



On the distance between the expressions of a permutation

Marc Autord, Patrick Dehornoy

► To cite this version:

Marc Autord, Patrick Dehornoy. On the distance between the expressions of a permutation. 2009. hal-00362372

HAL Id: hal-00362372

<https://hal.science/hal-00362372>

Preprint submitted on 18 Feb 2009

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.

ON THE DISTANCE BETWEEN THE EXPRESSIONS OF A PERMUTATION

MARC AUTORD AND PATRICK DEHORNOY

ABSTRACT. We prove that the combinatorial distance between any two reduced expressions of a given permutation of $\{1, \dots, n\}$ in terms of transpositions lies in $O(n^4)$, a sharp bound. Using a connection with the intersection numbers of certain curves in van Kampen diagrams, we prove that this bound is sharp, and give a practical criterion for proving that the derivations provided by the reversing algorithm of [Dehornoy, JPAA 116 (1997) 115-197] are optimal. We also show the existence of length ℓ expressions whose reversing requires $C\ell^4$ elementary steps.

This paper is about the various ways of expressing a permutation as a product of transpositions and the complexity of transforming one such expression into another. We consider both the absolute complexity (“combinatorial distance”), which deals with the minimal possible number of steps, and the more specific complexity (“reversing complexity”), which arises when one uses subword reversing, a certain prescribed strategy for transforming expressions.

Throughout the paper, we denote by $\llbracket 1, n \rrbracket$ the set $\{1, 2, \dots, n\}$, and by s_i the transposition that exchanges i and $i + 1$. A well known result—see for instance [8]—states that, if π is any permutation of $\llbracket 1, n \rrbracket$ and u, v are any two reduced (*i.e.*, minimal length) expressions of π in terms of s_1, \dots, s_{n-1} , then one can transform u into v only using the braid relations

$$\begin{aligned} \text{(I)} \quad & s_i s_j s_i = s_j s_i s_j \quad \text{with } |i - j| = 1, \\ \text{(II)} \quad & s_i s_j = s_j s_i \quad \text{with } |i - j| \geq 2. \end{aligned}$$

In this context, we define the *combinatorial distance* $\text{dist}(u, v)$ of u and v to be the minimal number of braid relations needed to transform u into v . The standard proof for the finiteness of $\text{dist}(u, v)$ relies on the so-called Exchange Lemma of Coxeter groups, and it leads to an exponential upper bound for $\text{dist}(u, v)$ in terms of n . The first aim of this paper is to establish a polynomial upper bound, namely a sharp degree 4 one. Using “*n-expression*” as a shorthand for “expression representing a permutation of $\llbracket 1, n \rrbracket$ ”, *i.e.*, involving letters about s_1, \dots, s_{n-1} only, we prove

Proposition 1. *There exist positive constants C_1, C_2 such that, for each n ,*

- *all equivalent reduced n -expressions u, v satisfy $\text{dist}(u, v) \leq C_1 n^4$,*
- *there exist equivalent n -expressions u, v satisfying $\text{dist}(u, v) \geq C_2 n^4$.*

(The values $C_1 = 1/2$ and $C_2 = 1/8$ are valid for n large enough.)

The methods we use are geometrical. For the upper bound, we consider some area in the n -strand braid diagram naturally associated with an n -expression. For the lower bound, we consider van Kampen diagrams and introduce certain curves

1991 *Mathematics Subject Classification.* 20B30, 05E15, 20F55, 20F36.

Key words and phrases. permutation, transposition, reduced expression, combinatorial distance, van Kampen diagram, subword reversing, braid relation.

called separatrices, which are associated with the names of the strands involved in the successive crossings.

The latter notion, which seems of independent interest, provides general criteria for proving that a van Kampen diagram (*i.e.*, in algebraic terms, a derivation by braid relations) is possibly optimal, *i.e.*, it involves the minimal number of braid relations. In particular, Proposition 1.11 below states that a sufficient condition for a van Kampen diagram to be optimal is that any two separatrices cross at most once in it.

In the second part, we address similar complexity issues in the particular case of subword reversing. This is a specific strategy that, given two equivalent expressions u, v , returns a derivation of v from u by means of braid relations, *i.e.*, equivalently, constructs a van Kampen diagram for the pair (u, v) . We observe on a simple counter-example that the reversing method need not be optimal, but we deduce from the above approach based on separatrices a simple optimality criterion, namely Proposition 2.5 that states that a sufficient condition for the reversing method to be optimal for some pair (u, v) is that the so-called reversing diagram for (u, v) contains no digon, *i.e.*, in algebraic terms, the reversing sequence from $\bar{u}v$ contains no ϵ -step—all technical terms are defined below.

Finally, we address the general question of the complexity of the reversing method. Frustratingly, the only upper bound we can establish at the moment is exponential—this does not contradict the polynomial upper bound of Proposition 1, since reversing need not be optimal. On the other hand, the optimal lower bound of Proposition 1 induces a similar lower bound in the case of reversing. What is more interesting is to consider the case of non-necessarily equivalent expressions. In that case, the reversing method still applies, and its complexity remains widely unknown. The relevant question is to determine the number $\text{compl}_{\curvearrowright}(u, v)$ of elementary steps when one starts with expressions u, v of length ℓ (independently from the index n). When the lower bound of Proposition 1 is translated in this language, it leads to a quadratic lower bound $\text{compl}_{\curvearrowright}(u, v) \geq C\ell^2$. This value is far from optimal.

Proposition 2. *There exists a positive constant C_3 such that, for each ℓ ,*

- *there exist length ℓ expressions u, v satisfying $\text{compl}_{\curvearrowright}(u, v) \geq C_3\ell^4$.*

(For ℓ large enough, we can take $C_3 = 4/3$.) At the moment, we do not know whether the above result is optimal.

It is likely that most results of this paper extend to finite Coxeter groups of other type. However, the arguments developed here heavily rely on specific properties of permutations, so how extending to Coxeter types other than A and B is not clear.

Remark. Most constructions developed in this paper in the case of permutations and their expressions in terms of transpositions can be extended to the case of positive braid and their decompositions in terms of Artin's generators σ_i . Technically, the case of permutations corresponds to the particular case of the so-called simple braids, which are the divisors of Garside's fundamental braid Δ_n in the braid monoid B_n^+ , see [7] or [6, Chap. 9]. Our reason for choosing the language of permutations here is that it is more widely accessible and it avoids introducing the general framework of braids whereas most results would involve simple braids exclusively. Indeed, it turns out that the worst cases known so far, in particular in

terms of subword reversing, always involve simple braids. We have no explanation for this phenomenon.

We finally mention that the results of Section 1 are mainly due to the second author, whereas those of Section 2 are mainly due to the first author.

1. THE COMBINATORIAL DISTANCE

Hereafter, a word u on the alphabet $\{s_1, \dots, s_{n-1}\}$ is generically called an n -expression. Two n -expressions are called *equivalent* if they represent the same permutation of $\llbracket 1, n \rrbracket$. Throughout the paper (in particular in view of the braid diagrams considered below), it is convenient that the product, both for words and for permutations, refers to reverse composition: uv means “ u first, then v ”. An n -expression u is called *reduced* if the permutation represented by u has no expression that is shorter than u .

If u and v are equivalent reduced n -expressions, then, as recalled above, one can transform u into v using the braid relations of types I and II, and we denote by $\text{dist}(u, v)$ the minimal number of braid relations needed to do it. In this section, we establish bounds for $\text{dist}(u, v)$ when n grows to infinity.

1.1. An upper bound result. We begin with an upper bound result. To this end, we introduce one distinguished reduced expression, called *normal*, for each permutation, and we define a strategy that transforms any reduced expression into the (unique) normal expression that represents the same braid.

For each n -expression u , we define D_u to be the n -strand braid diagram obtained by associating with the letter s_i the pattern

$$(1.1) \quad \begin{array}{ccccccc} 1 & 2 & & i & i+1 & & n \\ | & | & \dots & | & \times & | & \dots & | \end{array}$$

and by stacking from top to bottom the elementary patterns corresponding to the successive letters of u . When we speak of the p th strand in D_u , we refer to the strand that starts at the p th position from the left on the top line. It is well known that u is a reduced expression if and only if any two strands in D_u cross at most once [6, Chap. 9].

