

Reconfiguring Minimum Dominating Sets in Trees

Magdalena Lemańska¹ Paweł Żyliński²

¹Department of Probability and Biomathematics
Gdańsk University of Technology, 80-233 Gdańsk, Poland

²Institute of Informatics
University of Gdańsk, 80-308 Gdańsk, Poland

Abstract

We provide tight bounds on the diameter of γ -graphs, which are reconfiguration graphs of the minimum dominating sets of a graph G . In particular, we prove that for any tree T of order $n \geq 3$, the diameter of its γ -graph is at most $n/2$ in the single vertex replacement adjacency model, whereas in the slide adjacency model, it is at most $2(n-1)/3$. Our proof is constructive, leading to a simple linear-time algorithm for determining the optimal sequence of “moves” between two minimum dominating sets of a tree.

Submitted: May 2019	Reviewed: August 2019	Revised: September 2019	Reviewed: December 2019
Revised: January 2020	Accepted: January 2020	Final: January 2020	Published: February 2020
Article type: Regular Paper		Communicated by: A. Lubiw	

1 Introduction

For a vertex v of a (simple) graph $G = (V_G, E_G)$, its *neighborhood*, denoted by $N_G(v)$, is the set of all vertices adjacent to v . The cardinality of $N_G(v)$, denoted by $d_G(v)$, is termed the *degree of v* . A vertex of degree one is termed a *leaf*, and the only neighbor of a leaf is called its *support vertex* (or simply, its *support*). If a support vertex has at least two leaves as neighbors, we call it a *strong support*, otherwise it is a *weak support*. A set of vertices $D \subseteq V_G$ of G is *dominating* if every vertex in the set $V_G - D$ has a neighbor in D . The cardinality of a minimum dominating set in G is termed the *domination number of G* and denoted by $\gamma(G)$, and any minimum dominating set of G is referred to as a γ -set.

Over the years, researchers have published thousands of papers on domination in graphs, exploring the topic in a variety of contexts. In particular, quite recently, two closely related concepts of reconfiguration graphs of the minimum dominating sets were introduced. In both of these variants, for a given graph G , the vertex set of the reconfiguration graph is the collection of all γ -sets of G ; however, the difference lies in the adjacency concept. Namely, in the *single vertex replacement adjacency model*, introduced in 2008 by Subramanian and Sridharan [16], two γ -sets X and Y of G are adjacent if there are vertices $x \in X$ and $y \in Y$ such that $X - \{x\} = Y - \{y\}$, whereas in the *slide adjacency model*, introduced by Fricke et al. [5] in 2011, it is required that, in addition, $xy \in E_G$. The single vertex replacement adjacency model was further studied in [10, 14, 15], and the slide adjacency model was further studied in [2, 3, 4]. Finally, reconfiguration graphs for dominating sets that are not necessarily minimum or for other models of domination have also been considered, see for example [1, 7, 8, 9, 11, 12, 17].

Herein, we focus on reconfiguration graphs of trees. For simplicity of presentation, we shall assume that in the two aforementioned models, both the reconfiguration graphs are termed the γ -graphs and denoted by Γ_G because the model under consideration is always either clear from the context, or not relevant. In 2011, Fricke et al. [5] posed the following question (among others, just as interesting, some of them having been already solved completely, see [3, 4, 13]): *In the slide adjacency model, is $\text{diam}(\Gamma_T) = O(n)$ for any tree T of order n ?* The partial answer for so-called caterpillars with one leg and for trees of diameter at most five was given by Bień [2], and only in 2018, Edwards et al. [4] answered the question in an affirmative way for all trees.

Theorem 1 [4] *For any tree T of order n , $\text{diam}(\Gamma_T) \leq 2\gamma(T) \leq n$ in the single vertex replacement adjacency model, whereas in the slide adjacency model, $\text{diam}(\Gamma_T) \leq 2(2\gamma(T) - |S_T|) \leq 2(n - 2)$, where S_T is the set of support vertices in T .*

However, the upper bounds established in Theorem 1 are not tight; in the single vertex replacement adjacency model, the best lower bound is $n/2$ [4] (being attained by the corona graph of a tree [6]), whereas in the slide adjacency






Figure 1: The mixed tree M has two arc-separators, marked with blue and green, respectively.

model it is $2(n-1)/3$ [5] (being attained by the path of order $n = 3k+1$, $k \geq 1$). Therefore, in this paper, we undertake their study and close down these gaps. Namely, our result is the following theorem:

Theorem 2 For any tree T of order $n \geq 3$, we have $\text{diam}(\Gamma_T) \leq \gamma(T) - |S_T''| \leq n/2$ in the single vertex replacement adjacency model, whereas in the slide adjacency model, $\text{diam}(\Gamma_T) \leq \min\{2(\gamma(T) - |S_T''|) - |S_T'|, 2(n-1)/3\}$, where S_T' (resp. S_T'') is the set of weak (resp. strong) support vertices in T .

