Research report 130

THE SET OF MINIMAL BRAIDS IS CO-NP-COMPLETE

M S Paterson[1] & A A Razborov[2]

(RR130)

Braids can be represented as two-dimensional diagrams showing the crossings of strings or as words over the generators of a braid group. A minimal braid is one with the fewest crossings (or the shortest words) among all possible representations topologically equivalent to that braid. The main result of this paper is that the set of minimal braids is co-NP-complete.

1 Department of Computer Science
University of Warwick
Coventry CV4 7AL
United Kingdom

2 Steklov Mathematical Institute
Moscow
USSR

July 1988

The set of minimal braids is co-NP-complete[0]

by

M.S. Paterson[1]

and

A.A. Razborov[2]

Abstract

Braids can be represented as two-dimensional diagrams showing the crossings of strings or as words over the generators of a braid group. A minimal braid is one with the fewest crossings (or the shortest words) among all possible representations topologically equivalent to that braid. The main result of this paper is that the set of minimal braids is co-NP-complete.

1. Introduction

Algorithmic problems in braid groups have received much attention since [A1]. An algorithm for the word problem was given by Artin [A1][A2], and Garside solved the conjugacy problem [G]. More recently, with the increasing interest in complexity, these problems have been re-examined with regard to efficiency. Artin's algorithm involves generating a canonical form of length exponential in the length of the original word. Apparently the first polynomial-time algorithm for the word problem results from recent work of Thurston [Th]. Whether there exists a polynomial-time algorithm for the conjugacy problem seems to be unknown .

In his polynomial-time algorithm for the word problem Thurston produces a canonical form for braid group elements whose length is quadratic in the length of the original word. Neither his form nor Artin's one is a minimal word representing the given braid. This situation is not very common in group theory; e.g., for free groups, HNN-extensions, free products and so on, the known normal forms are minimal when the generators are chosen in a natural way.

The present paper provides some complexity intuition as to why such a form with 'nice' properties cannot exist for braid groups (unless P≠NP). We show that the set of minimal braids is co-NP-complete. This implies that (again, unless P≠NP) there is no polynomial algorithm to produce a minimal representation of a given braid.

---

[0] This work was carried out during a visit to Warwick University by the second author, supported by a Visiting Fellowship from the SERC of the UK. The first author is supported by a Senior Fellowship from the SERC.

[1] Dept. of Computer Science, Univ. of Warwick, Coventry, UK.

[2] Steklov Mathematical Institute, Moscow, USSR.

There are two quite different approaches to braid groups, geometric and algebraic (compare [A] and [M]); each has some advantages and disadvantages. In this paper we follow an intermediate course. We say as precisely as possible *what* should be calculated and *how*, but the calculations themselves are omitted whenever the result is clear from geometric intuition.

2. Definitions and Main Theorem

Throughout, letters q,r,..., z stand for words over an alphabet, letters a,b,c,... for symbols from this alphabet. |z| is the length of z, and Λ is the empty word. '=' stands for graphical equality (i.e. as words), '$\equiv$' for the equivalence of two words representing the same element in a group. S_n is the symmetric group on n symbols. The notation u∗v∗w is used to denote the corresponding occurrence of the word v in the context uvw.

The group B_n of braids with n strings has the following representation:

$$B_n = \langle \sigma_1,...,\sigma_{n-1} \mid \sigma_i\sigma_j = \sigma_j\sigma_i \text{ for } j > i+1,\ \sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}, \text{ for } 1 \le i \le n-2 \rangle. \quad (1)$$

The σ_i's are called the *standard generators* of B_n. A geometric picture of σ_i is given by:

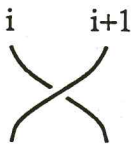

Fig. 1

We shall sometimes refer to standard generators as *positive crossings* and to their inverses as *negative crossings*. In the following development there is a special rôle for the initially leftmost string of the braids, which we shall call the *weft*. The other strings will be called *wires*. The problem we consider is presented in the style of Garey and Johnson [GJ].

Definition NON-MINIMAL BRAIDS ::

Instance: A braid group B, and a word w, in the standard generators of B.

Question: Is there a shorter word w' equivalent to w in B?

Main Theorem NON-MINIMAL BRAIDS is NP-complete.

Proof NON-MINIMAL BRAIDS is clearly in NP since the word problem for braids is solvable in polynomial time [Th?].

To show that this set is NP-hard we will use instances representing the following set of braids $\mathcal{F} = \bigcup_{r,m>1} \mathcal{F}_{r,m}$. For each r,m, $\mathcal{F}_{r,m} \subseteq B_{1+cm}$, where c is a parameter dependent on r and m to be chosen just below. The strings of B_{1+cm} are partitioned into the weft (the leftmost string), and m consecutive blocks of wires, to be called *cables*, consisting of c wires each. Each cable is labelled with a symbol from the alphabet $\Sigma = \{1,...,r\}$. It will be convenient to refer to cables labelled with i as *i-cables*, and to their wires as *i-wires*. The weft traverses the cables in s identical stages, each having r *levels* numbered sequentially from 1 to r, and each level consisting of t *loops*. We will choose these parameters as $s = 8m^2$, $c = rms$, and $t = m^6c^6$.

For each loop in level j, the weft starts on the left and travels all the way to the right, passing under i–cables where $i \leq j$, and over i–cables where $i < j$, then it returns passing under any i–cable where $i < j$, and over any i–cable where $i \geq j$, thus enfolding precisely the j–cables.

It is clear how to write the corresponding word in the standard generators; we denote this by x(q), where $q = a_1...a_m \in \Sigma^m$ and a_i is the label of the i'th cable.

