

Gray codes and symmetric chains

Petr Gregor^{*1}, Sven Jäger^{†2}, Torsten Mütze^{‡2}, Joe Sawada^{§3}, Kaja Wille^{¶2}

¹Department of Theoretical Computer Science and Mathematical Logic,
Charles University, Prague, Czech Republic

²Institut für Mathematik, Technische Universität Berlin, Germany

³School of Computer Science, University of Guelph, Canada

ABSTRACT. We consider the problem of constructing a cyclic listing of all bitstrings of length $2n+1$ with Hamming weights in the interval $[n+1-\ell, n+\ell]$, where $1 \leq \ell \leq n+1$, by flipping a single bit in each step. This is a far-ranging generalization of the well-known middle two levels problem (the case $\ell = 1$). We provide a solution for the case $\ell = 2$ and solve a relaxed version of the problem for general values of ℓ , by constructing cycle factors for those instances. Our proof uses symmetric chain decompositions of the hypercube, a concept known from the theory of posets, and we present several new constructions of such decompositions. In particular, we construct four pairwise edge-disjoint symmetric chain decompositions of the n -dimensional hypercube for any $n \geq 12$.

KEYWORDS: Gray code, Hamilton cycle, hypercube, poset, symmetric chain

1. INTRODUCTION

Gray codes are named after Frank Gray, a researcher at Bell Labs, who described a simple method to generate all 2^n bitstrings of length n by flipping a single bit in each step [Gra53], now known as the binary reflected Gray code. This code found widespread use, e.g., in circuit design and testing, signal processing and error correction, data compression etc.; many more applications are mentioned in the survey [Sav97]. The binary reflected Gray code is also implicit in the well-known *Towers of Hanoi* puzzle and the *Chinese ring* puzzle that date back to the 19th century. The theory of Gray codes has developed considerably in the last decades, and the term is now used more generally to describe an exhaustive listing of any class of combinatorial objects where successive objects in the list differ by a small amount. In particular, such generation algorithms have been developed for several fundamental combinatorial objects of interest for computer scientists, such as bitstrings, permutations, partitions, trees etc., all of which are covered in depth in the most recent volume of Knuth's seminal series *The Art of Computer Programming* [Knu11].

Since the discovery of the binary reflected Gray code, there has been continued interest in developing Gray codes for bitstrings of length n that satisfy various additional constraints. For instance, a Gray code with the property that each bit is flipped (almost) the same number of times was first constructed by Tootill [Too53]. Goddyn and Gvozdzjak constructed an n -bit Gray code in which any two flips of the same bit are almost n steps apart [GG03], which is best possible. These are only two examples of a large body of work on possible Gray code transition sequences; see also [BR96, SvZ08, DDGŠ13]. Savage and Winkler constructed a Gray code that generates all 2^n bitstrings such that all bitstrings with Hamming weight k appear before all bitstrings with weight

*E-Mail: gregor@ktiml.mff.cuni.cz

†E-Mail: jaeger@math.tu-berlin.de

‡E-Mail: muetze@math.tu-berlin.de

§E-Mail: jsawada@uoguelph.ca

¶E-Mail: wille@math.tu-berlin.de




FIGURE 1. (a) The 5-cube with the (standard) symmetric chain decomposition \mathcal{D}_0 , where the edges along the chains are highlighted by thick lines. (b) Building a cycle factor through the middle four levels of the 5-cube as explained in the proof of Theorem 2 with SCDs $\mathcal{D} := \mathcal{D}_0$ (black) and $\mathcal{D}' := \overline{\mathcal{D}_0}$ (red). The edges that are removed from \mathcal{D} and \mathcal{D}' are dotted, so the solid and dashed edges are the two matchings M and M' whose union forms the cycle factor. It has three cycles of lengths 4, 4 and 22, visiting all 30 bitstrings with Hamming weight in the interval $[1, 4]$.

$k + 2$, for each $0 \leq k \leq n - 2$ [SW95], where the *Hamming weight* of a bitstring is the number of its 1-bits. They used this construction to tackle the infamous *middle two levels problem*, which asks for a cyclic listing of all bitstrings of length $2n + 1$ with weights in the interval $[n, n + 1]$ by flipping a single bit in each step. This problem was raised in the 1980s and received considerable attention in the literature (a detailed historic account is given in [Müt16]). A general existence proof for such a Gray code for any $n \geq 1$ has been found only recently [Müt16, GMN17], and an algorithm for computing it using $\mathcal{O}(1)$ amortized time and $\mathcal{O}(n)$ space was subsequently presented in [MN17]. The starting point of this work is the following more general problem raised in [Sav93, GŠ10].

Problem M (middle 2ℓ levels problem). *For any $n \geq 1$ and $1 \leq \ell \leq n + 1$, construct a cyclic listing of all bitstrings of length $2n + 1$ with Hamming weights in the interval $[n + 1 - \ell, n + \ell]$ by flipping a single bit in each step.*

The special case $\ell = 1$ of Problem M is the middle two levels problem mentioned before. The case $\ell = n + 1$ is solved by the binary reflected Gray code discussed in the beginning. Moreover, the cases $\ell = n$ and $\ell = n - 1$ were settled in [EHH01, LS03] and [GŠ10], respectively.

A natural framework for studying such Gray code problems is the n -dimensional hypercube Q_n , or n -cube for short, the graph formed by all bitstrings of length n , with an edge between any two bitstrings that differ in exactly one bit. The 5-cube is illustrated in Figure 1 (a). The k th level of the n -cube is the set of all bitstrings with Hamming weight exactly k . In this terminology, Problem M asks for a Hamilton cycle in the subgraph of the $(2n + 1)$ -cube induced by the middle 2ℓ levels.

The most general version of this problem is whether the subgraph of the n -cube induced by all levels in an arbitrary weight interval $[a, b]$ has an (almost) Hamilton cycle. This was solved in [GM18] for all possible values of $n \geq 1$ and $0 \leq a \leq b \leq n$, except in the cases when the length n of the bitstrings is odd and the levels a and b are symmetric around the middle, which is exactly Problem M. For all other cases that paper provides algorithms that generate each bitstring in those Gray codes in constant time.

1.1. Our results. In this work we solve the case $\ell = 2$ of Problem M, i.e., we construct a cyclic listing of all bitstrings of length $2n + 1$ with Hamming weights in the interval $[n - 1, n + 2]$.

Theorem 1. *For any $n \geq 1$, the subgraph of the $(2n + 1)$ -cube induced by the middle four levels has a Hamilton cycle.*

Combining Theorem 1 with the results from [GM18] shows more generally that the subgraph of the n -cube induced by any four consecutive levels has an ‘almost’ Hamilton cycle.¹

As another partial result towards Problem M, we show that the subgraph of the $(2n + 1)$ -cube induced by the middle 2ℓ levels has a cycle factor. A *cycle factor* is a collection of disjoint cycles which together visit all vertices of the graph. In particular, a Hamilton cycle is a cycle factor consisting only of a single cycle. Note here that the existence of a cycle factor for general values of ℓ is not an immediate consequence of Hall’s theorem, which is applicable only for $\ell = 1$ and $\ell = n + 1$, as only in those cases all vertices of the underlying graph have the same degree.

Theorem 2. *For any $n \geq 1$ and $1 \leq \ell \leq n + 1$, the subgraph of the $(2n + 1)$ -cube induced by the middle 2ℓ levels has a cycle factor.*

Our proof of Theorem 2 is concise and illustrative, and it motivates the subsequent discussion, so we present it right now. It uses a well-known concept from the theory of partially ordered sets (posets), a so-called symmetric chain decomposition. Here we define this term for the n -cube using graph-theoretic language. A *symmetric chain* in Q_n is a path $(x_k, x_{k+1}, \dots, x_{n-k})$ in the n -cube where x_i is from level i for all $k \leq i \leq n - k$, and a *symmetric chain decomposition*, or SCD for short, is a partition of the vertices of Q_n into symmetric chains. For illustration, an SCD of Q_5 is shown in Figure 1 (a). We say that two SCDs are *edge-disjoint* if the corresponding paths in the graph Q_n are edge-disjoint, i.e., if there are no two consecutive vertices in one chain of the first SCD that are also contained in one chain of the second SCD. There is a well-known construction of two edge-disjoint SCDs in the n -cube for any $n \geq 1$ [SK79], which we will discuss momentarily.

Proof of Theorem 2. The proof is illustrated in Figure 1 (b). Consider two edge-disjoint SCDs \mathcal{D} and \mathcal{D}' in the $(2n + 1)$ -cube. Let \mathcal{R} and \mathcal{R}' be the chains obtained from \mathcal{D} and \mathcal{D}' , respectively, by restricting them to the middle 2ℓ levels, so chains that are longer than $2\ell - 1$ get shortened on both sides. As all chains in \mathcal{R} and \mathcal{R}' start and end at symmetric levels and the dimension $2n + 1$ is odd, all these paths have odd length (possible lengths are $1, 3, \dots, 2\ell - 1$). Therefore, by taking every second edge on every path from \mathcal{R} and \mathcal{R}' , we obtain two perfect matchings M and M' in the subgraph of the $(2n + 1)$ -cube induced by the middle 2ℓ levels. As the paths in \mathcal{R} and \mathcal{R}' are edge-disjoint, the matchings M and M' are also edge-disjoint. Therefore, the union of M and M' is the desired cycle factor. \square

This proof motivates the search for a large collection of pairwise edge-disjoint SCDs in the n -cube. We can then use any two of them to construct a cycle factor as described in the previous proof, and use this cycle factor as a starting point for building a Hamilton cycle. This two-step approach of building a Hamilton cycle via a cycle factor proved to be very successful for such problems (see e.g. [Joh09, Joh11, HRW12, Müt16, Hol17, SW18, MNW17]). Consequently, for the rest of this section we focus on edge-disjoint SCDs in the n -cube.

There is a well-known construction of an SCD for the n -cube that is best described by the following parenthesis matching approach pioneered by Greene and Kleitman [GK76]; see Figure 2. For any

¹If the four levels are not symmetric around the middle, then this subgraph of the n -cube has two partition classes of different sizes, and thus cannot have a Hamilton cycle. However, it was shown in [GM18] that in those cases the graph has a cycle that visits all vertices in the smaller partition class, and also a cyclic listing of all vertices in which only few transitions flip two instead of one bit, where ‘few’ means only as many as the difference in size between the two partition classes. Both of these notions are natural generalizations of a Hamilton cycle.




FIGURE 3. The edge-disjoint SCDs \mathcal{D}_0 (dashed vertical paths) and \mathcal{D}_1 (solid paths; chains of the same length are drawn with the same color) in the 6-cube. The bitstrings are drawn with white squares representing 0s and black squares representing 1s.