Notation. For a vertex v of a graph $G = (V_G, E_G)$, the *closed neighborhood* of v , denoted by $N_G[v]$, is the set $N_G(v) \cup \{v\}$, and for a subset $X \subseteq V_G$ of vertices, the *neighborhood* of X , denoted by $N_G(X)$, is defined to be $\bigcup_{v \in X} N_G(v)$, and the *closed neighborhood* of X , denoted by $N_G[X]$, is the set $N_G(X) \cup X$. Next, for a vertex $v \in X$, the *private neighborhood of v with respect to X* is the set $\text{pn}_G(v, X) = N_G[v] - N_G[X - \{v\}]$, that is, the set of vertices that are in the closed neighborhood of v , but are not in the closed neighborhood of any other vertex in X . A vertex in $\text{pn}_G(v, X)$ is referred as to a *private neighbor* of v (with respect to X), and private neighbor of v is *external* if it is distinct from v itself. The set of leaves, the set of weak supports, the set of strong supports, and the set of all supports of G are denoted by L_G , S_G' , S_G'' , and S_G , respectively.

For a mixed tree $M = (V_M, E_M, A_M)$, the sets of tails and heads of arcs in A_M are denoted by V_M° and V_M^\blacktriangleright , respectively (notice that $v \in V_M$ may be an element of both V_M° and V_M^\blacktriangleright). Next, let \mathcal{F}_M be the family of all maximal connected arc-free subgraphs of M , and let $R = (V_R, E_R) \in \mathcal{F}_M$ be a subgraph of M such that $V_R \cap V_M^\blacktriangleright = \emptyset$. Then the set $S_R = V_R \cap V_M^\circ$ is called an *arc-separator in M* , whereas the graph R itself — the *certificate graph of S_R* ; see Fig. 1 for an illustration. Observe that for any two distinct arc-separators S_1 and S_2 in M , we have $S_1 \cap S_2 = \emptyset$, and moreover, there is neither edge $uv \in E_M$ nor arc $(u, v) \in A_M$, nor arc $(v, u) \in A_M$ such that $u \in S_1$ and $v \in S_2$.

Observation 1 Every mixed tree possesses an arc-separator.

A rooted tree is a pair (T, r) , for simplicity denoted by T_r , where $T = (V_T, E_T)$ is a tree and $r \in V_T$ is a distinguished vertex termed the *root*. A vertex $x \in V_T$ is labelled an *ancestor* of a vertex y in T_r if x belongs to the



unique path joining y and r , and if, in addition, $xy \in E_T$, then x is a *parent* of y . Next, symmetrically, the terms *descendant* of x and *child* of x , respectively, are used to describe such a vertex y . Note that x is both an ancestor and a descendant of itself. Finally, we use $T_r(x)$ to describe the subtree of T_r induced by the descendants of x and rooted at x .

2 The proof of Theorem 2

The statement is trivially valid for the case of $\gamma(T) = 1$. Thus, assume now that $T = (V_T, E_T)$ is a tree of order $n \geq 4$, with $\gamma(T) \geq 2$. We start with a simple general lemma.

Lemma 1 *Let X and Y be two distinct minimal dominating sets of a graph G . If $X - \{x\} = Y - \{y\}$ for some $x \in X$ and $y \in Y$, then:*

- a) $1 \leq \text{dist}_G(x, y) \leq 2$ holds;
- b) *If the girth of G is at least five, that is, G is acyclic or the shortest cycle in G is of the length at least five, then $|\text{pn}(x, X) - \{x\}| \leq 1$ as well as $|\text{pn}(y, Y) - \{y\}| \leq 1$.*

Proof: (a) Because X and Y are minimal dominating sets of G and $X - \{x\} = Y - \{y\}$, we have that $\text{pn}_G(x, X) = \text{pn}_G(y, Y) \neq \emptyset$. Consequently, $N_G(x) \cap N_G(y) \neq \emptyset$, and hence $1 \leq \text{dist}_G(x, y) \leq 2$. (b) Next, if $|\text{pn}(x, X) - \{x\}| \geq 2$ or $|\text{pn}(y, Y) - \{y\}| \geq 2$, then G would have a cycle of length three or four, which is a contradiction. \square

