AUSTRALASIAN JOURNAL OF COMBINATORICS
Volume **81(3)** (2021), Pages 339–356

On reconfiguration graphs: an abstraction

Eileen Melville Beth Novick Svetlana Poznanović

Department of Mathematical Sciences
Clemson University
Clemson, SC 29634, U.S.A.
{rmelvil,nbeth,spoznan}@clemson.edu

Abstract

We address the structure of reconfiguration graphs by considering the questions of whether the classes of shortest path, coloring, and matroid independent set reconfiguration graphs are closed under taking union, Cartesian product, connected components, and Cartesian factors, expanding what is known. In order to give uniform proofs and highlight the similarities between these classes of graphs we introduce the notion of abstract reconfiguration spaces and show how certain properties of the space imply properties of the associated family of graphs.

1 Introduction

A reconfiguration graph for a given search problem has as vertices all feasible solutions to the problem, and two solutions are adjacent if and only if one can be obtained from the other by one application of a specific reconfiguration rule. The question of whether the reconfiguration graph is connected has been addressed in a number of settings, including vertex coloring [5, 8], list-edge coloring [17], clique, set cover, integer programming, matching, spanning tree, matroid bases [16], block puzzles [15], shortest path [19], vertex independent set [15, 16, 24, 20], and Boolean satisfiability [12]. Intimately related to this are the questions of deciding whether two feasible solutions are in the same connected component and finding paths between them [6, 10, 9, 16], as well as finding the diameter of the reconfiguration graph, see [4] for example. Graph theorists have also addressed the questions of classifying the reconfiguration graphs for certain problems such as the shortest path [1], k-coloring [3], and domination graphs [13]. While a lot of results have been established in this direction, full characterization of the structure of various classes of reconfiguration graphs has not been found.

Our initial motivation involved the shortest path reconfiguration problem, which is interesting in the following sense: While the shortest path problem is polynomially solvable, the corresponding reconfiguration problem is PSPACE-complete [19]. With

ISSN: 2202-3518

this in mind, the authors of [1] initiated the study of the structure of shortest path graphs.

A natural question for a class of graphs is to consider if it is closed under certain well-known graph operations. Graph products are especially well studied. Median graphs, of which hypercubes represent a special case, are closed under Cartesian product [25]. The closure under the Cartesian and other types of graph products of 1-perfectly orientable graphs was addressed in [14]. Graham's pebbling conjecture [11] is that the pebbling number of the Cartesian product of two graphs is bounded above by the product of the pebbling numbers of those graphs. Another natural operation is disjoint union. Obviously, planar graphs are closed under disjoint union, as are k-colorable graphs, bipartite graphs, and many more.

In [1] a very natural proof is provided that the class of shortest path graphs is closed under Cartesian product. A nice intuitive explanation of why k-coloring reconfiguration graphs are also closed under Cartesian product is given in [3] . With a bit of effort the authors of [1] showed that shortest path graphs are closed under disjoint union.

We were intrigued by the "reverse" questions: namely, are shortest path graphs closed under connected components and under Cartesian factors? In the process of determining that the answer to the first question is affirmative we realized that our result was in fact rather general. This lead to a new concept—an abstract reconfiguration space, the topic of the present paper. To illustrate the power of our generalization we explore three reconfiguration problems chosen because they represent, in a sense, a spectrum of possibilities. For $k \geq 4$, the vertex coloring problem is NP-complete, as is its corresponding reconfiguration problem. See [10] for an interesting discussion of the case $k = 3$. The matroid independence problem and its reconfiguration problem are both easy from a complexity point of view.

Thus, in the process of addressing the classification problem, four natural questions arise:

(Q1) Is a given class of reconfiguration graphs closed under disjoint union?

(Q2) Is a given class of reconfiguration graphs closed under Cartesian product?

(Q3) Is a connected component of a reconfiguration graph necessarily a reconfiguration graph for the same problem (e.g. shortest path, coloring, etc.)?

(Q4) Is a Cartesian factor of a reconfiguration graph necessarily a reconfiguration graph for the same problem?

In the case of the shortest path reconfiguration graphs, questions (Q1) and (Q2) have been answered affirmatively in [1]. In the case of k-coloring graphs, question (Q2) has been addressed affirmatively in [3]. The work in the present paper started with the goal to answer the remaining questions for these two classes of reconfiguration graphs. The initial proofs we discovered revealed similarities between the two classes which led us to define abstract reconfiguration spaces and prove properties about the closure under these four operations in a more general setting. In order to show that our setting of abstract reconfiguration spaces is not restrictive to just the two cases by which it was motivated we answer questions (Q1)-(Q4) for the collection of independent set reconfiguration graphs of matroids as well, where we apply the results for abstract reconfiguration spaces.

We see our results as primarily an important contribution to our understanding of the structure of shortest path reconfiguration graphs. Secondly, we propose that abstract reconfiguration sequences may yield results in the arena of reconfiguration more broadly.

We begin, in Section 2, by giving the complete definitions of the three reconfiguration spaces that we will analyze, as well as giving the necessary notation. In Section 3 we define abstract reconfiguration spaces and derive some properties about them. Then, in Section 4 we come back to shortest paths, coloring, and matroid independent set reconfiguration graphs and we answer questions (Q1)-(Q4) for each of them, applying the theory of abstract reconfiguration spaces whenever possible. Table 1 summarizes the results about the reconfiguration graphs under consideration.

	Shortest Path Graphs	Coloring Graphs	Matroid Indep't Set Graphs
Q1	Yes, [1]	No, Example 4.5	No, Section 2.3
Q2	Yes, [1]	Yes, [3]	Yes, Section 4.3
Q3	Yes, Theorem 4.2	No, Example 4.7	Vacuously true, Sec. 2.3
Q4	Yes, Theorem 4.4	Yes if connected, Thm. 4.10	Yes, Theorem 4.12

Table 1: Answering (Q1)–(Q4) for three classes of reconfiguration problems

2 Background

The *disjoint union* of graphs G and H will be denoted by $G+H$. It is the graph with vertex set (edge set, respectively) equal to the disjoint union of the vertex sets (edge sets, respectively) of G and H. The *Cartesian product* $G \square H$ of graphs G and H is a graph such that the vertex set of $G \square H$ is the Cartesian product $V(G) \times V(H)$; and two vertices (u, v) and (u', v') are adjacent in $G \square H$ if and only if either $u = u'$ and v is adjacent to v' in H, or $u' = v'$ and u is adjacent to v in G. H is a *Cartesian factor* of a graph G if there is a graph K such that $G = H \square K$.

