

Shortest paths between shortest paths and independent sets

Marcin Kamiński^{1*}, Paul Medvedev², and Martin Milanič^{3**}

¹Département d'Informatique, Université Libre de Bruxelles, Brussels, Belgium
Marcin.Kaminski@ulb.ac.be

²Department of Computer Science, University of Toronto, Toronto, Canada
pashadag@cs.toronto.edu

³FAMNIT and PINT, University of Primorska, Koper, Slovenia
martin.milanic@upr.si

Abstract. We study problems of reconfiguration of shortest paths in graphs. We prove that the shortest reconfiguration sequence can be exponential in the size of the graph and that it is NP-hard to compute the shortest reconfiguration sequence even when we know that the sequence has polynomial length. Moreover, we also study reconfiguration of independent sets in three different models and analyze relationships between these models, observing that shortest path reconfiguration is a special case of independent set reconfiguration in perfect graphs, under any of the three models. Finally, we give polynomial results for restricted classes of graphs (even-hole-free and P_4 -free graphs).

1 Introduction

One of the biggest impacts of algorithmic graph theory has been its usefulness in modeling real-world problems, where the domain of the problem is modeled as a graph and the constraints on the solution define feasible solutions. For example, consider the problem of routing a certain commodity between two nodes in a transportation network, using as few hops as possible. The transportation network can be modeled as a graph, each route can be modeled as a path, and the feasible solutions are all the shortest paths between the two nodes. Traditionally, the real-world user first defines a problem instance and then uses an algorithm to find a feasible solution which she then “implements” in the real world. However, some real-world situations do not follow this simple paradigm and are more dynamic because they allow the solution to “evolve” over time. For example, consider the situation where the commodity is already being transferred along a shortest route, but the operator has been instructed to use a different route, which is also a shortest path. She can physically switch the route only one node at a time, but does not wish to interrupt the transfer.

* Chargé de Recherches du FRS-FNRS.

** Supported in part by “Agencija za raziskovalno dejavnost Republike Slovenije”, research program P1-0285.

Thus, she would like to switch between the two routes in as few steps as possible, while maintaining a shortest path route at every intermediate step.

In general, this type of situation gives rise to a *reconfiguration* framework, where we consider an algorithmic problem \mathcal{P} and a way of transforming one feasible solution of an instance I of \mathcal{P} to another (*reconfiguration rule*). Given two feasible solutions s_1, s_k of I , we want to find a *reconfiguration sequence* s_1, \dots, s_k such that each s_i ($1 \leq i \leq k$) is a feasible solution of I , and the transition between s_i and s_{i+1} is allowed by the reconfiguration rule. An alternate definition is via the *reconfiguration graph*, where the vertices are the feasible solutions of I , and two solutions are adjacent if and only if one can be obtained from the other by the reconfiguration rule. The reconfiguration sequence is then a path between s_1 and s_k in the reconfiguration graph. We can then ask for the shortest reconfiguration sequence, or, in the *reconfigurability problem*, to simply check if the two solutions are *reconfigurable* (i.e., if such a sequence exists).

The reconfiguration framework has recently been applied in a number of settings, including vertex coloring [3, 4, 2, 1], list-edge coloring [12], clique, set cover, integer programming, matching, spanning tree, matroid bases [11], block puzzles [10], independent set [10, 11], and satisfiability [9]. In the well-studied vertex coloring problem, for example, we are given two k -colorings of a graph, and the reconfiguration rule allows to change the color of a single vertex. In a different example, we are given two independent sets, which we imagine to be two sets of tokens placed on the vertices, and the reconfiguration rule is to slide a single token along an edge (*token sliding*).

Though the complexities of each of the many reconfiguration problems may each be studied independently, a fundamental question is whether there exists any systematic relationship between the complexity of the original problem and that of its reconfigurability problem. To this end, current studies have revealed a pattern where most “natural” problems in P have their reconfigurability problems in P as well, while problems whose reconfigurability versions are at least NP-hard are NP-complete. For example, spanning tree, matching, and matroid problems in general (all in P) lead to polynomially solvable reconfigurability problems, while the reconfigurability of independent set, set cover, and integer programming (all NP-complete) are PSPACE-complete [11]. Another example is satisfiability, where Gopalan et al. [9] showed that reconfigurability instances arising from tight relations—a class for which it is easy to determine if the formula is satisfiable—can be solved in linear time; on the other hand, reconfigurability is PSPACE-complete for the class of formulas arising from non-tight relations.

