

On the central levels problem*

Petr Gregor^{†1}, Ondřej Mička^{‡1}, Torsten Mütze^{§1,2}

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

²Department of Computer Science, University of Warwick, United Kingdom

ABSTRACT. The *central levels problem* asserts that the subgraph of the $(2m + 1)$ -dimensional hypercube induced by all bitstrings with at least $m + 1 - \ell$ many 1s and at most $m + \ell$ many 1s, i.e., the vertices in the middle 2ℓ levels, has a Hamilton cycle for any $m \geq 1$ and $1 \leq \ell \leq m + 1$. This problem was raised independently by Buck and Wiedemann, Savage, Gregor and Škrekovski, and by Shen and Williams, and it is a common generalization of the well-known *middle levels problem*, namely the case $\ell = 1$, and classical binary Gray codes, namely the case $\ell = m + 1$. In this paper we present a general constructive solution of the central levels problem. Our results also imply the existence of optimal cycles through any sequence of ℓ consecutive levels in the n -dimensional hypercube for any $n \geq 1$ and $1 \leq \ell \leq n + 1$. Moreover, extending an earlier construction by Streib and Trotter, we construct a Hamilton cycle through the n -dimensional hypercube, $n \geq 2$, that contains the symmetric chain decomposition constructed by Greene and Kleitman in the 1970s, and we provide a loopless algorithm for computing the corresponding Gray code.

KEYWORDS: Gray code, Hamilton cycle, hypercube, middle levels, symmetric chain decomposition

1. INTRODUCTION

The *n -dimensional hypercube*, or *n -cube* for short, is the graph Q_n formed by all $\{0, 1\}$ -strings of length n , with an edge between any two bitstrings that differ in exactly one bit. This family of graphs has numerous applications in computer science and discrete mathematics, many of which are tied to famous problems and conjectures, such as the sensitivity conjecture of Nisan and Szegedy [NS94], recently proved by Huang [Hua19]; Erdős and Guys' crossing number problem [EG73] (see [FdFSV08]); Füredi's conjecture [Für85] on equal-size chain partitions (see [Tom15]); Shearer and Kleitman's conjecture [SK79] on orthogonal symmetric chain decompositions (see [Spi19]); the Ruskey-Savage problem [RS93] on matching extendability (see [Fin07, Fin19]), and the conjectures of Norine, and Feder and Subi on edge-antipodal colorings [Nor08, FS13], to name just a few.

The focus of this paper are Hamilton cycles in the n -cube and its subgraphs. A Hamilton cycle in a graph is a cycle that visits every vertex exactly once, and in the context of the n -cube, such a cycle is often referred to as a *Gray code*. Gray codes have found applications in signal processing, circuit testing, hashing, data compression, experimental design, binary counters, image processing, and in solving puzzles like the Towers of Hanoi or the Chinese rings; see Savage's survey [Sav97]. Gray codes are also fundamental for efficient algorithms to exhaustively generate combinatorial objects,

*An extended abstract of this paper appeared in the Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP 2020) [GMM20]. This work was supported by Czech Science Foundation grant GA19-08554S. Torsten Mütze is also supported by German Science Foundation grant 413902284.

[†]E-Mail: gregor@ktiml.mff.cuni.cz.

[‡]E-Mail: micka@ktiml.mff.cuni.cz.

[§]E-Mail: torsten.mutze@warwick.ac.uk.

a topic that is covered in depth in the most recent volume of Knuth’s seminal series ‘*The Art of Computer Programming*’ [Knu11].

To start with, it is an easy exercise to show that the n -cube has a Hamilton cycle for any $n \geq 2$. One such cycle is given by the classical *binary reflected Gray code* Γ_n [Gra53], defined inductively by $\Gamma_1 := 0, 1$ and $\Gamma_{n+1} := 0\Gamma_n, 1\Gamma_n^R$, where Γ^R denotes the reversal of the sequence Γ , and 0Γ or 1Γ means prefixing all strings in the sequence Γ by 0 or 1, respectively. For instance, this construction gives $\Gamma_2 = 00, 01, 11, 10$ and $\Gamma_3 = 000, 001, 011, 010, 110, 111, 101, 100$. The problem of finding a Hamilton cycle becomes considerably harder when we restrict our attention to subgraphs of the cube induced by a sequence of consecutive levels, where the k th level of Q_n , $0 \leq k \leq n$, is the set of all bitstrings with exactly k many 1s in them. One such instance is the famous *middle levels problem*, raised in the 1980s by Havel [Hav83] and independently by Buck and Wiedemann [BW84], which asks for a Hamilton cycle in the subgraph of the $(2m + 1)$ -cube induced by levels m and $m + 1$. This problem received considerable attention in the literature, and a construction of such a cycle for all $m \geq 1$ was provided only recently by Mütze [Müt16]. A much simpler construction was described subsequently by Gregor, Mütze, and Nummenpalo [GMN18].

There are several intuitive explanations why the middle levels problem is considerably harder than finding a Hamilton cycle in the entire cube. First of all, the entire $(n + 1)$ -cube has a simple inductive structure: It consists of two copies of Q_n plus a perfect matching connecting the two copies, and hence a Hamilton cycle in Q_{n+1} can be obtained by gluing together two cycles in these copies of Q_n via two matching edges, which can yield the cycle Γ_{n+1} from before. The subgraph of the $(2m + 1)$ -cube induced by the middle two levels, on the other hand, does not allow for such an inductive decomposition. Note also that as a consequence of the inductive construction of the binary reflected Gray code Γ_n , the number of times that each bit is flipped follows a highly skewed distribution $(2, 2, 4, 8, \dots, 2^{n-1})$. In stark contrast to this, along *any* Hamilton cycle in the middle levels graph, the number of times that each bit is flipped *must* follow the uniform distribution $(2C_m, 2C_m, \dots, 2C_m)$ with $C_m = \frac{1}{2m+1} \binom{2m+1}{m} = \frac{1}{m+1} \binom{2m}{m}$ the m th Catalan number, an observation that makes an easy induction proof unlikely.

1.1. Our results. In this paper we consider the *central levels problem*, a broad generalization of the middle levels problem: Does the subgraph of the $(2m + 1)$ -cube induced by the middle 2ℓ levels, i.e., by levels $m + 1 - \ell, \dots, m + \ell$, have a Hamilton cycle for any $m \geq 1$ and $1 \leq \ell \leq m + 1$? This problem was raised already in Buck and Wiedemann’s paper [BW84], and was reiterated independently by Savage [Sav93], Gregor and Škrekovski [GŠ10], and by Shen and Williams [SW19]. Clearly, the case $\ell = 1$ of the central levels problem is the aforementioned middle levels problem (solved in [Müt16]). Moreover, the case $\ell = 2$ was solved affirmatively in a paper by Gregor, Jäger, Mütze, Sawada, and Wille [GJM⁺21]. Also, the case $\ell = m + 1$ is established by the binary reflected Gray code Γ_{2m+1} . Furthermore, the case $\ell = m$ was solved by Buck and Wiedemann [BW84] and independently by El-Hashash and Hassan [EHH01], and in a more general setting by Locke and Stong [LS03], and the case $\ell = m - 1$ was settled in [GŠ10].

The main contribution of this paper is to solve the central levels problem affirmatively in full generality; see Figure 1 (a)–(d).

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

As the case $\ell = 1$ of Theorem 1 has been proved before, the proof of Theorem 1 presented in this paper assumes that $\ell \geq 2$. Nevertheless, this proof can be seen as a generalization of the earlier proofs [Müt16, GMN18] and [GJM⁺21] for the cases $\ell = 1$ and $\ell = 2$, respectively; see the remarks in Section 3.1 below.

The most general question in this context is to ask for a Hamilton cycle in Q_n that visits all vertices in *any* sequence of ℓ consecutive levels, i.e., the levels need not be symmetric around the middle, and the dimension n needs not be odd. Note however, that the n -cube is bipartite, and the partition classes are given by all even and odd levels, respectively. Consequently, any subgraph consisting of a sequence of ℓ consecutive levels is also bipartite, and to circumvent the imbalances that prevent the existence of a Hamilton cycle for general n and ℓ , we have to slightly generalize the notion of Hamilton cycles, and there are two reasonable such generalized notions. Firstly, a *saturating cycle* in a bipartite graph is a cycle that visits all vertices in the smaller partition class. If one partition class is empty, then the empty set is considered a saturating cycle, and if the smaller partition class has size 1, then a single edge is considered a saturating cycle. Secondly, a *tight enumeration* in a (bipartite) subgraph of the cube is a cyclic listing of all its vertices where the total number of bits flipped is exactly the number of vertices plus the difference in size between the two partition classes. If the graph has only a single vertex, this vertex is considered a tight enumeration. Clearly, if both partition classes have the same size, like in the central levels problem, then a saturating cycle and a tight enumeration are equal to a Hamilton cycle. In fact, all cases of this more general problem on saturating cycles and tight enumerations, except the cases of the central levels problem, were solved affirmatively already in [GM18], some of them conditional on a ‘yes’ answer to the central levels problem. Combining Theorem 1 with these previous results, we now also obtain an unconditional result for this more general question.

Corollary 2. *For any $n \geq 1$ and $1 \leq \ell \leq n + 1$, the subgraph of the n -cube induced by any sequence of ℓ consecutive levels has both a saturating cycle and a tight enumeration.*

Proof of Corollary 2. If $\ell = 1$, then all vertices are in the same partition class, and the other partition class is empty. It follows that the statement is trivially true for $\ell = 1$ and saturating cycles. To prove it for $\ell = 1$ and tight enumerations, first note that level 0 or level n consist only of a single vertex, which is trivially a tight enumeration. Otherwise we use a well-known result of Tang and Liu [TL73], who showed that for any $n \geq 2$ and $1 \leq k \leq n - 1$, there is a cyclic listing of all bitstrings of Q_n on level k such that any two consecutive strings differ in a transposition of 0 and 1. As two bits are flipped in each step, this is a tight enumeration. We now consider the case $\ell = 2$ for saturating cycles: If one of the two levels is level 0 or level n , which contains only a single vertex, then a single edge is a trivial saturating cycle. Otherwise the result follows from [MS17, Theorem 9].

All remaining cases for saturating cycles and tight enumerations are covered by Theorems 5 and 6 proved in [GM18], respectively. Specifically, part (iv) in both of these theorems is conditional on the validity of Theorem 1, which we establish in this paper. \square

The central levels problem studied in this paper is also closely related to another famous problem, which asks about Hamilton cycles in so-called bipartite Kneser graphs. The *bipartite Kneser graph* $H_{n,k}$, defined for any integers $k \geq 0$ and $n \geq 2k + 1$, is the bipartite graph whose vertex partition is given by all vertices on level k and $n - k$ of Q_n , with an edge between any two bitstrings that differ in exactly $n - 2k$ bits. That is, the edges in $H_{n,k}$ correspond to a level-monotone path in Q_n between a vertex on level k and a vertex on level $n - k$. It was shown that $H_{n,k}$ has a Hamilton cycle for all $k \geq 1$ and $n \geq 2k + 1$ in [MS17], completing a long line of previous partial results [Sim91, Hur94, Che00, Che03].

An essential tool in our proof of Theorem 1 are symmetric chain decompositions. This is a well-known concept from the theory of posets, which we now define specifically 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. It is well-known that the n -cube has an SCD for all $n \geq 1$, and the simplest explicit construction was given by Greene and

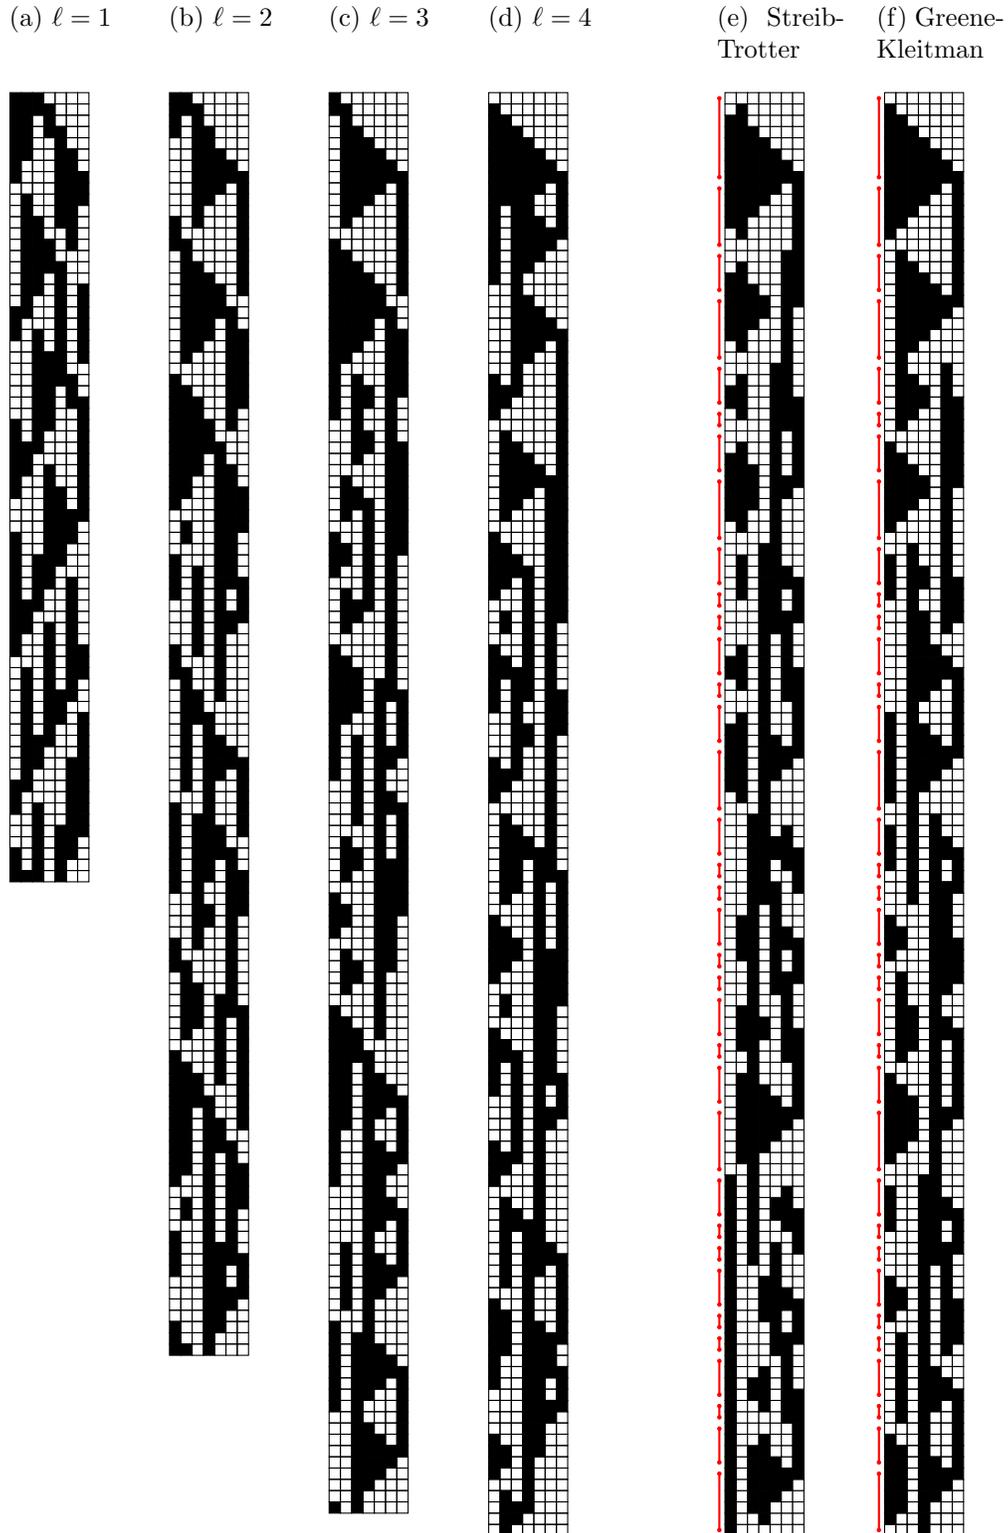


FIGURE 1. (a)-(d) The Hamilton cycles in $Q_{7,\ell}$ for $\ell = 1, 2, 3, 4$ constructed as in our proof of Theorem 1. (e) The Hamilton cycle in Q_7 containing an SCD obtained from the Streib-Trotter construction, with symmetric chains highlighted on the side. (f) The Hamilton cycle in Q_7 containing the Greene-Kleitman SCD obtained from our proof of Theorem 3. In this figure, 1-bits are drawn as black squares, 0-bits as white squares.

Kleitman [GK76] (see Section 2.2 below). Streib and Trotter [ST14] first investigated the interplay between SCDs and Hamilton cycles in the n -cube, and they described an SCD in Q_n that can be extended to a Hamilton cycle; see Figure 1 (e). Streib and Trotter’s SCD, however, is different from the aforementioned Greene-Kleitman SCD. In this paper, we extend Streib and Trotter’s result as follows; see Figure 1 (f).

Theorem 3. *For any $n \geq 2$, the Greene-Kleitman SCD can be extended to a Hamilton cycle in Q_n .*

The Greene-Kleitman SCD has found a large number of applications in the literature, e.g., to construct rotation-symmetric Venn diagrams [GKS04, RSW06], to solve different variants of the Littlewood-Offord problem [Kle65, Gri93] (see also [Bol86, Chap. 4]), or to learn monotone Boolean functions [Knu11, Sec. 7.2.1.6] (see also [dBvETK51, Aig73, WW77, SK79, Pik99]). Knowing that this SCD extends to a Hamilton cycle and that it is a crucial ingredient for solving the general central levels problem adds to this list of interesting properties and applications. Observe also that a Hamilton cycle that extends an SCD has the intriguing property that it minimizes the number of changes of direction from moving up to moving down, or vice versa, between consecutive levels in the cube. For comparison, the monotone paths constructed by Savage and Winkler [SW95] maximize these changes.

Motivated by these results and by the aforementioned conjecture of Ruskey and Savage [RS93] that every matching in Q_n extends to a Hamilton cycle, we raise the following brave conjecture.

Conjecture 4. *Every SCD can be extended to a Hamilton cycle in Q_n .*

Although every SCD of Q_n is the union of two matchings, there are matchings in Q_n that do not extend to an SCD; take for example the two edges obtained by starting at the vertices 0^n and 1^n and flipping the same bit. Consequently, an affirmative answer to Conjecture 4 would cover only some cases of the Ruskey-Savage conjecture.

