HAL
open science

New bounds and tractable instances for the transposition distance

Anthony Labarre

To cite this version:

Anthony Labarre. New bounds and tractable instances for the transposition distance. IEEE/ACM Transactions on Computational Biology and Bioinformatics, 2006, 3 (4), pp.380-394. 10.1109/TCBB.2006.56 . hal-00728947

HAL Id: hal-00728947

https://hal.science/hal-00728947

Submitted on 7 Sep 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

New Bounds and Tractable Instances for the Transposition Distance

Anthony Labarre

(Invited Paper)

Abstract

The problem of sorting by transpositions asks for a sequence of adjacent interval exchanges that sorts a permutation and is of the shortest possible length. The distance of the permutation is defined as the length of such a sequence. Despite the apparently intuitive nature of this problem, introduced in 1995 by Bafna and Pevzner, the complexity of both finding an optimal sequence and computing the distance remains open today. In this paper, we establish connections between two different graph representations of permutations, which allows us to compute the distance of a few non-trivial classes of permutations in linear time and space, bypassing the use of any graph structure. By showing that every permutation can be obtained from one of these classes, we prove a new tight upper bound on the transposition distance. Finally, we give improved bounds on some other families of permutations, and prove formulas for computing the exact distance of other classes of permutations, again in polynomial time.

Index Terms

Genome rearrangements, permutations, sorting by transpositions.

Anthony Labarre is funded by the "Fonds pour la Formation à la Recherche dans l'Industrie et dans l'Agriculture" (F.R.I.A.) and is with the Université Libre de Bruxelles, Département de Mathématique, CP 216, Service de Géométrie, Combinatoire et Théorie des Groupes, Boulevard du Triomphe, B-1050 Bruxelles, Belgium. E-mail: alabarre@ulb.ac.be.

New Bounds and Tractable Instances for the Transposition Distance

I. Introduction

The *genome rearrangement* problem ([1], [2]) can be formulated as that of finding a sequence of evolutionary events that transforms a given genome into another given one and is of the shortest possible length. The *distance* between the two genomes is the length of such a sequence.

The model we are interested in applies to the case where the order of genes is known and where all genomes share the same set and number of genes (without duplications), which allows us to represent them using *permutations*. Only one operation is taken into account here: biological *transpositions*, which consist in displacing a block of contiguous elements. It is easy to show that the induced distance is indeed a distance on the set of all permutations (i.e. it satisfies the three usual axioms), and that it is *left-invariant*: the distance between any two permutations π and σ of the same set equals the distance between $\sigma^{-1} \circ \pi$ and the *identity permutation* $\iota = (1\ 2\ \cdots\ n)$. We can therefore restrict our attention to the problem of *sorting permutations by transpositions*.

This problem was first introduced in 1995 by Bafna and Pevzner [3] (journal version in [4]), and the complexity of both sorting permutations and computing their distance, as well as the maximal value the latter can reach, are still open today. Several authors have proposed polynomial-time approximation algorithms (whose best approximation ratio has long been $\frac{3}{2}$ [4], [5], [6], until Elias and Hartman [7] recently proposed a new $\frac{11}{8}$-approximation) as well as heuristics (see [5], [8], [9], [10]).

In this paper, we establish connections between the common graph of a permutation and the "cycle graph" introduced in [4]. Use of the former was mentioned in [11] and led to a formula for computing another rearrangement distance in [12]. As we suspected, it proved fruitful for our problem too: the connections between the two graphs allowed us to compute the distance of a few non-trivial classes of permutations, bypassing the use of any graph structure, prove a new tight upper bound on the transposition distance, and improve that upper bound in some other cases.

This paper is organised as follows. In Section II, we review previous results and typical

notations. In Section III, we introduce a graph that we use in Section IV to provide a formula for computing the distance of some special permutations. In Section V, we use those permutations to derive an upper bound on the transposition distance of every permutation. Experimental data, comparisons and heuristic improvements of this bound are discussed in Section VI. We then turn to the study of other permutations in Sections VII and VIII, for which we can either compute the transposition distance or improve our upper bound on it. Finally, we discuss our results in Section IX and suggest some open questions of interest.

A preliminary version of this work was presented at the Fifth Workshop on Algorithms in Bioinformatics (WABI 2005), in Palma de Mallorca, Spain [13]. The main additions in this extended version consist, besides changes in the structure and presentation, in additional experimental data, three new sections (Sections VI, VII and VIII), and Appendices I and II.

II. Notations and Preliminaries

The *symmetric group* S_n is the set of all permutations of $\{1, 2, \ldots, n\}$; these are denoted by lower case Greek letters, typically $\pi = (\pi_1\ \pi_2\ \cdots\ \pi_n)$, with $\pi_i = \pi(i)$.

A. Transpositions and the Cycle Graph

Definition 2.1: For any π in S_n, the *transposition* $\tau(i, j, k)$ with $1 \leq i < j < k \leq n + 1$ applied to π exchanges the closed intervals determined respectively by i and $j - 1$ and by j and $k - 1$, transforming π into $\pi \circ \tau(i, j, k)$. So $\tau(i, j, k)$ is the following permutation:

$$\begin{pmatrix} 1 \cdots i-1 & \boxed{i \cdots j-1} & \boxed{j \cdots k-1} & k \cdots n \\ 1 \cdots i-1 & \boxed{j \cdots k-1} & \boxed{i \cdots j-1} & k \cdots n \end{pmatrix}.$$

Definition 2.2: The *cycle graph* of π in S_n is the bicoloured directed graph $G(\pi)$, whose vertex set $(\pi_0 = 0, \pi_1, \ldots, \pi_n, \pi_{n+1} = n + 1)$ is ordered by positions, and whose edge set consists of:

- *black* edges (π_i, π_{i-1}) for $1 \leq i \leq n + 1$;
- *grey* edges $(\pi_i, \pi_i + 1)$ for $0 \leq i \leq n$.

The set of black and grey edges decomposes in a single way into *alternating cycles*, i.e. cycles which alternate black and grey edges, and we note the number of such cycles $c(G(\pi))$. Fig. 1 shows an example of a cycle graph, together with its decomposition.

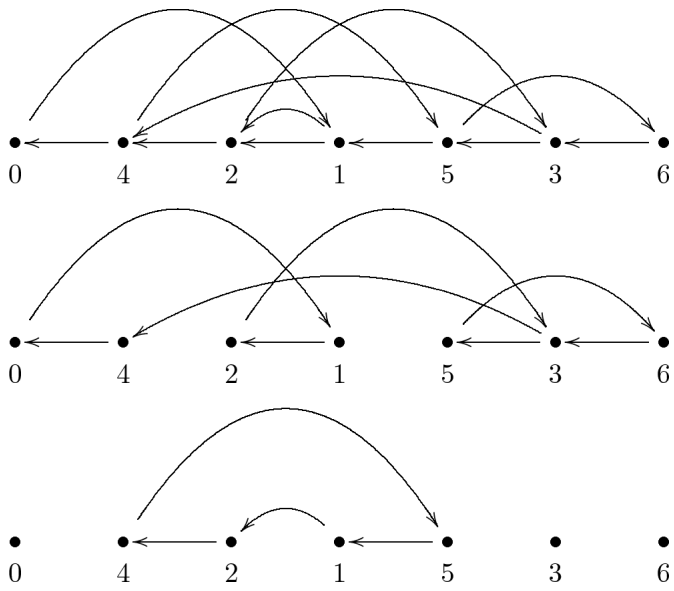

Fig. 1. The cycle graph of $(4\ 2\ 1\ 5\ 3)$ and its decomposition into two cycles

Definition 2.3: The *length* of an alternating cycle in G is the number of black edges it contains, and a k*-cycle* in G is an alternating cycle of length k.

Definition 2.4: A k-cycle in G is *odd* (resp. *even*) if k is odd (resp. even), and we note $c_{odd}(G(\pi))$ (resp. $c_{even}(G(\pi))$) the number of odd (resp. even) alternating cycles in $G(\pi)$.

Bafna and Pevzner [4] proved the following lower bound on the transposition distance, hereafter denoted by $d(\pi)$.

Theorem 2.1: [4] For all π in S_n:

$$d(\pi) \geq (n + 1 - c_{odd}(G(\pi)))/2.$$

Definition 2.5: A cycle in G is *unoriented* if it contains exactly one grey edge directed from left to right, and *oriented* otherwise.

For instance, the first cycle in the decomposition of the graph of Fig. 1 is oriented; the second is not. A transposition $\tau(i, j, k)$ is said to *act* on black edges coming out of vertices π_i, π_j and π_k in $G(\pi)$. By extension, a transposition acts on one cycle (resp. on two or three cycles) if all three black edges belong to that cycle (resp. to those two or three cycles).

Definition 2.6: For a permutation π, a k*-move* is a transposition τ such that $c(G(\pi \circ \tau)) = c(G(\pi)) + k$.

Lemma 2.1: [4] A transposition that acts on exactly two cycles in G is a 0-move.

Two alternating cycles can interact in several different ways, which we define below. To every alternating cycle C in a cycle graph G, associate an interval I_C defined by the minimum and maximum indices of the vertices that belong to C.

Definition 2.7: A cycle C_1 *contains* a cycle C_2 if $I_{C_1} \supset I_{C_2}$ and no black edge of C_1 belongs to I_{C_2}.

Definition 2.8: Two alternating cycles C_1, C_2 *cross* if they do not contain each other and at least one black edge of C_1 (resp. C_2) belongs to I_{C_2} (resp. I_{C_1}).

Definition 2.9: Two alternating cycles C_1, C_2 *interleave* if when reading the black edges of C_1 and C_2 from left to right, we alternately get a black edge from either cycle.

B. Reduced Permutations

Definition 2.10: For a permutation π, an ordered pair (π_i, π_{i+1}) is a *breakpoint* if $\pi_{i+1} \neq \pi_i+1$, and an *adjacency* otherwise. The number of breakpoints of π is denoted by $b(\pi)$.

Definition 2.11: A permutation π in S_n is *reduced* if $b(\pi) = n - 1$, $\pi_1 \neq 1$, and $\pi_n \neq n$.

Christie [5] shows that every permutation can be uniquely transformed into a reduced permutation without affecting its distance. The transformation of a permutation π into its reduced version $gl(\pi)$ consists in decomposing π into r *strips*, which are maximal intervals containing no breakpoint, then removing strip 1 if it begins with 1, strip r if it ends with n, replacing every other strip with its minimal element and finally, renumbering the resulting sequence so as to obtain a new permutation of a possibly smaller set. Since an adjacency is a 1-cycle in G, a reduced permutation can also be defined as one whose cycle graph has no 1-cycles[1].

