

Independent-set reconfiguration thresholds of hereditary graph classes

Citation for published version (APA):

de Berg, M., Jansen, B. M. P., & Mukherjee, D. (2018). Independent-set reconfiguration thresholds of hereditary graph classes. *Discrete Applied Mathematics*, 250, 165-182. <https://doi.org/10.1016/j.dam.2018.05.029>

Document license:

TAVERNE

DOI:

[10.1016/j.dam.2018.05.029](https://doi.org/10.1016/j.dam.2018.05.029)

Document status and date:

Published: 11/12/2018

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.



Independent-set reconfiguration thresholds of hereditary graph classes[☆]

Mark de Berg, Bart M.P. Jansen, Debankur Mukherjee^{*}

Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands



ARTICLE INFO

Article history:

Received 24 April 2017
 Received in revised form 30 April 2018
 Accepted 10 May 2018
 Available online 19 June 2018

Keywords:

Reconfiguration
 Independent set
 Token addition removal
 Token sliding

ABSTRACT

Traditionally, reconfiguration problems ask the question whether a given solution of an optimization problem can be transformed to a target solution in a sequence of small steps that preserve feasibility of the intermediate solutions. In this paper, rather than asking this question from an algorithmic perspective, we analyze the combinatorial structure behind it. We consider the problem of reconfiguring one independent set into another, using two different processes: (1) exchanging exactly k vertices in each step, or (2) removing or adding one vertex in each step while ensuring the intermediate sets contain at most k fewer vertices than the initial solution. We are interested in determining the minimum value of k for which this reconfiguration is possible, and bound these threshold values in terms of several structural graph parameters. For hereditary graph classes we identify structures that cause the reconfiguration threshold to be large.

© 2018 Elsevier B.V. All rights reserved.

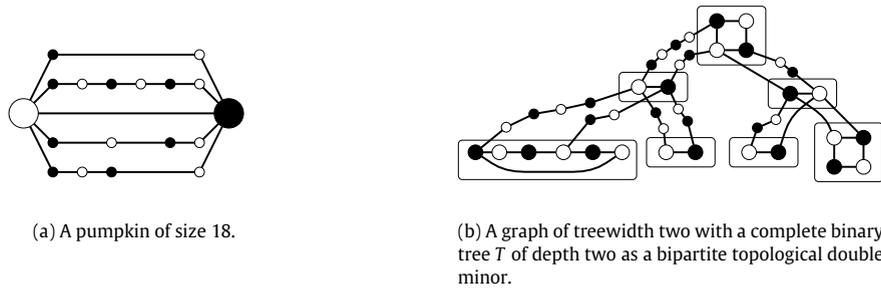
1. Introduction

Over the past decade, reconfiguration problems have drawn a lot of attention of researchers in algorithms and combinatorics [5,6,9,14,16–18,20,27,29,32]. In this framework, one asks the following question: Given two solutions I, J of a fixed optimization problem, can I be transformed into J by a sequence of small steps that maintain feasibility for all intermediate solutions? Such problems are practically motivated by the fact it may be impossible to adapt a new production strategy instantaneously if it differs too much from the strategy that is currently in use; changes have to be made in small steps, but production has to keep running throughout. From a theoretical perspective, the study of reconfiguration problems provides deep insights into the structure of the solution space. One of the well-studied examples is when the solution space consists of all the independent sets of a graph (optionally all having a prescribed size). In this case, three types of reconfiguration rules have been considered. These are naturally explained using *tokens* on vertices of the graph. In *Token Addition Removal* (TAR) [18,27], there is a token on every vertex of the initial independent set, and there is a buffer of tokens, initially empty. A step consists of removing a token from a vertex and placing it in the buffer, or placing a buffer token onto a vertex of the graph. The set of vertices with tokens must form an independent set at all times, and the goal is to move the tokens from the initial to the target independent set while ensuring the buffer size never exceeds a given threshold. In *Token Sliding* (TS) [16,20], a step consists of replacing one vertex v in the independent set by a neighbor of v (the token slides along an edge). In *Token Jumping* (TJ) [20] a step also consists of replacing a single vertex, but the newly added vertex need not have

[☆] An extended abstract of the paper appeared in 36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS), Chennai, India, 2016, doi:10.4230/LIPIcs.FSTTCS.2016.34.

^{*} Corresponding author.

E-mail addresses: m.t.d.berg@tue.nl (M. de Berg), b.m.p.jansen@tue.nl (B.M.P. Jansen), d.mukherjee@tue.nl (D. Mukherjee).



(a) A pumpkin of size 18.

(b) A graph of treewidth two with a complete binary tree T of depth two as a bipartite topological double minor.

Fig. 1. The bipartite structures responsible for large MTJ and TAR reconfiguration thresholds, respectively. A *pumpkin* consists of odd-length vertex-disjoint paths between two vertices. The special form of topological *minor* represents each vertex of the tree T by an edge or even cycle in G , and each edge of T by two odd-length paths connecting vertices in opposite partite sets in G .

any neighboring relation with the replaced vertex (the token jumps). Token jumping reconfiguration is equivalent to TAR reconfiguration with a buffer of size one.

These models have been analyzed in detail in the recent literature on algorithms [5,6,9,14,15,26], complexity theory [16,18,20,27], combinatorics [7,13], and even statistical physics [19,22,28]. See [17,29] for surveys on reconfiguration. It is known that the reconfiguration problem under all the above three rules is PSPACE-complete for general graphs, perfect graphs, and planar graphs [16,18,20]. The TJ and TAR reconfiguration problems are PSPACE-complete even for bounded bandwidth graphs [32] and NP-complete for bipartite graphs [25]. Also, TS reconfiguration is PSPACE-complete for bipartite graphs [25]. Further analyses on the complexity can be found in [5,6,9,14,26]. The constrained token moving problems are related to pebbling games that have been studied in the literature, with applications to robot motion planning [1,7,13,15].

As mentioned, the goal in reconfiguring independent sets is to go from one given independent I to another one J by a sequence of small steps. In the TS and TJ models, a step involves moving only a single token. This is ideal, but unfortunately reconfiguration is often impossible in the TS or TJ model. Reconfiguration in the TAR model is always possible if one makes the buffer size sufficiently large. However, having a large buffer size is undesirable. We are interested in determining the minimum buffer size that is sufficient to ensure any independent set in a given graph G can be reconfigured to any target independent set of the same size. We call this minimum the *TAR reconfiguration threshold* (precise definitions in Section 2). Our aim is to bound the threshold in terms of properties of the graph, and to identify the structures contained in hereditary graph classes that cause the thresholds to be large. We also generalize the TJ model to *Multiple Token Jumping* (MTJ), where in each step a prescribed number of tokens may be moved simultaneously. In the MTJ model, the question becomes: What is the minimum number of simultaneously jumping tokens needed to ensure any reconfiguration is possible? This quantity is called the *MTJ reconfiguration threshold*.

Our contribution. We provide upper and lower bounds on the MTJ and TAR reconfiguration thresholds in terms of several graph parameters. In general, our bounds apply to the reconfiguration thresholds of hereditary *graph classes*. The threshold of a graph class is the supremum of the threshold values of the graphs in that class: it is the smallest value k such that for any graph in the class, any source independent set I in that graph can be reconfigured into any target independent set J using steps of size k (for MTJ) or a buffer of size k (for TAR).

The MTJ reconfiguration threshold of graphs that are structurally very simple, may nevertheless be very large. For example, an even cycle with $2n$ vertices can be partitioned into two independent sets I and J of size n each. Any MTJ reconfiguration of I into J requires a jump of n vertices, and this is trivially sufficient. Since a cycle has a feedback vertex set (that is, a subset of vertices whose removal makes the graph acyclic, see Section 2) of size one, the MTJ threshold cannot be bounded in terms of the size of a minimum feedback vertex set. However, we prove that the threshold is upper-bounded by the size of a minimum vertex cover of G . Although this bound is tight in the worst case, there are many graph classes with a small MTJ threshold even though they require a large vertex cover. Trees for example have MTJ threshold at most one. We therefore introduce the notion of *pumpkin*, which consists of two nodes connected by at least two vertex-disjoint paths of odd length (Fig. 1(a)). The *size* of a pumpkin is the total number of vertices in the structure. We characterize the MTJ reconfiguration threshold of a hereditary graph class \mathcal{H} in terms of the size of the largest pumpkin it contains: the MTJ reconfiguration threshold is upper- and lower-bounded in terms of the largest pumpkin contained in a bipartite graph in \mathcal{H} .

TAR reconfiguration is more versatile than MTJ reconfiguration. In the concrete example of a $2n$ -cycle discussed above, its MTJ threshold is n while any pair of independent sets can be reconfigured in the TAR model using a buffer of size two. Moreover, we show that any graph that has a feedback vertex set of size k has TAR reconfiguration threshold at most $k + 1$, and reconfiguring one side of the complete bipartite graph $K_{n,n}$ to the other side shows that this is tight. Our main result concerning TAR reconfiguration states that the TAR reconfiguration threshold of any graph is upper-bounded by its pathwidth (see Section 2). Somewhat surprisingly, there are graphs of constant treewidth (treewidth 2 suffices) for which the TAR reconfiguration threshold is arbitrarily large. We also introduce the concept of *bipartite topological double minor* (BTD-minor), see Fig. 1(b), and show using an isoperimetric inequality that any hereditary graph class containing a graph having a

complete binary tree of depth d as a BTD-minor, has TAR reconfiguration threshold $\Omega(d)$. To the best of our knowledge, we are the first to present upper and lower bounds on MTJ and TAR reconfiguration threshold values in terms of other graph parameters.

We require the restriction to hereditary graph classes in some of our statements to be able to develop meaningful lower bounds on reconfiguration thresholds, as explained next. Let G be the disjoint union of $K_{n,n}$ and a graph H , and let I and J be the two partite sets of $K_{n,n}$. One can verify that I can be reconfigured to J by jumps of size at most one if and only if H has an independent set of size $n - 1$. Similarly, I can be TAR reconfigured to J using a buffer of size k if and only if H has an independent set of size $n - k$. Since the size of a maximum independent set is NP-complete to determine, there are no good characterizations of this quantity. When developing lower bounds on the threshold of a hereditary graph class Π , this issue disappears since the reconfiguration threshold of any class containing the graph G above, is at least as high as the threshold of H (which must be contained in Π if G is), which is n . The restriction to hereditary graph classes therefore enables us to focus our attention to reconfiguration problems where all vertices in the graph are contained in either the source or target independent set, thereby avoiding the obstacle that the reconfiguration threshold matches the size of a maximum independent set.

Applications. The MTJ and TAR reconfiguration thresholds play an important role in statistical physics and wireless communication networks. To understand the importance of the TAR reconfiguration threshold, consider the following process: In a graph G , nodes are trying to become *active* at some rate, independently of each other in a distributed manner. When a potential activation occurs at a node, it can only become active if none of its neighboring nodes are active at that moment. An active node deactivates at some rate independent of the other processes. At any point in time, the set of active nodes in this process forms an independent set of the graph. In statistical physics, this process is known as Glauber dynamics with *hard-core interaction*. This activity process on graphs has been studied extensively in the context of communication networks [23]. In this setup, the graph represents a communication network and calls arrive at each vertex. An arriving call at vertex v (say) can only be accepted if there is no ongoing call at any neighboring vertex of v , otherwise the call is lost. In this case, if a call is accepted at v , it is considered to be active. Further extensions of this model have gained a lot of attention in the context of CSMA algorithm for conflict-free transmission (see [31]) in wireless network. Here the graph represents the *interference graph* of wireless devices and the activation corresponds to information (data-packet) transmission, i.e., no two neighboring vertices can transmit simultaneously. Thus, the set of transmitting nodes at any point in time, forms an independent set of the graph. In all the above scenarios, it is interesting to study the time this activity process takes to reach a target independent set (active nodes), starting from some specific independent set, since this provides key insights about the average transmission delay and the temporal starvation [28,33]. Loosely speaking, when the activation rate is large, this time has been shown to depend crucially upon what we call the TAR reconfiguration threshold of the underlying graph (see [28] for an extended discussion). In particular, the mixing time of the Glauber dynamics on a graph increases exponentially with its TAR reconfiguration threshold, and hence the Glauber dynamics on the graph is fast mixing if and only if the TAR reconfiguration threshold is small. When the graph is bipartite, this dynamics has recently been considered in [10].

The MTJ reconfiguration threshold of a graph G can be interpreted in the following way. Consider the auxiliary graph, whose vertices correspond to size- s independent sets in G for some fixed s , with an edge between vertices representing sets I, J if $|I \setminus J| \leq k$. Then the MTJ reconfiguration threshold is at most k if and only if this graph auxiliary graph is connected for all s . The MTJ reconfiguration threshold therefore has applications in the parallel Glauber dynamics (PGD) [19,22], where the MTJ reconfiguration threshold provides the jump size required to make the underlying Markov process ergodic.

Organization. The succeeding sections are organized as follows. In Section 2 we provide graph-theoretic preliminaries. In Section 3 we provide a formal description of the two types of reconfiguration. In Section 4 we analyze MTJ reconfiguration. Section 5 deals with TAR reconfiguration.