TABLE 1. Known pairwise edge-disjoint SCDs in the n -cube for $n = 1, 2, \dots, 11$. The definitions of $\mathcal{X}_5, \mathcal{Y}_5, \mathcal{Z}_5$ and $\mathcal{X}_7, \mathcal{Y}_7$ are given in Section 3.3.

n	1	2	3	4	5	6	7	8	9	10	11
$\lfloor n/2 \rfloor + 1$	1	2	2	3	3	4	4	5	5	6	6
SCDs	\mathcal{D}_0	$\mathcal{D}_0, \overline{\mathcal{D}_0}$	$\mathcal{D}_0, \overline{\mathcal{D}_0}$	$\mathcal{D}_0, \overline{\mathcal{D}_0}, \mathcal{D}_1$	$\mathcal{X}_5, \mathcal{Y}_5, \mathcal{Z}_5$	$\mathcal{D}_0, \overline{\mathcal{D}_0}, \mathcal{D}_1, \overline{\mathcal{D}_1}$	$\mathcal{X}_7, \overline{\mathcal{X}_7}, \mathcal{Y}_7, \overline{\mathcal{Y}_7}$	$\mathcal{D}_0, \overline{\mathcal{D}_0}, \mathcal{D}_1, \overline{\mathcal{D}_1}$	$\mathcal{D}_0, \overline{\mathcal{D}_0}$	$\mathcal{D}_0, \overline{\mathcal{D}_0}, \mathcal{D}_1, \overline{\mathcal{D}_1}$	$\mathcal{D}_0, \overline{\mathcal{D}_0}$

Theorem 5 shows in particular that from k edge-disjoint SCDs in a hypercube of fixed dimension n , we obtain k edge-disjoint SCDs for infinitely many larger dimensions $2n, 3n, 4n, \dots$

We conjecture that the n -cube has $\lfloor n/2 \rfloor + 1$ pairwise edge-disjoint SCDs, but so far we only know that this holds for $n \leq 7$. Clearly, finding this many edge-disjoint SCDs would be best possible, as they use up all middle edges of the cube. Maximum sets of pairwise edge-disjoint SCDs in the n -cube we found for $n = 1, 2, \dots, 11$ are shown in Table 1, together with the aforementioned upper bound.

1.2. Related work. Apart from building Gray codes, symmetric chain decompositions have many other interesting applications, e.g., to construct rotation-symmetric Venn diagrams for n sets when n is a prime number [GKS04, RSW06], and to solve the Littlewood-Offord problem on sums of vectors [Bol86]. It would be very interesting to investigate how the new SCDs of the n -cube presented in this paper can be exploited for those and other applications.

A notion that is closely related to edge-disjoint SCDs is that of *orthogonal* chain decompositions, which were first considered by Shearer and Kleitman [SK79]. Two chain decompositions are called *orthogonal* if every pair of chains has at most one vertex in common, where one also allows chains that are not symmetric around the middle or chains that skip some levels. Shearer and Kleitman showed in their paper that \mathcal{D}_0 and $\overline{\mathcal{D}_0}$ are almost orthogonal (only the longest chains have two elements in common), and they conjectured that the n -cube has $\lfloor n/2 \rfloor + 1$ pairwise orthogonal chain decompositions where each decomposition consists of $\binom{n}{\lfloor n/2 \rfloor}$ many chains. Spink recently made some progress towards this conjecture, by showing that the n -cube has three orthogonal chain decompositions [Spi17].




FIGURE 4. The correspondence between bitstrings, lattice paths (left) and rooted trees (right).

Pikhurko showed via a parenthesis matching argument that all edges of the n -cube can be decomposed into symmetric chains [Pik99]. However, it is not clear whether these chains contain a subset that forms an SCD.

An interesting construction relating Hamilton cycles and SCDs in the n -cube was presented by Streib and Trotter in [ST14]. They inductively construct a Hamilton cycle in the n -cube for any $n \geq 2$ that can be partitioned into symmetric chains forming an SCD. This Hamilton cycle has the minimal number of ‘peaks’ where the differences in the Hamming weight change sign (and thus also the minimal number of corresponding ‘valleys’).

1.3. Outline of this paper. In Section 2 we introduce several definitions that will be used throughout this paper. In Section 3 we present the proofs of Theorems 3–5, and we describe the construction of the SCD \mathcal{D}_1 and of the SCDs in Q_5 and Q_7 referred to in Table 1. As it is somewhat technical, we defer the proof of Theorem 1 to Section 4. In Section 5 we present the results of computer experiments on the structure of the cycle factors through the middle 2ℓ levels of the $(2n + 1)$ -cube constructed as in the proof of Theorem 2. We conclude in Section 6 with some open problems.

2. PRELIMINARIES

We begin by introducing some terminology that is used throughout the following sections.

2.1. Bitstrings, lattice paths, and rooted trees. We use $L_{n,k}$ to denote the set of all bitstrings of length n with Hamming weight k , so this is exactly the k th level of Q_n . For any bitstring x , we write \bar{x} for its complement and $\text{rev}(x)$ for the reversed bitstring. For a set of bitstrings X and a bitstring x , we write $X \circ x$ for the set obtained by concatenating each bitstring from X with x .

We often interpret a bitstring x as a path in the integer lattice \mathbb{Z}^2 starting at the origin $(0, 0)$, where every 1-bit is interpreted as an \nearrow -step that changes the current coordinate by $(+1, +1)$ and every 0-bit is interpreted as a \searrow -step that changes the current coordinate by $(+1, -1)$; see Figure 4. Let $D_{n,k} \subseteq L_{n,k}$ denote the bitstrings that have the property that in every prefix, the number of 1s is at least as large as the number of 0s. We partition the set $D_{n,k}$ further into $D_{n,k}^{>0}$ and $D_{n,k}^{=0}$, according to whether this inequality is strict for all non-empty prefixes, or whether it holds with equality for at least one non-empty prefix, respectively. The empty bitstring $()$ therefore belongs to $D_{0,0}^{>0}$ and not to $D_{0,0}^{=0}$. We also define $D^{>0} := \bigcup_{n \geq k \geq 0} D_{n,k}^{>0}$, $D^{=0} := \bigcup_{n \geq k \geq 0} D_{n,k}^{=0}$ and $D := \bigcup_{k \geq 0} D_{2k,k}$. In terms of lattice paths, D corresponds to so-called *Dyck paths* that never move below the abscissa $y = 0$ and end at the abscissa. Similarly, $D^{>0}$ are paths that always stay strictly above the abscissa except at the origin, and $D^{=0}$ are paths that touch the abscissa at least once more. Any bitstring $x \in D^{=0}$ can be written uniquely as $x = (1, u, 0, v)$ with $u \in D$. We refer to this as the *canonical decomposition* of x . The set $D_{n,k}^-$ is defined similarly as $D_{n,k}$, but we require that in exactly one prefix, the number of 1s is strictly smaller than the number of 0s. That is, the lattice paths corresponding to $D_{n,k}^-$ move below the abscissa exactly once.

\nearrow -step is part of x ; otherwise x is left unmatched. It is straightforward to verify that these two partial mappings are inverse to each other, so they indeed define a matching between levels k and $k + 1$ of Q_n , which we denote by $M_{n,k}^i$.

The following properties of lexical matchings are immediate consequences of these definitions.

Lemma 6. *Let $0 \leq k \leq n - 1$ and $l := \max\{k, n - k - 1\}$. The lexical matchings defined before have the following properties.*

- (i) *For every $0 \leq i \leq l$, the matching $M_{n,k}^i$ saturates all vertices in the smaller of the two levels k and $k + 1$.*
- (ii) *The matchings $M_{n,k}^i$, $i = 0, 1, \dots, l$, form a partition of all edges of the subgraph of Q_n between levels k and $k + 1$.*
- (iii) *For every $0 \leq i \leq l$ we have $\overline{M_{n,k}^i} = M_{n,n-k-1}^{l-i}$ and $\text{rev}(M_{n,k}^i) = M_{n,k}^{l-i}$. Consequently, we have $\overline{\text{rev}(M_{n,k}^i)} = M_{n,n-k-1}^i$.*

3. PAIRWISE EDGE-DISJOINT SCDs

We proceed to prove Theorems 3–5.

3.1. Proof of Theorem 3.

Proof of Theorem 3. The structure of this proof is as follows: In the first part we define an SCD \mathcal{D}_1 in Q_n for even $n \geq 2$ via lattice paths, and we also give an equivalent definition of \mathcal{D}_0 in this framework. From this definition it will follow immediately that all chains in \mathcal{D}_1 are symmetric. In the second part of the proof, we show that all chains in \mathcal{D}_1 are disjoint and that they form a partition of all vertices of Q_n , proving that \mathcal{D}_1 is indeed an SCD. Moreover, we argue that \mathcal{D}_0 , $\overline{\mathcal{D}_0}$, \mathcal{D}_1 and $\overline{\mathcal{D}_1}$ are pairwise edge-disjoint. This is achieved by showing that the definitions of \mathcal{D}_0 and \mathcal{D}_1 via lattice paths are equivalent to taking the union of certain lexical matchings introduced in the previous section, and using the properties of these matchings stated in Lemma 6.

We now define \mathcal{D}_1 by interpreting the bitstrings of the vertices of Q_n as lattice paths. Before defining \mathcal{D}_1 , let us first give an equivalent definition of \mathcal{D}_0 for even $n \geq 2$ in this setting. For this, recall the definition of \mathcal{D}_0 given in the introduction and shown in the examples in Figure 1 (a) and Figure 3. To define the chain decomposition, we consider a vertex $x \in L_{n,n/2}$ in the middle level $n/2$, and we define the sequence of vertices reached from x when moving up the corresponding chain, and the sequence of vertices reached when moving down the chain. For this we consider the lattice path corresponding to the bitstring x . This lattice path ends at the coordinate $(n, 0)$ as the number of 0s equals the number of 1s. We now label a subsequence of \searrow -steps of this lattice path with integers $j = 1, 2, \dots$ according to the following procedure, see the top part of Figure 6 for an illustration:

- (a0) We place a marker at the rightmost highest point of x and set $j := 1$.
- (b0) If the marker is at height $h \geq 1$, we label the \searrow -step starting at the marker with j , and we move the marker to the starting point of the rightmost \searrow -step starting at height $h - 1$. We set $j := j + 1$ and repeat.
- (c0) If the marker is at height $h = 0$, we stop.

