

Reconfiguration of Satisfying Assignments and Subset Sums: Easy to Find, Hard to Connect*

Jean Cardinal¹, Erik D. Demaine², David Eppstein³,
Robert A. Hearn⁴, and Andrew Winslow⁵

¹ Université libre de Bruxelles (ULB), Brussels, Belgium
jcardin@ulb.ac.be

² Massachusetts Institute of Technology, Cambridge, MA, USA
edemaine@mit.edu

³ University of California, Irvine, CA, USA
eppstein@ics.uci.edu

⁴ bob@hearn.to

⁵ University of Texas Rio Grande Valley, Edinburg, TX, USA
andrew.winslow@utrgv.edu

Abstract. We consider the computational complexity of *reconfiguration* problems, in which one is given two combinatorial configurations satisfying some constraints, and is asked to transform one into the other using elementary transformations, while satisfying the constraints at all times. Such problems appear naturally in many contexts, such as model checking, motion planning, enumeration and sampling, and recreational mathematics. We provide hardness results for problems in this family, in which the constraints and operations are particularly simple.

More precisely, we prove the PSPACE-completeness of the following decision problems:

- Given two satisfying assignments to a planar monotone instance of Not-All-Equal 3-SAT, can one assignment be transformed into the other by single variable “flips” (assignment changes), preserving satisfiability at every step?
- Given two subsets of a set S of integers with the same sum, can one subset be transformed into the other by adding or removing at most three elements of S at a time, such that the intermediate subsets also have the same sum?
- Given two points in $\{0,1\}^n$ contained in a polytope P specified by a constant number of linear inequalities, is there a path in the n -hypercube connecting the two points and contained in P ?

These problems can be interpreted as reconfiguration analogues of standard problems in NP. Interestingly, the instances of the NP problems that appear as input to the reconfiguration problems in our reductions can be shown to lie in P. In particular, the elements of S and the coefficients of the inequalities defining P can be restricted to have logarithmic bit-length.

Keywords: Boolean Satisfiability, Subset Sum, Combinatorial Reconfiguration, PSPACE-completeness

* A preprint of this paper containing three omitted proofs is available on arXiv: <https://arxiv.org/abs/1805.04055>

1 Introduction

Many computational problems consist of deciding the existence of a combinatorial object subject to constraints expressible in algebraic or logical terms. We consider *reconfiguration* problems, in which one is given two objects satisfying a set of constraints, and the goal is to transform one into the other using simple reconfiguration moves such that all the constraints remain satisfied at every intermediate step. Such problems find applications in dynamic environments or reactive systems, in which solutions are required or designed to evolve, in accessibility problems in model checking, as well as in enumeration and sampling problems, in which connectivity of the search space plays a major role.

We focus on reconfiguration problems that are naturally derived from standard NP-complete problems. This line of inquiry seems to have begun with the Sliding Tokens problem, a reconfiguration version of Independent Set, by Hearn and Demaine [11], and has gained momentum with publications such as the extension of Schaefer’s dichotomy to the connectivity of Boolean satisfiability due to Gopalan et al. [9], and an overview of the complexity of reconfiguration problems by Ito et al. [16]. In the canonical example of Boolean satisfiability, one is given two satisfying assignments to connect by a sequence of variable assignment flips, such that the formula remains satisfied at every step. The study of this type of question also benefits from the interest of puzzle designers and recreational mathematicians; token-sliding problems, for instance, are related to the famous 15-puzzle, popular in the late 19th century [29]. Combinatorial reconfiguration now constitutes a quickly developing field with dedicated research groups and workshops (such as the combinatorial reconfiguration workshop held in Banff in January 2017). For a more thorough survey and history of this family of problems, we refer to van den Heuvel [13].

Reconfiguration of independent sets in graphs is among the most studied problem in this vein (see [4, 6, 8, 14] for recent results) and is relevant to our findings. In these problems, one is given a graph G and two independent sets of G of the same size k , and the goal is to transform one into the other using elementary operations, preserving independence at every step. The operations consist either of “token slides”, in which a vertex in the independent set is replaced by one of its neighbors [11], or of vertex additions and removals such that the size of the independent set is either k or $k - 1$ [16]. A third, related, model is that of “token jumping”, in which a vertex is replaced by another, so that the size remains unchanged [20]. In general, these reconfiguration problems are known to be PSPACE-complete.