2. Preliminaries

In this section we give the most important graph-theoretic definitions. Notions not defined here can be found in one of the textbooks [8,11]. A graph is a pair $G = (V, E)$, where V is the set of vertices, and E is the set of edges. We also use $V(G)$ and $E(G)$ to refer to the vertex and edge set of G , when convenient. All graphs we consider are finite, simple, and undirected. For $U \subseteq V$ we denote by $G - U$ the graph obtained from G by removing the vertices in U and their incident edges. A set $U \subseteq V$ is called an *independent set* of G , if $\{u, v\} \notin E$ for any $u, v \in U$. The *symmetric difference* of two sets U and U' is $U \Delta U' := (U \setminus U') \cup (U' \setminus U)$. A set $U \subseteq V$ is a *vertex cover* of G if every edge in E is incident with a vertex in U . The minimum cardinality of a vertex cover of G is denoted by $\text{vc}(G)$. A set $U \subseteq V$ is a *feedback vertex set* if $G - U$ is acyclic (a *forest*). The minimum cardinality of a feedback vertex set of G is denoted $\text{fvs}(G)$. For a vertex v , denote by $N_G(v)$ the set of its neighbors (excluding v itself). The *open* and *closed neighborhood* of a set $U \subseteq V$ are $N_G(U) := \bigcup_{s \in U} N_G(s) \setminus U$ and $N_G[U] := \bigcup_{s \in U} N_G(s) \cup U$, respectively. We omit the subscript when it is clear from the context. A graph $G' = (V', E')$ is said to be a *subgraph* of G , if $V' \subseteq V$, and $E' \subseteq E$. It is an *induced subgraph* of G if $V' \subseteq V$ and for any $u, v \in V'$ we have $\{u, v\} \in E$ if and only if $\{u, v\} \in E'$. The subgraph of G induced by $U \subseteq V$ is denoted $G[U]$. A *graph class* is a (possibly infinite) collection of graphs. A graph class Π is said to be *hereditary* if given any graph $G \in \Pi$, any induced subgraph of G belongs to the class Π as well. A graph is *bipartite* if its vertex set can be partitioned into two independent sets I and J , which are also called the

partite sets. We sometimes denote such a bipartite graph by $G = (I \cup J, E)$. A bipartite graph is *balanced* if $|I| = |J|$. A *matching* is a set of edges that do not share any endpoints. A matching *covers* a vertex v if it contains an edge incident on v . A matching is *perfect* if it covers all vertices. We will utilize the following well-known consequence of König’s theorem.

Fact 1 ([30, Corollary 16.7]). *Let $G = (I \cup J, E)$ be a bipartite graph. Then G has a matching covering I if and only if $|N(S)| \geq |S|$ for each $S \subseteq I$.*

A vertex v is a *cutvertex* in graph G if the removal of v increases the number of connected components. A graph is *biconnected* if it does not contain a cutvertex. Under this definition, the graph K_2 is biconnected. A *biconnected component* of G is a maximal biconnected subgraph of G .

Definition 1 ([8, §7.2]). A *path decomposition* of a graph $G = (V, E)$ is a sequence $\mathcal{P} = (X_1, X_2, \dots, X_r)$ of subsets of V called *bags*, satisfying the following conditions:

- (P1) $\bigcup_{i=1}^r X_i = V$. In other words, every vertex of G is in at least one bag.
- (P2) For every $\{u, v\} \in E$, there exists $l \in \{1, 2, \dots, r\}$ such that the bag X_l contains both u and v .
- (P3) For every $u \in V$, if $u \in X_i \cap X_k$ for some $i \leq k$, then $u \in X_j$ also for each j such that $i \leq j \leq k$. In other words, the indices of the bags containing u form an interval in $\{1, 2, \dots, r\}$.

The *width* of a path decomposition (X_1, \dots, X_r) is $\max_{1 \leq i \leq r} |X_i| - 1$. The *pathwidth* of G , denoted by $\text{pw}(G)$, is the minimum possible width of a path decomposition of G . A path-decomposition (X_1, X_2, \dots, X_r) of a graph G is *nice* if the following holds:

- (i) $X_1 = X_r = \emptyset$, and
- (ii) for every $i \in \{1, 2, \dots, r - 1\}$, there is either a vertex $v \notin X_i$ such that $X_{i+1} = X_i \cup \{v\}$, or there is a vertex $w \in X_i$ such that $X_{i+1} = X_i \setminus \{w\}$.

It is well-known (cf. [8, Lemma 7.2]) that every graph admits a nice path decomposition of width $\text{pw}(G)$. For any path decomposition $\mathcal{P} = (X_1, X_2, \dots, X_r)$ of $G = (V, E)$, and any vertex $v \in V$, define $l_{\mathcal{P}}(v) = \min\{i : v \in X_i\}$ and $r_{\mathcal{P}}(v) = \max\{i : v \in X_i\}$, i.e. $l_{\mathcal{P}}(v)$ and $r_{\mathcal{P}}(v)$ respectively denote the index of the first and last bag containing v . Note that if \mathcal{P} is nice, then $l_{\mathcal{P}}(\cdot)$ and $r_{\mathcal{P}}(\cdot)$ are injective maps over the set of vertices.

3. Definitions and basic facts for reconfiguration

In this section we formally define the two notions of reconfiguration and establish some basic facts.

Multiple token jump (MTJ). Given any two independent sets I and J , with $|I| = |J|$, we say that I can be *k-MTJ reconfigured* to J , if there exists a finite sequence of independent sets $(I = W_0, W_1, W_2, \dots, W_n, W_{n+1} = J)$ for some $n \geq 0$, such that for all $i \in \{0, \dots, n + 1\}$ the set W_i is independent in G , $|W_i| = |I| = |J|$, and $|W_{i+1} \setminus W_i| \leq k$. A step $W_i \rightarrow W_{i+1}$ in the reconfiguration process with $|W_i \setminus W_{i+1}| = k$ is called a *k-TJ move*. Given a graph $G = (V, E)$, define $\text{MTJ}(G, s)$ as the minimum value of k , such that any two independent sets of size s in G can be *k-MTJ reconfigured* to each other. Now define $\text{MTJ}(G) := \max_{1 \leq s \leq |V|} \text{MTJ}(G, s)$. Our goal is to characterize the value of $\text{MTJ}(G)$ in terms of certain parameters of the graph G . We call $\text{MTJ}(G)$ the *MTJ reconfiguration threshold* of the graph G . The MTJ reconfiguration threshold of a graph class \mathcal{I} is defined as $\text{MTJ}(\mathcal{I}) := \sup_{G \in \mathcal{I}} \text{MTJ}(G)$.

Token addition removal (TAR). Given any two independent sets I and J , with $|I| = |J|$, we say that I can be *k-TAR reconfigured* to J , if there exists a finite sequence of independent sets $(I = W_0, W_1, W_2, \dots, W_n, W_{n+1} = J)$ for some $n \geq 0$, such that W_i is independent in G , $|I| - |W_i| \leq k$, and $|W_{i-1} \Delta W_i| \leq 1$ for all $i \in \{0, \dots, n + 1\}$. We refer to the quantity $B_i := |I| - |W_i|$ as the *buffer size* at step i : the tokens that were on the initial independent set, and are not on the current independent set W_i , are placed in the buffer. Define $\text{TAR}(G, s)$ to be the smallest buffer size k such that any two independent sets of size s can be *k-TAR reconfigured* to each other. Define $\text{TAR}(G) := \max_{1 \leq s \leq |V|} \text{TAR}(G, s)$. As before, we call $\text{TAR}(G)$ the *TAR reconfiguration threshold* of the graph G , and extend the same terminology to graph classes \mathcal{I} by defining $\text{TAR}(\mathcal{I}) := \sup_{G \in \mathcal{I}} \text{TAR}(G)$.

Facts on reconfiguration. Observe that for any graph G , it holds that $\text{MTJ}(G) = 1$ if and only if $\text{TAR}(G) = 1$. In general, the TAR reconfiguration threshold is at most the MTJ reconfiguration threshold. Indeed, to see this, observe that each *k-TJ* move can be thought of as a sequence of $2k$ steps with maximum buffer size k . First, sequentially remove the k vertices that are jumping away, placing their tokens in the buffer; then sequentially place the buffer tokens on the k new vertices in the independent set.

Proposition 1. *Let G be a graph with independent sets I and J of equal size. If $I \setminus J$ can be *k-TAR reconfigured* (resp. *k-MTJ reconfigured*) to $J \setminus I$ in the graph $G[I \Delta J]$, then I can be *k-TAR reconfigured* (resp. *k-MTJ reconfigured*) to J in G .*

Proof. Consider a sequence of independent sets $(I \setminus J = W_0, \dots, W_{n+1} = J \setminus I)$ in $G[I \Delta J]$ that reconfigures $I \setminus J$ to $J \setminus I$. Since I and J are independent in G , no vertex of $I \Delta J$ is adjacent to a vertex of $I \cap J$. Hence $W'_i := W_i \cup (I \cap J)$ is an independent set in G for all i , and the sequence (W'_0, \dots, W'_{n+1}) reconfigures $(I \setminus J) \cup (I \cap J) = I$ to $(J \setminus I) \cup (I \cap J) = J$ in G . The step size and buffer size of this sequence in G are not greater than the corresponding values for the sequence in $G[I \Delta J]$, which completes the proof. \square

Proposition 1 shows that to upper-bound the TAR or MTJ reconfiguration threshold, it suffices to do so in balanced bipartite graphs where the source and target configurations are disjoint; note that $G[I\Delta J]$ is balanced bipartite and $I \setminus J$ and $J \setminus I$ are disjoint. We will frequently exploit this in our proofs. For any graph class Π , let Π_{bip} denote the set of bipartite graphs in Π . The following proposition shows that the reconfiguration threshold of a hereditary graph class is determined by the behavior of the bipartite graphs in the class. Note that for hereditary classes Π , the class Π_{bip} is also hereditary.

Proposition 2. For any hereditary graph class Π , we have $\text{MTJ}(\Pi) = \text{MTJ}(\Pi_{\text{bip}})$ and $\text{TAR}(\Pi) = \text{TAR}(\Pi_{\text{bip}})$.

Proof. The definitions of the thresholds imply that $\text{MTJ}(\Pi) \geq \text{MTJ}(\Pi_{\text{bip}})$ and $\text{TAR}(\Pi) \geq \text{TAR}(\Pi_{\text{bip}})$, since $\Pi \supseteq \Pi_{\text{bip}}$. For the reverse direction, assume that the reconfiguration threshold of Π_{bip} (in one of the models) is at most k and consider any graph $G \in \Pi$ with independent sets I and J of equal size. By **Proposition 1** the cost of reconfiguring I to J is bounded by the cost of reconfiguring $I \setminus J$ to $J \setminus I$ in $G[I\Delta J]$. Since $G[I\Delta J]$ is bipartite and Π is hereditary, we have $G[I\Delta J] \in \Pi_{\text{bip}}$, and hence the cost of reconfiguring in $G[I\Delta J]$ is at most k . So reconfiguring I to J can be done with cost at most k in this model. \square

4. Threshold for multiple token jump reconfiguration

We start our discussion of token jump reconfiguration by recalling the following known result.

Theorem 3 ([20, Theorem 7]). Let the graph $G = (V, E)$ be a forest. Then $\text{MTJ}(G) \leq 1$.

The intuition behind this result is that since a forest does not contain any cycle, one can start reconfiguring from the leaf nodes or the isolated vertices, each of which has at most one neighbor from the target configuration. For arbitrary graphs, the above procedure does not work since there may not be any leaves or isolated vertices. But if a graph G has a small vertex cover, then its MTJ reconfiguration threshold is again small, as shown by the following theorem. The main idea is the following. Suppose we have tokens on independent set I , and wish to move them to independent set J , with I and J disjoint. If S is a vertex cover of the graph with $|I \cap S| \geq |J \cap S|$, then we can simultaneously jump all tokens from $S \cap I$ to arbitrary vertices of $J \setminus S$. Since vertices outside S only have neighbors inside S , by jumping all S -tokens away we arrive at a new independent set that brings us closer to J . By repeating this process, and observing that we can also do the reconfiguration “backwards” (moving tokens from J to I), we prove the theorem by induction.

Theorem 4. Let $G = (V, E)$ be a graph. Then $\text{MTJ}(G) \leq \max(\text{vc}(G), 1)$.

Proof. We prove the theorem using induction on the number n of vertices in G . For $n = 1$ the claim is trivially true, so consider a graph $G = (V, E)$ with source and target independent sets I and J of equal size and $|V| > 1$. Our induction hypothesis is that any graph G' with less than $|V|$ vertices has MTJ reconfiguration threshold upper-bounded by $\max(\text{vc}(G'), 1)$.

By **Proposition 1** it is enough to show that in the graph $G[I\Delta J]$ induced by $I\Delta J$, starting from $I' = I \setminus J$, one can construct a sequence of MTJ moves to reach the configuration $J' = J \setminus I$ with step-size at most $\max(\text{vc}(G), 1)$. Let $V' := I\Delta J$. Note that $\text{vc}(G[I\Delta J]) \leq \text{vc}(G)$, and let $S \subseteq V'$ be a vertex cover of $G[I\Delta J]$ of cardinality at most $\text{vc}(G)$. If S is a vertex cover of $G[I\Delta J]$, then there is no edge between any two vertices of $V' \setminus S$. Assume without loss of generality that $|I' \cap S| \geq |J' \cap S|$ (otherwise swap the role of I' and J' , which does not affect reconfigurability). We distinguish three cases.

Case 1. If the vertex cover is empty ($S = \emptyset$), then $G[I\Delta J]$ has no edges. Consequently, all vertex subsets in the graph are independent, and we can reconfigure I' to J' by jumping one token at a time. By **Proposition 1**, this implies I can be reconfigured to J in G using jumps of size 1.

Case 2. Suppose that $|I'| = |J'| \leq \text{vc}(G)$. Then we can jump the tokens from I' onto J' in a single step of size at most $\text{vc}(G)$, and complete the argument using **Proposition 1**.

