NONCROSSING PARTITION LATTICES FROM PLANAR CONFIGURATIONS

STELLA COHEN, MICHAEL DOUGHERTY, ANDREW D. HARSH, AND SPENCER PARK MARTIN

Abstract. The lattice of noncrossing partitions is well-known for its wide variety of combinatorial appearances and properties. For example, the lattice is rank-symmetric and enumerated by the Catalan numbers. In this article, we introduce a large family of new noncrossing partition lattices with both of these properties, each parametrized by a configuration of n points in the plane.

Introduction

For each subset P of the complex plane, a *noncrossing partition* of P is a way of dividing P into subsets with pairwise disjoint convex hulls. The collection of all noncrossing partitions of P, denoted $\mathrm{NC}(P)$, is a partially ordered set under refinement. When P is the vertex set for a convex n-gon, $\mathrm{NC}(P)$ is the classical noncrossing partition lattice NC_n introduced by Kreweras [Kre72]. Among other things, Kreweras showed that the size of NC_n is counted by the combinatorially ubiquitous Catalan numbers $C_n = \frac{1}{n+1}\binom{2n}{n}$ and, more specifically, the number of lattice elements with rank k is the Narayana number $N_{n,k} = \frac{1}{n}\binom{n}{k}\binom{n}{k-1}$. Since $N_{n,k} = N_{n,n-k}$, this further says that NC_n is a *rank-symmetric* lattice. In the fifty years since its definition, the noncrossing partition lattice has made countless appearances in algebraic and geometric combinatorics - see the survey articles [McC06] and [BBG$^+$19] for more information.

Returning to the more general case prompts a natural question: for which subsets $P \subset \mathbb{C}$ does the poset $\mathrm{NC}(P)$ have similar properties to NC_n? While some existing work studies the size of $\mathrm{NC}(P)$ in asymptotic and extremal cases (e.g. [SW06] [RW13]), similarities to NC_n seem uncommon in the literature. In our first main theorem, we introduce a convexity condition on P which guarantees that $\mathrm{NC}(P)$ has the same size as NC_n.

We say that a set of points $P \subset \mathbb{C}$ in general position has *Property* Δ_k if, for every convex subset $A \subseteq P$, the convex hull of A contains at most $|A| + k - 3$ elements of P in its interior.

Theorem A (Theorem 3.4)**.** *Let $P \subset \mathbb{C}$ be a set of n points with Property Δ_1. Then $\mathrm{NC}(P)$ is a rank-symmetric graded lattice, and the number of elements with rank k is the Narayana number $N_{n,k}$. In particular, $|\mathrm{NC}(P)| = C_n = |\mathrm{NC}_n|$.*

It is worth noting that if P contains a point which lies in the convex hull of the others, then $\mathrm{NC}(P)$ is not isomorphic to NC_n. Thus, Theorem A introduces a large new class of lattices with the same number of elements in each rank as the noncrossing partition lattice.

Date: June 13, 2023.

If the conditions on P are weakened to only require Property Δ_2, then the number of noncrossing partitions may increase compared to those in Theorem A. Nevertheless, some symmetry is preserved.

Theorem B (Theorem 5.4)**.** *Let $P \subset \mathbb{C}$ be a set of n points with Property Δ_2. Then $\mathrm{NC}(P)$ is a rank-symmetric graded lattice.*

Some of the techniques used in proving Theorems A and B can be interpreted in a stronger topological context. Recall that the *(unordered) configuration space of n points in $\mathbb{C}$* is the topological space of all n-tuples in $\mathbb{C}^n$ with distinct entries, considered up to permutations of the coordinates. Also, if P barely fails to be in general position (i.e. there is a single triple of collinear points in P) but otherwise satisfies Property Δ_k, we say that P satisfies the *weak Property Δ_k*.

Theorem C (Corollary 2.10)**.** *Let $k \in \{1, 2\}$. The set of all configurations which satisfy the weak Property Δ_k forms a connected subspace of the configuration space of n points in $\mathbb{C}$.*

We are unaware of any prior appearances of the space described in Theorem C. It would be interesting to know the homology of this space for each k and, in particular, whether it is a classifying space for the n-strand braid group.

The article is structured as follows. In Section 1, we introduce some background on posets and partitions, along with basic properties of $\mathrm{NC}(P)$. Section 2 concerns the transformation of configurations with Property Δ_k and includes the proof of Theorem C. We give the proof of Theorem A in Section 3, then introduce some technical details in Section 4 which help us to prove Theorem B in Section 5.

1. Noncrossing Partitions

To start, we establish some basic definitions and properties for partitions, posets, and configurations - see [Sta12] for a standard reference. Recall that a *partition* expresses a set S as the union of a collection of pairwise disjoint subsets of S (called *blocks*). The set of all partitions for a fixed set S forms a partially ordered set under refinement: one partition lies "below" another in the partial order if each block in the latter partition can be obtained as a union of blocks in the former. This partially ordered set is a *lattice* in the sense that each pair of elements has a unique meet and a unique join. Let $\Pi(S)$ denote the lattice of partitions for S; in the standard case where $S = \{1, \ldots, n\}$, the associated partition lattice is denoted Π_n.

The partition lattice $\Pi(S)$ is *bounded* in the sense that it contains a unique minimum element $\hat{0}$ (in which each block is a singleton) and a unique maximum element $\hat{1}$ (in which all of S belongs to the same block). The partition lattice is also *graded*: if $|S| = n$ and we let $bl(\pi)$ denote the number of blocks in a partition $\pi \in \Pi(S)$, then the map $\rho\colon \Pi(S) \to \mathbb{N}$ given by $\rho(\pi) = n - bl(\pi)$ is a *rank function* for this lattice. Note that the minimum $\hat{0}$ and maximum $\hat{1}$ have ranks 0 and $n-1$ respectively. The *atoms* and *coatoms* of this lattice are defined to be the elements of rank 1 and $n-2$ respectively.

Our main object of study in this article is a subposet of the partition lattice for a finite set of points in the complex plane.

Definition 1.1. Fix $P \subset \mathbb{C}$ with $|P| = n$. For any $A \subseteq P$, the *convex hull* of A, denoted $\textsc{Conv}(A)$, is the smallest convex subset of $\mathbb{C}$ which contains A. Note that $\textsc{Conv}(A)$ is a convex polygon with up to $|A|$ vertices. A partition of P is

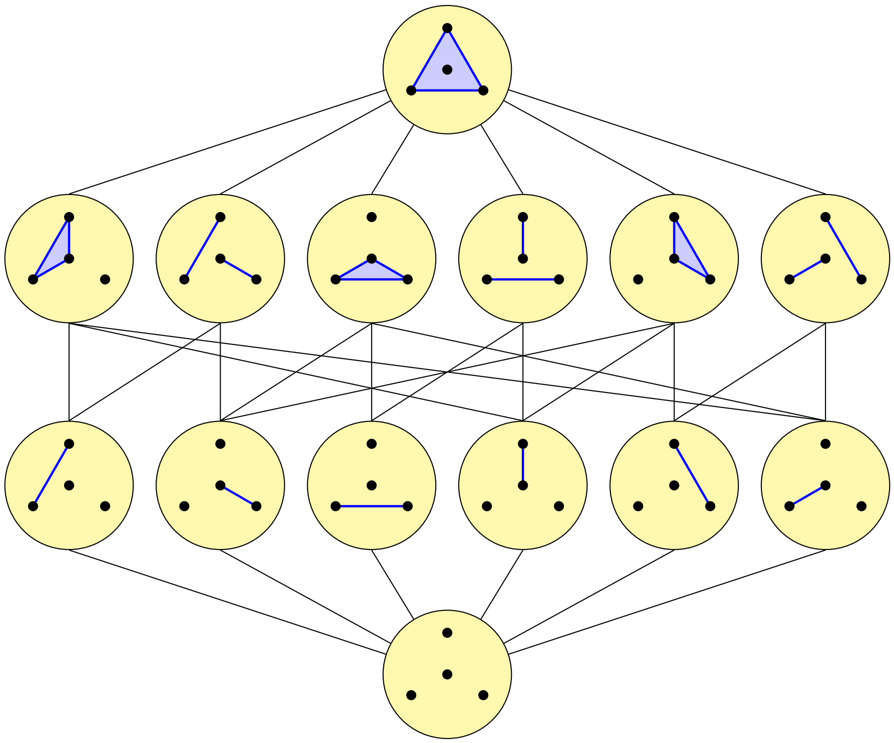

FIGURE 1. The lattice of noncrossing partitions for a particular arrangement of four points in $\mathbb{C}$

noncrossing if the convex hulls of its blocks are pairwise disjoint. The set of all noncrossing partitions for P forms a subposet of the partition lattice $\Pi(P)$, and we refer to this subposet as $\mathrm{NC}(P)$.

Example 1.2. Let $P = \{z_1, z_2, z_3, z_4\}$ be a set of points in $\mathbb{C}$ such that z_1, z_2, and z_3 form the vertices of a triangle which contains z_4 in its interior. Then every partition of P is noncrossing except for $\{\{z_1, z_2, z_3\}, \{z_4\}\}$, so the noncrossing partition lattice $\mathrm{NC}(P)$ has 14 elements, arranged according to the diagram in Figure 1.

As a poset, the noncrossing partitions of P inherit several useful properties from the larger partition lattice $\Pi(P)$.

Proposition 1.3. *If* $P \subset \mathbb{C}$ *with* $|P| = n$*, then* $\mathrm{NC}(P)$ *is a bounded graded lattice.*

Proof. Since the minimum and maximum elements of $\Pi(P)$ are noncrossing, we know that $\mathrm{NC}(P)$ is bounded. Furthermore, the rank function for $\Pi(P)$ descends to a rank function on the noncrossing partitions of P, so $\mathrm{NC}(P)$ is graded as well. To show that $\mathrm{NC}(P)$ is a lattice, we need only prove that $\mathrm{NC}(P)$ is a meet-semilattice (i.e. that each pair of elements in $\mathrm{NC}(P)$ has a unique meet) by a standard property of finite bounded posets [Sta12, Prop 3.3.1]. Indeed, if π_1 and π_2 are partitions in $\mathrm{NC}(P)$, then all refinements of π_1 and π_2 in the larger partition lattice $\Pi(P)$ are noncrossing as well. In particular, this means that the meet $\pi_1 \wedge \pi_2$ in $\Pi(P)$ is noncrossing and it follows that this partition is also the meet of π_1 and π_2 in $\mathrm{NC}(P)$. Therefore, $\mathrm{NC}(P)$ is a meet-semilattice and thus a lattice. $\square$

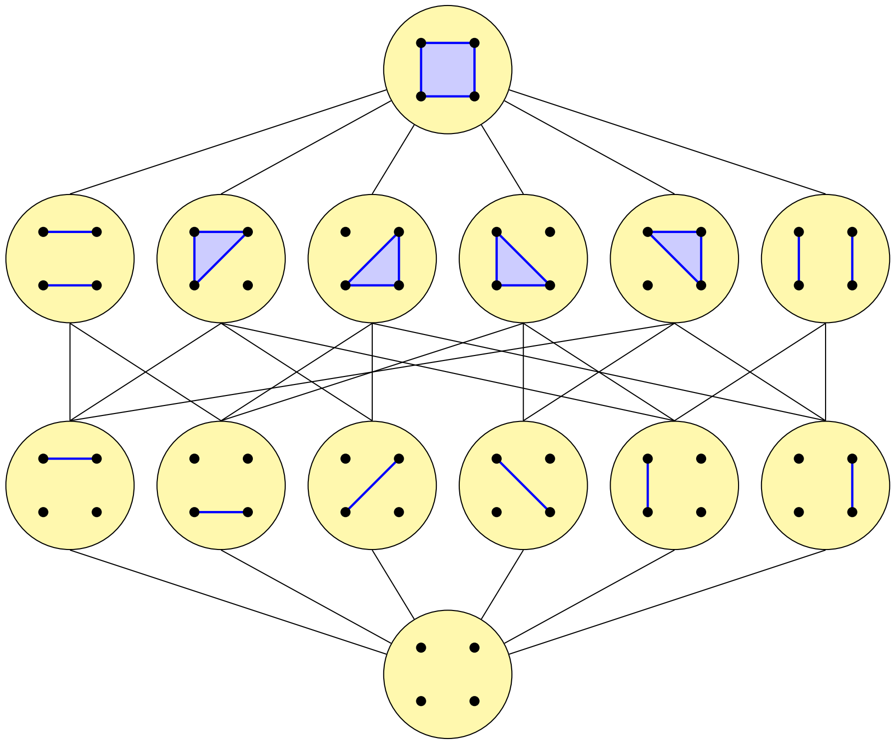

FIGURE 2. The classical noncrossing partition lattice NC_4

We close the section with a few important examples and some observations.

