Rainbow cycles in flip graphs*

Stefan Felsner, Linda Kleist, Torsten Mütze†, Leon Sering

Institut für Mathematik
TU Berlin, 10623 Berlin, Germany
{felsner,kleist,muetze,sering}@math.tu-berlin.de

Abstract. The flip graph of triangulations has as vertices all triangulations of a convex n-gon, and an edge between any two triangulations that differ in exactly one edge. An r-rainbow cycle in this graph is a cycle in which every inner edge of the triangulation appears exactly r times. This notion of a rainbow cycle extends in a natural way to other flip graphs. In this paper we investigate the existence of r-rainbow cycles for three different flip graphs on classes of geometric objects: the aforementioned flip graph of triangulations of a convex n-gon, the flip graph of plane trees on an arbitrary set of n points, and the flip graph of non-crossing perfect matchings on a set of n points in convex position. In addition, we consider two flip graphs on classes of non-geometric objects: the flip graph of permutations of $\{1, 2, \ldots, n\}$ and the flip graph of k-element subsets of $\{1, 2, \ldots, n\}$. In each of the five settings, we prove the existence and non-existence of rainbow cycles for different values of r, n and k.

1. Introduction

Flip graphs are fundamental structures associated with families of geometric objects such as triangulations, plane spanning trees, non-crossing matchings, partitions or dissections. A classical example is the flip graph $G_n^{\mathtt{T}}$ of the triangulations of a convex n-gon. The vertices of this graph are the triangulations of the n-gon, and two triangulations are adjacent whenever they differ by exactly one edge. In other words, moving along an edge of $G_n^{\mathtt{T}}$ corresponds to flipping the diagonal of a convex quadrilateral formed by two triangles. Figure 1 shows the graph $G_6^{\mathtt{T}}$.

A question that has received considerable attention is to determine the diameter of $G_n^{\mathtt{T}}$, i.e., the number of flips that is necessary and sufficient to transform any triangulation into any other; see the survey [BH09]. In a landmark paper [STT88], Sleator, Tarjan and Thurston proved that the diameter of $G_n^{\mathtt{T}}$ is $2n-10$ for sufficiently large n. Recently, Pournin [Pou14] gave a combinatorial proof that the diameter is $2n-10$ for all $n > 12$. A challenging algorithmic problem in this direction is to efficiently compute a minimal sequence of flips that transforms two given triangulations into each other; see [LZ98, Rog99, CM18]. These questions involving the diameter of the flip graph have also been investigated heavily when the n points are not in convex, but in general position; see e.g. [HOS96, HNU99, Epp10]. Moreover, apart from the diameter, many other properties of the flip graph $G_n^{\mathtt{T}}$ have been investigated, e.g., the eccentricities of its vertices [Pou19], its vertex-connectivity [HN99], its realizability as the graph of a convex polytope and the automorphism group [Lee89], its chromatic number [FMFPH$^+$09, BRSW18], and its genus [PP18a].

Another property of major interest is the existence of a Hamilton cycle in $G_n^{\mathtt{T}}$. This was first established by Lucas [Luc87] and a very nice and concise proof was given by Hurtado and Noy [HN99] (see [MP16, PP18b, HHMW19] for various generalizations). The reason for the interest in Hamilton cycles is that a Hamilton cycle in $G_n^{\mathtt{T}}$ corresponds to a so-called *Gray code*, i.e., an algorithm that allows to generate each triangulation exactly once, by performing only a single flip

*An extended abstract of this paper has been accepted for presentation at the 34th International Symposium on Computational Geometry, SoCG 2018.

†Torsten Mütze is also affiliated with Charles University, Faculty of Mathematics and Physics, and was supported by GACR grant GA 19-08554S, and by DFG grant 413902284.

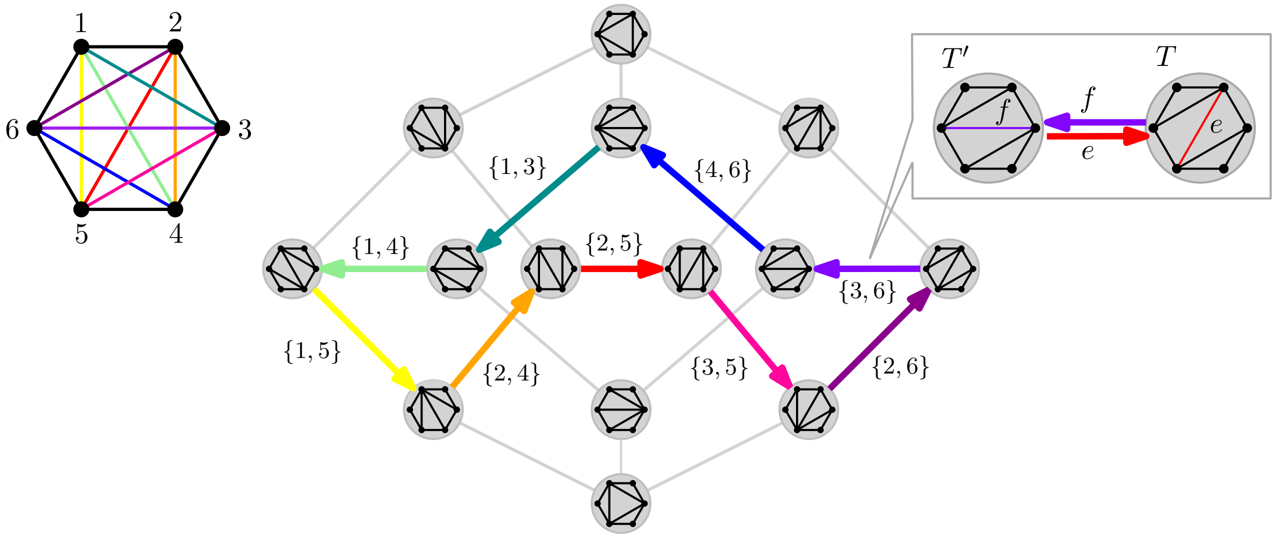

FIGURE 1. The flip graph of triangulations $G_n^{\mathtt{T}}$ of a convex n-gon for $n = 6$, and a rainbow cycle in this graph.

operation when moving to the next triangulation. In general, the task of a Gray code algorithm is to generate all objects in a particular combinatorial class, each object exactly once, by applying only a small transformation in each step, such as a flip in a triangulation. Combinatorial classes of interest include geometric configurations such as triangulations, plane spanning trees or non-crossing perfect matchings, but also classes without geometric information such as permutations, combinations, bitstrings etc. This fundamental topic is covered in depth in the most recent volume of Knuth's seminal series *The Art of Computer Programming* [Knu11], and in the classical books by Nijenhuis and Wilf [NW78, Wil89]. Here are some important Gray code results in the geometric realm: Hernando, Hurtado and Noy [HHN02] proved the existence of a Hamilton cycle in the flip graph of non-crossing perfect matchings on a set of $2m$ points in convex position for every even $m \geq 4$. Aichholzer et al. [AAHV07] described Hamilton cycles in the flip graphs of plane graphs on a general point set, for plane and connected graphs and for plane spanning trees on a general point set. Huemer et al. [HHNOP09] constructed Hamilton cycles in the flip graphs of non-crossing partitions of a point set in convex position, and for the dissections of a convex polygon by a fixed number of non-crossing diagonals.

As mentioned before, a Hamilton cycle in a flip graph corresponds to a cyclic listing of all objects in some combinatorial class, such that each object is encountered exactly once, by performing a single flip in each step. In this work we consider the *dual* problem: we are interested in a cyclic enumeration of some of the combinatorial objects, such that each flip operation is encountered exactly once. For instance, in the flip graph of triangulations $G_n^{\mathtt{T}}$, we ask for the existence of a cycle with the property that each inner edge of the triangulation appears (and disappears) exactly once. An example of such a cycle is shown in Figure 1. This idea can be formalized as follows. Consider two triangulations T and T' that differ in flipping the diagonal of a convex quadrilateral, i.e., T' is obtained from T by removing the diagonal e and inserting the other diagonal f. We view the edge between T and T' in the flip graph $G_n^{\mathtt{T}}$ as two arcs in opposite directions, where the arc from T to T' receives the label f, and the arc from T' to T receives the label e, so the label corresponds to the edge of the triangulation that enters in this flip; see the right hand side of Figure 1. Interpreting the labels as colors, we are thus interested in a directed cycle in the flip graph in which each color appears exactly once, and we refer to such a cycle as a *rainbow cycle*. More generally, for any integer $r \geq 1$, an *r-rainbow cycle* in $G_n^{\mathtt{T}}$ is a cycle in which each edge of the triangulation appears (and disappears) exactly r times.

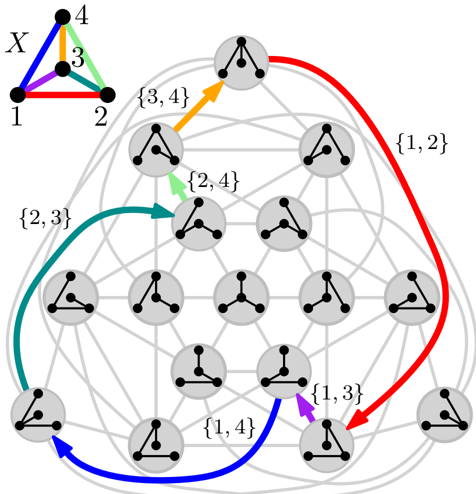

(a) Flip graph of plane spanning trees $G^{\mathtt{S}}_X$.

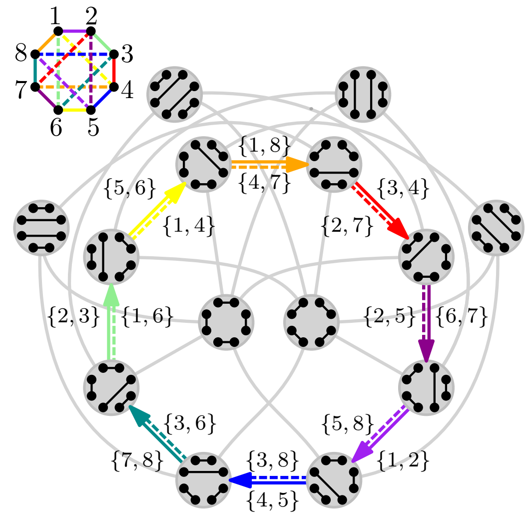

(b) Flip graph of non-crossing perfect matchings $G^{\mathtt{M}}_4$.

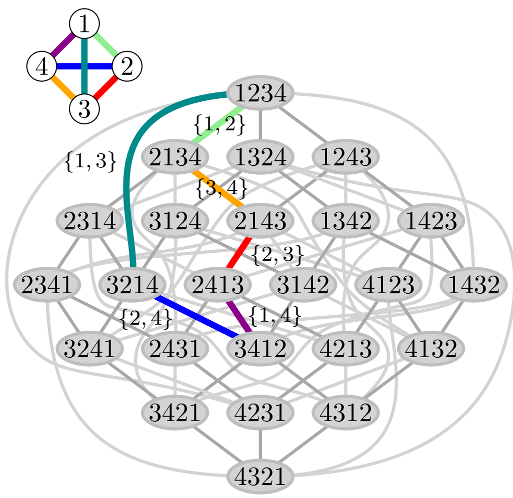

(c) Flip graph of permutations $G^{\mathtt{P}}_4$.

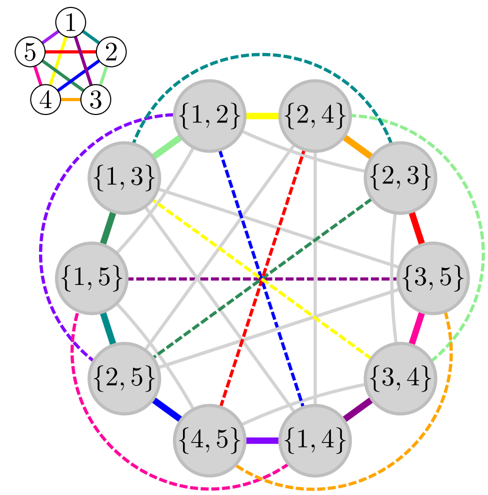

(d) Flip graph of subsets $G^{\mathtt{C}}_{5,2}$.

FIGURE 2. Examples of flip graphs with 1-rainbow cycles. In (d), two edge-disjoint rainbow Hamilton cycles in $G^{\mathtt{C}}_{5,2}$ are highlighted, one with bold edges and one with dashed edges.

Note that a rainbow cycle does not need to visit all vertices of the flip graph. Clearly, this notion of rainbow cycles extends in a natural way to all the other flip graphs discussed before; see Figure 2.

1.1. **Our results.** In this work we initiate the investigation of rainbow cycles in flip graphs for five popular classes of combinatorial objects. We consider three geometric classes: triangulations of a convex polygon, plane spanning trees on point sets in general position, and non-crossing perfect matchings on point sets in convex position. In addition, we consider two classes without geometric information: permutations of the set $[n] := \{1, 2, \ldots, n\}$, and k-element subsets of $[n]$. We proceed to present our results in these five settings in the order they were just mentioned. For the reader's convenience, all results are summarized in Table 1.

Table 1. Overview of results.

	flip graph			existence of r-rainbow cycle			
		vertices	arcs/edges	r	yes	no	
GEOMETRIC	$G_n^{\mathtt{T}}$	triangulations of convex n-gon	edge flip	1	$n \geq 4$		Thm. 1
				2	$n \geq 7$		
	$G_X^{\mathtt{S}}$	plane spanning trees on point set X in general position	edge flip	1,..., $\lvert X\rvert - 2$	$\lvert X\rvert \geq 3$		Thm. 2
	$G_m^{\mathtt{M}}$	non-crossing perfect matchings on $2m$ points in convex position	two edge flip	1	$m \in \{2, 4\}$	odd m, $m \in \{6, 8, 10\}$	Thm. 9
				2	$m \in \{6, 8\}$		
ABSTRACT	$G_n^{\mathtt{P}}$	permutations of $[n]$	transposition	1	$\lfloor n/2 \rfloor$ even	$\lfloor n/2 \rfloor$ odd	Thm. 14
	$G_{n,k}^{\mathtt{C}}$	k-subsets of $[n]$, $2 \leq k \leq \lfloor n/2 \rfloor$	element exchange	1	odd n and $k < n/3$	even n	Thm. 15
		2-subsets of $[n]$ for odd n		1	two edge-disjoint 1-rainbow Ham. cycles		

Our first result is that the flip graph of triangulations $G_n^{\mathtt{T}}$ defined in the introduction has a 1-rainbow cycle for $n \geq 4$ and a 2-rainbow cycle for $n \geq 7$ (Theorem 1 in Section 2).

Next, we consider the flip graph $G_X^{\mathtt{S}}$ of plane spanning trees on a point set X in general position; see Figure 2 (a). We prove that $G_X^{\mathtt{S}}$ has an r-rainbow cycle for any point set X with at least three points for any $r = 1, 2, \ldots, |X| - 2$ (Theorem 2 in Section 3).

We then consider the flip graph $G_m^{\mathtt{M}}$ of non-crossing perfect matchings on $2m$ points in convex position; see Figure 2 (b). We exhibit 1-rainbow cycles for $m = 2$ and $m = 4$ matching edges, and 2-rainbow cycles for $m = 6$ and $m = 8$. We also argue that there is no 1-rainbow cycle for $m \in \{6, 8, 10\}$, and none for any odd m. In fact, we believe that there are no 1-rainbow cycles in $G_m^{\mathtt{M}}$ for any $m \geq 5$. Our results for this setting are summarized in Theorem 9 in Section 4.

Next, we consider the flip graph $G_n^{\mathtt{P}}$ of permutations of $[n]$, where an edge connects any two permutations that differ in a transposition, i.e., in exchanging two entries at positions i and j; see Figure 2 (c). The edges of this graph are colored with the corresponding pairs $\{i, j\}$, and in a 1-rainbow cycle each of the $\binom{n}{2}$ possible pairs appears exactly once. We prove that $G_n^{\mathtt{P}}$ has a 1-rainbow cycle if $\lfloor n/2 \rfloor$ is even, and no 1-rainbow cycle if $\lfloor n/2 \rfloor$ is odd (Theorem 14 in Section 5).

Finally, we consider the flip graph $G_{n,k}^{\mathtt{C}}$ of k-element subsets of $[n]$, also known as (n, k)-combinations, where an edge connects any two subsets that differ in exchanging one element i for another element j, i.e., the symmetric difference of the subsets has cardinality two; see Figure 2 (d). The edges of this graph are colored with the corresponding pairs $\{i, j\}$, and in a 1-rainbow cycle each of the $\binom{n}{2}$ possible pairs appears exactly once. As $G_{n,k}^{\mathtt{C}}$ is isomorphic to $G_{n,n-k}^{\mathtt{C}}$, including the edge-coloring, we assume without loss of generality that $2 \leq k \leq \lfloor n/2 \rfloor$. We prove that $G_n^{\mathtt{C}}$ has a 1-rainbow cycle for every odd n and $k < n/3$, and we prove that it has no 1-rainbow cycle for any even n. The case $k = 2$ is of particular interest, as a 1-rainbow cycle in the flip graph $G_{n,2}^{\mathtt{C}}$ is a Hamilton cycle (both the number of subsets and the number of exchanges equal $\binom{n}{2}$). Moreover, we show that $G_{n,2}^{\mathtt{C}}$ even has two edge-disjoint 1-rainbow Hamilton cycles (for odd n). Our results in this setting are summarized in Theorem 15 in Section 6.

We conclude in Section 7 with some open problems.

1.2. **Related work.** Gray codes are named after Frank Gray, a physicist at Bell Labs, who in 1953 patented a simple scheme to generate all 2^n bitstrings of length n by flipping a single bit in each step. This classical inductive construction is now called the *binary reflected Gray code*; see [Wil89] or [Knu11]. Since its invention, there has been continued interest in developing binary Gray codes that satisfy various additional constraints, cf. the survey by Savage [Sav97]. The existence of a binary Gray code with the property that the bit-flip counts in each of the n coordinates are balanced, i.e., they differ by at most 2, was first established by Bakos [Ádá68, p. 28–37] (see also [BS96]). When n is a power of two, every bit appears (and disappears) exactly $1/2 \cdot 2^n/n =: r$ many times. This balanced Gray code therefore corresponds to an r-rainbow cycle in the corresponding flip graph. In this light, our results are a first step towards balanced Gray codes for other combinatorial classes. For 2-element subsets, we indeed construct perfectly balanced Gray codes.

The Steinhaus-Johnson-Trotter algorithm [Joh63, Tro62], also known as 'plain changes', is a method to generate all permutations of $[n]$ by adjacent transpositions $i \leftrightarrow i+1$. More generally, it was shown in [KL75] that all permutations of $[n]$ can be generated by any set of transpositions that form a spanning tree on the set of positions $[n]$. This is even possible under the additional constraint that in every second step the same transposition is applied [RS93].

The generation of (n, k)-combinations subject to certain restrictions on admissible exchanges $i \leftrightarrow j$ has been studied widely. Specifically, it was shown that all (n, k)-combinations can be generated with only allowing exchanges of the form $i \leftrightarrow i + 1$ [BW84, EHR84, Rus88], provided that n is even and k is odd, or $k \in \{0, 1, n-1, n\}$. The infamous *middle levels conjecture* asserts that all $(2k, k)$-combinations can be generated with only exchanges of the form $1 \leftrightarrow i$, and this conjecture has recently been proved in [Müt16, GMN18].

Rainbow cycles and paths have also been studied in graphs other than flip graphs. A well-known conjecture in this context due to Andersen [And89] asserts that every properly edge-colored complete graph on n vertices has a rainbow path of length $n-2$, i.e., a path that has distinct colors along each of its edges. Progress towards resolving this conjecture was recently made by Alon, Pokrovskiy and Sudakov [APS16], and Balogh and Molla [BM17].

2. Triangulations

In this section we consider a convex n-gon on points labeled clockwise by $1, 2, \ldots, n$, and we denote by $\mathcal{T}_n$ the set of all triangulations on these points. The graph $G_n^{\mathtt{T}}$ has $\mathcal{T}_n$ as its vertex set, and an arc (T, T') between any two triangulations T and T' that differ in exchanging the diagonal $e \in T$ of a convex quadrilateral formed by two triangles for the other diagonal $f \in T'$; see Figure 1. We refer to this operation as a *flip*, and we denote it by (e, f). Furthermore, we label the arc (T, T') with the edge f, so an arc is labeled with the edge that enters the triangulation in this flip. The set of arc labels of $G_n^{\mathtt{T}}$ is clearly $E_n := \{\{i, j\} \mid j - i > 1\} \setminus \{1, n\}$, and we think of these labels as colors. An r-rainbow cycle in $G_n^{\mathtt{T}}$ is a directed cycle along which every label from E_n appears exactly r times. Clearly, the length of an r-rainbow cycle equals $r|E_n| = r(\binom{n}{2} - n)$. For comparison, the number of vertices of $G_n^{\mathtt{T}}$ is the $(n-2)$-th Catalan number $\frac{1}{n-1}\binom{2n-4}{n-2}$. Given an r-rainbow cycle, the cycle obtained by reversing the orientation of all arcs is also an r-rainbow cycle, as every edge that appears r times also disappears r times. Here is an interesting interpretation of an r-rainbow cycle using the language of polytopes: The secondary polytope of the triangulations of a convex n-gon, called the *associahedron*, has the graph $G_n^{\mathtt{T}}$ as its skeleton, and the facets of this polytope are the triangulations with a fixed edge. Consequently, an r-rainbow cycle enters (and leaves) each facet of the associahedron exactly r times.

