

**Manuscript version: Author's Accepted Manuscript**

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

**Persistent WRAP URL:**

<http://wrap.warwick.ac.uk/118878>

**How to cite:**

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**Publisher's statement:**

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: [wrap@warwick.ac.uk](mailto:wrap@warwick.ac.uk).

# A constant-time algorithm for middle levels Gray codes\*

Torsten Mütze

Institut für Mathematik  
TU Berlin  
10623 Berlin, Germany  
muetze@math.tu-berlin.de

Jerri Nummenpalo

Department of Computer Science  
ETH Zürich  
8092 Zürich, Switzerland  
njerri@inf.ethz.ch

ABSTRACT. For any integer  $n \geq 1$ , a *middle levels Gray code* is a cyclic listing of all  $n$ -element and  $(n + 1)$ -element subsets of  $\{1, 2, \dots, 2n + 1\}$  such that any two consecutive sets differ in adding or removing a single element. The question whether such a Gray code exists for any  $n \geq 1$  has been the subject of intensive research during the last 30 years, and has been answered affirmatively only recently [T. Mütze. Proof of the middle levels conjecture. *Proc. London Math. Soc.*, 112(4):677–713, 2016]. In a follow-up paper [T. Mütze and J. Nummenpalo. An efficient algorithm for computing a middle levels Gray code. To appear in *ACM Transactions on Algorithms*, 2018] this existence proof was turned into an algorithm that computes each new set in the Gray code in time  $\mathcal{O}(n)$  on average. In this work we present an algorithm for computing a middle levels Gray code in optimal time and space: each new set is generated in time  $\mathcal{O}(1)$ , and the required space is  $\mathcal{O}(n)$ .

KEYWORDS: Middle levels conjecture, Gray code, Hamilton cycle

## 1. INTRODUCTION

Efficiently generating all objects in a particular combinatorial class such as permutations, subsets, partitions, trees, strings etc. is one of the oldest and most fundamental algorithmic problems. Such generation algorithms are used as building blocks in a wide range of practical applications; the survey [Sav97] lists numerous references. In fact, more than half of the most recent volume of Knuth’s seminal series *The Art of Computer Programming* [Knu11] is devoted to this fundamental subject. The ultimate goal for these problems is to come up with algorithms that generate each new object in constant time, entailing that consecutive objects may differ only in a constant amount. Such constant-time algorithms are known for several combinatorial classes, and many of these results are covered in the classical books [NW75, Wil89]. To mention some concrete examples, constant-time algorithms are known for the following problems:

- (1) generating all permutations of  $[n] := \{1, 2, \dots, n\}$  by adjacent transpositions [Joh63, Tro62, Der75, Sed77],
- (2) generating all subsets of  $[n]$  by adding or removing an element in each step [Gra53],
- (3) generating all  $k$ -element subsets of  $[n]$  by exchanging an element in each step [Ehr73, BER76, EHR84, EM84, Rus88],
- (4) generating all binary trees with  $n$  vertices by rotation operations [Zer85, Luc87, LRvBR93],
- (5) generating all spanning trees of a graph by exchanging an edge in each step [Cum66, Kam67, HH72].

In this paper we revisit the well-known problem of generating all  $n$ -element and  $(n + 1)$ -element subsets of  $[2n + 1]$  by adding or removing a single element in each step. In a computer these subsets are naturally represented by bitstrings of length  $2n + 1$ , with 1-bits at the positions of the elements

---

\*An extended abstract of this paper appeared in the Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017 [MN17].

contained in the set and 0-bits at the remaining positions. Consequently, the problem is equivalent to generating all bitstrings of length  $2n+1$  with weight  $n$  or  $n+1$ , where the *weight* of a bitstring is the number of 1s in it. We refer to such a Gray code as a *middle levels Gray code*. Clearly, a middle levels Gray code has  $N := \binom{2n+1}{n} + \binom{2n+1}{n+1} = 2^{\Theta(n)}$  many bitstrings in total, and the weight alternates between  $n$  and  $n+1$  in each step. The existence of a middle levels Gray code for any  $n \geq 1$  is asserted by the well-known *middle levels conjecture*, raised independently in the 80s by Havel [Hav83] and Buck and Wiedemann [BW84]. The conjecture has also been attributed to Dejter, Erdős, Trotter [KT88] and various others, it appears in the popular books [Win04, Knu11, DG12], and it is mentioned in Gowers' recent expository survey on Peter Keevash's work [Gow17]. The middle levels conjecture has attracted considerable attention over the last 30 years [Sav93, FT95, SW95, Joh04, DSW88, KT88, DKS94, HKRR05, GŠ10, MW12, SSS09, SA11], and a positive solution, i.e., an existence proof for a middle levels Gray code for any  $n \geq 1$ , has been announced only recently.

**Theorem 1** ([Müt16, GMN17]). *A middle levels Gray code exists for any  $n \geq 1$ .*

In a follow-up paper [MN18], this existence argument was turned into an algorithm for computing a middle levels Gray code.

**Theorem 2** ([MN18]). *There is an algorithm, which for a given bitstring of length  $2n+1$ ,  $n \geq 1$ , with weight  $n$  or  $n+1$  computes the next  $\ell \geq 1$  bitstrings in a middle levels Gray code in time  $\mathcal{O}(\ell n + n^2)$ .*

Clearly, the running time of this algorithm is  $\mathcal{O}(n)$  on average per generated bitstring for  $\ell = \Omega(n)$ . However, this falls short of the optimal  $\mathcal{O}(1)$  time bound one could hope for, given that in each step only a single bit needs to be flipped, which is a constant amount of change.

**1.1. Our results.** In this paper we present an algorithm for computing a middle levels Gray code in optimal time and space.

**Theorem 3.** *There is an algorithm, which for a given bitstring of length  $2n+1$ ,  $n \geq 1$ , with weight  $n$  or  $n+1$  computes the next  $\ell \geq 1$  bitstrings in a middle levels Gray code in time  $\mathcal{O}(\ell + n)$ .*

Clearly, the running time of this algorithm is  $\mathcal{O}(1)$  on average per generated bitstring for  $\ell = \Omega(n)$ , and the required initialization time  $\mathcal{O}(n)$  and the required space  $\mathcal{O}(n)$  are also optimal.

We implemented our new middle levels Gray code in C++, and we invite the reader to experiment with this code, which can be found on the authors' websites [www]. As a benchmark, we used this code to compute a middle levels Gray code for  $n = 19$  in 20 minutes on a standard desktop computer. This is by a factor of 72 faster than the 24 hours reported in [MN18] for the algorithm from Theorem 2, and by four orders of magnitude faster than the 164 days previously needed for a brute-force search [SA11]. Note that a middle levels Gray code for  $n = 19$  consists of  $N = 137.846.528.820 \approx 10^{11}$  bitstrings. For comparison, a program that simply counts from  $1, \dots, N$  and does nothing else was only by a factor of 5 faster (4 minutes) than our middle levels Gray code computation on the same hardware. Roughly speaking, we need about 5 machine instructions for producing the next bitstring in the Gray code.

We now also obtain efficient algorithms for a number of related Gray codes that have been constructed using Theorem 1 as an induction basis. These Gray codes consist of several combined middle levels Gray codes of smaller dimensions. Specifically, it was a long-standing problem to construct a Gray code that lists all  $k$ -element and  $(n-k)$ -element subsets of  $[n]$ , where  $n \geq 2k+1$ , by either adding or removing  $n-2k$  elements in each step. This was solved in [MS17], and using Theorem 3 this construction can now be turned into an efficient algorithm. Moreover, in [GM18] Theorem 3 is used to derive constant-time algorithms for generating minimum-change listings of all  $n$ -bit strings whose weight is in some interval  $[k, l]$ ,  $0 \leq k \leq l \leq n$ , a far-ranging generalization of the middle levels conjecture and the classical problems (2) and (3) mentioned before.

In this work we restrict our attention to computing one particular ‘canonical’ middle levels Gray code for any  $n \geq 1$ , even though we know from [Müt16] that there are double-exponentially many different ones (recall [MN18, Remark 3]).

**1.2. Making the algorithm loopless.** We shall see that most steps of our algorithm require only constant time  $\mathcal{O}(1)$  in the worst case to generate the next bitstring, but after every sequence of  $\Theta(n)$  such ‘fast’ steps, a ‘slow’ step which requires time  $\Theta(n)$  is encountered, yielding constant average time performance. Therefore, our algorithm could easily be transformed into a loopless algorithm, i.e., one with a  $\mathcal{O}(1)$  *worst case* bound for each generated bitstring, by introducing an additional FIFO queue of size  $\Theta(n)$  and by simulating the original algorithm such that during every sequence of  $d$  ‘fast’ steps,  $d - 1$  results are stored in the queue and only one of them is returned, and during the ‘slow’ steps the queue is emptied at the same speed. For this the constant  $d$  must be chosen so that the queue is empty when the ‘slow’ steps are finished. This idea of delaying the output to achieve a loopless algorithm is also used in [HR16] (see also [Sed77, Section 1]). Even though the resulting algorithm would indeed be loopless, it would still be slower than the original algorithm, as it produces every bitstring only after it was produced in the original algorithm, due to the delay caused by the queue and the additional instructions for queue management. In other words, the hidden constant in the  $\mathcal{O}(1)$  bound for the modified algorithm is higher than for the original algorithm, so this loopless algorithm is only of theoretical interest, and we will not discuss it any further.

