

Reconfiguration of Dominating Sets

Akira Suzuki^{1*}, Amer E. Mouawad^{2**}, and Naomi Nishimura^{2**}

¹ Graduate School of Information Sciences, Tohoku University
Aoba-yama 6-6-05, Aoba-ku, Sendai, 980-8579, Japan.
a.suzuki@ecei.tohoku.ac.jp

² David R. Cheriton School of Computer Science
University of Waterloo, Waterloo, Ontario, Canada.
{aabdou, nishi}@uwaterloo.ca

Abstract. We explore a reconfiguration version of the dominating set problem, where a dominating set in a graph G is a set S of vertices such that each vertex is either in S or has a neighbour in S . In a reconfiguration problem, the goal is to determine whether there exists a sequence of feasible solutions connecting given feasible solutions s and t such that each pair of consecutive solutions is adjacent according to a specified adjacency relation. Two dominating sets are adjacent if one can be formed from the other by the addition or deletion of a single vertex.

For various values of k , we consider properties of $D_k(G)$, the graph consisting of a vertex for each dominating set of size at most k and edges specified by the adjacency relation. Addressing an open question posed by Haas and Seyffarth, we demonstrate that $D_{\Gamma(G)+1}(G)$ is not necessarily connected, for $\Gamma(G)$ the maximum cardinality of a minimal dominating set in G . The result holds even when graphs are constrained to be planar, of bounded tree-width, or b -partite for $b \geq 3$. Moreover, we construct an infinite family of graphs such that $D_{\gamma(G)+1}(G)$ has exponential diameter, for $\gamma(G)$ the minimum size of a dominating set. On the positive side, we show that $D_{n-\mu}(G)$ is connected and of linear diameter for any graph G on n vertices with a matching of size at least $\mu + 1$.

1 Introduction

The *reconfiguration version* of a problem determines whether it is possible to transform one feasible solution s into a *target* feasible solution t in a step-by-step manner (a *reconfiguration*) such that each intermediate solution is also feasible. The study of such problems has received considerable attention in recent literature [8, 9, 13, 15, 16] and is interesting for a variety of reasons. From an algorithmic standpoint, reconfiguration models dynamic situations in which we seek to transform a solution into a more desirable one, maintaining feasibility during the process. Reconfiguration also models questions of evolution; it can represent

* Research supported by JSPS Grant-in-Aid for Scientific Research, Grant Numbers 24.3660 and 26730001.

** Research supported by the Natural Science and Engineering Research Council of Canada.

the evolution of a genotype where only individual mutations are allowed and all genotypes must satisfy a certain fitness threshold, i.e. be feasible. Moreover, the study of reconfiguration yields insights into the structure of the solution space of the underlying problem, crucial for the design of efficient algorithms. In fact, one of the initial motivations behind such questions was to study the performance of heuristics [9] and random sampling methods [4], where connectivity and other properties of the solution space play a crucial role. Even though reconfiguration gained popularity in the last decade or so, the notion of exploring the solution space of a given problem has been previously considered in numerous settings. One such example is the work of Mayr and Plaxton [18], where the authors consider the problem of transforming one minimum spanning tree of a weighted graph into another by a sequence of edge swaps.

Some of the problems for which the reconfiguration version has been studied include vertex colouring [1, 3, 4, 6, 5], list edge-colouring [14], list $L(2,1)$ -labeling [15], block puzzles [11], independent set [11, 13], clique, set cover, integer programming, matching, spanning tree, matroid bases [13], satisfiability [9], shortest path [2, 16], subset sum [12], dominating set [10, 19], odd cycle transversal, feedback vertex set, and hitting set [19]. For most **NP**-complete problems, the reconfiguration version has been shown to be **PSPACE**-complete [13, 14, 17], while for some problems in **P**, the reconfiguration question could be either in **P** [13] or **PSPACE**-complete [2].

The problem of transforming input s into input t can be viewed as the problem of determining if there is a path from s to t in a graph representing feasible solutions. Such a path is called a *reconfiguration sequence*. For the problem of dominating set, the *k-dominating graph*, defined formally in Section 2, consists of a node for each feasible solution and an edge for each pair of solutions that differ by a single vertex. Finding an s - t path in this graph (not included in the input) has been shown to be **W[2]**-hard [19], and hence not likely to yield even a fixed-parameter tractable algorithm [7].

Although having received less attention than the s - t path problem, other characteristics of the solution graph have been studied. Determining the diameter of the reconfiguration graph will result in an upper bound on the length of any reconfiguration sequence. For a problem such as colouring, one can determine the *mixing number*, the minimum number of colours needed to ensure that the entire graph is connected; such a number has been obtained for the problem of list edge-colouring on trees [14].