The idea of our proof of Theorem 2 is to treat a γ -set of the tree T as a set of k *tokens*, where $k = \gamma(T)$, that can be relocated within T , in discrete time steps, maintaining domination of the tree. Specifically, assume $V_T = \{1, 2, \dots, n\}$ and let D be the γ -set of T with the following property. When D is represented as the ordered k -tuple (v_1^D, \dots, v_k^D) of vertices in V_T , with $v_{i-1}^D < v_i^D$, $i \in [k] - \{1\}$ ¹, then the sequence $v_1^D \dots v_k^D$ is lexicographically the smallest one over the alphabet V_T , taken over all γ -sets of T . Next, let the k tokens, where $k = \gamma(T)$, be once labeled with identifying numbers $1, \dots, k$, which we shall refer to as Id_i , $i \in [k]$. Finally, let us initially locate these k tokens in such a way that the (unique) vertex occupied by the token Id_i is v_i^D , $i \in [k]$. Because the γ -graph of a tree T is connected [5], in both adjacency models, any sequence of consecutive (feasible) vertex replacements/slides (*moves*), starting from the set D and finishing at another γ -set of T , may be thought of as relocating our k -tokens, keeping their identifiers unchanged. In other words, we may uniquely associate any γ -set X of T with the ordered k -tuple (v_1^X, \dots, v_k^X) , where v_i^X is the vertex occupied by token Id_i . Following this convention, we observe that for any two (ordered) γ -sets X and Y of T , vertices X and Y are adjacent in the graph Γ_T if and only if for all but one $i \in [k]$, $v_i^X = v_i^Y$ holds. Next, for $i \in [k]$,

¹Herein, we use the convention that $[k]$ stands for the index set $\{1, 2, 3, \dots, k\}$, $k \geq 1$.

let V_T^i be the set of all vertices that can ever be occupied by token Id_i , that is, $V_T^i = \{v_i^X : X \text{ is a } \gamma\text{-set of } T\}$ (we emphasize that the set D defining the token labeling remains fixed).

Lemma 2 *For any $i \in [k]$, the relevant vertex sets V_T^i are the same in both adjacency models. In particular, the induced subgraph $T[V_T^i]$ is connected for any $i \in [k]$ (in both adjacency models).*

Proof: Due to the fact that every γ -graph in the slide adjacency model is a spanning subgraph of the relevant γ -graph in the single vertex replacement adjacency model [2], all we need is to argue that in the latter model, if X and Y are two adjacent γ -sets in the γ -graph of T , then a single move of a token in T from a vertex in X to a vertex in Y can be simulated by at most two subsequent moves of that token in the former model.

Let $X - \{x\} = Y - \{y\}$ for some $x \in X, y \in Y$. Assume without loss of generality that $\text{dist}_T(x, y) = 2$ (see Lemma 2). First, observe that the unique vertex $z \in N_T(x) \cap N_T(y)$ neither belongs to X nor to Y (otherwise, the set $X - \{x\} (= Y - \{y\})$ would be a smaller dominating set of T , which is a contradiction). Next, the minimality of X and Y combined with Lemma 1 implies that $\text{pn}(x, X) = \{z\} = \text{pn}(y, Y)$, and hence the set $Z = (X - \{x\}) \cup \{z\}$ is a γ -set of T , being adjacent to both X and Y in the γ -graph of T . Therefore, because $\text{dist}_T(x, z) = \text{dist}_T(z, y) = 1$, a single move of a token in T from x to y can be simulated by two subsequent moves of that token (from x to z and then from z to y) in the slide adjacency model, as required. \square

In the following sequence of lemmas we describe other properties of the sets V_T^i . These will be useful for the proof of Theorem 2.

Lemma 3 $V_T^i \cap V_T^j = \emptyset$ for any distinct $i, j \in [k]$ (in both adjacency models).

Proof: By Lemma 2, we may restrict ourselves only to the slide adjacency model. Suppose on the contrary that there exist distinct $i, j \in [k]$ such that $V_T^i \cap V_T^j \neq \emptyset$. Let Π be any (finite) walk in Γ_T starting at the γ -set D and traversing the edges of Γ_T until all vertices in $\cup_{t=1}^k V_T^t$ have been visited/occupied by tokens (tokens are moving with respect to the γ -sets visited along the walk); clearly, such a walk Π exists as Γ_T is connected [5]. Because $V_T^i \cap V_T^j \neq \emptyset$, there exist two γ -sets of T being adjacent along Π , say Y and Z , such that one of the tokens, say Id_a , is moved from a vertex of T , say y , and placed for the first time at another vertex of T , say z , that has already been visited by another token, say Id_b , with $b \neq a$. Let X be the γ -set of T with Id_b occupying vertex z for the first time along the walk Π . Consider now the rooted subtree $T' = T_z(y)$ of T_z , and, symmetrically, the rooted subtree $T'' = T_y(z)$ of T_y , see Fig. 2 for an illustration. From the choice of y and z , acyclicity of T and $\text{dist}_T(y, z) = 1$, it follows that:

- $Z \cap V_{T'}$ dominates all vertices in $V_{T'} - \{y\}$ and $|Z \cap V_{T'}| = |Y \cap V_{T'}| - 1$;

