Homomorphism Reconfiguration via Homotopy

Marcin Wrochna

Faculty of Mathematics, Informatics and Mechanics,
University of Warsaw, Warsaw, Poland
`m.wrochna@mimuw.edu.pl`

---- **Abstract** --

We consider the following problem for a fixed graph H: given a graph G and two H-colorings of G, i.e. homomorphisms from G to H, can one be transformed into the other by changing one color at a time, maintaining an H-coloring throughout. This is the same as finding a path in the $\mathrm{Hom}(G, H)$ complex. For $H = K_k$ this is the problem of finding paths between k-colorings, which was recently shown to be in P for $k \leq 3$ and PSPACE-complete otherwise (Bonsma and Cereceda 2009, Cereceda et al. 2011). We generalize the positive side of this dichotomy by providing an algorithm that solves the problem in polynomial time for any H with no C_4 subgraph. This gives a large class of constraints for which finding solutions to the Constraint Satisfaction Problem is NP-complete, but paths in the solution space can be found in polynomial time.

The algorithm uses a characterization of possible reconfiguration sequences (that is, paths in $\mathrm{Hom}(G, H)$), whose main part is a purely topological condition described in terms of the fundamental groupoid of H seen as a topological space.

1 Introduction

Reconfiguration

Reconfiguration is a framework in which we study how discrete structures, constrained in various ways, can be carefully transformed with small steps. This is often best described by finding paths in a solution graph, whose vertices are all solutions to a combinatorial problem and whose edges define the steps between solutions one is allowed to make.

For example, in k-Recoloring [3, 4, 8, 17], one is given two proper k-colorings of a graph G and the question is whether one can be transformed into the other by changing one color at a time, maintaining a proper coloring throughout. In other words, the solution graph has proper k-colorings as vertices (solutions) and edges (reconfigurations steps) between any two colorings that differ only at one vertex of G. Another well studied example is Token Jumping [12, 18], where the solutions are independent sets of some given size (seen as tokens on the graph's vertices) and a reconfiguration step removes one vertex from the set to add another (jumps one token to a different vertex). Yet another example is the reconfiguration of generalized SAT problems [11, 21, 24], where solutions are satisfying assignments of a given formula, and a reconfiguration step flips one variable of the assignment.
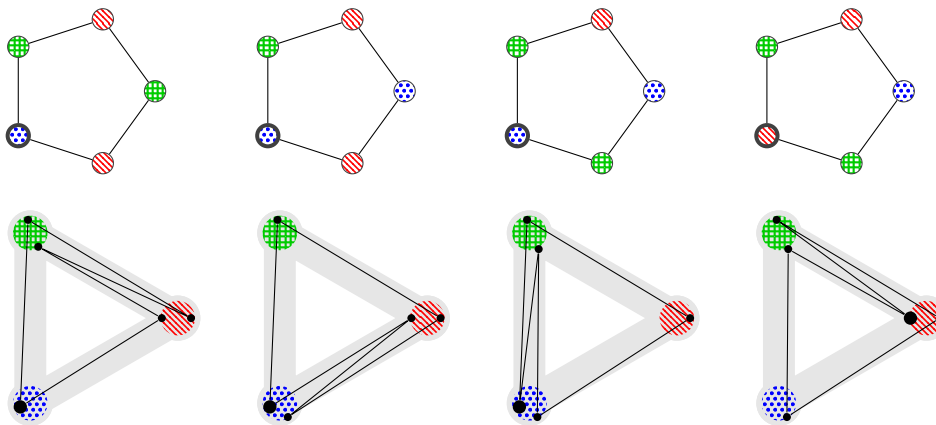
Figure 1 A sequence of 3-colorings of C_5 and the same sequence seen as mappings from C_5 to K_3 (a graph with three vertices: red, green, blue), that is, a path in $\mathrm{Hom}_1(C_5, K_3)$. One vertex of C_5 is thickened for clarity.

Homomorphisms

A homomorphism from a graph G to a graph H is a mapping $\sigma : V(G) \to V(H)$ such that edges are mapped to edges, that is, $uv \in E(G)$ implies $\sigma(u)\sigma(v) \in E(H)$. We also use the name *H-coloring*, especially when H is fixed in the context. Vertices of H are then called *colors*. Note that a K_k-coloring (so $H = K_k$ is the graph with all edges except loops) is the same as a (proper) k-coloring.

The solution graph $\mathrm{Hom}_1(G, H)$ is defined to be the graph with H-colorings of G as vertices and edges between any two H-colorings that differ in the color of only one vertex (Figure 1). For a fixed graph H, H-RECOLORING is the problem where given a graph G and two H-colorings of G we are asked whether they are connected by a path in $\mathrm{Hom}_1(G, H)$. SHORTEST H-RECOLORING asks whether there is such a path of at most some given length.

We use graph homomorphisms as a tool to explore how different constraints influence the complexity of reconfiguration. As our aim is to give more general statements about reconfiguration, they should be seen as a special case of Constraint Satisfaction Problems (CSPs), which can express a range of problems including k-COLORING, generalized SAT or INDEPENDENT SET (in weighted variants). However, graph homomorphisms already display many features of general CSPs and arise naturally in various situations. See [22] for an excellent survey and [16] for an in-depth book on the subject.

This approach allowed us to argue in [26] that the only notion of sparseness that can apply to (unparameterized) reconfiguration problems in general is treedepth. The reduction showed there explains why reconfiguration variants of easy combinatorial problems can be hard. In this paper we explore why reconfiguration variants of hard problems can be easy.

Motivations

The primary motivation for studying reconfiguration problems is to investigate the solution space of combinatorial problems, especially from the perspective of local search heuristics and random solution sampling. In particular, the success of Survey Propagation as a method for solving random Constraint Satisfaction Problems (CSPs) is connected to several conjectures about the structure of clusters of satisfying assignments (connected components in the solution graph) and frozen variables (variables/vertices that cannot change their value/color by any sequence of steps), see [1].