In previous work on reconfiguration of dominating sets, Haas and Seyfarth [10] considered the connectivity of the graph of solutions of size at most k , for various values of k relative to n , the number of vertices in the input graph G . They demonstrated that the graph is connected when $k = n - 1$ and G has at least two non-adjacent edges, or when k is one greater than the maximum cardinality of a minimal dominating set and G is non-trivially bipartite or chordal. They left as an open question, answered negatively here, whether the latter results could be extended to all graphs.

In this paper we extend previous work by showing in Section 3 that the solution graph is connected and of linear diameter for $k = n - \mu$ for any input graph with a matching of size least $\mu + 1$, for any nonnegative integer μ . In Section 4, we give a series of counterexamples demonstrating that $D_{\Gamma(G)+1}(G)$ is not guaranteed to be connected for planar graphs, graphs of bounded treewidth, or b -partite graphs for $b \geq 3$. In Section 5, we pose and answer a question about the diameter of $D_{\gamma(G)+1}(G)$ by showing that there is an infinite family of graphs of exponential diameter.

2 Preliminaries

We assume that each G is a simple, undirected graph on n vertices with vertex set $V(G)$ and edge set $E(G)$. The *diameter* of G is the maximum over all pairs of vertices u and v in $V(G)$ of the length of the shortest path between u and v .

A set $S \subseteq V(G)$ is a *dominating set* of G if and only if every vertex in $V(G) \setminus S$ is adjacent to a vertex in S . The minimum cardinality of any dominating set of G is denoted by $\gamma(G)$. Similarly, $\Gamma(G)$ is the maximum cardinality of any minimal dominating set in G .

For a vertex $u \in V(G)$ and a dominating set S of G , we say u is *dominated* by $v \in S$ if $u \notin S$ and u is adjacent to v . For a vertex v in a dominating set S , a *private neighbour* of v is a vertex dominated by v and not dominated by any other vertex in S ; the *private neighbourhood* of v is the set of private neighbours of v . A vertex v in a dominating set S is *deletable* if $S \setminus \{v\}$ is also a dominating set of G .

Fact 1 *A vertex v is deletable if and only if v has at least one neighbour in S and v has no private neighbour.*

Given a graph G and a positive integer k , we consider the k -*dominating graph* of G , $D_k(G)$, such that each vertex in $V(D_k(G))$ corresponds to a dominating set of G of cardinality at most k . Two vertices are adjacent in $D_k(G)$ if and only if the corresponding dominating sets differ by either the addition or the deletion of a single vertex; each such operation is a *reconfiguration step*. Formally, if A and B are dominating sets of G of cardinality at most k , then there exists an edge between A and B if and only if there exists a vertex $u \in V(G)$ such that $(A \setminus B) \cup (B \setminus A) = \{u\}$. We refer to vertices in G using lower case letters (e.g. u, v) and to the vertices in $D_k(G)$, and by extension their associated dominating sets, using upper case letters (e.g. A, B). We write $A \leftrightarrow B$ if there exists a path in $D_k(G)$ joining A and B . The following fact is a consequence of our ability to add vertices as needed to form B from A .

Fact 2 *For dominating sets A and B , if $A \subseteq B$, then $A \leftrightarrow B$ and $B \leftrightarrow A$.*

3 Graphs with a matching of size $\mu + 1$

Theorem 1. *For any nonnegative integer μ , if G has a matching of size at least $\mu + 1$, then $D_{n-\mu}(G)$ is connected for $n = |V(G)|$.*

Proof. For G a graph with matching $M = \{\{u_i, w_i\} \mid 0 \leq i \leq \mu\}$, we define $U = \{u_i \mid 0 \leq i \leq \mu\}$, $W = \{w_i \mid 0 \leq i \leq \mu\}$, and the set of *outsiders* $R = V(G) \setminus (U \cup W)$.

Using any dominating set S of G , we classify edges in M as follows: edge $\{u_i, w_i\}$, $0 \leq i \leq \mu$, is *clean* if neither u_i nor w_i is in S , *u-odd* if $u_i \in S$ but $w_i \notin S$, *w-odd* if $w_i \in S$ but $u_i \notin S$, *odd* if $\{u_i, w_i\}$ is *u-odd* or *w-odd*, and *even* if $\{u_i, w_i\} \subseteq S$. We use $\text{clean}(S)$ and $\text{odd}(S)$, respectively, to denote the numbers of clean and odd edges for S . Similarly, we let $u\text{-odd}(S)$ and $w\text{-odd}(S)$ denote the numbers of *u-odd* and *w-odd* edges for S . In the example graph shown in Figure 1, $\mu + 1 = 7$ and $R = \emptyset$. There is a single clean edge, namely $\{u_1, w_1\}$, three *w-odd* edges, two *u-odd* edges, and a single even edge.

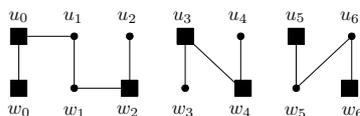


Fig. 1. Vertices in S are marked with squares.

