

Boundary braids

MICHAEL DOUGHERTY
JON MCCAMMOND
STEFAN WITZEL

The n -strand braid group can be defined as the fundamental group of the configuration space of n unlabeled points in a closed disk based at a configuration where all n points lie in the boundary of the disk. Using this definition, the subset of braids that have a representative where a specified subset of these points remain pointwise fixed forms a subgroup isomorphic to a braid group with fewer strands. We generalize this phenomenon by introducing the notion of boundary braids. A boundary braid is a braid that has a representative where some specified subset of the points remains in the boundary cycle of the disk. Although boundary braids merely form a subgroupoid rather than a subgroup, they play an interesting geometric role in the piecewise Euclidean dual braid complex defined by Tom Brady and the second author. We prove several theorems in this setting, including the fact that the subcomplex of the dual braid complex determined by a specified set of boundary braids metrically splits as the direct metric product of a Euclidean polyhedron and a dual braid complex of smaller rank.

20F36; 20F65

Braids and braid groups play an important role throughout mathematics, in part because of the multiple ways in which they can be described. In this article, we view the n -strand braid group BRAID_n as the fundamental group of the unordered configuration space of n distinct points in the closed unit disk \mathbb{D} , based at an initial configuration P where all n points lie in the boundary of \mathbb{D} . Requiring the points indexed by $B \subseteq \{1, \dots, n\}$ to remain fixed defines a parabolic subgroup $\text{FIX}_n(B)$ which is isomorphic to $\text{BRAID}_{n-|B|}$.

We introduce an extension of this idea. A (B, \cdot) -boundary braid is a braid that has a representative where the points indexed by B remain in the boundary of the disk but need not be fixed (see Definition 2.13). Our first main result is that the subgroup $\text{FIX}_n(B)$ has a canonical complement $\text{MOVE}_n(B, \cdot)$ in the set $\text{BRAID}_n(B, \cdot)$ of (B, \cdot) -boundary braids that gives rise to a unique decomposition (see Section 13):

Theorem A Let $B \subseteq \{1, \dots, n\}$ and let β be a (B, \cdot) -boundary braid in BRAID_n . Then there are unique braids $\text{FIX}^B(\beta)$ in $\text{FIX}_n(B)$ and $\text{MOVE}^B(\beta)$ in $\text{MOVE}_n(B, \cdot)$ such that

$$\beta = \text{FIX}^B(\beta) \text{MOVE}^B(\beta).$$

We call the elements of $\text{FIX}_n(B)$ *fix braids* and the elements of $\text{MOVE}_n(B, \cdot)$ *move braids*.

Associated to the braid group BRAID_n is the *dual braid complex*, introduced by Tom Brady [5] and denoted by $\text{CPLX}(\text{BRAID}_n)$. It is a contractible simplicial complex on which BRAID_n acts freely and cocompactly. Brady and the second author equipped $\text{CPLX}(\text{BRAID}_n)$ with a piecewise Euclidean *orthoscheme metric* and conjectured that it is $\text{CAT}(0)$ with respect to this metric [6]. They verified the conjecture for $n < 6$ and Haettel, Kielak and Schwer proved it for $n = 6$ [11]. We are interested in boundary braids as an approach to proving the conjecture. The sets of boundary braids, fix braids and move braids have induced subcomplexes in $\text{CPLX}(\text{BRAID}_n)$ with the following metric decomposition:

Theorem B Let $B \subseteq \{1, \dots, n\}$. The complex of (B, \cdot) -boundary braids decomposes as a metric direct product

$$\text{CPLX}(\text{BRAID}_n(B, \cdot)) \cong \text{CPLX}(\text{FIX}_n(B)) \times \text{CPLX}(\text{MOVE}_n(B, \cdot)).$$

The complex $\text{CPLX}(\text{MOVE}_n(B, \cdot))$ is \mathbb{R} times a Euclidean simplex and therefore $\text{CAT}(0)$. In particular, $\text{CPLX}(\text{BRAID}_n(B, \cdot))$ is $\text{CAT}(0)$ if and only if the smaller dual braid complex $\text{CPLX}(\text{FIX}_n(B)) \cong \text{CPLX}(\text{BRAID}_{n-|B|})$ is $\text{CAT}(0)$.

As part of our proof for Theorem B, we introduce a new type of configuration space for directed graphs and the broader setting of Δ -complexes. We refer to these as *orthoscheme configuration spaces* and explore their geometry for the case of oriented n -cycles. The other key element for proving Theorem B is a combinatorial study of noncrossing partitions associated to boundary braids.

Because the points indexed by B do not necessarily return to their original positions, either pointwise or as a set, boundary braids form a subgroupoid rather than a subgroup. More precisely, if we refer to a boundary braid where the points indexed by B move in the boundary to end at points indexed by B' as a (B, B') -boundary braid, then we see that a (B, B') -boundary braid can be composed with a (B', B'') -boundary braid to

produce a (B, B'') -boundary braid. The groupoid $\text{BRAID}_n(\cdot, \cdot)$ of boundary braids has subgroupoids $\text{FIX}_n(\cdot)$ and $\text{MOVE}_n(\cdot, \cdot)$ consisting of fix braids and move braids, respectively. We prove the following algebraic result (see Section 14):

Theorem C *The groupoid of boundary braids decomposes as a semidirect product*

$$\text{BRAID}_n(\cdot, \cdot) \cong \text{FIX}_n(\cdot) \rtimes \text{MOVE}_n(\cdot, \cdot).$$

The article is organized as follows. The first part, Sections 1–4, develops standard material about braid groups and their dual Garside structure in a way that suits our later applications. The second part, Sections 5–8, is concerned with complexes of ordered simplices. Specifically we show how to equip them with an orthoscheme metric and that a combinatorial direct product gives rise to a metric direct product. The final part, Sections 9–14, contains our work on boundary braids and the proofs of the main theorems.

Acknowledgements We thank the referee for their careful reading of the article and several helpful remarks.

Witzel was funded through the DFG project WI 4079/2. He also acknowledges the hospitality of UC Santa Barbara.

Part I Braids

Braids can be described in a variety of ways. In this part we establish the conventions used throughout the article, review basic facts about dual simple braids and the dual presentation for the braid group, and introduce the concept of a boundary braid.

In Section 1 we describe our specific view on braid groups, which is important for the definition of a boundary braid in Section 2. The correspondence between noncrossing partitions and dual simple braids is reviewed in Section 3. In Section 4 we study the parabolic subgroups in our context, they will support the fix part of a boundary braid.

1 Braid groups

In this article, braid groups are viewed as fundamental groups of certain configuration spaces.

Definition 1.1 (configuration spaces) Let X be a topological space, let n be a positive integer and let X^n denote the product of n copies of X whose elements are n -tuples $\vec{x} = (x_1, x_2, \dots, x_n)$ of elements $x_i \in X$. Alternatively, the elements of X^n can be thought of as functions from $[n]$ to X , where $[n]$ is the set $\{1, 2, \dots, n\}$. The *configuration space of n labeled points in X* is the subspace $\text{CONF}_n(X)$ of X^n of n -tuples with distinct entries, ie the subspace of injective functions. The *thick diagonal of X^n* is the subspace $\text{DIAG}_n(X) = \{(x_1, \dots, x_n) \mid x_i = x_j \text{ for some } i \neq j\}$ where this condition fails. Thus $\text{CONF}_n(X) = X^n - \text{DIAG}_n(X)$. The symmetric group acts on X^n by permuting coordinates and this action restricts to a free action on $\text{CONF}_n(X)$. The *configuration space of n unlabeled points in X* is the quotient space $\text{UCONF}_n(X) = (X^n - \text{DIAG}_n(X))/\text{SYM}_n$. Since the quotient map sends the n -tuple (x_1, \dots, x_n) to the n -element set $\{x_1, \dots, x_n\}$, we write $\text{SET}: \text{CONF}_n(X) \rightarrow \text{UCONF}_n(X)$ for this natural quotient map.

Since the topology of a configuration space only depends on the topology of the original space, the following lemma is immediate:

Lemma 1.2 (homeomorphisms) *A homeomorphism $X \rightarrow Y$ induces a homeomorphism $h: \text{UCONF}_n(X) \rightarrow \text{UCONF}_n(Y)$. In particular, for any choice of base-point $*$ in $\text{UCONF}_n(X)$, there is an induced isomorphism $\pi_1(\text{UCONF}_n(X), *) \cong \pi_1(\text{UCONF}_n(Y), h(*))$.*

Example 1.3 (configuration spaces) When X is the unit circle and $n = 2$, the space X^2 is a torus, $\text{DIAG}_2(X)$ is a $(1, 1)$ -curve on the torus, its complement $\text{CONF}_2(X)$ is homeomorphic to the interior of an annulus and the quotient $\text{UCONF}_2(X)$ is homeomorphic to the interior of a Möbius band.

Definition 1.4 (braids in \mathbb{C}) Let \mathbb{C} be the complex numbers with its usual topology and let $\vec{z} = (z_1, z_2, \dots, z_n)$ denote a point in \mathbb{C}^n . The thick diagonal of \mathbb{C}^n is a union of hyperplanes H_{ij} , with $i < j \in [n]$, called the *braid arrangement*, where H_{ij} is the hyperplane defined by the equation $z_i = z_j$. The configuration space $\text{CONF}_n(\mathbb{C})$ is the complement of the braid arrangement and its fundamental group is called the *n -strand pure braid group*. The *n -strand braid group* is the fundamental group of the quotient configuration space $\text{UCONF}_n(\mathbb{C}) = \text{CONF}_n(\mathbb{C})/\text{SYM}_n$ of n unlabeled points. In symbols,

$$\text{PBRAID}_n = \pi_1(\text{CONF}_n(\mathbb{C}), \vec{z}) \quad \text{and} \quad \text{BRAID}_n = \pi_1(\text{UCONF}_n(\mathbb{C}), Z),$$

where \vec{z} is some specified basepoint in $\text{CONF}_n(\mathbb{C})$ and $Z = \text{SET}(\vec{z})$ is the corresponding basepoint in $\text{UCONF}_n(\mathbb{C})$.

Remark 1.5 (short exact sequence) The quotient map SET is a covering map, so the induced map $\text{SET}_*: \text{PBRAID}_n \rightarrow \text{BRAID}_n$ on fundamental groups is injective. In fact, $\text{CONF}_n(\mathbb{C})$ is a regular cover of $\text{UCONF}_n(\mathbb{C})$, so the subgroup $\text{SET}_*(\text{PBRAID}_n) \subset \text{BRAID}_n$ is a normal subgroup and the quotient group $\text{BRAID}_n/\text{SET}_*(\text{PBRAID}_n)$ is isomorphic to the group SYM_n of covering transformations. The quotient map sends each braid to the permutation it induces on the n -element set used as the basepoint of BRAID_n , a map we define more precisely in the next section. We call this map PERM . These maps form a short exact sequence

$$(1-1) \quad \text{PBRAID}_n \xrightarrow{\text{SET}_*} \text{BRAID}_n \xrightarrow{\text{PERM}} \text{SYM}_n.$$

Example 1.6 ($n \leq 2$) When $n = 1$ the spaces $\text{UCONF}_1(\mathbb{C})$, $\text{CONF}_1(\mathbb{C})$ and \mathbb{C} are equal and contractible, and all three groups in (1-1) are trivial. When $n = 2$ the space $\text{CONF}_2(\mathbb{C})$ is \mathbb{C}^2 minus a copy of \mathbb{C}^1 , which retracts first to $\mathbb{C}^1 - \mathbb{C}^0$ and then to the circle \mathbb{S}^1 of unit-length complex numbers. The quotient space $\text{UCONF}_2(\mathbb{C})$ also deformation retracts to \mathbb{S}^1 and the map from $\text{CONF}_2(\mathbb{C})$ to $\text{UCONF}_2(\mathbb{C})$ corresponds to the map from \mathbb{S}^1 to itself sending z to z^2 . In particular, $\text{PBRAID}_2 \cong \text{BRAID}_2 \cong \mathbb{Z}$ and SET_* is the map that multiplies by 2 with quotient $\mathbb{Z}/2\mathbb{Z} \cong \text{SYM}_2$.

Convention 1.7 ($n > 2$) For the remainder of the article, we assume that the integer n is greater than 2, unless we explicitly state otherwise.

Let $\mathbb{D} \subset \mathbb{C}$ be the closed unit disk centered at the origin. Restricting to configurations of points that remain in \mathbb{D} does not change the fundamental group of the configuration space.

Proposition 1.8 (braids in \mathbb{D}) *The configuration space $\text{UCONF}_n(\mathbb{C})$ deformation retracts to the subspace $\text{UCONF}_n(\mathbb{D})$, so for any choice of basepoint Z in the subspace, $\pi_1(\text{UCONF}_n(\mathbb{D}), Z) = \pi_1(\text{UCONF}_n(\mathbb{C}), Z) = \text{BRAID}_n$.*

Proof Let $m(\vec{z}) = \max\{1, |z_1|, \dots, |z_n|\}$ for each $\vec{z} = (z_1, \dots, z_n) \in \mathbb{C}^n$ and note that m defines a continuous map from $\text{CONF}_n(\mathbb{C})$ to $\mathbb{R}_{\geq 1}$. The straight-line homotopy from the identity map on $\text{CONF}_n(\mathbb{C})$ to the map that sends \vec{z} to $(1/m(\vec{z}))\vec{z}$ is a deformation retraction from $\text{CONF}_n(\mathbb{C})$ to $\text{CONF}_n(\mathbb{D})$ and since $m(\vec{z})$ only depends on the entries of \vec{z} and not their order, this deformation retraction descends to one from $\text{UCONF}_n(\mathbb{C})$ to $\text{UCONF}_n(\mathbb{D})$. □

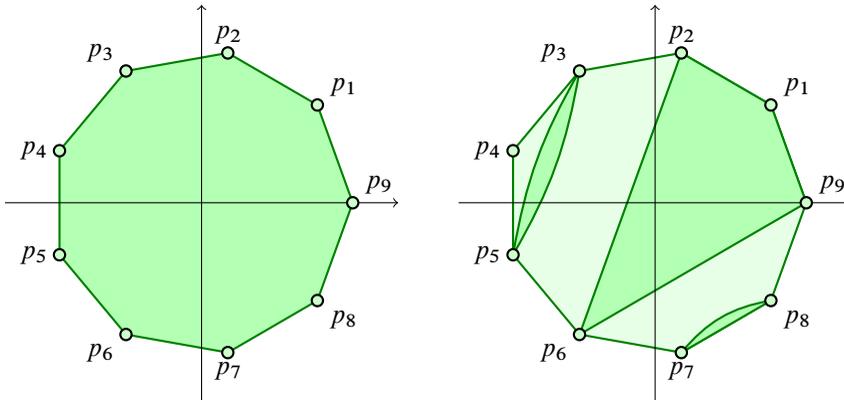


Figure 1: Left: the standard basepoint $P = \{p_1, p_2, \dots, p_9\}$ and the standard disk D for $n = 9$. Right: the standard subdisks D_A for A equal to $\{1, 2, 6, 9\}$, $\{3, 5\}$ and $\{7, 8\}$.

We have a preferred choice of basepoint and disk for BRAID_n .

Definition 1.9 (basepoints and disks) Let $\zeta = e^{2\pi i/n} \in \mathbb{C}$ be the standard primitive n^{th} root of unity and let p_i be the point ζ^i for all $i \in \mathbb{Z}$. Since $\zeta^n = 1$, the subscript i should be interpreted as an integer representing the equivalence class $i + n\mathbb{Z} \in \mathbb{Z}/n\mathbb{Z}$. In particular, we consider $p_{i-n} = p_i = p_{i+n} = p_{i+2n}$ without further comment. The *standard basepoint* for PBRAID_n is the n -tuple $\vec{p} = (p_1, p_2, \dots, p_n)$ and the *standard basepoint* for BRAID_n is the n -element set $P = \text{SET}(\vec{p}) = \{p_1, p_2, \dots, p_n\}$ of all n^{th} roots of unity. Let D be the convex hull of the points in P . Our standing assumption of $n > 2$ means that D is homeomorphic to the disk \mathbb{D} . We call D the *standard disk* for BRAID_n . See Figure 1.

The following result combines Lemma 1.2 and Proposition 1.8:

Corollary 1.10 (braids in D) A homeomorphism $\mathbb{D} \rightarrow D$ induces a homeomorphism of configuration spaces $h: \text{UCONF}_n(\mathbb{D}) \rightarrow \text{UCONF}_n(D)$. In particular, for any choice of basepoint Z in $\text{UCONF}_n(\mathbb{D})$, there is an induced isomorphism

$$\pi_1(\text{UCONF}_n(\mathbb{D}), Z) \cong \pi_1(\text{UCONF}_n(D), h(Z)) = \text{BRAID}_n.$$

Remark 1.11 (points in ∂D) When BRAID_n is viewed as the mapping class group of an n -times punctured disk, the punctures are not allowed to move into the boundary of the disk since doing so would alter the topological type of the punctured space.

When BRAID_n is viewed as the fundamental group of a configuration space of points in a closed disk, points are allowed in the boundary and we make extensive use of this extra flexibility.

Remark 1.12 (braid groups) By Corollary 1.10, the braid group BRAID_n is isomorphic to $\pi_1(\text{UCONF}_n(D), P)$, the fundamental group of the configuration space of n unlabeled points in the standard disk D based at the standard basepoint P . In the remainder of the article, we use the notation BRAID_n to refer to the specific group $\pi_1(\text{UCONF}_n(D), P)$.

2 Individual braids

This section establishes our conventions for describing individual braids and we introduce the concept of a boundary braid.

Definition 2.1 (representatives) Each braid $\alpha \in \text{BRAID}_n$ is a basepoint-preserving homotopy class of a path $f: [0, 1] \rightarrow \text{UCONF}_n(D, P)$ that describes a loop based at the standard basepoint P . We write $\alpha = [f]$ and say that the loop f represents α . We use Greek letters such as α , β and δ for braids and Roman letters such as f , g and h for their representatives.

Vertical drawings of braids in \mathbb{R}^3 typically have the $t = 0$ start at the top and the $t = 1$ end at the bottom. See Definition 2.3 for the details. As a mnemonic, we use superscripts for information about the start of a braid or a path and subscripts for information about its end.

Definition 2.2 (strands) Let $\alpha \in \text{BRAID}_n$ be a braid with representative f . A *strand of f* is a path in D that follows what happens to one of the vertices in P . There are two natural ways to name strands: by where they start and by where they end. The *strand that starts at p_i* is the path $f^i: [0, 1] \rightarrow D$ defined by the composition $f^i = \text{PROJ}_i \circ \tilde{f}^{\vec{p}}$, where the map $\tilde{f}^{\vec{p}}$ is the unique lift of the path f through the covering map $\text{SET}: \text{CONF}_n(D) \rightarrow \text{UCONF}_n(D)$ such that the lifted path starts at \vec{p} , ie $\tilde{f}^{\vec{p}}(0) = \vec{p}$, and $\text{PROJ}_i: \text{CONF}_n(D) \rightarrow D$ is projection onto the i^{th} coordinate. Similarly, the *strand that ends at p_j* is the path $f_j: [0, 1] \rightarrow D$ defined by the composition $f_j = \text{PROJ}_j \circ \tilde{f}_{\vec{p}}$ where $\tilde{f}_{\vec{p}}$ is the unique lift of the path f through the covering map $\text{SET}: \text{CONF}_n(D) \rightarrow \text{UCONF}_n(D)$ that ends at \vec{p} , ie $\tilde{f}_{\vec{p}}(1) = \vec{p}$. When the strand of f that starts at p_i ends at p_j , the path f^i is the same as the path f_j . We write f^i , f_j or f_j^i for this path

and we call it the (i, \cdot) -strand, the (\cdot, j) -strand or the (i, j) -strand of f depending on the information specified.

A braid representative is drawn by superimposing the graphs of its strands.

