FREELY BRAIDED ELEMENTS IN COXETER GROUPS

R.M. Green and J. Losonczy

Department of Mathematics and Statistics
Lancaster University
Lancaster LA1 4YF
England
E-mail: r.m.green@lancaster.ac.uk

Department of Mathematics
Long Island University
Brookville, NY 11548
USA
E-mail: losonczy@e-math.ams.org

Abstract. We introduce a notion of "freely braided element" for simply laced Coxeter groups. We show that an arbitrary group element w has at most $2^{N(w)}$ commutation classes of reduced expressions, where $N(w)$ is a certain statistic defined in terms of the positive roots made negative by w. This bound is achieved if w is freely braided. In the type A setting, we show that the bound is achieved only for freely braided w.

1. Introduction

A well-known result in the theory of Coxeter groups states that any two reduced expressions for the same element w of a Coxeter group are equivalent under the equivalence relation generated by braid relations. If w is such that any two of its reduced expressions are equivalent by short braid relations (i.e., iterated commutations of commuting generators), we call w "fully commutative," following Stembridge [8].

In this paper, we introduce and study "freely braided elements" for an arbitrary simply laced Coxeter group. This is a more general class of elements than the aforementioned fully commutative elements. The idea behind the definition is that although it may be necessary to use long braid relations in order to pass between two

1991 *Mathematics Subject Classification.* 20F55.
Key words and phrases. braid relation, commutation class, Coxeter group, root sequence.

reduced expressions for a freely braided element, the necessary long braid relations in a certain sense do not interfere with one another.

Every reduced expression for a Coxeter group element w determines a total ordering of the set of positive roots made negative by w. The resulting sequences, called root sequences, play a central role in this paper. We are particularly interested in triples of roots of the form $\{\alpha, \alpha+\beta, \beta\}$, where each root in the triple is made negative by w (inversion triples). An inversion triple that occurs consecutively in some root sequence for w will be called contractible, and the statistic $N(w)$ mentioned in the abstract above is the number of contractible inversion triples of w. We note that Fan and Stembridge have already shown in [5, Theorem 2.4] that an element of a simply laced Coxeter group is fully commutative if and only if it has no inversion triples.

In §4 of this paper, we prove that the number of commutation classes (short braid equivalence classes of reduced expressions) of w is bounded above by $2^{N(w)}$, and this bound is achieved if w is freely braided. Furthermore, a freely braided element w has a root sequence in which each contractible inversion triple occurs as a consecutive subword (Theorem 4.2.3). In §5, we prove that if the Coxeter group is of type A, then the $2^{N(w)}$ bound is achieved only for freely braided w. Our proof relies on the fact that every inversion triple is contractible in the type A setting.

Along the way, we give a short proof that the commutation graph of any element of a simply laced Coxeter group is bipartite, extending a result in [4].

Another possible approach to proving these results is to use Viennot's heaps of pieces [10]. These are certain labelled posets that can be associated to commutation classes of reduced expressions for elements of a Coxeter group (see [8, §1.2]). It can be shown that the dual of the heap of a reduced expression is isomorphic, as an abstract poset, to a certain poset arising naturally from the associated root sequence. However, we find it more convenient to argue directly with root sequences in this paper.

2. Preliminaries

2.1 Basic terminology and notation.

Let (W, S) be a Coxeter system, with Coxeter matrix $\{m(s,t)\}_{s,t\in S}$. It will be assumed throughout this paper that (W, S) is simply laced, so that $m(s,t) \in \{2, 3\}$

for all pairs of distinct $s, t \in S$. We refer to W itself as a "Coxeter group". The basic facts concerning Coxeter systems can be found in [3] or [6].

Let S^* be the free monoid generated by S. The elements of S^* will be written as finite sequences, e.g., $\mathbf{s} = (s_1, \ldots, s_n)$. We define the *length* of any $\mathbf{s} \in S^*$ to be the number of its entries.

There is a natural morphism of monoids $\phi : S^* \longrightarrow W$, given by $\phi(s_1, \ldots, s_n) = s_1 \cdots s_n$ (the empty sequence is mapped to the identity, e). We say that an element $\mathbf{s} \in S^*$ *represents* its image $w = \phi(\mathbf{s}) \in W$; furthermore, if the length of $\mathbf{s}$ is minimal among the lengths of all the sequences that represent w, then we say that $\mathbf{s}$ is *reduced*, and we call $\mathbf{s}$ a *reduced expression* for w. The *length* of w, denoted by $\ell(w)$, is then equal to the length of $\mathbf{s}$.

Let V be a vector space over $\mathbb{R}$ with basis $\{\alpha_s : s \in S\}$, and let B be the Coxeter form on V associated to (W, S). This is the symmetric bilinear form satisfying $B(\alpha_s, \alpha_t) = -\cos \frac{\pi}{m(s,t)}$ for all $s, t \in S$. We shall view V as the underlying space of a reflection representation of W, determined by the equalities $s\,\alpha_t = \alpha_t - 2B(\alpha_t, \alpha_s)\alpha_s$, for $s, t \in S$. The Coxeter form is preserved by W relative to this representation.

Define $\Phi = \{w\alpha_s : w \in W \text{ and } s \in S\}$. This is the root system of (W, S). Let Φ^+ be the set of all $\beta \in \Phi$ such that β is expressible as a linear combination of the α_s with nonnegative coefficients, and let $\Phi^- = -\Phi^+$. Then Φ equals the disjoint union $\Phi^+ \cup \Phi^-$ [6, §5.4]. The elements of Φ^+ (respectively, Φ^-) are called "positive" (respectively, "negative") roots. The α_s are often referred to as "simple" roots.