2.1 Shortest path reconfiguration graphs

For a graph G and two vertices a and b of G, we will denote by $S(G, a, b)$ the reconfiguration graph for the shortest paths between a and b. The vertices of $S(G, a, b)$ are the shortest paths between a and b and two vertices in $S(G, a, b)$ are connected by an edge if the paths in G differ in exactly one vertex. We say that a graph is a *shortest path reconfiguration graph*, or *shortest path graph*, if it is equal to $S(G, a, b)$ for some graph G and vertices a and b. The shortest path reconfiguration problem is noteworthy in that while efficient polynomial time algorithms exist for finding a shortest path in a graph, the problem of deciding if a shortest path reconfiguration graph is connected is PSPACE-complete [6]. The study of the structure of shortest path reconfiguration graphs was initiated in [1], where the questions (Q1) and (Q2) were answered.

Theorem 2.1 ([1]). *If H and K are shortest path graphs then $H + K$ and $H \,\square\, K$ are also shortest path graphs.*

Our main results include the affirmative answers to (Q3) and (Q4) for shortest path graphs. This is done in Section 4.1.

2.2 Coloring reconfiguration graphs

For a positive integer k and a graph G, we will denote by $\mathcal{C}_k(G)$ the *coloring reconfiguration graph* for the proper k-colorings of G. A proper k-coloring of G is an assignment of one of k colors, typically denoted $1, 2, \ldots, k$, to each vertex of G such that no two adjacent vertices are assigned the same color. These form the vertex set of $\mathcal{C}_k(G)$. Two colorings in $\mathcal{C}_k(G)$ are adjacent if they differ in the color of exactly one vertex in G.

The coloring graph arises in theoretical physics when studying the Glauber dynamics of an anti-ferromagnetic Potts model at zero temperature (see [2] and the references therein for the connections). The set of all proper k-colorings of a graph G forms the state space for a Markov chain with iterations given by randomly recoloring a randomly selected vertex of G. When this process exhibits rapid mixing, good estimates of the total number of proper k-colorings of G are obtained [18, 23, 27].

The connectivity of the k-coloring graph has been addressed in [8], while the decidability of whether two solutions are in the same connected component has been addressed in [7].

The positive answer to question (Q2) for k-coloring graphs follows from the following result.

Theorem 2.2 ([3]). *If $G = \sum_{i=1}^{n} G_i$ then $\mathcal{C}_k(G) = \square_{i=1}^{n} \mathcal{C}_k(G_i)$.*

In Section 4.2 we give examples that show that the answers to the questions (Q1), (Q3), and (Q4) are negative in general, but we show that a Cartesian factor of a k-coloring reconfiguration graph G is also a k-coloring graph if G is connected. We will make use of the the following property of coloring graphs that has already been observed in [3]. We include a proof for completeness.

Lemma 2.3. *The order of a k-coloring graph is either 1 or a multiple of k.*

Proof. The number of k-colorings of a graph is known to be a polynomial function of k and given by the so called chromatic polynomial [28]. The claim follows from the fact that the constant term of the chromatic polynomial is always 0, because there are no ways to color a graph using zero colors. □

2.3 Matroid independent set reconfiguration graphs

A matroid is a structure that generalizes the properties of independence in vector spaces, see [26]. It can be defined in several equivalent ways. Here we give the definition in terms of independent sets.

A *finite matroid* M is a pair $(E, \mathcal{I})$ where E is a finite set (called the ground set) and $\mathcal{I}$ is a nonempty family of subsets of E (called the independent sets) with the following properties:

1. (hereditary property) Every subset of an independent set is independent.
2. (augmentation property) If A and B are two independent sets and A is larger than B, then there exists $x \in A \backslash B$ such that $B \cup \{x\}$ is in $\mathcal{I}$.

The *direct sum* of two matroids $M_1 = (E_1, \mathcal{I}_1)$ and $M_2 = (E_2, \mathcal{I}_2)$ is the matroid $M = (E, \mathcal{I})$ with $E = E_1 \sqcup E_2$ and $\mathcal{I} = \{I_1 \cup I_2 \colon I_1 \in \mathcal{I}_1, I_2 \in \mathcal{I}_2\}$.

We denote by $\mathcal{G}(M)$ the reconfiguration graph for the independent set reconfiguration problem associated with the matroid M. The vertices of $\mathcal{G}(M)$ are the independent sets of M and two vertices are adjacent if their incidence vectors differ by one element. If $G = \mathcal{G}(M)$ for some matroid M, we refer to G as a *matroid independent set graph*. It follows from the hereditary property that $\mathcal{G}(M)$ is connected for every matroid M. The fact that all matroid independent set graphs are connected implies that the class is not closed under disjoint union. It is vacuously true that if $G + H$ is a matroid independent set graph, so are G and H. The fact the direct sum of matroids is a matroid implies that the Cartesian product of two matroid independent set graphs is an independent set graph. We will show in Section 4.3 that if $G \,\square\, H$ is a matroid independent set graph, then so are G and H. We note that a related concept, namely the basis graph of a matroid has been well studied, see for example [21, 22].

3 Abstracting reconfiguration graphs

In this section we introduce the concept of 'reconfiguration space', one of the goals being to give uniform treatment to the reconfiguration graphs from Section 2. The motivation for our definition is the fact that in each of the three cases that we consider, an instance of the problem corresponds to a collection of sequences. For example, in the case of shortest path reconfiguration, the shortest paths between a and b are sequences of vertices in the base graph. We will omit the vertices a and b in this list. For the coloring reconfiguration problem, we fix an ordering of the vertices in the base graph, say $v_1, \ldots, v_n$. Then a coloring c is given by the sequence $c(v_1), \ldots, c(v_n)$. Lastly, for the matroid independent set reconfiguration problem, we again fix the ordering of the elements in the ground set and an independent set is represented by a corresponding 0/1 incidence vector. In each case the sequences correspond to vertices in the reconfiguration graph. Note that the sequences all have the same length and they are adjacent in the reconfiguration graph if they differ in exactly one position.