Case 3. If the previous cases do not apply, we claim that $s := |I' \cap S| > 0$. Indeed, if $s = 0$, then since $|I' \cap S| \geq |J' \cap S|$ we would have $I' \cap S = J' \cap S = \emptyset$, implying that $S = \emptyset$ and that the first case applies. Moreover, we have $|J' \setminus S| \geq s$, otherwise

$$|J'| = |J' \cap S| + |J' \setminus S| \leq (\text{vc}(G) - s) + (s - 1) = \text{vc}(G) - 1, \tag{1}$$

and we are in the previous case. Now let Z be an arbitrary set of s vertices from $J' \setminus S$. Choose all the vertices in $I' \cap S$ and jump their tokens to Z , i.e., remove the vertices $I' \cap S$ from the independent set I' and replace them by Z to obtain I'' . The set I'' is independent because the fact that S is a vertex cover implies that the only neighbors of $Z \subseteq J' \setminus S$ belong to the set S , while I'' contains no vertex from S . Since $|I' \cap S| \leq |S| \leq \text{vc}(G)$, the step size of this move is at most $\text{vc}(G)$.

Consider the graph G' which is obtained from $G[I\Delta J]$ by removing Z and $I' \cap S$, which is again a balanced bipartite graph. Note that since the only neighbors of Z belong to I' (since G is bipartite and $Z \subseteq J'$), and belong to S (since S is a vertex cover and $Z \cap S = \emptyset$), it follows that G' contains no vertex that is a neighbor of Z in $G[I\Delta J]$. Consequently, the union of Z with any independent set in G' is independent in $G[I\Delta J]$. Since G' is smaller than G , by induction one can reconfigure $I' \setminus S$ to $J' \setminus S$ in the graph G' with steps of size at most $\text{vc}(G') \leq \text{vc}(G)$. Adding Z to each set in the corresponding reconfiguration sequence produces a sequence that reconfigures $(I' \setminus S) \cup Z$ to J' in $G[I\Delta J]$. By inserting the steps from I' to $(I' \setminus S) \cup Z$ at the front of this sequence, we obtain a reconfiguration from I' to J' with steps of size at most $\text{vc}(G)$ in $G[I\Delta J]$. By **Proposition 1** this implies that I can be reconfigured to J with steps of size $\text{vc}(G)$ in G , completing the proof for this case. \square

An even cycle of length $2n$ has MTJ reconfiguration threshold n . Since its vertex cover number is n , [Theorem 4](#) is best-possible. Long cycles are not the only graphs whose MTJ reconfiguration threshold equals half the size of the vertex set. Bistable graphs (defined below), of which the pumpkin structure defined in the introduction is a special case, also have this property.

4.1. MTJ reconfiguration threshold in terms of bistable rank

In this section we introduce the notion of *bistable graph*, derive several properties of bistable graphs, and use these to bound the MTJ reconfiguration threshold in terms of the size of the largest induced bistable subgraph. The resulting bounds on the MTJ reconfiguration threshold are tight, but can be hard to apply to specific graph classes: it may be difficult to estimate the size of the largest induced bistable graph, or even to determine whether a given graph is bistable or not. In [Section 4.2](#) we will therefore relate the size of the largest induced bistable subgraph to the size of the largest pumpkin subgraph. This will result in upper- and lower bounds on the MTJ reconfiguration threshold in terms of the largest pumpkin structure contained in the graph (class), which is arguably a more insightful parameter. The resulting bound will not be best-possible, however.

Definition 2 (*Bistable graphs*). A graph is called *bistable* if it is connected, bipartite, and has exactly two distinct maximum independent sets formed by the two partite sets in its unique bipartition. The *rank* of a bistable graph is defined as the size of its maximum independent sets.

Let $\text{bi}(G)$ denote the rank of the largest induced bistable subgraph of G . If G contains no induced bistable subgraphs (which can only occur if G has no edges), then we define $\text{bi}(G)$ to be one. For a graph class \mathcal{H} we define $\text{bi}(\mathcal{H}) := \sup_{G \in \mathcal{H}} \text{bi}(G)$.

The pumpkin shown in [Fig. 1\(a\)](#) forms an example of a bistable graph. We remark that in general, computing $\text{bi}(G)$ is an NP-hard problem.¹ As our interest here is in providing a characterization of the reconfiguration threshold for *restricted hereditary graph families* \mathcal{H} , in which the value of $\text{bi}(\mathcal{H})$ may be easy to compute depending on \mathcal{H} , this does not affect the relevance of our results.

[Lemma 5](#) connects bistable graphs to independent set reconfiguration. Consider the task of reconfiguring the J -partite set to the I -partite set in a balanced bipartite graph $G = (I \cup J, E)$. If we have a set $S \subseteq I$ such that $|S| \geq |N(S)|$, then one way to make progress in the reconfiguration is to select $|S|$ vertices from $N(S) \subseteq J$ and jump their tokens onto the vertices in S , resulting in a new independent set of the same size. The following lemma shows that when we consider a set S that is *minimal* with respect to being at least as large as its neighborhood, then the induced subgraph $G[N[S]]$ is bistable. Hence the cost of such a jump of $|S|$ vertices is bounded by $\text{bi}(G)$, which will allow us to bound the MTJ reconfiguration threshold.

Lemma 5. *Let $G = (I \cup J, E)$ be a balanced bipartite graph without isolated vertices and let $S \subseteq I$ be inclusion-wise minimal with the properties that $|S| \geq |N(S)|$ and S is not empty. Then $G[N[S]]$ is bistable.*

Proof. We have to show that the graph G' induced by the vertices from S and their neighborhood satisfies all conditions for being bistable. Since G is bipartite, G' is as well. Before proving the remaining properties, we establish the following claim.

Claim 1. *The graph G' has a matching covering S .*

Proof. Assume for a contradiction that no such matching exists. By [Fact 1](#), there is a set $S' \subseteq S$ with $|N_G(S')| = |N_{G'}(S')| < |S'|$. As G has no isolated vertices, we have $|S'| > 1$. Removing an arbitrary vertex from S' to obtain S'' then decreases the size of the set by at most one without increasing the neighborhood size. Hence $|N_G(S'')| \leq |S''|$, a contradiction to the minimality of S . \square

We now show that G' has all properties of a bistable graph.

Connectivity. Assume for a contradiction that G' is not connected. If G' has a connected component with vertex set C that contains at least as many I -vertices as J -vertices, then $S' := C \cap I$ is a strict subset of S with $|S'| \geq |C \cap J| \geq |N_{G'}(S')| = |N_G(S')|$, contradicting minimality of S . Otherwise, all connected components of G' have strictly more J -vertices than I -vertices. Since the J -vertices in G' form the neighborhood of S , this implies that $|N_G(S)| > |S|$, a contradiction to the choice of S . Hence G' is connected.

Balance. Since G' is connected, it has a unique bipartition and it is easy to verify that S is one of the partite sets: $G' = (S \cup J', E')$. Since there is a matching covering S ([Claim 1](#)) and all matching partners of vertices in S are distinct and belong to J' , we therefore have $|J'| \geq |S|$. We have $|J'| = |N_G(S)| \leq |S|$ by assumption on S , establishing that $|J'| = |S|$ which

¹ The classic [\[21\]](#) reduction from 3SAT to INDEPENDENT SET transforms a formula ϕ with n variables and m clauses of size exactly three, into a graph that has a vertex for each literal. There are edges between literals that belong to the same clause and between every pair of contradicting literals. The resulting graph H has an independent set of size m if and only if ϕ is satisfiable. If one forms a new graph G by taking two vertex-disjoint copies H_1, H_2 of H , making each vertex of H_1 adjacent to each vertex of H_2 , it is easy to see that $\text{bi}(G) \geq m$ if and only if ϕ is satisfiable, establishing NP-hardness.

proves G' is balanced. This implies that a matching in G' that saturates S (which exists by Claim 1) is in fact a perfect matching in G' .

Two maximum independent sets. Assume for a contradiction that G' has at least three maximum independent sets. Then there is a maximum independent set in G that is not equal to either of the two partite sets J' or S ; let X be such a maximum independent set. Since G' is bipartite and has a perfect matching M , the set X contains exactly one vertex from each matching edge in M . Now let $\hat{S} := X \cap I$. Since $X \neq J'$ by assumption, it follows that \hat{S} is not empty; since $X \neq S$, it is a proper subset of S . We show that $|N_G(\hat{S})| \leq |\hat{S}|$, contradicting minimality of S .

Let $M' \subseteq M$ denote the matching edges intersected by \hat{S} . Since X contains one vertex from each edge of M , for all edges in $M \setminus M'$ the J -endpoint of the edge belongs to the independent set X . So the J -endpoint of an edge in $M \setminus M'$ is not in the neighborhood of \hat{S} , as X is independent. Consequently, only the matching partners of \hat{S} can be in the neighborhood of \hat{S} , implying there are at most $|\hat{S}|$ such neighbors. Hence $|N_G(\hat{S})| \leq |\hat{S}|$; a contradiction. It follows that G' has at most two maximum independent sets. To see that it has exactly two, it suffices to observe that since G' has a perfect matching, both its partite sets are maximum independent sets.

This establishes that G' satisfies all conditions for being bistable and concludes the proof of Lemma 5. \square

Now, in the lemma below, we prove two key properties of bistable graphs. They will later be useful to relate the quantities $\text{PUM}(G)$ and $\text{BI}(G)$.

Lemma 6. *Let $G = (I \cup J, E)$ be a bistable graph. Then the following holds:*

1. G has a perfect matching covering I (and hence J).
2. G is biconnected.

Proof. 1. By König's theorem (cf. [30, Thm. 16.2]), the size of a maximum matching in the bipartite graph G equals the size of a minimum vertex cover in G . By Definition 2, the partite sets I and J are maximum independent sets and therefore have equal size. Since the complement of a maximum independent set is a minimum vertex cover, it follows that $V(G) \setminus I = J$ is a minimum vertex cover. Hence there is a matching of size $|J| = |I|$ in G , which is a perfect matching since it covers $2|J| = |V(G)|$ vertices.

2. Assume for a contradiction that G is not biconnected. Let v be a cutvertex and let M be a perfect matching in G , which exists by the previous property. Assume that $v \in I$; the argument for $v \in J$ is symmetric. Let u be the matching partner of v under M . Since v is a cutvertex, the graph $G - \{v\}$ consists of multiple connected components C_1, \dots, C_ℓ . Without loss of generality, assume that u is contained in component C_1 . For all components C_i with $i \geq 2$, the component contains the same number of I and J -vertices: for each vertex its matching partner in the opposite partite set belongs to the same component. For component C_1 the number of I -vertices is one smaller than the number of J -vertices, since the matching partner of u does not belong to C_1 . Consider the set S consisting of the J -vertices from C_1 along with the I -vertices of all other components C_2, \dots, C_ℓ of $G - \{v\}$. The set S is independent in $G - \{v\}$ since it consists of entire partite sets of different components of the bipartite graph. Since $v \notin S$ it follows that S is also independent in G . As S contains exactly one endpoint from each edge in M (the I -endpoint for matching edges intersecting a component C_i for $i \geq 2$, and the J -endpoint for the remaining matching edges) it follows that S is a maximum independent set in G that differs from I and J ; a contradiction to Definition 2. \square

Theorem 7. *For any graph G it holds that $\text{MTJ}(G) \leq \text{BI}(G)$. Moreover, if $G \neq K_1$, then there exists an induced subgraph G' of G with $\text{MTJ}(G') \geq \text{BI}(G) \geq \text{BI}(G')$.*

Proof. We first prove the lower bound on $\text{MTJ}(G)$. Note that if $G \neq K_1$ contains no induced bistable subgraphs, then G is a collection of isolated vertices, and in that case $\text{MTJ}(G) = \text{BI}(G) = 1$. Assume then that G contains a nonempty induced bistable subgraph $G' = (I' \cup J', E')$ of rank $\text{BI}(G)$. By Definition 2, the sets I' and J' are the only independent sets of size $\text{BI}(G)$ in G' . It follows that in any MTJ reconfiguration sequence from I' to J' , the set I' is immediately followed by J' which requires a jump of $|I'| = |J'| = \text{BI}(G)$ tokens simultaneously. Hence $\text{MTJ}(G') \geq \text{BI}(G)$.

We prove the upper bound on $\text{MTJ}(G)$ by induction on the size of the graph. If G consists of a single vertex, then there is a unique nonempty independent set, so $\text{MTJ}(G) = 0$. In the remainder, assume G has more than one vertex and let I and J be two independent sets in G of equal size. By Proposition 1 it suffices to prove that $I' := I \setminus J$ can be MTJ -reconfigured to $J \setminus I$ in the graph $G' := G[I \Delta J]$ with jumps of size at most $\text{BI}(G)$. Assume first that G' has no isolated vertices, and let $S \subseteq I \setminus J$ be an inclusion-wise minimal nonempty subset of $I \setminus J$ with the property that $|S| \geq |N_{G'}(S)|$. Such a set exists since G' is a balanced bipartite graph with partite sets $I \setminus J$ and $J \setminus I$, so the set $I \setminus J$ satisfies the stated condition (but may not yet be minimal). It is easy to verify that since G' has no isolated vertices and S is minimal, we have $|N_{G'}(S)| = |S|$. Now move all tokens from $N_{G'}(S)$ onto S in a single jump of size $|S|$. By Lemma 5, the graph $G'[N_{G'}(S)] = G[N_{G'}(S)]$ is a bistable induced subgraph of G of rank $|S|$, and therefore $\text{BI}(G) \geq |S|$ which shows that the size of the jump is sufficiently small. We may then invoke induction similarly as in the proof of Theorem 4 to complete the argument. If G' has an isolated vertex, then instead one can jump a token onto this isolated vertex and induct. This concludes the proof of Theorem 7. \square

The following corollary characterizes the MTJ reconfiguration threshold of hereditary graph classes. It follows directly from [Theorem 7](#). It applies to all graph classes except the one consisting only of the single graph K_1 with a single vertex, for which the reconfiguration threshold is zero but $\text{BI}(K_1) = 1$ by definition.

