NONCROSSING SET PARTITIONS

HENRI MÜHLE

We assume the reader to be familiar with partially ordered sets (or posets for short) and lattices. The notions used here are fairly standard. We recommend [7] for more background on posets and lattices, and hopefully as a source to answer all questions regarding notation or concepts that appear unclear. For a motivation and some historical background on the study of noncrossing partitions, we recommend Section 4.1 in [2].

Some exercises are scattered throughout the text, but can be recognized by the red box.

1. Set Partitions

1.1. **Definition and Representation.** In this section we start with a basic combinatorial treatment of all set partitions. The results presented here are well known, and should probably be considered folklore. The interested reader may consult [19, Chapter 1] for an exposition on the history of set partitions.

For simplicity, we use the abbreviation $[n] = \{1, 2, \ldots, n\}$ for any nonnegative integer n.

Definition 1.1

Let M be a finite set. A SET PARTITION of M is a family $\mathbf{x} = \{B_1, B_2, \ldots, B_m\}$ of non-empty subsets of M with the property that $M = B_1 \cup B_2 \cup \cdots \cup B_m$ and $B_i \cap B_j = \emptyset$ for $i, j \in [m]$ with $i \neq j$. The members of $\mathbf{x}$ are its BLOCKS, and we denote by $\mathrm{bl}(\mathbf{x}) = m$ the number of blocks of $\mathbf{x}$.

Since M is a finite set we can without loss of generality assume that $M = [n]$ for some nonnegative integer n. In order to avoid writing multiple set parentheses we use the following reduced notation: we order the blocks according to the value of their smallest element, and we list them one after another, separated by a vertical line. The elements in each block are also ordered linearly. Let us denote the family of all set partitions of an n-element set by Π_n. If the ground set matters, we also write Π_M.

Now let $\mathbf{x} \in \Pi_n$, and assume that there is a block $B \in \mathbf{x}$. If there are integers $i, j \in [n]$ with $i, j \in B$, then we write $i \sim_{\mathbf{x}} j$ as a shorthand.

There are several ways to graphically represent set partitions, and we want to emphasize two of them.

Definition 1.2

Let $n \geq 0$ and consider $\mathbf{x} \in \Pi_n$. The ARC DIAGRAM of $\mathbf{x}$ is constructed as follows: write the numbers $1, 2, \ldots, n$ from left to right on a horizontal line. Two integers i and j are connected by an arc (rising above this line) if and only if $i \sim_{\mathbf{x}} j$ and there is no $k \in [n]$ with $i < k < j$ and $i \sim_{\mathbf{x}} k$.

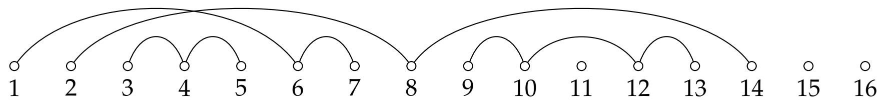

(A) An arc diagram.

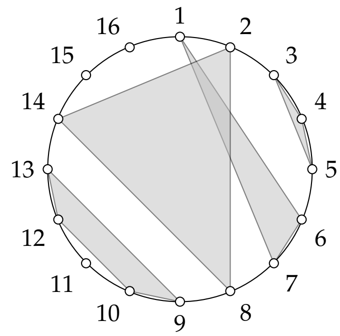

(B) A circle diagram.

FIGURE 1. Graphical representations of the set partition from Example 1.4.

In other words, we connect consecutive entries in a block by an arc. The connected components of the resulting graph are precisely the blocks of $\mathbf{x}$. Let us write $\mathscr{A}(\mathbf{x})$ for the arc diagram associated with $\mathbf{x}$.

DEFINITION 1.3

Let $n \geq 0$ and consider $\mathbf{x} \in \Pi_n$. The CIRCLE DIAGRAM of $\mathbf{x}$ is constructed as follows: write the numbers $1, 2, \ldots, n$ in clockwise order on a circle. For each $B \in \mathbf{x}$, draw the convex hull of B on this circle.

Let us write $\mathscr{C}(\mathbf{x})$ for the circle diagram of $\mathbf{x}$. Observe that if we remove in $\mathscr{C}(\mathbf{x})$ the circle and the edge connecting the smallest and the largest element of each block, then we obtain $\mathscr{A}(\mathbf{x})$ by straightening the vertices.

EXAMPLE 1.4

Consider $\mathbf{x} = \Bigl\{\{1,6,7\},\{2,8,14\},\{3,4,5\},\{9,10,12,13\},\{11\},\{15\},\{16\}\Bigr\} \in \Pi_{16}$. In reduced notation this would be $\mathbf{x} = 1\;6\;7 \mid 2\;8\;14 \mid 3\;4\;5 \mid 9\;10\;12\;13 \mid 11 \mid 15 \mid 16$. The arc diagram and the circle diagram of $\mathbf{x}$ are shown in Figure 1.

✎ EXERCISE 1

Let $\mathbf{x} \in \Pi_n$. Show that $\sim_{\mathbf{x}}$ is an equivalence relation on $[n]$. In other words, the set partitions of $[n]$ are in bijection with equivalence relations on $[n]$.

1.2. **Enumeration.** Let us now enumerate the set partitions of an n-element set.

Proposition 1.5

For $n \geq 0$ the cardinality of Π_n is given by the n^{th} Bell number, which is recursively defined by

$$B(n) = \sum_{k=0}^{n-1} \binom{n-1}{k} B(k).$$

The initial condition is $B(0) = 1$.

Proof. Consider a set partition $\mathbf{x} \in \Pi_n$ and let B denote its (unique) block containing 1. Let $\lvert B \rvert = k$ for $k \in [n]$. Then $\mathbf{x} \setminus B$ is a set partition of the $(n-k)$-element set $[n] \setminus B$. By induction, there are $B(n-k)$ possible set partitions having B as a block. Since $1 \in B$, there are $k-1$ vacant positions in B that could be filled with any of the $n-1$ elements, hence there are $\binom{n-1}{k-1}$ ways to choose a block of size k containing 1. By summing over k we obtain the desired recurrence

$$\begin{aligned} B(n) &= \sum_{k=1}^{n} \binom{n-1}{k-1} B(n-k) \\ &= \sum_{k=0}^{n-1} \binom{n-1}{n-1-k} B(k) \\ &= \sum_{k=0}^{n-1} \binom{n-1}{k} B(k). \end{aligned}$$

The last equality follows from the symmetry $\binom{n}{k} = \binom{n}{n-k}$ of the binomial coefficients.

We finish the proof by observing that there is a unique partition of the empty set. □

As it turns out the strategy of removing the block containg the number 1 from a set partition in order to set up a recursion will come in handy several times. We can in fact enumerate the set partitions of $[n]$ in a more refined way.

Proposition 1.6

For $n \geq 0$ and $k \in [n]$ the number of set partitions of $[n]$ with exactly k blocks is given by the k^{th} Stirling number of the second kind, which is recursively defined by

$$\left\{ n \atop k \right\} = k \left\{ n-1 \atop k \right\} + \left\{ n-1 \atop k-1 \right\}.$$

The initial conditions are $\left\{ 0 \atop 0 \right\} = 1$ and $\left\{ n \atop 0 \right\} = \left\{ 0 \atop n \right\} = 0$ for $n > 0$.

Proof. There are two possibilities for a set partition of $[n]$ with k blocks: it either contains $\{n\}$ as a block, or it does not.

There are $\left\{ n-1 \atop k-1 \right\}$ possibilities for the first case, since any set partition of $[n-1]$ into $k-1$ blocks can be extended to a set partition of $[n]$ into k blocks by adding the singleton block $\{n\}$.

In the second case, we can briefly forget about n and obtain a set partition of $[n-1]$ into k blocks. By induction, there are $\left\{ n-1 \atop k \right\}$ of these. There are k ways to reinsert n into one of the present blocks, which yields the result. □

COROLLARY 1.7

For $n \geq 0$ we have
$$B(n) = \sum_{k=0}^{n} \left\{ n \atop k \right\}.$$

The Stirling numbers of the second kind have the following explicit formula, which we will not derive here:
$$\left\{ n \atop k \right\} = \frac{1}{k!} \sum_{j=0}^{k} (-1)^{k-j} \binom{k}{j} j^n.$$

1.3. **Lattice Property.** Given a set partition of $[n]$ we may wonder how fine it is. Clearly the finest way to partition $[n]$ is by putting each element of $[n]$ in a single block. We call this the DISCRETE PARTITION, and denote it by $\mathbf{0}$. On the contrary, the coarsest partition of $[n]$ only has a single block, namely $[n]$ itself. We call this the FULL PARTITION, and denote it by $\mathbf{1}$. But where in this spectrum does an arbitrary set partition of $[n]$ sit?

DEFINITION 1.8

Let $n \geq 0$ and let $\mathbf{x}, \mathbf{x}' \in \Pi_n$ have $\mathbf{x} = \{B_1, B_2, \ldots, B_k\}$ and $\mathbf{x}' = \{B'_1, B'_2, \ldots, B'_{k'}\}$. The set partition $\mathbf{x}$ REFINES $\mathbf{x}'$ if for each $i \in [k]$ there exists some $i' \in [k']$ such that $B_i \subseteq B'_{i'}$. We also say that $\mathbf{x}'$ IS REFINED BY $\mathbf{x}$, and we write $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{x}'$.

The relation defined in Definition 1.8 is the DUAL REFINEMENT ORDER of Π_n. The attribute "dual" relates to the fact that the finest set partition of $[n]$ is minimal with respect to this order, meaning that we coarsen set partitions when going up in that order. Figure 2 shows the poset $(\Pi_4, \leq_{\mathrm{dref}})$.

✎ EXERCISE 2

Show that the relation $\leq_{\mathrm{dref}}$ is a partial order on Π_n, in which the trivial partition $\mathbf{0}$ is the least element and the full partition $\mathbf{1}$ is the greatest.
Show further that the function $\mathrm{rk}(\mathbf{x}) = n - \mathrm{bl}(\mathbf{x})$ is a rank function of $(\Pi_n, \leq_{\mathrm{dref}})$, i.e. it satisfies $\mathrm{rk}(\mathbf{0}) = 0$ and $\mathrm{rk}(\mathbf{y}) = \mathrm{rk}(\mathbf{x}) + 1$ for every cover relation $\mathbf{x} \lessdot_{\mathrm{dref}} \mathbf{y}$.

PROPOSITION 1.9

For $n \geq 0$ the poset $(\Pi_n, \leq_{\mathrm{dref}})$ is in fact a lattice.

Proof. This follows immediately from the fact that Π_n is a finite poset with a greatest element, and for any two set partitions $\mathbf{x}, \mathbf{x}' \in \Pi_n$ the intersection

$$\mathbf{x} \wedge_{\Pi} \mathbf{x}' = \{B \cap B' \mid B \in \mathbf{x}, B' \in \mathbf{x}', \text{ and } B \cap B' \neq \emptyset\} \tag{1}$$

is the greatest set partition that at the same time refines both $\mathbf{x}$ and $\mathbf{x}'$.

On the other hand, the supremum of $\mathbf{x}$ and $\mathbf{x}'$ is given by

$$\mathbf{x} \vee_{\Pi} \mathbf{x}' = \big\{C \mid C \text{ is a connected component of } \mathscr{A}(\mathbf{x}) \cup \mathscr{A}(\mathbf{x}')\big\}. \tag{2}$$

□

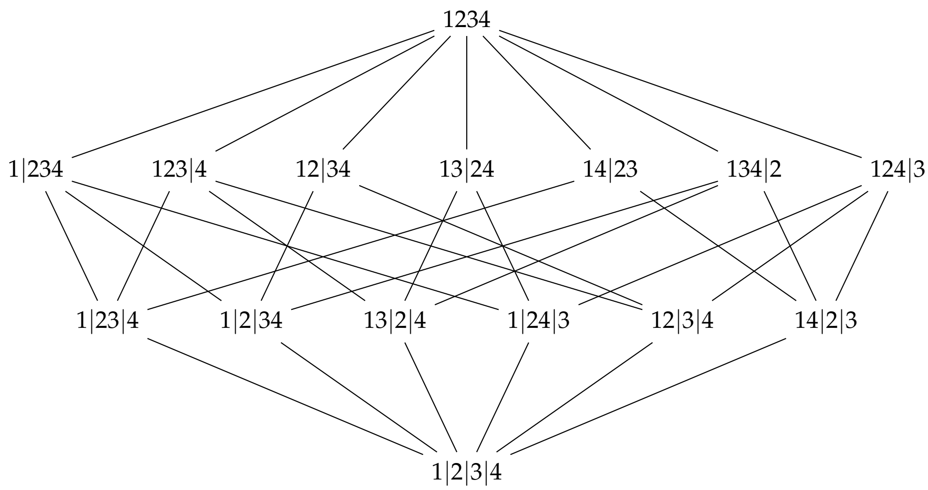

FIGURE 2. The poset $(\Pi_4, \leq_{\mathrm{dref}})$.

EXAMPLE 1.10

Let $\mathbf{x} = 1 \mid 2 \mid 3\,5\,7\,8 \mid 4 \mid 6$, and $\mathbf{x}' = 1\,3 \mid 2\,4 \mid 5\,6\,8 \mid 7$ be two set partitions of $[8]$. Their meet is $\mathbf{x} \wedge_{\Pi} \mathbf{x}' = 1 \mid 2 \mid 3 \mid 4 \mid 5\,8 \mid 6 \mid 7$. The graph $\mathbf{P}_{\mathbf{x},\mathbf{x}'}$ is

where the green edges make up the arc diagram of $\mathbf{x}$, and the blue edges make up the arc diagram of $\mathbf{x}'$. We see that $\mathbf{x} \vee_{\Pi} \mathbf{x}' = 1\,3\,5\,6\,7\,8 \mid 2\,4$.

✎ EXERCISE 3

Work out the details of the proof of Proposition 1.9.

1.4. **Möbius Function.** In this section we want to compute the values of the Möbius function on the poset $(\Pi_n, \leq_{\mathrm{dref}})$, and we approach this problem in a very combinatorial way. In fact, we use the following result of R. Stanley that relates the value of the Möbius function in a poset to the number of some particular maximal chains. Recall that an EDGE-LABELING is simply a function from the cover relations of a poset to the integers. This result was first stated for supersolvable lattices in [30, Theorem 1.2], and later for "admissible" lattices in [31, Corollary 3.3], but it essentially works already for arbitrary graded posets.

PROPOSITION 1.11: [6, Theorem 2.7]

Let $\mathcal{P}$ be a finite graded poset with least element $\hat{0}$ and greatest element $\hat{1}$. Assume that there exists an edge-labeling λ of $\mathcal{P}$ with the property that in each interval exists a unique rising maximal chain, and let $f(\mathcal{P};\lambda)$ denote the number of falling maximal chains of $\mathcal{P}$

with respect to λ. Then we have
$$\mu_{\mathcal{P}}(\hat{0},\hat{1}) = (-1)^{\mathrm{rk}(\mathcal{P})} f(\mathcal{P};\lambda).$$

Proof. In fact we prove a much stronger statement. Let $n = \mathrm{rk}(\mathcal{P})$ and let $S \subseteq [n-1]$. Define
$$P_S = \{x \in P \mid \mathrm{rk}(x) \in S\} \cup \{\hat{0},\hat{1}\},$$
and let $\mathcal{P}_S = (P_S, \leq)$. Let $\mathscr{M}$ denote the set of all maximal chains of $\mathcal{P}$, and let $\mathscr{M}_S$ denote the set of all maximal chains of $\mathcal{P}_S$. Now let $C = \{x_0, x_1, \ldots, x_n\} \in \mathscr{M}$, and define
$$\mathrm{Des}(C) = \bigl\{i \mid \lambda(x_{i-1}, x_i) > \lambda(x_i, x_{i+1})\bigr\} \subseteq [n-1].$$

The existence of a unique rising maximal chain $C_{x,y}$ in each interval $[x,y]$ gives rise to a map from $\theta : \mathscr{M}_S \to \mathscr{M}$. Let $S = \{i_1, i_2, \ldots, i_s\}$, and let $C = \{x_{i_1}, x_{i_2}, \ldots, x_{i_s}\} \in \mathscr{M}_S$ with $\mathrm{rk}(x_{i_j}) = i_j$. Define $\theta(C) = C_{\hat{0},x_{i_1}} \uplus C_{x_{i_1},x_{i_2}} \uplus \cdots \uplus C_{x_{i_s},\hat{1}} \in \mathscr{M}$. By definition we have $\mathrm{Des}\bigl(\theta(C)\bigr) \subseteq S$. Moreover, since $C = \theta(C) \cap P_S$ we conclude that θ is injective. If $C \in \mathscr{M}$ satisfies $\mathrm{Des}(C) \subseteq S$, then $\theta(C \cap P_S) = C$, which implies that θ is a bijection from $\mathscr{M}_S$ to $\{C \in \mathscr{M} \mid \mathrm{Des}(C) \subseteq S\}$.

If $f(\mathcal{P}_S;\lambda)$ denotes the number of maximal chains in $\mathscr{M}$ with descent set equal to S, then the reasoning in the previous paragraph implies
$$\bigl|\mathscr{M}_S\bigr| = \sum_{T \subseteq S} f(\mathcal{P}_T;\lambda).$$
Recall that Philip Hall's Theorem, see for instance [32, Proposition 3.8.5], implies
$$\mu_{\mathcal{P}}(\hat{0},\hat{1}) = \sum_{S \subseteq [n-1]} (-1)^{|S|+1}\bigl|\mathscr{M}_S\bigr|.$$
We therefore conclude using the Principle of Inclusion-Exclusion that
$$f(\mathcal{P}_S;\lambda) = \sum_{T \subseteq S} (-1)^{|S\setminus T|}\bigl|\mathscr{M}_T\bigr| = (-1)^{|S|+1}\sum_{T \subseteq S}(-1)^{|T|+1}\bigl|\mathscr{M}_T\bigr| = (-1)^{|S|+1}\mu_{\mathcal{P}_S}(\hat{0},\hat{1}).$$
If we plug in $S = [n-1]$, we obtain the claimed result. ☐

To simplify the task of computing the Möbius function of $(\Pi_n, \leq_{\mathrm{dref}})$, we recall that the Möbius function is multiplicative, and we observe that intervals in $(\Pi_n, \leq_{\mathrm{dref}})$ admit a nice decomposition as direct products of smaller partition lattices. For $X \subseteq [n]$ and $\mathbf{x} \in \Pi_n$, define the RESTRICTION of $\mathbf{x}$ to X by
$$\mathbf{x}_{|X} = \{B \cap X \mid B \in \mathbf{x}\}.$$

LEMMA 1.12

Let $n \geq 0$ and let $\mathbf{x}, \mathbf{y} \in \Pi_n$ with $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$. If $\mathrm{bl}(\mathbf{y}) = l$ and the i^{th} block of $\mathbf{y}$ is composed of k_i blocks of $\mathbf{x}$ for $i \in [l]$, then
$$[\mathbf{x},\mathbf{y}]_{\Pi} \cong \prod_{i=1}^{l} (\Pi_{k_i}, \leq_{\mathrm{dref}}).$$

Proof. Since $\mathbf{x} \leq_{\text{dref}} \mathbf{y}$ every block of $\mathbf{x}$ is contained in some block of $\mathbf{y}$. This means in particular that every block of $\mathbf{y}$ is the disjoint union of some blocks of $\mathbf{x}$. Let us write $\mathbf{x} = \{B_1, B_2, \ldots, B_k\}$ and $\mathbf{y} = \{C_1, C_2, \ldots, C_l\}$. Assume that $C_i = B_{i_1} \uplus B_{i_2} \uplus \cdots \uplus B_{i_{k_i}}$. Since the blocks of $\mathbf{y}$ are mutually disjoint, it suffices to focus on $\left[\mathbf{x}_{|C_i}, \mathbf{y}_{|C_i}\right]_\Pi$, where we have $\mathbf{y}_{|C_i} = C_i$, since $C_i \in \mathbf{y}$. This interval, however, is canonically isomorphic to $(\Pi_{k_i}, \leq_{\text{dref}})$, via the map $B_{i_j} \mapsto j$. $\square$

Let μ_{Π_n} denote the Möbius function of the lattice $(\Pi_n, \leq_{\text{dref}})$. We abbreviate $\mu(\Pi_n) = \mu_{\Pi_n}(\mathbf{0}, \mathbf{1})$.

Proposition 1.13

For $n \geq 1$ we have
$$\mu(\Pi_n) = (-1)^{n-1}(n-1)!.$$

Corollary 1.14

Let $n \geq 1$ and let $\mathbf{x}, \mathbf{y} \in \Pi_n$ with $\mathbf{x} \leq_{\text{dref}} \mathbf{y}$. If $\text{bl}(\mathbf{y}) = k$ and the i^{th} block of $\mathbf{y}$ is composed of k_i blocks of $\mathbf{x}$ for $i \in [l]$, then
$$\mu_{\Pi_n}(\mathbf{x}, \mathbf{y}) = (-1)^{\text{rk}(\mathbf{y}) - \text{rk}(\mathbf{x})}(k_1 - 1)!(k_2 - 1)! \cdots (k_l - 1)!.$$

Proof. This follows immediately from Lemma 1.12 and Proposition 1.13. Observe that $\text{bl}(\mathbf{x}) = l_1 + l_2 + \cdots + l_k$ and $\text{bl}(\mathbf{y}) = k$. $\square$

It remains to prove Proposition 1.13, and we prove this result with the help of a certain edge-labeling of $(\Pi_n, \leq_{\text{dref}})$. By construction, if $\mathbf{x} \lessdot_{\text{dref}} \mathbf{y}$, then there exist two blocks $B, B' \in \mathbf{x}$ such that $\mathbf{y} = \left(\mathbf{x} \setminus \{B, B'\}\right) \cup (B \cup B')$. Let us label this cover relation by

$$\lambda(\mathbf{x}, \mathbf{y}) = \max\{\min B, \min B'\}. \tag{3}$$

This labeling was first considered by I. Gessel, see [6, Example 2.9]. The following results are due to him.