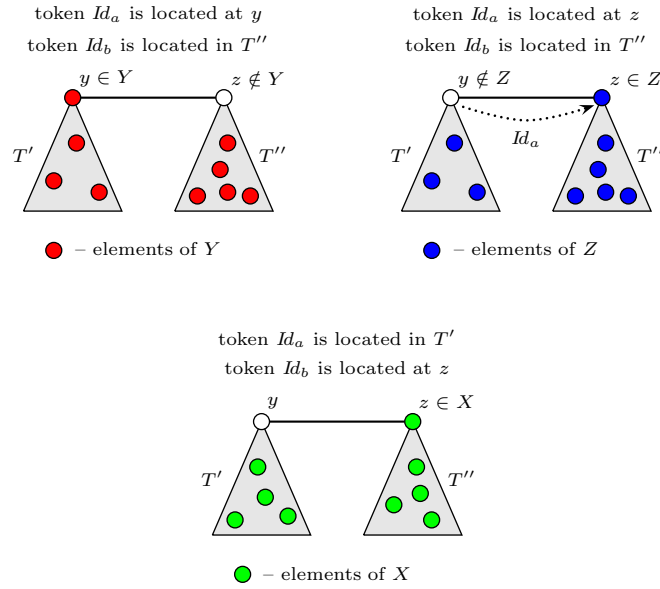


Figure 2: The set $S = (Z \cap V_{T'}) \cup (X \cap V_{T''})$ is a dominating set of T , and $|S| = \gamma(T) - 1$; notice that y may belong to X .

- $X \cap V_{T''}$ dominates all vertices in $V_{T''} \cup \{y\}$, and $|X \cap V_{T''}| = |Y \cap V_{T''}|$.

Consequently, because $V_T = V_{T'} \cup V_{T''}$ and $V_{T'} \cap V_{T''} = \emptyset$, the set $S = (Z \cap V_{T'}) \cup (X \cap V_{T''})$ is a dominating set of T with $|S| = \gamma(T) - 1$, which is a contradiction. \square

Lemma 4 For any $i \in [k]$, the distance between any two vertices in V_T^i is at most two in T (in both adjacency models).

Proof: By Lemma 2, we may again restrict ourselves only to the slide adjacency model. Suppose to the contrary that for some $i \in [k]$, there are two vertices $y, z \in V_T^i$ such that $\text{dist}_T(y, z) = 3$ (notice that in our supposition, we may, without loss of generality, restrict ourselves to vertices at the distance three because $T[V_T^i]$ is connected by Lemma 2). Let $\pi = v_0 v_1 v_2 v_3$ be the shortest path between $v_0 = y$ and $v_3 = z$ in T . Let Y and Z be two γ -sets of T such that token Id_l is located at vertex $v_0 (= y)$ and at vertex $v_3 (= z)$, respectively. Consider the rooted subtree $T' = T_{v_2}(v_1)$ of T_{v_2} and the rooted subtree $T'' = T_{v_1}(v_2)$ of T_{v_1} , see Fig. 3 for an illustration. Now, because T is a tree, $T[V_T^i]$ is connected (by Lemma 2), and $V_T^i \cap V_T^j = \emptyset$ for any distinct $i, j \in [k]$ (by Lemma 3), we observe that vertices $v_1, v_2 \notin Y$ and $v_1, v_2 \notin Z$. Consequently:

- $Z \cap V_{T'}$ dominates all vertices in $V_{T'}$ and $|Z \cap V_{T'}| = |Y \cap V_{T'}| - 1$;




Figure 3: The set $S = (Z \cap V_{T'}) \cup (Y \cap V_{T''})$ is a dominating set of T , and $|S| = \gamma(T) - 1$.

- $Y \cap V_{T''}$ dominates all vertices in $V_{T''}$.

Consequently, because $V_T = V_{T'} \cup V_{T''}$ and $V_{T'} \cap V_{T''} = \emptyset$, the set $S = (Z \cap V_{T'}) \cup (Y \cap V_{T''})$ is a dominating set of T with $|S| = \gamma(T) - 1$, which is a contradiction. \square

Lemma 5 *If $s \in S'_T$, then there exists $i_s \in [k]$ such that $V_T^{i_s} \subseteq \{s, l_s\}$, where l_s is the unique leaf adjacent to s in T , and so $\text{diam}(T[V_T^{i_s}]) \leq 1$ (in both adjacency models).*

Proof: By Lemma 2, we may focus only on the slide adjacency model. Let X be a γ -set of T such that $s \in X$ (clearly, such a γ -set exists) and let Id_{i_s} be the token located at vertex s . It follows from the minimality of X that no other token occupies the leaf l_s . Therefore, in order to move Id_{i_s} from s to a vertex distinct from the leaf l_s in T while maintaining domination of l_s , there must have already been located another token at s , together with Id_{i_s} , which contradicts Lemma 3. \square