Definition 2.12: Two permutations π and σ are *equivalent by reduction* if $gl(\pi) = gl(\sigma)$, which we also write as $\pi \equiv_r \sigma$.

Theorem 2.2: [5] For any two permutations π and σ: if $\pi \equiv_r \sigma$, then $d(\pi) = d(\sigma)$.

C. Toric Permutations

Eriksson et al. [14] introduced an equivalence relation on S_n, whose equivalence classes are called *toric permutations* and which we define using Hultman's notations [15].

[1]Note that $gl(\iota)$ is not defined, because every element would have to be removed.

Definition 2.13: The *circular permutation* obtained from a permutation π in S_n is $\pi^\circ = 0\ \pi_1\ \pi_2\ \cdots\ \pi_n$, with indices taken modulo $n+1$ so that $0 = \pi_0^\circ = \pi_{n+1}^\circ$.

This circular permutation can be read starting from any position, and the original linear permutation is reconstructed by taking the element following 0 as π_1 and removing 0. For x in $\{0, 1, 2, \ldots, n\}$, let $\overline{x}^m = (x+m) \pmod{n+1}$, and define the following operation on circular permutations:

$$m + \pi^\circ = \overline{0}^m\ \overline{\pi_1}^m\ \overline{\pi_2}^m\ \cdots\ \overline{\pi_n}^m.$$

Definition 2.14: For any π in S_n, the *toric permutation* $\pi_\circ^\circ$ is the set of permutations in S_n reconstructed from all circular permutations $m + \pi^\circ$ with $0 \leq m \leq n$.

Definition 2.15: Two permutations π, σ in S_n are *torically equivalent* if $\sigma \in \pi_\circ^\circ$ (or $\pi \in \sigma_\circ^\circ$), which we also write as $\pi \equiv_\circ^\circ \sigma$.

The following property is the main reason why toric permutations were introduced.

Lemma 2.2: [14] For all π, σ in S_n:

$$\pi \equiv_\circ^\circ \sigma \Rightarrow d(\pi) = d(\sigma).$$

Another interesting, related result has been proved by Hultman [15].

Lemma 2.3: [15] For all π in S_n and $0 \leq m \leq n$: every cycle in $G(\pi)$ is a cycle in $G(\sigma)$, where σ is the permutation obtained from $\pi^\circ + m$.

D. Known Upper Bounds

We conclude this section with all upper bounds on the transposition distance we know of.

Theorem 2.3: [4] For all π in S_n:

$$d(\pi) \leq n + 1 - c(G(\pi)). \tag{1}$$

Theorem 2.4: [4] For all π in S_n:

$$d(\pi) \leq 3(n + 1 - c_{odd}(G(\pi)))/4. \tag{2}$$

Theorem 2.5: [16] For all π in S_n:

$$d(\pi) \leq 3\ b(\pi)/4. \tag{3}$$

Theorem 2.6: [14] For all π in S_n:

$$d(\pi) \leq \begin{cases} \lceil 2n/3 \rceil & \text{if } n < 9 \ ; \\ \lfloor (2n-2)/3 \rfloor & \text{if } n \geq 9. \end{cases} \tag{4}$$

Elias and Hartman [7] proved upper bounds on the distance of three special classes of permutations[2].

Definition 2.16: A permutation π in S_n is *simple* if $G(\pi)$ contains no cycle of length greater than three.

Definition 2.17: A permutation π in S_n is a 2-*permutation* (resp. 3-*permutation*) if all cycles in $G(\pi)$ are of length 2 (resp. 3).

Note that a 2-permutation (resp. 3-permutation) only exists if $n+1$ can be divided by 4 (resp. 3).

Theorem 2.7: [7] For every simple permutation π in S_n which is neither a 2-permutation nor a 3-permutation:

$$d(\pi) \leq \lfloor (n+1)/2 \rfloor . \tag{5}$$

Theorem 2.8: [7] For every 2-permutation π in S_n :

$$d(\pi) \leq (n+1)/2. \tag{6}$$

Theorem 2.9: [7] For every 3-permutation π in S_n :

$$d(\pi) \leq 11 \left\lfloor \frac{n+1}{24} \right\rfloor + \left\lfloor \frac{3(\frac{n+1}{3} \bmod 8)}{2} \right\rfloor + 1. \tag{7}$$

III. Another Useful Graph

We introduce a slight variant of the well-known *graph of a permutation*.

Definition 3.1: The Γ-*graph* of a permutation π in S_n is the directed graph $\Gamma(\pi)$ with ordered vertex set $(\pi_1, \ldots, \pi_n)$ and edge set $\{(\pi_i, \pi_j) \mid \pi_i = j\}$.

Fig. 2 shows an example of a Γ-graph. If $C = (i_1, i_2, \ldots, i_k)$ is a cycle of π (i.e. π maps i_l onto i_{l+1} for $1 \leq l \leq k-1$ and i_k onto i_1), we obtain a cycle $(\pi_{i_1}, \pi_{i_2}, \ldots, \pi_{i_k})$, which we also denote C, in $\Gamma(\pi)$, and call it a k-*cycle*. The *length* of a cycle in Γ is therefore k.

[2]The definitions we give here are not the ones introduced by Hannenhalli and Pevzner [17] and Elias and Hartman [7], but we prove the equivalence between our definitions and theirs in Appendix II.

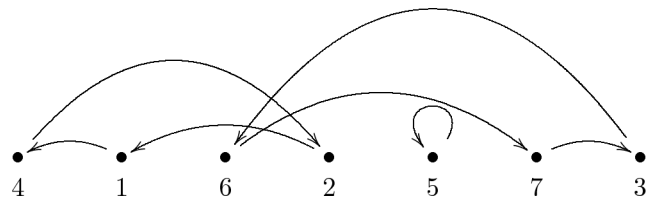

Fig. 2. The Γ-graph of the permutation $(4\ 1\ 6\ 2\ 5\ 7\ 3)$

Definition 3.2: A k-cycle in Γ is *increasing* (resp. *decreasing*) if $k \geq 3$ and its elements can be written as an increasing (resp. decreasing) sequence, and *non-monotonic* otherwise.

A cycle that is either increasing or decreasing is also referred to as *monotonic*. For instance, in Fig. 2, cycle $(4,\ 2,\ 1)$ is decreasing, cycle (5) is non-monotonic, and cycle $(3,\ 6,\ 7)$ is increasing. In a quite similar fashion to the parity of cycles defined in the context of G, a k-cycle in Γ is *odd* (resp. *even*) if k is odd (resp. even). Likewise, $c(\Gamma(\pi))$ denotes the number of cycles in $\Gamma(\pi)$, and $c_{odd}(\Gamma(\pi))$ (resp. $c_{even}(\Gamma(\pi))$) denotes the number of odd (resp. even) cycles in $\Gamma(\pi)$. Finally, note that Definitions 2.7, 2.8 and 2.9 naturally adapt to the Γ-graph.

IV. An Explicit Formula for Some Permutations

Definition 4.1: A *γ-permutation* is a reduced permutation that fixes even elements (thus n must be odd).

An example of a γ-permutation is $(3\ 2\ 1\ 4\ 7\ 6\ 9\ 8\ 5)$. We will show (Proposition 4.5) that the distance of such a permutation can be computed quickly, without the need of any graph structure.

Proposition 4.1: For every γ-permutation π in S_n:

$$\begin{cases} c_{even}(G(\pi)) &= 2\ c_{even}(\Gamma(\pi)); \\ c_{odd}(G(\pi)) &= 2\left(c_{odd}(\Gamma(\pi)) - \frac{n-1}{2}\right). \end{cases}$$

Proof: Each vertex π_i of $\Gamma(\pi)$, with i odd, is both the starting point of an edge (π_i, π_{j_1}) and the ending point of an edge (π_{j_2}, π_i). From our definitions, $\pi_i + 1$ is mapped onto itself, since it is even, and π_{j_1} precedes $\pi_i + 1$ in $\Gamma(\pi)$. In $G(\pi)$, those edges are each transformed, as explained below, into one sequence of two edges (grey-black for the first one, black-grey for the second one):

- (π_i, π_{j_1}) becomes $(\pi_i, \pi_i + 1), (\pi_i + 1, \pi_{j_1})$;

- (π_{j_2}, π_i) becomes $(\pi_i, \pi_{i-1}), (\pi_{i-1}, \pi_{j_2})$.

I.e. (π_i, π_{j_1}) is transformed in one of the following ways (depending on the relative positions of π_i and π_{j_1}):

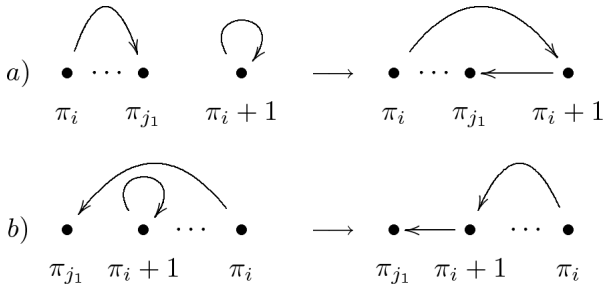

By definition of Γ, we know that $\pi_{j_2} = i$. Since $\pi_{i-1} = i - 1$, the edge (π_{j_2}, π_i) is transformed in one of the following ways (depending on the relative positions of π_i and π_{j_2}):

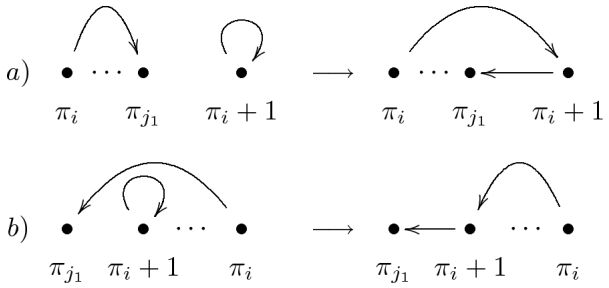

Therefore each k-cycle ($k \geq 2$) in $\Gamma(\pi)$ provides two alternating k-cycles in $G(\pi)$, one of which actually corresponds to the backwards course of the cycle in $\Gamma(\pi)$. Finally, 1-cycles in $\Gamma(\pi)$ are not preserved in $G(\pi)$, and there are $\frac{n-1}{2}$ of them. ■

The next observation follows naturally from our transformation.

Observation 4.1: For a γ-permutation π, the two alternating cycles C_1, C_2 in $G(\pi)$ that correspond to a k-cycle C in $\Gamma(\pi)$ interleave. Moreover:

1) if $k = 2$, then C_1 and C_2 are unoriented;
2) if C is monotonic, then either C_1 or C_2 is oriented;
3) if C is non-monotonic and $k \geq 4$, then both C_1 and C_2 are oriented.

