CAR-TR-728
CS-TR-3326
UMIACS-TR-94-92

CCR-93-07462
DACA76-92-C-0009
July 1994

Localization in Graphs

Samir Khuller
Department of Computer Science
Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742-3255

Balaji Raghavachari
Department of Computer Science
The University of Texas at Dallas
Richardson, TX 75083-0688

Azriel Rosenfeld
Computer Vision Laboratory
Center for Automation Research
University of Maryland
College Park, MD 20742-3275

Abstract

Navigation can be studied in a graph-structured framework in which the navigating agent (which we shall assume to be a point robot) moves from node to node of a "graph space". The robot can locate itself by the presence of distinctively labeled "landmark" nodes in the graph space. For a robot navigating in Euclidean space, visual detection of a distinctive landmark provides information about the direction to the landmark, and allows the robot to determine its position by triangulation. On a graph, however, there is neither the concept of direction nor that of visibility. Instead, we shall assume that a robot navigating on a graph can sense the distances to a set of landmarks.

Evidently, if the robot knows its distances to a sufficiently large set of landmarks, its position on the graph is uniquely determined. This suggests the following problem: given a graph, what are the fewest number of landmarks needed, and where should they be located, so that the distances to the landmarks uniquely determine the robot's position on the graph? This is actually a classical problem about metric spaces. A minimum set of landmarks which uniquely determine the robot's position is called a "metric basis", and the minimum number of landmarks is called the "metric dimension" of the graph. In this paper we present some results about this problem. Our main *new* result is that the metric dimension can be approximated in polynomial time within a factor of $O(\log n)$; we also establish some properties of graphs with metric dimension 2.

The support of the National Science Foundation Initiation Award CCR-93-07462, and that of the Advanced Research Projects Agency (ARPA Order No. 8459) and the U.S. Army Topographic Engineering Center under Contract DACA76-92-C-0009 is gratefully acknowledged.

1 Introduction

Consider a robot which is navigating in a space modeled by a graph, and which wants to know its current location. It can send a signal to find out how far it is from a set of fixed landmarks. We study the problem of computing the minimum number of landmarks required, and where they should be placed, such that the robot can always determine its location. The set of nodes where the landmarks are placed is called a *metric basis* of the graph, and the number of landmarks is called the *metric dimension* of the graph.

We associate "coordinates" with each node based on the distances from the node to the landmarks. Our goal is to pick just enough landmarks so that each node has a unique tuple of coordinates. For example, in Euclidean d-space, it is easy to show that any set of $d+1$ points in general position constitutes a metric basis.

Let $G = (V, E)$ be a connected, undirected graph. A "coordinate system" on G is defined as follows. We pick a set of nodes as the metric basis; each node in the basis corresponds to a landmark. For each landmark, the coordinate of each node $v \in V$ in the corresponding "dimension" is equal to the length of a shortest path from the landmark to v. Thus for a metric basis, each node has a vector of coordinates, a tuple of non-negative integers specifying the distances to that node from the nodes in the basis.

Definition 1 *The metric dimension of the graph G is denoted by $\beta(G)$.*

For example, a path has metric dimension 1, a cycle has metric dimension 2, and a complete graph on n nodes has metric dimension $n-1$.

We first note a simple property of shortest paths on graphs.

Proposition 1.1 *Let $G = (V, E)$ be an arbitrary graph. Let u, v and w be nodes of G and let $\{u, v\} \in E$. Let d be the length of a shortest path from u to w in G. Then the length of a shortest path from v to w is one of $\{d-1, d, d+1\}$.*

Related Work: The problem of finding the metric dimension of a graph was first studied by Harary and Melter [2]. They gave a characterization for the metric dimension of trees; their proof, however, has an error (more specifically, their proof of Lemma 1 has an error). We give a similar characterization for the metric dimension of trees, and we also give a characterization for the metric dimension of d-dimensional grid graphs. We then consider graphs having small metric dimension, and show that a graph has metric dimension 1 iff it is a path.