**1.3. Ingredients.** Our algorithm for computing a middle levels Gray code implements the strategy of the short proof of Theorem 1 presented in [GMN17]. In the most basic version, the algorithm computes several short cycles that together visit all bitstrings of length  $2n + 1$  with weight  $n$  or  $n + 1$ . We then modify a few steps of the algorithm so that these short cycles are joined to a Gray code that visits all bitstrings consecutively.

Let us briefly discuss the main differences between the algorithms from Theorem 2 and Theorem 3 and the improvements that save us a factor of  $n$  in the running time. At the lowest level, the algorithm from Theorem 2 consists of a rather unwieldy recursion, which for any given bitstring computes the next one in a middle levels Gray code. This recursion runs in time  $\Theta(n)$ , and therefore represents one of the bottlenecks in the running time. In addition, there are various high-level functions that are called every  $\Theta(n)$  many steps and run in time  $\Theta(n^2)$ , and which therefore also represent  $\Theta(n)$  bottlenecks. These high-level functions control which subsets of bitstrings are visited in which order, to make sure that each bitstring is visited exactly once.

To overcome these bottlenecks, we replace the recursion at the lowest level by a simple combinatorial technique, first proposed in [MSW18] and heavily used in the short proof of Theorem 1 presented in [GMN17]. This technique allows us to compute for certain ‘special’ bitstrings that are encountered every  $\Theta(n)$  many steps, a sequence of bit positions to be flipped during the next  $\Theta(n)$  many steps. Computing such a flip sequence can be done in time  $\Theta(n)$ , and when this is accomplished each subsequent step takes only constant time: We simply flip the precomputed positions one after the other, until the next ‘special’ bitstring is encountered and the flip sequence has to be recomputed. The high-level functions in the new algorithm are very similar as in the old one. We cut down their running time by a factor of  $n$  (from quadratic to linear) by using more sophisticated data structures and by resorting to well-known algorithms such as Booth’s linear-time algorithm [Boo80] for computing the lexicographically smallest rotation of a given string.

**1.4. Outline of this paper.** In Section 2 we introduce important definitions that will be used throughout the paper. In Section 3 we present our new middle levels Gray code algorithm. In Section 4 we prove the correctness of the algorithm, and in Section 5 we discuss how to implement it to achieve the claimed runtime and space bounds.

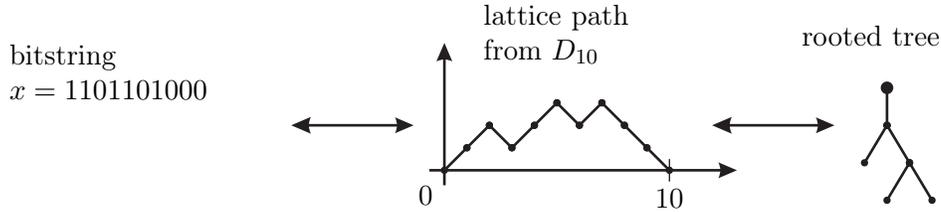


FIGURE 1. Bijections between bitstrings and lattice paths, and between Dyck paths and rooted trees (right).

## 2. PRELIMINARIES

*Operations on sequences and bitstrings.* We let  $(a_1, \dots, a_n)$  denote the sequence of integers  $a_1, \dots, a_n$ . We generalize this notation allowing  $a_i$  to be itself an integer sequence—in that case, if  $a_i = (b_1, \dots, b_m)$ , then  $(a_1, \dots, a_n)$  is shorthand for  $(a_1, \dots, a_{i-1}, b_1, \dots, b_m, a_{i+1}, \dots, a_n)$ . The empty integer sequence is denoted by  $()$ . For any sequence  $a$ , we let  $|a|$  denote its length. For any integer  $k \geq 0$  and any bitstring  $x$ , we write  $x^k$  for the concatenation of  $k$  copies of  $x$ . Moreover,  $\text{rev}(x)$  denotes the reversed bitstring, and  $\bar{x}$  denotes the bitstring obtained by taking the complement of every bit in  $x$ . We also define  $\overline{\text{rev}}(x) := \text{rev}(\bar{x}) = \text{rev}(x)$ . For any graph  $G$  whose vertices are bitstrings and any bitstring  $x$ , we write  $Gx$  for the graph obtained from  $G$  by appending  $x$  to all vertices.

*Bitstrings and lattice paths.* We let  $B_{n,k}$  denote the set of all bitstrings of length  $n$  with weight  $k$ . Any bitstring  $x \in B_{n,k}$  can be interpreted as a lattice path as follows; see Figure 1: We read  $x$  from left to right and draw a path in the integer lattice  $\mathbb{Z}_2$  that starts at the origin  $(0, 0)$ . For every 1-bit encountered in  $x$ , we draw an  $\nearrow$ -step that changes the current coordinate by  $(+1, +1)$ , and for every 0-bit encountered in  $x$ , we draw a  $\searrow$ -step that changes the current coordinate by  $(+1, -1)$ . Note that the resulting lattice path ends at the coordinate  $(n, 2k - n)$ . We let  $D_n$  denote the bitstrings from  $B_{2n,n}$  with the property that in every prefix, there are at least as many 1s as 0s. Moreover, we let  $D_n^-$  denote the bitstrings from  $B_{2n,n}$  that have this property for all but exactly one prefix. In terms of lattice paths,  $D_n$  are the paths that always stay above the abscissa  $y = 0$ , commonly known as *Dyck paths*, whereas  $D_n^-$  are the paths that move below the abscissa  $y = 0$  exactly once. It is well-known that  $|D_n| = |D_n^-|$  and that this quantity is given by the  $n$ th Catalan number. We also define  $D := \bigcup_{n \geq 0} D_n$ . Any  $x \in D$  can be written uniquely as  $x = 1u0v$  with  $u, v \in D$ . Similarly, any  $x \in D^-$  can be written uniquely as  $x = u01v$  with  $u, v \in D$ . We refer to this as the *canonical decomposition* of  $x$ ; see Table 1.

*Rooted trees.* An (*ordered*) *rooted tree* is a tree with a specified 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, every Dyck path  $x \in D_n$  can be interpreted as a rooted tree with  $n$  edges; see [Sta15] and Figure 1. We therefore refer to the elements of  $D_n$  also as rooted trees. Given a rooted tree  $x$ , the rotation operation  $\text{rot}(x)$  shifts the root to the leftmost child of the root; see Figure 6. In terms of bitstrings, if  $x = 1u0v$  is the canonical decomposition of  $x$ , then  $\text{rot}(x) = u1v0$ .

*The middle levels graph  $G_n$ .* We describe our algorithm to compute a middle levels Gray code using the language of graph theory. We let  $G_n$  denote the *middle levels graph*, which has all bitstrings of length  $2n+1$  with weight  $n$  or  $n+1$  as vertices, with an edge between any two bitstrings that differ in exactly one bit. Clearly, computing a middle levels Gray code is equivalent to computing a Hamilton cycle in  $G_n$ . We let  $H_n$  denote the graph obtained by considering the subgraph of  $G_n$  induced by

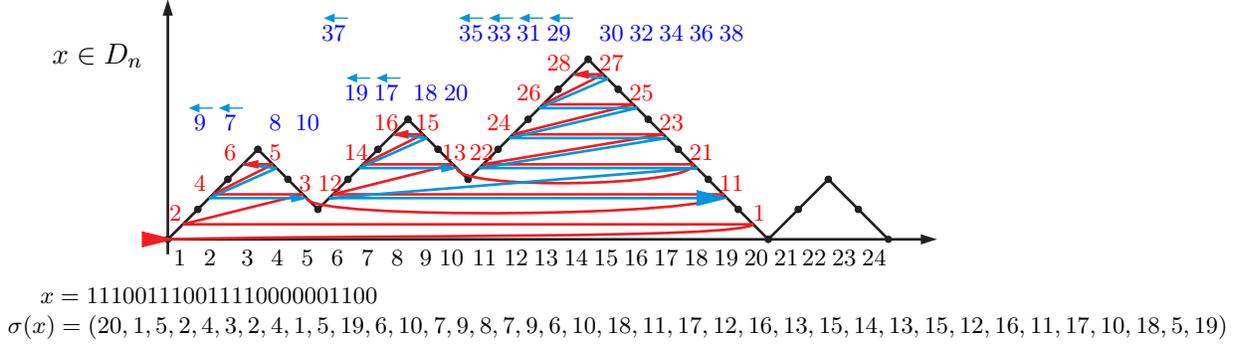


FIGURE 2. Illustration of the recursive computation of the flip sequence  $\sigma(x)$ . The number  $i = 1, 2, \dots, 38$  on the Dyck path indicates the position of  $x$  to be flipped in step  $i$ , i.e., the  $i$ th entry of  $\sigma(x)$ , where the little left-arrows act as modifiers that change the position drawn in the figure by  $-1$ . E.g., the 11th entry of  $\sigma(x)$  is the position of the  $\searrow$ -step of  $x$  marked with 11, so  $\sigma(11) = 19$ ; the 7th entry of  $\sigma(x)$  is the position of the  $\nearrow$ -step of  $x$  marked with 7 minus 1, so  $\sigma(7) = 3 - 1 = 2$ .

all vertices whose last bit equals 0, and by removing the last bit from every vertex. Note that  $G_n$  consists of a copy of  $H_n 0$ , a copy of  $\overline{\text{rev}}(H_n) 1$ , plus the matching  $M_n := \{(x0, x1) \mid x \in B_{2n,n}\}$ ; see Figure 5. The matching edges are the edges along which the last bit is flipped.