While finding paths in the solution graph is, for the above purposes, mostly a toy problem, it arises more directly in some settings. For example, the construction of Hearn and Demaine [12] allowed to show that many popular puzzles are PSPACE-complete [13, 20], and more interestingly, so is the problem of proof equivalence in a certain proof system [14], which answered an earlier question about normal forms of proofs.

Related work

H-Recoloring has been shown PSPACE-complete for $H = K_k$ where $k \geq 4$ by Bonsma and Cereceda [3], and in P for $k \leq 3$ by Cereceda et al. [8]. The latter result has been extended to show that Shortest K_3-Recoloring is also in P by Johnson et al. [17].

For CSPs in the Boolean domain, a dichotomy was shown by Gopalan et al. [11] – for a fixed set of Boolean constraints Γ (i.e., Boolean relations, clause types), the problem of finding paths in the solution graph of a SAT(Γ) instance is either in P or PSPACE-complete. In particular it is always in P when the corresponding satisfiability problem is in P (e.g. 2-SAT or Horn-SAT), but it is also in P for some Γ for which satisfiability is NP-complete. This was slightly corrected and extended to several similar problems by Schwerdtfeger [24], while a trichotomy was shown for the problem of finding *shortest* paths by Mouawad et al. [21]. Both [11] and [24] asked whether their results could be extended to larger domains. Our work can be seen as a step in this direction, but limited to only one symmetric relation of arity 2.

The corresponding extension of Schaefer's dichotomy [23] for satisfiability (deciding the existence of a solution) to CSPs with arbitrary finite domains is a long-standing open problem stated by Feder and Vardi [10]. They showed that the conjecture is equivalent when limited to one relation of arity 2 (digraph homomorphisms). Hell and Nešetřil proved the dichotomy in the case the relation is additionally assumed to be symmetric (graph homomorphism) [15]: the problem of deciding the existence of an H-coloring of a given graph is in P for H bipartite or containing a loop, and NP-complete otherwise.

Results

It is natural to ask whether the unexpected tractability of K_3-Recoloring (in light of 3-Colorability being NP-complete) is caused by the following property: whenever a vertex changes its color in a 3-coloring (e.g. from red to green), all of its neighbors must have one common color (blue).

We answer this in the positive considering the following definition: a graph H has the *monochromatic neighborhood property* if for every two colors $a, b \in V(H)$, the set of common neighbors $N_H(a) \cap N_H(b)$ contains at most one color. For graphs without loops this is equivalent to not containing C_4 (the cycle on 4 vertices) as a subgraph (not necessarily induced). This includes K_3 and all graphs of girth ≥ 5, for example. For graphs with loops allowed, this is equivalent to not containing C_4, K_3 with one loop added, nor K_2 with both loops added.

We show an algorithm that solves Shortest H-Recoloring in polynomial time for all H with the monochromatic neighborhood property. To achieve this, we characterize possible paths in $\mathrm{Hom}_1(G, H)$ by describing sequences of colors one vertex of G takes throughout an H-recoloring. We observe that the H-colorings of G correspond to continuous maps from G to H (seen as topological spaces) and that recoloring corresponds to a continuous transformation of these maps – a homotopy. This gives a topological condition on how a path in $\mathrm{Hom}_1(G, H)$ can look like. It turns out that to give a complete characterization we only need to add a simple parity condition and a condition freezing some easily found

vertices. Thus a combinatorial problem is reduced to describing possible homotopies, which we do with standard algebraic calculations. This gives an unexpected connection that might be interesting on its own.

In combinatorial algebraic topology

Reconfiguration of homomorphisms has already been studied independently in the field of combinatorial algebraic topology, though from a different angle. The notion of $\times$-homotopy of homomorphisms as defined by Dochtermann [9] is identical to reachability in the solution graph $\mathrm{Hom}_1(G, H)$. The solution graph $\mathrm{Hom}_1(G, H)$ arises as the subgraph of the exponential graph H^G induced by looped vertices and as the 1-skeleton of the Hom-complex $\mathrm{Hom}(G, H)$, see [2]. The Hom-complex, studied for its interesting categorical and topological properties, is a construction similar (and homotopy equivalent) to the clique complex of $\mathrm{Hom}_1(G, H)$. These definitions were introduced to provide lower bounds on the chromatic number of graphs, a notoriously hard problem. A typical theorem derived from such methods is that for loopless graphs G, H, if $\mathrm{Hom}_1(G, H)$ is connected for all G of degree at most d, then the chromatic number of H is at least $d/2$ (and is conjectured to be at least d) [5]. Studies have thus been mostly concerned with highly regular graphs for which the solution graph can be proved to be in some sense tightly connected.

We do not know of any prior work on the computational complexity of deciding $\times$-homotopy. More surprisingly, we do not know of any prior example where $\times$-homotopy (of homomorphisms) would be related to homotopy (of continuous mappings between graphs), more than through analogy.

A longer version of this paper, with full proofs, is available at `http://arxiv.org/abs/1408.2812`.

2 Preliminaries

An (undirected) graph G is a pair $(V(G), E(G))$ where $V(G)$ is a finite set of *vertices*, while $E(G)$ is the set of *edges* – unordered vertex pairs $\{u, v\}$, $u, v \in V(G)$, written uv for short. A *loop* is an edge uu for some vertex u. The neighborhood $N_G(v)$ is defined as $\{w \in V(G) \mid vw \in E(G)\}$. Hence $a \in N_H(a)$ iff H has a loop at a.

