A PATTERN THEOREM FOR RANDOM SORTING NETWORKS

OMER ANGEL, VADIM GORIN, AND ALEXANDER HOLROYD

Abstract. A sorting network is a shortest path from $12\cdots n$ to $n\cdots 21$ in the Cayley graph of the symmetric group S_n generated by nearest-neighbor swaps. A pattern is a sequence of swaps that forms an initial segment of some sorting network. We prove that in a uniformly random n-element sorting network, any fixed pattern occurs in at least cn^2 disjoint space-time locations, with probability tending to 1 exponentially fast as $n \to \infty$. Here c is a positive constant which depends on the choice of pattern. As a consequence, the probability that the uniformly random sorting network is geometrically realizable tends to 0.

1. Introduction

Let S_n be the group of all permutations $\sigma = (\sigma(1), \dots, \sigma(n))$ of $\{1, \dots, n\}$ with composition given by $(\sigma\tau)(i) = \sigma(\tau(i))$. We denote by σ_j the adjacent transposition or **swap** $(j\ j{+}1) = (1, \dots, j{+}1, j, \dots, n)$. A **sorting network** of **size** n is a sequence $(s_1, s_2, \dots, s_N)$ of $N := \binom{n}{2}$ integers with $0 < s_k < n$, such that the composition $\sigma_{s_1}\sigma_{s_2}\cdots\sigma_{s_N}$ equals the reverse permutation $(n, n-1, \dots, 1)$. We sometimes say that at time k a swap occurs at position s_k, and we illustrate a sorting network by a set of crosses with coordinates (k, s_k) for $k = 1, \dots, N$. (This is natural, since the crosses may be joined by horizontal lines to give a "wiring diagram" consisting of n polygonal lines whose order is reversed as we move from left to right; see Figure 1.)

Interest in sorting networks was initiated by Stanley, who proved in [St] that the number of sorting networks of size n is equal to the number of standard staircase-shape Young tableaux of size n, i.e. those with shape $(n-1, n-2, \dots, 1)$. Uniformly random sorting networks were introduced and studied by Angel, Holroyd, Romik, and Virag in [AHRV], giving rise to many striking results and conjectures.

Date: August 2011.

2010 *Mathematics Subject Classification.* 60C05; 05E10; 68P10.

Key words and phrases. Sorting network, random sorting, reduced word, pattern, Young tableau.

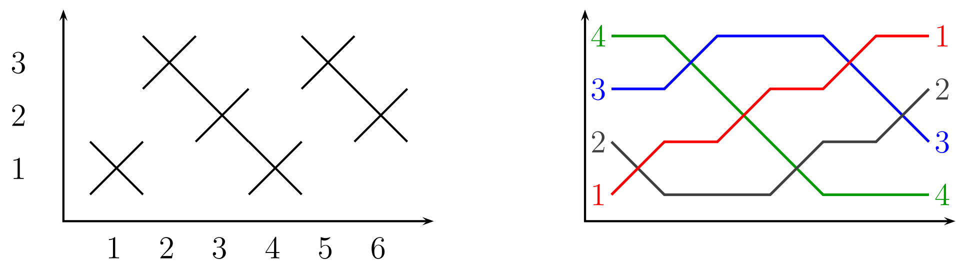

FIGURE 1. *Left:* the sorting network $(1, 3, 2, 1, 3, 2)$ of size 4, illustrated by crosses corresponding to its swaps. *Right:* the associated wiring diagram.

A **pattern** is any finite sequence of positive integers that is an initial segment of some sorting network. Thus for example, $(1, 2, 1)$ and $(4, 2)$ are patterns, but $(1, 1)$ and $(1, 2, 1, 2)$ are not. The **size** of a pattern is the minimum size of a sorting network that contains it as an initial segment, which is also one more than the maximal element in the pattern.

Let $\omega = (s_1, \dots, s_N)$ be a sorting network of size n and let γ be a pattern. Let $[i, j] \subseteq [1, N]$ and $[a, b] \subseteq [1, n-1]$, and consider the subsequence $t_1, \dots, t_\ell$ of $s_i, \dots, s_j$ consisting of precisely those elements lying in the interval $[a, b]$. We say that the pattern γ **occurs** at time interval $[i, j]$ and position $[a, b]$ (or simply at $[i, j] \times [a, b]$) if $\gamma = (t_1 - a + 1, \dots, t_k - a + 1)$, and no $k \in [i, j]$ has $s_k \in \{a - 1, b + 1\}$. In other words, the swaps in the space-time window $[i, j] \times [a, b]$ are precisely those of γ, after an appropriate shift in location, and there are no swaps at the two adjacent positions in this time interval. See Figure 2 for an example.