A MAXIMAL CHAIN of $(\Pi_n, \leq_{\text{dref}})$ is a sequence $C = (\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_{n-2})$ of $n-1$ set partitions with the property that $\mathbf{x}_0 = \mathbf{0}$, $\mathbf{x}_{n-2} = \mathbf{1}$, and $\mathbf{x}_{i-1} \lessdot_{\text{dref}} \mathbf{x}_i$ for $i \in [n-2]$. We say that C is RISING if its label sequence $\lambda(C) = \left(\lambda(\mathbf{x}_0, \mathbf{x}_1), \lambda(\mathbf{x}_1, \mathbf{x}_2), \ldots, \lambda(\mathbf{x}_{n-3}, \mathbf{x}_{n-2})\right)$ is strictly increasing. Conversely, we say that C is FALLING if $\lambda(C)$ is weakly decreasing.

✎ Exercise 4

Let $\mathbf{x}, \mathbf{y} \in \Pi_n$ with $\mathbf{x} \leq_{\text{dref}} \mathbf{y}$. Define $M_{\mathbf{x}} = \{\min B \mid B \in \mathbf{x}\}$ and $M_{\mathbf{y}} = \{\min B \mid B \in \mathbf{y}\}$. Show that for every maximal chain in $[\mathbf{x}, \mathbf{y}]_\Pi$ its label sequence is a permutation of $M_{\mathbf{x}} \setminus M_{\mathbf{y}}$.

The next two statements should be understood with respect to the labeling in (3).

Lemma 1.15

For $n \geq 1$ there exists a unique rising maximal chain in every interval of $(\Pi_n, \leq_{\text{dref}})$.

Proof. The statement is clearly true for $n \leq 2$. Now let $n > 2$. Let $\mathbf{x}, \mathbf{y} \in \Pi_n$ with $\mathbf{x} \leq_{\text{dref}} \mathbf{y}$. Let X be the set of labels of $[\mathbf{x}, \mathbf{y}]_\Pi$ as defined in Exercise 4. A chain is rising in $[\mathbf{x}, \mathbf{y}]_\Pi$ if and only if its label sequence is the identity permutation of X. Let $c = \min X$. We are done if we can show that there is a unique upper cover of $\mathbf{x}$ in $[\mathbf{x},\mathbf{y}]_\Pi$, say $\mathbf{x}'$, with $\lambda(\mathbf{x},\mathbf{x}') = c$. By construction $\mathbf{x}'$ exists.

Say that there is another upper cover $\mathbf{x}''$ of $\mathbf{x}$ in $[\mathbf{x},\mathbf{y}]_\Pi$ with $\lambda(\mathbf{x},\mathbf{x}'') = c$. This means that there is a block $B \in \mathbf{x}$ with $c \in B$, and two other blocks $B', B'' \in \mathbf{x}$ such that $\max\{\min B, \min B'\} = c = \max\{\min B, \min B''\}$. It follows further that there is some $C \in \mathbf{y}$ with $B \cup B' \cup B'' \subseteq C$. Then, however, we have $d = \max\{\min B', \min B''\} < c$ and $d \in X$, which contradicts the minimality of c.

By induction, there is a unique rising maximal chain in $[\mathbf{x}',\mathbf{y}]_\Pi$ which can be extended to a maximal chain of $[\mathbf{x},\mathbf{y}]_\Pi$, and we are done. $\square$

EXAMPLE 1.16

Let

$$\mathbf{x} = 1 \mid 2\ 14 \mid 3\ 4 \mid 5 \mid 6\ 7 \mid 8 \mid 9\ 13 \mid 10\ 12 \mid 11 \mid 15 \mid 16,$$
$$\mathbf{y} = 1\ 6\ 7 \mid 2\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 10\ 12\ 13 \mid 11 \mid 15 \mid 16.$$

It is quickly verified that $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$, and that $M_{\mathbf{x}} = \{1,2,3,5,6,8,9,10,11,15,16\}$ and $M_{\mathbf{y}} = \{1,2,3,9,11,15,16\}$. Therefore

$$M_{\mathbf{x}} \setminus M_{\mathbf{y}} = \{5,6,8,10\}.$$

The unique rising maximal chain in $[\mathbf{x},\mathbf{y}]_\Pi$ is $\{\mathbf{x}_0,\mathbf{x}_1,\mathbf{x}_2,\mathbf{x}_3,\mathbf{x}_4\}$ with

$$\mathbf{x}_0 = 1 \mid 2\ 14 \mid 3\ 4 \mid 5 \mid 6\ 7 \mid 8 \mid 9\ 13 \mid 10\ 12 \mid 11 \mid 15 \mid 16,$$
$$\mathbf{x}_1 = 1 \mid 2\ 14 \mid 3\ 4\ 5 \mid 6\ 7 \mid 8 \mid 9\ 13 \mid 10\ 12 \mid 11 \mid 15 \mid 16,$$
$$\mathbf{x}_2 = 1\ 6\ 7 \mid 2\ 14 \mid 3\ 4\ 5 \mid 8 \mid 9\ 13 \mid 10\ 12 \mid 11 \mid 15 \mid 16,$$
$$\mathbf{x}_3 = 1\ 6\ 7 \mid 2\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 13 \mid 10\ 12 \mid 11 \mid 15 \mid 16,$$
$$\mathbf{x}_4 = 1\ 6\ 7 \mid 2\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 10\ 12\ 13 \mid 11 \mid 15 \mid 16.$$

Let $\mathbf{a}_{i,j}$ denote the set partition whose only non-singleton block is $\{i,j\}$. The set $\{\mathbf{a}_{i,j} \mid 1 \leq i < j \leq n\}$ is the set of ATOMS of $(\Pi_n, \leq_{\mathrm{dref}})$, i.e. all elements of rank 1.

LEMMA 1.17

For $n \geq 1$ there exist $(n-1)!$ falling maximal chains in $(\Pi_n, \leq_{\mathrm{dref}})$.

Proof. The statement is clearly true for $n \leq 2$. Now let $n > 2$. In view of Exercise 4 a maximal chain can only be falling if its first cover relation is labeled by n. This is by construction only satisfied for maximal chains containing an atom $\mathbf{a}_{i,n}$ for $i \in [n-1]$. The interval $[\mathbf{a}_{i,n},\mathbf{1}]_\Pi$ is isomorphic to $(\Pi_{n-1}, \leq_{\mathrm{dref}})$ by virtue of Lemma 1.12. As in the proof of Lemma 1.15, any falling chain of $[\mathbf{a}_{i,n},\mathbf{1}]_\Pi$ can be extended to a falling chain of $(\Pi_n, \leq_{\mathrm{dref}})$, so that the claim follows by induction. $\square$

Proof of Proposition 1.13. Lemma 1.15 implies that the edge-labeling λ from (3) has the properties required by Proposition 1.11, and Lemma 1.17 implies that there are $(n-1)!$ falling maximal chains in the whole partition lattice $(\Pi_n, \leq_{\mathrm{dref}})$. This concludes the proof. $\square$

We conclude this section with the enumeration of *all* maximal chains in $(\Pi_n, \leq_{\mathrm{dref}})$.

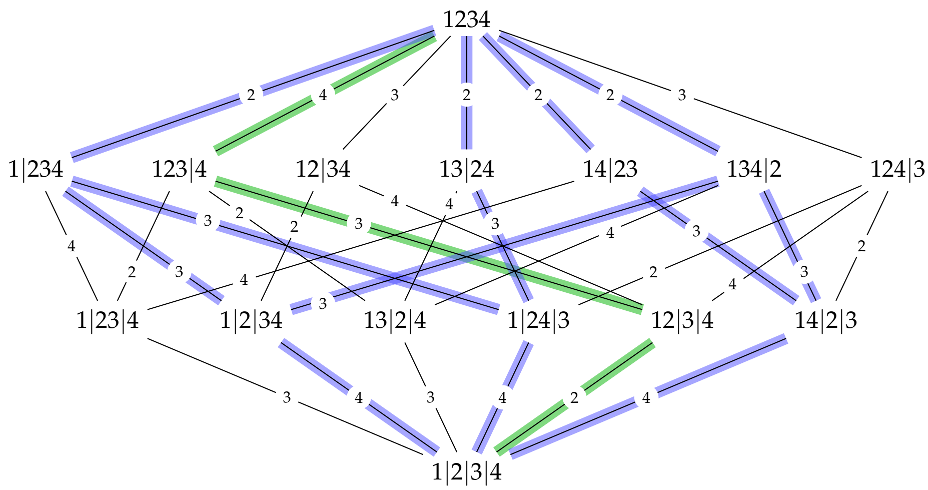

FIGURE 3. The poset $(\Pi_4, \leq_{\mathrm{dref}})$ with the edge-labeling λ. The unique rising maximal chain is marked in green, and the five falling maximal chains are marked in blue.

LEMMA 1.18

For $n \geq 1$ the number of maximal chains in $(\Pi_n, \leq_{\mathrm{dref}})$ is $\frac{n!(n-1)!}{2^{n-1}}$.

Proof. We proceed by induction on n. Observe that each maximal chain needs to pass through exactly one partition of the form $\mathbf{a}_{i,j}$. It is quickly verified that there are $\binom{n}{2}$ atoms, and for any atom $\mathbf{x}$, the interval $[\mathbf{x}, \mathbf{1}]_\Pi$ is isomorphic to $(\Pi_{n-1}, \leq_{\mathrm{dref}})$ by virtue of Lemma 1.12. Let c_n denote the number of maximal chains of $(\Pi_n, \leq_{\mathrm{dref}})$. It is quickly verified that $c_1 = 1$, and the previous reasoning yields

$$c_n = \binom{n}{2} c_{n-1} = \binom{n}{2} \frac{(n-1)!(n-2)!}{2^{n-2}} = \frac{n!(n-1)!}{2^{n-1}}.$$

□

2. NONCROSSING SET PARTITIONS

2.1. **Definition.** In this section we want to restrict our attention to a particular subset of all set partitions, namely those that are noncrossing. Most (if not all) of the results stated in this section were obtained first in the seminal paper by G. Kreweras [16].

DEFINITION 2.1

For $n \geq 0$ a set partition $\mathbf{x} \in \Pi_n$ is NONCROSSING if it does not contain four elements $i < j < k < l$ such that $i \sim_{\mathbf{x}} k$ and $j \sim_{\mathbf{x}} l$, but $i \not\sim_{\mathbf{x}} j$.

The attribute "noncrossing" comes from the fact that we can graphically represent such a set partition without crossings in the respective diagram. In a noncrossing set partition, no two arcs in its arc diagram cross, and equivalently, no two polygons in its circle diagram intersect. Let us denote the family of all noncrossing set partitions of $[n]$ by NC_n.

An easy consequence of Definition 2.1 is the fact that for $n \leq 3$ every set partition of $[n]$ is noncrossing. The smallest crossing set partition is $1\ 3 \mid 2\ 4$, which is also the only crossing set partition of $[4]$.

2.2. **Enumeration.** Analogously to Section 1.2, we want to enumerate noncrossing set partitions in two ways. First we count all of them.

PROPOSITION 2.2

For $n \geq 0$ the cardinality of NC_n is given by the n^{th} Catalan number, which is recursively defined by

$$\mathrm{Cat}(n) = \sum_{k=0}^{n-1} \mathrm{Cat}(k)\,\mathrm{Cat}(n-k-1). \tag{4}$$

The initial condition is $\mathrm{Cat}(0) = 1$.

Proof. Let $\mathbf{x} \in NC_n$. If $n = 0$, then there is a unique noncrossing set partition, namely the empty one. Otherwise there are two cases: either $\{1\}$ is a singleton block, or 1 is connected by an arc to some $k \in \{2, 3, \ldots, n\}$.

In the first case, we can remove the block $\{1\}$ and obtain a noncrossing set partition on an $(n-1)$-element set, of which by induction there exist $\mathrm{Cat}(n-1)$-many. In the second case, the arc between 1 and k breaks $\mathbf{x}$ into two pieces, one piece involving only the integers $\{2, 3, \ldots, k-1\}$ and another piece involving only the integers $\{k, k+1, \ldots, n\}$. (Note that the block containing 1 might have more elements than just 1 and k. However, any other element in this block must be larger than k.)

Since $\mathbf{x}$ is noncrossing these pieces are themselves (mutually disjoint) noncrossing set partitions on a $(k-2)$- and an $(n-k+1)$-element set, respectively. By induction the number of possibilities of these pieces is $\mathrm{Cat}(k-2)\,\mathrm{Cat}(n-k+1)$. We therefore obtain

$$\begin{aligned}\mathrm{Cat}(n) &= \mathrm{Cat}(n-1) + \sum_{k=2}^{n} \mathrm{Cat}(k-2)\,\mathrm{Cat}(n-k+1)\\ &= \mathrm{Cat}(n-1) + \sum_{k=0}^{n-2} \mathrm{Cat}(k)\,\mathrm{Cat}(n-k-1)\\ &= \sum_{k=0}^{n-1} \mathrm{Cat}(k)\,\mathrm{Cat}(n-k-1),\end{aligned}$$

since $\mathrm{Cat}(0) = 1$. $\square$

We will show bijectively in Section 2.6 that the Catalan numbers admit the following explicit form:

$$\mathrm{Cat}(n) = \frac{1}{n+1}\binom{2n}{n}.$$

Let us now count the noncrossing set partitions that have a fixed number of blocks.

PROPOSITION 2.3

For $n \geq 0$ and $k \in [n]$ the number of noncrossing set partitions of $[n]$ with exactly k blocks is given by the k^{th} Narayana number, which is recursively defined by

$$\mathrm{Nar}(n,k) = \mathrm{Nar}(n-1,k-1) + \sum_{i=0}^{n-2}\sum_{j=0}^{k-1} \mathrm{Nar}(i,j)\,\mathrm{Nar}(n-i-1,k-j).$$

The initial conditions are $\mathrm{Nar}(0,0) = 1$ and $\mathrm{Nar}(n,0) = 0$ for $n > 0$. Moreover, we have $\mathrm{Nar}(n,k) = 0$ for $k > n$.

Proof. Let $\mathbf{x} \in NC_n$ with $\mathrm{bl}(\mathbf{x}) = k$. If $n = 0$, then the empty partition is the only noncrossing set partition, which consists of zero blocks. For $n > 0$, we quickly observe that every noncrossing set partition has at least one and at most n blocks. We now proceed analogously to Proposition 2.2, and observe that there are two cases: either $\{1\}$ is a singleton block, or 1 is connected by an arc to some $i \in \{2,3,\ldots,n\}$.

In the first case, we can remove the block $\{1\}$ and obtain a noncrossing set partition on an $(n-1)$-element set with $k-1$ blocks, of which by induction exist $\mathrm{Nar}(n-1,k-1)$-many. In the second case, the arc between 1 and i breaks $\mathbf{x}$ into two pieces, one piece involving only the integers $\{2,3,\ldots,i-1\}$ and another piece involving only the integers $\{i,i+1,\ldots,n\}$.

Since $\mathbf{x}$ is noncrossing these pieces are themselves (mutually disjoint) noncrossing set partitions on a $i-2$- and an $n-i+1$-element set, respectively. This first piece has j blocks for some $j \in \{0,1,\ldots,k-1\}$. (Observe that $j = 0$ is only relevant in the case $i = 2$.) Since $\mathbf{x}$ has k blocks in total, it follows that the second piece needs to have exactly $k-j$ blocks. We therefore obtain

$$\begin{aligned}\mathrm{Nar}(n,k) &= \mathrm{Nar}(n-1,k-1) + \sum_{i=2}^{n}\sum_{j=0}^{k-1} \mathrm{Nar}(i-2,j)\,\mathrm{Nar}(n-i+1,k-j)\\ &= \mathrm{Nar}(n-1,k-1) + \sum_{i=0}^{n-2}\sum_{j=0}^{k-1} \mathrm{Nar}(i,j)\,\mathrm{Nar}(n-i-1,k-j).\end{aligned}$$

□

✎ EXERCISE 5

Verify Proposition 2.3 for $n = 5$ and $k = 3$. Draw the 20 noncrossing set partitions of $[5]$ with three blocks.

COROLLARY 2.4

For $n \geq 0$ we have

$$\mathrm{Cat}(n) = \sum_{k=0}^{n} \mathrm{Nar}(n,k).$$

We will prove bijectively in Section 2.6 that the Narayana numbers admit the following explicit form:

$$\mathrm{Nar}(n,k) = \frac{1}{n}\binom{n}{k}\binom{n}{k-1}.$$

2.3. **Lattice Property.** Since noncrossing set partitions are in particular set partitions, we can consider them under dual refinement order. More precisely, we may wonder what the subposet of $(\Pi_n, \leq_{\mathrm{dref}})$ induced by NC_n looks like. Figure 4 shows this poset for $n = 4$. We will see that it in fact inherits many nice properties. To that end let us define a map on Π_n that sends a set partition to the smallest noncrossing set partition that it refines.

DEFINITION 2.5

Let $\mathbf{x} \in \Pi_n$. Consider the graph whose vertices are the blocks of $\mathbf{x}$, and where there exists an edge between two blocks if and only if they are crossing. Let $\overline{\mathbf{x}}$ be the set partition whose blocks are given by the union over the elements of the connected components of this graph.

LEMMA 2.6

For $\mathbf{x} \in \Pi_n$ we have $\overline{\mathbf{x}} \in NC_n$.

Proof. Let $C, C' \in \overline{\mathbf{x}}$ with $i, k \in C$ and $j, l \in C'$ for some $i < j < k < l$. Observe that neither C nor C' can be blocks of $\mathbf{x}$, because they would have been joined in the process of creating $\overline{\mathbf{x}}$. It follows that either C or C' consists of a union of blocks of $\mathbf{x}$.

We present the case where $C \in \mathbf{x}$ and $C' = B'_1 \uplus B'_2 \uplus \cdots \uplus B'_s$, where $B'_r \in \mathbf{x}$ for $r \in [s]$, and $s > 1$. (The other cases are analogous.) Without loss of generality we may assume that $j \in B'_1$ and $l \in B'_2$. Since B'_1 and B'_2 are crossing in $\mathbf{x}$, we have $a, c \in B'_1$ and $b, d \in B'_2$ with $a < b < c < d$. In particular we may assume that $a = j$ and $d = l$. There are three cases: (i) $b < c < k$, (ii) $b < k < c$, and (iii) $k < b < c$, but in each case we find that B'_1 or B'_2 crosses with C. $\square$

EXAMPLE 2.7

Consider the set partition $\mathbf{x} = 1\ 6\ 7 \mid 2\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 10\ 12\ 13 \mid 11 \mid 15 \mid 16$ from Example 1.4 again. The associated graph is

$$\begin{array}{cccc} \{2,8,14\} & \{9,10,12,13\} & \{16\} & \\ | & & & \\ \{1,6,7\} & \{3,4,5\} & \{11\} & \{15\} \end{array}$$

As a consequence, we obtain $\overline{\mathbf{x}} = 1\ 2\ 6\ 7\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 10\ 12\ 13 \mid 11 \mid 15 \mid 16$.

✎ EXERCISE 6

Show that the map $\bar{\cdot} : \Pi_n \to \Pi_n$ is a closure operator with respect to $\leq_{\mathrm{dref}}$, i.e. it satisfies

(i) $\mathbf{x} \leq_{\mathrm{dref}} \overline{\mathbf{x}}$, (EXTENSITIVITY)
(ii) $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$ implies $\overline{\mathbf{x}} \leq_{\mathrm{dref}} \overline{\mathbf{y}}$, (MONOTONICITY)
(iii) $\overline{\overline{\mathbf{x}}} = \overline{\mathbf{x}}$. (IDEMPOTENCE)

PROPOSITION 2.8

For $n \geq 0$ the poset $(NC_n, \leq_{\mathrm{dref}})$ is in fact a lattice.

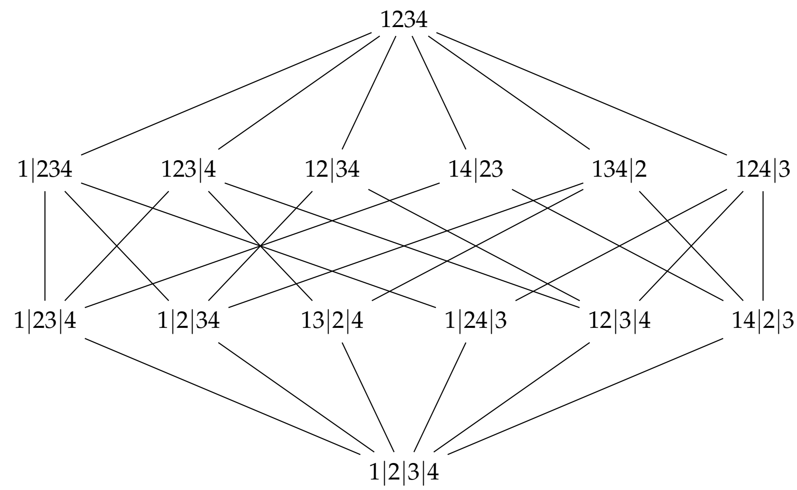

FIGURE 4. The poset $(NC_4, \leq_{\mathrm{dref}})$.