G and H in this paper are always connected undirected graphs with at least one edge. G is always assumed to have no loops. H can have loops, but is always assumed to have the monochromatic neighborhood property.

An *H-recoloring sequence* or *reconfiguration sequence* is a path in $\mathrm{Hom}_1(G, H)$, that is, a sequence of H-colorings of G in which consecutive colorings differ at one vertex. Consider a step of an H-recoloring sequence – a vertex $v \in V(G)$ changes color from $a \in V(H)$ to $b \in V(H)$. Since G is connected, loopless and has an edge, v has a neighbor, say $w \neq v$. As only v changes its color in the step, w has the same color, say $h \in V(H)$, before and after the step. The H-coloring before the step implies that $ha \in E(H)$, while the one after the step implies that $hb \in E(H)$. Thus $h \in N_H(a) \cap N_H(b)$. From the monochromatic neighborhood property we infer that $N_H(a) \cap N_H(b) = \{h\}$. We will often call h '*the color that all neighbors of v have during the step*' (that is, in the H-colorings before and after the step), without arguing its existence and uniqueness anymore.
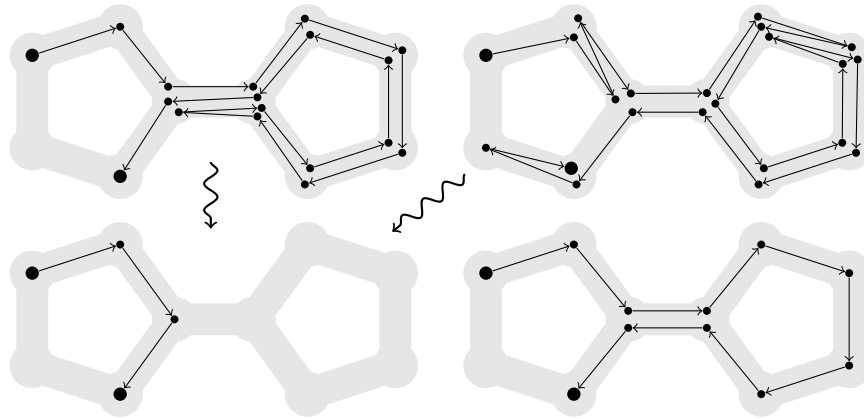
Figure 2 Examples of two walks (in a graph H on 10 vertices) which reduce to the same, bottom left one. The bottom right one is a different reduced walk; when its endpoints are fixed, it cannot be distorted as a curve to give any of the others.
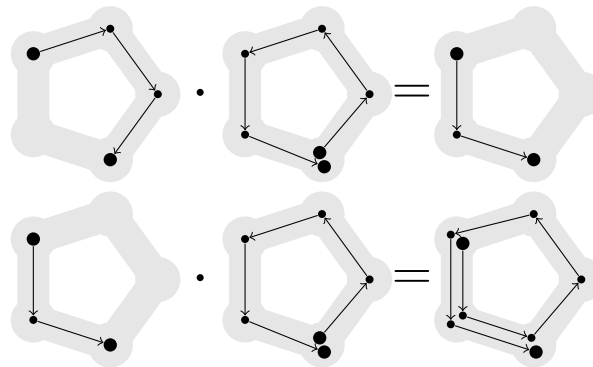
3 Fundamental groupoid

An *oriented edge* of a graph H is an ordered pair $e = (u, v)$ such that $\{u, v\}$ is an edge of H; we denote its initial vertex u as $\iota(e)$ and its target vertex v as $\tau(e)$. We write e^{-1} for $(\tau(e), \iota(e))$. A *walk* from u to v in a graph H is a sequence of oriented edges $e_1 e_2 \dots e_l$ of H such that $\iota(e_1) = u$, $\tau(e_l) = v$ and $\tau(e_i) = \iota(e_{i+1})$ for $i = 0, \dots, l - 1$. We write ε for an empty walk (the endpoints of ε will be irrelevant for us). The length of a walk is the number of edges in it. A walk W_1 from u to v can be concatenated to a walk W_2 from v to w to form a walk $W_1 W_2$ from u to w.

We call a walk *reduced* if it contains no two consecutive edges $e_i e_{i+1}$ such that $e_{i+1} = e_i^{-1}$. One can *reduce* a walk by removing any such two consecutive edges from the sequence. It can easily be seen that by iteratively reducing a walk W, one always gets the same reduced walk, which we therefore denote as $\overline{W}$, see Figure 2. For any two reduced walks $\overline{W_1}, \overline{W_2}$ such that W_2 starts where W_1 ends, we write $\overline{W_1} \cdot \overline{W_2}$ for $\overline{\overline{W_1}\, \overline{W_2}}$ and similarly one can observe that $\cdot$ is associative. For any walk $W = e_1 e_2 \dots e_l$ we write W^{-1} for the reversed walk $e_l^{-1} \dots e_2^{-1} e_1^{-1}$. Clearly $\overline{W} \cdot \overline{W}^{-1} = \varepsilon = \overline{W}^{-1} \cdot \overline{W}$ and $\varepsilon \cdot \overline{W} = \overline{W} \cdot \varepsilon = \overline{W}$. Therefore, the set of reduced walks of a graph forms together with the operations $\cdot$ and $()^{-1}$ a *groupoid* – a group, except for the binary operation being a partial function. This particular groupoid is called the *fundamental groupoid* $\pi(H)$ of H, see Figure 3.

A groupoid is in many ways similar to a group. Identities such as $(e \cdot f)^{-1} = f^{-1} \cdot e^{-1}$ known from group theory are easily reproved in groupoids. Using the fundamental groupoid as opposed to the better known fundamental group allows us to describe calculations much more uniformly, without the tedious change of base points, for example. Readers familiar with category theory may benefit from the view given by equivalent definitions: a groupoid is a category in which every morphism is invertible, and a group is a groupoid with one object. A brief survey on groupoids and their application in topology [6] and a more complete exposition [7] have been written by R. Brown.