Ito et al. [11] have conjectured that this relationship is not true in general, and that there exist problems in P which give rise, in a natural way, to NP-hard reconfigurability problems. Indeed, the problem of deciding whether two k -colorings are reconfigurable is PSPACE-complete for (i) bipartite graphs and $k \geq 4$, and (ii) planar graphs, for $4 \leq k \leq 6$ [1]. Clearly, 4-coloring of bipartite or planar graphs is in P. However, these are not “natural” problems in the sense that the colorings are not optimal. It is interesting to ask if there exists a “natural” problem in P whose reconfiguration version is NP-hard.

Another systematic relationship that has been pursued is between the complexity of a reconfigurability problem and the diameter of the reconfiguration graph. When the diameter is polynomial, a reconfiguration sequence is a trivial certificate for the reconfigurability of two instances, guaranteeing that the problem is in NP. However, current evidence further suggests that for reconfigurability problems that are solvable in polynomial time, the diameter is also polynomial. In the study of k -coloring, it was found that for $k \leq 3$, the reconfigurability problem is solvable in polynomial time and the diameter of the reconfiguration graph is at most quadratic in the number of vertices of the colored graph. For satisfiability, the formulas built from tight relations (whose reconfigurability is polynomial) lead to reconfiguration graphs with linear diameter [9]. We are not aware of any natural problems with the property that the diameter can be exponential while reconfigurability can be decided in polynomial time¹; however, such an example, if found, would indicate that the diameter cannot serve as a reliable indicator of the reconfigurability complexity.

In this paper, we introduce the reconfiguration version of the shortest path problem (Section 2), which arises naturally, such as in the routing example above. We show that the reconfiguration graph can have exponential diameter, implying that the shortest path reconfiguration problem probably breaks one of the two established patterns described above. On the one hand, if reconfigurability of shortest paths can be decided in polynomial time, then it is the first example of a reconfigurability problem in P with exponential diameter. On the other hand, if it is NP-hard, it is the first example of a “natural” problem in P whose reconfigurability version is NP-hard. For these reasons, we believe that shortest path reconfiguration is an important problem to study, not only for its practical application but also for our understanding of the systematic relationship between the hardness of a problem, the diameter of its reconfiguration graph, and the hardness of its reconfigurability problem. Towards this end, we give a non-trivial reduction from SAT to show that it is NP-hard to find the shortest reconfiguration sequence between two shortest paths

¹ For a very artificial one, consider the problem in which instances are n -bit words and two instances are adjacent if they differ by 1 modulo 2^n . The diameter of the reconfiguration graph is 2^{n-1} but all pairs of instances are reconfigurable.

(however, the complexity status of the reconfigurability problem remains open).

We also study reconfiguration of independent sets, where, unlike many other problems, there is more than one natural reconfiguration rule. In particular applications, for example, a threshold is specified that bounds the cardinality of the intermediate feasible solutions. Based on this idea, Ito et al. [11] considered an alternative to token sliding called *token addition and removal*, where one is allowed to either add or remove a token as long as there are at least $k - 1$ tokens at any given time, for some k . In this paper, we introduce *token jumping*, where one is allowed to move a single token to any other vertex. The token jumping reconfiguration graph is often easier to analyze than the token addition and removal one, since the cardinalities of two adjacent token sets are always the same. However, we show that the two models are polynomially equivalent, allowing for an easier way to analyze token addition and removal reconfiguration graphs (Section 3).

Finally, we show that reconfiguration of independent sets is a generalization of the reconfiguration of shortest paths; our hardness result for shortest paths then implies that it is NP-hard to find the shortest reconfiguration sequence between two independent sets, even in perfect graphs (Section 4). We also identify two restricted graph classes where reconfigurability is easy – even-hole-free graphs under token jumping, for which the reconfiguration graph is always connected, and P_4 -free graphs under token sliding (Section 5). Due to the space constraint, some proofs are omitted but can be found in the full version [13].

2 Shortest path reconfiguration

We define the reconfiguration rule for shortest paths in the natural way: two shortest (s, t) -paths are adjacent in the reconfiguration graph of shortest (s, t) -paths if and only if they differ, as sequences, in exactly one vertex.