### 3. THE ALGORITHM

Our algorithm to compute a Hamilton cycle in the middle levels graph  $G_n$  consists of several nested functions, and in the following we explain these functions from bottom to top. The low-level functions compute paths in  $G_n$ , and the high-level functions combine them to a Hamilton cycle.

**3.1. Computing paths in  $H_n$ .** In this section we describe a set of disjoint paths  $\mathcal{P}_n$  that together visit all vertices of the graph  $H_n$ . The starting vertices of these paths are the vertices  $x \in D_n$ , and in the following we describe a rule  $\sigma(x)$  that specifies the sequence of bit positions to be flipped along the path starting at  $x$ . To compute the flip sequence  $\sigma(x)$  for a given vertex  $x \in D_n$ , we interpret  $x$  as a Dyck path, and we alternately flip  $\searrow$ -steps and  $\nearrow$ -steps of this Dyck path (corresponding to 0s and 1s in  $x$ , respectively). Specifically, for  $x \in D_n$  we consider the canonical decomposition  $x = 1u0v$  and define  $a := 1$ ,  $b := |u| + 2$  and

$$\sigma(x) := (b, a, \sigma_{a+1}(u)) , \quad (1a)$$

where  $\sigma_a(x')$  is defined for any substring  $x' \in D$  of  $x$  starting at position  $a$  in  $x$  by considering the canonical decomposition  $x' = 1u'0v'$ , by defining  $b := a + |u'| + 1$  and by recursively computing

$$\sigma_a(x') := \begin{cases} () & \text{if } |x'| = 0 , \\ (b, a, \sigma_{a+1}(u'), a - 1, b, \sigma_{b+1}(v')) & \text{otherwise .} \end{cases} \quad (1b)$$

In words, the sequence  $\sigma(x)$  defined in (1a) first flips the  $\searrow$ -step immediately after the subpath  $u$  (at position  $b$  of  $x$ ), then the  $\nearrow$ -step immediately before the subpath  $u$  (at position 1), and then recursively steps of  $u$ . No steps of  $v$  are flipped at all. The sequence  $\sigma_a(x')$  defined in (1b) first flips the  $\searrow$ -step immediately after the subpath  $u'$  (at position  $b$  of  $x$ ), then the  $\nearrow$ -step immediately before the subpath  $u'$  (at position  $a$ ), then recursively steps of  $u'$ , then again the step immediately to the left of  $x'$  (which is not part of  $x'$ , hence the index  $a - 1$ ), then again the step immediately to the right of  $u'$  (at position  $b$ ), and finally recursively steps of  $v'$ ; see Figure 2. The recursion  $\sigma(x)$  has a straightforward combinatorial interpretation: We consider the Dyck subpaths of the Dyck path  $x$

TABLE 1. Paths  $\mathcal{P}_n$  in  $H_n$  for  $n = 3$  obtained from the flip sequences  $\sigma(x)$ ,  $x \in D_n$ . The gray boxes highlight the (possibly empty) substrings of  $x$  corresponding to the subpaths  $u$  or  $v$  in the canonical decomposition of  $x$ .

First vertex $x \in D_n$	Flip sequence $\sigma(x)$	Last vertex $x \in D_n^-$
1 1100 0	(6, 1, 5, 2, 4, 3, 2, 4, 1, 5)	1010 01
1 1010 0	(6, 1, 3, 2, 1, 3, 5, 4, 3, 5)	1100 01
1 10010	(4, 1, 3, 2, 1, 3)	10 01 10
1 01100	(2, 1)	0 11100
1 01010	(2, 1)	0 11010

with increasing height levels and from left to right on each level, and we process them in two phases. In phase 1, we flip the steps of each such subpath alternatingly between the rightmost and leftmost step, moving upwards. In phase 2, we flip the steps alternatingly between the leftmost and rightmost step, moving downwards again. We emphasize here that during the recursive computation of  $\sigma(x)$ , no steps of  $x$  are ever flipped, but we always consider the same Dyck path and its Dyck subpaths as function arguments.

Note that by the definition (1), we have  $|\sigma(x)| = 2|u| + 2$ , where  $x = 1u0v$  is the canonical decomposition. We let  $P_\sigma(x)$  denote the sequence of vertices obtained by starting at the vertex  $x$  and flipping bits one after the other according to the sequence  $\sigma(x)$ . The following properties were proved in [GMN17, Proposition 2].

- (i) For any  $x \in D_n$ ,  $P_\sigma(x)$  is a path in the graph  $H_n$ . Moreover, all paths in  $\mathcal{P}_n := \{P_\sigma(x) \mid x \in D_n\}$  are disjoint, and together they visit all vertices of  $H_n$ .
- (ii) For any first vertex  $x \in D_n$ , considering the canonical decomposition  $x = 1u0v$ , the last vertex of  $P_\sigma(x)$  is given by  $u01v \in D_n^-$ . Consequently, the sets of first and last vertices of the paths  $\mathcal{P}_n$  are given by  $D_n$  and  $D_n^-$ , respectively.

Table 1 shows the paths  $\mathcal{P}_n$  for  $n = 3$ .

**3.2. Flippable pairs.** To compute a Hamilton cycle in the middle levels graphs  $G_n$ , we apply small local modifications to certain pairs of paths from  $\mathcal{P}_n$ , giving us additional freedom in combining an appropriate set of paths to a Hamilton cycle. Specifically, we say that  $x, y \in D_n$  form a *flippable pair*  $(x, y)$ , if

$$\begin{aligned} x &= 110w0v \ , \\ y &= 101w0v \end{aligned} \tag{2}$$

for some  $w, v \in D$ . In terms of rooted trees, the tree  $y$  is obtained from  $x$  by removing the pending edge that leads to the leftmost child of the leftmost child of the root, and to make it the new leftmost child of the root; see Figure 3 (recall the correspondence between Dyck paths and trees explained in Section 2). We denote this operation by  $\tau$  and write  $y = \tau(x)$ . The preimage  $T_n \subseteq D_n$  of the mapping  $\tau$  are all trees with  $n$  edges of the form  $x = 110w0v$ , and the image of  $\tau$  are all trees of the form  $y = 101w0v$ , where  $w, v \in D$ ; see Figure 6. Note that these two sets are disjoint.

Evaluating the recursion (1) for the bitstrings in a flippable pair  $(x, y)$  as in (2) yields

$$\begin{aligned} \sigma(x) &= (b, 1, 3, 2, 1, 3, \sigma_4(w)) \ , \\ \sigma(y) &= (2, 1) \ , \end{aligned}$$

where  $b := |w| + 2$ . It follows that the paths  $P_\sigma(x)$  and  $P_\sigma(y)$  intersect a common 6-cycle  $C_6(x, y)$  in the graph  $H_n$  as shown in Figure 4. Specifically, this 6-cycle can be encoded by

$$C_6(x, y) = 1**w*v \ , \tag{3}$$

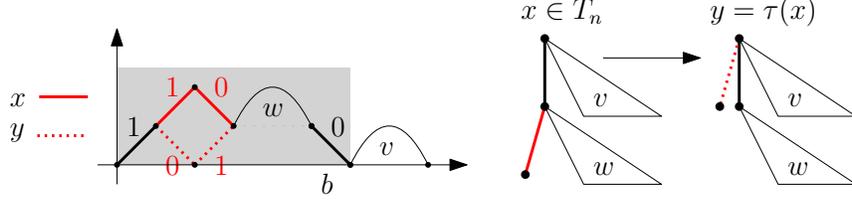


FIGURE 3. A flippable pair  $(x, y)$  and its Dyck path interpretation (left) and rooted tree interpretation (right).