Example 1.4. If $P \subset \mathbb{C}$ with $|P| = n$ such that each point in P lies on the boundary of $\textsc{Conv}(P)$ (i.e. P is in *convex position*), then $\mathrm{NC}(P)$ is isomorphic to the classical *noncrossing partition lattice* NC_n, initially defined by Kreweras [Kre72] - see Figure 2. In addition to the properties outlined in Proposition 1.3, Kreweras showed that the size of NC_n is equal to the *Catalan number* $C_n = \frac{1}{n+1}\binom{2n}{n}$ and, in particular, the number of partitions in NC_n with k blocks is the *Narayana number* $N_{n,k} = \frac{1}{n}\binom{n}{k}\binom{n}{k-1}$. For more information on the combinatorial significance of these connections, see [Sta15].

Noting that $N_{n,k} = N_{n,n-k}$ for all $1 \leq k \leq n$, one can see that NC_n is a *rank-symmetric* lattice. In fact, the classical noncrossing partition lattice is *self-dual* in the sense that it admits a bijection $\alpha \colon \mathrm{NC}_n \to \mathrm{NC}_n$ with the property that $\pi_1 \leq \pi_2$ if and only if $\alpha(\pi_2) \leq \alpha(\pi_1)$ [SU91]. However, this stronger condition is rarely held by $\mathrm{NC}(P)$ more generally.

Remark 1.5. Nica and Speicher showed in 1997 that intervals in the noncrossing partition lattice NC_n are isomorphic to products of smaller noncrossing partition lattices [NS97]. With this in mind, we note that if P is a set of n points of $\mathbb{C}$ in general position and if some point of P lies in the convex hull of the others, then $\mathrm{NC}(P)$ has an interval which is isomorphic to the lattice described in Example 1.2, which cannot be expressed as a product of noncrossing partition lattices. Therefore,

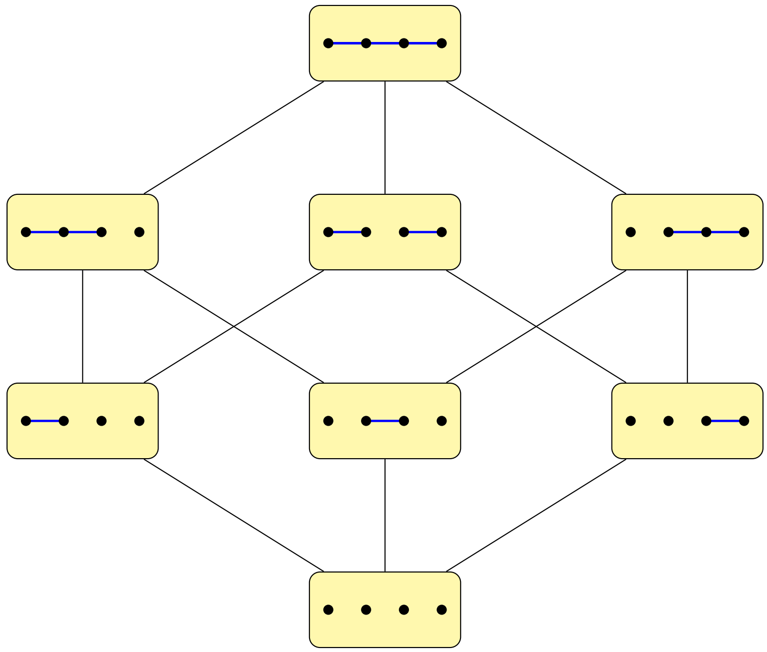

FIGURE 3. The Boolean lattice BOOL_{n-1} arises as the set of noncrossing partitions for a configuration of n collinear points.

$\text{NC}(P)$ is only isomorphic to NC_n if the elements of P form the vertices of a convex n-gon.

Example 1.6. If $P \subset \mathbb{C}$ with $|P| = n$ such that all points in P are collinear, then $\text{NC}(P)$ is isomorphic to the *Boolean lattice* BOOL_{n-1}, which is defined as the set of all subsets of a set with $n-1$ elements, partially ordered under inclusion. To see this, observe that each partition in $\text{NC}(P)$ is determined precisely by choosing a subset of the $n-1$ gaps between the n points; see Figure 3.

When $|P| = 4$, there are only four possibilities for $\text{NC}(P)$ (up to isomorphism), and three of them are depicted in the preceding figures. All three (indeed, all four) possess several useful lattice properties, including rank-symmetry, self-duality, and a simple counting formula. However, these properties do not always hold for larger sizes of P.

Example 1.7. If P consists of five points in general position with three points on the boundary of the convex hull and two points in the interior, then $|\text{NC}(P)| = 43$ (whereas $|\text{NC}_n| = 42$), although $\text{NC}(P)$ remains rank-symmetric. Furthermore, if P consists of six points in general position, arranged so that the three extremal points form an equilateral triangle and the three interior points form a shrunken equilateral triangle with the same center, then $\text{NC}(P)$ is not rank-symmetric: it has 15 atoms (rank 1) and coatoms (rank 4), but 55 elements at rank 2 and 57 elements at rank 3.

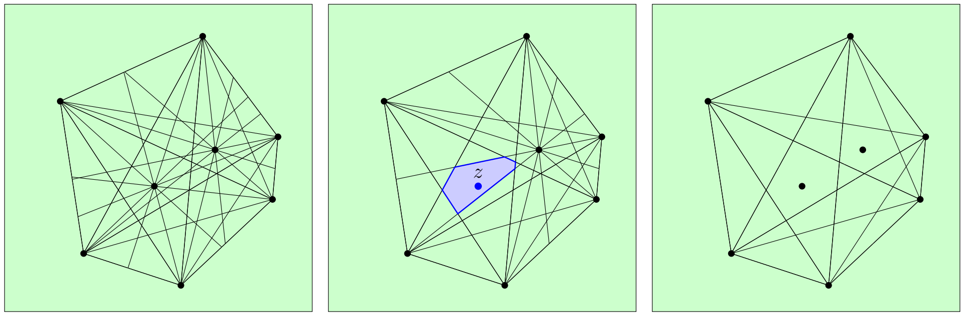

FIGURE 4. From left to right: an arrangement $\mathcal{A}$, the subarrangement $\mathcal{A}^z$ with the region R_z highlighted, and the subarrangement $\mathcal{A}^{\mathrm{ex}}$ of lines between pairs of extremal points. In each image, only the core of each line has been drawn.

2. CONFIGURATIONS

Before moving on to the main theorems, we introduce some tools for studying the geometry of planar configurations, by which we mean finite unordered sets of points in the Euclidean plane. We begin with some helpful terminology, partially inspired by [ER00]. Throughout this section, let P denote a configuration of n points in $\mathbb{C}$ in general position (i.e. no three points in P are collinear), unless otherwise specified.

Definition 2.1. Let $A \subseteq P$ and recall that $\textsc{Conv}(A)$ denotes the convex hull of the points in A. Define the *closure* $\overline{A}$ by $\textsc{Conv}(A) \cap P$ and the *interior* $\mathrm{int}(A)$ to be $\mathrm{int}(\textsc{Conv}(A)) \cap P$. We say that A is *convex* if $\mathrm{int}(A) \cap A$ is empty. Also, a point $p \in P$ is *internal* if $p \in \mathrm{int}(P)$ and *extremal* otherwise. We also define a technical condition: a configuration P in general position satisfies *Property* Δ_k if, for every convex subset A in P, the interior $\mathrm{int}(A)$ contains at most $(|A| - 3) + k$ points. Equivalently, P has Property Δ_k if P is in general position and each subset $B \subseteq P$ (not necessarily convex) has at most $\lfloor \frac{|B|-3+k}{2} \rfloor$ internal points. Finally, we say that P instead has the *weak Property* Δ_k if it satisfies the same convexity criteria, but has at most one instance of three collinear points.

It is worth noting that Property Δ_1 is equivalent to a simpler condition which is easier to check: for any $A \subseteq P$ with $|A| = 3$, we have $|\mathrm{int}(A)| \leq 1$. To see this, consider that each convex subset of k points in P forms the vertices of a convex k-gon, and any triangulation of this polygon consists of $k-2$ triangles; Property Δ_1 is equivalent to the requirement that each of those $k-2$ triangles has at most one point of P in its interior. Unfortunately, this does not generalize to Property Δ_k when $k > 1$.

If P satisfies Property Δ_k, then any small perturbation of P will also satisfy Property Δ_k since P is in general position. However, deformations which move a point in P across the line between two other points in P might not preserve this property. The main goal of this section is to provide some tools for moving points in P while preserving Property Δ_k. To start, we establish some terminology for the lines connecting points in P.

Definition 2.2. Each pair of distinct points in P determines a line in $\mathbb{C}$; let $\mathcal{A}$ denote the arrangement of the $\binom{n}{2}$ lines obtained in this way. If ℓ is the line obtained from the points z and w in P, then we say that z and w are the *endpoints* of ℓ and write $V(\ell) = \{z, w\}$. We also refer to the line segment between z and w as the *core* $c(\ell) = \textsc{Conv}(V(\ell))$ of ℓ. More generally, we write $V(\ell_1, \ldots, \ell_k)$ to mean the $2k$-element set of endpoints belonging to the lines $\ell_1, \ldots, \ell_k$ and we write $c(\ell_1, \ldots, \ell_k)$ to mean the convex hull $\textsc{Conv}(V(\ell_1, \ldots, \ell_k))$.

Definition 2.3. For each z in $\text{int}(P)$, let $\mathcal{A}^z$ denote the subarrangement of $\mathcal{A}$ obtained by deleting the lines which pass through z. Also, define $\mathcal{A}^{\text{ex}}$ to be the subarrangement of $\mathcal{A}$ which consists of all lines with two extremal endpoints, i.e. the intersection of all $\mathcal{A}^z$ for $z \in \text{int}(P)$. We associate two regions to each $z \in \text{int}(P)$: the connected component of $\mathbb{C} - \mathcal{A}^z$ containing z (which we denote R_z) and the connected component of $\mathbb{C} - \mathcal{A}^{\text{ex}}$ containing z (denoted R_z^{ex}). Note that each region is a convex polygon since it is a bounded subset of the plane determined by removing some number of half-planes. Finally, we say that a line in $\mathcal{A}$ is *separating* if it has points from $\text{int}(P)$ on either side of it, and a *boundary* line is one which contains a boundary edge for the convex hull $\textsc{Conv}(P)$.

For the sake of clarity, we will typically illustrate the line arrangement $\mathcal{A}$ by its intersection with the convex hull of P - see Figure 4 for an example.

Definition 2.4. A *move* is a bijection $m\colon P \to m(P)$ such that m fixes all of P except some element z, which is instead sent to a point $m(z)$ in the interior of a region adjacent to R_z. If ℓ is a line in the arrangement $\mathcal{A}$ which separates the regions R_z and $R_{m(z)}$, then we say that m *moves* z *across* ℓ. If both P and $m(P)$ satisfy Property Δ_k, we say that m is a Δ_k*-move*. Finally, note that m induces an isomorphism $m_*\colon \Pi(P) \to \Pi(m(P))$ by replacing z with $m(z)$ in each partition.

It is worth noting that for any $z \in P$, we can replace z with any other point in the region R_z without changing the isomorphism type of $\Pi(P)$, so moves on P can be described solely by the regions involved.

Definition 2.5. Let $z \in P$. If ℓ is a line in the arrangement $\mathcal{A}$ which contains a side of the region R_z, then we say that ℓ is *adjacent* to z. This line determines two open half-planes: $H^+_{z,\ell}$, which contains z, and $H^-_{z,\ell}$, which does not.

The following lemmas establish two useful cases of Δ_k-preserving moves.

Lemma 2.6. *If P has Property Δ_k and $m\colon P \to m(P)$ moves $z \in int(P)$ across a non-separating line in $\mathcal{A}^{ex}$, then $m(P)$ has Property Δ_k as well.*

Proof. Let $m\colon P \to m(P)$ be a move which takes z across a line ℓ in $\mathcal{A}^{\text{ex}}$, and let w_1 and w_2 be the endpoints of ℓ. Suppose for the sake of contradiction that $m(P)$ does not satisfy Property Δ_k; then there is a subset $A \subseteq m(P)$ with $|\text{int}(A)| > \lfloor \frac{|A|-3+k}{2} \rfloor$. Since P satisfies Property Δ_k, we know that $\overline{A}$ must contain $m(z)$ but not z.