Lemma 6 *If $s \in S''_T$, then there exists $i_s \in [k]$ such that $V_T^{i_s} = \{s\}$, and so $\text{diam}(T[V_T^{i_s}]) = 0$ (in both adjacency models).*

Proof: It follows by arguments analogous to those in the proof of Lemma 5. \square

We say that two (ordered) γ -sets $X = (v_1^X, \dots, v_k^X)$ and $Y = (v_1^Y, \dots, v_k^Y)$ of the given tree T are *inconsistent* at the coordinate $i \in [k]$ if $v_i^X \neq v_i^Y$; such a coordinate i itself, the vertices v_i^X and v_i^Y as well as the token Id_i are then also referred to as *inconsistent*, whereas the set $X - (X \cap Y)$ of all inconsistent vertices in Y (with respect to Y) is denoted by $\text{In}(X, Y)$, respectively.

Let X and Z be two (different) inconsistent γ -sets of the tree T (and so $\text{In}(X, Z) \neq \emptyset$), and let $M = (V_M, E_M, A_M)$ be the mixed tree, with the vertex set $V_M = V_T$, the edge set E_M and the arc set A_M , respectively, resulting from T by assigning the orientation to the edges (towards v_i^Z) on the shortest path between v_i^X and v_i^Z , for each $v_i^X \in \text{In}(X, Z)$. Let $R = (V_R, E_R)$ be the certificate graph of some arc-separator in M (such a graph R exists by Observation 1, and it is a subgraph of both T and M). We have a sequence of observations.

- (A) In the mixed tree M , all maximal directed paths are vertex-disjoint and of length of at most two (by combining Lemma 2, Lemma 3, and Lemma 4).
- (B) Therefore, $\text{In}(X, Z) = X - (X \cap Z) \subseteq V_M^\circ$ and $Z - (X \cap Z) \subseteq V_M^\blacktriangleright - V_M^\circ$, and thus $(Z - (X \cap Z)) \cap V_R = \emptyset$ (by the definition of a certificate graph); in other words, there is no inconsistent vertex in Z that belongs to V_R .
- (C) Finally, it follows from the definition of an arc-separator that if l is a leaf of R , then l is a leaf of T or $l = v_i^X$ ($\neq v_i^Z$) for some inconsistent coordinate $i \in [k]$. Notice that in the former case, $l = v_j^X = v_j^Z$ for some $j \in [k]$ may also hold.

Next, let U_d denote the set of all inconsistent vertices $v_i^X \in \text{In}(X, Z)$ such that $\text{dist}_T(v_i^X, v_i^Z) = d$; notice $d \in \{1, 2\}$ by Lemma 4. Observe that (see Fig. 4 for an illustration):

- (D) Because Z is a γ -set of T and $\text{dist}_T(v_i^X, v_i^Z) \geq 1$ for every $v_i^Z \in Z - (Z \cap X)$, the set $(Z \cap V_R) - U_2$ dominates all vertices in V_R , and so does the set $(X \cap V_R) - U_2$ (because $Z \cap V_R = X \cap V_R$ by the definition of an arc-separator). In other words, for the purpose of domination of R , vertices in the set $\{v_i^Z : v_i^X \in U_2\} \subseteq Z - (Z \cap X)$ are useless.
- (E) By similar arguments, the set $(Z \cap V_R) - \text{In}(X, Z)$ dominates all vertices in $V_R - \text{In}(X, Z)$, and so does the set $(X \cap V_R) - \text{In}(X, Z)$. In other words, no vertex in $U_1 (= \text{In}(X, Z) \cap (V_R - U_2))$ has an external private neighbor in V_R , that is, any such vertex may be required only to dominate itself in R .
- (F) Finally, $N_T(x_i) \cap (V_T - V_R) \subseteq N_T(z_i) \cap (V_T - V_R)$.

Consequently, tokens at inconsistent vertices in $\text{In}(X, Z) \cap V_R$ can be slid along the relevant arcs of M (recall that all maximal directed paths in M are vertex-disjoint), in a sequence, in total number $|U_1| + 2|U_2|$ of slides, to make all of them consistent, and the resulting set Y is a γ -set of T (by the properties




Figure 4: The set $(X \cap V_R) - U_2$ dominates all vertices in V_R , while the set $(X \cap V_R) - \text{In}(X, Z)$ dominates all vertices in $V_R - \text{In}(X, Z)$.

discussed above), with $|\text{In}(Y, Z)| < |\text{In}(X, Z)|$. Applying this approach repeatedly will eventually move all tokens from their initial positions v_1^X, \dots, v_k^X to the desired positions v_1^Z, \dots, v_k^Z , and — supported by Lemmas 4-6 — we may conclude that in the single vertex replacement adjacency model, the number of jumps is at most $\gamma(T) - |S_T''| \leq n/2$, and so $\text{diam}(\Gamma_T) \leq \gamma(T) - |S_T''| \leq n/2$ in this model, whereas in the slide adjacency model, the number of slides is at most $2(\gamma(T) - |S_T|) + |S_T'|$, and hence $\text{diam}(\Gamma_T) \leq 2(\gamma(T) - |S_T|) + |S_T'|$ in that model, as required.

