Star Transposition Gray Codes for Multiset Permutations

Petr Gregor
Department of Theoretical Computer Science and Mathematical Logic,
Charles University, Prague, Czech Republic

Torsten Mütze
Department of Computer Science, University of Warwick, Coventry, UK
Department of Theoretical Computer Science and Mathematical Logic,
Charles University, Prague, Czech Republic

Arturo Merino
Department of Mathematics, TU Berlin, Germany

Abstract

Given integers $k \geq 2$ and $a_1, \ldots, a_k \geq 1$, let $\boldsymbol{a} := (a_1, \ldots, a_k)$ and $n := a_1 + \cdots + a_k$. An $\boldsymbol{a}$-multiset permutation is a string of length n that contains exactly a_i symbols i for each $i = 1, \ldots, k$. In this work we consider the problem of exhaustively generating all $\boldsymbol{a}$-multiset permutations by star transpositions, i.e., in each step, the first entry of the string is transposed with any other entry distinct from the first one. This is a far-ranging generalization of several known results. For example, it is known that permutations ($a_1 = \cdots = a_k = 1$) can be generated by star transpositions, while combinations ($k = 2$) can be generated by these operations if and only if they are balanced ($a_1 = a_2$), with the positive case following from the middle levels theorem. To understand the problem in general, we introduce a parameter $\Delta(\boldsymbol{a}) := n - 2\max\{a_1, \ldots, a_k\}$ that allows us to distinguish three different regimes for this problem. We show that if $\Delta(\boldsymbol{a}) < 0$, then a star transposition Gray code for $\boldsymbol{a}$-multiset permutations does not exist. We also construct such Gray codes for the case $\Delta(\boldsymbol{a}) > 0$, assuming that they exist for the case $\Delta(\boldsymbol{a}) = 0$. For the case $\Delta(\boldsymbol{a}) = 0$ we present some partial positive results. Our proofs establish Hamilton-connectedness or Hamilton-laceability of the underlying flip graphs, and they answer several cases of a recent conjecture of Shen and Williams. In particular, we prove that the middle levels graph is Hamilton-laceable.

2012 ACM Subject Classification Mathematics of computing → Permutations and combinations; Mathematics of computing → Combinatorial algorithms; Mathematics of computing → Matchings and factors

Keywords and phrases Gray code, permutation, combination, transposition, Hamilton cycle

Digital Object Identifier 10.4230/LIPIcs.STACS.2022.34

Related Version Full proofs are provided in the preprint version [12] of this extended abstract.
Full Version: https://arxiv.org/abs/2108.07465

Funding This work was supported by Czech Science Foundation grant GA 19-08554S, and by German Science Foundation grant 413902284.
Arturo Merino: was also supported by ANID Becas Chile 2019-72200522.

Acknowledgements We thank Jiří Fink for inspiring brain-storming sessions in the early phases of this research project, and we thank Aaron Williams for interesting discussions about star transposition Gray codes.

1 Introduction

Permutations and combinations are two of the most fundamental classes of combinatorial objects. Specifically, *k-permutations* are all linear orderings of $[k] := \{1, \ldots, k\}$, and their number is $k!$. Moreover, *(α, β)-combinations* are all β-element subsets of $[n]$ where $n := \alpha + \beta$, and their number is $\binom{n}{\alpha} = \binom{n}{\beta}$. Permutations and combinations are generalized by so-called multiset permutations, and in this paper we consider the task of listing them such that any two consecutive objects in the list differ by particular transpositions, i.e., by swapping two elements. Such a listing of objects subject to a "small change" operation is often referred to as *Gray code* [27, 30]. One of the standard references for algorithms that efficiently generate various combinatorial objects, including permutations and combinations, is Knuth's book [19] (see also [25]).

1.1 Permutation generation

There is a vast number of Gray codes for permutation generation, most prominently the Steinhaus-Johnson-Trotter algorithm [18, 40], which generates all k-permutations by adjacent transpositions, i.e., swaps of two neighboring entries of the permutation; see Figure 1 (a). In this work, we focus on *star transpositions*, i.e., swaps of the first entry of the permutation with any later entry. An efficient algorithm for generating permutations by star transpositions was found by Ehrlich, and it is described as Algorithm E in Knuth's book [19, Section 7.2.1.2]; see Figure 1 (b). For any permutation generation algorithm based on transpositions, we can define the *transposition graph* as the graph with vertex set $[k]$, and an edge between i and j if the algorithm uses transpositions between the ith and jth entry of the permutation. Clearly, the transposition graph for adjacent transpositions is a path, whereas the transposition graph for star transpositions is a star (hence the name "star transposition"). In fact, Kompel'maher and Liskovec [20], and independently Slater [36], showed that all k-permutations can be generated for any transposition tree on $[k]$. Transposition Gray codes for permutations with additional restrictions were studied by Compton and Williamson [7] and by Shen and Williams [33].

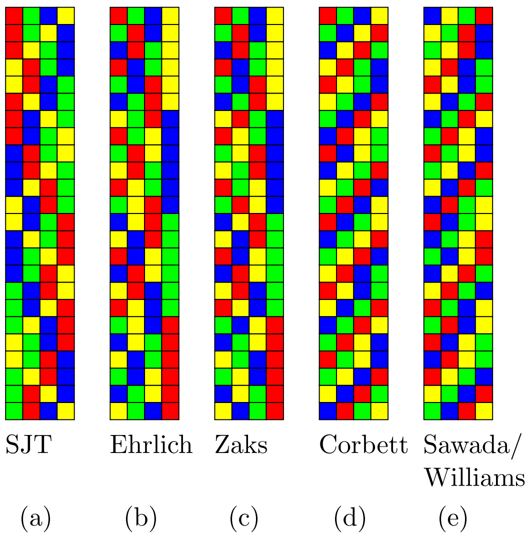

Figure 1 Gray codes for 4-permutations (SJT=Steinhaus-Johnson-Trotter).

Several known algorithms for permutation generation use operations other than transpositions. Specifically, Zaks [43] presented an algorithm for generating permutations by prefix reversals; see Figure 1 (c). Moreover, Corbett [8] showed that all k-permutations can be generated by cyclic left shifts of any prefix of the permutation by one position; see Figure 1 (d). Another notable result is Sawada and Williams' recent solution [32] of the Sigma-Tau problem, proving that all k-permutations can be generated by cyclic left shifts of the entire permutation by one position or transpositions of the first two elements; see Figure 1 (e).

All of the aforementioned results can be seen as explicit constructions of Hamilton paths in the Cayley graph of the symmetric group, generated by different sets of generators (transpositions, reversals, or shifts). It is an open problem whether the Cayley graph of the symmetric group has a Hamilton path for any set of generators [28]. This is a special case of the well-known open problem whether any connected Cayley graph has a Hamilton path, or even more generally, whether this is the case for any vertex-transitive graph [21].

1.2 Combination generation and the middle levels conjecture

In a computer, (α, β)-combinations can be conveniently represented by bitstrings of length $n := \alpha + \beta$, where the ith bit is 1 if the element i is in the set and 0 otherwise. For example, the $(5,3)$-combination $\{1,6,7\}$ is represented by the string 10000110.