In order to delineate between various classes of reconfiguration problems we introduce the term 'reconfiguration space'.

Definition 3.1. A *sequence set* S is a finite set of sequences each of the same length, called the *length* of S. A *reconfiguration space* $\mathcal{S}$ is a collection of sequence sets.

Definition 3.2. For a sequence set S in a reconfiguration space $\mathcal{S}$ we denote by G_S the corresponding *reconfiguration graph*. Namely, G_S is the graph whose vertices are the elements of S and there is an edge between two vertices if and only if the corresponding sequences differ in exactly one position. The set $\{G_S \colon S \in \mathcal{S}\}$ is called the *graph set* of $\mathcal{S}$ and will be denoted by $\{G_{\mathcal{S}}\}$. When the reconfiguration space $\mathcal{S}$ is understood, by a *reconfiguration graph* we mean a graph G so that G is in $\{G_{\mathcal{S}}\}$. If the sequences in a sequence set S corresponding to two adjacent vertices $\mathbf{a}$ and $\mathbf{a}'$ in G_S are $a_1 \dots a_{i-1}a_ia_{i+1} \dots a_n$ and $a_1 \dots a_{i-1}a_i'a_{i+1} \dots a_n$, we say that those vertices differ in the i-th *index*, or that i is the *difference index* of the edge $\mathbf{aa}'$.

We first prove some properties in the abstract context, namely Lemma 3.3 and Proposition 3.4, that were proven already in the context of shortest path reconfiguration graphs in [1].

Lemma 3.3. *Let G be a reconfiguration graph.*

(i) If p_1 and p_2 are two paths in G between vertices v_1 and v_2, and the difference index c appears exactly once in p_1, then it must appear in p_2.

(ii) If vertices v_1, v_2, and v_3 form a triangle in G, then all the edges v_1v_2, v_2v_3, and v_1v_2 must have the same difference index. Furthermore, if there exist edges v_1v_2 and v_2v_3 in G with the same difference index, then v_1v_3 is also an edge in G.

(iii) If $v_1v_2v_3v_4$ is a (not necessarily induced) 4-cycle in G then the edges v_1v_2 and v_3v_4 must have the same difference index.

Proof. (i) The fact that the difference index c appears exactly once in p_1 implies that the sequences v_1 and v_2 differ at index c. Therefore, any other path from v_1 to v_2 must account for the change in that index as well.

(ii) Note that v_1v_2 and $v_1v_3v_2$ are both paths between v_1 and v_2. By part (i), the difference index of v_1v_2 must also appear in $v_1v_3v_2$. By symmetry we conclude that no difference index can appear exactly once in the triangle, which means all three difference indices must be the same.

Suppose now that v_1v_2 and v_2v_3 are two different edges in G that share the same difference index. Then the sequences v_1, v_2, and v_3 all differ in the same position, which implies that v_1 and v_3 are connected by an edge (with the same difference index).

(iii) Here, v_1v_2 and $v_1v_4v_3v_2$ are both paths from v_1 to v_2, so, by part (i), they must both contain the difference index of v_1v_2. If the difference index of v_4v_3 is the same, then we are done, so we may assume without loss of generality that v_1v_4 has the same difference index as v_1v_2. Then, by part (ii), v_2v_4 is an edge in G and has the same difference index as v_1v_2. Then we may conclude that $v_2v_4v_3$ is another triangle, and thus has all edges sharing the same difference index. Thus v_1v_2 has the same difference index as v_2v_4 which has the same difference index as v_4v_3.

□

Proposition 3.4. *Let $G = H \,\Box\, K$ be a connected reconfiguration graph.*

(i) Any two edges in G of the form $(a_i, b)(a_j, b)$ and $(a_i, b')(a_j, b')$ have the same difference index.

(ii) Any two edges in G of the form $(a_i, b_j)(a_k, b_j)$ and $(a_\ell, b_m)(a_\ell, b_n)$ have distinct difference indices.

Proof. (i) Since G is connected, so are H and K. Therefore, there is a path $b = b_0, b_1, b_2, b_3, \ldots, b_\ell = b'$ in H from b to b'. Then for every $k \in \{1, \ldots, \ell\}$, $(a_i, b_{k-1}), (a_j, b_{k-1}), (a_j, b_k), (a_i, b_k)$ is a 4-cycle in G and, by part (iii) of Lemma 3.3, the edges $(a_i, b_{k-1})(a_j, b_{k-1})$ and $(a_i, b_k)(a_j, b_k)$ have the same difference index, which imples the claim.

(ii) By part (i), the edge $(a_\ell, b_m)(a_\ell, b_n)$ has the same difference index as $(a_i, b_m)(a_i, b_n)$. Also the edge $(a_i, b_j)(a_k, b_j)$ has the same difference index as$(a_i, b_m)(a_k, b_m)$. Suppose $(a_i, b_j)(a_k, b_j)$ and $(a_\ell, b_m)(a_\ell, b_n)$ have the same difference index. Then the edges $(a_i, b_m)(a_k, b_m)$ and $(a_i, b_m)(a_i, b_n)$ also have the same difference index. This by part (ii) of Lemma 3.3, implies that there is an edge between (a_k, b_m) and (a_i, b_n), which is a contradiction. Therefore, the edges $(a_i, b_j)(a_k, b_j)$ and $(a_\ell, b_m)(a_\ell, b_n)$ cannot have the same difference index.

□

Definition 3.5. The *product* of two sequence sets A and B in a reconfiguration space $\mathcal{S}$ contains all possible concatenations of a sequence in A and a sequence in B. Namely, their product is the set $AB = \{ab\colon a \in A, b \in B\}$. A reconfiguration space $\mathcal{S}$ is *closed under sequence concatenation* if for any sequence sets A and B in $\mathcal{S}$, their product AB is also in $\mathcal{S}$.