Proof. The proof works along the same lines as the proof of Proposition 1.9. The key observation is that the intersection of two noncrossing set partitions defined in (1) is again noncrossing. □

Proposition 2.8 suggests that $\mathbf{x} \wedge_{\Pi} \mathbf{y} = \mathbf{x} \wedge_{NC} \mathbf{y}$. In light of Exercise 6 it follows that

$$\mathbf{x} \vee_{NC} \mathbf{y} = \overline{\mathbf{x} \vee_{\Pi} \mathbf{y}}.$$

Consider for instance $\mathbf{x} = 1\ 3 \mid 2 \mid 4$ and $\mathbf{y} = 1 \mid 2\ 4 \mid 3$. We have $\mathbf{x} \vee_{\Pi} \mathbf{y} = 1\ 3 \mid 2\ 4$ and $\mathbf{x} \vee_{NC} \mathbf{y} = 1\,2\,3\,4 = \overline{1\,3 \mid 2\,4}$.

COROLLARY 2.9

For $n \geq 0$ the lattice $(NC_n, \leq_{\mathrm{dref}})$ is a meet-sublattice of $(\Pi_n, \leq_{\mathrm{dref}})$.

2.4. **Self-Duality.** In contrast to $(\Pi_n, \leq_{\mathrm{dref}})$ the noncrossing partition lattice has a striking property: it is (locally) self-dual. This means that every interval of $(NC_n, \leq_{\mathrm{dref}})$ is isomorphic to its dual. One way to observe the duality works by exhibiting a particular anti-automomorphism. In fact, we find a whole family of (anti)-automorphisms.

For $X \subseteq \mathbb{Z}$ and $a \in \mathbb{Z} \setminus \{0\}$ define the DILATION of X by $aX = \{ax \mid x \in X\}$, and define the TRANSLATION of X by $X + a = \{x + a \mid x \in X\}$. It is straightforward to generalize these definitions to families of sets, and therefore to set partitions.

DEFINITION 2.10

Let $\mathbf{x} \in NC_n$. The KREWERAS COMPLEMENT of $\mathbf{x}$ is the coarsest $\mathbf{y} \in NC_n$ such that the union $2\mathbf{x} - 1 \cup 2\mathbf{y}$ is a noncrossing set partition of $[2n]$. We usually write $K(\mathbf{x})$ instead of $\mathbf{y}$.

Example 2.11

Let $\mathbf{x} = 1\;2\;6\;7\;8\;14 \mid 3\;4\;5 \mid 9\;10\;12\;13 \mid 11 \mid 15 \mid 16 \in NC_{16}$. Its circle diagram is shown in Figure 5a. Figure 5c shows the circle diagram of $K(\mathbf{x})$, and Figure 5b shows the superposition of both diagrams, i.e. $2\mathbf{x} - 1 \cup 2K(\mathbf{x})$.

Lemma 2.12

For $\mathbf{x} \in NC_n$ we have $\mathbf{x} \wedge_{NC} K(\mathbf{x}) = \mathbf{0}$ and $\mathbf{x} \vee_{NC} K(\mathbf{x}) = \mathbf{1}$. Moreover, for $\mathbf{x}, \mathbf{y} \in NC_n$ we have $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$ if and only if $K(\mathbf{y}) \leq_{\mathrm{dref}} K(\mathbf{x})$.

Proof. The claim on join and meet is straightforward from the construction. So is the claim that the map K reverses order. Intuitively this is clear, because the smaller $\mathbf{x}$ is, the more space is there for the parts of $K(\mathbf{x})$. □

It is also straightforward from the construction that for each block $\{i_1, i_2, \ldots, i_s\} \in \mathbf{x}$ the block $\{i_1 - 1, i_2 - 1, \ldots, i_s - 1\}$ belongs to $K^2(\mathbf{x})$. (Addition is considered modulo n.) In other words, the circle diagram of $K^2(\mathbf{x})$ is a counterclockwise rotation of the circle diagram of $\mathbf{x}$ by an angle of $2\pi/n$. As a consequence the map K^{2n} is the identity, which implies that each of the maps K^i for $i \in [2n]$ is a bijection.

Proposition 2.13

For $n \geq 0$ the lattice $(NC_n, \leq_{\mathrm{dref}})$ is self-dual.

Proof. This follows from the fact that K is a order-reversing bijection. □

We can use the Kreweras complement to describe the structure of the intervals in $(NC_n, \leq_{\mathrm{dref}})$, and obtain a result analogous to Lemma 1.12. For a finite set M, let K_M denote the Kreweras complement in $(NC_M, \leq_{\mathrm{dref}})$

Lemma 2.14

Let $n \geq 0$ and let $\mathbf{x}, \mathbf{y} \in NC_n$ with $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$. We have

$$[\mathbf{x}, \mathbf{y}]_{NC} \cong \prod_{B \in \mathbf{y}} \prod_{X \in K_B(\mathbf{x}_{|B})} (NC_X, \leq_{\mathrm{dref}}).$$

Proof. It follows analogously to Lemma 1.12 that

$$[\mathbf{x}, \mathbf{y}]_{NC} \cong \prod_{B \in \mathbf{y}} [\mathbf{x}_{|B}, \mathbf{1}_{|B}]_{NC}.$$

Since the Kreweras complement is an anti-automorphism, we obtain

$$[\mathbf{x}_{|B}, \mathbf{1}_{|B}]_{NC} \cong [\mathbf{0}_{|B}, K_B(\mathbf{x}_{|B})]_{NC}.$$

We can once more decompose this as follows

$$[\mathbf{0}_{|B}, K_B(\mathbf{x}_{|B})]_{NC} \cong \prod_{X \in K_B(\mathbf{x}_{|B})} [\mathbf{0}_{|X}, \mathbf{1}_{|X}]_{NC} \cong \prod_{X \in K_B(\mathbf{x}_{|B})} (NC_X, \leq_{\mathrm{dref}}).$$

□

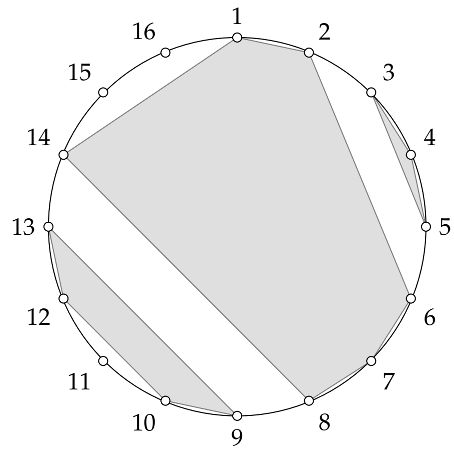

(A) A noncrossing set partition.

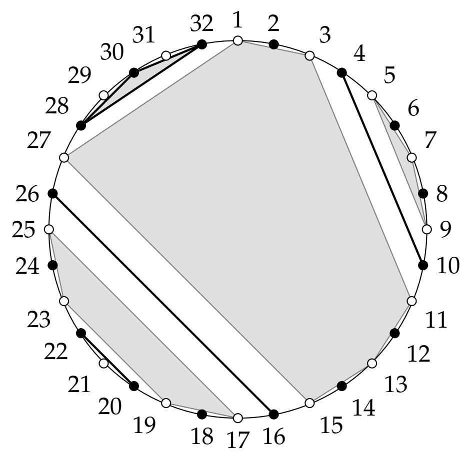

(B) The superposition of the diagrams in Figures 5a and 5c.

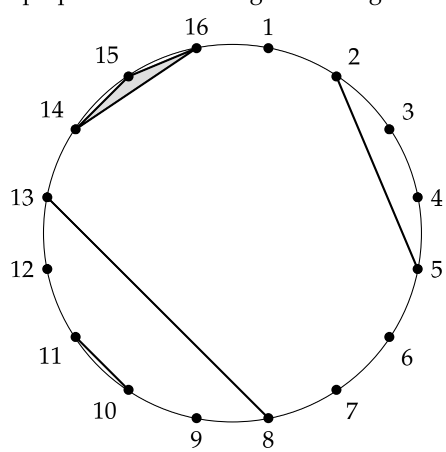

(C) The Kreweras complement of the noncrossing set partition in Figure 5a.

2.5. **Möbius Function.** It is immediate that the labeling of the set partition lattice defined in (3) restricts to a labeling of $(NC_n, \leq_{\text{dref}})$. But it does more than that: it also inherits the property that there is a unique rising maximal chain in each interval.

LEMMA 2.15

For $n \geq 1$ there exists a unique rising maximal chain in every interval of $(NC_n, \leq_{\text{dref}})$.

Proof. Let $\mathbf{x}, \mathbf{y} \in NC_n$, and let M be a maximal chain in $[\mathbf{x}, \mathbf{y}]_\Pi$ that does not belong to $[\mathbf{x}, \mathbf{y}]_{NC}$. There then exists a maximal consecutive sequence $\{\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_s\} \subseteq M$ of crossing set partitions. In other words, $\mathbf{z}_i \notin NC_n$ for $i \in [s]$ and $\mathbf{z}_i \lessdot_{\text{dref}} \mathbf{z}_{i+1}$ for $i \in [s-1]$ and there are $\mathbf{x}', \mathbf{y}' \in M \cap NC_n$ with $\mathbf{x}' \lessdot_{\text{dref}} \mathbf{z}_1$ and $\mathbf{z}_s \lessdot_{\text{dref}} \mathbf{y}'$.

We can thus find two crossing blocks $B, B' \in \mathbf{z}_1$, and one of the blocks is also present in $\mathbf{x}'$. Without loss of generality we have $i, k \in B$ and $j, l \in B'$ with $i < j < k < l$, and we can assume that $i = \min B$ and $j = \min B'$. If $B \in \mathbf{x}'$, then there must be two blocks $C, C' \in \mathbf{x}'$ with $B' = C \uplus C'$ such that $j = \min C$ and $l = \min C'$ without loss of generality. If $B' \in \mathbf{x}'$, then there must be two blocks C, C' with $B = C \uplus C'$ such that $i = \min C$ and $k = \min C'$ without loss of generality. In any case we have $\lambda(\mathbf{x}', \mathbf{z}_1) \geq k$.

Since $\mathbf{y}' \in NC_n$ there must be $j \in [k]$ and two blocks $D, D' \in \mathbf{z}_j$ with $B \subseteq D$ and $B' \subseteq D'$ such that $\mathbf{z}_{j+1} = \mathbf{z}_j \setminus \{D, D'\} \cup (D \cup D')$. (If $j = k$, then we understand $\mathbf{z}_{j+1} = \mathbf{y}'$.) It follows that $\min D \leq i$ and $\min D' \leq j$, which implies

$$\lambda(\mathbf{x}_j, \mathbf{x}_{j+1}) \leq j < k \leq \lambda(\mathbf{x}, \mathbf{x}_1).$$

It follows that M is not rising in $[\mathbf{x}, \mathbf{y}]_\Pi$.

As a consequence every rising chain in every interval of $(\Pi_n, \leq_{\text{dref}})$ belongs to the corresponding interval in $(NC_n, \leq_{\text{dref}})$. Lemma 1.15 implies that every interval of $(\Pi_n, \leq_{\text{dref}})$ has exactly one rising maximal chain, which concludes the proof. □

LEMMA 2.16

For $n \geq 1$ there exist $\mathrm{Cat}(n-1)$ falling maximal chains in $(NC, \leq_{\text{dref}})$.

Proof. For a maximal chain in $(NC_n, \leq_{\text{dref}})$ to be falling its last label needs to be 2, which can only be obtained if it contains a noncrossing set partition consisting of two blocks B_1 and B_2 with $1 \in B_1$ and $2 \in B_2$. Since B_1 and B_2 are noncrossing we conclude that $B_1 = \{1, k+1, k+2, \ldots, n\}$ and $B_2 = \{2, 3, \ldots, k\}$ for some $k \in \{2, 3, \ldots, n\}$, and we write $\mathbf{x}_k$ for the resulting set partition. Lemma 2.14 implies that $[\mathbf{0}, \mathbf{x}_k]_{NC} \cong (NC_{k-1}, \leq_{\text{dref}}) \times (NC_{n-k+1}, \leq_{\text{dref}})$.

Let f_n denote the number of falling maximal chains in $(NC_n, \leq_{\text{dref}})$. It is quickly checked by induction that the number of falling maximal chains in $(NC_{k-1}, \leq_{\text{dref}}) \times (NC_{n-k+1}, \leq_{\text{dref}})$ equals $f_{k-1} f_{n-k+1}$. (As in the proof of Proposition 1.13 we can relate the number of falling maximal chains to the value of the Möbius function between least and greatest element, and since the Möbius function on a direct product of posets equals the product of the Möbius functions on the factors, the claim follows.)

We thus obtain by induction and (4) that

$$f_n = \sum_{k=2}^{n} f_{k-1} f_{n-k+1} = \sum_{k=0}^{n-2} f_{k+1} f_{n-k-1} = \sum_{k=0}^{n-2} \mathrm{Cat}(k)\, \mathrm{Cat}\big((n-1)-k-1\big) = \mathrm{Cat}(n-1)$$

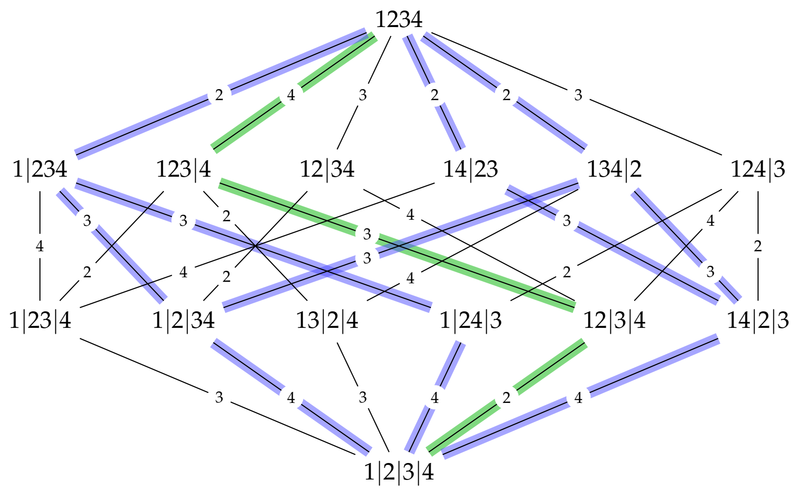

FIGURE 6. The lattice $(NC_4, \leq_{\text{dref}})$ with the edge-labeling λ. The unique rising maximal chain is marked in green, and the five falling maximal chains are marked in blue.

as desired. $\square$

We conclude this section with the computation of particular values of the Möbius function of $(NC_n, \leq_{\text{dref}})$. Let us abbreviate $\mu(NC_n) = \mu_{NC_n}(\mathbf{0}, \mathbf{1})$.

PROPOSITION 2.17

For $n \geq 1$ we have $\mu(NC_n) = (-1)^{n-1}\,\text{Cat}(n-1)$.

Proof. This follows from Lemmas 2.15 and 2.16 analogously to the proof of Proposition 1.13. $\square$

COROLLARY 2.18

Let $n \geq 1$ and let $\mathbf{x}, \mathbf{y} \in NC_n$ with $\mathbf{x} \leq_{\text{dref}} \mathbf{y}$. Then

$$\mu_{NC_n}(\mathbf{x}, \mathbf{y}) = (-1)^{\text{rk}(\mathbf{x}) - \text{rk}(\mathbf{y})} \prod_{B \in \mathbf{y}} \prod_{X \in K_{|B}(\mathbf{x}_{|B})} \text{Cat}(|X| - 1).$$

Figure 6 shows the lattice $(NC_4, \leq_{\text{dref}})$ together with the labeling λ from (3). The unique rising maximal chain is highlighted in green, and the five falling maximal chains are highlighted in blue.

2.6. **Chain Enumeration.** Let us now count the number of chains in $(NC_n, \leq_{\text{dref}})$ whose elements have given ranks. This result (and its corollaries) are due to P. Edelman [10], and they are purely combinatorial in nature.

A central result for this computation is the following Cycle Lemma. It was rediscovered several times in varying degrees of generality; see [8] for a historical account.

Lemma 2.19: The Cycle Lemma

Let $k, m, n \in \mathbb{N}$ with $m \geq kn$. For any sequence of $p_1 p_2 \cdots p_{m+n}$ of m boxes and n circles, exactly $m - kn$ out of the $m+n$ cyclic permutations $p_j p_{j+1} \cdots p_{m+n} p_1 p_2 \cdots p_{j-1}$ have the property that in every prefix the number of boxes is more than k times the number of circles.

Proof. Let us call a sequence k-DOMINATING if it has the property that in every prefix the number of boxes is more than k times the number of circles. Now arrange $p_1 p_2 \cdots p_{m+n}$ on a cycle. If we remove a sequence of k boxes followed by a circle from this cycle, then this process does not change the number of k-dominating sequences. This is because no k-dominating sequence can start with one of the removed figures, and removing this part from a sequence does not affect whether it is k-dominating or not. Since $m \geq kn$ we can always find such a sequence. The reduced sequence has $m-k$ boxes and $n-1$ circles, and we conclude that $m - k \geq k(n-1)$. We can therefore repeat this procedure until we are left with a sequence of $m - kn$ boxes. The positions of these boxes in the original sequence indicate the beginning of a k-dominating sequence. $\square$

Example 2.20

Let $k = 2$, $m = 5$, and $n = 2$. A sequence of five boxes and two circles is 2-dominating if there are at least three boxes before the first circle, and at least five boxes before the second circle. Consider the sequence $\square \bigcirc \square \square \bigcirc \square \square$. We can remove twice the sequence $\square \square \bigcirc$, and after doing so we are left with the fourth box (i.e. the second-to-last entry of the sequence). It follows that the cyclic permutation $\square \square \square \bigcirc \square \square \bigcirc$ is the only 2-dominating rearrangement of the sequence.

✎ Exercise 7

Find the three cyclic permutations of the above sequence that are 1-dominating.

Let us now apply the Cycle Lemma in the case $m = n+1$. For $b, n \in \mathbb{N}$ define

$$\sigma_b(n) = b\ b{+}1\ b{+}2\ \ldots\ n\ 1\ 2\ \ldots\ b{-}1.$$

We want to insert left and right parentheses (or boxes and circles) into this word, and the way this is done shall be determined by a $k+1$-tuple $(L; R_1, R_2, \ldots, R_k)$ of non-empty subsets of $[n]$ with

$$|L| = 1 + \sum_{i=1}^{k} |R_i|.$$

Given such a $k+1$-tuple we proceed as follows: for every $l \in L$ we put a left parenthesis to the left of l, and for each appearance of r in the sets R_i we put a right parenthesis to the right of r. Let us denote the resulting parenthesized string by $\hat{\sigma}_b(n)$. In view of the Cycle Lemma 2.19 there is a unique b such that $\hat{\sigma}_b(n)$ is WELL-PARENTHESIZED, i.e. it begins with a left parenthesis and every other left parenthesis closes properly.

As the last ingredient we consider a different representation of $\mathbf{x} = \{B_1, B_2, \ldots, B_k\} \in NC_n$, where we order the blocks (and their entries) with respect to a given string $\sigma_b(n)$ as follows: B_1 is the block containing b, and for $i > 1$ the block B_i is the block that contains the leftmost number in $\sigma_b(n)$ that is not contained in $\bigcup_{j=1}^{i-1} B_j$.

The basic idea of [10], which explains the combinatorics behind the chain enumeration in the noncrossing partition lattice, is the correspondence established in the following proposition, which at the same time provides an explicit form of the Narayana numbers defined earlier.

PROPOSITION 2.21

For $n \geq 1$ and $k \in [n]$ the number of noncrossing set partitions of $[n]$ with exactly k blocks is

$$\mathrm{Nar}(n,k) = \frac{1}{n}\binom{n}{k}\binom{n}{k-1}.$$

Proof. Let (L,R) be a pair of subsets of $[n]$ such that $\lvert L\rvert = \lvert R\rvert + 1 = k$. The Cycle Lemma 2.19 implies that there is a unique number $b \in [n]$ such that $\hat{\sigma}_b(n)$ is well-parenthesized.

We now construct a noncrossing set partition of $[n]$ from this string as follows. First we add a right parenthesis at the end of $\hat{\sigma}_b(n)$. Then we look for pairs of closing parentheses, and if the induced substring does not contain parentheses, then we take the elements of this induced substring and put them in a block. We now remove these elements together with their enclosing parentheses. Repeat until all parentheses are removed. Since $\hat{\sigma}_b(n)$ is well-parenthesized the resulting family of sets is a partition of $[n]$, and the construction ensures that this partition is noncrossing.

Conversely, let $\mathbf{x} \in NC_n$ with $\mathrm{bl}(\mathbf{x}) = k$, and fix $b \in [n]$. Order the blocks of $\mathbf{x}$ with respect to $\sigma_b(n)$. Let L be the set of the first numbers in each block, and let R be the set of last numbers in each block except B_1.