Garey and Johnson (unpublished result, cited in [1]) proved that the problem of finding the metric dimension of a general graph is NP-complete by a reduction from 3-dimensional matching. For completeness, we provide in the appendix a reduction from 3-SAT. By providing an approximation preserving reduction to the set cover problem, we then show that the metric dimension of a graph can be approximated in polynomial time within a factor of $O(\log n)$.

2 The metric dimensions of special graphs

2.1 Trees

In this section, we study the problem of computing the metric dimension of trees. We show that this problem can be solved efficiently in linear time. Let $T = (V, E)$ be an arbitrary tree on n

nodes. We will assume that T is not just a path; we will show later that the metric dimension of a path is 1.

Definition 2 *For each node $v \in V$ of a tree $T = (V, E)$, the* number of legs *at v, denoted by ℓ_v, is the number of components which are paths, created by the removal of v from T. A single isolated node is also considered to be a path.*

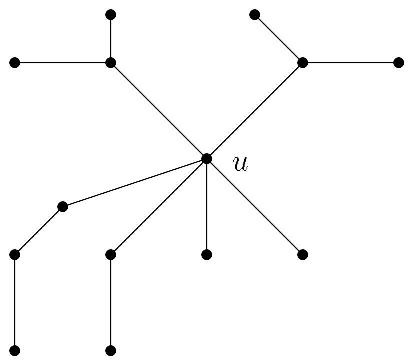

Figure 1: Example of a node with 4 legs.

For example, in Fig. 1, node u has 4 legs.

We now prove that the metric dimension of T, $\beta(T)$, is exactly

$$\sum_{v \in V : \ell_v > 1} \ell_v - 1.$$

The characterization obtained by Harary and Melter [2] is essentially the same, with a different proof.

We first obtain a lower bound on $\beta(T)$.

Lemma 2.1 *Let $T = (V, E)$ be a tree which is not a path. Then*

$$\beta(T) \geq \sum_{v \in V : \ell_v > 1} \ell_v - 1.$$

Proof. Consider any node v with $\ell_v > 1$. Observe that for any metric basis all but (at most) one of v's legs must have a landmark; otherwise the neighbors of v in those legs without landmarks have the same coordinates, making the configuration invalid. Therefore at least $\ell_v - 1$ landmarks must be placed on the legs of v. If T is not a path, the legs corresponding to different nodes (with at least two legs) are disjoint. Therefore the number of landmarks in any metric basis is at least the sum stated above. ∎

We now obtain an upper bound on $\beta(T)$ constructively.

Algorithm to place landmarks on a tree

1. Compute ℓ_v for each node v.
2. Each node v with $\ell_v > 1$ is allocated $\ell_v - 1$ landmarks. These landmarks are placed on all but one of the leaves associated with the legs of v.

It is easy to implement the above algorithm in linear time using a post-order traversal of the tree. Both steps of the algorithm can be completed in a single traversal of the tree. Also, the algorithm clearly uses the minimum number of landmarks necessary (as shown in Lemma 2.1). We now show that the algorithm generates a metric basis.

Lemma 2.2 *Let T be rooted arbitrarily. Any node v of degree greater than 2 has a descendant landmark.*

Proof. Let w be a deepest node in the subtree of v whose degree is greater than 2 (w may be the same as v). Then w has at least two legs (in the subtree of v) and at least one landmark is placed in the subtree of v. ∎

Lemma 2.3 *The above algorithm produces a valid configuration of landmarks for a given tree T (which is not a path) and uses $\sum_{q \in V}(\ell_q - 1)$ landmarks, where the sum is taken over those nodes with $\ell_q > 1$.*

Proof. Root the tree T at an arbitrary leaf r that has a landmark. We will show that for any pair of nodes u and v, there exists a landmark that distinguishes these two nodes.

Case 1 – u and v are at different distances from r: The landmark at r distinguishes u from v.