The following is a straightforward generalization of the results about shortest path (Theorem 2.1) and coloring reconfiguration graphs (Theorem 2.2).

Theorem 3.6. *Let $\mathcal{S}$ be a reconfiguration space. If $\mathcal{S}$ is closed under sequence concatenation then for any $G_A, G_B \in \{G_\mathcal{S}\}$ we have $G_A \,\Box\, G_B = G_{AB}$, and consequently the graph set $\{G_\mathcal{S}\}$ is closed under Cartesian product.*

Definition 3.7. Let S be a sequence set in a reconfiguration space $\mathcal{S}$. Let n be the length of S. We refer to $[n] = \{1, 2, \ldots, n\}$ as the *index set of* S. Let $J = \{i_1, i_2, \ldots, i_\ell\}$ be any non-empty subset of the index set of S with $i_1 < i_2 < \cdots < i_\ell$. Let $a = (a_1, a_2, \ldots, a_n) \in S$. Then by the *restriction* of a to J, denoted $a|_J$, we mean the subsequence $(a_{i_1}, a_{i_2}, \ldots, a_{i_\ell})$ of a. By the *restriction* of S to J, we mean the set $S|_J = \{a|_J \, : \, a \in S\}$. Note that the latter is a set: if the restriction results in multiple copies of some elements, we keep only one. If for all $S \in \mathcal{S}$, and all subsets J of the index set of S, the set $\{a|_J \, : \, a \in S\}$ is also in $\mathcal{S}$, we say that $\mathcal{S}$ is *closed under taking subsequences.*

Next we establish a property of reconfiguration spaces closed under taking subsequences. Specifically, we show that the collection of connected reconfiguration graphs arising in such a space are closed under taking Cartesian prime factors.

We first need another definition, a proposition, and a technical lemma.

Definition 3.8. Let S be a sequence set in a reconfiguration space $\mathcal{S}$ with index set $I = [n]$. Let $J \subseteq I$. If a is a sequence in $S_1 = S|_J$ and b is a sequence in $S_2 = S|_{I \setminus J}$ then we say that the *mix* of a and b, denoted $a \times_J b$, is the sequence of length n whose restriction to J is a and whose restriction to $I \setminus J$ is b. We will say that S is the *mix* of S_1 and S_2 with respect to $\{J, I \setminus J\}$, denoted $S_1 \times_J S_2$, if

$$S = \{a \times_J b \, : \, a \in S_1 \text{ and } b \in S_2\}.$$

We will use $a \times b$ and $S_1 \times S_2$ when J is understood.

Note that concatenation is a special kind of mixing. The following result is in the same vein as Theorem 3.6 and its proof is also straightforward.

Proposition 3.9. *If S is the mix of S_a and S_b, then $G_S = G_{S_a} \,\square\, G_{S_b}$.*

Lemma 3.10. *Let $G = H \,\square\, K$ be a connected reconfiguration graph. Let S be a sequence set with $G = G_S$. Then S is the mix of two collections of subsequences, say S_a and S_b, with $H \cong G_{S_a}$ and $K \cong G_{S_b}$.*

Proof. Let $H \,\square\, K$ be a connected reconfiguration graph, G_S, arising from a sequence set S. Let k be the length of S. By part (*ii*) of Proposition 3.4, the index set $I = [k]$ can be partitioned into $I_a = \{i_1, i_2, \ldots, i_m\}$ and $I_b = \{j_1, j_2, \ldots, j_r\}$, with $m + r = k$, where we let the former correspond to edges in $H \,\square\, K$ of the form $(a_i, b_j)(a_k, b_j)$, and the latter to edges of the form $(a_\ell, b_m), (a_\ell, b_n)$ together with any index that is not a difference index in G (note that because G is connected any such index has constant value in S). Let $S_a = S|_{I_a}$ and let $S_b = S|_{I_b}$.

<u>Claim</u>: $S = S_a \times S_b$, where the mixing is with respect to $\{I_a, I_b\}$.

Proof of Claim. By construction, $S \subseteq S_a \times S_b$. Assume for a contradiction that this set containment is strict. So there is a sequence $u_1 \in S_a$, and a sequence $v_2 \in S_b$, for which $u_1 \times v_2 \notin S$. The sequence u_1 is the restriction of some vertex $u = u_1 \times u_2$

in G_S to I_a. Similarly, v_2 is the restriction of some vertex $v = v_1 \times v_2$ in G_S to I_b. As u and v are both in $H \,\square\, K$, we can express u and v, respectively, as (a_i, b_j) and (a_ℓ, b_m). By the nature of Cartesian product, there is a shortest path in $H \,\square\, K$ of the form

$$(a_i, b_j), \ldots, (a_\ell, b_j), \ldots, (a_\ell, b_m),$$

where the subpath between (a_i, b_j) and (a_ℓ, b_j), uses only indices in I_a, while in the subpath between (a_ℓ, b_j) and (a_ℓ, b_m), all changes occur in I_b. It follows that $(a_\ell, b_j) = u_1 \times v_2$, a contradiction. We conclude that $S = V(H \,\square\, K) = S_a \times S_b$, proving the claim. □

It now follows from Proposition 3.9 that $G = G_{S_a} \,\square\, G_{S_b}$. Hence we have a second factorization for $G = H \,\square\, K$.

To see that $H \cong G_{S_a}$, we construct a function $f : S_a \mapsto V(H)$ as follows: Fix a specific sequence, say w, in S_b. For a given sequence $u \in S_a$, let $f_1(u)$ be the mix of u and w, $u \times w$, with respect to $\{I_a, I_b\}$. For a given sequence $s \in S$, let $f_2(s)$ be the unique vertex in $G = H \,\square\, K$ to which it corresponds. Lastly, for a given vertex $(a_i, b_j) \in H \,\square\, K$, define $f_3(a_i, b_j)$ to be a_i. Now, letting f be the composition of f_1, f_2, and f_3, we see that f is an edge-preserving bijection, establishing that $H \cong G_{S_a}$.