The following theorem summarizes the results of this section.

Theorem 1. *The flip graph of triangulations $G_n^{\mathtt{T}}$ has the following properties:*

(i) If $n \geq 4$, then $G_n^{\mathtt{T}}$ has a 1-rainbow cycle.
(ii) If $n \geq 7$, then $G_n^{\mathtt{T}}$ has a 2-rainbow cycle.

Proof. Let S_i be the star triangulation with respect to the point i, i.e., the triangulation where the point i has degree $n-1$. To transform S_1 into S_2 we can use the flip sequence

$$F_{1,n} := \big((\{1,3\},\{2,4\}),(\{1,4\},\{2,5\}),(\{1,5\},\{2,6\}),\ldots,(\{1,n-1\},\{2,n\})\big). \tag{1}$$

For any $i = 1,2,\ldots,n$, let $F_{i,n}$ denote the flip sequence obtained from $F_{1,n}$ by adding $i-1$ to all points on the right-hand side of (1). Here and throughout this proof addition is to be understood modulo n with $\{1,2,\ldots,n\}$ as representatives for the residue classes. Note that $F_{i,n}$ transforms S_i into S_{i+1} for any $i \in [n]$, and all the edges from E_n that are incident with the point $i+1$ appear exactly once during that flip sequence. Note also that $F_{i,n}$ has length $n-3$.

We begin proving (ii). The concatenation $(F_{1,n}, F_{2,n}, \ldots, F_{n,n})$ is a flip sequence which applied to S_1 leads back to S_1. Along the corresponding cycle C in $G_n^{\mathtt{T}}$, every edge from E_n appears exactly twice. Specifically, every edge $\{i,j\} \in E_n$ appears in the flip sequences $F_{i-1,n}$ and $F_{j-1,n}$. It remains to show that C is indeed a cycle, i.e., every triangulation appears at most once. For this observe that when applying $F_{i,n}$ to S_i, then for every $j = 1,2,\ldots,n-4$, in the j-th triangulation we encounter after S_i, the point i is incident with exactly $n-3-j$ diagonals, the point $i+1$ is incident with exactly j diagonals, while all other points are incident with at most two diagonals. We call these triangulations *bi-centered* with the two centers i and $i+1$. For $n \geq 8$, we have $\max_{1\leq j\leq n-4}\{n-3-j, j\} \geq 3$ and therefore we can determine at least one center k. The other center is either the point $k-1$ or the point $k+1$. Since only one of these two points is incident to some diagonal, we can identify it as the other center. Hence for any bi-centered triangulation encountered along C, we can uniquely reconstruct in which flip sequence $F_{i,n}$ it occurs. For $n = 7$ it can be verified directly that C is a 2-rainbow cycle.

It remains to prove (i). For $n \geq 4$, we obtain a 1-rainbow cycle in $G_n^{\mathtt{T}}$ by applying the flip sequence $X_n := (F_{3,4}, F_{4,5}, F_{5,6}, \ldots, F_{n-2,n-1}, F_{n-1,n}, F_{n,n})$ to the triangulation S_1. Note that X_n differs from X_{n-1} by replacing the terminal subsequence $F_{n-1,n-1}$ by $F_{n-1,n}$ and $F_{n,n}$; see Figure 3. By induction, this yields a sequence of length $\big(\binom{n-1}{2} - (n-1)\big) - (n-4) + 2(n-3) = \binom{n}{2} - n$. The fact that X_n produces a rainbow cycle follows by induction, by observing that applying X_{n-1} to S_1 in $G_n^{\mathtt{T}}$ yields a cycle along which every edge from E_{n-1} appears exactly once. Moreover, along this cycle the point n is not incident with any diagonals. The modifications described before to construct X_n from X_{n-1} shorten this cycle in $G_n^{\mathtt{T}}$ and extend it by a detour through triangulations where the point n is incident with at least one diagonal, yielding a cycle along which every edge from the following set appears exactly once:

$$\begin{aligned}
E_{n-1} \ \setminus \ & \big\{\{1,3\},\{1,4\},\ldots,\{1,n-2\}\big\} \\
\cup \ & \big\{\{n,2\},\{n,3\},\ldots,\{n,n-2\}\big\} \\
\cup \ & \big\{\{1,3\},\{1,4\},\ldots,\{1,n-2\},\{1,n-1\}\big\} \\
= E_{n-1} \ \cup \ & \big\{\{n,2\},\{n,3\},\ldots,\{n,n-2\}\big\} \ \cup \ \{1,n-1\} \ = \ E_n.
\end{aligned}$$

This shows that applying X_n to S_1 yields a 1-rainbow cycle in $G_n^{\mathtt{T}}$. $\square$

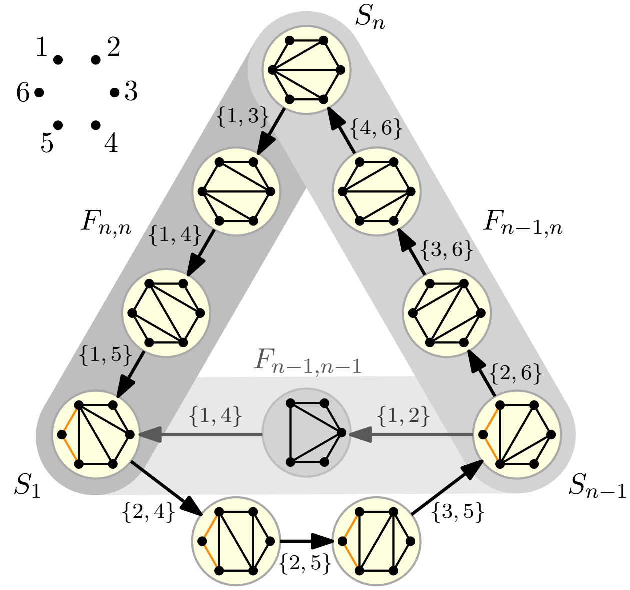

FIGURE 3. Illustration of the proof of Theorem 1 (i). The figure shows the inductive construction of a rainbow cycle for triangulations on 6 points from a rainbow cycle for 5 points.

3. SPANNING TREES

In this section we consider plane spanning trees on a set X of n points in general position, i.e., no three points are collinear. We use $\mathcal{S}_X$ to denote the set of all plane spanning trees on X. The graph G_X^{S} has $\mathcal{S}_X$ as its vertex set, and an arc (T,T') between any two spanning trees T and T' that differ in replacing an edge $e \in T$ by another edge $f \in T'$; see Figure 2 (a). We refer to this operation as a *flip*, and we denote it by (e,f). Furthermore, we label the arc (T,T') with the edge f, so an arc is labeled with the edge that enters the tree in this flip. Note that the entering edge f alone does not determine the flip uniquely (unlike for triangulations). Clearly, none of the two edges e and f can cross any of the edges in $T \cap T'$, but they may cross each other. The set of arc labels of G_X^{S} is clearly $E_X := \binom{X}{2}$, the set of all unordered pairs of points from X, and we think of these labels as colors. An r-rainbow cycle in G_X^{S} is a directed cycle along which every label from E_X appears exactly r times, so it has length $r\binom{n}{2}$.

The following theorem summarizes the results of this setting.

Theorem 2. *The flip graph of plane spanning trees G_X^{S} has the following properties:*

(i) For any point set X with $|X| \geq 3$ in general position, G_X^{S} has a 1-rainbow cycle.
(ii) For any point set X with $|X| \geq 4$ in general position and any $r = 2,3,\ldots,m$, where $m := |X|-1$ if $|X|$ is odd and $m := |X|-2$ if $|X|$ is even, G_X^{S} has an r-rainbow cycle.

3.1. **Proof of Theorem 2 (i).** We label the n points of X with integers $1,2,\ldots,n$ as follows; see Figure 4 (a). We first label an arbitrary vertex on the convex hull of X as point 1, and we then label the points from 2 to n in counter-clockwise order around 1 such that $\{1,2\}$ and $\{1,n\}$ are edges of the convex hull of X.

Given a graph G that has an edge e but that does not have an edge f, we write $G-e$ for the graph obtained from G by removing e, and we write $G+f$ for the graph obtained from G by adding f. Furthermore, for any subset $Y \subseteq X$ and point $i \in Y$ we write $S_i(Y)$ for the tree on Y that forms

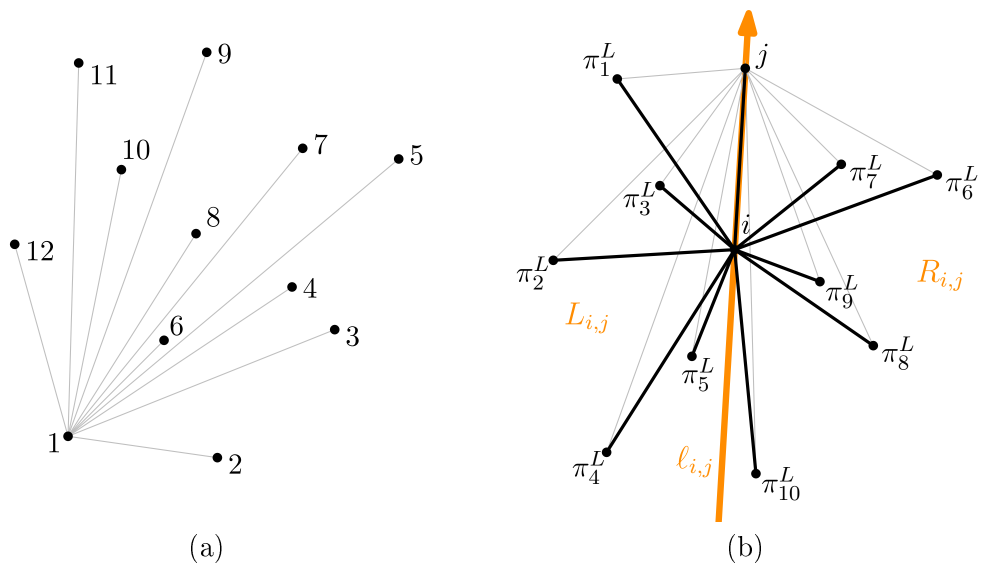

FIGURE 4. (a) Ordering of points $1, 2, \ldots, n$. (b) The point ordering π^L for the path from S_i to S_j.

a star with center vertex i. We write S_i for the star $S_i(X)$; see Figure 4 (b). For two distinct points $i, j \in [n]$, the directed line from i to j is denoted $\ell_{i,j}$. The direction allows us to distinguish the left and right half-plane. Let $L_{i,j}$ be the points of X strictly on the left and $R_{i,j}$ the points of X strictly on the right of the line $\ell_{i,j}$.

We will define two specific flip sequences $F^L_{i,j}$ and $F^R_{i,j}$ that transform the star S_i into the star S_j; see Figure 4 (b). Let τ^L be the sequence of all points k in $L_{i,j}$ ordered by decreasing clockwise angles (i, j, k). Similarly, let τ^R be the sequence of all points k in $R_{i,j}$ ordered by decreasing counter-clockwise angles (i, j, k). Let $\pi^L := (\tau^L, \tau^R)$ and $\pi^R := (\tau^R, \tau^L)$ be the concatenations of these two sequences.

The flip sequence $F^L_{i,j}$ is defined as

$$F^L_{i,j} := (f_1, f_2, \ldots, f_{n-1}) \quad \text{were} \quad f_k := \begin{cases} \big(\{i,j\}, \{j, \pi^L_1\}\big) & \text{if } k = 1, \\ \big(\{i, \pi^L_{k-1}\}, \{j, \pi^L_k\}\big) & \text{if } 2 \le k \le n-2, \\ \big(\{i, \pi^L_{n-2}\}, \{j, i\}\big) & \text{if } k = n-1. \end{cases}$$

The flip sequence $F^R_{i,j}$ is defined analogously, by using π^R instead of π^L.

Note that in both flip sequences, every edge from E_X that contains the point j appears exactly once. Furthermore, if $\{i, j\}$ is an edge of the convex hull of X, then either $L_{i,j}$ or $R_{i,j}$ is empty and therefore $F^L_{i,j} = F^R_{i,j}$. Otherwise, these flip sequences differ, as the first flip of $F^L_{i,j}$ adds an edge on the left of $\ell_{i,j}$, while the first flip of $F^R_{i,j}$ adds an edge on the right of $\ell_{i,j}$.

Clearly, each of the flip sequences $F^L_{i,j}$ and $F^R_{i,j}$ yields a path from S_i to S_j in the graph G^{S}_X, i.e., every flip adds an edge which is not in the tree and removes one which is in the tree, and the trees along the path are distinct plane spanning trees. We denote the paths from S_i to S_j in the graph G^{S}_X obtained from the flip sequences $F^L_{i,j}$ and $F^R_{i,j}$ by $P^L_{i,j}$ and $P^R_{i,j}$, respectively. We refer to the trees along these paths other than S_i and S_j as *intermediate trees*. Note that there are $n-2$ intermediate trees along each of the paths $P^L_{i,j}$ and $P^R_{i,j}$.

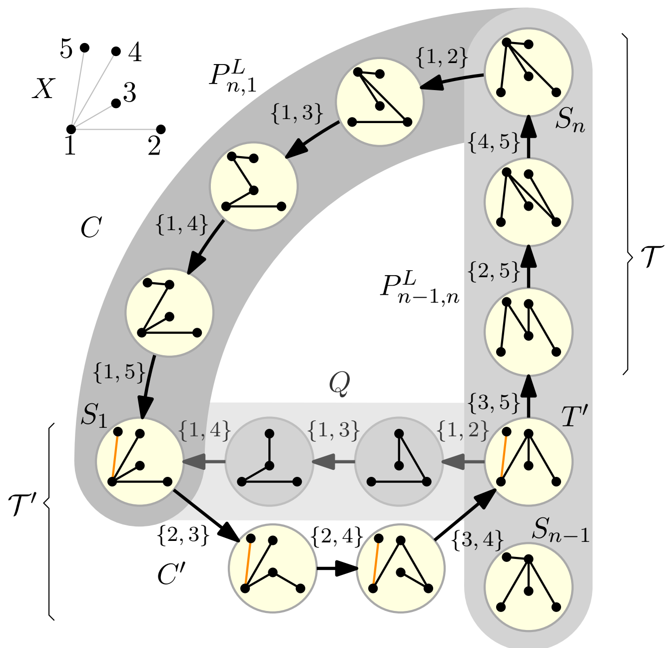

FIGURE 5. Illustration of the proof of Theorem 2 (i). In the induction step, the rainbow cycle from Figure 2 (a) on the point set $X \setminus \{5\} = [4]$ is extended to a rainbow cycle on the point set $X = [5]$.

Proof of Theorem 2 *(i).* We prove the following stronger statement by induction on n: For any point set $X = [n]$ of $n \geq 3$ points in general position, there is a 1-rainbow cycle in G^{S}_X that contains the subpath $P^L_{n,1}$. Recall that the edge $\{1, n\}$ lies on the convex hull of X and therefore $P^L_{n,1} = P^R_{n,1}$.

To settle the base case $n = 3$ we take the cycle (S_1, S_3, S_2).

The following inductive construction is illustrated in Figure 5. For the induction step let C' be the rainbow cycle for the point set $X' := X \setminus \{n\} = [n-1]$, and let Q be the subpath $P^L_{n-1,1} = P^R_{n-1,1}$ of C'. We obtain the desired rainbow cycle C for the point set $X = [n]$ as follows: We remove all intermediate trees on the path Q from C', and we add the edge $\{1, n\}$ to all remaining trees, so that all these trees are spanning trees on the point set X. The resulting path in G^{S}_X starts at $S_1(X) = S_1(X') + \{1, n\}$ and ends at $T' := S_{n-1}(X') + \{1, n\}$. Note that T' is the first intermediate tree on the path $P^L_{n-1,n}$, so we continue the cycle C from T' along this path until we reach the star $S_n(X)$, and from there we complete the cycle C along the path $P^L_{n,1} = P^R_{n,1}$ back to the star $S_1(X)$. By construction, C contains the required subpath.

We now argue that C does not visit any spanning tree twice. Let $\mathcal{T}'$ denote the set of trees on C' except the intermediate trees on Q, and let $\mathcal{T}$ denote the set of trees on the paths $P^L_{n-1,n}$ and $P^L_{n,1}$ except the trees $S_{n-1}(X)$, T' and $S_1(X)$. Note that all trees in $\mathcal{T}'$ contain the edge $\{1, n\}$ and the point n has degree 1, whereas all trees in $\mathcal{T}$ do not contain the edge $\{1, n\}$ or the point n has degree at least 2. It follows that $\mathcal{T}' \cap \mathcal{T} = \emptyset$. The intermediate trees on the path $P^L_{n-1,n}$ do not contain the edge $\{1, 2\}$, whereas the intermediate trees on the path $P^L_{n,1}$ do contain this edge, so no two intermediate trees of these paths are the same. We conclude that C does not visit any tree twice.

By construction, along the cycle C every edge from the following set appears exactly once:

$$\begin{aligned} E_{X'} &\setminus \big\{\{1,2\},\{1,3\},\dots,\{1,n-1\}\big\} \\ &\cup \big\{\{n,2\},\{n,3\},\dots,\{n,n-1\}\big\} \\ &\cup \big\{\{1,2\},\{1,3\},\dots,\{1,n-1\},\{1,n\}\big\} \\ &= E_{X'} \cup \big\{\{n,2\},\{n,3\},\dots,\{n,n-1\}\big\} \cup \{1,n\} \;=\; E_X. \end{aligned}$$

This shows that C is a 1-rainbow cycle in G_X^{S}. □

3.2. **Proof of Theorem 2 (ii).** The next lemma explicitly describes all intermediate trees along the paths $P_{i,j}^L$ and $P_{i,j}^R$. It is an immediate consequence of the definition of the flip sequences given in the previous section.

A *caterpillar* is a tree that has the property that when removing all leafs, the remaining graph is a path. We refer to any path that is obtained from a caterpillar by removing a number of leafs as a *central path* of the caterpillar. Note that all other vertices not on a central path are leafs of degree 1. Consequently, specifying the degree sequence of the vertices on a central path of a caterpillar describes the caterpillar uniquely. Note that a caterpillar may have several different central paths, e.g. the caterpillar with central path (a,b,c,d) and degree sequence $(1,4,4,1)$ can also be described via the central path (b,c) and the degree sequence $(4,4)$.

Lemma 3. *Let $n \geq 3$. For any two points $i,j \in [n]$ and any $1 \leq t \leq n-2$, the t-th intermediate tree on the path $P_{i,j}^L$ from S_i to S_j is a caterpillar and (i,π_t^L,j) is a central path of the caterpillar with degree sequence $(n-1-t,2,t)$. An analogous statement holds for all intermediate trees on the path $P_{i,j}^R$.*

The next lemma asserts that the intermediate trees along any two paths obtained from our flip sequences are all distinct.

Lemma 4. *Let $n \geq 6$. For any two paths $P \in \{P_{i,j}^L, P_{i,j}^R\}$ and $P' \in \{P_{i',j'}^L, P_{i',j'}^R\}$ with $\{i,j\} \neq \{i',j'\}$, all intermediate trees on P and P' are distinct. Equivalently, P and P' are internally vertex-disjoint paths in the graph G_X^{S}.*

Proof. Let $\pi \in \{\pi_{i,j}^L, \pi_{i,j}^R\}$ be the ordering of points corresponding to the path P. We argue that for any intermediate tree on P, it is possible to uniquely reconstruct the points i and j which form the center of the stars S_i and S_j that are the end vertices of the path P. Let T be the t-th intermediate tree on P, $1 \leq t \leq n-2$. By Lemma 3, the point π_t has degree 2 in T, whereas all other points π_a, $a \neq t$, are leafs of degree 1 in T, and $\deg(i) + \deg(j) = n-1 \geq 5$. If T has two points of degree at least 3, those must be i and j and we are done. Otherwise T has only one point with degree at least 3, we assume w.l.o.g. that it is point i. Then we can determine point j as the unique point with distance exactly 2 from i in T. Specifically, by Lemma 3, i and j are connected via π_t in T, and all other points are neighbors of either i or j, so they have distance 1 or 3 from i. This completes the proof. □

The proof of Theorem 2 (ii) is split into three parts. We first construct r-rainbow cycles for even values of r (Proposition 5), then for odd values of r (Proposition 7), and we finally settle some remaining small cases (Proposition 8).

