

Gray codes and symmetric chains

Petr Gregor

Department of Theoretical Computer Science and Mathematical Logic, Charles University,
Prague, Czech Republic
gregor@ktiml.mff.cuni.cz

Sven Jäger

Institut für Mathematik, Technische Universität Berlin, Germany
jaeger@math.tu-berlin.de

Torsten Mütze

Institut für Mathematik, Technische Universität Berlin, Germany
muetze@math.tu-berlin.de

Joe Sawada

School of Computer Science, University of Guelph, Canada
jsawada@uoguelph.ca

Kaja Wille

Institut für Mathematik, Technische Universität Berlin, Germany
wille@math.tu-berlin.de

Abstract

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

2012 ACM Subject Classification Mathematics of computing \rightarrow Combinatorics, Mathematics of computing \rightarrow Graph theory

Keywords and phrases Gray code, Hamilton cycle, hypercube, poset, symmetric chain

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.140

Category Regular submission

Related Version Full version available on *arxiv:1802.06021*

1 Introduction

Gray codes are named after Frank Gray, a researcher at Bell Labs, who described a simple method to generate all 2^n bitstrings of length n by flipping a single bit in each step [8], now known as the binary reflected Gray code. This code found widespread use, e.g., in circuit design and testing, signal processing and error correction, data compression, etc.; many more applications are mentioned in the survey [28]. The binary reflected Gray code is also implicit in the well-known *Towers of Hanoi* puzzle and the *Chinese ring* puzzle that date back to the 19th century. The theory of Gray codes has developed considerably




© Petr Gregor, Sven Jäger, Torsten Mütze, Joe Sawada, and Kaja Wille;
licensed under Creative Commons License CC-BY

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).

Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Don Sannella; Article No. 140,
pp. 140:1–140:14



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



in the last decades, and the term is now used more generally to describe an exhaustive listing of any class of combinatorial objects where successive objects in the list differ by a small amount. In particular, such generation algorithms have been developed for several fundamental combinatorial objects of interest for computer scientists, such as bitstrings, permutations, partitions, trees, etc., all of which are covered in depth in the most recent volume of Knuth's seminal series *The Art of Computer Programming* [20].

Since the discovery of the binary reflected Gray code, there has been continued interest in developing Gray codes for bitstrings of length n that satisfy various additional constraints. For instance, a Gray code with the property that each bit is flipped (almost) the same number of times was first constructed by Tootill [35]. Goddyn and Gvozđjak constructed an n -bit Gray code in which any two flips of the same bit are almost n steps apart [7], which is best possible. These are only two examples of a large body of work on possible Gray code transition sequences; see also [3, 5, 34]. Savage and Winkler constructed a Gray code that generates all 2^n bitstrings such that all bitstrings with Hamming weight k appear before all bitstrings with weight $k + 2$, for each $0 \leq k \leq n - 2$ [29], where the *Hamming weight* of a bitstring is the number of its 1-bits. They used this construction to tackle the infamous *middle two levels problem*, which asks for a cyclic listing of all bitstrings of length $2n + 1$ with weights in the interval $[n, n + 1]$ by flipping a single bit in each step. This problem was raised in the 1980s and received considerable attention in the literature (a detailed historic account is given in [22]). A general existence proof for such a Gray code for any $n \geq 1$ has been found only recently [12, 22], and an algorithm for computing it using $\mathcal{O}(1)$ amortized time and $\mathcal{O}(n)$ space was subsequently presented in [23]. The starting point of this work is the following more general problem raised in [13, 27].

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

The special case $\ell = 1$ of Problem M is the middle two levels problem mentioned before. The case $\ell = n + 1$ is solved by the binary reflected Gray code discussed in the beginning. Moreover, the cases $\ell = n$ and $\ell = n - 1$ were settled in [6, 21] and [13], respectively.

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

The most general version of this problem is whether the subgraph of the n -cube induced by all levels in an arbitrary weight interval $[a, b]$ has an (almost) Hamilton cycle. This was solved in [11] for all possible values of $n \geq 1$ and $0 \leq a \leq b \leq n$, except in the cases when the length n of the bitstrings is odd and the levels a and b are symmetric around the middle, which is exactly Problem M. For all other cases that paper provides algorithms that generate each bitstring in those Gray codes in constant time.