Definition 2.3 (drawings) Let $\alpha \in \text{BRAID}_n$ be a braid with representative f . A drawing of f is formed by superimposing the graphs of its strands inside the polygonal prism $[0, 1] \times D$. There are two distinct conventional embeddings of this prism into \mathbb{R}^3 . The complex plane containing D is identified with either the first two or the last two coordinates of \mathbb{R}^3 and the remaining coordinate indicates the value $t \in [0, 1]$ with the t -dependence arranged so that the $t = 0$ start of f is on the left or at the top and the $t = 1$ end of f is on the right or at the bottom. In the left-to-right orientation, for each $j \in [n]$, for each $t_0 \in [0, 1]$ and for each point $f_j(t_0) = z_0 = x_0 + iy_0 \in D \subset \mathbb{C}$ on the (\cdot, j) -strand, we draw the point $(t_0, x_0, y_0) \in \mathbb{R}^3$. In the top-to-bottom orientation, the same point on the (\cdot, j) -strand is drawn at $(x_0, y_0, 1 - t_0) \in \mathbb{R}^3$.

Definition 2.4 (multiplication) Let α_1 and α_2 be braids in BRAID_n with representatives f_1 and f_2 . The product $\alpha_1 \cdot \alpha_2$ is defined to be $[f_1 \cdot f_2]$, where $f_1 \cdot f_2$ is the concatenation of f_1 and f_2 . In the drawing of $f_1 \cdot f_2$ the drawing of f_1 is on the top or left of the drawing of f_2 , which is on the bottom or right. See Figure 2.

Definition 2.5 (permutations) A permutation of the set $[n]$ is a bijective correspondence between a left/top copy of $[n]$ and a right/bottom copy of $[n]$. Permutations are compactly described in disjoint cycle notation. A cycle such as $(1\ 2\ 3)$, for example, means that 1 on the left corresponds to 2 on the right, 2 on the left corresponds to 3 on the right, and 3 on the left corresponds to 1 on the right. Multiplication of permutations is performed by concatenating the correspondences left-to-right or top-to-bottom. The permutation τ of $[n]$ acts on $[n]$ from either the left or the right by following the correspondence: if i on the left corresponds to j on the right, then $i \cdot \tau = j$ and $i = \tau \cdot j$.

Definition 2.6 (permutation of a braid) The permutation of a braid α is the bijective correspondence of $[n]$ under which i on the left corresponds to j on the right if α has an (i, j) -strand. Note that the function $\text{PERM}(\alpha)$ only depends on the braid α and not on the representative f . The direction of the bijection $\text{PERM}(\alpha)$ is defined so that it is compatible with function composition, ie so that $\text{PERM}(\alpha_1 \cdot \alpha_2) = \text{PERM}(\alpha_1) \circ \text{PERM}(\alpha_2)$.

Information about how a braid permutes its strands can be used to distinguish different types of braids.

Definition 2.7 ((i, j) -braids) Let $\alpha \in \text{BRAID}_n$ be a braid with permutation $g = \text{PERM}(\alpha)$. We say that α is an (i, j) -braid if the strand that starts at p_i ends at p_j . In other words, α is an (i, j) -braid if and only if $g(j) = i$. When α is an (i, j) -braid and β is a (j, k) -braid, $\gamma = \alpha \cdot \beta$ is an (i, k) -braid and the inverse of α is a (j, i) -braid.

For our applications, we make use of a generating set for the braid group which is built out of braids we call rotation braids. Rotation braids are defined using special subsets of our standard basepoint P and subspaces of our standard disk D .

Definition 2.8 (subsets of P) For each nonempty $A \subset [n]$ of size k , we define $P_A = \{p_i \mid i \in A\} \subset P$ to be the subset of points indexed by the numbers in A . In this notation our original set P is $P_{[n]}$. Using this notation we can extend the notion of an (i, j) -braid. Let A and B be two subsets of $[n]$ of the same size and let $\alpha \in \text{BRAID}_n$ be a braid with permutation $g = \text{PERM}(\alpha)$. We say that α is an (A, B) -braid if every strand that starts in P_A , ends in P_B , ie if and only if $g(B) = A$.

Definition 2.9 (subdisks of D) For all distinct $i, j \in [n]$ let the edge e_{ij} be the straight line segment connecting p_i and p_j . For $k > 2$, let D_A be $\text{CONV}(P_A)$, the convex hull of the points in P_A and note that D_A is a k -gon homeomorphic to \mathbb{D} . We call this the *standard subdisk for $A \subset [n]$* . In this notation, our original disk D is $D_{[n]}$. For $k = 2$ and $A = \{i, j\}$, we define D_A so that it is also a topological disk. Concretely, we take two copies of the path along the edge $e = e_{ij}$ from p_i to p_j and then bend one or both of these copies so that they become injective paths from p_i to p_j with disjoint interiors which together bound a bigon inside of D . Moreover, when the edge e lies in the boundary of D we require that one of the two paths does not move, so that e itself is part of the boundary of the bigon. See Figure 1. For $k = 1$, we define D_A to be the single point $p_i \in P_A$, but note that this subspace is not a subdisk. The bending of the edges to form the bigons are chosen to be slight enough so that for all A and $B \subseteq [n]$ the standard disks D_A and D_B intersect if and only if the convex hulls $\text{CONV}(P_A)$ and $\text{CONV}(P_B)$ intersect.

We view the boundaries of these subdisks as directed graphs.

Definition 2.10 (boundary edges) When A has more than 1 element, we view ∂D_A , the topological boundary of the subdisk D_A , as having the structure of a directed graph. The vertex set is P_A and for every vertex p_i in P_A there is a directed edge that starts at p_i , proceeds along ∂D_A in a counterclockwise direction with respect to

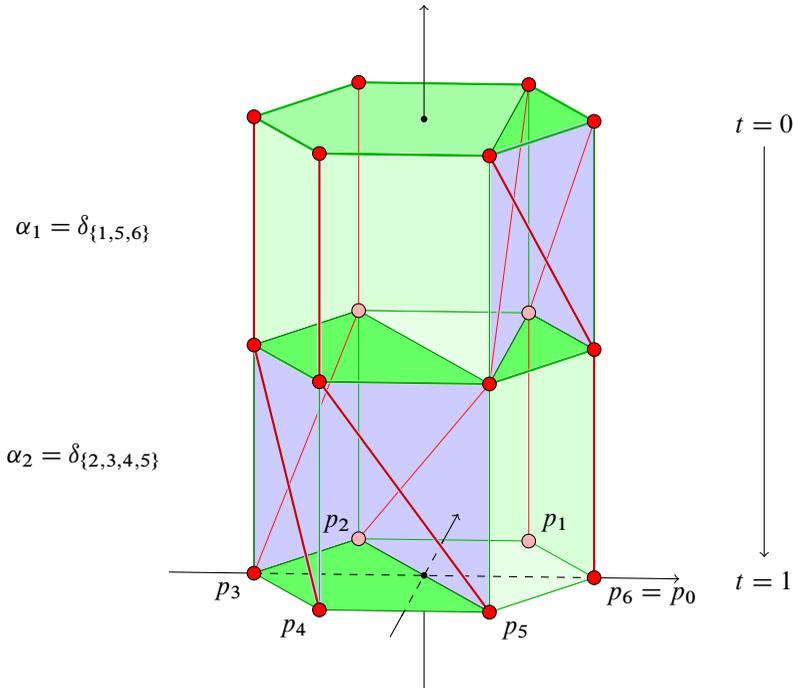


Figure 2: A drawing of $\alpha_1 \cdot \alpha_2$ where $\alpha_1 = \delta_{A_1}$ and $\alpha_2 = \delta_{A_2}$ are rotations with $A_1 = \{1, 5, 6\}$ and $A_2 = \{2, 3, 4, 5\}$.

the interior of D_A , and ends at the next vertex in P_A that it encounters. The edges of the graph ∂D_A are called the *boundary edges of D_A* . Note that edges and boundary edges are distinct concepts. An edge is unoriented and necessarily straight. A boundary edge is directed, it belongs to the boundary of a specific subdisk D_A and the path it describes might curve.

Definition 2.11 (rotation braids) For $A \subset [n]$ of size $k = |A| > 1$ we define an element $\delta_A \in \text{BRAID}_n$ that we call the *rotation braid of the vertices in P_A* . It is the braid represented by the path in $\text{UCONF}_n(D)$ that fixes the vertices in $P - P_A$ and where every vertex $p_i \in P_A$ travels in a counterclockwise direction along the oriented edge in the directed graph ∂D_A to the next vertex it encounters. If f is any representative of δ_A satisfying this description, we call f a *standard representative* of δ_A .

When $A = [n]$ we write δ instead of $\delta_{[n]}$ and when A has only a single element we let δ_A denote the identity element in BRAID_n . If $A = \{i, j\}$ and $e = e_{ij}$ is the edge connecting p_i and p_j , then we sometimes write δ_e to mean δ_A , the rotation of p_i

and p_j around the boundary of the bigon D_A . Note that if $A = \{i_1, i_2, \dots, i_k\} \subset [n]$ and $i_1 < i_2 < \dots < i_k$ is the natural linear order of its elements, then the bijection $\text{PERM}(\delta_A)$ is equal to the k -cycle (i_1, i_2, \dots, i_k) .

Example 2.12 (rotation braids) Figure 2 shows a drawing of the product of two rotation braids in BRAID_9 . The top braid α_1 is the rotation δ_{A_1} with $A_1 = \{1, 5, 6\}$ and $\text{PERM}(\alpha_1) = (1\ 5\ 6)$. The bottom braid α_2 is the rotation braid δ_{A_2} with $A_2 = \{2, 3, 4, 5\}$ and $\text{PERM}(\alpha_2) = (2\ 3\ 4\ 5)$. The product braid $\alpha = \alpha_1 \cdot \alpha_2$ is the rotation braid δ of all 6 vertices with $\text{PERM}(\alpha) = \text{PERM}(\alpha_1 \cdot \alpha_2) = \text{PERM}(\alpha_1) \circ \text{PERM}(\alpha_2) = (1\ 5\ 6)(2\ 3\ 4\ 5) = (1\ 2\ 3\ 4\ 5\ 6)$.

In this article, we focus on whether or not particular strands pass through the interior of the polygonal disk D . Those that remain in the boundary of D are called boundary strands.

Definition 2.13 (boundary braids) Let f be a representative of an (i, j) -braid $\alpha \in \text{BRAID}_n$. If the (i, j) -strand of f remains in the boundary ∂D , then it is a *boundary strand* of f and a *boundary parallel strand* of α . The linguistic shift from “boundary” to “boundary parallel” reflects the fact that while for many representatives of α , the (i, j) -strand will not remain in the boundary, it will always remain parallel to the boundary in a sense that can be made precise. When α has some representative in which its (i, j) -strand is a boundary strand, α is called a (i, j) -*boundary braid*. More generally, suppose α is a (B, C) -braid and there is a representative f of α such that every strand that starts in P_B is a boundary strand of f . We then call α a (B, C) -*boundary braid*.

Note that the definition of a (B, C) -boundary braid requires a single representative where all of these strands remain in ∂D . We will see in Section 11 that such a representative exists as soon as there are representatives which keep the (i, \cdot) -strand in the boundary for each $i \in B$.

Example 2.14 (boundary braids) In Figure 2 the rotation braid $\alpha_2 = \delta_{A_2} \in \text{BRAID}_6$ with $A_2 = \{2, 3, 4, 5\}$ is an (A_2, A_2) -braid but it is not an (A_2, A_2) -boundary braid since the $(5, 2)$ -strand passes through the interior of D . It is, however, a (B, C) -boundary braid with $B = \{1, 2, 3, 4, 6\}$ and $C = \{1, 3, 4, 5, 6\}$ since all five of the corresponding strands, ie the $(1, 1)$ -, $(2, 3)$ -, $(3, 4)$ -, $(4, 5)$ - and $(6, 6)$ -strands, remain in the boundary of D in its standard representative.

3 Dual simple braids

This section defines dual simple braids and the dual presentation of the braid group using the rotation braids from the previous section. We begin with the combinatorics of the noncrossing partition lattice. Recall that $P = P_{[n]} \subset \mathbb{C}$ denotes the set of n^{th} roots of unity.

Definition 3.1 (noncrossing partitions) A partition $\pi = \{A_1, \dots, A_k\}$ of the set $[n]$ is *noncrossing* when the convex hulls $\text{CONV}(P_{A_1}), \dots, \text{CONV}(P_{A_k})$ of the corresponding sets of points in P are pairwise disjoint. A partition is *irreducible* if it has exactly one block with more than one element. Since there is an obvious bijection between irreducible partitions and subsets of $[n]$ of size at least 2, we write π_A to indicate the irreducible partition whose unique nonsingleton block is A .

Definition 3.2 (noncrossing partition lattice) Let π and π' be noncrossing partitions of $[n]$. If each block of π is contained in some block of π' , then π is called a *refinement* of π' and we write $\pi \leq \pi'$. The set of all noncrossing partitions of $[n]$ under the refinement partial order has well-defined meets and joins and is called the *lattice of noncrossing partitions* NC_n .

The Hasse diagram for NC_4 is shown in Figure 3. The noncrossing partition lattice has a maximum partition with only one block and a minimum partition, also called the *discrete partition*, where each block contains a single element. We write NC_n^* for the poset of nondiscrete noncrossing partitions.

Definition 3.3 (rank function) The noncrossing partition lattice NC_n is a graded poset with a rank function. The *rank* of the noncrossing partition $\pi = \{A_1, A_2, \dots, A_k\}$ is $\text{rk}(\pi) = n - k$. In particular, the rank of the discrete partition is 0, the rank of the maximum partition is $n - 1$ and the rank of the irreducible partition π_A is $|A| - 1$.

For more about noncrossing partitions, see [14; 2; 15]. Using the rotation braids defined in Definition 2.11, there is a natural map from noncrossing partitions to braids.

Definition 3.4 (dual simple braids) Let $\pi = \{A_1, \dots, A_k\} \in \text{NC}_n$ be a noncrossing partition. The *dual simple braid* δ_π is defined to be the product of the rotation braids $\delta_{A_1} \cdots \delta_{A_k}$ and since the rotation braid of a singleton set is the trivial braid, the product only needs to be taken over the blocks of size at least 2. Moreover, because the standard subdisks D_{A_i} are pairwise disjoint, the rotation braids δ_{A_i} pairwise

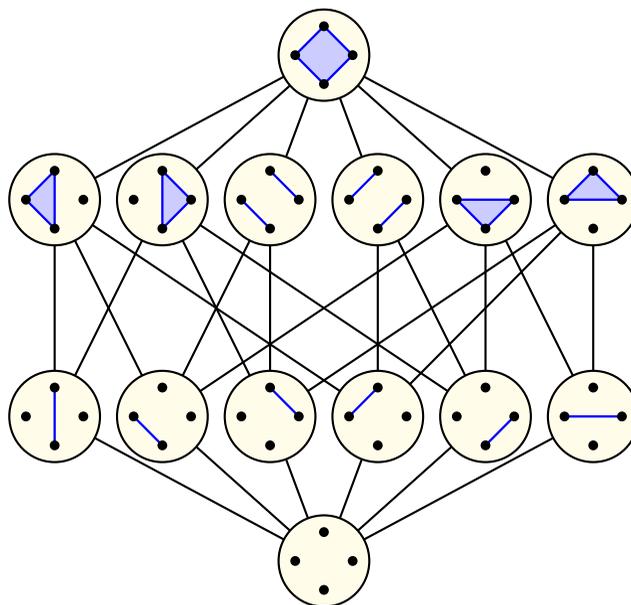


Figure 3: The noncrossing partition lattice NC_4 .

commute and the order in which they are multiplied is irrelevant. Finally note that for each $A \subseteq [n]$ of size at least 2, the irreducible partition π_A corresponds to the rotation braid δ_A . In accordance with notation for noncrossing partitions, we denote by $DS_n := \{\delta_\pi \mid \pi \in NC_n\}$ the set of dual simple braids and by $DS_n^* := \{\delta_\pi \mid \pi \in NC_n^*\}$ the set of nontrivial dual simple braids. We equip both sets with the order coming from NC_n ; see Figure 4.

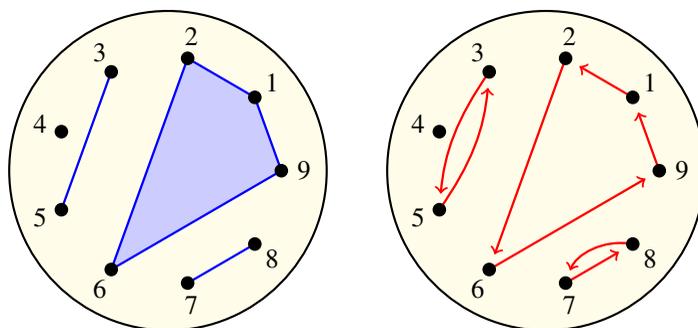


Figure 4: The noncrossing partition π with blocks $A_1 = \{1, 2, 6, 9\}$, $A_2 = \{3, 5\}$, $A_3 = \{4\}$ and $A_4 = \{7, 8\}$ on the left corresponds to the braid $\delta_\pi = \delta_{A_1} \delta_{A_2} \delta_{A_4}$ on the right.

Taking the transitive closure gives a partial order on all of BRAID_n . It has the following property:

Proposition 3.5 *The partial order \leq on BRAID_n is a left-invariant lattice order. The set DS_n is the interval $[1, \delta]$ with respect to this order. In particular, if $\sigma, \tau \in \text{NC}_n$ then $\sigma \leq \tau$ if and only if $\delta_\sigma^{-1}\delta_\tau$ is a dual simple braid if and only if $\delta_\tau\delta_\sigma^{-1}$ is a dual simple braid.*

Proof This can be seen from [4] but it is easier to quote results from [5]. Our dual simple braids are “(braids corresponding to) allowable elements” in [5]. That the order on DS_n is left-divisibility follows from [5, Lemma 3.10]. Consequently, taking the transitive closure is the same as taking an element to be ≥ 1 if and only if it is generated by dual simple braids. That this defines a left-invariant partial order follows from [5, Lemma 5.6]. That DS_n is the interval $[1, \delta]$ follows from the injectivity statement [5, Theorem 5.7]. □

There is a third poset, isomorphic to both NC_n and DS_n , which provides another useful perspective on the combinatorics of noncrossing partitions.

Definition 3.6 (noncrossing permutations) As described in Definition 3.4, the poset of dual simple braids is obtained via an injection of NC_n into BRAID_n . The composition of this injection with the PERM map is also injective and we refer to the image of the composition $\text{NC}_n \hookrightarrow \text{SYM}_n$ as the set of *noncrossing permutations*, denoted by NP_n . Refer to the noncrossing permutation corresponding to $\pi \in \text{NC}_n$ as σ_π . With the partial order induced by the noncrossing partition lattice, NP_n is isomorphic to both DS_n and NC_n .

To help the reader keep track of the notation adherent to NC_n and its counterparts within BRAID_n and SYM_n we provide a dictionary in Figure 5.

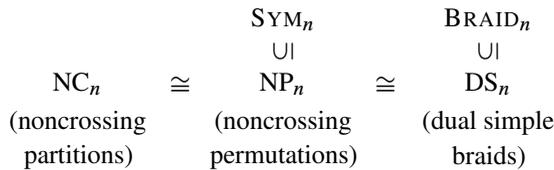


Figure 5: The names of noncrossing partitions as subsets of the symmetric group and the braid group.

The following proposition records standard facts about factorizations of dual simple braids into dual simple braids:

Proposition 3.7 (relations) *If $\pi, \pi' \in \text{NC}_n$ are noncrossing partitions with $\pi \leq \pi'$, then there exist unique $\pi_1, \pi_2 \in \text{NC}_n$ such that $\delta_{\pi_1} \delta_{\pi} = \delta_{\pi} \delta_{\pi_2} = \delta_{\pi'}$ in BRAID_n . Conversely, if $\pi_1, \pi_2, \pi_3 \in \text{NC}_n$ are noncrossing partitions such that $\delta_{\pi_1} \delta_{\pi_2} = \delta_{\pi_3}$ in BRAID_n , then $\pi_1 \leq \pi_3$, $\pi_2 \leq \pi_3$ and π_3 is the join of π_1 and π_2 in NC_n .*

Proof This follows from Theorem 3.7, Lemma 3.9 and Theorem 4.8 of [5]. \square

These relations are used to define the dual presentation of the braid group.