Case 2 – u and v are at the same depth and at least one of u or v has a (not necessarily proper) descendant w with degree greater than 2: By Lemma 2.2, w has a descendant landmark and this landmark distinguishes u from v.

Case 3 – u and v are at the same depth and neither has a descendant with degree greater than 2:

> *Case 3a – the path from u to v has only one node of degree greater than 2, namely $w = lca(u, v)$:* In this case, u and v are on different legs of w. Since w has at least two legs, it places landmarks on the leaves of all its legs but one. Hence at least one of these two legs receives a landmark, which distinguishes u from v.
>
> *Case 3b – there is a node x different from $w = lca(u, v)$ on the path from u to v and the degree of x is greater than 2:* The node x must have a descendant landmark which distinguishes u from v (note that u and v are at the same depth and hence w is equidistant from u and v).

∎

Theorem 2.4 *Let $T = (V, E)$ be a tree which is not a path. Then*

$$\beta(T) = \sum_{v \in V : \ell_v > 1} \ell_v - 1.$$

Proof. By Lemma 2.1, this sum is a lower bound on $\beta(T)$. Lemma 2.3 shows that the same sum is also an upper bound on $\beta(T)$. ∎

2.2 Grid graphs

We now study grid graphs formed by integer lattice points in a bounded d-dimensional space. Let us assume that the size of the grid is $D_1 \times D_2 \times \ldots \times D_d$.

Theorem 2.5 *The metric dimension of a d-dimensional grid ($d \geq 2$) is d.*

Proof. Assume we give each node a position vector which is its location in the integer lattice. We place the landmarks at the following positions. The landmark b_0 is kept at the origin $(0,0,\ldots,0)$. Let X_i be the node for which the i^{th} component of its position vector is D_i, with all other components being 0. The landmark $b_i, 1 \leq i \leq d-1$ is kept at node X_i.

We will now show that each node gets a unique coordinate tuple based on its distances from the set of landmarks. Let the distance of node v, with position vector $(x_1, x_2, \ldots, x_d)$, from landmark b_i be $d_i (0 \leq i \leq d-1)$. We get the following equations:

$$x_1 + x_2 + \ldots + x_d = d_0$$

$$(D_1 - x_1) + x_2 + \ldots + x_d = d_1$$

$$x_1 + (D_2 - x_2) + \ldots + x_d = d_2$$

$$x_1 + x_2 \ldots + (D_{d-1} - x_{d-1}) + x_d = d_{d-1}$$

It is not difficult to see that solving these equations yields a unique solution for the position vector of node v. Hence each node has distinct coordinates. We leave it for the reader to see why d is a lower bound on the metric dimension. □

3 Graphs with small metric dimension

We first investigate graphs that require only a few landmarks. We show that paths are the only graphs with $\beta = 1$. We then investigate a few properties of graphs with $\beta = 2$.

3.1 Graphs with metric dimension 1

Graphs that require only a single landmark are clearly simple in nature. We characterize them exactly.

Theorem 3.1 *A graph $G = (V, E)$ has $\beta = 1$ iff G is a path.*

Proof. We give a proof by contradiction. Let G be a graph with $\beta = 1$. Let the landmark node be vertex u of G. First observe that the degree of u is 1; otherwise the nodes adjacent to u will have the same coordinate of 1. Suppose G is not a path. Then it contains a node v whose degree is at least 3. Let $N = \{v_1, v_2, \ldots, v_k\}$ be the neighbors of v. Since there is only one landmark, every node has a single coordinate (distance from the landmark). Let d be the coordinate of v. By Proposition 1.1, the coordinates of each of the nodes in N is one of $\{d-1, d, d+1\}$. None of the nodes in N may be at a distance d from the landmark since the coordinate d is taken by v. Therefore, since $|N| \geq 3$, at least two nodes in N have the same coordinate. This is a contradiction because we assumed that $\beta(G) = 1$.