Note that if H has a loop at u, then $(u, u) = (u, u)^{-1}$ is an oriented edge and a reduced walk of length 1. On the other hand, $(u, u) \cdot (u, u) = \varepsilon$. (Hence $\pi(H)$ is strictly speaking not a *free groupoid*, see [19]).

Figure 3 Examples of · multiplication in the fundamental groupoid of $H = C_5$.

Topological interpretation

Let us comment on how this algebraic structure captures the topology of curves in the graph. When referring to topology, we give only informal interpretations, as the statements are not needed in any of the proofs.

A graph H without loops can be naturally associated with a topological space, constructed from copies of the $[0, 1]$ interval for each edge, with endpoints merged into vertices accordingly. A *curve* in this space is a continuous map $f : [0, 1] \to H$ (where H is meant as the topological space), not necessarily injective. Two curves f_0, f_1 are *homotopic* if one can be continuously transformed into the other, which means there is a set of functions $\{\phi_t : t \in [0, 1]\}$ such that $\phi_0 = f_0, \phi_1 = f_1$ and the mapping $\phi : (t, x) \mapsto \phi_t(x)$ is continuous as a function from $[0, 1] \times [0, 1]$ to H.

The fundamental groupoid fully describes curves up to homotopy. For any two vertices u, v of H, two curves f_0, f_1 are homotopic via a homotopy ϕ_t that fixes the endpoints ($\phi_t(0) = u, \phi_t(1) = v$ for all t) if and only if the corresponding reduced walks in $\pi(H)$ are equal. Considering only reduced walks that start and end in the same vertex v, we obtain a group, which is known as the fundamental group $\pi_1(H, v)$ of the topological space H. When no vertex is fixed, a closed curve is homotopic to another, via a homotopy such that $\phi_t(1) = \phi_t(0)$ for all t, if and only if the corresponding elements C_1, C_2 of $\pi(H)$ are *conjugate*, i.e. $C_2 = P^{-1} \cdot C_1 \cdot P$ for some $P \in \pi(H)$.

4 Vertex walks and realizability

Consider an H-recoloring sequence $S = \sigma_0, \dots, \sigma_l$ of G and any vertex $v \in V(G)$. We define $S(v)$ as the following walk in H:

- If $l = 0$, that is, S is empty, then $S(v) = \varepsilon$.
- If $l = 1$, that is, S contains a single reconfiguration step, then
  $$S(v) = \varepsilon \qquad \text{when } \sigma_0(v) = \sigma_1(v), \text{ and}$$
  $$S(v) = (\sigma_0(v), h)(h, \sigma_1(v)) \qquad \text{otherwise,}$$
  where h is the color that all neighbors of v have in σ_0 and σ_1.
- If $l > 1$, then $S(v) = S_1(v)S_2(v)\dots S_l(v)$, where $S_i(v)$ is the walk corresponding to the reconfiguration step from σ_i to σ_{i+1}.

For two H-colorings α, β of G and a vertex $q \in V(G)$, we call an element $Q \in \pi(H)$ *realizable* if there is an H-recoloring sequence $S = \sigma_0, \dots, \sigma_l$ such that $\sigma_0 = \alpha, \sigma_l = \beta$ and $\overline{S(q)} = Q$. We focus on the following question: which $Q \in \pi(H)$ are realizable? It is
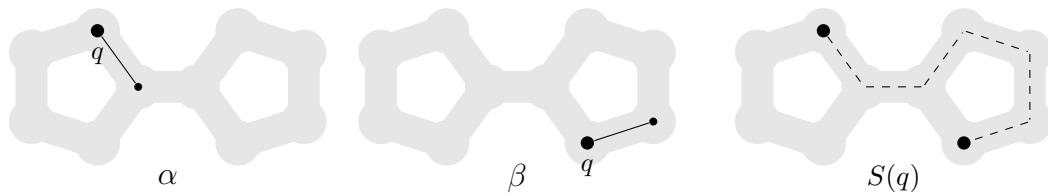
Figure 4 A realizable walk for $\alpha, \beta : K_2 \to H$ and q. Note the shortest walk from $\alpha(q)$ to $\beta(q)$ (of length 3) is not realizable because of parity.

immediate from the definition that Q must be a reduced walk from $\alpha(q)$ to $\beta(q)$ and have even length (notice that the parity of the length of walks is preserved by reducing), see Figure 4. This parity condition will be one of three conditions characterizing realizable walks.

If $W = (v_1, v_2)(v_2, v_3) \ldots (v_{n-1}, v_n)$ is a walk in G and α is an H-coloring of G, then $\alpha(W) = (\alpha(v_1), \alpha(v_2)) \ldots (\alpha(v_{n-1}), \alpha(v_n))$ is a walk in H.

5 Topological validity

With a homomorphism from G to H one can associate a continuous map from G to H (seen as topological spaces as described above). The following lemma is the key to the monochromatic neighborhood property of H.

▶ **Lemma 1.** *Let S be an H-recoloring sequence of G from α to β. Consider any walk W from vertex u to v in G. Then $\overline{S(v)} = \overline{\alpha(W)}^{-1} \cdot \overline{S(u)} \cdot \overline{\beta(W)}$.*

The proof follows directly from the monochromatic neighborhood property and induction. The statement can be understood as an algebraic expression of the fact that when an H-coloring is reconfigured into another, the corresponding continuous mappings can be continuously transformed into one another, and the walk $S(v)$ of each vertex v is (up to reduction) the curve traced by the point v (see Figure 5). Let us focus on corollaries.