Corollary 1. *For any hereditary graph class $\Pi \neq \{K_1\}$ it holds that $\text{MTJ}(\Pi) = \text{BI}(\Pi)$.*

Proof. By [Theorem 7](#) we have $\text{MTJ}(G) \leq \text{BI}(G) \leq \text{BI}(\Pi)$ for all graphs $G \in \Pi$, hence $\text{MTJ}(\Pi) \leq \text{BI}(\Pi)$. To prove the converse, consider an arbitrary graph $G \in \Pi$ having at least one edge. Then G contains an induced bistable subgraph $H = (I \cup J, E)$ of rank $\text{BI}(G)$, and since Π is hereditary we have $H \in \Pi$. Reconfiguring I to J in H requires a jump of size $|I| = |J| = \text{BI}(G)$ since those are the only two independent sets of that size. Hence $\text{MTJ}(\Pi) \geq \text{BI}(G)$ for all $G \in \Pi$ with at least one edge, showing that $\text{MTJ}(\Pi) \geq \text{BI}(\Pi)$ if Π contains at least one bistable graph. In the exceptional setting that Π contains no bistable graph, all graphs in Π are edgeless causing $\text{BI}(\Pi)$ to be one. Since $\Pi \neq \{K_1\}$, the graph consisting of two isolated vertices is contained in Π , which has reconfiguration threshold one. Hence the lower bound also holds in this case. \square

4.2. MTJ reconfiguration threshold in terms of pumpkin size

In this section we formally introduce the pumpkin structure described in the introduction. We relate pumpkins to bistable graphs to obtain bounds on the MTJ reconfiguration threshold in terms of the size of the largest pumpkin subgraph.

Definition 3 (Pumpkin). *A pumpkin is a graph consisting of two terminal vertices u and v linked by two or more vertex-disjoint paths with an odd number of edges, having no edges or vertices other than those on the paths. A path can consist of the single edge $\{u, v\}$. The size of the pumpkin is the total number of vertices.*

For a graph G we denote by $\text{PUM}(G)$ the size of the largest (not necessarily induced) subgraph isomorphic to a pumpkin that is contained in G , or zero if G contains no pumpkin. For a graph class Π we define $\text{PUM}(\Pi) := \sup_{G \in \Pi} \text{PUM}(G)$.

An example of a pumpkin structure is shown in [Fig. 1\(a\)](#). Observe that a pumpkin is a bipartite graph, since all cycles consist of two uv -paths of odd length and are therefore even. Furthermore, a pumpkin is a *balanced* bipartite graph: vertices u and v belong to different partite sets since their distance is odd, and on every (odd-length) uv -path in the structure there is an even number of interior vertices, which alternate between the two partite sets. It is not difficult to verify that the two partite sets are the only maximum independent sets in a pumpkin, leading to the following observation.

Observation 1. *Every pumpkin graph is bistable.*

The next theorem shows that the rank of the largest bistable induced subgraph of G can be upper-bounded in terms of the size of G 's largest pumpkin subgraph. For completeness, we remark that computing $\text{PUM}(G)$ is an NP-hard problem in general, just as computing $\text{BI}(G)$.² For the hereditary graph classes in which one wishes to bound the reconfiguration threshold, however, the problem may be tractable.

The proof of [Theorem 8](#) combines several properties of bistable graphs derived above. We first bound the length of a longest path in a bistable graph G in terms of its largest pumpkin, using the fact that a biconnected graph with a path of length more than L^2 has a cycle of length more than L [[12](#)], which forms a pumpkin since G is bipartite. So if we set $L := \text{PUM}(G)$, a bistable graph G cannot have a path of length more than L^2 . Consequently, if we build a DFS tree T rooted at an arbitrary vertex, its depth will be at most L^2 . We then prove that each vertex v in T has at most $L^3 + L^2$ children using the fact that G has a perfect matching, which in conjunction with the above bound on depth, yields the theorem.

Theorem 8. *For any bistable graph G , $\text{BI}(G) \leq f(\text{PUM}(G))$, where $f(k) = (k^3 + k^2)^{k^2+1} + 1$.*

Proof. Consider a bistable graph $G = (I \cup J, E)$. If G is acyclic, then any biconnected subgraph of G contains at most two vertices. There is a unique bistable graph with at most two vertices, which consists of a single edge and has rank one. Since any bistable graph is biconnected by [Lemma 6](#), we have $\text{BI}(G) \leq 1 = f(0) \leq f(\text{PUM}(G))$ if G is acyclic. In the remainder we assume that G contains a cycle, which implies that $\text{PUM}(G) \geq 1$ since any cycle in the bipartite graph G is even and forms a pumpkin.

For ease of notation, define $L := \text{PUM}(G)$. Construct a depth-first search (DFS) tree T of G , starting at an arbitrary vertex r which becomes the root of the tree. By the structure of the DFS process, we obtain the following property: if u and v are vertices of G that are adjacent in G , then u is an ancestor of v in T , or v is an ancestor of u . For $v \in T$, we use T_v to denote the subtree of T rooted at v . We will often use T_v to refer to the vertices in the tree as well.

Claim 2. *The depth of T is at most L^2 .*

² Reduce from the NP-complete problem HAMILTONIAN ST-PATH IN BIPARTITE GRAPHS (cf. [[4](#), Proposition 1]), which asks whether a bipartite graph H with bipartition $A \cup B = V(H)$ and vertices $s \in A, t \in B$ has a Hamiltonian path between s and t . Form graph G from three vertex-disjoint copies of H , adding one vertex s^* that is adjacent to the three copies of s , and one vertex t^* that is adjacent to the three copies of t . Then $\text{PUM}(G) = |V(G)|$ if and only if H has a Hamiltonian st -path.

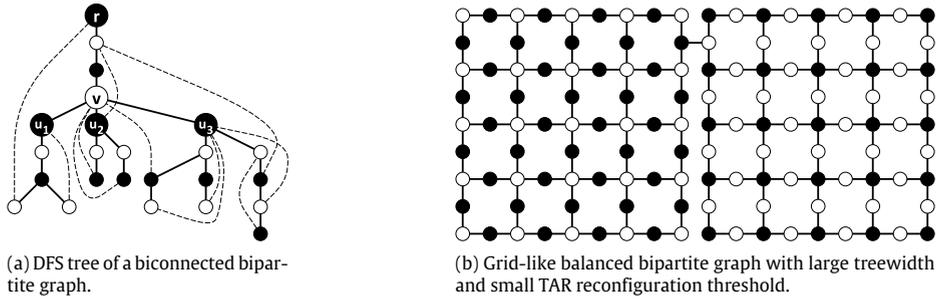


Fig. 2. 2(a) Depth-first search tree of a bipartite biconnected graph. Tree-edges are drawn solid, while the remaining edges of G are drawn with dotted lines. The three children u_1, u_2, u_3 of v induce subtrees of types A, B, and C, respectively. 2(b) Template for constructing graphs of large treewidth that can be TAR reconfigured with a buffer of size two. The treewidth is large due to the presence of a large grid minor.

Proof. Assume for a contradiction that there is a path from the root r of T to a leaf ℓ , consisting of more than L^2 edges. By Lemma 6, graph G is biconnected. The existence of a path of more than L^2 edges in a biconnected graph G is known [12, Theorem 1] to imply that G contains a simple cycle of length more than L . Since G is bipartite the cycle is even and forms a pumpkin: it splits into two odd paths. So $\text{PUM}(G) > L$, a contradiction. \lrcorner

If each vertex in T has at most $(L^3 + L^2)$ children, then the bound on the depth given by the previous claim implies that T (and therefore G) has at most $\sum_{i=0}^{L^2} (L^3 + L^2)^i = (L^3 + L^2)^{L^2+1} + 1$ vertices and therefore:

$$\text{BI}(G) \leq |V(G)| \leq (L^3 + L^2)^{L^2+1} + 1 \leq f(L) = f(\text{PUM}(G)). \tag{2}$$

To complete the proof, it therefore suffices to show that no vertex of T has more than $L^3 + L^2$ children. Assume for a contradiction that some vertex v exists with a larger number of children u_1, \dots, u_m , for some $m > L^3 + L^2$. By switching the roles of I and J if needed, we may assume that $v \in I$. For a vertex w , let M_w denote the set of its proper ancestors in T . We classify the children u_i for $i \in [m] = \{1, 2, \dots, m\}$ into three types:

- Type A: Some vertex of T_{u_i} has an edge in G to a vertex in $M_u \cap J$.
- Type B: Vertex u is not of type A and $|I \cap T_u| \neq |J \cap T_u|$.
- Type C: Vertex u is not of type A and $|I \cap T_u| = |J \cap T_u|$.

Observe that any child of v belongs to exactly one type.

Claim 3. *There are less than L^3 type-A children of v .*

Proof. Suppose there are at least L^3 type-A children of v , and assume these are numbered u_1, \dots, u_{L^3} . Each subtree T_{u_i} for $i \in [L^3]$ contains a vertex that has an edge in G to a proper ancestor of u_i in J , by our definition of types. Since $v \in I$, this cannot be v so that in fact it is also a proper ancestor of v and belongs to M_v .

Since T has depth at most L^2 by Claim 2, by the pigeon-hole principle there is a vertex $w \in M_v \cap J$, such that L subtrees among $T_{u_1}, \dots, T_{u_{L^3}}$ contain a vertex that is adjacent to w in G . From each such subtree T_{u_i} , we obtain a path in G from v to w whose internal vertices belong to T_{u_i} , by going from v to u_i , then to a neighbor of w in the subtree T_{u_i} using the tree edges, and ending with the edge to w . Applying this procedure to each of the L subtrees that connect to w yields at least L internally vertex-disjoint paths from v to w . Since G is bipartite and v and w belong to different partite sets, each path connecting v and w is of odd length. Hence this collection of L vertex-disjoint paths between v and w forms a pumpkin of size more than L : each of the L paths has at least one internal vertex, and together with v and w this gives size at least $L + 2$. This contradicts our choice of $L = \text{PUM}(G)$. \lrcorner

Claim 4. *There are at most $\text{DEPTH}(v) + 1$ type-B children of v .*

Proof. Since G is bistable, it has a perfect matching M by Lemma 6. By the properties of a DFS tree, for each vertex in T its neighbors in G are among its ancestors and descendants in T . The number of I and J -nodes in a subtree T_{u_i} rooted at a type-B child u_i are not equal. Since each vertex in T_{u_i} is assigned a unique neighbor in the other partite set by the perfect matching M , it follows that the matching partner of some vertex in T_{u_i} does not belong to T_{u_i} , and must therefore be a proper ancestor of u_i by the properties of DFS trees. Since there are at most $\text{DEPTH}(v) + 1$ ancestors of v to use as matching partners, and each type-B child uses a different ancestor as a matching partner for one of its vertices, the number of type-B subtrees is at most $\text{DEPTH}(v) + 1$. \lrcorner

Since the depth of T is at most L^2 by Claim 2, any vertex v that is not a leaf has depth at most $L^2 - 1$. Hence the number of type-B children of v is at most L^2 .

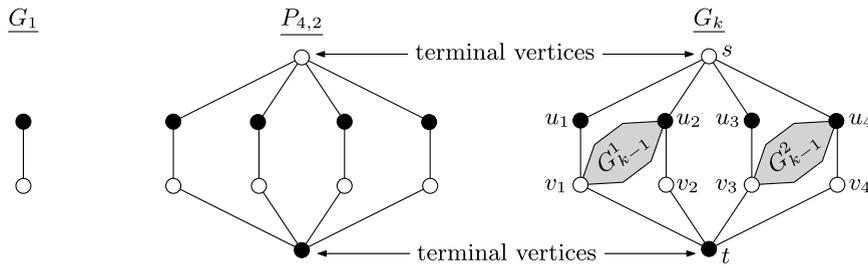


Fig. 3. Construction of a super-pumpkin.

Claim 5. No child of v is of type C.

Proof. Suppose there exists a type-C child u_i of v . Subtree T_{u_i} does not contain a vertex adjacent to $M_{u_i} \cap J$, else u_i would have been type A. Any G -neighbor of a vertex in T_{u_i} that is not contained in T_{u_i} is a proper ancestor of u_i , by the properties of DFS trees. Hence the set $J' = (J \setminus T_{u_i}) \cup (I \cap T_{u_i})$ forms an independent set in G . By the definition of type-C vertices, $|J \cap T_{u_i}| = |I \cap T_{u_i}|$, so that $|J'| = |J|$. This shows that J' is a maximum independent set distinct from I and J , contradicting the assumption that G is bistable. \square

The preceding claims show that no vertex of T has more than $L^2 + L^3$ children, which completes the proof of Theorem 8 using (2). \square

The following theorem is our main result on the MTJ reconfiguration threshold. It bounds the MTJ reconfiguration threshold of a hereditary graph class Π in terms of the maximum size of pumpkin subgraph of a graph in Π_{bip} . Recall that Π_{bip} contains the bipartite graphs in Π .

Theorem 9. For any hereditary graph class Π , the following holds:

$$g_1(\text{PUM}(\Pi_{\text{bip}})) \leq \text{MTJ}(\Pi) \leq g_2(\text{PUM}(\Pi_{\text{bip}})), \tag{3}$$

where $g_1, g_2 : \mathbb{N} \rightarrow \mathbb{N}$ are positive non-decreasing functions defined as $g_1(k) = k/2$ and $g_2(k) = (k^3 + k^2)^{k^2+1} + 1$. Moreover, for every graph G we have $\text{MTJ}(G) \leq g_2(\text{PUM}(G))$.

Proof. We combine the bounds on the MTJ reconfiguration threshold of Theorem 7, with the relation between pumpkins and bistable graphs of Theorem 8.