Fig. 3 illustrates Proposition 4.1 and Observation 4.1. We derive the following lower bound from Proposition 4.1 and Theorem 2.1.

Lemma 4.1: For every γ-permutation π in S_n, we have $d(\pi) \geq n - c_{odd}(\Gamma(\pi))$.

Proof: Straightforward. ■

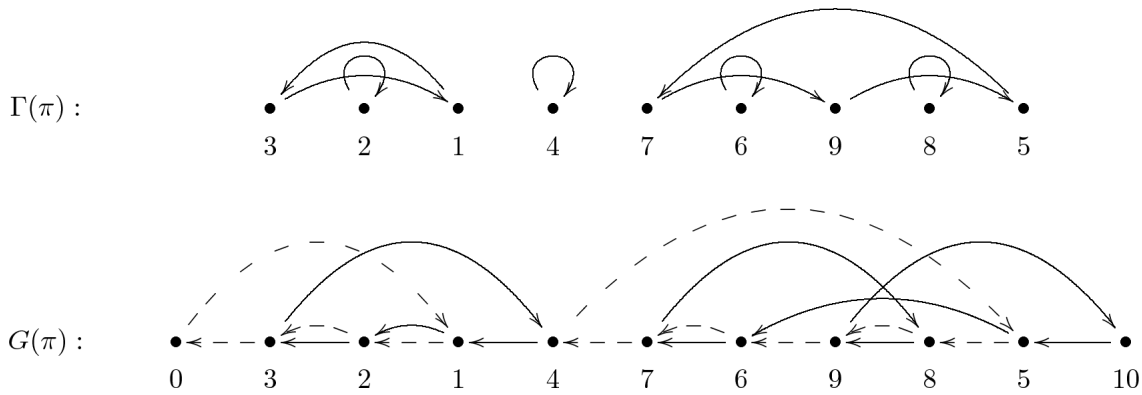

Fig. 3. Illustration of Proposition 4.1 and Observation 4.1

We first study γ-permutations such that Γ has only one "long" k-cycle (i.e. with $k > 1$), distinguishing between monotonic cycles and non-monotonic ones.

A. Monotonic Cycles

Definition 4.2: An α-*permutation* is a reduced permutation that fixes even elements and whose $\frac{n+1}{2}$ odd elements form one monotonic cycle in Γ, referred to as its *main cycle*.

An example of an α-permutation, for $n = 7$, is $(3\ 2\ 5\ 4\ 7\ 6\ 1)$. Note that for fixed n, there are only two α-permutations in S_n: one has an increasing main cycle, and the other has a decreasing main cycle. Therefore, the only other α-permutation, for $n = 7$, is $(7\ 2\ 1\ 4\ 3\ 6\ 5)$, which is the inverse of our above example.

Proposition 4.2: For every α-permutation π in S_n, we have

$$d(\pi) = n - c_{odd}(\Gamma(\pi)) = |C| - (|C| \bmod 2)\,,$$

where $|C| = \frac{n+1}{2}$ is the number of elements in its main cycle C.

Proof: Every α-permutation is a γ-permutation, so $d(\pi) \geq |C| - (|C| \bmod 2)$ (Lemma 4.1). Assume that C is increasing (a similar proof is easily obtained in the decreasing case), and consider transpositions $\tau_1(2, 4, n+1)$, $\tau_2(1, 3, n)$, $\tau_3(2, 3, n+1)$ and $\tau_4(1, 2, n+1)$. If $|C|$ is odd, then an optimal sorting sequence of length $|C| - 1$ for π is obtained by applying $\tau_2 \circ \tau_1$ exactly $\frac{|C|-1}{2}$ times. If $|C|$ is even, then an optimal sorting sequence of length $|C|$ for π is obtained by applying $\tau_2 \circ \tau_1$ exactly $\frac{|C|-2}{2}$ times, then τ_3 and finally τ_4. The proof that those sequences indeed sort π is given in Appendix I. ■

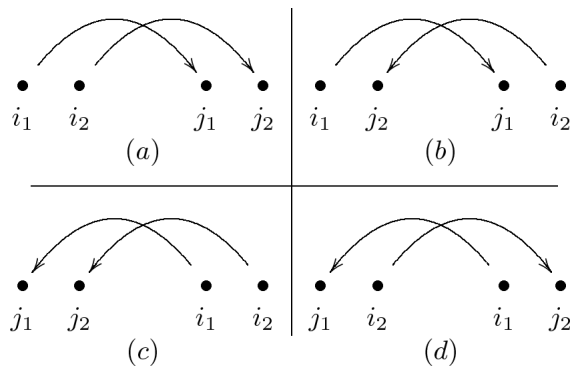

Fig. 4. The four possible configurations for crossing edges in $\Gamma(\pi)$

B. Non-Monotonic Cycles

Definition 4.3: A *β-permutation* is a reduced permutation that fixes even elements and whose odd elements form one non-monotonic cycle in Γ.

We now show that Proposition 4.2 still holds if the main cycle of Γ is non-monotonic. We use so-called *exchanges* in order to simplify the proofs, thus bypassing the construction of optimal sequences of transpositions.

Definition 4.4: An *exchange* $exc(i, j)$ is the permutation that exchanges elements in positions i and j, thus transforming every permutation π into the permutation $\pi \circ exc(i, j)$. So $exc(i, j)$ is the following permutation:

$$\begin{pmatrix} 1 \cdots i-1 \ \boxed{i} \ i+1 \cdots j-1 \ \boxed{j} \ j+1 \cdots n \\ 1 \cdots i-1 \ \boxed{j} \ i+1 \cdots j-1 \ \boxed{i} \ j+1 \cdots n \end{pmatrix}.$$

We only use exchanges of the form $exc(i, i+2k)$ with $k \geq 1$; such an exchange has the same effect as two transpositions, but the correspondence between those two types of operations is not that straightforward when exchanges are composed.

Definition 4.5: Two edges in $\Gamma(\pi)$ *cross* if the intervals determined by their endpoints do not contain each other and have a non-empty intersection.

Fig. 4 shows the four possible configurations for two crossing edges. Clearly, for every β-permutation π (except $(3\ 2\ 1)$), the main cycle of $\Gamma(\pi)$ contains crossing edges. We are going to transform π into a permutation σ that reduces to an α-permutation, by removing crossing edges using a certain sequence $\mathscr{E}$ of exchanges. This yields the following upper bound on the distance

of a β-permutation π:

$$d(\pi) \leq f(\mathscr{E}) + d(\sigma) \tag{8}$$

where $f(\mathscr{E})$ gives the minimum number of transpositions having the same effect on π as $\mathscr{E}$ does. Finding some σ is not difficult, but we have to find a σ such that our upper bound in (8) is minimised.

Eliminating a crossing can be done by making the ending point of one edge become the starting point of the one it crosses, and this will be achieved using a sequence of exchanges of the form described in the following proposition.

Proposition 4.3: For both sequences $\mathscr{E} = exc(i, i+2) \circ exc(i, i+4) \circ \cdots \circ exc(i, i+2t)$ and $\mathscr{F} = exc(i, i+2t) \circ \cdots \circ exc(i, i+4) \circ exc(i, i+2)$ of t exchanges:

$$f(\mathscr{E}) = f(\mathscr{F}) = t + (t \bmod 2).$$

Proof: Both sequences, when applied to the identity permutation, result in a permutation π which contains one long cycle and whose other cycles are all fixed points, since they are never affected by any exchange. If $t = 1$, then the long cycle is non-monotonic, and it is easily seen that $d(\pi) = 2$; otherwise, the long cycle of π is increasing in the case of $\mathscr{E}$, and decreasing in the case of $\mathscr{F}$. All elements before position i and after position $i + 2t$ are fixed, and removing them transforms π into $gl(\pi)$, which is an α-permutation whose main cycle has $t + 1$ elements. Therefore, by Theorem 2.2 and Proposition 4.2, we have:

$$\begin{aligned} d(\pi) &= t + 1 - ((t+1) \bmod 2) \\ &= t + (t \bmod 2) = f(\mathscr{E}) = f(\mathscr{F}). \end{aligned}$$

■

By a *path*, we mean a sequence of edges joining the ending point i of an edge to the starting point j of the edge it crosses, and such that the extremities of each edge in this path belong to the interval determined by i and j. Furthermore, we will refer to the elimination of this path as its *contraction*. Let us now compute the distance of β-permutations.

Proposition 4.4: For every β-permutation π in S_n, we have

$$d(\pi) = n - c_{odd}(\Gamma(\pi)) = |C| - (|C| \bmod 2)\,,$$

where $|C| = \frac{n+1}{2}$ is the number of elements in its main cycle C.

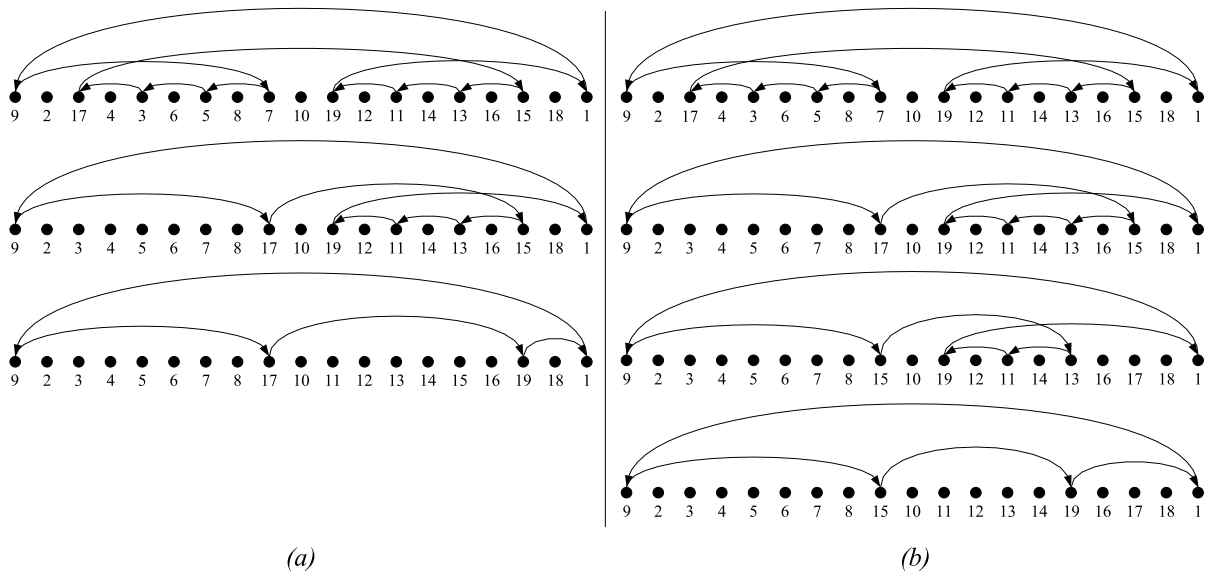

Fig. 5. Two ways of contracting paths in a β-permutation (1-cycles omitted for clarity)

Proof: Every β-permutation is a γ-permutation, so $d(\pi) \geq |C| - (|C| \bmod 2)$ (Lemma 4.1). If $\pi = (3\ 2\ 1)$ we are done; otherwise C contains at least one crossing.