We say that a pattern γ **occurs** R **times** in a sorting network ω if R is the maximum integer for which there exist pairwise disjoint rectangles $\{[i_r, j_r] \times [a_r, b_r]\}_{r=1}^{R}$ such that γ occurs at each. See Figure 3.

Theorem 1. *Fix any pattern γ of size k. There exist constants $c_1, c_2 > 0$ (depending on γ) such that for every $n \geq k$, the pattern γ occurs at least $c_1 n^2$ times in a uniformly random sorting network of size n, with probability at least $1 - e^{-c_2 n}$.*

We conjecture that the probability in Theorem 1 is in fact at least $1 - e^{-cn^2}$ for some $c = c(\gamma)$.

We will prove Theorem 1 by establishing a closely related result about uniformly random standard staircase-shape Young tableaux, and using a bijection due to Edelman and Greene [EG] between sorting networks and Young tableaux.

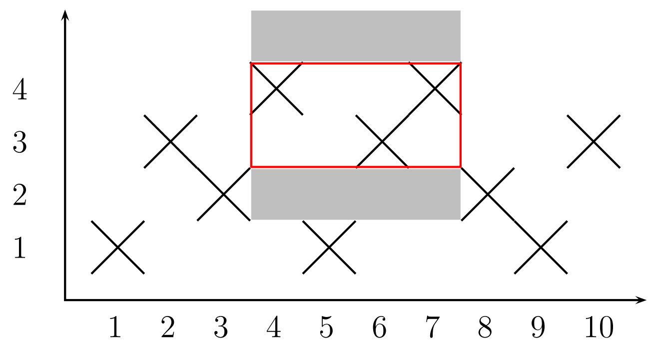

FIGURE 2. The pattern $(2, 1, 2)$ occurs in the sorting network $(1, 3, 2, 4, 1, 3, 4, 2, 1, 3)$ at time interval $[i, j] = [4, 7]$ and position $[a, b] = [3, 4]$. Note the requirement that the shaded regions contain no swaps.

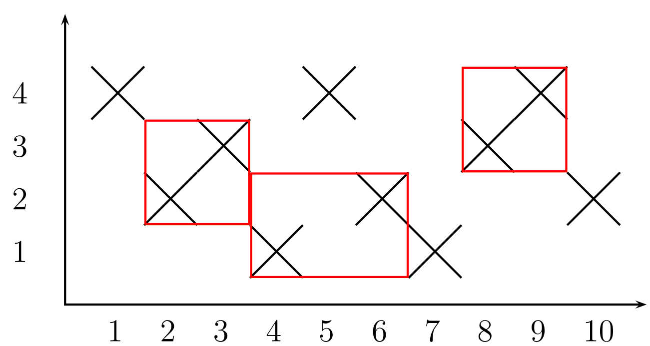

FIGURE 3. Pattern $(1, 2)$ occurs 3 times in the sorting network $(4, 2, 3, 1, 4, 2, 1, 3, 4, 2)$.

Write $\mathbb{N} = \{1, 2, \ldots\}$. A **Young diagram** λ is a set of the form $\{(i, j) \in \mathbb{N}^2 : 1 \leq j \leq \lambda_i\}$, where $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \cdots \geq 0$ are integers and $\sum_{i=1}^{\infty} \lambda_i =: |\lambda| < \infty$. The numbers λ_i are the **row lengths** of λ. In what follows we denote by $(\lambda_1, \lambda_2, \ldots)$ the Young diagram with row lengths $\lambda_1 \geq \lambda_2 \geq \ldots$. We call an element $x = (i, j) \in \lambda$ a **box**, and draw it as a unit square at location (i, j) (with the traditional convention that $(1, 1)$ is at the top left and the first coordinate is vertical). A **tableau** T of shape λ is a map from λ to the integers whose values are non-decreasing along rows and columns. We call $T(x)$ the **entry** assigned to box x. A **standard Young tableau** is a tableau T of shape λ such that the set of entries of T is $\{1, 2 \ldots, |\lambda|\}$. We are mostly interested in **standard staircase-shape Young tableaux** of size n, i.e. those with shape staircase Young diagram $(n - 1, n - 2, \ldots, 1)$.