Definition. Define $s_{j,i}$ to be $s_{j-1}s_{j-2}\dots s_{i+1}s_i$ for $j > i$, and to be the empty word ϵ for $j = i$. An n -expression is called *normal* if it has the form

$$(1.2) \quad s_{1,f(1)} s_{2,f(2)} \dots s_{n,f(n)},$$

for some function $f : \llbracket 1, n \rrbracket \rightarrow \llbracket 1, n \rrbracket$ satisfying $f(i) \leq i$ for each i .

(The first factor $s_{1,f(1)}$ is mentioned for symmetry, but is necessarily empty.)

Lemma 1.1. (i) *Every normal expression is reduced.*

(ii) *Each permutation of $\llbracket 1, n \rrbracket$ admits a unique normal n -expression.*

Proof. (i) Assume that u has the form (1.2). Then, for each i , the i th strand crosses over no j th strand with $j > i$ in the diagram D_u . Therefore, any two strands cross at most once in D_u , so u is a reduced expression.

(ii) For f satisfying $f(i) \leq i$ for each i , let n_f be the largest m satisfying $f(m) < m$, if any, and 0 otherwise, and let π_f be the permutation represented by the expression (1.2) associated with f . We prove that π_f determines f using induction on n_f . For $n_f = 0$, the only possibility is that $f(i) = i$ holds for each i , so

the permutation π_f is the identity. Assume now $n_f \geq 1$. By construction, we have $\pi_f(i) = i$ for $i \geq n_f$, and $\pi_f(n_f) = f(n_f)$. Hence π_f determines n_f and $f(n_f)$. Next, let f' be defined by $f'(i) = f(i)$ for $i < n_f$, and $f'(i) = i$ for $i \geq n_f$. Then we have $\pi_{f'} = \pi_f s_{n_f, f(n_f)}$, hence π_f determines $\pi_{f'}$. By construction, we have $n_{f'} \leq n_f - 1$. By induction hypothesis, $\pi_{f'}$ determines f' , hence so does π_f . Finally, f is determined by f' , n_f , and $f(n_f)$, hence by π_f . \square

For each (reduced) expression u , we denote by $\text{NF}(u)$ the unique normal expression that is equivalent to u . We shall now define a strategy for transforming u into $\text{NF}(u)$.

First, we concentrate on the last factor of the normal form. We have associated with every n -expression u a braid diagram D_u . We shall assume that the pattern (1.1) is drawn in a rectangle that has width $n - 1$ and height 1. So, if u is an n -expression of length ℓ , the diagram D_u is drawn in an $(n - 1) \times \ell$ grid. It includes $(n - 1)\ell$ squares of size 1, and it makes sense to count how many such squares lie on the left or on the right of a given strand.

Lemma 1.2. *For each n -expression u , define $a(u)$ to be the number of plain squares lying on the right of the n th strand in the diagram D_u . Then, for each reduced n -expression u , there exists an equivalent reduced expression $vs_{n,k}$, with v an $(n - 1)$ -expression, satisfying*

$$(1.3) \quad \text{dist}(u, vs_{n,k}) \leq a(u).$$

For an induction it is enough to establish


Claim. *If u is not of the form $vs_{n,k}$ with v an $(n - 1)$ -expression, there exists an n -expression u' satisfying $\text{dist}(u, u') = 1$ and $a(u') < a(u)$.*

Proof. Let p be the final position of the n th strand in D_u . The hypothesis implies $p < n$ holds as, otherwise, u itself would be an $(n - 1)$ -expression and it could be expressed as $us_{n,n}$. Then there exists a unique decomposition


$$u = v s_{n,i} s_j w$$

with v an $(n - 1)$ -expression, $i < n$, and $j \neq i - 1$: we consider the first block of crossings $s_{n,i}$ in which the n th strand is the front strand, and the hypothesis on u means that, after that block, there still remains at least one crossing s_j in which the n th strand is not the front strand, which means $j \neq i - 1$. We consider the various possible values of j with respect to i . First, $j = i - 1$ is excluded by hypothesis, whereas $j = i$ would contradict the hypothesis that u is reduced since the n th strand would cross the same strand twice. There remain two cases only.

Case 1: $|j - i| \geq 2$. Put $u' = vs_{n,i+1}s_j s_i w$. Then we have $\text{dist}(u, u') = 1$, and $a(u') = a(u) - 1$, as shown in the following diagrams, which compare the contributions of the factors $s_i s_j$ and $s_j s_i$ to the right hand side area of the n th strand (the n th strand is in bold, and the squares contributing to a are in grey)



Case 2: $j = i + 1$. By construction this may happen only for $i \leq n - 2$. Let $u' = vs_{n,i+2}s_is_js_iw$. Then we have again $\text{dist}(u, u') = 1$, and $a(u') = a(u) - 2$, as shown in the diagram



So the proof of the claim is complete, and the lemma follows. \square

Repeated uses of Lemma 1.2 lead to

Lemma 1.3. *For every reduced n -expression u of length ℓ , we have*

$$(1.4) \quad \text{dist}(u, \text{NF}(u)) \leq n(n-1)\ell/2.$$

Proof. We use induction on n . The result is obvious for $n \leq 3$. Assume $n \geq 4$. Assume that the last factor in $\text{NF}(u)$ is $s_{n,k}$. By Lemma 1.2, there exists a reduced $(n-1)$ -expression v of length at most ℓ —actually of length exactly $\ell - (n-k)$ —satisfying $\text{dist}(u, vs_{n,k}) \leq a(u)$. By construction, we have $a(u) \leq (n-1)\ell$. On the other hand, the uniqueness of the normal form implies $\text{NF}(u) = \text{NF}(v)s_{n,k}$. Hence, using the induction hypothesis, we find

$$\text{dist}(u, \text{NF}(u)) \leq (n-1)\ell + (n-1)(n-2)\ell/2 = n(n-1)\ell/2. \quad \square$$


We immediately deduce

Proposition 1.4. *For all equivalent n -expressions u, v of length ℓ , we have*

$$(1.5) \quad \text{dist}(u, v) \leq (n-1)(n-2)\ell.$$

As every reduced n -expression has length at most $n(n-1)/2$, the upper bound $O(n^4)$ of Proposition 1 follows.

Remark 1.5. Considering the rightmost strand in the argument of Lemma 1.2 is essential. Indeed, for each reduced expression u and each i such that $s_i u$ is not reduced, *i.e.*, such that the i th and the $i+1$ st strands cross in D_u , we can consider the area a_i of the domain bounded by the top line and the i th and $i+1$ st strands before they cross. It is natural to wonder whether u can be transformed into an equivalent expression $s_i v$ in such a way that the parameter a_i decreases at each step—thus obtaining a new proof of the Exchange Lemma. The answer is negative. Indeed, assume $u = s_1 s_3 s_2 s_1 s_3 s_2$. Then the second and third strands cross in D_u and $s_2 u$ is not reduced: u is equivalent to $s_2 s_3 s_2 s_1 s_2 s_3$. However, there is no way to apply a braid relation to u so as to decrease the area a_2 of the domain bounded by the second and third strands. Indeed, the two expressions at distance 1 from u are $u' = s_3 s_1 s_2 s_1 s_3 s_2$ and $u'' = s_1 s_3 s_2 s_3 s_1 s_2$, which satisfy $a_2(u) = a_2(u') = a_2(u'') = 9$, as shown in the diagrams



in each of which nine grey squares occur.

1.2. A lower bound result. We turn to the other direction, namely proving lower bounds on the combinatorial distance of two equivalent reduced expressions. To this end, we associate a name to each letter in a reduced expression and observe that applying one braid relation can only change the associated sequence of names by a limited amount.

Notation. Hereafter, we use $\llbracket 1, n \rrbracket^{(k)}$ for the set of all subsets of $\llbracket 1, n \rrbracket$ that have cardinality k (exactly), and $\llbracket 1, n \rrbracket^{(2,2)}$ for the set of all subsets of $\llbracket 1, n \rrbracket^{(2)}$ that have cardinality 2, *i.e.*, the set of all non-degenerate pairs of pairs in $\llbracket 1, n \rrbracket$.