One can associate to each $w \in W$ the set $\Phi(w) = \Phi^+ \cap w^{-1}(\Phi^-)$, which we call the *inversion set* of w. It is well known that $\Phi(w)$ has $\ell(w)$ elements and that w is uniquely determined by $\Phi(w)$. Given any reduced expression $(s_1, \ldots, s_n)$ for w, one has $\Phi(w) = \{r_1, \ldots, r_n\}$, where $r_1 = \alpha_{s_n}$ and $r_i = s_n \cdots s_{n-i+2}(\alpha_{s_{n-i+1}})$ for all $i \in \{2, \ldots, n\}$ [6, Exercise 1, §5.6]. Following [5], we form the sequence $\mathbf{r} = (r_1, \ldots, r_n)$, and call it the *root sequence of* $(s_1, \ldots, s_n)$, or a *root sequence for* w. Notice that any initial segment of a root sequence is also a root sequence for some element of W.

Let $w \in W$. Suppose that α, β are positive roots and a, b are positive integers such that $a\alpha + b\beta$ is a positive root. Then one has the following two implications: $\alpha, \beta \in \Phi(w) \Rightarrow a\alpha + b\beta \in \Phi(w)$ and $\alpha, \beta \notin \Phi(w) \Rightarrow a\alpha + b\beta \notin \Phi(w)$. This property of inversion sets is sometimes referred to as "biconvexity" (cf. [2, §3]).

2.2 Inversion triples.

Assume that (W, S) is of arbitrary simply laced type, unless stated otherwise.

Definition 2.2.1. Let $w \in W$. Any subset of $\Phi(w)$ of the form $\{\alpha, \alpha+\beta, \beta\}$ will be called an *inversion triple* of w. We say that an inversion triple T of w is *contractible* if there is a root sequence for w in which the elements of T appear consecutively (in some order). If the contractible inversion triples of w are pairwise disjoint, then w is said to be *freely braided.*

Remark 2.2.2. Let $\{\alpha, \alpha+\beta, \beta\}$ be an inversion triple of $w \in W$. It follows immediately from the final two paragraphs of §2.1 that $\alpha+\beta$ must occur between α and β in any root sequence for w.

The following example establishes the existence of non-contractible inversion triples.

Example 2.2.3. Suppose that (W, S) has a parabolic subgroup of type D_4, generated by $s_1, s_2, s_3, s_4 \in S$, where $s_2 s_i$ has order 3 for all $i \neq 2$, and $s_i s_j$ has order 2 whenever $i, j \in \{1, 3, 4\}$ are distinct. Write α_i for the simple root α_{s_i}. Consider the element $w = s_2 s_1 s_3 s_4 s_2 s_4 s_3 s_1 s_2$. The root sequence corresponding to the reduced expression $(s_2, s_1, s_3, s_4, s_2, s_4, s_3, s_1, s_2)$ is

$$(\alpha_2,\ \alpha_1+\alpha_2,\ \alpha_2+\alpha_3,\ \alpha_2+\alpha_4,\ \alpha_1+2\alpha_2+\alpha_3+\alpha_4,$$
$$\alpha_1+\alpha_2+\alpha_3,\ \alpha_1+\alpha_2+\alpha_4,\ \alpha_2+\alpha_3+\alpha_4,\ \alpha_1+\alpha_2+\alpha_3+\alpha_4).$$

Observe that α_2 belongs to precisely one inversion triple of w, namely

$$\{\alpha_2,\ \alpha_1+2\alpha_2+\alpha_3+\alpha_4,\ \alpha_1+\alpha_2+\alpha_3+\alpha_4\}.$$

However, there is no root sequence for w in which these roots appear consecutively. This can be seen as follows. By inspection of the above root sequence, α_2 is the only simple root in $\Phi(w)$. Hence, α_2 must appear first in any root sequence for w. The second root in any root sequence for w must be a linear combination of at most two simple roots; this rules out $\alpha_1+2\alpha_2+\alpha_3+\alpha_4$ as a possibility.

It will be shown in §5.1 that in the type A setting, every inversion triple is contractible.

3. Braid relations

In this section, (W, S) is assumed to be of arbitrary simply laced type. As in §2, we denote the associated Coxeter matrix by $\{m(s,t)\}_{s,t \in S}$.

3.1 More on root sequences.

The group W is by definition generated by the elements of S, subject only to the relations $(st)^{m(s,t)} = e$, for $s, t \in S$ [6, §5]. These relations can be restated as $s^2 = e$ for $s \in S$, $st = ts$ if $m(s,t) = 2$, and $sts = tst$ if $m(s,t) = 3$. Given any $s, t \in S$ and any nonnegative integer m, we write $(s,t)_m$ for the length m sequence $(s, t, s, \dots) \in S^*$. Let $\approx$ be the equivalence relation on S^* generated by the elementary relations $(s,t)_{m(s,t)} \approx (t,s)_{m(s,t)}$, for $s, t \in S$. When $m(s,t) \neq 1$, we call such an elementary relation a *braid relation*, qualifying it *short* or *long*, according as $m(s,t) = 2$ or 3.

Let $w \in W$. A well-known result, attributed variously to Matsumoto [7] and Tits [9], states that the equivalence class relative to $\approx$ of any reduced expression for w coincides with the set of all reduced expressions for w.

We use the term *commutation class* to refer to any equivalence class arising from the equivalence relation on S^* generated by the short braid relations.

Applying a braid relation to a given reduced expression corresponds to applying a permutation to the root sequence associated with that reduced expression. The following proposition makes this correspondence explicit.