First, note that the supposed bound on $|\text{int}(A)|$ precludes $m(z)$ from being an internal point of A. If it were internal, then A would necessarily contain w_1 and w_2, and A would therefore be a subset of $H^+_{m(z),\ell}$ since $z \notin \overline{A}$. However, $m(z)$ is the unique internal point of $m(P)$ in $H^+_{m(z),\ell}$ since we assumed ℓ was non-separating in the initial configuration P, so the inequality would not hold.

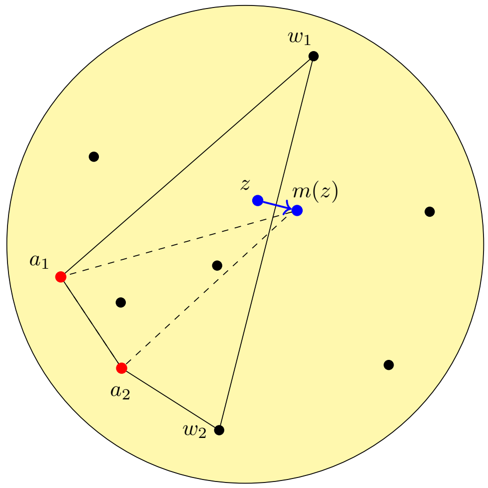

FIGURE 5. If $m\colon P \to m(P)$ is a move which takes z across a non-separating edge between extremal points w_1 and w_2, and if the triple $\{a_1, a_2, m(z)\}$ (indicated with dashed lines) has more than one point in its interior, then the quadrilateral $\{a_1, a_2, w_1, w_2\}$ (indicated with solid lines) has at least three points in its interior.

Thus, $m(z)$ is an extremal point of A. Define $B = (A - \{m(z)\}) \cup \{z, w_1, w_2\}$; then $|B| \leq |A| + 2$ and $|\mathrm{int}(B)| \geq |\mathrm{int}(A)| + 1$ (since z is an internal point for B but not A), and we can combine these to find the following chain of inequalities:

$$|\mathrm{int}(B)| \geq |\mathrm{int}(A)| + 1 > \left\lfloor \frac{|A| - 3 + k}{2} \right\rfloor + 1 \geq \left\lfloor \frac{|A| - 3 + k + 2}{2} \right\rfloor \geq \left\lfloor \frac{|B| - 3 + k}{2} \right\rfloor.$$

Since B is a subset of P, this contradicts our assumption that P has Property Δ_k - see Figure 5 for an example when $k = 1$. Therefore, $m(P)$ must satisfy Property Δ_k and we are done. □

Lemma 2.7. *If P has Property Δ_k and $m\colon P \to m(P)$ moves $z \in \mathit{int}(P)$ across a line in $\mathcal{A}^z$ with at least one internal endpoint, then $m(P)$ has Property Δ_k as well.*

Proof. Let ℓ be a line in $\mathcal{A}^z$ adjacent to z with endpoints w_1 and w_2, suppose that w_2 is an internal point of P, and let $m\colon P \to m(P)$ be the move which takes z across ℓ. As above, we suppose for the sake of contradiction that $m(P)$ does not satisfy Property Δ_k and can thus find a subset $A \subseteq m(P)$ with $|\mathrm{int}(A)| > \lfloor \frac{|A|-3+k}{2} \rfloor$ such that $\overline{A}$ contains $m(z)$ but not z.

Consider the three lines in $\mathcal{A}$ for which one endpoint is w_1 and the other belongs to the set $\{z, w_2, m(z)\}$. Since z, w_2, and $m(z)$ are internal points of P and both z and $m(z)$ are adjacent to ℓ, all three of these lines must pass through the same side of the polygon $\textsc{Conv}(P)$. Let $a_1, a_2 \in P$ be the extremal points which determine this side, where a_1 is on the same side of ℓ as z - see Figure 6 for an illustration.

Now, define $B = (A - \{m(z)\}) \cup \{z, w_1, w_2, a_1, a_2\}$. If $m(z) \in \mathrm{int}(A)$, then A must contain w_1 and w_2, and the fact that $z \notin A$ tells us that A does not contain any points in the half-plane $H^+_{z,\ell}$. Thus, in this case we have that $|B| \leq |A| + 2$ and $|\mathrm{int}(B)| \geq |\mathrm{int}(A)| + 1$ (since w_2 is internal for B but not A), which provides

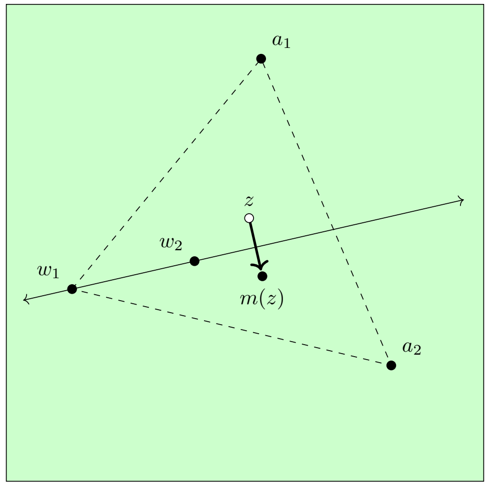

FIGURE 6. If w_2 is an internal point and m is a move which takes z across the line containing w_1 and w_2, then there are extremal points a_1 and a_2 such that the convex hull of a_1, a_2, and w_1 (depicted with dashed lines) has w_2, z, and $m(z)$ in its interior.

the same chain of inequalities as described in the proof of Lemma 2.6. Therefore, P does not satisfy Property Δ_k, which is a contradiction.

If $m(z)$ is instead an extremal point of A, then we see that $|B| \leq |A| + 4$ and $|\text{int}(B)| \geq |\text{int}(A)| + 2$, so we have a similar sequence of inequalities:

$$|\text{int}(B)| \geq |\text{int}(A)| + 2 > \left\lfloor \frac{|A| - 3 + k}{2} \right\rfloor + 2 \geq \left\lfloor \frac{|A| - 3 + k + 4}{2} \right\rfloor \geq \left\lfloor \frac{|B| - 3 + k}{2} \right\rfloor.$$

This also contradicts our assumption that P satisfies Property Δ_k, so we conclude that $m(P)$ must satisfy Property Δ_k. □

Note that the following lemma supposes only that P has Property Δ_2, so in particular it holds when P satisfies Property Δ_1 as well.

Lemma 2.8. *If P has Property Δ_2, then there is a point $z \in int(P)$ such that at least one side of the region R_z^{ex} belongs to a non-separating line in $\mathcal{A}^{ex}$.*

Proof. Let P be a configuration satisfying Property Δ_2 and let D be the convex hull of the regions R_z^{ex}, where z is an interior point of P. Let $\ell_1, \dots, \ell_k$ denote the lines (not necessarily in $\mathcal{A}$) which contain the k sides of D, arranged so that they appear in counter-clockwise order. If at least one ℓ_i contains a side of some region R_z^{ex}, then ℓ_i belongs to the arrangement $\mathcal{A}^{\text{ex}}$, and it follows that ℓ_i must be non-separating since all internal points lie on one side of it.

Suppose for the sake of contradiction that this is not the case. Then we can fix points $z_1, \dots, z_k \in \text{int}(P)$ such that for each i, the region R_{z_i} intersects the boundary of D at the point where the lines ℓ_i and ℓ_{i+1} (evaluated mod k) intersect. Since neither ℓ_i nor ℓ_{i+1} are in $\mathcal{A}^{\text{ex}}$, it follows that the half-planes $H^-_{z_i,\ell_i}$ and $H^-_{z_i,\ell_{i+1}}$ intersect in an unbounded region which contains an extremal point of P - see Figure 7 for an illustration.

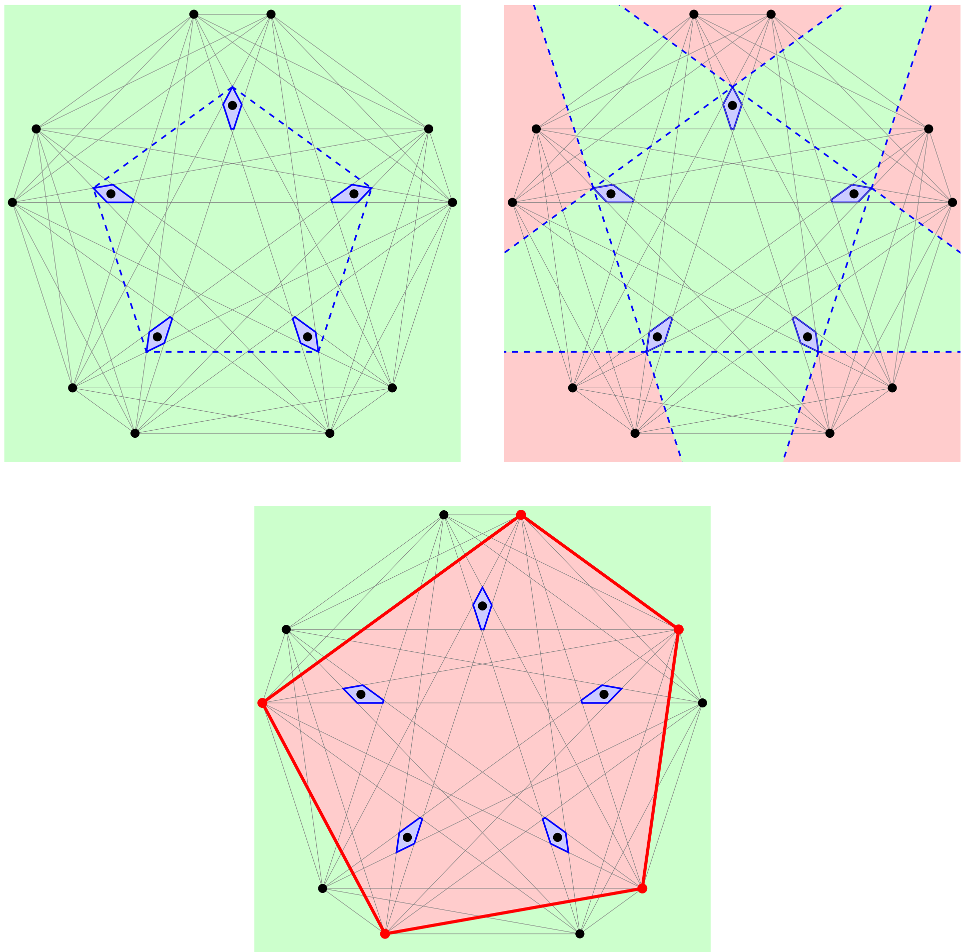

Figure 7. This configuration of 15 points has five internal points for which the corresponding regions have a convex hull (outlined with dashed blue lines in the upper left image) where no side of the convex hull contains a side of a region. By extending each side of the convex hull into a line, each of the five points determines a cone (shaded red in the upper right image) which contains at least one extremal point. Selecting one extremal point from each cone (highlighted red in the bottom image) yields a set of extremal points which contains the starting internal points and is at most as numerous, thus violating Property Δ_2.

If E is the set of k extremal points obtained in the manner above, then one can show that the convex hull $\textsc{Conv}(E)$ contains D. However, this means that $z_1, \ldots, z_k$ lie in the interior of E, which contradicts our assumption that P satisfies Property Δ_2. Therefore, at least one side of D must belong to a non-separating line in $\mathcal{A}^{\mathrm{ex}}$ and we are done. $\square$

Putting all of these tools together, we obtain a useful connectivity property.

Theorem 2.9. *Suppose P satisfies Property Δ_k, where $k \in \{1, 2\}$. Then there is a sequence of Δ_k-moves which transforms P into a convex configuration.*

Proof. We proceed by induction on the number of internal points for P. If P has no internal points, then P is already convex and we are done. Now, suppose the theorem holds for configurations with fewer internal points than P and define $\mathrm{nsnb}(P)$ to be the set of non-separating non-boundary lines in $\mathcal{A}^{\mathrm{ex}}$; we will prove that the theorem holds for P as well by a second induction on $|\mathrm{nsnb}(P)|$. If $|\mathrm{nsnb}(P)| = 0$, then each internal point of P is adjacent to a boundary line, and by Lemma 2.6, we can bring one of these internal points across a boundary line by a Δ_k-move, which reduces the number of internal points by one and allows us to apply the first inductive hypothesis.