Definition 3.8 (dual presentation) Let $S = \{s_{\pi} \mid \pi \in \text{NC}_n^*\}$ be a set indexed by the nontrivial noncrossing partitions and let R be the set of relations of the form $s_{\pi_1} s_{\pi_2} = s_{\pi_3}$ where such a relation is in R if and only if $\delta_{\pi_1} \delta_{\pi_2} = \delta_{\pi_3}$ holds in BRAID_n . The finite presentation $\langle S \mid R \rangle$ is called the *dual presentation of the n -strand braid group*.

The name reflects the following fact established by Brady in [5].

Theorem 3.9 (dual presentation) *The abstract group G defined by the dual presentation of the n -strand braid group is isomorphic to the n -strand braid group. Concretely, the function that sends $s_{\pi} \in S$ to the dual simple braid $\delta_{\pi} \in \text{BRAID}_n$ extends to a group isomorphism between G and BRAID_n .* \square

4 Parabolic subgroups

This section establishes properties of subgroups of BRAID_n indexed by noncrossing partitions of $[n]$. We begin by showing that two different configuration spaces have isomorphic fundamental groups.

Lemma 4.1 (isomorphic groups) *For a subset $A \subset [n]$ of size k , let $B = [n] - A$ and $D^B = D - P_B$. The natural inclusion map $D_A \hookrightarrow D^B$ extends to an inclusion map $h: \text{UCONF}_k(D_A) \hookrightarrow \text{UCONF}_k(D^B)$ and the induced map*

$$h_*: \pi_1(\text{UCONF}_k(D_A), P_A) \rightarrow \pi_1(\text{UCONF}_k(D^B), P_A)$$

is an isomorphism.

Proof When $k = 1$ both groups are trivial and there is nothing to prove. For each element $[f]$ in $\pi_1(\text{UCONF}_k(D^B), P_A)$, the path f can be homotoped so that it never leaves the subdisk D_A . One can, for example, modify f so that the configurations first radially shrink towards a point in the interior of D_A , followed by the original representative f on a rescaled version of D^B strictly contained in D_A , followed by a radial expansion back to the starting position. This shows that h_* is onto. Suppose $[f]$ and $[g]$ are elements in $\pi_1(\text{UCONF}_k(D_A), P_A)$ such that f and g are homotopic based paths in the bigger space $\text{UCONF}_k(D^B)$. A very similar modification can be done here so that the entire homotopy between f and g takes place inside the subdisk D_A , and this shows that h_* is injective. \square

We are interested in the images of these isomorphic groups inside the n -strand braid group.

Definition 4.2 (subgroups) Let A be a nonempty subset of $[n]$ of size k , $B = [n] - A$ and $D^B = D - P_B$. For each such A , we define a map from BRAID_k to BRAID_n , whose image is a subgroup we call BRAID_A . When $k = 1$, BRAID_k is trivial, the only possible map is the trivial map and BRAID_A is the trivial subgroup of BRAID_n . For $k > 1$, the subspace D_A is a disk, by Corollary 1.10 the group $\pi_1(\text{UCONF}_k(D_A), P_A)$ is isomorphic to BRAID_k , and, by Lemma 4.1, $\pi_1(\text{UCONF}_k(D^B), P_A)$ is also isomorphic to BRAID_k . Let $g: \text{UCONF}_k(D^B) \hookrightarrow \text{UCONF}_n(P)$ be the natural embedding that sends a set $U \in \text{UCONF}_k(D^B)$ to $g(U) = U \cup P_B \in \text{UCONF}_n(P)$ and note that $g(P_A) = P$. The group BRAID_A is the subgroup $g_*(\pi_1(\text{UCONF}_k(D^B), P_A))$.

Note that for every $A \subset [n]$, the rotation braid δ_A is an element of the subgroup BRAID_A . We are also interested in the braids that fix a subset of vertices in V .

Definition 4.3 (fixing vertices) Let α be a braid in BRAID_n represented by f . We say that f fixes the vertex $p_i \in P$ if the strand that starts at p_i is a constant path, ie $f^i(t) = p_i$ for all $t \in [0, 1]$. Similarly, f fixes $P_B \subset P$ if it fixes each $p_i \in P_B$ and a braid α fixes P_B if it has some representative f that fixes P_B . Let $\text{FIX}_n(B) = \{\alpha \in \text{BRAID}_n \mid \alpha \text{ fixes } P_B\}$. Since the special representatives can be concatenated and inverted while remaining special, $\text{FIX}_n(B)$ is a subgroup of BRAID_n .

The two constructions describe the same set of subgroups.

Lemma 4.4 ($\text{FIX}_n(B) = \text{BRAID}_A$) If A and B are nonempty sets that partition $[n]$, then the fixed subgroup $\text{FIX}_n(B)$ is equal to the subgroup BRAID_A .

Proof The map g described in Definition 4.2 shows that every braid in BRAID_A has a representative that fixes P_B . Thus $\text{BRAID}_A \subset \text{FIX}_n(B)$. Conversely, let α be a braid in $\text{FIX}_n(B)$ and let f be a representative of α that fixes P_B . Since the vertices in P_B are always occupied, f restricted to the strands that start in P_A is a loop in the space $\text{UCONF}_k(D^B)$. Thus α is in the subgroup $g_*(\pi_1(\text{UCONF}_k(D^B), P_A)) = \text{BRAID}_A$, which means that $\text{FIX}_n(B) \subset \text{BRAID}_A$ and the two groups are equal. \square

The subgroups of the form $\text{BRAID}_A = \text{FIX}_n(B)$ are used to construct the *dual parabolic subgroups* of BRAID_n .

Definition 4.5 (dual parabolic subgroups) Let $\pi = \{A_1, A_2, \dots, A_k\} \in \text{NC}_n$ be a noncrossing partition. We define the subgroup $\text{BRAID}_\pi \subset \text{BRAID}_n$ to be the internal direct product $\text{BRAID}_\pi = \text{BRAID}_{A_1} \times \dots \times \text{BRAID}_{A_k}$. These are pairwise commuting subgroups that intersect trivially because they are moving points around in disjoint standard subdisks D_{A_i} . We call BRAID_π a *dual parabolic subgroup*. The subgroup BRAID_A is an *irreducible dual parabolic* because it corresponds to the irreducible noncrossing partition π_A . And when $A = [n] - \{i\}$ we call BRAID_A a *maximal irreducible dual parabolic*.

The adjective “dual” is used to distinguish them from the standard parabolic subgroups associated with the standard presentation of BRAID_n , but the two collections of subgroups are closely related. To make the connection between them precise, we pause to discuss the standard presentation of BRAID_n and the standard parabolic subgroups derived from this presentation. We begin by recalling some of the basic relations satisfied by a pair of rotations indexed by edges in the disk D .

Definition 4.6 (basic relations) Let e and e' be two edges in D . Since they are straight line segments connecting vertices of the convex polygonal disk D , e and e' are either disjoint, share a common vertex or cross at some point in the interior of each edge. When e and e' are disjoint, the rotations δ_e and $\delta_{e'}$ commute, ie $\delta_e \delta_{e'} = \delta_{e'} \delta_e$. When e and e' share a common vertex, δ_e and $\delta_{e'}$ braid, ie $\delta_e \delta_{e'} \delta_e = \delta_{e'} \delta_e \delta_{e'}$. We call these commuting and braiding relations the *basic relations* of BRAID_n . When e and e' cross, no basic relation between δ_e and $\delta_{e'}$ is defined.

Artin showed that a small set of rotations indexed by edges in D is sufficient to generate BRAID_n and that the basic relations between them are sufficient to complete a presentation of BRAID_n .

Definition 4.7 (standard presentation) Consider the abstract group

$$(4-1) \quad G = \left\langle s_1, \dots, s_{n-1} \mid \begin{array}{l} s_i s_j = s_i s_j \quad \text{if } |i - j| > 1 \\ s_i s_j s_i = s_j s_i s_j \quad \text{if } |i - j| = 1 \end{array} \right\rangle.$$

This is the *standard presentation of the n -strand braid group* and $S = \{s_1, s_2, \dots, s_{n-1}\}$ is its *standard generating set*.

Theorem 4.8 [3] *The abstract group G defined by the standard presentation of the n -strand braid group is isomorphic to the n -strand braid group. Concretely, the function that sends s_i to $\delta_e \in \text{BRAID}_n$, where e is the edge connecting p_i and p_{i+1} , extends to an isomorphism between G and BRAID_n . □*

Standard parabolic subgroups are generated by subsets of S .

Definition 4.9 (standard parabolic subgroups) Let $S = \{s_1, s_2, \dots, s_{n-1}\}$ be the standard generating set for the abstract group $G \cong \text{BRAID}_n$. For any subset $S' \subset S$, the subgroup $\langle S' \rangle \subset G$ generated by S' is called a *standard parabolic subgroup*. The subsets of the form $S_{[i,j]} = \{s_\ell \mid i \leq \ell < j\}$ generate the *irreducible standard parabolic subgroups* of G . These subsets correspond to sets of edges forming a connected subgraph in the boundary of D .

We record two standard facts about the irreducible standard parabolic subgroups of the braid groups: (1) they are isomorphic to braid groups, and (2) they are closed under intersection.

Proposition 4.10 (isomorphisms) *Let $i, j \in [n]$ with $i < j$. Then the irreducible subgroup generated by $S_{[i,j]} = \{s_\ell \mid i \leq \ell < j\}$ is isomorphic to BRAID_k , where $k = j - i + 1$.*

Proof It is immediate from the standard presentation that BRAID_k maps onto $S_{[i,j]}$. That this map is injective follows from the solution of the word problem [3, Section 3]. □

Proposition 4.11 (intersections) *For all $S', S'' \subset S$, the intersection $\langle S' \rangle \cap \langle S'' \rangle$ is equal to the standard parabolic subgroup $\langle S' \cap S'' \rangle$. Moreover, when both $\langle S' \rangle$ and $\langle S'' \rangle$ are irreducible subgroups, so is $\langle S' \cap S'' \rangle$.*

Proof This is immediate from Proposition 4.10. □

For later use we record the following fact:

Lemma 4.12 *The map $\text{BRAID}_n \rightarrow \mathbb{Z}$ that takes δ_π to $\text{rk}(\pi)$ is the abelianization of BRAID_n .*

Proof The map is well-defined by Proposition 3.7 because if $\delta_{\pi_1} \delta_{\pi_2} = \delta_{\pi_3}$ then the rank function on BRAID_n satisfies $\text{rk } \pi_1 + \text{rk } \pi_2 = \text{rk } \pi_3$. The fact that it is the full abelianization is immediate from the standard presentation (4-1). \square

With this we end our digression on standard presentations and return to dual structure.

Lemma 4.13 (maximal dual parabolics) *The intersection of two maximal irreducible dual parabolic subgroups is an irreducible dual parabolic subgroup. In particular, for all $n > 0$ and for all $i, j \in [n]$,*

$$\text{FIX}_n(\{i, j\}) = \text{FIX}_n(\{i\}) \cap \text{FIX}_n(\{j\}).$$

Proof For every pair of vertices p_i and p_j one can pick a sequence $E = (e_1, \dots, e_{n-1})$ of edges in D so that together, in this order, they form an embedded path through all vertices of D starting at p_i and ending at p_j . Because the rotations δ_e for $e \in E$ satisfy the necessary basic relations (Definition 4.6), the function that sends $s_\ell \in S$ to the rotation δ_{e_ℓ} extends to group homomorphism from $g: G \rightarrow \text{BRAID}_n$. In fact, g is a group isomorphism since up to homeomorphism of \mathbb{C} this is just the usual isomorphism between the abstract group G and the braid group BRAID_n . Under this isomorphism the subgroup $\langle S_{[2,n]} \rangle$ is sent to the subgroup $\text{FIX}_n(\{i\})$, the subgroup $\langle S_{[1,n-1]} \rangle$ is sent to the subgroup $\text{FIX}_n(\{j\})$, and the subgroup $\langle S_{[2,n]} \rangle$ is sent to the subgroup $\text{FIX}_n(\{i, j\})$. Proposition 4.11 completes the proof. \square

Lemma 4.14 (relative maximal dual parabolics) *The intersection of two irreducible dual parabolic subgroups that are both maximal in a third irreducible dual parabolic subgroup is again an irreducible dual parabolic subgroup. In other words, for all $n > 0$ and for all $\{i\}, \{j\}, C \subset [n]$,*

$$\text{FIX}_n(C \cup \{i, j\}) = \text{FIX}_n(C \cup \{i\}) \cap \text{FIX}_n(C \cup \{j\}).$$

Proof When C is empty the statement is just Lemma 4.13 and when C is $[n]$ there is nothing to prove. When C is proper and nonempty, all three groups are contained in $\text{FIX}_n(C) = \text{BRAID}_A \cong \text{BRAID}_k$, where k is the size of $A = [n] - C$. There is a homeomorphism from D_A to the regular k -gon that sends vertices to vertices, so Lemma 1.2 shows that the assertion now follows by applying Lemma 4.13 to this k -gon. \square

Proposition 4.15 (arbitrary dual parabolics) *Every proper irreducible dual parabolic subgroup of BRAID_n is equal to the intersection of the maximal irreducible dual parabolic subgroups that contain it and, as a consequence, the collection of irreducible dual parabolics is closed under intersection. In other words, for all $n > 0$ and for every nonempty $B \subset [n]$,*

$$\text{FIX}_n(B) = \bigcap_{i \in B} \text{FIX}_n(\{i\})$$

and, as a consequence, for all nonempty $C, D \subset B$,

$$\text{FIX}_n(C \cup D) = \text{FIX}_n(C) \cap \text{FIX}_n(D).$$

Proof When B is a singleton the result is trivial, and when B has size 2 both claims are true by Lemma 4.13, so suppose that both claims hold for all subsets of size at most k with $k > 1$ and let B be a subset of size $k + 1$. If i and j are elements in B , and $C = B - \{i, j\}$, then $\text{FIX}_n(B) = \text{FIX}_n(C \cup \{i, j\})$, which is equal to $\text{FIX}_n(C \cup \{i\}) \cap \text{FIX}_n(C \cup \{j\})$ by Lemma 4.14. By applying the second inductive claim to the sets $C \cup \{i\}$ and $C \cup \{j\}$ and simplifying slightly, we can rewrite this as $\text{FIX}_n(C) \cap \text{FIX}_n(\{i\}) \cap \text{FIX}_n(\{j\})$. Applying the first inductive claim to the set C shows that first claim holds for B and the second claim for B follows as an immediate consequence. This completes the induction and the proof. \square

Part II Complexes

In this part, we study complexes built out of ordered simplices, specifically how they can be equipped with an orthoscheme metric. The abstract combinatorial complexes are described in Section 5 and the orthoscheme metric is introduced in Section 6. In Section 7 we discuss products in both contexts. Section 8 is concerned with a special kind of complex that will turn out to be important in Part III.

5 Ordered simplices

An ordered simplex is a simplex with a fixed linear ordering of its vertex set. Complexes built out of ordered simplices are often used as explicit models. Eilenberg and Steenrod, for example, use ordered simplicial complexes [9]. We follow Hatcher in using the more flexible Δ -complexes [12].

Definition 5.1 (ordered simplices) A k -simplex is the convex hull of $k + 1$ points p_0, p_1, \dots, p_k in general position in a sufficiently high-dimensional real vector space E . An ordered k -simplex is a k -simplex together with a fixed linear ordering of its $k + 1$ vertices. We write $\sigma = [p_0, p_1, \dots, p_k]$ for an ordered k -simplex σ with vertex set $\{p_0, p_1, \dots, p_k\} \subset E$ where the vertices are ordered left to right: $p_i < p_j$ in the linear order if and only if $i < j$ in the natural numbers. An isomorphism of ordered simplices is an affine bijection $[p_0, \dots, p_k] \rightarrow [p'_0, \dots, p'_k]$ that takes p_i to p'_i .

Let E be a vector space containing an ordered k -simplex σ . To facilitate computations, we establish a standard coordinate system on the smallest affine subspace of E containing σ which both identifies this subspace with \mathbb{R}^k and also reflects the linear ordering of its vertices.

Definition 5.2 (standard coordinates) Let $\sigma = [p_0, p_1, \dots, p_k]$ be an ordered k -simplex. We take the ambient vector space E to have origin p_0 and to be spanned by p_1, \dots, p_k . For each $i \in [k]$, let $\vec{v}_i = p_i - p_{i-1}$ be the vector from p_{i-1} to p_i , so that $\mathcal{B} = (\vec{v}_1, \vec{v}_2, \dots, \vec{v}_k)$ is an ordered basis for E . We call \mathcal{B} the *standard ordered basis* of σ . In this basis, $p_j = \sum_{i=1}^j \vec{v}_i$ and

$$\begin{aligned} x &= (x_1, x_2, \dots, x_k)_{\mathcal{B}} = x_1 \vec{v}_1 + \dots + x_k \vec{v}_k \\ &= (-x_1)p_0 + (x_1 - x_2)p_1 + \dots + (x_{k-1} - x_k)p_{k-1} + (x_k)p_k \\ &= (1 - x_1)p_0 + (x_1 - x_2)p_1 + \dots + (x_{k-1} - x_k)p_{k-1} + (x_k)p_k \end{aligned}$$

since p_0 is the origin. Since the coefficients in the last equation are barycentric coordinates on E , we see that $(x_1, x_2, \dots, x_k)_{\mathcal{B}}$ is in σ if and only if $1 \geq x_1 \geq x_2 \geq \dots \geq x_k \geq 0$. In particular, the facets of σ determine the $k + 1$ hyperplanes given by the equations $x_1 = 1$, $x_i = x_{i+1}$ for $i \in [k - 1]$ and $x_k = 0$. If σ and σ' are ordered simplices with oriented bases $(\vec{v}_1, \dots, \vec{v}_k)$ and $(\vec{v}'_1, \dots, \vec{v}'_k)$, the unique isomorphism $\sigma \rightarrow \sigma'$ takes $\sum_i \alpha_i \vec{v}_i$ to $\sum_i \alpha_i \vec{v}'_i$.

Example 5.3 (standard coordinates) Figure 6 shows an ordered 3-simplex σ . The vector \vec{v}_1 is from p_0 to p_1 , the vector \vec{v}_2 is from p_1 to p_2 and the vector \vec{v}_3 is from p_2 to p_3 . With respect to the ordered basis $\mathcal{B} = \{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$ with p_0 located at the origin, $p_1 = (1, 0, 0)$, $p_2 = (1, 1, 0)$ and $p_3 = (1, 1, 1)$.

Faces of ordered simplices are ordered by restriction. As such they have standard coordinates which can be described as follows.

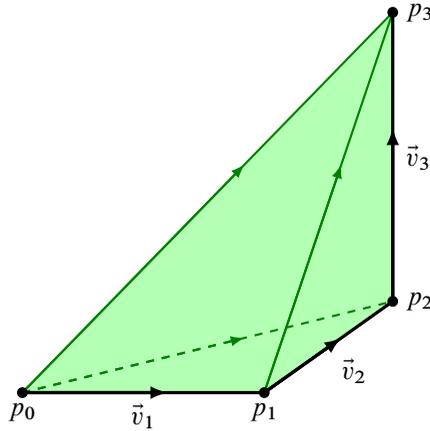


Figure 6: An ordered 3-simplex. In standard coordinates, p_0 is the origin, $p_1 = (1, 0, 0)$, $p_2 = (1, 1, 0)$ and $p_3 = (1, 1, 1)$ with respect to the ordered basis $\mathcal{B} = \{\vec{v}_1, \vec{v}_2, \vec{v}_3\}$.