2.1 Instances with exponential diameter

We now present a family of graphs G^k whose size is linear in k but the diameter of the reconfiguration graph is $\Omega(2^k)$. The graph G^1 contains vertices $\{x_i^1 \mid 1 \leq i \leq 7\} \cup \{y_i^1 \mid 1 \leq i \leq 6\} \cup \{s, t\}$ and edges $\{(x_i^1, y_i^1), (x_{i+1}^1, y_i^1), (y_i^1, t) \mid i \leq 6\} \cup \{(s, x_i^1) \mid 1 \leq i \leq 7\}$. The graph G^k is defined recursively with vertices $\{x_i^k \mid 1 \leq i \leq 7\} \cup \{y_i^k \mid 1 \leq i \leq 6\} \cup V(G^{k-1})$ and the edges $\{(x_i^k, y_i^k), (x_{i+1}^k, y_i^k) \mid i \leq 6\} \cup \{(y_i^k, x_j^{k-1}) \mid i \in \{1, 3, 5\}, j \leq 7\} \cup \{(y_2^k, x_1^{k-1}), (y_4^k, x_7^{k-1}), (y_6^k, x_1^{k-1})\} \cup E(G^{k-1} \setminus \{s\}) \cup \{(s, x_i^k) \mid 1 \leq i \leq 7\}$ (see Figure 1). Let $p_b^k = s, x_1^k, y_1^k, \dots, x_1^1, y_1^1, t$, and let $p_e^k = s, x_7^k, y_6^k, x_1^{k-1}, x_1^{k-1}, \dots, x_1^1, y_1^1, t$. We will consider the problem of reconfiguring p_b^k to p_e^k in G^k .




Fig. 1. The graph G^k for $k = 4$, where the reconfiguration distance between $p_b^k = s, x_1^k, y_1^k, \dots, x_1^1, y_1^1, t$ and $p_e^k = s, x_7^k, y_6^k, x_1^{k-1}, x_1^{k-1}, \dots, x_1^1, y_1^1, t$ is $\Theta(2^k)$. An edge with a circle end means that the vertex is connected to all the vertices in the next layer.

Lemma 1. *Let p be a shortest path in G^k that goes through y_1^k , and let q be a path that goes through y_6^k . Then the reconfiguration distance between p and q is at least $9(2^k - 1)$.*

Proof. We prove by induction on k , where the base case is clear. Let $\rho = p_1, \dots, p_n$ be the shortest reconfiguration sequence between p and q . First, let i' be the smallest integer such that $p_{i'+1}$ contains y_4^k , and let $i \leq i'$ be the smallest integer such that every path $p_i, \dots, p_{i'}$ contains y_3^k . By construction, we know that p_{i-1} , and hence p_i , contains y_1^{k-1} and $p_{i'+1}$, and hence $p_{i'}$, contains y_6^{k-1} . Hence, by the induction hypothesis, the length of this first phase, $i' - i + 1$, is at least $9(2^{k-1} - 1)$.

Next, let j' be the smallest integer such that $p_{j'+1}$ contains y_6^k , and let $j \leq j'$ be the smallest integer such that every path $p_j, \dots, p_{j'}$ contains y_5^k . By construction, we know that p_{j-1} , and hence p_j , contains y_6^{k-1} and $p_{j'+1}$, and hence $p_{j'}$, contains y_1^{k-1} . Hence, by the induction hypothesis, the length of this second phase, $j' - j + 1$, is at least $9(2^{k-1} - 1)$.

Observe from the graph construction that ρ must always visit y_{x-1}^k before visiting y_x^k , hence $i' < j$, and so the length of ρ is at least the sum of the two phases plus the moves of the first and second vertex necessary to percolate y_1^k down to y_6^k , proving the lemma. \square

On the other hand, there exists a reconfiguration sequence of length $11(2^k - 1)$ (see the full version [13] for a proof), giving the following theorem:

Theorem 1. *The reconfiguration distance in G^k between p_b^k and p_e^k is $\Theta(2^k)$.*

2.2 NP-Hardness of MIN-SPR

Given (G, s, t, p_b, p_e, k) , where p_b and p_e are shortest (s, t) -paths and k is an integer, the MIN-SPR problem is to determine whether there is a