Lower bound on MTJ. Consider a bipartite graph G in Π_{bip} . Then G contains a pumpkin subgraph on $\text{PUM}(G)$ vertices $S \subseteq V(G)$. Then $G[S] = (I \cup J, E)$ is a bipartite supergraph of a pumpkin, which is contained in Π_{bip} . Since any pumpkin is bistable by Observation 1, it follows that reconfiguring I to J in the pumpkin subgraph of $G[S]$ requires a jump of size $|I| = |J| = \text{PUM}(G)/2$. It is clearly no easier to reconfigure I to J in the supergraph $G[S] \in \Pi_{\text{bip}}$. Hence $\text{MTJ}(\Pi) \geq \text{PUM}(G)/2$ for all graphs G in Π_{bip} , giving the lower bound $\text{MTJ}(\Pi) \geq \text{PUM}(\Pi_{\text{bip}})/2$.

Upper bound on MTJ. By Proposition 2 note that it suffices to prove that $\text{MTJ}(\Pi_{\text{bip}}) \leq g_2(\text{PUM}(\Pi_{\text{bip}}))$. Consider an arbitrary graph $G \in \Pi_{\text{bip}}$. By Theorem 7, we have $\text{MTJ}(G) \leq \text{BI}(G) \leq g_2(\text{PUM}(G)) \leq g_2(\text{PUM}(\Pi_{\text{bip}}))$, where the second inequality follows from Theorem 8. Hence $\text{MTJ}(\Pi_{\text{bip}}) \leq g_2(\text{PUM}(\Pi_{\text{bip}}))$, concluding the proof. \square

While the upper bound of Theorem 9 has room for improvement, the following lemma shows that the exponential dependency on the pumpkin size in the upper bound is unavoidable.

Proposition 10. Let $\Pi_{\text{PUM}}(k) := \{G : \text{PUM}(G) \leq k\}$ be the class of all graphs G whose largest pumpkin subgraph has size at most k . Then $\text{MTJ}(\Pi_{\text{PUM}}(k)) = 2^{\Omega(k)}$.

Proof. For each $k \geq 1$ we will construct a graph G_k belonging to the class $\Pi_{\text{PUM}}(24k + 6)$ with $\text{MTJ}(G_k) = 2^{\Omega(k)}$. We will call G_k a super-pumpkin. It is defined recursively, as explained next.

Define a (4,2)-pumpkin, denoted $P_{4,2}$, to be a pumpkin whose two terminal vertices are connected by four paths with two interior vertices each. A super-pumpkin is now defined as follows. Like a regular pumpkin, it has two designated terminal vertices. The super-pumpkin G_1 consists of just a single edge, whose endpoints are its terminal vertices. The super-pumpkin G_k is obtained by gluing two copies of a super-pumpkin G_{k-1} – we will denote these copies by G_{k-1}^1 and G_{k-1}^2 – into a (4,2)-pumpkin $P_{4,2}$. This is done by identifying the terminal vertices of G_{k-1}^1 and G_{k-1}^2 with specific vertices of the (4,2)-pumpkin, as indicated in Fig. 3.

Note that $|G_k|$, the number of vertices of G_k , satisfies $|G_k| = 2|G_{k-1}| + 6$ with $|G_1| = 2$. Hence, $|G_k| = 2^k + 6 \sum_{i=0}^{k-2} 2^i = 2^k + 6(2^{k-1} - 1)$.

Claim 6. G_k has exactly two independent sets of size $|G_k|/2$, and these independent sets are disjoint.

Proof. The proof is by induction on k . It will be convenient to prove the following stronger claim on $\mathcal{I}(G_k)$, the set of all independent sets of G_k .

$\mathcal{I}(G_k)$ contains no independent set of size more than $|G_k|/2$ and exactly two independent sets of size $|G_k|/2$. These independent sets are disjoint, and one of them contains one terminal vertex of G_k while the other contains the other terminal vertex.

This claim trivially holds for $\mathcal{I}(G_1)$, so now consider $\mathcal{I}(G_k)$ for $k > 1$. Let s, t be the two terminal vertices of G_k , and label the other vertices of the pumpkin $P_{4,2}$ as u_1, \dots, u_4 and v_1, \dots, v_4 —see Fig. 3. We define $W := \{s, t, u_1, u_3, v_2, v_4\}$ to be the set of vertices in G_k that do not occur in G_{k-1}^1 or G_{k-1}^2 . We distinguish two types of independent sets in $\mathcal{I}(G_k)$.

Type 1. Independent sets I such that both G_{k-1}^1 and G_{k-1}^2 have $|G_{k-1}|/2$ vertices in I . Note that by the induction hypothesis we have $|\{u_2, v_1\} \cap I| = |\{u_4, v_3\} \cap I| = 1$ for any Type 1 independent set I . Moreover, the total number of vertices from I inside G_{k-1}^1 and G_{k-1}^2 is $2(|G_{k-1}|/2) = |G_k|/2 - 3$. We will argue that G_k has two Type 1 independent sets with $|G_k|/2$ vertices and with the required properties, and that all other Type 1 independent sets have less than $|G_k|/2$ vertices. To this end we distinguish three subtypes of Type 1.

- **Type 1(i).** Independent sets I with $u_2 \in I$ and $u_4 \in I$. By the induction hypothesis such independent sets I exist, and the choice of vertices of I inside G_{k-1}^1 and G_{k-1}^2 is fixed. Moreover, there is only one way to obtain an independent set I^* with $|G_k|/2$ vertices, namely by adding $\{u_1, u_3, t\}$ from W —all other selections from W give smaller independent sets.
- **Type 1(ii).** Independent sets I with $v_1 \in I$ and $v_3 \in I$. Again, the choice of vertices for I inside G_{k-1}^1 and G_{k-1}^2 is fixed, and there is only one way to obtain an independent set of $|G_k|/2$ vertices, this time by adding $\{s, v_2, v_4\}$. This independent set I^{**} is disjoint from I^* —this follows from the induction hypothesis and the fact $\{u_1, u_3, t\} \cap \{s, v_2, v_4\} = \emptyset$ —and it contains s while I^* contains t .
- **Type 1(iii).** Independent sets I with $u_2 \in I$ and $v_3 \in I$, or $v_1 \in I$ and $u_4 \in I$. Now at most two of the vertices from W can be in I , and so $|I| < |G_k|/2$.

Type 2. Independent sets I such that at least one of G_{k-1}^1 and G_{k-1}^2 has less than $|G_{k-1}|/2$ vertices in I . We will argue that all such independent sets have less than $|G_k|/2$ vertices.

Assume without loss of generality that G_{k-1}^1 has less than $|G_{k-1}|/2$ vertices in I . If G_{k-1}^2 also has less than $|G_{k-1}|/2$ vertices in I , then the total number of vertices from I in G_{k-1}^1 and G_{k-1}^2 is at most $2(|G_{k-1}|/2 - 1) = |G_k|/2 - 5$, and since at most four vertices can be selected from W we have $|I| < |G_k|/2$. If G_{k-1}^2 has $|G_{k-1}|/2$ vertices in I , then $|\{u_4, v_3\} \cap I| = 1$. Assume without loss of generality that $u_4 \in I$. Then s and v_4 are not in I . Since v_2 and t cannot be both in I , we conclude that we can select at most three vertices from W into I . This again implies that $|I| < |G_k|/2$.

Note that each $I \in \mathcal{I}(G_k)$ is of Type 1 or Type 2, since G_{k-1}^1 and G_{k-1}^2 cannot have more than $|G_{k-1}|/2$ vertices in I by the induction hypothesis. This finishes the proof of the claim. \square

Claim 6 implies that $\text{MTJ}(G_k) = |G_k|/2 = 2^{\Omega(k)}$: since there are only two independent sets of size $|G_k|/2$, say I and J , and these are disjoint, the only way to go from I to J is to remove all tokens from I and place them onto J . Next we bound the size of the largest pumpkin in G_k .

Claim 7. $\text{PUM}(G_k) \leq 24k + 6$.

Proof. Define d_{\max} to be the maximum degree in G_k and C_k to be the maximum length of any simple cycle in G_k . Then $\text{PUM}(G_k) \leq d_{\max} \cdot C_k/2$.

The following statement is easy to prove by induction: the degree of the terminal vertices in G_k is four, and the maximum degree of any other vertex in G_k is six. Hence, $d_{\max} = 6$ and so $\text{PUM}(G_k) \leq 3C_k$.

Next we argue that $C_k \leq 8k + 2$. To this end, define L_k to be the length (measured in number of vertices) of a longest simple path in G_k that ends at the two terminal vertices of G_k . Then $L_k = L_{k-1} + 4$ with $L_1 = 2$, and thus $L_k = 4k - 2$. We now prove that $C_k \leq 8k + 2$ by induction on k . We have $C_1 = 0$, so the statement is true for $k = 1$. Now suppose $k > 1$. Let \mathcal{C} be a simple cycle in G_k . If \mathcal{C} stays within one of the copies of G_{k-1} we have $|\mathcal{C}| \leq C_{k-1}$ by induction. Otherwise the maximum possible length for \mathcal{C} is obtained by taking a longest path from u_2 to v_1 in the first copy of G_{k-1} , a longest path from u_4 to v_3 in the second copy, and connecting them into a cycle using all six vertices in W , where W is defined as before. Hence,

$$C_k \leq \max(C_{k-1}, 2L_{k-1} + 6) = \max(C_{k-1}, 8k + 2).$$

It follows that $C_k \leq 8k + 2$. Hence, $\text{PUM}(G_k) \leq 3C_k \leq 24k + 6$. \square

This concludes the proof of Proposition 10. \square

5. Threshold for token addition removal reconfiguration

In this section we study the model of token additional removal. First observe that when G is a forest, we have $\text{MTJ}(G) \leq 1$ and therefore $\text{TAR}(G) \leq 1$ as well. Also, from Theorem 4 we get $\text{TAR}(G) \leq \max(\text{vc}(G), 1)$. But the inequality $\text{TAR}(G) \leq \text{MTJ}(G)$

tells us nothing about the behavior of the TAR reconfiguration threshold when the MTJ reconfiguration threshold is large. The next simple proposition immediately points towards this direction. Indeed observe that a large pumpkin (with large MTJ reconfiguration threshold) can have a small feedback vertex set; this happens for even cycles, for example. By first placing the tokens on a feedback vertex set into the buffer, and then reconfiguring towards the leaves of the resulting forest, we can bound the TAR reconfiguration threshold of G by $FVS(G) + 1$.

Proposition 11. *Let $G = (V, E)$ be a graph. Then $TAR(G) \leq FVS(G) + 1$.*

Proof. Let $G = (V, E)$ be a graph with a minimum feedback vertex set $S \subseteq V$ of size k , and let $I, J \subseteq V$ be independent sets of equal size. By Proposition 1 we can assume that $V = I \cup J$ and $I \cap J = \emptyset$. If $|I| = |J| \leq k$, then it is trivial to reconfigure I to J with a buffer of size at most k , by first moving all tokens from I into the buffer, and then onto J . In the remainder we assume $|I| = |J| > k$. Let S_I be a superset of $I \cap S$ of size k , and let S_J be a superset of $J \cap S$ of size k . Then the graph $G' := G - (S_I \cup S_J)$ is a subgraph of $G - S$ and is therefore acyclic since S is a feedback vertex set. By Theorem 3 it follows that $I' := I \setminus S_I$ can be MTJ reconfigured to $J' := J \setminus S_J$ in G' by jumps of size 1, which easily implies that I' can be TAR reconfigured to J' in G' using a buffer of size at most 1; let \mathcal{S} be a corresponding reconfiguration sequence. To reconfigure I to J in G , start by removing the tokens from the k vertices in S_I and place them in the buffer. Then apply the reconfiguration sequence \mathcal{S} to reconfigure I' to J' , using at most 1 extra buffer token. Finish by moving the k buffer tokens onto S_J to arrive at the independent set $J' \cup S_J = J$. \square

One can see that the above bound is tight, by considering the TAR reconfiguration threshold of a complete balanced bipartite graph. Indeed for $K_{n,n}$, the minimum size of a feedback vertex set is $n - 1$, and one can see that in order to include any one of the vertices of the target independent set the reconfiguration must pass through the empty set. This shows that the TAR reconfiguration threshold is also n .

5.1. TAR reconfiguration threshold in terms of pathwidth

As the main result of this section, we will show that the TAR reconfiguration threshold of a graph is upper-bounded in terms of its pathwidth. Before proving that statement, we present a structural lemma about path decompositions that will be useful in the proof.

Lemma 12. *Let $G = (I \cup J, E)$ be a bipartite graph with a nice path decomposition $\mathcal{P} = (X_1, \dots, X_r)$ of width k . Let $S \subseteq J$ such that $|N(S)| \leq |S|$ while no non-empty subset of S has this property. If we order the vertices in S as i_1, \dots, i_t such that $r_{\mathcal{P}}(i_1) < r_{\mathcal{P}}(i_2) < \dots < r_{\mathcal{P}}(i_t)$, then $|N(\{i_1, \dots, i_{t'}\})| < t' + k$ for all $1 \leq t' \leq t$.*

Intuitively, the lemma says the following. Suppose a set $S \subseteq J$ is inclusion-wise minimal with respect to being no smaller than its neighborhood. Then ordering S according to the right endpoints of the intervals representing S in the path decomposition, we are guaranteed that every prefix of S has a fairly small neighborhood compared to its size: the neighborhood size exceeds the size of the prefix by less than the pathwidth. Note that since the lemma deals with bipartite graphs only, no vertex of S can belong to the neighborhood of any prefix of S . The ordering of the vertices is uniquely defined since the path decomposition is nice. The bound of Lemma 12 is best-possible. Consider a complete bipartite graph $K_{n,n}$, with pathwidth n . In any optimal path decomposition, for $t' = 1$ the first vertex in the ordering has a neighborhood of size n and so $n < t' + n = 1 + n$, but a better bound is not possible.