By construction, every crossing in a braid diagram D_u involves two strands, each of which has an initial position that corresponds to an integer in $\llbracket 1, n \rrbracket$. By considering the initial positions of the strands that cross there, we associate with each instance of a letter s_i in an n -expression a well defined pair $\{p, q\}$ in $\llbracket 1, n \rrbracket^{(2)}$, hereafter called its *name*.

Definition. For each reduced n -expression u , we define $S(u)$ to be the sequence composed of the names of the successive letters in u .

So, formally, $S(u)$ is the sequence in $\llbracket 1, n \rrbracket^{(2)}$ recursively defined by $S(\epsilon) = \emptyset$ (the empty sequence) and, using \frown for concatenation,

$$S(u) = S(v) \frown (\{p, q\}),$$

assuming that $u = vs_i$ and the strands that finish at positions i and $i + 1$ in D_v are the p th and the q th ones, *i.e.*, start at positions p and q , respectively.

Example 1.6. Let $u_n = s_{1,1}s_{2,1}\dots s_{n,1}$. Then u_n is a reduced expression of the Coxeter element of \mathfrak{S}_n , *i.e.*, of the flip permutation ϕ that exchanges i and $n - i$ for each i . An easy induction gives

$$(1.6) \quad S(u_n) = (\{1, 2\}, \{1, 3\}, \{2, 3\}, \dots, \{n-2, n-1\}, \{1, n\}, \{2, n\}, \dots, \{n-1, n\}).$$

Symmetrically, let v_n be the expression obtained from u_n by reversing the order of the factors and flipping their entries, *i.e.*, $v_n = s_{n,1}s_{n,2}\dots s_{n,n-1}$. Then v_n is another reduced expression of ϕ , and we find

$$(1.7) \quad S(v_n) = (\{n-1, n\}, \dots, \{2, n\}, \{1, n\}, \{n-2, n-1\}, \dots, \{2, 3\}, \{1, 3\}, \{1, 2\}) :$$

so $S(v_n)$ is the sequence obtained by reversing the entries of $S(u_n)$.

By construction, if u is a reduced n -expression, the pair $\{p, q\}$ occurs in $S(u)$ if and only if the strands starting at positions p and q cross in the diagram D_u , hence if and only if $\{p, q\}$ is an inversion of the permutation represented by u . Hence, if u, v are equivalent reduced n -expressions, the pairs occurring in $S(u)$ and in $S(v)$ coincide, and $S(v)$ is a permuted image of $S(u)$.

We shall see now that comparing the sequences $S(u)$ and $S(v)$ leads to a lower bound on the combinatorial distance between u and v .

Definition. If S, S' are enumerations of $\llbracket 1, n \rrbracket^{(2)}$, we denote by $I_3(S, S')$ (*resp.* $I_{2,2}(S, S')$) the number of triples $\{p, q, r\}$ in $\llbracket 1, n \rrbracket^{(3)}$ (*resp.* the number of pairs of pairs $\{\{p, q\}, \{p', q'\}\}$ in $\llbracket 1, n \rrbracket^{(2,2)}$) such that the order of the pairs $\{p, q\}$, $\{p, r\}$, $\{q, r\}$ (*resp.* the order of $\{p, q\}$ and $\{p', q'\}$) is not the same in S and S' .

Proposition 1.7. *For all equivalent n -expressions u, v , we have*

$$(1.8) \quad \text{dist}(u, v) \geq I_3(S(u), S(v)) + I_{2,2}(S(u), S(v)).$$

More precisely, every derivation from u to v contains at least $I_3(S(u), S(v))$ relations of type I and $I_{2,2}(S(u), S(v))$ relations of type II.

Proof. We first consider the case $\text{dist}(u, v) = 1$, i.e., the case when v is obtained from u by applying one braid relation. Assume first that v is obtained by applying a type I relation $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$. Then there exists a unique triple $\{p, q, r\}$, namely the names of the three strands involved in the transformation, such that the sequence $S(v)$ is obtained from $S(u)$ by replacing the subsequence $\{p, q\}, \{p, r\}, \{q, r\}$ with $\{q, r\}, \{p, r\}, \{p, q\}$ —see Figure 1 below for an illustration. So the order of the three pairs arising from $\{p, q, r\}$ has changed between $S(u)$ and $S(v)$. On the other hand, any other triple in $\llbracket 1, n \rrbracket^{(3)}$ has at most two elements in common with $\{p, q, r\}$, and the order of the three pairs arising from that triple is the same in $S(u)$ and $S(v)$. Hence, we have $I_3(S(u), S(v)) = 1$ in this case. On the other hand, any pair of pairs in $\llbracket 1, n \rrbracket^{(2,2)}$ contains at most one of the three pairs $\{p, q\}, \{p, r\}, \{q, r\}$ and, therefore, the order of this pair is not changed from $S(u)$ and $S(v)$. Hence, we have $I_{2,2}(S(u), S(v)) = 0$ in this case.

Assume now that v is obtained by applying a type II relation $s_i s_j = s_j s_i$ with $|j - i| \geq 2$. Then there exists a unique pair of pairs $\{p, q\}, \{p', q'\}$ in $\llbracket 1, n \rrbracket^{(2,2)}$, namely the names of the four strands involved in the transformation, such that the sequence $S(v)$ is obtained from $S(u)$ by replacing the subsequence $\{p, q\}, \{p', q'\}$ with $\{p', q'\}, \{p, q\}$. So the order of the two considered pairs has changed between $S(u)$ and $S(v)$. On the other hand, any other pair of pairs in $\llbracket 1, n \rrbracket^{(2,2)}$ contains at most one of the two pairs $\{p, q\}, \{p', q'\}$, and its order is the same in $S(u)$ and $S(v)$. Hence, we have $I_{2,2}(S(u), S(v)) = 1$. Moreover, any triple in $\llbracket 1, n \rrbracket^{(3)}$ gives rise to a triple of pairs that contains at most one of $\{p, q\}, \{p', q'\}$. Hence the order of the three pairs in this triple has not changed from $S(u)$ and $S(v)$, and we have $I_3(S(u), S(v)) = 0$.

We conclude that, in every case, the quantity $I_3(S(u), S(v)) + I_{2,2}(S(u), S(v))$ changes by not more than one (and even by exactly one) when one braid relation is applied. This clearly implies (1.8). \square

We can now complete the proof of Proposition 1.

Corollary 1.8. *For each n , there exist reduced n -expressions u, v satisfying*

$$(1.9) \quad \text{dist}(u, v) \geq \frac{1}{8}n^4 + O(n^3).$$

Proof. Consider the expressions u_n, v_n of Example 1.6. As observed above, the sequences $S(u_n)$ and $S(v_n)$ are mirror images of one another. It follows that each triple $\{p, q, r\}$ in $\llbracket 1, n \rrbracket^{(3)}$ contributes 1 to the parameter I_3 , leading to

$$I_3(S(u_n), S(v_n)) = \#(\llbracket 1, n \rrbracket^{(3)}) = \binom{n}{3}.$$

Similarly, each pair of pairs $\{\{p, q\}, \{p', q'\}\}$ in $\llbracket 1, n \rrbracket^{(2,2)}$ contributes 1 to the parameter $I_{2,2}$, giving

$$I_{2,2}(S(u_n), S(v_n)) = \#(\llbracket 1, n \rrbracket^{(2,2)}) = \frac{1}{2} \binom{n}{2} \binom{n-2}{2} = 3 \binom{n}{4}. \quad \square$$

Owing to the last sentence in Proposition 1.7, we also deduce from the above computation the result that the number of type I braid relations occurring in every derivation of v from u is at least $n^3/6 + O(n^2)$.

In the context of Proposition 1.7, it is natural to wonder whether (1.8) is always an equality, as it turns out to be in simple cases.

Question 1.9. *Does the equality*

$$(1.10) \quad \text{dist}(u, v) = I_3(S(u), S(v)) + I_{2,2}(S(u), S(v))$$

hold for all equivalent n -expressions u, v ?

So far we have not been able to obtain any answer. In particular, computer tries failed to find expressions disproving the equality (1.10).