In that case, there is a path of t edges joining the two crossing edges; this path can be contracted by a sequence of t exchanges, sorting the elements belonging to that part of the cycle. For instance, in Case (a) of Fig. 4, it suffices to apply the sequence $exc(i_2, j_1) \circ \cdots \circ exc(i_2, i_2 + 4) \circ exc(i_2, i_2 + 2)$, and those t exchanges correspond to exactly $t + (t \bmod 2)$ transpositions (Proposition 4.3).

Once this path has been contracted, t vertices have been removed from C and this results in a permutation σ reducible to an α-permutation. Therefore:

$$\begin{aligned} d(\pi) &\leq d(\pi, \sigma) + d(\sigma) \\ &= t + (t \bmod 2) + |C| - t - ((|C| - t) \bmod 2) \\ &= |C| - (|C| \bmod 2)\,. \end{aligned}$$

If there are p paths of t_g edges each ($1 \leq g \leq p$), contracting them all "individually" takes $\sum_{g=1}^{p} t_g$ exchanges, or $\sum_{g=1}^{p}(t_g + (t_g \bmod 2))$ transpositions (Proposition 4.3). This can actually be improved by exchanging the last exchanged element in the first contracted path with the first element of the next path to contract, then continuing the contraction of the latter with

dependent exchanges as before, repeating the same process whenever need be. For instance, Fig. 5 shows two different transformations of a β-permutation into a permutation reducible to an α-permutation: scenario (a), which removes both crossings using two disjoint sequences, uses 3+3 exchanges = 8 transpositions (Proposition 4.3), whereas scenario (b), which removes both crossings using a single sequence, uses the same number of exchanges, but requiring only 6 transpositions this time.

Every β-permutation π whose Γ-graph contains p paths of t_g edges to contract ($1 \leq g \leq p$) can therefore be transformed into a permutation σ reducible to an α-permutation and such that $d(\pi, \sigma) = T + (T \bmod 2)$, where $T = \sum_{g=1}^{p} t_g$. The transformation removes T vertices from C, which yields the following upper bound:

$$\begin{aligned} d(\pi) &\leq d(\pi, \sigma) + d(\sigma) \\ &= T + (T \bmod 2) + |C| - T - ((|C| - T) \bmod 2) \\ &= |C| - (|C| \bmod 2) \end{aligned}$$

which equals the lower bound given above. ■

C. Distance of γ-Permutations

Each cycle in $\Gamma(\pi)$ can be sorted (by transpositions) individually, so that the resulting permutation has the same Γ-graph as π, except that one cycle has been transformed into fixed points. This strategy yields the following upper bound on $d(\pi)$.

Lemma 4.2: For every permutation π, consider its disjoint cycle decomposition $\Gamma(\pi) = C_1 \cup C_2 \cup \cdots \cup C_{c(\Gamma(\pi))}$. Denote $d(C)$ the minimum number of transpositions required to transform $C = (i_1, i_2, \ldots, i_k)$ into $(i_1), (i_2), \ldots, (i_k)$; then

$$d(\pi) \leq \sum_{i=1}^{c(\Gamma(\pi))} d(C_i). \tag{9}$$

We now show that (9) is tight for γ-permutations.

Proposition 4.5: For every γ-permutation π in S_n :

$$d(\pi) = n - c_{odd}(\Gamma(\pi)). \tag{10}$$

Proof: Denote $odd(\Gamma(\pi))$ (resp. $even(\Gamma(\pi))$) the set of odd (resp. even) cycles in $\Gamma(\pi)$; Lemma 4.2 and Propositions 4.2 and 4.4 yield

$$\begin{aligned} d(\pi) &\leq \sum_{i=1}^{c(\Gamma(\pi))} |C_i| - (|C_i| \bmod 2) \\ &= \sum_{C_{i_1} \in\ odd(\Gamma(\pi))} (|C_{i_1}| - 1) + \sum_{C_{i_2} \in\ even(\Gamma(\pi))} |C_{i_2}| \\ &= \sum_{i=1}^{c(\Gamma(\pi))} |C_i| - c_{odd}(\Gamma(\pi)). \end{aligned}$$

Since every element belongs to exactly one cycle, the last sum equals n and the proof follows from Lemma 4.1. ■

This proposition actually leads to a more general result.

Theorem 4.1: Every permutation π in S_n that reduces to a γ-permutation has distance

$$d(\pi) = n - c_{odd}(\Gamma(\pi)).$$

Moreover, every permutation σ with n odd and whose odd elements occupy odd positions and form an increasing subsequence modulo $n+1$ can be transformed in linear time into a permutation π such that $d(\sigma) = d(\pi) = n - c_{odd}(\Gamma(\pi))$.

Proof: Let π be a γ-permutation in S_n: transforming π into a permutation $\sigma \neq \pi$ such that $\sigma \equiv_r \pi$ is done by creating adjacencies in π, i.e. repeatedly adding an element e between π_i and π_{i+1} such that $e = \pi_i + 1$ or $e = \pi_{i+1} - 1$ (a subsequent renumbering of elements is of course required). Since either π_i or $\pi_i + 1$ is fixed (or possibly both, if this is not the first addition), adding e comes down to inserting a new 1-cycle in $\Gamma(\pi)$, and this increases both n and $c_{odd}(\Gamma(\pi))$ by 1 at each step, so (10) still holds (Theorem 2.2).

For the second category, note that $\pi^\circ \pm 1$ fixes all even elements and therefore falls into the category discussed above. The proof follows from Lemma 2.2. ■

V. A New Upper Bound

We now show that the right-hand side of (10) is an upper bound on the transposition distance. First we show why γ-permutations are so important.

Theorem 5.1: Every permutation π in S_n, except ι, can be obtained from a permutation σ in S_{n+k} that reduces to a γ-permutation.

Proof: If $\pi \neq \iota$ does not reduce to a γ-permutation, add a 1-cycle to $\Gamma(\pi)$ between every ordered pair (π_i, π_{i+1}) $(1 \leq i \leq n-1)$; then the resulting permutation σ in S_{n+k} reduces to a γ-permutation. The transformation can clearly be reverted, and this completes the proof. ■

Theorem 5.2: For all π in S_n:

$$d(\pi) \leq n - c_{odd}(\Gamma(\pi)). \tag{11}$$

Proof: If $\pi = \iota$, then the proof follows at once. Otherwise, let σ be the permutation from which π is obtained by removing k 1-cycles from $\Gamma(\sigma)$, as described in Theorem 5.1. The sorting strategy of Lemma 4.2, optimal for σ, still works for π, only it may not be optimal anymore. Moreover, Theorem 4.1 gives the distance of σ. Therefore:

$$\begin{aligned} d(\pi) \leq d(\sigma) &= n + k - c_{odd}(\Gamma(\sigma)) \\ &= n + k - c_{odd}(\Gamma(\pi)) - k \\ &= n - c_{odd}(\Gamma(\pi)). \end{aligned}$$

■

VI. Tests and Heuristic Improvements of the New Upper Bound

Table I shows the number of cases where (11) is at least as good as the bounds given in Section II. A first heuristic improvement can be obtained through torism.

Theorem 6.1: For all π in S_n:

$$d(\pi) \leq n - \max_{\sigma \in \pi_\circ^\circ} c_{odd}(\Gamma(\sigma)). \tag{12}$$

Proof: Straightforward from Theorem 5.2 and Lemma 2.2. ■

Experiments show (Table I) that (12) is a substantial improvement over (11), but it is hard to express or evaluate this improvement because the evolution of Γ under the toric equivalence relation does not seem easy to predict, whereas that of G is well known (Lemma 2.3). Note by the way that the other upper bounds cannot be lowered through torism, since neither the cycle graph structure nor the number of breakpoints will be affected.

A second heuristic improvement of (11) can be obtained through reduction.

TABLE I

COMPARISON OF THE NEW UPPER BOUNDS WITH PREVIOUS RESULTS

n	$n!$	$\lvert(11)\leq(1)\rvert$	$\lvert(12)\leq(1)\rvert$	$\lvert(11)\leq(2)\rvert$	$\lvert(12)\leq(2)\rvert$	$\lvert(11)\leq(3)\rvert$	$\lvert(12)\leq(3)\rvert$	$\lvert(11)\leq(4)\rvert$	$\lvert(12)\leq(4)\rvert$
1	1	1	1	1	1	1	1	1	1
2	2	2	2	1	1	1	1	1	1
3	6	6	6	2	2	1	1	6	6
4	24	19	24	8	11	8	9	15	21
5	120	101	112	45	60	24	36	31	54
6	720	529	671	304	451	49	73	495	703
7	5 040	3 837	4 654	2 055	3 318	722	1 336	1 611	3 574
8	40 320	28 354	37 209	17 879	27 486	3 094	5 957	4 355	9 864
9	362 880	257 844	336 744	104 392	259 195	60 871	132 801	10 243	21 610
10	3 628 800	2 469 217	3 280 815	430 164	1 244 002	361 659	931 584	485 154	1 376 134

n	$\lvert$simple permutations$\rvert$	$\lvert$2-permutations$\rvert$	$\lvert$3-permutations$\rvert$	$\lvert(11)\leq(5)\rvert$	$\lvert(12)\leq(5)\rvert$	$\lvert(11)\leq(6)\rvert$	$\lvert(12)\leq(6)\rvert$	$\lvert(11)\leq(7)\rvert$	$\lvert(12)\leq(7)\rvert$
1	1	-	-	1	1	-	-	-	-
2	2	-	1	1	1	-	-	1	1
3	6	1	-	6	6	1	1	-	-
4	16	-	-	11	16	-	-	-	-
5	48	-	12	19	30	-	-	12	12
6	204	-	-	32	50	-	-	-	-
7	876	21	-	369	192	9	21	-	-
8	3 636	-	464	749	1 530	-	-	10	20
9	18 756	-	-	1 433	2 781	-	-	-	-
10	105 480	-	-	2 678	4 896	-	-	-	-
11	561 672	1 485	38 720	46 342	112 364	87	281	862	2454

TABLE II

NUMBER OF CASES WHERE (13) OVERESTIMATES $d(\pi)$ BY Δ

n	$n!$	$\Delta = 0$	$\Delta = 1$	$\Delta = 2$	$\Delta = 3$	$\Delta = 4$
1	1	1	0	0	0	0
2	2	1	1	0	0	0
3	6	2	4	0	0	0
4	24	11	11	2	0	0
5	120	48	51	21	0	0
6	720	197	401	108	14	0
7	5 040	1 318	2 460	966	296	0
8	40 320	8 775	13 875	15 150	2 512	8
9	362 880	45 415	132 257	145 394	34 702	5 112

Theorem 6.2: For all $\pi \neq \iota$ in S_n, let $gl(\pi)$ denote its reduced version in S_m, where $m \leq n$; then

$$d(\pi) \leq m - \max_{\sigma \in (gl(\pi))^{\circ}_{\circ}} c_{odd}(\Gamma(\sigma)). \tag{13}$$

All other bounds can take advantage of this reduction as well, except for (1), (2) and (3). This time, we do not compare (13) with other bounds; instead, for $1 \leq i \leq 9$, we generate all permutations with their distance, and check how (13) overestimates their distance. Table II shows the results; for our range of experiments, it seems that (13) is a $\frac{3}{2}$-approximation.