Lemma 5.4 (facets) *Let $\sigma = [p_0, \dots, p_k]$ be an ordered simplex with ordered basis $\mathcal{B} = (\vec{v}_1, \dots, \vec{v}_k)$. The ordered basis \mathcal{B}' of the facet $\tau = [p_0, \dots, p_{i-1}, p_{i+1}, \dots, p_k]$ is*

$$\mathcal{B}' = \begin{cases} (\vec{v}_2, \dots, \vec{v}_k) & \text{if } 0 = i, \\ (\vec{v}_1, \dots, \vec{v}_{i-1}, \vec{v}_i + \vec{v}_{i+1}, \vec{v}_{i+2}, \dots, \vec{v}_k) & \text{if } 0 < i < k, \\ (\vec{v}_1, \dots, \vec{v}_{k-1}) & \text{if } i = k. \end{cases} \quad \square$$

In anticipation of Definition 6.2, the following is modeled on [7, Definition I.7.2]:

Definition 5.5 (Δ -complex) *Let $(\sigma_\lambda)_{\lambda \in \Lambda}$ be a family of ordered simplices with disjoint union $X = \bigcup (\sigma_\lambda \times \{\lambda\})$. Let \sim be an equivalence relation on X and let $K = X/\sim$. Let $p: X \rightarrow K$ be the quotient map and $p_\lambda: \sigma_\lambda \rightarrow K$, $x \mapsto p(x, \lambda)$, its restriction to σ_λ .*

We say that K is a Δ -complex if

- (1) the restriction of p_λ to the interior of σ_λ is injective;
- (2) for $\lambda \in \Lambda$ and every face τ of σ_λ there is a $\lambda' \in \Lambda$ and an isomorphism of ordered simplices $h: \tau \rightarrow \sigma_{\lambda'}$ such that $p_\lambda|_\tau = p_{\lambda'} \circ h$;
- (3) if $\lambda, \lambda' \in \Lambda$ and interior points $x \in \sigma_\lambda$ and $x' \in \sigma_{\lambda'}$ are such that $p_\lambda(x) = p_{\lambda'}(x')$ then there is an isomorphism of ordered simplices $h: \sigma_\lambda \rightarrow \sigma_{\lambda'}$ such that $p_{\lambda'}(h(x)) = p_\lambda(x)$ for $x \in \sigma_\lambda$.

We will usually regard Δ -complexes as equipped with a structure as above and refer to the simplices σ_λ as simplices of K . We will also make the identification in (2) implicit and regard faces of σ_λ as simplices of K .

Turning a simplicial complex into a Δ -complex means to orient the edges in a consistent way. A setting where a natural orientation exists is the following:

Proposition 5.6 (Cayley graphs and Δ -complexes) *Let G be a group and let $f: G \rightarrow (\mathbb{R}, +)$ be a group homomorphism. Let $S \subset G$ be a set of generators such that $f(s) > 0$ for every $s \in S$. The right Cayley graph $\Gamma = \text{Cay}(G, S)$ is a simplicial graph whose flag complex $X = \text{FLAG}(\Gamma)$ can be turned into a Δ -complex.*

Proof The right Cayley graph has no doubled edges because $f(g^{-1}) = -f(g)$ for $g \in G$, so that at most one of g and g^{-1} is in S . It has no loops because $f(1) = 0$, so that $1 \notin S$.

We define a relation \leq on G by declaring that $g \leq gs$ for $s \in S$ and taking the reflexive transitive closure. The homomorphism f guarantees that this is a partial order on G . Any two adjacent vertices are comparable, so the restriction to a simplex is a total order. □

6 Orthoschemes

The goal of this section is to equip certain Δ -complexes with a piecewise Euclidean metric.

Definition 6.1 (orthoscheme) *Let E be a Euclidean vector space and let $\sigma \subseteq E$ be a simplex. Then σ with the induced metric is called a *Euclidean simplex*. If σ is an ordered simplex and the associated ordered basis is orthogonal then σ is an *orthoscheme*. If it is an orthonormal basis then σ is a *standard orthoscheme*.*

Definition 6.2 (orthoscheme complex) *An *orthoscheme complex* is a Δ -complex where each simplex has been given the metric of an orthoscheme in such a way that the isomorphisms in the definition of a Δ -complex are isometries. It is equipped with the length pseudometric assigning to two points the infimal length of a piecewise affine path.*

Remark 6.3 Orthoscheme complexes are M_0 -simplicial complexes in the sense of [7, I.7.1], so we will not discuss the metric subtleties in detail. Our main interest concerning the metric is the behavior with respect to products, which is not among the subtleties.

Lemma 6.4 (orthoscheme complex structures and edge norms) *Let X be a Δ -complex. There is a one-to-one correspondence between orthoscheme complex structures on X and maps $\text{NORM}: \text{EDGES}(X) \rightarrow \mathbb{R}_{>0}$ that satisfy*

$$(6-1) \quad \text{NORM}([p_0, p_1]) + \text{NORM}([p_1, p_2]) = \text{NORM}([p_0, p_2])$$

for every 2-simplex $[p_0, p_1, p_2]$ of X .

Proof If X is equipped with a orthoscheme complex structure, defining

$$\text{NORM}([p, q]) = \|q - p\|^2$$

gives a map satisfying condition (6-1) (corresponding to the right angle in p_1).

Conversely, the squares of edge lengths of an orthoscheme need to satisfy (6-1) and this is the only requirement for a well-defined assignment. The unique isomorphism of ordered simplicial complexes between orthoschemes with same edge lengths is an isometry. Hence equipping the simplices of a Δ -complex with a orthoscheme metric satisfying (6-1) gives rise to an orthoscheme complex. \square

We can use this characterization of orthoscheme complexes to extend Proposition 5.6.

Proposition 6.5 (Cayley graphs and orthoschemes) *Let G be a group, $f: G \rightarrow (\mathbb{R}, +)$ be a group homomorphism, and S be a generating set of G with $f(s) > 0$ for every $s \in S$. The right Cayley graph $\Gamma = \text{Cay}(G, S)$ is a simplicial graph whose flag complex $X = \text{FLAG}(\Gamma)$ can be turned into an orthoscheme complex using f .*

Proof By Proposition 5.6 X is a Δ -complex and we claim that

$$\text{NORM}([g, gs]) := f(gs) - f(g) = f(s) > 0$$

satisfies (6-1).

Indeed,

$$\begin{aligned} \text{NORM}([g, gss']) &= f(gss') - f(g) \\ &= (f(gss') - f(gs)) + (f(gs) - f(g)) \\ &= \text{NORM}([gss', gs]) + \text{NORM}([gs, g]) \end{aligned}$$

since f is a homomorphism. \square

Definition 6.6 (dual braid complex) Let $S = DS_n^*$ be the set of nontrivial dual simple braids in the braid group BRAID_n . By Theorem 3.9, the set $S \subset \text{BRAID}_n$ generates the group and by Lemma 4.12 the abelianization map $f: \text{BRAID}_n \rightarrow \mathbb{Z}$ sends the nontrivial dual simple braid $\delta_\pi \in S$ to the positive integer $f(\delta_\pi) = \text{rk}(\pi)$. By Proposition 6.5, the flag complex $X = \text{FLAG}(\Gamma)$ of the simplicial graph $\Gamma = \text{Cay}(\text{BRAID}_n, S)$ can be turned into an orthoscheme complex using f to compute the norm of each edge. The resulting orthoscheme complex is called the *dual braid n -complex* and denoted by $\text{CPLX}(\text{BRAID}_n)$. Note that every edge of $\text{CPLX}(\text{BRAID}_n)$ is naturally labeled by an element of $S = DS_n^*$ or, equivalently, by a nondiscrete noncrossing partition. More generally, every simplex is naturally labeled by a chain of NC_n .

It is clear from the construction that BRAID_n acts freely on $\text{CPLX}(\text{BRAID}_n)$ and that $\text{CPLX}(\text{BRAID}_n)$ is covered by translates of the full subcomplex supported on DS_n , which is therefore a fundamental domain. The key feature, implying that

$$\text{BRAID}_n \backslash \text{CPLX}(\text{BRAID}_n)$$

is a classifying space for BRAID_n , is:

Theorem 6.7 [5] *The complex $\text{CPLX}(\text{BRAID}_n)$ is contractible.*

In fact, it was shown in [6] $\text{CPLX}(\text{BRAID}_n)$ is $\text{CAT}(0)$ when $n < 6$ and in [11] this was extended to the case $n = 6$.

7 Products

The main advantage of working with ordered simplices and Δ -complexes is that they admit well-behaved products.

Example 7.1 (products of simplices) The product of two 1-simplices is a quadrangle. It can be subdivided into two triangles in two ways but neither of these is distinguished. More generally, the product of two (positive-dimensional) simplices is not a simplex nor does it have a canonical simplicial subdivision.

In contrast, we will see that the product of two simplices whose vertices are totally ordered admits a canonical subdivision into chains. We start by looking at finite products of edges first, ie cubes.

Example 7.2 (subdivided cubes) Let \mathbb{R}^k be a k -dimensional real vector space with a fixed ordered basis $\mathcal{B} = \{\vec{v}_1, \dots, \vec{v}_k\}$. The *unit k -cube* CUBE_k in \mathbb{R}^k is the set of vectors where each coordinate is in the interval $[0, 1]$ and its vertices are the points where every coordinate is either 0 or 1. There is a natural bijection between the vertex set of CUBE_k and the set of all subsets of $[k]$: simply send each vertex to the set of indices of the coordinates where the value is 1. If we partially order the subsets of $[k]$ by inclusion (to form the Boolean lattice BOOL_k), this partially orders the vertices of CUBE_k , and by sending $B \subset [k]$ to the vector $\mathbf{1}_B = \sum_{i \in B} \vec{v}_i$, we obtain a convenient labeling for the vertices. At the extremes, we write $\mathbf{1} = \mathbf{1}_{[k]} = (1, 1, \dots, 1)$ and $\mathbf{0} = \mathbf{1}_\emptyset = (0, 0, \dots, 0)$.

Let \mathcal{H} be the collection of hyperplanes H_{ij} in \mathbb{R}^k defined by the equations $x_i = x_j$ for $i \neq j \in [k]$. There is a minimal cellular subdivision of CUBE_k for which $\text{CUBE}_k \cap H_{ij}$ is a subcomplex for all $i \neq j \in [k]$ and it is a simplicial subdivision. The subdivision has $k!$ top-dimensional simplices and the partial order on the vertices of CUBE_k is a linear order when restricted to each simplex. In particular, this subdivided k -cube is a Δ -complex.

Remark 7.3 When our selected basis \mathcal{B} is orthonormal, CUBE_k is a regular Euclidean unit cube and the top-dimensional simplices are orthoschemes. In other words, the $k!$ simplices in the simplicial structure for CUBE_k correspond to the $k!$ ways to take k steps from $\mathbf{0}$ to $\mathbf{1}$ in the coordinate directions. A 3-orthoscheme from the simplicial structure on CUBE_3 is shown in Figure 7.

Example 7.4 (products of subdivided cubes) The product of CUBE_k and CUBE_ℓ is naturally identified with the unit cube $\text{CUBE}_{k+\ell}$. Using the simplicial subdivision given in Example 7.2, we obtain a canonical Δ -complex structure for the product $\text{CUBE}_{k+\ell}$. Selecting top-dimensional simplices σ in CUBE_k and τ in CUBE_ℓ corresponds to a product of simplices $\sigma \times \tau$ in $\text{CUBE}_{k+\ell}$, which then inherits a simplicial subdivision from that of $\text{CUBE}_{k+\ell}$.

Since any ordered simplex can be considered as a top-dimensional simplex in the subdivision from Example 7.2, we can use the simplicial structure for $\text{CUBE}_{k+\ell}$ to describe the product of two ordered simplices as a Δ -complex.

Example 7.5 (product of ordered simplices) Let σ and τ be ordered simplices of dimension k and ℓ , respectively, with $\sigma = [v_0, v_1, \dots, v_k]$ and $\tau = [u_0, u_1, \dots, u_\ell]$.

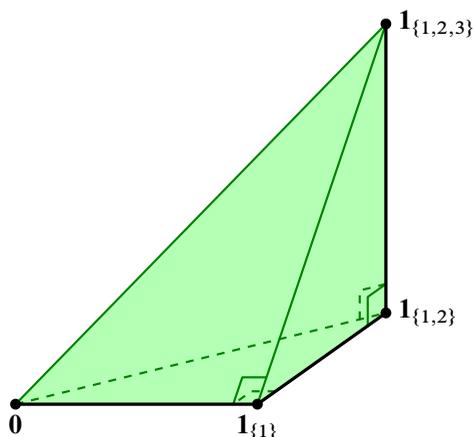


Figure 7: A 3-orthoscheme from $\mathbf{0}$ to $\mathbf{1} = \mathbf{1}_{\{1,2,3\}}$ inside CUBE_3 . The edges of the piecewise geodesic path are thicker and darker than the others.

In standard coordinates $\sigma \subset \mathbb{R}^k$ is the set of points $x = (x_1, x_2, \dots, x_k) \in \mathbb{R}^k$ satisfying the inequalities $1 \geq x_1 \geq x_2 \geq \dots \geq x_k \geq 0$. Similarly, $\tau \subset \mathbb{R}^\ell$ is the set of points $y = (y_1, y_2, \dots, y_\ell) \in \mathbb{R}^\ell$ satisfying the inequalities $1 \geq y_1 \geq y_2 \geq \dots \geq y_\ell \geq 0$. The product $\sigma \times \tau$ is the set of points $(x_1, \dots, x_k, y_1, \dots, y_\ell) \in \mathbb{R}^k \times \mathbb{R}^\ell$ satisfying both sets of inequalities. Let \mathcal{H} be the collection of $k \cdot \ell$ hyperplanes H_{ij} defined by the equations $x_i = y_j$, with $i \in [k]$ and $j \in [\ell]$. When we minimally subdivide the polytope $\sigma \times \tau$ so that for every $i \in [k]$ and every $j \in [\ell]$, $H_{ij} \cap (\sigma \times \tau)$ is a subcomplex of the new cell structure, then the new cell structure is a simplicial complex which contains $\binom{k+\ell}{k}$ simplices of dimension $k + \ell$. The points in the interiors of these top-dimensional simplices correspond to points $(x_1, \dots, x_k, y_1, \dots, y_\ell)$ where all $k + \ell$ coordinates are distinct and the simplex containing this point is determined by the xy -pattern of the coordinates when arranged in decreasing linear order. For example, if $k = 2, \ell = 1$ and $x_0 > x_1 > y_0 > x_2 > y_1$ then its pattern is $xxxyxy$ and all generic points with this pattern belong to the same top-dimensional simplex. The natural partial order on the vertices of $\sigma \times \tau$ is given by the rule $(v_{i_1}, u_{j_1}) \leq (v_{i_2}, u_{j_2})$ if and only if $i_1 \leq i_2$ and $j_1 \leq j_2$. This restricts to a linear order on each simplex in the new simplicial structure, which turns the result into an ordered simplicial complex.

Definition 7.6 (product of ordered simplices) Let σ and τ be ordered simplices. The decomposition of $\sigma \times \tau$ described in Example 7.5 is the *canonical decomposition*. We write $\sigma \boxtimes \tau$ to denote the Δ -complex that is $\sigma \times \tau$ with the canonical decomposition.

The construction described in Definition 7.6 is the natural generalization of partitioning the unit square in the first quadrant by the diagonal line where the two coordinates are equal. It readily generalizes to finite products of ordered simplices.

Example 7.7 (finite products) Let $\sigma_1, \sigma_2, \dots, \sigma_m$ be ordered simplices of dimension k_1, k_2, \dots, k_m , respectively, and view $\sigma_1 \times \sigma_2 \times \dots \times \sigma_m$ as a subset of $\mathbb{R}^{k_1 + \dots + k_m}$ with coordinates given by concatenating the standard ordered bases. Let \mathcal{H} be the finite collection of hyperplanes defined by an equation setting a canonical coordinate in one factor equal to a canonical coordinate in a different factor. The minimal subdivision of the product cell complex $X = \sigma_1 \times \sigma_2 \times \dots \times \sigma_m$ such that $H \cap X$ is a subcomplex in the new cell structure for every hyperplane $H \in \mathcal{H}$ is a simplicial complex with N simplices of dimension $k_1 + k_2 + \dots + k_m$, where N is the multinomial coefficient $\binom{k_1 + k_2 + \dots + k_m}{k_1, k_2, \dots, k_m}$. This illustrates that \boxtimes is associative.

The canonical subdivision of products of ordered simplices also readily extends to the product of Δ -complexes.

Definition 7.8 (products of Δ -complexes) Let X and Y be Δ -complexes. The product complex $X \times Y$ carries a canonical Δ -complex structure which can be described as follows. Let $p_\sigma: \sigma \rightarrow X$ and $p_\tau: \tau \rightarrow Y$ be simplices of X and Y . Then every simplex ρ in the canonical subdivision of $\sigma \times \tau$ is a simplex of $X \times Y$ via $p_\rho = (p_\sigma \times p_\tau)|_\rho$. We denote $X \times Y$ with this Δ -complex structure by $X \boxtimes Y$.

Note that when σ and τ are Euclidean simplices, their product $\sigma \times \tau$ inherits a Euclidean metric from the metric product of the Euclidean spaces containing them and when they are ordered Euclidean simplices, the ordered simplicial complex $\sigma \boxtimes \tau$ is constructed out of ordered Euclidean simplices.

The product described in Definition 7.6 is well-behaved when the factors are orthoschemes or standard orthoschemes.

Lemma 7.9 (products of orthoschemes) *If σ and τ are orthoschemes, then $\sigma \boxtimes \tau$ is an orthoscheme complex isometric to the metric product $\sigma \times \tau$. Moreover, when σ and τ are standard orthoschemes, then every top-dimensional simplex in $\sigma \boxtimes \tau$ is a standard orthoscheme.*

Proof Let σ and τ be ordered Euclidean simplices of dimension k and ℓ , respectively, and let $\sigma \boxtimes \tau$ be the simplicial decomposition of the Euclidean polytope $\sigma \times \tau$ into Euclidean simplices. If \mathcal{B}_1 and \mathcal{B}_2 are the standard ordered bases associated to σ

and τ , then, by Example 7.5, the Δ -complex $\sigma \boxtimes \tau$ is a subcomplex of the unit cube $\text{CUBE}_{k+\ell}$ in the ordered basis obtained by concatenating \mathcal{B}_1 and \mathcal{B}_2 . When \mathcal{B}_1 and \mathcal{B}_2 are both orthogonal, σ and τ are orthoschemes and the concatenated ordered bases produce a metric Euclidean cube for which the simplicial subdivision in Example 7.2 makes $\text{CUBE}_{k+\ell}$ into an orthoscheme complex. Hence, the subcomplex $\sigma \boxtimes \tau$ is an orthoscheme complex as well. The analogous result for standard orthoschemes follows by considering the case when \mathcal{B}_1 and \mathcal{B}_2 are both orthonormal. \square

As a consequence we get:

Proposition 7.10 (products of orthoscheme complexes) *If X and Y are orthoscheme complexes then $X \boxtimes Y$ is an orthoscheme complex isometric to the metric direct product $X \times Y$.* \square

8 Columns

In this section we describe a particularly useful type of orthoscheme complex, initially defined in [6].

Example 8.1 (orthoschemes and \mathbb{R}^k) Regard \mathbb{R} as an infinite linear graph with vertex set \mathbb{Z} and edges from i to $i+1$. Then \mathbb{R}^k is isometric to the k -fold product $\mathbb{R} \boxtimes \cdots \boxtimes \mathbb{R}$. This complex has vertex set \mathbb{Z}^k with simplices on vertices $\vec{x} \leq \vec{x} + \mathbf{1}_{B_1} \leq \cdots \leq \vec{x} + \mathbf{1}_{B_\ell}$ for $\emptyset \subsetneq B_1 \subseteq \cdots \subseteq B_\ell \subseteq [k]$. We call this the standard *orthoscheme tiling of \mathbb{R}^k* . It can also be viewed as the standard cubing of \mathbb{R}^k in which each k -cube has been given the simplicial subdivision described in Example 7.2.

Alternatively, the orthoscheme tiling of \mathbb{R}^k can be viewed as the cell structure of a simplicial hyperplane arrangement.

Definition 8.2 (types of hyperplanes) Consider the hyperplane arrangement consisting of two types of hyperplanes. The *first type* are those defined by the equations $x_i = \ell$ for all $i \in [k]$ and all $\ell \in \mathbb{Z}$. The *second type* are those defined by the equations $x_i - x_j = \ell$ for all $i \neq j \in [k]$ and all $\ell \in \mathbb{Z}$. When both types of hyperplanes are used, the resulting hyperplane arrangement partitions \mathbb{R}^k into its standard orthoscheme tiling.