The *special word* $w_0 = w_0(q)$, such that $w_0(q) \equiv x(q)$, describes the following wiring layout:

(i) Each 1–cable is passed over the other cables and accumulated in a block on the left, then the 2–cables are passed over the remaining cables and accumulated in a block to the right of the 1–block, and so on. In this process, no cables with the same labels ever cross each other and the cables are sorted into numerical order by label using right-over-left transpositions of cables. The weft remains on the extreme left.

(ii) The weft is brought under the 1–block, then wrapped around the whole block as a coil with t–1 turns leaving the block on the right. The weft continues, making a similar coil around the 2–block, and so on for each block in turn, finally returning over all the cables. This whole sequence is repeated s times.

(iii) The cable crossings of part (i) are now reversed, using left-over-right transpositions, to restore the original ordering of the cables.

If K transpositions of cables are needed to sort the cables in part (i) then, since a cable crossing uses c^2 wire crossings, we have:

$$|w_0| = 2Kc^2 + 2tmcs.$$

Suppose the minimal number of cable crossings to arrange the cables into n labelled blocks corresponding to some ordering of Σ is K', and consider the word w' corresponding to the following layout.

(i) Arrange the cables into blocks using $K'c^2$ wire crossings. The procedure here is similar to that in the standard word, except that the blocks are ordered according to the given ordering of Σ, and in each crossing the wire with the label which is earlier in this ordering is taken over that with the later label.

(ii) Visit the blocks in numerical order to make the coils, finishing at the left. This requires at most $(r+1+2t)mcs$ crossings.

(iii) Restore the original order of the cables by reversing the crossings used in (i).

If $K' < K$, then since $c = rms$, we have:

$$|w'| \leq 2K'c^2 + (r+1+2t)mcs < 2Kc^2 + 2tmcs$$

and the special word is not minimal.

If $K' = K$ then $\text{length}(w') \geq 2Kc^2 + 2tmcs = \text{length}(w_0)$ and the special word is no longer than w'.

Figure 2 shows $w_0(213123213)$, except that we have used $s = 2$ for clarity. This word is not minimal since it is better to arrange the cables in the pattern '222111333' for coiling.