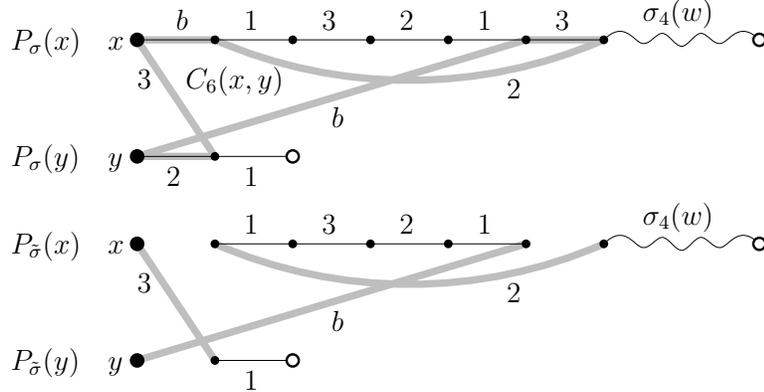


FIGURE 4. The paths  $P_\sigma(x)$  and  $P_\sigma(y)$  (black) for a flippable pair  $(x, y)$  traverse a common 6-cycle  $C_6(x, y)$  (gray). The symmetric difference yields paths  $P_{\tilde{\sigma}}(x)$  and  $P_{\tilde{\sigma}}(y)$  that have flipped end vertices. The numbers on the edges indicate the flipped bit positions.

where the six cycle vertices are obtained by substituting the three \*s by all six combinations of symbols from  $\{0, 1\}$  that use each symbol at least once. Consequently, taking the symmetric difference of the edge sets of  $P_\sigma(x)$  and  $P_\sigma(y)$  with the 6-cycle  $C_6(x, y)$  yields two paths on the same vertex set as  $P_\sigma(x)$  and  $P_\sigma(y)$ , but with interchanged end vertices. The resulting paths  $P_{\tilde{\sigma}}(x)$  and  $P_{\tilde{\sigma}}(y)$  have flip sequences

$$\begin{aligned} \tilde{\sigma}(x) &:= (3, 1) , \\ \tilde{\sigma}(y) &:= (b, 1, 2, 3, 1, 2, \sigma_4(w)) , \end{aligned} \quad (4)$$

and we refer to these two paths as *flipped paths* corresponding to the flippable pair  $(x, y)$ .

**3.3. The Hamilton cycle algorithm.** In this section we present our algorithm to compute a Hamilton cycle in the middle levels graph  $G_n$ . The Hamilton cycle is obtained by combining paths that are computed via the flip sequences  $\sigma$  and  $\tilde{\sigma}$ . We use the decomposition of  $G_n$  into  $H_n 0$ ,  $\overline{\text{rev}}(H_n) 1$ , plus the matching  $M_n$  discussed in Section 2; see Figure 5. By property (ii) from Section 3.1, the sets of first and last vertices of the paths  $\mathcal{P}_n$  are  $D_n$  and  $D_n^-$ , respectively. It is easy to see that these two sets are preserved under the mapping  $\overline{\text{rev}}$ . Together with property (i) from Section 3.1 it follows that

$$\mathcal{C}_n := \mathcal{P}_n 0 \cup \overline{\text{rev}}(\mathcal{P}_n) 1 \cup M'_n \quad (5)$$

with  $M'_n := \{(x_0, x_1) \mid x \in D_n \cup D_n^-\} \subseteq M_n$  is a 2-factor in the middle levels graph. A 2-factor in a graph is a collection of disjoint cycles that together visit all vertices of the graph. Note that along each of the cycles in the 2-factor, the paths from  $\mathcal{P}_n 0$  are traversed in forward direction, and the paths from  $\overline{\text{rev}}(\mathcal{P}_n) 1$  in backward direction. Observe also that the definition of flippable pairs given in Section 3.2 allows us to replace in the definition (5) any two paths  $P_\sigma(x)$  and  $P_\sigma(y)$

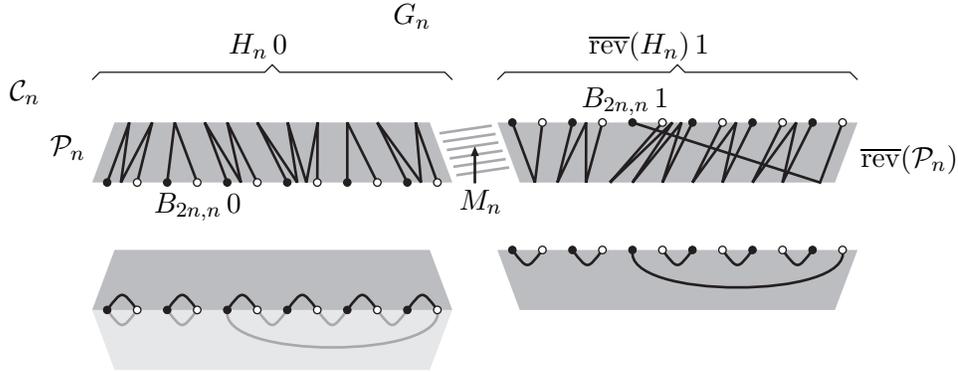


FIGURE 5. The top part shows the decomposition of the middle levels graph  $G_n$  and the definition (5). In this example, the 2-factor  $\mathcal{C}_n$  consists of three disjoint cycles that together visit all vertices of the graph. The sets  $D_n$  and  $D_n^-$  of first and last vertices of the paths  $\mathcal{P}_n$  and  $\overline{\text{rev}}(\mathcal{P}_n)$  are drawn in black and white, respectively. The bottom part shows a simplified drawing that depicts the structure of the 2-factor, which consists of two short cycles and one long cycle.

from  $\mathcal{P}_n$  for which  $(x, y)$  forms a flippable pair by the corresponding flipped paths  $P_{\overline{\sigma}}(x)$  and  $P_{\overline{\sigma}}(y)$ , yielding another 2-factor. Specifically, if the paths we replace lie on two different cycles, then the replacement will join the two cycles to one cycle. The final algorithm makes all those choices such that the resulting 2-factor consists only of a single cycle, i.e., a Hamilton cycle.

For the rest of the paper we will focus on proving that the algorithm `HAMCYCLE`, described in Algorithm 1, implies Theorem 3. `HAMCYCLE` is called with three input parameters:  $n$  determines the length  $2n + 1$  of the bitstrings,  $x$  is the starting vertex of the Hamilton cycle and must have weight  $n$  or  $n + 1$ , and  $\ell$  is the number of vertices to visit along the cycle. The variable  $y$  is the current vertex along the cycle, and the variable  $i$  counts the number of vertices that have already been visited. The calls `VISIT`( $y$ ) in lines H8, H11, H17 and H20 indicate where a function using our Hamilton cycle algorithm could perform further operations on the current vertex  $y$ . Each time a vertex along the cycle is visited, we increment  $i$  and check whether the desired number  $\ell$  of vertices has been visited (lines H9, H12, H18 and H21). We postpone the definition of the functions `INIT` and `ISFLIPTREE` called in lines H1 and H4 a little bit, and assume for a moment that the input vertex  $x$  of the middle levels graph  $G_n$  has the form  $x = z0$  with  $z \in D_n$ . In this case the variables  $y$  and  $i$  will be initialized to  $y := x$  and  $i := 1$  in line H1. Let us also assume that the return value of `ISFLIPTREE` called in line H4 is always `false`. With these simplifications, the algorithm `HAMCYCLE` computes exactly the 2-factor  $\mathcal{C}_n$  defined in (5) in the middle levels graph  $G_n$ . Indeed, one complete execution of the first for-loop corresponds to following one path from the set  $\mathcal{P}_n 0$  in the graph  $H_n 0$  starting at its first vertex and ending at its last vertex, and one complete execution of the second for-loop corresponds to following one path from the set  $\overline{\text{rev}}(\mathcal{P}_n) 1$  in the graph  $\overline{\text{rev}}(H_n) 1$  starting at its last vertex and ending at its first vertex. At the intermediate steps in lines H10 and H19, the last bit is flipped. These flips correspond to traversing an edge from the matching  $M'_n$ . The paths  $\mathcal{P}_n$  are computed in lines H5 and H14 using the recursion  $\sigma$ , and the resulting flip sequences are applied in the two inner for-loops (line H7 and H16). Note that if a path from the set  $\mathcal{P}_n$  has  $y \in D_n^-$  as a last vertex and if  $y = u01v$  is the canonical decomposition, then  $\overline{\text{rev}}$  maps the last vertex of the path  $P \in \mathcal{P}_n$  that has  $1 \overline{\text{rev}}(v) 0 \overline{\text{rev}}(u)$  as first vertex onto  $y$ . This is a consequence of property (ii) from Section 3.1, from which we obtain that the last vertex of  $P$  is  $\overline{\text{rev}}(v) 0 1 \overline{\text{rev}}(u)$ , and applying  $\overline{\text{rev}}$  to this vertex indeed yields  $y$ . From these observations and the definitions in lines H13, H14 and