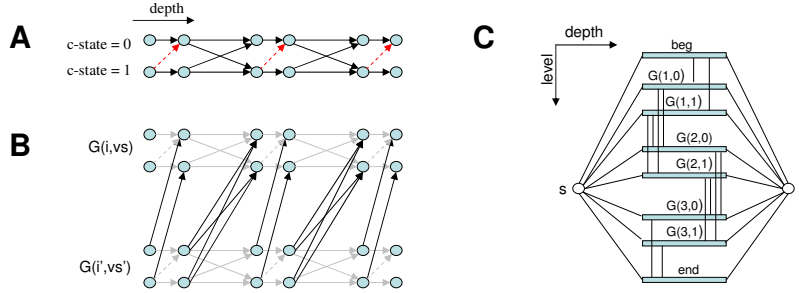


Fig. 2. The reduction from a formula ϕ to a graph G_ϕ for the case of three clauses and three variables. Panel A shows the internal connections of a gadget, with the potential formula edges that depend on ϕ given in red (dashed). Panel B shows the way we connect two given gadgets, while C shows the structure of the whole graph. Each of the rectangles represents a gadget, with the lines showing which parts are connected together.

reconfiguration sequence between p_b and p_e of length at most k . Let ϕ be a formula with variables x_1, \dots, x_n and clauses C_1, \dots, C_m . We will create an instance $(G_\phi, s, t, p_b, p_e, 2m(n+2))$ and show that ϕ is satisfiable if and only this instance is in MIN-SPR. For ease of presentation, the graph G_ϕ will be directed. However, our result holds for undirected graphs because the directed shortest (s, t) -paths in G_ϕ are exactly the shortest paths in the undirected version of G_ϕ .

For every variable x_i and its possible value $vs \in \{0, 1\}$, we build a gadget $G(i, vs)$. The vertex set is $\{v(i, vs, cs, j) \mid cs \in \{0, 1\}, 1 \leq j \leq 2m\}$. The values i, vs, cs , and j for a vertex are referred to as its *level*, *v-state*, *c-state*, and *depth*, and denoted by $l(v), vs(v), cs(v)$, and $d(v)$, respectively. For every $1 \leq j \leq 2m - 1$, and every cs , there is an edge from $v(i, vs, cs, j)$ to $v(i, vs, cs, j + 1)$. For all $1 \leq j \leq m - 1$, there is an edge from $v(i, vs, 0, 2j)$ to $v(i, vs, 1, 2j + 1)$, and from $v(i, vs, 1, 2j)$ to $v(i, vs, 0, 2j + 1)$. We also add edges, called formula edges, that are formula dependent. For all j , if $x_i = vs$ satisfies C_j , we add an edge from $v(i, vs, 1, 2j - 1)$ to $v(i, vs, 0, 2j)$. This gadget is shown in Figure 2A.

We now connect some of these gadgets together. The gadgets we connect are $G(i, vs)$ to $G(i + 1, 0)$ and to $G(i + 1, 1)$, for all $i \leq n - 1$ and all vs . Given two gadgets, $G(i, vs)$ and $G(i', vs')$, the meaning of connecting $G(i, vs)$ to $G(i', vs')$ is given as follows (shown in Figure 2BC). For all $j \leq 2m - 1$ and cs , there is an edge from $v(i', vs', cs, j - 1)$ to $v(i, vs, cs, j)$. Also, for all $j \leq m - 1$ and cs , there is an edge from $v(i', vs', cs, 2j)$ to $v(i, vs, 1 - cs, 2j + 1)$.

We next add a begin and end gadget to the graph, consisting of vertices beg_j and end_j , respectively, for $1 \leq j \leq 2m$. These are connected in a path, with edges (beg_j, beg_{j+1}) and (end_j, end_{j+1}) for $j \leq 2m - 1$. The level of the vertices in the begin (end) gadget is 0 ($i + 1$), the c-state is 0 (1), and

the depth of beg_j or end_j is j . For all vs , $j \leq 2m - 1$, there is an edge from $v(1, vs, 0, j)$ to beg_{j+1} , and from end_j to $v(n, vs, 1, j + 1)$.

