinct sets E_1, E_2 of edges of G , we have an edge $e = uv$ in the symmetric difference $E_1 \ominus E_2$ and thus one of the pairs (u, v) or (v, u) belongs to exactly one of $P(M, E_1)$ and $P(M, E_2)$. This implies that M is a strong distance-edge-monitoring set.

Conversely, let M be a strong distance-edge-monitoring set of G . Assume for a contradiction that it is not a vertex cover of G , so, there is an edge uv with $\{u, v\} \cap M = \emptyset$. Let E be the set of edges that are incident with u or with v (or both), and $E' = E \setminus \{uv\}$. Then, for any two vertices x, y with $x \notin \{u, v\}$, we have $d_{G-E}(x, y) = d_{G-E'}(x, y)$. In particular, this is true when $x \in M$. Thus, we have $P(M, E) = P(M, E')$, a contradiction. \square

Nevertheless, it could remain interesting to study an intermediate concept, where we wish to monitor all sets of edges of size at most a given constant. This has been done for example in the context of the metric dimension in [21].

References

- [1] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela and M. Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer, 1999.

- [2] E. Bampas, D. Bilò, G. Drovandi, L. Gualà, R. Klasing and G. Proietti. Network verification via routing table queries. *Journal of Computer and System Sciences* 81(1):234–248, 2015.
- [3] J. Baste, F. Beggas, H. Kheddouci and I. Sau. On the parameterized complexity of the Edge Monitoring problem. *Information Processing Letters* 121:39–44, 2017.
- [4] Z. Beerliova, F. Eberhard, T. Erlebach, A. Hall, M. Hoffmann, M. Mihalák and L. S. Ram. Network discovery and verification. *IEEE Journal on Selected Areas in Communications* 24(12):2168–2181, 2006.
- [5] Y. Bejerano and R. Rastogi. Robust monitoring of link delays and faults in IP networks. *IEEE/ACM Transactions on Networking* 14(5):1092–1103, 2006.
- [6] D. Bilò, T. Erlebach, M. Mihalák and P. Widmayer. Discovery of network properties with all-shortest-paths queries. *Theoretical Computer Science* 411(14–15):1626–1637, 2010.
- [7] J. Chen, B. Chor, M. Fellows, X. Huang, D. W. Juedes, I. A. Kanj and G. Xia. Tight lower bounds for certain parameterized NP-hard problems. *Information and Computation* 201(2):216–231, 2005.
- [8] L. Dall’Asta, J. I. Alvarez-Hamelin, A. Barrat, A. Vázquez and A. Vespignani. Exploring networks with traceroute-like probes: Theory and simulations. *Theoretical Computer Science* 355(1):6–24, 2006.
- [9] I. Dinur and D. Steurer. Analytical approach to parallel repetition. *Proceedings of the forty-sixth annual ACM symposium on Theory of computing, STOC 2014*:624–633, 2014.
- [10] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Springer, 2013.
- [11] L. Epstein, A. Levin and G. J. Woeginger. The (weighted) metric dimension of graphs: hard and easy cases. *Algorithmica* 72(4):1130–1171, 2015.
- [12] F. Foucaud, R. Klasing, M. Miller and J. Ryan. Monitoring the edges of a graph using distances. Proceedings of the 6th International Conference on Algorithms and Discrete Applied Mathematics (CALDAM 2020). *Lecture Notes in Computer Science* 12016:28–40, 2020.
- [13] M. R. Garey and D. S. Johnson. The rectilinear Steiner tree problem is NP-complete. *SIAM Journal on Applied Mathematics* 32(4):826–834, 1977.
- [14] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. Proceedings of the 19th IEEE International Conference on Computer Communications, INFOCOM’00, 2000, pp. 1371–1380.
- [15] F. Harary and R. A. Melter. On the metric dimension of a graph. *Ars Combinatoria* 2:191–195, 1976.
- [16] J. Hromkovič. *Algorithmics for Hard Problems - Introduction to Combinatorial Optimization, Randomization, Approximation, and Heuristics, Second Edition*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004.
- [17] D. S. Johnson, Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences* 9:256–278, 1974.
- [18] A. Kelenc, D. Kuziak, A. Taranenka and I. G. Yero. Mixed metric dimension of graphs. *Applied Mathematics and Computation* 314:429–438, 2017.
- [19] A. Kelenc, N. Tratnik and I. G. Yero. Uniquely identifying the edges of a graph: The edge metric dimension. *Discrete Applied Mathematics* 251:204–220, 2018.
- [20] L. Kellerhals and T. Koana. Parameterized complexity of geodetic set. Accepted for presentation at the 15th International Symposium on Parameterized and Exact Computation (IPEC 2020). <https://arxiv.org/abs/2001.03098>

- [21] T. Laihonen. The metric dimension for resolving several objects. *Information Processing Letters* 116(11):694–700, 2016.
- [22] P. Manuel, S. Klavžar, A. Xavier, A. Arokiaraj and E. Thomas. Strong edge geodetic problem in networks. *Open Mathematics* 15:1225–1235, 2017.
- [23] R. Niedermeier. *Invitation to fixed-parameter algorithms*. Oxford University Press, 2006.
- [24] O. R. Oellermann and J. Peters-Fransen. The strong metric dimension of graphs and digraphs. *Discrete Applied Mathematics* 155:356–364, 2007.
- [25] A. Sebő and E. Tannier. On metric generators of graphs. *Math. Oper. Res.* 29(2):383–393, 2004.
- [26] P. J. Slater. Leaves of trees. *Congressus Numerantium* 14:549–559, 1975.