Regarding the slide adjacency model and bounding the diameter of Γ_T in terms of the number of vertices, taking into account Lemmas 5 and 6, first observe that there are at least $|S_T| \geq 2$ tokens that require at most $|S_T|$ slides in total to make them consistent (recall that T is a tree of order at least four and $\gamma(T) \geq 2$). Next, if the number of slides to make a token Id_i consistent is equal to 2, then $|V_T^i| \geq 3$, and hence the number of such “expensive” tokens is at most $(|V_T| - 2|S_T|)/3 \leq (n - 4)/3$ (by Lemma 3). Therefore, a simple calculus shows that the maximum (total) number of slides is at most $2 + 2(n - 4)/3 = 2(n - 1)/3$, which finishes the proof of Theorem 2.

Remark. Let us note that the statements of Lemmas 3–4 cannot be carried over to the class of arbitrary graphs. As an example, consider the cycle $G = C_{3k+1}$ in which $V_G^i = V_G$ for any token Id_i (defined with respect to the γ -set D).

3 Algorithmic result

Observe that in the proof of Theorem 2, the relevant graph R can be extended and defined to be the union of the certificate graphs of an arbitrary number of (distinct) arc-separators in the mixed tree M . This is a core property that gives rise to a simple linear-time algorithm for determining the optimal sequence of jumps between two minimum dominating sets of a tree. The algorithm consists of three phases: pre-processing, assigning levels and final phase.

Pre-processing Phase. We identify pairs of vertices $(x_i, z_i) \in X \times Z$, each of which corresponds to the placement of the (unique by Lemma 3) token Id_i .

In that phase (see Fig. 5(a,b) for an illustration), we perform a DFS-based approach starting from a leaf $l \in L_T$, and for each vertex $v \in V_T$, we recursively determine the number n_v^X (resp. n_v^Z) of inconsistent vertices in the rooted subtree $T_l(v)$ that belong to a γ -set X (resp. to a γ -set Z). Notice that $|n_v^X - n_v^Z| \leq 1$ (because otherwise, vertex v must have been visited by two distinct tokens by Lemma 4 — a contradiction with Lemma 3). Next, using these data, starting from the same leaf l , the second pass of DFS is sufficient to identify the aforementioned pairs of vertices. More specifically, for the currently handled vertex v (in a post-order manner while performing DFS), assuming that the $i - 1$ pairs $(x, z) \in X \times Z$ has already been identified in all subtrees of $T_l(v)$ rooted at the children of v (if any), the following rules can be applied (they are exhaustive and distinct by Lemma 3 and Lemma 4).

- If $v \in X \cap Z$, then $x_i := v$ and $z_i := v$. (Notice that $n_v^X = n_v^Z$ in this case.)
- If $v \in X - Z$ and $n_v^X = n_v^Z$, then $x_i := v$, whereas z_i is assigned the unique non-associated yet vertex in $T_l(v)$ that belongs to Z .
- If $v \in Z - X$ and $n_v^X = n_v^Z$, then $z_i := v$, whereas x_i is assigned the unique non-associated yet vertex in $T_l(v)$ that belongs to X .
- Otherwise, continue: no vertices are associated, but if $v \in X$, then v is marked as “non-associated x ”, and if $v \in Z$ then it is marked as “non-associated z ”.

Assigning Levels Phase. We assigns levels to vertices/tokens in X . These levels will constitute the ordering that the tokens will move with respect to.

Let $M = (V_M, E_M, A_M)$ be the mixed tree defined in the proof of Theorem 2 (Section 2), resulting from T by assigning the orientation to the edges (towards z_i) on the shortest path between x_i and z_i , for each $x_i \in \text{In}(X, Z)$; see Fig. 6(a)