1.2. Efficient algorithms. We now discuss the algorithmic problem of efficiently computing the Hamilton cycles constructed in the proofs of Theorems 1 and 3. The ultimate goal for any Gray code problem is a *loopless* algorithm, a notion that was introduced by Ehrlich [Ehr73]. The running time of such an algorithm is $\mathcal{O}(1)$ per generated vertex, and the initialization time is ‘reasonable’, linear in n , say. Furthermore, the overall memory requirement should also be polynomial in n , ideally linear, and this *excludes* the space for storing the Hamilton cycle (cf. Ruskey’s [Rus03] ‘Don’t count the output principle’), which may not be needed for some applications.

For the Hamilton cycle in Theorem 3, we present such a loopless algorithm. Specifically, the algorithm uses $\mathcal{O}(1)$ time in every iteration to compute the next bitstring along the Hamilton cycle, its initialization time is $\mathcal{O}(n)$, and its space requirement is $\mathcal{O}(n)$. In this paper, we provide a pseudocode description of this algorithm, and we implemented it in C++, available for download and for demonstration on the Combinatorial Object Server [cos].

On the other hand, for the Hamilton cycles constructed in the proof of Theorem 1, we have no efficient generation algorithm, and it remains a challenging open problem to find one. While our proof of Theorem 1 is constructive and translates straightforwardly into an algorithm that computes the desired Hamilton cycle in time and space that are polynomial in the size of the graph (the middle 2ℓ levels of Q_n , $n := 2m + 1$), these quantities are clearly exponential in n . There are fundamental obstacles that prevent us from obtaining algorithms with polynomial bounds from our general proof of the central levels problem, as explained below (at the end of Section 7). The only two special cases of the central levels problem for which loopless algorithms are known, which were found prior to our work, are the binary reflected Gray code Γ_n [BER76] and the middle levels problem [MN17], i.e., the extreme cases $\ell = m + 1$ and $\ell = 1$, respectively.

1.3. Proof ideas. We first describe the ideas for proving Theorem 1. For any $m \geq 1$ we define $n := 2m + 1$, and for $1 \leq \ell \leq m + 1$ we let $Q_{n,\ell}$ denote the subgraph of Q_n induced by the middle 2ℓ levels. To prove that $Q_{n,\ell}$ has a Hamilton cycle for general m and ℓ , we combine and generalize the tools and techniques developed for the cases $\ell = 1$ and $\ell = 2$ in [GMN18] and [GJM⁺21], respectively. Our proof proceeds in two steps: In a first step, we construct a *cycle factor* in $Q_{n,\ell}$, i.e., a collection of disjoint cycles which together visit all vertices of $Q_{n,\ell}$. In a second step, we use local modifications to join the cycles in the factor to a single Hamilton cycle. Essentially, this technique reduces the Hamiltonicity problem in $Q_{n,\ell}$ to proving that a suitably defined auxiliary graph is connected, which is much easier.

In fact, the predecessor paper [GJM⁺21] already proved the existence of a cycle factor in $Q_{n,\ell}$, but this construction does not seem to yield a factor that would be amenable to analysis. In this paper, we therefore construct another cycle factor in $Q_{n,\ell}$, based on modifying the aforementioned Greene-Kleitman SCD of Q_n by the lexical matchings introduced by Kierstead and Trotter [KT88]. The resulting cycle factor in $Q_{n,\ell}$ has a rich structure, in particular the number of cycles and their lengths can be described combinatorially.

The simplest way to join two cycles C and C' from this factor to a single cycle is to consider a 4-cycle F that shares exactly one edge with each of the cycles C and C' (the other two edges of F must then go between C and C'), and to take the symmetric difference of the edge sets of $C \cup C'$ and of F , yielding a single cycle $(C \cup C') \triangle F$ on the same vertex set as $C \cup C'$. We refer to such a cycle F as a *flipping 4-cycle*. For example, if we interpret the binary reflected Gray code Γ_n as a cycle in Q_n , we see that $\Gamma_{n+1} = (0\Gamma_n \cup 1\Gamma_n^R) \triangle F$ where F is the 4-cycle $F = 0^{n+1}, 010^{n-1}, 110^{n-1}, 10^n$. In addition to flipping 4-cycles, we also use flipping 6-cycles, which intersect with the two cycles to be joined in a slightly more complicated way, albeit with the same effect of joining them to a single cycle. The most technical aspect of this part of the proof is to ensure that all flipping cycles used are edge-disjoint, so that the joining operations do not interfere with each other.

To prove Theorem 3, we proceed by induction from dimension n to $n + 2$, treating the cases of even and odd n separately. We first specify a particular ordering of all chains of the Greene-Kleitman SCD, and then show that this ordering admits a matching that alternately joins the bottom or top vertices of any two consecutive chains in our ordering. In fact, there is a close relation between our proofs of Theorem 1 and 3: The aforementioned construction of a cycle factor in $Q_{n,\ell}$ is particularly nice for $\ell = m + 1$, i.e., for the case where we consider the entire cube. Specifically, in this case our cycle factor contains *all* chains from the Greene-Kleitman SCD. These cycles can be joined to a single Hamilton cycle in such a way, so as to give exactly the aforementioned Hamilton cycle constructed for proving Theorem 3.

1.4. Outline of this paper. In Section 2 we build up the necessary preliminaries. Specifically, we introduce various Catalan bijections, the Greene-Kleitman SCD, and lexical matchings. In Section 3 we describe our construction of a cycle factor in $Q_{n,\ell}$, and we analyze its structure in Section 4, identifying two types of cycles, called short and long cycles. In Section 5 we describe the flipping 4-cycles that we use for attaching the short cycles to the long cycles. In Section 6 we describe the flipping 6-cycles that we use for joining the long cycles to a Hamilton cycle in $Q_{n,\ell}$. In Section 7 we put together all ingredients for proving Theorem 1, reducing the Hamiltonicity problem to a connectivity problem in a suitably defined auxiliary graph. In Section 8 we present our proof of Theorem 3 and the corresponding loopless algorithm. This part can be read independently of the previous parts, though it is closely related.

2. PRELIMINARIES

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

2.1. Bitstrings, lattice paths, and rooted trees. For any string x and any integer $k \geq 0$, we let x^k denote the concatenation of k copies of 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 0-bit is interpreted as a \setminus -step that changes the current coordinate by $(+1, -1)$ and every 1-bit is interpreted as an $/$ -step that changes the current coordinate by $(+1, +1)$; see Figure 2.

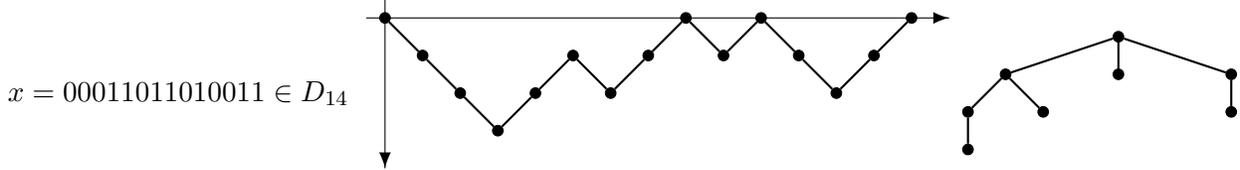


FIGURE 2. The correspondence between bitstrings (left), lattice paths (middle) and rooted trees (right).

Let D_{2k} denote the set of bitstrings with exactly k many 1s and k many 0s, such that in every prefix, the number of 0s is at least as large as the number of 1s. We also define $D := \bigcup_{k \geq 0} D_{2k}$. Note that $D_0 = \{\varepsilon\}$, where ε denotes the empty bitstring. In terms of lattice paths, D corresponds to so-called *Dyck paths* that never move above the line $y = 0$ and end on this line. If a lattice path x contains a substring $u \in D$, then we refer to this substring u as a *valley* in x . Any nonempty bitstring $x \in D$ can be written uniquely as $x = 0u1v$ and as $x = u'0v'1$ with $u, v, u', v' \in D$. We refer to this as the *left* and *right factorization* of x , respectively. The set D'_{2k} is defined similarly as D_{2k} , but we require that in exactly one prefix, the number of 0s is strictly smaller than the number of 1s. That is, the lattice paths corresponding to the bitstrings in D'_{2k} move above the line $y = 0$ exactly once.

An (*ordered*) *rooted tree* is a tree with a distinguished root vertex, and the children of each vertex have a specified left-to-right ordering. We think of a rooted tree as a tree embedded in the plane with the root on top, with downward edges leading from any vertex to its children, and the children appear in the specified left-to-right ordering. Using a standard Catalan bijection (see [Sta15]), every Dyck path $x \in D_{2k}$ can be interpreted as a rooted tree with k edges; see Figure 2. Given a rooted tree $x \in D$, we may rotate the tree to the right, yielding the tree \vec{x} ; see Figure 3. More precisely, \vec{x} is obtained from x by taking the leftmost child of the root of x as the new root. In terms of bitstrings, if x has the left factorization $x = 0u1v$ with $u, v \in D$, then \vec{x} has the right factorization $\vec{x} = u0v1$. Note that we have $\vec{D} = D$.

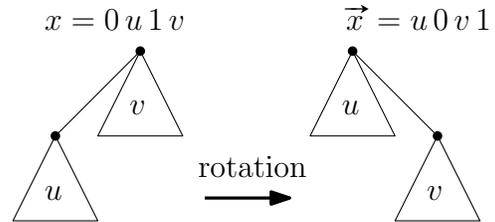


FIGURE 3. Definition of tree rotation.

2.2. The Greene-Kleitman SCD. We now describe Greene and Kleitman's [GK76] construction of an SCD in the n -cube; see Figure 4. For any vertex x of the n -cube, we interpret the 0s in x as opening brackets and the 1s as closing brackets. By matching closest pairs of opening and closing brackets in the natural way, the chain containing x is obtained by flipping the leftmost unmatched 0 to ascend the chain, or the rightmost unmatched 1 to descend the chain, until no more unmatched bits can be flipped. It is easy to see that this indeed yields an SCD of the n -cube for any $n \geq 1$. In the rest of the paper, we always work with this SCD due to Greene and Kleitman, and whenever referring to a chain, we mean a chain from this decomposition.

Each chain C of length h in Q_n can be encoded compactly as a string of length n over the alphabet $\{0, 1, *\}$ in the form

$$C = u_0 * u_1 * \cdots * u_{h-1} * u_h, \quad (1)$$

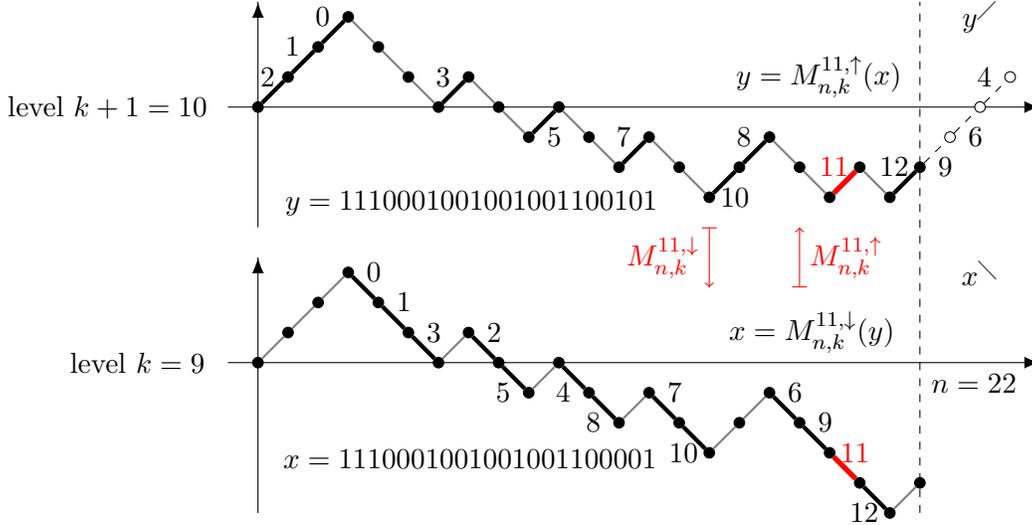


FIGURE 5. Definition of p -lexical matchings between levels 9 and 10 of Q_{22} , where steps flipped along the p -lexical edge are marked with p . Between those two levels, the vertex x is incident with p -lexical edges for each $p \in \{0, 1, \dots, 12\}$, and the vertex y is incident with p -lexical edges for each $p \in \{0, 1, \dots, 12\} \setminus \{4, 6, 9\}$.

and from right to left in each row. The partial mapping $M_{n,k}^{p,\uparrow}(x)$ is obtained by flipping the p th \setminus -step encountered in this fashion, where counting starts with $0, 1, \dots$, if this \setminus -step is part of the subpath x of $x \setminus$; otherwise x is left unmatched. Similarly, for any $x \in L_{n,k+1}$ we consider the lattice path $x \swarrow$ and scan it row-wise from top to bottom, and from left to right in each row. The partial mapping $M_{n,k}^{p,\downarrow}(x)$ is obtained by flipping the p th \swarrow -step encountered in this fashion if this \swarrow -step is part of the subpath x of $x \swarrow$; 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 , called the p -lexical matching, which we denote by $M_{n,k}^p$. We also define $M_n^p := \bigcup_{0 \leq k < n} M_{n,k}^p$, where we omit the index n whenever it is clear from the context. In the following, we will only ever use p -lexical edges for $p = 0, 1, 2$. For instance, it is well-known that taking the union of all 0-lexical edges, i.e., the set M^0 , yields exactly the Greene-Kleitman SCD [KT88]. This property is captured by the following lemma, together with several other explicit perfect matchings, consisting of $\{0, 1, 2\}$ -lexical edges between certain sets of vertices from our SCD; see Figure 6.

To state the lemma, for a set M of edges of Q_n and disjoint sets A, B of vertices, we let $M[A, B]$ denote the set of all edges of M between A and B . Moreover, for any vertex $x \in C_{h,i}^-$, $1 < i < h \leq n$, we consider the chain factorization $x = u_0 1 \cdots u_{i-2} 1 u_{i-1} 1 0 u_{i+1} \cdots 0 u_h$ with $u_0, \dots, u_h \in D$, and we define a neighbor $z(x)$ on the level below by

$$z(x) := \begin{cases} u_0 1 \cdots u_{i-2} 1 0 0 u_{i+1} \cdots 0 u_h & \text{if } u_{i-1} = \varepsilon, \\ u_0 1 \cdots u_{i-2} 1 0 v 0 w 1 0 u_{i+1} \cdots 0 u_h & \text{if } u_{i-1} = 0v1w \text{ with } v, w \in D. \end{cases}$$

Note that in the first case, $(x, z(x))$ is a 0-lexical edge, and in the second case, $(x, z(x))$ is a 2-lexical edge. The result of this operation can be written more compactly as

$$z(x) = u_0 1 \cdots u_{i-2} 1 0 \overrightarrow{u_{i-1}} 0 u_{i+1} \cdots 0 u_h. \quad (3)$$

Lemma 5. *For every $n \geq 3$, the following sets of edges $M[A, B]$ are perfect matchings in Q_n between the vertex sets A and B .*

- (i) $M^0[C_{h,i}, C_{h,i+1}]$ for every $0 \leq i < h \leq n$;
- (ii) $M^1[C_{1,0}^-, C_{1,1}^-]$, $M^1[C_{h,i}^+, C_{h+2,i+2}^-]$, and $M^1[C_{h,i}^+, C_{h+2,i}^-]$ for every $0 \leq i \leq h \leq n-2$;
- (iii) $Z^{02}[C_{h,i-1}^-, C_{h,i}^-]$ for every $1 < i < h \leq n$, where $Z^{02} := \{(x, z(x)) \mid x \in C_{h,i}^-\}$.

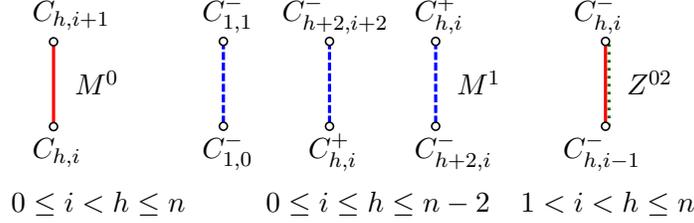


FIGURE 6. Perfect matchings described by Lemma 5. The $\{0, 1, 2\}$ -lexical edges are drawn solid, dashed, and dotted, respectively.

Proof. To prove (i), let $x \in C_{h,i}$ and consider the chain factorization (2) of x . By the definition of lexical matchings, the neighbor y of x on the level above reached via the 0-lexical edge is obtained by flipping the 0-bit after the valley u_i in x , so $y = u_0 1 \cdots u_i 1 u_{i+1} 0 u_{i+2} \cdots 0 u_h$. From this and (2) we conclude that $y \in C_{h,i+1}$ and that these edges reach all vertices in $C_{h,i+1}$.

To prove (ii), first consider the subcase $x \in C_{1,0}^-$ and the chain factorization $x = 0 u_1$ with $u_1 \in D$. Moreover, as $n \geq 3$, we have $u_1 \neq \varepsilon$, so we may consider the right factorization $u_1 = v 0 w 1$ with $v, w \in D$. By the definition of lexical matchings, the neighbor y of $x = 0 u_1 = 0 v 0 w 1$ on the level above reached via the 1-lexical edge is obtained by flipping the 0-bit after the valley v in x , so $y = 0 v 1 w 1 = u_0 1$ with $u_0 \in D$ defined by $\vec{u}_0 = u_1$. From this and (2) we conclude that $y \in C_{1,1}^-$ and that these edges reach all vertices in $C_{1,1}^-$.

We now consider the subcase $x \in C_{h,i}^+$ and the chain factorization (2) of x . We know that $u_i \neq \varepsilon$, so u_i has the right factorization $u_i = v 0 w 1$ with $v, w \in D$. The neighbor y of x on the level above reached via the 1-lexical edge is obtained by flipping the 0-bit after the valley v in x , so $y = u_0 1 \cdots u_{i-1} 1 v 1 w 1 0 u_{i+1} \cdots 0 u_h$. From this and (2) we conclude that $y \in C_{h+2,i+2}^-$ and that these edges reach all vertices in $C_{h+2,i+2}^-$. For the same x , let us now compute the neighbor y' of x on the level below reached via the 1-lexical edge. For this we consider the left factorization $u_i = 0 v' 1 w'$ with $v', w' \in D$. The vertex y' is reached by flipping the 1-bit after the valley v' in x , so $y' = u_0 1 \cdots u_{i-1} 1 0 v' 0 w' 0 u_{i+1} \cdots 0 u_h$. Similarly to before, we obtain that $y' \in C_{h+2,i}^-$ and that these edges reach all vertices in $C_{h+2,i}^-$, so the proof of part (ii) is complete.