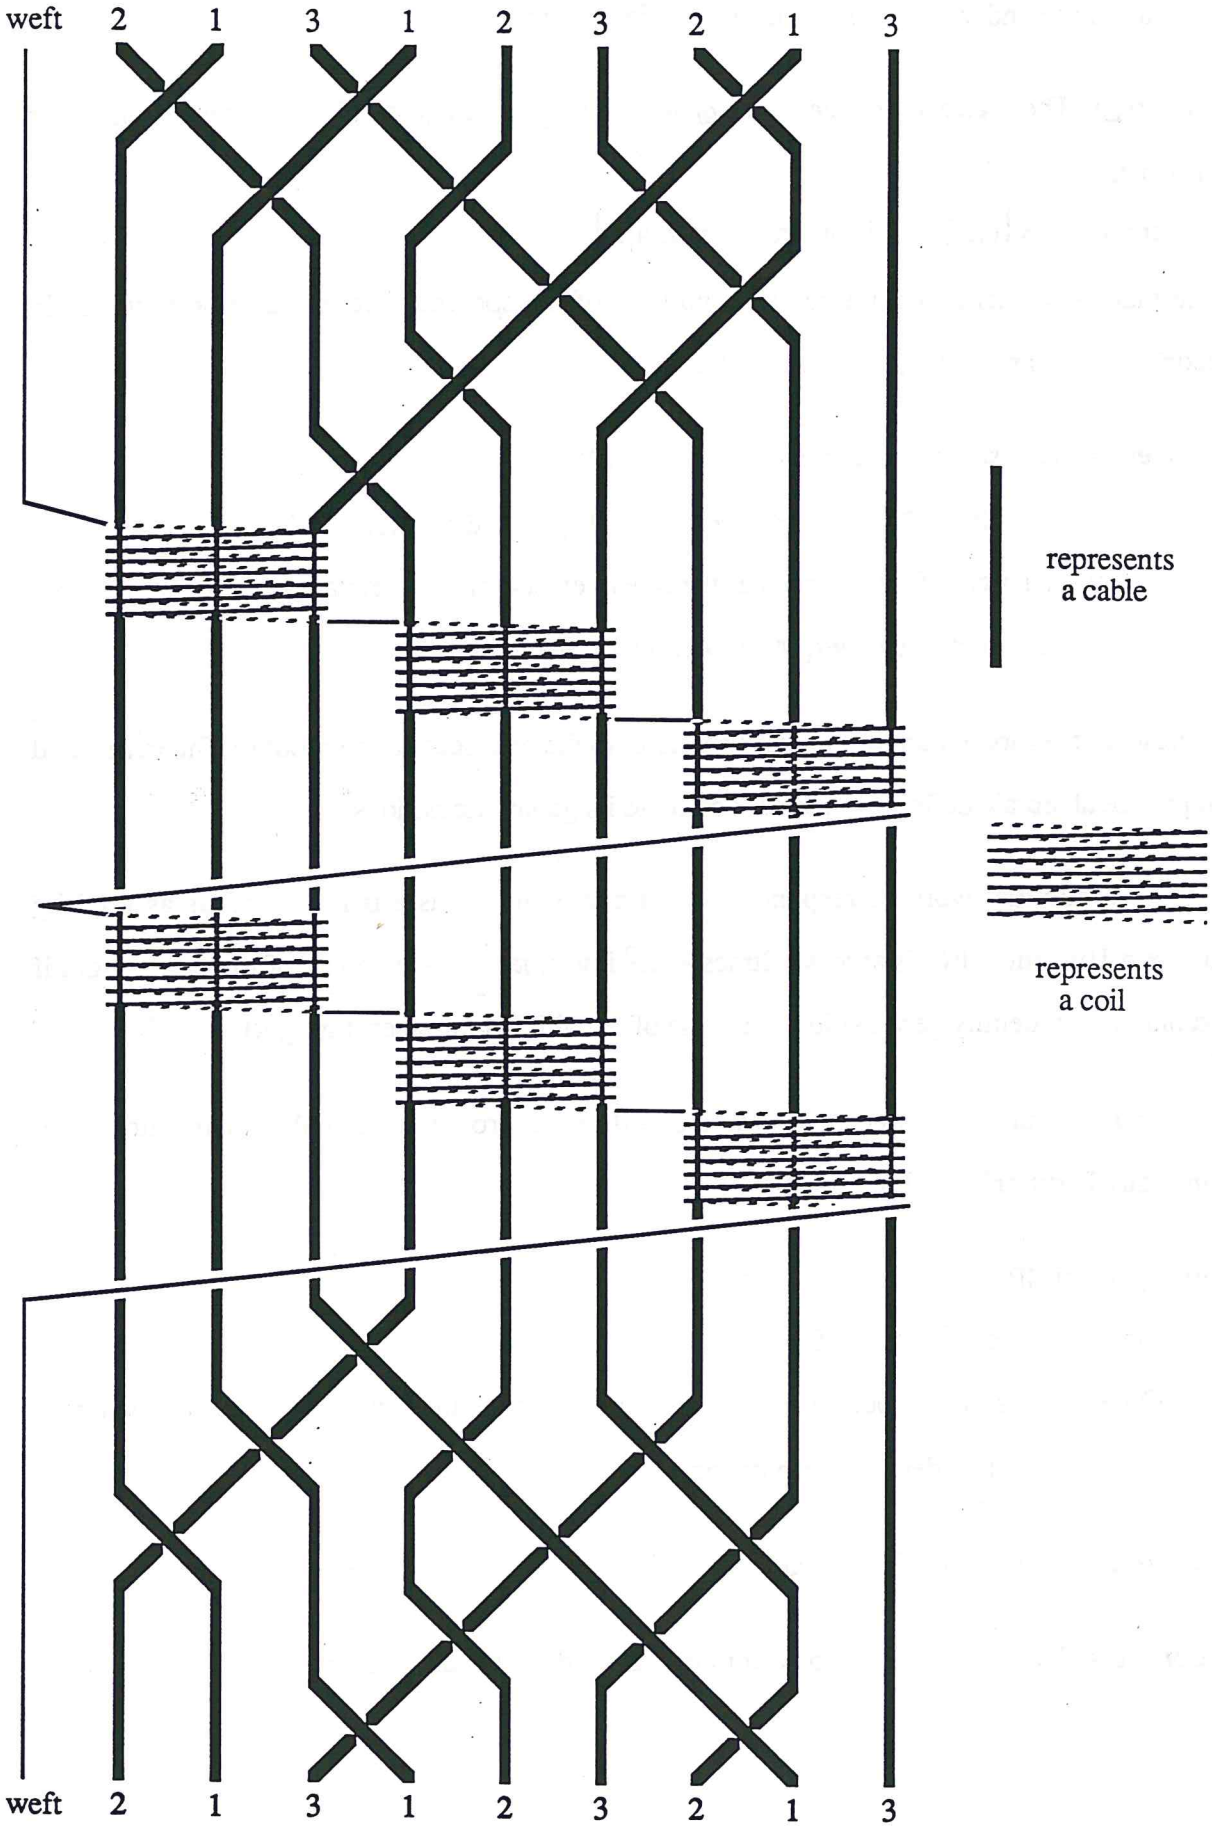
weft 2 1 3 1 2 3 2 1 3
represents a cable
represents a coil
weft 2 1 3 1 2 3 2 1 3

Fig. 2

We want to see under what conditions the special word is minimal.

Definition The *number of inversions of a string* $q = a_1...a_m \in \Sigma^m$ *with respect to* π is defined as:

$$inv(q, \pi) = |\{\langle i,j\rangle \mid i < j \text{ and } \pi(a_i) > \pi(a_j)\}|.$$

Note that $inv(q, \pi)$ is just the minimal number of transpositions required to permute q in accord with π, and $inv(q, \pi) \leq m(m-1)/2$.

Theorem 1 Let w be a word such that $w \equiv x(q)$. Then:

(i) w has at least $2tmcs$ positive crossings between the weft and wires;

(ii) if the length of w is minimal then there exists some permutation π of Σ such that w has at least $2c^2.inv(q, \pi)$ crossings between the wires.

To preserve the momentum of the proof we defer to the next section the proof of Theorem 1 and the precise algebraic definitions corresponding to its geometric notions.

If w is a minimal word corresponding to $q \in \Sigma^m$ and π is a permutation as assured by Theorem 1(ii), then $|w|$ is at least $2tmcs + 2c^2.inv(q, \pi)$. Hence w_0 is of minimal length if and only if the identity permutation is a value of π which minimizes $inv(q, \pi)$.

To complete our proof we show that the following problem, SNMP (Sorting does Not Minimally Partition), is NP-complete.

Definition SNMP ::

Instance: $q \in \Sigma^*$ where $\Sigma = \{1,...,r\}$.

Question: Is there a permutation π of $[1, ...,r]$ such that $inv(q,\pi) < inv(q,\iota)$, where ι is the identity permutation.