VII. PERFORATIONS OF α-PERMUTATIONS

After looking at γ-permutations, it is natural to wonder how deleting their fixed points affects their distance. A careful analysis allows us to further improve (11) in the case of α-permutations.

Note that deleting a 1-cycle in position i in Γ can be done by placing π_i just before $\pi_i + 1$ using a transposition, then removing the obtained adjacency and renumbering the other elements appropriately.

Definition 7.1: A *k-perforation* π in S_n of an α-permutation σ in S_{n+k} is a permutation obtained by removing $k \geq 1$ 1-cycles from $\Gamma(\sigma)$ and renumbering the remaining elements.

For instance, a 3-perforation of the α-permutation $(3\ 2\ 5\ 4\ 7\ 6\ 9\ 8\ 11\ 10\ 1)$ is $(3\ \boxed{2}\ 5\ 4\ 7\ \boxed{6}\ 9\ \boxed{8}\ 11\ 10\ 1)$ = $(2\ 4\ 3\ 5\ 6\ 8\ 7\ 1)$. Let us have a look at how the structure of G evolves when perforating an α-permutation.

Lemma 7.1: For every k-perforation π of an α-permutation σ in S_{n+k}:

$$c(G(\pi)) = c_{odd}(G(\pi)) = k$$

and $G(\pi)$ contains only non-crossing cycles, not containing each other, except for a large one containing all others.

Proof: Induction on k. The main cycle of $\Gamma(\sigma)$ is again assumed to be increasing, the decreasing case corresponding to σ^{-1} whose cycle graph has the same structure (see Hultman [15]). Recall that $n+k$ is odd, by definition of σ.

If $k=1$, let us remove some fixed element $\sigma_i = i$ (i is therefore even), by first applying transposition $\tau(i-1, i, i+1)$. This transposition acts on two interleaving cycles of same parity in G (Observation 4.1), and is therefore a 0-move (Lemma 2.1) transforming those cycles into a 1-cycle and an $(n+k)$-cycle, both odd. We now remove the adjacency, and get a permutation π with $c(G(\pi)) = c_{odd}(G(\pi)) = 1$.

For the induction, we again remove 1-cycles from Γ in two steps, by first applying all our transpositions, then removing k adjacencies. Since the thesis is assumed to hold for $k-1$ perforations, we start with the corresponding $(k-1)$-perforation π', and put back the $k-1$ adjacencies that needed to be deleted, thus obtaining a permutation π'' with $c(G(\pi'')) = c_{odd}(G(\pi'')) = 2(k-1)$. None of these cycles cross, and one of them contains all others. Let us now select some even π''_i we wish to remove, and apply the adequate transposition τ to make it adjacent to $\pi''_i + 1$. The odd alternating cycle to which this element belongs will be cut into three cycles: an adjacency (1-cycle) "framed" by two cycles.

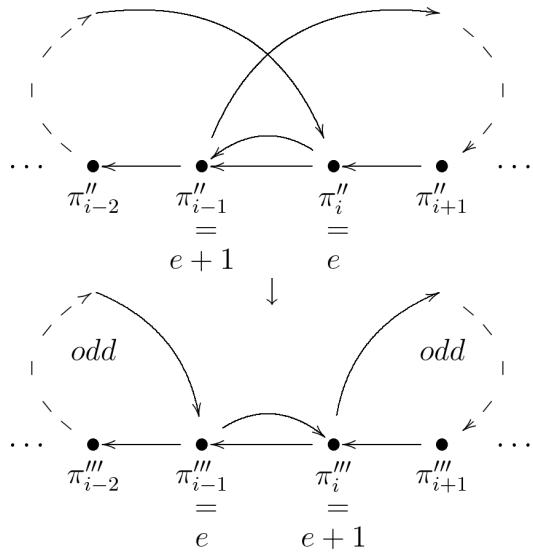

We need to prove that both framing cycles are odd, which comes down to showing that one of them is: indeed, the cycle we cut was odd, so the two cycles framing the new adjacency have

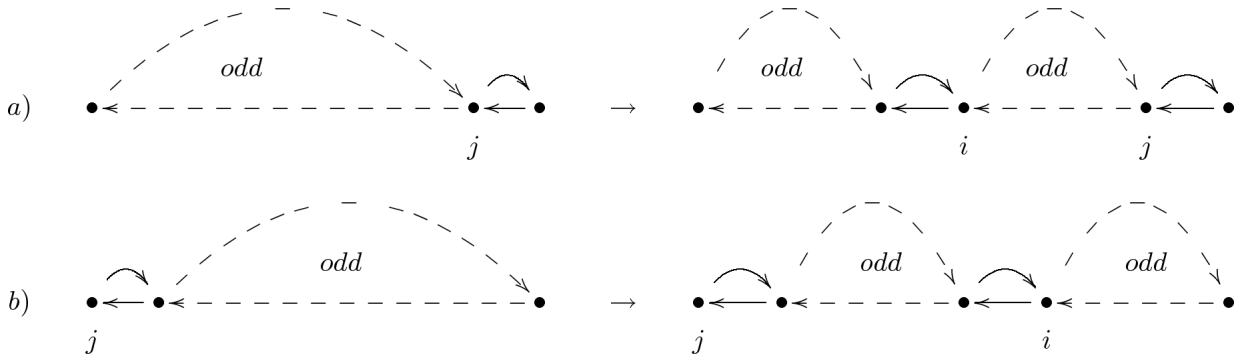

Fig. 6. Checking the parity of the new cycles in the proof of Lemma 7.1

the same parity. By induction, since at least one perforation has already been performed, there is an adjacency on the right-hand side or on the left-hand side of the cut cycle. This adjacency is caused by an even element in an odd position, and the number of black edges between an even position i and an odd position j is odd (Fig. 6 illustrates our claim). Therefore the three new cycles are odd, and we get the permutation $\pi''' = \pi'' \circ \tau$ with $c(G(\pi''')) = c_{odd}(G(\pi''')) = 2k$. We now remove k 1-cycles from $G(\pi''')$, and the proof follows. ■

This leads to a formula for computing the distance of such a permutation.

Corollary 7.1: For every k-perforation π of an α-permutation σ in S_{n+k}:

$$d(\pi) = n - c_{odd}(\Gamma(\pi)) - k + (|C| \bmod 2)\,,$$

where $|C| = \frac{n+k+1}{2}$ is the number of elements in its main cycle C.

Proof: Again, assume without loss of generality that C is increasing; Lemma 7.1 and Theorem 2.1 yield $d(\pi) \geq |C| - k$. It is easily seen that removing a fixed point from $\Gamma(\sigma)$ replaces the edge of length 2 that overhangs it in C with an edge of length 1, so C contains k edges of length 1 and $\frac{n+k+1}{2} - k - 1 = \frac{n-1-k}{2}$ edges of length 2 (the last one has length n). Using $\frac{n-1-k}{2}$ transpositions of the form $\tau(i, i+1, i+2)$, where i is the starting point of an edge of length 2, we transform π into $(2\ 3\ 4\ 5\ \cdots\ n+k-2\ n+k-1\ n+k\ 1)$ which is one transposition away from ι. We therefore apply

$$\frac{n-1-k}{2} + 1 = \frac{n+k+1}{2} - k = |C| - k$$

transpositions in order to sort π, which completes the proof since

$$\begin{aligned} d(\pi) &= |C| - k \\ &= |C| - k - (|C| \bmod 2) + (|C| \bmod 2) \\ &= d(\sigma) - k + (|C| \bmod 2) \\ &= n + k - c_{odd}(\Gamma(\sigma)) - k + (|C| \bmod 2) \\ &= n - c_{odd}(\Gamma(\sigma)) + (|C| \bmod 2) \\ &= n - c_{odd}(\Gamma(\pi)) - k + (|C| \bmod 2)\,. \end{aligned}$$

■

The next logical move, as in our analysis of γ-permutations, would be to consider perforations of β-permutations. However, counter-examples have been found that prevent us from proving an equivalent of Lemma 7.1 in the case of those permutations; for instance, consider the β-permutation $(7\ 2\ 13\ 4\ 3\ 6\ 5\ 8\ 15\ 10\ 9\ 12\ 11\ 14\ 1)$. Then the cycle graph of the 4-perforation $(7\ 2\ 13\ \boxed{4}\ 3\ \boxed{6}\ 5\ \boxed{8}\ 15\ 10\ 9\ 12\ 11\ \boxed{14}\ 1)$ = $(5\ 2\ 10\ 3\ 4\ 11\ 7\ 6\ 9\ 8\ 1)$ has only two cycles, both odd.

We can nevertheless still study permutations whose Γ-graph contains non-crossing cycles only. Fortunately, the exact distance of some subcases in that family can be computed; if not, we are nonetheless still able to improve (11).

Before tackling this general problem in the next section, we conclude this one with the particular case where all non-crossing long cycles are perforations of α-permutations, starting with the case shown in Fig. 7, where we do not allow containment of long cycles. In such a configuration, the 1-cycles between every pair of long cycles are referred to as the *separating* 1-cycles or, more concisely, the *separators*.

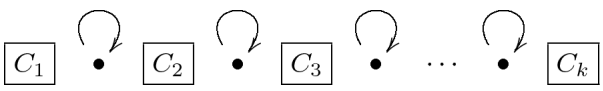

Fig. 7. A Γ-graph formed by sub-permutations separated by 1-cycles

Proposition 7.1: Let π in S_n be a permutation with $\Gamma(\pi)$ of the form shown in Fig. 7, where

C_i $(1 \leq i \leq k)$ is a k_i-perforation of an α-permutation; then

$$d(\pi) \;=\; n - c_{odd}(\Gamma(\pi)) - K + \sum_{i=1}^{k} (|C_i| \bmod 2)\,,$$

where $K = \sum_{i=1}^{k} k_i$ and $|C_i|$ is the number of elements in the main cycle of each perforation.