Next, suppose the claim holds for all configurations Q with $|\mathrm{int}(Q)| = |\mathrm{int}(P)|$ and $|\mathrm{nsnb}(Q)| < |\mathrm{nsnb}(P)|$. By Lemma 2.8, we know that there is an internal point z in P such that one side of the region R_z^{ex} is contained in a non-separating line in $\mathcal{A}^{\mathrm{ex}}$. This means that while z may not be adjacent to a non-separating line, there is a finite sequence of moves across lines with at least one internal endpoint which takes z to a region which is adjacent to a non-separating line in $\mathcal{A}^{\mathrm{ex}}$. Each move across lines with an internal endpoint preserves Property Δ_k by Lemma 2.7, after which we can perform a Δ_k-move across the non-separating line by Lemma 2.6. This new configuration has fewer non-separating non-boundary lines, so by the second inductive hypothesis, it can be further transformed via Δ_k-moves into a convex configuration and the proof is complete. □

From a topological perspective, Theorem 2.9 can be interpreted as a statement about the configuration space of n points in the plane. To do so, recall that the *weak Property* Δ_k is a slight weakening of Property Δ_k which allows for one instance of three collinear points.

Corollary 2.10 (Theorem C)**.** *Let $k \in \{1, 2\}$. The set of all configurations which satisfy the weak Property Δ_k forms a connected subspace of the configuration space of n points in $\mathbb{C}$.*

Applying a move to a configuration P will certainly affect the noncrossing partition lattice $\mathrm{NC}(P)$, and possibly even its isomorphism type. To close this section, we introduce a natural map between the larger partition lattices.

Definition 2.11. Let $m\colon P \to m(P)$ be a move which brings the point $z \in P$ across the line ℓ in $\mathcal{A}^z$ and let w_1 and w_2 be the two points of P on ℓ. The *block-switching map* $\mathrm{BS}_m\colon \Pi(P) \to \Pi(m(P))$ is defined for each $\pi \in \Pi(P)$ as follows: if w_1 and w_2 share a block in π and z belongs to a different, non-singleton block, then define $\mathrm{BS}_m(\pi)$ to be the result of removing $\{w_1, w_2\}$ and $\{m(z)\}$ from their respective blocks in $m_*(\pi)$ and swapping them; otherwise, define $\mathrm{BS}_m(\pi) = m_*(\pi)$. See Figure 8 for an illustration. Note that BS_m is a rank-preserving bijection, but not an isomorphism.

3. Property Δ_1 and Catalan Numbers

Our strategy for proving Theorem A is to show that if P satisfies Property Δ_1, then applying a Δ_1-move to P does not change the size of $\mathrm{NC}(P)$. From here, applying Theorem 2.9 completes the proof. We begin with a useful lemma and some terminology, then give the proof of Theorem A.

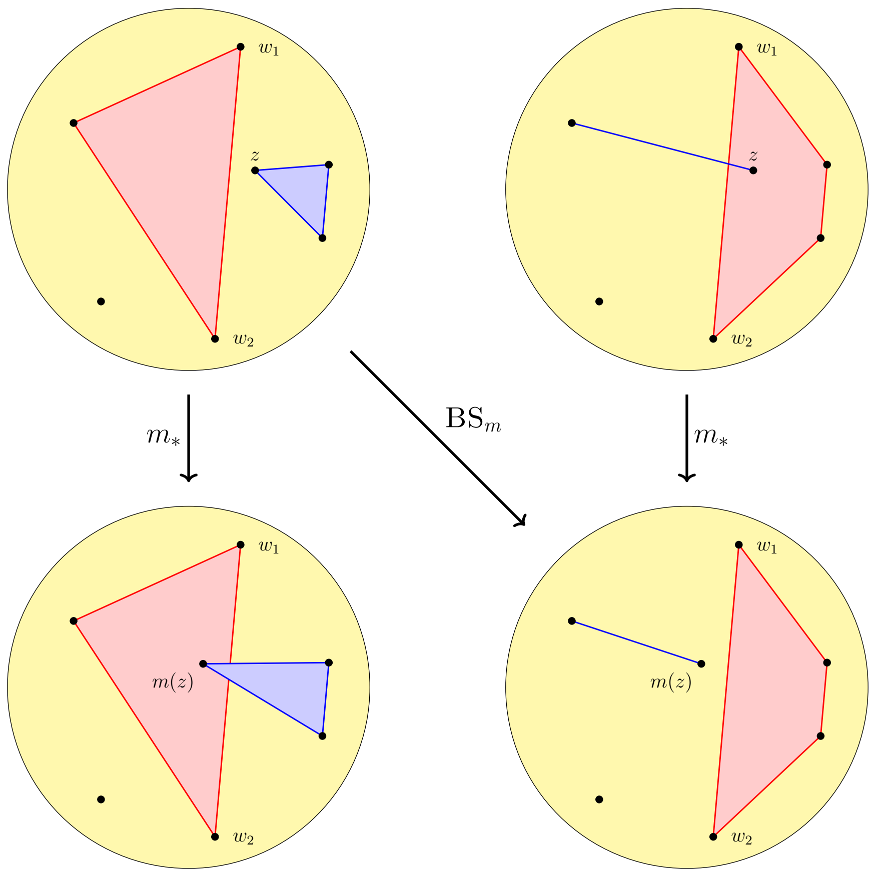

Figure 8. The block-switching map BS_m, compared to the induced map m_* for a fixed move m. In this example, BS_m takes an element of $\mathrm{NC}(P)$ to an element of $\mathrm{NC}(m(P))$.

Lemma 3.1. *Suppose P satisfies Property Δ_1 and let $z \in int(P)$. Then $R_z = R_z^{ex}$.*

Proof. We know by definition that R_z and R_z^{ex} are convex polygons with $R_z \subseteq R_z^{\mathrm{ex}}$, so we just need to show that each of its sides is a subset of a line in $\mathcal{A}^{\mathrm{ex}}$. Suppose that one of the sides for R_z is a subset of a line ℓ with endpoints u and v, where v (and possibly u as well) is internal. Then ℓ must not contain any other points in P (since P is assumed to be in general position), so it must eventually intersect the boundary of $\textsc{Conv}(P)$ in some edge between extremal vertices w_1 and w_2, and thus v lies within the triangle with vertex set $\{u, w_1, w_2\}$. Since ℓ was assumed to contain a side of R_z, we know that z must also be contained in the same triangle - see Figure 9 for an illustration. But this implies that P does not satisfy Property Δ_1, which is a contradiction. □

Remark 3.2. Let P be a configuration satisfying Property Δ_1, suppose that $z \in P$ is adjacent to a line ℓ in $\mathcal{A}^{\mathrm{ex}}$, and let π be a partition of P. Since P satisfies Property Δ_1, we know by Lemma 3.1 that ℓ contains an edge of the region R_z. In other words, there are no lines in the arrangement $\mathcal{A}^z$ which lie between ℓ and z.

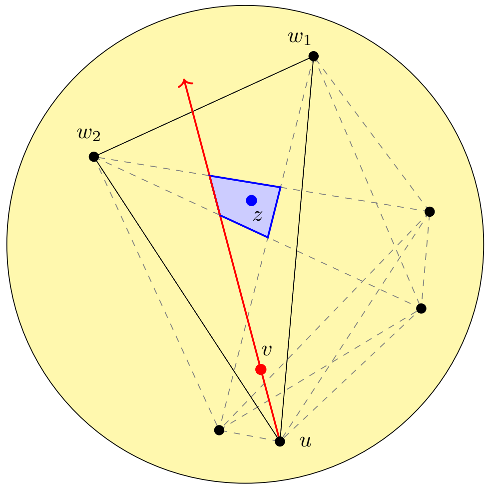

FIGURE 9. If the region R_z has a side which belongs to a line through an interior point v, then z and v must both belong to a common triangle.

This implies that if B_z and B_ℓ are the blocks in π containing z and the endpoints of ℓ respectively, then $\textsc{Conv}(B_z)$ and $\textsc{Conv}(B_\ell)$ are disjoint if and only if $B_z \cap H^-_{z,\ell} = \emptyset$ and $B_\ell \cap H^+_{z,\ell} = \emptyset$.

Definition 3.3. Let $m\colon P \to m(P)$ be a move. We say that $\pi \in \Pi(P)$ is *pre-m-noncrossing* if π is noncrossing, but its image $m_*(\pi) \in \Pi(m(P))$ is not. Similarly, we say that $\mu \in \Pi(m(P))$ is *post-m-noncrossing* if μ is noncrossing, but its preimage $m_*^{-1}(\mu) \in \Pi(P)$ is not. Then $\mathrm{NC}(m(P))$ can be obtained from $\mathrm{NC}(P)$ by the following procedure: remove partitions which are pre-m-noncrossing, apply m_* to all remaining elements, then add in the post-m-noncrossing partitions.

We are now ready to prove the main theorem of this section.

Theorem 3.4 (Theorem A)**.** *Let $P \subset \mathbb{C}$ be a configuration of n points which satisfies Property Δ_1. Then $\mathrm{NC}(P)$ is a rank-symmetric graded lattice, and the number of elements with rank k is the Narayana number $N_{n,k}$. In particular, $|\mathrm{NC}(P)| = C_n = |\mathrm{NC}_n|$.*

Proof. By Theorem 2.9, we know there is a sequence of Δ_1-moves which transforms P into a convex configuration, for which we know the lattice of noncrossing partitions is isomorphic to NC_n. Therefore, we need only show that if $m\colon P \to m(P)$ is a Δ_1-move, then the lattices $\mathrm{NC}(P)$ and $\mathrm{NC}(m(P))$ have the same number of elements in each rank.

Our strategy is to show that the block-switching map BS_m restricts to a rank-preserving bijection $\mathrm{NC}(P) \to \mathrm{NC}(m(P))$. Suppose that $m\colon P \to m(P)$ is a Δ_1-move which brings the point z across the line ℓ in $\mathcal{A}^{\mathrm{ex}}$, where the endpoints of ℓ are w_1 and w_2, and let $m(z) = y$.

Let π be a partition of P. Applying Remark 3.2 and the fact that $H^-_{y,\ell} = H^+_{z,\ell}$ and $H^+_{y,\ell} = H^-_{z,\ell}$, we have the following sequence of equivalences:

$$\begin{aligned}
\pi \text{ is pre-}m\text{-noncrossing} &\leftrightarrow \text{in } \pi\text{: } B_z \cap H^-_{z,\ell} \text{ and } B_\ell \cap H^+_{z,\ell} \text{ are empty;}\\
&\quad\ \text{either } B_y \cap H^-_{y,\ell} \text{ or } B_\ell \cap H^+_{y,\ell} \text{ is nonempty}\\
&\leftrightarrow \text{in } \mathrm{BS}_m(\pi)\text{: } B_\ell \cap H^-_{z,\ell} \text{ and } B_y \cap H^+_{z,\ell} \text{ are empty;}\\
&\quad\ \text{either } B_\ell \cap H^-_{y,\ell} \text{ or } B_z \cap H^+_{y,\ell} \text{ is nonempty}\\
&\leftrightarrow \text{in } \mathrm{BS}_m(\pi)\text{: } B_\ell \cap H^+_{y,\ell} \text{ and } B_y \cap H^-_{y,\ell} \text{ are empty;}\\
&\quad\ \text{either } B_\ell \cap H^+_{z,\ell} \text{ or } B_z \cap H^-_{z,\ell} \text{ is nonempty}\\
&\leftrightarrow \mathrm{BS}_m(\pi) \text{ is post-}m\text{-noncrossing}
\end{aligned}$$

Thus, the block-switching map BS_m induces a rank-preserving bijection between the pre-m-noncrossing partitions of P and the post-m-noncrossing partitions of $m(P)$, and we are done. $\square$

4. Skewers

Our proof of Theorem B follows a similar structure to that of Theorem A, but the weakening of our hypotheses from Property Δ_1 to Property Δ_2 means that Lemma 3.1 no longer holds. That is, it is possible that while some points in P are adjacent to non-separating lines in $\mathcal{A}$, none of these lines are in $\mathcal{A}^{\text{ex}}$. As a result, we must account for moves which take an interior point across a line with an interior endpoint, which in turn means we need to understand situations where a partition block overlaps with the line being moved across. With this in mind, we use this section to introduce the notion of "skewering" lines in $\mathcal{A}$ and explore the restrictions imposed by Property Δ_2. Throughout the rest of this section, let P be a configuration of n points in $\mathbb{C}$ and let $\mathcal{A}$ be the corresponding line arrangement.