Proposition 5. *Let X be a set of $n \geq 6$ points in general position. For any $r = 1,2,\dots,\lfloor (n-1)/2 \rfloor$, there is a $2r$-rainbow cycle in G_X^{S}.*

In the proof we will use a decomposition of the complete graph on n vertices into $\lfloor (n-1)/2 \rfloor$ Hamilton cycles (and a perfect matching for even n, which will not be used in the proof, though). Such a decomposition exists by Walecki's theorem; see [Als08].

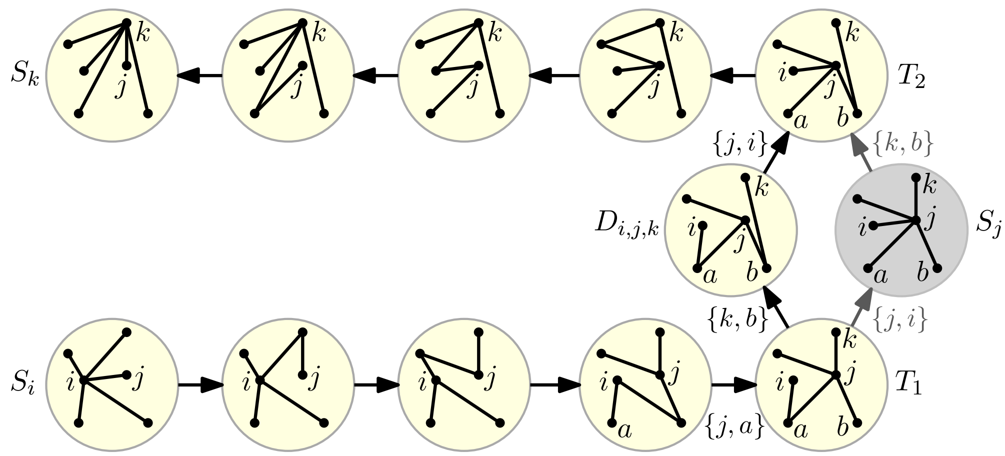

FIGURE 6. The path $P_{i,j} = P^L_{i,j}$ from S_i to S_j (bottom) and the path $P_{j,k} = P^R_{j,k}$ from S_j to S_k (top) and a detour around S_j via the detour tree $D_{i,j,k}$, in the case where $a \notin \{b,k\}$.

Proof. We apply Walecki's theorem to obtain a set $\mathcal{H}$ of r edge-disjoint Hamilton cycles in K_n, the complete graph on n vertices. We now consider the complete graph K_X on the point set $X = [n]$, and we map the Hamilton cycles in $\mathcal{H}$ onto K_X such that one Hamilton cycle $H_0 \in \mathcal{H}$ visits all points on the convex hull of X successively. We orient H_0 so that it visits the points on the convex hull in counter-clockwise order. If X is not in convex position, then there is a unique edge e on the convex hull of X that is not covered by H_0. If this edge e is contained in some other Hamilton cycle $H_1 \in \mathcal{H}$, then we orient H_1 so that the edge e is also traversed in counter-clockwise direction on the convex hull. Each of the remaining Hamilton cycles in $\mathcal{H}$ is oriented arbitrarily in one of the two directions. The union of these r oriented Hamilton cycles in K_X yields a directed graph with in-degree and out-degree equal to r at each point. We fix an arbitrary Eulerian cycle $\mathcal{E}$ in this graph. Note that $\mathcal{E}$ visits each point exactly r times, and it traverses all edges on the convex hull of X in counter-clockwise direction.

We define a directed closed tour C' in G^{S}_X, which possibly contains certain trees multiple times, by considering every triple of points (i,j,k) along $\mathcal{E}$. If $i \in R_{j,k}$, then we say that the triple (i,j,k) *takes a right-turn*, and if $i \in L_{j,k}$, then we say that the triple (i,j,k) *takes a left-turn*. If (i,j,k) takes a right-turn, then we add the path $P^L_{j,k}$ to C', and if (i,j,k) takes a left-turn, then we add the path $P^R_{j,k}$ to C'. Figure 6 shows an example of two concatenated paths. From Lemma 4 we know that the stars S_i, $i \in [n]$, are the only trees that are visited multiple times by C'. Specifically, each S_i is visited exactly r times by C'. Furthermore, every edge $\{i,j\} \in E_X$ appears on exactly $2r$ arcs of C', once on every path to S_i and once on every path to S_j. It follows that the tour C' has the $2r$-rainbow property.

To modify C' into a $2r$-rainbow cycle C we will use detours around most of the stars; see Figures 6 and 7. In the following we define the *detour tree* $D_{i,j,k}$ for those triples of points (i,j,k) on $\mathcal{E}$ where $\{j,k\}$ is not an edge of the convex hull of X. These detour trees are then used to replace all the stars of C' except a single occurrence of each S_j where j is on the convex hull. These replacements yield C.

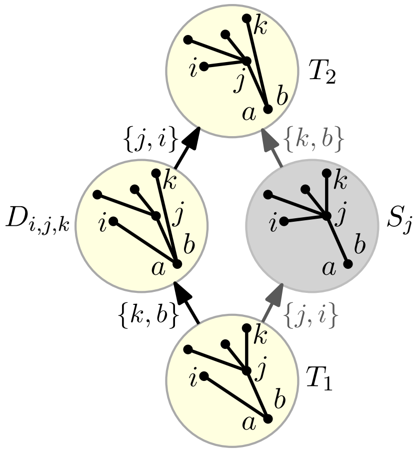

FIGURE 7. The detour trees $D_{i,j,k}$ for the cases $a = b$ (left) and $a = k$ (right).

To define $D_{i,j,k}$ consider the path $P_{i,j} \in \{P^L_{i,j}, P^R_{i,j}\}$ from S_i to S_j and the path $P_{j,k} \in \{P^L_{j,k}, P^R_{j,k}\}$ from S_j to S_k in C'. Let T_1 be the predecessor of S_j on $P_{i,j}$, and let T_2 be the successor of S_j on $P_{j,k}$. There are points $a, b \in [n]$ such that $S_j = T_1 - \{i, a\} + \{j, i\}$ and $T_2 = S_j - \{j, k\} + \{k, b\}$. The detour tree is defined as

$$D_{i,j,k} = T_1 - \{j, k\} + \{k, b\} = T_2 - \{j, i\} + \{i, a\}. \tag{2}$$

We denote the two relevant flips in this definition by $f_1 := (\{i, a\}, \{j, i\})$ and $f_2 := (\{j, k\}, \{k, b\})$. As $\{j, k\}$ does not lie on the convex hull of X, both half-planes $L_{j,k}$ and $R_{j,k}$ contain points from X. By our choice of the path from S_j to S_k based on the orientation of the triple (i, j, k), the points b and i lie in opposite half-planes. It follows that $b \neq i$ and that i, j, k, and b are four different points. Note also that the point a is different from i and j, but it may happen that $a = b$ or that $a = k$; see Figure 7.

Claim. The flip $f_2 = (\{j, k\}, \{k, b\})$ is a legal flip applicable to T_1 and the resulting graph $D_{i,j,k}$ is a plane spanning tree.

To establish this claim we check the following four properties:

- *$\{j, k\}$ is an edge of T_1.* T_1 contains all edges incident with j except the edge $\{i, j\}$ by Lemma 3. As $i \neq k$, the tree T_1 contains in particular the edge $\{j, k\}$.
- *$\{k, b\}$ is not an edge of T_1.* The only edge in T_1 that is not incident to j is the edge $\{i, a\}$. As $i \neq k$ and $i \neq b$, it follows that $\{k, b\}$ is not an edge of T_1.
- *$D_{i,j,k}$ is a spanning tree.* We only need to check that $D_{i,j,k}$ is connected, which is true as the two endpoints j and k of the removed edge $\{j, k\}$ are connected via the path (j, b, k) in $D_{i,j,k}$.
- *$D_{i,j,k}$ is plane.* The added edge $\{k, b\}$ does not cross any edges incident with j, otherwise it would not be the first added edge along $P_{j,k}$. It remains to show that $\{k, b\}$ does not cross the edge $\{i, a\}$ either. If $a = b$ or $a = k$, then there is no crossing and we are done, so we can assume that a is different from the four other points i, j, k, b. In the following we assume that the triple (i, j, k) takes a left-turn, the other case follows by symmetry. Recall from before that $b \in R_{j,k}$ and $i \in L_{j,k}$. Hence, the counter-clockwise order around j is $(b, k, i, \bar{k})$, where $\bar{k}$ denotes the antipodal direction of k. From the definition of $P^L_{i,j}$ and $P^R_{i,j}$ it follows that a is either the immediate predecessor or the immediate successor of i in the cyclic order of the points around j. If the counter-clockwise order around j is $(b, k, a, i, \bar{k})$ or $(b, k, i, a, \bar{k})$ then

there is a line through j which separates $\{k, b\}$ and $\{i, a\}$. In the remaining case a and b belong to $R_{i,j} \cap R_{j,k}$. By the definition of T_2 we also have $a \in R_{k,b}$ in this case, so the line through k and a separates $\{k, b\}$ and $\{i, a\}$.

This completes the proof of the claim.

Note that executing the flip operations f_1 and f_2 in the opposite order just changes the order in which the edges $\{j, i\}$ and $\{k, b\}$ appear, so the resulting tour C still has the $2r$-rainbow property. To show that C is a cycle, it remains to show that each detour tree $D_{i,j,k}$ is used only once in C. For this we first give an explicit description of the intermediate trees, which is an immediate consequence of the previous definitions.

Lemma 6. *Let $n \geq 6$. For any triple $(i, j, k) \in \mathcal{E}$, the intermediate tree $D_{i,j,k}$ defined in* (2) *has the following properties:*

(i) If $a \notin \{b, k\}$, then $D_{i,j,k}$ is a caterpillar and (a, j, b) is a central path of the caterpillar with degree sequence $(2, n-3, 2)$, where the unique leafs in distance 2 from j are the points i and k; see Figure 6.

(ii) If $a = b$, then $D_{i,j,k}$ is a caterpillar and $(j, a) = (j, b)$ is a central path of the caterpillar with degree sequence $(n-3, 3)$, where the unique leafs in distance 2 from j are the points i and k; see the left hand side of Figure 7.
Moreover, if (i, j, k) takes a left-turn, then the point $a = b$ lies within the sector $R_{i,j} \cap R_{j,k}$ and there is no point in the sector $R_{i,j} \cap R_{j,a}$. If (i, j, k) takes a right-turn, on the other hand, then an analogous statement holds with right and left half-planes interchanged.

(iii) If $a = k$, then $D_{i,j,k}$ is a caterpillar and $(j, b, a) = (j, b, k)$ is a central path of the caterpillar with degree sequence $(n-3, 2, 2)$, where k is the unique point in distance 2 from j, and i is the unique leaf incident with k; see the right hand side of Figure 7.

Since at most one of the triples (i, j, k) or (k, j, i) appears along $\mathcal{E}$, this lemma allows us to reconstruct the triple (i, j, k) from any given detour tree $D_{i,j,k}$. Only in the case $n = 6$ when the degree sequence of the central path of the caterpillar is $(3, 3)$ (case (ii) of the lemma), there is an ambiguity which of the two points of the central path is j. This ambiguity can be resolved by using the additional property mentioned in (ii) involving half-planes. It can be easily checked that if this condition holds for (i, j, k) as in the lemma, then it does not hold for $(x, a, y) = (x, b, y)$, where x and y are the unique leafs incident with j. Consequently, all detour trees $D_{i,j,k}$ included in C are distinct.

It remains to argue that each detour tree $D_{i,j,k}$ is distinct from all intermediate trees along any path $P_{i',j'} \in \{P^L_{i',j'}, P^R_{i',j'}\}$ from which C is built. By Lemma 3, the t-th intermediate tree T on $P_{i',j'}$ has a central path with degree sequence $(n-1-t, 2, t)$ for each $1 \leq t \leq n-2$. In cases (i) and (ii) of Lemma 6, comparing the degree sequences shows that $D_{i,j,k}$ must be different from T. In case (iii) there can only be a conflict if $t = 2$, as then we have $(n-1-t, 2, t) = (n-3, 2, 2)$. Matching the degree sequences in this case, we must have $(i', j') = (j, a) = (j, k)$. However, in the second intermediate spanning ($t = 2$) on $P_{i',j'} = P_{j,k}$ the point b has degree 1, whereas the point b has degree 2 in $D_{i,j,k}$. It follows that all detour trees $D_{i,j,k}$ are distinct from all intermediate trees.

This shows that C is indeed a $2r$-rainbow cycle in G^{S}_X, completing the proof of Proposition 5. $\square$

To construct a $(2r-1)$-rainbow cycle in G^{S}_X we slightly modify the construction from the previous proof. Specifically, we remove one Hamilton cycle from K_X before building the Eulerian cycle, which decreases the rainbow count by 2 for each edge from E_X. Instead we substitute the 1-rainbow cycle constructed in the proof of Theorem 2 (i), yielding a $(2r-1)$-rainbow cycle.

Proposition 7. *Let X be a set of $n \geq 6$ points in general position. For any $r = 2, 3, \ldots, \lfloor (n-1)/2 \rfloor$, there is a $(2r-1)$-rainbow cycle in G^{S}_X.*

Proof. The construction starts as in the proof of Proposition 5. We consider a set $\mathcal{H}$ of r edge-disjoint Hamilton cycles in the complete graph K_n, and we map the Hamilton cycles in $\mathcal{H}$ onto the complete graph K_X on the point set $X = [n]$ such that one Hamilton cycle H_0 visits all points on the convex hull of X successively. In addition, we perform the mapping so that H_0 contains the edges $\{n-1, n\}$, $\{n, 1\}$, and $\{1, 2\}$ (the first of these three edges does not necessarily lie on the convex hull, but the latter two edges do by our ordering of the points). We orient each of the Hamilton cycles as in the previous proof, so that all edges on the convex hull are oriented counter-clockwise. We now remove the Hamilton cycle H_0, yielding a set of directed Hamilton cycles $\mathcal{H}^- := \mathcal{H} \setminus \{H_0\}$, and build a Eulerian cycle $\mathcal{E}$ in this graph, which has in-degree and out-degree equal to $r - 1 \geq 1$ at each point. We fix one triple of the form $(i', 1, k')$ in $\mathcal{E}$. From the Eulerian cycle $\mathcal{E}$ we build a directed closed tour C' in G_X^{S} as in the previous proof, where for the special triple $(i', 1, k')$ we use the path $P_{1,k'}^R$, regardless of the orientation of this triple. Along the edges of this tour, every edge from E_X appears exactly $2r - 2$ times. We then modify C' into a cycle C by considering every triple $(i, j, k) \in \mathcal{E}$ except the special triple $(i', 1, k')$ and by replacing S_j by the corresponding detour tree $D_{i,j,k}$ if the edge $\{j, k\}$ does not lie on the convex hull of X (as before). For the special triple $(i', 1, k')$, we do not replace S_1. As in $\mathcal{H}^-$ no directed edge starting at 1 proceeds along an edge of the convex hull of X (H_0 uses the edge $\{1, 2\}$, and the edge $\{n, 1\}$ is oriented towards 1), all occurrences of S_1 in C' except for the single occurrence corresponding to the special triple $(i', 1, k')$ are replaced in C. Now let C^1 be the 1-rainbow cycle constructed as in the proof of Theorem 2 (i) starting with S_1. We replace the unique occurrence of S_1 in C by C^1 followed by $D_{n,1,k'}$, yielding a tour C^+; see Figure 8. We claim that C^+ is a $(2r-1)$-rainbow cycle in G_X^{S}. Clearly, C^+ has the $(2r-1)$-rainbow property, so we only need to show that C^+ is a cycle, i.e., no tree is visited more than once. By the arguments given in the proof of Proposition 5, it suffices to show that all trees on C^1 are distinct from the ones in $C \setminus \{S_1\}$.

We divide the cycle C^1 into segments according to its inductive construction in the proof of Theorem 2 (i). Specifically, for $c = 3, 4, \ldots, n$ we define

$$T_c := S_c([c]) + \{1, c+1\} + \{1, c+2\} + \cdots + \{1, n\};$$

see Figure 8. These are the spanning trees along which the cycle in the inductive construction described in Theorem 2 (i) is split in each step. We follow C^1 starting at S_1 and argue that each of the trees along the cycle is distinct from $C \setminus \{S_1\}$. The arguments are divided into cases (1)–(9), which are illustrated in Figure 8.

(1) The first tree on C^1 is S_1. As argued before, C contains only a single occurrence of S_1 which was replaced, so S_1 is unique in C^+.
(2) The tree T_3 is a caterpillar with central path $(1, 3, 2)$ with degree sequence $(n-2, 2, 1)$. This tree is different from any detour tree by Lemma 6. By Lemma 3, it can only be equal to the first intermediate tree on the path $P_{1,2}^L = P_{1,2}^R$. However, as the edge $\{1, 2\}$ is contained in H_0 and not in $\mathcal{E}$, the path $P_{1,2}^L$ is not part of C', and therefore T_3 does not occur in C.
(3) The successor of T_3 on C^1 is a caterpillar with central path $(1, 4, 2)$ with degree sequence $(n-3, 2, 2)$. By Lemma 6 this tree can only be equal to a detour tree $D_{i,j,k}$ as captured in case (iii) of the lemma, which would imply $j = 1$ and $k = 2$. However, no triple of the form $(i, j, k) = (i, 1, 2)$ is contained in $\mathcal{E}$. Moreover, by Lemma 3 this tree can only be equal to the second intermediate tree on the path $P_{1,2}$, which is not part of C as argued before.
(4) The tree T_c, $4 \leq c \leq n-2$, is a caterpillar with a central path $(1, c)$ with degree sequence $(n-c+1, c-1)$. It follows from Lemmas 3 and 6 that for $4 < c < n-2$ this tree is different from any tree on C. For $c = 4$, this tree can only be equal to a detour tree $D_{i,j,k}$ as captured in case (ii) of Lemma 6 where $a = b$, which would imply $(j, a) = (1, 4)$ and $\{i, k\} = \{2, 3\}$, or $(j, a) = (c, 1)$ and $\{i, k\} = \{5, 6\}$ for $n = 6$. In both cases we obtain a contradiction to the property that the point $a = b$ lies within the sector $R_{i,j} \cap R_{j,k}$ if the triple (i, j, k) takes a

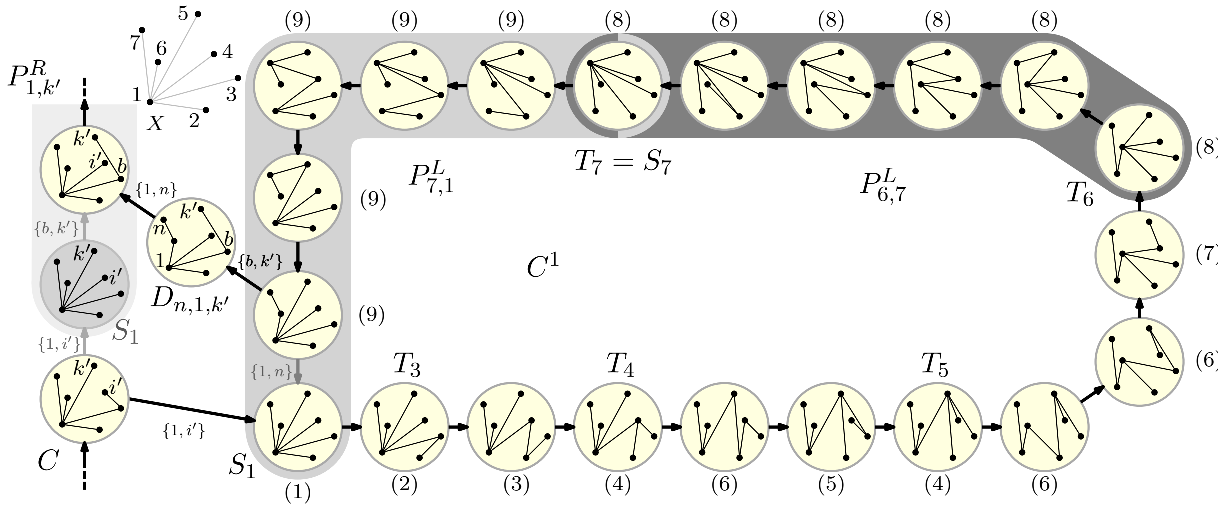

FIGURE 8. Modification of the $(2r-2)$-rainbow cycle C by including the 1-rainbow cycle C^1 to obtain a $(2r-1)$-rainbow cycle C^+ for a point set with $n = 7$ points.

left-turn, or in the sector $L_{i,j} \cap L_{j,k}$ if the triple takes a right-turn. For $c = n-2$, a similar reasoning shows that T_c is different from any trees on C.