Proof: Lemma 7.1 and Theorem 2.1 yield

$$d(\pi) \geq \frac{n+1-\sum_{i=1}^{k} k_i}{2} = \frac{n+1-K}{2}.$$

We have $n = k - 1 + \sum_{i=1}^{k} n_i$, where n_i is the number of elements of each perforation, and Lemma 4.2 and Corollary 7.1 yield

$$\begin{aligned} d(\pi) \leq \sum_{i=1}^{k} d(C_i) \;&=\; \sum_{i=1}^{k} |C_i| - k_i \\ &=\; \sum_{i=1}^{k} \frac{n_i + k_i + 1 - 2k_i}{2} \\ &=\; \frac{1}{2}\sum_{i=1}^{k} (n_i + 1 - k_i) \\ &=\; \frac{n+1-K}{2} \end{aligned}$$

The expression given in the thesis is obtained by replacing $d(C_i)$ with the expression provided by Corollary 7.1. ■

We now show that removing any subset of the separators in the case we just examined does not affect the distance. For any transposition τ and any permutation π, let $\Delta c_{odd}(\tau, G(\pi)) = c_{odd}(G(\pi \circ \tau)) - c_{odd}(G(\pi))$. The following lemma will be useful.

Lemma 7.2: Let $\tau = \tau(i, i+1, \pi^{-1}_{\pi_i+1})$, and let C_1, C_2 be two cycles in $G(\pi)$ which share vertex π_i as shown below:

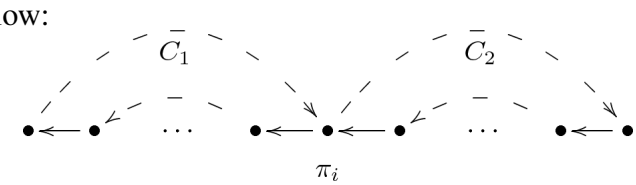

Then $\Delta c_{odd}(\tau, G(\pi)) = 2$ if both C_1 and C_2 are even, and 0 otherwise.

Proof: Fig. 8 shows the four cases. ■

Corollary 7.2: Let π be a permutation that satisfies the conditions of Proposition 7.1; then removing j $(1 \leq j \leq k-1)$ separators from $\Gamma(\pi)$ yields a permutation with the same distance.

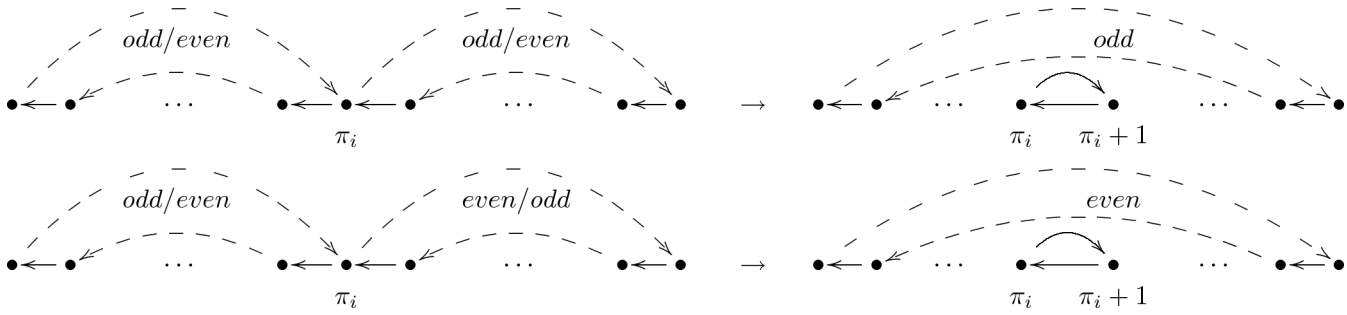

Fig. 8. Proof of Lemma 7.2

Proof: By Lemma 7.1, each C_i in $\Gamma(\pi)$ corresponds to a collection of alternating cycles in $G(\pi)$ wrapped in a large one, and all of them are odd. Every pair of consecutive "wrapping cycles" in $G(\pi)$ shares a vertex, which is the 1-cycle separating the corresponding long cycles in $\Gamma(\pi)$. By Lemma 7.2, deleting that separating cycle does not change the bounds obtained in Proposition 7.1, and the proof follows. ■

We refer to sub-permutations reducing to α-permutations as *α-cycles*. Similar arguments can be used to handle the case of cycles in Γ that contain other ones, so we have the following result.

Theorem 7.1: For every π in S_n whose Γ-graph contains only 1-cycles and k non-crossing perforations of α-cycles:

$$d(\pi) \quad = \quad n - c_{odd}(\Gamma(\pi)) - K + \sum_{i=1}^{k} (|C_i| \bmod 2)\,,$$

where K is the number of edges of length 1 in $\Gamma(\pi)$.

Proof: The formula follows from Proposition 7.1 and previous observations. The correspondence with 1-edges in $\Gamma(\pi)$ was observed in the proof of Corollary 7.1, and this is the only case where deleting a 1-cycle creates an edge of length 1. ■

It is less clear how exactly a perforation would be defined in the case of crossing cycles. Even less clear is the evolution of cycles in G when deleting fixed points in this situation: it depends on how the cycles cross and on their monotonicity. We can however prove some further results on permutations whose Γ-graph has no crossing cycles, which we do in the next section.

VIII. Non-Crossing Cycles in Γ

We consider permutations with a Γ-graph of the form shown in Fig. 7, and have a look at what happens in G and Γ when deleting separators. Depending on the parity of each long cycle, the deletion of separators can have various effects.

Proposition 8.1: Let π in S_n be a permutation with $\Gamma(\pi)$ of the form shown in Fig. 7, where C_i ($1 \leq i \leq k$) is one of the following:

- an α-permutation with an odd main cycle;
- a β-permutation with an odd main cycle;
- a perforation of an α-permutation.

Then deleting j separators ($1 \leq j \leq k-1$) transforms π into a permutation with the same distance.

Proof: By Propositions 4.1 and 4.5, we have $d(\pi) = \frac{n+1-c_{odd}(G(\pi))}{2}$. Each pair (C_i, C_{i+1}) yields a pair of alternating cycles (Observation 4.1 and Lemma 7.1) that share the separator as described in Lemma 7.2. This Lemma also implies that deleting the separator does not change the lower bound of Theorem 2.1, which is tight for π, because it will decrease both n and the number of odd alternating cycles by 1. So $d(\pi)$ is a lower bound on the distance of the resulting permutation, and since $d(\pi)$ is also an upper bound on that distance (Lemma 4.2), the proof follows. ■

Although we are unable to compute the exact distance when all large cycles are even (and are not perforations of α-permutations), we can still lower (11) in that case. In order to express this improved bound formally, we need to introduce the following graph.

Definition 8.1: Given a permutation π with $\Gamma(\pi)$ of the form shown in Fig. 7, the *contact graph* $H(\pi)$ is the undirected graph whose vertices are the long cycles in $\Gamma(\pi)$ and whose edges are $\{C_i, C_{i+1}\}$ if C_i and C_{i+1} are even and not separated by a 1-cycle in $\Gamma(\pi)$.

This graph uniquely decomposes into p connected components, which we denote $\mathcal{C}_1, \ldots, \mathcal{C}_p$. The following lemma will be useful.

Lemma 8.1: Let

$$\xi_k = (\underbrace{3\ 2\ 1}_{1}\ \underbrace{6\ 5\ 4}_{2}\ \cdots\ \underbrace{n\ n-1\ n-2}_{k});$$

then $d(\xi_k) \leq \left\lceil \frac{3k}{2} \right\rceil = \left\lceil \frac{n}{2} \right\rceil$.

Proof: Since ξ_k is a simple permutation (Fig. 9), the proof follows from Theorem 2.7. ■

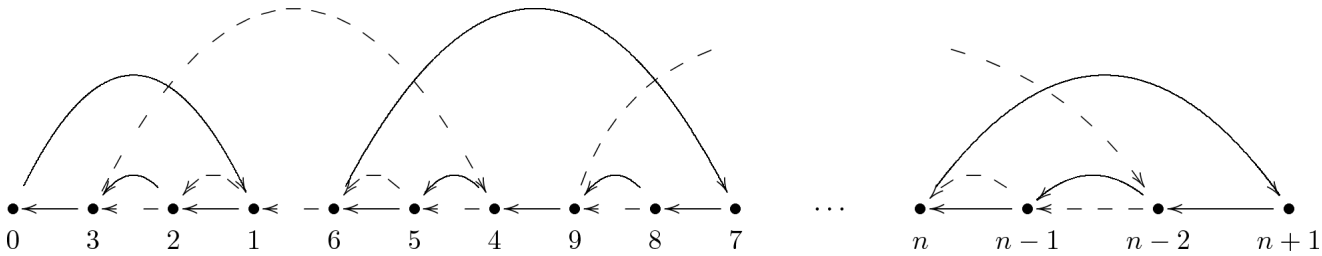

Fig. 9. The cycle graph of ξ_k

One way to sort ξ_k is to handle 2-cycles of $\Gamma(\xi_k)$ pairwise, i.e. partition ξ_k into $\lfloor \frac{k}{2} \rfloor$ sub-permutations of the form of ξ_2. Those can each be sorted optimally using three transpositions, and possibly, one last sub-permutation of the form of ξ_1 will require two transpositions. Note that this permutation is the general form of an example given by Christie [5] that shows how his improved lower bound on the transposition distance fails (meaning that even though it gives a larger value than the lower bound of Theorem 2.1, it still underestimates the true distance). Branch-and-bound seems however to indicate that the upper bound of Lemma 8.1 is the actual distance of ξ_k.

Proposition 8.2: Let π be a γ-permutation with $\Gamma(\pi)$ of the form shown in Fig. 7, where C_i $(1 \leq i \leq k)$ is either an α-permutation or a β-permutation with an even main cycle; then deleting j separators $(1 \leq j \leq k-1)$ transforms π into a permutation σ such that

$$d(\sigma) \leq d(\pi) - 2k + \sum_{i=1}^{p} \left\lceil \frac{3|\mathcal{C}_i|}{2} \right\rceil,$$

where $\mathcal{C}_i$ $(1 \leq i \leq p)$ is a connected component of $H(\sigma)$.