Finally, we add a s and t vertex to the graph, and make an edge from s to every depth 1 vertex, and from every depth $2m$ vertex to t . The depth of $s(t)$ is defined to be 0 ($2m + 1$). We call the resulting directed graph G_ϕ . Let $p_b = s, beg_1, \dots, beg_{2m}, t$ and $p_e = s, end_1, \dots, end_{2m}, t$ be two paths in this graph. Then, $(G_\phi, s, t, p_b, p_e, 2m(n + 2))$ is the instance of the MIN-SPR problem that we will consider here.

The intuition behind the reduction is that in order for the path to percolate down from p_b to p_e in a minimal number of steps, it must pass consecutively through exactly one of $G(i, 0)$ or $G(i, 1)$ for every variable x_i . The choice of which one corresponds to assigning x_i the corresponding value. Furthermore, each shortest path that goes through a gadget can visit the vertex at depth $2j$ with a c -state of 0 or 1. This corresponds to having the j^{th} clause satisfied or not. Initially, the path goes only through vertices with c -state 0, and the only way to switch the c -state at a given depth is via a formula edge. By going through a gadget $G(i, vs)$, there is an opportunity to use the formula edges to switch the c -state of all clauses that $x_i = vs$ would satisfy. In order to reach the final path p_e , the c -state of all the vertices must be 1, hence all the clauses must be satisfied.

First, we will show that the reduction is sound. Each edge (a, b) is considered to be either *odd* or *even*, depending on the parity of $d(a)$. We call edges that connect vertices on the same level (exactly those that belong to the same gadget) as *intra-level*, while the edges that connect vertices on different levels are called *inter-level*. We say that a reconfiguration sequence *visits* a vertex if there exists $p \in \rho$ that contains that vertex.

Fact 1 *Let $e = (a, b)$ be an edge in G_ϕ . The following follows directly from construction:*

1. $l(b) \leq l(a) \leq l(b) + 1$.
2. If e is an *intra-level odd edge*, $cs(a) = 0$ implies that $cs(b) = 0$.
3. If e is an *inter-level odd edge*, then $cs(a) = cs(b)$.
4. If e is a *non-formula odd edge*, then $cs(a) = cs(b)$.
5. If e is *intra-level*, then $vs(a) = vs(b)$.

These facts about G_ϕ capture most of the properties of the reduction that are needed to prove correctness of the following Lemmas (2, 3, 4, 5, and 6), whose proofs can be found in the full version [13].

Lemma 2. *The length of a reconfiguration sequence is at least $2m(n + 2)$. Moreover, each move in an sequence that has this length must either increase the c -state or the level of the switched vertex, but not both.*

Lemma 3. *No path can contain two vertices with the same level but different v -state.*

Lemma 4. *Suppose there exists a reconfiguration sequence ρ of length $2m(n+2)$. Then ρ visits at least one vertex at every level, and all the vertices that it visits at a given level have the same v-state.*

Suppose there exists a reconfiguration sequence ρ of length $2m(n+2)$. Lemma 4 allows us to build an assignment θ by assigning θ_i the v-state of the vertices of level i in ρ .

Lemma 5. *The assignment θ is satisfying for ϕ .*

We also show that the reduction is complete.

Lemma 6. *If ϕ is satisfiable, then there exists a reconfiguration sequence of length at most $2m(n+2)$.*

Combining Lemma 5 and Lemma 6 together with the fact that the reduction can be clearly done in polynomial time, we have the following theorem.

Theorem 2. *The MIN-SPR problem is NP-hard, even if k is polynomial in $|V(G)|$.*

3 Independent set reconfiguration: the models

We now turn our attention to reconfiguring independent sets. Consider the reconfiguration rule in which a move from one valid configuration to another is made by a *token jump*: moving a token from one vertex to an unoccupied vertex (not necessarily a neighbor of it), such that the resulting set is independent. The token sliding, token jumping, and token addition and removal reconfiguration rules give rise to the following three reconfigurability problems.

Token sliding (TS) / token jumping (TJ): Given a graph G and two independent sets A, B in G , determine if A can be reconfigured into B via a sequence of independent sets, each of which results from the previous one by a single token slide (for TS) or jump (for TJ).

Token addition and removal (TAR): Given a graph G , an integer k and two independent sets A, B in G , both of size $\geq k$, is there a way to transform A into B via independent sets, each of which results from the previous one by adding or removing one vertex of G , without ever going through an independent set of size less than $k-1$?