It suffices to show that for S an arbitrary dominating set of G such that $|S| \leq n - \mu$, $S \leftrightarrow N$ for $N = V(G) \setminus W$; N is clearly a dominating set as each vertex $w_i \in W = V(G) \setminus N$ is dominated by u_i . By Fact 2, for S' a dominating set of G such that $S' \supseteq S$ and $|S'| = n - \mu$, since S' is a superset of S , then $S \leftrightarrow S'$. The reconfiguration from S' to N can be broken into three stages. In the first stage, for a dominating set S_0 with no clean edges, we show $S' \leftrightarrow S_0$ by repeatedly decrementing the number of clean edges (u_i or w_i is added to the dominating set for some $0 \leq i \leq \mu$). In the second stage, for T_μ with μ *u-odd* edges and one even edge, we show $S_0 \leftrightarrow T_\mu$ by repeatedly incrementing the number of *u-odd* edges. Finally, we observe that deleting the single remaining element in $T_\mu \cap W$ yields $T_\mu \leftrightarrow N$.

In stage 1, for $x = \text{clean}(S')$, we show that $S' = S_x \leftrightarrow S_{x-1} \leftrightarrow S_{x-2} \leftrightarrow \dots \leftrightarrow S_0$ where for each $0 \leq j \leq x$, S_j is a dominating set of G such that $|S_j| = n - \mu$ and $\text{clean}(S_j) = j$. To show that $S_a \leftrightarrow S_{a-1}$ for arbitrary $1 \leq a \leq x$, we prove that there is a deletable vertex in some even edge and hence a vertex in a clean edge can be added in the next reconfiguration step. For $b = \text{odd}(S_a)$, the set E of vertices in even edges is of size $2((\mu + 1) - a - b)$. Since each vertex in E has a neighbour in S_a , if at least one vertex in E does not have a private neighbour, then E contains a deletable vertex (Fact 1).

The μ vertices in $V(G) \setminus S_a$ are the only possible candidates to be private neighbours. Of these, the b vertices of $V(G) \setminus S_a$ in odd edges cannot be private neighbours of vertices in E , as each is the neighbour of a vertex in $S_a \setminus E$ (the other endpoint of the edge). The number of remaining candidates, $\mu - b$, is smaller than the number of vertices in E ; $\mu \geq 2a + b$ as the vertices of $V(G) \setminus S_a$ must contain both endpoints of any clean edge and one endpoint for any odd edge.

Hence, there exists at least one deletable vertex in E . When we delete such a vertex and add an arbitrary endpoint of a clean edge, the clean edge becomes an odd edge and the number of clean edges decreases. We can therefore reconfigure from S_a to the desired dominating set, and by applying the same argument a times, to S_0 .

In the second stage we show that for $y = u\text{-odd}(S_0)$, $S_0 = T_y \leftrightarrow T_{y+1} \leftrightarrow T_{y+2} \leftrightarrow \dots \leftrightarrow T_\mu$ where for each $y \leq j \leq \mu$, T_j is a dominating set of G such that $|T_j| = n - \mu$, $\text{clean}(T_j) = 0$, and $u\text{-odd}(T_j) = j$. To show that $T_c \leftrightarrow T_{c+1}$ for arbitrary $y \leq c \leq \mu - 1$, we use a counting argument to find a vertex in an even edge that is in W and deletable; in one reconfiguration step the vertex is deleted, increasing the number of u -odd edges, and in the next reconfiguration step an arbitrary vertex in R or in a w -odd edge is added to the dominating set. We let $d = w\text{-odd}(T_c)$ (i.e. the number of w -odd edges for T_c) and observe that since there are c u -odd edges, d w -odd edges, and no clean edges, there exist $(\mu + 1) - c - d$ even edges. We define E_w to be the set of vertices in W that are in the even edges, and observe that each has a neighbour in T_c ; a vertex in E_w will be deletable if it does not have a private neighbour.

Of the μ vertices in $V(G) \setminus T_c$, only those in R are candidates to be private neighbours of vertices in E_w , as each vertex in an odd edge has a neighbour in T_c . As there are c u -odd edges and d w -odd edges, the total number of vertices in $R \cap V(G) \setminus T_c$ is $\mu - c - d$. Since this is smaller than the number of vertices in E_w , at least one vertex in E_w must be deletable. When we delete such a vertex from T_c and in the next step add an arbitrary vertex from the outsiders or w -odd edges, the even edge becomes a u -odd edge and the number of u -odd edges increases. Note that we can always find such a vertex since there are $\mu - c - d$ outsiders, d w -odd edges, and $c \leq \mu - 1$. Hence, we can reconfigure from T_c to T_{c+1} , and by $\mu - c$ repetitions, to T_μ . \square

Corollary 1 results from the length of the reconfiguration sequence formed in Theorem 1; reconfiguring to S' can be achieved in at most $n - \mu$ steps, and stages 1 and 2 require at most 2μ steps each, as $\mu \in O(n)$ is at most the numbers of clean and u -odd edges. Theorem 2 shows that Theorem 1 is tight.