(5) The predecessor of T_c, $5 \le c \le n-2$, is a caterpillar with a central path of the form $(1, c, x)$ with degree sequence $(n-c+1, c-2, 2)$, which by Lemmas 3 and 6 is distinct from any tree on C.

(6) For $4 \le c \le n-2$, all trees strictly between T_i and the predecessor of T_{c+1} on C^1, have diameter at least 5. Specifically, the k-th successor of T_i contains a path of the form $(n, 1, c+1, x, c, y)$. By Lemmas 3 and 6 all trees on C have diameter at most 4.

(7) The predecessor of T_{n-1} on C^1 has a central path of the form $(1, n-1, x)$ with degree sequence $(2, n-3, 2)$, where n is the unique leaf incident with 1 and $n-2$ is the unique leaf incident with x. By Lemma 6 this tree can only be equal to a detour tree $D_{i,j,k}$ as captured in case (i) of the lemma, which would imply $(i,j,k) = (n-2, n-1, n)$ or $(i,j,k) = (n, n-1, n-2)$. However, as the edge $\{n-1, n\}$ is not in $\mathcal{E}$, this triple does not occur in $\mathcal{E}$ either, so this detour tree is not in C.

(8) The path from T_{n-1} and $T_n = S_n(X)$ is by definition the path $P^L_{n-1,n}$ with the first tree $S_{n-1}(X \setminus \{n\})$ removed, and as the edge $\{n-1, n\}$ is not in $\mathcal{E}$, all trees on this path are distinct from the ones on C. The spanning tree $T_n = S_n$ is not in C, as the only edge on the convex hull of X incident with n, if such an edge is present in $\mathcal{H}^-$ at all, is oriented towards n, so the next edge along $\mathcal{E}$ is not a convex hull edge.

(9) The path from $T_n = S_n$ to the predecessor of S_1 is exactly $P^L_{n,1}$ with the last tree S_1 removed, and as the edge $\{n, 1\}$ is not in $\mathcal{E}$, all trees on this path are distinct from the ones on C. The detour tree $D_{n,1,k'}$ is distinct from all trees on C as argued in the proof of Proposition 5.

This completes the proof of Proposition 7. □

Proposition 8. *For any point set X with $n = 4$ points in general position, there is a 2-rainbow cycle in G^{S}_X. For any point set X with $n = 5$ points in general position and any $r \in \{2, 3, 4\}$, there is an r-rainbow cycle in G^{S}_X.*

The rainbow cycles for proving Proposition 8 can be constructed explicitly using slight variants of the methods described in the preceding proofs. This proof is deferred to the appendix.

Proof of Theorem 2 (ii). Combine Proposition 5, Proposition 7 and Proposition 8. □

4. Matchings

In this section we consider a set of $n = 2m$ points in convex position labeled clockwise by $1, 2, \dots, n$. Without loss of generality we assume that the points are distributed equidistantly on a unit circle centered at the origin. We use $\mathcal{M}_m$ to denote the set of all non-crossing perfect matchings with m edges on these points. The graph $G_m^{\mathtt{M}}$ has $\mathcal{M}_m$ as its vertex set, and an arc (M, M') between any two matchings M and M' that differ in exchanging two edges $e = \{a, b\} \in M$ and $f = \{c, d\} \in M$ for the edges $e' = \{a, c\}$ and $f' = \{b, d\} \in M'$; see Figure 2 (b). We refer to this operation as a *flip*. Furthermore, we label the arc (M, M') with the edges e' and f', so an arc is labeled with the edges that enter the matching in this flip. The set of arc labels of $G_m^{\mathtt{M}}$ is $E_m := \{\{i, j\} \mid i, j \in [n] \text{ and } j - i \text{ is odd}\}$. In this definition, the difference $j - i$ must be odd so that an even number of points lies on either side of the edge $\{i, j\}$. Every arc of $G_m^{\mathtt{M}}$ carries two such labels, and we think of these labels as colors. An r-rainbow cycle in $G_m^{\mathtt{M}}$ is a directed cycle along which every label in E_m appears exactly r times. As every arc is labeled with two edges, an r-rainbow cycle has length $r|E_m|/2 = rm^2/2$. The number of vertices of $G_m^{\mathtt{M}}$ is the m-th Catalan number $\frac{1}{m+1}\binom{2m}{m}$.

The following theorem summarizes the results of this section.

Theorem 9. *The flip graph of non-crossing perfect matchings $G_m^{\mathtt{M}}$, $m \geq 2$, has the following properties:*

(i) If m is odd, then $G_m^{\mathtt{M}}$ has no 1-rainbow cycle.
(ii) If $m \in \{6, 8, 10\}$, then $G_m^{\mathtt{M}}$ has no 1-rainbow cycle.
(iii) If $m \in \{2, 4\}$, then $G_m^{\mathtt{M}}$ has a 1-rainbow cycle, and if $m \in \{6, 8\}$, then $G_m^{\mathtt{M}}$ has a 2-rainbow cycle.

4.1. Proof of Theorem 9 (i).

Proof. A 1-rainbow cycle must have length $m^2/2$. For odd m, this number is not integral, so there can be no such cycle. □

4.2. Proof of Theorem 9 (ii).

In view of part (i) of Theorem 9, we assume for the rest of this section that the number of matching edges m is even.

The following definitions are illustrated in Figure 9. The *length* of a matching edge $e \in M$, denoted by $\ell(e)$, is the minimum number of other edges from M that lie on either of its two sides. Consequently, a matching edge on the convex hull has length 0, whereas the maximum possible length is $(m-2)/2$, so there are $m/2$ different edge lengths. Note that E_m contains exactly $n = 2m$ edges of each length.

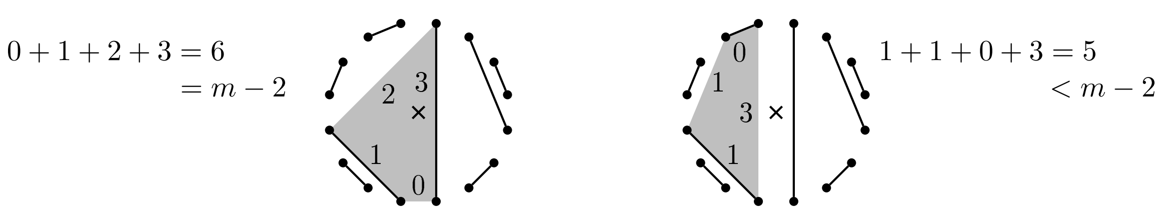

Figure 9. Examples of a centered 4-gon (left) and a non-centered 4-gon (right) for $m = 8$ matching edges. The numbers are the edge lengths.

We call a convex quadrilateral formed by four edges from E_m a *centered 4-gon*, if the sum of the edge lengths of the quadrilateral is $m - 2$. Note that this is the maximum possible value. We refer to a flip involving a centered 4-gon as a *centered flip*. Equivalently, a flip is centered if and only if the corresponding 4-gon contains the origin. Note that all flips in the rainbow cycle in Figure 2 (b) are centered flips. This is in fact not a coincidence, as shown by the following lemma.

Lemma 10. *All flips along an r-rainbow cycle in $G_m^{\mathtt{M}}$ must be centered flips.*

Proof. E_m contains exactly $n = 2m$ edges of each length $0, 1, \ldots, (m-2)/2$. Along an r-rainbow cycle C, exactly rn edges of each length appear and exactly rn edges of each length disappear. Consequently, the average length of all edges that appear or disappear along C is $(m-2)/4$. By definition, in a centered flip the average length of the four edges involved in the flip is exactly the same number $(m-2)/4$; whereas for a non-centered flip, it is strictly smaller. Therefore, C must not contain any non-centered flips. □

Lemma 10 suggests to restrict our search for rainbow cycles to the subgraph of $G_m^{\mathtt{M}}$ obtained by considering only arcs that represent centered flips. We denote this subgraph of $G_m^{\mathtt{M}}$ on the same vertex set $\mathcal{M}_m$ by H_m. This graph is shown in Figure 10 for $m = 6$ matching edges.

Proof of Theorem 9 *(ii).* The proof for the cases $m = 8$ and $m = 10$ is computer-based, and uses exhaustive search for a 1-rainbow cycle in each connected component of H_m. However, we are unable to prove this by hand. We proceed to show that there is no 1-rainbow cycle in H_m for $m = 6$. Unlike all our other non-existence proofs, this one is not a simple parity argument, but involves structural considerations.

Suppose there exists a 1-rainbow cycle C in H_m for $m = 6$. Clearly, C has length $m^2/2 = 18$. The graph H_6 has eight connected components; see Figure 10. Five of these components are trees, and do not contain any cycles, and three of them contain cycles. From those three cyclic components, two are isomorphic to each other and differ only by rotation of the matchings by $\pi/6$, so only one of the two is shown at the bottom left of Figure 10. The only cycles in these two components are of length 8 or 12, so they do not contain the desired rainbow cycle. It therefore remains to show that the third cyclic component F shown at the top right of Figure 10 has no rainbow cycle. This component has a cycle B of length 12 containing all matchings with a single edge of length 2. We refer to this cycle as the *base cycle* of F; it is drawn in the center of F in the figure. Removing the base cycle B from F leaves three other cycles of length 12, that differ only by rotation of the matchings in them. We call these three cycles *satellite cycles*, and we refer to the edges between the base cycle and the satellite cycles as *spokes*. Each satellite cycle is attached with 4 spokes to the base cycle, and the spokes are spaced equidistantly along both cycles, and no two spokes share any vertices.

There are three different types of centered 4-gons involved in the flips in F. Each type is characterized by the cyclic sequence of edge lengths of the 4-gon, where mirroring counts as the same 4-gon (this corresponds to reversing the sequence of edge lengths). These types are shown in Figure 11, and they denoted by a, b, and c.

In Figure 10, the edges corresponding to each flip type a, b, and c are drawn solid gray, solid black and dashed black, respectively. Note that all flips along the base cycle are of type a, all flips along spokes are of type b, and along each satellite cycle, one flip of type a alternates with two flips of type c.

Note that the set E_m of 36 edges that appear along C contains exactly 12 edges of each length from $\{0, 1, 2\}$. Since a flip of type a does not involve any edge of length 1, and flips b and c involve exactly two edges of length 1, we must perform in total 12 flips of types b or c along C (as we must have $24 = 2 \cdot 12$ appearances or disappearances of an edge of length 1), and hence exactly 6 flips of type a. Since the edges corresponding to flips of types b and c form stars in F, it follows that every sequence of three consecutive flips along C must contain at least one flip of type a. Combining these two observations shows that from the 18 flips along C, exactly every third must be of type a.

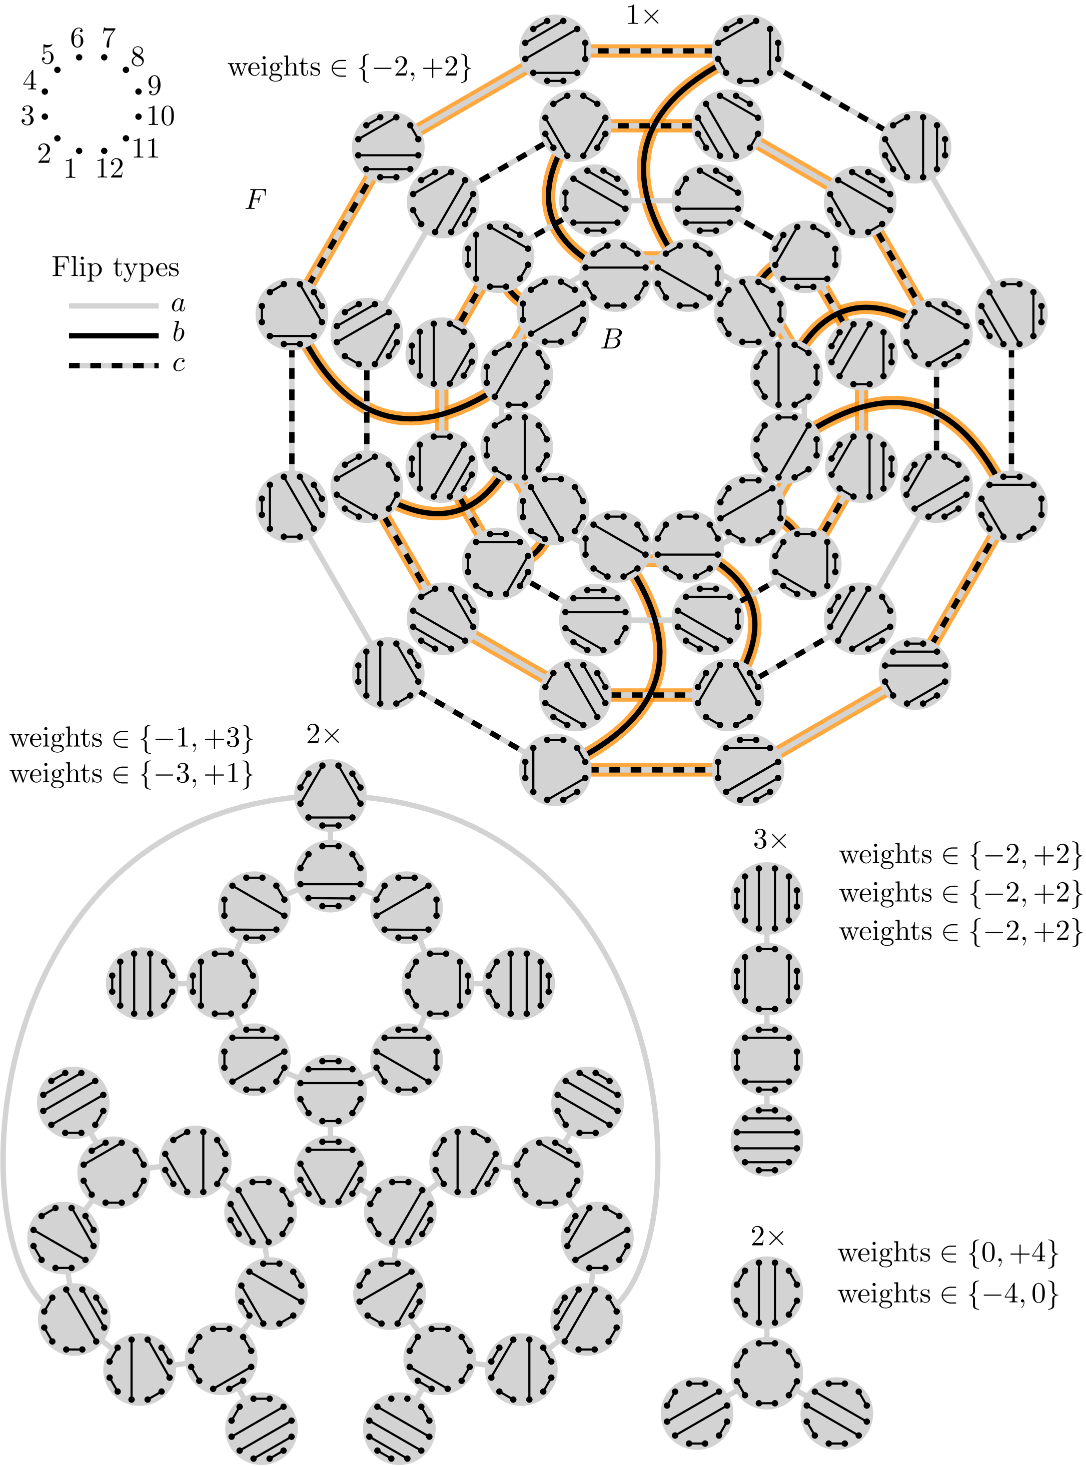

FIGURE 10. Illustration of the graph $H_6 \subseteq G_6^{\mathtt{M}}$. Some components of this graph are isomorphic to each other and differ only by rotation of the matchings by multiples of $\pi/6$. Only one representative for each component is shown, together with its multiplicity. The total number of matchings is the 6th Catalan number 132. The 2-rainbow cycle constructed in the proof of Theorem 9 (iii) is highlighted in the component F.

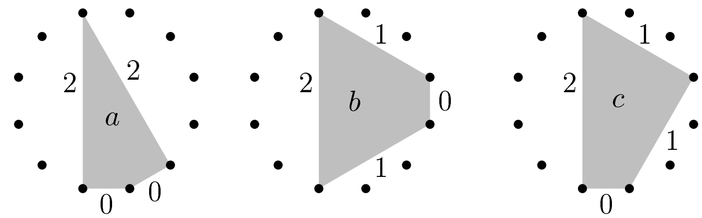

FIGURE 11. The three different types of centered 4-gons involved in the flips in F.

Therefore, C can be partitioned into six paths of length 3 such that each path either has the form $\beta := (b, a, b)$ or $\sigma := (c, a, c)$. A path of type β moves from a satellite cycle along a spoke to the base cycle, then uses one edge of the base cycle, and then returns via a spoke to another satellite cycle. A path of type σ moves along three consecutive edges of a satellite cycle, between the end vertices of two spokes. The rainbow cycle C can therefore be described by a cyclic sequence of 6 symbols from $\{\beta, \sigma\}$. In this sequence, any β must be followed by σ, as we must not traverse the same spoke twice. Moreover, at most three σ symbols can appear consecutively, otherwise we would close a cycle of length 12 along a satellite cycle. Without loss of generality, we may assume that the maximum length substring of consecutive σ symbols comes first in this sequence. This leaves the following possible patterns: $(\sigma, \beta, \sigma, \beta, \sigma, \beta)$, $(\sigma, \sigma, \beta, \sigma, \sigma, \beta)$ and $(\sigma, \sigma, \sigma, \beta, \sigma, \beta)$. A straightforward case analysis shows that for any cycle C following one of those patterns, one of the matching edges of length 2 appears and disappears twice, rather than only once. Consequently, F and H_6 contain no 1-rainbow cycle. $\square$

4.3. **Proof of Theorem 9 (iii).**

Proof. There are two non-crossing matchings with $m = 2$ edges, connected by two arcs in $G_m^{\mathtt{M}}$ that form a 1-rainbow cycle. For $m = 4$, a 1-rainbow cycle in $G_m^{\mathtt{M}}$ is shown in Figure 2 (b).

For $m = 6$, a 2-rainbow cycle in $G_m^{\mathtt{M}}$ can be constructed using the path P of length 6 between matchings M and M' in $G_m^{\mathtt{M}}$ depicted in Figure 12.

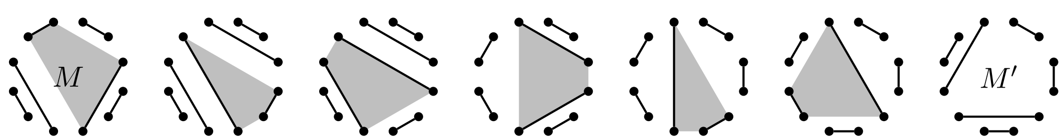

FIGURE 12. Definition of path P in $G_6^{\mathtt{M}}$ from M to M'.

The gray areas in the figure highlight the quadrilaterals involved in each flip from left to right. Note that M' differs from M by a clockwise rotation by an angle of $\alpha := 2\pi/6$. It is easy to check that repeating this flip sequence six times, rotating all flips by an angle of $\alpha \cdot i$ for $i = 0, 1, \ldots, 5$, yields a cycle C in $G_m^{\mathtt{M}}$. This cycle is highlighted in Figure 10. To verify that C is indeed a 2-rainbow cycle, consider the 12 matching edges that appear along the path P and are shown in Figure 13.

In the figure we differentiate the different lengths of the matching edges by three different line styles solid, dashed or dotted. Note that there are four edges from each length. It is straightforward to check that rotating this figure by $\alpha \cdot i$ for $i = 0, 1, \ldots, 5$ covers each matching edge from E_m exactly twice. Consequently, C is a 2-rainbow cycle in $G_6^{\mathtt{M}}$.

For $m = 8$, a 2-rainbow cycle in $G_m^{\mathtt{M}}$ can be constructed using the path P of length 8 between matchings M and M' depicted in Figure 14.

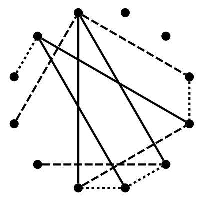

FIGURE 13. The 12 matching edges that appear along the path P in $G_6^{\mathtt{M}}$.

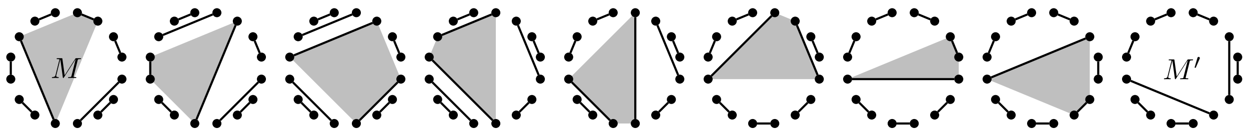

FIGURE 14. Definition of path P between matchings M and M' in $G_8^{\mathtt{M}}$.