---

**Algorithm 1:** HAMCYCLE( $n, x, \ell$ )

---

**Input:** An integer  $n \geq 1$ , a vertex  $x \in G_n$ , an integer  $\ell \geq 1$

**Result:** Starting from the vertex  $x$ , the algorithm visits the next  $\ell$  vertices on a Hamilton cycle in  $G_n$

```

H1  $(y, i) := \text{INIT}(n, x, \ell)$ 
H2 while true do
H3    $y^- := y_1 y_2 \dots y_{2n}$                                 /* ignore last bit of  $y$  */
H4   if  $((y^- \in T_n \text{ and } \text{ISFLIPTREE}(y^-)) \text{ or } (y^- \in \tau(T_n) \text{ and } \text{ISFLIPTREE}(\tau^{-1}(y^-)))$  then
       $s := \tilde{\sigma}(y^-)$                                        /* compute flip sequence  $\tilde{\sigma}$  ... */
H5   else  $s := \sigma(y^-)$                                    /* ... or  $\sigma$  */
H6   for  $j := 1$  to  $|s|$  do                                /* flip bits according to sequence  $s$  */
H7      $y_{s_j} := 1 - y_{s_j}$                                   /* flip bit at position  $s_j$  */
H8     VISIT( $y$ )
H9     if  $(i := i + 1) = \ell$  return
H10   $y_{2n+1} := 1$                                            /* flip last bit to 1 */
H11  VISIT( $y$ )
H12  if  $(i := i + 1) = \ell$  return
H13   $u01v := y_1 y_2 \dots y_{2n}$                             /* canonical decomposition of  $y_1 \dots y_{2n} \in D_n^-$  */
H14   $s := \sigma(1 \overleftarrow{\text{rev}}(v) 0 \overleftarrow{\text{rev}}(u))$               /* compute flip sequence  $\sigma$  */
H15  for  $j := |s|$  downto  $1$  do                          /* flip bits according to reverse sequence  $s$  */
H16     $y_{2n+1-s_j} := 1 - y_{2n+1-s_j}$                     /* flip bit at position  $2n+1-s_j$  */
H17    VISIT( $y$ )
H18    if  $(i := i + 1) = \ell$  return
H19   $y_{2n+1} := 0$                                            /* flip last bit to 0 */
H20  VISIT( $y$ )
H21  if  $(i := i + 1) = \ell$  return

```

---

H16 it follows that the paths in the second set on the right hand side of (5) are indeed traversed in backward direction (starting at the last vertex and ending at the first vertex).

We now explain the significance of the function ISFLIPTREE called in line H4 of our algorithm. This function interprets the current first vertex  $y^- \in T_n \subseteq D_n$  or  $\tau^{-1}(y^-) \in T_n$  as a rooted tree, and whenever it returns **true**, then instead of computing the flip sequence  $\sigma(y^-)$  in line H5, the algorithm computes the modified flip sequence  $\tilde{\sigma}(y^-)$  in line H4. Consequently, the function ISFLIPTREE controls which pairs of paths from  $\mathcal{P}_n$ , whose first vertices form a flippable pair, are replaced by the corresponding flipped paths, so that the resulting 2-factor in the middle levels graph  $G_n$  is a Hamilton cycle. Observe that these modifications only apply to the first set on the right hand side of (5), but not to the second set.

The function ISFLIPTREE therefore encapsulates the core ‘intelligence’ of our Hamilton cycle algorithm. We define this function and the function INIT in the next two sections. The correctness proof for the algorithm HAMCYCLE is provided in Section 4 below.

**3.4. The function ISFLIPTREE.** To define the Boolean function ISFLIPTREE, we need a few more definitions related to trees.

*Leaves, stars, and tree center.* Any vertex of degree 1 of a tree is called a *leaf*. In particular, the root of a rooted tree can be a leaf. We call a leaf *thin*, if its unique neighbor in the tree has degree 2.

A *star* is a tree in which all but at most one vertex are leaves. The *center* of a tree is the set of vertices that minimize the maximum distance to any other vertex. Any tree either has a unique center vertex, or two center vertices that are adjacent. All these notions apply to rooted trees, but also to abstract trees.

The following auxiliary function `ROOT` computes a canonically rooted version of a given rooted tree. Formally, for any tree  $x \in D_n$  and any integer  $r \geq 0$  the return value of `ROOT(rotr(x))` is the same rotated version of  $x$ . In the following functions, all comparisons between trees are performed using the bitstring representation.

*The function ROOT:* Given a tree  $x \in D_n$ , first compute its center vertex/vertices. If there are two centers  $c_1$  and  $c_2$ , then let  $x'$  be the tree obtained by rooting  $x$  so that  $c_1$  is the root and  $c_2$  its leftmost child, let  $x''$  be the tree obtained by rooting  $x$  so that  $c_2$  is the root and  $c_1$  its leftmost child, and return the tree from  $\{x', x''\}$  with the lexicographically smaller bitstring representation. If the center  $c$  is unique, then let  $y_1, y_2, \dots, y_k$  be the subtrees of  $x$  rooted at  $c$ . Consider the bitstring representations of these subtrees, and compute the lexicographically smallest rotation of the string  $(-1, y_1, -1, y_2, -1, \dots, -1, y_k)$  using Booth's algorithm [Boo80]. Here  $-1$  is an additional symbol that is lexicographically smaller than 0 and 1, ensuring that the lexicographically smallest string rotation starts at a tree boundary. Let  $\hat{x}$  be the tree obtained by rooting  $x$  at  $c$  such that the subtrees  $y_1, \dots, y_k$  appear exactly in this lexicographically smallest ordering, and return  $\hat{x}$ .

We are now ready to define the function `ISFLIPTREE`.

*The function ISFLIPTREE:* Given a tree  $x \in T_n$ , return `false` if  $x$  is a star. Otherwise compute  $\hat{x} := \text{ROOT}(x)$ . If  $x$  has a thin leaf, then let  $x'$  be the tree obtained by rotating  $\hat{x}$  until it has the form  $x' = 1100v$  for some  $v \in D$ . Return `true` if  $x = x'$ , and return `false` otherwise. If  $x$  has no thin leaf, then let  $x'$  be the tree obtained by rotating  $\hat{x}$  until it has the form  $x' = 1(10)^k 0v$  for some  $k \geq 2$  and  $v \in D$ . Return `true` if  $x = x'$  and if  $v = (10)^l$  implies that  $l \geq k$ , and return `false` otherwise.

**3.5. The function INIT.** It remains to define the function `INIT(n, x, ℓ)` called in line H1 of the algorithm `HAMCYCLE`. This function moves forward along the Hamilton cycle from the given vertex  $x$  in  $G_n$  and visits all vertices until the first vertex of the form  $z0$  with  $z \in D_n$  is encountered. We then initialize the current vertex as  $y := z0$ , and set the vertex counter  $i$  to the number of vertices visited along the cycle from  $x$  to  $y$ .

The first task is to compute, for the given vertex  $x$ , which path  $P_\sigma(y)0$  or  $\overline{\text{rev}}(P_\sigma(y))1$  the vertex  $x$  is contained in. With this information we can run essentially one iteration of the while-loop of the algorithm `HAMCYCLE`, after which we reach the first vertex of the form  $z0$  with  $z \in D_n$ .

If the last bit of  $x$  is 0, i.e.,  $x = z0$ , then we compute  $y \in D_n$  such that  $z \in P_\sigma(y)$  as follows: We consider the point(s) with lowest height on the lattice path  $z$ . If the lowest point of  $z$  is unique, then we partition  $z$  uniquely as

$$z = \begin{cases} u_1 0 u_2 0 \dots u_d 0 w 0 1 v_d 1 v_{d-1} 1 \dots v_1 1 v & \text{if } y \in B_{2n,n} \text{ ,} \\ u_1 0 u_2 0 \dots u_d 0 1 w 1 v_d 1 v_{d-1} 1 \dots v_1 1 v & \text{if } y \in B_{2n,n+1} \end{cases} \quad (6a)$$

for some  $d \geq 0$  and  $u_1, \dots, u_d, v_1, \dots, v_d, w \in D$ . If there are at least two lowest points of  $y$ , then we partition  $y$  uniquely as

$$z = \begin{cases} u_1 0 u_2 0 \dots u_d 0 1 w 0 v_d 1 v_{d-1} 1 \dots v_1 1 v & \text{if } y \in B_{2n,n} \text{ ,} \\ u_1 0 u_2 0 \dots u_d 1 w 0 1 v_d 1 v_{d-1} 1 \dots v_1 1 v & \text{if } y \in B_{2n,n+1} \end{cases} \quad (6b)$$

for some  $d \geq 0$  and  $u_1, \dots, u_d, v_1, \dots, v_d, w \in D$ . In all four cases, a straightforward calculation using the definition (1) shows that