Proposition 3.1.1. *Let $w \in W$, let $\mathbf{u}, \mathbf{v} \in S^*$ and let $s, s', s'' \in S$. Denote the length of $\mathbf{v}$ by k.*

(a) *Assume that $(\mathbf{u}, s, s', \mathbf{v})$ is a reduced expression for w, and let $\mathbf{r} = (r_i)$ be the root sequence of $(\mathbf{u}, s, s', \mathbf{v})$.*

(i) *Suppose $m(s,s') = 2$, so that $(\mathbf{u}, s', s, \mathbf{v})$ is also a reduced expression for w. Then the root sequence $\mathbf{r}'$ of $(\mathbf{u}, s', s, \mathbf{v})$ can be obtained from $\mathbf{r}$ by interchanging r_{k+1} and r_{k+2}, which are mutually orthogonal.*

(ii) *If r_{k+1} and r_{k+2} are orthogonal, then $m(s,s') = 2$.*

(b) *Assume that $(\mathbf{u}, s, s', s'', \mathbf{v})$ is a reduced expression for w, and let $\mathbf{r} = (r_i)$ be the root sequence of $(\mathbf{u}, s, s', s'', \mathbf{v})$.*

(i) *Suppose $s'' = s$, so that $m(s,s') = 3$ and $(\mathbf{u}, s', s, s', \mathbf{v})$ is also a reduced expression for w. Then the root sequence $\mathbf{r}'$ of $(\mathbf{u}, s', s, s', \mathbf{v})$ can be obtained from $\mathbf{r}$ by* *interchanging* r_{k+1} *and* r_{k+3}. *Furthermore, we have* $r_{k+1} + r_{k+3} = r_{k+2}$.

(ii) *If* $r_{k+1} + r_{k+3} = r_{k+2}$, *then* $s'' = s \neq s'$ *and* $m(s, s') = 3$.

Proof. We first prove the proposition under the additional hypothesis that $\mathbf{v}$ has length 0. For simplicity of notation, we denote the simple roots corresponding to s, s' and s'' by α, α' and α'', respectively.

Concerning (a), if $m(s, s') = 2$, then the reduced expression $(\mathbf{u}, s, s')$ has root sequence $(\alpha', \alpha, r_3, \dots)$ and $(\mathbf{u}, s', s)$ has root sequence $(\alpha, \alpha', r_3, \dots)$. Statement (i) immediately follows. For (ii), we observe that if the roots α' and $s'(\alpha)$ are orthogonal, then so are α and α', and hence $m(s, s') = 2$.

We turn to (b). If $s = s''$, then the expression $(\mathbf{u}, s, s', s)$ is reduced and has root sequence $(\alpha, \alpha + \alpha', \alpha', r_4, \dots)$; likewise, $(\mathbf{u}, s', s, s')$ is reduced and has root sequence $(\alpha', \alpha + \alpha', \alpha, r_4, \dots)$. This establishes (i). To prove (ii), we observe that in order for the equation

$$\alpha'' + s''s'(\alpha) = s''(\alpha')$$

to hold, we must have $s'' = s$; otherwise, the simple root α would lie in the support of the left hand side but not the right hand side (since $s \neq s'$). It follows that $m(s, s') = 3$ if $r_1 + r_3 = r_2$.

The general case (where $\mathbf{v}$ has length ≥ 0) follows from the above, using the definition of root sequence together with the fact that W acts linearly on V and preserves B. $\square$

Definition 3.1.2. Let $\mathbf{r}$ and $\mathbf{r}'$ be as in part (a)(i) (respectively, part (b)(i)) of Proposition 3.1.1. We call the passage from $\mathbf{r}$ to $\mathbf{r}'$ a *short braid move* (respectively, *long braid move*). We say that two root sequences are *commutation equivalent* if one can be transformed into the other by applying a (possibly empty) sequence of short braid moves. The set of all root sequences that are commutation equivalent to a given root sequence will be called the *commutation class* of that root sequence.

Remark 3.1.3. Let $w \in W$. The recipe that associates a root sequence to a reduced expression defines a bijection from the set of all reduced expressions for w to the set of all root sequences for w. By Proposition 3.1.1, this bijection is compatible with the application of both long and short braid moves. Hence, by the result of Matsumoto and Tits cited earlier, any root sequence for w may be transformed into any other by applying a (possibly empty) sequence of long and short braid moves.

Definition 3.1.4. Let $w \in W$ and let $\mathbf{r} = (r_1, \ldots, r_n)$ be any root sequence for w. Following [4], we introduce a partial ordering $\leq_{\mathbf{r}}$ of $\Phi(w)$ by stipulating that $r_i <_{\mathbf{r}} r_j$ if $i < j$ and the roots r_i and r_j are not orthogonal relative to B.

The following result is closely related to [5, Proposition 2.2].

Proposition 3.1.5. *Let $w \in W$ and let $\mathbf{r}$ and $\mathbf{r}'$ be root sequences for w with respective partial orderings $\leq_{\mathbf{r}}$ and $\leq_{\mathbf{r}'}$ of $\Phi(w)$. Then $\leq_{\mathbf{r}}$ and $\leq_{\mathbf{r}'}$ agree if and only if $\mathbf{r}$ and $\mathbf{r}'$ are commutation equivalent.*

Proof. For the "if" part, it is enough to treat the case where $\mathbf{r}$ and $\mathbf{r}'$ differ by a single short braid move. In such a situation, the conclusion follows immediately from Definition 3.1.4.

For the converse, write $\mathbf{r} = (r_1, \ldots, r_n)$ and $\mathbf{r}' = (r'_1, \ldots, r'_n)$. Let g be the permutation of $\{1, \ldots, n\}$ such that $r_i = r'_{g(i)}$ for all $i \in \{1, \ldots, n\}$. If $\leq_{\mathbf{r}}$ and $\leq_{\mathbf{r}'}$ agree, then, for all $i, j \in \{1, \ldots, n\}$ such that $r'_{g(i)}$ and $r'_{g(j)}$ are nonorthogonal, we have $i < j$ if and only if $g(i) < g(j)$. In particular, this means that if i is an index less than n such that $r'_{g(i)}$ is not orthogonal to $r'_{g(n)}$, then $g(i) < g(n)$. Hence, $r'_{g(n)}$ can be moved to the last entry of the root sequence $\mathbf{r}'$ using only short braid moves. Denote the resulting root sequence by $\mathbf{r}''$. Since any initial segment of a root sequence is also a root sequence, and since any element of W is uniquely determined by its inversion set, there is a $w' \in W$ such that the first $n-1$ entries of $\mathbf{r}$ and the first $n-1$ entries of $\mathbf{r}''$ are both root sequences for w'. The partial orders corresponding to these shorter root sequences agree. By induction, each of the shorter root sequences can be transformed into the other using only short braid moves. $\square$

3.2 More on inversion triples.

The proposition below characterizes contractible inversion triples. Its proof depends on the following basic fact from the theory of Coxeter groups: given $w \in W$ and $s \in S$, one has $\alpha_s \in \Phi(w)$ if and only if w has a reduced expression with s as its last entry (see [6, Proposition 5.7, Theorem 5.8]).