Theorem 2 SNMP is NP-complete.

Proof It is clearly in NP. We show it to be NP-hard by a chain of reductions.

A very similar problem, GROUPING BY SWAPPING, was shown NP-complete by T.D. Howell (unpublished manuscript, 1977, referred to in [GJ]), but we require here a stronger result.

It is a straightforward strengthening of Cook's Theorem that 3SAT remains NP-complete even when the input formula is always accompanied by some assignment which satisfies all but one of the clauses. This follows from the observation that any nondeterministic polynomial-time Turing machine can be modified to accept the same set in the same time but also to have a standard terminating rejection computation. Under suitable conventions the latter computation transforms to an assignment satisfying all but one of the clauses.

Our next step is to show the NP-hardness of the following set.

NON-MINIMAL FEEDBACK ARC SET ::

Instance: A directed graph G, and a subset S of arcs of G such that each circuit in G contains some arc of S, i.e. S is a *feedback arc set.*

Question: Is there a feedback arc set S' with $|S'| < |S|$?

Lemma 1 NON-MINIMAL FEEDBACK ARC SET is NP-complete.

Proof For any instance, F, of 3SAT with v variables and c clauses we define a corresponding directed graph, G_F. G_F has 4vc vertices, $a_{u,i}$, $b_{u,i}$, $A_{u,i}$ and $B_{u,i}$ for $1 \le u \le v$, $1 \le i \le c$. There are arcs $\langle a_{u,i}, b_{u,i}\rangle$, $\langle A_{u,i}, B_{u,i}\rangle$, $\langle b_{u,i}, A_{u,j}\rangle$, $\langle B_{u,i}, a_{u,j}\rangle$ for all $1 \le u \le v$, $1 \le i \le c$, $1 \le j \le c$ together with some further arcs corresponding to each clause.

We shall think of the arcs $\langle a_{u,i}, b_{u,i}\rangle$, $\langle A_{u,i}, B_{u,i}\rangle$ as corresponding to the potential occurrence of literals x_u and $-x_u$ respectively in the i'th clause. It can be verified that the only feedback arc sets of minimal size for the graph defined so far consist of the union over all u, $1 \le u \le v$, of either all the arcs corresponding to x_u or all those corresponding to $-x_u$. Thus there is a natural correspondence between these *potential feedback sets* and assignments to the variables. However G_F has in addition, for each clause C_i, $1 \le i \le c$, three arcs which connect the arcs corresponding to the literals of C_i into a circuit. For example if $C_i = \{x_3, -x_4, x_6\}$ the three arcs are $\langle b_{3,i}, A_{4,i}\rangle$, $\langle B_{4,i}, a_{6,i}\rangle$ and $\langle b_{6,i}, a_{3,i}\rangle$.

When one of the potential feedback sets described above is removed, the only possible circuits remaining in G_F are some 6-circuits corresponding to clauses. If the potential feedback set chosen corresponds to a satisfying assignment to F then all of these circuits will be broken. In this case we have a minimal feedback arc set of size cv. If there is no such satisfying assignment then at *least* one further arc must be removed. When we are given an assignment satisfying all but one of the clauses, at *most* one extra arc needs to be removed. Thus for those instances $\langle G_F, S\rangle$ where S corresponds to an assignment satisfying all but one clause of F, the satisfiability of F is equivalent to the non-minimality of S. $\square_{\{\text{Lemma 1}\}}$

We complete the proof that SNMP is NP-complete by reducing NON-MINIMAL FEEDBACK ARC SET to SNMP. Let $G = \langle V, A\rangle$ and suppose $\langle G, S\rangle$ is an instance of NON-MINIMAL FEEDBACK ARC SET. Since S is a feedback arc set for G, the vertex set V can be ordered so that every arc in A–S is a 'forward arc', i.e. of the form $\langle a, b\rangle$ where $a < b$ in this ordering. Without loss of generality suppose that $V = \{1,...,r\}$ where the natural ordering provides such an ordering. We construct a word $w_{G,S}$ over the alphabet V such that sorting $w_{G,S}$ with respect to V provides a minimal-transposition partitioning if and only if S is a minimal size set of feedback arcs.

Consider first an arbitrary palindrome, P, over V. It is easy to see that the partition of P with respect to any permutation of V requires exactly the same number of transpositions. If P is now modified by interchanging one pair of adjacent symbols, say ij is changed to ji, then any permutation of V where i precedes j will require two more transpositions than the others. A similar observation holds when several such interchanges are made. We therefore construct $w_{G,S}$ as follows.

Suppose that $A = \{a_1, ..., a_p\}$. For $1 \le j \le p$, let $e_j = uv$ and $E_j = vu$ where $a_j = \langle u, v\rangle$. Then:

$$w_{G,S} = e_1e_2...e_pe_p...e_2e_1$$

which is palindromic except for transpositions corresponding to each arc of A. Let q be the number of transpositions required to partition the palindromic string $e_1e_2...e_pE_p...E_2E_1$. As

observed above q is independent of the ordering chosen for the partition. For any permutation π of V, let $back(G, \pi) = \{\langle \pi(i), \pi(j)\rangle \in A \mid i > j\}$. Then

$$inv(w_{G,S}, \pi) = q - |A| + 2|back(G, \pi)|.$$

Hence,

$\langle G, S\rangle \in$ NON-MINIMAL FEEDBACK ARC SET

$\Leftrightarrow \exists \pi$ such that $|back(G, \pi)| < |S|$

$\Leftrightarrow \exists \pi$ such that $inv(w_{G,S}, \pi) < inv(w_{G,S}, \iota)$

$\Leftrightarrow w_{G,S} \in$ SNMP. $\square_{\{\text{Theorem 2}\}}$

This also completes the proof of the Main Theorem. $\square_{\{\text{Main Theorem}\}}$

3. Proving Theorem 1.

We first collect some well-known general definitions and facts about braid groups which will be used in the proof.