Flipping the \searrow -steps of x marked with $1, 2, \dots$ in this order yields the sequence of vertices reached from x when moving up the chain containing x . An analogous labeling procedure obtained by interchanging left and right, \searrow -steps and \nearrow -steps, and starting with ending points yields the sequence of vertices reached from x when moving down the chain containing x . We let $C_0(x)$ denote the corresponding chain. Observe that $C_0(x)$ is a symmetric chain, as the height of the marker decreases by 1 in each step, so the number of edges we move up from x equals the number of edges we move down

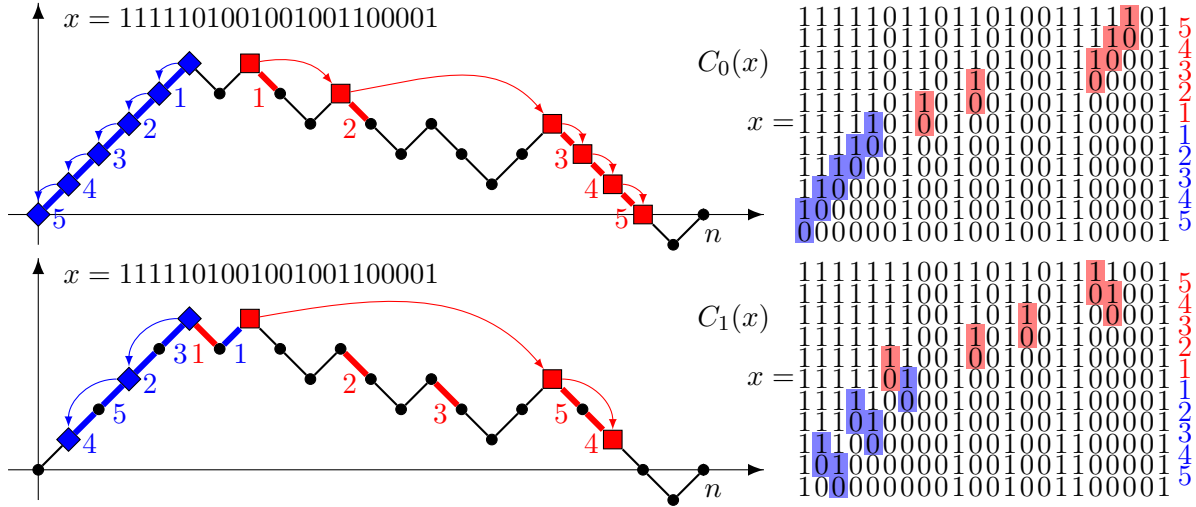


FIGURE 6. The labeling procedures that define the symmetric chains $C_0(x)$ (top) and $C_1(x)$ (bottom). The markers that define the upward and downward steps of the chains are drawn as a square and a diamond, respectively. The resulting chain $C_0(x)$ is the same as the one shown in Figure 2.

from x . It is easy to verify that the SCD \mathcal{D}_0 defined before via the parenthesis matching approach satisfies

$$\mathcal{D}_0 = \bigcup_{x \in L_{n,n/2}} C_0(x) .$$

To define \mathcal{D}_1 we modify the labeling procedure as follows, see the bottom part of Figure 6:

- (a1) We place a marker at the rightmost highest point of x and set $j := 1$. If there is a \searrow -step to the left of the marker starting at the same height, we label the nearest such step with 1 and set $j := 2$.
- (b1) If the marker is at height $h \geq 2$, we label the rightmost \searrow -step starting at height $h - 1$ with j . We consider all \searrow -steps starting at height $h - 2$ to the right of the labeled step and the \searrow -step starting at the marker, we label the second step from the right from this set with $j + 1$, and we move the marker to the starting point of the rightmost \searrow -step starting at height $h - 2$. We set $j := j + 2$ and repeat.
- (c1) If the marker is at height $h = 1$ or $h = 0$, we stop.

We let $C_1(x)$ denote the corresponding chain obtained by flipping bits according to this labeling rule and the corresponding symmetric rule obtained by interchanging left and right, \searrow -steps and \swarrow -steps, and starting with ending points. Observe that $C_1(x)$ is a symmetric chain, as the height of the marker decreases by 2 in each iteration (and we label two steps in each iteration) and the conditional marking with $j = 1$ in step (a1) occurs if and only if the highest point of x is unique, so the number of edges we move up from x equals the number of edges we move down from x . At this point it is not at all clear yet that the chains $C_1(x)$, $x \in L_{n,n/2}$, are disjoint, nor that they form a partition of the vertices of Q_n . This is what we will argue about next, which will prove that

$$\mathcal{D}_1 := \bigcup_{x \in L_{n,n/2}} C_1(x) \tag{1}$$

is actually an SCD of Q_n .




FIGURE 7. Unions of lexical matchings $M^i = M_{n,k}^i$ yielding edge-disjoint chain decompositions in Q_n for $n = 6$. The resulting chains in \mathcal{D}_0 and \mathcal{D}_1 in Q_6 are shown in Figure 3.

We do this by showing that \mathcal{D}_0 and \mathcal{D}_1 as defined before equal the union of certain lexical matchings defined in Section 2.2. By property (i) from Lemma 6, for any sequence $\mathbf{i} := (i_0, i_1, \dots, i_{n-1})$, $i_k \in \{0, 1, \dots, \max\{k, n - k - 1\}\}$, the union

$$\mathcal{D}_{\mathbf{i}} := \bigcup_{k=0}^{n-1} M_{n,k}^{i_k} \quad (2)$$

is a chain decomposition of Q_n . The resulting chains are not necessarily symmetric, though.

From the definitions in Section 2.2 it also follows that \mathcal{D}_0 equals the union of the 0-lexical matchings, and that for even $n \geq 2$, \mathcal{D}_1 as defined in (1) equals the union of the 1-lexical matchings; formally we have

$$\mathcal{D}_0 = \mathcal{D}_{(0,0,\dots,0)} = \bigcup_{k=0}^{n-1} M_{n,k}^0, \quad \mathcal{D}_1 = \mathcal{D}_{(1,1,\dots,1)} = \bigcup_{k=0}^{n-1} M_{n,k}^1.$$

Consequently, \mathcal{D}_0 and \mathcal{D}_1 are indeed chain decompositions, and by the definitions of \mathcal{D}_0 and \mathcal{D}_1 given at the beginning of the proof all chains in these decompositions are in fact symmetric. The fact that \mathcal{D}_0 , $\overline{\mathcal{D}_0}$, \mathcal{D}_1 and $\overline{\mathcal{D}_1}$ are pairwise edge-disjoint can be seen by applying property (iii) from Lemma 6 and by observing that by property (ii), $\mathcal{D}_{\mathbf{i}}$ and $\mathcal{D}_{\mathbf{j}}$ as defined in (2) are edge-disjoint if and only if the sequences \mathbf{i} and \mathbf{j} differ in every corresponding entry, see Figure 7.

This completes the proof. \square

Clearly, $\mathcal{D}_{(0,0,\dots,0)}$ as defined in (2) equals \mathcal{D}_0 for *every* $n \geq 1$, so the union of all 0-lexical matchings forms an SCD in any dimension. In contrast to that, the union of all 1-lexical matchings $\mathcal{D}_{(1,1,\dots,1)}$ only forms an SCD for *even* $n \geq 2$. Computer experiments show that for $n \in \{8, 10\}$ there is no union of lexical matchings $\mathcal{D}_{\mathbf{i}}$ for any sequence \mathbf{i} that forms an SCD that is edge-disjoint from \mathcal{D}_0 , $\overline{\mathcal{D}_0}$, \mathcal{D}_1 and $\overline{\mathcal{D}_1}$; recall Table 1.

Furthermore, taking unions of the so-called modular matchings introduced by Duffus, Kierstead, and Snevily [DKS94] does not yield SCDs in Q_n for $n = 5$ and $n = 7$, and only two edge-disjoint SCDs for $n = 6$.




FIGURE 8. Illustration of the proof of Theorem 5. Construction of two edge-disjoint SCDs in Q_5 from two edge-disjoint SCDs in Q_2 and two edge-disjoint SCDs in Q_3 . The chains of the SCD \mathcal{C}_1 of Q_5 as constructed in the proof are highlighted in gray.

3.2. Proof of Theorem 5.

Proof of Theorem 5. For the reader's convenience, this proof is illustrated in Figure 8. Assume that $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_k$ and $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_k$ denote k pairwise edge-disjoint SCDs of Q_a and Q_b , respectively. We will think of Q_{a+b} as the Cartesian product $Q_a \times Q_b$ of Q_a and Q_b . We show how to construct for every $i \in [k]$ an SCD \mathcal{C}_i of $Q_{a+b} = Q_a \times Q_b$ which uses only edges of the form $((u, v), (u', v'))$ where (u, u') is an edge from \mathcal{A}_i or (v, v') is an edge from \mathcal{B}_i . From this it follows that the SCDs $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$ are pairwise edge-disjoint.

The SCD \mathcal{C}_i of Q_{a+b} is defined as follows: The Cartesian products $A \times B$ of chains $A \in \mathcal{A}_i$ and $B \in \mathcal{B}_i$ partition the vertices of Q_{a+b} into two-dimensional grids. \mathcal{C}_i is obtained by partitioning each of those grids into symmetric chains in the natural way, see Figure 8 (cf. [dBvETK51]): Specifically, let $A = (x_1, \dots, x_\alpha)$ and $B = (y_1, \dots, y_\beta)$ be the vertices in the chains A and B from bottom to top. As A and B are symmetric, we know that $|x_1| + |x_\alpha| = a$ and $|y_1| + |y_\beta| = b$, where $|x|$ denotes the Hamming weight of the bitstring x . This implies that $|(x_1, y_1)| + |(x_\alpha, y_\beta)| = |x_1| + |y_1| + |x_\alpha| + |y_\beta| = a + b$, i.e., the bottom and top vertex of the grid $A \times B$ are on symmetric levels in Q_{a+b} . We may therefore decompose $A \times B$ into disjoint symmetric chains \mathcal{C}_j , $j = 1, 2, \dots, \min\{\alpha, \beta\}$, by setting

$$\mathcal{C}_j := ((x_1, y_j), (x_2, y_j), \dots, (x_{\alpha-j+1}, y_j), (x_{\alpha-j+1}, y_{j+1}), \dots, (x_{\alpha-j+1}, y_\beta)) . \quad (3)$$

□

Note that in the proof of Theorem 5 we have some degrees of freedom in partitioning the two-dimensional grids $A \times B$ into symmetric chains. If we perform this construction using $Q_{n+1} = Q_n \times Q_1$ for $n = 1, 2, \dots$ and always partition the grids according to (3), then the resulting SCD equals \mathcal{D}_0 . If instead we partition always according to

$$\mathcal{C}_j := ((x_j, y_1), (x_j, y_2), \dots, (x_j, y_{\beta-j+1}), (x_{j+1}, y_{\beta-j+1}), \dots, (x_\alpha, y_{\beta-j+1})) , \quad (4)$$

then we get the SCD $\overline{\mathcal{D}}_0$. The difference between (3) and (4) is whether in building \mathcal{C}_j we first move along the first coordinates, or first along the last coordinates.