$$y := 1 u_1 1 u_2 \dots 1 u_d 1 w 0 v_d 0 v_{d-1} 0 \dots v_1 0 v$$

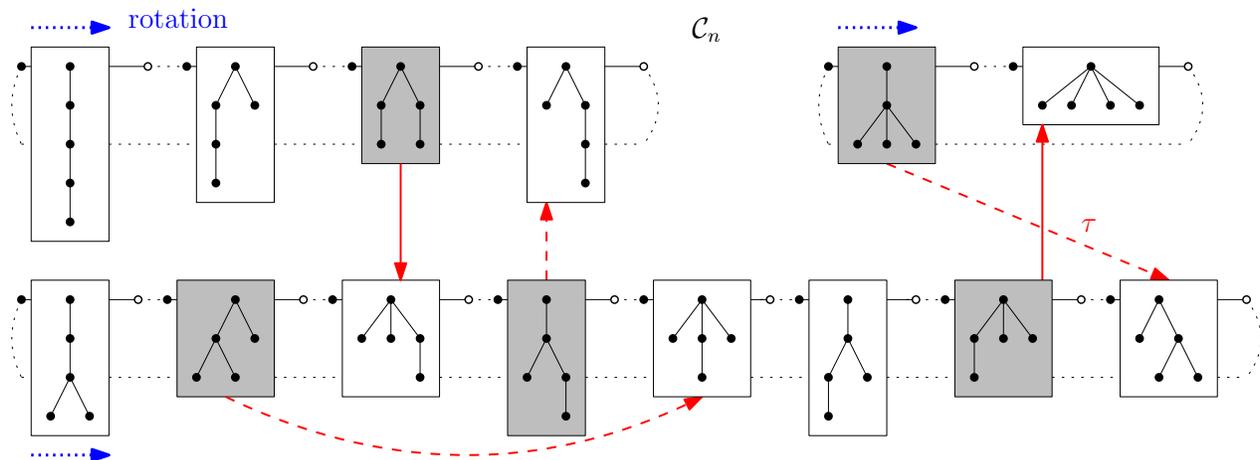


FIGURE 6. Tree rotation along the cycles of the 2-factor  $\mathcal{C}_n$ , the mapping  $\tau$  and the graph  $\mathcal{T}_n$  for  $n = 4$ . The highlighted trees form the set  $T_n$ , the preimage of  $\tau$ . The graph  $\mathcal{T}_n$  has the three cycles of  $\mathcal{C}_n$  as nodes, and it contains only the two solid directed edges (not the dashed ones), forming a spanning tree on the cycles of  $\mathcal{C}_n$ .

is indeed the first vertex of the path  $P_\sigma(y)$  that contains to the vertex  $z$ . In particular, if  $z \in D_n$ , then  $d = 0$  and  $y = z = 1w0v$ .

If the last bit of  $x$  is 1, i.e.,  $x = z1$ , then we compute  $y \in D_n$  such that  $z \in \overline{\text{rev}}(P_\sigma(y))$  by applying the previous steps to the vertex  $\overline{\text{rev}}(z)$ .

For more details how the function INIT is implemented, see our C++ implementation [www].

#### 4. CORRECTNESS OF THE ALGORITHM

The properties (i) and (ii) of the paths  $\mathcal{P}_n$  claimed in Section 3.1 were proved in [GMN17, Proposition 2]. Consequently, if we assume for a moment that the function ISFLIPTREE always returns **false**, then the arguments given in Section 3.3 show that the algorithm HAMCYCLE correctly computes the 2-factor  $\mathcal{C}_n$  defined in (5) in the middle levels graph  $G_n$ . Moreover, by the symmetric definition in line H4, for any flippable pair  $(x, y)$  either the flip sequence  $\sigma$  is applied to both  $x$  and  $y$ , or the modified flip sequence  $\tilde{\sigma}$  is applied to both  $x$  and  $y$ . Consequently, by the definition of flippable pairs given in Section 3.2, the algorithm HAMCYCLE correctly computes a 2-factor in the graph  $G_n$  for *any* Boolean function ISFLIPTREE on the set  $T_n$ . It remains to argue that for the particular Boolean function ISFLIPTREE defined in the previous section, our algorithm indeed computes a Hamilton cycle. For this we analyze the structure of the 2-factor  $\mathcal{C}_n$ , which is best described by yet another family of trees.

*Plane trees.* A *plane tree* is a tree with a specified cyclic ordering of the neighbors of each vertex. We think of a plane tree as a tree embedded in a plane, where the neighbors of each vertex appear exactly in the specified ordering in counterclockwise direction around that vertex; see Figure 7. Equivalently, plane trees arise as equivalence classes of rooted trees under rotation.

It was shown in [GMN17, Proposition 2] that for any cycle from  $\mathcal{C}_n$ , if we consider two consecutive vertices of the form  $x0$  and  $y0$  with  $x, y \in D_n$  on the cycle and the canonical decomposition of  $x = 1u0v$ , then we have  $y = u1v0$ . In terms of rooted trees, we have  $y = \text{rot}(x)$ . Consequently, the set of cycles of  $\mathcal{C}_n$  is in bijection with the equivalence classes of all rooted trees with  $n$  edges under rotation; see Figure 6. In particular, the number of cycles in the 2-factor is given by the number of plane trees with  $n$  edges.

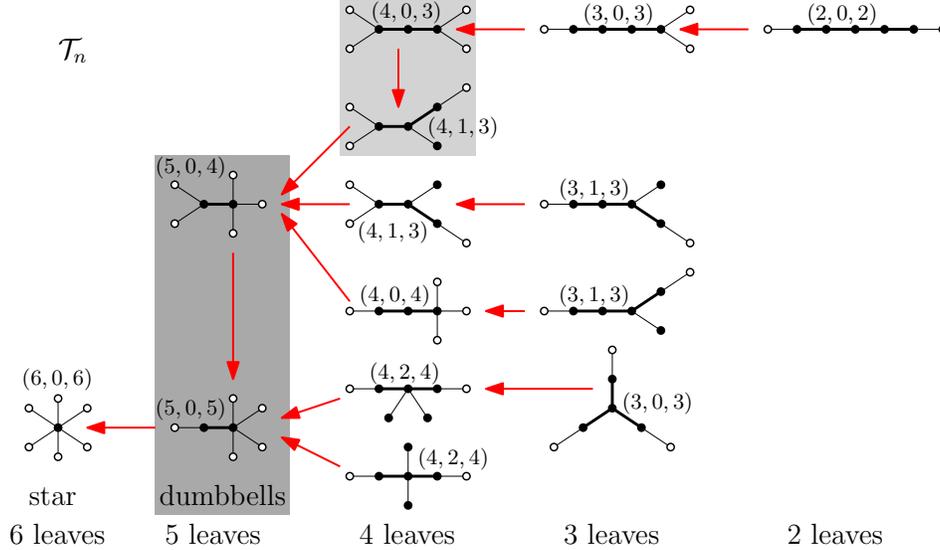


FIGURE 7. The graph  $\mathcal{T}_n$  for  $n = 6$ , where only the plane trees corresponding to each node are shown. The skeleton of each plane tree is drawn with bold edges. The terminal leaves, and the black leaves are non-terminal leaves. The triples assigned to each plane tree are the signatures (number of leaves, number of non-terminal leaves, maximum degree). The light-gray box highlights an edge of  $\mathcal{T}_n$  along which the number of leaves stays the same, but the number of non-terminal leaves goes up from 0 to 1. The dark-gray box highlights all dumbbells and an edge of  $\mathcal{T}_n$  along which the number of leaves and non-terminal leaves stays the same, but the maximum degree goes up from 4 to 5.

The definition of flippable pairs  $(x, y)$  given in Section 3.2 shows that if  $P_\sigma(x)$  and  $P_\sigma(y)$  are contained in two distinct cycles of  $\mathcal{C}_n$ , then replacing these two paths by the flipped paths  $P_{\bar{\sigma}}(x)$  and  $P_{\bar{\sigma}}(y)$  joins the two cycles to a single cycle on the same set of vertices; recall Figure 4. As mentioned before, this replacement operation is equivalent to taking the symmetric difference of the edge sets of  $P_\sigma(x)$  and  $P_\sigma(y)$  with the 6-cycle  $C_6(x, y)$  defined in (3). The following two properties were established in [GMN17, Proposition 3]: For any flippable pairs  $(x, y)$  and  $(x', y')$ , the 6-cycles  $C_6(x, y)$  and  $C_6(x', y')$  are edge-disjoint. Moreover, 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  $P_\sigma(x)$  are not interleaved, but one pair appears before the other pair along the path. Consequently, none of the 6-cycles used for these joining operations interfere with each other.

Using these observations, we define an auxiliary graph  $\mathcal{T}_n$  as follows. The nodes of  $\mathcal{T}_n$  are the equivalence classes of rooted trees with  $n$  edges under rotation, which can be interpreted as plane trees. For any flippable pair  $(x, y)$  for which  $\text{IsFLIPTREE}(x) = \text{true}$ , we add a directed edge from the equivalence class containing  $x$  to the equivalence class containing  $y = \tau(x)$  to the graph  $\mathcal{T}_n$ . The graph  $\mathcal{T}_n$  is shown in see Figure 6 and Figure 7 for  $n = 4$  and  $n = 6$ , respectively. By what we said before, the nodes of  $\mathcal{T}_n$  correspond to the cycles of the 2-factor  $\mathcal{C}_n$ , and the edges correspond to the flipped pairs of paths used by the algorithm  $\text{HAMCYCLE}$ . To complete the correctness proof for our algorithm, it thus remains to prove the following lemma.