The mapping $^{\#}: \{\sigma_1,...,\sigma_{n-1}\} \to S_n$, defined by

$$\sigma_i^{\#} = (i,i+1), \tag{2}$$

can be extended to a homomorphism $^{\#}: B_n \to S_n$. The kernel of $^{\#}$ is denoted by $\mathcal{A}_n$. Geometrically, $x^{\#}$ is the permutation on strings realized by a braid x, and $\mathcal{A}_n$ is the subgroup of those braids which ultimately return the strings to their initial order. Given an occurrence $u*a*v$ of a letter a from the alphabet $\{\sigma_1,...,\sigma_{n-1}\}^{\pm 1}$ in a word uav, denote by $\chi(u*a*v)$ the initial indices of the (unordered) pair of strings that cross at $u*a*v$. Formally, if $a = \sigma_k^{\pm 1}$ then

$$\chi(u*a*v) = \{(u^{\#})^{-1}(k), (u^{\#})^{-1}(k+1)\}. \tag{3}$$

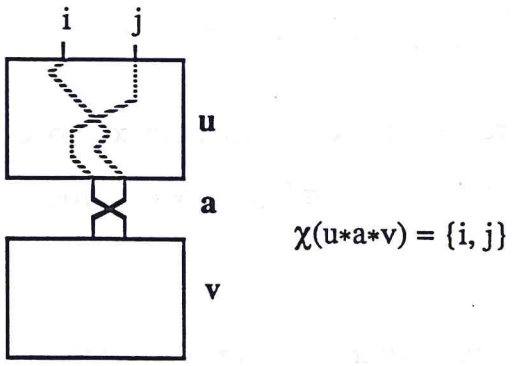

Fig. 3

Given $I \subseteq \{1,...,r\}$ and a word x we define a new word $\mu_I(x)$ in $|I| - 1$ generators. Geometrically, $\mu_I(x)$ is the result of 'dissolving' all strings not belonging to I. To give an

algebraic definition, we first introduce a mapping θ_I, on occurrences of the form $u*\sigma_k^{\pm 1}*v$ and taking values in the set $\{\sigma_1,...,\sigma_{|I|-1}, \Lambda\}$, defined by:

$$\theta_I(u*\sigma_k^{\varepsilon}*v) = \sigma_q^{\varepsilon} \text{ if } \chi(u*\sigma_k^{\varepsilon}*v) \subseteq I, \qquad (4)$$

$$= \Lambda \text{ otherwise,}$$

where q is the number of those i for which $1 \le i \le k$ and $(u^{\#})^{-1}(i) \in I$.

Geometrically, it is the crossing obtained from $u*\sigma_k^{\pm 1}*v$ by deleting strings numbered by $\{1,...,r\} \setminus I$; if at least one string forming $u*\sigma_k^{\pm 1}*v$ is deleted then this crossing is destroyed and the result is Λ. Now θ_I can be extended to arbitrary occurrences by:

$$\theta_I(u*a_1...a_n*v) = \prod_{1 \le i \le n} \theta_I(ua_1...a_{i-1}*a_i*a_{i+1}...a_n v).$$

Finally set

$$\mu_I(x) = \theta_I(*x*). \qquad (5)$$

The following three facts are clear from the geometry and can be checked by calculation.

Fact 1 $x \equiv y$ implies $\mu_I(x) \equiv \mu_I(y)$. □

Therefore μ_I can be regarded as a mapping from B_n to $B_{|I|}$.

Fact 2 μ_I restricted to $\mathcal{A}_n$ is a group homomorphism from $\mathcal{A}_n$ to $\mathcal{A}_{|I|}$. □

Fact 3 θ_I and χ 'commute', i.e., if $\chi(u*\sigma_k^{\varepsilon}*v) \subseteq I$ then :

$$\chi(\theta_I(*u*\sigma_k^{\varepsilon}v)*\theta_I(u*\sigma_k^{\varepsilon}*v)*\theta_I(u\sigma_k^{\varepsilon}*v*))$$

coincides with $\chi(u*\sigma_k^{\varepsilon}*v)$ renumbered relative to I. □

We need also three particular facts about the group B_3.

Lemma 2 Let x be a word over $\{\sigma_1,\sigma_2\}^{\pm 1}$ such that $x \equiv (\sigma_1\sigma_2^2\sigma_1)^d$ in B_3 for some d. If x contains a subword σ_ε^e where $\varepsilon \in \{1,2\}$, then x contains at least $(e-2)$ negative occurrences.

Proof We shall exploit the fact that the structure of B_3 (unlike larger braid groups) is very clear. Indeed, applying the automorphism

$$\sigma_1 = \Delta^{-1}a^{-2}, \ \sigma_2 = a\Delta, \ (a = \sigma_1^{-1}\sigma_2^{-1}, \ \Delta = \sigma_2\sigma_1\sigma_2)$$

To the presentation (1) with $n = 3$, we see that in the new generators B_3 has the form

$$B_3 = \langle \Delta, a \mid \Delta^{-2} = a^3 \rangle,$$

i.e., B_3 is a free product with amalgam (see e.g. [LS]). Each element $y \in B_3$ has a uniquely determined normal form:

$$y = \Delta^{2p} a^{\varepsilon_0} \Delta a^{\varepsilon_1} \Delta \ldots a^{\varepsilon_{r-1}} \Delta a^{\varepsilon_r},$$

where $\varepsilon_0, \varepsilon_r \in \{0,1,2\}$ and $\varepsilon_1, \ldots, \varepsilon_{r-1} \in \{1,2\}$. For instance,

$$\sigma_1 \equiv \Delta a,\ \sigma_2 \equiv a\Delta,\ \sigma_1^{-1} \equiv a^2\Delta,\ \sigma_2^{-1} \equiv \Delta a^2, \text{ and } x \equiv \Delta^{2d}(\Delta a^2 \Delta a^2)^d$$

exhibit the normal forms.