To prove part (iii), note that (2) and (3) imply that $z(x) \in C_{h,i-1}^-$ and that $C_{h,i-1}^-$ is precisely the image of $C_{h,i}^-$ under the mapping z . \square

3. CYCLE FACTOR CONSTRUCTION

We now construct a cycle factor $\mathcal{C}_{n,\ell}$ in the graph $Q_{n,\ell}$, $n = 2m + 1$, i.e., in the subgraph of the n -cube induced by the middle 2ℓ levels. Throughout this and the following sections we consider fixed $m \geq 1$ and $2 \leq \ell \leq m + 1$. We construct the cycle factor incrementally, starting with chains from the Greene-Kleitman SCD and adding $\{0, 1, 2\}$ -lexical edges between certain sets of vertices, see Figure 7. In the following, when referring to a subgraph given by a set of edges, we mean the subgraph of $Q_{n,\ell}$ induced by those edges. Moreover, we say that a chain is *short* if its length is at most $2\ell - 3$, i.e., if it does not span all levels of $Q_{n,\ell}$.

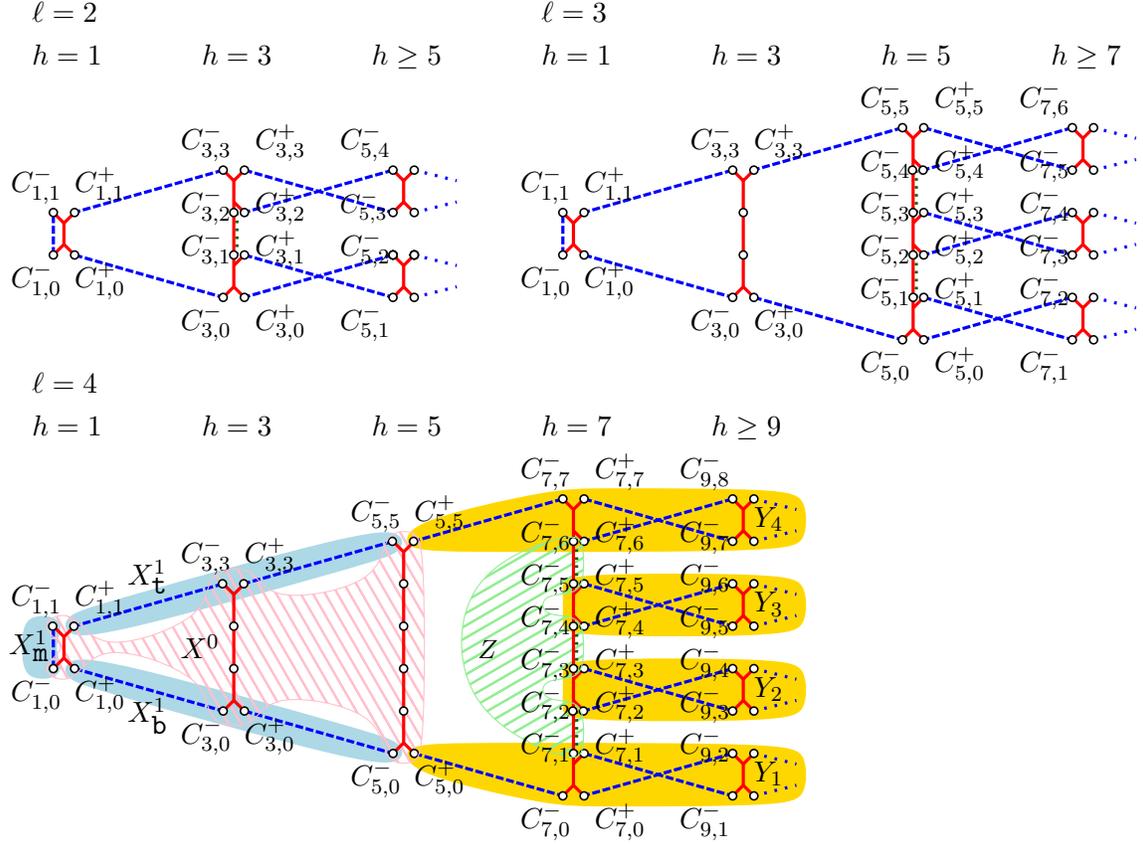


FIGURE 7. Illustration of the cycle factor $\mathcal{C}_{n,\ell}$ for $\ell = 2, 3, 4$. Each bullet represents an entire set of vertices, as specified in the figure, lines between them specify perfect matchings between these sets. The $\{0, 1, 2\}$ -lexical edges are drawn with solid, dashed, and dotted lines, respectively. In the bottom part, various sets of matching edges are highlighted.

Our construction starts by taking all those short chains, formally

$$X^0 := \bigcup_{0 \leq i < h \leq 2\ell - 3} M^0[C_{h,i}, C_{h,i+1}]; \quad (4a)$$

recall Lemma 5 (i). From Lemma 5 (ii) we know that 1-lexical edges perfectly match all bottom vertices of $[-+]$ -chains of length 1 with all top vertices of $[-+]$ -chains of length 1 along the edges

$$X_m^1 := M^1[C_{1,0}^-, C_{1,1}^-]. \quad (4b)$$

Furthermore, for $1 \leq h \leq 2\ell - 5$, 1-lexical edges perfectly match all top vertices of $[?+]$ -chains of length h with all top vertices of $[?-]$ -chains of length $h + 2$, and all bottom vertices of $[+?]$ -chains of length h with all bottom vertices of $[-?]$ -chains of length $h + 2$ along the edges

$$X_t^1 := \bigcup_{1 \leq h \leq 2\ell - 5} M^1[C_{h,h}^+, C_{h+2,h+2}^-], \quad X_b^1 := \bigcup_{1 \leq h \leq 2\ell - 5} M^1[C_{h,0}^+, C_{h+2,0}^-], \quad (4c)$$

respectively. Note that the only vertices of short chains that have degree 1 in the set

$$X := X^0 \cup X_m^1 \cup X_t^1 \cup X_b^1 \quad (4d)$$

are exactly the vertices of $C_{2\ell-3,2\ell-3}^+$ and $C_{2\ell-3,0}^+$; that is, the top vertices of $[?+]$ -chains of length $2\ell - 3$ and the bottom vertices of $[+?]$ -chains of length $2\ell - 3$.

Next, between every pair of consecutive levels of $Q_{n,\ell}$ we take all 0-lexical and 1-lexical edges that are not incident to a degree-2 vertex in X . Specifically, between these pairs of levels we take all 0-lexical edges from chains that are not short and all 1-lexical edges between chains that are not short. In addition, between the top two levels we take all 1-lexical edges between top vertices of $[?+]$ -chains of length $2\ell - 3$ and top vertices of $[?-]$ -chains of length $2\ell - 1$, and symmetrically, between the bottom two levels we take all 1-lexical edges between bottom vertices of $[+?]$ -chains of length $2\ell - 3$ and bottom vertices of $[-?]$ -chains of length $2\ell - 1$. Formally, these sets of edges are

$$Y_1 := Y'_1 \cup M^1[C_{2\ell-3,0}^+, C_{2\ell-1,0}^-], \quad Y_\ell := Y'_\ell \cup M^1[C_{2\ell-3,2\ell-3}^+, C_{2\ell-1,2\ell-1}^-], \quad Y_k := Y'_k \quad (5a)$$

for $1 < k < \ell$ where

$$Y'_k := \bigcup_{\substack{h \geq 2\ell-1 \\ i := (h - (2\ell-1))/2 + 2(k-1)}} M^0[C_{h,i}, C_{h,i+1}] \cup M^1[C_{h,i}^+, C_{h+2,i+2}^-] \cup M^1[C_{h,i+1}^+, C_{h+2,i+1}^-] \quad (5b)$$

for $1 \leq k \leq \ell$. Note that Y_1 and Y_ℓ contain all $\{0, 1\}$ -lexical edges between the bottom two levels or the top two levels of $Q_{n,\ell}$, respectively. We also define

$$Y := \bigcup_{1 \leq k \leq \ell} Y_k. \quad (5c)$$

As a consequence of these definitions and Lemma 5 (i) and (ii), the only vertices of $Q_{n,\ell}$ that have degree 1 in the set $X \cup Y$ are exactly the vertices of $C_{2\ell-1,i}^-$ for $1 \leq i \leq 2\ell - 2$. We thus add the edges

$$Z := \bigcup_{i=1,3,5,\dots,2\ell-3} Z^{02}[C_{2\ell-1,i}^-, C_{2\ell-1,i+1}^-] \quad (6)$$

defined in part (iii) of Lemma 5, which makes

$$\mathcal{C}_{n,\ell} := X \cup Y \cup Z \quad (7)$$

a cycle factor in the graph $Q_{n,\ell}$.

Note that if $\ell = 2$, then the sets X_t^1 and X_b^1 are empty and X^0 contains only chains of length 1; see the top left part of Figure 7. In the other extreme case $\ell = m + 1$, the set $Y \cup Z$ contains only a single path of length $n + 2$, namely the unique chain of length $2\ell - 1 = n$ with an additional 1-lexical edge from Y_1 and Y_ℓ attached on each side.

3.1. Comparison with previous constructions. Our cycle factor construction generalizes the construction for $\ell = 1$ presented in [Müt16, GMN18], which simply consisted in taking the union of all 0-lexical and 1-lexical edges between the middle two levels. It also generalizes the construction for $\ell = 2$ presented in [GJM⁺21], which also only used $\{0, 1, 2\}$ -lexical matchings. In fact, all these earlier papers actually used $\{m, m - 1, m - 2\}$ -lexical matching edges, but these are isomorphic to $\{0, 1, 2\}$ -lexical edges by reversing bitstrings. The earlier construction for $\ell = 2$ seemed rather arbitrary at the time, but now nicely fits into the general picture shown in Figure 7¹.

4. STRUCTURE OF CYCLES

In this section, we describe the structure of cycles in the factor $\mathcal{C}_{n,\ell}$ defined in (7) (where $n = 2m + 1$, $m \geq 1$, $2 \leq \ell \leq m + 1$). In particular, we give a combinatorial interpretation for certain vertices encountered along each cycle, allowing us to compute the number and lengths of some of the cycles. In the following, we call a cycle that contains a short $[--]$ -chain from the Greene-Kleitman SCD a *short cycle*, and any other cycle is called *long*. The key properties about short and long cycles are captured in Lemmas 6 and 11 below, respectively.

¹As the picture of this construction resembles a rocket, with the tip on the left and the boosters on the right, one might be tempted to consider this rocket science.

TABLE 1. Number of cycles in the factor $\mathcal{C}_{n,\ell}$, $n = 2m + 1$, for small values of m and ℓ . The first number in brackets counts short cycles of range $2\ell - 2$, the second number counts long cycles (of range 2ℓ). For fixed m and ℓ , the full distribution of cycles across ranges can be recovered by considering all other table entries on the same row. For instance, the full distribution of cycles of all ranges for $m = 8$ and $\ell = 9$ is [1430, 2002, 1638, 910, 350, 90, 14, 1] from 1430 short cycles of range 4 (through the middle 4 levels of the 17-cube), to 14 short cycles of range 16, to a single long cycle of range 18 (through all levels).

$m \setminus \ell$	1	2	3	4	5	6	7	8	9
1	1	1 [0,1]							
2	1	1 [0,1]	3 [2,1]						
3	2	1 [0,1]	6 [5,1]	10 [4,1]					
4	3	4 [0,4]	17 [14,3]	29 [14,1]	35 [6,1]				
5	6	6 [0,6]	46 [42,4]	93 [48,3]	118 [27,1]	126 [8,1]			
6	14	19 [0,19]	142 [132,10]	307 [165,10]	412 [110,5]	452 [44,1]	462 [10,1]		
7	34	49 [0,49]	446 [429,17]	1010 [572,9]	1438 [429,8]	1643 [208,5]	1704 [65,1]	1716 [12,1]	
8	95	150 [0,150]	1475 [1430,45]	3474 [2002,42]	5113 [1638,43]	6002 [910,22]	6337 [350,7]	6421 [90,1]	6435 [14,1]

For any $n = 2m + 1$, $m \geq 1$, we can compute the number of $[-]$ -chains of length h (by counting lattice paths using a reflection trick), and we thus obtain the number of short cycles of range $2r$ as

$$\binom{2m-1}{m-r+1} - \binom{2m-1}{m-r} = \frac{r-1}{m} \binom{2m}{m-r+1},$$

see Table 1. As we shall see, the structure of long cycles is more complicated in general, and we are not able to count them explicitly, with few exceptions: For $\ell = 1$ the number of all cycles of the factor is given by the number of plane trees with m edges (see [GMN18, Proposition 2]). For $\ell = 2$ the number of all (long) cycles is given by the number of plane trivalent trees with m internal vertices (see [GJM⁺21, Proposition 12]). For $\ell = m + 1$ there is exactly one long cycle, and so the total number of cycles in the factor is $\binom{2m-1}{m-1}$.

4.2. Long cycles. We now describe long cycles. First we show that each of the sets Y_k defined in (5) contains only paths, but no cycles, which will allow us to show that every long cycle has range 2ℓ and visits vertices from all 2ℓ levels. To describe the end vertices of the paths formed by the edge set Y_k , we use the following result shown in [GMN18, Proposition 2].

Lemma 8. *For any $d \geq 1$, the union of the 0- and 1-lexical matchings between levels d and $d + 1$ of Q_{2d} contains no cycles, but only paths, and the sets of first and last vertices of these paths are D_{2d} and D'_{2d} , respectively. Furthermore, for any path with first vertex $x \in D_{2d}$ and last vertex $y \in D'_{2d}$, if $x = u0v1$, with $u, v \in D$, is the right factorization of x , then $y = u10v$.*

Lemma 9. *For every $1 \leq k \leq \ell$, the edge set Y_k defined in (5) contains only paths, but no cycles. Furthermore, any path formed by the edges of Y_k has its first and last vertex in the following sets, and if its first vertex x has the form specified below, then its last vertex y has the form specified below:*

(i) *If $k = \ell$, then we have $x \in C_{2\ell-3,2\ell-3}^+$ and $y \in C_{2\ell-1,2\ell-2}^-$, and if*

$$x = u_0 1 \cdots u_{2\ell-4} 1 \underbrace{v 0 w 1}_{u_{2\ell-3}} \in C_{2\ell-3,2\ell-3}^+, \text{ then}$$

$$y = u_0 1 \cdots u_{2\ell-4} 1 v 1 0 w \in C_{2\ell-1,2\ell-2}^-.$$

(ii) If $k = 1$, then we have $x \in C_{2\ell-3,0}^+$ and $y \in C_{2\ell-1,1}^-$, and if

$$x = \underbrace{0v1w}_{u_0} 0u_1 \cdots 0u_{2\ell-3} \in C_{2\ell-3,0}^+, \text{ then}$$

$$y = v10w0u_1 \cdots 0u_{2\ell-3} \in C_{2\ell-1,1}^-.$$

(iii) If $1 < k < \ell$, then we have $x \in C_{2\ell-1,2k-1}^-$ and $y \in C_{2\ell-1,2k-2}^-$, and if

$$x = u_0 1 \cdots u_{2k-3} 1 u_{2k-2} 1 0 u_{2k} \cdots 0 u_{2\ell-1} \in C_{2\ell-1,2k-1}^-, \text{ then}$$

$$y = u_0 1 \cdots u_{2k-3} 1 0 u_{2k-2} 0 u_{2k} \cdots 0 u_{2\ell-1} \in C_{2\ell-1,2k-2}^-.$$

Proof. For any bitstring $x = x_1 \cdots x_n$ and any integers $\alpha \geq 0$ and $0 \leq \beta \leq n$ we define the mapping $\varrho_{\alpha,\beta}(x) := \varrho_{\alpha,\beta}(x_1 \cdots x_n) := 0^\alpha x_1 \cdots x_{n-\beta}$, i.e., this mapping prepends x with α many 0s, and removes the last β bits.

To prove (i), we define $\alpha := 2\ell - 3$ and we consider the edge set $\varrho_{\alpha,0}(Y_\ell)$ between levels d and $d+1$ of Q_{2d} , where $2d = n + \alpha$. From the definition of lexical matchings and (5), we see that $\varrho_{\alpha,0}(Y_\ell)$ is a set of 0- and 1-lexical matching edges in Q_{2d} . Moreover, as $\varrho_{\alpha,0}(C_{2\ell-3,2\ell-3}^+) \subseteq D_{2d}$ and $\varrho_{\alpha,0}(C_{2\ell-1,2\ell-2}^-) \subseteq D'_{2d}$, applying Lemma 8 in Q_{2d} shows that the edge set Y_ℓ contains no cycles and any path formed by those edges has its first and last vertex in the sets $C_{2\ell-3,2\ell-3}^+$ and $C_{2\ell-1,2\ell-2}^-$. Moreover, for a first vertex $x' := 0^\alpha x \in D_{2d}$ with x as in the lemma we have the right factorization $x' = 0^\alpha x = a 0 w 1$ with $a := 0^\alpha u_0 1 \cdots u_{2\ell-4} 1 v \in D$, so Lemma 8 shows that the last vertex of this path is indeed $y' = a 1 0 w = 0^\alpha y$ with y as above.

To prove (ii), consider the automorphism φ of Q_n that reverses and complements all bits. One can check that φ maps 0- and 1-lexical matchings between levels i and $i+1$ to 0- and 1-lexical matchings between levels $n-i$ and $n-i-1$ for all $0 \leq i < n$. Consequently, using (5) we obtain that $Y_1 = \varphi(Y_\ell)$. As moreover $C_{2\ell-3,0}^+ = \varphi(C_{2\ell-3,2\ell-3}^+)$ and $C_{2\ell-1,1}^- = \varphi(C_{2\ell-1,2\ell-2}^-)$ the claim follows from part (i).