We now show that if G is a path then $\beta(G) = 1$. Let a landmark be placed at one of the two ends of the path. It is easily verified that this is a metric basis of the graph. □

3.2 Graphs with metric dimension 2

Graphs with $\beta = 2$ have a richer structure. We study a few properties of such graphs. We show that these graphs contain neither K_5 nor $K_{3,3}$ as a subgraph. This might lead one to conjecture that such graphs have to be planar; but we will exhibit a non-planar graph with metric dimension 2.

Theorem 3.2 *A graph G with $\beta(G) = 2$ cannot have K_5 as a subgraph.*

Proof. Consider a graph G with K_5 as a subgraph. Let the nodes of the subgraph be $v_1, \ldots, v_5$. Suppose two landmarks are sufficient for G. Since every pair of nodes in $v_1, \ldots, v_5$ are adjacent in G, by Proposition 1.1, the first coordinate of these nodes must be one of $\{y, y+1\}$ for some integer y. Similarly the second coordinate of the nodes is one of $\{z, z+1\}$ for some z. With these coordinates, there are only four distinct coordinates for the five nodes, thus making the configurations of the landmarks invalid. ∎

Theorem 3.3 *A graph G with $\beta(G) = 2$ cannot have $K_{3,3}$ as a subgraph.*

Proof. Assume for contradiction that $K_{3,3}$ is present as a subgraph and that there is a metric basis of size two. All nodes have been given distinct coordinates. Let the nodes of $K_{3,3}$ be $\{v_1, v_2, v_3\}$ and $\{v_4, v_5, v_6\}$ with edges going across from one set of nodes to the other. Among these six nodes, let node v_4 have the smallest first coordinate and have coordinates (a, b). Nodes $\{v_1, v_2, v_3\}$ must all have first coordinate either a or $a + 1$.

1. Suppose all three are $a + 1$. The second coordinates must be $\{b-1, b, b+1\}$ (in some order). This forces the second coordinates of nodes v_5 and v_6 to be b. There is no way to assign distinct coordinates to nodes $\{v_4, v_5, v_6\}$.

2. Suppose all three are a. The second coordinates must be $\{b-1, b, b+1\}$ (in some order). There are two nodes with coordinates (a, b).

3. Suppose nodes v_1 and v_2 have first coordinate a, and node v_3 has first coordinate $a + 1$. Nodes v_1 and v_2 have their second coordinates $\{b-1, b+1\}$ in some order. Clearly the second coordinate of nodes v_5 and v_6 is b. There is no way to assign distinct coordinates to nodes $\{v_4, v_5, v_6\}$.

4. Suppose node v_1 has first coordinate a, and nodes v_2 and v_3 have first coordinate $a+1$. Node v_1 can be either $(a, b+1)$ or $(a, b-1)$.

   (a) Node $v_1 = (a, b+1)$. Nodes v_2 and v_3 have to choose their second coordinates. The choices are $\{b, b-1\}$ or $\{b, b+1\}$ or $\{b+1, b-1\}$. We consider each case separately.

      (i) The second coordinate of v_5 must be b. There is no choice for the first.

      (ii) In this case nodes v_5 and v_6 have to pick from $\{a, a+1\}$ for the first coordinate and $\{b, b+1\}$ for the second coordinate. Since there are a total of four distinct choices and nodes v_1, v_2 and v_3 have used up three of them we cannot assign coordinates to v_5 and v_6.

      (iii) The second coordinate of v_5 must be b. There is no choice for the first.

   (b) Node $v_1 = (a, b-1)$. Nodes v_2 and v_3 have to choose their second coordinates. The choices are $\{b, b-1\}$ or $\{b, b+1\}$ or $\{b+1, b-1\}$. We consider each case separately.

(i) The choices for nodes v_5 and v_6 are $\{a, a+1\}$ for the first coordinate and $\{b-1, b\}$ for the second coordinate. Since there are a total of four distinct choices and nodes v_1, v_2 and v_3 have used up three of them we cannot assign coordinates to nodes v_5 and v_6.