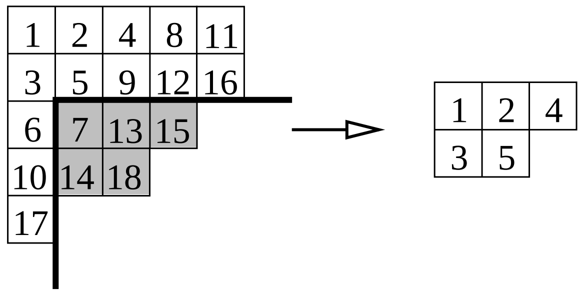

FIGURE 4. A standard Young tableau T of shape $\lambda = (5, 5, 4, 3, 1)$, the subdiagram $\lambda^{(3,2)}$, the subtableau $T^{(3,2)}$, and a standard Young tableau identically ordered with $T^{(3,2)}$.

For $(i, j), (k, \ell) \in \mathbb{N}^2$ we write $(i, j) \leq (k, \ell)$ if $i \leq k$ and $j \leq \ell$. For a Young diagram λ and a box $(i, j) \in \lambda$, we define the **subdiagram** $\lambda^{(i,j)}$ with top-left corner (i, j) by $\lambda^{(i,j)} := \{(k, \ell) \in \lambda : (k, \ell) \geq (i, j)\}$; clearly $\lambda^{(i,j)}$ is mapped to a Young diagram by the translation $(k, \ell) \mapsto (k - i + 1, \ell - j + 1)$. If T is a tableau of shape λ then we define the **subtableau** $T^{(i,j)}$ to be the restriction of T to $\lambda^{(i,j)}$, and we call $\lambda^{(i,j)}$ the **support** of $T^{(i,j)}$.

We say that two tableaux S and T of the same shape λ are **identically ordered** if for all $x, y \in \lambda$ we have $S(x) < S(y)$ if and only if $T(x) < T(y)$. Furthermore, if S and T are tableaux or subtableaux, and there is a translation θ that maps (bijectively) the support of S to the support of T, then we say that S and T are **identically ordered** if for all x, y in the support of S we have $S(x) < S(y)$ if and only if $T(\theta(x)) < T(\theta(y))$. Figure 4 illustrates the above notations.

Theorem 1 will be deduced from the following.

Theorem 2. *Let T be any standard staircase-shape Young tableau of size k. For some positive constants c_1', c_2' and c_3' (depending only on k), with probability at least $1 - e^{-c_3' n}$, a uniformly random standard staircase-shape Young tableau of size $n \geq k$ contains at least $c_1' n$ subtableaux with pairwise disjoint supports such that:*

(1) each is identically ordered with T;
(2) all their entries are greater than $N - c_2 n$.

As an application of Theorem 1 we prove that a uniformly random sorting network is not geometrically realizable in the following sense. Consider a set X of n points in $\mathbb{R}^2$ such that no two points from X lie on the same vertical line, no three points are collinear, and no two pairs of points define parallel lines. Label the points $1, \ldots, n$ from left to right (i.e. in order of their first coordinate). Let X_ϕ be the set obtained by rotating $\mathbb{R}^2$ by angle ϕ about the origin, and let σ_ϕ be the permutation found by reading the labels in X_ϕ from left to right. As ϕ increases from 0 to π, the permutation σ_ϕ changes via a sequence of swaps, which form a sorting network. Any sorting network that can be generated in this way is called **geometrically realizable**. (Such networks were called *stretchable* in [AHRV], but this term is used with a different meaning in [GR, GP]).

Goodman and Pollack [GP] gave an example of a sorting network of size 5 that is not geometrically realizable. On the other hand, in [AHRV], it was conjectured (on the basis of strong experimental and heuristic evidence) that a uniformly random sorting network is with high probability *approximately* geometrically realizable, in the sense that its distance to some random geometrically realizable network tends to zero in probability (in a certain natural metric). The conjectures of [AHRV] would also imply that, for fixed m, the sorting network obtained by observing only m randomly chosen particles from a uniformly random sorting network of size $n \geq m$ is with high probability geometrically realizable as $n \to \infty$. (The conjectures also imply that these size-m networks have a limiting distribution as $n \to \infty$, as well as providing a precise description of the limit. Certain aspects of the latter prediction were verified rigorously in [AH].) However, we prove that with high probability a uniformly random sorting network is *not* itself geometrically realizable.