Proposition 3.2.1. *Let $w \in W$ and let $T = \{\alpha, \alpha + \beta, \beta\}$ be an inversion triple of w. The following are equivalent:*

(i) *T is contractible;*

(ii) *there is a root sequence* $\mathbf{r}$ *for* w *such that* $\alpha+\beta$ *covers* α *or* β *relative to* $\leq_{\mathbf{r}}$;

(iii) *there is a root sequence* $\mathbf{r}$ *for* w *such that* α *or* β *covers* $\alpha+\beta$ *relative to* $\leq_{\mathbf{r}}$.

Proof. The implications (i) $\Rightarrow$ (ii) and (i) $\Rightarrow$ (iii) are immediate from the definition of contractible inversion triple.

To prove (ii) $\Rightarrow$ (i), it is enough by symmetry to treat the case where there is a root sequence $\mathbf{r}$ for w such that $\alpha+\beta$ covers α relative to $\leq_{\mathbf{r}}$. We may assume that α and $\alpha+\beta$ appear consecutively in $\mathbf{r}=(r_1,\ldots,r_n)$ (apply a sequence of short braid moves to $\mathbf{r}$ if necessary). In other words, we may assume the existence of an integer k satisfying $r_k=\alpha$ and $r_{k+1}=\alpha+\beta$. Let $(s_1,\ldots,s_n)$ be the reduced expression corresponding to $\mathbf{r}$. Parse this reduced expression as $(\mathbf{u}, s_{n-k}, s_{n-k+1}, \mathbf{v})$, and denote the images of $\mathbf{u}$ and $\mathbf{v}$ in W by u and v, respectively. For notational simplicity, write s for s_{n-k} and t for s_{n-k+1}. Then $\alpha=v^{-1}(\alpha_t)$ and $\alpha+\beta=v^{-1}t(\alpha_s)$. Note for future reference that $t(\alpha_s)=s(\alpha_t)=\alpha_s+\alpha_t$.

Contractibility of T will follow if we can produce a root sequence for ust of the form $(v(\alpha),\ v(\alpha+\beta),\ v(\beta),\ \ldots)=(\alpha_t,\ t(\alpha_s)=\alpha_s+\alpha_t,\ \alpha_s,\ \ldots)$. Observe that

$$v^{-1}(\alpha_s)=v^{-1}t(\alpha_s)-v^{-1}(\alpha_t)=(\alpha+\beta)-\alpha=\beta\in\Phi(w),$$

hence $\alpha_s\in\Phi(ust)$, and consequently $\alpha_t=st(\alpha_s)\in\Phi(u)$. Thus, u has a reduced expression with last entry t, and this gives us a reduced expression for ust of the form $(\ldots,t,s,t)$. The latter corresponds to a root sequence with the required property.

The proof of (iii) $\Rightarrow$ (i) proceeds along similar lines. We start with a root sequence $\mathbf{r}=(r_1,\ldots,r_n)$ for w for which there is an index k satisfying $r_k=\alpha+\beta$ and $r_{k+1}=\beta$. We parse the reduced expression $(s_1,\ldots,s_n)$ corresponding to $\mathbf{r}$ as $(\mathbf{u}, s_{n-k}, s_{n-k+1}, \mathbf{v})$, denote the images of $\mathbf{u}$ and $\mathbf{v}$ in W as u and v, and write s for s_{n-k} and t for s_{n-k+1}. Then $\alpha+\beta=v^{-1}(\alpha_t)$ and $\beta=v^{-1}t(\alpha_s)$.

We will be finished if we can show that v has a reduced expression with first entry s, say $(s,\mathbf{x})$. For then, letting x denote the image in W of $\mathbf{x}$, we would have

$$x^{-1}(\alpha_s)=-x^{-1}sts(\alpha_t)=-v^{-1}ts(\alpha_t)=v^{-1}(\alpha_t)-v^{-1}t(\alpha_s)=(\alpha+\beta)-\beta=\alpha,$$

showing that $\alpha,\alpha+\beta,\beta$ appear consecutively in some root sequence for w. An equivalent requirement is for $v^{-1}(\alpha_s)$ to be a negative root, and this follows because

$$v^{-1}(\alpha_s)=v^{-1}(\alpha_s+\alpha_t)-v^{-1}(\alpha_t)=v^{-1}t(\alpha_s)-v^{-1}(\alpha_t)=\beta-(\alpha+\beta)=-\alpha.$$

The proof is complete. □

The next proposition will be needed in §4.1.

Proposition 3.2.2. *Let $w \in W$. Let $\alpha, \beta \in \Phi(w)$ be distinct and nonorthogonal. Suppose that there exist root sequences $\mathbf{r}$ and $\mathbf{r}'$ for w such that $\alpha \leq_{\mathbf{r}} \beta$ and $\beta \leq_{\mathbf{r}'} \alpha$. Then there is a contractible inversion triple of w containing both α and β.*

Proof. The relation $\alpha \leq_{\mathbf{r}} \beta$ implies that in $\mathbf{r}$, the root α appears to the left of β; likewise, in $\mathbf{r}'$ the root β appears to the left of α. Now, by Remark 3.1.3, $\mathbf{r}$ can be transformed into $\mathbf{r}'$ by a sequence of long and short braid moves. Since α and β are nonorthogonal, they cannot both be involved in the same short braid move. Hence, there is a long braid move in our sequence of moves that involves both α and β. This implies the existence of a contractible inversion triple of w containing α and β. □

4. The statistic $N(w)$ and free braidedness

Throughout this section, (W, S) will be of arbitrary simply laced type.