**Lemma 4.** *For any  $n \geq 1$ , the graph  $\mathcal{T}_n$  is a spanning tree.*

*Proof.* For the reader's convenience, the following definitions are illustrated in Figure 7. All these notions apply to rooted trees and to plane trees. The *skeleton* of a tree is the tree obtained by removing all leaves. A leaf of a tree is called *terminal*, if it is adjacent to a leaf of the skeleton.

A *dumbbell* is a tree in which all but exactly two vertices are leaves. Equivalently, a dumbbell has a skeleton consisting of a single edge, or a dumbbell is a tree with  $n$  edges and  $n - 1$  leaves. Note that any tree that is not a star has a skeleton with at least one edge, and any tree that is neither a star nor a dumbbell has a skeleton with at least two edges.

Consider a directed edge in the graph  $\mathcal{T}_n$  that arises from two rooted trees  $(x, y)$  with  $y = \tau(x)$  and  $\text{ISFLIPTREE}(x) = \text{true}$ . By the definition of the function  $\text{ISFLIPTREE}$ ,  $x$  is not a star, so in particular we have  $n \geq 3$ . Moreover, exactly one of the following two conditions holds. Case (a):  $x = 1100v$  for some  $v \in D$ . In this case, we have  $y = \tau(x) = 1010v$ , and as  $v$  is nonempty by the condition  $n \geq 3$ , we obtain that  $y$  has one more leaf than  $x$ . Case (b):  $x = 1(10)^k 0v$  for some  $k \geq 2$  and  $v \in D$ , and  $v = (10)^l$  implies that  $l \geq k$ . In this case, as  $x$  is not a star, the subtree  $v$  has at least one edge, so the root of  $x$  is not a leaf. Moreover, all leaves of  $x$  in the leftmost subtree are terminal leaves. We distinguish two subcases. Case (b1):  $v$  is not a star rooted at the center, i.e.,  $x$  is not a dumbbell. In this case, one of the terminal leaves in the leftmost subtree of  $x$  becomes a non-terminal leaf in  $y = \tau(x) = 101(10)^{k-1}0v$ , so  $y$  and  $x$  have the same number of leaves, but  $y$  has one more non-terminal leaf. Case (b2):  $v$  is a star rooted at the center, i.e.,  $x$  is a dumbbell and  $v = (10)^l$  for some  $l \geq k$ . In this case, the root of  $x$  is a vertex that has maximum degree  $l + 1$ , whereas the root of the dumbbell  $y = \tau(x) = 101(10)^{k-1}0(10)^l$  has degree  $l + 2$ , so both dumbbells  $y$  and  $x$  have the same number of leaves and non-terminal leaves, but the maximum degree of  $y$  is one higher than that of  $x$ .

To every plane tree  $T$  we therefore assign a *signature*  $s(T) := (\ell, t, d)$ , where  $\ell$  is the number of leaves of  $T$ ,  $t$  is the number of non-terminal leaves of  $T$ , and  $d$  is the maximum degree of  $T$ ; see Figure 7. By our observations from before, for any directed edge  $(T, T')$  between two plane trees  $T$  and  $T'$  in  $\mathcal{T}_n$ , when comparing  $s(T) := (\ell, t, d)$  and  $s(T') := (\ell', t', d')$ , then either  $\ell < \ell'$  in case (a) from before, or  $\ell = \ell'$  and  $t < t'$  in case (b1) from before, or  $(\ell, t) = (\ell', t')$  and  $d < d'$  in case (b2) from before. It follows that  $\mathcal{T}_n$  does not have any cycles where all edges have the same orientation. In particular,  $\mathcal{T}_n$  has no loops.

Moreover, as a consequence of the initial canonical rooting performed by the call to  $\text{ROOT}$ , the function  $\text{ISFLIPTREE}$  returns  $\text{true}$  for at most one tree from each equivalence class of rooted trees under rotation. This implies that each node of  $\mathcal{T}_n$  has out-degree at most 1. Consequently,  $\mathcal{T}_n$  does not have a cycle where the edges have different orientations, as such a cycle would have a node with out-degree 2. Combining these observations shows that  $\mathcal{T}_n$  is acyclic.

It remains to prove that  $\mathcal{T}_n$  is connected. For this we show how to move from any plane tree  $T$  along the edges of  $\mathcal{T}_n$  to the star with  $n$  edges; see Figure 7. We assume that  $T$  is not the star with  $n$  edges, so in particular  $n \geq 3$ . If  $T$  has a thin leaf, then we can clearly root  $T$  so that the rooted tree has the form  $1100v$  for some  $v \in D$ . Consequently, there exists an edge in  $\mathcal{T}_n$  that leads from  $T$  to a tree  $T'$  that has one more leaf than  $T$ . If  $T$  has no thin leaf, then we can root  $T$  so that the rooted tree has the form  $1(10)^k 0v$  for some  $k \geq 2$  and  $v \in D$ , and  $v = (10)^l$  implies that  $l \geq k$ . To see this we distinguish two cases. If  $T$  is not a dumbbell, then the skeleton of  $T$  has at least two edges, so rooting  $T$  at any vertex in distance 1 from a leaf of the skeleton yields a rooted tree of the form  $1(10)^k 0v$  for some  $k \geq 2$  and  $v \in D$  where  $v$  is not a star rooted at the center. On the other hand, if  $T$  is a dumbbell, then we can root  $T$  at a vertex of maximum degree so that the rooted tree has the form  $1(10)^k v$  for some  $k \geq 2$  and  $v = (10)^l$  with  $l \geq k$ . Consequently, there exists an edge in  $\mathcal{T}_n$  that leads from  $T$  to a tree  $T'$  that has the same number of leaves, and either one more non-terminal leaf, or the same number of non-terminal leaves, but a maximum degree that is one higher than that of  $T$ . We repeat this argument, following directed edges of  $\mathcal{T}_n$ , until we arrive at the star with  $n$  edges.

We have shown that  $\mathcal{T}_n$  is acyclic and connected, so it is indeed a spanning tree.  $\square$

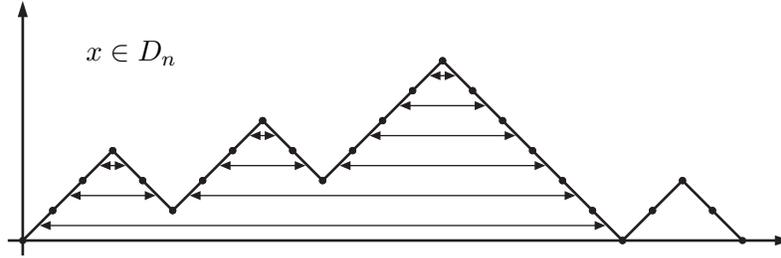


FIGURE 8. Auxiliary pointers to compute the flip sequence  $\sigma(x)$  for any  $x \in D_n$  in time  $\mathcal{O}(n)$ . The lattice path  $x$  is the same as in Figure 2, and the resulting sequence  $\sigma(x)$  is shown in that figure.

## 5. RUNNING TIME AND SPACE REQUIREMENTS

**5.1. Running time.** For any  $x \in D_n$ , the flip sequence  $\sigma(x)$  can be computed in linear time. To achieve this, we precompute an array of bidirectional pointers below the Dyck subpaths of  $x$  between corresponding pairs of an  $\nearrow$ -step and  $\searrow$ -step on the same height; see Figure 8. Using these pointers, each canonical decomposition operation encountered in the recursion (1) can be performed in constant time, so that the overall running time of the recursion is  $\mathcal{O}(n)$ . Clearly, the sequence  $\tilde{\sigma}(x)$  can also be computed in time  $\mathcal{O}(n)$  by modifying the sequence  $\sigma(x)$  as described in (4) in constantly many positions. Obviously, the functions  $\overline{\text{rev}}(u)$  and  $\overline{\text{rev}}(v)$  called in line H14 can also be computed in time  $\mathcal{O}(n)$ .

To compute the functions `ROOT` and `ISFLIPTREE`, we first convert the given bitstring  $x \in D_n$  to a tree in adjacency list representation, which can clearly be done in time  $\mathcal{O}(n)$ ; recall the correspondence between Dyck paths and rooted trees from Figure 1. The adjacency list representation allows us to compute the center vertex/vertices in linear time by removing leaves in rounds until only a single vertex or a single edge is left (see [Ski08]). Moreover, it allows us to perform each rotation operation `rot` in constant time, and a full tree rotation in time  $\mathcal{O}(n)$ . Booth's algorithm to compute the lexicographically smallest string rotation also runs in linear time [Boo80]. This gives the bound  $\mathcal{O}(n)$  for the time spent in the functions `ROOT` and `ISFLIPTREE`.