Theorem 3. *The probability that a uniformly random sorting network of size n is geometrically realizable tends to zero as n tends to infinity.*

While our proof yields an exponential (in n) bound on the probability that a uniform sorting network of size n is geometrically realizable, we believe the probability is even $O(e^{-cn^2})$.

The paper is organized as follows. In Section 2 we recall basic definitions and the Edelman-Greene bijection between sorting networks and standard Young tableaux. In Sections 3 and 4 we prove some auxiliary lemmas about Young tableaux and sequences of random variables, respectively. In Section 5 we prove Theorem 2 and then deduce Theorem 1 as a corollary. Finally, in Section 6 we prove Theorem 3.

2. Sorting networks and Young tableaux

Edelman and Greene [EG] introduced a bijection between sorting networks of size n and standard staircase-shape Young tableaux of size n, i.e. of shape $(n-1, n-2, \dots, 1)$. We describe it in a slightly modified version that is more convenient for us.

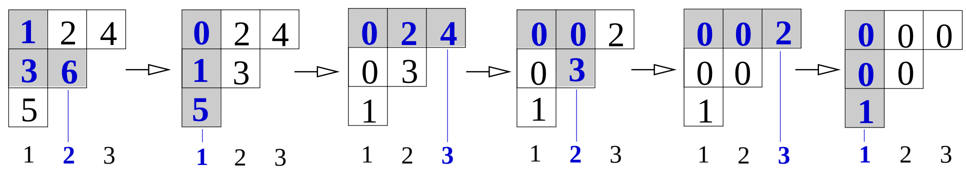

FIGURE 5. A standard staircase-shape Young tableau, sliding paths (shaded) and the sequence of tableaux in the Edelman–Greene bijection. Here $n = 4$ and the corresponding sorting network is $(2, 1, 3, 2, 3, 1)$. Vertical lines show the correspondence between the positions of maximal entries in the tableaux and numbers $s_1, \dots, s_N$ of the sorting network.

Given a standard staircase-shape Young tableaux T of size n, we construct a sequence of integers $s_1, \dots, s_N$ as follows. Set $T_1 = T$ and repeat the following for $t = 1, 2, \dots, N$.

(1) Let $x = (n - j, j)$ be the location of the maximal entry in the tableau T_t. Set $s_t = j$.
(2) Compute the sliding path, which is a sequence $x_1, x_2, \dots, x_\ell$, such that $x_1 = x$ and for $i = 1, 2, \dots$ we define x_{i+1} to be the box among $\{x_i - (1, 0), x_i - (0, 1)\}$ with larger entry in T_t, with the convention that $T_t(x) = 0$ for every x outside the staircase Young diagram of size n. Let ℓ be the minimal i such that $T_t(x_i) = 0$.
(3) Perform the sliding, i.e. define the tableau T_{t+1} as follows. Set $T_{t+1}(x_i) = T_t(x_{i+1})$ for $i = 1, \dots, \ell - 1$ and set $T_{t+1}(y) = T_t(y)$ for all boxes y of the staircase Young diagram of size n not belonging to $\{x_1, \dots, x_{\ell-1}\}$.

An example of this procedure is shown in Figure 5. Edelman and Greene [EG] proved that the resulting sequence of numbers is indeed a sorting network, and furthermore that the algorithm provides a bijection between standard staircase-shape Young tableaux and sorting networks.

Now we fix n, consider the set of all sorting networks of this size and equip it with the uniform measure. The Edelman–Greene bijection maps this measure to the uniform measure on the set of all standard staircase-shape Young tableaux of size n.