3.3. Proof of Theorem 4. We begin by constructing the SCDs in Q_5 and Q_7 mentioned in Table 1.




FIGURE 9. Illustration of the three edge-disjoint SCDs in N_5 (a) and four edge-disjoint SCDs in N_7 (b). The names of the SCDs correspond to the ones used in Table 1. In N_7 , two pairs of SCDs are complementary.

Lemma 7. Q_5 contains three pairwise edge-disjoint SCDs, Q_7 contains four pairwise edge-disjoint SCDs, and this is best possible.

Proof. We consider the graph Q_n with the two vertices in the outermost levels 0 and n removed, and we identify all bitstrings that differ only by rotation into so-called *necklaces*. The resulting graph N_n is a multigraph version of the cover graph of the *necklace poset*. Specifically, the multiplicity of the edges in N_n corresponds to the number of ways a bit from a necklace can be flipped to reach the corresponding adjacent necklace. E.g., in N_5 the necklace $x := 10000$ has two edges leading to $y := 11000$, as we can flip the second or the fifth bit in x to reach y . This way, a necklace on level k has $n - k$ edges going up, and k edges going down, like the vertices in Q_n . The multigraphs N_5 and N_7 are shown in Figure 9. If n is prime, then every SCD in N_n corresponds to an SCD in Q_n , by turning each chain from N_n into n chains in Q_n obtained by rotating a representative of each necklace in all possible ways. Moreover, one of the chains of length $n - 2$ needs to be extended by the all-zero and all-one bitstring to a chain of length n in Q_n . Observe that in this way, k edge-disjoint SCDs in N_n give rise to k edge-disjoint SCDs in Q_n .

As $n = 5$ and $n = 7$ are prime, we thus obtain three edge-disjoint SCDs in Q_5 from the SCDs in N_5 shown in Figure 9 (a), and four edge-disjoint SCDs in Q_7 from the SCDs in N_7 shown in Figure 9 (b). These SCDs use up all middle edges, so this is best possible. \square

Proof of Theorem 4. For $n = 7$ the statement follows from Lemma 7. For $n \geq 13$ we apply Theorem 5 to Q_{n-7} and Q_7 , using the four edge-disjoint SCDs in Q_{n-7} given by Theorem 3 (note that $n - 7 \geq 6$), and the four edge-disjoint SCDs in Q_7 given by Lemma 7. \square

4. THE MIDDLE FOUR LEVELS PROBLEM

In this section we prove Theorem 1. The proof proceeds similarly as the proof of the middle two levels problem [Müt16, GMN17]. First, we construct a cycle factor \mathcal{C}_{2n+1} in the middle four levels of Q_{2n+1} , and we then modify the cycles in the factor locally to join them to a Hamilton cycle. Specifically, to join two cycles C and C' from our cycle factor, we consider a suitable 6-cycle and take the symmetric difference between their edge sets, such that the result is a single cycle on the same vertex set as C and C' , see Figure 15. This process is iterated until all cycles are joined to a single Hamilton cycle. This technique reduces the problem of proving that the middle four levels of Q_{2n+1} have a Hamilton cycle to the problem of proving that a suitably defined auxiliary graph

\mathcal{H}_{n+1} has a spanning tree, which is much easier. This section is organized as follows: We first define the cycle factor \mathcal{C}_{2n+1} and analyze its structure. We then introduce the 6-cycles for the joining operations, and finally show that they can be used to join the cycles of the factor to a Hamilton cycle in the desired fashion.

4.1. Construction of the cycle factor \mathcal{C}_{2n+1} . As we are not able to analyze the cycle factor in the middle four levels of Q_{2n+1} arising from the proof of Theorem 2 (see Table 2 in Section 5), we start with a different construction. To construct the cycle factor \mathcal{C}_{2n+1} in the middle four levels of Q_{2n+1} , we use the i -lexical matchings defined in Section 2.2. Specifically, we take the union of all n -lexical and $(n+1)$ -lexical matching edges between the upper two levels $n+1$ and $n+2$ and between the lower two levels $n-1$ and n , plus certain carefully chosen edges E from the $(n-2)$ -lexical, $(n-1)$ -lexical and n -lexical matching between the middle levels n and $n+1$. Formally, we set

$$\mathcal{C}_{2n+1} := (M_{2n+1,n+1}^n \cup M_{2n+1,n+1}^{n+1}) \cup (M_{2n+1,n-1}^n \cup M_{2n+1,n-1}^{n+1}) \cup E, \quad (5)$$

where the set of edges E is defined in (8) below. By this definition and by Lemma 6 (i) and (ii), all vertices in the outer levels $n-1$ and $n+2$ have degree two in the subgraph \mathcal{C}_{2n+1} , and we will choose E so that all vertices in the inner levels n and $n+1$ have degree two as well.

To define the set E , we consider the union of the matchings between the upper two levels

$$\mathcal{P} := M_{2n+1,n+1}^n \cup M_{2n+1,n+1}^{n+1}. \quad (6)$$

The following result was shown in [GMN17, Proposition 2].

Lemma 8. *For any $n \geq 1$, the set \mathcal{P} defined in (6) is a set of paths (without any cycles), and the sets of start and end vertices of these paths are $D_{2n+1,n+1}^{\bar{0}}$ and $D_{2n+1,n+1}^-$, respectively. Furthermore, for any path with first vertex $x \in D_{2n+1,n+1}^{\bar{0}}$ and last vertex $y \in D_{2n+1,n+1}^-$, if $x = (1, u, 0, v)$ is the canonical decomposition of x , then $y = (u, 0, 1, v)$.*

The first vertices of the paths \mathcal{P} , denoted by $F(\mathcal{P})$, are the vertices from level $n+1$ covered by the matching $M_{2n+1,n+1}^n$ and not by $M_{2n+1,n+1}^{n+1}$. Similarly, the last vertices of the paths \mathcal{P} , denoted by $L(\mathcal{P})$, are covered by the latter matching but not by the former. We let $I(\mathcal{P})$ denote the set of vertices in level $n+1$ covered by neither of the two matchings. We refer to those vertices as *isolated*.

We let $f := \overline{\text{rev}}$ denote the automorphism of Q_{2n+1} that flips all bits and reverses them, i.e., $f(x_1, x_2, \dots, x_{2n+1}) = (\bar{x}_{2n+1}, \dots, \bar{x}_2, \bar{x}_1)$. Using the abbreviation (6) and Lemma 6 (iii), we may rewrite the definition (5) equivalently as

$$\mathcal{C}_{2n+1} := \mathcal{P} \cup f(\mathcal{P}) \cup E. \quad (7)$$

For any set X of bitstrings and bitstrings a, b we let ${}_aX_b$ denote the subset of the bitstrings from X that have the prefix a and the suffix b . So \mathcal{P}_0 and \mathcal{P}_1 are the sets of paths from \mathcal{P} with fixed last bit equal to 0 or 1, respectively. Note that $\mathcal{P} = \mathcal{P}_0 \cup \mathcal{P}_1$, as neither the n -lexical matching nor the $(n+1)$ -lexical matching between levels $n+1$ and $n+2$ uses any edges along which the last bit is flipped. We start with the following observations, illustrated in Figure 10.

Lemma 9. *For any $n \geq 1$, the sets of first, last and isolated vertices of the paths \mathcal{P} defined in (6) are given by*

$$\begin{aligned} (i) \quad & F(\mathcal{P}) = D_{2n+1,n+1}^{\bar{0}}, \quad L(\mathcal{P}) = D_{2n+1,n+1}^-, \quad I(\mathcal{P}) = D_{2n+1,n+1}^{\bar{0}}, \\ (ii) \quad & F(\mathcal{P}_0) = D_{2n,n+1}^{\bar{0}} \circ 0, \quad F(\mathcal{P}_1) = D_{2n,n}^{\bar{0}} \circ 1. \end{aligned}$$

Proof. The first two statements in part (i) are given by Lemma 8. The statement $I(\mathcal{P}) = D_{2n+1,n+1}^{\bar{0}}$ follows by considering for which $x \in L_{2n+1,n+1}$ the lattice path x^\uparrow (recall the definition from Section 2.2) has its n th and $(n+1)$ th \searrow -step (in the counting from top to bottom and from right to




FIGURE 10. Lattice path interpretation of vertex sets in level $n+1$ of Q_{2n+1} involved in the construction. Heavier (black) steps denote coordinates with fixed bits 0 or 1. Fat steps denote coordinates flipped along the edges of E : solid (red) by E^n , dashed (green) by E^{n-1} , dash-dotted (blue) by E^{n-2} .

left in each row, starting from 0) in the last two \searrow -steps at coordinates $2n+2$ and $2n+3$ that were added to x . One can easily observe that this happens if and only if $x \in D_{2n+1, n+1}^{>0}$. Part (ii) follows immediately from the definitions and from part (i) by considering the lattice paths before the fixed last bit. \square

We now define the set E of edges for the cycle factor \mathcal{C}_{2n+1} between levels n and $n+1$ of Q_{2n+1} so that each vertex of $F(\mathcal{P})$ and $L(\mathcal{P})$ will be incident with exactly one edge of E , each vertex of $I(\mathcal{P})$ will be incident with exactly two edges of E , and all other vertices in level $n+1$ are not incident with any edges from E , see Figure 11. Effectively, adding the edges from E makes all degrees in the subgraph \mathcal{C}_{2n+1} in level $n+1$ equal to two. We then show that adding the edges from E also makes all degrees in level n equal to two, so that \mathcal{C}_{2n+1} is indeed a cycle factor.