Corollary 1. *The diameter of $D_{n-\mu}(G)$ is in $O(n)$ for G a graph with a matching of size $\mu + 1$.*

Theorem 2. *For any nonnegative integer μ , there exists a graph G_μ with a matching of size μ such that $D_{n-\mu}(G_\mu)$ is not connected.*

Proof. Let G_μ be a path on $n = 2\mu$ vertices. Clearly, G_μ has μ disjoint edges, $n - \mu = 2\mu - \mu = \mu$, and $D_{n-\mu}(G_\mu) = D_\mu(G_\mu)$. We let S be a dominating set of G_μ such that $|S| \geq \mu + 1$. At least one vertex in S must have all its neighbors in S and is therefore deletable. It follows that $\Gamma(G_\mu) = \mu$ and $D_{n-\mu}(G_\mu) = D_\mu(G_\mu) = D_{\Gamma(G_\mu)}(G_\mu)$ which is not connected by the result of Haas and Seyffarth [10, Lemma 3]. \square

4 $D_{\Gamma(G)+1}(G)$ may not be connected

In this section we demonstrate that $D_{\Gamma(G)+1}(G)$ is not connected for an infinite family of graphs $G_{(d,b)}$ for all positive integers $b \geq 3$ and $d \geq 2$, where graph $G_{(d,b)}$ is constructed from $d+1$ cliques of size b . We demonstrate using the graph $G_{(4,3)}$ as shown in part (a) of Figure 2, consisting of fifteen vertices partitioned into five cliques of size 3: the *outer clique* C_0 , consisting of the top, left, and right *outer vertices* o_1 , o_2 , and o_3 , and the four *inner cliques* C_1 through C_4 , ordered from left to right. We use $v_{(i,1)}$, $v_{(i,2)}$, and $v_{(i,3)}$ to denote the top, left, and right vertices in clique C_i , $1 \leq i \leq 4$. More generally, a graph $G_{(d,b)}$ has $d+1$ b -cliques C_i for $0 \leq i \leq d$. The clique C_0 consists of outer vertices o_j for $1 \leq j \leq b$, and for each inner clique C_i , $1 \leq i \leq d$ and each $1 \leq j \leq b$, there exists an edge $\{o_j, v_{(i,j)}\}$.

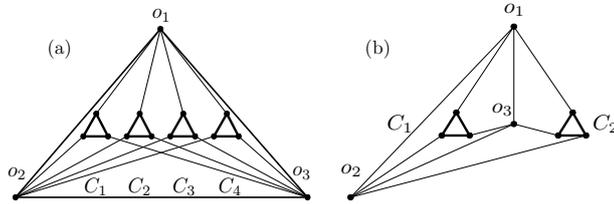


Fig. 2. Counterexamples for (a) general and (b) planar graphs.

For any $1 \leq j \leq b$ a dominating set does not contain o_j , then the vertices $v_{(i,j)}$ of the inner cliques must be dominated by vertices in the inner cliques (hence Fact 3). In addition, the outer vertex o_j can be dominated only by another outer vertex or some vertex $v_{(i,j)}$, $1 \leq i \leq d$ (hence Fact 4).

Fact 3 *Any dominating set that does not contain all of the outer vertices must contain at least one vertex from each of the inner cliques.*

Fact 4 *Any dominating set that does not contain any outer vertex must contain at least one vertex of the form $v_{(\cdot,j)}$ for each $1 \leq j \leq b$.*

Lemma 1. *For each graph $G_{(d,b)}$ as defined above, $\Gamma(G_{(d,b)}) = d + b - 2$.*

Proof. We first demonstrate that there is a minimal dominating set of size $d + b - 2$, consisting of $\{v_{(1,j)} \mid 2 \leq j \leq b\} \cup \{v_{(i,1)} \mid 2 \leq i \leq d\}$. The first set dominates $b - 1$ of the outer vertices and the first inner clique and the second set dominates o_1 and the rest of the inner cliques. The dominating set is minimal, as the removal of any vertex $v_{(1,j)}$, $2 \leq j \leq b$, would leave vertex o_j with no neighbour in the dominating set and the removal of any $v_{(i,1)}$, $2 \leq i \leq d$, would leave $\{v_{(i,j)} \mid 1 \leq j \leq b\}$ with no neighbour in the dominating set.

By Fact 3, any dominating set that does not contain all outer vertices must contain at least one vertex in each of the d inner cliques. Since the outer vertices

form a minimal dominating set, any other minimal dominating set must contain at least one vertex from each of the inner cliques.