1.1 Our results

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

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




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

Combining Theorem 1 with the results from [11] shows more generally that the subgraph of the n -cube induced by any four consecutive levels has an ‘almost’ Hamilton cycle.¹

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

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

Our proof of Theorem 2 is concise and illustrative, and it motivates the subsequent discussion, so we present it right now. It uses a well-known concept from the theory of partially ordered sets (posets), a so-called symmetric chain decomposition. Here we define this term for the n -cube using graph-theoretic language. A *symmetric chain* in Q_n is a path $(x_k, x_{k+1}, \dots, x_{n-k})$ in the n -cube where x_i is from level i for all $k \leq i \leq n - k$, and a *symmetric chain decomposition*, or SCD for short, is a partition of the vertices of Q_n into symmetric chains. For illustration, an SCD of Q_5 is shown in Figure 1 (a). We say that two SCDs are *edge-disjoint* if the corresponding paths in the graph Q_n are edge-disjoint, i.e., if there are no two consecutive vertices in one chain of the first SCD that are also contained in one chain of the second SCD. There is a well-known construction of two edge-disjoint SCDs in the n -cube for any $n \geq 1$ [31], which we will discuss momentarily.

¹ If the four levels are not symmetric around the middle, then this subgraph of the n -cube has two partition classes of different sizes, and thus cannot have a Hamilton cycle. However, it was shown in [11] that in those cases the graph has a cycle that visits all vertices in the smaller partition class, and also a cyclic listing of all vertices in which only few transitions flip two instead of one bit, where ‘few’ means only as many as the difference in size between the two partition classes. Both of these notions are natural generalizations of a Hamilton cycle.



Figure 2 The parenthesis matching approach for constructing the symmetric chain containing a bitstring x , yielding the symmetric chain decomposition \mathcal{D}_0 . The highlighted bits are the leftmost unmatched 0 and the rightmost unmatched 1 in each bitstring.

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

This proof motivates the search for a large collection of pairwise edge-disjoint SCDs in the n -cube. We can then use any two of them to construct a cycle factor as described in the previous proof, and use this cycle factor as a starting point for building a Hamilton cycle. This two-step approach of building a Hamilton cycle via a cycle factor proved to be very successful for such problems (see e.g. [15, 16, 17, 18, 22, 24, 30]). Consequently, for the rest of this section we focus on edge-disjoint SCDs in the n -cube.


There is a well-known construction of an SCD for the n -cube that is best described by the following parenthesis matching approach pioneered by Greene and Kleitman [9]; see Figure 2. 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 move up the chain, or the rightmost unmatched 1 to move down 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 $n \geq 1$. We denote this standard SCD by \mathcal{D}_0 ; it is shown in Figure 1 (a) for $n = 5$.

By taking complements, we obtain another SCD, which we denote by $\overline{\mathcal{D}}_0$. It is not hard to see that \mathcal{D}_0 and $\overline{\mathcal{D}}_0$ are in fact edge-disjoint for any $n \geq 1$ [31]. Figure 1 (b) shows both SCDs for $n = 5$, and how they are used for building a cycle factor. Apart from this standard construction, we are not aware of any other construction of an SCD in the n -cube, even though there are several different ways to describe the same SCD (see e.g. [1, 4, 36]).

Our next result is a simple construction of another SCD in the n -cube for *even* values of $n \geq 2$, which we call \mathcal{D}_1 . This construction is based on lattice paths and will be explained in Section 4 below. It has the additional feature that $\mathcal{D}_0, \overline{\mathcal{D}}_0, \mathcal{D}_1$ and $\overline{\mathcal{D}}_1$ are pairwise edge-disjoint for $n \geq 6$.

► **Theorem 3.** *For any even $n \geq 6$, the n -cube contains four pairwise edge-disjoint SCDs.*

Figure 3 shows the SCDs \mathcal{D}_0 and \mathcal{D}_1 in Q_6 . Their complements $\overline{\mathcal{D}}_0$ and $\overline{\mathcal{D}}_1$ are not shown for clarity. Note that four edge-disjoint SCDs are best possible for Q_6 , as they use up all edges incident with the middle level.