The hyperplanes of the first type define the standard cubing of \mathbb{R}^k and the hyperplanes of the second type are closely related to the Coxeter complex of the affine symmetric group.

Definition 8.3 (affine symmetric group) The Euclidean Coxeter group of type \tilde{A}_{k-1} is also called the *affine symmetric group* $\widetilde{\text{SYM}}_k$. It is generated by orthogonal reflections in the hyperplanes of the second kind.

Remark 8.4 The spherical Coxeter group of type A_{k-1} , the symmetric group, is generated by reflections in hyperplanes of the second type for which $\ell = 0$. Since the roots $e_i - e_j$ are perpendicular to the vector $\mathbf{1}$, both the symmetric group and the affine symmetric group act on the $(k-1)$ -dimensional space $\mathbf{1}^\perp$.

Definition 8.5 (Coxeter shapes and columns) The hyperplane arrangement that consists solely of the hyperplanes of the second type restricted to any hyperplane H defined by the equation $\langle x, \mathbf{1} \rangle = r$ for some $r \in \mathbb{R}$ partitions $H \cong \mathbb{R}^{k-1}$ into a reflection tiling by Euclidean simplices whose shape is encoded in the extended Dynkin diagram of type \tilde{A}_{k-1} . We call the isometry type of this Euclidean simplex the *Coxeter shape* or *Coxeter simplex* of type \tilde{A}_{k-1} and when the subscript is clear from context it is often omitted to improve clarity. When this hyperplane arrangement is not restricted to a hyperplane orthogonal to the vector $\mathbf{1}$, the closure of a connected component of the complementary region is an unbounded infinite column that is a metric product $\sigma \times \mathbb{R}$, where σ is a Coxeter simplex of type \tilde{A} and \mathbb{R} is the real line. We call these the *columns of \mathbb{R}^k* .

One consequence of this column structure is that the standard orthoscheme tiling of \mathbb{R}^k partitions the columns of \mathbb{R}^k into a sequence of orthoschemes. We begin with an explicit example.

Example 8.6 (column in \mathbb{R}^3) Let \mathcal{C} be the unique column of \mathbb{R}^3 that contains the 3-simplex shown in Figure 7. The column \mathcal{C} is defined by the inequalities $x_1 \geq x_2 \geq x_3 \geq x_1 - 1$ and its sides are the hyperplanes defined by the equations $x_1 - x_2 = 0$, $x_2 - x_3 = 0$ and $x_1 - x_3 = 1$. The vertices of \mathbb{Z}^3 contained in this column form a sequence $\{v_\ell\}_{\ell \in \mathbb{Z}}$ where the order of the sequence is determined by the inner product of these points with the special vector $\mathbf{1} = (1^3) = (1, 1, 1)$. Concretely, the vertex v_ℓ is the unique point in $\mathbb{Z}^3 \cap \mathcal{C}$ such that $\langle v_\ell, \mathbf{1} \rangle = \ell \in \mathbb{Z}$. The vectors in this case are $v_{-1} = (0, 0, -1)$, $v_0 = (0, 0, 0)$, $v_1 = (1, 0, 0)$, $v_2 = (1, 1, 0)$, $v_3 = (1, 1, 1)$, $v_4 = (2, 1, 1)$ and so on. Successive points in this list are connected by unit-length edges in coordinate directions and this turns the full list into a spiral of edges. Traveling up the spiral, the edges cycle through the possible directions in a predictable order. In this case they travel one unit step in the positive x -direction, y -direction, z -direction,

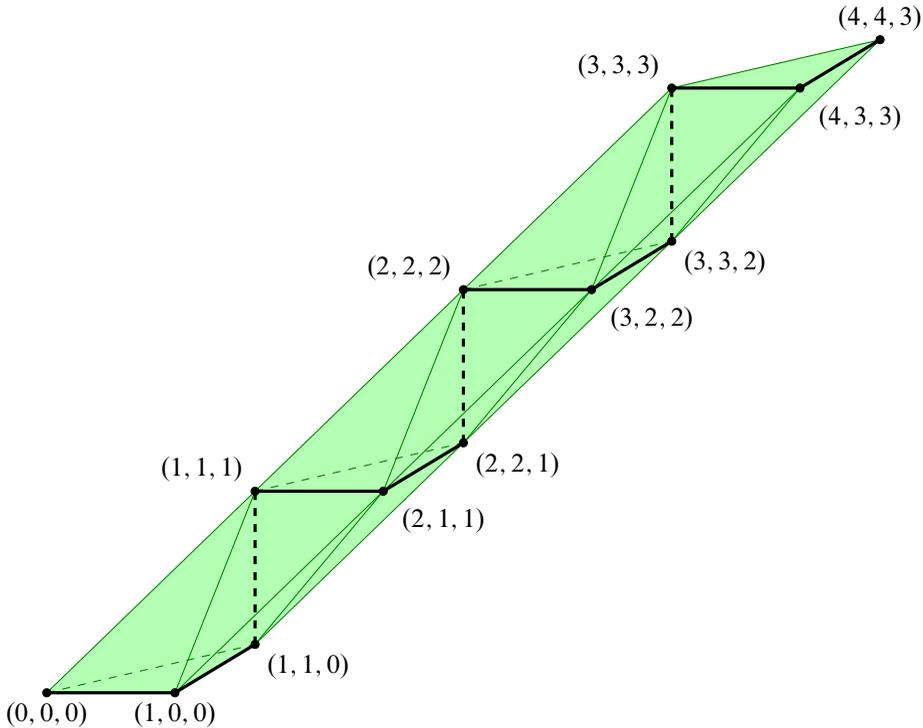


Figure 8: A portion of the column in \mathbb{R}^3 that contains the orthoscheme shown in Figure 7. The edges of the spiral are thicker and darker than the others; see Example 8.6.

x -direction, y -direction, z -direction and so on. Any 3 consecutive edges in the spiral have a standard 3-orthoscheme as its convex hull and the union of these individual orthoschemes is the convex hull of the full spiral, which is also the full column \mathcal{C} . See Figure 8. Metrically \mathcal{C} is $\sigma \times \mathbb{R}$, where σ is an equilateral triangle, ie the Coxeter simplex of type \tilde{A}_2 .

Columns in \mathbb{R}^k have many of the same properties.

Definition 8.7 (columns in \mathbb{R}^k) A column \mathcal{C} of \mathbb{R}^k can be defined by inequalities of the form

$$(8-1) \quad x_{\pi_1} + a_{\pi_1} \geq x_{\pi_2} + a_{\pi_2} \geq \dots \geq x_{\pi_k} + a_{\pi_k} \geq x_{\pi_1} + a_{\pi_1} - 1,$$

where (π_1, \dots, π_k) is a permutation of the integers $(1, \dots, k)$ and $a = (a_1, \dots, a_k)$ is a point in \mathbb{Z}^k . The vertices of \mathbb{Z}^k contained in \mathcal{C} form a sequence $\{v_\ell\}_{\ell \in \mathbb{Z}}$, where

the order of the sequence is determined by the inner product of these points with the vector $\mathbf{1} = (1, 1, \dots, 1)$. Concretely, the vertex v_ℓ is the unique point in $\mathbb{Z}^k \cap \mathcal{C}$ such that $\langle v_\ell, \mathbf{1} \rangle = \ell \in \mathbb{Z}$. Successive points in this list are connected by unit-length edges in coordinate directions and this turns the full list into a spiral of edges. Traveling up the spiral, the edges cycle through the possible directions in a predictable order based on the list $(\pi_1, \pi_2, \dots, \pi_k)$. Any k consecutive edges in the spiral have a standard k -orthoscheme as its convex hull and the union of these individual orthoschemes is the convex hull of the full spiral, which is also the full column \mathcal{C} . Metrically, \mathcal{C} is $\sigma \times \mathbb{R}$ where σ is a Coxeter simplex of type \tilde{A}_{k-1} . Since the full column is a convex subset of \mathbb{R}^k , it is a CAT(0) space.

Definition 8.8 (dilated columns) If the -1 in the final inequality of (8-1) defining a column in \mathbb{R}^k is replaced by a $-\ell$ for some positive integer ℓ , then the shape described is a *dilated column*, ie a dilated version of a single column. As a metric space, a dilated column is a metric direct product of the real line and a Coxeter shape of type \tilde{A} dilated by a factor of ℓ and is also a CAT(0) space. As a cell complex, a dilated column is the union of ℓ^{k-1} ordinary columns of \mathbb{R}^k tiled by orthoschemes.

Some of these dilated columns are of particular interest.

Definition 8.9 ((k, n) -dilated columns) Let $n > k > 0$ be positive integers and let \mathcal{C} be the full subcomplex of the orthoscheme tiling of \mathbb{R}^k restricted to the vertices of \mathbb{Z}^k that satisfy the strict inequalities

$$x_1 < x_2 < \dots < x_k < x_1 + n.$$

We call \mathcal{C} the (k, n) -dilated column in \mathbb{R}^k . A point $x \in \mathbb{Z}^k$ is in \mathcal{C} if and only if its coordinates are strictly increasing in value from left to right and the gap between the first and the last coordinate is strictly less than n . To see that the subspace \mathcal{C} really is a dilated column of \mathbb{R}^k , note that it is defined by the weak inequalities

$$x_1 - 1 \leq x_2 - 2 \leq \dots \leq x_k - k \leq x_1 - (k + 1) + n.$$

There is a natural bijection between the sets of integer vectors satisfying these two sets of inequalities that uses the usual combinatorial trick for converting between statements about strictly increasing integer sequences and statements about weakly increasing ones. From the weak inequalities we see that the (k, n) -dilated column \mathcal{C} is an $(n-k)$ -dilation of an ordinary column and thus a union of $(n-k)^{k-1}$ ordinary columns.

Example 8.10 ((2, 6)–dilated column) When $k = 2$ and $n = 6$, the defining inequalities are $x < y < x + 6$ and a portion of the (2, 6)–dilated column \mathcal{C} is shown in Figure 9. The meaning of the vertex labels used in the figure are explained in Example 10.4. Note that \mathcal{C} is metrically an ordinary column dilated by a factor of 4, its cell structure is a union of $(6 - 2)^{2-1} = 4$ ordinary columns, and it is defined by the weak inequalities $x + 1 \leq y$ and $y \leq x + 5$ or, equivalently, $x - 1 \leq y - 2 \leq x - 3 + 6$.

Part III Boundary braids

In this part we systematically study boundary braids, working toward the main results stated in the introduction. We begin by introducing orthoscheme configuration spaces (Section 9) and describing the specific case of an oriented n –cycle (Section 10). These model the move braid part of boundary braids. A key technical result, proven in Section 11, is that strands that are individually boundary parallel are simultaneously boundary parallel. In Section 12 the main decomposition result is obtained for dual simple boundary braids. It is then extended in to general boundary braids and complexes in Section 13. Section 14 explains groupoid decomposition.

9 Configuration spaces

In this section we introduce a new combinatorial model for the configuration space of k points in a directed graph and, more generally, k points in an orthoscheme complex. In contrast to the configuration spaces for graphs used by Abrams and Ghrist, which are cubical [1; 10], our models are simplicial.

Definition 9.1 (products of graphs) Let Γ be a metric simplicial graph with oriented edges of unit length. Note that Γ can be regarded either as an ordered simplicial complex or as a cubical complex and we can form direct products of several copies of Γ in either context. The resulting spaces will be naturally isometric but their cell structures differ. We denote by $\text{PROD}_k(\Gamma, \square)$ and $\text{PROD}_k(\Gamma, \square)$ the orthoscheme product and cubical product, respectively, of k copies of Γ .

Example 9.2 (orthoscheme product spaces) If Γ is an oriented edge of unit length then $\text{PROD}_k(\Gamma, \square)$ is a unit k –cube while $\text{PROD}_k(\Gamma, \square)$ is the simplicial subdivision of the k –cube described in Example 7.2. If Γ is \mathbb{R} subdivided into edges of unit length then $\text{PROD}_k(\Gamma, \square)$ is the standard cubing of \mathbb{R}^k while $\text{PROD}_k(\Gamma, \square)$ is the standard orthoscheme tiling of \mathbb{R}^k described in Example 8.1.

Recall from Definition 1.1 that the (topological) configuration space of k points in Γ is $\Gamma^k - \text{DIAG}_k(\Gamma)$, where $\text{DIAG}_k(\Gamma)$ is the thick diagonal. To obtain a combinatorial configuration space, we take the full subcomplex supported on this subset with respect to either of the above cell structures. For the cubical structure this was first done by Abrams [1]. With the simplicial cell structure in place, our definition is completely analogous.

Definition 9.3 (orthoscheme configuration spaces) Let Γ be a metric simplicial graph with oriented edges of unit length. The *orthoscheme configuration space of k labeled points in an oriented graph Γ* is the full subcomplex $\text{CONF}_k(\Gamma, \square)$ of $\text{PROD}_k(\Gamma, \square)$ supported on $\text{PROD}_k(\Gamma, \square) - \text{DIAG}_k(\Gamma)$. Thus, a closed orthoscheme of $\text{PROD}_k(\Gamma, \square)$ lies in $\text{CONF}_k(\Gamma, \square)$ if and only if it is disjoint from $\text{DIAG}_k(\Gamma)$. The *orthoscheme configuration space of k unlabeled points* is $\text{UCONF}_k(\Gamma, \square) = \text{CONF}_k(\Gamma, \square) / \text{SYM}_k$.

Remark 9.4 (open questions) Since the simplicial structure for $\text{PROD}_k(\Gamma, \square)$ is a refinement of the cubical structure of $\text{PROD}_k(\Gamma, \square)$, the orthoscheme configuration space $\text{CONF}_k(\Gamma, \square)$ lies between the topological configuration space $\text{CONF}_k(\Gamma)$ and the cubical one, $\text{CONF}_k(\Gamma, \square)$. It is therefore interesting to compare it to either, specifically to determine under which conditions two of them are homotopy equivalent. In Example 10.2 we will see that the orthoscheme configuration and the cubical configuration space are generally not homotopy equivalent. Cubical configuration spaces are known to be nonpositively curved [1]. We do not know whether the same is true of orthoscheme configuration spaces. However, in the next section we will see that they are in the most basic case where Γ is a single oriented cycle.

10 Points on a cycle

For the purposes of this article, we are primarily interested in the orthoscheme configuration spaces of a single oriented cycle. In this section, we treat this case in detail.

Definition 10.1 (oriented cycles) An *oriented n -cycle* is a directed graph Γ_n with vertices indexed by the elements of $\mathbb{Z}/n\mathbb{Z}$ and a directed unit-length edge from i to $i + 1$ for each $i \in \mathbb{Z}/n\mathbb{Z}$. In illustrations we draw an oriented n -cycle so that it is the boundary cycle of a regular n -gon in the plane with its edges oriented counterclockwise. The graph Γ_n can be viewed as $\mathbb{R}/n\mathbb{Z}$. Similarly, the orthoscheme product space

$\text{PROD}_k(\Gamma_n, \square)$ is a k -torus $\mathbb{R}^k / (n\mathbb{Z})^k$, where \mathbb{R}^k carries the orthoscheme structure described in Example 8.1.

For the rest of the section Γ_n will denote an oriented n -cycle.

Example 10.2 (cubical vs orthoscheme) In both the cubical and orthoscheme cell structure of $(\Gamma_n)^n$ the only vertices not on the thick diagonal $\text{DIAG}_n(\Gamma_n)$ are the n -tuples where each entry is a distinct vertex of Γ_n . These form a single SYM_n -orbit. In the cubical structure no edge avoids $\text{DIAG}_n(\Gamma_n)$, so $\text{CONF}_n(\Gamma_n, \square)$ is a discrete space consisting of $n!$ points and $\text{UNCONF}_n(\Gamma_n, \square)$ is a single point.

In the orthoscheme structure, there are edges that are disjoint from $\text{DIAG}_n(\Gamma_n)$. These correspond to the motion where all n points rotate around the n -cycle simultaneously in the same oriented direction and they are longest edges in the top-dimensional orthoschemes. No other simplices avoid $\text{DIAG}_n(\Gamma_n)$. Thus $\text{CONF}_n(\Gamma_n, \square)$ has $(n-1)!$ connected components each of which is an oriented n -cycle. The unordered configuration space $\text{UNCONF}_n(\Gamma_n, \square)$ is a circle consisting of a single vertex and a single edge.

This illustrates that the cubical and the orthoscheme configuration spaces are generally not homotopy equivalent. Note that in this example, the topological configuration spaces are homotopy equivalent to the orthoscheme versions. However, reversing the orientation of a single edge makes the orthoscheme configuration spaces equal to the cubical ones and therefore not homotopy equivalent to the topological ones.

The purpose of the present section is the following result:

Proposition 10.3 (points on a cycle and curvature) *Let k and n be integers with $0 < k \leq n$. Each connected component of the universal cover of $\text{CONF}_k(\Gamma_n, \square)$ is isomorphic, as an orthoscheme complex, to the (k, n) -dilated column and is therefore $\text{CAT}(0)$. In particular, $\text{CONF}_k(\Gamma_n, \square)$ and $\text{UNCONF}_k(\Gamma_n, \square)$ are nonpositively curved.*

Proof First recall that $\text{PROD}_k(\tilde{\Gamma}_n, \square) \cong \widetilde{\text{PROD}}_k(\Gamma_n, \square)$ is \mathbb{R}^k with the structure described in Example 8.1 (where the tilde denotes the universal cover on both sides). Let C be the subcomplex obtained by removing the hyperplanes of the form $x_i - x_j = \ell$ with $i \neq j \in [k]$ and $\ell \in n\mathbb{Z}$ and taking the full subcomplex. Since these hyperplanes descend to the thick diagonal, we see that $\text{CONF}_k(\Gamma_n) \cong C / (n\mathbb{Z})^k$.

Notice that these hyperplanes include the ones used to define the (k, n) -dilated column in \mathbb{R}^k (Definition 8.9). Thus one connected component of C is a (k, n) -dilated column

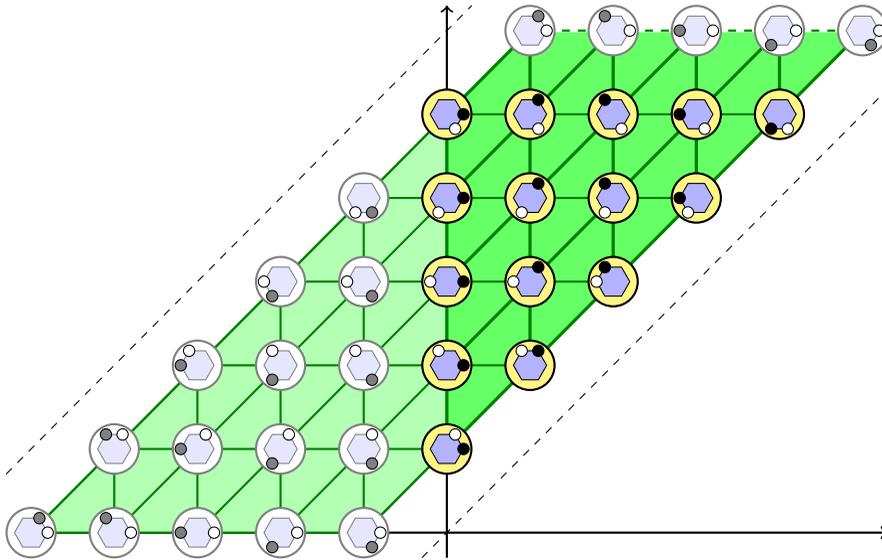


Figure 9: A portion of the $(2, 6)$ -diluted column, ie the full subcomplex of the orthoscheme complex of \mathbb{R}^2 on the vertices satisfying the strict inequalities $x < y < x + 6$. The vertex labels and the shaded regions are used to construct simplicial configuration spaces for 2 labeled points in a 6-cycle and for 2 unlabeled points in a 6-cycle.

in \mathbb{R}^k . Since SYM_k permutes the connected components, each component is a dilated column. Thus each component of C is $\text{CAT}(0)$ and, in particular, is simply connected. Since both $(n\mathbb{Z})^k$ and SYM_k act freely on C , both $C \rightarrow \text{CONF}_k(\Gamma_n)$ and $C \rightarrow \text{UCONF}_k(\Gamma_n)$ are covering maps and thus the configuration spaces $\text{CONF}_k(\Gamma_n)$ and $\text{UCONF}_k(\Gamma_n)$ are both nonpositively curved. \square

We now give two examples which illustrate Proposition 10.3.