Reconfiguration problems for graph colorings followed, and have a large dedicated body of results as well [1, 3, 5, 7, 17]. Again, many such problems are known to be PSPACE-complete. Reconfiguration problems for shortest paths [2, 19], vertex covers [18], dominating sets [10], and Steiner trees [24] have also been considered.

As discussed by van den Heuvel [13], the question of the relation between the complexity of the existence problem (of a satisfying assignment, for instance) and that of the reconfiguration problem is intriguing. In many early examples,

reconfiguration problems in P are obtained from existence problems that are in P , and many PSPACE-hardness proofs follow the lines of the NP-hardness proof of the corresponding satisfiability problem. In the Schaefer-type dichotomy theorem established by Gopalan et al. [9], all satisfiability problems in P yield a reconfiguration problem in P as well. In some cases, the satisfiability problem is NP-complete while the reconfiguration problem is in P . (For example, this is the case for 1-in-3 SAT, whose reconfiguration problem is trivial.) Examples in which the existence problem is in P , but the reconfiguration problem is PSPACE-complete can also be found. Prominent examples are reconfiguration of shortest paths [2] and reconfiguration of 4-colorings of bipartite and planar graphs [3]. Our results provide further examples of such a situation.

Our results. We give hardness results for reconfiguration problems involving solutions of special families of Boolean satisfiability problems, subset sum and knapsack problems, and, more generally, 0-1 linear programming problems.

In Section 2, we prove that the problem of reconfiguring satisfying assignments to a planar monotone instance of Not-All-Equal 3-SAT by single variable flips is PSPACE-complete. Interestingly, the planar Not-All-Equal 3-SAT problem is in P . If we further restrict to monotone instances, the reconfiguration problem is equivalent to reconfiguration of 2-colorings of 3-uniform hypergraphs with planar vertex-edge incidence graphs.

In Section 3, we consider the Subset Sum reconfiguration problem, that is, reconfiguration of subsets of a set of integers with the same sum. For this, we need to be able to perform elementary moves involving three elements of the set. We show that this problem is again PSPACE-complete.

Finally, in Section 4, we prove the PSPACE-completeness of the problem of finding a path between two points of the hypercube that is constrained to lie within a polytope. We show that the hardness result holds even if the number of inequalities defining the polytope is $O(1)$, and the coefficients involved are polynomial.

2 Planar NAE 3-SAT Reconfiguration

In this section, we give new results on the reconfiguration problems for a variant of Boolean satisfiability.

Definition 1 (Boolean Satisfiability Reconfiguration Problem). *Given an instance of a Boolean satisfiability problem and two satisfying assignments s and t , does there exist a sequence of satisfying assignments s_1, s_2, \dots, s_k such that $s_1 = s$, $s_k = t$, and for all $i \in [k - 1]$, s_{i+1} can be obtained from s_i by a single variable flip?*

Such problems (also referred to as the $s - t$ -connectivity problems for Boolean satisfiability) have been considered extensively before [9, 22, 23, 28, 26, 27]. Here we investigate the complexity of the reconfiguration versions of Boolean satisfiability problems in which the *variable-clause incidence graph* is planar. The variable-clause incidence graph of a CNF formula is a bipartite graph in whose set of

vertices is the union of the set of clauses and the set of variables of the formula, and a variable vertex is adjacent to a clause vertex if the variable appears in the clause, in either positive or negative form. The *planar 3-SAT problem* is the 3-SAT problem restricted to instances with a planar variable-clause incidence graph. It has long been known that planar 3-SAT is NP-complete [21].

In the *NAE 3-SAT problem*, satisfying assignments are forbidden from containing clauses in which all literals have the same value. Hence in a satisfying assignment, every clause has exactly two literals with the same value. In an instance of *Monotone NAE 3-SAT*, all literals appearing in the clauses are positive.

Monotone NAE 3-SAT is equivalent to 2-coloring 3-uniform hypergraphs, and known to be NP-complete from Schaefer's dichotomy theorem. We consider instances of *Planar NAE 3-SAT*, where the variable-clause incidence graph is planar. In 1988, Moret proved the surprising result that Planar NAE 3-SAT is in P by reducing the problem to that of finding a maximum cut in a planar graph [25]. We prove:

Theorem 1. *Planar Monotone NAE 3-SAT Reconfiguration is PSPACE-complete.*

It is interesting to observe that the problem is PSPACE-complete despite the satisfiability problem lying in P. The proof relies on the Nondeterministic Constraint Logic framework of Hearn and Demaine [11, 12].