Remark 1.10. A result similar to Proposition 1.7 can be obtained by simply counting the inversion number $I(S(u), S(v))$ of the sequences $S(u)$ and $S(v)$, *i.e.*, the total number of pairs of pairs whose order is changed. One easily checks that $I(S(u), S(v))$ is changed by three when a type I relation is applied, and by one in the case of a type II relation, thus leading to


$$\text{dist}(u, v) \geq I(S(u), S(v))/3.$$

Also, it can be mentionned that using areas in braid diagrams as in Section 1.1 can also lead to lower bounds on the combinatorial distance. Indeed, it is easy to check that applying one braid relation can change such areas by a bounded factor K only, leading to inequalities of the generic form

$$\text{dist}(u, v) \geq |\text{area}(D_u) - \text{area}(D_v)|/K.$$

1.3. Van Kampen diagrams. Proposition 1.7 and the approach of Section 1.2 leads to a nice geometric criterion for proving that a derivation between two reduced expressions u, v of a permutation π is possibly optimal, *i.e.*, that it realizes the combinatorial distance $\text{dist}(u, v)$ between u and v .

Assume that u, v are equivalent expressions. A *van Kampen diagram* for (u, v) is a planar connected diagram \mathcal{K} consisting of finitely many adjacent tiles of the two types



and such that the boundary of \mathcal{K} consists of two paths labeled u and v . Because of the orientation of the edges of the tiles, such a diagram has exactly one initial vertex (source) and one terminal vertex (sink), and two boundary paths labeled u and v from the source to the sink.

It is standard—see for instance [9] or [6]—that, if u, v are reduced expressions, then v can be derived from u using braid relations if and only if there exists a van Kampen diagram for (u, v) . See Figure 2 for an example. More precisely, if v can be derived from u using N braid relations, then there exists a van Kampen diagram for (u, v) that contains N tiles, and conversely. So, we have the natural notion of an optimal (or minimal) van Kampen diagram:

Definition. Assume that \mathcal{K} is a van Kampen diagram for (u, v) . We say that \mathcal{K} is *optimal* if the number of tiles in \mathcal{K} equals the combinatorial distance $\text{dist}(u, v)$.

In other words, a van Kampen diagram is declared optimal if there exists no smaller (*i.e.*, containing less tiles) diagram with the same boundary. Our aim in the sequel will be to describe criteria for recognizing that a van Kampen diagram is possibly optimal. The first idea is to use names. Assume that \mathcal{K} is a van Kampen diagram for a pair (u, v) . Then we can unambiguously attribute a *name* with every tile of \mathcal{K} . First, we attribute a name to each edge of \mathcal{K} . Let e be such an edge. It follows from the definition of a van Kampen diagram that there exists at least one path that connects the source of \mathcal{K} to any vertex V , and at least one path that connects V to the sink of \mathcal{K} . Hence there is a path γ that connects the source of \mathcal{K} to its sink and contains e . Then the successive labels of the edges of γ make an expression w , which is certainly equivalent to u (and v) since, by construction, there is a van Kampen diagram for (u, w) , namely the subdiagram of \mathcal{K} bounded by γ and the path labeled u . Then, as in Section 1.2, we attribute a name to each letter of w , and copy these names on the corresponding edges of γ . In this way, we have given a name to e which is a certain pair in $\llbracket 1, n \rrbracket^{(2)}$. We claim that this name only depends on e , and not on the choice of the path γ . As any two paths γ, γ' correspond to equivalent expressions w, w' , it is enough to consider the case when w and w' are deduced from one another using one braid relation, and the result is then obvious. Then, we observe that the names occurring on the edges of an elementary tile can be of two types only, namely those displayed in Figure 1.




FIGURE 1. Giving names to the tiles in a van Kampen diagram: the left tile is called $\{p, q, r\}$, the right one is called $\{\{p, q\}, \{p', q'\}\}$.

Definition. (See Figures 1 and 2.) Assume that \mathcal{K} is a van Kampen diagram for (u, v) . We define the *name* of a tile in \mathcal{K} as follows:

- for a type I tile, it is the (unique) triple $\{p, q, r\}$ of $\llbracket 1, n \rrbracket^{(3)}$ such that the names of the border edges are $\{p, q\}$, $\{p, r\}$, and $\{q, r\}$;
- for a type II tile, it is the (unique) pair $\{\{p, q\}, \{p', q'\}\}$ of $\llbracket 1, n \rrbracket^{(2,2)}$ such that the names of the border edges are $\{p, q\}$ and $\{p', q'\}$.

Then the optimality criterion is

Proposition 1.11. *A van Kampen diagram in which any two tiles have different names is optimal.*

Proof. The proof of Proposition 1.7 shows that, if \mathcal{K} contains exactly one tile named $\{p, q, r\}$, then the order of $\{p, q\}$, $\{p, r\}$, $\{q, r\}$ has changed between $S(u)$ and $S(v)$. Hence, under the hypothesis, if \mathcal{K} contains N hexagons with pairwise distinct names, then we have $I_3(S(u), S(v)) \geq N$. Similarly, if \mathcal{K} contains exactly one tile named $\{\{p, q\}, \{p', q'\}\}$, then the order of $\{p, q\}$ and $\{p', q'\}$ has changed between $S(u)$ and $S(v)$. So, if \mathcal{K} contains N' squares with pairwise distinct names, then we have $I_{2,2}(S(u), S(v)) \geq N'$. Applying Proposition 1.7, we deduce $\text{dist}(u, v) \geq N + N'$, hence $\text{dist}(u, v) = N + N'$. \square





FIGURE 2. Two van Kampen diagrams for the pair $(s_1s_2s_1s_3s_2s_1, s_3s_2s_3s_1s_2s_3)$: a tiling by squares and hexagons witnessing the equivalence of the two expressions. Indexed by pairs, the separatrices are drawn in grey, inducing a name for every tile. In the top diagram, no two tiles share the same name, hence the diagram is optimal, and the combinatorial distance is 6; by contrast, two tiles in the bottom diagrams share the name $\{\{1,4\}, \{2,3\}\}$, hence the diagram cannot be optimal.

1.4. Separatrices. Having given names to the edges of a van Kampen diagram, we can now draw curves that connect the edges with the same name.

Definition. Assume that \mathcal{K} is a van Kampen diagram for (u, v) , and $\{p, q\}$ is an inversion of the permutation π represented by u and v , i.e., we have $(q - p)(\pi(q) - \pi(p)) < 0$. The $\{p, q\}$ -separatrix in \mathcal{K} is the curve $\Sigma_{p,q}$ obtained by connecting the middle points of the (diametrically opposed) edges named $\{p, q\}$ inside each tile, oriented so that the edges of u are crossed first, and those of v last.

If $\{p, q\}$ is not an inversion of the involved permutation, we may consider that $\Sigma_{p,q}$ still exists, but it lies outside the diagram and cuts no edge.

So, by definition, separatrices are obtained by connecting patterns of the form



To make $\Sigma_{p,q}$ unique, we might require in addition that the curve we choose inside each tile is the image of the median lines under the affine transformation that maps a regular polygon to the considered tile (provided the latter is convex).

Then we have the following optimality criterion involving separatrices.

Proposition 1.12. *A van Kampen diagram in which any two separatrices cross at most once is optimal.*

Proof. By construction, for each triple $\{p, q, r\}$ in $\llbracket 1, n \rrbracket^{(3)}$, the only place where $\Sigma_{p,q}$ and $\Sigma_{p,r}$ may intersect is a type I tile named $\{p, q, r\}$ and, conversely, the three separatrices appearing in a type I tile pairwise cross one another. Similarly, for each $\{\{p, q\}, \{p', q'\}\}$ in $\llbracket 1, n \rrbracket^{(2,2)}$, the only place where $\Sigma_{p,q}$ and $\Sigma_{p',q'}$ may intersect is a type II tile named $\{\{p, q\}, \{p', q'\}\}$, and, conversely, the two separatrices do cross in such a type II tile. Summarizing, the number of times two separatrices $\Sigma_{p,q}$ and $\Sigma_{p',q'}$ cross is exactly the number of type I tiles named $\{p, q, p', q'\}$ if the set $\{p, q, p', q'\}$ has three elements, and the number of type II tiles named $\{\{p, q\}, \{p', q'\}\}$ if it has four elements. Hence all tiles in \mathcal{K} have pairwise different names if and only if any two separatrices cross at most once in \mathcal{K} , and we apply Proposition 1.11. \square

2. SUBWORD REVERSING

Up to now, we considered the combinatorial distance between two expressions u, v representing some permutation, *i.e.*, the minimal number of braid relations needed to transform u into v . We now address a related, but different question, namely the number of braid relations needed to transform u into v when one uses the specific strategy called *subword reversing* [3, 4]. The latter is known to solve the word problem, *i.e.*, to provide a step-by-step transformation of u into v by means of braid relations when this is possible, *i.e.*, when u and v are equivalent reduced expressions.