Proof: Instead of removing separators directly, we first apply some transpositions on π. Each sub-α-permutation can be sorted "incompletely" using the optimal sorting sequence of Proposition 4.2, without the last two transpositions. A similar process can be applied to sub-β-permutations, which first require a transformation as depicted in the proof of Proposition 4.4. By reduction, the resulting permutation has a Γ-graph of the form shown in Fig. 7, where each C_i' is now of the form of ξ_1. Let us now remove a subset of j separators $(1 \leq j \leq k-1)$ from that permutation; this will diminish the number of components in its contact graph, thus creating sub-permutations of the form of ξ_k. The following upper bound is obtained from Lemmas 4.2

and 8.1:

$$\begin{aligned} d(\sigma) &\leq \sum_{i=1}^{p} d(\mathcal{C}_i) \\ &\leq \sum_{i=1}^{p} \left(\sum_{C_j \in \mathcal{C}_i} (d(C_j) - 2) + d(\xi_{|\mathcal{C}_i|}) \right) \\ &= \sum_{i=1}^{p} \sum_{C_j \in \mathcal{C}_i} (d(C_j) - 2) + \sum_{i=1}^{p} d(\xi_{|\mathcal{C}_i|}) \\ &\leq d(\pi) - 2k + \sum_{i=1}^{p} \left\lceil \frac{3|\mathcal{C}_i|}{2} \right\rceil . \end{aligned}$$

■

An easy particular case of this proposition is when all separators are deleted; in that case, $d(\sigma) \leq d(\pi) - \left\lceil \frac{k}{2} \right\rceil$. There remains one case to deal with, which encompasses both previous Propositions.

Proposition 8.3: Let π be a γ-permutation with $\Gamma(\pi)$ of the form shown in Fig. 7, where C_i ($1 \leq i \leq k$) is one of the following:

- an α-permutation or a β-permutation with an even or an odd main cycle;
- a perforation of an α-permutation.

Then deleting j separators ($1 \leq j \leq k-1$) transforms π into a permutation σ such that

$$d(\sigma) \leq d(\pi) - 2k + \sum_{i=1}^{p} \left\lceil \frac{3|\mathcal{C}_i|}{2} \right\rceil ,$$

where $\mathcal{C}_i$ ($1 \leq i \leq p$) is a connected component of $H(\sigma)$.

Proof: As hinted by Lemma 7.2 and confirmed by previous results, the only case in which deleting a separator affects the distance of the resulting permutation is when that deletion occurs between two even cycles. This means that Proposition 8.2 naturally generalises to the case where some cycles are allowed to be odd, because deleting separators adjacent to at least one long odd cycle will not modify the distance of the resulting permutation. By the same arguments as those used in that Proposition's proof, we obtain the same upper bound on the distance of the resulting permutation, and α-permutations, β-permutations as well as perforations of the former kind can be handled individually in σ as was already done in π. ■

We conclude with the case where we allow containment and perforation of α-cycles.

Theorem 8.1: For all π in S_n with $\Gamma(\pi)$ containing only non-crossing α-cycles that are odd or perforated (possibly both) and 1-cycles, we have

$$d(\pi) \;=\; n - c_{odd}(\Gamma(\pi)) - K + \sum_{i=1}^{k} (|C_i| \bmod 2)\,,$$

where C_i $(1 \leq i \leq k)$ are the long cycles in $\Gamma(\pi)$ and K is the number of edges of length 1.

Proof: Suppose that every pair of consecutive long cycles in $\Gamma(\pi)$ is separated by a 1-cycle; since each long cycle is odd or a perforation of an α-permutation, the corresponding alternating cycles in $G(\pi)$ are all odd (Proposition 4.1 and Lemma 7.1). Therefore removing any subset of the separators cannot affect the distance (Lemma 7.2), so the strategy of Lemma 4.2 remains optimal and the proof follows from Theorem 7.1. ■

IX. Conclusions

We have exhibited connections between two different graph representations of permutations, one of which is a well-known object in combinatorics and the other one is the traditional structure used in the problem of sorting permutations by transpositions. Those connections allowed us to derive a formula for computing the distance of a non-trivial class of permutations, which we called γ-permutations. Showing how γ-permutations could be used to generate all others, we were able to prove that our formula is an upper bound on the transposition distance of every permutation. A more involved analysis of the operation used to obtain other permutations from this class allowed us to describe three additional interesting families of permutations: more instances for which our bound is tight; instances for which our bound is not tight, but for which we found other formulas to compute their distance; and finally, instances for which we can lower our upper bound without a guarantee that the obtained formula gives the exact distance.

It should be noted that (10) gives the distance of more permutations than the ones characterised in Theorem 4.1: among the other permutations for which (10) still holds are 1-perforations of α-permutations with an odd main cycle (Corollary 7.1), permutations obtained by concatenating such configurations, whether they are separated (Proposition 7.1) or not (Corollary 7.2), and permutations characterised in Proposition 8.1. Our results can also be used as upper bounds in some cases where cycles cross, for which it seems difficult to give an accurate formula or a more precise upper bound.

A few questions remain open. Although we now have a large quantity of permutations whose distance is computable in polynomial time, there are still some instances for which we have no clear answer yet. Among those are perforations of β-permutations, and permutations whose Γ-graph contains only crossing cycles and do not reduce to γ-permutations. Is it possible to compute their distance in polynomial time, or to show it is NP-hard to do it? Can an improved upper bound be given as well?

An obviously related question is that of finding the *diameter*, i.e. the maximal value the transposition distance can reach. Using permutations whose distance we know, can we give an improved upper bound on the distance of permutations that do not belong to these families, and therefore improve the upper bound of Theorem 2.9?

APPENDIX I
ON THE SEQUENCES OF PROPOSITION 4.2

Consider the following transpositions:

$$\begin{cases} \tau_1 &= \tau(2,4,n+1); \\ \tau_2 &= \tau(1,3,n); \\ \tau_3 &= \tau(2,3,n+1); \\ \tau_4 &= \tau(1,2,n+1). \end{cases}$$

Proposition 10.1: For every α-permutation π in S_n whose main cycle C is odd and increasing, the sequence

$$(\tau_2 \circ \tau_1)^{\frac{|C|-1}{2}}$$

sorts π.

Proof: Induction on $|C|$. The base case is $\pi = (3\ 2\ 5\ 4\ 1)$; we have $\pi \circ \tau_1 = (3\ 4\ 1\ 2\ 5)$, and $(3\ 4\ 1\ 2\ 5) \circ \tau_2 = \iota$.

For the induction, the permutation to sort is

$$\pi = (3\ 2\ 5\ 4\ 7\ 6\ \cdots\ n-2\ n-3\ n\ n-1\ 1).$$

Applying τ_1 to π transforms it into

$$(3\ 4\ 7\ 6\ \cdots\ n-2\ n-3\ n\ n-1\ 1\ 2\ 5)$$

to which we apply τ_2, thus transforming it into

$$(7\ 6\ \cdots\ n-2\ n-3\ n\ n-1\ 1\ 2\ 3\ 4\ 5).$$

Reducing the latter permutation merges the last five elements into a new element called 1, and subtracts 4 to every other element. It is then clear that, if σ is the permutation for which our induction hypothesis is true, then $\pi \circ \tau_1 \circ \tau_2 \equiv_r \sigma$, and this completes the proof. ■

Proposition 10.2: For every α-permutation π in S_n whose main cycle C is even and increasing, the sequence

$$\tau_4 \circ \tau_3 \circ \left(\tau_2 \circ \tau_1\right)^{\frac{|C|-2}{2}}$$

sorts π.

Proof: Similar to that of Proposition 10.1, with base case $\pi = (3\ 2\ 5\ 4\ 7\ 6\ 1)$. ■

APPENDIX II
ON THE DEFINITION OF SIMPLE PERMUTATIONS, 2-PERMUTATIONS, AND 3-PERMUTATIONS

A *signed permutation* is a permutation whose elements can be either positive or negative. Denote $S_n^{\pm}$ the group of permutations of $\{\pm 1, \pm 2, \ldots, \pm n\}$. It is not mandatory for a signed permutation to have negative elements, so $S_n \subset S_n^{\pm}$. The following graph was introduced by Bafna and Pevzner [18] in the context of sorting permutations by reversals.

Definition 11.1: Given a signed permutation π in $S_n^{\pm}$, transform it into an unsigned permutation π' in S_{2n} by replacing π_i with the sequence $(2\pi_i - 1, 2\pi_i)$ if $\pi_i > 0$, or $(2|\pi_i|, 2|\pi_i| - 1)$ if $\pi_i < 0$, for $1 \leq i \leq n$. The *breakpoint graph* of π' is the undirected bicoloured graph $BG(\pi')$ with ordered vertex set $(\pi'_0 = 0, \pi'_1, \pi'_2, \ldots, \pi'_{2n}, \pi'_{2n+1} = 2n+1)$ and whose edge set consists of:

- black edges $\{\pi'_{2i}, \pi'_{2i+1}\}$ for $0 \leq i \leq n$;
- grey edges $\{\pi'_{2i}, \pi'_{2i} + 1\}$ for $0 \leq i \leq n$.

We show that for every signed permutation π with no negative element, the cycle graph $G(\pi)$ is equivalent to the breakpoint graph $BG(\pi')$. By equivalent, we mean that every alternating cycle in $G(\pi)$ is an alternating cycle in $BG(\pi')$, and that the "topological" relations between the cycles are the same; for instance, if two cycles cross in either graph, then they also cross in the other one.

Theorem 11.1: For all π in S_n: $G(\pi) \equiv BG(\pi')$.

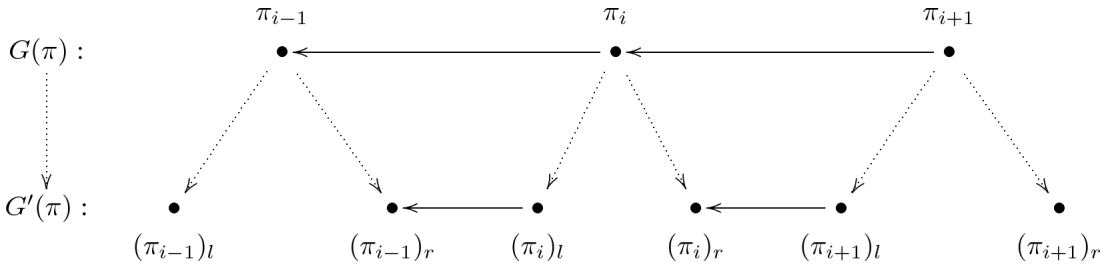

Fig. 10. Mapping of the black edges in the transformation of $G(\pi)$ into $BG(\pi')$; here, $G'(\pi)$ is a graph that will be isomorphic to $BG(\pi')$ once the orientation of edges is removed

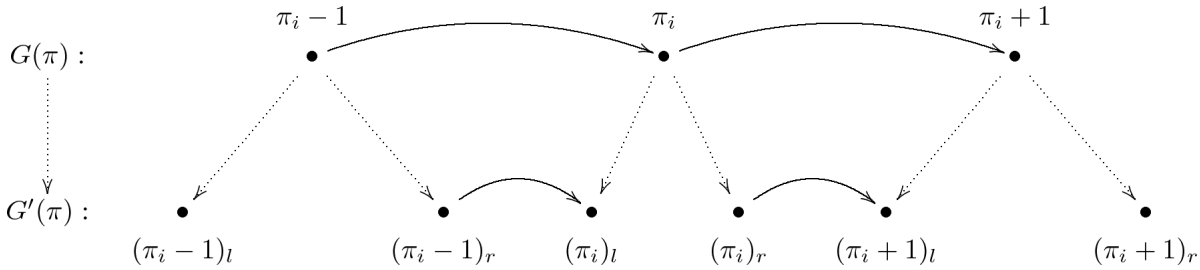

Fig. 11. Mapping of the grey edges in the transformation of $G(\pi)$ into $BG(\pi')$

Proof: We show that either graph can be constructed by transforming the other one without affecting its features. Intuitively, transforming $G(\pi)$ into $BG(\pi')$ is done by spacing black edges in $G(\pi)$ and removing the orientation; conversely, transforming $BG(\pi')$ into $G(\pi)$ is done by orienting edges in $BG(\pi')$, then merging every consecutive pair of vertices that are not connected by a black edge.