Fig. 1. Monotone NAE clauses implementing the AND (top) and OR (bottom) nodes in a constraint graph. Only one of the three versions of the gadget is used, depending on the encoding used for the edge orientations.

Nondeterministic Constraint Logic (NCL). In nondeterministic constraint logic, a *constraint graph* is an edge- and node-weighted graph. A *configuration* of such a graph is an orientation of its edges, and an orientation is *legal* provided that the sum of the weights of edges pointing to a node is at least the weight of this node. In what follows, we will further restrict to graphs in which all node weights equal 2, and edges have weights either 1 or 2. The latter are referred to as red and blue edges, respectively. Furthermore, we only have two types of nodes: AND nodes with one blue and two red incident edges, and OR nodes with three blue incident edges. It was proved that the framework retains all of its expressive power, even under these restrictions [11]. The names of the two node types come from the interpretation of the incoming weight constraint: a configuration is legal if and only if (i) for all AND nodes, the blue edge is not outgoing unless both red edges are incoming, (ii) for all OR nodes, at least one edge is incoming.

Definition 2 (C2C Problem). *Given a constraint graph and two legal configurations C_1 and C_2 , can C_2 be obtained from C_1 by flipping one edge at a time, so that all intermediate configurations are also legal?*

Theorem 2 ([11]). *The C2C problem is PSPACE-complete, even if the constraint graph is restricted to be planar.*

Sketch of Proof of Theorem 1. As a warmup, we first consider the known reduction from the planar C2C problem to planar 3-SAT reconfiguration. Given a planar constraint graph, we define one Boolean variable per edge. When considering an edge x incident to a node, we denote by x_{in} the literal corresponding to the orientation of x towards the node, and the opposite literal by x_{out} . For a given AND node with a blue incident edge x and two red incident edges y and z , we add the two clauses $(x_{\text{in}} \vee y_{\text{in}})$ and $(x_{\text{in}} \vee z_{\text{in}})$, forcing both y_{in} and z_{in} to be true whenever x_{in} is false. For a given OR node with three incident blue edges x , y , and z , we add the single clause $(x_{\text{in}} \vee y_{\text{in}} \vee z_{\text{in}})$. The resulting variable-clause incidence graph is planar whenever the initial constraint graph is. This reduction is due to Sarah Eisenstat,⁶ and is also alluded to by Gopalan et al. [9].

The hardness proof for the planar monotone NAE 3-SAT reconfiguration problem is more involved. The AND and OR nodes are also translated into monotone NAE clauses involving variables representing edges of the constraint graph, but also other additional variables. Furthermore, the translation may use different sets of clauses depending on whether the variable representing the edge corresponds to the incoming or outgoing orientation. The reduction is summarized on Figure 1. In this figure, triangles represent monotone NAE clauses. The symbols F and T represent variables whose values have been fixed by a “rigid” gadget. The complete proof is provided in the full version of the paper on arXiv.

⁶ MIT Course 6.890, “Algorithmic Lower Bounds: Fun with Hardness Proofs” (Fall ’14), Lecture 17.

3 Subset Sum Reconfiguration

We now consider the reconfiguration problem for the well-known subset sum problem.

Definition 3 (Subset Sum Problem). *Given an integer x and a set of integers $S = \{a_1, a_2, \dots, a_n\}$, does there exist a subset $A \subseteq [n]$ such that $\sum_{i \in A} a_i = x$?*

If we restrict our reconfiguration steps to involve only a single element of S , the reconfiguration problem is trivial, as no single such move can maintain the same sum. We therefore consider more general reconfiguration steps. We say that a set of integers A_1 can be *k-move reconfigured* into a second set of integers A_2 whenever the symmetric difference of A_1 and A_2 has cardinality at most k .

Definition 4 (k-move Subset Sum Reconfiguration Problem). *Given two solutions A_1 and A_2 to an instance of the subset sum problem, can A_2 be obtained by repeated k-move reconfiguration, beginning with A_1 , so that all intermediate subsets are also solutions?*

The problem remains trivial for $k = 2$, since any removed element must be replaced by itself. For $k = 3$, we prove the following theorem.