First, we see that in a given instance of H-Recoloring, the reduced vertex walk $\overline{S(q)}$ of one vertex q in a solution S determines all other vertex walks up to reductions (in other words, up to homotopy). This is why we can focus on the realizability of one element $Q \in \pi(H)$ instead of a whole recoloring sequence. If we decide that Q is realizable, we will later use this lemma to completely recover a recoloring sequence S such that $\overline{S(q)} = Q$.

Second, note that the statement doesn't depend on how we choose walks in G. This means that for every pair of walks in G with equal endpoints, there is a topological condition on how solutions look like.

▶ **Definition 2.** *Let α, β be two H-colorings of G and let q be a vertex of G. Let $Q \in \pi(H)$. We say Q is topologically valid for α, β, q if for every vertex w and every two walks W_1, W_2 from q to w in G we have $\overline{\alpha(W_1)}^{-1} \cdot Q \cdot \overline{\beta(W_1)} = \overline{\alpha(W_2)}^{-1} \cdot Q \cdot \overline{\beta(W_2)}$.*

▶ **Corollary 3.** *Let α, β be two H-colorings of G and let q be a vertex of G. If $Q \in \pi(H)$ is realizable for α, β, q then Q is topologically valid for α, β, q.*

Using the groupoid structure, this condition can easily be rephrased in terms of single closed walks, instead of pairs of walks with equal starts and equal ends.

▶ **Lemma 4.** *Let α, β be two H-colorings of G and let q be a vertex of G. Then $Q \in \pi(H)$ is* topologically valid *for α, β, q if and only if for every closed walk C from q to q we have $\overline{\beta(C)} = Q^{-1} \cdot \overline{\alpha(C)} \cdot Q$.*
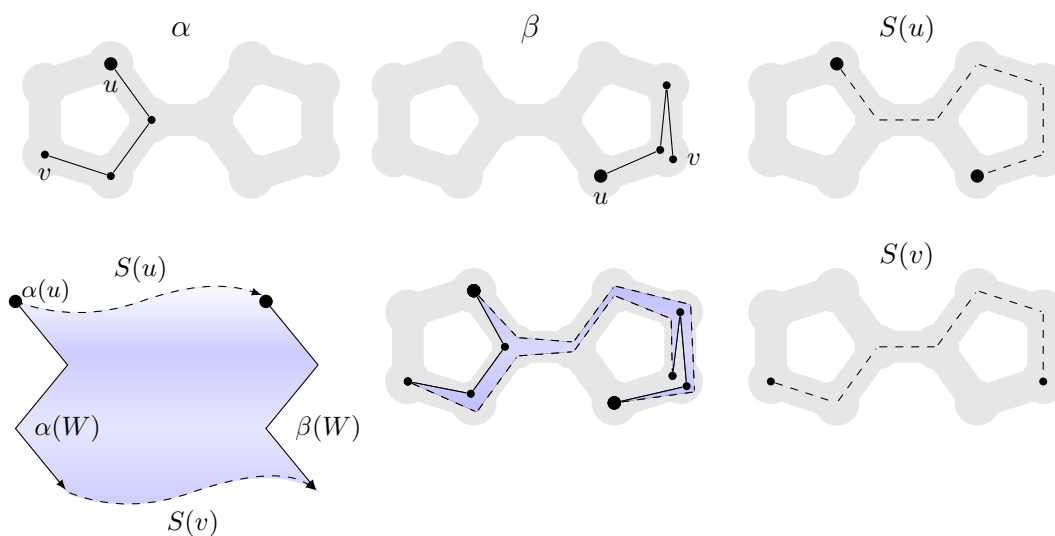
Figure 5 If α can be transformed to β by reconfiguration, then it can also be transformed by a homotopy. Restricting this homotopy to a walk W from u to v in G gives a map $\phi : [0,1] \times [0,|W|] \to H$ such that $\phi(0, \cdot) = \alpha(W)$ and $\phi(1, \cdot) = \beta(W)$, $\phi(\cdot, 0) = \overline{S(u)}$ and $\phi(\cdot, |W|) = \overline{S(v)}$. Since ϕ is a continuous mapping of a rectangle to H, it's boundary, and so the image of it's boundary, can be contracted to a point. This is the meaning of the equality $\overline{\alpha(W)}^{-1} \cdot \overline{S(u)} \cdot \overline{\beta(W)} \cdot \overline{S(v)}^{-1} = \varepsilon$.

The name is motivated by the following fact: Q is *topologically valid* for α, β, q if and only if there is a homotopy continuously transforming α to β such that q traces the curve Q throughout this transformation ($\phi_0 = \alpha$, $\phi_1 = \beta$ and the image of $t \mapsto \phi_t(q)$ is Q).

We note that the existence of a topologically valid walk for α, β implies in particular that cycles have the same homotopy class in α as in β, which algebraically is expressed in the fact that they have the same conjugacy class: $\overline{\beta(C)} = Q^{-1} \cdot \overline{\alpha(C)} \cdot Q$ for some $Q \in \pi(H)$. This generalizes one of the conditions of a characterization of 3-colorings connected by recoloring sequences given in [8] (Theorem 7 (C2)).

Another implication is that, intuitively, for each cycle C in G such that $\overline{\alpha(C)} \neq \varepsilon$, two solutions can differ only in the number of times vertex walks wind around $\overline{\alpha(C)}$. This is analyzed in more detail in the full version of the paper, but we will see this reflected when describing the set of realizable walks (e.g. Figure 8).

6 Tight closed walks and frozen vertices

We have seen that any realizable walk must have even length and must be topologically valid. There is one more obstruction to reconfiguration: closed walks whose sequences of colors are in a sense tightly stretched around H. In an H-coloring α of G, a vertex v of G is called *frozen* if for every H-recoloring sequence from α the resulting H-coloring β has $\beta(v) = \alpha(v)$. A closed walk $C = e_1 e_2 \dots e_l$ is *cyclically reduced* if it is reduced and also $e_l \neq e_1^{-1}$. In other words, repeating C gives an infinite reduced walk. A closed walk C is *α-tight* if $\alpha(C)$ is cyclically reduced. We show that vertices on tight closed walks are frozen.

