1 Introduction

Let Π be a *circular sequence* of $\{1,2,\ldots,n\}$. It means, $\Pi = \left(\pi_0,\ldots,\pi_{\binom{n}{2}}\right)$ is a sequence of permutations of $\{1,2,\ldots,n\}$ such that π_0 is the identity permutation $\{1,2,\ldots,n\}$ and $\pi_{\binom{n}{2}}$ is the reverse permutation $\{n,n-1,\ldots,1\}$. Any two consecutive permutations differ by exactly one transposition of two elements in adjacent positions and every these transpositions have defined weight determined by the position of swaped elements. We are interested in case, when the weight in the middle (for even n) is zero and increases by one (na obe strany)– figure (OBRAZEK). Our aim is finding circular sequence with the smallest cost, where the cost of circular sequence is sum of weights its permutations.

Circular sequences are also connected with geometrical problems. nasi volby vah se tyka nasledujici:

Let S be a set of n points in general position in the plane. Every pair of points a, b of this set determine a line, which divides S into two parts, let call them S_1 and S_2 (FIGURE). Now define difference of configuration S as:

$$D(S) = \sum_{a,b\in S, a\neq b} |s_1 - s_2|,$$

where s_1 resp. s_2 denote cardinality of S_1 resp. S_2 (sometimes we will use next notation $D(ab) = |s_1 - s_2|$). In other words, it is sum of differences between the numbers of points in the two open halfplanes determined by the line ab (in particular, if the configuration has an even number of points, then the difference of each pair of point is even). Our aim is determine the following function

$$F(n) = \min_{|S|=n} D(S).$$

(opet radsi i slovne)

The connection between geometrical difference problem and circular sequences is folowing: Let S be configuration of n points as above. Let L be directed line that is not orthogonal to any of the lines defined by pairs of points in S. We label the points in S as $p_1, p_2, \ldots, p_n$, according to the order in which their orthogonal projctions appear along L. Now suppose that we start rotating L counterclockwise, say. The ordering of the projections changes exactly at the positions where L passes through a position orthogonal to the line defined by some pair of points a, b in S. At the time the projection change occurs, a and b are adjacent in the ordering and the ordering changes by transposing a and b. Thus, if we keep track of all permutations of the projections as the line L is rotated by 180°, we obtain a circular sequence Π.

Now we show, that cost of circular sequence get in this way is $\frac{1}{4}D(S)$. While we are rotating line L, in one moment projections of two adjacent points a a b are same and this projection is exactly on intersection of lines L and ab. Let m_l be number of points on the left from a and b in this projection, m_r number of points on the right and m minimum of m_l and m_r. Number of points on the other side is exactly $n - m - 2$, so $D(ab) = n - 2m - 2$. On the other side, weight of transposition a and b is $\frac{n}{2} - k - 1$ (ie. $\frac{D(ab)}{2}$), because transposition on the most left (right) position is $\frac{n}{2} - 1$ and weight of swap a with b are decreased by exactly m numbers before (after). So we can get cost of Π simply by summing over all lines between two points in S and divide it by 2. In $D(S)$ we count every line exactly twice and this implies claimed relation between $D(S)$ and cost of its circular sequence.

Upper bound for function $F(n)$

This upper bound follows from known mercedes construction of many halving lines.

popis mercedese:

The mercedes configuration M is defined in recursive way. First we have four points in non-convex position. We scaled them and triplied (ztrojime) (precizneji) So we have configuration of $4 \cdot 3^k$ points (in k step). Let $n := 4 \cdot 3^{k-1}$ be size of each ot the three parts of the trefoil. Now we want to compute difference of this mercedes configuration.

First we express difference of configuration C of n points in next way

$$D(C) := \sum_{a \in C} \sum_{b \in C, b \neq a} D(ab),$$

where $D(ab)$ is difference of pair of points a, b mentioned above. In particular, every line contribute twice to total difference of the configuration. Mercedes configuration yields the reccurence:

$$D(M_k) = 3 \cdot D(M_{k-1}) + 6 \cdot \sum_{i=1}^{n} \sum_{j=1}^{n} (|2i + 2j - 2 - 3n|),$$

where M_k is mercedes configuration after k steps. Difference of M_0 (difference of the non-convex four-point configuration) is 12. FIGURE The double sum $6 \cdot \sum_{i=1}^{n} \sum_{j=1}^{n} (|2i + 2j - 2 - 3n|)$ in the reccurence above simplifies to the cubic polynomial $\frac{13}{2}n^3 - 2n$, so the reccurence can be expressed as