Proof of Lemma 12. First observe that in a graph with a path decomposition of width $k = 0$ there can be no edges. Then the only vertex-minimal set S satisfying the assumptions is an isolated vertex, for which the claim trivially holds. In the remainder we assume $k \geq 1$. For $t' = t$ we have $\{i_1, \dots, i_{t'}\} = S$, and by assumption $|N(S)| \leq |S|$. So for $t' = t$ the claim in the lemma holds trivially for any $k \geq 1$. Assume for a contradiction that there is some $t' < t$ such that:

$$|N(\{i_1, \dots, i_{t'}\})| - t' \geq k. \tag{4}$$

We partition $T := S \cup N(S)$ into three disjoint subsets to derive some structural properties that will lead to a contradiction.

- (i) $T_1 := \{v \in S \cup N(S) : r_{\mathcal{P}}(v) \leq r_{\mathcal{P}}(i_{t'})\}$, the set of all vertices in $S \cup N(S)$ that are not contained in any of the bags after the bag with index $r_{\mathcal{P}}(i_{t'})$.
- (ii) $T_2 := \{v \in S \cup N(S) : l_{\mathcal{P}}(v) > r_{\mathcal{P}}(i_{t'})\}$, the set of all vertices in $S \cup N(S)$ that are not contained in any of the bags before or including the bag with index $r_{\mathcal{P}}(i_{t'})$.
- (iii) $T_3 := \{v \in S \cup N(S) : l_{\mathcal{P}}(v) \leq r_{\mathcal{P}}(i_{t'}) < r_{\mathcal{P}}(v)\}$, the set of all vertices in $S \cup N(S)$ that are contained in some bags before or including the bag with index $r_{\mathcal{P}}(i_{t'})$ and also in some bag after it.

Observe that $(T_1 \cap S) \cup (T_2 \cap S) \cup (T_3 \cap S)$ is a partition of S , and that $T_1 \cap S = \{i_1, \dots, i_{t'}\}$.

Claim 8. $|T_3| \leq k$.

Proof. From property (P3) in the definition of path decomposition we know that $T_3 \subseteq X_\ell$ for $\ell = r_{\mathcal{P}}(i_{t'})$. Now $|X_\ell| \leq k + 1$ since the width of \mathcal{P} is at most k , and we know that $i_{t'} \in X_\ell \setminus T_3$. Therefore we have $|T_3| \leq |X_\ell \setminus \{i_{t'}\}| \leq k$. \square

For the remainder of the proof we distinguish two cases.

Case 1: $T_2 \cap S = \emptyset$. Then $(T_1 \cap S) \cup (T_3 \cap S)$ is a partition of S . By Claim 8 we have $|T_3 \cap S| \leq k$, and therefore

$$|S| = |T_1 \cap S| + |T_3 \cap S| \leq |T_1 \cap S| + k = t' + k. \tag{5}$$

Claim 9. In Case 1 we have $N(\{i_{t'+1}, \dots, i_t\}) \subseteq N(\{i_1, \dots, i_{t'}\})$.

Proof. Assume for a contradiction that $v \in N(\{i_{t'+1}, \dots, i_t\}) \setminus N(\{i_1, \dots, i_{t'}\})$. By (4) we have $|N(\{i_1, \dots, i_{t'}\})| \geq t' + k$, and the existence of v shows that

$$|N(S)| = |N(\{i_1, \dots, i_t\})| > |N(\{i_1, \dots, i_{t'}\})| \geq t' + k \geq |S|,$$

by (5). But this contradicts the starting assumption that $|N(S)| \leq |S|$. \square

Claim 10. In Case 1 we have $|T_3 \cap S| \geq k$, implying that $T_3 \subseteq S$ and $|T_3 \cap S| = k$.

Proof. Suppose that $|T_3 \cap S| < k$. Then:

$$\begin{aligned} |N(S)| &\geq |N(\{i_1, \dots, i_{t'}\})| && \text{since } S \supseteq \{i_1, \dots, i_{t'}\} \text{ and } S \text{ is independent,} \\ &\geq k + t' && \text{by (4),} \\ &= |T_1 \cap S| + k && \text{since } |T_1 \cap S| = t', \\ &> |T_1 \cap S| + |T_3 \cap S| && \text{by the assumption } k > |T_3 \cap S|, \\ &= |S| && \text{since } T_2 \cap S = \emptyset, \end{aligned}$$

contradicting the precondition to the lemma. It follows that $|T_3 \cap S| \geq k$. Since $|T_3 \cap S| \leq |T_3| \leq k$ by Claim 8, it follows that $|T_3 \cap S| = k$ and that all vertices of T_3 belong to S . \square

Let $\ell := r_{\mathcal{P}}(i_{t'})$. Since the path decomposition is nice there is only one vertex (i.e., $i_{t'}$) that occurs in X_ℓ but not after X_ℓ . So $X_\ell = \{i_{t'}\} \cup T_3$, and Claim 10 implies that no vertex of $N(S)$ occurs in X_ℓ since $X_\ell = \{i_{t'}\} \cup T_3 \subseteq S$. Claim 9 shows that all neighbors of $i_{t'+1}, \dots, i_t$ are also neighbor to some vertex of the prefix $i_1, \dots, i_{t'}$. Since $i_1, \dots, i_{t'}$ are ordered by increasing right endpoint of the intervals representing them in the decomposition, all neighbors of $i_{t'+1}, \dots, i_t$ therefore have to occur in a bag with index at most $r_{\mathcal{P}}(i_{t'})$, and since X_ℓ contains no vertex of $N(S)$, by (P3) it follows that no vertex of $N(S)$ occurs in a bag with index ℓ or later. Since $X_\ell = \{i_{t'}\} \cup T_3$ and $|T_3 \cap S| = k$ by the previous claim, there are $k + 1$ vertices in X_ℓ . Since the size difference of consecutive bags in a nice path decomposition is exactly one, and no bag has size more than $k + 1$ since the width is k , it follows that $X_{\ell-1} = X_\ell \setminus \{v\}$ for some vertex $v \in \{i_{t'}\} \cup T_3 \subseteq S$. Since no vertex of $N(S)$ occurs in bag X_ℓ or after, and v does not occur in $X_{\ell-1}$ or earlier, it follows that v does not occur in a bag together with a vertex of $N(S)$. By the definition of path decomposition, this implies that v has no neighbor in $N(S)$; since $v \in S$ and S is an independent set (it is a subset of a partite set of a bipartite graph), this implies that v is an isolated vertex in G . But since $1 \leq t' < t = |S|$, the set $S' := \{v\}$ is a nonempty strict subset of S for which $0 = |N(S')| \leq |S'| = 1$, contradicting the precondition to the lemma. This concludes the proof of Case 1.

Case 2: $T_2 \cap S \neq \emptyset$. We continue the proof of Lemma 12 for the case that $T_2 \cap S \neq \emptyset$. We will show that $T_2 \cap S$ is a nonempty strict subset of S with $|T_2 \cap S| \geq |N(T_2 \cap S)|$. This will contradict our assumption that S is inclusion-wise minimal with the property that $|S| \geq |N(S)|$. Now let us denote $|T_3 \cap I| = k_I$ and $|T_3 \cap J| = k_J$. Note that $T_3 \cap J = T_3 \cap S$, and observe from Claim 8 that

$$k_I + k_J \leq k. \tag{6}$$

Recall from the choice of $r_{\mathcal{P}}(i_{t'})$ that $|T_1 \cap S| = |\{i_1, \dots, i_{t'}\}| = t'$. Since $S = (T_1 \cup T_2 \cup T_3) \cap S$, and the T_i 's are mutually disjoint, we have:

$$\begin{aligned} |S| &= |T_1 \cap S| + |T_2 \cap S| + |T_3 \cap S| \\ &= t' + |T_2 \cap S| + k_J. \end{aligned}$$

Therefore,

$$|T_2 \cap S| = |S| - k_J - t'. \tag{7}$$

Also note that

$$\begin{aligned} |N(S)| &= |N((T_1 \cup T_2 \cup T_3) \cap S)| \\ &\geq |N((T_1 \cup T_2) \cap S)| \\ &= |N(T_1 \cap S)| + |N(T_2 \cap S)| - |N(T_1 \cap S) \cap N(T_2 \cap S)|. \end{aligned} \tag{8}$$

Now observe that any vertex which is a neighbor of some vertex in $T_1 \cap S$ and some vertex in $T_2 \cap S$, must be both in some bag with index at most $r_{\mathcal{P}}(i_{t'})$ (to meet $T_1 \cap S$) and in some bag with index strictly more than $r_{\mathcal{P}}(i_{t'})$ (to meet $T_2 \cap S$). This implies that $N(T_1 \cap S) \cap N(T_2 \cap S) \subseteq T_3 \cap I$. Therefore

$$|N(T_1 \cap S) \cap N(T_2 \cap S)| \leq |T_3 \cap I| = k_t. \tag{9}$$

Hence, (8) and (9) yield

$$|N(T_2 \cap S)| \leq |N(S)| - |N(\{i_1, \dots, i_{t'}\})| + k_t. \tag{10}$$

Therefore, combining (7) and (10) we get

$$\begin{aligned} |T_2 \cap S| - |N(T_2 \cap S)| &\geq |S| - k_j - t' - |N(S)| + |N(\{i_1, \dots, i_{t'}\})| - k_t \\ &= (|S| - |N(S)|) + (|N(\{i_1, \dots, i_{t'}\})| - t') - (k_t + k_j) \\ &\geq 0 \quad \text{from our hypothesis about } S, \text{ and Equation (4) and (6).} \end{aligned}$$

So $T_2 \cap S$ is a nonempty strict subset of S satisfying the key property, contradicting that S is inclusion-wise minimal. This completes the proof of Lemma 12. \square

Using Lemma 12 we bound the TAR reconfiguration threshold in terms of pathwidth.

Theorem 13. *Let $G = (V, E)$ be a graph. Then $\text{TAR}(G) \leq \max(\text{PW}(G), 1)$.*

Proof. We prove this theorem using induction on the number of vertices. As before, it is enough to consider $G = (V, E)$ and assume that the initial and target independent sets I and J respectively are such that $|I| = |J|$, $I \cup J = V$ and $I \cap J = \emptyset$. We will show that $\text{PW}(G) \leq k$ implies that $\text{TAR}(G) \leq k$, using induction on the number of vertices n . For $n = 1$, the statement is trivially true. Now fix any $k \geq 1$, and assume the induction hypothesis that any graph G with n vertices satisfying $\text{PW}(G) \leq k$ has $\text{TAR}(G) \leq k$.

Assume G is a graph of $n + 1$ vertices having pathwidth at most k . Let S be an inclusion-minimal subset of J for which $|S| \geq |N(S)|$. Such a set exists since $|J| = |I| \geq |N(J)|$. We will show that if we reconfigure the set S in a suitable order by moving tokens from $N(S)$ onto S , then the buffer size will not grow beyond k . There are enough vertices in S to accommodate all tokens on $N(S)$, and afterward we will invoke induction.

We first deal with a special case. If $S = \{v\}$ is a singleton set, then it has degree at most one since $|S| \geq |N(S)|$. Move the token from the neighbor u of v (or from an arbitrary vertex u , if v has no neighbors) into the buffer, and then onto v . By induction there exists a TAR reconfiguration from $I \setminus \{u\}$ to $J \setminus \{v\}$ in $G - \{u, v\}$ using a buffer of size at most $\max(\text{PW}(G - \{u, v\}), 1) \leq \max(\text{PW}(G), 1)$. When inserting the token move from u onto v at the beginning of this sequence, we get a TAR reconfiguration from I to J with the desired buffer size. In the remainder of the proof we can therefore assume $|S| \geq 2$. This implies that $|S| = |N(S)|$: if $|S| > |N(S)|$ and $|S| \geq 2$, then we can remove a vertex v from S to obtain $|S \setminus \{v\}| \geq |N(S \setminus \{v\})|$ for the nonempty set $S \setminus \{v\}$, contradicting minimality.

Let $\mathcal{P} = (X_1, X_2, \dots, X_r)$ be a nice path decomposition of width at most k . If G has no edges, then S is a singleton set containing an isolated vertex. Since we already covered that case, we know G has at least one edge, so any path decomposition has width $k \geq 1$. Enumerate the vertices of S as i_1, \dots, i_m such that $r_{\mathcal{P}}(i_1) < \dots < r_{\mathcal{P}}(i_m)$. Hence the vertices are ordered by increasing rightmost endpoint of the interval of bags containing it.

In order to describe the reconfiguration procedure we suitably group several TAR reconfiguration steps together as one step in the algorithm. In particular, one reconfiguration step in the algorithm described below will consist of a run of successive removals of nodes, followed by a single node addition.

We use the notion of a *buffer set* B_t at the t th step of the reconfiguration, such that $|B_t|$ will correspond to the number of tokens in the buffer at any particular time, and $\max_t |B_t| + 1$ will correspond to the maximum buffer size of the corresponding TAR reconfiguration sequence. The buffer set is a subset of vertices, showing where the tokens in the buffer came from. At time step $t = 0$, define $W_0 = I$ to be the independent set of vertices with a token, and let the buffer set B_0 be empty. We will define intermediate independent sets W_i and buffer sets B_i representing the grouped reconfiguration steps. The algorithm stops when W_m contains all vertices in S ; we will then invoke the induction hypothesis to finish the sequence. From the sequence (W_0, W_1, \dots, W_m) one obtains a formal reconfiguration sequence as defined in Section 3 by inserting “transitioning independent sets” in between W_i and W_{i+1} for all i . From W_i , repeatedly remove one vertex until arriving at $W_{i+1} \setminus W_i$, and then add the single vertex of $W_{i+1} \setminus W_i$ to the resulting set.