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

■ **Table 1** Known pairwise edge-disjoint SCDs in the n -cube for $n = 1, 2, \dots, 11$. The definitions of $\mathcal{X}_5, \mathcal{Y}_5, \mathcal{Z}_5$ and $\mathcal{X}_7, \mathcal{Y}_7$ are given in Section 4.3.

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

For odd values of n , we can still construct four edge-disjoint SCDs in the n -cube (except in a few small cases). However, the construction is not as direct and explicit as for even n .

► **Theorem 4.** *For $n = 7$ and any odd $n \geq 13$, the n -cube contains four pairwise edge-disjoint SCDs.*

For odd n , we can combine any two of the four edge-disjoint SCDs in the n -cube guaranteed by Theorem 4 to a cycle factor in the middle 2ℓ levels, as explained before, yielding in total $\binom{4}{2} = 6$ distinct cycle factors, four of which are non-isomorphic. To prove Theorem 4, we construct four edge-disjoint SCDs in the 7-cube in an ad hoc fashion and then apply the following product construction.

► **Theorem 5.** *If Q_a and Q_b each contain k pairwise edge-disjoint SCDs, then Q_{a+b} contains k pairwise edge-disjoint SCDs.*

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

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

1.2 Related work

Apart from building Gray codes, symmetric chain decompositions have many other interesting applications, e.g., to construct rotation-symmetric Venn diagrams for n sets when n is a

prime number [14, 26], and to solve the Littlewood-Offord problem on sums of vectors [2].

A notion that is closely related to edge-disjoint SCDs is that of *orthogonal* chain decompositions, which were first considered by Shearer and Kleitman [31]. Two chain decompositions are called *orthogonal* if every pair of chains has at most one vertex in common, where one also allows chains that are not symmetric around the middle or chains that skip some levels. Shearer and Kleitman showed in their paper that \mathcal{D}_0 and $\overline{\mathcal{D}_0}$ are almost orthogonal (only the longest chains have two elements in common), and they conjectured that the n -cube has $\lfloor n/2 \rfloor + 1$ pairwise orthogonal chain decompositions where each decomposition consists of $\binom{n}{\lfloor n/2 \rfloor}$ many chains. Spink recently made some progress towards this conjecture, by showing that the n -cube has three orthogonal chain decompositions [32].

Pikhurko showed via a parenthesis matching argument that all edges of the n -cube can be decomposed into symmetric chains [25]. However, it is not clear whether these chains contain a subset that forms an SCD. Another interesting construction relating Hamilton cycles and SCDs in the n -cube was presented by Streib and Trotter [33]. They inductively construct a Hamilton cycle in the n -cube for any $n \geq 2$ that can be partitioned into symmetric chains forming an SCD. This Hamilton cycle has the minimal number of ‘peaks’ where the differences in the Hamming weight change sign.

1.3 Outline of this paper

In Section 2 we introduce several definitions that will be used throughout this paper. In Section 3 we sketch the main ideas for proving Theorem 1. The full proof is omitted due to space constraints and can be found in the preprint [10]. In Section 4 we present the proofs of Theorems 3–5, and we describe the construction of the SCD \mathcal{D}_1 and of the SCDs in Q_5 and Q_7 referred to in Table 1. We conclude in Section 5 with some open problems.

2 Preliminaries

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

2.1 Bitstrings and lattice paths

We use $L_{n,k}$ to denote the set of all bitstrings of length n with Hamming weight k , so this is exactly the k th level of Q_n . For any bitstring x , we write \bar{x} for its complement and $\text{rev}(x)$ for the reversed bitstring. We often interpret a bitstring x as a path in the integer lattice \mathbb{Z}^2 starting at the origin $(0,0)$, where every 1-bit is interpreted as an \nearrow -step that changes the current coordinate by $(+1, +1)$ and every 0-bit is interpreted as a \searrow -step that changes the current coordinate by $(+1, -1)$; see Figure 4. Note that for any bitstring x , the inverted bitstring \bar{x} corresponds to mirroring the lattice path horizontally, and the inverted and reversed bitstring $\overline{\text{rev}}(x)$ corresponds to mirroring the lattice path vertically.