In the 1980s, Buck and Wiedemann [3] conjectured that all (α, α)-combinations can be generated by star transpositions for every $\alpha \geq 1$, i.e., in every step we swap the first bit of the bitstring representation with a later bit. Figure 2 (a) shows such a star transposition Gray code for $(4,4)$-combinations. Buck and Wiedemann's conjecture was raised independently by Havel [15], as a question about the existence of a Hamilton cycle through the middle two levels of the $(2\alpha-1)$-dimensional hypercube. This conjecture became known as *middle levels conjecture*, and it attracted considerable attention in the literature and made its way into popular books [9, 42], until it was answered affirmatively by Mütze [23]; see also [13].

Similarly to permutations, there are also many known methods for generating general (α, β)-combinations that use operations other than star transpositions, see [5, 11, 17, 29, 38]. In particular, (α, β)-combinations can be generated by adjacent transpositions if and only if $\alpha = 1$, $\beta = 1$, or α and β are both odd [3, 10, 26].

1.3 Multiset permutations

Shen and Williams [34] proposed a far-ranging generalization of the middle levels conjecture that connects permutations and combinations. Their conjecture is about multiset permutations. For integers $k \geq 2$ and $a_1, \ldots, a_k \geq 1$, an $(a_1, \ldots, a_k)$-*multiset permutation* is a string over the alphabet $\{1, \ldots, k\}$ that contains exactly a_i occurrences of the symbol i. We refer to the sequence $\boldsymbol{a} := (a_1, \ldots, a_k)$ as the *frequency vector*, as it specifies the frequency of each symbol. The length of a multiset permutation is $n_{\boldsymbol{a}} := a_1 + \cdots + a_k$, and if the context is clear we omit the index and simply write $n = n_{\boldsymbol{a}}$. If all symbols appear equally often, i.e., $a_1 = \cdots = a_k = \alpha$, we use the abbreviation $\alpha^k := (a_1, \ldots, a_k)$. For example 123433153 is a $(2,1,4,1)$-multiset permutation, and 331232142144 is a 3^4-multiset permutation.

Clearly, multiset permutations are a generalization of permutations and combinations. Specifically, k-permutations are 1^k-multiset permutations, and (α, β)-combinations are (α, β)-multiset permutations (up to shifting the symbol names $1, 2 \mapsto 0, 1$). Stachowiak [37] showed that $(a_1, \ldots, a_k)$-multiset permutations can be generated by adjacent transpositions if and only if at least two of the a_i are odd.

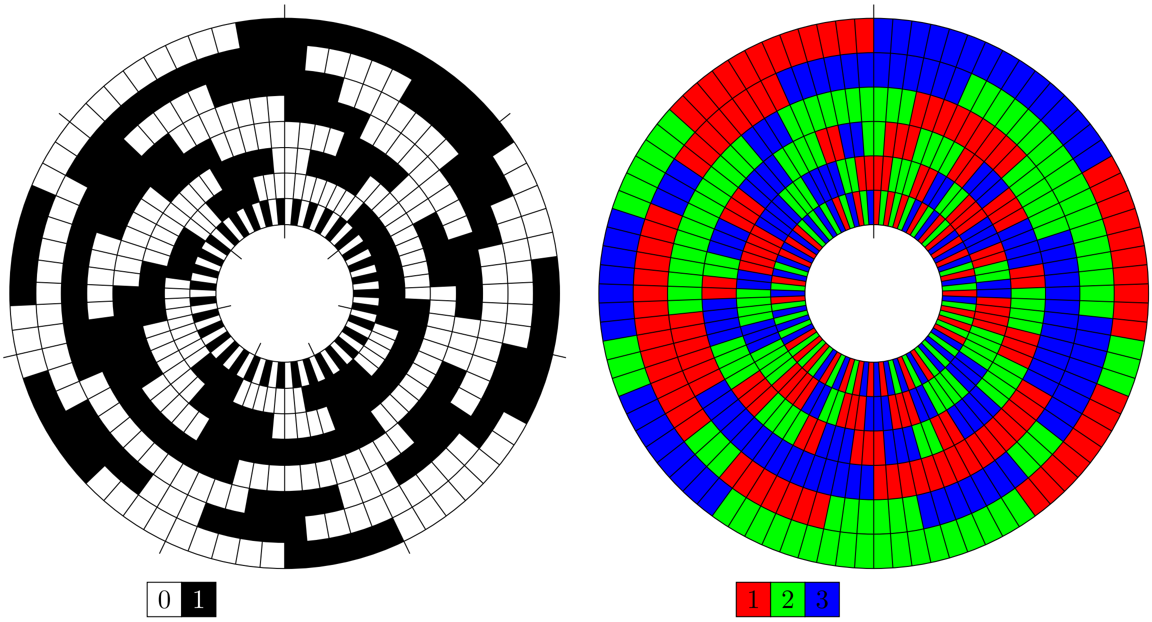

(a) $(a_1, a_2) = (4, 4)$

(b) $(a_1, a_2, a_3) = (2, 2, 2)$

Figure 2 Star transposition Gray codes for (a) (4, 4)- and (b) (2, 2, 2)-multiset permutations. The strings are arranged in clockwise order, starting at 12 o'clock, with the first entry on the inner track, and the last entry on the outer track. As every star transposition changes the first entry, the color on the inner track changes in every step.

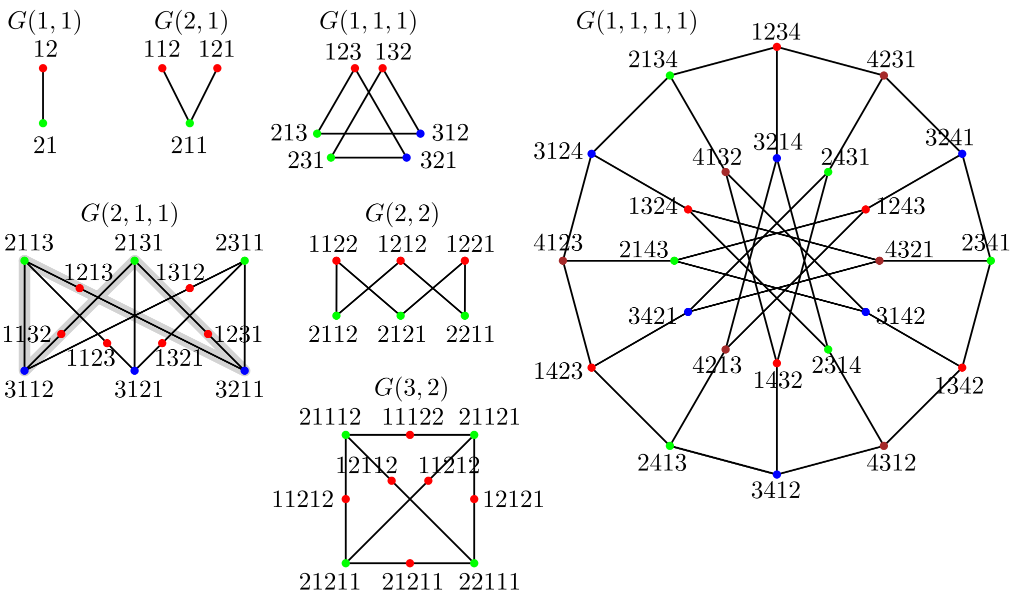

Figure 3 Star transposition graphs $G(\boldsymbol{a})$ for several small multiset permutations $\boldsymbol{a}$. Vertices are colored according to the first entry of the multiset permutations, and these color classes form independent sets. In $G(2, 1, 1)$, an odd cycle is highlighted.