To prove (iii), let $x \in C_{2\ell-1,2k-1}^-$ be as in the lemma. Moreover, let $b := 0 u_{2k} \cdots 0 u_{2\ell-1}$ be the suffix of x , let β be the length of b , and let $\alpha := 2k - 3$. Consider the vertex $z := u_0 1 \cdots u_{2k-3} 0 u_{2k-2} 1 b$, which is joined to x via a 1-lexical edge. This edge, however, is not in Y_k , as we have $z \in C_{2\ell-3,2k-3}^+$ from (2). Adding all those edges to Y_k yields a larger set of edges $Y' := Y_k \cup M^1[C_{2\ell-3,2k-3}^+, C_{2\ell-1,2k-1}^-]$. We now consider the edge set $\varrho_{\alpha,\beta}(Y')$ between levels d and $d+1$ of Q_{2d} , where $2d = n + \alpha - \beta$. Similarly to before, $\varrho_{\alpha,\beta}(Y')$ is a set of 0-lexical and 1-lexical matching edges in Q_{2d} . Moreover, as $\varrho_{\alpha,\beta}(C_{2\ell-3,2k-3}^+) \subseteq D_{2d}$ and $\varrho_{\alpha,\beta}(C_{2\ell-1,2k-2}^-) \subseteq D'_{2d}$, applying Lemma 8 in Q_{2d} shows that the edge set Y' , and hence Y_k , contains no cycles and any path formed by the edges of Y' has its first and last vertex in the sets $C_{2\ell-3,2k-3}^+$ and $C_{2\ell-1,2k-2}^-$. Moreover, for the first vertex $z' := \varrho_{\alpha,\beta}(z) \in D_{2d}$ with z from before we have the right factorization $z' = a 0 u_{2k-2} 1$ with $a := 0^\alpha u_0 1 \cdots u_{2k-3} \in D$, so Lemma 8 shows that the last vertex of this path is $y' := a 1 0 u_{2k-2}$, and as $y' = \varrho_{\alpha,\beta}(y)$ with y as in the lemma, the claim is proved. \square

The next lemma describes the effect of walking along a path in the set $Y \cup Z$ from its first to its last vertex.

Lemma 10. *The union of edges $Y \cup Z$ with Y and Z as defined in (5) and (6) contains only paths, but no cycles, and the sets of first and last vertices of these paths are $C_{2\ell-3,2\ell-3}^+$ and $C_{2\ell-3,0}^+$, respectively. Furthermore, for any path with first vertex*

$$x = u_0 1 u_1 \cdots 1 \underbrace{v 0 w 1}_{u_{2\ell-3}} \in C_{2\ell-3,2\ell-3}^+,$$

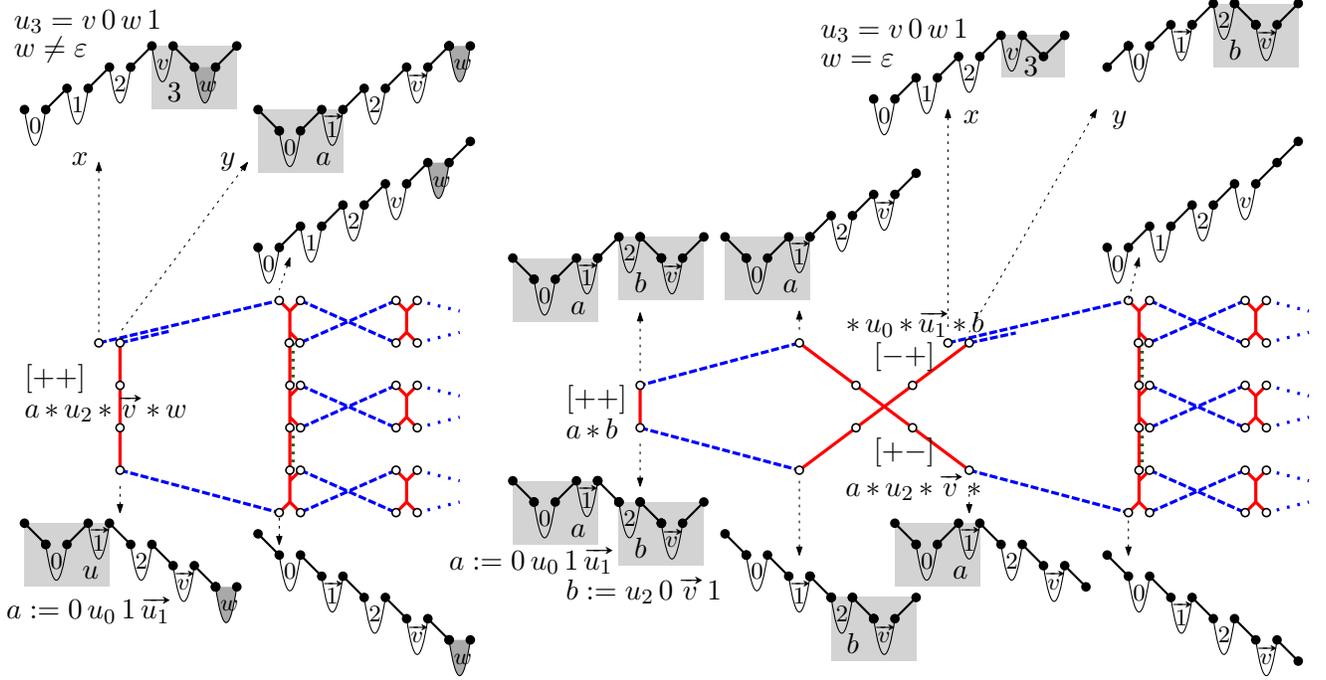


FIGURE 9. Structure of long cycles as described by Lemma 11, for case (i) ($w \neq \varepsilon$) on the left and for case (ii) ($w = \varepsilon$) on the right.

the corresponding last vertex is

$$y = 0 u_0 1 \vec{u}_1 0 u_2 0 \vec{u}_3 0 u_4 0 \cdots \vec{u}_{2\ell-5} 0 u_{2\ell-4} 0 \vec{v} 0 w \in C_{2\ell-3,0}^+.$$

Proof. As the edges in Z perfectly match the last vertices of paths formed by Y_k with the first vertices of paths formed by Y_{k-1} for all $2 < k \leq \ell$, and the last vertices of paths formed by Y_2 with the last vertices of paths formed by Y_1 , this follows by iterating Lemma 9 and by using (3), where the Z -edges cause the rotation of every second valley $u_1, u_3, \dots, u_{2\ell-5}, v$ in x . \square

The next lemma describes how the paths formed by the edges $Y \cup Z$ interact with the short chains in X that are not contained in a short cycle (recall the last part of Lemma 6) to form long cycles. We describe this interaction locally, in terms of how consecutive vertices from the set $C_{2\ell-3,2\ell-3}^+$ on a long cycle look like, which is enough for our purposes.

Lemma 11. *Each long cycle in the factor $\mathcal{C}_{n,\ell}$ defined in (7) has range 2ℓ and visits vertices from all 2ℓ levels. Moreover, let $x = u_0 1 u_1 \cdots 1 u_{2\ell-4} 1 \underbrace{v 0 w 1}_{u_{2\ell-3}} \in C_{2\ell-3,2\ell-3}^+$ be a vertex from the cycle and let y be the next vertex from $C_{2\ell-3,2\ell-3}^+$ on the cycle after x .*

- (i) *If $w \neq \varepsilon$, then we have $y = a 1 u_2 1 \vec{u}_3 1 u_4 1 \cdots \vec{u}_{2\ell-5} 1 u_{2\ell-4} 1 \vec{v} 1 w$, where $a := 0 u_0 1 \vec{u}_1$, and between x and y the cycle traverses the $[++]$ -chain $a * u_2 * \vec{u}_3 * u_4 * \cdots * \vec{u}_{2\ell-5} * u_{2\ell-4} * \vec{v} * w$ of length $2\ell - 3$.*
- (ii) *If $w = \varepsilon$, then we have $y = 1 u_0 1 \vec{u}_1 1 u_2 1 \vec{u}_3 1 \cdots \vec{u}_{2\ell-5} 1 b$, where $b := u_{2\ell-4} 0 \vec{v} 1$, and between x and y the cycle traverses the $[+-]$ -chain $a * u_2 * \vec{u}_3 * \cdots * \vec{u}_{2\ell-5} * u_{2\ell-4} * \vec{v} * w$ of length $2\ell - 3$, the $[++]$ -chain $a * u_2 * \vec{u}_3 * \cdots * \vec{u}_{2\ell-5} * b$ of length $2\ell - 5$, and the $[-+]$ -chain $* u_0 * \vec{u}_1 * u_2 * \cdots * \vec{u}_{2\ell-5} * b$ of length $2\ell - 3$, where $a := 0 u_0 1 \vec{u}_1$, plus two additional edges connecting these chains.*

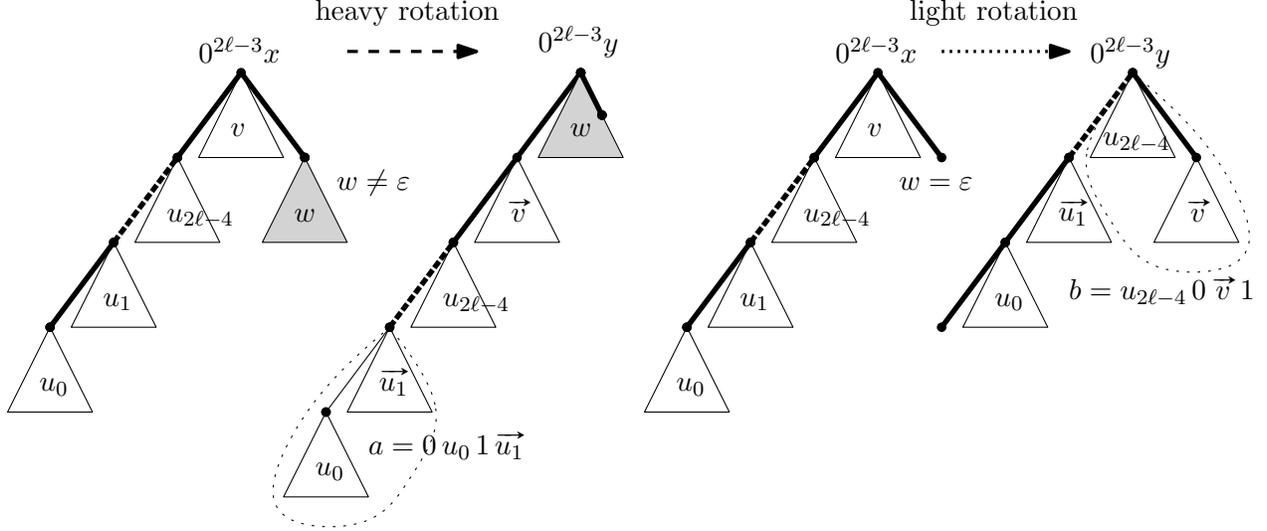


FIGURE 10. Interpretation of the transformations stated in parts (i) and (ii) of Lemma 11 in terms of rooted trees, called heavy and light rotations, shown on the left and right hand side, respectively. The spine edges are drawn bold.

Proof. The first part of the lemma follows immediately from the last part of Lemma 6 and the first part of Lemma 10. To prove parts (i) and (ii) consider Figure 9, and observe that a path in $Y \cup Z$ starting from x as in the lemma has last vertex as specified by Lemma 10, and then the cycle continues with a path in X on short chains as claimed until it arrives at the claimed vertex y (use (4) and the definition of 1-lexical matchings). \square

It is illuminating to interpret the transformations captured by parts (i) and (ii) of Lemma 11 in terms of rooted trees; see Figure 10. If we prepend $0^{2\ell-3}$ to the vertices x and y as in the lemma, then the resulting bitstrings $0^{2\ell-3}C_{2\ell-3,2\ell-3}^+ =: \mathcal{T}_{n,\ell}$ can be interpreted as rooted trees on $m + \ell - 1$ edges with a root of degree at least 2 and the leftmost leaf in depth at least $2\ell - 3$. We refer to the operations $0^{2\ell-3}x \mapsto 0^{2\ell-3}y$ with x and y as in parts (i) or (ii) of Lemma 11 as a *heavy rotation*, or a *light rotation*, respectively. Moreover, we refer to the subtree $0^{2\ell-3}1^{2\ell-3}01$ of any tree from $\mathcal{T}_{n,\ell}$ as its *spine*. Combinatorially, a heavy rotation is an inverse rotation at the root, plus a rotation of every second subtree attached to the spine. A light rotation moves all subtrees attached to the spine one spine vertex to the right, and it also applies a rotation to every second such subtree. These operations define an equivalence relation on $\mathcal{T}_{n,\ell}$ whose equivalence classes correspond to long cycles. Unfortunately, in general there seems to be no ‘nice’ combinatorial interpretation of these equivalence classes, unless in some special cases like $\ell \in \{1, 2, m + 1\}$; recall the remarks from the end of Section 4.1. The number of long cycles determined experimentally for small values of m and ℓ can be found in Table 1.

5. FLIPPING 4-CYCLES

In this section we describe how to modify the cycle factor $\mathcal{C}_{n,\ell}$ (where $n = 2m + 1$, $m \geq 1$, $2 \leq \ell \leq m + 1$), so that every short cycle is joined to some long cycle. The remaining task, solved in the next sections, will then be to join the long cycles to a single Hamilton cycle. The modifications via 4-cycles exploit relations between pairs of short chains of the Greene-Kleitmann SCD that were first used by Griggs, Killian, and Savage [GKS04], and by Killian, Ruskey, Savage, and Weston [KRSW04] with the purpose of constructing symmetric Venn diagrams.

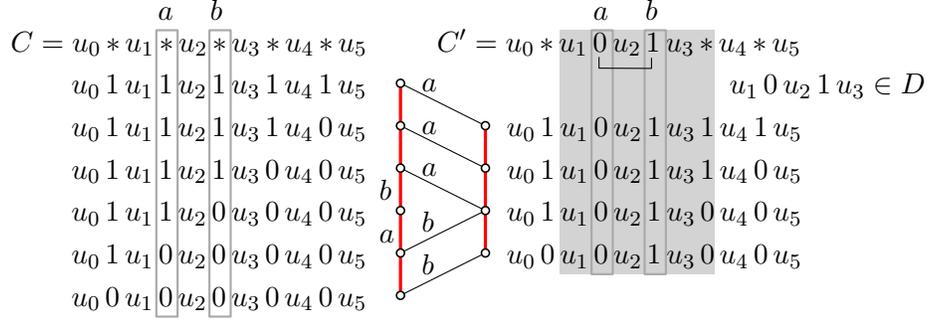


FIGURE 11. Flipping 4-cycles between a chain C and one of its children C' . The labels a and b along edges denote positions of flipped bits.

A *flipping 4-cycle* between two vertex-disjoint paths P, P' is a 4-cycle F that shares exactly one edge with each of the two paths. Note that the symmetric difference of the edge sets $(P \cup P') \triangle F$ gives path Q, Q' on the same vertex set with flipped end vertices: Specifically, if P, P' are $xy, x'y'$ -paths, then Q, Q' are $xy', x'y$ -paths. Thus, if P, P' are subpaths of two cycles C, C' , then $(C \cup C') \triangle F$ is a single cycle.

Given a chain C of length h , we say that a chain C' of length $h - 2$ is a *child* of C , if C' is obtained from C by replacing two consecutive $*$ s in C by 0 and 1, respectively; see Figure 11. We refer to this as *matching two $*$ s in C* . We also call C a *parent* of C' . Note that C has exactly $h - 1$ many children, obtained by matching any two consecutive of the h many $*$ s in total. Moreover, the number of parents of C' is given by the number of outermost matched 01-pairs in C' , which can be up to $(n - (h - 2))/2$. A straightforward verification (see Figure 11 for an example) yields the following lemma.

Lemma 12. *For every chain C of length h and its child C' obtained by matching two $*$ s at positions a, b , there are exactly $h - 2$ many flipping 4-cycles between C and C' , each using a distinct edge of C and C' , except the two consecutive edges of C that flip the coordinates a and b .*

The preceding lemma provides us enough options for selecting multiple edge-disjoint flipping 4-cycles as follows.

Lemma 13. *For every chain C of length $h \geq 5$ and any edge f of C , there are $h - 1$ edge-disjoint flipping 4-cycles between C and each of its $h - 1$ children that all avoid the edge f .*

We can think of f as a forbidden edge that will be used to connect C further to one of its parent chains.

Proof. We may assume w.l.o.g. that $C = *^h$, as any other chain of length h is obtained from this chain by inserting valleys between and around the $*$ s; recall (1). We label the edges on C by $1, \dots, h$ from bottom to top, and we let $p \in \{1, \dots, h\}$ be the index of the edge f to be avoided. We label all $h - 1$ children of C from 1 to $h - 1$, such that the i th child is obtained by matching the i th and $(i + 1)$ st $*$ in C . To the i th child, $1 \leq i \leq h - 1$, we assign the edge

$$e(i) := \begin{cases} \mu(i + 2, h) & \text{if } \mu(i + 2, h) < p, \\ \mu(i + 3, h) & \text{otherwise,} \end{cases}$$

on C , where $\mu(x, h)$ is the representative of the residue class of x modulo h from the set $\{1, \dots, h\}$. Since $e(i) \notin \{i, i + 1\}$ by the definition of $e(i)$ and the assumption $h \geq 5$, by Lemma 12 there is a flipping 4-cycle between C and its i th child that uses the edge $e(i)$ on C . Moreover, as the

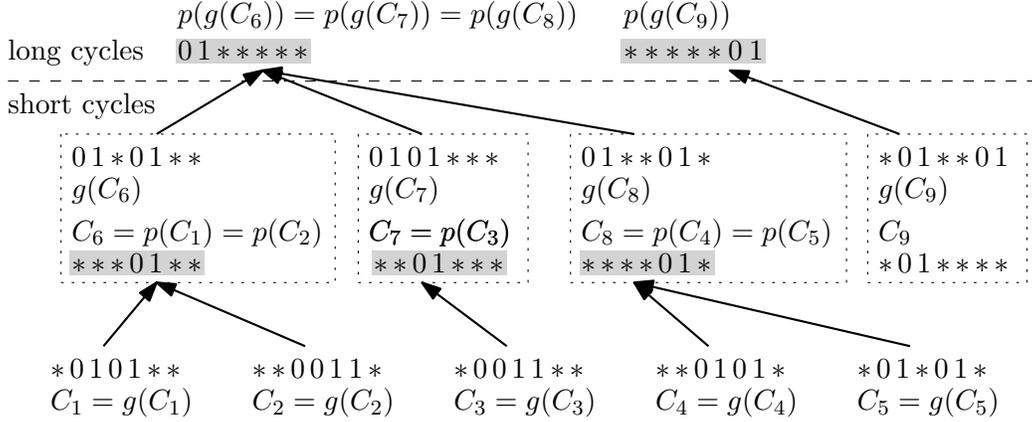


FIGURE 12. For $n = 2m + 1 = 7$ and $\ell = 4$ we have 9 short cycles represented by the $[-]$ -chains C_1, \dots, C_9 of length 3 or $5 = 2\ell - 3$. The dotted boxes indicate chains C_i with $C_i \neq g(C_i)$, but with C_i and $g(C_i)$ on the same short cycle. Arrows connect $g(C_i)$ with $p(g(C_i))$. The top two chains are contained in a long cycle (in this case, in the same cycle). The roots of the trees mentioned in the proof of Lemma 14 are highlighted in gray.

mapping $e: \{1, \dots, h-1\} \rightarrow \{1, \dots, p-1, p+1, \dots, h\}$ is a bijection, these flipping 4-cycles for all $1 \leq i \leq h-1$ are edge-disjoint, and they all avoid f . \square

Recall that by Lemma 6, short cycles correspond to $[-]$ -chains of length $3 \leq h \leq 2\ell - 3$. For each short cycle represented by such a chain C of length h , we now specify a chain $g(C)$ on the same cycle, called a *gluing chain*, and we also specify a parent chain $p(g(C))$ of $g(C)$ that belongs to a cycle of a bigger range; that is, a short cycle with a $[-]$ -chain of length $h+2$ or a long cycle. These definitions are illustrated in Figure 12. Specifically, for a short cycle represented by a $[-]$ -chain $C = *u_1 * \dots * u_{h-1} *$ of length $3 \leq h \leq 2\ell - 3$ we define

$$g(C) := \begin{cases} C & \text{if } h < 2\ell - 3, \\ a * u_3 * \dots * u_{h-1} * & \text{if } h = 2\ell - 3 \text{ and } u_1 = \varepsilon, \\ * u_1 * \dots * u_{h-3} * b & \text{if } h = 2\ell - 3 \text{ and } u_1 \neq \varepsilon, \end{cases} \quad (8)$$

where $a := 0u_1 1u_2 \in D$ and $b := u_{h-2} 0u_{h-1} 1 \in D$. Note that in the second case, $g(C)$ is a $[+-]$ -chain of length $h-2$, and in the third case, $g(C)$ is a $[-+]$ -chain of length $h-2$. By Lemma 6, in each case the chain $g(C)$ belongs to the same short cycle as C . To define $p(g(C))$, let $0 < i < h$ be the smallest index such that $u_i \neq \varepsilon$. We let $p(g(C))$ be the parent of $g(C)$ obtained by replacing the leftmost matched 01-pair in u_i by two *s. That is, given the left factorization $u_i = 0v 1w$ with $v, w \in D$, then $p(g(C))$ is obtained from $g(C)$ by replacing u_i with $*v * w$. Observe that in all three cases of (8), the chain $p(g(C))$ has the same type as $g(C)$ ($[-]$, $[+-]$, or $[-+]$, respectively). As a consequence of Lemmas 6 and 11, the chain $p(g(C))$ therefore belongs to a short cycle of bigger range if $h < 2\ell - 3$ and to a long cycle if $h = 2\ell - 3$.