An analogous argument shows that $K \cong G_{S_b}$. □

Note that the statement of Lemma 3.10 did not specify that S_a and S_b are reconfiguration sets in $\mathcal{S}$. But, if $\mathcal{S}$ is closed under taking subsequences, then S_a and S_b are indeed in $\mathcal{S}$, and in this case our desired result follows easily:

Theorem 3.11. *If $\mathcal{S}$ is a reconfiguration space closed under taking subsequences then the collection of connected graphs in $\{G_\mathcal{S}\}$ is closed under taking Cartesian factors.*

Proof. Let $G = H \,\square\, K$ be a connected reconfiguration graph in a graph set $\{G_\mathcal{S}\}$, where $\mathcal{S}$ is closed under taking subsequences. Then, by Lemma 3.10, H and K are also in $\{G_\mathcal{S}\}$. □

Our results in this section address questions (Q2) and (Q4). We answer questions (Q1) and (Q3) for each of our example reconfiguration spaces separately in the next section.

4 Answering our four questions

4.1 Shortest path graphs

Recall that for a shortest path reconfiguration graph G_S, the sequences in the *shortest path sequence set* S correspond to shortest a, b-paths in the base graph, where we do not include a, b in the list. Assuming that all vertices in the base graph contribute to a shortest a, b-path, the base graph is completely determined by S. For clarity, we will say that a sequence set S is *connected* if G_S is connected. Furthermore, we say that a subset S' of a sequence set S is a *connected component* of S if $G_{S'}$ is a connected component of G_S. If the length of S is n and $1 \leq i \leq n$ with x a value at index i, then we denote by $S_i(x)$ the set of all subsequences in $S|_{[i]}$, that terminate in the value x at index i. When $S_i(x)$ is connected (not connected) we say that S_i is *connected (not connected) at the index value* x.

In this section we answer questions (Q3) and (Q4) affirmatively for the class of shortest path reconfiguration graphs. For (Q3), we unfortunately, could not rely on a general theory for reconfiguration spaces. Specifically, we note that if S is a shortest path sequence set then the presence of a sequence having values u and v at indices $i-1$ and i, respectively, indicates an edge uv in the base graph. Hence for every sequence s in $S_{i-1}(u)$, there is a sequence in $S_i(v)$ obtained from s by appending the value v at index i. We call this phenomenon the *add-an-edge property.* In fact, it is straightforward to show that this property characterizes shortest path sequence sets. Any abstraction in the context of general sequence sets that we could think of to answer (Q3) required the add-an-edge property and thus the resulting class would be very similar to the one of shortest path graphs. The answer to (Q4), on the other hand, follows as an application of Theorem 3.11 and our affirmative answer to (Q3).

The following lemma shows that while keeping the reconfiguration graph the same, we can change the corresponding base graph so that the shortest path reconfiguration graphs of certain important subgraphs of the base graph are connected.

Lemma 4.1. *If S is a shortest path sequence set of length n and size m, then there exists a shortest path sequence set S' of the same length and size, so that $G_S \cong G_{S'}$ and so that the following* connectivity property *holds:*

> *For $1 \leq i \leq n$, the sequence set $S'_i(x)$ is connected for every value x at index i.*

Proof. Fix a counter-example, S, of size m and length n. Of all sequence sets with reconfiguration graph G_S, assume we have chosen S so that the smallest index, i, for which the connectivity property fails is largest possible. Assume further that of all such sequence sets, S is chosen so that the number of index values at which S_i is not connected is smallest as well. Clearly $i > 1$, as $G_{S_1(x)} = K_1$ for all values x at index 1. For a contradiction, we construct a sequence set S' of size m, of length n, with $G_S = G_{S'}$, with S'_k connected for all values at index k provided $k < i$, but with S'_i failing to be connected at one fewer index values than does S_i. Let x be a value at index i for which S_i fails to be connected.

(i) If $a = (a_1, a_2, \ldots, a_{i-1}, x)$ and $b = (b_1, b_2, \ldots, b_{i-1}, x)$ are two sequences in $S_i(x)$ with $a_{i-1} = b_{i-1}$, then a and b are in the same connected component of $S_i(x)$.

Proof of (i). Let $a_{i-1} = b_{i-1} = y$. Since $S_{i-1}(y)$ is connected there is a path in $G_{S_{i-1}(y)}$ between $a|_{[i-1]}$ and $b|_{[i-1]}$. Since S is a shortest path sequence set it has the add-an-edge property. Hence we may append the value x to each sequence in this path in $G_{S_{i-1}(y)}$ thereby attaining a path in $G_{S_i(x)}$ between a to b. This establishes (i).

Let r be the number of connected components of $S_i(x)$ and denote them as $H_1, H_2, \ldots, H_r$.

(ii) The set of index values at $i-1$ in $S_i(x)$ can be partitioned into blocks $X_1, X_2, \ldots, X_r$ so that for $1 \leq \ell \leq r$, a sequence in $S_i(x)$ is in the component H_ℓ if and only if its index at $i-1$ is in X_ℓ. The proof is immediate from (i).

To construct S' we introduce r new values, $x_1, x_2, \ldots, x_r$ and for each sequence $s \in S$ having value x at i, we replace x with x_ℓ if index $i-1$ of s is in X_ℓ. See Figure 1 for an illustration. We note that the length of S' is n and the size of S' is m.

(iii) For $k < i$, the sequence set S'_k is connected at every index value; if y is an index value at i but $y \notin \{x_1, x_2, \ldots, x_r\}$, then $S'_i(y)$ is connected if and only if $S_i(y)$ is connected. Indeed, in replacing each occurrence of the index value x at index i with the appropriate value x_ℓ, none of the sequence sets in the statement of (iii) are changed at all.

(iv) The smallest index for which S'_i fails to be connected for some index value is greater than or equal to i, and the number of index values at which S'_i fails to be connected is one less than that of S_i. The proof is immediate from (iii) and the fact that, by construction, each $S'_i(x_\ell)$ is connected.

(v) $G_{S'} \cong G_S$. Indeed, in replacing each occurrence of the index value x at index i with the appropriate value x_ℓ, no new adjacencies are created. On the other hand, let a and b be adjacent sequences in S and let a' and b' be their new versions in S'. Let k be the unique index at which a and b differ. If $k \neq i$, then $a|_{[i]}$ and $b|_{[i]}$ are either identical or adjacent and hence in the same connected component in S_i whence the index value of a' and b' at i remain equal. If $k = i$, then at most one of $\{a, b\}$ has index value x at i whence a' and b' also differ at index i only. In either case, a' is adjacent to b'.