Shen and Williams [34] conjectured that all α^k-multiset permutations can be generated by star transpositions, for any $\alpha \geq 1$ and $k \geq 2$. We state their conjecture in terms of Hamilton cycles in a suitably defined graph, as follows. We write $\Pi(\boldsymbol{a}) = \Pi(a_1, \ldots, a_k)$ for the set of all $(a_1, \ldots, a_k)$-multiset permutations. Moreover, we let $G(\boldsymbol{a}) = G(a_1, \ldots, a_k)$ denote the graph on the vertex set $\Pi(\boldsymbol{a}) = \Pi(a_1, \ldots, a_k)$ with an edge between any two multiset permutations that differ in a star transposition, i.e., in swapping the first entry of the multiset permutation with any entry at positions $2, \ldots, n$ that is distinct from the first one. Figure 3 shows various examples of the graph $G(\boldsymbol{a})$. When denoting specific multiset permutations we sometimes omit commas and brackets for brevity, for example $1312214 \in \Pi(3, 2, 1, 1)$.

▶ **Conjecture 1** ([34]). *For any $\alpha \geq 1$ and $k \geq 2$, the graph $G(\alpha^k)$ has a Hamilton cycle.*

In this and the following statements, the single edge $G(1, 1)$ is also considered a cycle, as it gives a cyclic Gray code. Note that $G(a_1, \ldots, a_k)$ is vertex-transitive if and only if $a_1 = \cdots = a_k =: \alpha$. In this case, Conjecture 1 is an interesting instance of the aforementioned conjecture of Lovász [21] on Hamilton paths in vertex-transitive graphs.

Evidence for Conjecture 1 comes from the results mentioned in Sections 1.1 and 1.2 on generating permutations by star transpositions and the solution of the middle levels conjecture, respectively, formulated in terms of the graph $G(\boldsymbol{a})$ below. These known results settle the boundary cases $\alpha = 1$ and $k \geq 2$, and $\alpha \geq 1$ and $k = 2$, respectively, of Conjecture 1.

▶ **Theorem 2** (Ehrlich; [20]; [36]). *For any $k \geq 2$, the graph $G(1^k)$ has a Hamilton cycle.*

▶ **Theorem 3** ([23, 13]). *For any $\alpha \geq 1$, the graph $G(\alpha, \alpha)$ has a Hamilton cycle.*

In their paper, Shen and Williams also provided an ad-hoc solution for the first case of their conjecture that is not covered by Theorems 2 and 3, namely a Hamilton cycle in $G(2, 2, 2)$, which is displayed in Figure 2 (b).

We approach Conjecture 1 by tackling the following even more general question: For which frequency vectors $\boldsymbol{a} = (a_1, \ldots, a_k)$ does the graph $G(\boldsymbol{a})$ have a Hamilton cycle? By renaming symbols, we may assume w.l.o.g. that the entries of the vector $\boldsymbol{a}$ are non-increasing, i.e.,

$$a_1 \geq a_2 \geq \cdots \geq a_k. \tag{1}$$

We can thus think of the vector $\boldsymbol{a}$ as an integer partition of n.

2 Our results

For any $i \in [n]$ and $c \in [k]$, we write $\Pi(\boldsymbol{a})^{i,c}$ for the set of all multiset permutations from $\Pi(\boldsymbol{a})$ whose ith symbol equals c. Note that every star transposition changes the first entry; see the inner track of each of the two wheels in Figure 2. As a consequence, $G(\boldsymbol{a})$ is a k-partite graph with partition classes $\Pi(\boldsymbol{a})^{1,1}, \ldots, \Pi(\boldsymbol{a})^{1,k}$; see Figure 3. Moreover, the partition class $\Pi(\boldsymbol{a})^{1,1}$ is a largest one because of (1). This k-partition of the graph $G(\boldsymbol{a})$ is a potential obstacle for the existence of Hamilton cycles and paths. Specifically, if one partition class is larger than all others combined, then there cannot be a Hamilton cycle, and if the size difference is more than 1, then there cannot be a Hamilton path.

We capture this by defining a parameter $\Delta(\boldsymbol{a})$ for any integer partition $\boldsymbol{a} = (a_1, \ldots, a_k)$ as

$$\Delta(\boldsymbol{a}) := n - 2a_1 = -a_1 + \sum_{i=2}^{k} a_i. \tag{2}$$

We will see that if $\Delta(\boldsymbol{a}) < 0$, then the partition class $\Pi(\boldsymbol{a})^{1,1}$ of the graph $G(\boldsymbol{a})$ is larger than all others combined, excluding the existence of Hamilton cycles. On the other hand, if $\Delta(\boldsymbol{a}) \geq 0$, then every partition class of the graph $G(\boldsymbol{a})$ is at most as large as all others combined (equality holds if $\Delta(\boldsymbol{a}) = 0$), which does not exclude the existence of a Hamilton cycle. The cases with $\Delta(\boldsymbol{a}) = 0$ lie on the boundary between the two regimes, and they are the hardest in terms of proving Hamiltonicity. These cases can be seen as generalizations of the middle levels conjecture, namely the case $\boldsymbol{a} = (\alpha, \alpha)$ captured by Theorem 3, which also satisfies $\Delta(\boldsymbol{a}) = 0$.

▶ **Theorem 4.** *For any integer partition $\boldsymbol{a} = (a_1, \ldots, a_k)$ with $\Delta(\boldsymbol{a}) < 0$ the graph $G(\boldsymbol{a})$ does not have a Hamilton cycle, and it does not have a Hamilton path unless $\boldsymbol{a} = (2, 1)$.*

For $k = 2$ symbols, the condition $\Delta(\boldsymbol{a}) < 0$ is equivalent to $a_1 > a_2$, i.e., there is no star transposition Gray code for "unbalanced" combinations.

We now discuss the cases $\Delta(\boldsymbol{a}) \geq 0$. Our first main goal is to reduce all cases with $\Delta(\boldsymbol{a}) > 0$ to cases with $\Delta(\boldsymbol{a}) = 0$. For doing so, it is helpful to consider stronger notions of Hamiltonicity. Specifically, we consider Hamilton-connectedness and Hamilton-laceability, which have been heavily studied (see [1, 2, 6, 14, 16, 35]). A graph is called *Hamilton-connected* if there is a Hamilton path between any two distinct vertices. A bipartite graph is called *Hamilton-laceable* if there is a Hamilton path between any pair of vertices from the two partition classes. In general, the graphs $G(\boldsymbol{a})$ are not bipartite, so we say that $G(\boldsymbol{a})$ is *1-laceable* if there is a Hamilton path between any vertex in $\Pi(\boldsymbol{a})^{1,1}$ and any vertex not in $\Pi(\boldsymbol{a})^{1,1}$, i.e., between any vertex with first symbol 1 and any vertex with first symbol distinct from 1.

This approach is inspired by the following result of Tchuente [39], who strengthened Theorem 2 considerably.

▶ **Theorem 5** **([39]).** *For any $k \geq 4$, the graph $G(1^k)$ is Hamilton-laceable.*

The key insight is that proving a stronger property makes the proof easier and shorter, because the inductive statement is more powerful and flexible; see Section 3.2 below. Encouraged by this, we raise the following conjecture about graphs $G(\boldsymbol{a})$ with $\Delta(\boldsymbol{a}) = 0$. It is another natural and far-ranging generalization of the middle levels conjecture, which we support by extensive computer experiments and by proving some special cases.