We define $L'(\mathcal{P}) \subseteq L(\mathcal{P})$ as the set of all vertices $x \in L(\mathcal{P}) = D_{2n+1, n+1}^-$ such that $x = (u, 0, 1, v)$ with some $u \in D$ and $v \in D^{=0}$ (otherwise $v \in D^{>0}$). Similarly, $F'(\mathcal{P}) \subseteq F(\mathcal{P})$ is the set of vertices $x \in F(\mathcal{P}) = D_{2n+1, n+1}^{=0}$ such that $x = (1, u, 0, v, 1, w)$ with $u, v, w \in D$ and $v \neq ()$. For a set of edges M and a set of vertices X , we let $M[X]$ denote the set of edges from M incident with X . We then define

$$E := E^n \cup E^{n-1} \cup E^{n-2}, \quad (8a)$$

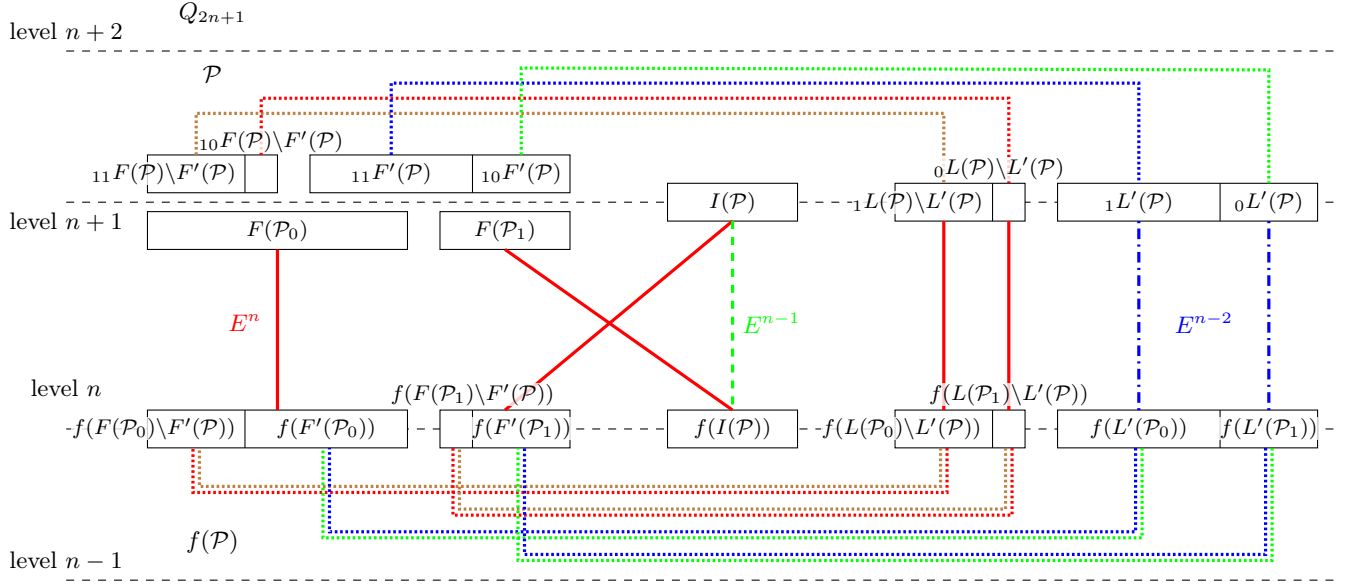


FIGURE 11. Vertex sets involved in constructing the cycle factor \mathcal{C}_{2n+1} in the middle four levels of Q_{2n+1} . The relevant vertex sets are drawn by rectangular boxes, where the width of the boxes represents the size of the set, drawn to scale for large values of n . The vertex sets and their images under the automorphism $f = \bar{r}\bar{e}\bar{v}$ are aligned above each other. The paths $\mathcal{P} = M_{2n+1,n+1}^n \cup M_{2n+1,n+1}^{n+1}$ between levels $n+1$ and $n+2$ and $f(\mathcal{P})$ between levels $n-1$ and n connecting various pairs of end vertices from the vertex sets are indicated by dotted lines. The intermediate vertices of these paths are not shown in the figure. The additional edges between levels n and $n+1$ from the sets $E^i \subseteq M_{2n+1,n}^i$, $i = n, n-1, n-2$, connecting those paths to a cycle factor are drawn by solid (red), dashed (green) and dash-dotted (blue) lines, respectively.

where

$$E^n := M_{2n+1,n}^n [F(\mathcal{P}) \cup (L(\mathcal{P}) \setminus L'(\mathcal{P})) \cup I(\mathcal{P})] , \quad (8b)$$

$$E^{n-1} := M_{2n+1,n}^{n-1} [I(\mathcal{P})] , \quad (8c)$$

$$E^{n-2} := M_{2n+1,n}^{n-2} [L'(\mathcal{P})] . \quad (8d)$$

To prove that \mathcal{C}_{2n+1} as defined in (7) is indeed a cycle factor, we now consider the sets of vertices in level n that are covered by the edges from E , see Figure 11.

Lemma 10. *For any $n \geq 1$, the edges from E defined in (8) match the following sets of vertices in levels $n+1$ and n of Q_{2n+1} .*

- (i) *The sets $F(\mathcal{P}_0)$ and $f(F(\mathcal{P}_0))$, $F(\mathcal{P}_1)$ and $f(I(\mathcal{P}))$, $I(\mathcal{P})$ and $f(F(\mathcal{P}_1))$ are matched by edges from E^n .*
- (ii) *The sets $I(\mathcal{P})$ and $f(I(\mathcal{P}))$ are matched by edges from E^{n-1} .*
- (iii) *The sets $L(\mathcal{P}) \setminus L'(\mathcal{P})$ and $f(L(\mathcal{P}) \setminus L'(\mathcal{P}))$ are matched by edges from E^n , and the sets $L'(\mathcal{P})$ and $f(L'(\mathcal{P}))$ are matched by edges from E^{n-2} .*

Proof. It suffices to show that in each of the cases, the edges from E join a vertex from one of the sets X in level $n+1$ to a vertex from the corresponding set Y in level n (i.e., we show that these edges form an injection from X to Y). The fact that they form a surjection follows by applying




FIGURE 12. Heavy and light tree rotations (left), and the pull operation (right).

the same argument to $f(\mathcal{C}_{2n+1})$, using that f is an involution and that $f(M_{2n+1,n}^i) = M_{2n+1,n}^i$ for $i = n, n-1, n-2$ by Lemma 6 (iii).

To prove the first two statements in (i), consider a vertex $x \in F(\mathcal{P}) = D_{2n+1,n+1}^=0$ in level $n+1$ and the corresponding lattice path (recall Lemma 9 (i)). The edge of E_n flips the last \nearrow -step starting at the abscissa, i.e., this edge joins $x = (u, 1, v)$ where $u, v \in D$ are uniquely determined with the vertex $y = (u, 0, v) = f(f(v), 1, f(u))$. If $x \in F(\mathcal{P}_0)$, i.e., $v \neq ()$, then we have $y \in f(F(\mathcal{P}_0))$ as $(f(v), 1, f(u)) \in D_{2n,n+1}^=0 \circ 0$ (recall Lemma 9 (ii)). If $x \in F(\mathcal{P}_1)$, i.e., $v = ()$, then we have $y \in f(I(\mathcal{P}))$ as $(1, f(u)) \in D_{2n+1,n+1}^{>0}$ (recall Lemma 9 (i)). The remaining claims can be shown by similar calculations with the help of Lemma 9. We omit the details. \square

Lemma 10 allows us to conclude that \mathcal{C}_{2n+1} as defined in (7) with the edge set E defined in (8) is indeed a cycle factor in the subgraph of Q_{2n+1} induced by the middle four levels, as every vertex in the four levels is covered by exactly two edges.

4.2. Structure of the cycle factor \mathcal{C}_{2n+1} . We now analyze the structure of the cycle factor \mathcal{C}_{2n+1} defined in (7). Observe from Figure 11 that on each cycle of \mathcal{C}_{2n+1} all paths from \mathcal{P} are visited in the same orientation from the first vertices $F(\mathcal{P})$ to the corresponding last vertices from $L(\mathcal{P})$. The following lemma shows for a given path from \mathcal{P} on a cycle C of \mathcal{C}_{2n+1} which path from \mathcal{P} is encountered next on the cycle C . To state the lemma we introduce a bit of notation.

It is convenient here to identify the paths from \mathcal{P} by their first vertices, so by Lemma 8 this is the set $D_{2n+1,n+1}^=0$. By appending an additional 0-bit to the bitstrings $D_{2n+1,n+1}^=0$, we obtain Dyck paths of length $2n+2$ with exactly $n+1$ upsteps and $n+1$ downsteps that touch the abscissa at least three times (in the origin $(0,0)$, at $(2n+2,0)$ and at some intermediate point). It turns out that the structure of the cycle factor \mathcal{C}_{2n+1} can be described most conveniently by interpreting the Dyck paths $D_{2n+1,n+1}^=0 \circ 0$ as rooted trees as described in Section 2.1 and illustrated in Figure 4. We introduce the abbreviation $\mathcal{T}_{n+1} := D_{2n+1,n+1}^=0 \circ 0$ for these trees. Note that they have exactly $n+1$ edges and the root has degree at least two.

The following definitions are illustrated in Figure 12. We say that a tree from \mathcal{T}_{n+1} is *left-light* if the leftmost child of the root is a leaf; otherwise, it is *left-heavy*. Analogously, we define *right-light* or *right-heavy* trees by considering the rightmost child of the root. A *left-rotation* of a tree moves the root to the leftmost child of the root. In terms of bitstrings, this operation changes $(1, u, 0, v)$ with $u, v \in D$ to $(u, 1, v, 0)$. A *right-rotation* is the inverse operation of a left-rotation. Given a left-heavy tree $x = (1, u, 0, s, 1, v_3, 0)$, $u, v_3 \in D$, where $s = ()$ or $s = (v_1, 1, v_2, 0)$ with $v_1, v_2 \in D$, a *heavy rotation* is a left-rotation of the tree plus a right-rotation of the subtree s , i.e., the resulting tree y is given by $y = (u, 1, 1, v_3, 0, 0)$ if $s = ()$ and $y = (u, 1, 1, v_1, 0, v_2, 1, v_3, 0, 0)$ otherwise. Note that the resulting tree y is right-heavy. Given a left-light tree $x = (1, 0, s, 1, v_3, 0)$, $v_3 \in D$, where $s = ()$ or $s = (v_1, 1, v_2, 0)$ with $v_1, v_2 \in D$, a *light rotation* is a right-rotation of the tree plus a right-rotation of the (possibly empty) subtree s plus detaching the pending edge that leads to the leftmost child of the root of x and reattaching it as a rightmost child of the new root, i.e., the resulting tree y is given by $y = (1, 0, v_3, 1, 0)$ if $s = ()$ and $y = (1, 1, v_1, 0, v_2, 0, v_3, 1, 0)$ otherwise. Note that the resulting tree y is right-light. To any given tree from \mathcal{T}_{n+1} , we can either apply a heavy or a light rotation, depending on whether the tree is left-heavy or left-light, respectively. We refer to this mapping on \mathcal{T}_{n+1} as ρ . Analogously, ρ^{-1} applies either an inverse heavy or an inverse light rotation depending on whether the tree is right-heavy or right-light, respectively.

The following lemma asserts that the sequence of rooted trees corresponding to first vertices of paths from \mathcal{P} that are encountered when following a cycle from our factor \mathcal{C}_{2n+1} corresponds to repeatedly applying ρ , i.e., either applying a heavy rotation or a light rotation. In other words, the cycles of \mathcal{C}_{2n+1} are in bijection with equivalence classes of rooted trees from \mathcal{T}_{n+1} when iterating the mapping ρ .