Any word in $\{\Delta, \Delta^{-1}, a\}$ (i.e. with only positive occurrences of a) can be reduced to its normal form by successive operations of transferring an occurrence of $\Delta^{\pm 2}$ to the left end, cancellation of Δ with Δ^{-1}, and *a-replacement*, where a^3 is replaced by Δ^{-2}. For any such word z, define $w(z)$ to be the number of occurrences of a in z. Note that, for the normal forms given above,

$$w(\Delta a) = w(a\Delta) = 1;\ w(a^2\Delta) = w(\Delta a^2) = 2 \text{ and } w(\Delta^{2d}(\Delta a^2 \Delta a^2)^d) = 4d.$$

Let g be the total number of negative crossings in x. Using the obvious homomorphism η: $B_3 \to \langle \mathbb{Z}, + \rangle$, given by $\eta(\sigma_1) = \eta(\sigma_2) = 1$, we find that the number of positive crossings must be $4d+g$. Substituting the normal forms for σ_1, σ_2, σ_1^{-1} and σ_2^{-1} in x, we obtain a longer (possibly reducible) word x' with $w(x') = 4d+3g$. x' contains a subword $(a\Delta)^{e-1}a$, from the image of the subword σ_ε^e. But the reduced normal form of x' ($\equiv$x) contains a's only in occurrences of a^2, hence all the e single occurrences of a in x' must be either cancelled or consolidated when going from x' to its normal form. Therefore in doing so, one needs at least (e-2) a–cancellations. However, any such cancellation decreases the value of w by three, so the total number of cancellations does not exceed g. This completes the proof. $\square_{\{\text{Lemma 2}\}}$

The following claim is proved in a similar way. We count occurrences of a^2 in x' rather than of single a's.

Lemma 3 Let x be a word over $\{\sigma_1, \sigma_2\}^{\pm 1}$ such that $x \equiv (\sigma_1 \sigma_2^{2t} \sigma_1^{2t-1})^s$ or $x \equiv (\sigma_1^{2t-1} \sigma_2^{2t} \sigma_1)^s$ for some t,s. If x contains a positive subword with e alternations between σ_1 and σ_2, then x contains at least $(e-2s)/3$ negative occurrences.

Proof Beginning as in the proof of Lemma 2, we obtain a (possibly reducible) word x', with $w(x') = 4ts+3g$, and $x' \equiv x$. Each of the e occurrences of $\sigma_1\sigma_2$ or $\sigma_2\sigma_1$ yields directly or indirectly an adjacent pair of a's in x', but the normal form for x has just $2s$ occurrences of a^2. Since adjacencies between a's are only removed (at most three at a time) by the replacement of a^3 by Δ^{-2} with the consequent reduction of w by three, we deduce that $e-2s \leq 3g$. $\square_{\{\text{Lemma 3}\}}$

We need now a deeper fact about subwords of those words considered in Lemma 3.

Lemma 4 Let x be a word over $\{\sigma_1, \sigma_2\}^{\pm 1}$ such that $x \equiv (\sigma_1\sigma_2^{2t}\sigma_1^{2t-1})^s$ or $x \equiv (\sigma_1^{2t-1}\sigma_2^{2t}\sigma_1)^s$ for some t,s. If x contains a subword u with the following properties:

(i) the function χ considered on all one-letter occurrences in u misses at least one of the three possible values, i.e., some pair does not cross within u, and

(ii) for some $\varepsilon \in \{1, 2\}$, u contains fewer than $2t$ occurrences of $\sigma_\varepsilon^{\pm 1}$,

then x contains at least $|u|/4 - 3t + 1$ negative occurrences.

Proof Let x' be the word obtained from x by performing all 'free' cancellations, i.e., cancellations of $\sigma_\varepsilon^{-1}\sigma_\varepsilon$ or $\sigma_\varepsilon\sigma_\varepsilon^{-1}$. The image u' of u in x' possesses both properties (i) and (ii), and $|u'| \geq |u| - 2g$ where g is the number of negative occurrences in x. Moreover, u' cannot contain any occurrences of the form $\sigma_1^{\delta_1}\sigma_2^{\delta_2}\sigma_1^{\delta_3}$ or $\sigma_2^{\delta_1}\sigma_1^{\delta_2}\sigma_2^{\delta_3}$, ($\delta_i \in \{-1,1\}$), because these would contradict (i).