To see that S' is a shortest path sequence set, we note that

(vi) S' satisfies the add-an-edge property. Indeed, suppose that for some $2 \leq i \leq n$, u and v are values at indices $i-1$ and i for some sequence a in S'. Let s be a sequence in $S'_{i-1}(u)$. We need to show that by appending v to s we get a sequence in $S'_i(v)$. If neither u nor v is contained in $\{x_1, \ldots, x_r\}$ then the add-an-edge property is 'inherited' from S. If v is contained in this set, say $v = x_j$, then, in S, the sequence s can be extended to a path b whose ith index is x. By the induction hypothesis, $S_{i-1}(u)$ is connected and, therefore, the sequences $a|_{[i-1]}$ and s extended by x are in the same component of $S_i(x)$. That means that the sequences $a|_{[i-1]}$ and s extended by x_j are both in $S'_i(x_j)$. The case when $u = x_j$ is even simpler.

This completes the proof of the lemma. □

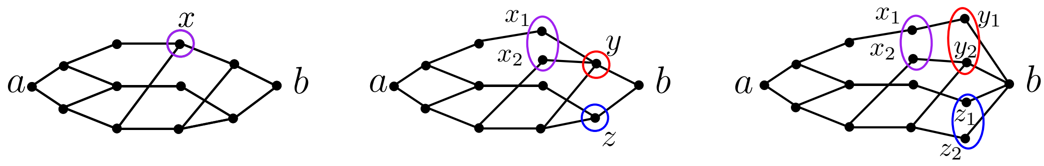

Figure 1: Successive changes to base graph as described in the proof of Lemma 4.1

Theorem 4.2. *If G is a shortest path reconfiguration graph, and H is a connected component of G then H is a shortest path reconfiguration graph.*

Proof. If G is a shortest path reconfiguration graph, then there is a shortest path sequence set S for which $G = G_S$. Let n be the length of S. By Lemma 4.1, we may choose S so that the sequence set $S_n(x)$ is connected for all values x at index n. Let H be a connected component of G. Let $S[H]$ be the subset of S that induces H, i.e., $G_{S[H]} = H$. Say there are r distinct values of S at index n and denote them as $z_1, z_2, \ldots, z_r$. Note that S is the disjoint union of r connected subsets:

$$S = \bigcup_{j=1}^{r} S_n(z_j).$$

Furthermore, as $S[H]$ is itself a connected subset of S, it follows that for $1 \leq j \leq r$, $S_n(z_j)$ is either contained in or disjoint from $S[H]$. Let $J = \{j \in [r] : S_n(z_j) \text{ is disjoint from } S[H]\}$. Construct a subset S' of S by removing all sequences whose nth index is in $\{z_j : j \in J\}$. Then S' is a sequence set with $G_{S'} = H$. To see that S' is a shortest path sequence set, we note that removing all sequences ending in z_j is equivalent to removing the corresponding vertex in the base graph. □

Next we show that shortest path reconfiguration spaces enjoy the property of being closed under taking subsequences and thus we can apply Theorem 3.11.

Proposition 4.3. *Let $\mathcal{S}$ be the reconfiguration space of shortest path graphs. Let $S = S(G, a, b)$ be an element in $\mathcal{S}$ of length k and let $J \subseteq [k]$. Then $S|_J \in \mathcal{S}$.*

Proof. It suffices to establish the result for $J = I \setminus \{j\}$ for some $j \in I$. So let $j \in I$ and let $S' = S|_J$, where $J = I \setminus \{j\}$. We construct a graph G' from G so that $a, b \in V(G')$ and so that S' is the set of shortest a, b paths in G'. Specifically, start with G and do the following to construct G':

(a) If $2 \leq j \leq k-1$, then add all edges uw for which the index of u is $j-1$, the index of w is $j+1$ and there is at least one vertex, v, 'between' u and w at index j, i.e. with uv and vw both edges of G. If $j = 1$, add the edge au for all vertices u at index 2. If $j = k$, add the edge ub for all vertices u at index $k-1$.

(b) Delete all vertices originally at index j.

To complete the proof, note that, by construction,

$$a_1 - a_2 - \cdots - a_{j-1} - a_{j+1} - \cdots - a_k$$

is a shortest a, b path in G' if and only if

$$a_1 - a_2 - \cdots - a_{j-1} - a_j - a_{j+1} - \cdots - a_k$$

is a shortest a, b path in G for some a_j. □

Theorem 4.4. *If $H \,\square\, K$ is a shortest path graph, then H and K are shortest path graphs.*

Proof. If $H \,\square\, K$ is connected then the claim follows from Proposition 4.3 and Theorem 3.11. Otherwise, assume that H and K split into connected components as $H = H_1 + \cdots + H_k$ and $K = K_1 + \cdots + K_\ell$. Then

$$H \,\square\, K = \sum_{\substack{1 \le i \le k \\ 1 \le j \le \ell}} H_i \,\square\, K_j$$

and by Theorem 4.2, each component $H_i \square K_j$ is a shortest path graph. Since $H_i \square K_j$ is connected we conclude that all H_i, $1 \le i \le k$ and K_j, $1 \le j \le \ell$ are shortest path graphs, and hence, by Theorem 2.1, so are H and K. □

4.2 Coloring graphs

In this section we address the questions (Q1) - (Q4) for coloring graphs. As we pointed out in Section 1, (Q2) has been answered affirmatively in [3]. The disjoint union of two k-coloring graphs is not necessarily a k-coloring graph as can be seen from the following example.

Example 4.5. Take G to be a k-coloring graph that does not have K_k as a factor. Note that K_k is the k-coloring graph of the graph with 1 vertex. However, we claim that the disjoint union $G + K_k$ is not a k-coloring graph. To see this, suppose otherwise, i.e., $G + K_k$ is a k-coloring graph of a base graph B. Then K_k corresponds to colorings of B that are only different in the color of a fixed vertex v (Lemma 3.3 (ii)). Since v can assume all k colors without any other changes in the coloring, v is isolated. By Theorem 2.2,

$$G + K_k = \mathcal{C}_k(B \setminus v) \,\square\, K_k,$$

which implies that G has K_k as a factor. Since this is a contradiction with the starting assumption, we see that $G + K_k$ is not a k-coloring graph.