Lemma 11. *For any $n \geq 1$, any cycle C of the cycle factor \mathcal{C}_{2n+1} defined in (7), and any vertex $x \in F(\mathcal{P}) = D_{2n+1, n+1}^{-0}$ let y be the next vertex from $F(\mathcal{P})$ on the cycle C encountered after x . Let $x = (1, u, 0, v)$, $u \in D$, be the canonical decomposition of x . If $u \neq ()$, i.e., $x \in {}_{11}F(\mathcal{P})$, then we have*

$$y = \begin{cases} (u, 1, 1, v_3, 0) & \text{if } v = (1, v_3) \text{ with } v_3 \in D, \text{ i.e., } x \in {}_{11}F(\mathcal{P}) \setminus F'(\mathcal{P}) , \\ (u, 1, 1, v_1, 0, v_2, 1, v_3, 0) & \text{if } v = (v_1, 1, v_2, 0, 1, v_3) \text{ with } v_1, v_2, v_3 \in D, \text{ i.e., } x \in {}_{11}F'(\mathcal{P}) . \end{cases}$$

In terms of rooted trees, $(y, 0) \in \mathcal{T}_{n+1}$ is obtained from $(x, 0) \in \mathcal{T}_{n+1}$ by a heavy rotation.

If $u = ()$, i.e., $x \in {}_{10}F(\mathcal{P})$, then we have

$$y = \begin{cases} (1, 0, v_3, 1) & \text{if } v = (1, v_3) \text{ with } v_3 \in D, \text{ i.e., } x \in {}_{10}F(\mathcal{P}) \setminus F'(\mathcal{P}) , \\ (1, 1, v_1, 0, v_2, 0, v_3, 1) & \text{if } v = (v_1, 1, v_2, 0, 1, v_3) \text{ with } v_1, v_2, v_3 \in D, \text{ i.e., } x \in {}_{10}F'(\mathcal{P}) . \end{cases}$$

In terms of rooted trees, $(y, 0) \in \mathcal{T}_{n+1}$ is obtained from $(x, 0) \in \mathcal{T}_{n+1}$ by a light rotation.

Proof. We follow the cycle C in Figure 11 from some vertex $x \in F(\mathcal{P})$ until we encounter the next vertex $y \in F(\mathcal{P})$. The case when $u \neq ()$ is shown in Figure 13. In this case, following the path of \mathcal{P} from $x \in {}_{11}F(\mathcal{P})$ leads to the vertex $x' := (u, 0, 1, v) \in {}_1L(\mathcal{P})$ by Lemma 8. Then, depending on whether $x' \in L(\mathcal{P}) \setminus L'(\mathcal{P})$ or $x' \in L'(\mathcal{P})$, we continue along the cycle via an edge from E^n or E^{n-2} , respectively, from level $n+1$ to level n . This corresponds to the two subcases distinguished by v and shown in the left and right column in Figure 13. By Lemma 10, we get to a vertex $z' \in f(L(\mathcal{P}_0))$ that is last on some path of $f(\mathcal{P}_0)$. By following this path backwards using Lemma 8 we get to the corresponding first vertex $z \in f(F(\mathcal{P}_0))$. Using again Lemma 10, we then traverse an edge from E^n to go from from level n to level $n+1$ where we encounter the vertex $y \in F(\mathcal{P}_0)$. Observe in Figure 13 that y has exactly the claimed form.

The case when $u = ()$ is depicted in Figure 14 and can be verified analogously. Note that in this case we in addition visit a vertex from $I(\mathcal{P})$ and from $f(I(\mathcal{P}))$ before coming to y .

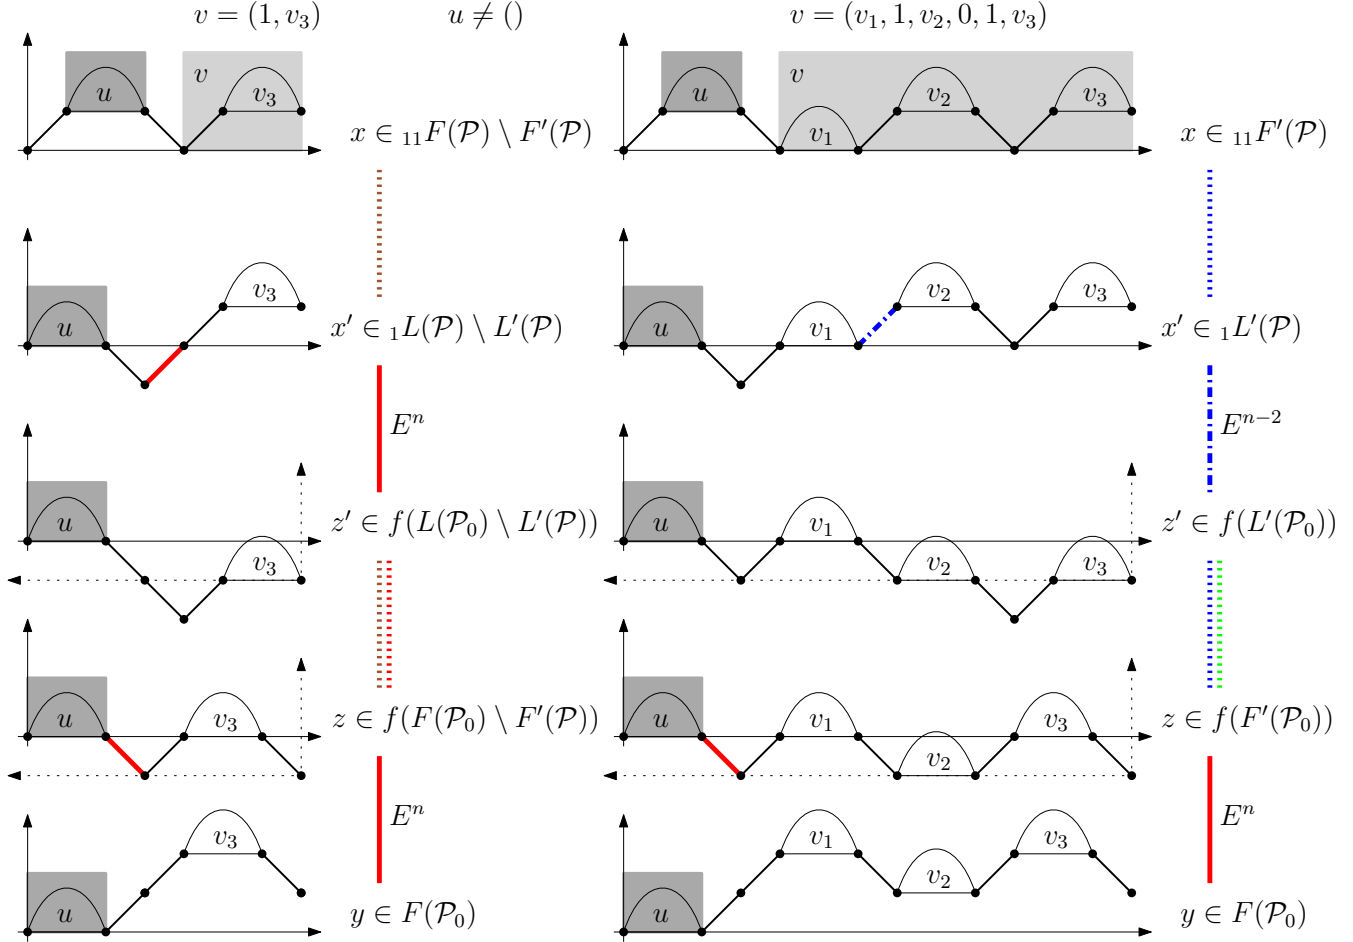


FIGURE 13. Proof of the first part of Lemma 11. The two columns represent the subpaths from x to y on a cycle from our factor for the two cases of v . Applying $f = \overline{\text{rev}}$ can be interpreted as reading the lattice path backwards, and this is indicated by the dotted coordinate systems.

□

The number of cycles in the factor \mathcal{C}_{2n+1} is 1, 1, 1, 4, 6, 19, 49, 150, 442, 1424 for $n = 1, 2, \dots, 11$. This sequence matches the first entries of [OEI17], i.e., the number of plane trivalent trees with n internal vertices. A *plane trivalent tree* is a tree where every vertex has degree 1 or 3, and the neighbors of each vertex have a specified cyclic ordering. We establish the correspondence between cycles from our factor and plane trivalent trees in the following lemma. This special family of trees appears only in this lemma and its proof. All subsequent arguments use again the set of rooted trees \mathcal{T}_{n+1} introduced before.

Lemma 12. *For any $n \geq 1$, the cycles from the factor \mathcal{C}_{2n+1} defined in (7) are in bijection with the set of plane trivalent trees with n internal vertices.*

Proof. We first define, for any $x \in D$, binary trees $\ell(x)$ and $r(x)$. If $x = ()$, then $\ell(x)$ and $r(x)$ consist only of a single vertex. Otherwise we write x uniquely as $x = (1, u, 0, v) = (u', 1, v', 0)$ with $u, v, v', v' \in D$, and then $\ell(x)$ consists of a root with left child $\ell(u)$ and right child $r(v)$, and $r(x)$ consists of a root with left child $\ell(u')$ and right child $r(v')$. Given any vertex $x \in F(\mathcal{P}) = D_{2n+1, n+1}^0$, we map the bitstring $x' := (x, 0)$ to a trivalent tree $\tau(x')$ rooted at one of its n internal vertices




FIGURE 14. Proof of the second part of Lemma 11. The two columns represent the subpaths from x to y on a cycle from our factor for the two cases of v . Applying $f = \overline{\text{rev}}$ can be interpreted as reading the lattice path backwards, and this is indicated by the dotted coordinate systems.

as follows. We first write x' uniquely in the form $x' = (1, u, 0, v, 1, w, 0)$ with $u, v, w \in D$, and we define $\tau(x')$ as the tree that consists of a root with left child $\ell(u)$, middle child $r(v)$ and right child $r(w)$. One can show that under this bijection τ between $D_{2n+1, n+1}^0 \circ 0$ and trivalent trees rooted at one of their n internal vertices, the mapping ρ corresponds to rotating the root of the trivalent tree to the leftmost child until the root is again an internal vertex. Consequently, by Lemma 11 cycles

from \mathcal{C}_{2n+1} correspond to equivalence classes of rooted trivalent trees whose root is one of their n internal vertices under this rotation operation. Obviously, these are exactly plane trivalent trees with n internal vertices. \square

4.3. Flippable pairs. In this section we define certain 6-cycles in the graph Q_{2n+1} between levels $n+1$ and $n+2$ that can be used to join pairs of cycles from our factor \mathcal{C}_{2n+1} as described in the beginning of this section (see Figure 15). Let us emphasize here that all modifications of the cycle factor happen only between the two upper levels $n+1$ and $n+2$.