$$D(M_k) = 3 \cdot D(M_{k-1}) + \frac{13}{2}n^3 - 2n.$$

Substituting $N := 3n = 4 \cdot 3^k$ for the total number of points (i.e. M_k is mercedes configuration on N points), here is a closed-form solution for the difference of an N-point mercedes configuration M_k. For simplicity denote $D(N) := D(M_k)$. Then $D(N) = 3 \cdot D(\frac{N}{3}) + \frac{13}{54}N^3 - \frac{2}{3}N$. From computation we derive that the leading term is $\frac{13}{48}$, which corresponds after appropriate scaling to the $\frac{13}{24}k^3$ upper bound for the cost of an algorithm reversing $2k$ points with weights $k-1, k-2, \dots, -1, 0, 1, \dots, k-1$.

2 Preliminaries

2.1 The Problem

Given n and a non-negative *cost-vector* w of length $n-1$ we seek the cheapest sequence of steps reversing the permutation $(1, 2, \dots n)$.

A *step* is a tarnsposition of two elements on neighbour positions. The *cost* of a single step swapping positions i and $i+1$ is w_i.

Any sequence of steps reversing the permutation $(1, 2, \dots n)$ is called a *startegy* and we denote it by $\mathcal{S}$. The total cost of $\mathcal{S}$ is the sum of individual costs of all the steps. We denote the total cost $\mathtt{c}_w(\mathcal{S})$.

We always assume the cost-vector to be non-negative, [1] as any negative element of w would make to be the optimum $-\infty$. We call a cost-vector w *symmetric* if $w_i = w_{n-i}$ for all $i = 1 \dots n-1$.

In the following sections we consider the given problem under various assumptions on the cost-vectors.

Definition 1. Let $\mathtt{opt}_w(n)$ be the cost of the cheapest strategy reversing $(1, 2 \dots n)$ with costs-vector w. Note that there is always a reversing strategy of finite cost and that this minimum is well-defined when w is non-negative.

2.2 Optimal Strategies

In this section we consider some general properties of strategies and the function `opt`. The first lemma allows the restriction to shortest strategies of known length for any cost-vector w.

Lemma 2 (Swapping Lemma)**.** For any non-negative cost-vector w and any strategy $\mathcal{S}$ where two elements a and b swap more than once there is a shorter strategy $\mathcal{S}'$ where a and b swap exactly once and the cost of $\mathcal{S}'$ is at most the cost of $\mathcal{S}$.

[1] Later we show that forcing the solution to be also shortest would solve this problem, but then the problem can be converted to one with non-negative costs. **ODKAZ**

Proof. Every two elements have to swap odd number of times as they have to end up in a reversed order. If we have an optimal strategy where some two elements a and b swap more than once, we can ommit two of these swaps and still get the same result. $\square$

Corollary 3. There is always an optimal strategy of length $\binom{n}{2}$. This strategy swaps every two elements exactly once. If w is positive, then every optimal strategy has this length.

Proof. No strategy can be shorter as the resulting inverse permutation has exactly $\binom{n}{2}$ inversions and every step adds at most one inversion.

By the pidgeonhole principle, any longer strategy must contain a multiple switch of some elements a and b as there are only $\binom{n}{2}$ pairs. Therefore it can be shorthened by the previous Lemma. Any longer strategy is non-optimal with a positive vector w as in this case the shortening strictly decreases its cost.

Note that the longer strategies might still be optimal if all the extra swaps happen at a zero cost. $\square$

The two following lemmas give lower bounds for any linear arbitrary combination of cost-vectors with known lower bounds. The first lemma gives an equality, while in the second the inequality is sharp for some cost-vectors. **PROVE**

Lemma 4 (Multiplication Lemma)**.** Let w be a cost-vector and let $\lambda > 0$ be a real number. Then $\lambda\texttt{opt}_w(n) = \texttt{opt}_{\lambda w}(n)$ and the same strategies are optimal.

Proof. This follows from the observation that for any strategy $\mathcal{S}$ holds:

$$\lambda\texttt{opt}_w(\mathcal{S}) = \texttt{opt}_{\lambda w}(\mathcal{S}).$$

$\square$

Lemma 5 (Summing Lemma)**.** Let w_1 and w_2 be cost-vectors and let $w = w_1 + w_2$. Then

$$\texttt{opt}_w(n) \geq \texttt{opt}_{w_1}(n) + \texttt{opt}_{w_2}(n).$$

If the same strategy $\mathcal{S}$ is optimal for both cost-vectors w_1 and w_2 then $\mathcal{S}$ is also optimal for w and an equality holds.