Example 10.4 (2 labeled points in a 6-cycle) Figure 9 shows a portion of the infinite strip that is the $(2, 6)$ -diluted column in \mathbb{R}^2 . When this strip is quotiented by the portion of the $(6\mathbb{Z})^2$ -action on \mathbb{R}^2 that stabilizes this strip, its vertices can be labeled by two labeled points in a hexagon. The black dot indicates the value of its x -coordinate mod 6 and the white dot indicates the value of its y -coordinate mod 6. The rightmost vertex of the hexagon corresponds to $0 \pmod 6$ and the residue classes proceed in a counterclockwise fashion. The five hexagons on the y -axis, for example, have x -coordinate equal to $0 \pmod 6$ and y -coordinate ranging from 1 to $5 \pmod 6$. One component of the labeled orthoscheme configuration space $\text{CONF}_2(X, \square)$ is an

annulus formed by identifying the top and bottom edges of the region shown according to their vertex labels. Actually, in this case there is only $(2 - 1)! = 1$ component, so the annulus is the full labeled orthoscheme configuration space.

Example 10.5 (2 unlabeled points in a 6-cycle) The unlabeled orthoscheme configuration space is formed by further quotienting the labeled orthoscheme configuration space to remove the distinction between black and white dots. In particular, the 5 vertices shown on the horizontal line $y = 6$ are identified with the 5 vertices on the vertical line $x = 0$. This identification can be realized by the glide reflection sending (x, y) to $(y, x + 2)$, a map which also generates the unlabeled stabilizer of the $(2, 6)$ -dilated column. The heavily shaded region is a fundamental domain for this \mathbb{Z} -action and the unlabeled orthoscheme configuration space is then formed by identifying its horizontal and vertical edges with a half-twist forming a Möbius strip. The heavily shaded labels are the preferred representatives of the vertices in the quotient.

11 Boundary braids

We now come to our main object of study, boundary braids. The goal of this section is a key technical result saying that if certain strands of a braid can individually be realized as boundary parallel strands then they can be realized as boundary parallel strands simultaneously.

Lemma 11.1 (boundary parallel rotation braids) *Let $A \subseteq [n]$ not be a singleton and define*

$$B = \{b \mid b \notin A \text{ or } b + 1 \in A\}.$$

Then δ_A is a (B, \cdot) -boundary braid but not a (b, \cdot) -boundary braid for any $b \in [n] - B$.

Proof For $A = \emptyset$ and $A = [n]$ the statement is clear so we assume $2 \leq |A| < n$ from now on.

The first statement is straightforward by considering the standard representative of δ_A , given by constant-speed parametrization of each strand along the boundary of the subdisk D_A .

For the second statement, let $b \in [n] - B$. Then $b \in A$ but $b + 1 \notin A$. Fix the standard representative f of δ_A . Let $c = (b - 1) \cdot \text{PERM}(\delta_A)$, meaning that the strand f^{b-1} ends in c . Thus $c \in \{b - 1, b\}$. We consider a disk $U \subseteq I \times D$ that is bounded by the following four paths in $\partial(I \times D)$: the strand f_c^{b-1} , the strand f_{b+1}^{b+1} , the straight

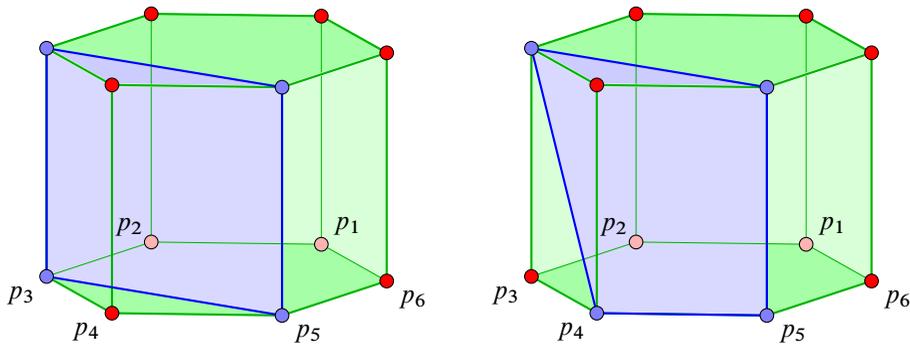


Figure 10: The two possible shapes of the disk U when $n = 6$ and $b = 4$. In either case, the strand f^4 must cross U an odd number of times.

line in $\{0\} \times D$ connecting $(0, p_{b-1})$ and $(0, p_{b+1})$, and the straight line in $\{1\} \times D$ connecting $(1, p_c)$ and $(1, p_{b+1})$; see Figure 10 for an example. If $c = b - 1$, then the disk U separates $(0, p_b)$ from all other vertices of the form $(1, p_i)$ except for $(1, p_b)$, but the strand f^b cannot end in b since $b \in A$. If $c = b$, then U separates $(0, p_b)$ from all other vertices. In either case, the strand f^b starts on one side of the disk and ends on the other side, and thus it intersects the disk (up to isotopy transversely) an odd number of times. Since the parity of the number of transverse intersections is preserved under homotopy of strands and strands which remain in the boundary have no such intersections, we may conclude that the (b, \cdot) -strand is not boundary parallel in any representative for δ_A . \square

Definition 11.2 (wrapping number) Let β be a (b, \cdot) -boundary braid and let f be a representative for which the image of f^b lies in the boundary of D_n . If we view the boundary of D_n as an n -fold cover $\pi: \partial D_n \rightarrow \mathbb{S}^1$ of the standard cell structure for \mathbb{S}^1 with one vertex and one edge, then boundary paths in ∂D_n that start and end at vertices of D_n may be considered as lifts of loops in \mathbb{S}^1 . More concretely, let $\varphi: \mathbb{R} \rightarrow \partial D_n$ be a covering map such that $\varphi(i) = p_i$ for each $i \in \mathbb{Z}$. Let \tilde{f}^b be any lift of f^b via this covering and define the *wrapping number* of the (b, c) -strand of f to be $w_f(b, c) = \tilde{f}^b(1) - \tilde{f}^b(0)$.

A slightly different description is as follows. Consider the diagram

$$\begin{array}{ccc}
 [0, 1] & \xrightarrow{f^b} & \partial D_n \\
 \downarrow /0 \sim 1 & & \downarrow \pi \\
 \mathbb{S}^1 & \xrightarrow{f_*^b} & \mathbb{S}^1
 \end{array}$$

where the left map is the quotient map that identifies 0 and 1. The map f_*^b is defined by commutativity of the diagram. Then the wrapping number of f^b is the winding number of f_*^b .

Notice that the wrapping number is n times what one would reasonably define as the winding number.

Lemma 11.3 *The wrapping number is well-defined.*

Proof Let f be a representative of the (b, \cdot) -boundary braid β for which f^b lies in the boundary; we temporarily denote the wrapping number by $w_f(b, \cdot)$ to indicate the presumed dependence on our choice of representative. If f and f' both represent β , then $f' \cdot f^{-1}$ is a representative for the trivial braid with $w_{f' \cdot f^{-1}}(b, b) = w_{f'}(b, \cdot) - w_f(b, \cdot)$. It therefore suffices to show that each strand in every representative of the trivial braid has wrapping number zero.

Now, let f be a representative of the trivial braid 1 for which f^b lies in the boundary, and suppose that the wrapping number $w_f(b, b) \neq 0$. If $f^{b'}$ is another strand in f , then we may obtain a map to the pure braid group PBRAID_2 by forgetting all strands except f^b and $f^{b'}$. The image β' of the trivial braid under this map may be written as an even power of δ_2 since δ_2^2 generates PBRAID_2 , and the wrapping numbers can then be related as $w_1(b, b) = \frac{1}{2}n w_{\beta'}(b, b)$. However, it is clear from the procedure of forgetting strands that the resulting braid in PBRAID_2 is trivial, and since every braid in PBRAID_2 has both strands boundary parallel, we know that $w_{\beta'}(b, b)$ is zero, and thus so is $w_1(b, b)$. Therefore, every representative for the trivial braid has trivial wrapping numbers, and we are done. \square

Lemma 11.4 *If β and γ are braids in BRAID_n such that β is a (b, b') -boundary braid and γ is a (b', \cdot) -boundary braid, then $\gamma\beta$ is a (b, \cdot) -boundary braid with wrapping number $w_{\gamma\beta}(b) = w_\beta(b) + w_\gamma(b')$.*

Proof If f and g are representatives of β and γ , respectively, such that $f_{b'}^b$ and $g^{b'}$ are boundary strands, then fg is a representative of $\beta\gamma$ such that $(fg)^b$ is a boundary strand. For this representative it is clear that the wrapping numbers add up in the described way. \square

Lemma 11.5 *Let $B \subseteq [n]$ and $\beta \in \text{BRAID}_n(B, \cdot)$. Then $\beta \in \text{FIX}_n(B)$ if and only if $w_\beta(b) = 0$ for all $b \in B$.*

Proof If $\beta \in \text{FIX}_n(B)$, then there is a representative f of β in which each (b, \cdot) -strand is fixed and thus $w_\beta(b, \cdot) = 0$.

For the other direction, we begin with the case that $B = \{b\}$. Let f be a representative of $\beta \in \text{BRAID}_n(B, \cdot)$ with the (b, \cdot) -strand in the boundary of D_n , and suppose that $w_\beta(b, \cdot) = 0$. Then the strand f^b begins and ends at the vertex p_b , and there is a homotopy $f(t)$ of f which moves every other strand off the boundary without changing f^b . That is, $f^{b'}(t) \notin \partial D_n$ whenever $b' \in [n] - \{b\}$ and $0 < t < 1$. After performing this homotopy, we note that $f(1)$ is a representative of β in which the (b, b) -strand has wrapping number 0 and there are no other strands in the boundary. Thus, there is a homotopy of this strand to the constant path, and therefore $\beta \in \text{FIX}_n(\{b\})$.

More generally, if $B \subseteq [n]$, then the set of braids $\beta \in \text{BRAID}_n(B, \cdot)$ with $w_\beta(b, \cdot) = 0$ for all $b \in B$ lie in the intersection of the fixed subgroups $\text{FIX}_n(\{b\})$. By Proposition 4.15, this is equal to $\text{FIX}_n(B)$ and we are done. □

Lemma 11.6 *Let b_1, \dots, b_k be integers satisfying $0 < b_1 < \dots < b_k \leq n$ and suppose that $\beta \in \text{BRAID}_n$ is a (b_i, \cdot) -boundary braid for every i . Then*

$$b_1 + w_\beta(b_1, \cdot) < b_2 + w_\beta(b_2, \cdot) < \dots < b_k + w_\beta(b_k, \cdot) < b_1 + w_\beta(b_1, \cdot) + n.$$

Proof It suffices to prove that

$$b_i + w_\beta(b_i, \cdot) < b_j + w_\beta(b_j, \cdot) < b_i + w_\beta(b_i, \cdot) + n$$

whenever $i < j$ or, in other words, that

$$w_\beta(b_j, \cdot) - w_\beta(b_i, \cdot) \in \{b_i - b_j + 1, \dots, b_i - b_j + n - 1\}.$$

As a first case, suppose both $w_\beta(b_i, \cdot)$ and $w_\beta(b_j, \cdot)$ are divisible by n . Then forgetting all but the (b_i, b_i) - and (b_j, b_j) -strands of β yields a pure braid $\beta' \in \text{PBRAID}_2$ which can be expressed as $\beta' = \delta_2^{2\ell}$ for some $\ell \in \mathbb{Z}$ since $\text{PBRAID}_2 = \langle \delta_2^2 \rangle$. Then

$$w_\beta(b_h, b_h) = \frac{1}{2}n w_{\beta'}(b_h, b_h) = n\ell$$

for each $h \in \{i, j\}$ and since every two-strand braid has simultaneously boundary parallel strands with equal wrapping numbers, we conclude that $w_\beta(b_i, b_i) = w_\beta(b_j, b_j)$. Therefore, $w_\beta(b_j, b_j) - w_\beta(b_i, b_i) = 0$, which satisfies the inequalities above.

For the general case, define

$$\gamma = \beta \delta_n^{-w_\beta(b_i, \cdot)}$$

and observe that $w_\gamma(b_i, \cdot) = 0$. Note that $w_\gamma(b_j, \cdot)$ is not congruent to $b_i - b_j \pmod n$; if it were, then the (b_i, \cdot) - and (b_j, \cdot) -strands of γ would terminate in the same vertex. Let e then be the representative of $w_\gamma(b_j, \cdot)$ modulo n that lies in the interval $\{b_i - b_j + 1, \dots, b_i - b_j + n - 1\}$. Then

$$\alpha = \gamma \delta_{[n]-\{b_i\}}^{-e}$$

has both its (b_i, \cdot) - and its (b_j, \cdot) -strand boundary parallel with wrapping numbers

$$w_\alpha(b_i, \cdot) = 0 \quad \text{and} \quad w_\alpha(b_j, \cdot) \equiv 0 \pmod n.$$

It follows from the case initially considered that the congruence on the right is actually an equality.

Tracing back, we see that $w_\gamma(b_j, \cdot) = e$ and

$$w_\beta(b_j, \cdot) - w_\beta(b_i, \cdot) = e \in \{b_i - b_j + 1, \dots, b_i - b_j + n - 1\},$$

as claimed. □

In what follows we will see that the inequalities given above are sharp in the sense that any tuple of numbers satisfying the hypotheses for Lemma 11.6 can be realized as the wrapping numbers for a braid.

Lemma 11.7 *Let b_1, \dots, b_k and w_1, \dots, w_k be integers satisfying $0 < b_1 < \dots < b_k \leq n$ and*

$$b_1 + w_1 < b_2 + w_2 < \dots < b_k + w_k < b_1 + w_1 + n.$$

There is a braid $\beta \in \text{BRAID}_n$ such that β is a (B, \cdot) -boundary braid for $B = \{b_1, \dots, b_k\}$ with $w_\beta(b_i, \cdot) = w_i$ for each i .

Proof First let $w = \min\{w_1, \dots, w_k\}$ and note that for any $\beta' \in \text{BRAID}_n(B, \cdot)$ the braid $\beta = \beta' \delta^w$ has wrapping numbers $w_\beta(b, \cdot) = w_{\beta'}(b, \cdot) + w$ by Lemma 11.4. It therefore suffices to show the claim in the case where $w = 0$ and thus all w_i are in $\{0, \dots, n - 2\}$ by the presumed inequalities. We assume this from now on.

Now the proof is by induction on $v = \max\{w_1, \dots, w_k\}$, the case $v = 0$ being trivial. Let $B_0 = \{b_i \mid w_i = 0\}$ and $B_{\geq 1} = \{b_i \mid w_i \geq 1\}$. We claim that there is no $b \in B_0$ with $b - 1 \in B_{\geq 1}$. If there were, then necessarily $b = b_i$ and $b - 1 = b_{i-1}$ for some index i (with the understanding that $b_0 = b_k$ and $b_{-1} = b_{k-1}$), but then

$$b_{i-1} + w_{i-1} \geq b_{i-1} + 1 = b_i = b_i + w_i$$

and this violates the assumption.

Let β' be a braid satisfying the claim for

$$b'_i = \begin{cases} b_i & \text{if } w_i = 0, \\ b_i + 1 & \text{if } w_i \geq 1, \end{cases}$$

and $w'_i = \max\{w_i - 1, 0\}$, where we note that such a braid exists by the induction hypothesis. Let $A = B_{\geq 1} \cup \{b + 1 \mid b \in B_{\geq 1}\}$. We claim that $\beta = \delta_A \beta'$ is as needed. Indeed, β is a (B, \cdot) -boundary braid by Lemma 11.1, and it has the wrapping numbers

$$w_\beta(b, \cdot) = \begin{cases} w_{\beta'}(b + 1, \cdot) + 1 & \text{if } b \in B_{\geq 1}, \\ w_{\beta'}(b, \cdot) & \text{if } b \in B_0. \end{cases}$$

Thus $w_\beta(b_i, \cdot) = w_i$ for every i . □

Proposition 11.8 *Let $\beta \in \text{BRAID}_n$. All boundary parallel strands of β are simultaneously boundary parallel. That is, if β is a (b, \cdot) -boundary braid for every $b \in B$, then it is a (B, \cdot) -boundary braid.*

Proof Let $B \subseteq [n]$ and suppose $\beta \in \text{BRAID}_n$ is a (b, \cdot) -boundary braid for each $b \in B$. If we write $B = \{b_1, \dots, b_k\}$ with $0 < b_1 < b_2 < \dots < b_k \leq n$, then the wrapping numbers $w_\beta(b_i, \cdot)$ satisfy the inequalities given by Lemma 11.6. Therefore by Lemma 11.7 there is a (B, \cdot) -boundary braid $\gamma \in \text{BRAID}_n$ with the same wrapping numbers as β . By Lemma 11.4, $\beta\gamma^{-1}$ is a (b, b) -boundary braid with $w_{\beta\gamma^{-1}}(b, b) = 0$ for each $b \in B$. Applying Lemma 11.5, we see that $\beta\gamma^{-1} \in \text{FIX}_n(B)$ and, in particular, it is a (B, B) -boundary braid. It follows that $\beta = (\beta\gamma^{-1})\gamma$ is a (B, \cdot) -boundary braid, as well. □

As a consequence of Proposition 11.8 we obtain the following proposition, which we state in analogy with Proposition 4.15:

Corollary 11.9 *Intersections of sets of boundary braids are sets of boundary braids. Concretely, if $B \subseteq [n]$, then*

$$\text{BRAID}_n(B, \cdot) = \bigcap_{b \in B} \text{BRAID}_n(\{b\}, \cdot)$$

and, equivalently,

$$\text{BRAID}_n(C \cup D, \cdot) = \text{BRAID}_n(C, \cdot) \cap \text{BRAID}_n(D, \cdot)$$

for any $C, D \subseteq [n]$.

12 Dual simple boundary braids

We start studying boundary braids in more detail by exploring the poset of boundary braids that are also dual simple.

Definition 12.1 (boundary braids) Let $B \subseteq [n]$. We denote the subposet of DS_n consisting of (B, \cdot) -boundary braids by $DS_n(B, \cdot)$. Notice that if $\beta \in DS_n$, then the wrapping numbers satisfy $w_\beta(b, \cdot) \in \{0, 1\}$ for each $b \in B$.

Definition 12.2 (boundary partitions) Let $B \subseteq [n]$. We say that a noncrossing partition $\pi \in NC_n$ is a (B, \cdot) -boundary partition if each $b \in B$ either shares a block with $b + 1$ (modulo n) or forms a singleton block $\{b\} \in \pi$. We denote by $NC_n(B, \cdot)$ the poset of all (B, \cdot) -boundary partitions.

Definition 12.3 (boundary permutations) Let $B \subseteq [n]$. We say that $\sigma \in NP_n$ is a (B, \cdot) -boundary permutation if for all $b \in B$, $b \cdot \sigma \in \{b, b + 1\}$ (modulo n). We denote the sets of (B, \cdot) -boundary permutations by $NP_n(B, \cdot)$.

These definitions fit together in the expected way.

Proposition 12.4 Let $B \subseteq [n]$. The natural identifications between NC_n , DS_n , and NP_n restrict to isomorphisms

$$DS_n(B, \cdot) \cong NC_n(B, \cdot) \cong NP_n(B, \cdot).$$

Proof Fix $B \subseteq [n]$. Let $\pi \in NC_n(B, \cdot)$ and consider the corresponding dual simple braid $\delta_\pi \in DS_n$. It is

$$\delta_\pi = \prod_{\substack{A \in \pi \\ |A| \geq 2}} \delta_A.$$

It is clear that δ_π is a (B, \cdot) -boundary braid if and only if each δ_A is. By Lemma 11.1 this is the case if $b \in A$ implies $b + 1 \in A$ for every $b \in B$ and every (nonsingleton) A . This matches the definition for π to be a (B, \cdot) -boundary partition.