4.1 The map F_w.

Let $w \in W$. Fix an arbitrary antisymmetric relation $\preceq$ on $\Phi(w)$ with the property that any two roots in $\Phi(w)$ are comparable relative to $\preceq$.

Let $\mathcal{C}(w)$ and $\mathcal{I}(w)$ denote the set of commutation classes of root sequences for w and the set of contractible inversion triples of w, respectively. We define a map

$$F_w : \mathcal{C}(w) \longrightarrow \{0, 1\}^{\mathcal{I}(w)},$$

depending on $\preceq$, as follows. Let $C \in \mathcal{C}(w)$ and let $\leq_C$ be the partial ordering of $\Phi(w)$ determined by C; this is well-defined by Proposition 3.1.5. Given any $\{\alpha, \alpha+\beta, \beta\} \in \mathcal{I}(w)$, we define $F_w(C)(\{\alpha, \alpha+\beta, \beta\})$ to be 0 if α and β are in the same relative order with respect to $\leq_C$ and $\preceq$, and otherwise we define $F_w(C)(\{\alpha, \alpha+\beta, \beta\})$ to be 1. (Note that α and β are comparable under $\leq_C$.)

Theorem 4.1.1. *The map F_w is injective.*

Proof. Let C and C' be distinct root sequence commutation classes of w, and let $\leq_C$ and $\leq_{C'}$ be the respective partial orderings of $\Phi(w)$. By Proposition 3.1.5, these partial orderings do not agree. Hence, there exist nonorthogonal roots $\alpha, \beta \in \Phi(w)$

such that $\alpha <_C \beta$ and $\beta <_{C'} \alpha$ (note that any two nonorthogonal roots in $\Phi(w)$ are comparable under any given partial ordering arising from a commutation class of w). We now invoke Proposition 3.2.2 to deduce that α and β belong to some contractible inversion triple T of w. It is clear from the definition of F_w that $F_w(C)(T) \neq F_w(C')(T)$. $\square$

Theorem 4.1.1 has the following immediate

Corollary 4.1.2. *Every $w \in W$ has at most $2^{N(w)}$ commutation classes, where $N(w)$ denotes the number of contractible inversion triples of w.* $\square$

Remark 4.1.3. In [4], Elnitsky exhibits, in types A, B and D, a bijection from the set of all commutation classes of an arbitrary group element to a set of rhombic tilings of a polygon determined by that element. Bédard constructs in [1] a bijection from the set of all commutation classes of an arbitrary Weyl group element w to a certain set of functions from $\Phi(w)$ to the set of positive integers.

4.2 Freely braided elements and commutation classes.

In this section, we prove that the freely braided elements of Definition 2.2.1 achieve the bound of Corollary 4.1.2.

Given any $v \in V$, we denote by $v^\perp$ the subspace $\{w \in V : B(v, w) = 0\}$ of V.

Lemma 4.2.1. *Let $w \in W$ have length at least 3, and let $\alpha_1, \alpha_2, \alpha_3 \in \Phi(w)$ be distinct. Suppose that $\mathbf{r}$ and $\mathbf{r}'$ are root sequences for w such that α_2 appears between α_1 and α_3 in $\mathbf{r}$ but not in $\mathbf{r}'$. Then either $\alpha_2^\perp \cap \{\alpha_1, \alpha_3\} \neq \emptyset$ or else there is a contractible inversion triple of w containing α_2 and one of α_1, α_3.*

Proof. By Remark 3.1.3, it is possible to transform $\mathbf{r}$ into $\mathbf{r}'$ using short and long braid moves. Consideration of such a sequence of moves yields the result. $\square$

Lemma 4.2.2. *Let $w \in W$ be freely braided, and let $T = \{\alpha, \alpha + \beta, \beta\}$ be a contractible inversion triple of w. Suppose that*

$$(\ldots, \alpha, \gamma_1, \gamma_2, \ldots, \gamma_k, \alpha + \beta, \delta_1, \delta_2, \ldots, \delta_m, \beta, \ldots)$$

is a root sequence for w. Then each root γ_i is orthogonal to α, and each root δ_i is orthogonal to β.

Proof. First assume toward a contradiction that γ_1 is not orthogonal to α. We apply Lemma 4.2.1 with $\alpha_1 = \alpha$, $\alpha_2 = \gamma_1$ and $\alpha_3 = \alpha + \beta$. Because T is contractible, there is a root sequence for w in which γ_1 does not appear between the roots α and $\alpha+\beta$. Further, since w is freely braided, and since γ_1 does not lie in T, it cannot be the case that γ_1 belongs to a contractible inversion triple containing either α or $\alpha+\beta$. Lemma 4.2.1 then implies $\gamma_1^\perp \cap \{\alpha, \alpha+\beta\} \neq \emptyset$. But γ_1 is not orthogonal to α by assumption; hence, $\gamma_1 \perp (\alpha+\beta)$. This in turn implies that γ_1 is not orthogonal to β. Apply Lemma 4.2.1 again with $\alpha_1 = \alpha$, $\alpha_2 = \gamma_1$ and $\alpha_3 = \beta$, and use the same reasoning to show $\gamma_1 \perp \beta$. This is a contradiction.

An induction on k now shows that the γ_i are all orthogonal to α, and an analogous argument handles the δ_i. $\square$

The following theorem and its corollary provide some justification for the terminology "freely braided element".

Theorem 4.2.3. *Let $w \in W$ be freely braided. Then there is a root sequence* $\mathbf{r}$ *for w such that the roots in any given contractible inversion triple of w appear consecutively in* $\mathbf{r}$. *Furthermore, every root sequence for w is commutation equivalent to such a root sequence* $\mathbf{r}$.