Proof. Lorem ipsum dolor sit amet. $\square$

3 Symmetric Bimonotone Cost-Vectors

The first special case to consider is the case of symmetric bimonotone cost-vectors.

Definition 6. A cost-vector w is *bimonotonic* if the sequence $\{w_i\}_1^{n-1}$ is first non-increasing and then non-decreasing.

The following lemma has been proved by László Lovász et al. [Lo04] and has a deep connection to k-sets.

Lemma 7 ((1, 0, 1)-sequences)**.** Let $w = (1, 1, \dots 1, 0, 0, \dots 0, 1, 1, \dots 1)$ be a symmetric cost-vector consisting of k ones, $n - 2k - 1$ zeros and k ones in that order. If $k < n/3$ then $\mathsf{c}_w(n) = \mathsf{c}_w(\mathcal{S}_M(n) = 3\binom{k+1}{2}$.

Definition. The *Mercedes strategy* for reversing permutations of length n denoted by $\mathcal{S}_M(n)$ arises from the Mercedes construction. ...

Corollary. The Mercedes strategy $\mathcal{S}_M(n)$ is optimal for any (symmetric?) (bimonotone?) cost-vector w such that $w_{n/3} = w_{n/3+1} = \dots = w_{2n/3} = min_i(w_i)$.

To do: prove: symmetric? bimonotone?

4 Increasing Cost Vectors

4.1 General Observations

We say that a swap of a number m with a number m' is called a *step.* During this step, m does one *semistep* and m' does another one. Our approach is to examine the trajectories of individual numbers and the semisteps that they make in a particular strategy. Formally, by a *trajectory* of a number m in a strategy $\mathcal{S}$ we mean the sequence of positions occupied by m during $\mathcal{S}$. The sequence of m's left/right semisteps is called the *tour* of m. Note that two distinct trajectories of m may define the same tour, because m does not change position in every step.

Regardless of the strategy, each number m has a unique hypothetic *ideal trajectory* that goes from the original position m straight to the final position $n-m+1$. It is easy to observe that in any reversing strategy, each number's trajectory includes the ideal one (perhaps interleaved with other positions), and the same is true for tours.

Let s be a semistep of a number m from position p to p' in a strategy $\mathcal{S}$. If a semistep from p to p' belongs to m's ideal tour and if s is the first semistep of m from p to p' in $\mathcal{S}$, we call s a *necessary semistep* in $\mathcal{S}$. Otherwise, we call it an *unnecessary semistep* in $\mathcal{S}$.

Lemma 8. Any optimal reversing strategy performs $n^2 - n$ semisteps, out of which $\lfloor n^2/2 \rfloor$ are necessary and $\lceil n^2/2 - n \rceil$ are unnecessary semisteps.

Proof. It follows from Corollary 3 that there are $\binom{n}{2}$ steps, hence $n^2 - n$ semisteps.

The number of necessary semisteps done by a number m is the distance between the initial position m and the final position $n - m + 1$. By summing this over all numbers, we get

$$\sum_{m=1}^{\lfloor \frac{n}{2} \rfloor}(n - 2m + 1) + \sum_{m=\lfloor \frac{n}{2} \rfloor + 1}^{n} (2m - n - 1).$$

If n is odd, this yields $n^2/2 - 1/2$, and for n even, the result is $n^2/2$, as desired. The remaining semisteps are the unnecessary semisteps, so their number is simply the difference of n and $\lfloor n^2/2 \rfloor$. $\square$

4.2 Increasing and Non-decreasing Weights

There is a simple strategy which we call the Bubble strategy, $\mathcal{S}_B$:

for $j = 1$ to $n - 1$ **do**
  for $i = 1$ to $n - j$ **do**
    swap the number at i-th position with number at $(i + 1)$-th position
  end for
end for

Proposition 9. The cost $\mathsf{c}_w(\mathcal{S}_B)$ is $\sum_{j=1}^{n-1}(n - j)w_j$.

Proof. The cost of j-th cycle is $\sum_{i=1}^{n-j} w_i$. The total cost over all cycles is

$$\sum_{j=1}^{n-1}\sum_{i=1}^{n-j} w_i = \sum_{j=1}^{n-1}(n - j)w_j.$$

$\square$

Theorem 10. Let w be a non-decreasing cost vector. Then the Bubble strategy is optimal, and the cost of any optimal strategy is $\sum_{j=1}^{n-1}(n - j)w_j$.