We say that two vertices $x, y \in F(\mathcal{P}) = D_{2n+1, n+1}^=0$ form a *flippable pair* (x, y) , if x and y have the form

$$\begin{aligned} x &= (1, u_1, 1, u_2, \dots, 1, u_d, 1, 1, 0, w, 0, v_d, 0, v_{d-1}, 0, \dots, v_1, 0, v_0) , \\ y &= (1, u_1, 1, u_2, \dots, 1, u_d, 1, 0, 1, w, 0, v_d, 0, v_{d-1}, 0, \dots, v_1, 0, v_0) \end{aligned} \quad (9)$$

with $d \geq 0$ and $u_1, \dots, u_d, v_1, \dots, v_d, w, (v_0, 0) \in D$. Recall that $(x, 0)$ and $(y, 0)$ can be viewed as rooted trees from \mathcal{T}_{n+1} . If (x, y) is a flippable pair, then the tree $(y, 0)$ is obtained from the tree $(x, 0)$ by moving a pending edge from a vertex in the leftmost subtree to its predecessor. We refer to this as a *pull operation*, see Figure 12.

Any 6-cycle between levels $n+1$ and $n+2$ of Q_{2n+1} can be uniquely encoded as a string of length $2n+1$ over $\{0, 1, *\}$ with n many 1s, $n-2$ many 0s and three $*$ s. The 6-cycle corresponding to this string is obtained by substituting the three $*$ s by all six combinations of at least two different symbols from $\{0, 1\}$. We define a set of 6-cycles \mathcal{S}_{2n+1} between levels $n+1$ and $n+2$ of Q_{2n+1} consisting of all 6-cycles

$$C_6(x, y) := (u_1, 0, u_2, 0, \dots, u_d, 0, 1, *, *, w, *, v_d, 1, v_{d-1}, 1, \dots, v_1, 1, v_0) \quad (10)$$

for a flippable pair (x, y) , $x, y \in D_{2n+1, n+1}^=0$, as in (9).

The following result was proved in [GMN17, Proposition 3]. For any $x \in D_{2n+1, n+1}^=0$, we write $P(x)$ for the path from the set \mathcal{P} defined in (6) that starts at the vertex x .

Lemma 13. *For any $n \geq 1$, the 6-cycles $C_6(x, y) \in \mathcal{S}_{2n+1}$ defined in (10) have the following properties:*

- (i) *Let (x, y) be a flippable pair. The 6-cycle $C_6(x, y)$ intersects $P(x)$ in two non-incident edges and it intersects $P(y)$ in a single edge. Moreover, the symmetric difference of the edge sets of the two paths $P(x) \in \mathcal{P}$ and $P(y) \in \mathcal{P}$ with the 6-cycle $C_6(x, y)$ gives two paths $P'(x)$ and $P'(y)$ on the same set of vertices as $P(x)$ and $P(y)$, interconnecting x with the last vertex of $P(y)$, and y with the last vertex of $P(x)$, respectively.*
- (ii) *For any flippable pairs (x, y) and (x', y') , the 6-cycles $C_6(x, y)$ and $C_6(x', y')$ are edge-disjoint.*
- (iii) *For any flippable pairs (x, y) and (x, y') , the two pairs of edges that the two 6-cycles $C_6(x, y)$ and $C_6(x, y')$ have in common with the path $P(x)$ are not interleaved, but one pair appears before the other pair along the path.*

4.4. Proof of Theorem 1. With Lemma 11 and 13 in hand, we are now ready to prove Theorem 1.

Proof of Theorem 1. Let \mathcal{C}_{2n+1} and \mathcal{S}_{2n+1} be the cycle factor and the set of 6-cycles defined in Section 4.1 and 4.3, respectively.

Consider two cycles $C, C' \in \mathcal{C}_{2n+1}$ containing paths $P, P' \in \mathcal{P}$ with first vertices $x, y \in D_{2n+1, n+1}^=0$, respectively, such that (x, y) is a flippable pair. By Lemma 13 (i), the symmetric difference of the edge sets $C \cup C'$ and $C_6(x, y)$ forms a single cycle on the same vertex set as $C \cup C'$, i.e., this joining operation reduces the number of cycles in the factor by one, see Figure 15. Recall that in terms of rooted trees, $(y, 0) \in \mathcal{T}_{n+1}$ is obtained from $(x, 0) \in \mathcal{T}_{n+1}$ by a pull operation, see Figure 12.




FIGURE 15. Two cycles from our factor joined by taking the symmetric difference with a 6-cycle. The paths $P(x)$ and $P(y)$ from the set \mathcal{P} (solid black) lying on the two cycles traverse the 6-cycle $C_6(x, y)$ (solid gray) as shown. The symmetric difference yields paths $P'(x)$ and $P'(y)$ that have flipped end vertices.

We repeat this joining operation until all cycles in the factor are joined to a single Hamilton cycle. For this purpose we define an auxiliary graph \mathcal{H}_{n+1} whose nodes represent the cycles in the factor \mathcal{C}_{2n+1} and whose edges connect pairs of cycles that can be connected to a single cycle with such a joining operation that involves a 6-cycle from the set \mathcal{S}_{2n+1} . Formally, the node set of \mathcal{H}_{n+1} is given by partitioning the set \mathcal{T}_{n+1} into equivalence classes under the mapping ρ . By Lemma 11, the nodes of \mathcal{H}_{n+1} therefore indeed correspond to the cycles in the factor \mathcal{C}_{2n+1} . Specifically, each rooted tree $(x, 0) \in \mathcal{T}_{n+1} = D_{2n+1, n+1}^{\neq 0} \circ 0$ belonging to some node of \mathcal{H}_{n+1} corresponds to the first vertex x of some path $P \in \mathcal{P}$ such that P lies on the cycle corresponding to that node. For every flippable pair (x, y) , $x, y \in D_{2n+1, n+1}^{\neq 0}$, we add the edge to \mathcal{H}_{n+1} that connects the node containing the tree $(x, 0)$ to the node containing the tree $(y, 0)$. By our initial argument, such a flippable pair yields a 6-cycle $C_6(x, y)$ that can be used to join the two corresponding cycles to a single cycle. As mentioned before, the 6-cycle lies entirely between levels $n + 1$ and $n + 2$.

To complete the proof of Theorem 1, it therefore suffices to prove that the auxiliary graph \mathcal{H}_{n+1} is connected. Indeed, if \mathcal{H}_{n+1} is connected, then we can pick a spanning tree in \mathcal{H}_{n+1} , corresponding to a collection of 6-cycles $\mathcal{T}_{2n+1} \subseteq \mathcal{S}_{2n+1}$, such that the symmetric difference between the edge sets $\mathcal{C}_{2n+1} \Delta \mathcal{T}_{2n+1}$ forms a Hamilton cycle in the middle four levels of Q_{2n+1} . Of crucial importance here are properties (ii) and (iii) in Lemma 13, which ensure that whatever subset of 6-cycles we use in this joining process, they will not interfere with each other, guaranteeing that each 6-cycle indeed reduces the number of cycles by one, as desired.

At this point we reduced the problem of proving that the middle four levels of Q_{2n+1} have a Hamilton cycle to showing that the auxiliary graph \mathcal{H}_{n+1} is connected, which is much easier. Indeed, all we need to show is that any rooted tree from \mathcal{T}_{n+1} can be transformed into any other tree from \mathcal{T}_{n+1} by a sequence of heavy rotations, light rotations, pulls and their inverse operations. Recall that heavy and light rotations correspond to following the same cycle from \mathcal{C}_{2n+1} (staying at the same node in \mathcal{H}_{n+1}), and a pull corresponds to a joining operation (traversing an edge in \mathcal{H}_{n+1} to another node). For this we show that any rooted tree $x \in \mathcal{T}_{n+1}$ can be transformed into the special tree $s := (1, 1, 0, 1, 0, \dots, 1, 0, 0, 1, 0) \in \mathcal{T}_{n+1}$, i.e., a right-light tree with a root of degree two and a star as its left subtree. To achieve this we distinguish three cases, see Figure 16.

- (a) x is left-light and right-light. By a light rotation we obtain a right-light tree with a root of degree two. We then repeatedly pull pending edges in the left subtree towards the left child of the root until we end up at the tree s .
- (b) x is left-heavy. First, by a sequence of pulls and inverse pulls we transform the left subtree of x into a path of length $p \geq 2$. Then, by exactly $p - 2$ heavy rotations we arrive at a tree whose left subtree is a path on two edges. Another heavy rotation yields a left-light tree with a root of degree two. Next, performing a light rotation yields a left-light and right-light tree, and we then continue as in case (a).




FIGURE 16. Transformation of trees from \mathcal{T}_{n+1} into the tree s in the proof of Theorem 1. Shaded subtrees are non-empty.

(c) x is left-light and right-heavy. By a single inverse heavy rotation we obtain a tree that is left-heavy, and we then continue as in case (b).

This shows that \mathcal{H}_{n+1} is connected, and thus completes the proof. \square

5. COMPUTER EXPERIMENTS

The numbers of cycles in the factor through the middle 2ℓ levels of the $(2n+1)$ -cube constructed as in the proof of Theorem 2 using the two edge-disjoint SCDs $\mathcal{D} := \mathcal{D}_0$ and $\mathcal{D}' := \overline{\mathcal{D}_0}$ are shown in Table 2 for $n = 1, 2, \dots, 12$. Note that the number of cycles in the middle two levels (the first column $\ell = 1$) seems to grow faster than the number of cycles in the middle four levels (the second column $\ell = 2$).

If instead we perform the proof of Theorem 2 with the two edge-disjoint SCDs constructed via our product construction, we obtain another set of cycle factors. Specifically, denoting by $\mathcal{D}_0(a)$ the SCD \mathcal{D}_0 in the a -cube for $a \in \{2, 3\}$, then $\mathcal{D} := \mathcal{D}_0(3) \times \mathcal{D}_0(2)^{(n-3)/2}$ and $\mathcal{D}' := \overline{\mathcal{D}_0(3)} \times \overline{\mathcal{D}_0(2)}^{(n-3)/2}$ are two edge-disjoint SCDs in the $(2n+1)$ -cube, where \times denotes the product operation described in the proof of Theorem 5, and the exponent denotes the $(n-3)/2$ -fold such product. Using those SCDs, we obtain cycle factors whose cycle lengths are shown in Table 3.

Being able to analyze the structure of these cycle factors might be the key towards a general solution of Problem M, using one of these cycle factors as a starting point, and then modifying it locally to join the cycles to a single Hamilton cycle. Recall that the proof of Theorem 1 used the very same strategy, but it started from a different cycle factor tailored to the case $\ell = 2$ whose structure we were able to analyze.