Theorem 3. *The 3-move subset sum reconfiguration problem is strongly PSPACE-complete.*

The problem is *strongly* PSPACE-complete, meaning that it remains PSPACE-complete when the input integer set is given in unary. The corresponding instances of the subset sum problem can be solved in polynomial time using dynamic programming. This is another example of a reconfiguration problem that is PSPACE-complete, despite the underlying decision problem lying in P. We first note that the problem is contained in PSPACE. A simple proof can be found in the full version of the paper on arXiv. As for hardness, the reduction is done in two steps. First, from the Sliding Tokens problem to the Exact Cover reconfiguration problem (Lemma 1), then to the 3-move Subset Sum reconfiguration problem (Theorem 3).

Definition 5 (Token Slide Reconfiguration). *Given two independent sets I_1, I_2 of a graph $G = (V, E)$, I_1 can be reconfigured into I_2 via a token slide provided $(I_1 - I_2) \cup (I_2 - I_1) = \{v_1, v_2\}$ and $\{v_1, v_2\} \in E$.*

Observe that a token slide corresponds to changing the selection of a vertex $v_1 \in I_1$ to a neighboring vertex $v_2 \in I_2$, possible exactly when v_1 is the only vertex in I_1 among v_1, v_2 , and their neighbors.

Definition 6 (Sliding Tokens Problem). *Given two independent sets I_1, I_2 , can I_1 be reconfigured into I_2 via repeated token slides?*

An *exact cover* is a set cover that covers every element exactly once.

Definition 7 (Exact Cover Split and Merge Reconfiguration). *Given a set \mathcal{S} of subsets of a set U , and two exact covers $C_1, C_2 \subseteq \mathcal{S}$, C_1 can be reconfigured into C_2 via a split (and C_2 can be reconfigured into C_1 via a merge) provided that there exist $S_1, S_2, S_3 \subseteq \mathcal{S}$ with $C_1 - C_2 = S_1$ and $C_2 - C_1 = \{S_2, S_3\}$.*

Since C_1, C_2 are exact covers, $S_1 = S_2 \cup S_3$ and $S_2 \cap S_3 = \emptyset$.

Definition 8 (Exact Cover Reconfiguration Problem). *Given a set \mathcal{S} of subsets of a set U , can C_1 be reconfigured into C_2 via repeated splits and merges?*

Recall that a set \mathcal{S} of subsets of a set U can be considered as a hypergraph $G = (U, \mathcal{S})$, where each element of U is a vertex and each element of \mathcal{S} is a hyperedge. We say that a hypergraph is k -colorable whenever we can assign one of k colors to each vertex such that no two vertices in a hyperedge have the same color.

Lemma 1. *The exact cover reconfiguration problem is PSPACE-hard for instances that are 23-colorable hypergraphs.*

Proof. The proof of Theorem 23 of [11] establishes that the sliding tokens problem is PSPACE-hard on 3-regular graphs (see Section 3.2 of [3] for further discussion). A trivial modification of the proof suffices to prove that a *labeled* variant of the sliding tokens problem, where each token has a unique label, is also PSPACE-hard. The reduction is from this variant. The following describes an input instance of the labeled sliding tokens problem:

- $G = (V, E)$, a 3-regular graph.
- T , a set of labeled tokens.
- $p_1 : T \rightarrow V$, a function mapping each labeled token to a vertex placement in the starting configuration.
- $p_2 : T \rightarrow V$, a function mapping each labeled token to a vertex placement in the ending configuration.

Also, $I_1 = \{p_1(t) : t \in T\}$ and $I_2 = \{p_2(t) : t \in T\}$ are independent sets of size $|T| \leq |V|$.

Output U and \mathcal{S} . The output exact cover instance has a set U consisting of two types of elements: *vertices* $v_1, v_2, \dots, v_{|V|}$ and *tokens* $t_1, t_2, \dots, t_{|T|}$. That is, $U = \{v_1, v_2, \dots, v_{|V|}\} \cup \{t_1, t_2, \dots, t_{|T|}\}$.

For each pair of adjacent vertices $v_i, v_j \in V$, the set consisting of these two vertices and their neighbors is called a *slide set*, denoted $S_{i,j}$. The output set \mathcal{S} of subsets of U contains the following subsets for every pair of adjacent vertices v_i, v_j and token t_k :

- All subsets of $S_{i,j} - \{v_i\}$ and $S_{i,j} - \{v_j\}$.
- $\{v_i, t_k\}$ and $\{v_j, t_k\}$.
- $S_{i,j} \cup \{t_k\}$.