Investigating the optimality of the reversing method, *i.e.*, comparing the number of braid relations used by subword reversing and the combinatorial distance, is a natural question. Little is known at the moment. Here, we shall establish several partial results. We observe in particular that the reversing method need not be optimal but, using the geometric criteria of Sections 1.3 and 1.4, we characterize cases when reversing turns out to be optimal. Also, in a slightly different context of non-equivalent initial expressions, we shall prove that there exist length ℓ expressions u, v whose reversing require using $O(\ell^4)$ braid relations, which is not redundant with Proposition 1.

2.1. Subword reversing and reversing diagrams. We recall that an n -expression is a word on the alphabet $\{s_1, \dots, s_{n-1}\}$. Hereafter, we introduce a second alphabet $\{\bar{s}_1, \dots, \bar{s}_{n-1}\}$ which is a formal copy of the previous one. A word over the extended alphabet $\{s_1, \dots, s_{n-1}, \bar{s}_1, \dots, \bar{s}_{n-1}\}$ will be called an *extended n -expression*.

For each extended expression u , we denote by \bar{u} the extended expression obtained from u by reversing the order of the letters and exchanging s_i and \bar{s}_i everywhere. So, for instance, we have $\overline{s_1 s_2} = \bar{s}_2 \bar{s}_1$.

We now introduce a binary relation \curvearrowright (or rewrite rule) on extended expressions.

Definition. If w, w' are extended expressions, we declare that $w \curvearrowright w'$ holds if w' is obtained from w by replacing some subword $\bar{s}_i s_j$ either by $s_j s_i \bar{s}_j \bar{s}_i$ (if $|i - j| = 1$ holds), or by $s_j \bar{s}_i$ (if $|i - j| \geq 2$ holds), or by the empty word ϵ (if $j = i$ holds). A finite sequence (w_0, \dots, w_N) is called a *reversing sequence* if $w_k \curvearrowright w_{k+1}$ holds for each k . Finally, we say that w *reverses* to w' if there exists a reversing sequence (w_0, \dots, w_N) satisfying $w_0 = w$ and $w_N = w'$.

The principle of reversing is to push the letters \bar{s}_i to the right, and the letters s_j to the left, until no subword $\bar{s}_i s_j$ remains. For instance, $\bar{s}_1 \bar{s}_2 \bar{s}_1 s_2 s_1 s_2$ reverses to the empty word, as we have

$$(*) \quad \bar{s}_1 \bar{s}_2 \bar{s}_1 s_2 s_1 s_2 \curvearrowright \bar{s}_1 \bar{s}_2 s_2 s_1 \bar{s}_2 \bar{s}_1 s_1 s_2 \curvearrowright \bar{s}_1 s_1 \bar{s}_2 \bar{s}_1 s_1 s_2 \curvearrowright \bar{s}_2 \bar{s}_1 s_1 s_2 \curvearrowright \bar{s}_2 s_2 \curvearrowright \epsilon,$$

where, at each step, the subword that will be reversed has been written in bold.

As will become clear below, a reversing sequence from $\bar{u}v$ to the empty word provides a distinguished way of transforming u into v by means of braid relations and, therefore, there exists an associated van Kampen diagram. Now, it follows from the particular definition of subword reversing that the associated van Kampen diagrams have specific properties, namely they can essentially be drawn on a rectangular grid, a specific point that will be important in the sequel.

These diagrams that are essentially van Kampen diagrams will be called *reversing diagrams*. As a van Kampen diagram, a reversing diagram consists of edges labeled by letters s_i . The specific point is that, in a reversing diagram, all edges are either horizontal right-oriented edges and vertical down-oriented edges, and that, in addition to the latter, there may exist ϵ -labeled arcs. Assume that (w_0, \dots, w_N) is a reversing sequence, hence a sequence of extended expressions containing both types of letters s_i and \bar{s}_j . First, we draw a connected path indexed by the successive letters of w_0 by attaching a horizontal arrow $\xrightarrow{s_i}$ with each letter s_i , and a vertical arrow \downarrow_{s_i} with each letter \bar{s}_i . Then, we inductively complete the diagram as follows. Assume that one goes from w_{k-1} to w_k by reversing some subword $\bar{s}_i s_j$. By induction hypothesis, the latter subword $\bar{s}_i s_j$ corresponds to an open pattern $s_i \downarrow \xrightarrow{s_j}$ in the diagram. Then we complete that pattern with new arrows, according to the rule

$$\begin{array}{c} s_i \downarrow \xrightarrow{s_j} \\ \rightarrow \quad s_i \downarrow \begin{array}{c} \xrightarrow{s_j} \\ \downarrow_{s_i} \end{array} s_j \quad \text{for } |i-j|=1, \quad s_i \downarrow \xrightarrow{s_j} s_i \quad \text{for } |i-j| \geq 2, \quad s_i \downarrow \xrightarrow{s_j} \epsilon \quad \text{for } i=j, \\ \text{(type I)} \quad \text{(type II)} \quad \text{(type III)} \end{array}$$

with the convention that ϵ -labeled dotted arcs, dotted arcs in the diagrams, are subsequently ignored. For instance, the reversing diagram associated with the above reversing sequence $(*)$ is displayed in Figure 3.

Remark 2.1. Because of the ϵ -labelled arcs, the above patterns are not the most general ones appearing in a reversing diagram. The general patterns are actually

$$\begin{array}{c} s_i \downarrow \xrightarrow{s_j} \\ \rightarrow \quad s_i \downarrow \begin{array}{c} \xrightarrow{s_j} \\ \downarrow_{s_i} \end{array} s_j \quad \text{for } |i-j|=1, \quad s_i \downarrow \xrightarrow{s_j} s_i \quad \text{for } |i-j| \geq 2, \quad s_i \downarrow \xrightarrow{s_j} \epsilon \quad \text{for } i=j. \\ \text{(type I)} \quad \text{(type II)} \quad \text{(type III)} \end{array}$$




FIGURE 3. Reversing diagram associated with the sequence $(*)$, i.e., the reversing diagram of the pair $(s_1s_2s_1, s_2s_1s_2)$. Metric aspects are ignored, so the reversing diagram is always considered up to a piecewise affine deformation.

In this way, we associate with every reversing sequence a reversing diagram. Conversely, it is easy to see that, starting with a diagram as above, we can recover a (not necessarily unique) reversing sequence by reading the labels of the various paths going from the bottom-left corner to the top-right corner, and using the convention that a vertical s_i -labeled edge contributes \bar{s}_i .

By construction, subword reversing may be applied to any initial word consisting of letters s_i and \bar{s}_i , and not only to words of the form $\bar{u}v$ where u and v are equivalent reduced expressions. So there is no ambiguity in the following notion.

Definition. If u, v are expressions, the *reversing diagram for (u, v)* is the reversing diagram starting with the word $\bar{u}v$.

By construction, the reversing diagram for the pair (u, v) starts with a vertical down-oriented path labeled u and a horizontal right-oriented path labeled v starting from a common source. For instance, the diagram of Figure 3 is the reversing diagram for the pair $(s_1s_2s_1, s_2s_1s_2)$.