In the light of the previous example, it is interesting to see how difficult it is to find a k-coloring graph that does not have K_k as a factor. The next proposition shows that it is not difficult at all.

Proposition 4.6. *The k-coloring graph of B has K_k as a factor if and only if B has an isolated vertex.*

Proof. If B has an isolated vertex, its k-coloring graph contains K_k as a factor by Theorem 2.2. For the converse, suppose that the k-coloring graph of B has K_k as a factor. Then, in particular, it contains K_k as a subgraph. Fix one occurrence of K_k as a subgraph. As discussed before, the k colorings correspond to changes in one vertex, which must be isolated. □

Unlike the case of path reconfiguration graphs, $G+H$ may be a k-coloring graph, but as can be seen from the next example, G and H need not be. In fact, in the next example even the connected components of a k-coloring graph need not be k-coloring graphs if $k > 1$.

Example 4.7. Take for example, G and H to be empty graphs (no edges) whose total number of vertices is $k!$. Then $G+H$ is the k-coloring graph of K_k, but if $|G|$ and $|H|$ are not divisible by k, then by Lemma 2.3, neither G nor H is a k-coloring graph.

The class of k-coloring graphs is not closed under taking Cartesian factors as can be seen from the following example.

Example 4.8. $G \,\square\, H$ may be a k-coloring graph, but G and H need not be. Let $D(k)$ be the largest power of k that divides $k!$, and let e_j be the empty graph on j vertices. Then $e_{k!}$ is a k-coloring graph, but its Cartesian factor e_r where $r = \frac{k!}{D(k)}$ is not by Lemma 2.3.

One will notice that the graphs in the previous example are not connected. Indeed, as we will show in Theorem 4.10, the class of connected coloring graphs is closed under taking Cartesian factors. However, the following infinite class of examples shows that for the k-coloring reconfiguration problem, the reconfiguration space is *not* closed under taking subsequences and, therefore, we cannot apply Theorem 3.11.

Example 4.9. Let G be the complete bipartite graph $K_{k,1}$, also known as a k-star. List the vertices of G as $v_1, v_2, \dots, v_k, v_{k+1}$, where v_{k+1} is the unique vertex of degree k. Let S be the set of sequences of proper k-colorings of G. Restricting S to the first k indices yields the set of all multisequences $(a_1, a_2, \dots, a_k)$ taken from $\{1, 2, \dots, k\}$ with at least two entries the same. The resulting set, say S', is of cardinality $k^k - k!$ and cannot be the colorings of any graph: Such a graph would necessarily be edgeless as at least one member of S' would color the ends of any edge identically. But the coloring set of $\overline{K_n}$ is the collection of all k^k colorings.

Although we cannot apply Theorem 3.11 directly, we are able to apply Lemma 3.10 to establish our next result.

Theorem 4.10. *If $H \,\square\, K$ is a connected k-coloring graph, then H and K are k-coloring graphs.*

Proof. Let $H \,\square\, K = \mathcal{C}_k(G)$ for some base graph G and k a positive integer. Assume $H \,\square\, K$ is connected. Let S be the sequence set of the k-colorings of $H \,\square\, K$. (Note that this means G_S, the reconfiguration graph associated with S, is our coloring graph, $H \,\square\, K$.) The index set, I, of our coloring graph corresponds to the vertex set of the base graph. To emphasize the fact that for a coloring graph each difference index corresponds to a vertex in the base graph we call such vertices in G *difference vertices.*

By Lemma 3.10, S is the mix of two subsequences, S_a and S_b where the mix is with respect to disjoint index sets, say I_a and I_b. It is an easy exercise to show that for a connected coloring graph every vertex is a difference vertex. Hence, by the proof of Lemma 3.10, I_a and I_b correspond to a partition $\{V_A, V_B\}$ of $V(G)$, where difference vertices in V_A correspond to adjacencies in $H \,\square\, K$ of the form $(a_i, b_\ell) \sim (a_j, b_\ell)$, and difference vertices in V_B to adjacencies of the form $(a_\ell, b_i) \sim (a_\ell, b_j)$. Also by Lemma 3.10, we have $H \cong G_{S_a}$ and $K \cong G_{S_b}$.

Claim: G_{S_a} and G_{S_b} are the coloring graphs of the subgraphs induced by the vertex sets V_A and V_B, respectively.

To establish the claim, by Theorem 2.2 it suffices to show that there is no edge in the base graph with one end in V_A and the other in V_B. To show that this is true, assume for a contradiction that $(v_a, v_b) \in E(G)$ with $v_a \in V_A$ and $v_b \in V_B$. Consider some coloring C of G. Then $C = (C_A, C_B)$ with $C_A \in G_{S_a}$ and $C_B \in G_{S_b}$. Let c_1 be the color of v_a in C and let c_2 be the color of v_b. Now consider the (proper) coloring C' of G obtained by interchanging the roles of c_1 and c_2. We have $C' = (C'_A, C'_B)$ for some $C'_A \in G_{S_a}$ and $C'_B \in G_{S_b}$. Because $H \,\square\, K = G_{S_a} \,\square\, G_{S_b}$, by the nature of Cartesian product, it follows that (C'_A, C_B) is also a coloring of G. But the color of v_a in C'_A is c_2 and the color of v_b in C_B is also c_2, contradicting our assumptions that C is a proper coloring and that $v_a \sim v_b$. This completes the proof of the claim and of the theorem. □

4.3 Matroid independent set reconfiguration graphs

As we have observed in Section 2.3, (Q1) is false and (Q3) is vacuously true because all matroid independent set graphs are connected. The answer to (Q2) is affirmative because Theorem 3.6 applies immediately as evidenced by the concept of direct sum of two matroids. We next show that if $H \,\square\, K$ is an independent set graph, then so are H and K.