Output C_1 and C_2 . The starting configuration C_1 is the union of $\{\{v_i\} : v_i \in V - I_1\}$ and, for every $v_i \in I_1$, a set $\{v_i, t_k\}$ with a distinct t_k . Similarly, the ending configuration C_2 is the union of $\{\{v_i\} : v_i \in V - I_2\}$ and, for every $v_i \in I_2$, a set $\{v_i, t_k\}$ with a distinct t_k .

23-colorability of (U, \mathcal{S}) . Since G is 3-regular, G^3 has degree at most 21. So G can be 22-colored such that no two vertices of distance at most 3 (i.e. in a common slide set) have the same color. Such a coloring ensures that no pair of vertices in a common set in \mathcal{S} share a color. Coloring the tokens in T a distinct (23rd) color then gives a coloring of U such that no pair of elements of a common set share the same color.

High-level idea. The subsets containing exactly one vertex and token (e.g., $\{v_i, t_k\}$) represent the presence of the token t_k on vertex v_i . Subsets consisting of a slide set and token (e.g., $S_{i,j} \cup \{t_k\}$) represent the presence of a “mid-slide” token between v_i and v_j .

Sliding a token t_k from v_i to v_j is simulated by first merging $\{v_i, t_k\}$ and $S_{i,j} - \{v_i\}$ into $S_{i,j} \cup \{t_k\}$, and then splitting this set into $S_{i,j} - \{v_j\}$ and $\{v_j, t_k\}$. This sequence enforces the absence of tokens on neighbors of v_i and v_j , and the presence of a token on v_i or v_j , but not both. Before a merge-split sequence, additional splits and merges of token-less sets may be needed to obtain $S_{i,j} - \{v_i\}$.

Bijection between configurations. Call a configuration C of the output Exact Cover Reconfiguration instance *maximally split* if every C in C contains exactly one vertex and up to one token. The following defines a function f_{red} from token arrangements to maximally split covers:

- Each token-less vertex corresponds to a set $\{v_i\}$ in the cover.
- Each token t_k placed at v_i corresponds to a set $\{v_i, t_k\}$ in the cover.

Notice that f_{red} is a bijection and $f_{\text{red}}(p_1) = C_1$, $f_{\text{red}}(p_2) = C_2$.

Reduction structure. The remainder of the proof is devoted to proving the following claim: a token arrangement p' is reachable from a token arrangement p if and only if $f_{\text{red}}(p')$ is reachable from $f_{\text{red}}(p)$ via splits and merges.

Both directions are proved inductively. That is, we consider only “adjacent” configurations. We also assume that the starting token arrangement $p : T \rightarrow V$ has $\{p(t) : t \in T\}$ independent.

Sliding tokens reachability \Rightarrow exact cover reachability. Let p be a token arrangement that can be reconfigured into p' via a token slide from v_i to v_j . Then $f_{\text{red}}(p')$ can be reached from $f_{\text{red}}(p)$ via the following sequence of merges and splits:

1. Repeatedly merge token-less vertex sets to form $S_{i,j} - \{v_i\}$.
2. Merge $S_{i,j} - \{v_i\}$ and $\{v_i, t_k\}$ into $S_{i,j} \cup \{t_k\}$.
3. Split $S_{i,j} \cup \{t_k\}$ into $S_{i,j} - \{v_j\}$ and $\{v_j, t_k\}$.
4. Repeatedly split the token-less vertex set $S_{i,j} - \{v_j\}$ into single vertex sets.

Exact cover reachability \Rightarrow *sliding tokens reachability*. For each exact cover configuration C in the output instance, at least one maximally split configuration is reachable from C via a sequence of splits. Call the set of all such configurations the *sploot set* of C , denoted $\text{sploot}(C)$.

Let C, C' be maximally split configurations such that C can be reconfigured into C' and C_{inter} is the first configuration reached such that $\text{sploot}(C_{\text{inter}}) \neq \{C\}$. By induction, assume $C' \in \text{sploot}(C_{\text{inter}})$.

Since splits and token-less merges do not add elements to a sploot set, C_{inter} is obtained by merging two sets, one of which contains a token. Since the only token-containing sets that can be merged are those of the form $\{v_i, t_k\}$, C_{inter} is obtained by merging $\{v_i, t_k\}$ and $S_{i,j} - \{v_i, t_k\}$ to obtain $S_{i,j} \cup \{t_k\}$ for some v_i, v_j , and t_k . Notice that it may be the case that $S_{i,j} = S_{i',j'}$ for other pairs i', j' .