(ii) The second coordinate of v_5 must be b. There is no choice for the first.

(iii) The second coordinate of v_5 must be b. The first coordinate is forced to be $a+1$. There is no choice for node v_6.

□

Theorem 3.4 *There are non-planar graphs with metric dimension 2.*

Proof. We give an example of a non-planar graph whose metric dimension is 2 (Fig. 2). It is easily verified that two landmarks are sufficient for this graph (place the landmarks on nodes u and v). A K_5 homeomorph of the graph is shown using bold lines, thus showing that the graph is non-planar. □

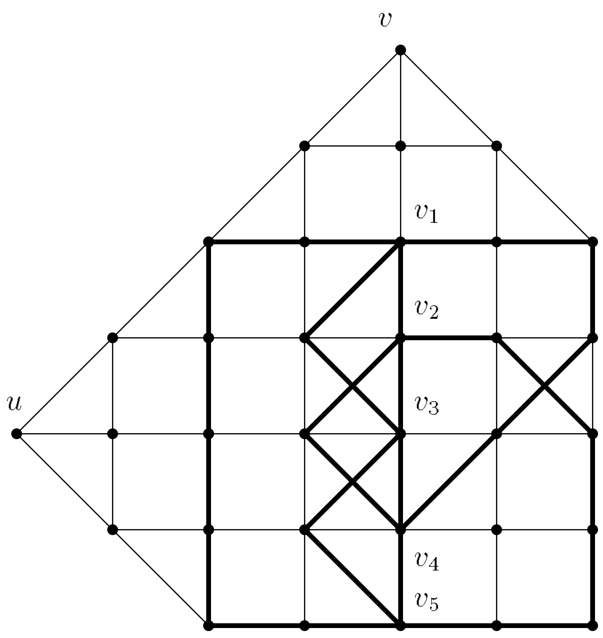

Figure 2: A non-planar graph with $\beta = 2$.

The following theorem captures a few other properties of graphs with metric dimension 2.

Theorem 3.5 *Let $G = (V, E)$ be a graph with metric dimension 2 and let $\{a, b\} \subset V$ be a metric basis in G. The following are true:*

1. *There is a unique shortest path P between a and b.*
2. *The degrees of a and b are at most 3.*
3. *Every other node on P has degree at most 5.*

Proof. Suppose there were two shortest paths P_1 and P_2 between a and b. Consider the nodes nearest to a in which P_1 and P_2 differ, i.e., distinct nodes u and v on the two shortest paths which are both equidistant from a. It is easy to verify that u and v have exactly the same coordinates, contradicting the fact that the placement of landmarks on the nodes a and b is valid. Hence the shortest path between a and b is unique.

Let the coordinates of a be $(0, x)$. All neighbors of a have 1 as their first coordinate. Therefore, applying Proposition 1.1, the coordinates of the neighbors of a must be one of $(1, x-1)$, $(1, x)$ or $(1, x+1)$. Hence the degree of a is at most 3, and analogously for b.

Consider any other node w on the shortest path between a and b. Let its coordinate be (p, q). Clearly $p + q = x$, the distance between a and b. The degree of w is at most 5 since there are no nodes with coordinates $(p-1, q-1)$ or $(p-1, q)$ or $(p, q-1)$. ☐

The following theorem gives a lower bound on the diameter of a graph with metric dimension 2.

Theorem 3.6 *Let $G = (V, E)$ be a graph with metric dimension 2. Let D be the diameter of G. Then $|V| \leq D^2 + 2$ (and hence $D = \Omega(\sqrt{n})$).*

Proof. Consider any valid configuration of two landmarks on G. Since the diameter of G is D, each coordinate of G is an integer between 0 and D. Only the two nodes on which landmarks were placed have one coordinate 0. Each of the remaining nodes must get a unique coordinate from one of D^2 possibilities. Therefore G has at most $D^2 + 2$ nodes. ☐

Remark: The bound on $|V|$ above can be slightly refined based on some of the observations we made earlier. The landmarks (nodes on which landmarks were placed) have degree at most 3; hence at most six nodes have one coordinate 1. Thus $|V| \leq (D-1)^2 + 8$.