Definition 4.1. For each pair of distinct lines $\ell_1, \ell_2 \in \mathcal{A}$, the intersection $\ell_1 \cap \ell_2$ can be classified into one of four different types (without loss of generality):

- if $\ell_1 \cap \ell_2 = \emptyset$, then ℓ_1 and ℓ_2 are *parallel*;
- if $\ell_1 \cap \ell_2$ lies in $c(\ell_1) \cap c(\ell_2)$, then ℓ_1 and ℓ_2 *intersect internally*;
- if $\ell_1 \cap \ell_2$ lies in neither $c(\ell_1)$ nor $c(\ell_2)$, then ℓ_1 and ℓ_2 *intersect externally*;
- if $\ell_1 \cap \ell_2$ lies in $c(\ell_2)$ but not $c(\ell_1)$, then ℓ_1 *skewers* ℓ_2.

As a useful shorthand, we write $\ell_1 \dashv \ell_2$ to mean that ℓ_1 skewers ℓ_2. When this is the case, note that the convex hull $c(\ell_1, \ell_2)$ contains one of the endpoints of ℓ_1 as an internal point - we refer to this as the *link vertex* for the skewer. Finally, a *skewering sequence* is a collection of lines $\ell_1, \ldots, \ell_k$ in $\mathcal{A}$ with $\ell_1 \dashv \ell_2 \dashv \cdots \dashv \ell_k$.

Definition 4.2. Suppose that P satisfies Property Δ_2. A *skewering tree* is a subset $T \subset \mathcal{A}$, together with two additional pieces of data - a special element $\ell_0 \in T$ called the *initial line* and a chosen closed half-plane bounded by ℓ_0, which we call the *positive side* of ℓ_0 - with the following properties:

- no two elements of T intersect internally;
- for all $\ell \in T$ with $\ell \neq \ell_0$, there is a unique line $\ell' \in T$ which skewers ℓ;
- no element of P is the link vertex for more than one skewer.

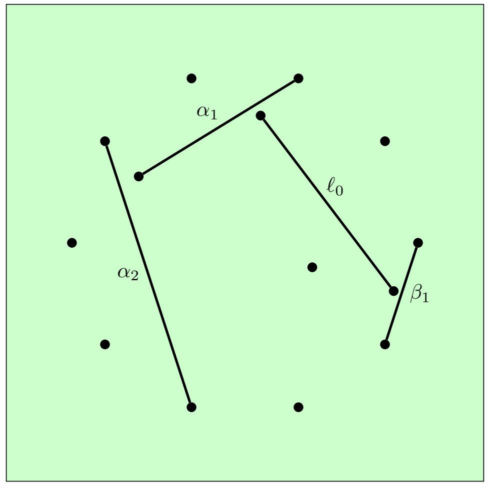

FIGURE 10. A skewering tree for a configuration of 14 points which satisfies Property Δ_2, but not Property Δ_1. The tree consists of two skewering sequences: $\ell_0 \dashv \alpha_1 \dashv \alpha_2$ and $\ell_0 \dashv \beta_1$. For the sake of visual clarity, only the cores of the four lines have been drawn.

We can also build a skewering tree inductively as follows: begin with an initial line $\ell_0 \in \mathcal{A}$, select one of the two half-planes bounded by ℓ_0 to be the positive side, and define $T = \{\ell_0\}$. Next, either stop here or add a line from $\mathcal{A}$ to T which is skewered by another element of T such that the requirements above remain satisfied. Repeat this process and stop at any point; the resulting set T is a skewering tree. See Figure 10 for an example. In either construction, we refer to the non-initial lines in T which do not skewer any other lines as *leaves*. Finally, we say that a skewering tree is *maximal* if it is not properly contained in any other skewering tree in $\mathcal{A}$.

The following lemma places fairly strong restrictions on the planar structure of skewering trees in configurations which satisfy Property Δ_2.

Lemma 4.3. *Suppose P satisfies Property Δ_2, let $T \subset \mathcal{A}$ be a skewering tree, and let ℓ, ℓ_1, and ℓ_2 be distinct lines in T such that ℓ is a leaf and $\ell_1 \dashv \ell_2$. Then ℓ does not intersect the convex hull $c(\ell_1, \ell_2)$.*

Proof. First, suppose that ℓ is the last element in a skewering sequence which includes ℓ_1 and ℓ_2, i.e. that there is a skewering sequence $\ell_0 \dashv \alpha_1 \dashv \cdots \alpha_k$ such that $\ell = \alpha_k$, and $\alpha_i = \ell_1$ and $\alpha_{i+1} = \ell_2$ for some i. The set $V(\{\ell_0, \alpha_1, \ldots, \alpha_k\})$ then consists of $2(k+1)$ points, of which at least k are internal: one for each non-leaf element in the skewering sequence. If ℓ were to intersect the convex hull $c(\ell_1, \ell_2)$, then one of the endpoints of ℓ would lie in the interior of the convex hull $c(\ell_1, \ell_2, \ell)$, but that would mean that the $2(k+1)$-element set described above has $k+1$ internal points, which contradicts our assumption that P satisfies Property Δ_2.

On the other hand, suppose that ℓ is in a different skewering sequence than ℓ_1 and ℓ_2. That is, suppose that T contains the sequences $\ell_0 \dashv \alpha_1 \dashv \cdots \dashv \alpha_k$ and $\ell_0 \dashv \beta_1 \dashv \cdots \dashv \beta_m$, where $\alpha_k = \ell$, and $\beta_i = \ell_1$ and $\beta_{i+1} = \ell_2$ for some i. Then $V(\{\ell_0, \alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_m\})$ is a set of $2(k+m+1)$ points in P, of which at least $k+m$ are internal. Once again, if ℓ intersected the convex hull $c(\ell_1, \ell_2)$, this

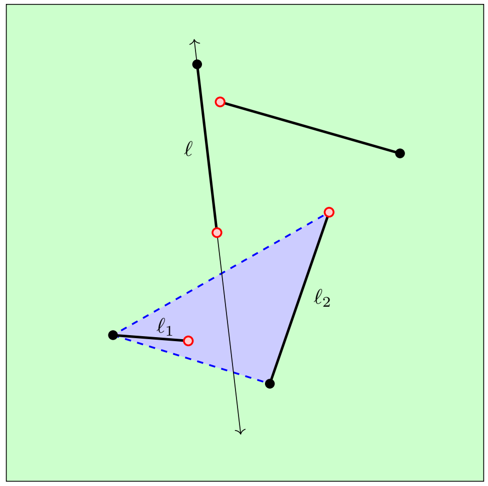

FIGURE 11. If ℓ intersects the convex hull of ℓ_1 and ℓ_2 (denoted in blue), then the set of all endpoints has half of its elements in the interior (shown here as unfilled red dots).

would imply that our set of $2(k+m+1)$ points has $k+m+1$ internal points, which again would violate Property Δ_2. In both cases, we see that ℓ does not intersect $c(\ell_1, \ell_2)$. □

Note that one may "prune" a skewering tree by iteratively removing leaves, and the result remains a skewering tree at each step. Thus, Lemma 4.3 further shows that the non-leaf elements of a skewering tree are similarly constrained.

Definition 4.4. The union of all lines in a skewering tree T forms a subset of the plane with the structure of an unbounded graph, consisting of vertices, line segments, rays, and lines (one may equivalently view this as a graph embedding on the 2-dimensional sphere, viewed as the stereographic projection of the plane); let Γ_T denote the unbounded graph obtained from this by removing any ray which does not contain the core of its corresponding line in T. We refer to Γ_T as the *planar realization* of T, observing that Γ_T is an acyclic graph, i.e. a tree. Note that while a pair of rays in Γ_T might overlap, Lemma 4.3 implies that no ray in Γ_T has a transverse intersection with a line segment.

Definition 4.5. For each non-initial line ℓ in T, the core $c(\ell)$ is contained in the boundary for exactly one region of the complement $\mathbb{C} - \Gamma_T$ (since the other side of ℓ contains the core of the line which skewers it). The *cell* associated to ℓ, denoted C_ℓ, is the union of the interior of this region together with the line segment or ray which contains $c(\ell)$. Note that C_ℓ is a convex (possibly unbounded) subset of the plane which contains exactly one side of its boundary. The initial line ℓ_0 bounds two regions and thus corresponds to two cells: $C^+_{\ell_0}$, which contains the core $c(\ell_0)$ and belongs to the positive side of ℓ_0, and $C^-_{\ell_0}$, which does not. See Figure 12 for an illustration. For any cell $C^\pm_\ell$, we write $V(C^\pm_\ell)$ to mean the intersection of P with $C^\pm_\ell$.

Lemma 4.6. *The cells of a skewering tree are pairwise disjoint.*

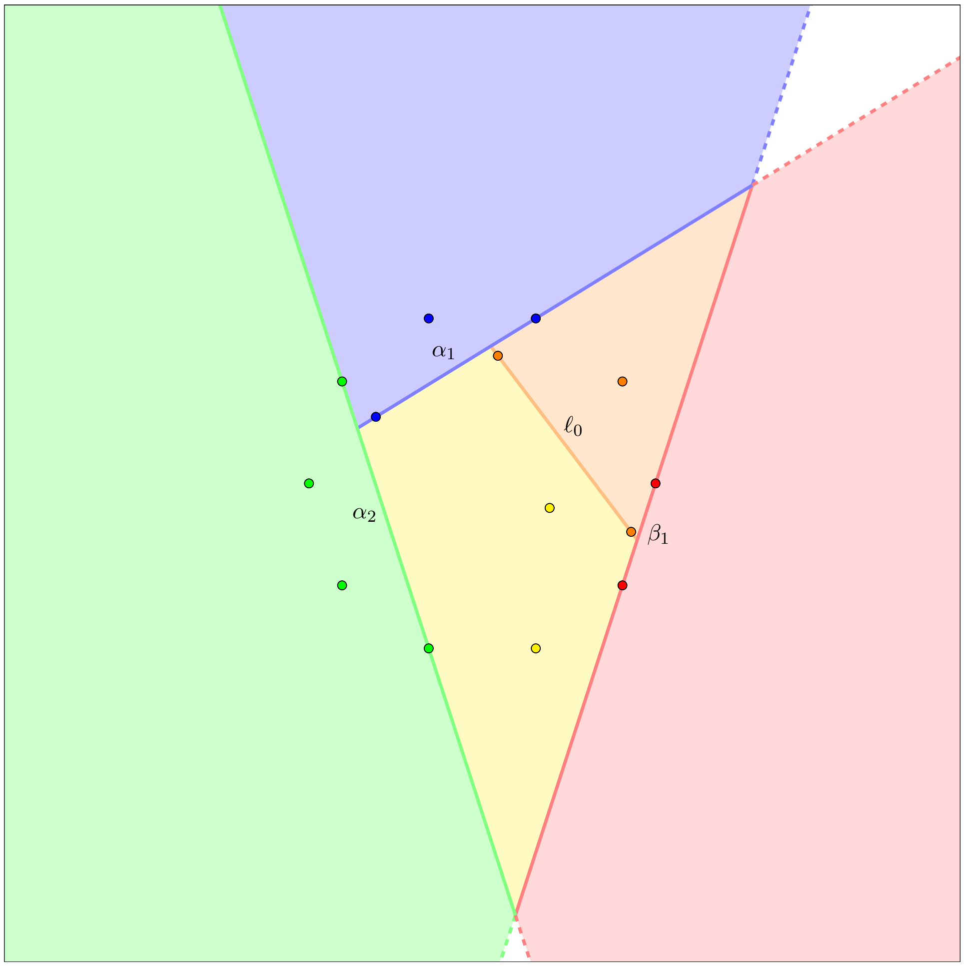

FIGURE 12. Cells associated to the skewering tree in Figure 10. In this case, the positive side of the initial line ℓ_0 is chosen to be the upper-right side; the corresponding cell is shaded orange.

Proof. The interior of a cell for a skewering tree is a connected component of the complement $\mathbb{C} - \Gamma_T$, so it follows that the interiors of two cells overlap if and only if the interiors are identical. What remains to be shown is that no connected component of $\mathbb{C} - \Gamma_T$ belongs to the cores of two different lines in T.

First, we consider lines α_i and α_j which come from the same skewering sequence $\alpha_0 \dashv \alpha_1 \dashv \cdots \dashv \alpha_k$, where $0 \leq i < j \leq k$. Note that the union

$$c(\alpha_i, \alpha_{i+1}) \cup c(\alpha_{i+1}, \alpha_{i+2}) \cup \cdots \cup c(\alpha_{j-2}, \alpha_{j-1})$$

is a connected subset of the plane and by Lemma 4.3, it cannot intersect α_j. Thus, this subset lies on one side of α_j (the side which contains the cell C_{α_i}), while the cell C_{α_j} lies on the other. Thus, the cells associated to α_i and α_j are disjoint.