We say that A and B are TS- (TAR-, TJ-) reconfigurable if they belong to the same connected component of the TS- (TAR-, TJ-) reconfiguration graph. We now establish the equivalence between the TJ and TAR problems (the proof can be found in the full version [13]).

Theorem 3. *Two independent sets A and B of size s in a graph G are TJ-reconfigurable if and only if they are TAR-reconfigurable with parameter $k = s$. Moreover, $\text{dist}_{\text{TAR}}(A, B) = 2\text{dist}_{\text{TJ}}(A, B)$, and there exists an algorithm that, given a reconfiguration sequence between two independent sets in one of these two models outputs a reconfiguration sequence connecting the two sets in the other model in time polynomial in the length of the sequence. The algorithm maps every shortest TAR-sequence to a shortest TJ-sequence, and vice versa.*

Theorem 3 immediately implies that results holding for the TAR model can be transferred to the TJ model. New results can also be derived via this relationship:

Corollary 1. *There exists a polynomial-time algorithm for the TJ problem in line graphs.*

Proof. By Theorem 3, the TJ problem in line graphs is polynomially reducible to the TAR problem in line graphs. Due to the correspondence between matchings in a graph and independent sets in its line graph, the problem is polynomially equivalent to the MATCHING RECONFIGURATION problem. For a polynomial-time algorithm for this problem, see Ito et al. [11]. \square

4 Hardness of independent set reconfiguration

TS, TAR, and TJ reconfiguration problems are all PSPACE-complete in general graphs. For the TS problem, this was announced in [10] (see also [1]) without an explicit proof. For the TAR problem, this was shown by Ito et al. [11]. In fact, their proof uses only token slides (which are done by token additions and removals), also implying that TS is PSPACE-complete. Theorem 3 immediately implies PSPACE-completeness for the TJ model. The hardness of these reconfigurability problems of course implies the hardness of the more difficult problems of finding the length of the shortest reconfiguration sequence. However, in this section we show that these related variants remain NP-hard even when the graph is restricted to be perfect.

We use a reduction from the reconfiguration of shortest paths. Given a graph G and two vertices s and t of G at a distance k apart, let G_1 denote the graph obtained from G by deleting from it all the vertices and edges not appearing on any shortest (s, t) -path. For $i \in \{0, 1, \dots, k\}$, let D_i be the set of vertices in G_1 at distance i from s and $k - i$ from t . The graph G' is the graph obtained from G_1 by turning every set D_i into a clique, and complementing the edges of G_1 between every pair of consecutive layers D_i and D_{i+1} . Formally, $V(G') = V(G_1)$ and $E(G') = \{uv : u \neq v, \exists i \text{ such that } u, v \in D_i\} \cup \bigcup_{i=0}^{k-1} \{uv : u \in D_i, v \in D_{i+1}, uv \notin E(G_1)\}$. The

idea of the construction is that there is a bijective correspondence between shortest (s, t) -paths in G and independent sets of size $k + 1$ in G' . This gives the following theorem (the proof is straightforward):

Theorem 4. *For every graph G , there is a polynomially computable length-preserving bijection (length-doubling for TAR) between shortest reconfiguration sequences in the shortest path reconfiguration graph of G and those in the TS- (TJ-, TAR-) reconfiguration graph for G' .*

The following corollary is a direct consequence of Theorem 2 and the fact that graph G' contains no odd holes or their complements and hence is perfect [5]. Recall that a *hole* in a graph is a chordless cycle with at least four vertices, and a hole is even (odd) if it has an even (odd) number of vertices.

Corollary 2. *Let G be a perfect graph, A, B two independent sets in G , and k an integer. It is NP-hard to determine if there exists a reconfiguration sequence of length at most k between A and B in the TS, TJ, and TAR models, even if k is polynomial in $|V(G)|$.*

5 Positive results for independent set reconfiguration

In this section we identify two restrictions on the input graphs which make the reconfigurability of independent sets easy to solve.

5.1 Even-hole-free graphs in the TJ model

We will show that two token sets of the same size in any even-hole-free graph are TJ-reconfigurable. Given a graph G and two independent sets A and B in G of the same size, the *Piran graph* $\Pi(A, B)$ of A and B is the subgraph of G induced by the vertex set $(A \setminus B) \cup (B \setminus A)$. The following simple lemma gives a sufficient condition under which it is always possible to jump a token from A to B .