Before we prove the theorem, we need several definitions and lemmas. Assume that a number m left a position i during step s_{start}, and then m first returned to i during step s_{end}. We call the pair $(s_{\text{start}}, s_{\text{end}})$ a *zigzag*, a the sequence of semisteps starting by s_{start} and ending with s_{end} is called a *detour of* m. Note that m goes either to the right during s_{start} and to the left during

s_{end}, or vice versa. We call these cases a *right zigzag* and *right detour*, or a *left zigzag* and *left detour*, respectively.

Since weights of transpositions only increase rightwards, one may observe that right detours are undesirable. Indeed, we have the following proposition.

Proposition 11. In an optimal reversing strategy for an increasing cost vector w, no number makes any right detour. If w is non-decreasing, then there exists an optimal strategy without any right detours.

Proof. Let w be a non-decreasing cost vector and let $\mathcal{S}$ be a reversing strategy optimal for w. We modify $\mathcal{S}$ so that all right detours are destroyed without increasing the total cost.

Let $\mathcal{D}$ be a right detour of a number m starting with step s_{start} and ending with s_{end}. Let $a_1, \dots, a_j$ be all the numbers that are positioned to the left of m before s_{start}. We say that $\mathcal{D}$ is *extreme* if none of the numbers $a_1, \dots, a_j$ does a right detour between s_{start} and s_{end}. Observe that for any right detour, there is at least one extreme right detour to the left of it; hence, we can without loss of generality assume that $\mathcal{D}$ is extreme. Further, we assume that $\mathcal{D}$ is minimal in inclusion, i.e., only the s_{start} and s_{end} steps change the position of m.

Let b and c be the numbers swapped with m during s_{start} and s_{end}, respectively. Then $c \in \{a_1, \dots, a_j\}$, otherwise c would be swapped with m more than once. Clearly, $b \notin \{a_1, \dots, a_j\}$. Let the permutation before s_{start} be

$$a_1, a_2, \dots, a_l, c, a_{l+2}, \dots, a_{j-1}, a_j, m, b, \dots$$

From s_{start} to s_{end}, the number c moves $j-l$ positions to the right. By the extremality of $\mathcal{D}$, the number c does not do any semistep to the left (otherwise, there is a right detour). Hence between s_{start} and s_{end}, the number c does exactly $j-l$ semisteps, all of them to the right. By a slight abuse of notation, we denote them $s_1, \dots, s_{j-l}$.

Lemma 12. During the steps $s_1, \dots, s_{j-l}$, the number c was swapped with $a_{l+2}, a_{l+3}, \dots, a_{j-1}, a_j, b$, in this order.

Proof. Throughout this proof we consider only swaps that are done from s_{start} to s_{end}. Suppose that, in the step s_{j-l}, the number c was swapped with a number x.

Since c never goes left and m never goes right after s_{start}, no number can get between m and c except for $a_{l+2}, \dots, a_j, b$. Hence, if x is not b, then $x \in \{a_{l+2}, \dots, a_j\}$, and it must have moved to the right so as to get to its position before s_{j-l}. But during s_{j-l}, the number x goes left, so the tour of

x contains a right detour. A contradiction with the extremality of $\mathcal{D}$. So c was swapped with b during s_{j-l}.

The rest of the claim follows by induction using the same arguments for the numbers a_j down to a_{l+2}. $\square$

We continue the detour-destroying proof by reordering the strategy steps to create a new strategy $\mathcal{S}'$. It follows from Lemma 12 that none of $a_{l+2}, \ldots,$ a_j, b was swapped with anything before its swap with c. Also, neither m nor b is involved in any of $s_1, \ldots, s_{j-l-1}$. Hence we can move these steps back in time and do them before the beginning of the detour $\mathcal{D}$.

Then, instead of s_{start}, we swap c and m, because they are now neighbors. We perform all the steps that remain before s_{end} with the role of c and m exchanged. After that, instead of s_{end} we do nothing, because c and m have already been swapped. Now the permutation is the same as it was after s_{end}, and the rest of the strategy is the same as in $\mathcal{S}$.