Now, we consider lines α_i and β_j which come from distinct skewering sequences $\alpha_0 \dashv \alpha_1 \dashv \cdots \dashv \alpha_k$ and $\beta_0 \dashv \beta_1 \dashv \cdots \dashv \beta_m$, where $\alpha_0 = \ell_0 = \beta_0$ and $i, j \geq 1$.

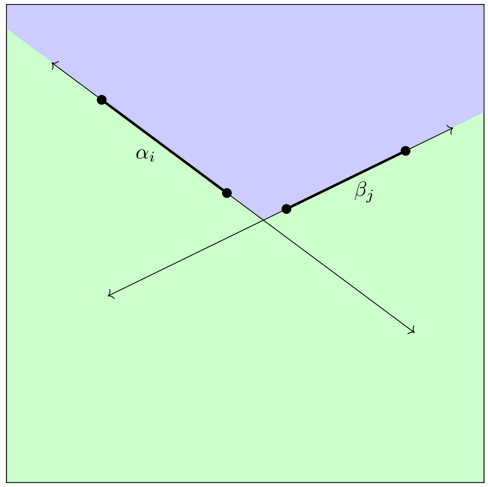

FIGURE 13. In the complement of the lines α_i and β_j, the region which has the cores of α_i and β_j in its boundary (shaded blue here) must contain the cores of lines which skewer α_i and β_j, which means that the cells C_{α_i} and C_{β_j} lie outside the shaded region.

Similar to the previous case, we observe that

$$\left(\bigcup_{t=1}^{i-1} c(\alpha_{t-1}, \alpha_t)\right) \cup \left(\bigcup_{t=1}^{j-1} c(\beta_{t-1}, \beta_t)\right)$$

is a connected subset of the plane and by Lemma 4.3, it must be wholly contained in one of the four components of the complement $\mathbb{C} - (\alpha_i \cup \beta_j)$. In particular, it must belong to the unique component which contains the core of α_i and the core of β_j in its boundary. Thus, this region of $\mathbb{C} - (\alpha_i \cup \beta_j)$ does not contain the cells for either α_i or β_j, which means that the two cells are disjoint - see Figure 13 for an illustration. □

Lemma 4.7. *Let P be a configuration of n points which satisfies Property Δ_2. Then the cells of a skewering tree cover the elements of P.*

Proof. Let $z \in P$, let T be a skewering tree and let Ω be the connected component of $\mathbb{C} - \Gamma_T$ which contains z. Suppose for the sake of contradiction that Ω does not belong to a cell of T. This immediately rules out the possibility that Ω touches only a single line in T, since that line would necessarily be a leaf and thus Ω would belong to the cell associated to that line. The remaining case to consider is that that there are distinct lines $\ell_1, \ell_2 \in T$ such that

(1) ℓ_1 and ℓ_2 bound adjacent sides of Ω, and
(2) Ω lies in the unique component of $\mathbb{C} - (\ell_1 \cup \ell_2)$ which has neither $c(\ell_1)$ nor $c(\ell_2)$ in its boundary.

Notice that this implies that the endpoints of ℓ_1, the endpoints of ℓ_2, and z form a 5-element subset of P where ℓ_i and ℓ_j each have one (non-link) endpoint in the interior of the convex hull. Similar to the proof of Lemma 4.6, we consider two cases according to whether ℓ_1 and ℓ_2 belong to the same skewering sequence or not.

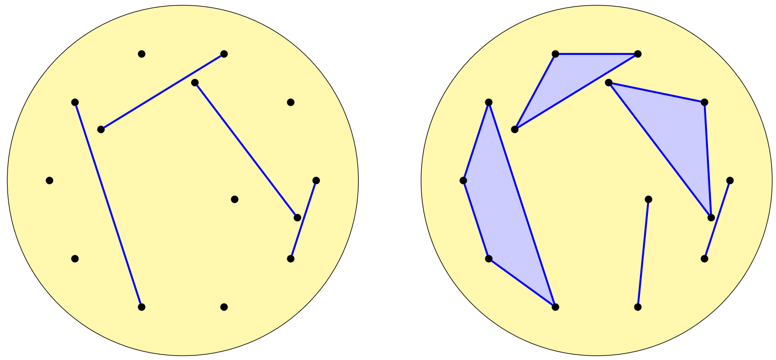

FIGURE 14. The minimum and maximum elements for the interval associated to the skewering tree drawn in Figure 10

First, suppose T contains a skewering sequence $\alpha_0 \dashv \alpha_1 \dashv \cdots \dashv \alpha_k$ such that $\ell_1 = \alpha_i$ and $\ell_2 = \alpha_j$, where $0 \leq i < j \leq k$ and let A be the set containing both z and the endpoints of $\alpha_i, \alpha_{i+1}, \ldots, \alpha_j$. Then $|A| = 2(j-i+1)+1$ and we know that $j-i$ points in A are link vertices for skewers involving other lines in A. Together with the two points mentioned above, we have that A is a set of $2(j-i+1)+1$ points where $j-i+2$ of them are internal, which contradicts our assumption that P satisfies Property Δ_2.

Next, suppose that T contains skewering sequences $\alpha_0 \dashv \alpha_1 \dashv \cdots \dashv \alpha_k$ and $\beta_0 \dashv \beta_1 \dashv \cdots \dashv \beta_m$ such that $\ell_1 = \alpha_i$ and $\ell_2 = \beta_j$ for some $i, j \geq 1$, and let B be the set containing z together with the endpoints of $\ell_0, \alpha_1, \ldots, \alpha_i, \beta_1, \ldots, \beta_j$. Then $|B| = 2(i+j+1)+1$, and at least $i+j$ elements of B are internal due to being link vertices of skewers in this set. By the same reasoning as in the previous case, we know that $i+j+2$ of the $2i+2j+3$ points in B are internal, which again contradicts our assumption that P satisfies Property Δ_2.

Since both cases lead to a contradiction of Property Δ_2, we conclude that z must in fact belong to a cell of the skewering tree T, and this completes the proof. $\square$

We now conclude this section by defining a subposet of $\mathrm{NC}(P)$ associated to each skewering tree.

Definition 4.8. Let P be a configuration of n points satisfying Property Δ_2 and let T be a skewering tree for P. We define the *skewering interval* $\mathrm{NC}(P,T)$ to be the subposet of $\mathrm{NC}(P)$ consisting of all partitions π satisfying the following conditions:

(1) for each $\ell \in T$, the endpoints of ℓ share a block in π;
(2) endpoints of distinct lines in T belong to distinct blocks in π;
(3) each block in π has a convex hull which lies in a cell of T.

As the name suggests, $\mathrm{NC}(P,T)$ is an interval in $\mathrm{NC}(P)$. Let $\hat{0}_T$ be the partition in $\mathrm{NC}(P)$ for which the only non-singleton blocks are $V(\ell)$ for each $\ell \in T$. Let $\hat{1}_T$ be the partition where two points in P belong to the same block if and only if they belong to the same cell of T (note that this requires us to choose a positive side for ℓ_0 before discussing the skewering interval). Then $\mathrm{NC}(P,T)$ is the interval $[\hat{0}_T, \hat{1}_T]$. See Figure 14.

We close the section with some features of skewering intervals which will be useful in the proof of Theorem B. To begin, we provide a product decomposition for skewering intervals.

Definition 4.9. Let ℓ be a boundary line in the arrangement $\mathcal{A}$ corresponding to P. Define π_ℓ to be the partition of P in which the two endpoints of ℓ share a block, while every other block is a singleton. Further, let $\mathrm{NC}_\ell(P)$ denote the interval $[\pi_\ell, \hat{1}]$ in $\mathrm{NC}(P)$.

Lemma 4.10. *Let T be a skewering tree for P. Then the skewering interval* $\mathrm{NC}(P,T)$ *is isomorphic to the product*

$$\left(\prod_{\substack{\ell\in T\\ \ell\neq\ell_0}} \mathrm{NC}_\ell(V(C_\ell))\right) \times \mathrm{NC}_{\ell_0}(V(C_{\ell_0}^+)) \times \mathrm{NC}(V(C_{\ell_0}^-)).$$

Proof. This follows immediately from Definitions 4.8 and 4.9. □

Next, we note that each skewering interval is "centered" in the sense that the rank of its minimum element is equal to the corank of its maximum.

Definition 4.11. Let π_0 and π_1 be elements of $\mathrm{NC}(P)$ with $\pi_0 \leq \pi_1$. We say that the interval $[\pi_0, \pi_1]$ is *centered* if $\rho(\pi_0) + \rho(\pi_1) = |P| - 1$, or equivalently if $bl(\pi_0) + bl(\pi_1) = |P| + 1$.

Lemma 4.12. *Each skewering interval* $\mathrm{NC}(P,T)$ *is a centered subposet of* $\mathrm{NC}(P)$.

Proof. Let T be a skewering tree for P. Then the partition $\hat{0}_T$ has rank $|T|$ and $\hat{1}_T$ has rank $|P| - (|T| + 1)$, so the interval $[\hat{0}_T, \hat{1}_T]$ is centered. □

Finally, we examine the ways in which skewering intervals can intersect.

Lemma 4.13. *If P has Property Δ_2, then distinct maximal skewering trees with the same initial line and choice of positive side yield disjoint skewering intervals.*

Proof. Let T and T' be distinct maximal skewering trees and suppose that the partition π is contained in both $\mathrm{NC}(P,T)$ and $\mathrm{NC}(P,T')$. In other words, both $\hat{0}_T \leq \pi \leq \hat{1}_T$ and $\hat{0}_{T'} \leq \pi \leq \hat{1}_{T'}$. Since T and T' are distinct and maximal, there must be lines α, ℓ, and ℓ' such that $\alpha, \ell \in T$, $\alpha, \ell' \in T'$, and α skewers both ℓ and ℓ'. Note that by Lemma 4.3, ℓ and ℓ' cannot skewer one another, and by definition of a skewering tree, ℓ and ℓ' do not intersect internally. Therefore the two lines are either parallel or have an external intersection.

Let $H^-_{\alpha,\ell}$ and $H^-_{\alpha,\ell'}$ denote the closed half-planes bounded by ℓ and ℓ' respectively which do not include the core of α. Since α skewers ℓ, we know that the cell C_ℓ (and therefore the core $c(\ell)$) must belong to $H^-_{\alpha,\ell}$; the analogous statement holds for ℓ'. By combining the inequalities above, we see that $\hat{0}_T \leq \hat{1}_{T'}$, which means that the core $c(\ell)$ must be contained in the cell $C_{\ell'}$, which implies that $c(\ell)$ belongs to $H^-_{\alpha,\ell'}$. Similarly, the fact that $\hat{0}_{T'} \leq \hat{1}_T$ tells us that $c(\ell')$ is contained in $H^-_{\alpha,\ell}$.

Combining all of the above, we know that the intersection $H^-_{\alpha,\ell} \cap H^-_{\alpha,\ell'}$ includes both $c(\ell)$ and $c(\ell')$, but not $c(\alpha)$. However, by the same reasoning used in the proof of Lemma 4.6, this is precisely the region in which the core of α must be placed

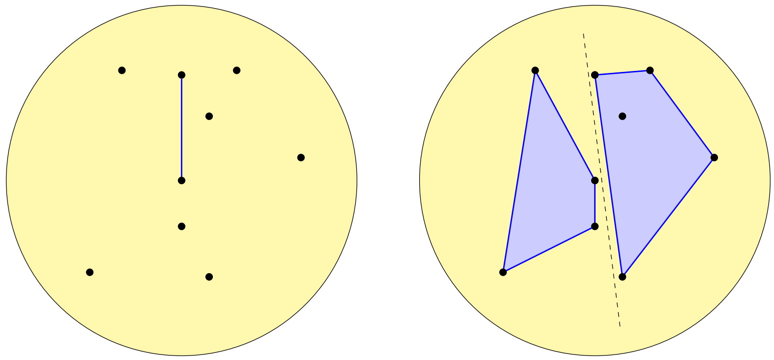

FIGURE 15. An atom in $\mathrm{NC}(P)$ with the corresponding coatom

for it to skewer both ℓ and ℓ'. Therefore, we have a contradiction, so the skewering intervals for T and T' must be disjoint. □

5. Property Δ_2 and Rank Symmetry

We now turn our attention to Theorem B, in which we demonstrate that if P satisfies Property Δ_2, then $\mathrm{NC}(P)$ is rank-symmetric. As a first step, one could write an explicit bijection between the atoms and coatoms of $\mathrm{NC}(P)$, where P is an arbitrary configuration of n points. This was previously demonstrated by Razen and Welzl in the special case where P is in general position [RW13] and is straightforward to generalize. In short: each atom is determined by a single line in $\mathcal{A}$, and by slightly rotating this line counterclockwise about the midpoint of its core, we obtain a line which divides P into two pieces, thus producing a coatom in $\mathrm{NC}(P)$ - see Figure 15 for an illustration.