By construction, a reversing diagram becomes a van Kampen diagram when all ϵ -arcs are collapsed—but we do *not* do it, and keep the diagrams as they are now, insisting that they are drawn in a rectangular grid. A reversing diagram contains three types of tiles:

- type I tiles, which are hexagons, and correspond to type I braid relations,
- type II tiles, which are squares, and correspond to type II braid relations,
- type III tiles, which are digons, and correspond to free group relations $\bar{s}_i s_i = 1$.

The latter will be called *trivial*.

The connection between subword reversing and the problem of recognizing equivalent reduced expressions of a permutation is given by the following result.

Proposition 2.2. [3] (i) *For all reduced expressions u, v , there exists a unique pair of reduced expressions u', v' such that $\bar{u}v$ reverses to $v'\bar{u}'$.*

(ii) *Two reduced expressions u, v represent the same permutation if and only if $\bar{u}v$ reverses to the empty word.*

In the case of (ii), one implication is clear: by construction, a reversing sequence from $\bar{u}v$ to the empty word gives a reversing diagram for (u, v) that concludes with ϵ -arcs everywhere on the bottom and the right, hence, after collapsing the ϵ -edges, it gives a van Kampen diagram for (u, v) , thus proving that u and v are equivalent. The converse implication is not obvious, as not every van Kampen diagram comes from a reversing diagram. The specific point is that, in a reversing diagram, two edges at most start from any vertex, a property that fails in the top diagram of Figure 2: so that diagram is certainly not associated with a reversing. The proof of

Proposition 2.2—which actually extends to arbitrary braids—relies on the so-called Garside theory of braids [7, 6].

In terms of reversing diagrams, the situation is as follows. In the particular case when u and v are equivalent reduced expressions, the reversing diagram for (u, v) finishes for ϵ -labeled arcs everywhere, and collapsing all these ϵ -labeled arcs yields a van Kampen diagram for (u, v) . In general, if u, v are arbitrary expressions, then the reversing diagram for (u, v) is still finite, and it finishes with arrows forming a word of the form $v'\overline{u'}$ where u', v' are two expressions that need not be empty. Then collapsing all the ϵ -labeled arcs yields a van Kampen diagram for the pair (uv', vu') . It can be shown that, if u and v are reduced, then uv' and vu' are reduced and equivalent—the latter point is obvious as, by construction, the reversing diagram for (u, v) provides a van Kampen diagram for (uv', vu') —and the permutation represented by uv' and vu' is the least upper bound of the permutations represented by u and v with respect to the weak order of \mathfrak{S}_n [1].

In every case, the number of nontrivial tiles in the reversing diagram or, equivalently, the number of nontrivial steps in an associated reversing sequence, is well defined.

Definition. Assume that u, v are reduced expressions. The *reversing complexity* of (u, v) , denoted $\text{compl}_\curvearrowright(u, v)$, is the number of nontrivial tiles in the reversing diagram for (u, v) .

Equivalently, the reversing complexity $\text{compl}_\curvearrowright(u, v)$ is the number of nontrivial steps in a reversing sequence from \overline{uv} to a word of the form $v'\overline{u'}$. By the above remarks, a reversing diagram for (u, v) with N nontrivial tiles provides a van Kampen diagram for (uv', vu') with N tiles, where $v'\overline{u'}$ is the final word of the reversing process. So we always have

$$(2.1) \quad \text{dist}(uv', vu') \leq \text{compl}_\curvearrowright(u, v).$$

In particular, when we start with equivalent reduced expressions u, v , we have

$$(2.2) \quad \text{dist}(u, v) \leq \text{compl}_\curvearrowright(u, v),$$

since, in this case, the final expressions u' and v' are empty by Proposition 2.2(ii).

We shall now discuss the converse inequality, *i.e.*, the question of whether subword reversing, viewed as a particular strategy for finding derivations between equivalent expressions of a permutation, is efficient, or even possibly optimal.

Proposition 2.3. *Subword reversing is not always optimal: There exist equivalent reduced expressions u, v satisfying $\text{dist}(u, v) < \text{compl}_\curvearrowright(u, v)$.*

Proof. Consider the 4-expressions $u = s_1s_2s_1s_3s_2s_1$ and $v = s_3s_2s_3s_1s_2s_3$, two expressions of the flip permutation ϕ_4 of \mathfrak{S}_4 . Together with Proposition 1.11, the top diagram in Figure 2 gives $\text{dist}(u, v) = 6$. On the other hand, the reversing diagram of Figure 4 gives $\text{compl}_\curvearrowright(u, v) = 8$. \square

2.2. An optimality criterion. Despite the negative result of Proposition 2.3, experiments show that subword reversing is often an efficient strategy. What we do now is to establish sufficient criteria for recognizing that reversing is possibly optimal. Of course, we shall say that a reversing diagram \mathcal{D} is *optimal* if the van Kampen diagram obtained by collapsing the ϵ -labeled arcs in \mathcal{D} is optimal, *i.e.*, if it realizes the combinatorial distance between the boundary expressions.





FIGURE 4. Reversing diagram for $\bar{s}_1 \bar{s}_2 \bar{s}_3 \bar{s}_1 \bar{s}_2 \bar{s}_1 s_3 s_2 s_3 s_1 s_2 s_1$. There are eight nontrivial tiles (four type I hexagons and four type II squares), giving $\text{compl}_{\sim}(s_1 s_2 s_1 s_3 s_2 s_1, s_3 s_2 s_3 s_1 s_2 s_1) = 8$. Collapsing the ϵ -edges in the above diagram yields the bottom van Kampen diagram of Figure 2. The failure of optimality is witnessed by the two intersections of the separatrices $\Sigma_{1,4}$ and $\Sigma_{2,3}$.

Proposition 2.4. *A reversing diagram containing no digon, i.e., containing only tiles of type I and II, is optimal.*

Proof. Assume that \mathcal{D} is the reversing diagram for (u, v) . By definition, the separatrices of \mathcal{D} start from the edges corresponding to u , i.e., here, from the left. Then an induction on the number of tiles shows that only the following orientations may appear in the tiles of \mathcal{D} .



Hence every vertical edge of \mathcal{D} is crossed by a separatrix from left to right, and every horizontal edge is crossed from bottom to top. Moreover, we see that digons are the only tiles that can change the orientation of a separatrix from horizontal to vertical—whereas only hexagons can change the orientation from vertical to horizontal. Also we see that, if two separatrices cross in a tile, then, when entering that tile, at least one of them is vertical.


Now assume that two separatrices Σ, Σ' cross at least twice in \mathcal{D} . Two cases may occur, according to whether the names of these separatrices involve three or four integers. Assume first that there exist p, q, r satisfying $\Sigma = \Sigma_{p,q}$ and $\Sigma' = \Sigma_{p,r}$. Then Σ and Σ' can cross only in hexagons named $\{p, q, r\}$, in which case, putting $\Sigma'' = \Sigma_{q,r}$, the separatrices Σ and Σ'' , as well as Σ' and Σ'' , cross too. Let H_1 (resp. H_2) be the first (resp. second) hexagon where Σ, Σ' and Σ'' cross. Without loss of generality, we may assume that Σ is horizontal when entering H_1 (hence also when exiting it) and that Σ' is vertical when exiting H_1 (hence also when entering it). Then Σ' remains above Σ (measured from the bottom of the diagram)—with Σ''

lying in between—until they enter the hexagon H_2 . Therefore, the only possibility is that Σ' is horizontal when entering H_2 (from the left), whereas Σ is vertical when entering H_2 (from the bottom). This is possible only if a digon changes the orientation of Σ from horizontal to vertical between H_1 and H_2 .


Assume now that there exist p, q, p', q' satisfying $\Sigma = \Sigma_{p,q}$ and $\Sigma' = \Sigma_{p',q'}$. Then Σ and Σ' can cross only in squares named $\{\{p, q\}, \{p', q'\}\}$. Let S_1 (resp. S_2) be the first (resp. second) square where Σ and Σ' cross. Without loss of generality, we can assume that Σ is horizontal after (and before) S_1 , and Σ' is vertical. Then, as in the case of hexagons, Σ' remains above Σ until they enter the square S_2 . Therefore, Σ' is horizontal when entering S_2 (from the left), whereas Σ is vertical when entering S_2 (from the bottom). Hence a digon changes the orientation of Σ from horizontal to vertical between S_1 and S_2 .