We now show that flipping 4-cycles between all gluing chains and their selected parents can be chosen to be pairwise edge-disjoint.

Lemma 14. *There is a set $\mathcal{F}_{n,\ell}$ of pairwise edge-disjoint flipping 4-cycles, each between $g(C)$ and $p(g(C))$ for all short cycles C of the cycle factor $\mathcal{C}_{n,\ell}$ defined in (7).*

In this lemma we identify a short cycle C with its corresponding $[-]$ -chain of length $3 \leq h \leq 2\ell - 3$.

Proof. We consider the directed graph on all short chains as nodes and arcs from $g(C)$ to $p(g(C))$ for all short cycles C , and in this graph we consider all nontrivial components; see Figure 12. As the chain $p(g(C))$ is longer than $g(C)$, these components are trees that are oriented towards a set of roots, and by the definition of the mappings g and p , these roots are $[- -]$ -, $[+ -]$ - and $[- +]$ -chains of length $2\ell - 3$. We select one flipping 4-cycle into $\mathcal{F}_{n,\ell}$ for every arc in each of these trees, and the selection is done by processing each tree independently, starting at the root and descending towards its leaves, along parent-child pairs of chains. We show that this can be done so that all selected 4-cycles are pairwise edge-disjoint. Clearly, these 4-cycles can only share an edge on a chain $g(C)$ or $p(g(C))$ for some short cycle C .

First, observe that $g(C)$ has length 1 only if $\ell = 3$ and C has length $h = 3 = 2\ell - 3$. In this case the arc from $g(C)$ to $p(g(C))$ is an entire tree, so we can take the single flipping 4-cycle between $g(C)$ and $p(g(C))$, which exists by Lemma 12, into $\mathcal{F}_{n,\ell}$.

For the rest of the proof we assume that $\ell \geq 4$, so the nodes of all trees are chains of length at least 3, and all interior nodes are chains of length at least 5. Let C denote the current chain, which is the root of some currently unprocessed subtree, and let $h \geq 5$ be its length. By Lemma 13, we may select edge-disjoint flipping 4-cycles to all children of C in the tree, one for each incoming arc to C , all avoiding the edge f of C that was previously chosen for a flipping 4-cycle between C and $p(C)$. We add to the set $\mathcal{F}_{n,\ell}$ all those 4-cycles, and we proceed recursively in each subtree below C . \square

6. FLIPPING 6-CYCLES

In this section we define certain 6-cycles between the top two levels in $Q_{n,\ell}$ (where $n = 2m + 1$, $m \geq 1$, $2 \leq \ell \leq m + 1$); that is, between levels $m + \ell - 1$ and $m + \ell$ in Q_n , that can be used to join pairs of long cycles from the cycle factor $\mathcal{C}_{n,\ell}$ defined in (7). Recall that by Lemma 9, in the top two levels the cycle factor $\mathcal{C}_{n,\ell}$ consists of paths formed by the edges in Y_ℓ , and these paths can be identified by their first vertices $C_{2\ell-3,2\ell-3}^+$. In the following, whenever referring to paths in Y_ℓ we mean maximal paths, i.e., the components formed by the edges in Y_ℓ (and not proper subpaths of these).

Recall that a flipping 4-cycle for a pair of paths has one edge in common with each of these paths, and taking the symmetric difference of the edge sets of the paths and the 4-cycle yields two paths on the same vertex set with flipped end vertices. A flipping 6-cycle works very similarly, albeit its intersection pattern with the two paths is more complicated; see Figure 13. Specifically, such a flipping 6-cycle shares *two* non-incident edges with one of the paths, and one edge with the other path. The precise conditions are stated in Lemma 15 below. Let us emphasize that the flipping

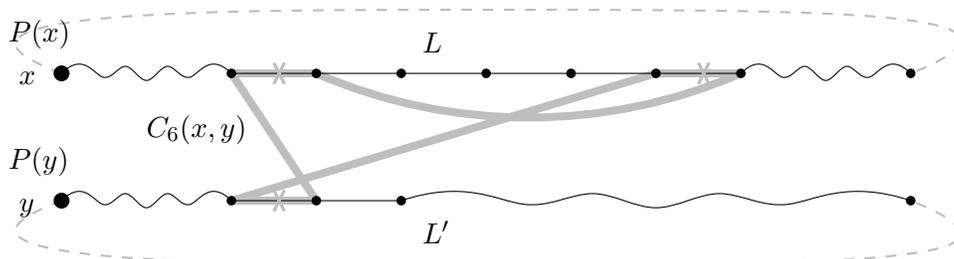


FIGURE 13. Two long cycles L, L' from the cycle factor $\mathcal{C}_{n,\ell}$ joined by taking the symmetric difference with a flipping 6-cycle. The paths $P(x)$ and $P(y)$ from the set Y_ℓ (solid black) lying on the two long cycles traverse the 6-cycle $C_6(x,y)$ (solid gray) as shown. The symmetric difference yields paths $P'(x)$ and $P'(y)$ that have flipped end vertices.

6-cycles we use are by definition edge-disjoint with all flipping 4-cycles considered in the previous section, as the 4-cycles are all between levels $m - \ell + 2$ and $m + \ell - 1$ of Q_n , so none of them uses any edges from the top two levels.

We say that two vertices $x, y \in C_{2\ell-3, 2\ell-3}^+$ form a *flippable pair* (x, y) , if $0^{2\ell-3}x$ and $0^{2\ell-3}y$ have the form

$$\begin{aligned} 0^{2\ell-3}x &= u'_0 0 u_1 \cdots 0 u_d 0 u_{d+1} 0 1 1 v_d 1 \cdots v_2 1 v_1 1, \\ 0^{2\ell-3}y &= u'_0 0 u_1 \cdots 0 u_d 0 u_{d+1} 1 0 1 v_d 1 \cdots v_2 1 v_1 1 \end{aligned} \quad (9)$$

with $d \geq 0$ and $u'_0, u_1, \dots, u_{d+1}, v_1, \dots, v_d \in D$.

Recall from Section 4 that $0^{2\ell-3}x$ and $0^{2\ell-3}y$ can be viewed as rooted trees from $\mathcal{T}_{n,\ell}$, where all but one spine edge is contained in the subtree u'_0 . If (x, y) is a flippable pair, then the tree $0^{2\ell-3}y$ is obtained from the tree $0^{2\ell-3}x$ by moving a pending edge from a vertex in the rightmost subtree to its predecessor. Specifically, the pending edge (q, r) must form the rightmost subtree of a vertex q in the rightmost subtree of $0^{2\ell-3}x$, and this edge is removed from q and reattached to the predecessor p of q to become the subtree directly right of the edge (p, q) . We refer to this as a *pull operation*; see Figure 14.

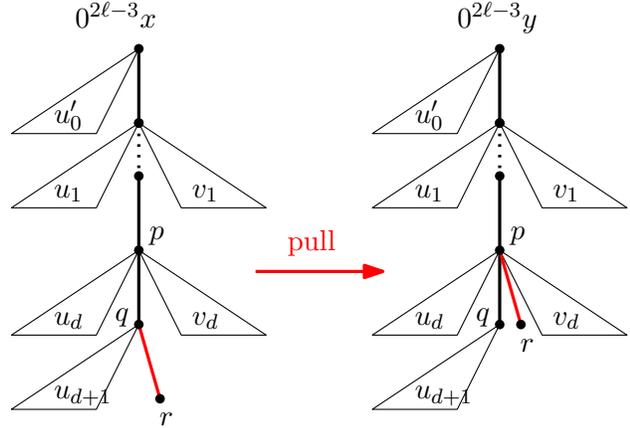


FIGURE 14. The pull operation on rooted trees.

Any 6-cycle between the top two levels of $Q_{n,\ell}$ can be uniquely encoded as a string of length n over $\{0, 1, *\}$ with $m + \ell - 2$ many 1s, $m - \ell$ 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 $\{0, 1\}$. For any flippable pair (x, y) as in (9), let u_0 be such that $u'_0 = 0^{2\ell-3}u_0$, and define the *flipping 6-cycle*

$$C_6(x, y) := u_0 1 u_1 \cdots 1 u_d * u_{d+1} ** 1 0 v_d 0 \cdots v_2 0 v_1. \quad (10)$$

Moreover, we let $\mathcal{S}_{n,\ell}$ denote the set of all those 6-cycles, obtained as the union of $C_6(x, y)$ over all flippable pairs (x, y) .

The following result was proved in [GMN18, Proposition 3]. For any $x \in C_{2\ell-3, 2\ell-3}^+$, we write $P(x)$ for the path from Y_ℓ that starts at the vertex x .

Lemma 15. *The 6-cycles $C_6(x, y) \in \mathcal{S}_{n,\ell}$ 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 $P(x)$ and $P(y)$ 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 $P(y)$, connecting 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.*

7. PROOF OF THEOREM 1

With Lemmas 11, 14 and 15 in hand, we are now ready to prove Theorem 1. Recall that $\mathcal{T}_{n,\ell} = 0^{2\ell-3}C_{2\ell-3,2\ell-3}^+$ is the set of rooted trees on $m + \ell - 1$ edges with a root of degree at least 2 and the leftmost leaf in depth at least $2\ell - 3$. Each $x \in \mathcal{T}_{n,\ell}$ has the form $x = 0^{2\ell-3}u_0 1 \cdots u_{2\ell-4} 1 u_{2\ell-3}$ with $u_0, \dots, u_{2\ell-3} \in D$, $u_{2\ell-3} \neq \varepsilon$. Considering the right factorization $u_{2\ell-3} = v 0 w 1$ with $v, w \in D$, we say that x is *right-empty* if $w = \varepsilon$, and we say that x is *right-full* if $u_0 = \cdots = u_{2\ell-4} = v = \varepsilon$ (and hence $w \neq \varepsilon$). In terms of trees, w is the subtree rooted at the rightmost child of the root of x .

Proof of Theorem 1. The case $\ell = 1$ of the theorem was proved in [Müt16], so we now assume that $\ell \geq 2$. Let $m \geq 1$, $n := 2m + 1$, $2 \leq \ell \leq m + 1$, and consider the subgraph $Q_{n,\ell}$ of the n -cube induced by the middle 2ℓ levels.

Let $\mathcal{C}_{n,\ell}$ be the cycle factor in the graph $Q_{n,\ell}$ defined in (7), let $\mathcal{F}_{n,\ell}$ be the set of flipping 4-cycles from Lemma 14, and let $\mathcal{S}_{n,\ell}$ be the set of flipping 6-cycles defined in (10). By the choice of $\mathcal{F}_{n,\ell}$ in Lemma 14, the symmetric difference $\mathcal{C}_{n,\ell} \Delta \mathcal{F}_{n,\ell}$ of the edge sets of $\mathcal{C}_{n,\ell}$ and $\mathcal{F}_{n,\ell}$ is a cycle factor with the property that every cycle traverses a path from the set Y_ℓ of edges between the top two levels of $Q_{n,\ell}$ (levels $m + \ell - 1$ and $m + \ell$ in Q_n), so to complete the proof it is enough to show that long cycles can be joined to a single cycle via flipping 6-cycles from $\mathcal{S}_{n,\ell}$.

Consider two long cycles L, L' containing paths $P(x), P(y) \in Y_\ell$ with first vertices $x, y \in C_{2\ell-3,2\ell-3}^+$, respectively, such that (x, y) is a flippable pair. By Lemma 15 (i), $(L \cup L') \Delta C_6(x, y)$ forms a single cycle on the same vertex set as $L \cup L'$, i.e., this joining operation reduces the number of long cycles in the factor by one; see Figure 13. Recall that in terms of rooted trees, $0^{2\ell-3}y \in \mathcal{T}_{n,\ell}$ is obtained from $0^{2\ell-3}x \in \mathcal{T}_{n,\ell}$ by a pull operation; see Figure 14. We repeat this joining operation until all long cycles in $\mathcal{C}_{n,\ell}$ are joined to a single cycle. For this purpose we define an auxiliary graph $\mathcal{H}_{n,\ell}$ whose nodes are the equivalence classes of trees from $\mathcal{T}_{n,\ell}$ under heavy and light rotations; see Figure 10. By Lemma 11, the nodes of this graph correspond to maximal sets of paths from Y_ℓ that lie on the same long cycle. Moreover, for the edge set of $\mathcal{H}_{n,\ell}$ we take all pairs of sets that contain the two trees that form a flippable pair (differing by a pull operation).

To complete the proof of Theorem 1, it therefore suffices to prove that the auxiliary graph $\mathcal{H}_{n,\ell}$ is connected. Indeed, if $\mathcal{H}_{n,\ell}$ is connected, then we can pick a spanning tree in $\mathcal{H}_{n,\ell}$ corresponding to a collection of flipping 6-cycles $\mathcal{S}'_{n,\ell} \subseteq \mathcal{S}_{n,\ell}$, such that the symmetric difference $\mathcal{C}_{n,\ell} \Delta (\mathcal{F}_{n,\ell} \cup \mathcal{S}'_{n,\ell})$ forms a Hamilton cycle in $Q_{n,\ell}$. Of crucial importance here are properties (ii) and (iii) in Lemma 15, which ensure that whatever subset of flipping 6-cycles we use in this joining process, they will not interfere with each other, guaranteeing that each flipping 6-cycle indeed reduces the number of long cycles by one, as desired. Recall also that flipping 4-cycles and flipping 6-cycles are edge-disjoint by definition, so they do not interfere with each other either.

At this point we reduced the problem of proving that $Q_{n,\ell}$ has a Hamilton cycle to showing that the auxiliary graph $\mathcal{H}_{n,\ell}$ is connected, which is much easier. Indeed, all we need to show is that any rooted tree from $\mathcal{T}_{n,\ell}$ can be transformed into any other tree from $\mathcal{T}_{n,\ell}$ by a sequence of heavy rotations, light rotations, pulls and their inverse operations (actually, we shall only use light rotations and pulls in our proof). Recall that heavy and light rotations correspond to following the same long cycle, and a pull corresponds to a joining operation. For this we show that any rooted tree $x \in \mathcal{T}_{n,\ell}$ can be transformed into the special tree $s := 0^{2\ell-3} 1^{2\ell-3} 0 0 1 \cdots 0 1 0 \in \mathcal{T}_{n,\ell}$, i.e., a right-full tree with a star rooted at the center as the rightmost child of the root. To achieve this, we distinguish three cases; see Figure 15.

- (a) x is *right-full*. We pull edges within the subtree rooted at the rightmost child of the root until this subtree is a star rooted at the center, which produces the desired tree s .

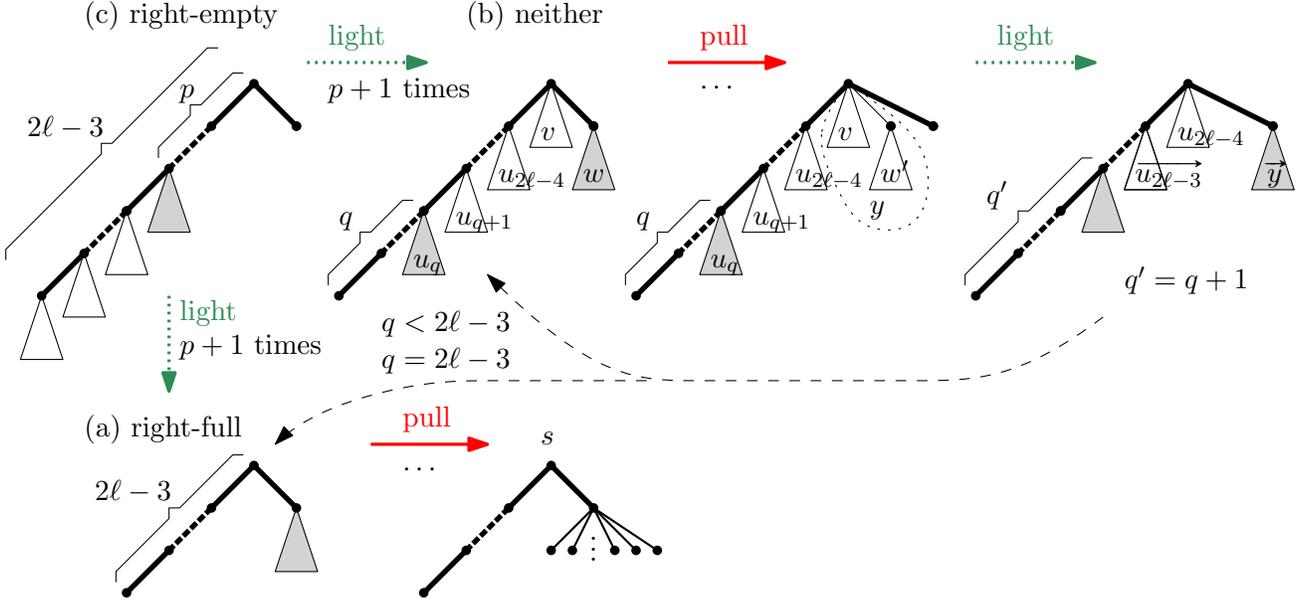


FIGURE 15. Transformation of trees from $\mathcal{T}_{n,\ell}$ into the tree s in the proof of Theorem 1. Gray subtrees are nonempty, white subtrees are possibly empty. Spine edges are drawn bold.