▶ **Conjecture 6.** *For any integer partition $\boldsymbol{a} = (a_1, \ldots, a_k)$ with $\Delta(\boldsymbol{a}) = 0$ the graph $G(\boldsymbol{a})$ is Hamilton-1-laceable, unless $\boldsymbol{a} = (2, 2)$.*

The exceptional graph $G(2, 2)$ mentioned in this conjecture is a 6-cycle; see Figure 3. Assuming the validity of this conjecture, we settle all cases $G(\boldsymbol{a})$ with $\Delta(\boldsymbol{a}) > 0$ in the strongest possible sense. While being a conditional result, the main purpose of this theorem is to reduce all cases $\Delta(\boldsymbol{a}) \geq 0$ to the boundary cases $\Delta(\boldsymbol{a}) = 0$.

▶ **Theorem 7.** *Conditional on Conjecture 6, for any integer partition $\boldsymbol{a} = (a_1, \ldots, a_k)$ with $\Delta(\boldsymbol{a}) > 0$ the graph $G(\boldsymbol{a})$ is Hamilton-connected, unless $\boldsymbol{a} = (1, 1, 1)$ or $\boldsymbol{a} = 1^k$ for $k \geq 4$, and possibly unless $\boldsymbol{a} = (\alpha, \alpha, 1)$ for $\alpha \geq 3$.*

The dependence of Theorem 7 on Conjecture 6 can be captured more precisely. Specifically, $G(\boldsymbol{a})$ with $\Delta(\boldsymbol{a}) > 0$ is shown to be Hamilton-connected, assuming that $G(\boldsymbol{b})$ with $\Delta(\boldsymbol{b}) = 0$ is Hamilton-1-laceable for all integer partitions $\boldsymbol{b}$ that are majorized componentwise by $\boldsymbol{a}$.

The three exceptions mentioned in Theorem 7 are well understood: Specifically, $G(1, 1, 1)$ is a 6-cycle; see Figure 3. Furthermore, $G(1^k)$ for $k \geq 4$ is Hamilton-laceable by Theorem 5. Lastly, we will show that $G(\alpha, \alpha, 1)$ for $\alpha \geq 3$ satisfies a variant of Hamilton-laceability, which also guarantees a Hamilton cycle. In fact, we believe that $G(\alpha, \alpha, 1)$ is Hamilton-connected, but we cannot prove it.

We provide the following evidence for Conjecture 6. First of all, with computer help we verified that $G(\boldsymbol{a})$ is indeed Hamilton-1-laceable for all integer partitions $\boldsymbol{a} \neq (2,2)$ with $\Delta(\boldsymbol{a}) = 0$ that satisfy $n \leq 8$, i.e., for $\boldsymbol{a} \in \{(1,1),(2,1,1),(3,3),(3,2,1),(3,1,1,1),(4,4),(4,3,1),(4,2,2),(4,2,1,1),(4,1,1,1,1)\}$. Furthermore, we prove the case of $k = 2$ symbols unconditionally. Note that for $k = 2$, Hamilton-1-laceability is the same as Hamilton-laceability. Recall that $G(\alpha,\alpha)$ is isomorphic to the subgraph of the $(2\alpha-1)$-dimensional hypercube induced by the middle two levels, so the following result is a considerable strengthening of Theorem 3, the middle levels theorem.

▶ **Theorem 8.** *For any $\alpha \geq 3$, the graph $G(\alpha,\alpha)$ is Hamilton-laceable.*

We also have the following (unconditional) result for $k = 3$ symbols.

▶ **Theorem 9.** *For any $\alpha \geq 2$, the graph $G(\alpha,\alpha-1,1)$ has a Hamilton cycle.*

Lastly, we consider integer partitions $\boldsymbol{a} = (a_1,\ldots,a_k)$, $k \geq 3$, with $\Delta(\boldsymbol{a}) \geq 0$ and an upper bound on the part size, i.e., $a_1 \leq \alpha$ for some constant α. By the remarks after Theorem 7, the inductive proof of the theorem for such integer partitions only relies on Conjecture 6 being satisfied for integer partitions with the same upper bound on the part size. For any fixed bound α, there are only finitely many such partitions with $\Delta(\boldsymbol{a}) = 0$ that can be checked by computer. For example, for $\alpha = 4$ these are $\boldsymbol{a} \in \{(2,1,1),(3,2,1),(3,1,1,1),(4,3,1),(4,2,2),(4,2,1,1),(4,1,1,1,1)\}$. This yields the following (unconditional) result.

▶ **Theorem 10.** *For $\alpha \in \{2,3,4\}$ and any integer partition $\boldsymbol{a} = (a_1,\ldots,a_k)$ with $\Delta(\boldsymbol{a}) > 0$ and $a_1 = \alpha$, the graph $G(\boldsymbol{a})$ is Hamilton-connected.*

In words, Theorem 10 settles all integer partitions $\boldsymbol{a}$ whose largest part is at most 4. In particular, this settles the cases $\alpha \in \{2,3,4\}$ and $k \geq 2$ of Shen and Williams' Conjecture 1 in a rather strong sense.

3 Proof ideas

In this section, we give a high-level overview of the main ideas and techniques used in our proofs. No formal proofs of our results are presented in this extended abstract due to space constraints, but they can be found in the preprint [12].

3.1 The case $\Delta(\boldsymbol{a}) < 0$

The main idea for proving Theorem 4 is that if $\Delta(\boldsymbol{a}) < 0$, then the partition class $\Pi(\boldsymbol{a})^{1,1}$ of the graph $G(\boldsymbol{a})$ is larger than all others combined, which excludes the existence of a Hamilton cycle. To exclude the existence of a Hamilton path, we show that the size difference is strictly more than 1, unless $\boldsymbol{a} = (2,1)$. Note that the graph $G(2,1)$ is the path on three vertices, so in this case the size difference is precisely 1. These arguments are based on straightforward algebraic manipulations involving multinomial coefficients.

3.2 The case $\Delta(\boldsymbol{a}) > 0$

To prove Theorem 7, it is convenient to think of an integer partition $\boldsymbol{a} = (a_1,\ldots,a_k)$ as an infinite non-increasing sequence $(a_1,a_2,\ldots)$, with only k nonzero entries at the beginning. Given two such integer partitions $\boldsymbol{a} = (a_1,a_2,\ldots)$ and $\boldsymbol{b} = (b_1,b_2,\ldots)$, we write $\boldsymbol{b} \prec \boldsymbol{a}$ if $b_i \leq a_i$ for all $i \geq 1$. Integer partitions with the partial order $\prec$ form a lattice, which is the