2.2 Lexical matchings

We now introduce certain matchings between two consecutive levels of the hypercube, which were first described by Kierstead and Trotter [19]. Originally, these matchings were defined and analyzed for the graph between the middle two levels of the $(2n+1)$ -cube in an attempt to tackle the middle two levels problem. We first generalize them to the n -cube for arbitrary n and an arbitrary pair of consecutive levels k and $k+1$. For $i \in \{0, 1, \dots, n-1\}$ the *i -lexical matching* is defined as follows; see Figure 4. We interpret a bitstring x as a lattice path,




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


and we let x^\uparrow denote the lattice path that is obtained by appending \searrow -steps to x until the resulting path ends at height -1 . If x ends at a height less than -1 , then $x^\uparrow := x$. Similarly, we let x^\downarrow denote the lattice path obtained by appending \nearrow -steps to x until the resulting path ends at height $+1$. If x ends at a height more than $+1$, then $x^\downarrow := x$. We define the matching by two partial mappings $M_{n,k}^{i,\uparrow}: L_{n,k} \rightarrow L_{n,k+1}$ and $M_{n,k}^{i,\downarrow}: L_{n,k+1} \rightarrow L_{n,k}$ defined as follows: For any $x \in L_{n,k}$ we consider the lattice path x^\uparrow and scan it row-wise from top to bottom, and from right to left in each row. The partial mapping $M_{n,k}^{i,\uparrow}(x)$ is obtained by flipping the i th \searrow -step encountered in this fashion, where counting starts with $0, 1, \dots$, if this \searrow -step is part of x ; otherwise x is left unmatched. Similarly, for any $x \in L_{n,k+1}$ we consider the lattice path x^\downarrow and scan it row-wise from top to bottom, and from left to right in each row. The partial mapping $M_{n,k}^{i,\downarrow}(x)$ is obtained by flipping the i th \nearrow -step encountered in this fashion if this \nearrow -step is part of x ; otherwise x is left unmatched. It is straightforward to verify that these two partial mappings are inverse to each other, so they indeed define a matching between levels k and $k + 1$ of Q_n , which we denote by $M_{n,k}^i$.

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

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

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

Property (i) holds as in the smaller of the two levels k and $k + 1$, no steps are appended to the lattice paths when computing the i -lexical matching. Property (ii) holds as every edge in Q_n corresponds to some flip in the lattice path corresponding to the end vertex of the



■ **Figure 5** Joining two cycles C_1 and C_2 (black) from our cycle factor by taking the symmetric difference with a 6-cycle C' (gray).

edge in the smaller of the two levels k and $k + 1$. Property (iii) follows from the observation that complementing a bitstring corresponds to mirroring the lattice path horizontally, and reverting a bitstring corresponds to mirroring the lattice path horizontally and vertically.

3 The middle four levels problem

In this section we outline the main steps for proving Theorem 1. The proof proceeds similarly as the proof of the middle two levels problem [12, 22]. In a first step, we construct a cycle factor in the middle four levels of the $(2n + 1)$ -cube, and in a second step we modify the cycles in the factor locally to join them to a Hamilton cycle. The cycle factor is constructed by taking the union of the following edge sets: all n -lexical and $(n + 1)$ -lexical matching edges between the upper two levels $n + 1$ and $n + 2$ and between the lower two levels $n - 1$ and n as well as certain carefully chosen edges from the $(n - 2)$ -lexical, the $(n - 1)$ -lexical, and the n -lexical matching between the middle two levels n and $n + 1$. The most technical step here is to choose an appropriate set of edges between the middle two levels, so that the resulting subgraph has degree two at every vertex. When this is accomplished, we define a set of 6-cycles between levels $n + 1$ and $n + 2$ such that any two of these 6-cycles are edge-disjoint and every such 6-cycle C' intersects with two cycles C_1 and C_2 from our cycle factor as shown in Figure 5. Consequently, taking the symmetric difference of the edge sets of C_1 , C_2 , and C' results in a single cycle on the same vertex set as C_1 and C_2 . We repeat this joining process until we end up with a single Hamilton cycle. In this process, we exploit that all of the 6-cycles used for the joining are edge-disjoint, and that on any cycle of the factor, no pairs of edges that two 6-cycles have in common with this cycle are interleaved, so there are never any conflicts between them. The main advantage of this two-step approach to proving Hamiltonicity is that it effectively reduces the problem of proving that a graph has a Hamilton cycle to the problem of proving that a suitably defined auxiliary graph is connected, which is much easier. All details of this proof can be found in [10].