As in the proof of Lemmas 2 and 3, substitute the normal forms for $\sigma_1^{\pm 1}$ and $\sigma_2^{\pm 1}$ in x' and u'. If x'' and u'' are the resulting words, the same arguments as before show that the total number of a–replacements in x'' when going to its normal form does not exceed g. But we know that u' does not contain any occurrences of the forms $\sigma_\varepsilon^{-1}\sigma_\varepsilon$, $\sigma_\varepsilon\sigma_\varepsilon^{-1}$, $\sigma_1^{\pm 1}\sigma_2^{\pm 1}\sigma_1^{\pm 1}$ or $\sigma_1^{\pm 1}\sigma_2^{\pm 1}\sigma_1^{\pm 1}$. It is easy to see that these are the only kinds of occurrence which could result in a cancellation of a–syllables in u'. Therefore we have only consolidations between a–syllables when going from u'' to its normal form u''', and, in particular, u''' contains at least $|u|/2 - g$ a–syllables. No more than $g+2$ of them can be affected within x''. Therefore u''' and x''', the normal forms of u'' and x'', contain a common piece p with at least

$|u|/2 - 2g - 2$ a–syllables. Assuming that $g < |u|/4 - 3t - 1$, we find that p contains more than $6t - 4$ a–syllables.

But $x''' \equiv x$, and so $x''' = \Delta^{2e}A$, for some e, where A is a subword of the periodic word $(a^2\Delta(a\Delta)^{2t-2})^{\infty}$. Therefore p contains at least 3 a-syllables of the form a^2. Let $p = p_1a^2\Delta(a\Delta)^{2t-2}a^2\Delta(a\Delta)^{2t-2}a^2\Delta p_2$. The middle occurrence of a^2 in u''' resulted from an occurrence of $\sigma_1^{\delta}\sigma_2^{\delta}$ or $\sigma_2^{\delta}\sigma_1^{\delta}$ in u'. Without loss of generality we can assume that it was obtained from $u_1'*\sigma_1^{\delta}\sigma_2^{\delta}*u_2'$. But then, to get the piece $a^2\Delta(a\Delta)^{2t-2}$, we would require $u_2' = \sigma_2^{\delta(2t-2)}u_3'$, because any occurrence $\sigma_2^{\delta}\sigma_1^{\delta}$ results in a^2 in u_1. So, u' (and therefore also u) contains at least 2t occurrences of $\sigma_2^{\pm 1}$, and similarly for $\sigma_1^{\pm 1}$. This contradicts (ii) and the Lemma is proved. $\square_{\{\text{Lemma 4}\}}$

For your convenience we repeat the statement of Theorem 1.

Theorem 1 Let w be a word such that $w \equiv x(q)$. Then:

(i) w has at least 2tmcs positive crossings between the weft and wires;

(ii) if the length of w is minimal then there exists some permutation π of Σ such that w has at least $2c^2.\text{inv}(q, \pi)$ crossings between the wires.

Proof of Theorem 1(i) It is evident that, for all i, $\mu_{\{\text{weft},i\}}(x(q)) \equiv \sigma_1^{2ts} \in B_2$. (1) implies that $B_2 \cong \mathbb{Z}$, and hence the (possibly reducible) word $\mu_{\{\text{weft},i\}}(w)$, equivalent to σ_1^{2ts}, contains at least 2ts positive occurrences of σ_1. By (4), (5) all these occurrences came from positive occurrences of the form $u*a*v$ with $\chi(u*a*v) = \{\text{weft}, i\}$. Summing over all i, we find that the number of positive occurrences $u*a*v$ for which $\text{weft} \in \chi(u*a*v)$ is at least 2tmcs. $\square_{\{\text{Theorem 1(i)}\}}$

Proof of Theorem 1(ii) Set $n = mc$. We have already seen that there exists a w_0 such that $w_0 \equiv x(q)$ and $|w_0| < n^2 + 2tsn$, therefore

$$|w| < n^2 + 2tsn. \tag{6}$$

From (6) and Theorem 1(i), we obtain:

$$(\text{number of negative occurrences in } w) < n^2. \tag{7}$$

Let T be the number of crossings between wires in *different* cables. We shall actually prove a stronger result than that stated in Theorem 1(ii), namely that for some $\pi \in \Sigma$, $T \geq 2c^2.\mathrm{inv}(q,\pi)$. Pick at random a system of representatives $i_1,\ldots,i_m$ in the cables, one wire per cable. Then the expectation of the total number of crossings among $i_1,\ldots,i_m$ is T/c^2. Choose such a system $i_1,\ldots,i_m$ that the number of crossings is minimal and apply $\mu_{\{\mathrm{weft},i_1,\ldots,i_m\}}$ to w. Then our problem is reduced to the case $c = 1$ with the difference that $\mu_{\{\mathrm{weft},i_1,\ldots,i_m\}}(w)$ can be non-minimal. However (7) still holds (with the original value of c) for $\mu_{\{\mathrm{weft},i_1,\ldots,i_m\}}(w)$. So, assume $c = 1$ and a word w is given such that $w \equiv x(q)$ and (7) holds.

If w contains at least m^2 wire crossings then the result is proved. Otherwise, w contains a piece u ($w = puq$) with no wire crossings within u such that $|u| \geq 2ts/m$, because $|w| \geq 2tsm$ by Theorem 1(i) with $c = 1$. To complete our proof it is sufficient to show that, for all k, $p^{\#}$ arranges the k–wires into a consecutive block, because we can then take the permutation of the alphabet induced by $p^{\#}$ as π.

Start by choosing a wire (say j_0) such that the weft has at least $2ts/m^2$ crossings with this wire within u.

Claim If $a_j \neq a_{j_0}$ $(1 \leq j \leq m)$ then within u the weft has at least $2t - 1$ crossings with the wire j.