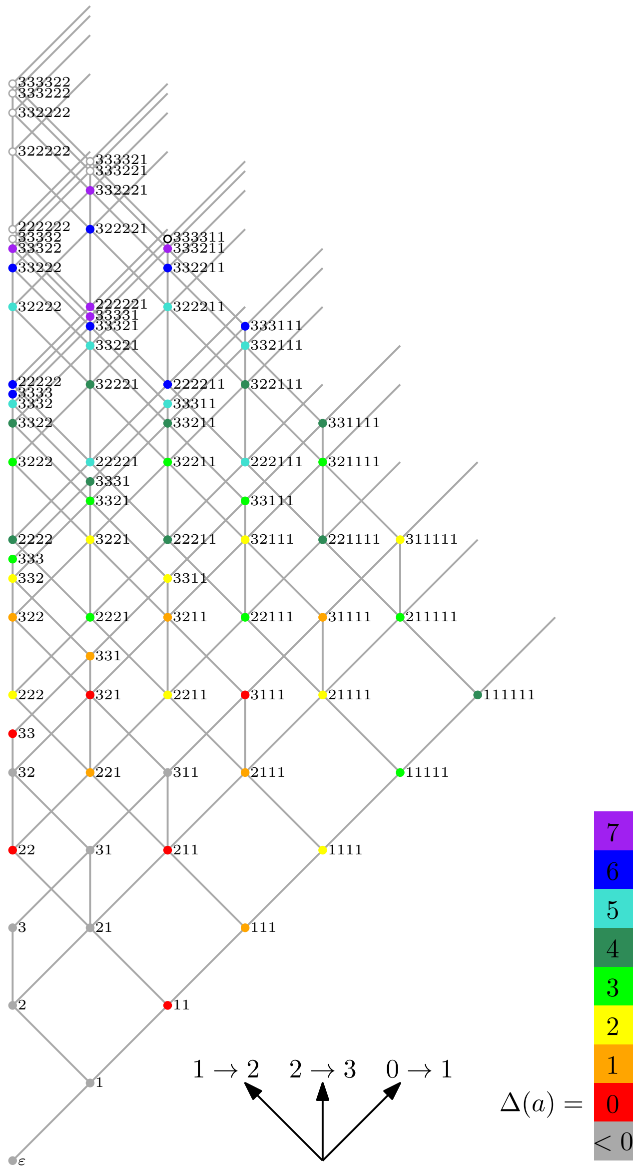

Figure 4 The lattice of integer partitions $\boldsymbol{a} = (a_1, \ldots, a_k)$ with largest part $a_1 \leq 3$. The coordinates are projected into three dimensions depending on which value is increased.

sublattice of the infinite lattice $\mathbb{N}^{\mathbb{N}}$ cut out by the hyperplanes defined by (1); see Figure 4. The cover relations in this lattice are given by decrementing any of the a_i for which $a_i > a_{i+1}$. We write $\boldsymbol{b} \prec \boldsymbol{a}$ for partitions $\boldsymbol{a} \neq \boldsymbol{b}$ if $\boldsymbol{b} \prec \boldsymbol{a}$ and there is no $\boldsymbol{c} \notin \{\boldsymbol{a}, \boldsymbol{b}\}$ with $\boldsymbol{b} \prec \boldsymbol{c} \prec \boldsymbol{a}$.

In this lattice of integer partitions, the hyperplane defined by $\Delta(\boldsymbol{a}) = 0$ separates the cases where Hamiltonicity is impossible, which lie on the side of the hyperplane where $\Delta(\boldsymbol{a}) < 0$ (Theorem 4), from the cases where Hamiltonicity can be established more easily, which lie on the side of the hyperplane where $\Delta(\boldsymbol{a}) > 0$ (Theorem 7). The cases $\Delta(\boldsymbol{a}) = 0$ on the hyperplane are the hardest ones (Conjecture 6).

Our proof of Theorem 7 proceeds by induction in this partition lattice and establishes the Hamiltonicity of $G(\boldsymbol{a})$ by using the Hamiltonicity of $G(\boldsymbol{b})$ for all integer partitions $\boldsymbol{b} \prec \boldsymbol{a}$, where Conjecture 6 serves as the base case of the induction. This is based on the observation that fixing one of the symbols at positions $2, \ldots, n$ in $G(\boldsymbol{a})$ yields subgraphs that are isomorphic to $G(\boldsymbol{b})$ for $\boldsymbol{b} \prec \boldsymbol{a}$.

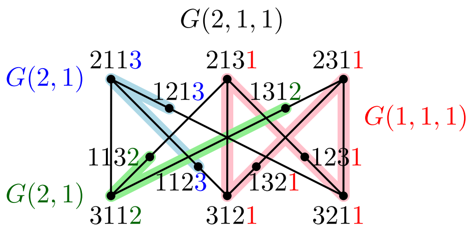

Figure 5 Decomposing $G(2,1,1)$ into three subgraphs by fixing the last symbol.

Specifically, for any $\boldsymbol{a} = (a_1, \ldots, a_k)$, $i = 2, \ldots, n$, and $c \in [k]$, the subgraph of $G(\boldsymbol{a})$ induced by the vertex set $\Pi(\boldsymbol{a})^{i,c}$ is isomorphic to $G(\boldsymbol{b})$ where $\boldsymbol{b}$ is the partition obtained from $\boldsymbol{a}$ by decreasing a_i by 1 (and possibly sorting the resulting numbers non-increasingly); see Figure 5.

Moreover, for any $\boldsymbol{b} \prec \boldsymbol{a}$ we have $\Delta(\boldsymbol{b}) = \Delta(\boldsymbol{a}) - 1$ or $\Delta(\boldsymbol{b}) = \Delta(\boldsymbol{a}) + 1$. In particular, if $\Delta(\boldsymbol{a}) > 0$, then we have $\Delta(\boldsymbol{b}) \geq 0$. For example, the vertex set of $G(\boldsymbol{a})$ for $\boldsymbol{a} = (3,2,2)$ ($\Delta(\boldsymbol{a}) = 1$) can be partitioned into one copy of $G(\boldsymbol{b})$ for $\boldsymbol{b} = (2,2,2)$ (if the fixed symbol is $c = 1$; $\Delta(\boldsymbol{b}) = 2$) and two copies of $G(\boldsymbol{b}')$ for $\boldsymbol{b}' = (3,2,1)$ (if the fixed symbol is $c = 2$ or $c = 3$; $\Delta(\boldsymbol{b}') = 0$). Therefore, we may construct a Hamilton path in $G(3,2,2)$ by gluing together paths in each of these three subgraphs which exist by induction.

While conceptually simple, implementing this idea incurs considerable technical obstacles, in particular for some of the graphs $G(\boldsymbol{a})$ with $\Delta(\boldsymbol{a}) = 1$, i.e., instances that are very close to the hyperplane $\Delta(\boldsymbol{a}) = 0$. The proof is split into several interdependent lemmas, and it is the technically most demanding part of our work.

Theorem 10 follows immediately from the inductive proof of Theorem 7 and by settling finitely many cases with computer help.

To further illustrate the ideas outlined before, we close this section by reproducing Tchuente's proof of Theorem 5.