We now consider any dominating set S of size at least $d + b - 1$ containing one vertex for each inner clique and show that it is not minimal. If S contains at least one outer vertex, we can find a smaller dominating set by removing all but the outer vertex and one vertex for each inner clique, yielding a total of $d + 1 < d + b - 1$ vertices (since $b \geq 3$). Now suppose that S consists entirely of inner vertices. By Fact 4, S contains at least one vertex of the form $v_{(\cdot, j)}$ for each $1 \leq j \leq b$. Moreover, for at least one value $1 \leq j' \leq b$, there exists more than one vertex of the form $v_{(\cdot, j')}$ as $d + b - 1 > b$. This allows us to choose b vertices of the form $v_{(\cdot, j)}$ for each $1 \leq j \leq b$ that dominate at least two inner cliques as well as all outer vertices. By selecting one member of S from each of the remaining $d - 2$ inner cliques, we form a dominating set of size $d + b - 2 < d + b - 1$, proving that S is not minimal. \square

Theorem 3. *There exists an infinite family of graphs such that for each G in the family, $D_{\Gamma(G)+1}(G)$ is not connected.*

Proof. For any positive integers $b \geq 3$ and $d \geq 2$, we show that there is no path between dominating sets A to B in $D_{d+b-1}(G_{(d,b)})$, where A consists of the vertices in the outer clique and B consists of $\{v_{(i,\ell)} \mid 1 \leq i \leq d, 1 \leq \ell \leq b, i \equiv \ell \pmod{b}\}$.

By Fact 3, before we can delete any of the vertices in A , we need to add one vertex from each of the inner cliques, resulting in a dominating set of size $d + b = \Gamma(G_{(d,b)}) + 2$. As there is no such vertex in our graph, there is no way to connect A and B . \square

Each graph $G_{(d,b)}$ constructed for Theorem 3 is a b -partite graph; we can partition the vertices into b independent sets, where the j th set, $1 \leq j \leq b$ is defined as $\{v_{(i,j)} \mid 1 \leq i \leq d\} \cup \{o_i \mid 1 \leq i \leq d, i \equiv j + 1 \pmod{b}\}$. Moreover, we can form a tree decomposition of width $2b - 1$ of $G_{(d,b)}$, for all positive integers $b \geq 3$ and $d \geq b$, by creating bags with the vertices of the inner cliques and adding all outer vertices to each bag.

Corollary 2. *For every positive integer $b \geq 3$, there exists an infinite family of graphs of tree-width $2b - 1$ such that for each G in the family, $D_{\Gamma(G)+1}(G)$ is not connected, and an infinite family of b -partite graphs such that for each G in the family, $D_{\Gamma(G)+1}(G)$ is not connected.*

Theorem 3 does not preclude the possibility that when restricted to planar graphs or any other graph class that excludes $G_{(d,b)}$, $D_{\Gamma(G)+1}(G)$ is connected. However, the next corollary follows directly from the fact that $G_{(2,3)}$ is planar (part (b) of Figure 2).

Corollary 3. *There exists a planar graph G for which $D_{\Gamma(G)+1}(G)$ is not connected.*

5 On the diameter of $D_k(G)$

In this section, we obtain a lower bound on the diameter of the k -dominating graph of a family of graphs G_n . We describe G_n in terms of several component subgraphs, each playing a role in forcing the reconfiguration of dominating sets.

A *linkage gadget* (part (a), Figure 3) consists of five vertices, the *external vertices* (or endpoints) e_1 and e_2 , and the *internal vertices* i_1 , i_2 , and i_3 . The external vertices are adjacent to each internal vertex as well as to each other; the following results from the internal vertices having degree two:

Fact 5 *In a linkage gadget, the minimum dominating sets of size one are $\{e_1\}$ and $\{e_2\}$. Any dominating set containing an internal vertex must contain at least two vertices. Any dominating set in a graph containing m vertex-disjoint linkage gadgets with all internal vertices having degree exactly two must contain at least one vertex in each linkage gadget.*

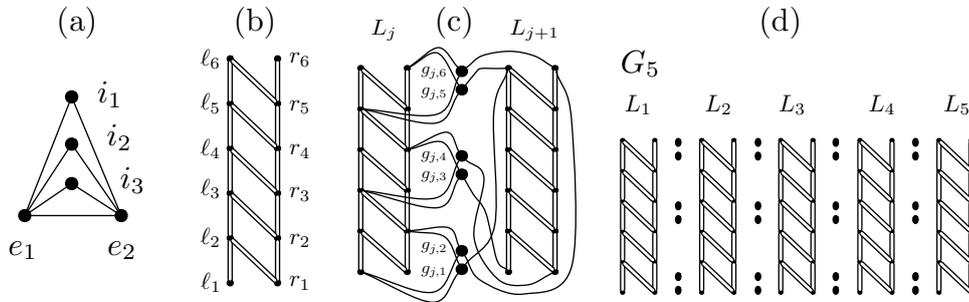


Fig. 3. Parts of the construction.