Proof of Claim Apply $\mu_{\{\mathrm{weft},j,j_0\}}$ to the word w. We obtain a word of the form $(\sigma_1\sigma_2^{2t}\sigma_1^{2t-1})^s$ or $(\sigma_1^{2t-1}\sigma_2^{2t}\sigma_1)^s$. $\theta_{\{\mathrm{weft},j,j_0\}}(p*u*q)$ satisfies property (i) of Lemma 4, because j and j_0 do not intersect within u. But by (7) the total number of negative occurrences in $\mu_{\{\mathrm{weft},j,j_0\}}(w)$ is less than n^2 and $n^2 < ts/(2m^2) - 3t + 1 = t + 1$. So the conclusion of Lemma 4 does not hold and therefore property (ii) must fail. Hence there are at least 2t occurrences of $\sigma_\varepsilon^{\pm 1}$ in $\theta_{\{\mathrm{weft},j,j_0\}}(p*u*q)$ for any $\varepsilon \in \{1,2\}$. Since $\theta_{\{\mathrm{weft},j,j_0\}}(p*u*q)$ does not contain any (j,j_0)–crossings, then if $p^{\#}(j_0) < p^{\#}(j)$, χ of any occurrences $\sigma_1^{\pm 1}$ is $\{\mathrm{weft},j_0\}$, otherwise χ of any occurrence $\sigma_2^{\pm 1}$ is $\{\mathrm{weft}, j\}$. The claim is proved. $\square_{\{\mathrm{Claim}\}}$

Assume now that k ($1 \le k \le r$) is fixed. By the Claim and the observation that $2ts/m^2 \ge 2t$, we can choose a k–wire i such that there are at least 2t crossings in w between i and the weft. By (7), there exists a segment v, where $w = pu_1vu_2q$, containing at least $2n^4$ ($=2t/n^2$) crossings between the weft and wire i, and no negative crossings. We claim that $p^{\#}$, and therefore $(pu_1)^{\#}$ because there are no wire crossings within $p*u*q$, takes all the wires associated with k to a consecutive block. Suppose not. Then there exists a k'–wire j inside the k–block such that $k' \neq k$. Choose also in the k–block a k–wire i' such that j lies between i and i'. Consider two cases.

CASE 1. There exists a subsegment v' of v, ($w = p'v'q'$), containing at least $n^2 + 2$ crossings of the weft with i, but no crossings of the weft with i'.

Since wires do not intersect each other within u, $p'^{\#}(j)$ lies between $p'^{\#}(i)$ and $p'^{\#}(i')$, as it was for p and pu_1. We now apply the mapping $\theta = \theta_{\{weft,\, i,\, i'\}}$ to the word w and find that:

$$\mu_{\{weft,\, i,\, i'\}}(w) = \theta(*w*) \equiv (\sigma_1\sigma_2{}^2\sigma_1)^{ts}.$$

From our knowledge of $p'*v'*q'$, the only possible crossings in $\theta(p'*v'*q')$ are between the weft and wire i. Therefore $\theta(p'*v'*q') = \sigma_\varepsilon{}^s$ where $s \ge n^2 + 2$ and $\varepsilon \in \{1, 2\}$. This yields a contradiction with (7) and Lemma 2.

CASE 2. There is no such subsegment as in Case 1. This implies that the sequence X of pairs representing successive crossings of v contains at least $4n^2 - 4$ alternations of {weft, i} and {weft, i'}. Since $(pu_1)^{\#}(j)$ lies between $(pu_1)^{\#}(i)$ and $(pu_1)^{\#}(i')$, between any such alternating pair in X there exists at least one occurrence of {weft, j}. So, there are at least $4n^2 - 4$ alternations of {weft, i} and {weft, j}. Applying the mapping $\mu_{\{weft,\, i,\, j\}}$, we find that:

(i) $\mu_{\{weft,\, i,\, j\}}(w)$ has at least $4n^2 - 4$ alternations of σ_1, σ_2;

and (ii) $\mu_{\{weft,\, i,\, j\}}(w) \equiv (\sigma_1{}^{2t-1}\sigma_2{}^{2t}\sigma_1)^s$ or $(\sigma_1\sigma_2{}^{2t}\sigma_1{}^{2t-1})^s$.

Because $4n^2 - 4 \ge 3n^2 + 2s$, this contradicts Lemma 3.

This contradiction with the assumption that there exists a k'-wire lying within the k-block shows that for any k, $p^{\#}$ takes all the k–wires to a consecutive block. As observed above,

this completes the proof of Theorem 1(ii) since the required permutation π is just that induced by $p^{\#}$. $\square_{\{\text{Theorem 1(ii)}\}}$

References

[A1] Artin E., 'Theorie der Zöpfe', *Abh.math.Semin. Hamburg Univ.*, **4** (1925) 47-72.

[A2] Artin E., 'Theory of braids', *Annals of Mathematics*, **48** (1947) 101-126.

[GJ] Garey M.R., Johnson D.S., Computers and Intractability. A Guide to the Theory of NP-Completeness, (Freeman, San Francisco, 1979) (рус. пер. Гэри М., Джонсон Д., Вычислительные машины и труднорешаемые задачи, М., Мир, 1981).

[G] Garside F.A., 'The braid group and other groups', *Quarterly J. of Mathematics*, (Oxford Second Series) **20** (1969) 235-254 (рус. пер. Математика, 1970, 14:4, 117-132.)

[LS] Lyndon R.G., Shupp P.E., Combinatorial Group Theory, (Springer-Verlag, 1977) (рус. пер. Линдон Р., Шупп П., Комбинаторная теория групп, М., Мир, 1981).

[M] Марков А.А., 'Основы алгебраической теории кос', *Труды Математического Института им. В. А. Стеклова*, т. 16 (1945) (Russian with English summary).

[Th] Thurston W.P., 'Finite state algorithms for the braid groups', preliminary draft, 1988.