Lemma 4.11. *Let $\mathcal{M}$ be the reconfiguration space of matroid independent set graphs. Then $\mathcal{M}$ is closed under taking subsequences.*

Proof. Let S be a reconfiguration set in $\mathcal{I}$ and let k be the the common length of the sequences in S, i.e., the number of elements in the ground set $E = \{1, 2, \dots, k\}$ of the corresponding matroid $M = (E, \mathcal{I})$. Let $E' = \{i_1, i_2, \cdots, i_\ell\}$ be a subset of $\{1, \dots, k\}$. Let $\mathcal{I}' = \{\mathcal{I} \cap \mathcal{E}' \colon \mathcal{I} \in \mathcal{I}\}$. Then the pair $M' = (E', \mathcal{I}')$ is known to be a matroid, called the restriction of M to E'. The set of sequences that corresponds to the independent sets of M is clearly $S|_{E'}$ and thus $S|_{E'} \in \mathcal{I}$. □

Theorem 4.12. *If $H \,\square\, K$ is a matroid independent set reconfiguration graph, then H and K are matroid independent set reconfiguration graphs.*

Proof. If $H \,\square\, K$ is an independent set reconfiguration graph than it is connected. Therefore, Lemma 4.11 and Theorem 3.11 imply that H and K are both independent set reconfiguration graphs. □

References

[1] J. Asplund, K. Edoh, R. Haas, Y. Hristova, B. Novick and B. Werner, Reconfiguration graphs of shortest paths, *Discrete Math.* 341 (10) (2018), 2938–2948.

[2] L. Beaudin, J. Ellis-Monaghan, G. Pangborn and R. Shrock, A little statistical mechanics for the graph theorist, *Discrete Math.* 310 (13-14) (2010), 2037–2053.

[3] J. Beier, J. Fierson, R. Haas, H. M. Russell and K. Shavo, Classifying coloring graphs, *Discrete Math.* 339 (8) (2016), 2100–2112.

[4] M. Bonamy, M. Johnson, I. Lignos, V. Patel and D. Paulusma, On the diameter of reconfiguration graphs for vertex colourings, *Electron. Notes Discrete Math.* 38 (2011), 161–166.

[5] M. Bonamy, M. Johnson, I. Lignos, V. Patel and D. Paulusma, Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs, *J. Comb. Optim.* 27 (1) (2014), 132–143.

[6] P. Bonsma, The complexity of rerouting shortest paths, *Theor. Comput. Sci.* 510 (2013), 1–12.

[7] P. Bonsma and L. Cereceda, Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances, *Theor. Comput. Sci.* 410 (50) (2009), 5215–5226.

[8] L. Cereceda, J. van den Heuvel and M. Johnson, Connectedness of the graph of vertex-colourings, *Discrete Math.* 308 (5-6) (2008), 913–919.

[9] L. Cereceda, J. van den Heuvel and M. Johnson, Mixing 3-colourings in bipartite graphs, *European J. Combin.* 30 (7) (2009), 1593–1606.

[10] L. Cereceda, J. van den Heuvel and M. Johnson, Finding paths between 3-colorings, *J. Graph Theory* 67 (1) (2011), 69–82.

[11] F. R. Chung, Pebbling in hypercubes, *SIAM J. Discrete Math.* 2 (4) (1989), 467–472.

[12] P. Gopalan, P. G. Kolaitis, E. Maneva and C. H. Papadimitriou, The connectivity of Boolean satisfiability: computational and structural dichotomies, *SIAM J. Comput.* 38 (6) (2009), 2330–2355.

[13] R. Haas and K. Seyffarth, The k-dominating graph, *Graphs Combin.* 30 (3) (2014), 609–617.

[14] T. R. Hartinger and M. Milanič, 1-perfectly orientable graphs and graph products, *Discrete Math.* 340 (7) (2017), 1727–1737.

[15] R. A. Hearn and E. D. Demaine, PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation, *Theor. Comput. Sci.* 343 (1-2) (2005), 72–96.

[16] T. Ito, E. D. Demaine, N. J. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara and Y. Uno, On the complexity of reconfiguration problems, *Theor. Comput. Sci.* 412 (12-14) (2011), 1054–1065.

[17] T. Ito, M. Kaminski and E. D. Demaine, Reconfiguration of list edge-colorings in a graph, *Discrete Appl. Math.* 160 (15) (2012), 2199–2207.

[18] M. Jerrum, A very simple algorithm for estimating the number of k-colorings of a low-degree graph, *Random Structures Algorithms* 7 (2) (1995), 157–165.

[19] M. Kamiński, P. Medvedev and M. Milanič, Shortest paths between shortest paths, *Theor. Comput. Sci.* 412 (39) (2011), 5205–5210.

[20] M. Kamiński, P. Medvedev and M. Milanič, Complexity of independent set reconfigurability problems, *Theor. Comput. Sci.* 439 (2012), 9–15.

[21] S. B. Maurer, Matroid basis graphs I, *J. Combin. Theory Ser. B* 14 (3) (1973), 216–240.

[22] S. B. Maurer, Matroid basis graphs II, *J. Combin. Theory Ser. B* 15 (2) (1973), 121–145.

[23] M. Molloy, The Glauber dynamics on colorings of a graph with high girth and maximum degree, *SIAM J. Comput.* 33 (3) (2004), 721–737.

[24] A. E. Mouawad, N. Nishimura, V. Raman, N. Simjour and A. Suzuki, On the parameterized complexity of reconfiguration problems, *Algorithmica* 78 (1) (2017), 274–297.

[25] H. M. Mulder, The interval function of a graph, *Mathematical Centre Tracts* 132, Mathematisch Centrum, Amsterdam, 1980.

[26] J. G. Oxley, Matroid theory, Second Ed., *Oxford Graduate Texts in Mathematics* 21, Oxford University Press, Oxford, 2011.

[27] E. Vigoda, Improved bounds for sampling colorings, *J. Math. Phys.* 41 (3) (2000), 1555–1569.

[28] D. B. West, *Introduction to Graph Theory*, Prentice Hall, Inc., Upper Saddle River, NJ, 1996.

(Received 26 May 2020; revised 6 Sep 2021)