This establishes a bijection between the set of pairs (L,R) with $L, R \subseteq [n]$ and $\lvert L\rvert = \lvert R\rvert + 1 = k$ and the set of pairs $(\mathbf{x}, b)$ for $\mathbf{x} \in NC_n$ with $\mathrm{bl}(\mathbf{x}) = k$ and $b \in [n]$. We thus obtain the equality

$$\binom{n}{k}\binom{n}{k-1} = n \cdot \mathrm{Nar}(n,k),$$

which proves the proposition. $\square$

EXAMPLE 2.22

Consider once more $\mathbf{x} = 1\ 2\ 6\ 7\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 10\ 12\ 13 \mid 11 \mid 15 \mid 16 \in NC_{16}$, and pick $b = 12$. The reordering of $\mathbf{x}$ with respect to $\sigma_{12}(16)$ looks as follows

$$\mathbf{x} = 12\ 13\ 9\ 10 \mid 14\ 1\ 2\ 6\ 7\ 8 \mid 15 \mid 16 \mid 3\ 4\ 5 \mid 11,$$

and we obtain $L = \{3, 11, 12, 14, 15, 16\}$ and $R = \{5, 8, 11, 15, 16\}$. The corresponding well-parenthesized string is

$$\hat{\sigma}_{12}(16) = (12\ 13(14(15)(16)1\ 2(3\ 4\ 5)6\ 7\ 8)9\ 10(11),$$

and we reobtain $\mathbf{x}$.

The next step is to enumerate chains in $(NC_n, \leq_{\mathrm{dref}})$ whose elements have given ranks. More precisely, let $(t_1, t_2, \ldots, t_k)$ be a tuple of integers with $0 < t_1 < t_2 < \cdots < t_k < n-1$. Let $\mathcal{N}_n(t_1, t_2, \ldots, t_k)$ denote the number of chains $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k\}$ in $(NC_n, \leq_{\mathrm{dref}})$ with $\mathrm{rk}(\mathbf{x}_i) = t_i$ for $i \in k$.

PROPOSITION 2.23

Let $t_0 = 0$ and $t_{k+1} = n-1$, and define $s_i = t_i - t_{i-1}$ for $i \in [k+1]$. Then

$$\mathcal{N}_n(t_1, t_2, \ldots, t_k) = \frac{1}{n}\binom{n}{s_1}\binom{n}{s_2}\cdots\binom{n}{s_{k+1}}.$$

Proof. The proof uses the bijective idea introduced in Proposition 2.21. This time, however, we consider $k+1$-tuples of subsets of $[n]$ instead of pairs.

Indeed, let $(L; R_1, R_2, \ldots, R_k)$ be such a $k+1$-tuple of subsets of $[n]$ with the property that $\lvert L\rvert = n - s_1$ and $\lvert R_i\rvert = s_{i+1}$ for $i \in [k]$. Note that

$$\sum_{i=1}^{k}\lvert R_i\rvert = \sum_{i=1}^{k} s_{i+1} = \sum_{i=1}^{k}(t_{i+1} - t_i) = t_{k+1} - t_1 = n - 1 - t_1 = \lvert L\rvert - 1.$$

The Cycle Lemma 2.19 implies that there is a unique $b \in [n]$ such that the parenthesized string $\hat{\sigma}_b(n)$ coming from $(L; R_1, R_2, \ldots, R_k)$ is well-parenthesized. Let $\mathbf{x}_1$ denote the noncrossing set partition that is constructed from $\hat{\sigma}_b(n)$. We then have $\mathrm{bl}(\mathbf{x}_1) = \lvert L\rvert$, and therefore $\mathrm{rk}(\mathbf{x}_1) = t_1$. Now for every $r \in R_1$ we remove the first right parenthesis to the right of r and the corresponding left parenthesis. The resulting string is still well-parenthesized, and we thus obtain a noncrossing set partition $\mathbf{x}_2$, which has $\lvert L\rvert - \lvert R_1\rvert$ blocks, and thus $\mathrm{rk}(\mathbf{x}_2) = t_2$. It is also guaranteed by construction that $\mathbf{x}_1 \leq_{\mathrm{dref}} \mathbf{x}_2$. If we continue this process until all parentheses are removed we have obtained a chain $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k\}$ with $\mathrm{rk}(\mathbf{x}_i) = t_i$ for $i \in [k]$.

Conversely, let $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k\}$ be a chain of noncrossing set partitions with $\mathrm{rk}(\mathbf{x}_i) = t_i$ for $i \in [k]$ and fix some $b \in [n]$. As in Proposition 2.21 the pair $(\mathbf{x}_k, b)$ has an associated pair (L_k, R_k) with

$$\lvert R_k\rvert + 1 = \lvert L_k\rvert = \mathrm{bl}(\mathbf{x}_k) = n - t_k.$$

Now we construct a triple $(L_{k-1}; R_{k-1}, R_k)$ from (L_k, R_k) as follows. For each block $B_i \in \mathbf{x}_k$ suppose that it is broken into the blocks $B_{i_1}, B_{i_2}, \ldots, B_{i_s}$ of $\mathbf{x}_{k-1}$, where the order suggested by the indices comes from the ordering of $\mathbf{x}_{k-1}$ with respect to $\sigma_b(n)$. Let F_i consist of the first elements of the blocks $B_{i_2}, B_{i_3}, \ldots, B_{i_s}$, and let E_i consist of the last elements of these blocks. Let

$$R_{k-1} = \bigcup_{i=1}^{\mathrm{bl}(\mathbf{x}_k)} E_i \quad\text{and}\quad L_{k-1} = L_k \cup \bigcup_{i=1}^{\mathrm{bl}(\mathbf{x}_k)} F_i.$$

It follows that

$$\lvert R_{k-1}\rvert = \mathrm{bl}(\mathbf{x}_{k-1}) - \mathrm{bl}(\mathbf{x}_k) = \mathrm{rk}(\mathbf{x}_k) - \mathrm{rk}(\mathbf{x}_{k-1}) = t_k - t_{k-1} = s_k,$$

and $\lvert L_{k-1}\rvert = \lvert L_k\rvert + s_k$. It is clear that $(L_{k-1}; R_{k-1}, R_k)$ corresponds exactly to $\mathbf{x}_{k-1}$. We repeat this process until we have reached $\mathbf{x}_1$, and we obtain a $k+1$-tuple $(L_1; R_1, R_2, \ldots, R_k)$ with the desired properties.

This bijective correspondence proves the following equality:

$$\binom{n}{s_1}\binom{n}{s_2}\cdots\binom{n}{s_{k+1}} = n \cdot \mathcal{N}_n(t_1, t_2, \ldots, t_k),$$

since the left side enumerates precisely the possible $k+1$-tuples $(L; R_1, R_2, \ldots, R_k)$. (Note that $\binom{n}{n-s_1} = \binom{n}{s_1}$.) This concludes the proof. ∎

EXAMPLE 2.24

Consider the following four noncrossing set partitions:

$$\begin{aligned}
\mathbf{x}_1 &= 1 \mid 2 \mid 3\,5 \mid 4 \mid 6 \mid 7 \mid 8 \mid 9 \mid 10 \mid 11 \mid 12 \mid 13 \mid 14 \mid 15 \mid 16,\\
\mathbf{x}_2 &= 1\,14 \mid 2\,8 \mid 3\,5 \mid 4 \mid 6\,7 \mid 9\,12\,13 \mid 10 \mid 11 \mid 15 \mid 16,\\
\mathbf{x}_3 &= 1\,2\,8\,14 \mid 3\,5 \mid 4 \mid 6\,7 \mid 9\,10\,12\,13 \mid 11 \mid 15 \mid 16,\\
\mathbf{x}_4 &= 1\,2\,6\,7\,8\,14 \mid 3\,4\,5 \mid 9\,10\,12\,13 \mid 11 \mid 15 \mid 16,
\end{aligned}$$

and fix $b = 12$. The reordering of $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4$ with respect to $\sigma_{12}(16)$ looks as follows:

$$\begin{aligned}
\mathbf{x}_1 &= 12 \mid 13 \mid 14 \mid 15 \mid 16 \mid 1 \mid 2 \mid 3\,5 \mid 4 \mid 6 \mid 7 \mid 8 \mid 9 \mid 10 \mid 11,\\
\mathbf{x}_2 &= 12\,13\,9 \mid 14\,1 \mid 15 \mid 16 \mid 2\,8 \mid 3\,5 \mid 4 \mid 6\,7 \mid 10 \mid 11,\\
\mathbf{x}_3 &= 12\,13\,9\,10 \mid 14\,1\,2\,8 \mid 15 \mid 16 \mid 3\,5 \mid 4 \mid 6\,7 \mid 11,\\
\mathbf{x}_4 &= 12\,13\,9\,10 \mid 14\,1\,2\,6\,7\,8 \mid 15 \mid 16 \mid 3\,4\,5 \mid 11,
\end{aligned}$$

We have seen in Example 2.22 that $L_4 = \{3, 11, 12, 14, 15, 16\}$ and $R_4 = \{5, 8, 11, 15, 16\}$. The first block of $\mathbf{x}_4$ is also contained in $\mathbf{x}_3$, and thus $E_1^{(3)} = F_1^{(3)} = \emptyset$. The second block $\{14, 1, 2, 6, 7, 8\}$ of $\mathbf{x}_4$ is broken into two blocks in $\mathbf{x}_3$, namely $\{14, 1, 2, 8\}$ and $\{6, 7\}$. We get $E_2^{(3)} = \{7\}$ and $F_2^{(3)} = \{6\}$. The third, fourth and sixth block of $\mathbf{x}_4$ are singletons and thus left intact, which yields $E_3^{(3)} = E_4^{(3)} = E_6^{(3)} = F_3^{(3)} = F_4^{(3)} = F_6^{(3)} = \emptyset$. The fifth block $\{3, 4, 5\}$ is broken again in two blocks $\{3, 5\}$ and $\{4\}$. We obtain $E_5^{(3)} = F_5^{(3)} = \{4\}$. We thus obtain $L_3 = \{3, 4, 6, 11, 12, 14, 15, 16\}$ and $R_3 = \{4, 7\}$.
In the same manner we obtain $L_2 = \{2, 3, 4, 6, 10, 11, 12, 14, 15, 16\}$ and $R_2 = \{8, 10\}$, as well as $L_1 = \{1, 2, 3, 4, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16\}$ and $R_1 = \{1, 7, 8, 9, 13\}$.
The resulting 5-tuple is $(L_1; R_1, R_2, R_3, R_4)$, which induces the parenthesization

$$\hat{\sigma}_{12}(16) = (12(13)(14(15)(16)(1)(2(3(4)5)(6(7))(8)))(9)(10)(11),$$

which corresponds to $\mathbf{x}_1$. If we now remove the parentheses indicated by R_1, we obtain

$$\hat{\sigma}_{12}(16) = (12\ 13(14(15)(16)1(2(3(4)5)(6\ 7)8))9(10)(11),$$

which corresponds to $\mathbf{x}_2$. If we remove the parentheses indicated by R_2, we obtain

$$\hat{\sigma}_{12}(16) = (12\ 13(14(15)(16)1\ 2(3(4)5)(6\ 7)8)9\ 10(11),$$

which corresponds to $\mathbf{x}_3$. Finally, we remove the parentheses indicated by R_3, and we obtain

$$\hat{\sigma}_{12}(16) = (12\ 13(14(15)(16)1\ 2(3\ 4\ 5)6\ 7\ 8)9\ 10(11),$$

which corresponds to $\mathbf{x}_4$.

As an immediate corollary we obtain the number of maximal chains in $(NC_n, \leq_{\text{dref}})$.

COROLLARY 2.25

For $n \geq 1$ the number of maximal chains in $(NC_n, \leq_{\text{dref}})$ is n^{n-2}.

Proof. The desired quantity is precisely $\mathscr{N}_n(1,2,\ldots,n-2)$, so that by Proposition 2.23 we obtain

$$\begin{aligned}\mathscr{N}_n(1,2,\ldots,n-2) &= \frac{1}{n}\underbrace{\binom{n}{1}\binom{n}{1}\cdots\binom{n}{1}}_{n-1\text{ times}}\\ &= \frac{1}{n}n^{n-1}\\ &= n^{n-2}.\end{aligned}$$

□

We can also use the construction from Proposition 2.23 to count the number $\mathscr{Z}_n(m)$ of multichains of length $m-1$ in $(NC_n, \leq_{\mathrm{dref}})$.

PROPOSITION 2.26

For $m,n\geq 1$ we have

$$\mathscr{Z}_n(m) = \frac{1}{n}\binom{mn}{n-1}.$$

Proof. Let us start with an $n-1$-element set $D\subseteq[mn]$. We write $D = S_1\uplus S_2\uplus\cdots\uplus S_m$, where

$$S_j = \bigl\{i \mid i\in D \text{ and } (j-1)n+1\leq i\leq jn\bigr\}.$$

We can turn each of these sets into subsets of $[n]$ by considering

$$S'_j = \bigl\{i \mid (j-1)n+i\in S_j\bigr\}.$$

Let $\bar{S}_1 = [n]\setminus S_1$. We then have

$$\lvert\bar{S}_1\rvert = n - \lvert S_1\rvert = 1 + \sum_{i=2}^{m}\lvert S'_i\rvert.$$

Then the m-tuple $(\bar{S}_1; S'_2, S'_3,\ldots,S'_m)$ then induces a chain $(\mathbf{x}_1,\mathbf{x}_2,\ldots,\mathbf{x}_{m-1})$ together with some $b\in[n]$ via the construction from Proposition 2.23. (Note that if $S'_{i+1}=\emptyset$ for $i\in[m-2]$, then $\mathbf{x}_i=\mathbf{x}_{i+1}$.)

Conversely pick $m-1$ noncrossing set partitions satisfying $\mathbf{x}_1\leq_{\mathrm{dref}}\mathbf{x}_2\leq_{\mathrm{dref}}\cdots\leq_{\mathrm{dref}}\mathbf{x}_{m-1}$ and $b\in[n]$. Consider the chain $\{\mathbf{x}_{i_1},\mathbf{x}_{i_2},\ldots,\mathbf{x}_{i_k}\}$, where i_j is the first occurrence of $\mathbf{x}_{i_j}$ in the given multichain. Via the construction of Proposition 2.23 we obtain a $k+1$-tuple $(L;R_1,R_2,\ldots,R_k)$ of subsets of $[n]$. From this we obtain an m-tuple $(\bar{S}_1;S'_2,S'_3,\ldots,S'_m)$ of subsets of $[n]$ by setting $\bar{S}_1=L$ and $S'_{i_j+1}=R_j$ and $S'_j=\emptyset$ for the remaining $j\in[m]$. Let D be the set constructed from this m-tuple as in the beginning of this proof. We have

$$\lvert D\rvert = \sum_{i=1}^{m}\lvert S_i\rvert = n-\lvert\bar{S}_1\rvert + \sum_{i=2}^{m}\lvert S'_i\rvert = n-\lvert L\rvert + \sum_{i=1}^{k}\lvert R_i\rvert = n-\lvert L\rvert+\lvert L\rvert-1 = n-1.$$

We have thus found a bijection between the $(n-1)$-element subsets of $[mn]$ and pairs (X,b) where X is a multichain of $(NC_n,\leq_{\mathrm{dref}})$ with $m-1$ elements and some $b\in[n]$. We obtain the equality

$$\binom{mn}{n-1} = n\cdot\mathscr{Z}_n(m),$$

which concludes the proof. $\square$

EXAMPLE 2.27

Let us compute multichain of length 3 in $(NC_6, \leq_{\mathrm{dref}})$ given by the set $D = \{4, 8, 10, 12, 23\}$. We obtain the partition

$$S_1 = \{4\}, \quad S_2 = \{8, 10, 12\}, \quad S_3 = \emptyset, \quad S_4 = \{23\}.$$

We further obtain

$$S'_1 = \{4\}, \quad S'_2 = \{2, 4, 6\}, \quad S'_3 = \emptyset, \quad S'_4 = \{5\},$$

and $\bar{S}_1 = \{1, 2, 3, 5, 6\}$. The parenthesization $\sigma_b(6)$ induced by $(\bar{S}_1; S'_2, S'_3, S'_4)$ is well-parenthesized for $b = 1$, and we obtain

$$\hat{\sigma}_1(6) = (1(2)(3\,4)(5)(6).$$

The construction from the proof of Proposition 2.23 yields the elements

$$\mathbf{x}_1 = 1 \mid 2 \mid 3\,4 \mid 5 \mid 6,$$
$$\mathbf{x}_2 = \mathbf{x}_3 = 1\,2\,3\,4\,6 \mid 5.$$

Since every element of NC_n itself constitutes a multichain consisting of one element, the cardinality of $(NC_n, \leq_{\mathrm{dref}})$ thus equals $\mathscr{Z}_n(2)$. We therefore obtain an explicit form for the Catalan numbers from Proposition 2.26.

COROLLARY 2.28

For $n \geq 1$ the cardinality of NC_n is given by

$$\mathrm{Cat}(n) = \frac{1}{n+1}\binom{2n}{n}.$$

Proof. We have seen in Proposition 2.2 that the cardinality of NC_n is given by $\mathrm{Cat}(n)$, and Proposition 2.26 implies that this number equals

$$\mathscr{Z}_n(2) = \frac{1}{n}\binom{2n}{n-1} = \frac{1}{n+1}\binom{2n}{n}.$$

$\square$

3. Applications

In this section we present a few situations, where noncrossing set partitions pop up sort of unexpectedly.

3.1. The Moment-Cumulant Formula. In probability theory moments and cumulants of a random variable X are two very basic concepts. While the n^{th} moment is simply the expectation of X^n, the cumulants are the coefficients of the logarithmic transform of the moment generating function of X. More precisely, if X is a random variable with probability distribution μ_X, then we define the n^{TH} MOMENT of X by

$$m_n(X) = E(X^n) = \int t^n d\mu_X(t).$$

The (EXPONENTIAL) MOMENT GENERATING FUNCTION of X is simply

$$M_X(t) = \sum_{n\geq 0} m_n(X)\frac{t^n}{n!} = E(\exp^{tX}).$$

Observe that $M_X(t)$ is essentially the Fourier transform of μ_X. The CUMULANTS of X are now the coefficients of the logarithmic transform of $M_X(t)$, i.e.

$$C_X(t) = \log M_X(t) = \sum_{n\geq 0} c_n(X)\frac{t^n}{n!}.$$

The relation between moments and cumulants is given by the Moment-Cumulant Formula.

THEOREM 3.1: [27]

For a random variable X and $n \geq 1$ we have

$$m_n(X) = \sum_{\mathbf{x}\in\Pi_n} \prod_{B\in\mathbf{x}} c_{|b|}(X). \tag{5}$$

Two observations are imminent. Firstly, moments and cumulants determine each other, i.e. if two random variables have the same moments, then they also have the same cumulants and vice versa. Secondly, we can easily compute the first moments in terms of the first cumulants. We obtain:

$$\begin{aligned} m_1(X) &= c_1(X),\\ m_2(X) &= c_1^2(X) + c_2(X),\\ m_3(X) &= c_1^3(X) + 3c_1(X)c_2(X) + c_3(X). \end{aligned}$$

Now we can of course recursively compute the cumulants in terms of the moments. One of the crucial insights of [27] is that this can be done much more convenient with the help of the Möbius Inversion Formula, see [26, Proposition 2], on the partition lattice so that we obtain

$$\begin{aligned} c_n(X) &= \sum_{\mathbf{x}\in\Pi_n} \mu_\Pi(\mathbf{x},\mathbf{1}) \prod_{B\in\mathbf{x}} m_{|B|}(X)\\ &= \sum_{\mathbf{x}\in\Pi_n} (-1)^{\mathrm{bl}(\mathbf{x})-1}(\mathrm{bl}(\mathbf{x})-1)! \prod_{B\in\mathbf{x}} m_{|B|}, \end{aligned} \tag{6}$$

where the last equality follows from Corollary 1.14. We can therefore immediately compute

$$c_3(X) = 2m_1^3(X) - 3m_1(X)m_2(X) + m_3(X)$$

without explicitly knowing $c_1(X)$ and $c_2(X)$.

Cumulants have another nice property: they are linear on independent random variables. Recall that if X and Y are independent random variables their joint expectation equals the product of the single expectiations, i.e. $E(XY) = E(X)E(Y)$, and this relation extends to the moment generating function: $M_{XY}(t) = M_X(t)M_Y(t)$. Since $C_{XY}(t) = \log M_{XY}(t)$, we conclude that $c_n(XY) = c_n(X) + c_n(Y)$ whenever X and Y are independent.

Now let us lift this setting to something more algebraic. It is well known that we can formally add and scale random variables, which motivates the following setup. Let $\mathcal{A}$ be a unital $\mathbb{R}$-vector space, i.e. a $\mathbb{R}$-vector space with a unit element 1. Fix a linear functional φ on $\mathcal{A}$ with $\varphi(1) = 1$, and call the pair $(\mathcal{A}, \varphi)$ a (FORMAL) PROBABILITY SPACE. The elements of $\mathcal{A}$ are then called (FORMAL) RANDOM VARIABLES. We then define moments and cumulants as before, by setting $m_n(x) = \varphi(x^n)$ for $x \in \mathcal{A}$, and by defining $c_n(x)$ via (5). If we want to model classical random variables, then the formula for the expectation of two independent random variables essentially requires $\mathcal{A}$ to be a commutative vector space.