A *ladder* (part (b) of Figure 3, linkages shown as double edges) is a graph consisting of twelve *ladder vertices* paired into six *rungs*, where rung i consists of the vertices l_i and r_i for $1 \leq i \leq 6$, as well as the 45 internal vertices of fifteen linkage gadgets. Each linkage gadget is associated with a pair of ladder vertices, where the ladder vertices are the external vertices in the linkage gadget. The fifteen pairs are as follows: ten *vertical pairs* $\{l_i, l_{i+1}\}$ and $\{r_i, r_{i+1}\}$ for $1 \leq i \leq 5$, and five *cross pairs* $\{l_{i+1}, r_i\}$ for $1 \leq i \leq 5$. For convenience, we refer to vertices l_i , $1 \leq i \leq 6$ and the associated linkage gadgets as the *left side of the ladder* and to vertices r_i , $1 \leq i \leq 6$ and the associated linkage gadgets as the *right side of the ladder*, or collectively as the *sides of the ladder*.

The graph G_n consists of n ladders L_1 through L_n and $n - 1$ sets of *gluing vertices*, where each set consists of three *clusters* of two vertices each. For $l_{j,i}$ and $r_{j,i}$, $1 \leq i \leq 6$, the ladder vertices of ladder L_j , and $g_{j,1}$ through $g_{j,6}$ the gluing vertices that join ladders L_j and L_{j+1} , we have the following connections for $1 \leq j \leq n - 1$:

- Edges connecting the *bottom cluster* to the bottom two rungs of ladder L_j and the top rung of ladder L_{j+1} : $\{\ell_{j,1}, g_{j,1}\}, \{\ell_{j,1}, g_{j,2}\}, \{r_{j,2}, g_{j,1}\}, \{r_{j,2}, g_{j,2}\}, \{\ell_{j+1,6}, g_{j,1}\}, \{r_{j+1,6}, g_{j,2}\}$.
- Edges connecting the *middle cluster* to the middle two rungs of ladder L_j and the bottom rung of ladder L_{j+1} : $\{\ell_{j,3}, g_{j,3}\}, \{\ell_{j,3}, g_{j,4}\}, \{r_{j,4}, g_{j,3}\}, \{r_{j,4}, g_{j,4}\}, \{\ell_{j+1,1}, g_{j,3}\}, \{r_{j+1,1}, g_{j,4}\}$.
- Edges connecting the *top cluster* to the top two rungs of ladder L_j and the top rung of ladder L_{j+1} : $\{\ell_{j,5}, g_{j,5}\}, \{\ell_{j,5}, g_{j,6}\}, \{r_{j,6}, g_{j,5}\}, \{r_{j,6}, g_{j,6}\}, \{\ell_{j+1,6}, g_{j,5}\}, \{r_{j+1,6}, g_{j,6}\}$.

Figure 3 parts (c) and (d) show details of the construction of G_n ; they depict, respectively, two consecutive ladders and G_5 , both with linkages represented as double edges. When clear from context, we sometimes use single subscripts instead of double subscripts to refer to the vertices of a single ladder.

We let $\mathcal{D} = \{\{\ell_{(j,2i-1)}, \ell_{(j,2i)}\}, \{r_{(j,2i-1)}, r_{(j,2i)}\} \mid 1 \leq i \leq 3, 1 \leq j \leq n\}$ denote a set of $6n$ pairs in G_n ; the corresponding linkage gadgets are vertex-disjoint. Then Fact 5 implies the following:

Fact 6 *Any dominating set S of G_n must contain at least one vertex of each of the linkage gadgets for vertical pairs in the set \mathcal{D} and hence is of size at least $6n$; if S contains an internal vertex, then $|S| > 6n$.*

Choosing an arbitrary external vertex for each vertical pair does not guarantee that all vertices on the side of a ladder are dominated; for example, the set $\{\ell_i \mid i \in \{1, 4, 5\}\}$ does not dominate the internal vertices in the vertical pair $\{\ell_2, \ell_3\}$. Choices that do not leave such gaps form the set $\mathcal{C} = \{\mathcal{C}_i \mid 1 \leq i \leq 4\}$ where $\mathcal{C}_1 = \{1, 3, 5\}$, $\mathcal{C}_2 = \{2, 3, 5\}$, $\mathcal{C}_3 = \{2, 4, 5\}$, and $\mathcal{C}_4 = \{2, 4, 6\}$.

Fact 7 *In any dominating set S of size $6n$ and in any ladder L in G_n , the restriction of S to L must be of the form S_i for some $1 \leq i \leq 7$, as illustrated in Figure 4.*

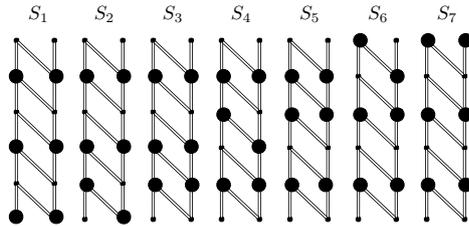


Fig. 4. Minimum dominating sets for G_1 .