- (b) x is neither right-full nor right-empty. Let q be the distance between the left end vertex of the spine and the closest nonempty subtree attached to the spine. That is, q is the smallest index i such that $u_i \neq \varepsilon$ if such u_i exists, and $q = 2\ell - 3$ if $u_0 = \dots = u_{2\ell-4} = \varepsilon$ (as x is not right-full, $v \neq \varepsilon$ in this case). First, we repeatedly pull the single edge from the rightmost leaf of x all the way towards the root, arriving at a right-empty tree. After a single light rotation we obtain a tree that is not right-empty and that is either right-full (if $q = 2\ell - 3$), and then we conclude the argument as in case (a), or in which the distance q' between the left end vertex of the spine and the closest nonempty subtree has increased by 1 (if $q < 2\ell - 3$), i.e., we have $q' = q + 1$. Repeating the argument from case (b) hence terminates after at most $2\ell - 2$ repetitions.
- (c) x is right-empty. Let p be the distance between the root and the closest nonempty subtree attached to the spine. After exactly $p + 1$ light rotations, we are in case (a) or (b).

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

In our construction of the Hamilton cycle $\mathcal{C}_{n,\ell} \triangleq (\mathcal{F}_{n,\ell} \cup \mathcal{S}'_{n,\ell})$ described in the previous proof, the choice of the set of flipping 6-cycles $\mathcal{S}'_{n,\ell}$ depends on the global structure of long cycles. As explained in Section 4.2, our Lemma 11 does not provide any information about the global structure of these cycles, which is the main obstacle that prevents us from translating our construction to a polynomial-time algorithm for computing the next vertex on the Hamilton cycle, i.e., to an algorithm that only has local information about the current vertex in deciding which bit to flip next. Without explicitly constructing all long cycles, which may take time and space that are exponential in n , we do not even know how many 6-cycles should be selected into the set $\mathcal{S}'_{n,\ell}$.

8. PROOF OF THEOREM 3

For any chain C , we let $|C|$ denotes its length, i.e., the number of $*$ s in C . For any chain C with $|C| \geq 2$, we let $f(C)$ and $\ell(C)$, respectively, denote the chains obtained by replacing the first two $*$ s or the last two $*$ s in C by 0 and 1. Note that if $|C| \geq 2$, then we have $f(\ell(*C*)) = \ell(f(*C*))$. For

any chain C , we denote its bottom end vertex by $b(C)$ and its top end vertex by $t(C)$. Recall that $b(C)$ and $t(C)$ are obtained from C by replacing all *s by 0s or 1s, respectively.

Our goal is to order the chains of the Greene-Kleitman SCD in Q_n , $n \geq 2$, so that any consecutive pair of chains can be joined, alternatingly at their top ends or bottom ends, to a Hamilton cycle. Formally, given an SCD in Q_n , a *cycle ordering* is a sequence C_1, \dots, C_k of the chains of this SCD for which there exists a set of edges E of Q_n such that $C_1 \cup C_2 \cup \dots \cup C_k \cup E$ is a Hamilton cycle that traverses the chains in this order, either from bottom to top or vice versa, and moreover if $|C_i| = 0$ for some $i \in \{1, \dots, k\}$ then the two edges from E incident with this one-vertex chain must have their other end vertices in distinct levels. The following simple but powerful lemma shows that the direction in which each chain is traversed along the Hamilton cycle (upwards or downwards) is determined only by the chain length.

Lemma 16. *Let Λ_n be a cycle ordering of chains of an SCD in Q_n , $n \geq 2$. Then in the corresponding Hamilton cycle, any two chains C and C' with $|C| \equiv |C'| \pmod{4}$ are traversed in the same direction.*

This lemma holds for arbitrary SCDs, in particular for the SCD arising from the Greene-Kleitman construction.

Proof. This is a consequence of the following observations: After traversing a chain C with $|C| \geq 2$, the next chain in the cycle ordering has either length $|C| - 2$ or $|C| + 2$, and is traversed in the opposite direction in both cases. Similarly, if $|C| = 1$, the next chain in the cycle ordering has either length $|C| = 1$ or $|C| + 2 = 3$, and is traversed in the same direction or the opposite direction, respectively. Moreover, if $|C| = 0$, then the previous and next chain have length $|C| + 2 = 2$, and by the condition that the two edges of the cycle incident with this one-vertex chain must have their other end vertices in distinct levels, both chains of length 2 are traversed in the same direction. \square

We now define a cycle ordering Λ_n , $n \geq 2$, for the Greene-Kleitman SCD; see Figure 16 for illustration. The corresponding Hamilton cycle is oriented so that it traverses the longest chain $*^n$, which will be the first in the ordering Λ_n , from bottom to top. Our construction works inductively, and the induction step goes from n to $n + 2$, with separate rules for even and odd n . The base cases are $n = 0$ and $n = 1$, for which the entire cube consists only of a single vertex and a single edge, respectively, so for these cases the notion of a cycle ordering is not defined.

For even n , we define $\Lambda_0 := \varepsilon$, and for $n \geq 0$ and given $\Lambda_n =: C_1, \dots, C_N$ we define $\Lambda_{n+2} := \rho(\Lambda_n) = \rho(C_1), \dots, \rho(C_N)$ with

$$\rho(C) := \begin{cases} \lambda(C) & \text{if } |C| \equiv n \pmod{4}, \\ \lambda(C)^R & \text{if } |C| \not\equiv n \pmod{4}, \end{cases} \quad (11a)$$

and

$$\lambda(C) := \begin{cases} *C*, f(*C*), f(\ell(*C*)), \ell(*C*) & \text{if } |C| \geq 2, \\ *C*, 0C1 & \text{if } |C| = 0. \end{cases} \quad (11b)$$

We call the chains of $\lambda(C)$ arising from C the *descendants of C* . Essentially this rule replaces each chain C in Λ_n by its descendants $\lambda(C)$, where the order of descendants can be reversed, indicated by the superscript R , depending on the length of C modulo 4.

For odd n , we define $\Lambda_1 := *$, and for $n \geq 1$ and given Λ_n we define $\Lambda_{n+2} := \rho(\Lambda_n)$, where ρ is as before and

$$\lambda(C) := \begin{cases} *C*, \ell(*C*), \ell(f(*C*)), f(*C*) & \text{if } |C| \geq 3, \\ *C*, \ell(*C*), f(*C*) & \text{if } |C| = 1. \end{cases} \quad (11c)$$

Lemma 17. *Λ_n contains every chain of the Greene-Kleitman SCD exactly once.*

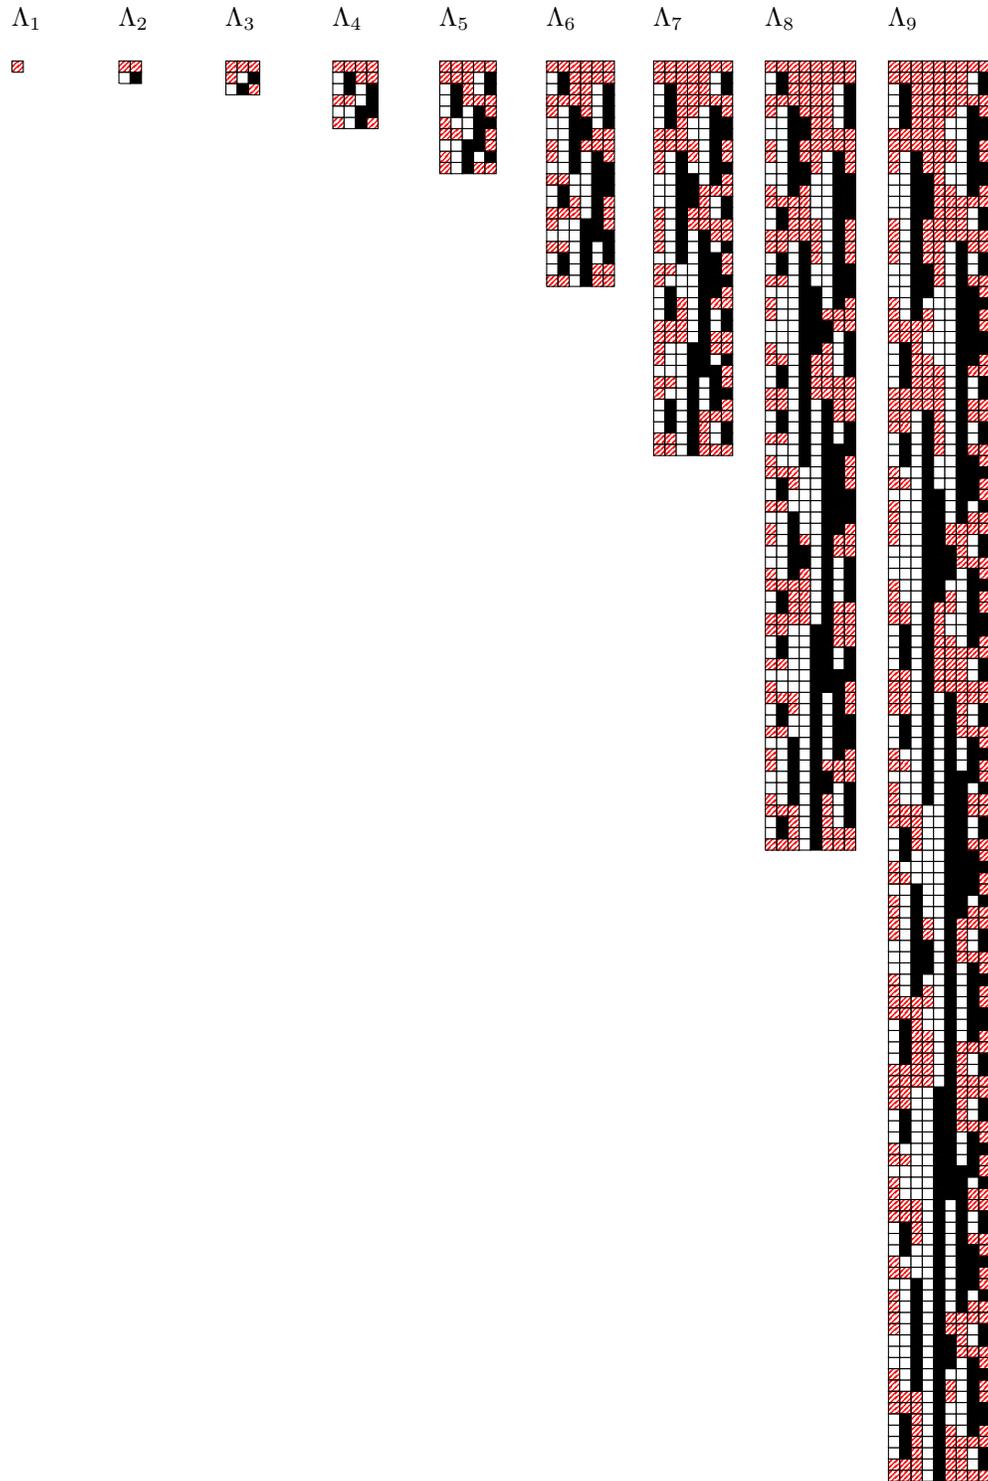


FIGURE 16. The cycle orderings Λ_n of Greene-Kleitman chains for $n = 1, \dots, 9$. In this figure, 1-bits are drawn as black squares, 0-bits as white squares, and *s as hatched squares. Figure 1 (f) is obtained by traversing the chains of Λ_7 up and down alternatingly.

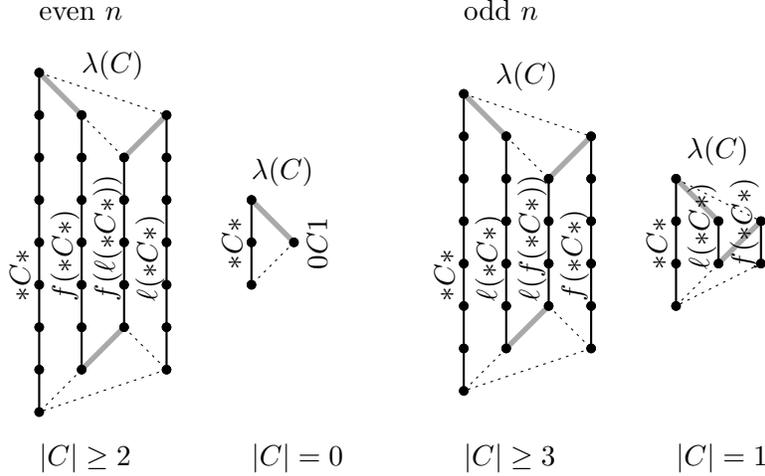


FIGURE 17. Connections between top and bottom ends of the descendants $\lambda(C)$ of a chain C , as guaranteed by Lemmas 18 and 22. Bold gray edges are the connections used along the Hamilton cycle. Dotted edges are present but not used.

Proof. Observe that for even n , if $|C| \geq 2$ in (11b), then $*C*$ is a $[- -]$ -chain of length $|C| + 2$, $f(*C*)$ is a $[+ -]$ -chain of length $|C|$, $\ell(*C*)$ is a $[- +]$ -chain of length $|C|$, and $f(\ell(*C*))$ is a $[+ +]$ -chain of length $|C| - 2$, and if $|C| = 2$, then the lattice path $f(\ell(*C*)) \in D$ touches the line $y = 0$ at least three times. If $|C| = 0$, on the other hand, then $*C*$ is a $[- -]$ -chain of length 2, $0C1$ is a $[+ +]$ -chain of length 1, and the lattice path $0C1 \in D$ touches the line $y = 0$ exactly twice. Overall, the inductive rule (11b) produces only distinct chains, and every chain of the Greene-Kleitman SCD in Q_{n+2} is produced. A similar argument works for odd n , showing that Λ_n indeed contains every chain of the Greene-Kleitman SCD exactly once. \square

To complete the proof of Theorem 3, it remains to show that any two consecutive chains in Λ_n can be joined by an edge between their top ends or bottom ends alternatingly. For this we need the following simple lemmas that guarantee these connecting edges.

Lemma 18. *For any $n \geq 2$ and any chain C with $|C| \geq 2$, the chains C and $f(C)$, and the chains C and $\ell(C)$ are connected both at their top and bottom ends in Q_n .*

All the connecting edges between top and bottom ends among the descendants of a chain guaranteed by Lemma 18 are shown in Figure 17.

Proof. By symmetry, it suffices to prove the lemma for the chains C and $f(C)$. As $|C| \geq 2$, we may consider the first two $*$ s in C and write $C = u*v*C'$ with $u, v \in D$, which yields $f(C) = u0v1C'$. From this we obtain that $b(C) = u0v0b(C')$ and $b(f(C)) = u0v1b(C')$, so $b(C)$ and $b(f(C))$ differ exactly in the bit after the valley v . Similarly, we have $t(C) = u1v1t(C')$ and $t(f(C)) = u0v1t(C')$, so $t(C)$ and $t(f(C))$ differ exactly in the bit after the valley u . \square

The next two lemmas are illustrated in the top part of Figure 18.

Lemma 19. *For any $n \geq 2$ and any two chains C, C' connected at their bottom ends in Q_n , we have that $*C*$ and $*C'*$ are connected at their bottom ends in Q_{n+2} .*

Proof. Note that $b(*C*) = 0b(C)0$ and $b(*C'*) = 0b(C')0$. Consequently, as C and C' are connected at their bottom ends, $b(C)$ and $b(C')$ differ in exactly one bit, implying that $b(*C*)$ and $b(*C'*)$ also differ in exactly one bit. \square

Lemma 20. For any $n \geq 2$ and any chain C with $|C| \geq 2$ in Q_n , we have that $\ell(*C*)$ and $\ell(*f(C)*)$ are connected at their bottom ends in Q_{n+2} . Specifically, if $|C| = 2$ and $C = u*v*w$ with $u, v, w \in D$, then we have

$$\begin{aligned}\ell(*C*) &= *u*v0w1, \\ \ell(*f(C)*) &= 0u0v1w1.\end{aligned}\tag{12a}$$

If $|C| \geq 3$ and $C = C'*u$ with $u \in D$ and $|C'| \geq 2$, then we have

$$\begin{aligned}\ell(*C*) &= *C'0u1, \\ \ell(*f(C)*) &= *f(C')0u1.\end{aligned}\tag{12b}$$

Consequently, in the first case, the chains $\ell(*C*)$ and $\ell(*f(C)*)$ differ in exactly three positions, and in the second case, they differ in exactly two positions.

Proof. First of all, the relations (12a) and (12b) can be verified directly using the definition of f and ℓ . From (12a) we obtain that $b(\ell(*C*))$ and $b(\ell(*f(C)*))$ differ exactly in the bit after v . From (12b) we obtain that $b(\ell(*C*)) = 0b(C')0u1$ and $b(\ell(*f(C)*)) = 0b(f(C'))0u1$, and as the substrings $b(C')$ and $b(f(C'))$ differ in exactly one bit by Lemma 18, this also holds for the entire strings. \square

Note that if $|C| \geq 3$ (not if $|C| = 2$), then by (12b) the two chains mentioned in Lemma 20 are also connected at their top ends, but this connection is irrelevant for us.

Proof of Theorem 3 (even n). We show that Λ_n , $n \geq 2$ even, defined in (11b) is a cycle ordering of the Greene-Kleitman chains, by proving that any consecutive pair of chains is connected at their top or bottom ends alternatingly, starting with the first chain $*^n$ of length n that is traversed from bottom to top. We will also establish the following additional property P: For any two consecutive chains C and C' connected at their top ends, we either have $C = f(C')$ or $f(C) = C'$. These invariants can easily be checked for the induction base case $n = 2$, which is given by $\Lambda_2 = **,01$.