This algebraic approach, however, also makes sense for $\mathcal{A}$ noncommutative, and is essentially the starting point for the theory of free probability due to D. Voiculescu [36]. In that case, we need a substitute for the concept of independent random variables, since the old definition relies crucially on the commutativity of $\mathcal{A}$. This is done as follows. Two (formal) random variables $x, y \in \mathcal{A}$ are FREELY INDEPENDENT (or simply FREE) if $\varphi(xy) = 0$ whenever $\varphi(x) = 0 = \varphi(y)$. It then requires some work to show that there is a family of quantities associated with $x \in \mathcal{A}$, the FREE CUMULANTS of x, which behave linearly on free random variables. In other words, for $x \in \mathcal{A}$ we can define its free cumulants $\kappa_n(x)$ such that the implication

$$\text{if } x \text{ and } y \text{ are free, then } \kappa_n(xy) = \kappa_n(x) + \kappa_n(y)$$

holds. The exact details of this construction are, however, beyond the scope of this manuscript. For our purposes, the most interesting connection is the following result due to R. Speicher.

THEOREM 3.2: [28]

For a (non-commutative) random variable x and $n \geq 1$ we have

$$m_n(x) = \sum_{\mathbf{x} \in NC_n} \prod_{B \in \mathbf{x}} \kappa_{|B|}(x). \tag{7}$$

In particular, we can once more compute the free cumulants from the moments by means of the Möbius Inversion Formula, which we apply this time on the lattice of noncrossing partitions instead on the lattice of all set partitions:

$$\begin{aligned} \kappa_n(x) &= \sum_{\mathbf{x} \in NC_n} \mu_{NC}(\mathbf{x}, \mathbf{1}) \prod_{B \in \mathbf{x}} m_{|B|}(x) \\ &= \sum_{\mathbf{x} \in NC_n} \prod_{B \in K(\mathbf{x})} (-1)^{|B|-1} \operatorname{Cat}(|B|-1) \prod_{B \in \mathbf{x}} m_{|B|}(x), \end{aligned} \tag{8}$$

where the last equality follows from Corollary 2.18.

In fact, for any number sequence m_n, we can define (free) cumulants c_n (resp. κ_n) via (6) (resp. (8)). A very surprising example of free cumulants appears in the representation theory of the symmetric group [5], where they can be used to express the asymptotic behavior of certain characters.

✎ EXERCISE 8

What is the number of labeled simple graphs on $[n]$, and what is the number of labeled trees on $[n]$? Compute the number $\bar{g}_c(4)$ of connected labeled simple graphs on $[4]$ via (6), and the number $\bar{f}(4)$ of labeled forests on $[4]$ via (5).

3.2. **Connected Components of Positroids.** Another instance of (free) cumulants can be encountered in matroid theory.

Let E be a finite set, and let $\mathcal{B} \subseteq \wp(E)$ be a non-empty family of subsets of E. The pair $\mathbb{M} = (E, \mathcal{B})$ is a MATROID on E if it satisfies the following BASIS EXCHANGE AXIOM: if $B_1, B_2 \in \mathcal{B}$ and $b_1 \in B_1 \setminus B_2$, then there exists $b_2 \in B_2 \setminus B_1$ such that $B_1 \setminus \{b_1\} \cup \{b_2\} \in \mathcal{B}$.

The elements of $\mathcal{B}$ are the BASES of $\mathbb{M}$, and it can be shown that they all have the same size. A subset $X \subseteq E$ is INDEPENDENT if there is some $B \in \mathcal{B}$ with $X \subseteq B$, otherwise it is DEPENDENT.

For two matroids $\mathbb{M} = (E, \mathcal{B})$ and $\mathbb{M}' = (E', \mathcal{B}')$ we define the DIRECT SUM of $\mathbb{M}$ and $\mathbb{M}'$ to be the matroid $\mathbb{M} \oplus \mathbb{M}' = (E \uplus E', \mathcal{B} \uplus \mathcal{B}')$, where $\mathcal{B} \uplus \mathcal{B}' = \{B \uplus B' \mid B \in \mathcal{B} \text{ and } B' \in \mathcal{B}'\}$. A matroid is CONNECTED if it cannot be written as a direct sum of two smaller matroids. Otherwise, we can write $\mathbb{M} = \bigoplus_{i=1}^{s} \mathbb{M}_i$ for connected matroids $\mathbb{M}_i$; the CONNECTED COMPONENTS of $\mathbb{M}$.

We have the following relation between matroids and set partitions, which is due to H. Whitney [37, Theorem 19].

PROPOSITION 3.3

Let $\mathbb{M} = (E, \mathcal{B})$ be a matroid, and define a relation $\sim$ on E by setting $a \sim b$ whenever there are two bases $B, B' \in \mathcal{B}$ such that $B' = B \setminus \{a\} \cup \{b\}$. This relation is an equivalence relation on E, and its equivalence classes are the connected components of $\mathbb{M}$.

In order to prove Proposition 3.3 we need some more notation. A dependent set $X \subseteq E$ for which every proper subset is independent is a CIRCUIT. Let us denote the set of circuits of $\mathbb{M}$ by $\mathcal{C}(\mathbb{M})$. It is well known that the set of circuits satisfy the following axioms [21, Section 1.1]:

(C1) $\emptyset \notin \mathcal{C}(\mathbb{M})$;

(C2) if $X, Y \in \mathcal{C}(\mathbb{M})$ and $X \subseteq Y$, then $X = Y$;

(C3) if $X_1, X_2 \in \mathcal{C}(\mathbb{M})$ and $e \in X_1 \cap X_2$, then there exists
$$f \in X_2 \setminus X_1 \text{ and } Y \in \mathcal{C}(\mathbb{M}) \text{ such that } f \in Y \subseteq (X_1 \cup X_2) \setminus \{e\}.$$

Moreover, every family of subsets of E satisfying (C1)–(C3) is the set of circuits of some matroid. We need the following simple observation.

LEMMA 3.4: [21, Proposition 1.1.6]

Let $X \subseteq E$ be independent, and let $e \in E \setminus X$ be such that $X \cup \{e\}$ is dependent. Then there exists a unique circuit $C \subseteq X \cup \{e\}$ containing e.

Proof. Let C be a circuit in $X \cup \{e\}$, which must exist since this set is dependent. If $e \notin C$, then $C \subseteq X$ contradicting the assumption that X was independent. It follows that any circuit in $X \cup \{e\}$ needs to contain e. Suppose there is another circuit $C' \subseteq X \cup \{e\}$. Then $C \cap C' \neq \emptyset$, and with (C3) we can find a circuit $D \subseteq (C \cup C') \setminus \{e\}$. But then, $D \subseteq X$, which is a contradiction. □

Define a map $\gamma : E \to \wp(E)$ by

$$\gamma(e) = \{e\} \cup \big\{ f \in E \mid \text{there is } C \in \mathcal{C}(\mathbb{M}) \text{ such that } e, f \in C \big\}.$$

LEMMA 3.5: [1, Proposition 7.2]

For $a, b \in E$ holds $a \sim b$ if and only if $a \in \gamma(b)$.

Proof. Assume first that $a \sim b$. By definition there are bases $B, B' \in \mathcal{B}$ such that $B' = B \setminus \{a\} \cup \{b\}$. Since $b \notin B$ Lemma 3.4 implies that there exists a unique circuit $C \subseteq B \cup \{b\}$ with $b \in B$. If $a \notin C$, then $C \subseteq B'$, which contradicts the assumption that B' is a basis. Hence $a \in C$ and thus $a \in \gamma(b)$.

Conversely let $a \in \gamma(b)$ and let $C \in \mathcal{C}(\mathbb{M})$ with $a, b \in C$. Consider the contraction $\mathbb{M}/C$, which is the matroid on $E \setminus C$ given by the bases $\big\{ B \setminus C \mid B \in \mathcal{B} \text{ such that } |B \cap C| \text{ is maximal}\big\}$. Let D be a basis of $\mathbb{M}/C$. Then by definition $D \cup C$ is a circuit of $\mathbb{M}$, and $B = D \cup C \setminus \{a\}$ and $B' = D \cup C \setminus \{b\}$ are both bases of $\mathbb{M}$ with $B = B' \setminus \{a\} \cup \{b\}$, which implies $a \sim b$. $\square$

Proof of Proposition 3.3. In view of Lemma 3.5 it suffices to show that γ is reflexive, symmetric and transitive. The first two properties are immediate from the definition. Suppose that $a, b, c \in E$ with $a \in \gamma(b)$ and $b \in \gamma(c)$. There exist circuits $C_1, C_2 \in \mathscr{C}(\mathbb{M})$ with $a \in C_1$ and $c \in C_2$ such that $C_1 \cap C_2 \neq \emptyset$. Moreover, choose C_1 and C_2 among all such circuits with the property that $\lvert C_1 \cup C_2 \rvert$ is minimal.

Assume that there is no circuit in $\mathbb{M}$ containing both a and c. Thus $C_1 \neq C_2$, and in particular $a \in C_1 \setminus C_2$. If we pick $d \in C_1 \cap C_2$, then (C3) yields a circuit $C_3 \subseteq (C_1 \cup C_2) \setminus \{d\}$ with $a \in C_3$. By assumption C_3 cannot contain c. If $C_3 \subseteq C_1$, then (C2) implies $C_3 = C_1$, which contradicts $d \notin C_3$. We can thus find an element $e \in C_2 \setminus C_1$ with $e \in C_3$. If we apply (C3) once more, we obtain a circuit $C_4 \subseteq (C_2 \cup C_3) \setminus \{e\}$ with $c \in C_4$. Again we see that $C_4 \not\subseteq C_2$ so that the intersection $C_4 \cap (C_3 \setminus C_2)$ is non-empty. It follows that $C_4 \cap C_1 \neq \emptyset$, and we recall that $a \in C_1$ and $c \in C_4$. We have, however, that $C_1 \cup C_4 \subseteq (C_1 \cup C_2) \setminus \{e\}$, and thus $\lvert C_1 \cup C_4 \rvert < \lvert C_1 \cup C_2 \rvert$, which contradicts the choice of C_1 and C_2. We have thus shown that a circuit containing a and c must exist, and therefore $a \in \gamma(c)$.

Moreover, we have seen that for any $a, b \in E$ with $a \not\sim b$ there does not exist a circuit containing both a and b. If $E_1, E_2, \ldots, E_s$ are the equivalence classes of $\sim$, we conclude that $\mathcal{C}(\mathbb{M})$ can be partitioned into sets $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_s$ such that for all $i \in [s]$ the set $\mathcal{C}_i$ is a family of subsets of E_i satisfying (C1)–(C3), and therefore defines a matroid $\mathbb{M}_i = (E_i, \mathcal{B}_i)$ such that $\mathcal{C}(\mathbb{M}_i) = \mathcal{C}_i$. It follows further that $\mathbb{M}_i$ is connected and $\mathbb{M} = \mathbb{M}_1 \oplus \mathbb{M}_2 \oplus \cdots \oplus \mathbb{M}_s$. $\square$

Recall for instance from Exercise 1 that equivalence relations correspond bijectively to set partitions. Therefore, any matroid has a canonically associated set partition. In other words, if $m(n)$ denotes the number of matroids on $[n]$ and $m_c(n)$ denotes the number of connected matroids on $[n]$, then the numbers $m(n)$ and $m_c(n)$ satisfy (5). In fact, as described in [1] we can find a meaningful subclass of matroids such that we recover (7).

A major source of matroids arise in the following way. Fix a field $\mathbb{K}$, and let $A \in \mathrm{Mat}_{\mathbb{K}}(d, n)$ be a $d \times n$ matrix of rank d over $\mathbb{K}$. If we denote the columns of A by $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$, then the subsets $B \subseteq [n]$ for which $\{\mathbf{a}_i \mid i \in B\}$ is a linear basis of $\mathbb{K}^d$ are the bases of the REPRESENTABLE matroid $\mathbb{M}(A)$.

Recall that a MINOR of A is the determinant of some $k \times k$ submatrix of A. From now on, let $\mathbb{K} = \mathbb{R}$, and choose A in such a way that all its maximal minors are nonnegative. We call such matrices TOTALLY NONNEGATIVE. The matroid associated with a totally nonnegative matrix is a

POSITROID. This class of matroids was introduced by A. Postnikov in [23] and further studied for instance in [20]. A key observation is that the property of being a positroid depends on the order of the columns of A (or equivalently on the order of the ground set).

EXAMPLE 3.6

Fix $a_1 \leq a_2 \leq \cdots \leq a_n$, and define $\mathbf{a}_j = (1, a_j, a_j^2, \ldots, a_j^{d-1})^{\perp}$ for $j \in [n]$. The matrix $A = (\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n)$ is totally nonnegative, since every maximal minor is a Vandermonde determinant. More precisely, if $I = \{i_1, i_2, \ldots, i_d\}$, then the maximal minor $\Delta_I(A)$ induced by the columns in I is given by

$$\Delta_I(A) = \prod_{1 \leq s < t \leq d} (a_{i_t} - a_{i_s}),$$

which is nonnegative by assumption.

Here are a few basic observations on positroids. For $k, l \in [n]$ define the CYCLIC INTERVAL $[k, l]$ by

$$[k, l] = \begin{cases} \{k, k+1, \ldots, l\}, & \text{if } k \leq l, \\ \{k, k+1, \ldots, n, 1, \ldots, l\}, & \text{if } k > l. \end{cases}$$

LEMMA 3.7: [1, Lemma 3.3]

Let $\mathbb{M}$ be a positroid on $[n]$. For any $a \in [n]$ we have that $\mathbb{M}$ is also a positroid on the cyclic interval $[a, a-1]$.

Proof. Let A be the underlying $d \times n$ matrix of $\mathbb{M}$, and suppose its columns are $\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_n$ in that order. Consider the matrix A' given by the column vectors $\mathbf{a}_2, \ldots, \mathbf{a}_n, (-1)^{d-1}\mathbf{a}_1$. Let $I = \{i_1, i_2, \ldots, i_d\}$ be a d-element subset of $[n]$, and denote by

$$\Delta_I(A) = \det(\mathbf{a}_{i_1}, \mathbf{a}_{i_2}, \ldots, \mathbf{a}_{i_d})$$

the $d \times d$ minor of A induced by I, and define $I' = \{i_1 - 1, i_2 - 1, \ldots, i_d - 1\}$. If $i_1 = 1$, then $i_1 - 1 = n$. We then see that $\Delta_I(A) = \Delta_{I'}(A')$ whenever $1 \notin I$. Otherwise we can assume that $i_1 = 1$, and Laplace's Determinant Formula lets us compute $\Delta_I(A)$ by expanding along the first column:

$$\Delta_I(A) = \sum_{j=1}^{d} (-1)^{1+j} a_{1,j} \det(\mathbf{a}_{i_2}, \ldots, \mathbf{a}_{i_d}),$$

where $a_{1,j}$ is the j^{th} entry of $\mathbf{a}_1$. Analogously we obtain $\Delta_{I'}(A')$ by expanding along the last column:

$$\Delta_{I'}(A') = \sum_{j=1}^{d} (-1)^{d+j} a'_{1,j} \det(\mathbf{a}_{i_2}, \ldots, \mathbf{a}_{i_d}),$$

where $a'_{1,j}$ is the j^{th} entry of the last column of A', which by construction equals $(-1)^{d-1}a_{1,j}$. Since $(-1)^{d+j}a'_{1,j} = (-1)^{d+j+d-1}a_{1,j} = (-1)^{j+1}a_{1,j}$, we obtain $\Delta_I(A) = \Delta_{I'}(A')$. Consequently, $\mathbb{M}(A')$ is a positroid, which coincides with $\mathbb{M}$ after cyclically shifting the ground set. Thus $\mathbb{M}$ is a positroid on the cyclic interval $[2, 1]$, and the claim follows by iterating this argument. $\square$

PROPOSITION 3.8: [1, Proposition 3.4]

Let $k, l \in [n]$ and suppose that $\mathbb{M}_1$ is a positroid on the cyclic interval $[k+1, l]$ and M_2 is a positroid on the cyclic interval $[l+1, k]$. Then $\mathbb{M}_1 \oplus \mathbb{M}_2$ is a positroid on $[n]$.

Proof. In view of Lemma 3.7 it suffices to consider the case $l = n$ and $k < l$. Let $\mathbb{M}_1$ be a positroid on $[k]$ given by the totally nonnegative matrix A_1, and let $\mathbb{M}_2$ be positroid on $[k+1, n]$ given by the totally nonnegative matrix A_2. The block diagonal matrix $A = \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix}$ is by construction totally nonnegative, since every maximal minor of A is a block diagonal matrix, and the determinant of a block diagonal matrix is the product of the determinants of the blocks. Moreover, A represents the matroid $\mathbb{M}_1 \oplus \mathbb{M}_2$, which is thus a positroid. $\square$

We also have the converse of Proposition 3.8. The proof of this statement, however, requires a few too many new notions, so we omit here.

PROPOSITION 3.9: [1, Proposition 7.4]

Let $\mathbb{M} = \mathbb{M}_1 \oplus \mathbb{M}_2$ be a positroid on $[n]$, where $\mathbb{M}_1$ and $\mathbb{M}_2$ are connected positroids. Then the ground sets of $\mathbb{M}_1$ and $\mathbb{M}_2$ are cyclic intervals of $[n]$.

The main result of this section is the following.

THEOREM 3.10: [1, Theorem 7.6]

Let $\mathbb{M}$ be a positroid on $[n]$ and let $E_1, E_2, \ldots, E_s$ be the ground sets of the connected components of $\mathbb{M}$. Then $\mathbf{x}_{\mathbb{M}} = \{E_1, E_2, \ldots, E_s\}$ is a noncrossing set partition of $[n]$.
Conversely, let $\mathbf{x} = \{E_1, E_2, \ldots, E_s\} \in NC_n$, and let $\mathbb{M}_i$ be a connected positroid on E_i for every $i \in [s]$. Then $\mathbb{M} = \bigoplus_{i=1}^{s} \mathbb{M}_i$ is a positroid.

Proof. Let $\mathbb{M} = \bigoplus_{i=1}^{s} \mathbb{M}_i$ be a positroid. Proposition 3.3 implies that $\mathbf{x}_{\mathbb{M}} = \{E_1, E_2, \ldots, E_s\} \in \Pi_n$. If $\mathbf{x}_{\mathbb{M}}$ is crossing, then there are two parts E_a and E_b and $i < j < k < l$ such that $i, k \in E_a$ and $j, l \in E_b$. It follows that neither E_a nor E_b is a connected interval of $[n]$. If we restrict $\mathbb{M}$ to $E_a \uplus E_b$ we obtain a positroid $\mathbb{M}' = \mathbb{M}_a \oplus \mathbb{M}_b$, see [1, Proposition 3.5], which contradicts Proposition 3.9.

Conversely, let $\mathbf{x} = \{E_1, E_2, \ldots, E_s\} \in NC_n$. We proceed by induction on s, where the induction base $s = 1$ holds trivially. The assumption that $\mathbf{x}$ is noncrossing ensures that there is some block which is a cyclic interval of $[n]$. Without loss of generality we can put $E_s = [k, l]$. Then $\mathbf{x} \setminus \{E_s\} \in NC_{[l+1,k-1]}$, and the induction hypothesis ensures that $\mathbb{M}' = \bigoplus_{i=1}^{s-1} M_i$ is a positroid on the cyclic interval $[l+1, k-1]$. Moreover, we have assumed that $\mathbb{M}_s$ is a connected positroid on the cyclic interval $[k, l]$. Proposition 3.8 now implies that $\mathbb{M}' \oplus \mathbb{M}_s$ is a positroid. $\square$

In particular, if $p(n)$ denotes the number of positroids on $[n]$, and $p_c(n)$ denotes the number of connected positroids on $[n]$, then the numbers $p(n)$ and $p_c(n)$ satisfy (7).

✎ EXERCISE 9

According to [1, Theorem 10.4] we have $p(n) = \sum_{k=0}^{n} \frac{n!}{k!}$. Use this and (8) to compute $p_c(4)$.

3.3. **Exceptional Sequences in the Category of Representations of a Path.** In this section we outline how the lattice of noncrossing set partitions arises as a poset on certain families of representations of a path. Our exposition follows [35] and [25, Section 4].

Let Q be a directed graph on vertex set $[n]$ and fix a field $\mathbb{K}$. A Q-REPRESENTATION V is an assignment of a finite-dimensional $\mathbb{K}$-vector space V_i to each vertex i together with a linear map $V_\alpha : V_i \to V_j$ for every oriented edge $\alpha : i \to j$ in Q. For two Q-representations V and W we define a MORPHISM from V to W to be a collection of linear maps $f_i : V_i \to W_i$ for all $i \in [n]$ such that for all edges $\alpha : i \to j$ we have $W_\alpha \circ f_i = f_i \circ V_\alpha$. Let $\mathrm{Hom}(V, W)$ denote the set of all morphisms from V to W. In fact, $\mathrm{Hom}(V, W)$ is itself a $\mathbb{K}$-vector space. These definitions give rise to the CATEGORY OF Q-REPRESENTATIONS denoted by $\mathrm{rep}\, Q$.

For two Q-representations V and W their DIRECT SUM $V \oplus W$ is defined via $(V \oplus W)_i = V_i \oplus W_i$ and $(V \oplus W)_\alpha = V_\alpha \oplus W_\alpha$. We call V INDECOMPOSABLE if it is not isomorphic to the direct sum of two non-zero representations of Q. Let $\mathrm{ind}\,\mathrm{rep}\, Q$ denote the set indecomposable Q-representations. Two Q-representations V and W are ORTHOGONAL if $\dim \mathrm{Hom}(V, W) = 0$. An EXCEPTIONAL SEQUENCE is a family of pairwise orthogonal, indecomposable Q-representation, and we denote the set of exceptional sequences of $\mathrm{rep}\, Q$ by $\mathrm{Exc}(\mathrm{rep}\, Q)$.