Such a merge allows two kinds of splits:

- Splitting $S_{i,j}$ into $S_{i,j} - \{v_i, t_k\}$ (to obtain the previous configuration, with sploot set $\{C\}$).
- Splitting $S_{i,j}$ into $S_{i',j'} - \{v'_j, t_k\}$, where $S_{i,j} = S_{i',j'}$ (to obtain a new configuration with sploot set $\{C'\}$, where C' is identical to C , except that C' contains $\{v'_j, t_k\}, \{v'_i\}$ instead of $\{v_i, t_k\}, \{v_j\}$).

Since $S_{i,j} - \{v_i\}, \{v_j, t_k\} \in C$, the token arrangement p with $f_{\text{red}}(p) = C$ has no tokens on vertices in $S_{i,j}$ except for token t_k on v_i . Since $S_{i,j} \cup S_{i',j'} = S_{i,j}$ contains all neighbors of v_i, v_j, v'_i, v'_j , the token arrangement obtained by moving the location of t_k in p from v_i to v_j, v'_i , or v'_j is an independent set.

So all that remains is to prove that there are a sequence of slides moving t_k from v_i to v'_j via vertices in $\{v_i, v_j, v'_i, v'_j\}$. Since $S_{i,j} = S_{i',j'}$, $v'_i, v'_j \in S_{i,j}$ and so either $v_i \in \{v'_i, v'_j\}$, or there is an edge $\{v_i, v'_i\}$ or $\{v_i, v'_j\} \in E$. So t_k can slide from v_i to either v'_i or v'_j (via 0 or 1 slides), and then from v'_i or v'_j to v'_j (via 0 or 1 slides). \square

We are now ready to prove the main result of this section.

Theorem 3. *The 3-move subset sum reconfiguration problem is strongly PSPACE-complete.*

Proof. The reduction is from the Exact Cover reconfiguration problem for instances that are 23-colorable induced hypergraphs, proved PSPACE-hard by Lemma 1. Observe that every 3-move subset sum reconfiguration is either a *merge*, where a_i and a_j are replaced by $a_i + a_j$, or a *split*, where $a_i + a_j$ is replaced by a_i and a_j . Each set split or merge will correspond to a 3-move split or merge, respectively, in the output instance.

Output numbers and sum. A function $f : U \rightarrow \mathbb{N}$ maps each element of the universe U of the input exact cover reconfiguration problem to a positive integer, and the numbers in the output 3-move subset sum reconfiguration instance are $\{\sum_{a \in S} f(a) : S \in \mathcal{S}\}$ and the output target sum is $\sum_{a \in U} f(a)$.

Elements of U are partitioned according to their colors $1, 2, \dots, 23$ and (arbitrarily) labeled $a_1, a_2, \dots, a_{|U|}$. The function f maps a color- j element a_i to $i \cdot 2^{100j \lceil \log_2(|U|) \rceil}$. In binary, this mapping consists of the binary encoding of i followed by $100j \lceil \log_2(|U|) \rceil$ zeros.

Output size. The output instance consists of $|\mathcal{S}|$ numbers, each between 0 and $|U| \cdot 2^{100 \cdot 23 \lceil \log_2(|U|) \rceil} = O(|U|^2)$. So the output sum is $O(|U|^3)$. Thus the output instance, encoded in unary, has length $O(|\mathcal{S}||U|^2 + |U|^3)$, i.e. polynomial in the input instance.

Correctness. A reconfiguration in both the exact cover and 3-move subset sum problems involves splitting or merging elements. Thus it suffices to prove that the function f yields a one-to-one mapping $g : \mathcal{S} \rightarrow \mathbb{N}$ given by $g(S) = \sum_{a \in S} f(a)$.

Recall that the function f maps each element $a_i \in U$ to a value based upon the color of a_i . The sums of the outputs of f for all elements of all colors 1 to $j-1$ is at most $2^{100(j-1) \lceil \log_2(|U|) \rceil} \cdot |U|^2 \leq 2^{(100j-98) \lceil \log_2(|U|) \rceil}$ while the output of f for any element of any color j or larger is at least $2^{100j \lceil \log_2(|U|) \rceil} \geq 2^{98} \cdot 2^{(100j-98) \lceil \log_2(|U|) \rceil}$.