For the induction step consider $n \geq 2$ to be even, and assume that Λ_n is a cycle ordering satisfying property P. By Lemma 18, the descendants $\lambda(C)$ for any chain C from Λ_n can be joined as shown on the left hand side of Figure 17, so we only need to check the connections between the first and last chains among consecutive groups of descendants. Indeed, if C and C' are consecutive in Λ_n and joined at their bottom ends, then C is traversed from top to bottom and C' from bottom to top in the Hamilton cycle; see the top left part of Figure 18. Consequently, by Lemma 16, we have $|C| \not\equiv |*^n| = n \pmod{4}$ and $|C'| \equiv n \pmod{4}$, i.e., by (11a) the sequence Λ_{n+2} contains $\lambda(C)^R$ and $\lambda(C')$, and indeed, the bottom vertex of the last chain of $\lambda(C)^R$, namely $*C*$, is connected to the bottom vertex of the first chain of $\lambda(C')$, namely $*C'*$, by Lemma 19.

Similarly, if C and C' are consecutive in Λ_n and joined at their top ends, then C is traversed from bottom to top and C' from top to bottom in the Hamilton cycle; see the top right part of Figure 18. Consequently, by Lemma 16, we have $|C| \equiv n \pmod{4}$ and $|C'| \not\equiv n \pmod{4}$, i.e., by (11a) the sequence Λ_{n+2} contains $\lambda(C)$ and $\lambda(C')^R$, and indeed, the bottom vertex of the last chain of $\lambda(C)$, namely $\ell(*C*)$, is connected to the bottom vertex of the first chain of $\lambda(C')^R$, namely $\ell(*C'*)$, using that by property P we have either $C = f(C')$ or $f(C) = C'$, so we can invoke Lemma 20. Moreover, property P still holds for Λ_{n+2} by the definition (11b) (note that if $|C| = 0$, then we have $0C1 = f(*C*)$). \square

To prove Theorem 3 for odd n , we need two additional lemmas, illustrated at the bottom part of Figure 18. Lemma 21 is the ‘dual’ version of Lemma 20, with f replaced by ℓ and vice versa, which we need as the definitions (11b) and (11c) in the first case differ in the ordering of descendants.

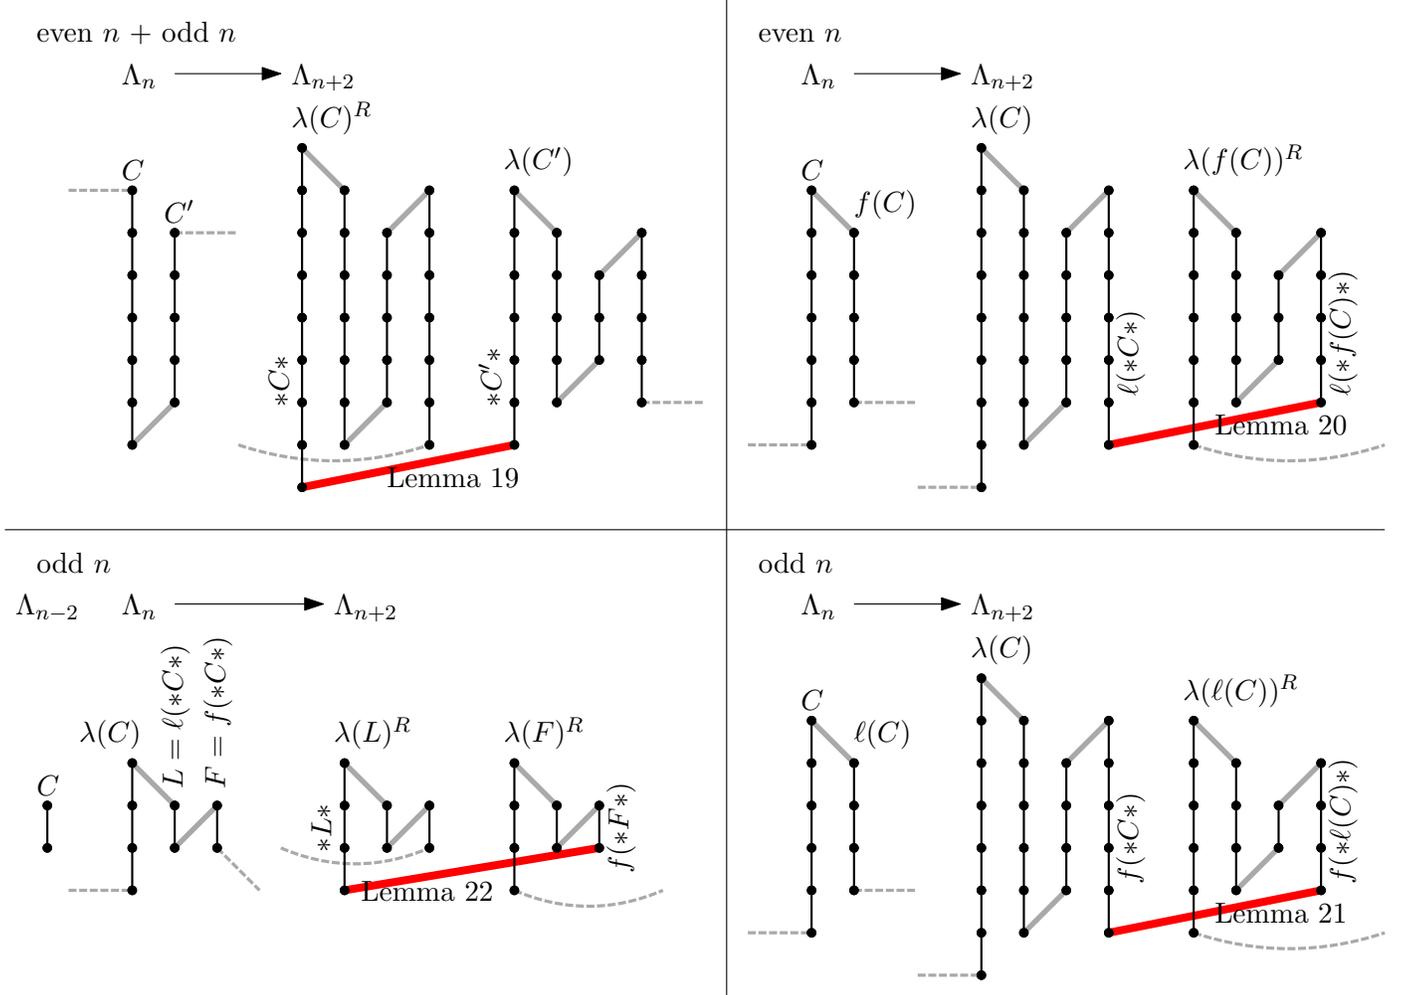


FIGURE 18. Joining of the descendants of two consecutive chains from Λ_n in the induction step $n \rightarrow n + 2$, via the thick edges guaranteed by Lemmas 19–22. The dashed edges are connections to preceding and subsequent chains on the Hamilton cycle.

Lemma 22 deals with the special case of consecutive chains of length 1, a situation that never occurs for even n .

Lemma 21. *For any $n \geq 3$ and any chain C with $|C| \geq 3$ in Q_n , we have that $f(*C*)$ and $f(*\ell(C)*)$ are connected at their bottom ends in Q_{n+2} . Specifically, if $C = u * C'$ with $u \in D$ and $|C'| \geq 2$, then we have*

$$\begin{aligned} f(*C*) &= 0 u 1 \ C' \ *, \\ f(*\ell(C)*) &= 0 u 1 \ \ell(C') \ *, \end{aligned} \tag{13}$$

*i.e., the chains $f(*C*)$ and $f(*\ell(C)*)$ differ in exactly two positions.*

Proof. The relation (13) can be verified directly using the definition of f and ℓ . From (13) we obtain that $b(f(*C*)) = 0 u 1 b(C') 0$ and $b(f(*\ell(C)*)) = 0 u 1 b(\ell(C')) 0$, and as the substrings $b(C')$ and $b(\ell(C'))$ differ in exactly one bit by Lemma 18, this also holds for the entire strings. \square

By (13), the two chains mentioned in Lemma 21 are also connected at their top ends, but this connection is irrelevant for us.

Lemma 22. *For any odd $n \geq 3$ and any chain C with $|C| = 1$ in Q_{n-2} , we have that the bottom end of $L := \ell(*C*)$ is joined to the top end of $F := f(*C*)$ in Q_n , and moreover $*L*$ and $f(*F*)$ are connected at their bottom ends in Q_{n+2} . Specifically, if $C = u*v$ with $u, v \in D$, then we have*

$$\begin{aligned} L &= *u0v1, & *L* &= **u0v1*, \\ F &= 0u1v*, & f(*F*) &= 00u1v1*, \end{aligned} \tag{14}$$

*i.e., both L and F , as well as $*L*$ and $f(*F*)$ differ in exactly three positions.*

Proof. The relation (14) can be verified directly using the definition of f and ℓ . From (14) we obtain that $b(L) = 0u0v1$ and $t(F) = 0u1v1$, so these two strings differ exactly in the bit after u . Similarly, we obtain that $b(*L*) = 00u0v10$ and $b(f(*F*)) = 00u1v10$, so these two strings also differ exactly in the bit after u . \square

Proof of Theorem 3 (odd n). We show that Λ_n , $n \geq 3$ odd, defined in (11c) is a cycle ordering of the Greene-Kleitman chains, by proving that any consecutive pair of chains is connected at their top or bottom ends alternatingly, starting with the first chain $*^n$ of length n that is traversed from bottom to top. We will also establish the following additional properties P' and Q: Property P' asserts that for any two consecutive chains C and C' connected at their top ends, we either have $C = \ell(C')$ or $\ell(C) = C'$. Property Q asserts that for any two consecutive chains B and B' of length 1 in Λ_n , there is a chain C of length 1 in Λ_{n-2} such that $\{B, B'\} = \{\ell(*C*), f(*C*)\}$, i.e., both B and B' are descendants of C (recall the second case of (11c)). Property Q implies in particular that no more than two chains of length 1 appear consecutively in Λ_n . These invariants can easily be checked for the induction base case $n = 3$, which is given by $\Lambda_3 = ***, *01, 01*$.

For the induction step consider $n \geq 3$ to be odd, and assume that Λ_n is a cycle ordering satisfying properties P' and Q. By Lemmas 18 and 22, the descendants $\lambda(C)$ for any chain C from Λ_n can be joined as shown on the right hand side of Figure 17, so we only need to check the connections between the first and last chains among consecutive groups of descendants. Indeed, if C and C' are consecutive in Λ_n and joined at their bottom ends, then C is traversed from top to bottom and C' from bottom to top in the Hamilton cycle; see the top left part of Figure 18. Consequently, by Lemma 16, we have $|C| \not\equiv |*^n| = n \pmod{4}$ and $|C'| \equiv n \pmod{4}$, i.e., by (11a) the sequence Λ_{n+2} contains $\lambda(C)^R$ and $\lambda(C')$, and indeed, the bottom vertex of the last chain of $\lambda(C)^R$, namely $*C*$, is connected to the bottom vertex of the first chain of $\lambda(C')$, namely $*C'*$, by Lemma 19.

Similarly, if C and C' are consecutive in Λ_n and joined at their top ends, then C is traversed from bottom to top and C' from top to bottom in the Hamilton cycle; see the bottom right part of Figure 18. Consequently, by Lemma 16, we have $|C| \equiv n \pmod{4}$ and $|C'| \not\equiv n \pmod{4}$, i.e., by (11a) the sequence Λ_{n+2} contains $\lambda(C)$ and $\lambda(C')^R$, and indeed, the bottom vertex of the last chain of $\lambda(C)$, namely $f(*C*)$, is connected to the bottom vertex of the first chain of $\lambda(C')^R$, namely $f(*C'*)$, using that by property P' we have either $C = \ell(C')$ or $\ell(C) = C'$, so we can invoke Lemma 21.

The last case to consider are two consecutive chains B and B' of length 1 in Λ_n , where the bottom end of B is joined to the top end of B' if $n \equiv 3 \pmod{4}$ or the other way round if $n \equiv 1 \pmod{4}$ (recall Lemma 16); see the bottom left part of Figure 18. We only consider the case $n \equiv 3 \pmod{4}$, as the other case is symmetric. By property Q, we know that $B = \ell(*C*) =: L$ and $B' = f(*C*) =: F$ for some chain C of length 1 in Λ_{n-2} . By (11a), the sequence Λ_{n+2} contains $\lambda(B)^R$ and $\lambda(B')^R$, and indeed, the bottom vertex of the last chain of $\lambda(B)^R = \lambda(L)^R$, namely $*L*$, is connected to the bottom vertex of the first chain of $\lambda(B')^R = \lambda(F)^R$, namely $f(*F*)$, by Lemma 22. Moreover, properties P' and Q still hold for Λ_{n+2} by the definition (11c). \square

8.1. Loopless algorithm. The following is an immediate consequence of the lemmas we established.

Theorem 23. *For any $n \geq 1$, the ordering of chains Λ_n defined in (11) is a 3-Gray code, i.e., any two consecutive chains, viewed as strings over the alphabet $\{0, 1, *\}$, differ in at most three positions.*

Proof. Note that by the definitions (11b) and (11c), any two consecutive chains in $\lambda(C)$ differ in exactly two positions. From this the theorem follows by induction on n , using Lemmas 19–22. \square

We now describe an algorithm which for a given dimension n computes the sequence of chains in the cycle ordering Λ_n defined in (11). Each chain is represented as a string of length n over the alphabet $\{0, 1, *\}$ (recall Figure 16), and whenever the next chain has been produced by the algorithm, it is visited in the [Visit] step. By Theorem 23, at most three entries of this string change between two visits. One can easily turn this algorithm into a loopless algorithm for computing the entire Hamilton cycle, flipping a single bit in each step, by replacing the [Visit] step by a loop that moves up and down the vertices on the current chain alternately.

For even n , a loopless implementation of this Gray code is given as Algorithm C, i.e., the algorithm requires only $\mathcal{O}(1)$ time between any two consecutive [Visit] steps. The corresponding variant of the algorithm for odd n is very similar and can be found in the Appendix. An implementation of both the even and odd case in C++ is available for download and for demonstration on the Combinatorial Object Server [cos]. The initialization required by this algorithm is $\mathcal{O}(n)$, and the used space is $\mathcal{O}(n)$. Table 2 shows the execution of this algorithm for the case $n = 6$.

Algorithm C (*Loopless Gray code for Greene-Kleitman chains in Q_n for even $n = 2m$, $m \geq 1$*). The input of the algorithm is an integer $m \geq 1$, and the produced chains in the cycle ordering Λ_{2m} are visited in step C2, where the current chain is stored in the array $C = c_1 \dots c_n$ with $c_i \in \{0, 1, *\}$. For maintaining the chain efficiently, the algorithm uses auxiliary arrays $p_1 \dots p_n$, $s_0 \dots s_n$, and $t_1 \dots t_{n+1}$, where if $c_i \in \{0, 1\}$, then p_i gives the position of the 0 or 1 matched to c_i , and if $c_i = *$, then s_i is the position of the closest $*$ to the right of c_i and t_i is the position of the closest $*$ to the left of c_i . The algorithm also maintains several additional arrays to track the recursive structure of Λ_n . Specifically, the chain constructed in dimension $2i$, $i = 1, \dots, m$, consists of the middle $2i$ entries of C . The array $d_0 \dots d_m$ is used to determine the dimension $2d_m$ in which the next construction step is performed. This array simulates a stack that generates the transition sequence of the Gray code, with d_m being the value on the top of the stack, an idea first used by Bitner, Ehrlich, and Reingold [BER76]. In addition, the algorithm maintains arrays $l_1 \dots l_m$, $b_1 \dots b_m$, $o_1 \dots o_m$, $q_1 \dots q_m$, where l_i is the length of the chain in dimension $2i$, $b_i \in \{+, -\}$ indicates whether we are currently in the first case of (11b) (+) or the second case (-), $o_i \in \{+, -\}$ indicates whether we are in the first case of (11a) (+) or the second case (-), and $q_i \in \{f, \ell, f^{-1}, \ell^{-1}\}$ specifies the next operation on the chain in dimension $2i$.

- C1.** [Initialize] Set $c_i \leftarrow *$ for $1 \leq i \leq n$, $s_i \leftarrow i + 1$ for $0 \leq i \leq n$, and $t_i \leftarrow i - 1$ for $1 \leq i \leq n + 1$. Also set $l_i \leftarrow 2i$ for $1 \leq i \leq m$, $b_1 \leftarrow -$ and $b_i \leftarrow +$ for $2 \leq i \leq m$, $o_i \leftarrow +$ and $q_i \leftarrow f$ for $1 \leq i \leq m$, $d_i \leftarrow i$ for $0 \leq i \leq m$.
- C2.** [Visit] Visit the chain $c_1 \dots c_n$.
- C3.** [Select dimension] Set $i \leftarrow d_m$. Terminate if $i = 0$.
- C4.** [Perform operation] Depending on the value of q_i , branch to one of the following four cases.
 - C41.** [Apply ℓ] If $q_i = \ell$, call **last**(i).
 - C42.** [Apply ℓ^{-1}] If $q_i = \ell^{-1}$, call **last** $^{-1}$ (i).
 - C43.** [Apply f] If $q_i = f$, then if $s_{m-i+1} \leq m + i$, call **first**(i), otherwise call **last** $^{-1}$ ($i + 1$), **first**(i), and **last**($i + 1$).
 - C44.** [Apply f^{-1}] If $q_i = f^{-1}$, then if $i = m$ or $c_{m-i} = *$, call **first** $^{-1}$ (i), otherwise call **last** $^{-1}$ ($i + 1$), **first** $^{-1}$ (i), and **last**($i + 1$).
- C5.** [Select next step] If $q_i \in \{f, \ell\}$, set $l_i \leftarrow l_i - 2$, otherwise set $l_i \leftarrow l_i + 2$. Also set $d_m \leftarrow m$. If $b_i = -$, then go to C6, else if $q_i \neq f^{-1}$, then go to C7, otherwise go to C8.