4 Pairwise edge-disjoint SCDs

We proceed to prove Theorems 3–5.

4.1 Proof of Theorem 3

To prove Theorem 3, we first give an equivalent definition of the SCD \mathcal{D}_0 defined in the introduction via the parenthesis matching approach, which is valid only for even values of $n \geq 2$; recall Figure 1 (a) and Figure 2.

For even $n \geq 2$, we consider a vertex $x \in L_{n, n/2}$ in the middle level $n/2$ of Q_n , and we define the sequence of vertices reached from x when moving up the corresponding chain, and the sequence of vertices reached when moving down the chain. For this we consider the




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

lattice path corresponding to the bitstring x . This lattice path ends at the coordinate $(n, 0)$ as the number of 0s equals the number of 1s. We now label a subsequence of \searrow -steps of this lattice path with integers $j = 1, 2, \dots$ according to the following procedure; see the top part of Figure 6 for an illustration:


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

Flipping the \searrow -steps of x marked with $1, 2, \dots$ in this order yields the sequence of vertices reached from x when moving up the chain containing x . An analogous labeling procedure obtained by interchanging left and right, \searrow -steps and \swarrow -steps, and starting with ending points yields the sequence of vertices reached from x when moving down this chain. We denote this chain by $C_0(x)$. Observe that $C_0(x)$ is a symmetric chain, as the height of the marker decreases by 1 in each step, so the number of edges we move up from x equals the number of edges we move down from x . It is easy to verify that the SCD \mathcal{D}_0 defined before via the parenthesis matching approach satisfies $\mathcal{D}_0 = \bigcup_{x \in L_{n,n/2}} C_0(x)$.

Proof of Theorem 3. We first define a set \mathcal{D}_1 of chains in Q_n for even values of $n \geq 2$ via a labeling rule similar to the rule for \mathcal{D}_0 described before. From this definition it follows immediately that all chains in \mathcal{D}_1 are symmetric. We then use an equivalent characterization of \mathcal{D}_0 and \mathcal{D}_1 as the unions of certain lexical matchings to show that the chains in \mathcal{D}_1 form a partition of all vertices of Q_n , proving that \mathcal{D}_1 is an SCD, and that $\mathcal{D}_0, \overline{\mathcal{D}_0}, \mathcal{D}_1,$ and $\overline{\mathcal{D}_1}$ are pairwise edge-disjoint.

For even $n \geq 2$, we consider a vertex $x \in L_{n,n/2}$ in the middle level of Q_n . We interpret it as a lattice path, and label some of its \searrow -steps as follows; see the bottom part of Figure 6:

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



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

step and the \searrow -step starting at the marker, we label the second step from the right from this set with $j + 1$, and we move the marker to the starting point of the rightmost \searrow -step starting at height $h - 2$. We set $j := j + 2$ and repeat.

(c1) If the marker is at height $h = 1$ or $h = 0$, we stop.

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

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

is actually an SCD of Q_n .

By property (i) from Lemma 6, for any sequence $\mathbf{i} := (i_0, i_1, \dots, i_{n-1})$ of indices $i_k \in \{0, 1, \dots, \max\{k, n - k - 1\}\}$ the union