Note that M' differs from M by a counter-clockwise rotation by an angle of $\alpha := 2\pi/8$. It is easy to check that repeating this flip sequence eight times, rotating all flips by an angle of $\alpha \cdot i$ for $i = 0, 1, \ldots, 7$, yields a cycle C in $G_m^{\mathtt{M}}$. To verify that C is indeed a 2-rainbow cycle, consider the 16 matching edges that appear along the path P shown in Figure 15.

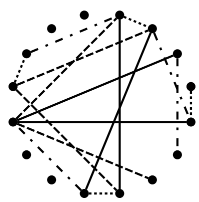

FIGURE 15. The 16 matching edges that appear along the path P in $G_8^{\mathtt{M}}$.

The different lengths of the matching edges are visualized by four different line styles solid, dashed, dash dotted, and dotted. Note that there are four edges from each length. It is straightforward to check that rotating this figure by $\alpha \cdot i$ for $i = 0, 1, \ldots, 7$ covers each matching edge from E_m exactly twice. Consequently, C is a 2-rainbow cycle in $G_8^{\mathtt{M}}$. $\square$

4.4. **Structure of the graph H_m.** In this section we prove that the graph H_m has at least $m-1$ connected components (Theorem 13 below).

The following definitions and Lemma 11 are illustrated in Figure 16. Consider a matching $M \in \mathcal{M}_m$ and one of its edges $e \in M$, and let i and j be the endpoints of e so that the origin lies to the right of the ray from i to j. We define the *sign of the edge* e as

$$\operatorname{sgn}(e) := \begin{cases} +1 & \text{if } i \text{ is odd,} \\ -1 & \text{if } i \text{ is even.} \end{cases}$$

Moreover, we define the *weight* of the matching M as

$$w(M) := \sum_{e \in M} \operatorname{sgn}(e) \cdot \ell(e).$$

Note that rotating a matching by π/m changes the weight by a factor of -1.

The following lemma is an immediate consequence of these definitions.

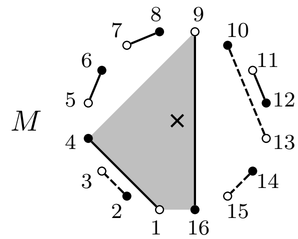

$$\begin{aligned} w(M) &= +(1+0+0+3+0) \\ &\quad -(0+1+0) \\ &= 3 \end{aligned}$$

FIGURE 16. Illustration of the weight of a matching with $m = 8$ edges. Edges with sign $+1$ are drawn solid, edges with sign -1 are drawn dashed. Applying the flip indicated in the figure changes the weight by $-(1+3) - (2+0) = -6 = -(m-2)$.

Lemma 11. *A centered 4-gon has two edges with positive sign and two edges with negative sign, and the pairs of edges with the same sign are opposite to each other. Consequently, applying a centered flip to any matching from $\mathcal{M}_m$ changes its weight by $-(m-2)$ if the two edges with negative sign appear in this flip, or by $+(m-2)$ if the two edges with positive sign appear in this flip. Flips of these two kinds must alternate along any sequence of centered flips.*

Proof. The first part of the lemma is an immediate consequence of the definitions given before. To see that flips that change the weight by $-(m-2)$ or $+(m-2)$ must alternate along any sequence of centered flips, note that in any matching $M \in \mathcal{M}_m$, all edges that are visible from the origin have the same sign, and any flip must change this sign. $\square$

The next lemma shows that the weight of a matching lies in a specific interval.

Lemma 12. *Given any matching $M \in \mathcal{M}_m$, we have*

$$w(M) \in [-(m-2), m-2] := \big\{ -(m-2), -(m-2)+1, \ldots, m-3, m-2 \big\}.$$

For any integer c in this set, there is a matching $M \in \mathcal{M}_m$ with $w(M) = c$.

The weights of all matchings with $m = 6$ edges are shown in Figure 10.

Proof. We fix a matching $M \in \mathcal{M}_m$ throughout the proof. For every odd $k \in [n]$ we consider the ray r_k from the origin $(0,0)$ through the point k; see Figure 17. This yields a set $C(r_k)$ of points $x = r_k \cap e$ where the ray r_k crosses any matching edge $e \in M$ in its interior. We define the *sign* of this crossing point x as the sign of the matching edge involved, i.e., $\mathrm{sgn}(x) := \mathrm{sgn}(e)$. Moreover, we define the *weight* of the ray r_k as the sum of signs of all crossing points from $C(r_k)$ along the ray, i.e., $w(r_k) := \sum_{x \in C(r_k)} \mathrm{sgn}(x)$. Let C be the union of all these crossing points between rays and matching edges.

Note that the number of rays that cross a fixed edge $e \in M$ is equal to $\ell(e)$, implying that

$$w(M) = \sum_{x \in C} \mathrm{sgn}(x) = \sum_{k \in [n] \text{ odd}} w(r_k). \tag{3}$$

We claim that if we follow any ray r_k, then the signs of any two consecutive crossing points x and y along the ray alternate, i.e., $\mathrm{sgn}(x) + \mathrm{sgn}(y) = 0$. To see this let $\{i, j\}$ and $\{p, q\}$ denote the edges from M causing the crossings x and y, respectively, such that i and p lie to the left of the ray r_k, and j and q lie to the right. Moreover, let A be the set of points from $[n]$ between i and p, and let B be the points between j and q; see Figure 17. Observe that all points from A must be matched within the set, and the same holds for all points within B. This is because if there was a matching edge between A and B, then it would cross r_k between x and y. It follows that $|A|$ and $|B|$ are

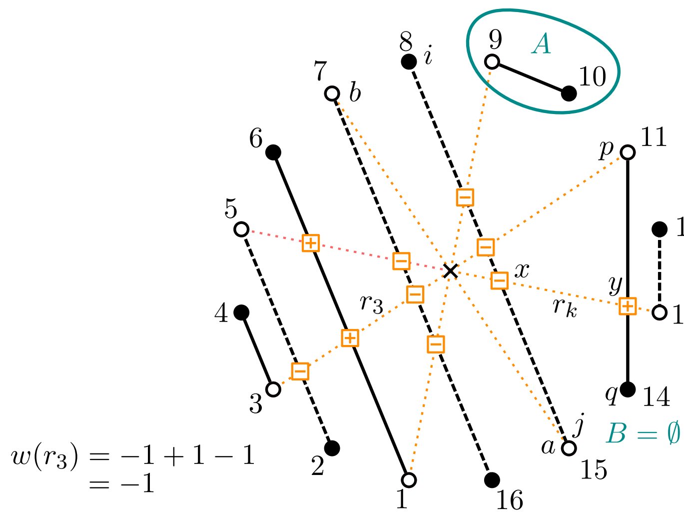

FIGURE 17. Illustration of the proof of Lemma 12.

even, and hence the distance between i and p, and the distance between j and k are both odd. This implies that the edges $\{i, j\}$ and $\{p, q\}$ have opposite signs, proving the claim.

An immediate consequence of this claim is that for any ray r_k we have $w(r_k) \in \{-1, 0, 1\}$. Moreover, there are at least two matching edges visible from the origin and therefore there are two odd points a and b for which r_a and r_b do not cross any matching edges, i.e., $w(r_a) = w(r_b) = 0$. It follows that two of the summands on the right hand side of (3) are 0, and the remaining $m-2$ summands are from the set $\{-1, 0, 1\}$. This proves the first part of the lemma.

To prove the second part of the lemma, consider a matching that has exactly two edges whose lengths sum up to c that are both visible from the origin $(0, 0)$ and $m-2$ matching edges of length 0. $\square$

Motivated by Lemma 11 and Lemma 12, we partition the set of all matchings $\mathcal{M}_m$ into sets $\mathcal{M}_{m,c}$, $c \in [-(m-2), (m-2)]$, where $\mathcal{M}_{m,c}$ contains all matchings with weight exactly c. Moreover, we define $\mathcal{M}^+_{m,c} := \mathcal{M}_{m,c} \cup \mathcal{M}_{m,c-(m-2)}$ for $c = 0, 1, \ldots, m-2$. These two lemmas imply the following structural result about the graph H_m.

Theorem 13. *For any even $m \geq 2$, the subgraph H_m of $G^{\mathtt{M}}_m$ that uses only centered flips has no edges between any two partition classes $\mathcal{M}^+_{m,c}$, $c = 0, 1, \ldots, m-2$, and therefore at least one connected component in each partition class, in total at least $m-1$ connected components.*

As Figure 10 shows, the subgraph of H_m induced by a partition class $\mathcal{M}^+_{m,c}$ is not necessarily connected, i.e., the number of connected components of H_m may exceed $m-1$. For instance, the subgraph of H_6 induced by the partition class $\mathcal{M}^+_{6,2}$ has four connected components and the total number of components of H_6 is eight.

We observed empirically for $m \in \{2, 4, 6, 8\}$ that for any $c \in [-(m-2), m-2]$, the number of matchings in $\mathcal{M}_{m,c}$ is given by

$$|\mathcal{M}_{m,c}| = \begin{cases} 2 & \text{if } c = 0, \\ N_1(m, |c|+1)/2 & \text{if } |c| \geq 1, \end{cases} \tag{4}$$

where $N_r(m,k)$ are the *generalized Narayana numbers*, defined as

$$N_r(m,k) = \frac{r+1}{m+1}\binom{m+1}{k}\binom{m-r-1}{k-1}$$

for integers $r \geq 0$ and $0 \leq k \leq m-r$. The quantity $N_r(m,k)$ counts Dyck paths in the integer lattice $\mathbb{Z}^2$ starting at the origin with m upsteps $(+1,+1)$ and $m-r$ downsteps $(+1,-1)$ with exactly k peaks. The Dyck path property means that such a path never moves below the abscissa.

Unfortunately, we are not able to prove (4) in general. Proving this relation would allow us to exactly compute the cardinalities of the partition classes $\mathcal{M}^+_{m,c}$ referred to in Theorem 13.

5. Permutations

In this section, we consider the set of permutations Π_n of $[n]$. We specify a permutation $\pi \in \Pi_n$ as $\pi = (\pi(1), \pi(2), \ldots, \pi(n))$. The graph G_n^{P} has Π_n as its vertex set, and an edge $\{\pi, \rho\}$ between any two permutations π and ρ that differ in exactly one transposition between the entries at positions i and j; see Figure 2 (c). We label the edge $\{\pi, \rho\}$ of G_n^{P} with the transposition $\{i,j\}$, and we think of these labels as colors. A 1-rainbow cycle in G_n^{P} is an undirected cycle along which every transposition appears exactly once, so it has length $\binom{n}{2}$. In this section we only consider 1-rainbow cycles, and we simply refer to them as rainbow cycles. Note that the number of vertices of G_n^{P} is $n!$.

The following theorem summarizes the results of this section.

Theorem 14. *The flip graph of permutations G_n^{P}, $n \geq 2$, has the following properties:*

(i) If $\lfloor n/2 \rfloor$ is odd, then G_n^{P} has no rainbow cycle.
(ii) If $\lfloor n/2 \rfloor$ is even, then G_n^{P} has a rainbow cycle.

Proof. The graph G_n^{P} is bipartite, since the parity changes in each step, so a cycle of length $\binom{n}{2}$ cannot exist when this number is odd, which happens exactly when $\lfloor n/2 \rfloor$ is odd. This proves (i).

To prove (ii) we assume that $\lfloor n/2 \rfloor$ is even, i.e., $n = 4\ell$ or $n = 4\ell + 1$ for some integer $\ell \geq 1$. We prove these cases by induction. As the graph G_n^{P} is vertex-transitive, it suffices to specify a sequence of $\binom{n}{2}$ transpositions that yields a rainbow cycle. We refer to such a sequence as a *rainbow sequence for* Π_n. A rainbow sequence of transpositions can be applied to any vertex in G_n^{P}, yielding a rainbow cycle. To settle the induction base $\ell = 1$, consider the rainbow sequence

$$R_4 := \big(\{1,2\}, \{3,4\}, \{2,3\}, \{1,4\}, \{2,4\}, \{1,3\}\big).$$

Applying R_4 to the permutation 1234 yields $C_4 := (1234, 2134, 2143, 2413, 3412, 3214)$, a rainbow cycle in G_4^{P}, where we omit brackets and commas in denoting these single-digit permutations. Note that applying the transposition $\{1,3\}$ to the last permutation in C_4 yields the first one. This rainbow cycle is depicted in Figure 18 (a).

We can interpret every transposition in a rainbow sequence for Π_n as an edge in K_n, yielding an ordering $1, 2, \ldots, \binom{n}{2}$ of the edges of K_n. In Figure 18, these edge orderings are shown in the top row. Whether an ordering of the edges of K_n corresponds to a valid rainbow sequence for Π_n or not can be decided as follows: Without loss of generality we start at the identity permutation $\mathrm{id}_n := (1, 2, \ldots, n)$, we apply the transpositions given by the edge ordering one after the other, checking that each permutation is encountered at most once and that the final permutation is again id_n.

For the induction step, we assume that we are given a rainbow sequence R_n for Π_n, $n = 4\ell$, and construct rainbow sequences for $\Pi_{n+1} = \Pi_{4\ell+1}$ and for $\Pi_{n+4} = \Pi_{4(\ell+1)}$. To this end we consider the effect of replacing a transposition $t = \{i,j\}$, $i < j$, in R_n by the sequence of three transpositions $\widehat{t}(n+1) := (\{i, n+1\}, \{i,j\}, \{j, n+1\})$. Note that the effect of this modification on the entries at

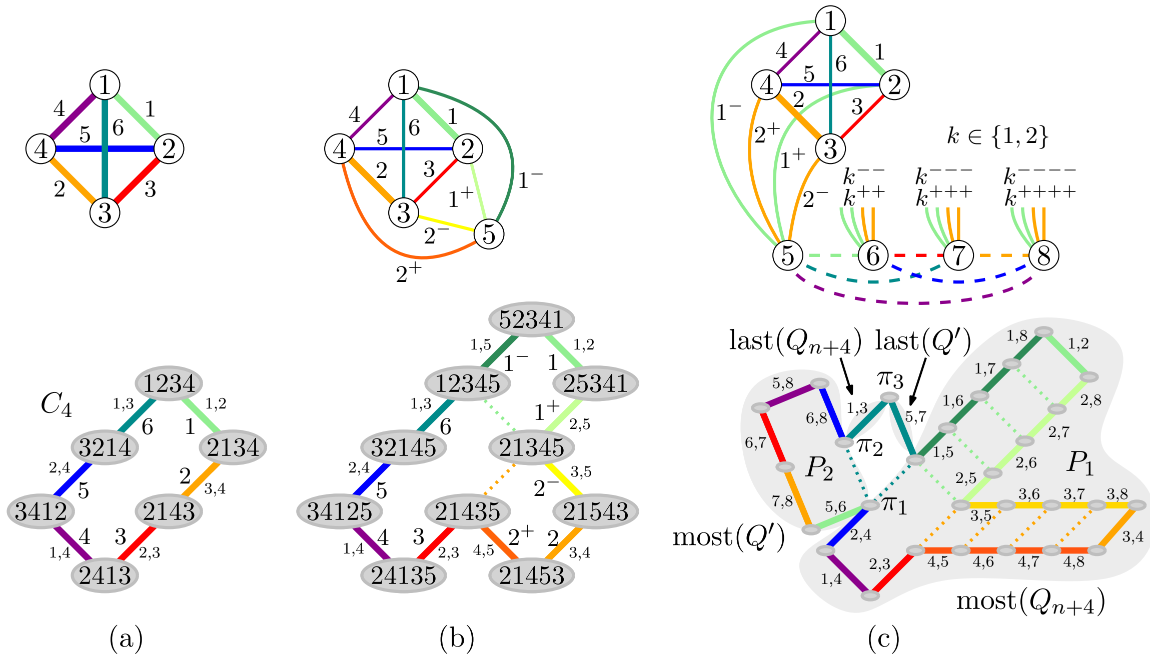

FIGURE 18. Illustration of the proof of Theorem 14: (a) induction base for $n = 4$, (b) induction step $n \to n+1$, (c) induction step $n \to n+4$ (where $n = 4\ell$). The rainbow cycles are shown in the bottom row, the corresponding edge orderings of K_n in the top row. In the top parts of (b) and (c), the bold edges mark the transpositions $t_i = \{2i-1, 2i\}$, $i = 1, 2, \ldots, n/2$, and in the edge orderings the edges labeled with $-$ and $+$ superscripts appear consecutively in the order k^-, k, k^+ or $k^-, k^{--}, k^{---}, k^{----}, k, k^{++++}, k^{+++}, k^{++}, k^+$, respectively, for $k \in \{1, 2\}$. Single-digit permutations and transpositions are denoted without brackets and commas.

positions i, j and $n+1$ of the permutation is the same, the only difference is that in the modified sequence also the transpositions $\{i, n+1\}$ and $\{j, n+1\}$ are used. This simple observation is the key to the following inductive construction.

We first assume that $n = 4\ell$, $\ell \geq 1$, and given a rainbow sequence R_n for Π_n we construct a rainbow sequence R_{n+1} for $\Pi_{n+1} = \Pi_{4\ell+1}$. For this we consider the transpositions $t_i := \{2i-1, 2i\}$, $i = 1, 2, \ldots, n/2$, in R_n, and to construct R_{n+1} we replace each t_i by the triple $\widehat{t}_i(n+1)$; see Figure 18 (b). This yields a sequence R_{n+1} of length $\binom{n}{2} + 2 \cdot \frac{n}{2} = \binom{n+1}{2}$. To see that R_{n+1} is indeed a rainbow sequence, consider applying it to the identity permutation id_{n+1}. The permutations visited along the cycle are exactly the ones encountered when applying R_n to id_n with the element $n+1$ added at the last position, plus n additional permutations where the symbol $n+1$ appears exactly once at each of the positions $1, 2, \ldots, n$.

We now assume that $n = 4\ell$, $\ell \geq 1$, and given a rainbow sequence R_n for Π_n we construct a rainbow sequence R_{n+4} for $\Pi_{n+4} = \Pi_{4(\ell+1)}$. For this we consider the transpositions $t_i := \{2i-1, 2i\}$, $i = 1, 2, \ldots, n/2$, in R_n. We assume without loss of generality that the last transposition in R_n is not one of the t_i (otherwise shift R_n cyclically). We then define $Q_n := R_n$ and construct auxiliary sequences Q_{n+1}, Q_{n+2}, Q_{n+3}, Q_{n+4} as follows: For $k \in \{1, 2, 3, 4\}$, Q_{n+k} is obtained from Q_{n+k-1} by replacing each t_i, $i = 1, 2, \ldots, n/2$, in Q_{n+k-1} by the triple $\widehat{t}_i(n+k)$. Now let Q' be the sequence of transpositions on the elements $\{n+1, n+2, n+3, n+4\}$ obtained from R_4 by adding n to all elements. The sequence R_{n+4} is constructed by suitably interleaving Q_{n+4} and Q'. Specifically, for any sequence $x = (x_1, x_2, \ldots, x_{k+1})$ we define $\mathrm{most}(x) := (x_1, x_2, \ldots, x_k)$ and $\mathrm{last}(x) := x_{k+1}$, so $x = (\mathrm{most}(x), \mathrm{last}(x))$. We then define

$$R_{n+4} := \big(\mathrm{most}(Q_{n+4}), \mathrm{most}(Q'), \mathrm{last}(Q_{n+4}), \mathrm{last}(Q')\big);$$

see Figure 18 (c). By construction, the last transposition in R_n is also the last transposition in Q_{n+4}. This yields a sequence R_{n+4} of length $\binom{n}{2} + 8 \cdot \frac{n}{2} + 6 = \binom{n+4}{2}$. We now verify that R_{n+4} is indeed a rainbow sequence for Π_{n+4}, by applying it to the identity permutation id_{n+4}. Let P_1 be the set of permutations encountered when applying $\mathrm{most}(Q_{n+4})$ to id_{n+4} and let $\pi_1 := \mathrm{most}(Q_{n+4})(\mathrm{id}_{n+4})$ be the last permutation on this cycle before reaching id_{n+4} again. Furthermore, let P_2 be the set of permutations encountered when applying $\mathrm{most}(Q')$ to π_1, but with π_1 removed, and let $\pi_2 := \mathrm{most}(Q')(\pi_1)$ be the last permutation on this cycle before reaching π_1 again. Moreover, let $\pi_3 := \mathrm{last}(Q_{n+4})(\pi_2)$; see Figure 18 (c). Note that $P_1 \cup P_2 \cup \{\pi_3\}$ is the set of all visited permutations. The same argument as for R_{n+1} before, applied four times, shows that no permutation in P_1 is encountered twice when applying Q_{n+4} to id_{n+4}. As R_4 is a rainbow sequence, no permutation in P_2 is encountered twice when applying Q' to π_1. Note that $\pi_1 = (\pi', n+1, n+2, n+3, n+4)$ with $\pi' := \mathrm{most}(R_n)(\mathrm{id}_n) \in \Pi_n$, and as R_n applied to id_n forms a rainbow cycle, π' is not a prefix of any permutation in $P_1 \setminus \{\pi_1\}$, but it is a prefix of every permutation in P_2. It follows that $P_1 \cap P_2 = \emptyset$. Finally, π_3 has id_n as a prefix, which shows that $\pi_3 \notin P_2$. Moreover, $\pi_3 \neq \mathrm{id}_{n+4}$, and id_n is not a prefix of any permutation in $P_1 \setminus \{\mathrm{id}_{n+4}\}$, proving that $\pi_3 \notin P_1$. This completes the proof of (ii). □

6. Subsets

In this section we consider the set of all k-element subsets of $[n]$, denoted by $C_{n,k} := \binom{[n]}{k}$, sometimes called (n,k)-combinations. The graph $G^{\mathsf{C}}_{n,k}$ has $C_{n,k}$ as its vertex set, and an edge $\{A, B\}$ between any two sets A and B that differ in exchanging an element x for another element y, i.e., $A \setminus B = \{x\}$ and $B \setminus A = \{y\}$; see Figure 2 (d). We label the edge $\{A, B\}$ of $G^{\mathsf{C}}_{n,k}$ with the transposition $A \triangle B = \{x, y\} \in C_{n,2}$, and we think of these labels as colors. A 1-rainbow cycle in $G^{\mathsf{C}}_{n,k}$ is an undirected cycle along which every transposition appears exactly once, so it has length $\binom{n}{2}$. In this section we only consider 1-rainbow cycles, and we simply refer to them as rainbow cycles. The number of vertices of $G^{\mathsf{C}}_{n,k}$ is clearly $\binom{n}{k}$. Consequently, a rainbow cycle for $k = 2$ is in fact a Hamilton cycle, i.e., a Gray code in the classical sense. As $G^{\mathsf{C}}_{n,k}$ and $G^{\mathsf{C}}_{n,n-k}$ are isomorphic, including the edge labels, we will assume without loss of generality that $k \leq \lfloor n/2 \rfloor$. Also note that for $k = 1$, the number of vertices of $G^{\mathsf{C}}_{n,k}$ is only n, which is strictly smaller than $\binom{n}{2}$ for $n > 3$, so we will also assume that $k \geq 2$.