6. OPEN PROBLEMS

We conclude with some interesting open problems.

- Understanding the structure of the cycle factors presented in Table 2 and Table 3 is an important step towards a general solution of Problem M. What is the number and length of cycles in these factors in terms of n and ℓ ? Is there a combinatorial interpretation of those numbers?
- Another promising line of attack towards Problem M is to generalize the construction of the cycle factor for $\ell = 2$ via lexical matchings defined in (5) in Section 4 to larger values of ℓ , and to analyze the resulting factor.
- What properties does our new SCD \mathcal{D}_1 have, in addition to being edge-disjoint from \mathcal{D}_0 ? Can we exploit this construction with respect to other applications, e.g., Venn diagrams? Are there other explicit constructions of SCDs in the n -cube, different from \mathcal{D}_0 , \mathcal{D}_1 , and their complements?

TABLE 2. Number of cycles of the factor in the middle 2ℓ levels of the $(2n+1)$ -cube for $n = 1, 2, \dots, 12$ and $1 \leq \ell \leq n+1$ arising from the proof of Theorem 2 using the two edge-disjoint SCDs $\mathcal{D} := \mathcal{D}_0$ and $\mathcal{D}' := \overline{\mathcal{D}_0}$.

n	$\ell = 1$	2	3	4	5	6	7	8	9	10	11	12	13
1	1	2											
2	2	3	6										
3	3	6	19	24									
4	6	10	58	95	102								
5	12	20	181	350	419	428							
6	26	39	552	1246	1644	1749	1760						
7	73	74	1633	4292	6263	6974	7127	7140					
8	146	138	4750	14560	23380	27344	28546	28751	28766				
9	360	300	13500	48892	86156	105890	113477	115290	115559	115576			
10	1408	552	37716	163624	314960	406559	448446	461034	463696	464033	464052		
11	2412	1138	103998	547614	1145771	1551226	1763481	1838964	1859347	1863014	1863431	1863452	
12	10204	2068	284316	1836489	4156230	5892150	6904696	7315848	7448880	7479282	7484252	7484753	7484776

TABLE 3. Number of cycles of the factor in the middle 2ℓ levels of the $(2n+1)$ -cube for $n = 1, 2, \dots, 12$ and $1 \leq \ell \leq n+1$ arising from the proof of Theorem 2 using the two edge-disjoint SCDs $\mathcal{D} := \mathcal{D}_0(3) \times \mathcal{D}_0(2)^{(n-3)/2}$ and $\mathcal{D}' := \overline{\mathcal{D}_0(3)} \times \overline{\mathcal{D}_0(2)}^{(n-3)/2}$.

n	$\ell = 1$	2	3	4	5	6	7	8	9	10	11	12	13
1	1	2											
2	2	3	4										
3	3	8	11	12									
4	10	22	34	39	40								
5	24	68	109	132	139	140							
6	80	213	362	456	494	503	504						
7	239	700	1225	1600	1779	1836	1847	1848					
8	802	2336	4222	5676	6466	6770	6850	6863	6864				
9	2638	7980	14740	20324	23662	25140	25617	25724	25739	25740			
10	9052	27618	52064	73330	87068	93839	96378	97084	97222	97239	97240		
11	31186	96904	185628	266344	321857	351676	364231	368320	369319	369492	369511	369512	
12	109460	343438	667320	972989	1194550	1322256	1381274	1403006	1409266	1410630	1410842	1410863	1410864

- We conjecture that the n -cube has $\lfloor n/2 \rfloor + 1$ pairwise edge-disjoint SCDs, but so far we can only establish the existence of four edge-disjoint SCDs for all $n \geq 12$. The main difficulty here is that we are missing a simple criterion like Hall's matching condition guaranteeing the existence of an SCD. Even finding five edge-disjoint SCDs in the n -cube for some small fixed n would be interesting, as this solution would extend to infinitely many larger values of n by Theorem 5. Beyond that, it would be very nice to construct more than constantly many edge-disjoint SCDs in the n -cube as n grows. The smallest instance where more than four edge-disjoint SCDs may exist is the 8-cube, recall Table 1. For $n = 9$ and $n = 11$ we only know two pairwise edge-disjoint SCDs, which is unsatisfying.

REFERENCES

- [Aig73] M. Aigner. Lexicographic matching in Boolean algebras. *J. Combin. Theory Ser. B*, 14:187–194, 1973.
- [Bol86] B. Bollobás. *Combinatorics: set systems, hypergraphs, families of vectors and combinatorial probability*. Cambridge University Press, Cambridge, 1986.
- [BR96] B. Bultena and F. Ruskey. Transition restricted Gray codes. *Electron. J. Combin.*, 3(1):Paper 11, 11 pp., 1996.
- [dBvETK51] N. de Bruijn, C. van Ebbenhorst Tengbergen, and D. Kruyswijk. On the set of divisors of a number. *Nieuw Arch. Wiskunde (2)*, 23:191–193, 1951.

- [DDGŠ13] D. Dimitrov, T. Dvořák, P. Gregor, and R. Škrekovski. Linear time construction of a compressed Gray code. *European J. Combin.*, 34(1):69–81, 2013.
- [DKS94] D. A. Duffus, H. A. Kierstead, and H. S. Snevily. An explicit 1-factorization in the middle of the Boolean lattice. *J. Combin. Theory Ser. A*, 65(2):334–342, 1994.
- [EHH01] M. El-Hashash and A. Hassan. On the Hamiltonicity of two subgraphs of the hypercube. In *Proceedings of the Thirty-second Southeastern International Conference on Combinatorics, Graph Theory and Computing (Baton Rouge, LA, 2001)*, volume 148, pages 7–32, 2001.
- [GG03] L. Goddyn and P. Gvozdjak. Binary Gray codes with long bit runs. *Electron. J. Combin.*, 10:Paper 27, 10 pp., 2003.
- [GK76] C. Greene and D. J. Kleitman. Strong versions of Sperner’s theorem. *J. Combin. Theory Ser. A*, 20(1):80–88, 1976.
- [GKS04] J. Griggs, C. E. Killian, and C. D. Savage. Venn diagrams and symmetric chain decompositions in the Boolean lattice. *Electron. J. Combin.*, 11(1):Paper 2, 30 pp., 2004.
- [GM18] P. Gregor and T. Mütze. Trimming and gluing Gray codes. *Theoret. Comput. Sci.*, 714:74–95, 2018.
- [GMN17] P. Gregor, T. Mütze, and J. Nummenpalo. A short proof of the middle levels theorem. *arXiv:1710.08249*, Oct 2017.
- [Gra53] F. Gray. Pulse code communication, 1953. March 17, 1953 (filed Nov. 1947). U.S. Patent 2,632,058.
- [GŠ10] P. Gregor and R. Škrekovski. On generalized middle-level problem. *Inform. Sci.*, 180(12):2448–2457, 2010.
- [Hol17] A. E. Holroyd. Perfect snake-in-the-box codes for rank modulation. *IEEE Trans. Inform. Theory*, 63(1):104–110, 2017.
- [HRW12] A. E. Holroyd, F. Ruskey, and A. Williams. Shorthand universal cycles for permutations. *Algorithmica*, 64(2):215–245, 2012.
- [Joh09] J. R. Johnson. Universal cycles for permutations. *Discrete Math.*, 309(17):5264–5270, 2009.
- [Joh11] J. R. Johnson. An inductive construction for Hamilton cycles in Kneser graphs. *Electron. J. Combin.*, 18(1):Paper 189, 12 pp., 2011.
- [Knu11] D. E. Knuth. *The Art of Computer Programming. Vol. 4A. Combinatorial Algorithms. Part 1*. Addison-Wesley, Upper Saddle River, NJ, 2011.
- [KT88] H. A. Kierstead and W. T. Trotter. Explicit matchings in the middle levels of the Boolean lattice. *Order*, 5(2):163–171, 1988.
- [LS03] S. Locke and R. Stong. Problem 10892: Spanning cycles in hypercubes. *Amer. Math. Monthly*, 110:440–441, 2003.
- [MN17] T. Mütze and J. Nummenpalo. A constant-time algorithm for middle levels Gray codes. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 2238–2253, 2017.
- [MNW17] T. Mütze, J. Nummenpalo, and B. Walczak. Sparse Kneser graphs are Hamiltonian. To appear in the *Proceedings of the 50th Annual ACM Symposium on the Theory of Computing, STOC 2018*, *arXiv:1711.01636*, Nov 2017.
- [Müt16] T. Mütze. Proof of the middle levels conjecture. *Proc. Lond. Math. Soc.*, 112(4):677–713, 2016.
- [OEI17] The On-Line Encyclopedia of Integer Sequences, Sequence A001683. <http://oeis.org>, 2017.
- [Pik99] O. Pikhurko. On edge decompositions of posets. *Order*, 16(3):231–244 (2000), 1999.
- [RSW06] F. Ruskey, C. D. Savage, and S. Wagon. The search for simple symmetric Venn diagrams. *Notices Amer. Math. Soc.*, 53(11):1304–1312, 2006.
- [Sav93] C. D. Savage. Long cycles in the middle two levels of the Boolean lattice. *Ars Combin.*, 35(A):97–108, 1993.
- [Sav97] C. D. Savage. A survey of combinatorial Gray codes. *SIAM Rev.*, 39(4):605–629, 1997.
- [SK79] J. Shearer and D. J. Kleitman. Probabilities of independent choices being ordered. *Stud. Appl. Math.*, 60(3):271–276, 1979.
- [Spi17] H. Spink. Orthogonal symmetric chain decompositions of hypercubes. *arXiv:1706.08545*, June 2017.
- [ST14] N. Streib and W. T. Trotter. Hamiltonian cycles and symmetric chains in Boolean lattices. *Graphs Combin.*, 30(6):1565–1586, 2014.
- [Sta15] R. P. Stanley. *Catalan numbers*. Cambridge University Press, New York, 2015.
- [SvZ08] I. N. Suparta and A. J. van Zanten. A construction of Gray codes inducing complete graphs. *Discrete Math.*, 308(18):4124–4132, 2008.
- [SW95] C. D. Savage and P. Winkler. Monotone Gray codes and the middle levels problem. *J. Combin. Theory Ser. A*, 70(2):230–248, 1995.
- [SW18] J. Sawada and A. Williams. A Hamilton path for the sigma-tau problem. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 568–575, 2018.
- [Too53] G. C. Tootill. *Proceedings IEE, Part B Supplement*, 103, 1953.

- [WW77] D. E. White and S. G. Williamson. Recursive matching algorithms and linear orders on the subset lattice. *J. Combin. Theory Ser. A*, 23(2):117–127, 1977.