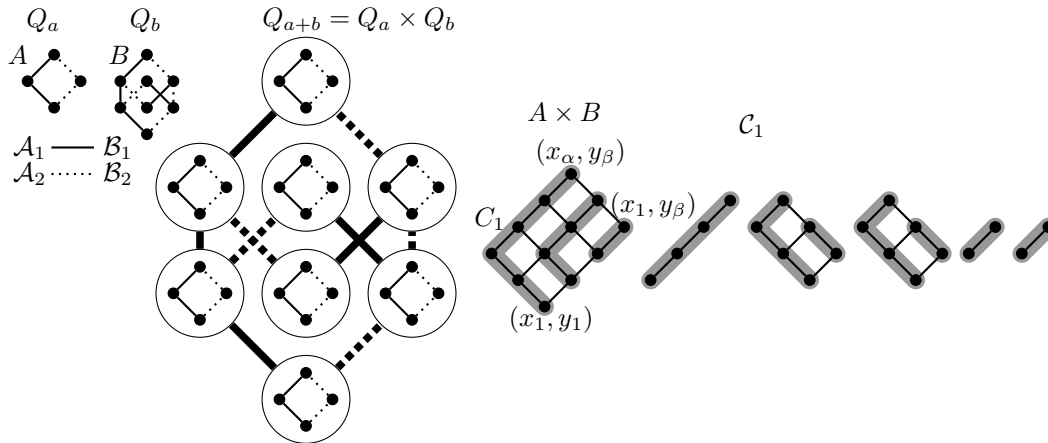
$$\mathcal{D}_{\mathbf{i}} := \bigcup_{k=0}^{n-1} M_{n,k}^{i_k} \tag{2}$$

is a chain decomposition of Q_n . The resulting chains are not necessarily symmetric, though. From the definitions in Section 2.2 it also follows that \mathcal{D}_0 equals the union of the 0-lexical matchings, and that for even $n \geq 2$, \mathcal{D}_1 as defined in (1) equals the union of the 1-lexical matchings; formally we have

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

Consequently, \mathcal{D}_1 is indeed a chain decomposition, and by the definition of \mathcal{D}_1 via the labeling procedure, all chains in this decomposition are symmetric, so \mathcal{D}_1 is indeed an SCD. The fact that \mathcal{D}_0 , $\overline{\mathcal{D}_0}$, \mathcal{D}_1 , and $\overline{\mathcal{D}_1}$ are pairwise edge-disjoint can be seen by applying property (iii) from Lemma 6 and by observing that by property (ii), $\mathcal{D}_{\mathbf{i}}$ and $\mathcal{D}_{\mathbf{j}}$ as defined in (2) are edge-disjoint if and only if the sequences \mathbf{i} and \mathbf{j} differ in every position; see Figure 7. ◀

Clearly, $\mathcal{D}_{(0,0,\dots,0)}$ as defined in (2) equals \mathcal{D}_0 for every $n \geq 1$, so the union of all 0-lexical matchings forms an SCD in any dimension. In contrast to that, the union of all 1-lexical matchings $\mathcal{D}_{(1,1,\dots,1)}$ only forms an SCD for even $n \geq 2$.



■ **Figure 8** Illustration of the proof of Theorem 5. Construction of two edge-disjoint SCDs in Q_5 from two edge-disjoint SCDs in Q_2 and two edge-disjoint SCDs in Q_3 . The chains of the SCD C_1 of Q_5 as constructed in the proof are highlighted in gray.

4.2 Proof of Theorem 5

Proof of Theorem 5. For the reader’s convenience, this proof is illustrated in Figure 8. Let $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_k$ and $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_k$ denote k pairwise edge-disjoint SCDs of Q_a and Q_b , respectively. We will think of Q_{a+b} as the Cartesian product $Q_a \times Q_b$ of Q_a and Q_b . We show how to construct for every $i \in [k]$ an SCD C_i of $Q_{a+b} = Q_a \times Q_b$ which uses only edges of the form $((u, v), (u', v'))$ where (u, u') is an edge from \mathcal{A}_i or (v, v') is an edge from \mathcal{B}_i . From this it follows that the SCDs C_1, C_2, \dots, C_k are pairwise edge-disjoint.

The SCD C_i of Q_{a+b} is defined as follows: The Cartesian products $A \times B$ of chains $A \in \mathcal{A}_i$ and $B \in \mathcal{B}_i$ partition the vertices of Q_{a+b} into two-dimensional grids. C_i is obtained by partitioning each of those grids into symmetric chains in the natural way; see Figure 8 (cf. [4]): Specifically, let $A =: (x_1, \dots, x_\alpha)$ and $B =: (y_1, \dots, y_\beta)$ be the vertices in the chains A and B from bottom to top. As A and B are symmetric, we know that $|x_1| + |x_\alpha| = a$ and $|y_1| + |y_\beta| = b$, where $|x|$ denotes the Hamming weight of the bitstring x . This implies that $|(x_1, y_1)| + |(x_\alpha, y_\beta)| = |x_1| + |y_1| + |x_\alpha| + |y_\beta| = a + b$, i.e., the bottom and top vertex of the grid $A \times B$ are on symmetric levels in Q_{a+b} . We may therefore decompose $A \times B$ into disjoint symmetric chains C_j , $j = 1, 2, \dots, \min\{\alpha, \beta\}$, by setting