Given a standard Young tableau T of shape λ with $|\lambda| = M$ we define a sequence of Young diagrams by

$$\lambda^i = \{x \in \lambda : T(x) \le M - i\}.$$

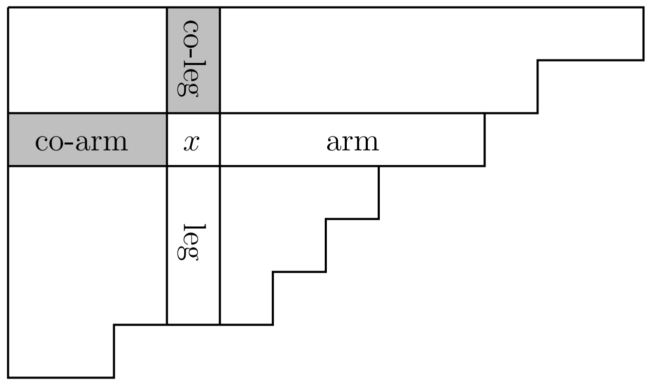

FIGURE 6. Hook (clear) and co-hook (shaded) in a Young diagram.

Thus $\lambda = \lambda^0 \supset \lambda^1 \supset \cdots \supset \lambda^M = \emptyset$, and $\lambda^i \setminus \lambda^{i+1}$ consists of the single box $T^{-1}(|\lambda| - i)$. If T is a uniformly random standard Young tableau of shape λ, then conditional on $\lambda^i, \lambda^{i-1}, \ldots, \lambda^0$, the restriction of T to λ^i is uniformly random. Thus the sequence of diagrams described above is a Markov chain.

3. Some properties of Young tableaux

In this section we present a fundamental result about Young diagrams (the hook formula) and deduce some of its consequences.

When drawing pictures of Young diagrams we adopt as usual the convention that the first coordinate i (the row index) increases downwards while the second coordinate j (the column index) increases from left to right. Given a Young diagram λ, its **transposed diagram** λ' is obtained by reflecting λ with respect to diagonal $i = j$. The column lengths of λ are the row lengths of λ'.

For any box $x = (i, j)$ of a Young diagram λ, its **arm** is the collection of $\lambda_i - j$ boxes to its right: $\{(i, j') \in \lambda : j' > j\}$. The **leg** of x is the set $\{(i', j) \in \lambda : i' > i\}$ of $\lambda'_j - i$ boxes below it. The union of the box x, its arm and its leg is called the **hook** of x. The number of boxes in the hook is called the **hook length** and is denoted by $h(x)$. The **co-arm** is the set $\{(i, j') \in \lambda : j' < j\}$; the **co-leg** is the set $\{(i', j) \in \lambda : i' < i\}$, and their union (which does not include x) is called the **co-hook** and denoted by $\mathcal{C}(x)$. See Figure 6. Finally, a **corner** of a Young diagram λ is a box $x \in \lambda$ such that $h(x) = 1$, or equivalently such that $\lambda \setminus \{x\}$ is also a Young diagram.

The dimension $\dim(\lambda)$ of a Young diagram λ is defined as the number of standard Young tableaux of shape λ (thus named because it is the dimension of the corresponding irreducible representations of the symmetric group).

Lemma 4 (Hook formula; [FRT])**.** *The dimension* $\dim(\lambda)$ *satisfies*

$$\dim(\lambda) = \frac{|\lambda|!}{\prod_{x\in\lambda} h(x)}.$$

See e.g. [FRT] or [M] for a proof.

Corollary 5. *Let T be a uniformly random standard Young tableau of shape λ, and let x be a corner of λ. The location $T^{-1}(|\lambda|)$ of the largest entry is distributed as follows.*

$$\mathbb{P}\big(T^{-1}(|\lambda|) = x\big) = \frac{\dim(\lambda \setminus \{x\})}{\dim(\lambda)} = \frac{1}{|\lambda|}\prod_{z\in\mathcal{C}(x)} \frac{h(z)}{h(z)-1}.$$

(Note that $h(z) > 1$ for any box in the co-hook $\mathcal{C}(x)$, so the right side is finite.)

Proof. This is immediate from Lemma 4. □

Lemma 6. *Fix $\ell > 0$. Let a Young diagram λ be a subset of the staircase Young diagram of size n, and let $x = (i, j)$ be a corner of λ with $i, j \geq n/3$ and $n - i - j \leq \ell$. Let T be a uniformly random standard Young tableau of shape λ. We have*

$$\mathbb{P}\big(T(x) = |\lambda|\big) \geq \frac{c}{n},$$

where c is a constant depending only on ℓ.