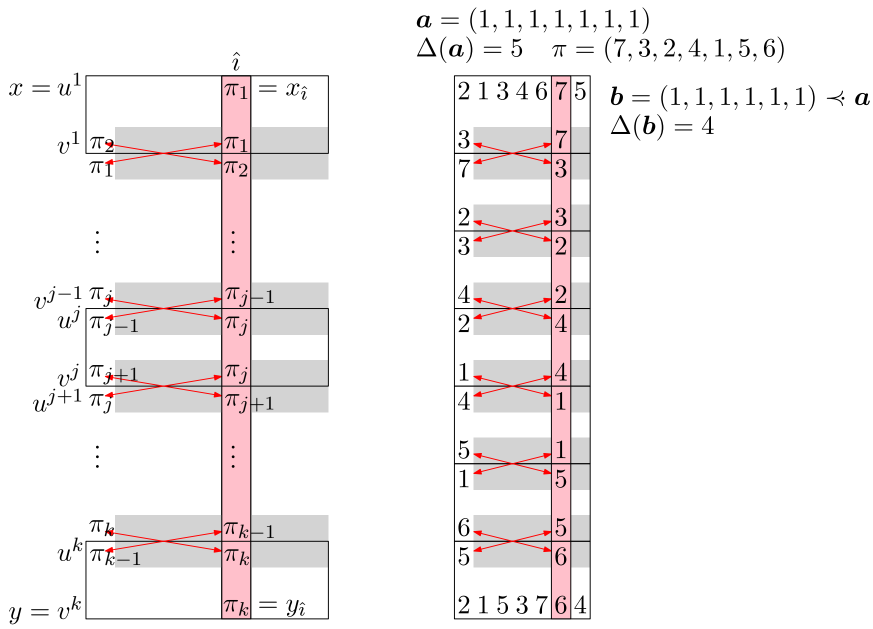

Figure 6 Illustration of the proof of Theorem 5. The left hand side shows the general schematic partitioning of the graph $G(1^k)$ into blocks, each of which is a copy of $G(1^{k-1})$, by fixing symbols at position $\hat{\imath}$. The right hand side shows a concrete example for $k = 7$.

Proof of Theorem 5. To prove that $G(1^k)$ is Hamilton-laceable, we proceed by induction on k. The induction basis $k = 4$ can be checked by straightforward case analysis. For the induction step, we assume that $G(1^{k-1})$, $k \geq 5$, is Hamilton-laceable, and we prove that $G(1^k)$ is also Hamilton-laceable. Note that $1^{k-1} \prec 1^k$ and $\Delta(1^k) = k - 2 = \Delta(1^{k-1}) + 1$. The following arguments are illustrated in Figure 6. The partition classes of the graph $G(1^k)$ are given by the parity of the permutations, i.e., by the number of inversions. Therefore, we consider two distinct permutations x and y of $[k]$ with opposite parity, and we need to show how to connect them by a Hamilton path in $G(1^k)$. As $x \neq y$, there is a position $\hat{\imath} > 1$ in which x and y differ, i.e., $x_{\hat{\imath}} \neq y_{\hat{\imath}}$, and this is the position that we will fix to different symbols. Specifically, we choose a permutation π of $[k]$ such that $\pi_1 = x_{\hat{\imath}}$ and $\pi_k = y_{\hat{\imath}}$. The permutation π captures the order in which we will fix symbols at position $\hat{\imath}$. We then choose permutations u^j, v^j of $[k]$, $j = 1, \ldots, k$, satisfying $u^1 = x$, $v^k = y$, and such that u^j is obtained from v^{j-1} by a star transposition of the symbol π_j at position 1 with the symbol π_{j-1} at position $\hat{\imath}$, for all $j = 2, \ldots, k$. Moreover, we choose u^j and v^j such that the parity of the number of inversions after removing the symbol π_j is opposite. Specifically, for u^j this parity is the same as for $u^1 = x$ if and only if π_j has the same parity as π_1, and for v^j this parity is the same as for $v^k = y$ if and only if π_j has the same parity as π_k. For each $j = 1, \ldots, k$, we now consider the permutations whose $\hat{\imath}$th entry equals π_j (formally, this is the set $\Pi(1^k)^{\hat{\imath},\pi_j}$). Clearly, the subgraph of $G(1^k)$ induced by these permutations is isomorphic to $G(1^{k-1})$. In other words, by removing the $\hat{\imath}$th entry and renaming entries to $1, \ldots, k-1$, we obtain permutations of $[k-1]$. Consequently, by induction there is a path in $G(1^k)$ that visits all permutations whose $\hat{\imath}$th entry equals π_j and that connects u^j to v^j. The concatenation of those k paths obtained by induction is the desired Hamilton path in $G(1^k)$ from x to y, which completes the induction step. ◀

Note that the constraints imposed on the permutations π and u^j, v^j, $j = 1, \ldots, k$, in this proof are very mild, and leave a lot of room for modifications to construct many different Hamilton paths, possibly so as to satisfy some additional conditions.

3.3 The case $\Delta(\boldsymbol{a}) = 0$

Our proofs of Theorems 8 and 9 build on ideas introduced in the papers [13, 22].

Specifically, the first step in proving Theorem 8 is to build a cycle factor in the graph $G(\alpha, \alpha)$, i.e., a collection of disjoint cycles in the graph that together visit all vertices. We then choose vertices x and y from the two partition classes of the graph that we want to connect by a Hamilton path. In this we can take into account automorphisms of $G(\alpha, \alpha)$, i.e., for proving laceability only certain pairs of vertices x and y in the two partition classes have to be considered. In the next step, we join a small subset of cycles from the factor, including the ones containing x and y, to a short path between x and y. This is achieved by taking the symmetric difference of the edge set of the cycle factor with a carefully chosen path P from x to y that alternately uses edges on one of the cycles from the factor and edges that go between two such cycles; see Figure 7 (a)+(b). In the last step, we join the remaining cycles of the factor to the path between x and y, until we end with a Hamilton path from x to y. Each such joining is achieved by taking the symmetric difference of the cycle factor with a suitably chosen 6-cycle; see Figure 7 (b)+(c).

It was shown in [13] that the cycles of the aforementioned cycle factor in $G(\alpha, \alpha)$ are bijectively equivalent to plane trees with α vertices, and the joining operations via 6-cycles can be interpreted combinatorially as local change operations between two such plane trees. To prove Theorem 9, we first generalize the construction of this cycle factor in the graph $G(\alpha, \alpha)$ to a cycle factor in any graph $G(\boldsymbol{a})$, $\boldsymbol{a} = (a_1, \ldots, a_k)$, with $\Delta(\boldsymbol{a}) = 0$. It turns out that the cycles of this generalized factor can be interpreted combinatorially as *vertex-labeled*

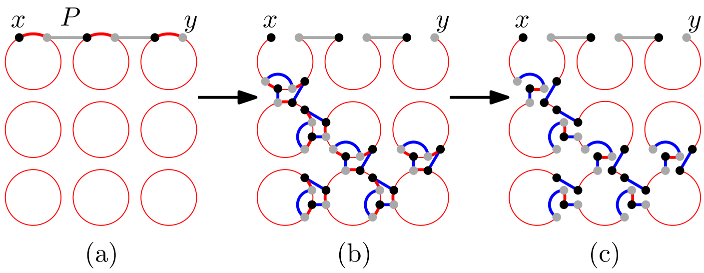

Figure 7 Strategy of the proof of Theorem 8.

plane trees, where exactly a_i vertices have the label i for $i = 2, \ldots, k$. If $\boldsymbol{a} = (\alpha, \alpha)$, then all vertex labels are the same, so we can consider the trees as unlabeled. Moreover, the joining 6-cycles in $G(\alpha, \alpha)$ generalize nicely to joining 12- or 6-cycles in $G(\boldsymbol{a})$ with $\Delta(\boldsymbol{a}) = 0$, and they correspond to local change operations involving sets T of labeled plane trees with $|T| \in \{2, 3, 5, 6\}$, depending on the location of vertex labels. For proving Theorem 9, we consider the special case $\boldsymbol{a} = (\alpha, \alpha - 1, 1)$, i.e., exactly one vertex in the plane trees is labeled differently from all other vertices, and we show that there is a choice of joining cycles so that the symmetric difference with the cycle factor yields a Hamilton cycle in the graph $G(\boldsymbol{a})$.