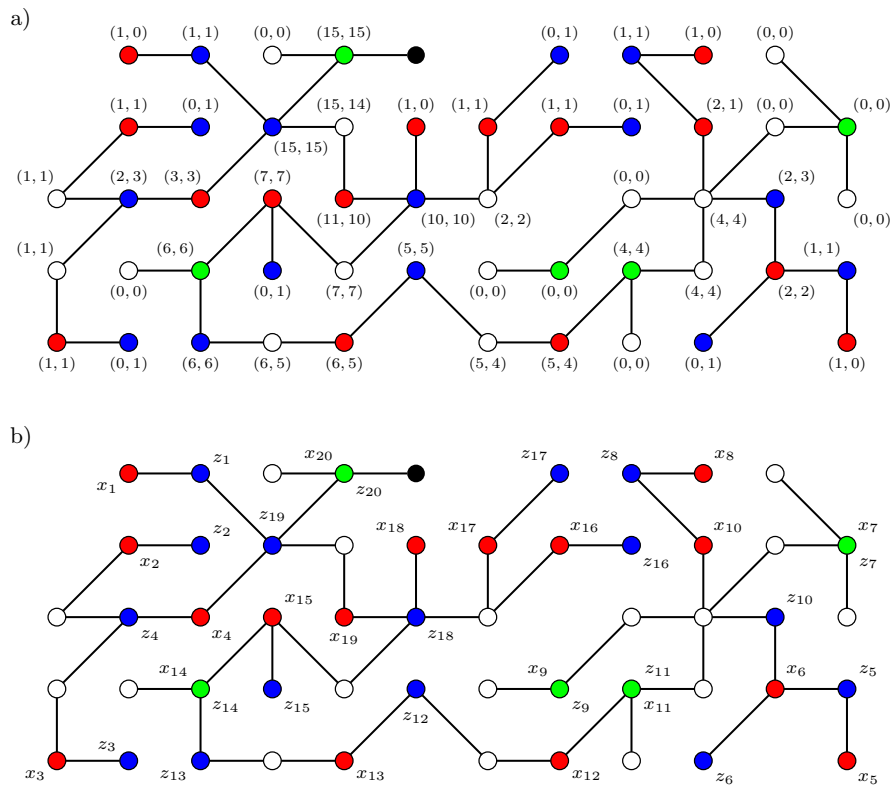


Figure 5: Pre-processing Phase. A tree T with $\gamma(T) = 20$ and the γ -sets X and Z of T : $X - Z$ is marked red, $Z - X$ is marked blue, and $X \cap Z$ is marked green. (a) Determining the numbers n_v^X and n_v^Z (depicted as pairs (n_v^X, n_v^Z)), starting at the black leaf). (b) Identifying the pairs $(x_i, z_i) \in X \times Z$; herein, children of a vertex are visited in a counterclockwise manner, with respect to the given plane embedding of T .




Figure 6: Assigning Levels Phase. (a) The mixed tree M . (b) The Hasse diagram $H = (\text{In}(X, Z), A)$ of $\langle \text{In}(X, Z), \prec \rangle$.

for an illustration. Recall that all directed paths in M are vertex-disjoint and of a length of at most two. Define the partially ordered set $\langle \text{In}(X, Z), \prec \rangle$, where for two distinct $x_i, x_j \in \text{In}(X, Z)$, $\prec x_j$ if and only if there is no arc-free path between x_i and x_j in M and all the arcs on the (unique) path between x_i and x_j are oriented towards x_j . Next, consider the transitive reduction $H = (\text{In}(X, Z), A)$ of $\langle \text{In}(X, Z), \prec \rangle$ in the form of the Hasse diagram, with the layers L_1, \dots, L_t , where $t \leq \gamma(T)$ (notice that because T is a tree, such a transitive reduction exists); see Fig. 6(b). These layers define now the labeling of inconsistent vertices in X : if $x_i \in L_k$, then x_i is assigned the level k . Observe that H is not necessarily connected, but it is a directed forest, that is, its underlying undirected graph is a forest (because T is a tree). Moreover, it can be computed, together with the layers L_1, \dots, L_t , in linear time by applying the third pass of a DFS-based approach on the tree T .

Final Phase. We move tokens from x_i to z_i with respect to the increasing order of the assigned levels to inconsistent vertices.

Before we proceed with the correctness proof of our 3-phase algorithm, let us point out that it was not our intention to optimize the number of DFS-phases in our algorithm. Therefore, we believe that with respect to this criterion, some improvement is possible, and we eventually conclude our paper with the following theorem.

Theorem 3 *Given two γ -sets X and Z of a tree T , an optimal sequence of jumps through which X can be transformed into Z can be computed in linear time (in both adjacency models).*

Proof: For a level $l \in \{1, \dots, t\}$, let Y_{l+1} denote the set resulting from moving all tokens in L_l to the relevant vertices in Z . It follows from the definition/construction that $Y_{t+1} = Z$, and for each $l \in \{1, \dots, t-1\}$, $L_{l+1} \subseteq Y_{l+1}$ and $\text{In}(Y_{l+1}, Z) = \text{In}(X, Z) - \bigcup_{i=1}^l L_i$.

Due to the fact that L_1 is the set of minimal elements in $\langle \text{In}(X, Z), \prec \rangle$, L_1 is the sum of a number of arc-separators in the mixed tree $M_1 = M$ (exploited in Phase 2 and defined in the proof of Theorem 2). Consequently, it follows from the proof of Theorem 2 (i.e., the arguments from the paragraph just after Lemma 6) that the set Y_2 , resulting from moving tokens located at inconsistent vertices in L_1 towards the relevant vertices in Y_2 (and so in Z), in any order, is a γ -set of T .