▶ **Lemma 5.** *Let α be an H-coloring of G and let C be an α-tight walk in G. Then all vertices of C are frozen in α.*
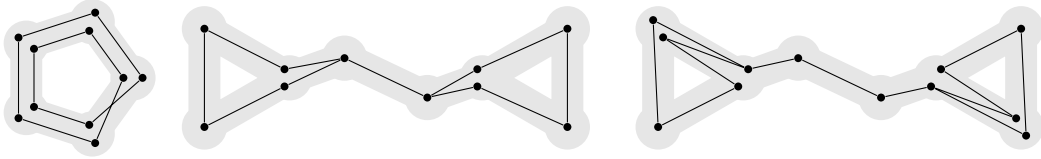
Figure 6 In the left example a C_5-coloring of C_{10} is shown. there is a tight closed walk visiting all 10 vertices of the cycle in order. Hence no reconfiguration step is possible. Similarly in the middle example, there is a tight closed walk going around one 5-cycle, along the bridge, around the other 5-cycle and back along the bridge. In the right example no closed walk is tight, but the 4 middle vertices are frozen.

Tight closed walks can be found by starting from some vertex and exploring walks W such that $\alpha(W)$ is reduced. If they are arbitrarily long, then they contain the same oriented edge twice, giving an α-tight closed walk. We can answer whether an oriented edge is reachable from itself by a non-trivial reduced walk in time $\mathcal{O}(|E(G)|)$. The (potentially infinite) prefix tree of reduced walks starting from one vertex gives a generalization of the layer construction of [8] used to characterize frozen vertices when $H = K_3$. For other H, frozen vertices can arise in other situations (Figure 6), but it turns out we won't need to identify them.

Finding a frozen vertex v allows us to limit potentially realizable walks Q to a single one. Since $S(v) = \varepsilon$, Lemma 1 implies $Q = \overline{S(q)} = \overline{\alpha(W)}^{-1} \cdot \overline{S(v)} \cdot \overline{\beta(W)} = \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)}$.

▶ **Corollary 6.** *Let $Q \in \pi(H)$ be realizable for α, β – H-colorings of G and $q \in V(G)$. Let C be an α-tight closed walk in H. Then for every vertex v on C and every walk W from v to q we have $Q = \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)}$.*

7 Characterization of realizable walks

In the following theorem we show that there are no more conditions for a walk to be realizable than the three we described: even length, topological validity and frozenness of tight closed walks. This is very unexpected – the fact that edges are actually discrete and cannot be stretched arbitrarily, for example, turns out to imply no further obstructions to reconfiguration (it only restricts the input H-colorings).

▶ **Theorem 7.** *Let α, β be two H-colorings of G. Consider any vertex q of G and let $Q \in \pi(H)$ be a reduced walk in H from $\alpha(q)$ to $\beta(q)$. Then Q is realizable for α, β, q if and only if*
- *Q is topologically valid for α, β, q,*
- *Q has even length,*
- *if there is an α-tight walk, then for any walk W from it to q, $Q = \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)}$.*
Furthermore, there is an algorithm working in time $\mathcal{O}(|V(G)|^2 + |V(G)| \cdot |Q| + |E(H)|)$ that given G, H, α, β and a walk Q, checks whether Q satisfies the above conditions and if so, outputs a recoloring sequence such that $S(q) = Q$ and $S(v)$ is reduced for all $v \in V(G)$.

See Figure 7 for an example on how the conditions apply. While proving that realizability implies the conditions is easy, the proof of the converse spans 2 pages. The idea is that Lemma 1 gives vertex walk in a consistent way thanks to the first condition, they correspond to recoloring steps thanks to the second condition, and those can be ordered to give a recoloring sequence thanks to the third condition. In particular, the proof gives an effective way to obtain from a realizable $Q \in \pi(H)$ a recoloring sequence in which every vertex walk is reduced.
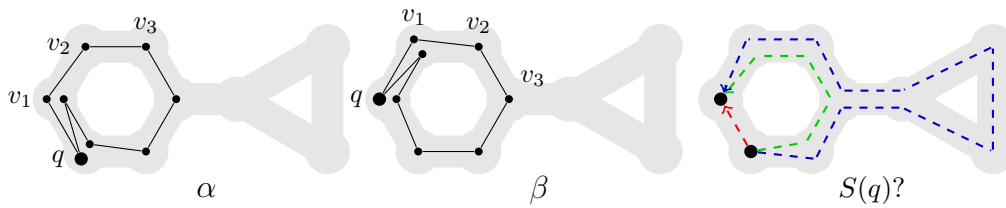
Figure 7 Even though no vertex is frozen, α cannot be reconfigured to β. The red and green walks are not realizable for α, β, q because of parity. The dark blue walk has good parity, but is not topologically valid (imagine pulling q along it).

▶ **Corollary 8.** *Let α, β be two H-colorings of G. Let $S = \sigma_0, \ldots, \sigma_l$ be a shortest H-recoloring sequence between $\sigma_0 = \alpha$ and $\sigma_l = \beta$. Then $S(v)$ is reduced for each vertex v.*

Proof. Suppose $S(q)$ is not reduced for some q. Let $Q = \overline{S(q)}$. By the above theorem we know from one side that Q is realizable. From the other side we obtain a solution S' such that $S'(v) = \overline{S'(v)}$ for all v and $S'(q) = Q$. By Lemma 1, this implies $S'(v) = \overline{S(v)}$, for all v. But $\overline{S(v)}$ is always no longer than $S(v)$, and $\overline{S(q)}$ is strictly shorter than $S(q)$. Since the number of recoloring steps is equal to half the sum of lengths of all $S(v)$, S was not shortest. ◀

8 An algorithm

The following lemma follows from well-known calculations in the fundamental groupoid, which we recall in the appendix of the full version of the paper.