4 Approximating the metric dimension of a graph

In this section, we show that the metric dimension of a graph can be approximated in polynomial time within a factor of $O(\log n)$. We show that there is an approximation preserving reduction from the problem of finding $\beta(G)$ to the set cover problem. We can then use the $O(\log n)$ factor approximation algorithm for the set cover problem [3, 4] to obtain an approximation algorithm for the metric dimension problem.

Theorem 4.1 *Given an arbitrary graph $G = (V, E)$, $\beta(G)$ can be approximated within a factor of $O(\log n)$ in polynomial time.*

Proof. We construct an instance of the set cover problem from G. The intuition is that every pair of distinct nodes must be distinguished by some landmark. We can easily compute all those pairs of nodes that are distinguished by placing a landmark on a given node. The metric dimension problem is that of finding a set of nodes of minimum cardinality such that every pair of nodes is distinguished by some node in this set.

The elements of the universe (in the set cover problem) correspond to pairs of nodes of G, $\{u, v\} : u \neq v$. For each node $v \in V$, we place the set of all pairs of nodes which are distinguished by placing a landmark at v into a single subset S_v. Therefore there are $\binom{n}{2}$ elements and n subsets in the set cover problem ($|V| = n$). It is easily verified that there is a set cover of size k iff there exists a metric basis of size k in G. Finding a set cover within a factor of $O(\log n)$ therefore yields the same approximation for the metric dimension problem. ☐

Acknowledgements

We would like to thank Bob Melter for pointers to previous literature regarding metric dimensions of graphs, and for providing us with a copy of [2]. Thanks also to Christine Piatko for faxing me a copy of [2], and to David Johnson for faxing us the NP-completeness proof cited in [1].

References

[1] M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, 1979.

[2] F. Harary and R.M. Melter, The metric dimension of a graph, *Ars Combinatoria*, pp. 191–195, 1976.

[3] D.S. Johnson, Approximation algorithms for combinatorial problems, *J. Comp. Sys. Sci.*, 9:256-278, 1974.

[4] L. Lovász, On the ratio of optimal integral and fractional covers, *Disc. Math.*, 13:383-390, 1975.

Appendix A NP-hardness in general graphs

We now show that the problem of finding the metric dimension of an arbitrary graph is NP-hard.

Theorem Appendix A.1 *Given an arbitrary graph $G = (V, E)$ and an integer k, deciding whether $\beta(G) \leq k$ is NP-complete.*

Proof. The problem is clearly in NP. We give the NP-hardness proof by a reduction from 3-SAT.

Consider an arbitrary input to 3-SAT, a formula F with n variables and m clauses. Let the variables be $x_1, \ldots, x_n$ and the clauses be $C_1, \ldots, C_m$.

For each variable x_i we construct a gadget as follows (see Fig. 3):

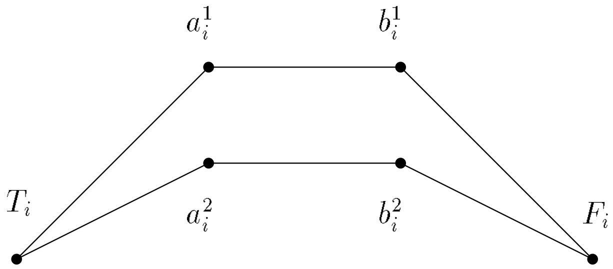

Figure 3: Gadget for a variable

The nodes T_i and F_i are the "true" and "false" ends of the gadget. The gadget is attached to the rest of the graph only through these nodes.