$$C_j := ((x_1, y_j), (x_2, y_j), \dots, (x_{\alpha-j+1}, y_j), (x_{\alpha-j+1}, y_{j+1}), \dots, (x_{\alpha-j+1}, y_\beta)) . \quad \blacktriangleleft$$

4.3 Proof of Theorem 4

We begin by constructing the SCDs in Q_5 and Q_7 mentioned in Table 1.

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

Proof. We consider the graph Q_n with the two vertices in the outermost levels 0 and n removed, and we identify all bitstrings that differ only by rotation into so-called *necklaces*. The resulting graph N_n is a multigraph version of the cover graph of the *necklace poset*. Specifically, the multiplicity of the edges in N_n corresponds to the number of ways a bit from a necklace can be flipped to reach the corresponding adjacent necklace. E.g., in N_5 the necklace $x := 10000$ has two edges leading to $y := 11000$, as we can flip the second



■ **Figure 9** Illustration of the three edge-disjoint SCDs in N_5 (a) and four edge-disjoint SCDs in N_7 (b). The names of the SCDs correspond to the ones used in Table 1.

or the fifth bit in x to reach y . This way, a necklace on level k has $n - k$ edges going up, and k edges going down, like the vertices in Q_n . The multigraphs N_5 and N_7 are shown in Figure 9. If n is prime, then every SCD in N_n corresponds to an SCD in Q_n , by turning each chain from N_n into n chains in Q_n , obtained by rotating a representative of each necklace in all possible ways. Moreover, one of the chains of length $n - 2$ needs to be extended by the all-zero and all-one bitstring to a chain of length n in Q_n . Observe that in this way, k edge-disjoint SCDs in N_n give rise to k edge-disjoint SCDs in Q_n .

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

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

5 Open problems

- Understanding the structure of the cycle factors constructed as in the proof of Theorem 2 is an important step towards a general solution of Problem M. We performed some computer experiments in this direction; see [10]. What is the number and length of cycles in these factors? Is there a combinatorial interpretation of those numbers?
- Are there other explicit constructions of SCDs in the n -cube, different from \mathcal{D}_0 , \mathcal{D}_1 , and their complements?
- We conjecture that the n -cube has $\lfloor n/2 \rfloor + 1$ pairwise edge-disjoint SCDs. The main difficulty here is that we are missing a simple criterion like Hall's matching condition guaranteeing the existence of an SCD. Even finding five edge-disjoint SCDs in the n -cube for some small fixed n would be interesting, as this solution would extend to infinitely many larger values of n by Theorem 5. Beyond that, it would be very nice to construct more than constantly many edge-disjoint SCDs in the n -cube as n grows.