But the same argument can be inductively (successively) applied to all the γ -sets Y_l and Z , $l \in \{1, 2, \dots, t\}$, and the partially ordered set $\langle \text{In}(Y_l, Z), \prec \rangle$, defined now with respect to Y_l and Z . Namely, observe that L_{l+1} is the set of minimal elements in $\langle \text{In}(Y_{l+1}, Z), \prec \rangle$, which implies that L_{l+1} is the sum of a number of arc-separators in the relevant mixed tree M_l (defined now with respect to Y_l and Z). Consequently, it follows from the proof of Theorem 2 that the set Y_{l+1} is a γ -set of T for each $l \in \{1, \dots, t\}$. Therefore, moving tokens with respect to the increasing order of the assigned levels to inconsistent vertices constitutes a feasible optimal reconfiguration of the γ -set X into the γ -set Z .

Finally, with respect to the complexity issue, all we need is to observe that all three phases can clearly be accomplished in linear time. \square

Acknowledgements

We would like to thank the referees for their remarkable suggestions and comments on our manuscript.

References

- [1] S. Alikhani, D. Fatehi, and S. Klavžar. On the structure of dominating graphs. *Graphs and Combinatorics*, 33(4):665–672, 2017. doi:10.1007/s00373-017-1792-5.
- [2] A. Bień. Gamma graphs of some special classes of trees. *Annales Mathematicae Silesianae*, 29(1):25–34, 2015. doi:10.1515/amsil-2015-0003.
- [3] E. Connelly, S. Hedetniemi, and K. Hutson. A note on γ -graphs. *AKCE International Journal of Graphs and Combinatorics*, 8(1):23–31, 1999.
- [4] M. Edwards, G. MacGillivray, and S. Nasserar. Reconfiguring minimum dominating sets: the γ -graph of a tree. *Discussiones Mathematicae Graph Theory*, 38(3):703–716, 2018. doi:10.7151/dmgt.2044.
- [5] G. Fricke, S. Hedetniemi, S. Hedetniemi, and K. Hutson. γ -graphs of graphs. *Discussiones Mathematicae Graph Theory*, 31(3):517–531, 2011.
- [6] R. Frucht and F. Harary. On the corona of two graphs. *Aequationes Mathematicae*, 4(1-2):322–325, 1970. doi:10.1007/BF01844162.
- [7] R. Haas and K. Seyffarth. The k -dominating graph. *Graphs and Combinatorics*, 30(3):609–617, 2014. doi:10.1007/s00373-013-1302-3.
- [8] R. Haas and K. Seyffarth. Reconfiguring dominating sets in some well-covered and other classes of graphs. *Discrete Mathematics*, 340(8):1802–1817, 2017. doi:10.1016/j.disc.2017.03.007.
- [9] A. Haddadan, T. Ito, A. E. Mouawad, N. Nishimura, H. Ono, A. Suzuki, and Y. Tebbal. The complexity of dominating set reconfiguration. *Theoretical Computer Science*, 651:37–49, 2016. doi:10.1016/j.tcs.2016.08.016.
- [10] S. Lakshmanan and A. Vijayakumar. The gamma graph of a graph. *AKCE International Journal of Graphs and Combinatorics*, 7(1):53–59, 2010.
- [11] C. Mynhardt and L. Teshima. A note on some variations of the γ -graph. Available at [arXiv.org/abs/1707.02039](https://arxiv.org/abs/1707.02039).
- [12] C. Mynhardt, L. Teshima, and R. Roux. Connected k -dominating graphs. *Discrete Mathematics*, 342(1):145–151, 2019. doi:10.1016/j.disc.2018.09.006.
- [13] G. Rote. The maximum number of minimal dominating sets in a tree. In T. M. Chan, editor, *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1201–1214, 2019.
- [14] N. Sridharan, S. Amutha, and S. B. Rao. Induced subgraphs of gamma graphs. *Discrete Mathematics, Algorithms and Applications*, #1350012, 2013. doi:10.1142/S1793830913500122.



- [15] N. Sridharan and K. Subramanian. Trees and unicyclic graphs are γ -graphs. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 69(1):231–236, 2009.
- [16] K. Subramanian and N. Sridharan. γ -graph of a graph. *Bulletin of Kerala Mathematics Association*, 5:17–34, 2008.
- [17] A. Suzuki, A. E. Mouawad, and N. Nishimura. Reconfiguration of dominating sets. *Journal of Combinatorial Optimization*, 32(4):1182–1195, 2016. doi:10.1007/s10878-015-9947-x.