So, in any case, two separatrices may cross twice only if there is a digon in \mathcal{D} . Then we apply Proposition 1.12. \square

The previous result can be improved by showing that some digons are harmless and can be ignored. Indeed, consider a pattern of the form $\bar{s}_i s_{i+1} s_i$. Then we have $\bar{s}_i s_{i+1} s_i \curvearrowright s_{i+1} s_i \bar{s}_{i+1} \bar{s}_i s_i \curvearrowright s_{i+1} s_i \bar{s}_{i+1}$, corresponding to an hexagon with an

appended digon  in the diagram. Let us introduce two new types

of hexagonal tiles, namely, for $|i - j| = 1$,




Using such tiles amounts to replacing two adjacent tiles with one unique tile of the new type, but they do not change anything in the rest of the diagram. In this way, we obtain a new type of reversing diagrams that we call *compacted*. Note that the compacted diagram associated with an initial expression $\bar{u}v$ need not be unique, as there may be several ways of grouping the tiles, see Figure 5. However, as the situation after a tile of type I' or I'' is exactly the same as the situation after the corresponding type I tile followed by a type III tile, the number of nontrivial tiles is the same in any diagram associated with a given initial pair (u, v) .

The expected improvement of Proposition 2.4 is

Proposition 2.5. *A compacted reversing diagram containing no digon, i.e., containing only tiles of types I, I', I'', and II, is optimal.*

Proof. The new tiles of type I' and I'' do not change the orientation of separatrices. Indeed, the corresponding possibilities are






FIGURE 5. Two slightly different ways of compacting the reversing diagram of Figure 4.

None of these tiles changes the horizontal-vertical orientation of the separatrices, and, therefore, their appearing in a reversing diagram does not affect the argument used in the proof of Proposition 2.4. \square

An application of the above criterion will be mentioned in Remark 2.11 below.

2.3. Upper bounds. Very little is known about the reversing complexity in general. In particular, the following conjecture, which is the natural counterpart of Proposition 1.4, remains open at the moment—see [2] for partial results.

Conjecture 2.6. *For all n -expressions u, v of length ℓ , the reversing complexity $\text{compl}_{\curvearrowright}(u, v)$ lies in $O(n^2\ell)$.*

Even the weaker result of the reversing complexity being polynomial is not known. By adapting the method used for [5, Prop. 3], one comes up with the weak result that, if u, v are length ℓ expressions, then $\text{dist}(u, v) \leq C \cdot 81^\ell$ holds for some constant C . Using a careful analysis of separatrices, one can obtain the following improvement.

Proposition 2.7. *If u, v are expressions of length ℓ , then we have $\text{dist}(u, v) \leq C \cdot 9^\ell$ for some constant C .*

As the argument is complicated and, at the same time, the result seems far from optimal, we skip the proof and refer to [2] for details.

Remark. The index n does not appear in Proposition 2.7. This reflects the fact that, although the maximal reversing complexity between two n -expressions of length ℓ increases with n and ℓ , it does not increase indefinitely: if we denote by $N(n, \ell)$ the maximal reversing complexity between two n -expressions of length ℓ , then, for each ℓ , the value of $N(n, \ell)$ is constant for $n \geq 2\ell$. This is due to the fact that, when n is too large with respect to ℓ , the indices of the transpositions s_i occurring in an n -expression cannot cover the whole of $\llbracket 1, n \rrbracket$ and commutation relations of type II occur. Here again, we refer to [2] for more details.

2.4. A lower bound for the reversing complexity. The n -expressions used in the proof of Proposition 1.7 to establish the inequality (1.9) have length $\ell = n(n-1)/2$ so that, in this way, we obtain for infinitely many values of ℓ equivalent reduced expressions of length ℓ that satisfy $\text{dist}(u, v) \geq \ell(\ell-1)/2$, hence, a fortiori,

$$(2.3) \quad \text{compl}_{\curvearrowright}(u, v) \geq \frac{\ell(\ell-1)}{2}.$$

Question 2.8. *Can one construct a sequence (u_ℓ, v_ℓ) of pairwise equivalent reduced expressions of length ℓ such that $\text{compl}_\curvearrowright(u_\ell, v_\ell)$ is more than quadratic in ℓ ?*

We leave Question 2.8 open, but we now address another related question and establish a result that illustrates how complicated the reversing process may be.

Proposition 2.9. *For each ℓ there exist expressions u, v of length ℓ satisfying*

$$\text{compl}_\curvearrowright(u, v) \geq \frac{4}{3}\ell^4$$

for ℓ large enough.

We begin with an auxiliary lemma. Hereafter we write $w \curvearrowright^k w'$ if there is a length k reversing sequence from w to w' , not counting trivial steps of type III.

Lemma 2.10. *For $i, p \geq 1$, put $a_{i,p} = s_{i+p-1}s_{i+p-2}\dots s_i$, $b_{i,p} = s_i s_{i+1} \dots s_{i+p-1}$, $c_{i,p} = a_{i,p}a_{i+1,p}$, and $d_{i,p} = b_{i+1,p}b_{i,p}$. Then, for all i, p , we have*

$$(2.4) \quad \overline{b_{i,p}} a_{i+1,p} \curvearrowright^{N_p} a_{i,p+1} \overline{b_{i,p+1}},$$

$$(2.5) \quad \overline{d_{i,p}} c_{i+2,p} \curvearrowright^{N'_p} c_{i,p+2} \overline{d_{i,p+2}},$$

with $N_p = p^2 + p - 1$ and $N'_p = 4p^2 + 8p - 3$.




FIGURE 6. Proof of Relation 2.5, here with $i = 1$ and $p = 3$; the grey rectangles correspond to Relation 2.4.

Proof. For (2.4) we use induction on p . The case $p = 1$ is

$$\overline{b_{i,1}} b_{i+1,1} \curvearrowright^1 a_{i,2} \overline{b_{i,2}},$$

a restatement of $\overline{s_i s_{i+1}} \curvearrowright^1 s_{i+1} s_i \overline{s_{i+1} s_i}$. Assume $p \geq 2$. Applying the induction hypothesis once, plus one reversing step of type I and one step of type III—or one

type I' step instead—and $2p - 1$ steps of type II, we obtain

$$\begin{aligned}
\overline{b_{i,p}} a_{i+1,p} &= \overline{s_{i+p-1}} \overline{\mathbf{b_{i,p-1} s_{i+p}}} a_{i+1,p-1} \\
&\curvearrowright^{p-1} \overline{s_{i+p-1}} s_{i+p} \overline{\mathbf{b_{i,p-1} a_{i+1,p-1}}} \\
&\curvearrowright^{N_{p-1}} \overline{\mathbf{s_{i+p-1} s_{i+p}}} a_{i,p} \overline{b_{i,p}} \\
&\curvearrowright^1 s_{i+p} s_{i+p+1} \overline{s_{i+p}} \overline{\mathbf{s_{i+p+1} a_{i+1,p}}} \overline{b_{i,p}} \\
&\curvearrowright^0 s_{i+p} s_{i+p+1} \overline{\mathbf{s_{i+p} a_{i+1,p-1}}} \overline{b_{i,p}} \\
&\curvearrowright^p s_{i+p} s_{i+p+1} a_{i+1,p-1} \overline{s_{i+p}} \overline{b_{i,p}} = a_{i,p+1} \overline{b_{i,p+1}},
\end{aligned}$$

where, in each case, the factors that are about to be reversed are marked in bold. We deduce $N_p = N_{p-1} + 2p = p^2 + p - 1$.