Let V, W be Q-representations. We say that V is a SUBREPRESENTATION of W if V_i is a subspace of W_i for each $i \in [n]$ and for $\alpha : i \to j$ the map V_α is induced from the inclusions of V_i and V_j into W_i and W_j, respectively. The inclusion maps form an injective morphism from V to W.

If V is a subrepresentation of W, then we define the QUOTIENT REPRESENTATION W/V by $(W/V)_i = W_i/V_i$ and $(W/V)_\alpha = W_\alpha/V_\alpha$. The quotient maps form a surjective morphism from W to W/V. Consequently $0 \to V \to W \to W/V \to 0$ is a short exact sequence.

For three Q-representations U, V, W we say that W is an EXTENSION of U by V if there is a subrepresentation of W which is isomorphic to V such that the corresponding quotient representation W/V is isomorphic to U. Let $\mathrm{Ext}(U, V)$ denote the set of extensions of U by V (up to equivalence). An extension is TRIVIAL if there is a morphism from W to V which is the identity on V.

LEMMA 3.11: [35, Lemma 3.1]

If W is a trivial extension of U by V, then $W \cong U \oplus V$. In other words, the short exact sequence $0 \to V \to W \to U \to 0$ is split.

Proof. Let $s \in \mathrm{Hom}(W, V)$ which is the identity on V, and let g be the quotient map from W to U. Then $s \oplus g \in \mathrm{Hom}(W, V \oplus U)$, and it acts as an isomorphism on every vertex. It therefore is an isomorphism of Q-representations. □

If X is a set of Q-representations, let $\mathrm{Ext}(X)$ be the smallest set of Q-representations that contains X and is closed under extensions. Define a partial order on $\wp(\mathrm{rep}\, Q)$ by $X \sqsubseteq Y$ if and only if $\mathrm{Ext}(X) \subseteq \mathrm{Ext}(Y)$.

For the remainder of this section, let Q be the directed path with n vertices, i.e. the directed graph on $[n]$ with directed edges (i, j) whenever $j = i + 1$. We then write $\mathbb{A}_n$ instead of Q. Our goal is to prove the following result, which is a special case of [12, Theorem 1.1].

THEOREM 3.12

For $n \geq 1$ there is an explicit bijection $\tau : NC_n \to \mathrm{Exc}\big(\mathrm{rep}\, \mathbb{A}_{n-1}\big)$. Moreover, for $\mathbf{x}, \mathbf{y} \in NC_n$ we have $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$ if and only of $\tau(\mathbf{x}) \sqsubseteq \tau(\mathbf{y})$.

Since the exceptional sequences of $\mathbb{A}_n$ consist of indecomposable representations, we need to understand what these look like. For $i, j \in [n]$ with $i \leq j$ define E^{ij} to be the $\mathbb{A}_n$-representation which assigns a one-dimensional vector space to each $p \in [n]$ with $i \leq p \leq j$, and where we put identity maps between successive one-dimensional vector spaces, and zero maps elsewhere. For simplicity, we denote the identity maps by "1", and the zero maps by "0".

PROPOSITION 3.13: [35, Proposition 5.1]

An $\mathbb{A}_n$-representation is indecomposable if and only if it is isomorphic to E^{ij} for some $i, j \in [n]$ with $i \leq j$.

Proof. We first show that E^{ij} is indecomposable. Indeed assume that $E^{ij} = U \oplus V$ for some non-zero $\mathbb{A}_n$-representations U, V. Let $p \in [n]$ with $i \leq p \leq j$. By construction $(E^{ij})_p$ is one-dimensional, which implies that either U_p or V_p is zero. Moreover, since U and V are non-zero we can choose $p < j$ in such a way that either U_p and V_{p+1} are zero, or V_p and U_{p+1} are zero. If $\alpha : p \to p+1$, then in both cases U_α and V_α are zero, which implies that $(U \oplus V)_\alpha$ is zero. By construction, however, we have that $(E^{ij})_\alpha$ is non-zero, which is a contradiction.

Conversely, let V be an indecomposable $\mathbb{A}_n$-representation, and write f_p instead of $V_{p\to p+1}$ for $p \in [n-1]$. Let i be minimal such that $V_i \neq 0$, and pick $t \in V_i$. Let j be maximal such that $f_{j-1} \cdots f_{i+1} f_i(t) \neq 0$. If T denotes the subrepresentation of V generated by t, then we see that T_p is one-dimensional for $i \leq p \leq j$ and zero otherwise. In particular, T is isomorphic to E^{ij}.

Let $\iota_p : T_p \to V_p$ denote the inclusion map, and define a map $s_j : V_j \to T_j$ such that $s_j \circ \iota_j$ is the identity. For $p \in [n]$ with $i \leq p < j$ define s_p inductively such that $s_p \circ \iota_p$ is the identity and $f_p \circ s_p = s_{p+1} \circ f_p$. Moreover, if $i > 1$, then $V_{i-1} = 0$ by the minimality of i. We conclude $f_{i-1} = 0$ which implies $f_{i-1} \circ s_{i-1} = s_i \circ f_{i-1}$. If $j < n$, then the maximality of j implies that $f_j f_{j-1} \cdots f_i(t) = 0$. Hence the restriction of f_j to T_j is zero, which implies $f_j \circ s_j = s_{j+1} \circ f_j$. Consequently $s \in \mathrm{Hom}(V, T)$, and $V \in \mathrm{Ext}(V/T, T)$. Lemma 3.11 implies that $V \cong T \oplus V/T$. Since V is indecomposable and T is non-zero, we conclude $V \cong T \cong E^{ij}$. □

We thus have
$$\operatorname{ind}\operatorname{rep}\mathbb{A}_n = \{E^{ij} \mid 1 \leq i \leq j \leq n\}.$$
A consequence of Proposition 3.13 is that any exceptional sequence of $\operatorname{rep}\mathbb{A}_{n-1}$ consists of E^{ij}'s. It remains to determine when two indecomposables of $\operatorname{rep}\mathbb{A}_{n-1}$ are orthogonal.

PROPOSITION 3.14: [35, Proposition 6.1]

For $i, j, k, l \in [n]$ the space $\mathrm{Hom}(E^{ij}, E^{kl})$ is either zero- or one-dimensional. It is one-dimensional if and only if $k \leq i \leq l \leq j$.

Proof. Let $V_p = (E^{ij})_p$ and $W_p = (E^{kl})_p$ for $p \in [n]$, and let $f_p : V_p \to V_{p+1}$ and $g_p : W_p \to W_{p+1}$ for $p \in [n-1]$. Assume that there is a morphism $s \in \mathrm{Hom}(E^{ij}, E^{kl})$, and let $s_p : V_p \to W_p$ denote the restrictions.

By construction, V_i generates the whole $\mathbb{A}_n$-representation E^{ij}, so that s is determined by s_i. If $i < k$ or $i > l$, then $\dim W_i = 0$, so that s_i must be zero, which forces s to be zero. We thus have $\dim \mathrm{Hom}(E^{ij}, E^{kl}) = 0$ in that case.

Therefore let $k \leq i \leq l$. We can then choose s_i to be one-dimensional, since it is a map between one-dimensional vector spaces.

If $j < l$ we have $f_j = 0$ and $g_j = 1$, and $\dim V_{j+1} = 0$ and $\dim W_{j+1} = 1$, and hence $s_{j+1} = 0$. Consequently $f_j \circ s_{j+1} = 0$, and since s is a morphism we need to have $s_j \circ g_j = 0$, which can only happen if $s_j = 0$. If we repeat this process, we see that s needs to be zero, which implies $\dim \mathrm{Hom}(E^{ij}, E^{kl}) = 0$ in that case as well.

Finally, if $l \leq j$, we can construct s in such a way that it is non-zero, which yields $\dim \mathrm{Hom}(E^{ij}, E^{kl}) = 1$. □

We obtain the first part of Theorem 3.12 as a corollary. Let $\mathbf{x} \in \Pi_n$, and define its ARC SET by

$$\mathrm{Arc}(\mathbf{x}) = \bigl\{(i,j) \mid i < j \text{ and } i \sim_{\mathbf{x}} j \text{ and for all } i < k < j \text{ we have } i \not\sim_{\mathbf{x}} k\bigr\}.$$

LEMMA 3.15

For $\mathbf{x} \in \Pi_n$ we have $\bigl|\mathrm{Arc}(\mathbf{x})\bigr| = \mathrm{rk}(\mathbf{x})$.

Proof. Let $\mathbf{x} = \{B_1, B_2, \ldots, B_k\}$ with $|B_i| = b_i$ for $i \in [k]$. By definition, the block B_i contributes $b_i - 1$ arcs to $\mathrm{Arc}(\mathbf{x})$, and we obtain

$$\bigl|\mathrm{Arc}(\mathbf{x})\bigr| = \sum_{i=1}^{k} (b_i - 1) = n - k = \mathrm{rk}(\mathbf{x}).$$

□

Let us consider the map $\tau : \Pi_n \to \wp\bigl(\mathrm{ind}\,\mathrm{rep}\,\mathbb{A}_{n-1}\bigr)$ that sends $(i,j) \in \mathrm{Arc}(\mathbf{x})$ to E^{ij-1}.

COROLLARY 3.16

Let $n \geq 1$ and $\mathbf{x} \in \Pi_n$. Then $\tau(\mathbf{x}) \in \mathrm{Exc}\bigl(\mathrm{rep}\,\mathbb{A}_{n-1}\bigr)$ if and only if $\mathbf{x} \in NC_n$.

Proof. By construction $\tau(\mathbf{x})$ is a set of E^{ij}'s. Proposition 3.14 implies that $\dim \mathrm{Hom}(E^{ij}, E^{kl}) = 1$ if and only if $k \leq i \leq l \leq j$ which is by definition the case if and only if $\mathbf{x}$ is crossing. □

EXAMPLE 3.17

Let $n = 4$. The directed graph $\mathbb{A}_3$ is

$$\underset{1}{\circ} \longrightarrow \underset{2}{\circ} \longrightarrow \underset{3}{\circ}$$

and in view of Proposition 3.13 its indecomposable representations are

$$E^{11}, \quad E^{12}, \quad E^{13}, \quad E^{22}, \quad E^{23}, \quad E^{33}.$$

Proposition 3.14 tells us that

$$\begin{aligned}
\dim \operatorname{Hom}(E^{12}, E^{11}) &= \dim \operatorname{Hom}(E^{13}, E^{11}) = 1,\\
\dim \operatorname{Hom}(E^{12}, E^{22}) &= \dim \operatorname{Hom}(E^{22}, E^{23}) = 1,\\
\dim \operatorname{Hom}(E^{33}, E^{13}) &= \dim \operatorname{Hom}(E^{33}, E^{23}) = 1,\\
\dim \operatorname{Hom}(E^{23}, E^{12}) &= \dim \operatorname{Hom}(E^{13}, E^{12}) = \dim \operatorname{Hom}(E^{23}, E^{13}) = 1.
\end{aligned}$$

It follows that any exceptional sequence of $\mathbb{A}_3$-representations can have at most three elements. Clearly, the empty set and any one-element set of indecomposable representations is an exceptional sequence. We have excluded nine of the 15 two-element sets of indecomposable representations above, and we can check that there is a unique exceptional sequence of size three, namely $\{E^{11}, E^{22}, E^{33}\}$, which yields $\mathrm{Cat}(4) = 14$ exceptional sequences.
More precisely, the map τ is given by:

$$\begin{array}{lll}
\tau(1 \mid 2 \mid 3 \mid 4) = \emptyset, & \tau(1\,2 \mid 3 \mid 4) = \{E^{11}\}, & \tau(1\,3 \mid 2 \mid 4) = \{E^{12}\}\\
\tau(1\,4 \mid 2 \mid 3) = \{E^{13}\}, & \tau(1 \mid 2\,3 \mid 4) = \{E^{22}\}, & \tau(1 \mid 2\,4 \mid 3) = \{E^{23}\}\\
\tau(1 \mid 2 \mid 3\,4) = \{E^{33}\}, & \tau(1\,2\,3 \mid 4) = \{E^{11}, E^{22}\}, & \tau(1\,2\,4 \mid 3) = \{E^{11}, E^{23}\}\\
\tau(1\,3\,4 \mid 2) = \{E^{12}, E^{33}\}, & \tau(1 \mid 2\,3\,4) = \{E^{22}, E^{33}\}, & \tau(1\,2 \mid 3\,4) = \{E^{11}, E^{33}\}\\
\tau(1\,4 \mid 2\,3) = \{E^{13}, E^{22}\}, & \tau(1\,2\,3\,4) = \{E^{11}, E^{22}, E^{33}\}. &
\end{array}$$

In order to prove the second part of Theorem 3.12, we need to understand the potential extensions of the E^{ij}'s.

PROPOSITION 3.18: [35, Proposition 6.2]

There exists a non-trivial extension of E^{ij} by E^{kl} if and only if $i+1 \leq k \leq j+1 \leq l$. In that case, any such non-trivial extension is isomorphic to $E^{il} \oplus E^{kj}$, where E^{kj} is zero when $k = j+1$.

Proof. Let $W \in \operatorname{Ext}(E^{ij}, E^{kl})$, which by definition means that we have a short exact sequence

$$0 \longrightarrow E^{kl} \xrightarrow{f} W \xrightarrow{g} E^{ij} \longrightarrow 0,$$

where f is injective and g is surjective, and $\operatorname{Im}(f) = \ker(g)$. We have the following picture.

$$\begin{tikzcd}
{} \arrow[r] & (E^{kl})_{i-1} \arrow[r, "v_{i-1}"] \arrow[d, "f_{i-1}"] & (E^{kl})_i \arrow[r, "v_i"] \arrow[d, "f_i"] & \cdots \arrow[r, "v_{j-1}"] & (E^{kl})_j \arrow[r, "v_j"] \arrow[d, "f_j"] & (E^{kl})_{j+1} \arrow[r] \arrow[d, "f_{j+1}"] & {}\\
{} \arrow[r] & W_{i-1} \arrow[r, "w_{i-1}"] \arrow[d, "g_{i-1}"] & W_i \arrow[r, "w_i"] \arrow[d, "g_i"] & \cdots \arrow[r, "w_{j-1}"] & W_j \arrow[r, "w_j"] \arrow[d, "g_j"] & W_{j+1} \arrow[r] \arrow[d, "g_{j+1}"] & {}\\
{} \arrow[r] & 0 \arrow[r, "u_{i-1}"] & (E^{ij})_i \arrow[r, "u_i"] & \cdots \arrow[r, "u_{j-1}"] & (E^{ij})_j \arrow[r, "u_j"] & 0 \arrow[r] & {}
\end{tikzcd}$$

Pick $t \in W_i$ such that $g(t) \neq 0$, and let T be the subrepresentation of W generated by t, and let $\iota : T \to W$ be the inclusion map. It follows that the image of t in $(E^{ij})_p$ is non-zero for $i \leq p \leq j$, which implies that $\dim T_p = 1$ for these p. If $\dim T_{j+1} = 0$, then $T \cong E^{ij}$, and we conclude $\iota \circ g = 1$. It follows that W is trivial, and Lemma 3.11 implies $W = E^{ij} \oplus E^{kl}$.

Thus W can be non-trivial only if $\dim T_{j+1} = 1$, which requires $k \leq j+1 \leq l$. Assume in addition that $k \leq i$. Since $\dim W_{j+1} \neq 0$, we can find a non-zero $r \in W_{j+1}$ with $r = w_j \cdots w_i(t)$. Since $\dim(E^{ij})_{j+1} = 0$, we conclude that $r \in \ker(g_{j+1}) = \operatorname{Im}(f_{j+1})$. Since $k \leq i$ we can find $x \in (E^{kl})_i$ such that $f_{j+1} \circ (v_j \cdots v_i)(x) = r$. Since $f_i(x) \in \operatorname{Im}(f_i) = \ker(g_i)$, and $t \notin \ker(g_i)$, we have that $t' = t - f_i(x) \neq 0$. Since the arrows in the above diagram commute, we obtain

$$w_j \cdots w_i(t') = w_j \cdots w_i(t) - w_j \cdots w_i\bigl(f_i(x)\bigr) = r - f_{j+1} \circ (v_j \cdots v_i)(x) = r - r = 0.$$

Consequently, the subrepresentation of W generated by t' is isomorphic to E^{ij}, and we conclude as in the first part of the proof that W is trivial.

Finally, let $i+1 \leq k \leq j+1 \leq l$. If we pick some $t \in W_i$ such that $w_j \cdots w_i(t) \neq 0$, then the representation generated by t is isomorphic to E^{il}. Analogously to the proof of Proposition 3.13 we conclude that $W = E^{il} \oplus Z$ for some representation Z. It is then straightforward to verify that
$Z \cong \begin{cases} E^{kj}, & \text{if } k < j+1,\\ 0, & \text{if } k = j+1.\end{cases}$ □

EXAMPLE 3.19

Let us continue Example 3.17. We conclude from Proposition 3.18 that non-trivial extensions exist only in the following cases:

$$\begin{aligned} &E^{12} \in \operatorname{Ext}(E^{11}, E^{22}), \quad E^{13} \in \operatorname{Ext}(E^{11}, E^{23}), \quad E^{13} \oplus E^{22} \in \operatorname{Ext}(E^{12}, E^{23}),\\ &E^{13} \in \operatorname{Ext}(E^{12}, E^{33}), \quad E^{23} \in \operatorname{Ext}(E^{22}, E^{33}). \end{aligned}$$

Let us now conclude the proof of Theorem 3.12.

Proof of Theorem 3.12. We have seen in Corollary 3.16 that $\tau : NC_n \to \operatorname{Exc}(\operatorname{rep} \mathbb{A}_{n-1})$ is a bijection. It thus remains to show that τ sends $\leq_{\mathrm{dref}}$ to $\sqsubseteq$.

Let $\mathbf{x}, \mathbf{y} \in NC_n$ with $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$. If $\operatorname{Arc}(\mathbf{x}) \subseteq \operatorname{Arc}(\mathbf{y})$, then $\tau(\mathbf{x}) \subseteq \tau(\mathbf{y})$, and consequently $\operatorname{Ext}\bigl(\tau(\mathbf{x})\bigr) \subseteq \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$. Otherwise there is an arc $(i,l) \in \operatorname{Arc}(\mathbf{x})$ which is broken into a sequence of arcs $(j_0, j_1), (j_1, j_2), \ldots, (j_{s-1}, j_s) \in \operatorname{Arc}(\mathbf{y})$ for $s \geq 2$, where $j_0 = i$ and $j_s = l$. For any $i \in [s-1]$ we thus have $E^{j_i j_{i+1}-1}, E^{j_{i+1} j_{i+2}-1} \in \tau(\mathbf{y})$. Since $j_0 < j_1 < \cdots < j_s$ we have $j_i + 1 \leq j_{i+1} \leq j_{i+1} \leq j_{i+2} - 1$, and Proposition 3.18 implies $E^{j_i j_{i+2}-1} \in \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$. Repeated application yields $E^{il-1} \in \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$, and we conclude $\tau(\mathbf{x}) \subseteq \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$, which implies $\operatorname{Ext}\bigl(\tau(\mathbf{x})\bigr) \subseteq \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$.

Conversely suppose that $\operatorname{Ext}\bigl(\tau(\mathbf{x})\bigr) \subseteq \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$, and assume that $(i,l) \in \operatorname{Arc}(\mathbf{x})$. We thus have $E^{il-1} \in \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$. If $E^{il-1} \in \tau(\mathbf{y})$, then $(i,l) \in \operatorname{Arc}(\mathbf{x})$. Otherwise E^{il-1} arises as an extension of two $\mathbb{A}_{n-1}$-representations, in which case Proposition 3.18 forces $E^{ij_1-1}, E^{j_1 l} \in \operatorname{Ext}\bigl(\tau(\mathbf{y})\bigr)$. (Since E^{il-1} is indecomposable, it needs to arise as an extension of two indecomposable representations.) We repeat this process until we find a sequence $E^{ij_1-1}, E^{j_1 j_2-1}, \ldots, E^{j_s l-1} \in \tau(\mathbf{y})$, which by construction yields $(i, j_1), (j_1, j_2), \ldots, (j_s, l) \in \operatorname{Arc}(\mathbf{y})$. We conclude that $i \sim_{\mathbf{y}} l$, and thus $\mathbf{x} \leq_{\mathrm{dref}} \mathbf{y}$. □

Figure 7 shows $\Bigl(\operatorname{Exc}\bigl(\operatorname{rep} \mathbb{A}_{n-1}\bigr), \sqsubseteq\Bigr)$.

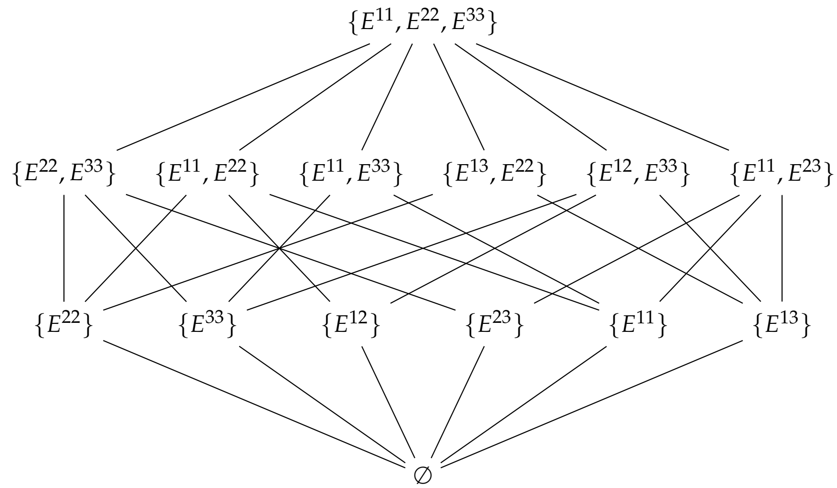

FIGURE 7. The lattice $\Big(\mathrm{Exc}\big(\mathrm{rep}\,\mathbb{A}_{n-1}\big), \sqsubseteq\Big)$.