Proof. Fact 6 implies that the only choices for the left (right) vertices are $\{\ell_i \mid i \in \mathcal{C}_j\}$ ($\{r_i \mid i \in \mathcal{C}_j\}$) for $1 \leq j \leq 4$. The sets S_i , $1 \leq i \leq 7$, are the only combinations of these choices that dominate all the internal vertices in the cross pairs. \square

We say that ladder L_j is in state S_i if the restriction of the dominating set to L_j is of the form S_i , for $1 \leq j \leq n$ and $1 \leq i \leq 7$.

The exponential lower bound in Theorem 4 is based on counting how many times each ladder is modified from S_1 to S_7 or vice versa; we say ladder L_j undergoes a *switch* for each such modification. We first focus on a single ladder.

Fact 8 *For S a dominating set of G_1 , a vertex $v \in S$ is deletable if and only if either v is the internal vertex of a linkage gadget one of whose external vertices is in S , or for every linkage gadget containing v as an external vertex, either the other external vertex is also in S or all internal vertices are in S .*

Lemma 2. *In $D_{\gamma(G_1)+1}(G_1)$ there is a single reconfiguration sequence between S_1 and S_7 , of length 12.*

Proof. We define P to be the path in the graph corresponding to the reconfiguration sequence $S_1 \leftrightarrow S_1 \cup \{\ell_2\} \leftrightarrow S_2 \leftrightarrow S_2 \cup \{r_2\} \leftrightarrow S_3 \leftrightarrow S_3 \cup \{\ell_4\} \leftrightarrow S_4 \leftrightarrow S_4 \cup \{r_4\} \leftrightarrow S_5 \leftrightarrow S_5 \cup \{\ell_6\} \leftrightarrow S_6 \leftrightarrow S_6 \cup \{r_6\} \leftrightarrow S_7$ and demonstrate that there is no shorter path between S_1 and S_7 .

By Facts 7 and 6, G_1 has exactly seven dominating sets of size six, and any dominating set S of size seven contains two vertices from one vertical pair d in \mathcal{D} and one from each of the remaining five. The neighbours of S in $D_{\gamma(G_1)+1}(G_1)$ are the vertices corresponding to the sets S_i , $1 \leq i \leq 7$, obtained by deleting a single vertex of S . The number of neighbours is thus at most two, depending on which, if any, vertices in d are deletable.

If at least one of the vertices of S in d is an internal vertex, then at most one vertex satisfies the first condition in Fact 8. Thus, for S to have two neighbours, there must be a ladder vertex that satisfies the second condition of Fact 8, which by inspection of Figure 4 can be seen to be false.

If instead d contains two ladder vertices, in order for S to have two neighbours, the four ladder vertices on the side containing d must correspond to the union of two of the sets in \mathcal{C} . There are only three such unions, $\mathcal{C}_1 \cup \mathcal{C}_2$, $\mathcal{C}_2 \cup \mathcal{C}_3$, and $\mathcal{C}_3 \cup \mathcal{C}_4$, which implies that the only pairs with common neighbours are $\{S_i, S_{i+1}\}$ for $1 \leq i \leq 6$, as needed to complete the proof. \square

For $n > 2$, we cannot reconfigure ladders independently from each other, as we need to ensure that all gluing vertices are dominated. For consecutive ladders L_j and L_{j+1} , any cluster that is not dominated by L_j must be dominated by L_{j+1} ; the bottom, middle, and top clusters are not dominated by any vertex in S_2 , S_4 , and S_6 , respectively.

Fact 9 *In any dominating set S of G_n , for any $1 \leq j < n$, if L_j is in state S_2 , then L_{j+1} is in state S_7 ; if L_j is in state S_4 , then L_{j+1} is in state S_1 ; and if L_j is in state S_6 , then L_{j+1} is in state S_7 .*

Lemma 3. *For any reconfiguration sequence in which L_j and L_{j+1} are initially both in state S_1 , if L_j undergoes p switches then L_{j+1} must undergo at least $2p + 1$ switches.*

Proof. We use a simple counting argument. When $p = 1$, the result follows immediately from Fact 9 since L_j can only reach state S_7 if L_{j+1} is reconfigured from S_1 to S_7 to S_1 and finally back to S_7 . After the first switch of L_j , both ladders are in state S_7 .

For any subsequent switch of L_j , L_j starts in state S_7 because for L_j to reach S_1 from S_2 or to reach S_7 from S_6 , by Fact 9 L_{j+1} must have been in S_7 . Since by definition L_j starts in S_1 or S_7 , to enable L_j to undergo a switch, L_{j+1} will have to undergo at least two switches, namely S_7 to S_1 and back to S_7 . \square

Theorem 4. *For S a dominating set of G_n such that every ladder of G_n is in state S_1 and T a dominating set of G_n such that every ladder of G_n is in state S_7 , the length of any reconfiguration sequence between S and T is at least $12(2^{n+1} - n - 2)$.*

Proof. We first observe that Lemma 2 implies that the switch of any ladder requires at least twelve reconfiguration steps; since the vertex associated with a dominating set containing a gluing vertex will have degree at most one in the k -dominating graph, there are no shortcuts formed.