The following theorem summarizes the results of this section.

Theorem 15. *Let $n \geq 4$ and $2 \leq k \leq \lfloor n/2 \rfloor$. The flip graph of subsets $G^{\mathsf{C}}_{n,k}$ has the following properties:*

(i) If n is even, then $G^{\mathsf{C}}_{n,k}$ has no rainbow cycle.
(ii) If n is odd and $k = 2$, then $G^{\mathsf{C}}_{n,2}$ has a rainbow Hamilton cycle.
(iii) If n is odd and $k = 2$, then $G^{\mathsf{C}}_{n,2}$ has two edge-disjoint rainbow Hamilton cycles.
(iv) If n is odd and $3 \leq k < n/3$, then $G^{\mathsf{C}}_{n,k}$ has a rainbow cycle.

With the help of a computer we found even more than two edge-disjoint rainbow Hamilton cycles in $G^{\mathsf{C}}_{n,2}$ for odd n; see Table 2. Moreover, we firmly believe the $G^{\mathsf{C}}_{n,k}$ also has a rainbow cycle for $n/3 \leq k \leq \lfloor n/2 \rfloor$, but we are not able to prove this.

6.1. **Proof of Theorem 15 (i).** We start to show the non-existence of rainbow cycles in the case that n is even.

B_i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	dist(B_i)	dist$(B_i \triangle B_{i+1}) = \lvert d_i \rvert$
B_1	×																×	1	3
B_2	×		×															2	5
B_3	×														×			3	$7 = \ell - 1$
B_4	×				×													4	1
B_5	×					×												5	$6 = \ell - 2$
B_6	×											×						6	4
B_7	×							×										7	2
B_8	×									×								8	
$\sigma(B_1)$	×	×																	dist$(B_\ell \triangle \sigma(B_1))$ $= 8 = \ell$

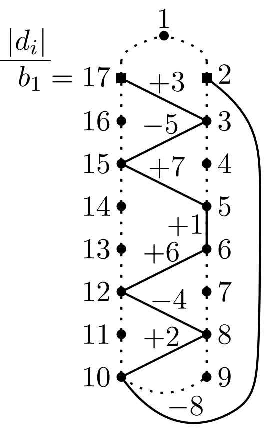

FIGURE 19. A rainbow block for $\ell = 8$, corresponding to case (6a). A cross in row B_i and column j indicates that $j \in B_i$. On the right hand side, the sequence $(b_1, b_2, \ldots, b_\ell, 2)$ for this block is depicted as a path.

Proof of Theorem 15 *(i).* Note that, for a fixed element $x \in [n]$, there are $n-1$ transpositions involving x. If x is in a set along a rainbow cycle and such a transposition is applied, then the next set along the cycle does not contain x, and vice versa. In a rainbow cycle we return to the starting set and use each of these transpositions exactly once, so $n-1$ must be even, or equivalently, n must be odd. $\square$

6.2. **Proof of Theorem 15 (ii).** For the rest of this section we assume that n is odd, i.e., $n = 2\ell+1$ for some integer $\ell \geq 2$. To prove parts (ii)–(iv) of Theorem 15, we construct rainbow cycles using a *rainbow block*. To introduce this notion, we need some definitions. For any set $A \subseteq [n]$ we let $\sigma(A)$ denote the set obtained from A by adding 1 to all elements, modulo n with $\{1, 2, \ldots, n\}$ as residue class representatives. Moreover, for any pair $\{x, y\} \in C_{n,2}$, we define $\mathrm{dist}(\{x, y\}) := \min\{y - x, x - y\} \in [\ell]$ where the differences are also taken modulo n.

We call a sequence $B = (B_1, B_2, \ldots, B_\ell)$ of subsets $B_i \in C_{n,k}$ a *rainbow block* if

$$C(B) := \big(B, \sigma^1(B), \sigma^2(B), \ldots, \sigma^{2\ell}(B)\big) \tag{5}$$

is a rainbow cycle in $G^{\mathsf{C}}_{n,k}$. Note that the sequence $C(B)$ has the correct length $\ell \cdot (2\ell+1) = \binom{n}{2}$. By definition, a rainbow cycle built from a rainbow block is highly symmetric. In the following proofs we will formulate various sufficient conditions guaranteeing that B is a rainbow block, and construct B such that those conditions are satisfied.

Proof of Theorem 15 *(ii).* Let $n = 2\ell + 1$ for some integer $\ell \geq 2$. We define a sequence $B = (B_1, B_2, \ldots, B_\ell)$ of pairs $B_i \in C_{n,2}$ such that the following conditions hold:

(a) $B_i = \{1, b_i\}$ for $i \in [\ell]$ with $3 \leq b_i \leq n$ and $b_1 = n$,
(b) $\{\mathrm{dist}(B_i) \mid i \in [\ell]\} = [\ell]$, and
(c) $\{\mathrm{dist}(B_i \triangle B_{i+1}) \mid i \in [\ell-1]\} \cup \{\mathrm{dist}(B_\ell \triangle \sigma(B_1)\} = [\ell]$.

Figure 19 shows a sequence B satisfying these conditions for $\ell = 8$.

We claim that a sequence B satisfying (a)–(c) is a rainbow block, i.e., the cycle $C = C(B)$ defined in (5) is a rainbow cycle. This can be seen as follows: By (a) and (5), any two consecutive sets in C differ in exactly one transposition. Here we use that $B_\ell \triangle \sigma(B_1) = \{1, b_\ell\} \triangle \{1, 2\} = \{b_\ell, 2\}$ and that $B = \sigma^{2\ell+1}(B)$. To prove that C is a cycle in $G^{\mathsf{C}}_{n,2}$, it remains to argue that all pairs in C are distinct. Note that for every fixed value of $d \in [\ell]$, there are exactly n different pairs $A \in C_{n,2}$ with $\mathrm{dist}(A) = d$. Therefore, by (b) and (5), every pair $A \in C_{n,2}$ appears exactly once in C. We now argue that C is a rainbow cycle, i.e., that every transposition appears exactly once along the cycle C. The argument here is very similar. For every fixed value of $d \in [\ell]$, there are exactly n different transpositions $T \in C_{n,2}$ with $\mathrm{dist}(T) = d$. Therefore, by (c) and (5), every transposition $T \in C_{n,2}$ appears exactly once along C. This proves that C is indeed a rainbow cycle in $G^{\mathsf{C}}_{n,2}$.

It remains to show how to construct a rainbow block satisfying the conditions (a)–(c). For this it suffices to define values for the elements b_i, $i = 2, 3, \ldots, \ell$; recall that $B_i = \{1, b_i\}$ and $b_1 = n$. For even ℓ we define

$$(d_1, d_2, \ldots, d_{\ell-1}) := \begin{cases} \big(+3, -5, +7, \ldots, -(\ell-3), +(\ell-1), +1, +(\ell-2), -(\ell-4), \ldots, -4, +2\big) & \text{if } \ell \equiv 0 \bmod 4, \\ \big(+3, -5, +7, \ldots, +(\ell-3), -(\ell-1), -1, -(\ell-2), +(\ell-4), \ldots, -4, +2\big) & \text{if } \ell \equiv 2 \bmod 4, \end{cases} \tag{6a}$$

and for odd ℓ we define

$$(d_1, d_2, \ldots, d_{\ell-1}) := \begin{cases} \big(+3, -5, +7, \ldots, +(\ell-2), -\ell, -1, -(\ell-3), +(\ell-5), \ldots, +4, -2\big) & \text{if } \ell \equiv 1 \bmod 4, \\ \big(+3, -5, +7, \ldots, -(\ell-2), +\ell, +1, +(\ell-3), -(\ell-5), \ldots, +4, -2\big) & \text{if } \ell \equiv 3 \bmod 4. \end{cases} \tag{6b}$$

Using that $b_1 = n$, we then define for all $i \in [\ell-1]$

$$b_{i+1} := b_i + d_i = b_1 + \sum_{1 \le j \le i} d_j \mod n. \tag{7}$$

This definition yields a sequence $B = (B_1, B_2, \ldots, B_\ell)$ where $B_i = \{1, b_i\}$ for $i \in [\ell]$. It can be easily verified using (6) and (7) that this sequence satisfies condition (a), and that $\mathrm{dist}(B_i) = i$ for all $i \in [\ell]$, proving (b). Moreover, we have

$$b_\ell = \begin{cases} \ell + 2 & \text{if } \ell \text{ is even}, \\ \ell + 1 & \text{if } \ell \text{ is odd}. \end{cases} \tag{8}$$

From these definitions it also follows that $\mathrm{dist}(B_i \triangle B_{i+1}) = \mathrm{dist}(\{b_i, b_{i+1}\}) = |d_i|$ for all $i \in [\ell-1]$. For even ℓ, it therefore follows from (6a) that the set $\{\mathrm{dist}(B_i \triangle B_{i+1}) \mid i \in [\ell-1]\}$ contains all numbers $\{1, 2, \ldots, \ell\}$ except ℓ. On the other hand, for odd ℓ, it follows from (6b) that this set contains all numbers $\{1, 2, \ldots, \ell\}$ except $\ell - 1$. These missing numbers are contributed by

$$\mathrm{dist}(B_\ell \triangle \sigma(B_1)) = \mathrm{dist}(\{b_\ell, 2\}) \overset{(8)}{=} \begin{cases} \ell & \text{if } \ell \text{ is even}, \\ \ell - 1 & \text{if } \ell \text{ is odd}, \end{cases}$$

so the sequence B indeed satisfies (c). This proves that B is a rainbow block, so the cycle $C(B)$ defined in (5) is a rainbow cycle in $G^{\mathsf{C}}_{n,2}$. □

As each pair and each transposition of a rainbow cycle C in $G^{\mathsf{C}}_{n,2}$ correspond to an edge of K_n, such a rainbow cycle has a nice interpretation as an iterative two-coloring of the edges of K_n; see Figure 20. One color class are the pairs along the cycle, and the other color class are the transpositions along the cycle. At each point we consider two consecutive pairs $\{x, y\}$ and $\{x, z\}$ in C and the transposition $\{y, z\}$ between them. This means that the edges $\{x, y\}$ and $\{x, z\}$ in K_n corresponding to the pairs share the vertex x, and the edge $\{y, z\}$ corresponding to the transposition goes between the other end vertices of these two edges, as depicted on the left hand side of Figure 20. The rainbow cycle shown on the right hand side of the figure is $C(B)$ with the rainbow block B as defined in the preceding proof of Theorem 15 (ii) for $n = 7$.

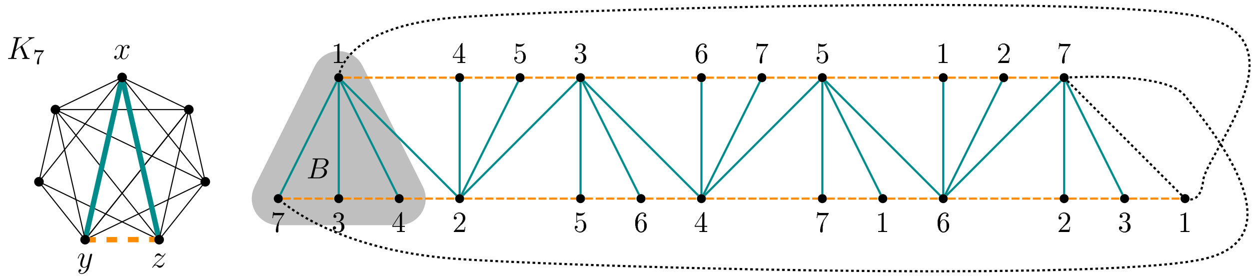

FIGURE 20. Interpretation of the rainbow cycle constructed in the proof of Theorem 15 (ii) as an iterative two-coloring of the edges of K_n for $n = 7$ ($\ell = 3$). The pairs in the rainbow block $B = (\{1, 7\}, \{1, 3\}, \{1, 4\})$ are highlighted in gray. In the drawing on the right, each vertex of K_n appears multiple times; different copies must be identified. In particular, the leftmost and rightmost edge in the drawing are identified as on a Möbius strip, indicated by the dotted lines. Solid edges represent pairs and dashed edges represent transpositions along the rainbow cycle. The solid edges form a caterpillar, and the dashed edges a Eulerian cycle in K_n.

6.3. **Proof of Theorem 15 (iii).** We now extend the method from the proof in the previous section and show that it even yields two edge-disjoint rainbow Hamilton cycles in $G^{\mathsf{C}}_{n,2}$.

We call a sequence of numbers $d = (d_1, d_2, \ldots, d_\ell)$ with $-\ell \leq d_i \leq \ell$ a *rainbow sequence*, if the sequence $B = (B_1, B_2, \ldots, B_\ell)$ with pairs $B_i = \{1, b_i\}$, $i \in [\ell]$, with b_i as defined in (7), satisfies the conditions (a)–(c) in the proof of Theorem 15 (ii) and if $\mathrm{dist}(\{b_\ell, 2\}) = |d_\ell|$. The last entry d_ℓ of the sequence is determined by the previous entries and by condition (c). In particular, we have $\{|d_i| \mid i \in [\ell]\} = [\ell]$. Recall that the elements d_i of a rainbow sequence are the increments/decrements by which b_{i+1} differs from b_i for $i = 1, 2, \ldots, \ell - 1$, and by which $\sigma(1) = 2$ differs from b_ℓ. In the previous proof we showed that such a rainbow sequence d gives rise to a rainbow cycle $C(B)$. We let $C(d)$ denote the rainbow Hamilton cycle defined in (5) for the rainbow block $B = B(d)$ defined via the rainbow sequence d.

Proof of Theorem 15 *(iii).* Observe that if $d = (d_1, d_2, \ldots, d_\ell)$ is any rainbow sequence, then the reversed sequence $\mathrm{rev}(d) := (d_\ell, d_{\ell-1}, \ldots, d_1)$ is also a rainbow sequence different from d. In particular, the number of rainbow sequences is always even.

To prove the theorem, we will show that for any rainbow sequence d, the rainbow Hamilton cycles $C(d)$ and $C(d')$ with $d' := \mathrm{rev}(d)$ are edge-disjoint cycles in the graph $G^{\mathsf{C}}_{n,2}$. We let b_i and b'_i, $i = 1, 2, \ldots, \ell$, denote the values defined in (7) for the sequences d_i and d'_i, respectively. As $d' = \mathrm{rev}(d)$ we clearly have

$$\begin{aligned} d_i &= d'_{\ell-i+1}, & 1 \leq i \leq \ell, \\ b_i &= n - b'_{\ell-i+2} + 2, & 2 \leq i \leq \ell. \end{aligned} \tag{9a}$$

Consider any pair $A \in C_{n,2}$ visited by the cycle $C(d)$, and let $\alpha, \beta \in [n]$ and $i \in [\ell]$ be such that $A = \{\alpha, \beta\} = \sigma^{\alpha-1}(\{1, b_i\})$. Note that α, β and i are uniquely determined by (5), and α is not necessarily smaller than β. Let $A^-, A^+ \in C_{n,2}$ be the pairs that precede and that follow the pair A on the cycle $C(d)$, respectively. See Figure 21 for an illustration.

□

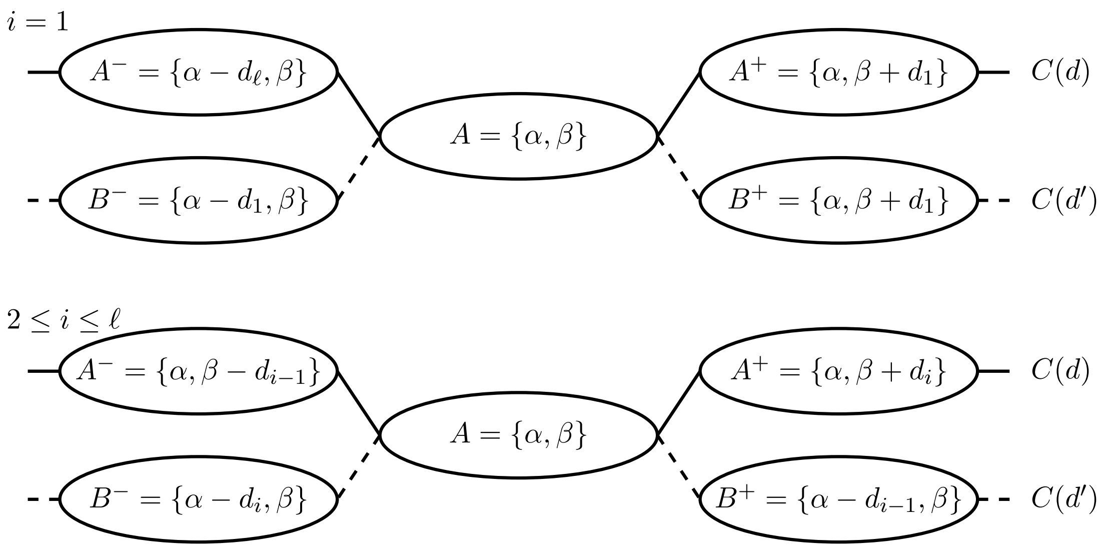

FIGURE 21. Illustration of the proof of Theorem 15 (iii). The figure shows the predecessors and successors of a pair $A = \{\alpha, \beta\} \in C_{n,2}$ on the cycles $C(d)$ (solid edges) and $C(d')$ (dashed edges) in the cases $i = 1$ (top) and $2 \leq i \leq \ell$ (bottom).

By (5) and (7) we have

$$A^- = \begin{cases} \{\alpha - d_\ell \bmod n, \beta\} & \text{if } i = 1, \\ \{\alpha, \beta - d_{i-1} \bmod n\} & \text{if } 2 \leq i \leq \ell, \end{cases} \tag{10a}$$

$$A^+ = \{\alpha, \beta + d_i \bmod n\}. \tag{10b}$$

We now consider the pairs B^- and B^+ that precede and follow the pair A on the cycle $C(d')$, respectively. By (5), (7) and (9a) we have

$$B^- = \{\alpha - d'_{\ell-i+1} \bmod n, \beta\} = \{\alpha - d_i \bmod n, \beta\}, \tag{11a}$$

$$B^+ = \begin{cases} \{\alpha, \beta + d_\ell \bmod n\} & \text{if } i = 1, \\ \{\alpha + d'_{\ell-i+2} \bmod n, \beta\} = \{\alpha + d_{i-1} \bmod n, \beta\} & \text{if } 2 \leq i \leq \ell. \end{cases} \tag{11b}$$

Using that $d_1 \neq d_\ell$ it follows immediately from (10) and (11) that the edge sets $\{\{A^-, A\}, \{A, A^+\}\}$ and $\{\{B^-, A\}, \{A, B^+\}\}$ are disjoint (see Figure 21), implying that $C(d)$ and $C(d')$ are edge-disjoint cycles, as claimed. This completes the proof.