▶ **Lemma 9.** *Let α, β be H-colorings of G and q a vertex of G. Consider the set $\Pi \subseteq \pi(H)$ of topologically valid walks for α, β, q. One of the following holds:*
0. $\Pi = \emptyset$.
1. $\Pi = \{Q\}$ *for some $Q \in \pi(H)$.*
2. $\Pi = \{R^n \cdot P \mid n \in \mathbb{Z}\}$ *for some $R, P \in \pi(H)$.*
3. Π *contains all reduced walks from $\alpha(q)$ to $\beta(q)$.*
*Furthermore, there is an algorithm working in time $\mathcal{O}(|E(G)| \cdot |V(G)| + |E(H)|)$ that given G, H, α, β, q outputs which case holds and outputs Q or R, P in cases 1, 2.*

Intuitively, the algorithm finds an appropriate set of at most $E(G)$ closed walks in G and reduce the problem to a system of equations given by Lemma 4 for those closed walks. Consider a closed walk C and what this walk maps to, up to homotopy. If this is non-trivial, that is, $\overline{\alpha(C)} \neq \varepsilon$, then one can show that two topologically valid walks may differ only in the number of times they wind around $\overline{\alpha(C)}$ (or the shortest root R such that $R^n = \overline{\alpha(C)}$), see Figure 8. We solve the system of equations to get one solution P, if it exists. If two closed walks wind around different roots, this implies P is the only solution. If all non-trivial roots are equal R, then case (2.) of the lemma holds. If all roots are trivial, then case (3.) holds. Thus the problem reduces to calculating and comparing roots of the chosen closed walks.

The remaining constraints for realizability don't change the picture much, as stated in the next theorem.

▶ **Theorem 10.** *Let α, β be H-colorings of G and q a vertex of G. Consider the set $\Pi' \subseteq \pi(H)$ of realizable walks for α, β, q. One of the following holds:*
0. $\Pi' = \emptyset$.

$$\alpha \qquad \beta \qquad R \qquad P \qquad R^{-1} \cdot P$$
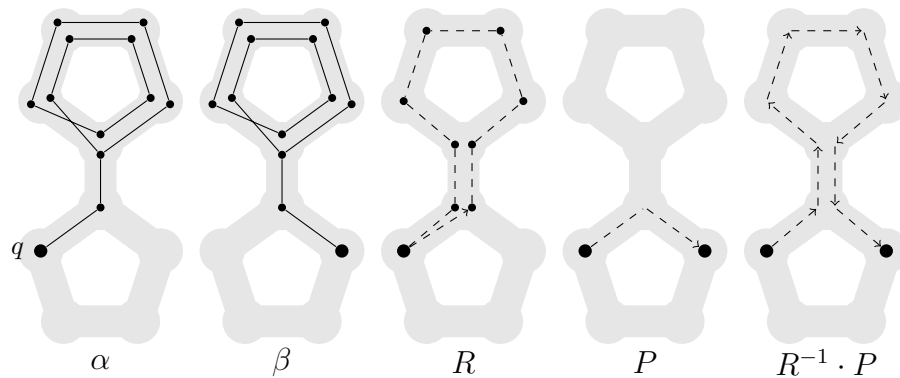
Figure 8 There is essentially only one closed walk C from q to q in this example, which winds around the root R $(\overline{\alpha(C)} = R^2)$. The topologically valid paths are exactly $\{R^n \cdot P \mid n \in \mathbb{Z}\}$ (think about deforming α by pulling q: one can pull it once or more around the top cycle by rotating all of α, but this is impossible for the bottom cycle if we are to reach β).

1. $\Pi' = \{Q\}$ *for some* $Q \in \pi(H)$.
2. $\Pi' = \{R^n \cdot P \mid n \in \mathbb{Z}\}$ *for some* $R, P \in \pi(H)$.
3. Π' *contains all reduced walks of even length from* $\alpha(q)$ *to* $\beta(q)$.
Furthermore, there is an algorithm working in time $\mathcal{O}(|E(G)|^2 + |E(H)|)$ *that given* α, β, q
outputs in polynomial time which case holds and outputs Q *or* R, P *in cases 1, 2.*

The algorithm of Theorem 10 simply runs the algorithm of Lemma 9 to handle the topological condition. To handle the parity condition it discards Q if it has odd length; replaces P with $R \cdot P$ if both R and P are odd; replaces R with R^2 if odd; returns an empty set if P is odd and R is even. To handle the tight closed walk condition, it searches for such walks as described before Corollary 6 and restricts the set to the single walk implied by this corollary, if it applies. To solve H-RECOLORING, that is, to decide whether the set of realizable walks is non-empty, we have to additionally check whether the set of even reduced walks from $\alpha(q)$ to $\beta(q)$ in H is empty. It is empty if and only if H is bipartite and $\alpha(q), \beta(q)$ are on different sides of a bipartition, which is easily checked in time $\mathcal{O}(|E(H)|)$.

▶ **Corollary 11.** *Let H be a graph (possibly with loops) with the monochromatic neighborhood property. Then H-RECOLORING is in P.*

Shortest recoloring sequences can also be found in polynomial time using Theorem 10, but this requires some more care, mostly in the case we need to find a reduced walk of even length that will correspond to a shortest solution.