Proof. Let $\mathbf{r}'$ be an arbitrary root sequence for w, and let $T = \{\alpha, \alpha+\beta, \beta\}$ be a contractible inversion triple of w. By Remark 2.2.2 and symmetry, we may assume that in $\mathbf{r}'$, the root α appears to the left of $\alpha+\beta$, which in turn appears to the left of β. According to Lemma 4.2.2 and Proposition 3.1.1(a), there is a sequence of short braid moves that transforms $\mathbf{r}'$ into a root sequence $\mathbf{r}''$ for w in which α, $\alpha+\beta$ and β appear consecutively in that order.

The conclusion of the theorem follows from the above by an induction on the number of contractible inversion triples. The crucial point to notice is that if T' is a consecutively occurring contractible inversion triple in $\mathbf{r}'$, then the sequence of short braid moves of the preceding paragraph may be chosen so that the roots occurring in T' remain consecutive in $\mathbf{r}''$. This can be arranged because w is freely braided. $\square$

Corollary 4.2.4. *If $w \in W$ is freely braided, then w has precisely $2^{N(w)}$ commutation classes of root sequences, where $N(w)$ is the number of contractible inversion triples of w.*

Proof. In view of Corollary 4.1.2, it suffices to show that w has at least $2^{N(w)}$

commutation classes. Theorem 4.2.3 guarantees the existence of a root sequence $\mathbf{r}$ for w in which the elements of each contractible inversion triple occur consecutively; moreover, the contractible inversion triples in $\mathbf{r}$ are disjoint as w is freely braided. Consequently, F_w is surjective, and the conclusion follows. □

The above corollary translates via Remark 3.1.3 into an equivalent statement about commutation classes of reduced expressions for w.

4.3 Commutation graphs.

Let $w \in W$. The *commutation graph* of w, $G(w)$, is defined to be the graph with vertex set $\mathcal{C}(w)$ and with edge set consisting of all pairs $\{C, C'\}$ of commutation classes with the following property: there exist representatives $\mathbf{r} \in C$ and $\mathbf{r}' \in C'$ such that $\mathbf{r}$ and $\mathbf{r}'$ differ only by a single long braid move.

We define the *parity* of each commutation class $C \in \mathcal{C}(w)$ to be the number $(-1)^{M(C)}$, where $M(C) = \sum_{T \in \mathcal{I}(w)} F_w(C)(T)$. Note that the parity of a commutation class depends on the map F_w, which in turn depends on the relation $\preceq$.

Proposition 4.3.1. *Let $w \in W$. Then the graph $G(w)$ is bipartite.*

Proof. Let $\{C, C'\}$ be an edge of $G(w)$. By the definition of edge and Remark 2.2.2, there is precisely one $T \in \mathcal{I}(w)$ such that $F_w(C)(T) \neq F_w(C')(T)$. Hence, the classes C and C' have unequal parity. □

Remark 4.3.2. Let Γ be the Coxeter graph of (W, S). If Γ itself is bipartite, then Proposition 4.3.1 can be established using a simple argument counting the number of occurrences of each of the two types of generators in a reduced expression for w.

5. Type A

For the remainder of this paper, we shall assume that (W, S) is of type A. It will be convenient to work with a particular realization of (W, S), under which W is the group $\mathcal{S}_n$ of permutations of $\{1, \ldots, n\}$, and S is the set of simple transpositions $s_1 = (1, 2), \ldots, s_{n-1} = (n-1, n)$. Let $\varepsilon_1, \ldots, \varepsilon_n$ be the standard basis of $\mathbb{R}^n$, and let $\langle \cdot, \cdot \rangle$ be the inner product on $\mathbb{R}^n$ satisfying $\langle \varepsilon_i, \varepsilon_j \rangle = \delta_{ij}$. Allow $\mathcal{S}_n$ to act on $\mathbb{R}^n$ by permuting the indices of $\varepsilon_1, \ldots, \varepsilon_n$, and set $\alpha_i = \varepsilon_i - \varepsilon_{i+1}$ $(i = 1, \ldots, n-1)$. Then the vectors $\alpha_1, \ldots, \alpha_{n-1}$ may be regarded as the simple roots corresponding to $s_1, \ldots, s_{n-1}$, and the form $\frac{1}{2}\langle \cdot, \cdot \rangle$ may be regarded as the Coxeter form B. With these identifications, we have $\Phi = \{\varepsilon_i - \varepsilon_j : i \neq j\}$ and $\Phi^+ = \{\varepsilon_i - \varepsilon_j : i < j\}$.

Let $w \in \mathcal{S}_n$. By the "1-line notation" for w, we shall mean the string of numbers $w(1) \cdots w(n)$. Let (i,j) be a (not necessarily simple) transposition in $\mathcal{S}_n$, with $i < j$. Note that the 1-line notation for the product $w(i,j)$ can be obtained from that for w by interchanging the numbers in positions i and j. Further, we have $\ell(w(i,j)) < \ell(w) \Leftrightarrow w(i) > w(j) \Leftrightarrow \varepsilon_i - \varepsilon_j \in \Phi(w)$.

The set of inversion triples of any permutation $w \in \mathcal{S}_n$ can be read off of the 1-line notation for w. Specifically, a triple $\{\varepsilon_i - \varepsilon_j, \varepsilon_i - \varepsilon_k, \varepsilon_j - \varepsilon_k\}$ of positive roots is an inversion triple of w if and only if $w(i) > w(j) > w(k)$.

5.1 Contractibility in type A.

The following proposition stands in contrast to Example 2.2.3.

Proposition 5.1.1. *Let $w \in \mathcal{S}_n$. Then every inversion triple of w is contractible.*

Proof. We argue by induction on $m = \ell(w)$. The assertion is vacuously true if $0 \leq m < 3$. Suppose that $m \geq 3$ and that w has at least one inversion triple. Let $T = \{\varepsilon_i - \varepsilon_j, \varepsilon_i - \varepsilon_k, \varepsilon_j - \varepsilon_k\}$ be an inversion triple of w.