With the help of a computer we determined all rainbow sequences d for $\ell = 1, 2, \ldots, 7$, as shown in Table 2, and we computed the maximum number of Hamilton cycles of the form $C(d)$ that are pairwise edge-disjoint.

As the degree of all vertices of $G^{\mathsf{C}}_{n,2}$ is $2(n-2)$, the number of edge-disjoint Hamilton cycles in the graph is at most $n - 2 = 2\ell - 1$. If this upper bound is matched, then we obtain a 2-factorization of $G^{\mathsf{C}}_{n,2}$ into edge-disjoint rainbow Hamilton cycles. As the table shows, this happens only in the trivial case $\ell = 1$, but it almost happens for $\ell = 2$ and $\ell = 6$. In fact, the following theorem shows that $n - 3$ edge-disjoint rainbow Hamilton cycles is the maximum we can achieve when considering only cycles of the form $C(d)$ for rainbow sequences d.

Theorem 16. *For any $\ell \geq 2$, if D is a set of rainbow sequences such that $C(d)$ and $C(d')$ are edge-disjoint cycles in $G^{\mathsf{C}}_{n,2}$, $n = 2\ell + 1$, for all $d, d' \in D$ with $d \neq d'$, then $|D| \leq n - 3$.*

Table 2. Number of rainbow sequences d and maximum number of pairwise edge-disjoint rainbow Hamilton cycles of the form $C(d)$ in the graph $G^{\mathsf{C}}_{n,2}$, $n = 2\ell + 1$, for $\ell = 1, 2, \ldots, 7$. For $\ell = 6$, rainbow sequences yielding ten edge-disjoint cycles are shown on the right. Only five sequences are shown explicitly, the other five are obtained by reversal. For $\ell \in \{1, 2, 6\}$ (printed in bold), $G^{\mathsf{C}}_{n,2}$ has a 2-factorization where all but one 2-factor are a rainbow Hamilton cycle.

ℓ	n	#rainbow sequences d	#edge-disjoint HCs $C(d)$	upper bound $n-2=2\ell-1$
1	3	1	**1**	**1**
2	5	2	**2**	**3**
3	7	2	2	5
4	9	6	2	7
5	11	14	5	9
6	13	80	$\boxed{\mathbf{10}}$	**11**
7	15	304	10	13

$\{d^i, \mathrm{rev}(d^i) \mid i \in [5]\}$ where

$d^1 := (-1, -2, -4, +5, -3, -6)$
$d^2 := (-2, +5, +4, +3, -1, +6)$
$d^3 := (+3, +6, -2, +4, -1, +5)$
$d^4 := (-3, +6, +5, +1, +2, +4)$
$d^5 := (-4, +3, -1, -6, +2, -5)$

Proof. The first pair B_1 in a rainbow block is the same $B_1 = \{1, b_1\}$, $b_1 = n$, for all rainbow sequences $d \in D$, recall condition (a) from the proof of Theorem 15 (ii). Moreover, the second pair B_2 has the form $\{1, b_2\}$ with $3 \leq b_2 \leq n$. By condition (c) we must have $b_2 \neq b_1$, so $3 \leq b_2 \leq n-1$. This leaves only $n-3$ different possible values for b_2, and the different rainbow sequences $d \in D$ must all have different values b_2, otherwise the corresponding cycles $C(d)$ would have the first edge in common. □

Note that even relaxing condition (a) to $2 \leq b_i \leq n$ in the proof of Theorem 15 (ii) does not change the conclusion of Theorem 16, as the value $b_i = 2$ is also forbidden by condition (b). However, Theorem 16 does not rule out the existence of a 2-factorization of $G^{\mathsf{C}}_{n,2}$ into rainbow Hamilton cycles that are *not* of the form $C(d)$ for some rainbow sequence d.

6.4. **Proof of Theorem 15 (iv).** In this section we describe rainbow cycles in $G^{\mathsf{C}}_{n,k}$ for $k \geq 3$.

Proof of Theorem 15 *(iv).* Let $n = 2\ell + 1$ for some integer $\ell \geq 2$.

As before, we construct a rainbow cycle in $G^{\mathsf{C}}_{n,k}$ using a rainbow block. Specifically, we will define a sequence $B = (B_1, B_2, \ldots, B_\ell)$ of subsets $B_i \in C_{n,k}$ such that the following conditions hold:

(a) $B_i = [k-1] \cup \{b_i\}$ for $i \in [\ell]$ with $k+1 \leq b_i \leq n$ and $b_1 = n$,
(b) the numbers $b_1, b_2, \ldots, b_\ell$ are all distinct, and
(c) $\{\mathrm{dist}(B_i \triangle B_{i+1}) \mid i \in [\ell - 1]\} \cup \{\mathrm{dist}(B_\ell \triangle \sigma(B_1))\} = [\ell]$.

We claim that a sequence B satisfying (a)–(c) is a rainbow block, i.e., the cycle $C = C(B)$ defined in (5) is a rainbow cycle. The proof of this fact is very much analogous to the argument in the proof of part (ii) of the theorem, so we omit it here.

An example of a rainbow block for $\ell = 8$ and $k = 4$ is shown in Figure 22.

We interpret a sequence B satisfying these conditions as a path on the vertex set $[n]$ as follows. Given an edge $\{x, y\}$ of the path, we refer to the quantity $\mathrm{dist}(\{x, y\}) \in [\ell]$ as the *length* of this edge. We say that an edge is *short* if its length is at most $k-1$, and it is *long* if its length is at least k.

Given a sequence B satisfying the conditions (a)–(c), then the sequence $(b_1, b_2, \ldots, b_\ell, k)$ is a simple path of length ℓ on the vertex set $[k, n] := \{k, k+1, \ldots, n\} \subseteq [n]$ that starts at the vertex $b_1 = n$,

B_i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	d_i
B_1	×	×	×														×	5
B_2	×	×	×		×													1
B_3	×	×	×			×												8
B_4	×	×	×											×				6
B_5	×	×	×					×										4
B_6	×	×	×									×						3
B_7	×	×	×						×									2
B_8	×	×	×								×							7
$\sigma(B_1)$	×	×	×	×														

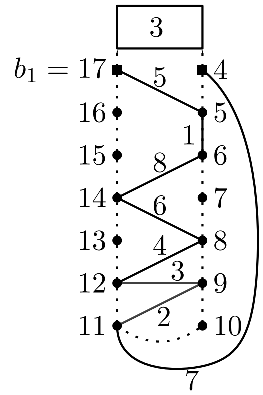

FIGURE 22. A rainbow block for $\ell = 8$ and $k = 4$. On the right hand side, the sequence $(b_1, b_2, \ldots, b_\ell, k)$ for this block is depicted as a path starting with $b_1 = 17$. The numbers d_i are the edge lengths along the path.

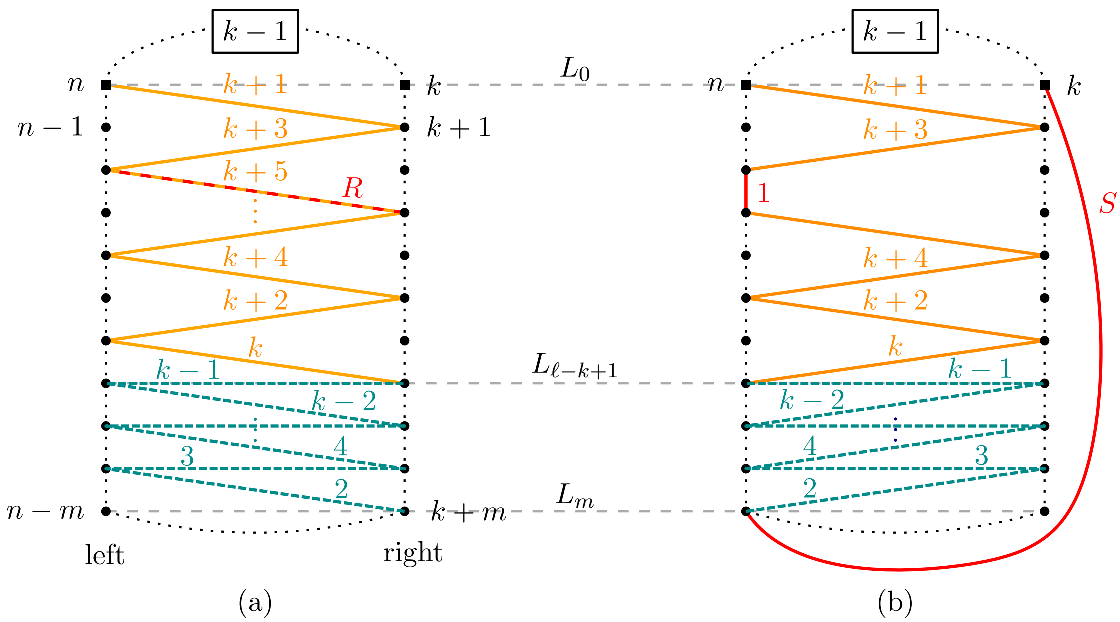

FIGURE 23. (a) A zigzag path for even ℓ and even k ($\ell = 14$, $k = 8$). The long edges are drawn with bold solid lines, the short edges with bold dashed lines.
(b) A modified zigzag path ($R = k+5 = 13$ and $S = 11$; this is not a rainbow path).

ends at the vertex k, and that has the property that along the path every edge length from the set $[\ell]$ appears exactly once; see the right hand side of Figure 22. We refer to such a path as a *rainbow path*.

The following definitions are illustrated in Figure 23.

We define $L_i := \{n-i, k+i\}$ and we say that the vertices in L_i belong to *level* i where $i = 0, 1, \ldots, m$ and $m := \ell - \lceil k/2 \rceil$. We refer to the edge $\{n-i, k+i\}$ as a *level edge*, and to any edge of the form $\{i, i+1\}$ as a *cycle edge*. For any two consecutive levels L_i and L_{i+1}, we call the edges $\{n-i, k+i+1\}$ and $\{k+i, n-i-1\}$ *diagonal edges* between these two levels. Note that these two diagonal edges have the same length. We refer to the vertices $\{n-i \mid i = 0, 1, \ldots, m\}$ and $\{k+i \mid i = 0, 1, \ldots, m\}$ as *left* and *right vertices*, respectively. Note that every level contains exactly

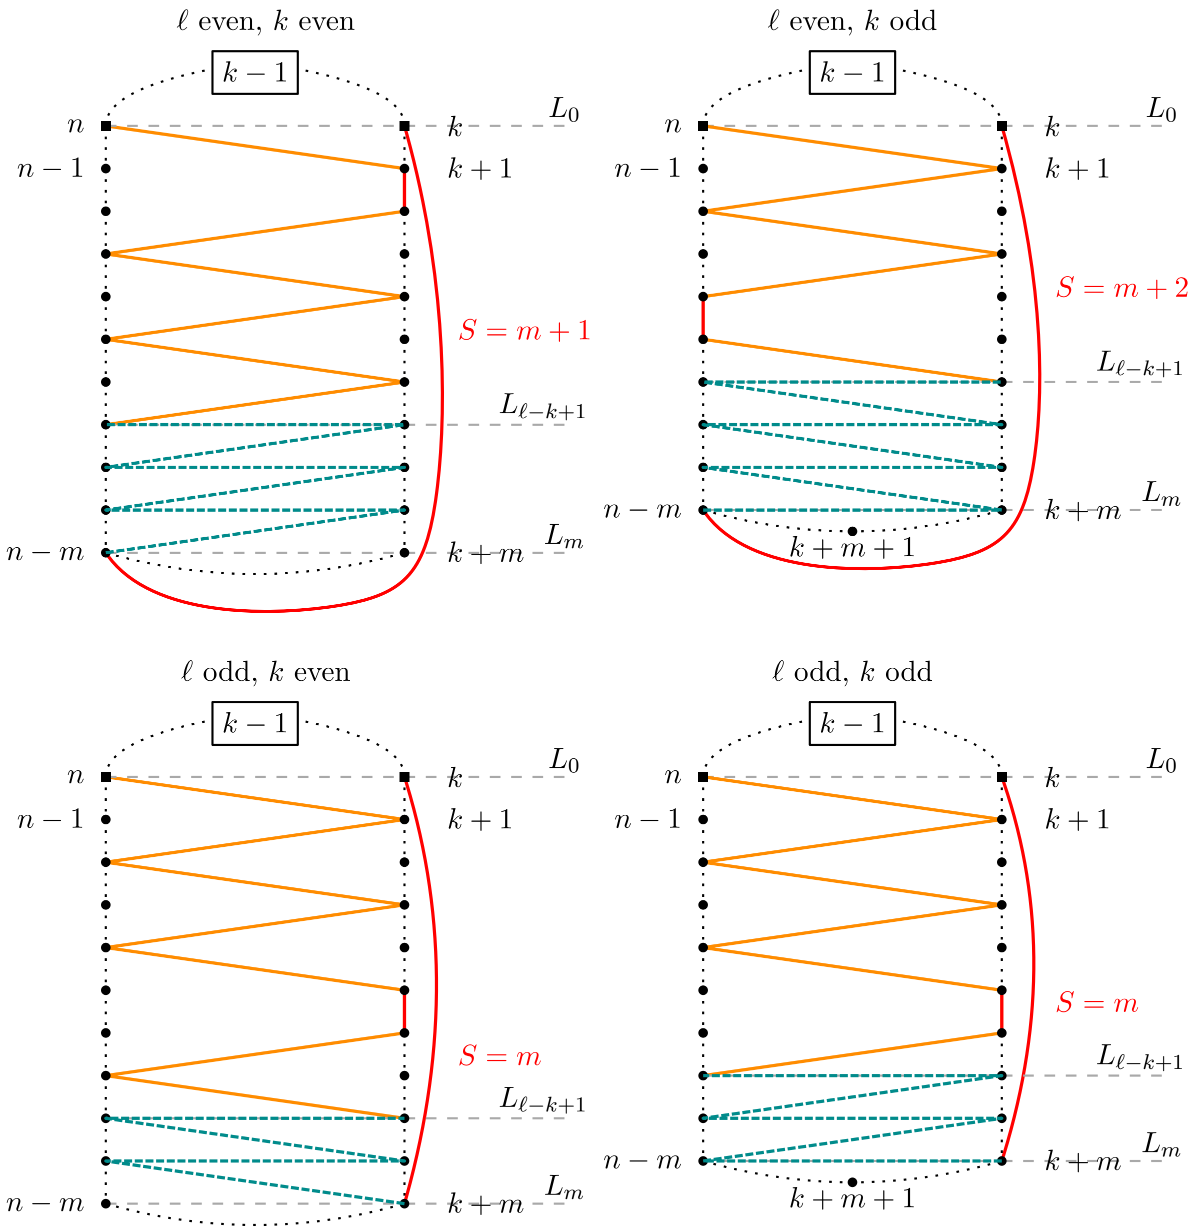

Figure 24. Rainbow paths obtained from our construction for all four cases of even/odd ℓ and k (specifically, these examples are for $\ell=14$ and $k\in\{8,9\}$, and for $\ell=13$ and $k\in\{6,7\}$).

two vertices, a left and a right vertex. For odd k, there is a unique vertex $(k+m+1)\in[k,n]$ which is not assigned to any level, and which is neither a left nor a right vertex.

In our figures, we display the vertices in $[k-1]=[n]\setminus[k,n]$ in a box at the top, and the vertices in levels $L_0,L_1,\ldots,L_m$ from top to bottom with vertices in the same level drawn at the same vertical position, all left vertices are drawn to the left of all right vertices. The vertex $(k+m+1)$ for odd k is shown below the level L_m.

We construct a rainbow path in two steps; see Figure 23. We first define a *zigzag path*. This is a path of length $\ell-1$ that starts at the vertex n and ends at a vertex in the last level L_m. The path first alternates between left and right vertices in the levels $L_0,L_1,\ldots,L_{\ell-k+1}$ along diagonal edges, and it then alternates between left and right vertices in the levels $L_{\ell-k+1}, \dots, L_m$ along level edges and diagonal edges alternatingly, starting with a level edge.

Observe that in the first part the zigzag path uses long edges, namely every length from the set $\{k, k+1, \dots, \ell\}$ exactly once. Note that as $n = 2\ell + 1$ is odd, the parity of the edge lengths changes after using the longest or second-longest edge. In the second part the zigzag path uses short edges, namely every length from the set $\{2, 3, \dots, k-1\}$ exactly once. Note also that a zigzag path ends at a right vertex if ℓ is even and at a left vertex if ℓ is odd.

We now modify the zigzag path as follows; see Figures 23 and 24. We replace one diagonal edge of some length R by a cycle edge of length 1 between the same two levels, and exchange each vertex after this edge for the other vertex in the same level (in our figures, this corresponds to mirroring the second part of the path after the modification at a vertical line). In addition, we extend the resulting path with one additional edge, making it a path of length ℓ, with an edge of length S that leads from the new end vertex in level L_m to the vertex k. It is easy to check that

$$S = \begin{cases} m & \text{if } \ell \text{ is odd,} \\ m+1 & \text{if } \ell \text{ is even and } k \text{ is even,} \\ m+2 & \text{if } \ell \text{ is even and } k \text{ is odd;} \end{cases}$$

see Figure 24. In any case we have $S \geq m$. Observe that if the edge of length S in a zigzag path is a diagonal, then choosing $R := S$ in this construction yields a rainbow path, as every edge length from the set $[\ell]$ is used exactly once.

As all long edges along a zigzag path are diagonal edges, requiring that $S \geq k$ is enough to ensure that the edge of length S is a diagonal edge. This condition is satisfied, as the last inequality in the estimate

$$S \geq m = \ell - \lceil k/2 \rceil \geq \ell - k/2 - 1/2 \overset{!}{>} k - 1$$

is equivalent to our assumption $k < n/3 = (2\ell + 1)/3$. This completes the proof. $\square$

We remark that there are values of n and k such there is no rainbow cycle of the form $C(B)$ with a rainbow block B satisfying the conditions (a)–(c) from the previous proof. For instance, there is no rainbow block of this form for $\ell = k = 4$ and $\ell = k = 8$. However, the slightly more general rainbow blocks shown in Figure 25 and Figure 26 work in these cases. These blocks do not satisfy conditions (a) and (b) stated in the proof, but they satisfy condition (c).

B_i	1	2	3	4	5	6	7	8	9	
B_1	×	×	×						×	4
B_2	×	×					×		×	3
B_3	×	×	×				×			2
B_4	×	×	×		×					1
$\sigma(B_1)$	×	×	×	×						

FIGURE 25. A rainbow block for $\ell = k = 4$.

7. Open problems

For all the combinatorial classes considered in this paper, it would be very interesting to exhibit r-rainbow cycles for larger values of r (recall Table 1), in particular for the flip graphs of permutations and subsets. Ultimately, the aim is to cover the entire range of possible values of r, all the way up to rainbow Hamilton cycles, though in some cases there are divisibility issues, namely if the number of vertices of the flip graph is not a multiple of the number of edge labels (=colors). Another natural

B_i	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	
B_1	×	×	×	×	×	×	×										×	8
B_2	×	×	×	×	×	×									×		×	7
B_3	×	×	×	×	×	×	×								×			6
B_4	×	×	×	×	×	×	×		×									5
B_5	×	×	×	×	×	×	×							×				3
B_6	×	×	×	×	×	×	×				×							2
B_7	×	×	×	×	×	×	×						×					1
B_8	×	×	×	×	×	×	×					×						4
$\sigma(B_1)$	×	×	×	×	×	×	×	×										

FIGURE 26. A rainbow block for $\ell = k = 8$.

next step is to investigate rainbow cycles in other flip graphs, e.g., for non-crossing partitions of a convex point set or for dissections of a convex polygon (see [HHNOP09]).

We believe that the flip graph of non-crossing perfect matchings $G_m^{\mathtt{M}}$ has no 1-rainbow cycle for any $m \geq 5$. This is open for the even values of $m \geq 12$. Moreover, the subgraph H_m of $G_m^{\mathtt{M}}$ restricted to centered flips (see Figure 10) is a very natural combinatorial object with many interesting properties that deserve further investigation. What is the number of connected components of H_m and what is their size? Which components are trees and which components contain cycles? As a starting point, it would be very nice to prove the conjectured formula (4) for the number of matchings with a certain weight, and to understand the connection to the generalized Narayana numbers.

We conjecture that the flip graph of subsets $G_{n,k}^{\mathtt{C}}$ has a 1-rainbow cycle for all $2 \leq k \leq n-2$. This is open for $n/3 \leq k \leq 2n/3$. In view of Theorem 15 (iii) we ask: does $G_{n,2}^{\mathtt{C}}$ have a factorization into $n-2$ edge-disjoint rainbow Hamilton cycles?