To reconfigure from S to T , ladder L_1 must undergo at least one switch. By Lemma 3, ladder L_2 will undergo at least $3 = 2^2 - 1$ switches, hence $2^j - 1$ switches for ladder L_j , $1 \leq j \leq n$. Since each switch requires twelve steps, the total number of steps is thus at least $12 \sum_{i=1}^n (2^i - 1) = 12(2^{n+1} - n - 2)$. \square

Corollary 4. *There exists an infinite family of graphs such that for each graph G_n in the family, $D_{\gamma(G_n)+1}(G_n)$ has diameter $\Omega(2^n)$.*

6 Conclusions and future work

In answering Haas and Seyffarth's question concerning the connectivity of $D_k(G)$ for general graphs and $k = \Gamma(G) + 1$, we have demonstrated infinite families of planar, bounded treewidth, and b -partite graphs for which the k -dominating graph is not connected. It remains to be seen whether k -dominating graphs are connected for graphs more general than non-trivially bipartite graphs or chordal graphs, and whether $D_{\Gamma(G)+2}(G)$ is connected for all graphs. It would also be useful to know if there is a value of k for which $D_k(G)$ is guaranteed not to have exponential diameter. Interestingly, for our connectivity and diameter examples, incrementing the size of the sets by one is sufficient to break the proofs.

References

1. Marthe Bonamy and Nicolas Bousquet. Recoloring bounded treewidth graphs. In *Proc. of the 7th Latin-American Algorithms, Graphs, and Optimization Symp.*, 2013.

2. P. Bonsma. The complexity of rerouting shortest paths. In *Proc. of Mathematical Foundations of Computer Science*, pages 222–233, 2012.
3. P. Bonsma and L. Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theor. Comput. Sci.*, 410(50):5215–5226, 2009.
4. L. Cereceda, J. van den Heuvel, and M. Johnson. Connectedness of the graph of vertex-colourings. *Discrete Math.*, 308(56):913–919, 2008.
5. L. Cereceda, J. van den Heuvel, and M. Johnson. Finding paths between 3-colorings. *J. of Graph Theory*, 67(1):69–82, 2011.
6. Luis Cereceda, Jan van den Heuvel, and Matthew Johnson. Mixing 3-colourings in bipartite graphs. *European J. of Combinatorics*, 30(7):1593–1606, 2009.
7. R. G. Downey and M. R. Fellows. *Parameterized complexity*. Springer-Verlag, New York, 1997.
8. Gerd Fricke, Sandra Mitchell Hedetniemi, Stephen T. Hedetniemi, and Kevin R. Hutson. γ -Graphs of Graphs. *Discussiones Mathematicae Graph Theory*, 31(3):517–531, 2011.
9. P. Gopalan, P. G. Kolaitis, E. N. Maneva, and C.H. Papadimitriou. The connectivity of boolean satisfiability: computational and structural dichotomies. *SIAM J. on Computing*, 38(6):2330–2355, 2009.
10. R. Haas and K. Seyffarth. The k -Dominating Graph. *Graphs and Combinatorics*, March 2013. Online publication.
11. R. A. Hearn and E. D. Demaine. PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theor. Comput. Sci.*, 343(1-2):72–96, 2005.
12. T. Ito and E. D. Demaine. Approximability of the subset sum reconfiguration problem. In *Proc. of the 8th Annual Conf. on Theory and Applications of Models of Computation*, pages 58–69, 2011.
13. T. Ito, E. D. Demaine, N. J. A. Harvey, C. H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theor. Comput. Sci.*, 412(12-14):1054–1065, 2011.
14. T. Ito, M. Kamiński, and E. D. Demaine. Reconfiguration of list edge-colorings in a graph. *Discrete Applied Math.*, 160(15):2199–2207, 2012.
15. T. Ito, K. Kawamura, H. Ono, and X. Zhou. Reconfiguration of list $L(2,1)$ -labelings in a graph. In *Proc. of the 23rd Int. Symp. on Algorithms and Computation*, pages 34–43, 2012.
16. M. Kamiński, P. Medvedev, and M. Milanič. Shortest paths between shortest paths. *Theor. Comput. Sci.*, 412(39):5205–5210, 2011.
17. M. Kamiński, P. Medvedev, and M. Milanič. Complexity of independent set reconfigurability problems. *Theor. Comput. Sci.*, 439:9–15, June 2012.
18. E. W. Mayr and C. G. Plaxton. On the spanning trees of weighted graphs. *Combinatorica*, 12(4):433–447, 1992.
19. A. E. Mouawad, N. Nishimura, V. Raman, N. Simjour, and A. Suzuki. On the parameterized complexity of reconfiguration problems. In *Proc. of the 8th Int. Symp. on Parameterized and Exact Computation*, pages 281–294, 2013.