It was shown in [GMN17, Proposition 2] that the distance between any two consecutive vertices of the form  $x_0$  and  $y_0$  with  $x, y \in D_n$  on a cycle of the 2-factor  $\mathcal{C}_n$  is exactly  $4n + 2$ . Moreover, replacing a path  $P_\sigma(x)$  in the first set on the right hand side of (5) by the path  $P_{\tilde{\sigma}}(x)$  does not change this distance; recall Figure 4. It follows that in each iteration of the while-loop of our algorithm `HAMCYCLE`, exactly  $4n + 2$  vertices are visited. Combining this with the time bounds  $\mathcal{O}(n)$  derived for the functions  $\sigma$ ,  $\tilde{\sigma}$ ,  $\overline{\text{rev}}$  and `ISFLIPTREE` that are called once or twice during each iteration of the while-loop, we conclude that the while-loop takes time  $\mathcal{O}(\ell + n)$  to visit  $\ell$  vertices of the Hamilton cycle.

The function `INIT` takes time  $\mathcal{O}(n)$ , as the partitions (6) can be computed in linear time, and as we visit at most linearly many vertices in this function (every path in  $\mathcal{P}_n$  has only length  $\mathcal{O}(n)$ ).

Combining the time bounds  $\mathcal{O}(n)$  for the initialization phase and the time  $\mathcal{O}(\ell + n)$  spent in the while-loop, we obtain the claimed overall bound  $\mathcal{O}(\ell + n)$  for the algorithm `HAMCYCLE`.

**5.2. Space requirements.** Throughout our algorithm, we only store constantly many bitstrings of length  $2n$  and rooted trees with  $n$  edges, proving that the entire space needed is  $\mathcal{O}(n)$ .

## REFERENCES

- [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.
- [Boo80] K. S. Booth. Lexicographically least circular substrings. *Inform. Process. Lett.*, 10(4-5):240–242, 1980.
- [BW84] M. Buck and D. Wiedemann. Gray codes with restricted density. *Discrete Math.*, 48(2-3):163–171, 1984.
- [Cum66] R. L. Cummins. Hamilton circuits in tree graphs. *IEEE Trans. Circuit Theory*, CT-13:82–90, 1966.
- [Der75] N. Dershowitz. A simplified loop-free algorithm for generating permutations. *Nordisk Tidskr. Informationsbehandling (BIT)*, 15(2):158–164, 1975.
- [DG12] P. Diaconis and R. Graham. *Magical mathematics*. Princeton University Press, Princeton, NJ, 2012. The mathematical ideas that animate great magic tricks, With a foreword by Martin Gardner.
- [DKS94] D. A. Duffus, H. A. Kierstead, and H. S. Snevily. An explicit 1-factorization in the middle of the Boolean lattice. *J. Combin. Theory Ser. A*, 65(2):334–342, 1994.
- [DSW88] D. Duffus, B. Sands, and R. Woodrow. Lexicographic matchings cannot form Hamiltonian cycles. *Order*, 5(2):149–161, 1988.
- [Ehr73] G. Ehrlich. Loopless algorithms for generating permutations, combinations, and other combinatorial configurations. *J. Assoc. Comput. Mach.*, 20:500–513, 1973.
- [EHR84] P. Eades, M. Hickey, and R. C. Read. Some Hamilton paths and a minimal change algorithm. *J. Assoc. Comput. Mach.*, 31(1):19–29, 1984.
- [EM84] P. Eades and B. McKay. An algorithm for generating subsets of fixed size with a strong minimal change property. *Inform. Process. Lett.*, 19(3):131–133, 1984.
- [FT95] S. Felsner and W. T. Trotter. Colorings of diagrams of interval orders and  $\alpha$ -sequences of sets. *Discrete Math.*, 144(1-3):23–31, 1995. Combinatorics of ordered sets (Oberwolfach, 1991).
- [GM18] P. Gregor and T. Mütze. Trimming and gluing Gray codes. *Theoret. Comput. Sci.*, 714:74–95, 2018.
- [GMN17] P. Gregor, T. Mütze, and J. Nummenpalo. A short proof of the middle levels theorem. *arXiv:1710.08249*, Oct 2017.
- [Gow17] W. T. Gowers. Probabilistic combinatorics and the recent work of Peter Keevash. *Bull. Amer. Math. Soc. (N.S.)*, 54(1):107–116, 2017.
- [Gra53] F. Gray. Pulse code communication, 1953. March 17, (filed Nov. 1947). U.S. Patent 2,632,058.
- [GŠ10] P. Gregor and R. Škrekovski. On generalized middle-level problem. *Inform. Sci.*, 180(12):2448–2457, 2010.
- [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.
- [HH72] C. A. Holzmann and F. Harary. On the tree graph of a matroid. *SIAM J. Appl. Math.*, 22:187–193, 1972.
- [HKRR05] P. Horák, T. Kaiser, M. Rosenfeld, and Z. Ryjáček. The prism over the middle-levels graph is Hamiltonian. *Order*, 22(1):73–81, 2005.
- [HR16] F. Herter and G. Rote. Loopless Gray code enumeration and the tower of Bucharest. In *8th International Conference on Fun with Algorithms, FUN 2016, June 8-10, 2016, La Maddalena, Italy*, pages 19:1–19:19, 2016.
- [Joh63] S. M. Johnson. Generation of permutations by adjacent transposition. *Math. Comp.*, 17:282–285, 1963.
- [Joh04] J. R. Johnson. Long cycles in the middle two layers of the discrete cube. *J. Combin. Theory Ser. A*, 105(2):255–271, 2004.
- [Kam67] T. Kamae. The existence of a Hamilton circuit in a tree graph. *IEEE Trans. Circuit Theory*, CT-14:279–283, 1967.
- [Knu11] D. E. Knuth. *The Art of Computer Programming. Vol. 4A. Combinatorial Algorithms. Part 1*. Addison-Wesley, Upper Saddle River, NJ, 2011.
- [KT88] H. A. Kierstead and W. T. Trotter. Explicit matchings in the middle levels of the Boolean lattice. *Order*, 5(2):163–171, 1988.
- [LRvBR93] J. M. Lucas, D. Roelants van Baronaigien, and F. Ruskey. On rotations and the generation of binary trees. *J. Algorithms*, 15(3):343–366, 1993.
- [Luc87] J. M. Lucas. The rotation graph of binary trees is Hamiltonian. *J. Algorithms*, 8(4):503–535, 1987.
- [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, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 2238–2253, 2017.
- [MN18] T. Mütze and J. Nummenpalo. Efficient computation of middle levels Gray codes. To appear in *ACM Transaction on Algorithms*. An extended abstract appeared in the Proceedings of the European Symposium on Algorithms (ESA 2015), 2018.
- [MS17] T. Mütze and P. Su. Bipartite Kneser graphs are Hamiltonian. *Combinatorica*, 37(6):1207–1219, 2017.
- [MSW18] T. Mütze, C. Standke, and V. Wiechert. A minimum-change version of the Chung-Feller theorem for Dyck paths. *European J. Combin.*, 69:260–275, 2018.

- [Müt16] T. Mütze. Proof of the middle levels conjecture. *Proc. London Math. Soc.*, 112(4):677–713, 2016.
- [MW12] T. Mütze and F. Weber. Construction of 2-factors in the middle layer of the discrete cube. *J. Combin. Theory Ser. A*, 119(8):1832–1855, 2012.
- [NW75] A. Nijenhuis and H. S. Wilf. *Combinatorial algorithms*. Academic Press, New York-London, 1975. Computer Science and Applied Mathematics.
- [Rus88] F. Ruskey. Adjacent interchange generation of combinations. *J. Algorithms*, 9(2):162–180, 1988.
- [SA11] M. Shimada and K. Amano. A note on the middle levels conjecture. *arXiv:0912.4564*, September 2011.
- [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.
- [Sed77] R. Sedgewick. Permutation generation methods. *Comput. Surveys*, 9(2):137–164, 1977.
- [Ski08] S. Skiena. *The Algorithm Design Manual (2. ed.)*. Springer, 2008.
- [SSS09] I. Shields, B. Shields, and C. D. Savage. An update on the middle levels problem. *Discrete Math.*, 309(17):5271–5277, 2009.
- [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.
- [Tro62] H. Trotter. Algorithm 115: Perm. *Commun. ACM*, 5(8):434–435, August 1962.
- [Wil89] H. S. Wilf. *Combinatorial algorithms: an update*, volume 55 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1989.
- [Win04] P. Winkler. *Mathematical puzzles: a connoisseur's collection*. A K Peters, Ltd., Natick, MA, 2004.
- [www] currently <http://www.math.tu-berlin.de/~muetze> and <http://people.inf.ethz.ch/njerri>.
- [Zer85] D. Zerling. Generating binary trees using rotations. *J. Assoc. Comput. Mach.*, 32(3):694–701, 1985.