Suppose $C_j = y_j^1 \vee y_j^2 \vee y_j^3$, where y_j^k is a literal in clause C_j. For each such clause C_j we construct a gadget as follows (see Fig. 4).

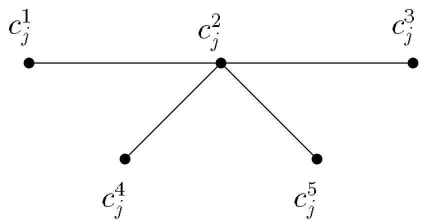

Figure 4: Gadget for a clause

We now show the connections between the clause and variable gadgets.

If a variable x_i occurs as a positive literal in clause C_j, we add the edges $\{T_i, c_j^1\}, \{F_i, c_j^1\}$ and $\{F_i, c_j^3\}$. If it occurs in C_j as a negative literal, we add the same edges, except we replace $\{F_i, c_j^3\}$ by $\{T_i, c_j^3\}$. Fig. 5 shows the edges added thus corresponding to the clause $C_j = x_1 \vee \bar{x}_2 \vee x_3$. We call these the *truth testing* edges.

For all k such that neither x_k nor $\bar{x}_k$ occur in C_j, add the following edges to the generated graph: $\{T_k, c_j^1\}, \{T_k, c_j^3\}, \{F_k, c_j^1\}, \{F_k, c_j^3\}$. The reason why these edges are added is that no matter what value is assigned to x_k (corresponds to placing a landmark at an appropriate location), this

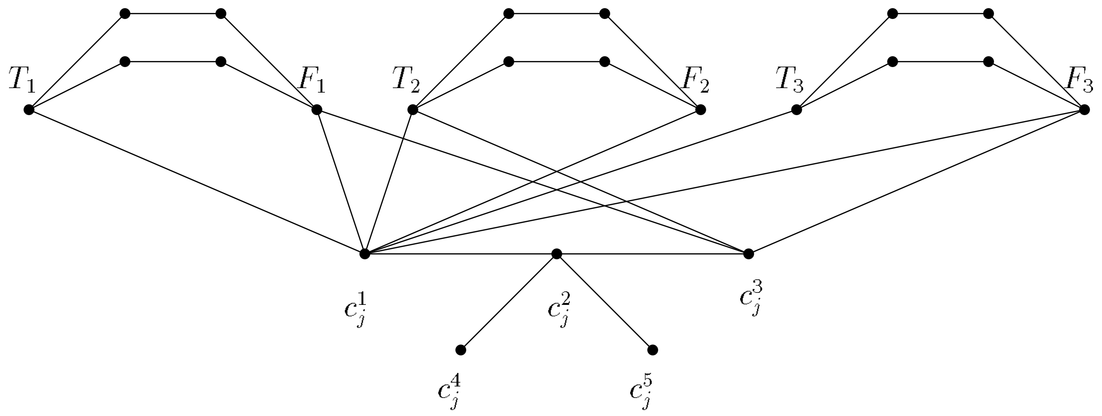

Figure 5: Clause $C_j = x_1 \vee \bar{x}_2 \vee x_3$.

gives identical coordinates to both c_j^1 and c_j^3 in the gadget corresponding to clause C_j. We call these the *neutralizing edges.*

Thus the graph G that is constructed from the formula F with n variables and m clauses has $6n + 5m$ nodes. The edges of G are variable gadget edges, clause gadget edges, truth testing edges and neutralizing edges. It is clear that given F, G can be easily constructed in polynomial time.

We will now prove that F is satisfiable if and only if the metric dimension of G is exactly $n+m$.

We will first note a few useful properties of G.