The computation for (2.5) is illustrated in Figure 6. Using (2.4) four times, plus $4p + 1$ type II steps, we obtain:

$$\begin{aligned}
\overline{d_{i,p}} c_{i+2,p} &= \overline{b_{i,p}} \overline{\mathbf{b_{i+1,p} a_{i+2,p}}} a_{i+3,p} \\
&\curvearrowright^{N_p} \overline{b_{i,p}} a_{i+1,p+1} \overline{b_{i+1,p+1}} a_{i+3,p} \\
&= \overline{\mathbf{b_{i,p} s_{i+p+1}}} a_{i+1,p} \overline{b_{i+1,p+1}} a_{i+3,p} \\
&\curvearrowright^p s_{i+p+1} \overline{\mathbf{b_{i,p} a_{i+1,p}}} \overline{b_{i+1,p+1}} a_{i+3,p} \\
&\curvearrowright^{N_p} s_{i+p+1} a_{i,p+1} \overline{b_{i,p+1}} \overline{b_{i+1,p+1}} a_{i+3,p} \\
&= a_{i,p+2} \overline{b_{i,p+1}} \overline{b_{i+2,p}} \overline{\mathbf{s_{i+1} a_{i+3,p}}} \\
&\curvearrowright^p a_{i,p+2} \overline{b_{i,p+1}} \overline{\mathbf{b_{i+2,p} a_{i+3,p}}} \overline{s_{i+1}} \\
&\curvearrowright^{N_p} a_{i,p+2} \overline{b_{i,p+1}} a_{i+3,p+1} \overline{b_{i+3,p+1}} \overline{s_{i+1}} \\
&= a_{i,p+2} \overline{b_{i+1,p}} \overline{\mathbf{s_i s_{i+p+2} a_{i+2,p}}} \overline{b_{i+2,p+2}} \\
&\curvearrowright^{p+1} a_{i,p+2} \overline{\mathbf{b_{i+1,p} s_{i+p+2}}} a_{i+2,p} \overline{s_i} \overline{b_{i+1,p+1}} \\
&\curvearrowright^p a_{i,p+2} s_{i+p+2} \overline{\mathbf{b_{i+1,p} a_{i+2,p}}} \overline{s_i} \overline{b_{i+1,p+1}} \\
&\curvearrowright^{N_p} a_{i,p+2} s_{i+p+2} a_{i+1,p+1} \overline{b_{i+1,p+1}} \overline{s_i} \overline{b_{i+1,p+1}} = c_{i,p+2} \overline{d_{i,p+2}},
\end{aligned}$$

leading to $N'_p = 4p^2 + 8p - 3$. \square

We can now establish Proposition 2.9.

Proof of Proposition 2.9. (See Figure 7.) We put

$$u_\ell = s_{2\ell} s_{2\ell-2} \dots s_2 \quad \text{and} \quad v_\ell = s_1 s_3 \dots s_{2\ell-1},$$

and analyze the reversing of $\overline{u_\ell v_\ell}$. The latter consists of three sequences of elementary steps. First, $\ell(\ell - 2)/2$ steps of type II lead to

$$\overline{s_2} s_1 \overline{s_4} s_3 \dots \overline{s_{2\ell}} s_{2\ell-1}.$$

Then, ℓ type I steps lead to $s_1 s_2 \overline{s_1} \overline{s_2} s_3 s_4 \overline{s_3} \overline{s_4} \dots s_{2\ell-1} s_{2\ell} \overline{s_{2\ell-1}} \overline{s_{2\ell}}$, which is

$$c_{1,1} \overline{d_{2,1}} c_{3,1} \overline{d_{4,1}} \dots c_{2\ell-1,1} \overline{d_{2\ell,1}}.$$

From there, we apply (2.5) repeatedly: after $\ell - 1$ applications, we obtain

$$c_{1,1} \cdot c_{1,3} \overline{d_{2,3}} c_{3,3} \overline{d_{4,3}} \dots c_{2\ell-3,3} \overline{d_{2\ell-2,3}} \cdot \overline{d_{2\ell,1}};$$

after $\ell - 2$ more applications, we obtain

$$c_{1,1} c_{1,3} \cdot c_{1,5} \overline{d_{2,5}} c_{3,5} \overline{d_{4,5}} \dots c_{2\ell-5,5} \overline{d_{2\ell-4,5}} \cdot \overline{d_{2\ell-2,3}} \overline{d_{2\ell,1}},$$

and so on. After using (2.5) $\ell(\ell - 1)/2$ times, we finally obtain

$$c_{1,1} c_{1,3} \dots c_{1,2\ell+1} \overline{d_{2,2\ell-1}} \overline{d_{4,2\ell-3}} \dots \overline{d_{2\ell-2,3}} \overline{d_{2\ell,1}}.$$

A careful bookkeeping shows that the total number of reversing steps involved in the process is $(8\ell^4 - 23\ell^2 + 9\ell + 12)/6$, hence $\Theta(\ell^4)$ as announced. \square




FIGURE 7. Proof of Proposition 2.9, here for $\ell = 4$; each grey rectangle corresponds to applying Relation (2.5), hence contains a number of elementary tiles that lies in $O(\ell^2)$.

Remark 2.11. At the expense of using one type I' tile for the proof of Relation 2.4, no type III tile is used throughout the above constructions. Using Proposition 2.5, we conclude that the reversing diagram we obtained gives an optimal van Kampen diagram, *i.e.*, it realizes the combinatorial distance between the boundary words, here $u_\ell c_{1,1} c_{1,3} \dots c_{1,2\ell+1}$ and $v_\ell d_{2\ell,1} d_{2\ell-2,3} \dots, d_{4,2\ell-3} d_{2,2\ell-1}$.

With Proposition 2.9, we prove that $\text{compl}_\sim(u, v)$ can be quartic in the length of u and v . We conjecture this lower bound is also an upper bound, but have no proof of this result so far. The problem is that we have no control on the number of hexagons and digons that may occur in a reversing diagram. There is a quadratic upper bound on the lengths of the final expressions u', v' that may arise from some initial pair (u, v) of length ℓ expressions, but this does not directly lead to a bound on the number of type I reversing steps used (the only ones that increase the length) because some subsequent type III steps might erase the letters so created.

Remark 2.12. As mentioned in the introduction, most results of this paper extend to positive braids. For instance, the optimality criterion of Proposition 1.12 extends to positive braids at the expense of adding a notion of rank in the definition of separatrices: in the braid diagram associated with a permutation, *i.e.*, with a simple braid, any two strands cross at most once, and we introduce one separatrix $\Sigma_{p,q}$ only. For the case of arbitrary positive braids, we should introduce several separatrices for pairs of strands that cross more than one time, $\Sigma_{p,q}^{(k)}$ being associated with the k th intersection of the strands p and q . As for subword reversing, it works in the general braid case exactly as in the case of simple braids, *i.e.*, of permutations. Experiments show that the worst cases in terms of complexity arise with simple braids. So we have no better result in the general braid case than in the particular permutation case.

REFERENCES

- [1] A. Björner and F. Brenti, *Combinatorics of Coxeter groups*, Graduate Texts in Math., vol. 231, Springer; 2005.
- [2] M. Autord, *Aspects algorithmiques du retournement de mot*, PhD Thesis, Université de Caen; 2009.
- [3] P. Dehornoy, *Groups with a complemented presentation*, J. Pure Appl. Algebra **116** (1997) 115–137.
- [4] P. Dehornoy, *On completeness of word reversing*, Discrete Math. **225** (2000) 93–119.
- [5] P. Dehornoy & B. Wiest, *On word reversing in braid groups*, Int. J. Algebra Comput. **16**(5) (2006) 931–947.
- [6] D. Epstein, with J. Cannon, D. Holt, S. Levy, M. Paterson & W. Thurston, *Word Processing in Groups*, Jones & Bartlett Publ. (1992).
- [7] F.A. Garside, *The braid group and other groups*, Quart. J. Math. Oxford **20-78** (1969) 235–254.
- [8] J.E. Humphreys, *Reflection Groups and Coxeter Groups*, Cambridge Univ. Texts; 1989.
- [9] R.C. Lyndon and P.E. Schupp, *Combinatorial Group Theory*, Springer-Verlag; 1977, reprinted in 2001.

M.A., LABORATOIRE DE MATHÉMATIQUES NICOLAS ORESME, UNIVERSITÉ DE CAEN, 14032 CAEN, FRANCE

E-mail address: `marc.autord@math.unicaen.fr`

P.D. (corresponding author), LABORATOIRE DE MATHÉMATIQUES NICOLAS ORESME, UNIVERSITÉ DE CAEN, 14032 CAEN, FRANCE

E-mail address: `dehornoy@math.unicaen.fr`

URL: `//www.math.unicaen.fr/~dehornoy`