- C6.** [Reached last descendant in case $|C| = 0$] Set $d_i \leftarrow d_{i-1}$ and $d_{i-1} \leftarrow i - 1$. If $o_i = +$, then set $l_i \leftarrow 2$, otherwise set $l_i \leftarrow 4$. Also set $b_i \leftarrow +$, $o_i \leftarrow -o_i$, and $q_i \leftarrow f$, and go back to C2.
- C7.** [Descendants in case $|C| \geq 2$ remaining] If $q_i \in \{\ell, \ell^{-1}\}$ set $q_i \leftarrow f^{-1}$, otherwise set $q_i \leftarrow \ell$ if $o_i = +$ and $q_i \leftarrow \ell^{-1}$ if $o_i = -$. Go back to C2.
- C8.** [Reached last descendant in case $|C| \geq 2$] Set $d_i \leftarrow d_{i-1}$, $d_{i-1} \leftarrow i - 1$, and $j \leftarrow d_i$. If $j > 0$ and $q_j \in \{f, \ell\}$, set $l_i \leftarrow l_i - 2$, otherwise set $l_i \leftarrow l_i + 2$. If $l_i = 0$, set $b_i \leftarrow -$ and $q_i \leftarrow f^{-1}$, else if $l_i = 2$ and $o_i = -$, set $b_i \leftarrow -$ and $q_i \leftarrow f$, otherwise set $q_i \leftarrow f$. Also set $o_i \leftarrow -o_i$, and go back to C2.

The algorithm uses the following four auxiliary functions, which implement the functions $f, f^{-1}, \ell, \ell^{-1}$ on the chain C ; see Figure 19. The parameter $i \in \{1, \dots, m\}$ determines the middle $2i$ entries of C , namely $c_{m-i+1}, \dots, c_{m+i}$, that the function works on.

- **first**(i): Set $\alpha \leftarrow m - i + 1$, $\beta \leftarrow s_\alpha$, and $\gamma \leftarrow s_\beta$. Then set $c_\alpha \leftarrow 0$, $c_\beta \leftarrow 1$, $p_\alpha \leftarrow \beta$, $p_\beta \leftarrow \alpha$, $s_{\alpha-1} \leftarrow \gamma$, and $t_\gamma \leftarrow \alpha - 1$.
- **first** $^{-1}$ (i): Set $\alpha \leftarrow m - i + 1$, $\beta \leftarrow p_\alpha$, and $\gamma \leftarrow s_{\alpha-1}$. Then set $c_\alpha \leftarrow *$, $c_\beta \leftarrow *$, $s_{\alpha-1} \leftarrow \alpha$, $s_\alpha \leftarrow \beta$, $s_\beta \leftarrow \gamma$, $t_\gamma \leftarrow \beta$, $t_\beta \leftarrow \alpha$, and $t_\alpha \leftarrow \alpha - 1$.
- **last**(i): Set $\alpha \leftarrow m + i$, $\beta \leftarrow t_\alpha$, and $\gamma \leftarrow t_\beta$. Then set $c_\beta \leftarrow 0$, $c_\alpha \leftarrow 1$, $p_\beta \leftarrow \alpha$, $p_\alpha \leftarrow \beta$, $s_\gamma \leftarrow \alpha + 1$, and $t_{\alpha+1} \leftarrow \gamma$.
- **last** $^{-1}$ (i): Set $\alpha \leftarrow m + i$, $\beta \leftarrow p_\alpha$, and $\gamma \leftarrow t_{\alpha+1}$. Then set $c_\beta \leftarrow *$, $c_\alpha \leftarrow *$, $s_\gamma \leftarrow \beta$, $s_\beta \leftarrow \alpha$, $s_\alpha \leftarrow \alpha + 1$, $t_{\alpha+1} \leftarrow \alpha$, $t_\alpha \leftarrow \beta$, and $t_\beta \leftarrow \gamma$.

Note that the two subcases in steps **C43** and **C44** are captured by the relations (12b) and (12a) in Lemma 20, respectively.

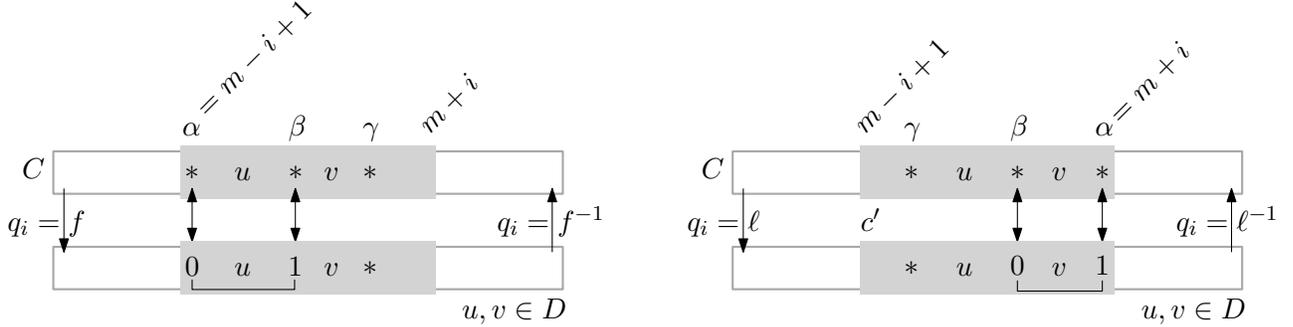


FIGURE 19. Illustration of the four auxiliary functions used in Algorithm C.

9. ACKNOWLEDGEMENTS

We thank Jiří Fink for several valuable discussions about symmetric chain decompositions, and for feedback on an earlier draft of this paper. We also thank the anonymous reviewers whose suggestions helped improving the presentation.

REFERENCES

- [Aig73] M. Aigner. Lexicographic matching in Boolean algebras. *J. Combin. Theory Ser. B*, 14:187–194, 1973.
- [BER76] J. Bitner, G. Ehrlich, and E. Reingold. Efficient generation of the binary reflected Gray code and its applications. *Comm. ACM*, 19(9):517–521, 1976.
- [Bol86] B. Bollobás. *Combinatorics: set systems, hypergraphs, families of vectors and combinatorial probability*. Cambridge University Press, Cambridge, 1986.
- [BW84] M. Buck and D. Wiedemann. Gray codes with restricted density. *Discrete Math.*, 48(2-3):163–171, 1984.
- [Che00] Y. Chen. Kneser graphs are Hamiltonian for $n \geq 3k$. *J. Combin. Theory Ser. B*, 80(1):69–79, 2000.

TABLE 2. Protocol of Algorithm C for $n = 2m = 6$. Empty entries have unchanged values compared to the previous iteration/row. Struck through values have changed twice in some iteration of the algorithm.

	C	p	s	t	l	b	o	q	d	
1	*****		1 2 3 4 5 6 7	0 1 2 3 4 5 6	2 4 6	--+	+++	f f f	0 1 2 3	
2	0 1	2 1	3	0	4			ℓ		
3	0 1	6 5	7	4	2			f^{-1}		
4	**		1	2	2		-	f	2	
5	0 1	3 2	4	1	2			ℓ	3	
6	0 1	4 1	7	0	0			ℓ^{-1}		
7	**		5	0 6	2			f^{-1}		
8	* *		1 5	4	2		-	f	2	
9	0 1	5 4	6	1	0			f^{-1}	3	
10	0 1	6 1	7	0	2		+	-	2	
11	**0	6 3	1 2 7	2	0		-	-	1 3	
12	0 1	2 1	7	0	0			ℓ^{-1}		
13	* *		3 6	0 3 6	2			f^{-1}		
14	**		1 3	2	2		-	+	f	2 1
15	0 0 1	5 4 3 2	6	1	2		+	-	0 3	
16	0 1	6 1	7	0	2		+	-	2	
17	** 0	6 5	1 2 7	2	4		+	+	f	1 0 3
18	0 1	2 1	7	0	0			ℓ^{-1}		
19	**		5	0 5 6	2			f^{-1}		
20	**		1 5	2	6			f	2 0	

- [Che03] Y. Chen. Triangle-free Hamiltonian Kneser graphs. *J. Combin. Theory Ser. B*, 89(1):1–16, 2003.
- [cos] The Combinatorial Object Server: <http://www.combos.org/chains>.
- [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.
- [EG73] P. Erdős and R. K. Guy. Crossing number problems. *Amer. Math. Monthly*, 80:52–58, 1973.
- [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.
- [Ehr73] G. Ehrlich. Loopless algorithms for generating permutations, combinations, and other combinatorial configurations. *J. Assoc. Comput. Mach.*, 20:500–513, 1973.
- [FdFSV08] L. Faria, C. M. H. de Figueiredo, O. Sýkora, and I. Vrt’o. An improved upper bound on the crossing number of the hypercube. *J. Graph Theory*, 59(2):145–161, 2008.
- [Fin07] J. Fink. Perfect matchings extend to Hamilton cycles in hypercubes. *J. Combin. Theory Ser. B*, 97(6):1074–1076, 2007.
- [Fin19] J. Fink. Matchings extend into 2-factors in hypercubes. *Combinatorica*, 39(1):77–84, 2019.
- [FS13] T. Feder and C. Subi. On hypercube labellings and antipodal monochromatic paths. *Discrete Appl. Math.*, 161(10-11):1421–1426, 2013.
- [Für85] Z. Füredi. Problem session. In *Kombinatorik geordneter Mengen*, Oberwolfach, BRD, 1985.
- [GJM⁺21] P. Gregor, S. Jäger, T. Mütze, J. Sawada, and K. Wille. Gray codes and symmetric chains. To appear in *J. Combin. Theory Ser. B*. Preprint available at *arXiv:1802.06021*, 2021.
- [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.
- [GMM20] P. Gregor, O. Mička, and T. Mütze. On the central levels problem. In A. Czumaj, A. Dawar, and E. Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020*, volume 168 of *LIPICs*, pages 60:1–60:17. Schloss Dagstuhl, 2020.
- [GMN18] P. Gregor, T. Mütze, and J. Nummenpalo. A short proof of the middle levels theorem. *Discrete Analysis*, 2018:8:12 pp., 2018.

- [Gra53] F. Gray. Pulse code communication, 1953. March 17, 1953 (filed Nov. 1947). U.S. Patent 2,632,058.
- [Gri93] J. R. Griggs. On the distribution of sums of residues. *Bull. Amer. Math. Soc. (N.S.)*, 28(2):329–333, 1993.
- [GŠ10] P. Gregor and R. Škrekovski. On generalized middle-level problem. *Inform. Sci.*, 180(12):2448–2457, 2010.
- [Hav83] 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.
- [Hua19] H. Huang. Induced subgraphs of hypercubes and a proof of the sensitivity conjecture. *Ann. of Math. (2)*, 190(3):949–955, 2019.
- [Hur94] G. Hurlbert. The antipodal layers problem. *Discrete Math.*, 128(1-3):237–245, 1994.
- [Kle65] D. J. Kleitman. On a lemma of Littlewood and Offord on the distribution of certain sums. *Math. Z.*, 90:251–259, 1965.
- [Knu11] D. E. Knuth. *The Art of Computer Programming. Vol. 4A. Combinatorial Algorithms. Part 1*. Addison-Wesley, Upper Saddle River, NJ, 2011.
- [KRSW04] C. E. Killian, F. Ruskey, C. D. Savage, and M. Weston. Half-simple symmetric Venn diagrams. *Electron. J. Combin.*, 11(1):Research Paper 86, 22, 2004.
- [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 Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2238–2253. SIAM, Philadelphia, PA, 2017.
- [MS17] T. Mütze and P. Su. Bipartite Kneser graphs are Hamiltonian. *Combinatorica*, 37(6):1207–1219, 2017.
- [Müt16] T. Mütze. Proof of the middle levels conjecture. *Proc. Lond. Math. Soc.*, 112(4):677–713, 2016.
- [Nor08] S. Norine. Edge-antipodal colorings of cubes. The Open Problem Garden. Available at http://www.openproblemgarden.org/op/edge_antipodal_colorings_of_cubes, 2008.
- [NS94] N. Nisan and M. Szegedy. On the degree of Boolean functions as real polynomials. *Comput. Complexity*, 4(4):301–313, 1994. Special issue on circuit complexity (Barbados, 1992).
- [Pik99] O. Pikhurko. On edge decompositions of posets. *Order*, 16(3):231–244 (2000), 1999.
- [RS93] 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.
- [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.
- [Rus03] F. Ruskey. Combinatorial generation. Book draft, 2003.
- [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.
- [Sim91] J. Simpson. Hamiltonian bipartite graphs. In *Proceedings of the Twenty-second Southeastern Conference on Combinatorics, Graph Theory, and Computing (Baton Rouge, LA, 1991)*, volume 85, pages 97–110, 1991.
- [SK79] J. Shearer and D. J. Kleitman. Probabilities of independent choices being ordered. *Stud. Appl. Math.*, 60(3):271–276, 1979.
- [Spi19] H. Spink. Orthogonal symmetric chain decompositions of hypercubes. *SIAM J. Discrete Math.*, 33(2):910–932, 2019.
- [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.
- [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.
- [SW19] X. S. Shen and A. Williams. A middle levels conjecture for multiset permutations with uniform-frequency. Williams College Technical Report CSTR-201901. Available at <http://tmuetze.de/papers/sw1.pdf>, 2019.
- [TL73] D. Tang and C. Liu. Distance-2 cyclic chaining of constant-weight codes. *IEEE Trans. Computers*, C-22:176–180, 1973.
- [Tom15] I. Tomon. On a conjecture of Füredi. *European J. Combin.*, 49:1–12, 2015.
- [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.

APPENDIX A. LOOPLESS ALGORITHM FOR ODD n

Algorithm C' (*Loopless Gray code for Greene-Kleitman chains in Q_n for odd $n = 2m + 1$, $m \geq 0$*). The input of the algorithm is an integer $m \geq 0$, and the produced chains in the cycle ordering Λ_{2m+1} are visited in step C'2. The algorithm uses the same data structures and auxiliary functions as Algorithm C before. The entries $b_i \in \{+, -\}$ now indicate whether we are in the first case of (11c) (+) or the second case (-). Moreover, the entries of the array q may now also attain two other values in addition to $\{f, f^{-1}, \ell, \ell^{-1}\}$, namely $g := \ell^{-1} \circ f$ and $g^{-1} := f^{-1} \circ \ell$, to account for the transitions between the last two chains in the second case of (11c).

- C'1.** [Initialize] Set $c_i \leftarrow *$ for $1 \leq i \leq n$, $s_i \leftarrow i + 1$ for $0 \leq i \leq n$, and $t_i \leftarrow i - 1$ for $1 \leq i \leq n + 1$. Also set $l_i \leftarrow 2i + 1$ for $1 \leq i \leq m$, $b_i \leftarrow +$ for $1 \leq i \leq m$, $o_i \leftarrow +$ and $q_i \leftarrow \ell$ for $1 \leq i \leq m$, $d_i \leftarrow i$ for $0 \leq i \leq m$.
- C'2.** [Visit] Visit the chain $c_1 \dots c_n$.
- C'3.** [Select dimension] Set $i \leftarrow d_m$. Terminate if $i = 0$.
- C'4.** [Perform operation] Depending on the value of q_i , branch to one of the following six cases.
- C'41.** [Apply ℓ] If $q_i = \ell$, call $\text{last}(i)$.
- C'42.** [Apply ℓ^{-1}] If $q_i = \ell^{-1}$, call $\text{last}^{-1}(i)$.
- C'43.** [Apply f] If $q_i = f$, call $\text{first}(i)$.
- C'44.** [Apply f^{-1}] If $q_i = f^{-1}$, call $\text{first}^{-1}(i)$.
- C'45.** [Apply ℓ^{-1} , then f] If $q_i = g = \ell^{-1} \circ f$, then if $i = m$ call $\text{last}^{-1}(i)$ and $\text{first}(i)$, otherwise call $\text{last}^{-1}(i)$, $\text{first}(i)$, and $\text{first}(i + 1)$.
- C'46.** [Apply f^{-1} , then ℓ] If $q_i = g^{-1} = f^{-1} \circ \ell$, then if $i = m$ call $\text{first}^{-1}(i)$ and $\text{last}(i)$, otherwise call $\text{first}^{-1}(i + 1)$, $\text{first}^{-1}(i)$, and $\text{last}(i)$.
- C'5.** [Select next step] If $q_i \in \{f, \ell\}$, set $l_i \leftarrow l_i - 2$, else if $q_i \in \{f^{-1}, \ell^{-1}\}$, set $l_i \leftarrow l_i + 2$. Also set $d_m \leftarrow m$. If $b_i = -$ and $q_i \in \{\ell, g^{-1}\}$, then go to C6, else if $b_i = -$ and $q_i \in \{g, \ell^{-1}\}$, then go to C7, else if $q_i \neq \ell^{-1}$, then go to C8, otherwise go to C9.
- C'6.** [Descendants in case $|C| = 1$ remaining] If $q_i = \ell$ set $q_i \leftarrow g$, otherwise set $q_i \leftarrow \ell^{-1}$. Go back to C2.
- C'7.** [Reached last descendant in case $|C| = 1$] Set $d_i \leftarrow d_{i-1}$, $d_{i-1} \leftarrow i - 1$, and $j \leftarrow d_i$. If $j > 0$ and $q_j \in \{f, \ell, g\}$, set $l_i \leftarrow l_i - 2$, otherwise set $l_i \leftarrow l_i + 2$. If $l_i = 1$, set $q_i \leftarrow g^{-1}$, else if $l_i = 3$ and $q_{i-1} = \ell^{-1}$, set $b_i \leftarrow +$, $o_i \leftarrow -$ and $q_i \leftarrow \ell$, else if $l_i = 3$ and $q_{i-1} \neq \ell^{-1}$, set $q_i \leftarrow \ell$, otherwise set $b_i \leftarrow +$, $q_i \leftarrow \ell$ and $o_i \leftarrow -o_i$. Go back to C2.
- C'8.** [Descendants in case $|C| \geq 3$ remaining] If $q_i \in \{f, f^{-1}\}$ set $q_i \leftarrow \ell^{-1}$, otherwise set $q_i \leftarrow f$ if $o_i = +$ and $q_i \leftarrow f^{-1}$ if $o_i = -$. Go back to C2.
- C'9.** [Reached last descendant in case $|C| \geq 3$] Set $d_i \leftarrow d_{i-1}$, $d_{i-1} \leftarrow i - 1$, and $j \leftarrow d_i$. If $j > 0$ and $q_j \in \{f, \ell, g\}$, set $l_i \leftarrow l_i - 2$, otherwise set $l_i \leftarrow l_i + 2$. If $l_i = 1$, set $b_i \leftarrow -$ and $q_i \leftarrow g^{-1}$, else if $l_i = 3$ and $o_i = -$, set $b_i \leftarrow -$ and $q_i \leftarrow \ell$, otherwise set $q_i \leftarrow \ell$. Also set $o_i \leftarrow -o_i$, and go back to C2.

Note that the two subcases in steps **C'45** and **C'46** are captured by the left two equations and the right two equations in (14) in Lemma 22, respectively.