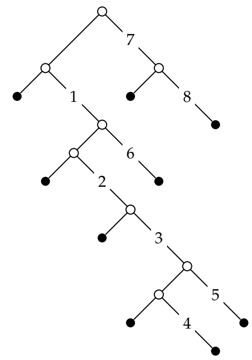

FIGURE 8. A full binary tree with 17 nodes. The right-edges are labeled according to when they are first encountered in depth-first search.

4. BIJECTIVE COMBINATORICS OF NC_n

In this section we introduce some other CATALAN OBJECTS, i.e. families of combinatorial objects that are counted by the Catalan numbers. To date more than 200 such Catalan objects have been found. A comprehensive exposition on the history of the Catalan numbers, and an extensive list of Catalan objects and their interactions is [33]. We also refer to [29] for an early draft of this book and [22] for a historical account on Catalan numbers.

It is the purpose of this section to describe bijections between noncrossing set partitions and five popular families of Catalan objects. For most of the objects under consideration the fact that they are enumerated by the Catalan numbers can be easily established using the recurrence relation from Proposition 2.2. We therefore present a bijection between these objects and noncrossing set partitions. Each of these bijections has appeared in the literature before, and we give the appropriate references in the appropriate place.

4.1. **Binary Trees.**

The first family of Catalan objects that we investigate are full binary trees.

DEFINITION 4.1

Let $n \geq 0$. A FULL BINARY TREE on $2n+1$ nodes is a tree in which every internal node has two children.

Denote by $\mathcal{T}_n$ the set of all full binary trees with $2n+1$ nodes. Observe that any $\tau \in \mathcal{T}_n$ has $2n$ edges, n internal nodes, and $n+1$ leaves. Moreover, since every inner node has exactly two children, one of them is the LEFT CHILD, and the other the RIGHT CHILD. In a natural way, these children are connected by a LEFT and a RIGHT edge, and induce a LEFT and a RIGHT SUBTREE. Figure 8 shows a full binary tree with 17 nodes.

PROPOSITION 4.2

For $n \geq 0$ we have $\lvert\mathcal{T}_n\rvert = \mathrm{Cat}(n)$.

We prove Proposition 4.2 bijectively by exhibiting an explicit bijection from $\mathcal{T}_n$ to NC_n.

THEOREM 4.3

For $n \geq 0$ there is an explicit bijection from $\mathcal{T}_n$ to NC_n.

Proof. Let $\tau \in \mathcal{T}_n$. Let us walk around τ depth-first, i.e. at every node we first visit the left subtree, and then the right subtree. We label the right-edges of τ in the order we encounter them. Every inner node of τ that does not have a left parent starts a run of right-edges, and the labels along these runs in τ are blocks of a partition $\mathbf{x}_\tau \in \Pi_n$. If there are integers $i < j < k < l$ such that i and k belong to the same block of $\mathbf{x}_\tau$, then the run containing j must lie strictly to the left of the run containing i (and k), while the run containing l lies strictly to the right. It follows that j and l belong to different blocks of $\mathbf{x}_\tau$, and thus $\mathbf{x}_\tau \in NC_n$. The map $\tau \mapsto \mathbf{x}_\tau$ is clearly injective.

Conversely let $\mathbf{x} \in NC_n$. We order the blocks of $\mathbf{x}$ increasingly. Let $B \in \mathbf{x}$ have $B = \{i_1, i_2, \ldots, i_k\}$ for some $k > 0$. Construct a sequence of k edges labeled by $i_1, i_2, \ldots, i_k$, and add a left-child to each but the last node. This is certainly a full binary tree. For every $j \in [k]$ with $i_{j+1} - i_j > 1$ there must be a block of $\mathbf{x}$ containing $i_j + 1$. Attach the tree corresponding this block to the left-child of the vertex whose right-edge is labeled i_{j+1}. If this construction at some point produces a tree whose right-edges are labeled by $1, 2, \ldots, s$ for some $s < n$, then we append this tree to the left-child of the root of the tree constructed from the block containing $s + 1$. Since all the trees coming from the blocks of $\mathbf{x}$ are full binary trees, we eventually reach a full binary tree $\tau_{\mathbf{x}} \in \mathcal{T}_n$. It is straightforward to verify that the map $\mathbf{x} \mapsto \tau_{\mathbf{x}}$ is injective.

Moreover, it follows from the construction that the two maps are mutual inverses, which concludes the proof. □

The full binary trees with $2n + 1$ nodes are clearly in bijection with binary trees on n nodes, by simply cutting off the leaves. If we move the label on a right edge to the corresponding inner node, then the bijection described in Theorem 4.3 has for instance appeared in [25, Theorem 4.4.3.1] before.

EXAMPLE 4.4

Consider $\mathbf{x} = 1\;6 \mid 2\;3\;5 \mid 4 \mid 7\;8$. The full binary trees corresponding to the blocks of $\mathbf{x}$ are the following.

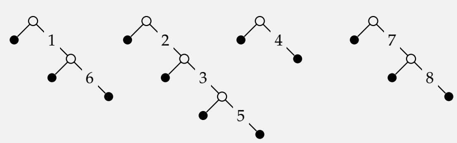

If we assemble these trees as described in the bijective proof of Theorem 4.3 we obtain the full binary tree in Figure 8.

EXERCISE 10

Construct the full binary tree with 33 nodes that is the image of $\mathbf{x} = 1\;2\;6\;7\;8\;14 \mid 3\;4\;5 \mid 9\;10\;12\;13 \mid 11 \mid 15 \mid 16$ under the bijection from Theorem 4.3.

4.2. **Dyck Paths.** Another well-known example of a family of combinatorial objects counted by the Catalan numbers is the set of Dyck paths.

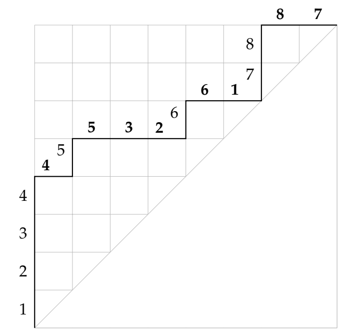

FIGURE 9. A Dyck path of semilength 8. Here we have labeled the up-steps, and matched them with their corresponding right-steps.

DEFINITION 4.5

Let $n \geq 0$. A lattice path in $\mathbb{Z} \times \mathbb{Z}$ from $(0,0)$ to (n,n) is a DYCK PATH of semilength n if it consists of $2n$ steps which are either of the form $(0,1)$ (so-called UP-STEPS) or $(1,0)$ (so-called RIGHT-STEPS), and if it stays weakly above the diagonal $x = y$.

Let us denote the set of all Dyck paths of semilength n by $\mathcal{D}_n$. Figure 9 shows a Dyck path of semilength 16.

PROPOSITION 4.6

For $n \geq 0$ we have $\lvert\mathcal{D}_n\rvert = \mathrm{Cat}(n)$.

Our proof of Proposition 4.6 establishes an explicit bijection between NC_n and $\mathcal{D}_n$, which probably first appeared in [9, Appendix E.6].

THEOREM 4.7: [9, Appendix E.6]

For $n \geq 0$ there is an explicit bijection from $\mathcal{D}_n$ to NC_n.

Proof. Let $\mathfrak{p} \in \mathcal{D}_n$. Label the up-steps of $\mathfrak{p}$ by $1, 2, \ldots, n$ in the order they occur. Now, from the i^{th} upstep, we shoot a laser with slope 1, and we label the first right-step that we hit with this beam by i. The labels of consecutive runs of right-steps of $\mathfrak{p}$ certainly form a partition $\mathbf{x}_{\mathfrak{p}}$ of $[n]$. Fix $i < j < k < l$ such that i and k are labels of the same consecutive run of right-steps, and say that j and l belong to some other consecutive runs of right-steps. By construction the y-coordinate of the run containing j must be smaller than the y-coordinate of the run containing i (and k), and the y-coordinate of the run containing l must be larger. Therefore, j and l cannot belong to the same run of right-steps, and must lie in different blocks of $\mathbf{x}_{\mathfrak{p}}$, which thus satisfies $\mathbf{x}_{\mathfrak{p}} \in NC_n$. Moreover, the map $\mathfrak{p} \mapsto \mathbf{x}_{\mathfrak{p}}$ is clearly injective.

Conversely let $\mathbf{x} \in NC_n$. Order the blocks of $\mathbf{x}$ reversely, i.e. starting with the largest element, and proceeding down to the smallest element. Now order the blocks of $\mathbf{x}$ according to their largest elements in order. Suppose that $\mathbf{x} = \{B_1, B_2, \ldots, B_k\}$ with $i_j = \max B_j$ for $j \in [k]$. We thus have $i_1 < i_2 < \cdots < i_k = n$. Let $i_0 = 0$, and construct a lattice path $\mathfrak{p}_{\mathbf{x}}$ of semilength n by drawing

$i_j - i_{j-1}$ up-steps followed by $\lvert B_j\rvert$ right-steps for every $j \in [k]$. In $\mathfrak{p}_{\mathbf{x}}$ the number of right-steps is $\sum_{j=1}^{k}\lvert B_j\rvert = n$, and the number of up-steps is $\sum_{j=0}^{k}(i_j - i_{j-1}) = i_k = n$. Moreover, $\sum_{j=1}^{s}\lvert B_j\rvert \leq i_s$ for every $s \in [k]$. (If we use more than i_s numbers to compose the first s blocks, one of the blocks must contain a number bigger than i_s, so that one of the values $i_1, i_2, \ldots, i_{s-1}$ must be bigger than i_s, contradicting the way we have ordered the blocks of $\mathbf{x}$.) It follows that $\mathfrak{p}_{\mathbf{x}} \in \mathcal{D}_n$, and the map $\mathbf{x} \mapsto \mathfrak{p}_{\mathbf{x}}$ is injective.

It is easily seen that the two maps described here are mutual inverses, which concludes the proof. □

The bijective proof of Proposition 4.6 is illustrated in Figure 9. There are many nice proofs of Proposition 4.6. Let us close this section with a proof that uses the Cycle Lemma 2.19.

Cyclic Proof of Proposition 4.6. Let $\mathfrak{p} \in \mathcal{D}_n$. We can represent $\mathfrak{p}$ by a word $w_{\mathfrak{p}}$ on the alphabet $\{U, R\}$, where we follow the path from $(0,0)$ to (n,n) and record a U whenever we see an up-step, and we record an R whenever we see a right-step. The word $w_{\mathfrak{p}}$ clearly has length $2n$, and it contains n-times the letter U and n-times the letter R. Moreover, every prefix of $w_{\mathfrak{p}}$ has the property that it contains at least as many letters U as it contains letters R. Append an additional letter U to the front of $w_{\mathfrak{p}}$ to obtain a word $\bar{w}_{\mathfrak{p}}$ with the property that every prefix of $\bar{w}_{\mathfrak{p}}$ contains more letters U than it contains letters R. Record for later that $\mathfrak{p}$ is uniquely determined by $\bar{w}_{\mathfrak{p}}$.

Now consider all words of length $2n+1$ with $n+1$ letters U and n letters R. Declare two such words equivalent if one can be obtained by cyclically shifting the letters of the other. Since each word has $2n+1$ letters, every equivalence class has $2n+1$ elements, and the total number of such words is $\binom{2n+1}{n+1}$.

The Cycle Lemma 2.19 implies that in each equivalence class there exists a unique 1-dominating word, and in view of the first part of this proof the Dyck paths of semilength n correspond bijectively to these 1-dominating words. We have just shown that

$$\lvert\mathcal{D}_n\rvert = \frac{1}{2n+1}\binom{2n+1}{n+1} = \frac{(2n+1)!}{(2n+1)(n+1)!n!} = \frac{(2n)!}{(n+1)!n!} = \frac{1}{n+1}\binom{2n}{n} = \mathrm{Cat}(n).$$

□

✎ EXERCISE 11

Construct the Dyck path of semilength 16 that is the image of $\mathbf{x} = 1\ 2\ 6\ 7\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 10\ 12\ 13 \mid 11 \mid 15 \mid 16$ under the bijection from Theorem 4.7.

4.3. **Nonnesting Set Partitions.** There is in fact a second family of set partitions that is enumerated by the Catalan numbers. Instead of crossings, however, we now forbid nestings.

DEFINITION 4.8

For $n \geq 0$ a set partition $\mathbf{x} \in \Pi_n$ is NONNESTING if it does not contain four elements $i < j < k < l$ such that $i \sim_{\mathbf{x}} l$ and $j \sim_{\mathbf{x}} k$ but $i \not\sim_{\mathbf{x}} j$.

The justification for this name comes from the fact that a set partition is nonnesting if and only if in its arc diagram no two arcs nest. Let us denote the set of all nonnesting set partitions of $[n]$ by NN_n.

Analogously to the crossing case, all set partitions of $[n]$ for $n \leq 3$ are nonnesting. The smallest nesting set partition is $1\,4 \mid 2\,3$, which is also the only nesting set partition of $[4]$.

The following bijection from NC_n to NN_n was first described in [4], and later refined and extended in [13] and [11]. Other bijections were for instance given in [18] or [3].

Let $\mathbf{x} \in \Pi_n$ with $\mathbf{x} = \{B_1, B_2, \ldots, B_k\}$. As usual, suppose that $\min B_i < \min B_j$ whenever $i < j$. For $i \in [k]$ define $m_i = \min B_i$ and $b_i = \lvert B_i \rvert$. Let $M_{\mathbf{x}} = \{m_1, m_2, \ldots, m_k\}$ as in Exercise 4, and let $b(\mathbf{x}) = (b_1, b_2, \ldots, b_k)$. Observe that $m_1 = 1$.

THEOREM 4.9: [4, Theorem 3.1]

For $\mathbf{x} \in NN_n$ there exists a unique $\mathbf{x}' \in NC_n$ with $M_{\mathbf{x}} = M_{\mathbf{x}'}$ and $b(\mathbf{x}) = b(\mathbf{x}')$.

Proof. For the proof we explicitly construct the noncrossing set partition $\mathbf{x}'$ from $M_{\mathbf{x}}$ and $b(\mathbf{x})$. We start with a collection of k *chains*, whose lengths are given by $b(\mathbf{x})$. Here a chain of length b is a sequence of $b - 1$ consecutive arcs, where the labels of the start and end points of these arcs are not yet determined.

The construction then works inductively as follows. We first place the chain of length b_1, and label its first vertex by 1. In fact we may label the first $m_2 - 1$ vertices by $1, 2, \ldots, m_2 - 1$. Then we place the chain of length b_2 such that its first vertex becomes the m_2^{nd} total vertex, and we do not introduce crossings. In particular, this chain is placed entirely below the arc of the first chain starting with the $(m_2 - 1)^{\mathrm{st}}$ vertex. We proceed in the same way with the remaining chains, until all vertices are labeled. Observe that this procedure is uniquely determined.

The inverse map starts with a noncrossing set partition $\mathbf{x}'$ and creates a nonnesting set partition $\mathbf{x}$ from $M_{\mathbf{x}'}$ and $b_{\mathbf{x}'}$. It works verbatim, except that we try to avoid nestings instead of crossings. Again, the steps in this procedure are uniquely determined. □

EXAMPLE 4.10

Consider $\mathbf{x} = 1\,6 \mid 2\,3\,5 \mid 4 \mid 7\,8 \in NC_8$. We have

$$M_{\mathbf{x}} = \{1, 2, 4, 7\} \quad \text{and} \quad b(\mathbf{x}) = (2, 3, 1, 2).$$

We now successively place the four chains of lengths $2, 3, 1$, and 2 while avoiding nestings.

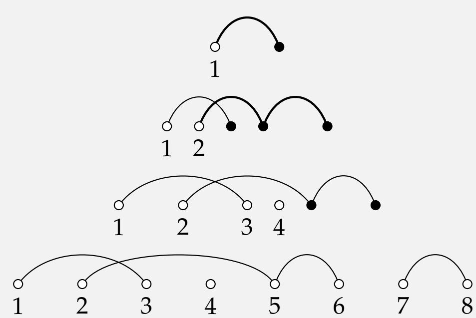

We thus obtain the nonnesting set partition $\mathbf{x}' = 1\,3 \mid 2\,5\,6 \mid 4 \mid 7\,8$ as the image of $\mathbf{x}$ under the bijection from Theorem 4.9.

There is yet another perspective on nonnesting set partitions: these can in fact be viewed as antichains in a particular poset.

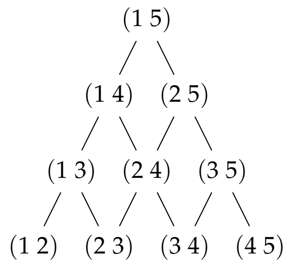

FIGURE 10. The triangular poset Δ_5.

DEFINITION 4.11

Let $n \geq 1$, and define $T_n = \{(i,j) \mid 1 \leq i < j \leq n\}$. Consider the partial order on T_n, where $(i,j) \leq (k,l)$ if and only if $i \geq k$ and $j \leq l$. We call the resulting poset $\Delta_n = (T_n, \leq)$ the TRIANGULAR POSET of order n.

Figure 10 shows the triangular poset of order 5.

THEOREM 4.12

For $n \geq 0$, the collection of arcs of $\mathbf{x} \in NN_n$ is an antichain of Δ_n, and every antichain of Δ_n arises in this fashion.

Proof. Let $X \subseteq T_n$ be a collection of pairs of the form (i,j) for $1 \leq i < j \leq n$. If $|X| \leq 1$, then X is certainly an antichain in Δ_n, and the set partition whose arcs are determined by X is certainly nonnesting.

Now suppose that $|X| \geq 2$. Then we can find $(i,j),(k,l) \in X$. Observe that if $i = k$ or $j = l$, then (i,j) and (k,l) are comparable in Δ_n, and in that case they do not correspond to arcs of a set partition. Let us therefore without loss of generality assume that $i < k$. There are two possibilities:

(i) $i < j, k < l$. It follows that (i,j) and (k,l) are not comparable in Δ_n, and the arcs in the corresponding set partition do not nest. (In this case, the relation between j and k does not matter.)

(ii) $i < k < l < j$. It follows that $(k,l) \leq (i,j)$ and the arc corresponding to (k,l) nests in the arc corresponding to (i,j). □

PROPOSITION 4.13

For $n \geq 0$ we have $|NN_n| = \mathrm{Cat}(n)$. Moreover, the number of antichains of Δ_n is given by $\mathrm{Cat}(n)$.

Proof. This follows from Proposition 2.2 and Theorems 4.9 and 4.12 □

It turns out that the nonnesting set partitions share quite some enumerative features with the noncrossing set partitions. On a structural level, however, the posets $(NC_n, \leq_{\mathrm{dref}})$ and $(NN_n, \leq_{\mathrm{dref}})$ are quite different. In fact, for $n \geq 5$ the poset $(NN_5, \leq_{\mathrm{dref}})$ is not self-dual and no longer a lattice; for $n \geq 6$ it is not even graded anymore.

The reason for the failure of the lattice property is that the intersection of two nonnesting set partitions need not be nonnesting anymore. Take for instance $\mathbf{x} = 1\,2\,4\,5 \mid 3$ and $\mathbf{x}' = 1\,3\,5 \mid 2\,4$. Both set partitions are nonnesting, but $\mathbf{x} \wedge_{\Pi} \mathbf{x}' = 1\ 5 \mid 2\ 4 \mid 3$, which is unfortunately nesting. Also there are two nonnesting, mutually incomparable set partitions that refine $\mathbf{x} \wedge_{\Pi} \mathbf{x}'$, namely $\mathbf{y} = 1\,5 \mid 2 \mid 3 \mid 4$ and $\mathbf{y}' = 1 \mid 2\,4 \mid 3 \mid 5$.

If we on the other hand consider the nonnesting set partition $\mathbf{x} = 1\ 4 \mid 2\ 5 \mid 3\ 6$, then all its upper covers in $(\Pi_n, \leq_{\text{dref}})$ are nesting. These set partitions are $\mathbf{x}_1 = 1\,2\,4\,5 \mid 3\,6$, $\mathbf{x}_2 = 1\,3\,4\,6 \mid 2\,5$ and $\mathbf{x}_3 = 1\,4 \mid 2\,3\,5\,6$.

It is an intriguing open question whether some other properties of $(NC_n, \leq_{\text{dref}})$, for instance the existence of an edge-labeling which has a unique rising chain per interval is true in $(NN_n, \leq_{\text{dref}})$.

✎ EXERCISE 12

Construct the nonnesting partition of $[16]$ that is the image of $\mathbf{x} = 1\ 2\ 6\ 7\ 8\ 14 \mid 3\ 4\ 5 \mid 9\ 10\ 12\ 13 \mid 11 \mid 15 \mid 16$ under the bijection from Theorem 4.9.

Let us conclude this section with an intriguing observation. Let r_i denote the number elements of rank i on Δ_n, and let $e_i = n - r_i$. We then see that $e_i = i$ for $i \in [n-1]$.