1) starting with $G(\pi)$: split each vertex π_i ($1 \leq i \leq n$) into two unconnected vertices $(\pi_i)_l$, $(\pi_i)_r$ (one to the left and one to the right), and rename π_0 (resp. π_{n+1}) into $(\pi_0)_r$ (resp. $(\pi_{n+1})_l$). Black edge (π_i, π_{i-1}) is mapped onto a new black edge $((\pi_i)_l, (\pi_{i-1})_r)$, as shown in Fig. 10. Similarly, grey edge $(\pi_i, \pi_i{+}1)$ is mapped onto a new grey edge $((\pi_i)_r, (\pi_i{+}1)_l)$, as shown in Fig. 11. Finally, rename $(\pi_i)_l$ (resp. $(\pi_i)_r$) into $2\pi_i - 1$ (resp. $2\pi_i$) and remove the orientation of edges. This results in $BG(\pi')$, since:
   a) each black edge (π_i, π_{i-1}) is mapped onto a black edge $\{(\pi_i)_l, (\pi_{i-1})_r\} = \{2\pi_i - 1, 2\pi_{i-1}\}$;
   b) each grey edge $(\pi_i, \pi_i{+}1)$ is mapped onto a grey edge $\{(\pi_i)_r, (\pi_i{+}1)_l\} = \{2\pi_i, 2\pi_i{+}1\}$.

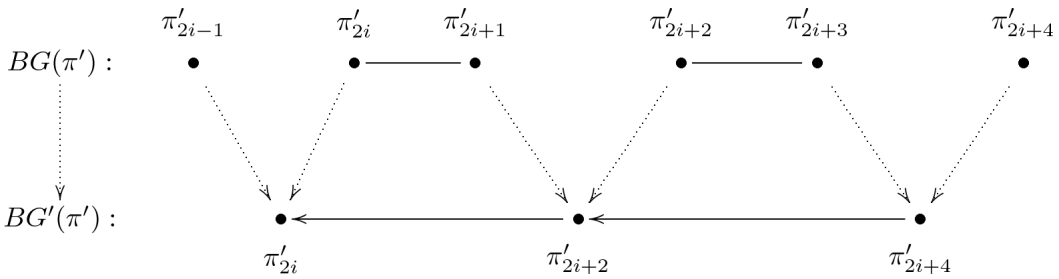

Fig. 12. Mapping of the black edges in the transformation of $BG(\pi')$ into $G(\pi)$; here, $BG'(\pi')$ is a graph isomorphic to $G(\pi)$

2) starting with $BG(\pi')$: since π' comes from some permutation π with no negative element, for all $1 \leq i \leq n$, we have $\pi'_{2i} = 2\pi_i$ and $\pi'_{2i-1} = 2\pi_i - 1$. This implies that alternating cycles in $BG(\pi')$ can be followed starting from the leftmost vertex of a black edge, then following a grey edge that will take us to the rightmost vertex of the next black edge. Therefore, adding an orientation to all edges that corresponds to this course will result in a collection of directed alternating cycles that can be followed using the direction of the arrows, and this orientation is obtained by transforming grey edge $\{\pi'_{2i}, \pi'_{2i} + 1\}$ into $(\pi'_{2i}, \pi'_{2i} + 1)$, and black edge $\{\pi'_{2i}, \pi'_{2i+1}\}$ into (π'_{2i+1}, π'_{2i}).
   Next, for $1 \leq i \leq n$, merge vertices π'_{2i-1} and π'_{2i} into vertex π'_{2i}, and rename vertex π'_{2n+1} into π'_{2n+2}; black edge (π'_{2i+1}, π'_{2i}) is mapped onto a new black edge (π'_{2i+2}, π'_{2i}).
   Finally, replace π'_{2i} with π_i, for $0 \leq i \leq n + 1$. This results in $G(\pi)$, since:
   a) each black edge $\{\pi'_{2i}, \pi'_{2i+1}\}$ is mapped onto a black edge $(\pi'_{2(i+1)}, \pi'_{2i}) = (\pi_{i+1}, \pi_i)$;
   b) each grey edge $\{\pi'_{2i}, \pi'_{2i} + 1\}$ is mapped onto a grey edge $(\pi_i, \pi_i + 1)$.

■

As in the case of the cycle graph, the *length* of a cycle in a breakpoint graph is the number of black edges it contains.

Definition 11.2: [17] A permutation π in $S_n^{\pm}$ is *simple* if $BG(\pi')$ does not contain a cycle of length greater than three.

Definition 11.3: [17] A permutation π in $S_n^{\pm}$ is a 2-*permutation* (resp. 3-*permutation*) if all cycles in $BG(\pi')$ are of length 2 (resp. 3).

Corollary 11.1: For every π in S_n, Definition 2.16 (resp. Definition 2.17) and Definition 11.2 (resp. Definition 11.3) are equivalent.

Proof: Straightforward from Theorem 11.1. ■

REFERENCES

[1] J. Meidanis and J. Setubal, *Introduction to Computational Molecular Biology*. Brooks-Cole, 1997.

[2] P. A. Pevzner, *Computational molecular biology*. Cambridge, MA: MIT Press, 2000.

[3] V. Bafna and P. A. Pevzner, "Sorting permutations by transpositions," in *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. San Francisco, CA: ACM/SIAM, Jan. 1995, pp. 614–623.

[4] ——, "Sorting by transpositions," *SIAM Journal on Discrete Mathematics*, vol. 11, no. 2, pp. 224–240 (electronic), May 1998.

[5] D. A. Christie, "Genome rearrangement problems," Ph.D. dissertation, University of Glasgow, Scotland, Aug. 1998.

[6] T. Hartman and R. Shamir, "A simpler and faster 1.5-approximation algorithm for sorting by transpositions," *Information and Computation*, vol. 204, no. 2, pp. 275–290, Feb. 2006.

[7] I. Elias and T. Hartman, "A 1.375-approximation algorithm for sorting by transpositions," in *Proceedings of the Fifth Workshop on Algorithms in Bioinformatics*, ser. Lecture Notes in Bioinformatics, R. Casadio and G. Myers, Eds., vol. 3692. Mallorca, Spain: Springer-Verlag, Oct. 2005, pp. 204–214.

[8] S. A. Guyer, L. S. Heath, and J. P. Vergara, "Subsequence and run heuristics for sorting by transpositions," in *Fourth DIMACS Algorithm Implementation Challenge*, Rutgers University, Aug. 1995.

[9] J. P. C. Vergara, "Sorting by bounded permutations," Ph.D. dissertation, Virginia Polytechnic Institute, Blacksburg, Virginia, USA, Apr. 1997.

[10] M. E. M. T. Walter, L. R. A. F. Curado, and A. G. Oliveira, "Working on the problem of sorting by transpositions on genome rearrangements," in *Proceedings of the Fourteenth Annual Symposium on Combinatorial Pattern Matching*, ser. Lecture Notes in Computer Science. Berlin: Springer, 2003, vol. 2676, pp. 372–383.

[11] Z. Dias and J. Meidanis, "An alternative algebraic formalism for genome rearrangements," *Comparative Genomics: Empirical and Analytical Approaches to Gene Order Dynamics, Map Alignment and the Evolution of Gene Families*, vol. 1, pp. 213–223, 2000.

[12] ——, "Genome rearrangements distance by fusion, fission, and transposition is easy," in *Proceedings of the Eighth International Symposium on String Processing and Information Retrieval*. Laguna de San Rafael, Chile: IEEE Computer Society Press, Nov. 2001, pp. 250–253.

[13] A. Labarre, "A new tight upper bound on the transposition distance," in *Proceedings of the Fifth Workshop on Algorithms in Bioinformatics*, ser. Lecture Notes in Bioinformatics, R. Casadio and G. Myers, Eds., vol. 3692. Mallorca, Spain: Springer-Verlag, Oct. 2005, pp. 216–227.

[14] H. Eriksson, K. Eriksson, J. Karlander, L. Svensson, and J. Wästlund, "Sorting a bridge hand," *Discrete Mathematics*, vol. 241, no. 1-3, pp. 289–300, 2001, selected papers in honor of Helge Tverberg.

[15] A. Hultman, "Toric permutations," Master's thesis, Dept. of Mathematics, KTH, Stockholm, Sweden, 1999.

[16] Z. Dias, J. Meidanis, and M. E. M. T. Walter, "A new approach for approximating the transposition distance," in *Proceedings of the Seventh International Symposium on String Processing and Information Retrieval*. La Coruña, Spain: IEEE Computer Society Press, Sept. 2000, pp. 199–208.

[17] S. Hannenhalli and P. A. Pevzner, "Transforming cabbage into turnip: Polynomial algorithm for sorting signed permutations by reversals," *Journal of the ACM*, vol. 46, no. 1, pp. 1–27, Jan. 1999.

[18] V. Bafna and P. A. Pevzner, "Genome rearrangements and sorting by reversals," in *Proceedings of the Thirty-Fourth Annual Symposium on Foundations of Computer Science*. Palo Alto, Los Alamitos, CA: ACM/SIAM, 1993, pp. 148–157.

Anthony Labarre obtained a Master's degree in Computer Science in 2004 and a DEA degree (Diplôme d'Études Approfondies) in Sciences in 2005, both at the Université Libre de Bruxelles, Brussels, Belgium. He is currently working on a PhD thesis. His research interests include genome rearrangements, phylogenetic networks, and enumerative combinatorics.