Thus if a pair of sets $S_1, S_2 \subseteq \mathcal{S}$ have $S_1 \neq S_2$, then their color- j elements differ, this difference cannot be made up by adding or removing elements of colors 1 to $j-1$ (values too small) or colors $j+1$ to 23 (values too large). Thus if $S_1 \neq S_2$, then $g(S_1) \neq g(S_2)$. \square

4 Reconfiguration Problems and Paths in Hypercubes

The n -hypercube is the graph with vertex set $\{0,1\}^n$ such that two vertices are adjacent whenever their coordinates differ by exactly one component. In this section, we consider the following abstraction of reconfiguration problems involving subsets.

Definition 9 (Constrained Hypercube Path). *Given two vertices s, t of the n -hypercube, both contained in a polytope $P := \{x \in \mathbb{R}^n : Ax \leq b\}$ for some $A = (a_{ij}) \in \mathbb{Z}^{d \times n}$ and $b \in \mathbb{Z}^d$, does there exist a path from s to t in the hypercube, all vertices of which lie in P ?*

The constrained hypercube path problem can be seen as a reconfiguration analogue of the *0-1 integer linear programming* (0-1 ILP) satisfiability problem, which simply asks for the existence of a 0-1 point in the inside P , and is a standard NP-complete problem from Karp's list. (Note that this problem is distinct from the 0-1 ILP Reconfiguration problem defined in Ito et al. [16]: in the latter, a solution must optimize some objective function, while we are only concerned with satisfiability.)

The subset sum problem is the question of the existence of a 0-1 point in a polytope consisting of a subspace of dimension $n-1$, hence defined by two linear constraints with the same coefficients. Similarly, the knapsack (decision) problem involves exactly two linear constraints, and the Knapsack reconfiguration problem can be cast as a special case of the constrained hypercube path problem where $d=2$. The definitions are as follows.

Definition 10 (Knapsack Problem). *Given integers ℓ and u and two sets of integers $S = \{a_1, a_2, \dots, a_n\}$ and $W = \{w_1, w_2, \dots, w_n\}$, does there exist a subset $A \subseteq [n]$ such that $\sum_{i \in A} a_i \geq \ell$ and $\sum_{i \in A} w_i \leq u$?*

Definition 11 (Knapsack Reconfiguration Problem). *Given two solutions A_1 and A_2 to an instance of the knapsack problem, can A_2 be obtained by repeated 1-move reconfiguration, beginning with A_1 , so that all intermediate subsets are also solutions?*

Demaine and Ito considered the knapsack reconfiguration problem in the case where $S = W$ [15]. They proved that the problem was NP-hard, and gave an approximation algorithm for finding a reconfiguration sequence in which the intermediate steps satisfy one of the constraints only up to some multiplicative factor. Whether the knapsack reconfiguration problem is PSPACE-complete is a tantalizing open question. Characterizing the complexity of the knapsack reconfiguration problem implies understanding the complexity of the constrained hypercube path problem for bounded values of d . We do not settle the former question, but provide an answer to the latter. The proof of Theorem 4 uses techniques from the proof of Theorem 3, and can be found in full version of the paper on arXiv.

Theorem 4. *The Constrained Hypercube Path problem is PSPACE-complete, even when $d = O(1)$.*

Acknowledgements. This work was initiated at the 32nd Bellairs Winter Workshop on Computational Geometry, January 27-February 3, 2017. We thank the other participants of the workshop for a productive and positive atmosphere.

References

1. Bonamy, M., Johnson, M., Lignos, I., Patel, V., Paulusma, D.: Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs. *J. Comb. Optim.* 27(1), 132–143 (2014)
2. Bonsma, P.S.: The complexity of rerouting shortest paths. *Theor. Comput. Sci.* 510, 1–12 (2013)
3. Bonsma, P.S., Cereceda, L.: Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theor. Comput. Sci.* 410(50), 5215–5226 (2009)
4. Bonsma, P.S., Kaminski, M., Wrochna, M.: Reconfiguring independent sets in claw-free graphs. In: 14th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT). pp. 86–97 (2014)
5. Bonsma, P.S., Mouawad, A.E., Nishimura, N., Raman, V.: The complexity of bounded length graph recoloring and CSP reconfiguration. In: 9th International Symposium on Parameterized and Exact Computation (IPEC). pp. 110–121 (2014)
6. Bonsma, P.S.: Independent set reconfiguration in cographs and their generalizations. *Journal of Graph Theory* 83(2), 164–195 (2016)
7. Cereceda, L., van den Heuvel, J., Johnson, M.: Finding paths between 3-colorings. *Journal of Graph Theory* 67(1), 69–82 (2011)
8. Demaine, E.D., Demaine, M.L., Fox-Epstein, E., Hoang, D.A., Ito, T., Ono, H., Otachi, Y., Uehara, R., Yamada, T.: Linear-time algorithm for sliding tokens on trees. *Theor. Comput. Sci.* 600, 132–142 (2015)