4 Open questions

We conclude this paper with the following open questions.

- We believe that the ideas outlined in Section 3.3 to prove that $G(\alpha, \alpha - 1, 1)$ has a Hamilton cycle are in principle suitable to prove that $G(\boldsymbol{a})$ has a Hamilton cycle for all $\boldsymbol{a}$ with $\Delta(\boldsymbol{a}) = 0$, which would be an important first step towards a proof of Conjecture 6. In particular, the construction of the cycle factor and gluing tuples based on vertex-labeled plane trees are fully general. The main difficulty in combining these ingredients lies in the fact that some gluing cycles join more than two cycles from the factor (namely 3, 5, or 6 cycles) to a single cycle, and in this case the resulting interactions between different gluing cycles seem to be hard to control.
- We conjecture that $G(\alpha, \alpha, 1)$ for $\alpha \geq 2$ is Hamilton-connected, just as all other graphs $G(\boldsymbol{a})$ with $\Delta(\boldsymbol{a}) > 0$ covered by Theorem 7. However, we are unable to prove this based on Conjecture 6. With computer help, we verified that $G(\alpha, \alpha, 1)$ is Hamilton-connected for $\alpha = 2, 3, 4$. Proving this in general would streamline our proof of Theorem 7 considerably, as it would make several auxiliary lemmas redundant. To prove that $G(\alpha, \alpha, 1)$ is Hamilton-connected, it may help to establish a Hamiltonicity property for graphs $G(\boldsymbol{a})$ with $\Delta(\boldsymbol{a}) = 0$ and $k \geq 3$ that is stronger than 1-laceability. Specifically, in addition to a Hamilton path between any vertex in $\Pi(\boldsymbol{a})^{1,1}$ and any vertex not in $\Pi(\boldsymbol{a})^{1,1}$, we may also ask for a Hamilton path between any two distinct vertices in $\Pi(\boldsymbol{a})^{1,1}$. We checked by computer whether $G(\boldsymbol{a})$ has this stronger property for $\boldsymbol{a} \in \{(2, 1, 1), (3, 2, 1), (3, 1, 1, 1), (4, 3, 1), (4, 2, 2), (4, 2, 1, 1), (4, 1, 1, 1, 1)\}$, and it was satisfied in all cases except for $\boldsymbol{a} = (2, 1, 1)$.
- While our proofs are constructive, they are far from yielding efficient algorithms for computing the corresponding Gray codes. Ideally, one would like algorithms whose running time is polynomial in n per generated multiset permutation of length n. Such algorithms are known for the Hamilton cycles mentioned in Theorems 2 and 3, see [19, Section 7.2.1.2] and [24], respectively. An interesting direction could be to explore greedy algorithms for generating multiset permutations by star transpositions, which may yield much simpler constructions to start with, cf. [41, 4, 31].

- Knuth raised the question whether there are star transposition Gray codes for (α, α)-combinations whose flip sequence can be partitioned into $2\alpha - 1$ blocks, such that each block is obtained from the previous one by adding $+1$ modulo $2\alpha - 1$. This problem is a strengthening of the middle levels conjecture, and it was answered affirmatively in [22]. The Gray code for $(4, 4)$-combinations shown in Figure 2 (a) has such a 7-fold cyclic symmetry. We can ask more generally: Are there star transposition Gray codes for multiset permutations whose flip sequence can be partitioned into $n - 1$ blocks, such that each block is obtained from the previous one by adding $+1$ modulo $n - 1$? Figure 8 shows an ad-hoc solution for $(2, 2, 2)$-multiset permutations with 5-fold cyclic symmetry.

- A more general version of the problem considered in this paper is the following: We consider an alphabet $\{1, \ldots, k\}$ of size $k \geq 2$, and frequencies $a_1, \ldots, a_k \geq 1$ that specify that symbol i appears exactly a_i times for all $i = 1, \ldots, k$. Moreover, there is an additional integer parameter s with $1 \leq s \leq k - 1$ that has the following significance. The objects to be generated are all pairs (S, x), where S is a set or string of s distinct symbols, and x is a string of the remaining $n - s$ symbols, where $n := a_1 + \cdots + a_k$. A star transposition swaps one symbol from S with one symbol from x that is currently not in S, and the question is whether there is a star transposition Gray code for all those objects.

  Note that multiset permutations considered in this paper are the special case when $s = 1$. Andrea Sportiello suggested this problem with $a_1 = \cdots = a_k = \alpha$ (uniform frequency) and S being a set as a generalization of the middle levels conjecture ($s = 1$, $k = 2$, $a_1 = a_2 = \alpha$). Moreover, Ajit A. Diwan suggested this problem with $a_1 = \cdots = a_k = \alpha$ (uniform frequency) and a set S of size $s = k - 1$. Note that the uniform frequency case is particularly interesting, as the underlying flip graph for this problem is vertex-transitive if and only if $a_1 = \cdots = a_k$ (recall Lovász' conjecture [21]).

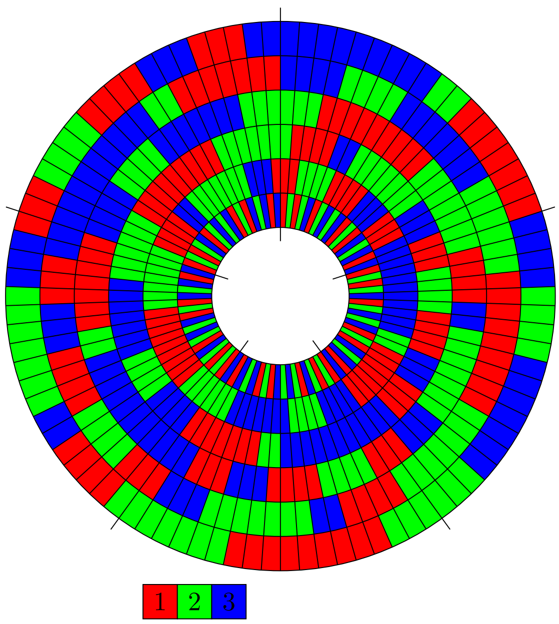

Figure 8 Star transposition Gray code for $(2, 2, 2)$-multiset permutations with 5-fold cyclic symmetry.