Now let σ_π be the permutation corresponding to π . Note that $b \cdot \sigma_\pi = b$ if and only if $\{b\}$ is a block of π and that $b \cdot \sigma_\pi = b + 1$ if and only if b and $b + 1$ lie in the same block. From this it is clear that σ_π is a (B, \cdot) -boundary permutation if and only if π is a (B, \cdot) -boundary partition. \square

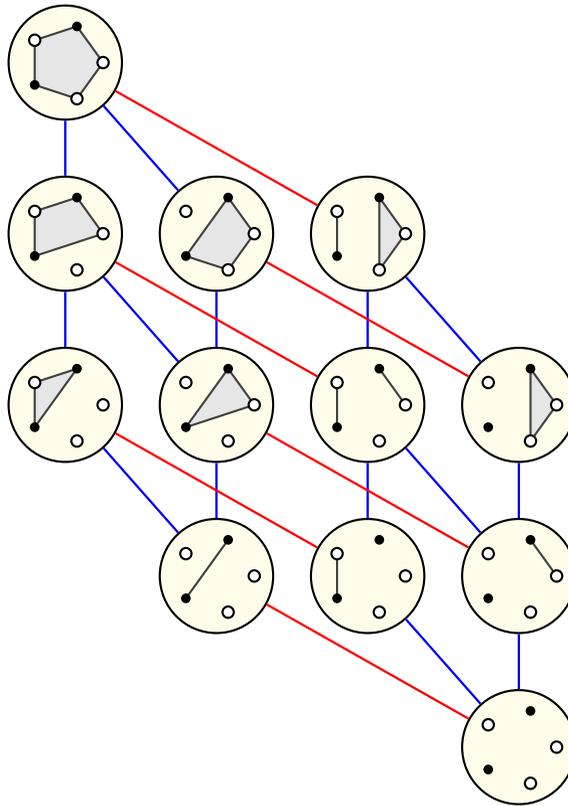


Figure 11: The poset of boundary partitions $NC_5(B, \cdot)$, where the upper-right vertex of each noncrossing partition is labeled by 1 and elements of $B = \{2, 4, 5\}$ are labeled by a white dot. The blue and red edge colors serve to illustrate the direct product structure described in Proposition 12.16.

Example 12.5 Let $B = \{2, 4, 5\} \subseteq [5]$. Then $NC_5(B, \cdot)$ is a subset of NC_5 with 12 elements, depicted in Figure 11.

We now define two maps on the posets described above. The first map takes a braid to one fixing B while the second takes it to a canonical braid with the same behavior on the specified boundary strands.

Definition 12.6 (FIX^B) For $B \subseteq [n]$ we define the map

$$FIX^B: NC_n \rightarrow NC_n(B, \cdot), \quad \pi \mapsto \{A - B \mid A \in \pi\} \cup \{\{b\} \mid b \in B\}.$$

Thus $FIX^B(\pi)$ is obtained from π by making each $b \in B$ a singleton block. We also denote by FIX^B the corresponding maps $NP_n \rightarrow NP_n(B, \cdot)$ and $DS_n \rightarrow DS_n(B, \cdot)$.

We call an element in the image of FIX^B a B -fix partition, braid or permutation, and we refer to the entire image by the shorthand notation $\text{FIX}(\text{NC}_n(B))$. We adopt similar notation for the analogous settings of $\text{NP}_n(B, \cdot)$ and $\text{DS}_n(B, \cdot)$.

Lemma 12.7 *Let $B \subseteq [n]$ and let $\pi \in \text{NC}_n$.*

- (1) $\text{FIX}^B(\pi)$ is the maximal B -fix element below π .
- (2) FIX^B preserves order.
- (3) $\text{FIX}^B(\pi) \leq \pi$.
- (4) FIX^B is idempotent, ie $(\text{FIX}^B)^2 = \text{FIX}^B$.
- (5) If $\alpha \in \text{DS}_n$ is B -fix and $\alpha\beta \in \text{DS}_n$, then $\text{FIX}^B(\alpha\beta) = \alpha \text{FIX}^B(\beta)$.

Proof The first statement is clear from the definition and the second, third and fourth statements follow from it. In the fifth statement $\alpha \text{FIX}^B(\beta) \leq \alpha\beta$ and $\alpha \text{FIX}^B(\beta)$ is B -fix, so $\alpha \text{FIX}^B(\beta) \leq \text{FIX}^B(\alpha\beta)$ by (1). Conversely, $\alpha^{-1} \text{FIX}^B(\alpha\beta) \leq \beta$ and $\alpha^{-1} \text{FIX}^B(\alpha\beta)$ is B -fix, so $\alpha^{-1} \text{FIX}^B(\alpha\beta) \leq \text{FIX}^B(\beta)$ by (1). □

Lemma 12.8 $\text{FIX}(\text{NC}_n(B)) = \text{DS}_n \cap \text{FIX}_n(B)$.

Proof A braid $\delta_\pi \in \text{DS}_n$ lies in $\text{FIX}(\text{NC}_n(B))$ if and only if every $b \in B$ is a singleton block of π if and only if $\delta_\pi \in \text{FIX}_n(B)$. □

Definition 12.9 (MOVE^B) For $B \subseteq [n]$ we define the map

$$\text{MOVE}^B: \text{DS}_n(B, \cdot) \rightarrow \text{DS}_n(B, \cdot)$$

by the equation $\delta_\pi = \text{FIX}^B(\delta_\pi)\text{MOVE}^B(\delta_\pi)$. We also denote by MOVE^B the corresponding maps $\text{NC}_n(B, \cdot) \rightarrow \text{NC}_n(B, \cdot)$ and $\text{NP}_n(B, \cdot) \rightarrow \text{NP}_n(B, \cdot)$. We refer to the image of MOVE^B by the shorthand $\text{MOVE}(\text{NC}_n(B, \cdot))$, with analogous notation for $\text{NP}_n(B, \cdot)$ and $\text{DS}_n(B, \cdot)$.

Lemma 12.10 *The map MOVE^B is well-defined, ie $\text{FIX}^B(\delta_\pi)^{-1}\delta_\pi$ is a dual simple braid for each $\pi \in \text{NC}_n(B, \cdot)$.*

Proof We know from Lemma 12.7 that $\text{FIX}^B(\pi) \leq \pi$. Thus by Proposition 3.7 there is a π' such that $\text{FIX}^B(\delta_\pi)\delta_{\pi'} = \delta_\pi$. Then $\text{MOVE}^B(\delta_\pi) = \delta_{\pi'}$. □

Example 12.11 Let $\sigma = (1\ 2\ 3\ 4\ 5\ 6)(7\ 8\ 9)$ and $B = \{2, 4, 5, 7\}$. Then we have $\text{FIX}^B(\sigma) = (1\ 3\ 6)(8\ 9)$ and $\text{MOVE}^B(\sigma) = (2\ 3)(4\ 5\ 6)(7\ 8)$. See Figure 12.

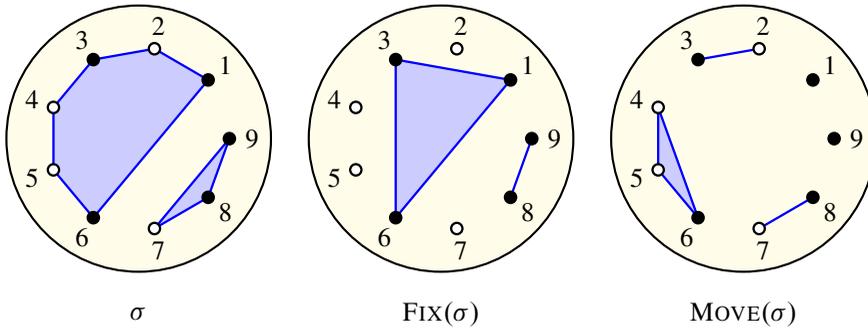


Figure 12: The noncrossing partitions corresponding to σ , $\text{FIX}^B(\sigma)$ and $\text{MOVE}^B(\sigma)$ as described in Example 12.11. The white dots form the set B .

Although the output of MOVE^B is less easily described than that of FIX^B , both maps satisfy many similar properties. Mirroring Lemma 12.7, we now describe several properties of the MOVE^B map.

Lemma 12.12 *Let $B \subseteq [n]$ and let $\pi \in \text{NC}_n(B, \cdot)$.*

- (1) $\text{MOVE}^B(\pi)$ is the minimal $\pi' \leq \pi$ such that $w_{\delta_{\pi'}}(b, \cdot) = w_{\delta_{\pi}}(b, \cdot)$ for each $b \in B$.
- (2) MOVE^B preserves order.
- (3) $\text{MOVE}^B(\pi) \leq \pi$.
- (4) MOVE^B is idempotent, ie $(\text{MOVE}^B)^2 = \text{MOVE}^B$.

Proof Suppose that $\delta_{\pi'}$ is a dual simple braid with $\pi' \leq \pi$ and wrapping numbers which satisfy $w_{\delta_{\pi'}}(b, \cdot) = w_{\delta_{\pi}}(b, \cdot)$ for all $b \in B$. Then $\delta_{\pi} \delta_{\pi'}^{-1}$ is a dual simple braid by Proposition 3.5 and is B -fix by Lemma 12.8. Hence, by Lemma 12.7, $\delta_{\pi} \delta_{\pi'}^{-1} \leq \text{FIX}^B(\delta_{\pi})$ and, by rearranging, we have that $\text{FIX}^B(\delta_{\pi})^{-1} \delta_{\pi} \leq \delta_{\pi'}$ and thus $\text{MOVE}^B(\delta_{\pi}) \leq \delta_{\pi'}$.

The second statement follows since for each $b \in B$, the wrapping number $w_{\beta}(b, \cdot)$ is monotone with respect to $\text{NC}_n(B, \cdot)$. The third and fourth statements are immediate from the first. □

The map MOVE^B is multiplicative in the following sense:

Lemma 12.13 *Let $B \subseteq [n]$. Let $\beta \in \text{DS}_n(B, B')$ and $\beta' \in \text{DS}_n(B', \cdot)$ be such that $\beta\beta' \in \text{DS}_n(B, \cdot)$. Then $\text{MOVE}^B(\beta\beta') = \text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta')$.*

Proof By Lemma 12.12(2), we know that $\text{MOVE}^B(\beta) \leq \text{MOVE}^B(\beta\beta')$ and thus by Proposition 3.5, $\text{MOVE}^B(\beta)^{-1}\text{MOVE}^B(\beta\beta')$ is a dual simple braid. By Lemma 11.4, this braid has B' -indexed wrapping numbers which are equal to those of $\text{MOVE}^{B'}(\beta')$. Hence, by Lemma 12.12, we know that

$$\text{MOVE}^{B'}(\beta') \leq \text{MOVE}^B(\beta)^{-1} \text{MOVE}^B(\beta\beta').$$

Equivalently, we have

$$\text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta') \leq \text{MOVE}^B(\beta\beta')$$

in the partial order on BRAID_n , and thus $\text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta')$ is a dual simple braid. Since this braid has the same B -indexed wrapping numbers as $\text{MOVE}^B(\beta\beta')$, another application of Lemma 12.12(1) tells us that

$$\text{MOVE}^B(\beta\beta') \leq \text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta').$$

Combining the above inequalities, we have

$$\text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta') \leq \text{MOVE}^B(\beta\beta') \leq \text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta')$$

and therefore $\text{MOVE}^B(\beta\beta') = \text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta')$. □

We now study the structure of $\text{FIX}(\text{NC}_n(B))$ and $\text{MOVE}(\text{NC}_n(B, \cdot))$ inside $\text{NC}_n(B, \cdot)$.

Remark 12.14 (minima and maxima) The identity braid is clearly the minimal element for both $\text{FIX}(\text{DS}_n(B))$ and $\text{MOVE}(\text{DS}_n(B, \cdot))$. Since FIX^B and MOVE^B are order-preserving maps, the maximal elements are $\text{FIX}^B(\delta)$ and $\text{MOVE}^B(\delta)$, respectively.

Lemma 12.15 Let $B \subseteq [n]$. Then $\text{FIX}^B(\text{MOVE}^B(\delta_\pi))$ is the identity braid for all $\pi \in \text{NC}_n(B, \cdot)$. In particular, the intersection of $\text{FIX}(\text{NC}_n(B))$ and $\text{MOVE}(\text{NC}_n(B, \cdot))$ contains only the discrete partition.

Proof Let $\pi \in \text{NC}_n(B, \cdot)$ be arbitrary. Then

$$\text{FIX}^B(\text{MOVE}^B(\delta_\pi)) = \text{FIX}^B(\text{FIX}^B(\delta_\pi)^{-1}\delta_\pi) = \text{FIX}^B(\delta_\pi)^{-1}\text{FIX}^B(\delta_\pi) = 1$$

by Lemma 12.7(5). The second claim follows directly from the wrapping number characterizations of $\text{FIX}(\text{NC}_n(B))$ and $\text{MOVE}(\text{NC}_n(B, \cdot))$ given in Definition 12.6 and Lemma 12.17. □

We now prove the main result of this section.

Proposition 12.16 *Let $B \subseteq [n]$. Then $\text{NC}_n(B, \cdot)$ is isomorphic to the direct product of the subposets $\text{FIX}(\text{NC}_n(B))$ and $\text{MOVE}(\text{NC}_n(B, \cdot))$.*

Proof Since FIX^B and MOVE^B are order-preserving maps on $\text{NC}_n(B, \cdot)$, the map which sends π to the element

$$(\text{FIX}^B(\pi), \text{MOVE}^B(\pi)) \in \text{FIX}(\text{NC}_n(B)) \times \text{MOVE}(\text{NC}_n(B, \cdot))$$

is order-preserving as well.

Suppose that $\text{FIX}^B(\pi) = \text{FIX}^B(\pi')$ and $\text{MOVE}^B(\pi) = \text{MOVE}^B(\pi')$. Then by definition of MOVE^B we have

$$\text{FIX}^B(\delta_\pi)^{-1} \delta_\pi = \text{FIX}^B(\delta_{\pi'})^{-1} \delta_{\pi'}$$

and thus $\delta_\pi = \delta_{\pi'}$, so the map is injective.

To see surjectivity, let $\pi \in \text{FIX}(\text{NC}_n(B))$ and $\pi' \in \text{MOVE}(\text{NC}_n(B, \cdot))$ be arbitrary and let $\beta = \delta_\pi \delta_{\pi'}$. Note that in the partial order on BRAID_n obtained by extending that of NC_n ,

$$\beta = \text{FIX}^B(\delta_\pi) \text{MOVE}^B(\delta_{\pi'}) \leq \text{FIX}^B(\delta_n) \text{MOVE}^B(\delta_n) \leq \delta_n$$

and thus β is a dual simple braid. Then

$$\text{FIX}^B(\beta) = \text{FIX}^B(\delta_\pi) \text{FIX}^B(\delta_{\pi'}) = \text{FIX}^B(\delta_\pi) = \delta_\pi$$

by Lemmas 12.7(5), 12.15 and 12.7(4), showing also that $\text{MOVE}^B(\beta) = \delta_{\pi'}$.

It remains to see that incomparable elements are mapped to incomparable elements. So suppose π and π' are such that $\text{FIX}^B(\pi) \leq \text{FIX}^B(\pi')$ and $\text{MOVE}^B(\pi) \leq \text{MOVE}^B(\pi')$. Then

$$\delta_\pi = \text{FIX}^B(\delta_\pi) \text{MOVE}^B(\delta_\pi) \leq \text{FIX}^B(\delta_{\pi'}) \text{MOVE}^B(\delta_\pi) \leq \text{FIX}^B(\delta_{\pi'}) \text{MOVE}^B(\delta_{\pi'}) \leq \delta_{\pi'}$$

by Proposition 3.7, so π and π' are comparable as well. □

We close by proving the following extension of Lemma 12.15.

Lemma 12.17 *Let $B \subsetneq [n]$. An element $\beta \in \text{MOVE}(\text{DS}_n(B, \cdot))$ is uniquely determined by the tuple $(w_b(\beta))_{b \in B}$. In particular, $\text{MOVE}(\text{DS}_n(B, B'))$ contains at most one element.*

Proof Suppose δ_π and $\delta_{\pi'}$ are dual simple braids in $DS_n(B, \cdot)$ with the property that $w_{\delta_\pi}(b, \cdot) = w_{\delta_{\pi'}}(b, \cdot)$ for each $b \in B$. Each wrapping number is either 0 or 1, and these can be characterized within π and π' by the fact that for each $b \in B$, either $\{b\}$ is a singleton or b and $b + 1$ (modulo n) share a block. This property is preserved under common refinement, so $\delta_{\pi \wedge \pi'}$ has the same B -indexed wrapping numbers as δ_π and $\delta_{\pi'}$. By definition, $\delta_{\pi \wedge \pi'} \leq \delta_\pi$ and $\delta_{\pi \wedge \pi'} \leq \delta_{\pi'}$, and we know by Lemma 12.12(2) that $\text{MOVE}^B(\delta_{\pi \wedge \pi'}) \leq \text{MOVE}^B(\delta_\pi)$ and $\text{MOVE}^B(\delta_{\pi \wedge \pi'}) \leq \text{MOVE}^B(\delta_{\pi'})$. Finally, since all three of these braids have the same B -indexed wrapping numbers, we may conclude by Lemma 12.12(1) that

$$\text{MOVE}^B(\delta_\pi) = \text{MOVE}^B(\delta_{\pi \wedge \pi'}) = \text{MOVE}^B(\delta_{\pi'})$$

and we are done. □

13 The complex of boundary braids

In this section, we describe the subcomplex of the dual braid complex which is determined by the set of boundary braids.

Definition 13.1 (complex of boundary braids) Let $B \subseteq [n]$. The *complex of (B, \cdot) -boundary braids*, denoted by $\text{CPLX}(\text{BRAID}_n(B, \cdot))$, is the full subcomplex of the complex $\text{CPLX}(\text{BRAID}_n)$ supported on $\text{BRAID}_n(B, \cdot)$.

The boundary strands of a (B, \cdot) -boundary braid define a path in the configuration space of $|B|$ points in ∂P . The following lemma is the combinatorial version of this statement. Recall from Definition 6.6 that we regard the edges of $\text{CPLX}(\text{BRAID}_n(B, \cdot))$ as labeled by elements of NC_n^* . Note also that the edge from $\beta \in \text{BRAID}_n(B, B')$ to $\beta' \in \text{BRAID}_n(B, B'')$ carries a label in $\text{NC}_n(B', B'')$.

Lemma 13.2 Let $B \subseteq [n]$. There is a surjective map

$$\text{BDRY}^B: \text{CPLX}(\text{BRAID}_n(B, \cdot)) \rightarrow \text{UCONF}_{|B|}(\Gamma_n, \square)$$

that takes $\beta \in \text{BRAID}_n(B, B')$ to B' .

Proof Let $\beta \in \text{BRAID}_n(B, B')$. An edge out of β labeled $\pi \in \text{NC}_n(B', \cdot)$ is taken to the edge of $\text{UCONF}_{|B|}(\Gamma_n, \square)$ out of B' that keeps $b \in B'$ fixed if it forms a singleton block of π and that moves it to $b + 1$ if it does not. The (boundary partition) condition

that b and $b + 1$ share a block of π for every $b \in B'$ ensures that this is compatible with how vertices are mapped. It also ensures that if $b + 1 \in B$ as well, then $b + 1$ also moves, and so the edge actually exists in $\text{UCONF}_{|B|}(\Gamma_n, \square)$.

To verify surjectivity we show that for any edge e from B' to B'' in $\text{UCONF}_{|B|}(\Gamma_n, \square)$ and any $\beta \in \text{BRAID}_n(B, B')$ there does in fact exist a $\pi \in \text{NC}_n(B', B'')$ such that the edge out of β labeled π is taken to e . If $B = B' = [n]$ this is achieved by the maximal element $\pi = \{[n]\}$. Otherwise the fact that the edge e exists means that for every interval $\{i, \dots, j\}$ (modulo n) of B either that same interval or the interval $\{i + 1, \dots, j + 1\}$ is in B' . The needed partition π is the one whose nonsingleton blocks are the intervals $\{i, \dots, j + 1\}$ where the second possibility happens. \square

Our goal is to show that BDRY^B is in fact a trivial bundle. To do so, we use the local decomposition results from Section 12 to obtain a splitting. More precisely, we want to construct a map SPLIT^B that makes the diagram

$$(13-1) \quad \begin{array}{ccc} \widetilde{\text{UCONF}}_{|B|}(\Gamma_n, \square) & \overset{\text{SPLIT}^B}{\dashrightarrow} & \text{CPLX}(\text{BRAID}_n(B, \cdot)) \\ & \searrow \text{COV} & \swarrow \text{BDRY}^B \\ & \text{UCONF}_{|B|}(\Gamma_n, \square) & \end{array}$$

commute (where COV denotes the covering map).