For $t \geq 1$, the transition from $t - 1$ to t is obtained as follows. Let u_t be an arbitrary vertex from $B_{t-1} \cup (N(i_t) \cap W_{t-1})$. Intuitively, at step t we take the token from u_t (in the buffer set or on a neighbor of i_t) and move it onto vertex i_t , causing u_t to disappear from the buffer and adding i_t to the independent set. To ensure the resulting set is independent, tokens on neighbors of i_t are moved into the buffer beforehand. Observe that the above step is valid only if $B_{t-1} \cup (N(i_t) \cap W_{t-1})$ is nonempty. Below in Claim 11 we show that due to the choice of S , this is indeed the case for all $t \leq m$. Formally, we obtain the following:

Algorithm (Reconfiguration for Graphs with Small Pathwidth). Initialize with $B_0 = \emptyset$ and $W_0 = I$. We recursively define B_t and W_t for $t \geq 1$.

1. The neighbors of i_t that have tokens (i.e. that are in the current independent set) are removed from the previous independent set W_{t-1} , making room to add i_t to the new independent set: $W_t = (W_{t-1} \setminus N(i_t)) \cup \{i_t\}$.
2. The neighbors of i_t belonging to the previous independent set W_t move to the buffer, while u_t is removed from the buffer since its token has moved onto i_t :

$$B_t = (B_{t-1} \cup (N(i_t) \cap W_{t-1})) \setminus \{u_t\}. \tag{11}$$

As mentioned earlier, a step from W_t to W_{t+1} can be thought of as a sequence of successive removals of the nodes $N(i_{t+1}) \cap W_t$, and then addition of the node i_{t+1} . During this successive TAR reconfiguration sequence corresponding to the step W_t to W_{t+1} , the maximum buffer size is given by $|B_{t+1}| + 1$, since the buffer size will be $|B_t \cup (N(i_{t+1}) \cap W_t)|$ just before the buffer token from u_{t+1} is moved onto i_{t+1} . Therefore, the maximum buffer size in the entire TAR reconfiguration sequence starting from W_0 and ending at W_m is given by $\max_{0 \leq t \leq m} |B_t| + 1$. Also, at the end of the algorithm, all vertices from the set S will be in the independent set, and no vertex in the buffer set. This can be seen by observing the following. Initially all tokens were on the vertices belonging to the set $N(S) \subseteq I$, since $S \subseteq J$. At each step of the algorithm essentially one token is selected from $N(S)$ as long as the number of such tokens is positive, and is placed on some vertex in S . Now since $|S| \geq |N(S)|$, all the tokens in $N(S)$ must eventually exhaust before the algorithm terminates placing one token at each vertex of S . For the validity of the above algorithm we claim the following, which in turn also characterizes the size of the buffer set at all intermediate time steps.

Claim 11. For all $1 \leq t \leq m$ we have that $B_{t-1} \cup (N(i_t) \cap W_{t-1})$ is nonempty, and that $|B_t| = |N(\{i_1, \dots, i_t\})| - t$.

Proof. Suppose on the contrary that there exists $t' \leq m$, such that $B_{t'-1} \cup (N(i_{t'}) \cap W_{t'-1})$ is empty for the first time. If $t' = 1$, then $B_{t'-1} \cup (N(i_{t'}) \cap W_{t'-1})$ is empty, and in particular $N(i_{t'}) = \emptyset$, so that $i_{t'} = i_1$ is an isolated vertex. But since $|S| \geq 2$ by our argument above, it follows that $S' = \{i_1\}$ is a nonempty strict subset with $|S'| \geq |N(S')|$; a contradiction. So in the remainder we consider $t' > 1$. We show that, for all $t < t'$, $|B_t| = |N(\{i_1, \dots, i_t\})| - t$. Using this, we prove that $2 \leq t' \leq m$ leads to a contradiction.

Observe that for any $t < t'$, after the t th step of the algorithm, the total number of distinct vertices that have been added to the buffer set is given by $|N(\{i_1, \dots, i_t\})|$. Furthermore, for all $t'' \leq t < t'$, the set $B_{t''-1} \cup (N(i_{t''}) \cap W_{t''-1})$ has always been nonempty. This implies that at each step, precisely one token has been removed from the buffer, thus reducing the size of the buffer set by moving a buffer token onto a vertex that is added to the independent set. Therefore, in total t times the size of the buffer set reduces by one. Since initially the buffer set was empty, for any $t < t'$ we have $|B_t| = |N(\{i_1, \dots, i_t\})| - t$.

Since we have assumed that $B_{t'-1} \cup (N(i_{t'}) \cap W_{t'-1})$ is empty, we know $B_{t'-1}$ is empty, and therefore from the above argument $|B_{t'-1}| = |N(\{i_1, \dots, i_{t'-1}\})| - (t' - 1) = 0$.

Defining $S' := \{i_1, \dots, i_{t'-1}\} \subseteq S$, we have $|N(S')| \leq |S'|$. Since $t' \geq 2$ the set S' is nonempty, contradicting the minimality of S . This proves the first part of the claim. Since the buffer does not become empty until after step t , the given argument then also proves the second part of the claim. \square

Note that in particular $|B_m| = |N(\{i_1, \dots, i_m\})| - m = |N(S)| - |S| = 0$; the buffer empties for the first time only after reconfiguring the whole set.

It remains to show that throughout the process the buffer size will not grow beyond k , i.e. $|B_t| \leq k - 1$, for all $t \leq m$. Claim 11 (ii) implies that $\max_{t \leq m} |B_t| \geq k$ if and only if $\exists t \leq m$ such that $|N(\{i_1, \dots, i_t\})| - t \geq k$, which is not possible due to Lemma 12. This then ensures that throughout the algorithm, the buffer size will never exceed k .

Since the buffer set empties out after reconfiguring the set S , after the execution of the algorithm, $W_m \cap J = S$ and $W_m \cap I \subset V \setminus (S \cup N(S))$. Now define $G' = G - (S \cup N(S))$, $I' = I \cap W_m$, and $J' = J \setminus S$. Observe that G' has pathwidth at most k , and $|I'| = |I \cap W_m| = |I| - |S| = |J'|$. Furthermore, since S is non-empty, $|V(G')| \leq n$. By the induction hypothesis, there exists a TAR reconfiguration sequence from I' to J' in G' using a buffer of size at most k . Since $N(S)$ is not in G' , any independent set in G' remains to be an independent set in G when augmented with the set S . Therefore we can first apply the given reconfiguration from $N(S)$ to S , followed by the reconfiguration from I' to J' , to reconfigure I to J with a buffer of size at most k . \square

Observe by considering a complete balanced bipartite graph on $2n$ vertices $K_{n,n}$, that in general the above bound is tight. Indeed, from [3] we know that $K_{n,n}$ has pathwidth equal to n , and as explained earlier, the TAR reconfiguration threshold is also n .

5.2. Obstructions to TAR reconfigurability

Having proved Theorem 13, it is natural to ask whether pathwidth in some sense characterizes the TAR reconfiguration threshold: does large pathwidth of a graph imply that its TAR reconfiguration threshold is large? This is not the case: the pathwidth of a complete binary tree is proportional to its depth [24], but its reconfiguration threshold is one by Theorem 3.

We now identify a graph structure which forces the TAR reconfiguration threshold to be large. First we formally introduce the special type of minor illustrated in Fig. 1(b).

Definition 4 (*Bipartite Topological Double Minor*). Let $G = (I \cup J, E)$ be a bipartite graph and let H be an arbitrary graph. Then H is a *bipartite topological double minor* of G , if one can assign to every $v \in V(H)$ a subgraph $\varphi(v)$ of G , which is either an edge or an even cycle in G , and one can assign to each edge $e = \{u, v\} \in E(H)$ a pair of odd-length paths $\psi_1(e), \psi_2(e)$ in G , such that the following holds:

- For any $u, v \in V(H)$ with $u \neq v$ the subgraphs $\varphi(u)$ and $\varphi(v)$ are vertex-disjoint.
- For any $v \in V(H)$ no vertex of $\varphi(v)$ occurs as an interior vertex of a path $\psi_1(e)$ or $\psi_2(e)$, for any $e \in E(H)$.
- For any $e, e' \in E(H)$ the paths $\psi_1(e)$ and $\psi_2(e')$ are internally vertex-disjoint.
- For any $e = \{u, v\} \in E(H)$ the paths $\psi_1(e)$ and $\psi_2(e)$ both have one endpoint in $\varphi(v)$ and one endpoint in $\varphi(u)$.
- For any $v \in V(H)$ and edge $\{u, v\} \in E(H)$, the attachment points of $\psi_1(e)$ and $\psi_2(e)$ in $\varphi(v)$ belong to different partite sets.

The triple $(\varphi, \psi_1, \psi_2)$ is a *BTD-minor model* of H in G . For an edge $e \in E(H)$ we define $\psi'_1(e), \psi'_2(e) \subseteq V(G)$ as the *interior* vertices of the paths $\psi_1(e)$ and $\psi_2(e)$, which may be \emptyset if the path consists of a single edge.

Intuitively, H occurs as a bipartite topological double minor (or *BTD-minor*) if each vertex of H can be realized by an edge or even cycle, and every edge of H can be realized by two odd-length paths that connect an I -vertex of $\varphi(v)$ to a J -vertex of $\varphi(u)$ and the other way around, in such a way that these structures are vertex-disjoint except for the attachment of paths to cycles. The definition easily extends to bipartite graphs whose bipartition is not given, since a BTD-minor is contained within a single connected component of the graph, which has a unique bipartition.

Proposition 14. *Let $G = (I \cup J, E)$ be a bipartite graph having a connected graph H as a BTD-minor model $(\varphi, \psi_1, \psi_2)$, such that each vertex of G is in the image of φ, ψ_1 , or ψ_2 . Then G has a perfect matching with $|I| = |J|$ edges, and for any independent set W in G :*

1. For each vertex v of H we have $|W \cap \varphi(v)| \leq |\varphi(v)|/2$.
2. For each edge e of H and $i \in \{1, 2\}$ we have $|W \cap \psi'_i(e)| \leq |\psi'_i(e)|/2$.

For a maximum independent set W , equality holds in all cases.

Proof. To see that G has a perfect matching, observe that each $\varphi(v)$ for $v \in V(H)$ is either an edge or an even cycle, which can be covered completely by a matching consisting of edges from $\varphi(v)$. For each $e \in E(H)$ and $i \in \{1, 2\}$ there is an even number of interior vertices on the path $\psi_i(e)$, since the path has odd length. The interior vertices $\psi'_i(e)$ can therefore also be covered completely by a matching of edges among $\psi'_i(e)$. Since each vertex of G is in the image of φ or $\psi_{1,2}$, the sets $\varphi(v)$ together with sets $\psi'_i(e)$ for $e \in E(H)$ and $i \in \{1, 2\}$ cover $V(G)$. Since these sets are vertex-disjoint by Definition 4, they form a partition of G . By the preceding argument, this implies G has a perfect matching M where no edge crosses the described partition of $V(G)$. This matching has size $|I| = |J|$ since G is bipartite.

Now we prove the two claimed properties. If W is an independent set, it contains at most one endpoint from each edge in M and therefore contains at most half the vertices of each $\varphi(v \in V(H))$ and each $\psi'_i(e \in E(H))$. Any independent set W achieving equality for all these sets has size $|I| = |J|$ and is therefore maximum. \square

For a bipartite graph G , let $\text{TREEMINOR}(G)$ denote the largest integer k for which G contains a complete binary tree of depth k as a BTD-minor. For a class of bipartite graphs Π we define $\text{TREEMINOR}(\Pi) := \sup_{G \in \Pi} \text{TREEMINOR}(G)$.

Theorem 15. *There exists a real constant $c > 0$ such that any hereditary graph class Π satisfies $\text{TAR}(\Pi) \geq c \cdot \text{TREEMINOR}(\Pi_{\text{bip}})$.*

Proof. As before, we consider a balanced bipartite graph $G \in \Pi_{\text{bip}}$ with bipartition $V(G) = I \cup J$ that has a complete binary tree T of depth d as a BTD-minor. Since the graph class is hereditary, for the lower bound, we consider only the subgraph of G induced by $\bigcup_{v \in V(T)} \{\varphi(v)\} \cup (\bigcup_{e \in E(T)} \{\psi_1(e) \cup \psi_2(e)\})$, and without loss of generality, we shall refer to it as G itself.

Fact 2 ([2]). *There is a universal constant $c_1 > 0$ such that if T is a complete binary tree of depth d , then $\max_{1 \leq i \leq |V(T)|} \min_{S \subseteq V(T); |S|=i} |N_T(S)| \geq c_1 \cdot d$.*

The above implies that there exists $i_0 \leq |V(T)|$, such that any size- i_0 subset of $V(T)$ has a neighborhood of size at least $c_1 \cdot d$. Let $I \cup J$ be the unique bipartition of the connected graph G , and consider an arbitrary TAR reconfiguration sequence from I and J . In this sequence ($I = W_0, W_1, \dots, W_t = J$) of independent sets in G , look at the reconfiguration step when for the first time there exists $S \subseteq V(T)$ with $|S| = i_0$, such that the intermediate independent set W at that step contains $\bigcup_{v \in S} (\varphi(v) \cap J)$, and for all $v \notin S$ it satisfies $(\varphi(v) \cap W \cap J) \subsetneq (\varphi(v) \cap J)$. We will prove that $|J| - |W| \geq c_1 \cdot d$, implying that from the initial independent set of $|I| = |J|$ tokens, at least $c_1 \cdot d$ tokens must reside in the buffer.

To prove the theorem, consider the intermediate independent set W , and the set $S \subseteq V(T)$ with $|S| = i_0$ satisfying the above criteria. The following claim shows that for each vertex in $N_T(S)$, the independent set W uses at least one vertex fewer than the maximum independent set J does.