So the new strategy $\mathcal{S}'$ is the following:

1. all steps of $\mathcal{S}$ before s_{start},
2. $s_1, \ldots, s_{j-l-1}$,
3. swap the pair (c, m),
4. all the remaining steps of $\mathcal{S}$ from s_{start} to the one before s_{end} (with the role of m and c exchanged),
5. all steps of $\mathcal{S}$ after s_{end}.

By the arguments above, $\mathcal{S}'$ is a correct reversing strategy. Now we compare the costs of $\mathcal{S}$ and $\mathcal{S}'$. We have removed s_{end} and introduced a new swap (c, m). The remaining swaps were reordered, but their positions were preserved, and costs as well. Say that the swap s_{end} occured at position i, and its cost is w_i. Then the swap (c, m) ocurred at position $i-1$, so it costs w_{i-1}. Hence the total cost was not increased, and if $w_{i-1} < w_i$, it was even decreased.

Our aim is to show that $\mathcal{S}'$ is better, meaning that it contains "less" right detours. Indeed, we prove that $\mathcal{S}'$ contains less right zigzags.

By doing the above modification to $\mathcal{S}$, the trajectories of numbers were changed. But it is straightforward to verify that the tours of all numbers except for b and m are preserved, hence the number of their zigzags remains the same.

XXX The tours of m and b

Clearly, if the strategy $\mathcal{S}$ contains any right detours, there is at least one right zigzag. Hence, if the cost vector w is increasing, we perform the above modification and get a contradiction with the optimality of $\mathcal{S}$.

If w is non-decreasing, then applying the above modification to an optimal strategy $\mathcal{S}$ preserves the cost while decreasing the number of right zigzags by one. By applying it repeatedly, we obtain an optimal strategy without any right detours. $\square$

The following lemma is an immediate corollary of Proposition 11.

Lemma 13. Let m be a number. In a trajectory of m without right detours, all left semisteps precede all right ones. Furthermore, in the trajectory of m, there is at most one left detour and that is situated to the left of the leftmost position of the ideal trajectory of m.

Lemma 14. Let w be a non-decreasing cost vector. In an optimal strategy for w, every number first goes leftwards to the first position, and then rightwards to its final one.

Proof. We count the lengths of longest possible detours of each number. By Corollary 13, the detour of a number j must be positioned to the left of j's ideal trajectory, i.e., it must be within the $\min(j-1, n-j)$ leftmost positions. And since all the left steps precede all right ones, j can go only once in each direction, and the length of j's detour is at most twice the number of these positions.

Assume that n is even, and let $n = 2k$. By summing up the maximal detour lengths, we get

$$\sum_{j=1}^{k} 2(j-1) + \sum_{j=k+1}^{2k} 2(2k-j) = 2k(k-1) = n^2/2 - n.$$

If $n = 2k+1$, then the sum of maximal detour lengths is

$$\sum_{j=1}^{k} 2(j-1) + \sum_{j=k+1}^{2k+1} 2(2k+1-j),$$

which is

$$k^2 - k + k^2 + k = 2k^2 = \frac{n^2}{2} - n + \frac{1}{2}.$$

Since in both cases we are summing integers, the resulting expressions are also integral and equal to $\lceil n^2/2 - n \rceil$. But by Lemma 8, that is the number of unnecessary semisteps present in any optimal algorithm. Hence, the detour of every number must be the unique longest possible one. **XXX proc** $\square$

Proof of Theorem 10. We compare the trajectories of numbers in Bubble to the trajectories described by Lemma 14.

It is then easy to see that, in the Bubble strategy, a number m first goes one step left per each of the first $m-1$ cycles, then it reaches the first position and in the m-th cycle, it is bubbled to its final position, $n-m+1$. After that, no step involves that position, so m does not move. Hence, the trajectory of m is the same as the one described by Lemma 14.

We have shown that the number trajectories of Bubble and any optimal strategy are identical. Therefore, the costs of all semisteps in any optimal strategy sum up to the same cost as for Bubble. Hence Bubble is optimal, and the cost of Bubble is $\sum_{j=1}^{n-1}(n-j)w_j$ by Proposition 9. $\square$

References

[Lo71] L. Lovász: *On the number of halving lines*, Ann. Univ. Sci. Budapest, Eőtvős, Sect. Math. 14:107–108, 1971.

[Lo04] L. Lovász, K. Vesztergombi, U. Wagner, E. Welzl: *Convex quadrilaterals and k-sets*, Towards a Theory of Geometric Graphs, (J. Pach, Ed.), AMS Contemporary Mathematics 342 (2004), 139–148.