There is nothing special about the bound $\frac{n}{3}$ on i, j – the lemma and proof hold as long as $i, j \geq \varepsilon n$, though the constant in the resulting bound tends to 0 as $\varepsilon \to 0$.

Proof of Lemma 6. The box $(i - k, j)$ of the co-hook has hook length $\lambda_{i-k} - j + k + 1 \leq n - i - j + 2k + 1 \leq \ell + 2k + 1$. Similarly the box $(i, j - k)$ has hook length at most $\ell + 2k + 1$. It follows that

$$\begin{aligned}\mathbb{P}(T^{-1}(|\lambda|) = x) &= \frac{1}{|\lambda|}\prod_{k<i}\frac{h(k,j)}{h(k,j)-1}\prod_{k<j}\frac{h(i,k)}{h(i,k)-1} \\ &\geq \frac{1}{n^2}\bigg(\prod_{k<n/3}\frac{\ell+2k+1}{\ell+2k}\bigg)^2.\end{aligned}$$

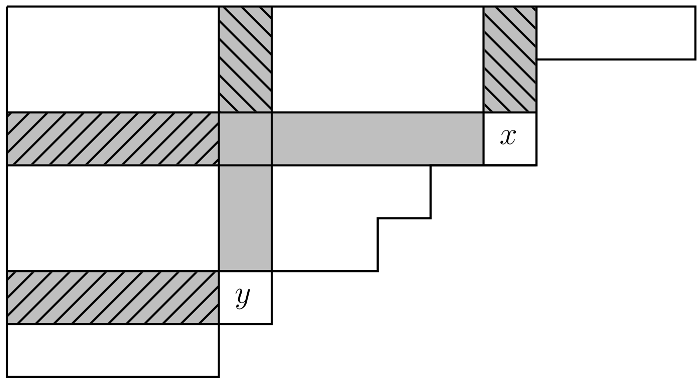

FIGURE 7. The co-hooks $\mathcal{C}(x)$ and $\mathcal{C}(y)$ (shaded) and the matched parts of the co-arms and co-legs (hatched in matching directions).

(Here we used that the factors are all decreasing in h, greater than 1, and that $i, j \geq n/3$.) It is now easy to estimate

$$\left(\prod_{k<n/3} \frac{\ell+2k+1}{\ell+2k}\right)^2 \geq \left(\prod_{k<n/3} \frac{\ell+2k+1}{\ell+2k}\right)\left(\prod_{k<n/3} \frac{\ell+2k+2}{\ell+2k+1}\right) = \frac{\ell+2\lfloor n/3\rfloor+2}{\ell+2} > cn$$

for some $c = c(\ell)$. □

Lemma 7. *Let T be a uniformly random standard Young tableau of shape λ, let x and y be two corners of λ and $\ell = \|x-y\|_\infty$. Then*

$$\frac{\mathbb{P}(T^{-1}(|\lambda|) = x)}{\mathbb{P}(T^{-1}(|\lambda|) = y)} \leq (\ell+1)(2\ell+1).$$

For our application all we need is a bound of the form $C(\ell)$ on this ratio, though we note that the bound we get is close to optimal for a tableau of shape $(n+1, n, \ldots, n)$ with $\ell+1$ rows, for large n.

Proof of Lemma 7. To compare the expressions from Corollary 5 for x and y, let us introduce a partial matching between $\mathcal{C}(x)$ and $\mathcal{C}(y)$. We match boxes of the co-arm of x and the co-arm of y if they are in the same column. We match boxes of the co-leg of x and the co-leg of y if they are in the same row. All boxes inside the rectangle with opposite vertices x and y remain unmatched (see Figure 7).

Writing $x = (i_1, j_1)$ and $y = (i_2, j_2)$ without loss of generality assume that $i_1 < i_2$ and $j_1 > j_2$. Clearly, if $z \in \mathcal{C}(x)$ and $z' \in \mathcal{C}(y)$ are a matched pair, then $h(z') = h(z) \pm s$, where $s = i_2 - i_1 + j_1 - j_2$ and the sign is plus if the box z belongs to the co-leg of x and minus otherwise. Let $M(x)$, $U(x)$ be the matched and unmatched parts of the co-hooks and similarly for y. We have

$$\frac{\mathbb{P}(T^{-1}(|\lambda|)=x)}{\mathbb{P}(T^{-1}(|\lambda|)=y)} = \frac{\prod_{z\in U(x)} \frac{h(z)}{h(z)-1}}{\prod_{z\in U(y)} \frac{h(z)}{h(z)-1}} \times \prod_{z\in M(x)} \frac{\left(\frac{h(z)}{h(z)-1}\right)}{\left(\frac{h(z)\pm s}{h(z)-1\pm s}\right)}, \tag{1}$$

where the choice of the sign $\pm$ depends on whether a box z belongs to the co-arm or the co-leg of x.