References

1. B. Alspach and Y. Qin. Hamilton-connected Cayley graphs on Hamiltonian groups. *European J. Combin.*, 22(6):777–787, 2001. doi:10.1006/eujc.2001.0456.
2. T. Araki. Hyper Hamiltonian laceability of Cayley graphs generated by transpositions. *Networks*, 48(3):121–124, 2006. doi:10.1002/net.20126.
3. M. Buck and D. Wiedemann. Gray codes with restricted density. *Discrete Math.*, 48(2-3):163–171, 1984. doi:10.1016/0012-365X(84)90179-1.
4. B. Cameron, J. Sawada, and A. Williams. A Hamilton cycle in the k-sided pancake network, 2021. arXiv:2103.09256.
5. P. Chase. Combination generation and graylex ordering. *Congr. Numer.*, 69:215–242, 1989. Eighteenth Manitoba Conference on Numerical Mathematics and Computing (Winnipeg, MB, 1988).
6. C. C. Chen and N. F. Quimpo. On strongly Hamiltonian abelian group graphs. In *Combinatorial mathematics, VIII (Geelong, 1980)*, volume 884 of *Lecture Notes in Math.*, pages 23–34. Springer, Berlin-New York, 1981.
7. R. C. Compton and S. G. Williamson. Doubly adjacent Gray codes for the symmetric group. *Linear and Multilinear Algebra*, 35(3-4):237–293, 1993. doi:10.1080/03081089308818261.
8. P. F. Corbett. Rotator graphs: An efficient topology for point-to-point multiprocessor networks. *IEEE Transactions on Parallel and Distributed Systems*, 3:622–626, 1992.
9. P. Diaconis and R. Graham. *Magical mathematics*. Princeton University Press, Princeton, NJ, 2012. The mathematical ideas that animate great magic tricks, With a foreword by Martin Gardner.
10. 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. doi:10.1145/2422.322413.
11. P. Eades and B. McKay. An algorithm for generating subsets of fixed size with a strong minimal change property. *Inform. Process. Lett.*, 19(3):131–133, 1984. doi:10.1016/0020-0190(84)90091-7.
12. P. Gregor, A. Merino, and T. Mütze. arXiv:2108.07465.
13. P. Gregor, T. Mütze, and J. Nummenpalo. A short proof of the middle levels theorem. *Discrete Analysis*, 2018:8:12 pp., 2018.
14. F. Harary and M. Lewinter. Hypercubes and other recursively defined Hamilton laceable graphs. *Congr. Numer.*, 60:81–84, 1987. Eighteenth Southeastern International Conference on Combinatorics, Graph Theory, and Computing (Boca Raton, Fla., 1987).
15. I. Havel. Semipaths in directed cubes. In *Graphs and other combinatorial topics (Prague, 1982)*, volume 59 of *Teubner-Texte Math.*, pages 101–108. Teubner, Leipzig, 1983.
16. S.-Y. Hsieh, G.-H. Chen, and C.-W. Ho. Hamiltonian-laceability of star graphs. *Networks*, 36(4):225–232, 2000. doi:10.1002/1097-0037(200012)36:4<225::AID-NET3>3.0.CO;2-G.
17. T. Jenkyns and D. McCarthy. Generating all k-subsets of $\{1 \cdots n\}$ with minimal changes. *Ars Combin.*, 40:153–159, 1995.
18. S. Johnson. Generation of permutations by adjacent transposition. *Math. Comp.*, 17:282–285, 1963.
19. D. E. Knuth. *The Art of Computer Programming. Vol. 4A. Combinatorial Algorithms. Part 1.* Addison-Wesley, Upper Saddle River, NJ, 2011.
20. V. L. Kompel'maher and V. A. Liskovec. Sequential generation of permutations by means of a transposition basis. *Kibernetika (Kiev)*, 3:17–21, 1975.
21. L. Lovász. Problem 11. In *Combinatorial Structures and Their Applications (Proc. Calgary Internat. Conf., Calgary, Alberta, 1969)*. Gordon and Breach, New York, 1970.
22. A. I. Merino, O. Mička, and T. Mütze. On a combinatorial generation problem of Knuth. In D. Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 735–743. SIAM, 2021. doi:10.1137/1.9781611976465.46.

23 T. Mütze. Proof of the middle levels conjecture. *Proc. Lond. Math. Soc.*, 112(4):677–713, 2016. doi:10.1112/plms/pdw004.

24 T. Mütze and J. Nummenpalo. A constant-time algorithm for middle levels Gray codes. *Algorithmica*, 82(5):1239–1258, 2020. doi:10.1007/s00453-019-00640-2.

25 A. Nijenhuis and H. Wilf. *Combinatorial algorithms.* Academic Press, New York-London, 1975. Computer Science and Applied Mathematics.

26 F. Ruskey. Adjacent interchange generation of combinations. *J. Algorithms*, 9(2):162–180, 1988. doi:10.1016/0196-6774(88)90036-3.

27 F. Ruskey. Combinatorial Gray code. In M.-Y. Kao, editor, *Encyclopedia of Algorithms*, pages 342–347. Springer, 2016.

28 F. Ruskey and C. D. Savage. Hamilton cycles that extend transposition matchings in Cayley graphs of S_n. *SIAM J. Discrete Math.*, 6(1):152–166, 1993. doi:10.1137/0406012.

29 F. Ruskey and A. Williams. The coolest way to generate combinations. *Discrete Math.*, 309(17):5305–5320, 2009. doi:10.1016/j.disc.2007.11.048.

30 C. D. Savage. A survey of combinatorial Gray codes. *SIAM Rev.*, 39(4):605–629, 1997. doi:10.1137/S0036144595295272.

31 J. Sawada and A. Williams. Greedy flipping of pancakes and burnt pancakes. *Discrete Appl. Math.*, 210:61–74, 2016. doi:10.1016/j.dam.2016.02.005.

32 J. Sawada and A. Williams. Solving the sigma-tau problem. *ACM Trans. Algorithms*, 16(1):Art. 11, 17 pp., 2020. doi:10.1145/3359589.

33 X. S. Shen and A. Williams. A 'hot potato' Gray code for permutations. *Electronic Notes in Discrete Mathematics*, 44:89–94, 2013. doi:10.1016/j.endm.2013.10.014.

34 X. S. Shen and A. Williams. A k-ary middle levels conjecture. In *Proceedings of the 23rd Thailand-Japan Conference on Discrete and Computational Geometry, Graphs, and Games*, 2021.

35 G. J. Simmons. Almost all n-dimensional rectangular lattices are Hamilton-laceable. In *Proceedings of the Ninth Southeastern Conference on Combinatorics, Graph Theory, and Computing (Florida Atlantic Univ., Boca Raton, Fla., 1978)*, Congress. Numer., XXI, pages 649–661. Utilitas Math., Winnipeg, Man., 1978.

36 P. J. Slater. Generating all permutations by graphical transpositions. *Ars Combin.*, 5:219–225, 1978.

37 G. Stachowiak. Hamilton paths in graphs of linear extensions for unions of posets. *SIAM J. Discrete Math.*, 5(2):199–206, 1992. doi:10.1137/0405016.

38 D. Tang and C. Liu. Distance-2 cyclic chaining of constant-weight codes. *IEEE Trans. Computers*, C-22:176–180, 1973.

39 M. Tchuente. Generation of permutations by graphical exchanges. *Ars Combin.*, 14:115–122, 1982.

40 H. F. Trotter. Algorithm 115: Perm. *Commun. ACM*, 5(8):434–435, 1962. doi:10.1145/368637.368660.

41 A. Williams. The greedy Gray code algorithm. In *Algorithms and Data Structures - 13th International Symposium, WADS 2013, London, ON, Canada, August 12-14, 2013. Proceedings*, pages 525–536, 2013. doi:10.1007/978-3-642-40104-6_46.

42 P. Winkler. *Mathematical puzzles: a connoisseur's collection.* A K Peters, Ltd., Natick, MA, 2004.

43 S. Zaks. A new algorithm for generation of permutations. *BIT*, 24(2):196–204, 1984. doi:10.1007/BF01937486.