Consider the situation where there is a positive integer $l < n$ such that $\{l, l+1\} \cap \{i, j, k\} = \emptyset$ and $w(l) > w(l+1)$. Then $\ell(ws_l) < \ell(w)$, and T is an inversion triple of ws_l. By the inductive hypothesis, T is contractible relative to ws_l. If $(r_1, \ldots, r_{m-1})$ is any root sequence for ws_l in which the elements of T appear consecutively, then $(\alpha_l, s_l(r_1), \ldots, s_l(r_{m-1}))$ is a root sequence for w in which the elements of T appear consecutively, since s_l fixes the elements of T pointwise.

We may therefore assume that $w(l) < w(l+1)$ for every positive integer $l < n$ such that $\{l, l+1\} \cap \{i, j, k\} = \emptyset$.

Suppose that $i < j - 1$. Since $w(i) > w(j)$, the previous paragraph implies that either $w(i) > w(i+1)$ or $w(j-1) > w(j)$. Suppose for the moment that $w(i) > w(i+1)$. Then $\ell(ws_i) < \ell(w)$, and the set $T' = \{\varepsilon_{i+1} - \varepsilon_j, \varepsilon_{i+1} - \varepsilon_k, \varepsilon_j - \varepsilon_k\}$ is an inversion triple of ws_i. By the inductive hypothesis, T' is contractible relative to ws_i. If $(r_1, \ldots, r_{m-1})$ is any root sequence for ws_i in which the elements of T' appear consecutively, then $(\alpha_i, s_i(r_1), \ldots, s_i(r_{m-1}))$ is a root sequence for w in which the elements of T appear consecutively.

Similar reasoning handles the case where $w(j-1) > w(j)$. The induction step is therefore proved for $i < j - 1$.

If $j < k - 1$, then one argues in a similar way. It remains to deal with the case where $j = i+1$ and $k = j+1$. Here, the element w has a reduced expression of the form $(\ldots, s_i, s_{i+1}, s_i)$, and it is clear that the first three entries of the corresponding root sequence belong to T. $\square$

Let $w \in \mathcal{S}_n$ and let $v \in \mathcal{S}_k$, where $k \leq n$. Suppose that there do not exist $i_1, \ldots, i_k \in \{1, \ldots, n\}$ with $i_1 < \cdots < i_k$ such that the numbers $w(i_1), \ldots, w(i_k)$ are in the same relative order as $v(1), \ldots, v(k)$. Then one says that w "avoids the pattern" $v(1) \cdots v(k)$. It can be deduced from Proposition 5.1.1 and the discussion that precedes it that a given permutation is freely braided if and only if it avoids each of the four patterns $3421, 4231, 4312$ and 4321.

5.2 Converse to Corollary 4.2.4 in type A.

Let $w \in \mathcal{S}_n$. Let $\mathcal{T}$ be a nonempty subset of $\mathcal{I}(w)$, the set of all (contractible) inversion triples of w. We say that the map F_w of §4.1 *separates* $\mathcal{T}$ if every map from $\mathcal{T}$ to $\{0, 1\}$ is the restriction of some element of $F_w(\mathcal{C}(w))$. Clearly, if F_w fails to separate a nonempty subset of $\mathcal{I}(w)$, then F_w is not surjective.

The following theorem is a partial converse to Corollary 4.2.4. Its proof relies on Proposition 5.1.1, which, as was observed in Example 2.2.3, does not hold for all simply laced Coxeter systems.

Theorem 5.2.1. *Suppose that $w \in \mathcal{S}_n$ has $2^{N(w)}$ inversion triples. Then w is freely braided.*

Proof. Assume toward a contradiction that there is a non-freely-braided $w \in \mathcal{S}_n$ with $2^{N(w)}$ inversion triples (all of which must be contractible by Proposition 5.1.1). In other words, w is non-freely-braided and F_w is surjective. Let $\varepsilon_i - \varepsilon_j$ be an element of $\Phi(w)$ that belongs to at least two inversion triples of w, and assume that $j - i$ is as large as possible relative to this property. Let T and T' be distinct inversion triples of w containing $\varepsilon_i - \varepsilon_j$. By symmetry, there are six cases to consider.

Case 1. $\varepsilon_j - \varepsilon_k \in T$ and $\varepsilon_j - \varepsilon_l \in T'$, where $i < j < k < l$.
Observe that $\varepsilon_k - \varepsilon_l \notin \Phi(w)$. Otherwise, the root $\varepsilon_i - \varepsilon_l$ would belong to at least two inversion triples of w, namely $\{\varepsilon_i - \varepsilon_k, \varepsilon_i - \varepsilon_l, \varepsilon_k - \varepsilon_l\}$ and $\{\varepsilon_i - \varepsilon_j, \varepsilon_i - \varepsilon_l, \varepsilon_j - \varepsilon_l\}$, contradicting our choice of $\varepsilon_i - \varepsilon_j$.

By Proposition 5.1.1 and Remark 2.2.2, there is a root sequence for w of the form

$$(\ldots,\varepsilon_i-\varepsilon_j,\varepsilon_i-\varepsilon_k,\varepsilon_j-\varepsilon_k,\ldots,\varepsilon_i-\varepsilon_l,\ldots,\varepsilon_j-\varepsilon_l,\ldots) \text{ or}$$

$$(\ldots,\varepsilon_j-\varepsilon_l,\ldots,\varepsilon_i-\varepsilon_l,\ldots,\varepsilon_i-\varepsilon_j,\varepsilon_i-\varepsilon_k,\varepsilon_j-\varepsilon_k,\ldots).$$

Since $\varepsilon_k-\varepsilon_l \notin \Phi(w)$, and since $\varepsilon_i-\varepsilon_k$ and $\varepsilon_i-\varepsilon_l$ are not orthogonal, these last two roots maintain their positions relative to one another in every root sequence for w. If $\varepsilon_i-\varepsilon_k$ lies always to the left of $\varepsilon_i-\varepsilon_l$, then $\varepsilon_i-\varepsilon_j$ cannot at the same time lie to the left of $\varepsilon_i-\varepsilon_k$ and to the right of $\varepsilon_i-\varepsilon_l$. A similar statement can be made if $\varepsilon_i-\varepsilon_k$ lies always to the right of $\varepsilon_i-\varepsilon_l$. It follows that F_w does not separate $\{T,T'\}$.