Lemma 13.3 *The map SPLIT^B in (13-1) exists. It is characterized (modulo deck transformations) by the property that if an edge in its image is labeled by $\pi \in \text{NC}_n(B', \cdot)$ then $\text{MOVE}^{B'}(\pi) = \pi$.*

Proof We will use the shorthands $\widetilde{\text{UCONF}}$, UCONF and CPLX . Suppose there is an edge from V' to V'' in $\widetilde{\text{UCONF}}$. Under COV it maps to an edge from B' to B'' in UCONF . If β' is a vertex above B' (with respect to BDRY^B) then any vertex $\beta'' = \beta' \delta_\pi$ with $\pi \in \text{NC}_n(B', B'')$ has the property that $\text{BDRY}^B(\beta'') = B''$. Our characterization states that if $\beta' = \text{SPLIT}^B(V')$ then $\text{SPLIT}^B(V'')$ should be the β'' with $\pi \in \text{MOVE}^{B'}(\text{NC}_n(B', B''))$, which is unique by Lemma 12.17. Similarly, if $\text{SPLIT}^B(V'')$ has already been defined, this uniquely characterizes $\text{SPLIT}^B(V')$.

If we choose a base vertex $V_0 \in \widetilde{\text{UCONF}}$ above B , declare that $\text{SPLIT}^B(V_0)$ is the vertex labeled by the identity braid and extend the definition according to the above rule, then we get a map that is defined everywhere since $\widetilde{\text{UCONF}}$ is connected (by edge paths).

It remains to see that this map is well-defined, ie that extensions along different edge paths agree. Since $\widetilde{\text{UCONF}}$ is simply connected, it suffices to check this along 2-simplices. This amounts to the requirement that if $\beta \in \text{BRAID}_n(B, B')$, $\delta_{\pi'} \in \text{MOVE}^{B'}(\text{DS}_n(B', B''))$ and $\delta_{\pi''} \in \text{MOVE}^{B''}(\text{DS}_n(B'', \cdot))$ are such that $\delta_{\pi'}\delta_{\pi''} \in \text{DS}_n(B', \cdot)$, then $\text{MOVE}^{B'}(\delta_{\pi'}\delta_{\pi''}) = \delta_{\pi'}\delta_{\pi''}$. This is true by Lemma 12.13. \square

Definition 13.4 (move complex) Let $B \subseteq [n]$. We denote the image of SPLIT^B by $\text{CPLX}(\text{MOVE}_n(B, \cdot))$ and call it the *move complex* associated to B . Its vertex set is denoted by $\text{MOVE}_n(B, \cdot)$.

Corollary 13.5 Let $B \subseteq [n]$. The corestriction of SPLIT to the move complex $\text{CPLX}(\text{MOVE}_n(B, \cdot))$ is an isomorphism. In particular, the move complex is a $\text{CAT}(0)$ subcomplex of the dual braid complex.

Proof The corestriction of SPLIT to $\text{CPLX}(\text{MOVE}_n(B, \cdot))$ is a covering map by Lemma 13.3. We need to show that it is injective. To see this, recall from Proposition 10.3 that $\widetilde{\text{UCONF}}_{|B|}(\Gamma_n, \square)$ is isomorphic to a dilated column. Let (b_1, \dots, b_k) with $0 < b_1 < \dots < b_k < n$ be a basepoint above B in the dilated column. Note that the edge from (b_1, \dots, b_k) to $(b_1 + \varepsilon_1, \dots, b_k + \varepsilon_k)$, with $(\varepsilon_i)_i \in \{0, 1\}^k$ is taken by SPLIT to a (B, \cdot) -boundary braid with wrapping numbers $(\varepsilon_1, \dots, \varepsilon_k)$. It follows more generally that coordinates of the dilated column relative to the basepoint correspond to wrapping numbers. In particular, SPLIT is injective. \square

If we develop the image of FIX^B in a similar way to how we just developed MOVE^B , we encounter a familiar structure.

Definition 13.6 (fix complex) The *fix complex* $\text{CPLX}(\text{FIX}_n(B))$ is the full subcomplex of $\text{CPLX}(\text{BRAID}_n)$ supported on $\text{FIX}_n(B)$.

Note that $\text{FIX}_n(B)$ is a parabolic subgroup and $\text{CPLX}(\text{FIX}_n(B))$ is an isomorphic copy of $\text{CPLX}(\text{BRAID}_{n-|B|})$. The fix complex does indeed relate to the decomposition of $\text{NC}_n(B, \cdot)$ in a similar way as the move complex:

Lemma 13.7 The edges out of $\beta \in \text{FIX}_n(B)$ that lie in $\text{CPLX}(\text{FIX}_n(B))$ are precisely those labeled by elements of $\text{FIX}(\text{NC}_n(B))$. In particular, the fiber of BDRY^B over B' is a union over translates of $\text{CPLX}(\text{FIX}_n(B'))$.

Proof The first claim is just Lemma 12.8. For the second claim note that the edges out of $\beta \in \text{BRAID}_n(B, B')$ that are collapsed to a point are precisely those labeled by $\text{FIX}^{B'}(\text{NC}_n(B', \cdot))$. Thus the fiber of BDRY^B over B' is the union over the $\beta \text{FIX}_n(B')$ for $\beta \in \text{BRAID}_n(B, B')$. \square

Proposition 13.8 *Let $B, B' \subseteq [n]$ and let $\beta \in \text{BRAID}_n(B, B')$ be arbitrary. There are unique braids $\text{FIX}^B(\beta) \in \text{FIX}_n(B)$ and $\text{MOVE}^B(\beta) \in \text{MOVE}_n(B, B')$ such that*

$$\beta = \text{FIX}^B(\beta) \text{MOVE}^B(\beta).$$

Moreover:

- (1) $\text{MOVE}^B(\text{MOVE}^B(\beta)) = \text{MOVE}^B(\beta)$.
- (2) $\text{MOVE}^B(\beta)^{-1} = \text{MOVE}^{B'}(\beta^{-1})$.
- (3) If $\beta' \in \text{BRAID}_n(B', \cdot)$ then $\text{MOVE}_n(\beta\beta') = \text{MOVE}_n^B(\beta) \text{MOVE}_n^{B'}(\beta')$.

Remark 13.9 By Lemmas 13.3 and 13.7, the braids $\text{FIX}^B(\beta)$ and $\text{MOVE}^B(\beta)$ coincide with the definitions in Section 12 if $\beta \in \text{DS}_n(B, \cdot)$.

Proof Uniqueness amounts to the statement that $\text{FIX}_n(B) \cap \text{MOVE}_n(B, \cdot) = \{1\}$, which follows from Corollary 13.5.

Let $\beta \in \text{BRAID}_n(B, B')$. To define $\text{MOVE}^B(\beta)$, consider the diagram (13-1) and let $V \in \widetilde{\text{UCONF}}_{|B|}(\Gamma_n, \square)$ be the base vertex with $\text{SPLIT}^B(V) = 1$. Let p be an edge path from 1 to β in $\text{CPLX}(\text{BRAID}_n(B, \cdot))$, let $q = \text{BDRY}^B(p)$ and let \tilde{q} be the path starting in V and covering q . We take $\text{MOVE}^B(\beta)$ to be the endpoint of q . The properties (2) and (3) follow from the corresponding properties of paths. Commutativity of (13-1) shows that if we did the same construction with β replaced by $\text{MOVE}^B(\beta)$, we would again end up at $\text{MOVE}^B(\beta)$, thus proving (1).

Putting $\text{FIX}^B(\beta) = \beta \text{MOVE}^B(\beta)^{-1}$, it remains to verify that $\text{FIX}^B(\beta) \in \text{FIX}_n(B)$. We compute

$$\begin{aligned} \text{MOVE}(\text{FIX}^B(\beta)) &= \text{MOVE}^B(\beta) \text{MOVE}^{B'}(\text{MOVE}^B(\beta)^{-1}) \\ &= \text{MOVE}^B(\beta) \text{MOVE}^B(\text{MOVE}^B(\beta))^{-1} \\ &= \text{MOVE}^B(\beta) \text{MOVE}^B(\beta)^{-1} \\ &= 1. \end{aligned}$$

This means that a path from 1 to $\text{FIX}^B(\beta)$ is mapped to a null-homotopic path in the complex $\text{UCONF}_{|B|}(\Gamma_n, \square)$. Thus $\text{FIX}^B(\beta)$ lies in the same component of the fiber of BDRY^B over B as 1, which by Lemma 13.7 is $\text{FIX}_n(B)$. \square

Lemma 13.10 *If $\beta \in \text{BRAID}_n(B, B')$ and $\beta' \in \text{BRAID}_n(B', \cdot)$, then*

$$\text{FIX}^B(\beta\beta') = \text{FIX}^B(\beta) \text{FIX}^{B'}(\beta')^{\text{MOVE}^B(\beta)^{-1}},$$

where we use the shorthand x^y to mean the conjugation $y^{-1}xy$.

Proof By Proposition 13.8, we have on one hand

$$\beta\beta' = \text{FIX}^B(\beta\beta') \text{MOVE}^B(\beta\beta') = \text{FIX}^B(\beta\beta') \text{MOVE}^B(\beta) \text{MOVE}^{B'}(\beta'),$$

and on the other hand

$$\beta\beta' = \text{FIX}^B(\beta) \text{MOVE}^B(\beta) \text{FIX}^{B'}(\beta') \text{MOVE}^{B'}(\beta').$$

Solving for $\text{FIX}^B(\beta\beta')$ proves the claim. □

We are now ready to prove the main result of this article.

Theorem 13.11 *Let $B \subseteq [n]$. The map*

$$\varphi: \text{BRAID}_n(B, \cdot) \rightarrow \text{FIX}_n(B) \times \text{MOVE}_n(B, \cdot), \quad \beta \mapsto (\text{FIX}^B(\beta), \text{MOVE}^B(\beta)),$$

induces an isomorphism of orthoscheme complexes

$$\text{CPLX}(\text{BRAID}_n(B, \cdot)) \cong \text{CPLX}(\text{FIX}_n(B)) \boxtimes \text{CPLX}(\text{MOVE}_n(B, \cdot))$$

which is, in particular, an isometry.

Proof The map φ is a bijection by Proposition 13.8.

For each $\beta \in \text{BRAID}_n(B, B')$, we may restrict the domain of φ to the subcomplex of simplices in $\text{CPLX}(\text{BRAID}_n(B, \cdot))$ with minimum vertex β and the image of φ to the subcomplex of simplices in the orthoscheme product

$$\text{CPLX}(\text{FIX}_n(B)) \boxtimes \text{CPLX}(\text{MOVE}_n(B, \cdot))$$

with minimum vertex labeled by $(\text{FIX}^B(\beta), \text{MOVE}^B(\beta))$. It suffices for us to show that this restriction of φ is an isomorphism of orthoscheme complexes, and since both complexes are flag complexes, we need only check this on the 1–skeleton.

The edges leaving β are parametrized by $\text{NC}_n(B', \cdot)$, which by Proposition 12.16 is isomorphic to

$$(13-2) \quad \text{FIX}^{B'}(\text{NC}_n(B', \cdot)) \times \text{MOVE}^{B'}(\text{NC}_n(B', \cdot)).$$

The edges out of $(\text{FIX}^B(\beta), \text{MOVE}^B(\beta))$ are parametrized by

$$(13-3) \quad \text{FIX}^B(\text{NC}_n(B, \cdot)) \times \text{MOVE}^{B'}(\text{NC}_n(B', \cdot)).$$

Recall that for each $\pi \in \text{NC}_n(B', \cdot)$ we have

$$\text{MOVE}^B(\beta\delta_\pi) = \text{MOVE}^B(\beta) \text{MOVE}^{B'}(\delta_\pi)$$

by Proposition 13.8 and

$$\text{FIX}^B(\beta\delta_\pi) = \text{FIX}^B(\beta) \text{FIX}^{B'}(\delta_\pi)^{\text{MOVE}^B(\beta)^{-1}}.$$

by Lemma 13.10. This shows that φ indeed induces an isomorphism between the posets (13-2) and (13-3), namely it is the identity on the second factor and conjugation by $\text{MOVE}^B(\beta)$ on the first. Since these posets are isomorphic, the given restriction of φ is an isomorphism of orthoscheme complexes and, by Proposition 7.10, this isomorphism is an isometry. □

Corollary 13.12 *Let $B \subseteq [n]$. If the complex $\text{CPLX}(\text{BRAID}_{n-|B|})$ is $\text{CAT}(0)$, then $\text{CPLX}(\text{BRAID}_n(B, \cdot))$ is $\text{CAT}(0)$ as well.*

Proof By Theorem 13.11, $\text{CPLX}(\text{BRAID}_n(B, \cdot))$ is isomorphic to the metric direct product of $\text{CPLX}(\text{FIX}_n(B))$, which is isomorphic to the dual braid complex $\text{CPLX}(\text{BRAID}_{n-|B|})$, and $\text{CPLX}(\text{MOVE}_n(B, \cdot))$, which is $\text{CAT}(0)$ by Corollary 13.5. The claim therefore follows from [7, Exercise II.1.16(2)]. □

14 The groupoid of boundary braids

We close with a more algebraic view on the results of the last section. We refer the reader to [13, Chapter 12; 8, Chapter II] for basic background on groupoids.

Definition 14.1 *The groupoid of boundary braids has as objects the finite subsets of $[n]$. The morphisms from B to B' are $\text{BRAID}_n(B, B')$ if $|B| = |B'|$ and empty otherwise. Composition is composition of braids.*

Remark 14.2 *To be precise one should say that morphisms are represented by boundary braids as a braid may at the same time be a (B, B') -boundary braid and a (C, C') -boundary braid, and thus represent two different morphisms. Since a morphism is uniquely determined by the braid and either its source or its target, we trust that no confusion will arise from this imprecision.*

Parabolic subgroups form a subgroupoid in a trivial way.

Definition 14.3 *The groupoid of fix braids has as its objects the finite subsets of $[n]$. The morphisms from B to B' are $\text{FIX}_n(B)$ if $B = B'$ and are empty otherwise.*

The groupoid of fix braids is normal in the following sense:

Lemma 14.4 *If $\beta \in \text{BRAID}_n(B, B')$ and $\beta' \in \text{FIX}_n(B)$, then $\beta^{-1}\beta'\beta \in \text{FIX}_n(B')$.*

Proof If β is a (b, b') -boundary braid then $w_\beta(b) = -w_{\beta^{-1}}(b')$. The claim now follows from Lemmas 11.4 and 11.5. \square

The corresponding quotient morphism is the map BDRY^B from Lemma 13.2 that takes a (B, B') -boundary braid to a (B, B') -path in the fundamental groupoid of $\text{UCONF}_{|B|}(\Gamma_n, \square)$. The upshot of the last section is that this map splits with image move braids.

Definition 14.5 *The groupoid of move braids has as objects finite subsets of $[n]$. The morphisms from B to B' are the braids $\text{MOVE}_n(B, B')$, which are images under SPLIT^B of (B, B') -paths.*

It follows from Proposition 13.8(2)–(3) that this is indeed a subgroupoid. Now the algebraic conclusion can be formulated as follows; see [16, Section 4] for a discussion of semidirect products.

Theorem 14.6 *The groupoid $\text{BRAID}_n(\cdot, \cdot)$ is a semidirect product*

$$\text{FIX}_n(\cdot) \rtimes \text{MOVE}_n(\cdot, \cdot).$$

Specifically,

- (1) every $\beta \in \text{BRAID}_n(B, B')$ decomposes uniquely as $\beta = \varphi\mu$ with $\varphi \in \text{FIX}_n(B)$ and $\mu \in \text{MOVE}_n(B, B')$;
- (2) if $\varphi\mu = \mu'\varphi'$ with $\mu, \mu' \in \text{MOVE}_n(B, B')$, $\varphi \in \text{FIX}_n(B)$ and $\varphi' \in \text{FIX}_n(B')$, then $\mu = \mu'$.

Proof The first statement is Proposition 13.8. For the second note that $\text{BDRY}^B(\varphi\mu) = \text{BDRY}^B(\mu'\varphi')$ since elements of $\text{FIX}_n(\cdot)$ are mapped trivially under BDRY^B . It follows that $\text{MOVE}^B(\mu'\varphi') = \text{MOVE}^B(\varphi\mu) = \mu$. \square

References

- [1] **A D Abrams**, *Configuration spaces and braid groups of graphs*, PhD thesis, University of California, Berkeley (2000) Available at <https://search.proquest.com/docview/304583880>
- [2] **D Armstrong**, *Generalized noncrossing partitions and combinatorics of Coxeter groups*, Mem. Amer. Math. Soc. 949, Amer. Math. Soc., Providence, RI (2009) MR

- [3] **E Artin**, *Theorie der Zöpfe*, Abh. Math. Sem. Univ. Hamburg 4 (1925) 47–72 MR
- [4] **J Birman, KH Ko, S J Lee**, *A new approach to the word and conjugacy problems in the braid groups*, Adv. Math. 139 (1998) 322–353 MR
- [5] **T Brady**, *A partial order on the symmetric group and new $K(\pi, 1)$'s for the braid groups*, Adv. Math. 161 (2001) 20–40 MR
- [6] **T Brady, J McCammond**, *Braids, posets and orthoschemes*, Algebr. Geom. Topol. 10 (2010) 2277–2314 MR
- [7] **MR Bridson, A Haefliger**, *Metric spaces of non-positive curvature*, Grundle. Math. Wissen. 319, Springer (1999) MR
- [8] **P Dehornoy, F Digne, E Godelle, D Krammer, J Michel**, *Foundations of Garside theory*, EMS Tracts in Math. 22, Eur. Math. Soc., Zürich (2015) MR
- [9] **S Eilenberg, N Steenrod**, *Foundations of algebraic topology*, Princeton Univ. Press (1952) MR
- [10] **R Ghrist**, *Configuration spaces and braid groups on graphs in robotics*, from “Knots, braids, and mapping class groups” (J Gilman, W W Menasco, X-S Lin, editors), AMS/IP Stud. Adv. Math. 24, Amer. Math. Soc., Providence, RI (2001) 29–40 MR
- [11] **T Haettel, D Kielak, P Schwer**, *The 6–strand braid group is CAT(0)*, Geom. Dedicata 182 (2016) 263–286 MR
- [12] **A Hatcher**, *Algebraic topology*, Cambridge Univ. Press (2002) MR
- [13] **P J Higgins**, *Notes on categories and groupoids*, Van Nostrand Reinhold Math. Stud. 32, Van Nostrand Reinhold, London (1971) MR
- [14] **J McCammond**, *Noncrossing partitions in surprising locations*, Amer. Math. Monthly 113 (2006) 598–610 MR
- [15] **RP Stanley**, *Enumerative combinatorics, I*, Wadsworth & Brooks, Monterey, CA (1986) MR
- [16] **S Witzel**, *Classifying spaces from Ore categories with Garside families*, Algebr. Geom. Topol. 19 (2019) 1477–1524 MR

Department of Mathematics and Statistics, Swarthmore College

Swarthmore, PA, United States

Department of Mathematics, University of California, Santa Barbara

Santa Barbara, CA, United States

Mathematical Institute, Gießen University

Gießen, Germany

mdoughe1@swarthmore.edu, jon.mccammond@math.ucsb.edu,

stefan.witzel@math.uni-giessen.de

<https://www.swarthmore.edu/NatSci/mdoughe1/>, <http://web.math.ucsb.edu/~jon.mccammond/>,

<https://www.math.uni-bielefeld.de/~switzel/>

Received: 23 January 2019

Revised: 3 February 2020