PROPOSITION 4.14

For $n \geq 0$ we have

$$\mathrm{Cat}(n+1) = \prod_{i=1}^{n} \frac{e_i + e_n + 2}{e_i + 1}.$$

Proof. This is a simple computation:

$$\begin{aligned}\prod_{i=1}^{n} \frac{e_i + e_n + 2}{e_i + 1} &= \frac{n+3}{2} \cdot \frac{n+4}{3} \cdot \cdots \cdot \frac{2n+2}{n+1}\\ &= \frac{(2n+2)!}{(n+2)!(n+1)!}\\ &= \frac{1}{n+2}\binom{2n+2}{n+1}\\ &= \mathrm{Cat}(n+1).\end{aligned}$$

□

The numbers e_i defined above are the EXPONENTS of the symmetric group $\mathfrak{S}_{n+1}$, and the connection between these exponents and the rank-numbers in Δ_n was first observed by A. Shapiro, R. Steinberg, and B. Kostant, see [15], (in a much more general framework).

4.4. **312-Avoiding Permutations.** Now we turn our attention to permutations avoiding the pattern 312. For this recall that $\mathfrak{S}_n$ is the symmetric group of degree n, i.e. the group of all permutations of $[n]$.

DEFINITION 4.15

Let $n \geq 0$. A permutation $\pi \in \mathfrak{S}_n$ AVOIDS the pattern 312 if there exist no three integers $i < j < k$ such that $\pi(j) < \pi(k) < \pi(i)$.

Let $\mathfrak{S}_n(312)$ denote the set of all 312-avoiding permutations of $\mathfrak{S}_n$. The number of 312-avoiding permutations was first determined by D. Knuth, see [14, Exercise 2.2.1.4].

PROPOSITION 4.16: [14, Exercise 2.2.1.4]

For $n \geq 0$ we have $\lvert\mathfrak{S}_n(312)\rvert = \mathrm{Cat}(n)$.

We want to prove Proposition 4.16 bijectively, and therefore need to establish a few more concepts. A DESCENT of a permutation $\pi \in \mathfrak{S}_n$ is an integer $i \in [n]$ such that $\pi(i) > \pi(i+1)$. Let $\mathrm{Des}(\pi)$ denote the set of descents of π, and let $\mathrm{des}(\pi) = \lvert\mathrm{Des}(\pi)\rvert$. Moreover, for a finite set X we denote by $\wp(X)$ its power set.

LEMMA 4.17

For $n \geq 0$ the map $\mathrm{Des} : \mathfrak{S}_n \to \wp([n])$ is injective.

Proof. We prove this result by induction on the number of descents. If $\mathrm{des}(\pi) = 0$, then π must be the identity. Now consider $\pi, \pi' \in \mathfrak{S}_n$ with $\mathrm{Des}(\pi) = X = \mathrm{Des}(\pi')$ and suppose that $\mathrm{des}(\pi) = \mathrm{des}(\pi') = k > 0$. Let $i = \min X$, and consider the permutations σ, σ' that arise from π and π', respectively, by exchanging the i^{th} and the $(i+1)^{\text{st}}$ letter. We have $\mathrm{des}(\sigma) = \mathrm{des}(\sigma') = k-1 < k$, and $\mathrm{Des}(\sigma) = X \setminus \{i\} = \mathrm{Des}(\sigma')$. By induction we find $\sigma = \sigma'$, and thus $\pi = \pi'$. □

Following [24] we define a NONCROSSING ARC DIAGRAM as follows: we write n points in order on a vertical line with 1 at the bottom, and we may connect two vertices i and j with $i < j$ by an arc, which is a monotone curve that passes either to the left or to the right of every point strictly between i and j. No two arcs of the same diagram may intersect except perhaps at their endpoints, and no two arcs of the same diagram may have the same lower or upper endpoints. Let $\mathrm{Arc}(n)$ denote the set of all noncrossing arc diagrams with n points.

Each permutation $\pi \in \mathfrak{S}_n$ has a PERMUTATION DIAGRAM defined by drawing n dots in an $n \times n$-box, where the i^{th} dot is placed at coordinate $\bigl(i, \pi(i)\bigr)$. Now we connect the i^{th} dot with the $(i+1)^{\text{st}}$ dot if and only if $i \in \mathrm{Des}(\pi)$. If we now move the dots to the left on a vertical line, where we allow lines to bend, but not to pass through dots or other lines, then we obtain a noncrossing arc diagram, which we denote by $\delta(\pi)$.

THEOREM 4.18: [24, Theorem 3.1]

For $n \geq 0$ the map $\delta : \mathfrak{S}_n \to \mathrm{Arc}(n)$ is a bijection.

Proof. It follows from Lemma 4.17 that δ is injective. Let us now describe the inverse map. Given an arc diagram $D \in \mathrm{Arc}(n)$, denote its connected components by $C_1, C_2, \ldots, C_k$. These connected components are either single dots or sequences of dots connected by a single curve. Then C_i is a LEFT COMPONENT if there is no component that is strictly to its left in the drawing of D. Since D is noncrossing, we can always find a left component of D. (See the proof of [24, Proposition 3.2] for a detailed explanation. If we have two arcs, we can say one of them is LEFT OF the other, when it is drawn further to the left than the other. It then remains to show that this relation is acyclic.) We now inductively construct a permutation from D. We order the left components by their minimal elements, and we remove the left component which comes first in this order, and write down its entries in reverse order. Then we remove this left component from D, and repeat the process, until

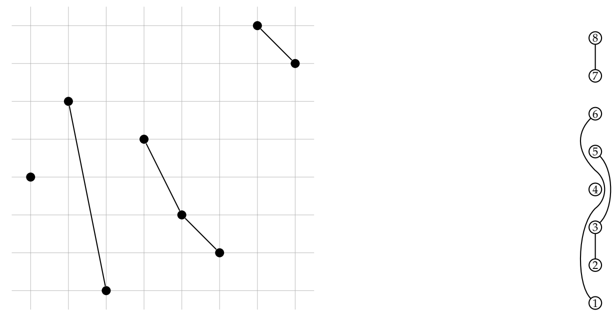

(A) The permutation diagram of $\pi = 46153287$ with marked descents.

(B) The noncrossing arc diagram $\delta(\pi)$ for $\pi = 46153287$.

FIGURE 11. The map from permutations to noncrossing arc diagrams.

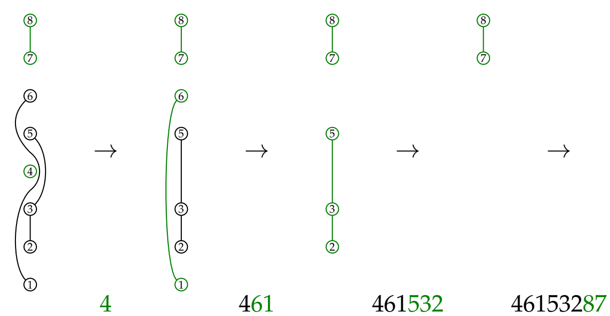

FIGURE 12. How to construct a permutation from a noncrossing arc diagram. The available left components at each step are marked in green.

we have no components left. It is thus ensured that we obtain a permutation π_D of $[n]$, since each $k \in [n]$ belongs to exactly one connected component of D, and elements that have an outgoing edge in D are by construction descents of π_D. It is then clear that $\delta(\pi_D) = D$, and we are done. $\square$

EXAMPLE 4.19

Let us illustrate the map δ from Theorem 4.18 and its inverse with an example. Consider the permutation $\pi = 46153287 \in \mathfrak{S}_8$. We have $\mathrm{Des}(\pi) = \{2,4,5,7\}$. Its permutation diagram with marked descents is displayed in Figure 11a, and the noncrossing arc diagram $\delta(\pi)$ is shown in Figure 11b. Figure 12 illustrates the inverse map.

We conclude this section with the observation that the map δ establishes a bijection between $\mathfrak{S}_n(312)$ and NC_n.

Lemma 4.20

Let $n \geq 0$ and $\pi \in \mathfrak{S}_n$. We have $\pi \notin \mathfrak{S}_n(312)$ if and only if there exist integers i, k with $i+1 < k$ such that $\pi_{i+1} < \pi_k < \pi_i$.

Proof. If such integers i, k exist, then we clearly have $\pi \notin \mathfrak{S}_n(312)$. Conversely, suppose that $\pi \notin \mathfrak{S}_n(312)$. By definition there exist integers $i < j < k$ with $\pi_j < \pi_k < \pi_i$. We proceed by induction on $j - i$. If $j - i = 1$, we are done. Otherwise we can find j' with $i < j' < j$. There are two possible cases: if $\pi_{j'} < \pi_k$, then we have $\pi_{j'} < \pi_k < \pi_i$ and $j' - i < j - i$ so that we conclude the claim by induction. Otherwise $\pi_{j'} > \pi_k$, but then we obtain $\pi_j < \pi_k < \pi_{j'}$, and $j - j' < j - i$ and we conclude the claim by induction once again. □

Theorem 4.21

For $n \geq 0$ the map δ restricts to a bijection from $\mathfrak{S}_n(312)$ to NC_n.

Proof. This proof relies on two simple observations. Firstly, the set of noncrossing set partitions are clearly in bijection with the set of noncrossing arc diagrams in which no arc passes to the left of a dot. (Rotate the diagram by 90 degrees, and reflect horizontally.) Secondly, an arc passes to the left of a dot in $\delta(\pi)$ if and only if π has a descent at i such that there is some $k > i+1$ with $\pi(i+1) < \pi(k) < \pi(i)$. Lemma 4.20 implies that this is equivalent to $\pi \notin \mathfrak{S}_n(312)$ so that δ^{-1} is in fact surjective. □

✎ Exercise 13

Find the permutation $\pi \in \mathfrak{S}_{16}(312)$ whose image $\delta(\pi)$ is our running example $\mathbf{x} =$ 1 2 6 7 8 14 | 3 4 5 | 9 10 12 13 | 11 | 15 | 16.

Proof of Proposition 4.16. This follows from Proposition 2.2 and Theorem 4.21. □

4.5. **Triangulations.** For $n \geq 3$ consider a regular n-gon P_n, whose vertices are labeled clockwise by $1, 2, \ldots, n$. A DIAGONAL is an edge between two vertices i and j, where $|j - i| > 1$. Without loss of generality we can identify the diagonals of P_n with pairs of integers (i, j) where $i < j$. It is quickly verified that P_n admits $\frac{n(n-3)}{2}$ diagonals. Two diagonals (i, k) and (j, l) are CROSSING if $i < j < k < l$.

Definition 4.22

Let $n \geq 3$. A TRIANGULATION of P_n is a maximal set of pairwise noncrossing diagonals.

Let $\Delta(P_n)$ denote the set of triangulations of P_n. We can quickly check by induction that any triangulation of P_n consists of $n - 3$ diagonals.

Proposition 4.23

For $n \geq 1$ we have $\left|\Delta(P_{n+2})\right| = \mathrm{Cat}(n)$.

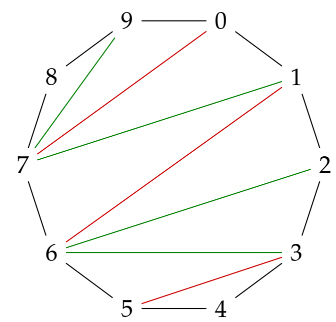

FIGURE 13. A colored triangulation of P_{10}.

For our bijective proof of Proposition 4.23 we make use of the coloring method of diagonals belonging to a given triangulation $T \in \Delta(P_{n+2})$ that was described in [17, Section 6]. Firstly, label the vertices of P_{n+2} by $0, 1, \ldots, n+1$, and fix a diagonal (i, j) in T. If we remove this diagonal from T, then we obtain a quadrilateral $Q(T)$ with vertices i, j, k, l. In particular, (k, l) is the other diagonal in $Q(T)$. If we traverse the vertices of $Q(T)$ in order, then we color (i, j) green if j is the vertex of $Q(T)$ with the biggest label, and we color it red otherwise. A colored triangulation of P_{10} is shown in Figure 13.

LEMMA 4.24: [17, Section 6]

For $n \geq 1$, any set of less than $n-1$ diagonals of P_{n+2} can be completed in a unique way to a triangulation of P_{n+2} by adding only green diagonals.

Proof. Let D be a set of diagonals of P_{n+2} with $|D| < n-1$. Since any triangulation of P_{n+2} consists of $n-1$ diagonals, we can find a polygon P' with more than three vertices whose interior does not contain a diagonal of D. Denote the vertices of P' by $i_1, i_2, \ldots, i_k$ in increasing order. Any diagonal of the form (i_1, i_j) for $j \in \{3, 4, \ldots, k-1\}$ must be red in any triangulation extending D. This follows, since the quadrilateral containing (i_1, i_j) must either contain the vertices i_2, i_{j+1} or i_{j-1}, i_k. In any case, i_j is not the biggest vertex in this quadrilateral. Since we only allow D to be completed by green edges, we conclude that any such triangulation must contain the diagonal (i_2, i_k). If we iterate this process, we see that we see that there is a unique way to complete D by green edges. $\square$

A consequence of Lemma 4.24 is that every triangulation of P_{n+2} is uniquely determined by its set of red diagonals. The bijection from $\Delta(P_{n+2})$ and NC_n appears for example in [34, Section 8].

THEOREM 4.25: [34, Section 8]

For $n \geq 1$ there is an explicit bijection from $\Delta(P_{n+2})$ to NC_n.

Proof. Let $T \in \Delta(P_{n+2})$. Remove the external edges of T and all green diagonals. For every red diagonal (i, j) move the vertex i a little bit counterclockwise, and move the vertex j a little bit clockwise so that i and j sit above the diagonal (i, j). If we now remove the vertices 0 and $n+1$, we are left with a partition $\mathbf{x}_T \in \Pi_n$. By construction (since the diagonals of T are mutually noncrossing), we have $\mathbf{x}_T \in NC_n$.

Conversely, let $\mathbf{x} \in NC_n$. If $n = 1$, then $\mathbf{x} = 1$ and we map it to the empty triangulation of P_3. Otherwise, let B be the unique block of $\mathbf{x}$ containing 1, and suppose that $k = \max B$. Draw the vertices $0, 1, \ldots, n+1$ clockwise on a circle, and draw a red diagonal from 0 to $k+1$. Then $\mathbf{x} \setminus B$ is (isomorphic to) a noncrossing set partition of $[n']$ for some $n' < n$, and we can inductively insert the remaining red diagonals. Lemma 4.24 implies that we can complete this set of red diagonals in a unique way to a triangulation $T_{\mathbf{x}} \in \Delta(P_{n+2})$.

Again we see quickly that the two maps are mutual inverses, which concludes the proof. ☐

EXAMPLE 4.26

Consider $\mathbf{x} = 1\,6 \mid 2\,3\,5 \mid 4 \mid 7\,8 \in NC_8$. We inductively add the following separating lines into P_{10}.

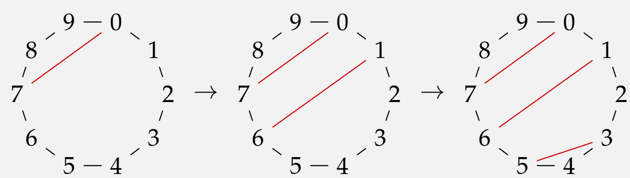

The unique way to complete this partial triangulation to a full triangulation of P_{10} by green edges yields the triangulation shown in Figure 13.

✎ EXERCISE 14

Construct the triangulation of P_{18} that corresponds to the noncrossing partition $\mathbf{x} = 1\,2\,6\,7\,8\,14 \mid 3\,4\,5 \mid 9\,10\,12\,13 \mid 11\,15\,16$.

REFERENCES

[1] Federico Ardila, Felipe Rincón, and Lauren Williams, *Positroids and Non-Crossing Partitions*, Transactions of the American Mathematical Society **368** (2016), 337–363.

[2] Drew Armstrong, *Generalized Noncrossing Partitions and Combinatorics of Coxeter Groups*, Memoirs of the American Mathematical Society **202** (2009).

[3] Drew Armstrong, Christian Stump, and Hugh Thomas, *A Uniform Bijection between Nonnesting and Noncrossing Partitions*, Transactions of the American Mathematical Society **365** (2013), 4121–4151.

[4] Christos A. Athanasiadis, *On Noncrossing and Nonnesting Partitions for Classical Reflection Groups*, The Electronic Journal of Combinatorics (1998).

[5] Philippe Biane, *Representations of Symmetric Groups and Free Probability*, Advances in Mathematics **138** (1998), 126–181.

[6] Anders Björner, *Shellable and Cohen-Macaulay Partially Ordered Sets*, Transactions of the American Mathematical Society **260** (1980), 159–183.

[7] Brian A. Davey and Hilary A. Priestley, *Introduction to Lattices and Order*, Cambridge University Press, Cambridge, 2002.

[8] Nachum Dershowitz and Shmuel Zaks, *The Cycle Lemma and Some Applications*, European Journal of Combinatorics **11** (1990), 35–40.

[9] Emeric Deutsch, *Dyck Path Enumeration*, Discrete Mathematics **204** (1999), 167–202.

[10] Paul H. Edelman, *Chain Enumeration and Non-Crossing Partitions*, Discrete Mathematics **31** (1980), 171–180.

[11] Alex Fink and Benjamin I. Giraldo, *Bijections between Noncrossing and Nonnesting Partitions for Classical Reflection Groups*, Portugaliae Mathematicae **67** (2009), 369–401.

[12] Colin Ingalls and Hugh Thomas, *Noncrossing Partitions and Representations of Quivers*, Composition Mathematica **145** (2009), 1533–1562.

[13] Anisse Kasraoui and Jiang Zeng, *Distribution of Crossings, Nestings and Alignments of Two Edges in Matchings and Partitions*, The Electronic Journal of Combinatoics **13** (2006).
[14] Donald E. Knuth, *The Art of Computer Programming, Vol. 1, Fundamental Algorithms*, 3rd ed., Addison-Wesley, Reading, MA, 1973.
[15] Bertram Kostant, *The Principal Three-Dimensional Subgroup and the Betti Numbers of a Complex Simple Lie Group*, American Journal of Mathematics **81** (1959), 973–1032.
[16] Germain Kreweras, *Sur les partitions non croisées d'un cycle*, Discrete Mathematics **1** (1972), 333–350.
[17] Carl W. Lee, *The Associahedron and Triangulations of the n-gon*, European Journal of Combinatorics **10** (1989), 551–560.
[18] Ricardo Mamede, *A Bijection between Noncrossing and Nonnesting Partitions of Type A, B and C*, Contributions to Discrete Mathematics **6** (2011), 70–90.
[19] Toufik Mansour, *Combinatorics of Set Partitions*, CRC Press, Boca Raton, FL, 2013.
[20] Suho Oh, *Positroids and Schubert Matroids*, Journal of Combinatorial Theory (Series A) **118** (2011), 2426–2435.
[21] James G. Oxley, *Matroid Theory*, Oxford University Press, Oxford, 1992.
[22] Igor Pak, *History of Catalan Numbers* (2014), available at arXiv:1408.5711.
[23] Alexander Postnikov, *Total Positivity, Grassmannians, and Networks* (2006), available at arXiv:0609764.
[24] Nathan Reading, *Noncrossing Arc Diagrams and Canonical Join Representations*, SIAM Journal on Discrete Mathematics **29** (2015), 736–750.
[25] Claus Michael Ringel, *The Catalan Combinatorics of the Hereditary Artin Algebras* (Alex Martsinkovsky, Gordana Todorov, and Kiyoshi Igusa, eds.), Contemporary Mathematics, vol. 673, American Mathematical Society, Providence, RI, 2016.
[26] Gian-Carlo Rota, *On the Foundations of Combinatorial Theory I: Theory of Möbius Functions*, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete **2** (1964), 340–368.
[27] Terence P. Speed, *Cumulants and Partition Lattices*, Australian & New Zealand Journal of Statistics **25** (1983), 378–388.
[28] Roland Speicher, *Multiplicative Functions on the Lattice of Non-Crossing Partitions and Free Convolution*, Mathematische Annalen **298** (1994), 611–628.
[29] Richard P. Stanley, *Catalan Addendum. New Problems for Enumerative Combinatorics, Vol. 2*. http://math.mit.edu/~rstan/ec/catadd.pdf.
[30] ______, *Supersolvable Lattices*, Algebra Universalis **2** (1972), 197–217.
[31] ______, *Finite Lattices and Jordan-Hölder Sets*, Algebra Universalis **4** (1974), 361–371.
[32] ______, *Enumerative Combinatorics, Vol. 1*, 2nd ed., Cambridge University Press, Cambridge, 2011.
[33] ______, *Catalan Numbers*, Cambridge University Press, Cambridge, 2015.
[34] Hugh Thomas, *Tamari Lattices and Noncrossing Partitions in Type B*, Discrete Mathematics **306** (2006), 2711–2723.
[35] ______, *The Tamari Lattice as it Arises in Quiver Representations* (Folkert Müller-Hoissen, Jean Marcel Pallo, and Jim Stasheff, eds.), Progress in Mathematics, vol. 299, Springer, Basel, 2012.
[36] Dan Voiculescu, *Free Probability Theory: Random Matrices and von Neumann Algebras*, Proceedings of the International Congress of Mathematicians 1994 (1995), 227–241.
[37] Hassler Whitney, *On the Abstract Properties of Linear Dependence*, American Journal of Mathematics **57** (1935), 509–533.