8. Acknowledgements

We thank Manfred Scheucher for his quick assistance in running computer experiments that helped us to find rainbow cycles in small flip graphs. We also thank the anonymous referee of this paper who pointed out several interesting closely related references.

References

[AAHV07] O. Aichholzer, F. Aurenhammer, C. Huemer, and B. Vogtenhuber. Gray code enumeration of plane straight-line graphs. *Graphs Combin.*, 23(5):467–479, 2007.

[Ádá68] A. Ádám. *Truth functions and the problem of their realization by two-terminal graphs.* Akadémiai Kiadó, Budapest, 1968.

[Als08] B. Alspach. The wonderful Walecki construction. *Bull. Inst. Combin. Appl.*, 52:7–20, 2008.

[And89] L. D. Andersen. Hamilton circuits with many colours in properly edge-coloured complete graphs. *Mathematica Scandinavica*, 64:5–14, 1989.

[APS16] N. Alon, A. Pokrovskiy, and B. Sudakov. Random subgraphs of properly edge-coloured complete graphs and long rainbow cycles. *arXiv:1608.07028*, Aug 2016.

[BH09] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.

[BM17] J. Balogh and T. Molla. Long rainbow cycles and Hamiltonian cycles using many colors in properly edge-colored complete graphs. *arXiv:1706.04950*, Jun 2017.

[BRSW18] L. A. Berry, B. Reed, A. Scott, and D. R. Wood. A logarithmic bound for the chromatic number of the associahedron. *arXiv:1811.08972*, Nov 2018.

[BS96] G. S. Bhat and C. D. Savage. Balanced Gray codes. *Electron. J. Combin.*, 3(1):Research Paper 25, approx. 11 pp., 1996.

[BW84] M. Buck and D. Wiedemann. Gray codes with restricted density. *Discrete Math.*, 48(2-3):163–171, 1984.

[CM18] S. Cleary and R. Maio. Edge conflicts do not determine geodesics in the associahedron. *SIAM J. Discrete Math.*, 32(2):1003–1015, 2018.

[EHR84] P. Eades, M. Hickey, and R. C. Read. Some Hamilton paths and a minimal change algorithm. *J. Assoc. Comput. Mach.*, 31(1):19–29, 1984.

[Epp10] D. Eppstein. Happy endings for flip graphs. *J. Comput. Geom.*, 1(1):3–28, 2010.

[FMFPH+09] R. Fabila-Monroy, D. Flores-Peñaloza, C. Huemer, F. Hurtado, J. Urrutia, and D. R. Wood. On the chromatic number of some flip graphs. *Discrete Math. Theor. Comput. Sci.*, 11(2):47–56, 2009.

[GMN18] P. Gregor, T. Mütze, and J. Nummenpalo. A short proof of the middle levels theorem. To appear in *Discrete Analysis. arXiv:1710.08249*, 2018.

[HHMW19] E. Hartung, H. P. Hoang, T. Mütze, and A. Williams. Combinatorial generation via permutation languages. I. Fundamentals. *arXiv:1906.06069*, 2019.

[HHN02] C. Hernando, F. Hurtado, and M. Noy. Graphs of non-crossing perfect matchings. *Graphs Combin.*, 18(3):517–532, 2002.

[HHNOP09] C. Huemer, F. Hurtado, M. Noy, and E. Omaña-Pulido. Gray codes for non-crossing partitions and dissections of a convex polygon. *Discrete Appl. Math.*, 157(7):1509–1520, 2009.

[HN99] F. Hurtado and M. Noy. Graph of triangulations of a convex polygon and tree of triangulations. *Comput. Geom.*, 13(3):179–188, 1999.

[HNU99] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete Comput. Geom.*, 22(3):333–346, 1999.

[HOS96] S. Hanke, T. Ottmann, and S. Schuierer. The edge-flipping distance of triangulations. *Journal of Universal Computer Science*, 2(8):570–579, 1996.

[Joh63] S. Johnson. Generation of permutations by adjacent transposition. *Math. Comp.*, 17:282–285, 1963.

[KL75] V. L. Kompel'makher and V. A. Liskovets. Sequential generation of arrangements by means of a basis of transpositions. *Kibernetica*, 3:17–21, 1975.

[Knu11] D. E. Knuth. *The art of computer programming. Vol. 4A. Combinatorial algorithms. Part 1.* Addison-Wesley, Upper Saddle River, NJ, 2011.

[Lee89] C. W. Lee. The associahedron and triangulations of the n-gon. *European J. Combin.*, 10(6):551–560, 1989.

[Luc87] J. M. Lucas. The rotation graph of binary trees is Hamiltonian. *J. Algorithms*, 8(4):503–535, 1987.

[LZ98] M. Li and L. Zhang. Better approximation of diagonal-flip transformation and rotation transformation. In *Computing and Combinatorics, 4th Annual International Conference, COCOON '98, Taipei, Taiwan, R.o.C., August 12-14, 1998, Proceedings*, pages 85–94, 1998.

[MP16] T. Manneville and V. Pilaud. Graph properties of graph associahedra. *Sém. Lothar. Combin.*, 73:Art. B73d, 31, [2014-2016].

[Müt16] T. Mütze. Proof of the middle levels conjecture. *Proc. Lond. Math. Soc. (3)*, 112(4):677–713, 2016.

[NW78] A. Nijenhuis and H. S. Wilf. *Combinatorial algorithms.* Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, second edition, 1978. For computers and calculators, Computer Science and Applied Mathematics.

[Pou14] L. Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014.

[Pou19] L. Pournin. Eccentricities in the flip-graphs of convex polygons. To appear in *Journal of Graph Theory*, 2019.

[PP18a] H. Parlier and B. Petri. The genus of curve, pants and flip graphs. *Discrete Comput. Geom.*, 59(1):1–30, 2018.

[PP18b] H. Parlier and L. Pournin. Once punctured disks, non-convex polygons, and pointihedra. *Ann. Comb.*, 22(3):619–640, 2018.

[Rog99] R. O. Rogers. On finding shortest paths in the rotation graph of binary trees. In *Proceedings of the Thirtieth Southeastern International Conference on Combinatorics, Graph Theory, and Computing (Boca Raton, FL, 1999)*, volume 137, pages 77–95, 1999.

[RS93] F. Ruskey and C. Savage. Hamilton cycles that extend transposition matchings in Cayley graphs of S_n. *SIAM J. Discrete Math.*, 6(1):152–166, 1993.

[Rus88] F. Ruskey. Adjacent interchange generation of combinations. *J. Algorithms*, 9(2):162–180, 1988.

[Sav97] C. Savage. A survey of combinatorial Gray codes. *SIAM Rev.*, 39(4):605–629, 1997.

[STT88] D. D. Sleator, R. E. Tarjan, and W. P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.

[Tro62] H. F. Trotter. Algorithm 115: Perm. *Commun. ACM*, 5(8):434–435, 1962.

[Wil89] H. S. Wilf. *Combinatorial algorithms: an update*, volume 55 of *CBMS-NSF Regional Conference Series in Applied Mathematics.* Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1989.

Appendix A. Proof of Proposition 8

Proof of Proposition 8. For $n = 4$ points there are two possible configurations to consider: either all four points are in convex position, or three points are on the convex hull and one is in the interior. Solutions for both configurations are shown in Figure 27.

For $n = 5$ points we have to consider three different configurations, and in each case we have to construct r-rainbow cycles for each $r \in \{2, 3, 4\}$. Even though the arguments in the proofs of Proposition 5 and 7 do not work for $n = 5$, we can still use the described constructions with some minor modifications. Figure 28 shows such r-rainbow cycles for five points in convex position. Figure 29 covers the case that four points are on the convex hull and one point is in the interior. Figure 30 covers the case that three points are on the convex hull and two points are in the interior. $\square$

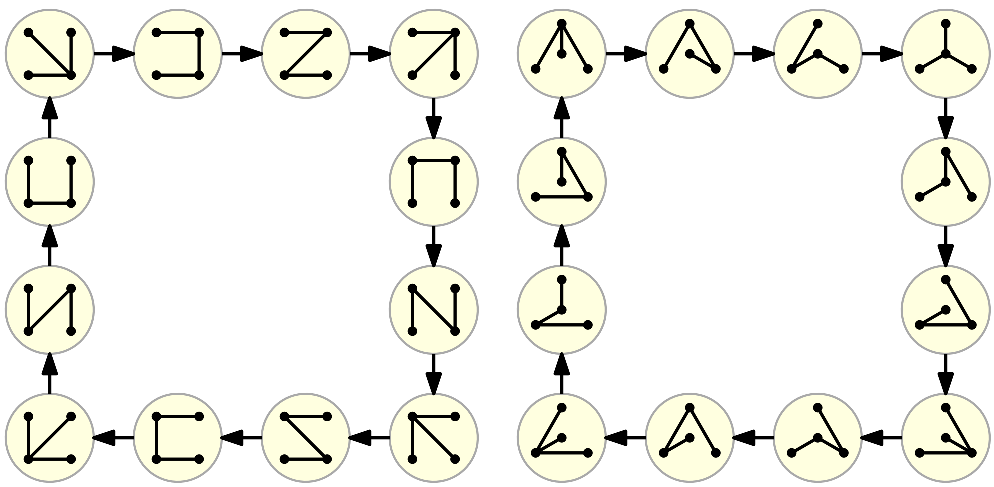

Figure 27. The figure shows 2-rainbow cycles for both configurations of $n = 4$ points.

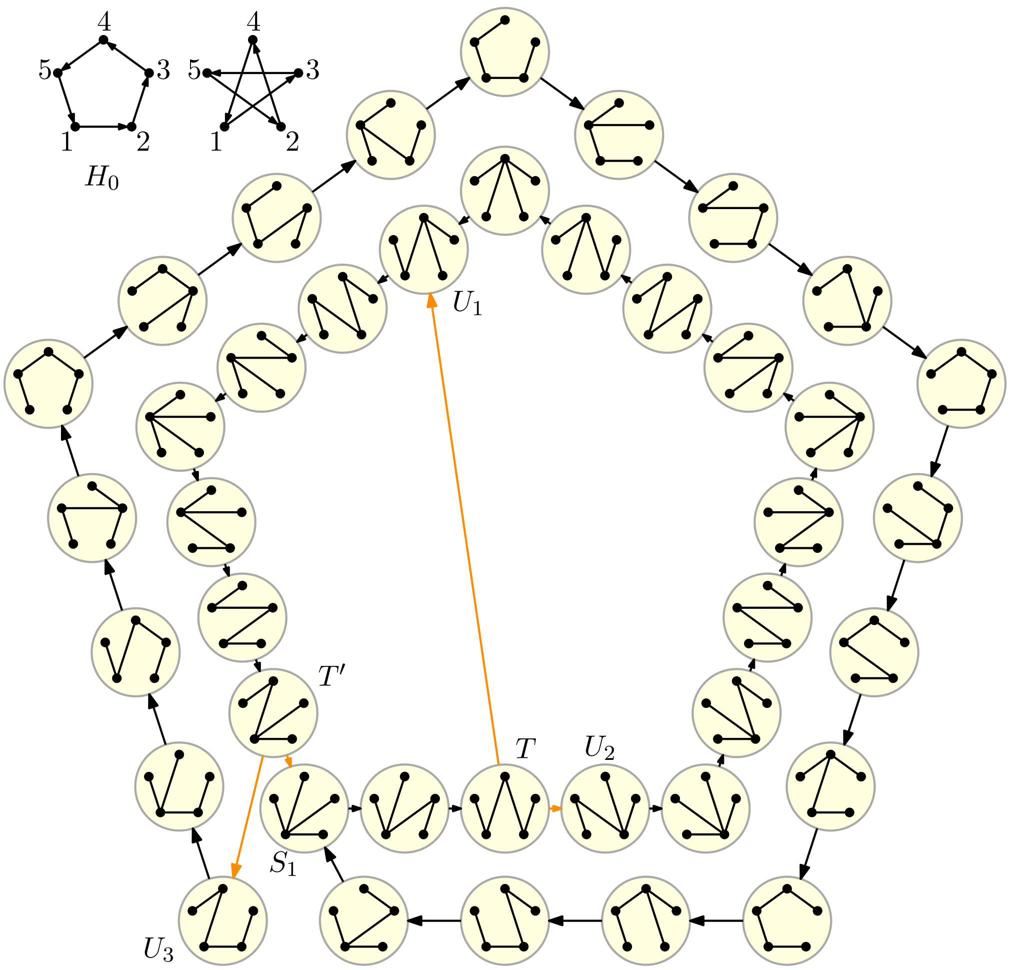

FIGURE 28. Illustration of 1-, 2-, 3-, and 4-rainbow cycles in G_X^{S}, where X is a set of $n = 5$ points in convex position. All four cycles contain S_1, but they use different arcs from T and T'. The 1-rainbow cycle uses the arcs (T, U_1) and (T', S_1). The 2-rainbow cycle uses the arcs (T, U_2) and (T', S_1). The 3-rainbow cycle uses the arcs (T, U_1) and (T', U_3). Finally, the 4-rainbow cycle uses the arcs (T, U_2) and (T', U_3).
The Hamilton cycles in K_X and their orientation used in the construction as in the proofs of Proposition 5 and 7 are shown in the upper left corner.

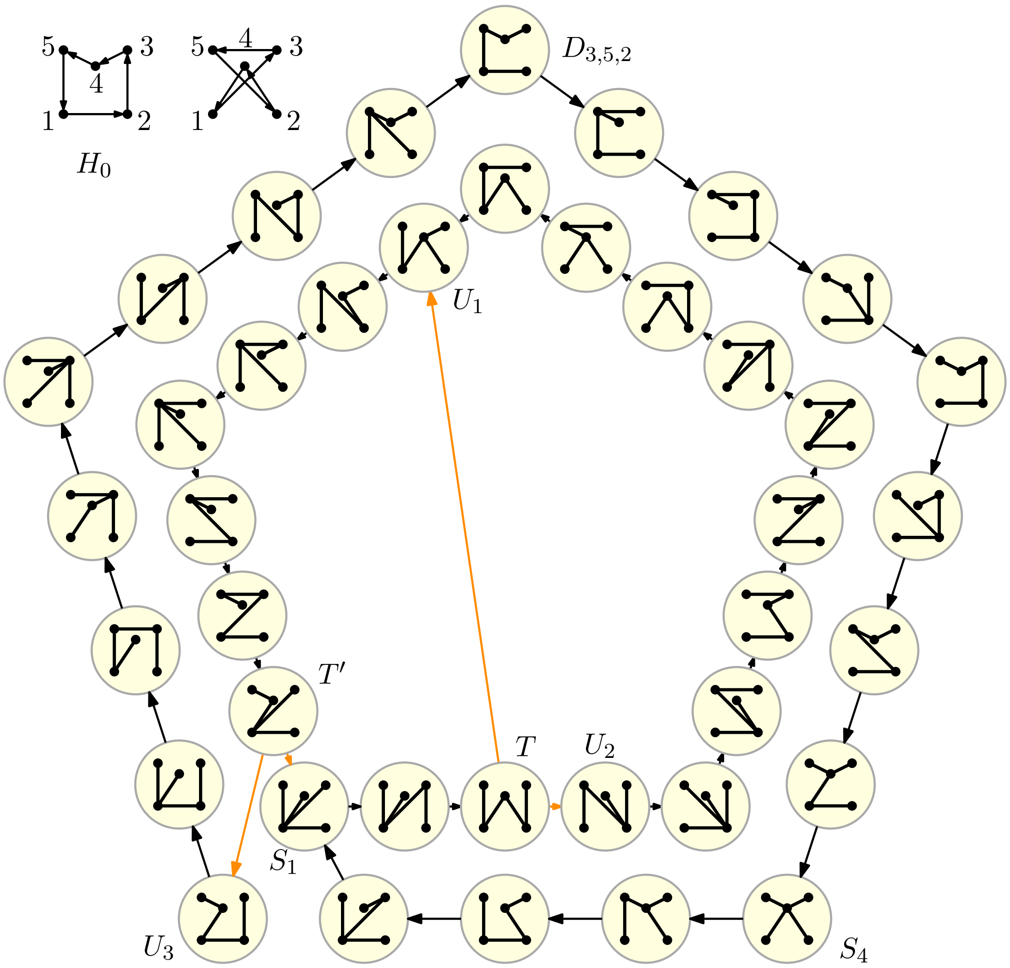

FIGURE 29. Illustration of 1-, 2-, 3-, and 4-rainbow cycles in G_X^{S}, where X is a set of $n = 5$ points with exactly four points on the convex hull. All four cycles contain S_1, but they use different arcs from T and T'. The 1-rainbow cycle uses the arcs (T, U_1) and (T', S_1). The 2-rainbow cycle uses the arcs (T, U_2) and (T', S_1). The 3-rainbow cycle uses the arcs (T, U_1) and (T', U_3). Finally, the 4-rainbow cycle uses the arcs (T, U_2) and (T', U_3). Note that in contrast to the construction in the proof of Proposition 5, we replaced $D_{2,4,1}$ at the bottom right by S_4, since $D_{2,4,1} = D_{3,5,2}$, and S_4 does not occur anywhere else.

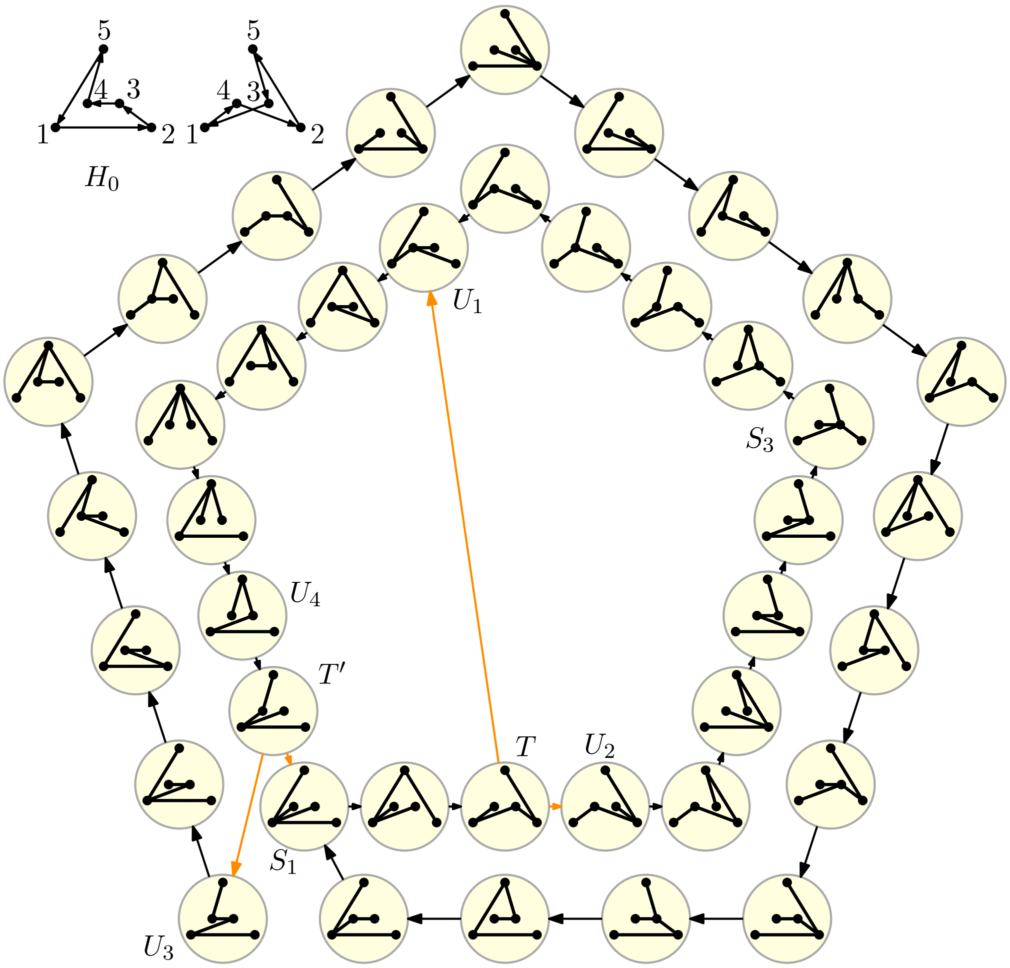

FIGURE 30. Illustration of 1-, 2-, 3-, and 4-rainbow cycles in G^{S}_X, where X is a set of $n = 5$ points with exactly three points on the convex hull. All cycles contain S_1, but they use different arcs from T and T'. The 1-rainbow cycle uses the arcs (T, U_1) and (T', S_1). The 2-rainbow cycle uses the arcs (T, U_2) and (T', S_1). The 3-rainbow cycle uses the arcs (T, U_1) and (T', U_3). Finally, the 4-rainbow cycle uses the arcs (T, U_2) and (T', U_3). Note that $D_{2,3,4}$ would coincide with the tree U_4 and is therefore replaced by S_3, which does not occur anywhere else.