Lemma 7. *Let A and B be two independent sets of the same size in a graph G . If the Piran graph $\Pi(A, B)$ is even-hole-free then there exists a token in $B \setminus A$ with at most one neighbor in $A \setminus B$.*

Proof. The Piran graph is bipartite, and as such, it does not contain odd cycles. If in addition, $\Pi(A, B)$ is even-hole-free, then it must be a forest. Since $|A \setminus B| = |B \setminus A|$, the number of edges in $\Pi(A, B)$ is in fact at most $|A \setminus B| + |B \setminus A| - 1 = 2|B \setminus A| - 1$. Therefore there exists a vertex in $B \setminus A$ with at most one neighbor in $A \setminus B$. \square

A consequence of Lemma 7 is the following result (a proof of which can be found in the full version [13]).

Theorem 5. *Let A and B be two independent sets of the same size in a graph G . If the Piran graph $\Pi(A, B)$ is even-hole-free, then A and B are TJ-reconfigurable. Moreover, there exists an algorithm running in time $O(|A|)$ that (if the Piran graph is even-hole-free) finds a shortest TJ-path between the two sets.*

The class of even-hole-free graphs includes the well known class of chordal graphs (hence also trees and interval graphs). The structure of even-hole-free graphs is understood and membership in this class can be decided in polynomial time [6]. Notice that if the input graph is even-hole-free, so is the Piran graph. Due to Theorem 5 we can easily solve the TJ reconfiguration problem for the class of even-hole-free graphs. Interestingly, determining the complexity of computing the maximum size of an independent set in an even-hole-free graph is, to the best of our knowledge, an open problem.

The example of the claw $G = K_{1,3}$ with leaves $\{v_1, v_2, v_3\}$, and the independent sets $A = \{v_1, v_2\}$, $B = \{v_1, v_3\}$, shows that the analogue of Theorem 5 does not hold for the TS model for the whole class of even-hole-free graphs. We leave it as an open question to determine whether the analogue holds for the class of (claw, even-hole)-free graphs.

5.2 P_4 -free graphs in the TS model

In this subsection we give a polynomial time algorithm to solve the TS problem in P_4 -free graphs. P_4 -free graphs (also known as cographs) are graphs without an induced subgraph isomorphic to a 4-vertex path. A polynomial-time algorithm for token sliding in P_4 -free graphs can be developed based on the following well-known characterization of P_4 -free graphs [7]: a graph G is P_4 -free if and only if for every induced subgraph F of G with at least two vertices, either F or the complement to F is disconnected. A *co-component* of a graph $G = (V, E)$ is the subgraph of G induced by the vertex set of a connected component of the complementary graph $\bar{G} = (V, \{uv \mid u, v \in V, u \neq v, uv \notin E\})$.

Theorem 6. *The TS problem is solvable in time $O(|V| + |E|)$ if the input graph $G = (V, E)$ is P_4 -free. Moreover, a shortest reconfiguration sequence, if it exists, can be found in time $O(|V| + |E|)$.*

Proof. We claim that Algorithm 1 below solves the TS problem on P_4 -free graphs.

The correctness of the algorithm is straightforward, using the above-mentioned characterization of P_4 -free graphs [7]. Using the result of Corneil et al. [8] showing that the decomposition of a P_4 -free graph $G = (V, E)$ into one-vertex graphs by means of taking components or co-components can be computed in time $O(|V| + |E|)$ [8], it is also easy

Algorithm 1 TS-reconfiguration of independent sets in P_4 -free graphs

Input: A P_4 -free graph $G = (V, E)$ and two independent sets A, B .