Let us bound the right side of (1). First note that all the boxes in the co-arm of x and all the boxes in the co-leg of y are matched. The product over $z \in U(y)$ is at least 1. Next, there are at most ℓ unmatched boxes of the co-arm of x and their hook lengths are distinct. Consequently

$$\prod_{z\in U(x)} \frac{h(z)}{h(z)-1} \le \prod_{m=2}^{\ell+1} \frac{m}{m-1} = \ell+1.$$

Turning to the last product in (1), a matched pair of boxes from the co-arms contributes to (1) the factor

$$\frac{\left(\frac{h(z)}{h(z)-1}\right)}{\left(\frac{h(z)-s}{h(z)-1-s}\right)},$$

which is easily seen to be less than 1.

Finally, every matched pair of boxes from the co-legs contributes to (1) the factor

$$\frac{\left(\frac{h(z)}{h(z)-1}\right)}{\left(\frac{h(z)+s}{h(z)-1+s}\right)} = 1 + \frac{s}{(h(z)-1)(h(z)+s)},$$

This is greater than 1 for any $h(z)$. As z varies over a co-leg of x, the values of $h(z)$ are distinct. Consequently, the contribution from the matched boxes from the co-legs is bounded from above by

$$\prod_{m=2}^{\infty} \frac{\left(\frac{m}{m-1}\right)}{\left(\frac{m+s}{m-1+s}\right)} = \lim_{r\to\infty} \frac{r}{r+s}\,(s+1) = s+1 \le 2\ell+1.$$

Multiplying all the aforementioned inequalities we get the required estimate. □

4. Sequences of random variables

Recall that a real-valued random variable Y **stochastically dominates** another real-valued random variable Z if and only if there exist a probability space Ω and two random variables $\widetilde{Y}, \widetilde{Z}$ defined on Ω, such that $\widetilde{Y} \overset{d}{=} Y$ and $\widetilde{Z} \overset{d}{=} Z$, and $\widetilde{Y} \geq \widetilde{Z}$ almost surely.

Lemma 8. *Let $X_1, \dots, X_N$ be random variables taking values in $\{1, \dots, m, \infty\}$ such that a.s. each $a \in [1, m]$ appears exactly r times. Let A_i be events, and define the filtration $\mathcal{F}_i = \sigma(X_1, \dots, X_i, A_1, \dots, A_{i-1})$. Assume $\mathbb{P}(A_i \mid \mathcal{F}_i) \geq p$ a.s. for some $p > 0$ and all i. Let G_a be the event*

$$G_a = \bigcap_{i=1}^{N} \big(\{X_i \neq a\} \cup A_i\big),$$

that is that A_i occurs whenever $X_i = a$. Then $\sum_{a=1}^{m} 1_{G_a}$ stochastically dominates the binomial random variable $\mathrm{Bin}(m, p^r)$.

To clarify the lemma, it helps to think of having m counters initialized at 0. At each step, a counter is selected (or no counter, signified by $X_i = \infty$), and that counter is advanced with probability at least p. The event G_a is that the ath counter is advanced every time it is selected. Then after every counter has been selected r times, the number of counters with the highest possible value r stochastically dominates a $\mathrm{Bin}(m, p^r)$ random variable. Note that the order in which counters are selected may depend arbitrarily on the past selections and advances. While this lemma seems intuitively clear and perhaps even obvious, the precise assumptions on the dependencies among the events and variables make the proof slightly delicate.

Proof of Lemma 8. First, we want to extend the probability space, and define events $A'_i \subseteq A_i$ and a finer filtration $\mathcal{F}'_i$ in such a way that $\mathbb{P}(A'_i \mid \mathcal{F}'_i) = p$ for all i.