References

- 1 M. Aigner. Lexicographic matching in Boolean algebras. *J. Combin. Theory Ser. B*, 14:187–194, 1973.
- 2 B. Bollobás. *Combinatorics: set systems, hypergraphs, families of vectors and combinatorial probability*. Cambridge University Press, Cambridge, 1986.
- 3 B. Bultena and F. Ruskey. Transition restricted Gray codes. *Electron. J. Combin.*, 3(1):Paper 11, 11 pp., 1996.
- 4 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.
- 5 D. Dimitrov, T. Dvořák, P. Gregor, and R. Škrekovski. Linear time construction of a compressed Gray code. *European J. Combin.*, 34(1):69–81, 2013.
- 6 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.
- 7 L. Goddyn and P. Gvozdjak. Binary Gray codes with long bit runs. *Electron. J. Combin.*, 10:Paper 27, 10 pp., 2003.
- 8 F. Gray. Pulse code communication, 1953. March 17, 1953 (filed Nov. 1947). U.S. Patent 2,632,058.
- 9 C. Greene and D. J. Kleitman. Strong versions of Sperner’s theorem. *J. Combin. Theory Ser. A*, 20(1):80–88, 1976.
- 10 P. Gregor, S. Jäger, T. Mütze, J. Sawada, and K. Wille. Gray codes and symmetric chains. *arXiv:1802.06021*. Preprint version of the present article, 2018.
- 11 P. Gregor and T. Mütze. Trimming and gluing Gray codes. *Theoret. Comput. Sci.*, 714:74–95, 2018.
- 12 P. Gregor, T. Mütze, and J. Nummenpalo. A short proof of the middle levels theorem. To appear in *Discrete Analysis*. *arXiv:1710.08249*, 2018.
- 13 P. Gregor and R. Škrekovski. On generalized middle-level problem. *Inform. Sci.*, 180(12):2448–2457, 2010.
- 14 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.
- 15 A. E. Holroyd. Perfect snake-in-the-box codes for rank modulation. *IEEE Trans. Inform. Theory*, 63(1):104–110, 2017.
- 16 A. E. Holroyd, F. Ruskey, and A. Williams. Shorthand universal cycles for permutations. *Algorithmica*, 64(2):215–245, 2012.
- 17 J. R. Johnson. Universal cycles for permutations. *Discrete Math.*, 309(17):5264–5270, 2009.
- 18 J. R. Johnson. An inductive construction for Hamilton cycles in Kneser graphs. *Electron. J. Combin.*, 18(1):Paper 189, 12 pp., 2011.
- 19 H. A. Kierstead and W. T. Trotter. Explicit matchings in the middle levels of the Boolean lattice. *Order*, 5(2):163–171, 1988.
- 20 D. E. Knuth. *The Art of Computer Programming. Vol. 4A. Combinatorial Algorithms. Part 1*. Addison-Wesley, Upper Saddle River, NJ, 2011.
- 21 S. Locke and R. Stong. Problem 10892: Spanning cycles in hypercubes. *Amer. Math. Monthly*, 110:440–441, 2003.
- 22 T. Mütze. Proof of the middle levels conjecture. *Proc. Lond. Math. Soc.*, 112(4):677–713, 2016.
- 23 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.

- 24 T. Mütze, J. Nummenpalo, and B. Walczak. Sparse Kneser graphs are Hamiltonian. To appear in *Proceedings of the 50th Annual ACM Symposium on the Theory of Computing (STOC 2018)*. *arXiv:1711.01636*, 2018.
- 25 O. Pikhurko. On edge decompositions of posets. *Order*, 16(3):231–244 (2000), 1999.
- 26 F. Ruskey, C. D. Savage, and S. Wagon. The search for simple symmetric Venn diagrams. *Notices Amer. Math. Soc.*, 53(11):1304–1312, 2006.
- 27 C. D. Savage. Long cycles in the middle two levels of the Boolean lattice. *Ars Combin.*, 35(A):97–108, 1993.
- 28 C. D. Savage. A survey of combinatorial Gray codes. *SIAM Rev.*, 39(4):605–629, 1997.
- 29 C. D. Savage and P. Winkler. Monotone Gray codes and the middle levels problem. *J. Combin. Theory Ser. A*, 70(2):230–248, 1995.
- 30 J. Sawada and A. Williams. A Hamilton path for the sigma-tau problem. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 568–575, 2018.
- 31 J. Shearer and D. J. Kleitman. Probabilities of independent choices being ordered. *Stud. Appl. Math.*, 60(3):271–276, 1979.
- 32 H. Spink. Orthogonal symmetric chain decompositions of hypercubes. *arXiv:1706.08545*, June 2017.
- 33 N. Streib and W. T. Trotter. Hamiltonian cycles and symmetric chains in Boolean lattices. *Graphs Combin.*, 30(6):1565–1586, 2014.
- 34 I. N. Suparta and A. J. van Zanten. A construction of Gray codes inducing complete graphs. *Discrete Math.*, 308(18):4124–4132, 2008.
- 35 G. C. Tootill. The use of cyclic permuted codes in relay counting circuits. *Proceedings IEE, Part B Supplement*, 103, 1956.
- 36 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.