Lemma Appendix A.2 *Let x_i be an arbitrary variable in F. In any metric basis, at least one of the nodes $\{a_i^1, a_i^2, b_i^1, b_i^2\}$ must have a landmark on it.*

Proof. Suppose none of these nodes has a landmark. Since these variables are not connected to any node other than the ones shown in the variable gadget (Fig. 3), symmetry implies that a_i^1 and a_i^2 have exactly the same coordinates. This contradicts the statement of the lemma that the placement of the landmarks is valid. □

Lemma Appendix A.3 *Let C_j be an arbitrary clause in F. In any metric basis, at least one of the nodes $\{c_j^4, c_j^5\}$ must have a landmark on it.*

Proof. If there is no landmark on either of these nodes, due to symmetry, these two nodes have exactly the same coordinates. This implies that the placement of landmarks is invalid. □

Corollary Appendix A.4 *The metric dimension of G is at least $m + n$.*

Lemma Appendix A.5 *If F is satisfiable, the metric dimension of G is $m + n$.*

Proof. We know that the metric dimension is at least $m + n$. We now exhibit a metric basis of size $m + n$ based on a satisfying assignment of F.

Fix a satisfying assignment for F. For each clause C_j, place a landmark on c_j^4. For each variable x_i, if its value is `true`, place a landmark on a_i^1; otherwise place a landmark on b_i^1.

We now show that this is a metric basis. The only sets of nodes for which we need to show that they have distinct coordinates are pairs of nodes of the form $\{c_j^1, c_j^3\}$ — end nodes of the same clause gadget. For any other pair of nodes, it is easy to find a landmark which distinguishes between them.

For any clause C_j, we show that c_j^1 and c_j^3 have different coordinates if landmarks were placed based on a satisfying assignment as above. Suppose C_j is satisfied by the variable x_i, a variable occurring as a positive literal in C_j and has the value `true` in the assignment (the case when x_i occurs as a negative literal in C_j and has the value `false` is symmetric). Corresponding to x_i being `true`, we placed a landmark on a_i^1. ¿From this landmark, c_j^1 is at distance 2, while c_j^3 is at distance 3. Thus all nodes have distinct coordinates and therefore we have a metric basis of size $m+n$. □

Lemma Appendix A.6 *If the metric dimension of G is $m+n$, then F is satisfiable.*

Proof. Consider any metric basis of size $m+n$ in G. By Lemmas Appendix A.2 and Appendix A.3, we know that in any metric basis, at least one landmark must be placed within each variable and each clause gadget. Since there are exactly $m+n$ landmarks, there is exactly one landmark per variable and one landmark per clause.

We now set an assignment of the variables as follows. For each variable x_i, if the landmark on its gadget is on either a_i^1 or a_i^2, set x_i to be `true`. Otherwise set x_i to be `false`. We will now show that this yields a satisfying assignment for F.

Consider an arbitrary clause C_j. We will show that at least one of its literals is true. The main idea is in tracing which landmark distinguishes between c_j^1 and c_j^3 and showing that the corresponding variable assignment satisfies C_j.

For each clause C_k, without loss of generality, one landmark is placed on c_k^4. If $j=k$, both c_j^1 and c_j^3 are at distance 2 from c_k^4. If $j \neq k$, then due to the neutralizing edges c_j^1 and c_j^3 are at distance 4 from c_k^4. Therefore none of these landmarks distinguish c_j^1 from c_j^3.

For any variable x_p which does not occur in C_j, the landmark in the variable gadget of x_p is at distance 2 from each of c_j^1 and c_j^3. Therefore the only landmark that could distinguish between c_j^1 and c_j^3 must be on a variable x_q which occurs in C_j. Due to the manner in which we have added truth testing edges, such a landmark distinguishes between the two nodes only if one of the following two statements holds.

1. x_q occurs as a positive literal in C_j and a landmark is placed on either a_q^1 or a_q^2; in this case x_q is set to `true`.

2. x_q occurs as a negative literal in C_j and a landmark is placed on either b_q^1 or b_q^2; in this case x_q is set to `false`.

In either case, the setting of x_q is such that it satisfies C_j. □

Lemmas Appendix A.5 and Appendix A.6 together complete the reduction from SAT to the metric dimension problem. This completes the proof of Theorem Appendix A.1. □