9. Gopalan, P., Kolaitis, P.G., Maneva, E.N., Papadimitriou, C.H.: The connectivity of Boolean satisfiability: Computational and structural dichotomies. *SIAM J. Comput.* 38(6), 2330–2355 (2009)
10. Haddadan, A., Ito, T., Mouawad, A.E., Nishimura, N., Ono, H., Suzuki, A., Tebbal, Y.: The complexity of dominating set reconfiguration. *Theor. Comput. Sci.* 651, 37–49 (2016)
11. Hearn, R.E., Demaine, E.D.: PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theoretical Computer Science* 343(1–2), 72–86 (2005)
12. Hearn, R.E., Demaine, E.D.: *Games, Puzzles, & Computation*. A. K. Peters/CRC Press (2009)
13. van den Heuvel, J.: The complexity of change. In: Blackburn, S.R., Gerke, S., Wildon, M. (eds.) *Surveys in Combinatorics*. London Mathematical Society Lecture Note Series (2013)
14. Hoang, D.A., Uehara, R.: Sliding tokens on a cactus. In: *27th International Symposium on Algorithms and Computation (ISAAC)*. pp. 37:1–37:26 (2016)
15. Ito, T., Demaine, E.D.: Approximability of the subset sum reconfiguration problem. *Journal of Combinatorial Optimization* 28(3), 639–654 (2014)
16. Ito, T., Demaine, E.D., Harvey, N.J.A., Papadimitriou, C.H., Sideri, M., Uehara, R., Uno, Y.: On the complexity of reconfiguration problems. *Theoretical Computer Science* 412, 154–165 (2011)
17. Ito, T., Kaminski, M., Demaine, E.D.: Reconfiguration of list edge-colorings in a graph. *Discrete Applied Mathematics* 160(15), 2199–2207 (2012)
18. Ito, T., Nooka, H., Zhou, X.: Reconfiguration of vertex covers in a graph. *IEICE Transactions* 99-D(3), 598–606 (2016)
19. Kaminski, M., Medvedev, P., Milanic, M.: Shortest paths between shortest paths. *Theor. Comput. Sci.* 412(39), 5205–5210 (2011)
20. Kaminski, M., Medvedev, P., Milanic, M.: Complexity of independent set reconfigurability problems. *Theor. Comput. Sci.* 439, 9–15 (2012)
21. Lichtenstein, D.: Planar satisfiability and its uses. *SIAM Journal on Computation* 11, 329–343 (1982)
22. Makino, K., Tamaki, S., Yamamoto, M.: On the Boolean connectivity problem for Horn relations. *Discrete Applied Mathematics* 158(18), 2024–2030 (2010)
23. Makino, K., Tamaki, S., Yamamoto, M.: An exact algorithm for the Boolean connectivity problem for k-CNF. *Theor. Comput. Sci.* 412(35), 4613–4618 (2011)
24. Mizuta, H., Ito, T., Zhou, X.: Reconfiguration of steiner trees in an unweighted graph. In: *27th International Workshop on Combinatorial Algorithms (IWOCA)*. pp. 163–175 (2016)
25. Moret, B.M.E.: Planar NAE3SAT is in P. *ACM SIGACT News* 19(2), 51–54 (1988)
26. Mouawad, A.E., Nishimura, N., Pathak, V., Raman, V.: Shortest reconfiguration paths in the solution space of Boolean formulas. In: *42nd International Colloquium on Automata, Languages, and Programming (ICALP)*. pp. 985–996 (2015)
27. Scharpfenecker, P.: On the structure of solution-graphs for Boolean formulas. In: *20th International Symposium on Fundamentals of Computation Theory (FCT)*. pp. 118–130 (2015)
28. Schwerdtfeger, K.W.: A computational trichotomy for connectivity of Boolean satisfiability. *JSAT* 8(3/4), 173–195 (2014)
29. Slocum, J., Sonneveld, D.: *The 15 Puzzle: How it Drove the World Crazy*. The Slocum Puzzle Foundation (2006)