Claim 12. Consider an edge $e = \{u, v\} \in E(T)$ with $u \in S$ and $v \notin S$, and let $Q_{e,v} \subseteq V(G)$ denote the vertices in $\varphi(v) \cup \psi'_1(e) \cup \psi'_2(e)$. The following holds:

$$|W \cap Q_{e,v}| < |J \cap Q_{e,v}| = \frac{|Q_{e,v}|}{2}. \quad (12)$$

Proof. By Proposition 14, the maximum independent set J contains exactly half the vertices of $Q_{e,v}$. If $|W \cap \psi'_i(e)| < |\psi'_i(e)|/2$ for some $i \in \{1, 2\}$, then we are done: by Proposition 14 the set W contains fewer vertices from $\psi'_i(e)$ than the maximum independent set J does, and this cannot be compensated within the other parts of the structure since J contains half the vertices there and no independent set contains more. In the remainder, we can assume that W contains exactly half the vertices from $\psi'_1(e)$ and $\psi'_2(e)$. Then the following are true:

- (i) All J -nodes of $\varphi(u)$ are in W (by our choice of W and since $u \in S$).
- (ii) Some J -node of $\varphi(v)$ is not in W (by our choice of W and since $v \notin S$).
- (iii) Some I -node of $\varphi(v)$ is not in W . To see this, let $i \in \{1, 2\}$ such that $\psi_i(e)$ is an odd-length path from a J -node in $\varphi(u)$ to an I -node in $\varphi(v)$, which exists by Definition 4, and orient it in that direction. Since the first vertex on the path is a J -node in $\varphi(u)$, it is contained in W as shown above. Hence the second vertex on the path, the first interior vertex, is not in W . Since exactly half the interior vertices from $\psi_i(e)$ belong to W , every other interior vertex from $\psi_i(e)$ is in W . Since the path has an even number of interior vertices and the first interior vertex is not in W , the last interior vertex must be in W . But this prevents its I -node neighbor in $\varphi(v)$ from being in W .

Therefore, since $\varphi(v)$ is either an edge or an even cycle, we have $|W \cap \varphi(v)| < |\varphi(v)|/2$ by observing the following: the only independent sets in $\varphi(v)$ of size $|\varphi(v)|/2$ are $\varphi(v) \cap I$ and $\varphi(v) \cap J$, but $\varphi(v) \cap W$ is not equal to either of these sets since it avoids a J -node and an I -node. Hence $|W \cap \varphi(v)| < |\varphi(v)|/2 = |J \cap \varphi(v)|$, and Proposition 14 shows that this cannot be compensated in other parts of the minor model, implying $|W \cap Q_{e,v}| < |J \cap Q_{e,v}|$. \square

Using Claim 12 we finish the proof of Theorem 15. For each $v \in N_T(S)$, pick an edge $e = \{u, v\}$ such that $u \in S$. By Claim 12 the set W contains less than half the vertices of $Q_{e,v}$, while the maximum independent set J contains exactly half. Since the sets $Q_{e,v}$ considered for different vertices $v \in N_T(S)$ are disjoint, while Proposition 14 shows that from the other pieces of the minor model W cannot use more vertices than J does, it follows that $|W| \leq |J| - |N_T(S)| \leq |J| - c_1 \cdot d$. Hence the buffer contains at least $c_1 \cdot d$ tokens. \square

6. Conclusion

In this paper we considered two types of reconfiguration rules for independent set, involving simultaneously jumping tokens and reconfiguration with a buffer. For both models, we derived tight bounds on the corresponding reconfiguration thresholds in terms of several graph parameters like the minimum vertex cover size, the minimum feedback vertex set size, and the pathwidth. Many results in the literature concerning the parameter pathwidth can be extended to hold for the parameter treewidth as well. This is not the case here; the upper bound on the TAR reconfiguration threshold in terms of pathwidth (Theorem 13) cannot be strengthened to treewidth, since one can make arbitrarily deep complete binary trees as BTD -minors in bipartite graphs of treewidth only two (see Fig. 1(b)). On the other hand, there are bipartite graphs of large treewidth with TAR reconfiguration threshold two (Fig. 2(b)). To characterize the TAR reconfiguration threshold one therefore needs to combine graph connectivity (as measured by the width parameters) with notions that constrain the parity of the connections in the graph. This is precisely why we introduced BTD -minors. We conjecture that the TAR reconfiguration threshold can also be upper-bounded in terms of the depth of the largest complete binary tree BTD -minor, i.e., the converse of Theorem 15 holds, in the sense that any hereditary graph class having a large TAR reconfiguration threshold must contain a graph having a complete binary tree of large depth as a BTD -minor. Our belief is based partially on the fact that a BTD -minor model of a deep complete binary tree is arguably the simplest graph of large pathwidth and feedback vertex number. Resolving this conjecture is our main open problem.

Acknowledgments

This research was financially supported by The Netherlands Organization for Scientific Research (NWO) through TOP-GO grant 613.001.012, Gravitation Networks grant 024.002.003 and a Veni grant ‘Frontiers in Parameterized Preprocessing’ (639.021.437). DM thanks Sem Borst for his comments on the motivation for the reconfiguration thresholds.

References

- [1] A. Adler, M. de Berg, D. Halperin, K. Solovey, Efficient multi-robot motion planning for unlabeled discs in simple polygons, in: Algorithmic Foundations of Robotics XI: Selected Contributions of the Eleventh International Workshop on the Algorithmic Foundations of Robotics, 2015, pp. 1–17. http://dx.doi.org/10.1007/978-3-319-16595-0_1.
- [2] B.V.S. Bharadwaj, L.S. Chandran, Bounds on isoperimetric values of trees, Discrete Math. 309 (4) (2009) 834–842. <http://dx.doi.org/10.1016/j.disc.2008.01.021>.

- [3] H.L. Bodlaender, A partial k -arboretum of graphs with bounded treewidth, *Theoret. Comput. Sci.* 209 (1–2) (1998) 1–45. [http://dx.doi.org/10.1016/S0304-3975\(97\)00228-4](http://dx.doi.org/10.1016/S0304-3975(97)00228-4).
- [4] H.L. Bodlaender, B.M.P. Jansen, S. Kratsch, Kernel bounds for path and cycle problems, *Theoret. Comput. Sci.* 511 (2013) 117–136. <http://dx.doi.org/10.1016/j.tcs.2012.09.006>.
- [5] P. Bonsma, Independent set reconfiguration in cographs, in: *Graph-Theoretic Concepts in Computer Science: 40th International Workshop, WG 2014, Nouan-le-Fuzelier, France, June 25–27, 2014, Revised Selected Papers, 2014*, pp. 105–116. http://dx.doi.org/10.1007/978-3-319-12340-0_9.
- [6] P. Bonsma, M. Kamiński, M. Wrochna, Reconfiguring independent sets in claw-free graphs, in: *Algorithm Theory – SWAT 2014: 14th Scandinavian Symposium and Workshops, Copenhagen, Denmark, July 2–4, 2014, Proceedings, 2014*, pp. 86–97. http://dx.doi.org/10.1007/978-3-319-08404-6_8.
- [7] G. Călinescu, A. Dumitrescu, J. Pach, Reconfigurations in graphs and grids, *SIAM J. Discrete Math.* 22 (1) (2008) 124–138. <http://dx.doi.org/10.1137/060652063>.
- [8] M. Cygan, F.V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, S. Saurabh, *Parameterized Algorithms*, Springer Publishing Company, Incorporated, 2015. <http://dx.doi.org/10.1007/978-3-319-21275-3>.
- [9] E.D. Demaine, M.L. Demaine, E. Fox-Epstein, D.A. Hoang, T. Ito, H. Ono, Y. Otachi, R. Uehara, T. Yamada, Linear-time algorithm for sliding tokens on trees, *Theoret. Comput. Sci.* 600 (2015) 132–142. <http://dx.doi.org/10.1016/j.tcs.2015.07.037>.
- [10] F. Den Hollander, F.R. Nardi, S. Taati, Metastability of hard-core dynamics on bipartite graphs. <https://arxiv.org/abs/1710.10232v2>.
- [11] R. Diestel, *Graph Theory*, Springer-Verlag Berlin Heidelberg, 2010.
- [12] G.A. Dirac, Some theorems on abstract graphs, *Proc. Lond. Math. Soc.* s3-2 (1) (1952) 69–81. <http://dx.doi.org/10.1112/plms/s3-2.1.69>.
- [13] A. Dumitrescu, J. Pach, Pushing squares around, *Graphs Combin.* 22 (1) (2006) 37–50. <http://dx.doi.org/10.1007/s00373-005-0640-1>.
- [14] E. Fox-Epstein, D.A. Hoang, Y. Otachi, R. Uehara, Sliding token on bipartite permutation graphs, in: *Algorithms and Computation: 26th International Symposium, ISAAC 2015, Nagoya, Japan, December 9–11, 2015, Proceedings, 2015*, pp. 237–247. http://dx.doi.org/10.1007/978-3-662-48971-0_21.
- [15] G. Goralý, R. Hassin, Multi-color pebble motion on graphs, *Algorithmica* 58 (3) (2010) 610–636. <http://dx.doi.org/10.1007/s00453-009-9290-7>.
- [16] R.A. Hearn, E.D. Demaine, PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation, *Theoret. Comput. Sci.* 343 (1–2) (2005) 72–96. <http://dx.doi.org/10.1016/j.tcs.2005.05.008>.
- [17] J. van den Heuvel, The complexity of change, in: S. Blackburn, S. Gerke, M. Wildon (Eds.), *Surveys in Combinatorics 2013*, Cambridge University Press, 2013, pp. 127–160. <http://dx.doi.org/10.1017/CBO9781139506748.005>. URL <http://ebooks.cambridge.org/ref/id/CBO9781139506748A011>.
- [18] T. Ito, E.D. Demaine, N.J. Harvey, C.H. Papadimitriou, M. Sideri, R. Uehara, Y. Uno, On the complexity of reconfiguration problems, *Theoret. Comput. Sci.* 412 (12–14) (2011) 1054–1065. <http://dx.doi.org/10.1016/j.tcs.2010.12.005>.
- [19] L. Jiang, M. Leconte, J. Ni, R. Srikant, J. Walrand, Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling, *IEEE Trans. Inform. Theory* 58 (10) (2012) 6541–6555. <http://dx.doi.org/10.1109/TIT.2012.2204032>.
- [20] M. Kamiński, P. Medvedev, M. Milanić, Complexity of independent set reconfigurability problems, *Theoret. Comput. Sci.* 439 (2012) 9–15. <http://dx.doi.org/10.1016/j.tcs.2012.03.004>.
- [21] R.M. Karp, Reducibility among combinatorial problems, in: *Proceedings of a Symposium on the Complexity of Computer Computations, Held March 20–22, 1972, At the IBM Thomas J. Watson Research Center, Yorktown Heights, New York*, in: *The IBM Research Symposia Series*, Plenum Press, New York, 1972, pp. 85–103.
- [22] P. Kataria, R. Roy, Parallel Glauber dynamics-based scheduling in static wireless grid networks with polynomially fading interference, *Electron. Lett.* 51 (23) (2015) 1948–1950. <http://dx.doi.org/10.1049/el.2014.3385>.
- [23] F.P. Kelly, Loss networks, *Ann. Appl. Probab.* 1 (3) (1991) 319–378. <http://dx.doi.org/10.1214/aoap/1177005872>.
- [24] N.G. Kinnarsley, M.A. Langston, Obstruction set isolation for the gate matrix layout problem, *Discrete Appl. Math.* 54 (2–3) (1994) 169–213. [http://dx.doi.org/10.1016/0166-218X\(94\)90021-3](http://dx.doi.org/10.1016/0166-218X(94)90021-3).
- [25] D. Lokshtanov, A.E. Mouawad, The complexity of independent set reconfiguration on bipartite graphs, in: *Proc. SODA '18, 2018*, pp. 185–195. URL <https://dl.acm.org/citation.cfm?id=3175279>.
- [26] D. Lokshtanov, A.E. Mouawad, F. Panolan, M.S. Ramanujan, S. Saurabh, Reconfiguration on sparse graphs, in: *Algorithms and Data Structures – 14th International Symposium, WADS 2015, Victoria, BC, Canada, August 5–7, 2015 Proceedings, 2015*, pp. 506–517. http://dx.doi.org/10.1007/978-3-319-21840-3_42.
- [27] A.E. Mouawad, N. Nishimura, V. Raman, N. Simjour, A. Suzuki, On the parameterized complexity of reconfiguration problems, in: *Parameterized and Exact Computation: 8th International Symposium, IPEC 2013, Sophia Antipolis, France, September 4–6, 2013, Revised Selected Papers, 2013* pp. 281–294. http://dx.doi.org/10.1007/978-3-319-03898-8_24.
- [28] F.R. Nardi, A. Zocca, S.C. Borst, Hitting time asymptotics for hard-core interactions on grids, *J. Stat. Phys.* 162 (2) (2016) 522–576. <http://dx.doi.org/10.1007/s10955-015-1391-x>.
- [29] N. Nishimura, Introduction to reconfiguration, *Algorithms* 11 (4) (2018) 52. <http://dx.doi.org/10.3390/a11040052>. URL <http://www.mdpi.com/1999-4893/11/4/52>.
- [30] A. Schrijver, *Combinatorial Optimization, Polyhedra and Efficiency*. Springer, Berlin, 2003.
- [31] X. Wang, K. Kar, Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks, in: *Proc. IEEE INFOCOM '05, vol. 1, 2005*, pp. 23–34. <http://dx.doi.org/10.1109/INFCOM.2005.1497875>. URL <http://ieeexplore.ieee.org/document/1497875/>.
- [32] M. Wrochna, Reconfiguration in bounded bandwidth and tree-depth, *J. Comput. Syst. Sci.* 93 (2018) 1–10. <http://dx.doi.org/10.1016/j.jcss.2017.11.003>.
- [33] A. Zocca, S. Borst, J. van Leeuwen, F. Nardi, Delay performance in random-access grid networks, *Perform. Eval.* 70 (10) (2013) 900–915. <http://dx.doi.org/10.1016/j.peva.2013.08.019>. URL <http://linkinghub.elsevier.com/retrieve/pii/S0166531613000977>.