One could hypothetically prove Theorem B by extending the map above to a rank-reversing bijection from $\mathrm{NC}(P)$ to itself, but this seems intractable in general. Instead, our proof technique is similar to that of Theorem A. First, we require a technical lemma regarding the interval $\mathrm{NC}_\ell(P)$ introduced in Definition 4.9. For the remainder of this section, let P denote a configuration of n points in $\mathbb{C}$ which satisfies Property Δ_2.

Definition 5.1. Let ℓ be a boundary line in the arrangement $\mathcal{A}$ corresponding to P. We say that P *satisfies Property* Δ_1 *relative to* ℓ if each $B \subseteq P$ with $V(\ell) \subseteq B$ has at most $\lfloor \frac{|B|-2}{2} \rfloor$ internal points.

The motivation for the preceding definition comes from its appearance in certain skewering trees for configurations which satisfy Property Δ_2.

Lemma 5.2. *Let T be a skewering tree for P with initial line ℓ_0. Suppose that ℓ_1 is a line in T with $\ell_0 \dashv \ell_1$, and that y is a point in P which lies in the convex hull $c(\ell_0, \ell_1)$ on the non-positive side of ℓ_0. Then $V(C^+_{\ell_0})$ satisfies Property Δ_1 relative to ℓ_0, and for each $\ell \in T$ with $\ell \neq \ell_0$, $V(C_\ell)$ satisfies Property Δ_1 relative to ℓ.*

Proof. We prove both claims by contradiction. To start, suppose that $V(C^+_{\ell_0})$ does not satisfy Property Δ_1 relative to ℓ_0. Then there is a subset $B \subseteq V(C^+_{\ell_0})$ which contains the endpoints of ℓ_0 such that B has more than $\lfloor \frac{|B|-2}{2} \rfloor$ internal points. If we define $B' = B \cup V(\ell_1) \cup \{y\}$, then the internal points of B, together with y and one endpoint of ℓ_0, are all internal points of B'. Therefore, B' has more than

$$\left\lfloor \frac{|B|-2}{2} \right\rfloor + 2 = \left\lfloor \frac{|B|+2}{2} \right\rfloor = \left\lfloor \frac{|B'|-1}{2} \right\rfloor$$

internal points, which violates the assumption that P satisfies Property Δ_2. Thus $V(C^+_{\ell_0})$ satisfies Property Δ_1 relative to ℓ_0.

Similarly, let $\ell \in T$ with $\ell \neq \ell_0$ and suppose that there is a subset $B \subseteq V(C_\ell)$ which contains the endpoints of ℓ such that B has more than $\lfloor \frac{|B|-2}{2} \rfloor$ internal points. If $\ell = \ell_1$, then we can define $B' = B \cup V(\ell_0) \cup \{y\}$ and observe by the same reasoning as above that B' has more than $\lfloor \frac{|B'|-1}{2} \rfloor$ internal points, which provokes a contradiction. Suppose instead that $\ell \neq \ell_1$. Then there is a skewering sequence $\alpha_1 \dashv \cdots \dashv \alpha_k \dashv \ell$ such that either $\ell_0 \dashv \alpha_1$ or $\ell_0 \dashv \ell_1 \dashv \alpha_1$. In either case, define

$$B' = B \cup V(\ell_0) \cup V(\ell_1) \cup \{y\} \cup V(\alpha_1) \cup \cdots \cup V(\alpha_k)$$

and observe that the internal points of B' include all the internal points of B, as well as one endpoint of each α_i, one endpoint of ℓ_0, and y. Therefore, the number of internal points in B' is more than

$$\left\lfloor \frac{|B|-2}{2} \right\rfloor + k + 3 = \left\lfloor \frac{|B|+2k+4}{2} \right\rfloor = \left\lfloor \frac{|B'|-1}{2} \right\rfloor,$$

which violates Property Δ_2. Thus, $V(C_\ell)$ must have Property Δ_1 relative to ℓ. $\square$

Lemma 5.3. *Suppose that P satisfies Property Δ_1 relative to the boundary line ℓ and that for each proper subset $Q \subset P$, the poset of noncrossing partitions* $\mathrm{NC}(Q)$ *is rank-symmetric. Then* $\mathrm{NC}_\ell(P)$ *is rank-symmetric.*

Proof. We proceed by induction on the number of internal points of P. First, if P has no internal points, then since ℓ was assumed to be a boundary line, we can see that $\mathrm{NC}_\ell(P)$ is isomorphic to NC_{n-1}, which is rank-symmetric. Now, suppose that the claim is true for any configuration with up to $k-1$ internal points which satisfies the lemma's hypotheses, and let P be a configuration with k internal points such that for all proper subsets $Q \subset P$, we know that $\mathrm{NC}(Q)$ is rank-symmetric. Let u and v be the endpoints of ℓ, and let P^v denote the complement $P - \{v\}$. We will compare the interval $[\pi_\ell, \hat{1}] \subset \mathrm{NC}(P)$ to the noncrossing partition lattice $\mathrm{NC}(P^v)$, which we know is rank-symmetric by assumption.

Define the map $\phi_v \colon \mathrm{NC}_\ell(P) \to \mathrm{NC}(P^v)$ by removing v from each partition in the domain and observe that ϕ_v is always injective, but typically not surjective. Our goal is to show that $\mathrm{NC}_\ell(P)$ is rank-symmetric; since $\mathrm{NC}(P^v)$ is assumed to be rank-symmetric and ϕ_v is injective, it suffices to show that the complement $\mathrm{NC}(P^v) - \phi_v(\mathrm{NC}_\ell(P))$ is a union of centered, disjoint, rank-symmetric intervals.

Let $W \subset P - \{u, v\}$ be the set of all points w with the property that $\mathrm{int}(\{u, v, w\})$ is nonempty and note that since P satisfies Property Δ_1 relative to ℓ, we must have $|\mathrm{int}(\{u, v, w\})| = 1$. Then $\mathrm{NC}(P^v) - \phi_v(\mathrm{NC}_\ell(P))$ is the collection of all partitions σ of P^v with a block which contains both u and a point $w \in W$, but not the unique point in $\mathrm{int}(\{u, v, w\})$. Rephrasing this characterization in the language of skewering trees, let $\mathcal{T}$ be the collection of skewering trees for P^v such that the initial line ℓ_0 has endpoints u and w for some $w \in W$ and the positive side of ℓ_0 is chosen to be the one which does not include v; then a partition $\sigma \in \mathrm{NC}(P^v)$ lies in $\mathrm{NC}(P^v) - \phi_v(\mathrm{NC}_\ell(P))$ if and only if it belongs to the skewering interval for some skewering tree in $\mathcal{T}$. Together with Lemmas 4.12 and 4.13, this tells us that the skewering intervals for trees in $\mathcal{T}$ form a collection of centered and disjoint intervals whose union is $\mathrm{NC}(P^v) - \phi_v(\mathrm{NC}_\ell(P))$. All that remains is to show that each such skewering interval is rank-symmetric.

Fix a skewering tree $T \in \mathcal{T}$ and consider the skewering interval $\mathrm{NC}(P^v, T)$. By Lemma 4.10, we have a poset isomorphism

$$\mathrm{NC}(P^v, T) \cong \left(\prod_{\substack{\alpha \in T \\ \alpha \neq \ell_0}} \mathrm{NC}_\alpha(P^v \cap C_\alpha) \right) \times \mathrm{NC}_{\ell_0}(P^v \cap C_{\ell_0}^+) \times \mathrm{NC}(P^v \cap C_{\ell_0}^-).$$

By our assumption that every proper subset of P has a rank-symmetric lattice of noncrossing partitions, we know that $\mathrm{NC}(P^v \cap C_{\ell_0}^-)$ is rank-symmetric. For the middle term in the product, we claim that the set $P^v \cap C_{\ell_0}^+$ satisfies Property Δ_1 relative to the boundary line ℓ_0. To see this, recall that the endpoints of ℓ_0 are the extremal point u and some $w \in W$, and define w' be the unique point of P which lies in the interior of $\textsc{Conv}(\{u, v, w\}$. If there were points $x, x', x'' \in P^v \cap C_{\ell_0}^+$ such that both x' and x'' lie in the convex hull $\textsc{Conv}(\{u, w, x\})$, then $\{u, v, w, w', x, x', x''\}$ would be a set of seven points, at least three of which are internal (w', x', and x''). This would violate our assumption that P has Property Δ_1 relative to ℓ, so it must be the case that $P^v - C_{\ell_0}^+$ satisfies Property Δ_1 relative to ℓ_0, and by our inductive hypothesis, we know that $\mathrm{NC}_{\ell_0}(P^v \cap C_{\ell_0}^+)$ is rank-symmetric.

Generalizing this argument, we now show that $P^v \cap C_\alpha$ satisfies Property Δ_1 relative to α for each non-initial line α in T. Let $\alpha_1 \dashv \cdots \dashv \alpha_m$ be a skewering sequence in T such that $\alpha_1 = \ell_0$ and $\alpha_m = \alpha$. If there are points $x, x', x'' \in P^v \cap C_\alpha$ such that x' and x'' lie in the triangle formed by x and the endpoints of α, and if w and w' are defined as above, then

$$\{x, x', x'', w', v\} \cup V(\alpha_1) \cup \cdots \cup V(\alpha_m)$$

is a set of $2m + 5$ points, of which at least $m + 2$ points must be internal - see Figure 16 for an illustration. Since this set contains both u and v, this violates our assumption that P satisfies Property Δ_1 relative to ℓ, so we may conclude that $P^v \cap C_\alpha$ satisfies Property Δ_1 relative to α, as desired. By the inductive hypothesis, $\mathrm{NC}_\alpha(P^v \cap C_\alpha)$ is rank-symmetric.

Finally, we have that $\mathrm{NC}(P^v, T)$ is a product of rank-symmetric posets and is therefore rank-symmetric itself, which completes the proof. □

Theorem 5.4 (Theorem B)**.** *Let $P \subset \mathbb{C}$ be a set of n distinct points which satisfies Property Δ_2. Then $\mathrm{NC}(P)$ is a rank-symmetric graded lattice.*

Proof. We proceed by induction on the number of internal points for P. When P has no internal points, $\mathrm{NC}(P)$ is isomorphic to the classical noncrossing partition lattice NC_n, which is rank-symmetric. Now, suppose that every configuration with fewer than $|\mathrm{int}(P)|$ internal points has a rank-symmetric lattice of noncrossing partitions; we will show that $\mathrm{NC}(P)$ is rank-symmetric as well. By Theorem 2.9, there is a sequence of Δ_2-moves which transforms P into a convex configuration, for which the lattice of noncrossing partitions is isomorphic to NC_n and thus rank-symmetric. The only remaining step is to prove that if $m\colon P \to m(P)$ is a Δ_2-move, then $\mathrm{NC}(P)$ is rank-symmetric if and only if $\mathrm{NC}(m(P))$ is rank-symmetric.

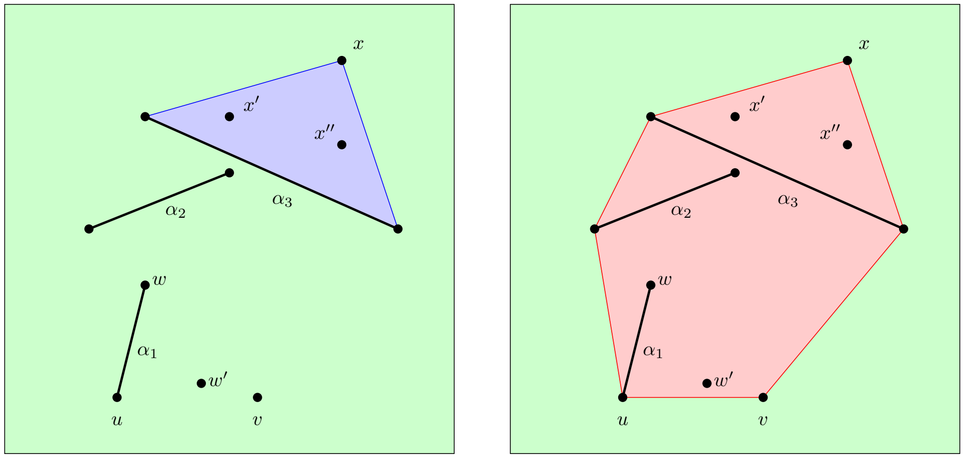

FIGURE 16. If $\alpha_1 \dashv \alpha_2 \dashv \alpha_3$ is a skewering sequence for P and the points in the cell for α_3 do not satisfy Property Δ_1 relative to α_3 (illustrated on the left in blue), then there is a set of points containing u and v (illustrated on right in red) which demonstrate that P does not satisfy Property Δ_1 relative to the line with endpoints u and v.