▶ **Theorem 12.** *Let H be a graph (possibly with loops) with the monochromatic neighborhood property. Then SHORTEST H-RECOLORING is in P.*

9 Conclusions and future work

Our result generalizes the algorithm for K_3-RECOLORING of [8] and recovers many of its features in a more general and perhaps more intuitive setting. When limited to $H = K_3$, we may observe that there is essentially only one possible root R for closed walks in H (a 3-cycle), and all the reduced walks in H with the same endpoints differ only by powers of R. This can be used to show that either no walk is topologically valid (that is, α, β are not homotopic), or all are. In the latter case, it suffices to find tight cycles and either there is

one, immediately implying the vertex walks in the only 3-recoloring sequence; or there is none, in which case any vertex walk for one vertex can be realized in a solution, as long as it has even length. In other words, either no walk is realizable, or $Q = \varepsilon$ is realizable for some (frozen) vertex, or for any vertex q, the walk from $\alpha(q)$ to $\beta(q)$ of length 0 or 2 is realizable. In particular, we don't need to find frozen vertices or do any of the calculations in Lemma 9, it suffices to run $|V(G)| + 1$ times the simpler algorithm of Theorem 7.

We note that none of the proofs used any structural properties of H. If we consider H-recoloring for any graph H, but only allow recoloring a vertex if all of its neighbors have one common color (in other words, a reconfiguration step is allowed only when the homotopy class of the mapping doesn't change), the same results will follow.

An obvious question is how far can our results be extended to more general CSPs – to the asymmetric case, to multiple constraints, to hypergraphs (relations of arbitrary arity)? Is there any connection with the tractable cases of generalized SAT problems? Another question is whether the problems of graph homomorphism reconfiguration exhibit a dichotomy. For which graphs H is H-Recoloring in P or PSPACE-complete? Some basic reductions are given in the author's master thesis [25]. Finally, it could be interesting to explore the implications of the monochromatic neighborhood property for the whole Hom complex.

References

1   Dimitris Achlioptas, Amin Coja-Oghlan, and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. *Random Structures and Algorithms*, 38(3):251–268, 2011.

2   Eric Babson and Dmitry N. Kozlov. Complexes of graph homomorphisms. *Israel Journal of Mathematics*, 152(1):285–312, 2006.

3   Paul Bonsma and Luis Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theor. Comput. Sci.*, 410(50):5215–5226, 2009.

4   Paul Bonsma and Amer E. Mouawad. The complexity of bounded length graph recoloring. *arXiv*, 1404.0337, 2014.

5   Graham R Brightwell and Peter Winkler. Graph homomorphisms and long range action. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 63:29–48, 2004.

6   Ronald Brown. From groups to groupoids: a brief survey. *Bull. London Math. Soc*, 19(2):113–134, 1987.

7   Ronald Brown. Topology and groupoids. 2006.

8   Luis Cereceda, Jan van den Heuvel, and Matthew Johnson. Finding paths between 3-colorings. *Journal of Graph Theory*, 67(1):69–82, 2011.

9   Anton Dochtermann. Hom complexes and homotopy theory in the category of graphs. *European J. Combin.*, 30(2):490–509, 2009.

10 Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory. *SIAM Journal on Computing*, 28(1):57–104, 1998.

11   Parikshit Gopalan, Phokion G. Kolaitis, Elitza N. Maneva, and Christos H. Papadimitriou. The connectivity of boolean satisfiability: Computational and structural dichotomies. *SIAM Journal on Computing*, 38(6):2330–2355, 2009.

12   Robert A. Hearn and Erik D. Demaine. PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theoretical Computer Science*, 343(1-2):72–96, 2005.

13   Robert A. Hearn and Erik D. Demaine. *Games, puzzles and computation.* A K Peters, 2009.

14   Willem Heijltjes and Robin Houston. No proof nets for MLL with units: Proof equivalence in MLL is PSPACE-complete. In *Proceedings of the CSL-LICS 2014 Joint Meeting*, pages 50:1–50:10. ACM, 2014.

15   Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48(1):92–110, 1990.

16   Pavol Hell and Jaroslav Nešetřil. *Graphs and homomorphisms*, volume 28. Oxford University Press, Oxford, 2004.

17   Matthew Johnson, Dieter Kratsch, Stefan Kratsch, Viresh Patel, and Daniël Paulusma. Colouring reconfiguration is fixed-parameter tractable. *arXiv*, 1403.6347, 2014.

18   Marcin Kamiński, Paul Medvedev, and Martin Milanič. Complexity of independent set reconfigurability problems. *Theoretical Computer Science*, 439:9–15, 2012.

19   Jin Ho Kwak and Roman Nedela. Graphs and their coverings. `http://www.savbb.sk/~nedela/graphcov.pdf`, 2007.

20   Rahul Mehta. 2048 is (PSPACE) hard, but sometimes easy. *arXiv*, 1408.6315, 2014.

21   Amer E. Mouawad, Naomi Nishimura, Vinayak Pathak, and Venkatesh Raman. Shortest reconfiguration paths in the solution space of boolean formulas. *arXiv*, 1404.3801, 2014.

22   Jaroslav Nešetril. Homomorphisms of structures (concepts and highlights). *Physics and Theoretical Computer Science: From Numbers and Languages to (Quantum) Cryptography Security*, 7:295, 2007.

23   Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of STOC '78*, pages 216–226, New York, NY, USA, 1978. ACM.

24   Konrad W. Schwerdtfeger. A computational trichotomy for connectivity of boolean satisfiability. *arXiv*, 1312.4524, 2013.

25   Marcin Wrochna. *Homomorphism reconfiguration in general graphs*, chapter 4. 2014. Master thesis, `http://mimuw.edu.pl/~mw290715/thesis.pdf`.

26   Marcin Wrochna. Reconfiguration in bounded bandwidth and treedepth. *arXiv*, 1405.0847, 2014.