Output: A shortest (A, B) -path in the TS-graph, if one exists, NO otherwise.

```
1: if  $|V(G)| = 1$  then return the trivial TS-path if  $|A| = |B|$ , and NO otherwise.
2: if  $G$  is disconnected, with connected components  $C_1, \dots, C_m$  then
3:   if there is an  $i \in \{1, \dots, m\}$  such that  $|A \cap C_i| \neq |B \cap C_i|$  then return NO.
4:   else solve the problem recursively for the connected components  $C_1, \dots, C_m$ 
       with respective token sets  $(A \cap C_1, B \cap C_1), \dots, (A \cap C_m, B \cap C_m)$ .
5:   if one of the outputs is NO then return NO.
6:   else merge the corresponding  $(A \cap C_i, B \cap C_i)$ -paths into an  $(A, B)$ 
       TS-path  $P$ , return  $P$ .
7: else
8:   if  $|A| = |B| = 1$  then return an  $(A, B)$  TS-path corresponding to a shortest
        $(A, B)$ -path in  $G$ .
9:   else
10:    if  $A$  and  $B$  are in the same co-component of  $G$  then solve the problem
        for  $A$  and  $B$  recursively on that co-component and return the output.
11:    else return NO.
```

to see that the algorithm can be implemented so that it runs in linear time. \square

Theorem 6 can be used to prove that the TS problem is solvable in polynomial time if the input graph is (claw, paw)-free (recall that the claw is $K_{1,3}$ and the paw is the graph obtained from the claw by adding one edge). This is due to the observation that the only connected (claw, paw)-free graph containing an induced P_4 are (long enough) paths and cycles.

6 Concluding remarks

In this paper, we studied the reconfiguration variants of the shortest path and independent set problems. We believe that the major open problem is to determine the complexity of deciding whether two shortest paths are reconfigurable. If the problem is NP-hard, then it will be the first example of an efficiently solvable reconfigurability problem with reconfiguration graphs of large diameter. If the problem is polynomially solvable, then it will be the first example of a “natural” problem in P whose reconfigurability version is NP-hard.

Acknowledgements. We are grateful to Paul Bonsma, Takehiro Ito and Daniel Pellicer for interesting and fruitful discussions. We also thank Paul Bonsma for pointing us to problems in P whose reconfiguration versions are PSPACE-complete.

References

1. Paul S. Bonsma and Luis Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theor. Comput. Sci.*, 410(50):5215–5226, 2009.
2. Paul S. Bonsma, Luis Cereceda, Jan van den Heuvel, and Matthew Johnson. Finding paths between graph colourings: Computational complexity and possible distances. *Electronic Notes in Discrete Mathematics*, 29:463–469, 2007.
3. Luis Cereceda, Jan van den Heuvel, and Matthew Johnson. Connectedness of the graph of vertex-colourings. *Discrete Mathematics*, 308(5-6):913–919, 2008.
4. Luis Cereceda, Jan van den Heuvel, and Matthew Johnson. Mixing 3-colourings in bipartite graphs. *European Journal of Combinatorics*, 30:1593–1606, 2009.
5. Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. *Ann. of Math.*, 164:51–229, 2006.
6. Michele Conforti, Gérard Cornuéjols, Ajtai Kapoor, and Kristina Vušković. Even-hole-free graphs part II: Recognition algorithm. *J. Graph Theory*, 40:238–266, 2002.
7. Derek G. Corneil, H. Lerchs, and L. Stewart Burlingham. Complement reducible graphs. *Discrete Applied Mathematics*, 3(3):163–174, 1981.
8. Derek G. Corneil, Yehoshua Perl, and Lorna K. Stewart. A linear recognition algorithm for cographs. *SIAM J. Comput.*, 14(4):926–934, 1985.
9. Parikshit Gopalan, Phokion G. Kolaitis, Elitza N. Maneva, and Christos H. Papadimitriou. The connectivity of Boolean satisfiability: Computational and structural dichotomies. *SIAM J. Comput.*, 38(6):2330–2355, 2009.
10. Robert A. Hearn and Erik D. Demaine. PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theor. Comput. Sci.*, 343(1-2):72–96, 2005.
11. Takehiro Ito, Erik D. Demaine, Nicholas J. A. Harvey, Christos H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. In *ISAAC*, volume 5369 of *Lecture Notes in Computer Science*, pages 28–39. Springer, 2008.
12. Takehiro Ito, Marcin Kamiński, and Erik D. Demaine. Reconfiguration of list edge-colorings in a graph. In *WADS*, volume 5664 of *Lecture Notes in Computer Science*, pages 375–386. Springer, 2009.
13. Marcin Kaminski, Paul Medvedev, and Martin Milanic. Shortest paths between shortest paths and independent sets. *CoRR*, abs/1008.4563, 2010.