In the proof of Theorem A, we showed that when P satisfies Property Δ_1, the block-switching map BS_m is a rank-preserving bijection between pre-m-noncrossing partitions of P and post-m-noncrossing partitions of $m(P)$. If P is only assumed to satisfy Property Δ_2, the block switching map might not be a bijection; there may be partitions $\pi \in \mathrm{NC}(P)$ such that $\mathrm{BS}_m(\pi)$ is not in $\mathrm{NC}(m(P))$, or elements $\sigma \in \mathrm{NC}(m(P))$ such that $\mathrm{BS}_m^{-1}(\sigma)$ does not lie in $\mathrm{NC}(P)$. To complete the proof, we must show that these two collections are rank-symmetric subposets of $\mathrm{NC}(P)$ and $\mathrm{NC}(m(P))$ respectively. By symmetry, it suffices to examine the first collection.

Suppose that the move m takes a point $z \in P$ across a line $\ell \in \mathcal{A}$ with endpoints $e(\ell) = \{w_1, w_2\}$ and let $F_m(P)$ denote the noncrossing partitions of P which fail to be accounted for by the block-switching map BS_m; that is,

$$F_m(P) = \{\pi \in \mathrm{NC}(P) \mid \mathrm{BS}_m(\pi) \notin \mathrm{NC}(m(P))\}.$$

If $F_m(P)$ is empty, then there is nothing to prove. Otherwise, for each $\pi \in F_m(P)$, there are points $x, y \in P$ such that $y, z \in \mathrm{int}(\{w_1, w_2, x\})$, and these five points are divided into three distinct blocks of π (each of which might contain other points) as follows: x and z belong to one block, w_1 and w_2 belong to another, and y lies in a third - see Figure 17. Using this characterization, we will prove that (1) $F_m(P)$ is a union of skewering intervals, (2) each of these skewering intervals is centered and rank-symmetric, and (3) intersections of the skewering intervals are centered and rank-symmetric.

More concretely, let $y_1, \ldots, y_k$ be the points in P such that for each $i \in \{1, \ldots, k\}$, there is some $x_i \in P$ such that w_1, w_2, and x_i form a triangle with interior points z and y_i (and nothing else, since P satisfies Property Δ_2). When there is more than one choice for x_i, we select the unique option where all other possible choices lie in the half-plane bounded by x_i and z which does not contain y_i. Then $F_m(P)$

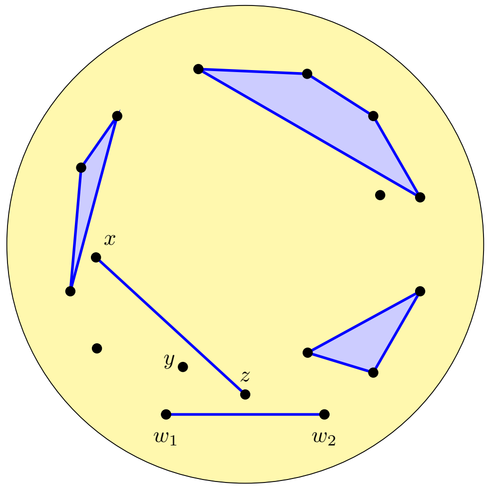

FIGURE 17. For the depicted partition π of P, if m moves z across the line containing w_1 and w_2, then $\mathrm{BS}_m(\pi)$ is not an element of $\mathrm{NC}(m(P))$ since it has w_1, w_2, and x together in a single block without the interior point y.

consists of all partitions in $\mathrm{NC}(P)$ where there exists an i such that z and x_i belong to the same block, w_1 and w_2 belong to a different block, and y_i belongs to a third block.

Let ℓ_i denote the line with endpoints x_i and z. Since z is adjacent to the line ℓ, we know that ℓ_i must skewer ℓ. The characterization above can thus be rephrased: the partition π belongs to $F_m(P)$ if and only if π lies in a skewering interval $\mathrm{NC}(P, T_i)$ for some skewering tree T_i with initial line ℓ_i such that $\ell \in T$ and y_i is on the non-positive side of ℓ_i. It is straightforward to see that each such skewering interval is a subset of $F_m(P)$, so we may conclude that $F_m(P)$ is a union of skewering intervals.

Each skewering interval in the union is centered by Lemma 4.12, which implies that $F_m(P)$ is itself a centered subposet of $\mathrm{NC}(P)$. By Lemma 4.10, each skewering interval $\mathrm{NC}(P, T_i)$ decomposes into a product of posets; the factor $\mathrm{NC}(V(C^-_{\ell_i}))$ is rank-symmetric by our inductive hypothesis and terms of the form $\mathrm{NC}_\ell(V(C_\ell))$ and $\mathrm{NC}_{\ell_i}(V(C^+_{\ell_i}))$ are rank-symmetric by Lemmas 5.2 and 5.3. Since the product of rank-symmetric posets is itself rank-symmetric, we may conclude that the interval $\mathrm{NC}(P, T_i)$ is rank-symmetric.

We are now at the final step: examining how the skewering intervals which make up $F_m(P)$ can intersect. First, note that if T_i and T'_i are two distinct skewering trees in the union which both use ℓ_i as the initial line, we know by Lemma 4.13 that $\mathrm{NC}(P, T_i) \cap \mathrm{NC}(P, T'_i)$ is empty. Next, consider two skewering trees T_i and T_j which appear in the union with $i \neq j$. The intersection $\mathrm{NC}(P, T_i) \cap \mathrm{NC}(P, T_j)$ is nonempty precisely when the cells $C^+_{\ell_i}$ for T_i and $C^+_{\ell_j}$ for T_j have an intersection which excludes both y_i and y_j (note that this precludes the possibility of nonempty triple intersections for skewering intervals). When this is the case, the two skewering intervals intersect in the interval $[\hat{0}_{T_i} \vee \hat{0}_{T_j}, \hat{1}_{T_i} \wedge \hat{1}_{T_j}]$; this is not a skewering interval, but it has many of the associated properties. By similar arguments to those given in Lemma 4.3, Lemma 4.6 and Lemma 4.7, we can see that the minimum element

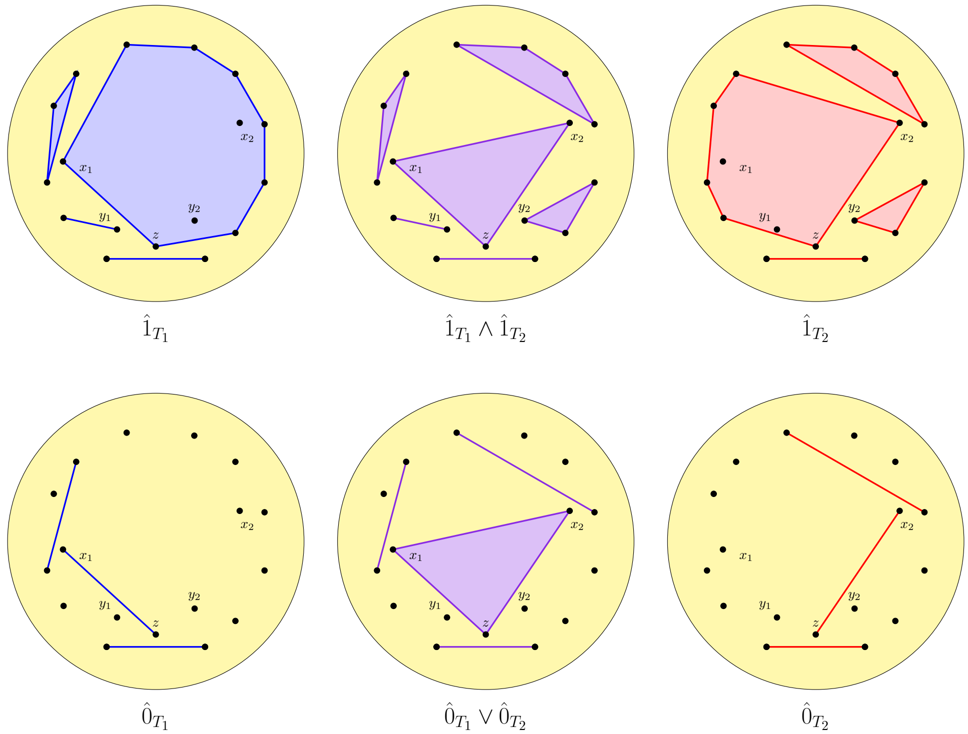

FIGURE 18. On the left and right, the minimum and maximum elements for a skewering interval are given. The intersection of these two intervals is another interval, albeit one which does not arise from a skewering tree; its minimum and maximum elements are displayed in the center.

$\hat{0}_{T_i} \vee \hat{0}_{T_j}$ consists of an edge for each non-initial line in T_i and T_j, together with the triangle with vertices y_i, y_j, and z. This triangle has three cells associated to it: one containing z_i, one containing z_j, and one which is the intersection of $C^+_{\ell_i}$ and $C^+_{\ell_j}$ (thus containing the triangle itself). Each other edge has a well-defined cell in the same way as a typical skewering tree does. The maximal element $\hat{1}_{T_i} \wedge \hat{1}_{T_j}$ is constructed using these cells in the same manner as for skewering trees - see Figure 18 for an illustration. Putting this all together, $[\hat{0}_{T_i} \vee \hat{0}_{T_j}, \hat{1}_{T_i} \wedge \hat{1}_{T_j}]$ admits a decomposition similar to the one described in Lemma 4.10, so by Lemmas 5.2 and 5.3, this interval is centered and rank-symmetric.

In summary, we have shown that $F_m(P)$ is a union of centered and rank-symmetric intervals in $\mathrm{NC}(P)$, whose pairwise intersections are themselves centered and rank-symmetric and whose k-wise intersections are empty when $k > 2$. Therefore, $F_m(P)$ is centered and rank-symmetric, which completes the proof. $\square$

ACKNOWLEDGMENTS

We are grateful to Swarthmore College for support during the summer of 2022, when the majority of our research took place.

References

[BBG+19] B. Baumeister, K.-U. Bux, F. Götze, D. Kielak, and H. Krause, *Non-crossing partitions*, Spectral structures and topological methods in mathematics, EMS Ser. Congr. Rep., EMS Publ. House, Zürich, 2019, pp. 235–274.

[ER00] P. H. Edelman and V. Reiner, *Counting the interior points of a point configuration*, Discrete Comput. Geom. **23** (2000), no. 1, 1–13. MR 1727120

[Kre72] G. Kreweras, *Sur les partitions non croisées d'un cycle*, Discrete Math. **1** (1972), no. 4, 333–350.

[McC06] Jon McCammond, *Noncrossing partitions in surprising locations*, Amer. Math. Monthly **113** (2006), no. 7, 598–610.

[NS97] Alexandru Nica and Roland Speicher, *A "Fourier transform" for multiplicative functions on non-crossing partitions*, J. Algebraic Combin. **6** (1997), no. 2, 141–160. MR 1436532

[RW13] Andreas Razen and Emo Welzl, *On the number of crossing-free partitions*, Comput. Geom. **46** (2013), no. 7, 879–893. MR 3061450

[Sta12] Richard P. Stanley, *Enumerative combinatorics. Volume 1*, second ed., Cambridge Studies in Advanced Mathematics, vol. 49, Cambridge University Press, Cambridge, 2012.

[Sta15] ———, *Catalan numbers*, Cambridge University Press, New York, 2015.

[SU91] Rodica Simion and Daniel Ullman, *On the structure of the lattice of noncrossing partitions*, Discrete Math. **98** (1991), no. 3, 193–206.

[SW06] Micha Sharir and Emo Welzl, *On the number of crossing-free matchings, cycles, and partitions*, SIAM J. Comput. **36** (2006), no. 3, 695–720. MR 2263008

Email address: `scohen2@swarthmore.edu`

Department of Mathematics & Statistics, Swarthmore College, Swarthmore, PA 19081

Email address: `doughemj@lafayette.edu`

Department of Mathematics, Lafayette College, Easton, PA 18042

Email address: `aharsh1@swarthmore.edu`

Department of Mathematics & Statistics, Swarthmore College, Swarthmore, PA 19081

Email address: `smartin4@swarthmore.edu`

Department of Mathematics & Statistics, Swarthmore College, Swarthmore, PA 19081