Case 2. $\varepsilon_k-\varepsilon_i \in T$ and $\varepsilon_l-\varepsilon_i \in T'$, where $k<l<i<j$.
One argues as in Case 1.

Case 3. $\varepsilon_k-\varepsilon_i \in T$ and $\varepsilon_j-\varepsilon_l \in T'$, where $k<i<j<l$.
Note that $\varepsilon_k-\varepsilon_j \in T \subseteq \Phi(w)$. Since both $\varepsilon_k-\varepsilon_i$ and $\varepsilon_j-\varepsilon_l$ also lie in $\Phi(w)$, the root $\varepsilon_k-\varepsilon_j$ belongs to at least two inversion triples of w, namely $\{\varepsilon_k-\varepsilon_j,\varepsilon_k-\varepsilon_l,\varepsilon_j-\varepsilon_l\}$ and $\{\varepsilon_k-\varepsilon_i,\varepsilon_k-\varepsilon_j,\varepsilon_i-\varepsilon_j\}$. This contradicts our choice of $\varepsilon_i-\varepsilon_j$.

Case 4. $\varepsilon_i-\varepsilon_k \in T$ and $\varepsilon_j-\varepsilon_l \in T'$, where $i<k<j<l$.
We have $\varepsilon_k-\varepsilon_j \in T \subseteq \Phi(w)$. Hence, $\varepsilon_k-\varepsilon_l=(\varepsilon_k-\varepsilon_j)+(\varepsilon_j-\varepsilon_l) \in \Phi(w)$. It follows that $\varepsilon_i-\varepsilon_l$ belongs to at least two inversion triples of w, namely T' and $\{\varepsilon_i-\varepsilon_k,\varepsilon_i-\varepsilon_l,\varepsilon_k-\varepsilon_l\}$, contradicting our choice of $\varepsilon_i-\varepsilon_j$.

Case 5. $\varepsilon_k-\varepsilon_i \in T$ and $\varepsilon_l-\varepsilon_j \in T'$, where $k<i<l<j$.
One argues as in Case 4.

Case 6. $\varepsilon_i-\varepsilon_k \in T$ and $\varepsilon_l-\varepsilon_j \in T'$, where $i<k<l<j$.
Here, we have $T=\{\varepsilon_i-\varepsilon_k,\varepsilon_i-\varepsilon_j,\varepsilon_k-\varepsilon_j\}$ and $T'=\{\varepsilon_i-\varepsilon_l,\varepsilon_i-\varepsilon_j,\varepsilon_l-\varepsilon_j\}$. Suppose first that $\varepsilon_k-\varepsilon_l \notin \Phi(w)$. Then the roots $\varepsilon_i-\varepsilon_k$ and $\varepsilon_i-\varepsilon_l$, which are not orthogonal, maintain their positions relative to one other in every root sequence for w. It follows that F_w does not separate $\{T,T'\}$.

Finally, suppose that $\varepsilon_k-\varepsilon_l \in \Phi(w)$. Then the sets $T''=\{\varepsilon_k-\varepsilon_l,\varepsilon_k-\varepsilon_j,\varepsilon_l-\varepsilon_j\}$ and $T'''=\{\varepsilon_i-\varepsilon_k,\varepsilon_i-\varepsilon_l,\varepsilon_k-\varepsilon_l\}$ are both inversion triples of w. We claim that F_w does not separate $\{T,T'',T'''\}$. To see this, observe that if C is a commutation class of w relative to which $\varepsilon_i-\varepsilon_k$ lies to the left of $\varepsilon_k-\varepsilon_j$ (this determines $F_w(C)(T)$) and $\varepsilon_k-\varepsilon_j$ lies to the left of $\varepsilon_k-\varepsilon_l$ (this determines $F_w(C)(T'')$), then $\varepsilon_i-\varepsilon_k$ lies to the left of $\varepsilon_k-\varepsilon_l$ (so that $F_w(C)(T''')$ is also determined). $\square$

Note added in proof. The authors have found a proof that the converse of Corollary 4.2.4 holds for arbitrary simply laced Coxeter systems.

References

[1] R. Bédard, *On commutation classes of reduced words in Weyl groups*, Europ. J. Combinatorics **20** (1999), 483–505.

[2] A. Björner, *Orderings of Coxeter groups*, Contemporary Math., vol. 34, 1984, pp. 175–195.

[3] N. Bourbaki, *Groupes et algèbres de Lie, Chapitres 4, 5, et 6*, Masson, Paris, 1981.

[4] S. Elnitsky, *Rhombic tilings of polygons and classes of reduced words in Coxeter groups*, J. Combin. Theory Ser. A **77** (1997), 193–221.

[5] C.K. Fan and J.R. Stembridge, *Nilpotent orbits and commutative elements*, J. Algebra **196** (1997), 490–498.

[6] J.E. Humphreys, *Reflection Groups and Coxeter Groups*, Cambridge University Press, Cambridge, 1990.

[7] H. Matsumoto, *Générateurs et relations des groupes de Weyl généralisés*, C. R. Acad. Sci. Paris **258** (1964), 3419–3422.

[8] J.R. Stembridge, *On the fully commutative elements of Coxeter groups*, J. Algebraic Combin. **5** (1996), 353–385.

[9] J. Tits, *Le problème des mots dans les groupes de Coxeter*, Ist. Naz. Alta Mat. (1968), *Sympos. Math.*, vol. 1, Academic Press, London, 1969, pp. 175–185.

[10] G.X. Viennot, *Heaps of pieces, I: basic definitions and combinatorial lemmas*, Combinatoire Énumérative, ed. G. Labelle and P. Leroux, Springer-Verlag, Berlin, 1986, pp. 321–350.