Let Ω be our original probability space and let μ be our original probability measure. For $i = 1, 2, \dots, N$ let $\mathcal{E}^i$ be the set of all elementary events in the finite σ–algebra $\mathcal{F}_i$ that have non-zero probabilities (with respect to μ). For any $E \in \mathcal{E}^i$ let Ω_i^E denote the probability space $\{0, 1\}$ with probability measure μ_i^E such that $\mu_i^E(1) = p/\mathbb{P}(A_i \mid E)$. Our new probability space Ω' is the product of Ω and all Ω_i^E:

$$\Omega' = \Omega \times \prod_{i=1}^{N} \prod_{E \in \mathcal{E}^i} \Omega_i^E.$$

In other words, an element of Ω' is a pair (ω, f), where $\omega \in \Omega$ and f is a function from $\bigsqcup_i \mathcal{E}^i$ to $\{0,1\}$ (here $\bigsqcup$ denotes set-theoretic disjoint union, so $\bigsqcup_i \mathcal{E}^i := \cup_i \{(E,i) : E \in \mathcal{E}_i\}$). We equip Ω' with the probability measure μ' which is the direct product of μ and the measures μ_i^E:

$$\mu' = \mu \times \prod_{i=1}^{N} \prod_{E \in \mathcal{E}^i} \mu_i^E.$$

In what follows we do not distinguish between a random variable $X(\omega)$ defined on Ω and the random variable $X(\omega, f) := X(\omega)$ defined on Ω'. In the same way we identify any event A of Ω with $\widetilde{A} := \{(\omega, f) \in \Omega' : \omega \in A\} \subseteq \Omega'$. In what follows all the probabilities are understood with respect to μ'.

For any $E \in \bigsqcup_i \mathcal{E}^i$ let f^E denote the random variable on Ω' given by

$$f^E(\omega, f) = f(E).$$

Now for any $E \in \mathcal{E}^i \subseteq \bigsqcup_j \mathcal{E}^j$ set

$$B_i^E := \{(\omega, f) \in \Omega' \mid \omega \in E,\ f(E) = 1\} = E \cap \{f^E = 1\}.$$

Denote

$$B_{(i)} = \bigcup_{E \in \mathcal{E}^i} B_i^E$$

and let $A_i' = A_i \cap B_{(i)}$.

Let us introduce a filtration on Ω':

$$\mathcal{F}_i' = \sigma\Big(X_1, \dots, X_i, A_1, \dots, A_{i-1}, \{f^E\}\Big),$$

where E runs over all elements of $\bigsqcup_{j=1}^{i-1} \mathcal{E}^j$.

Note that $A_i' \in \mathcal{F}_{i+1}'$. We claim that $\mathbb{P}(A_i' \mid \mathcal{F}_i') = p$ for every i. Indeed, this immediately follows from the definition of A_i' and the fact that A_i' is independent of all f^E for $E \in \bigsqcup_{j=1}^{i-1} \mathcal{E}^i$.

Moreover, consider any sequence of stopping times $1 \le \tau_1 < \cdots < \tau_\ell \le N$ (w.r.t. the filtration $\mathcal{F}'$). We claim that $\mathbb{P}\left(\bigcap_{i \le \ell} A_{\tau_i}'\right) = p^\ell$. The proof is a simple induction in ℓ. For $\ell = 1$ we have

$$\mathbb{P}(A'_{\tau_1}) = \sum_{i=1}^{N} \mathbb{P}(A'_i \cap \{\tau_1 = i\})$$
$$= \sum_{i=1}^{N} \mathbb{P}(\tau_1 = i)\mathbb{P}(A'_i \mid \tau_1 = i) \overset{(*)}{=} \sum_{i=1}^{N} \mathbb{P}(\tau_1 = i) \cdot p = p,$$

where in the equality $(*)$ we used that $\mathbb{P}(A'_i \mid \mathcal{F}'_i) = p$ and $\{\tau_1 = i\} \in \mathcal{F}'_i$. Now assume that our statement is true for $\ell = h - 1$. Then for $\ell = h$ we have

$$\mathbb{P}\left(\bigcap_{i=1}^{h} A'_{\tau_i}\right) = \sum_{j=1}^{N} \mathbb{P}(\tau_1 = j)\mathbb{P}(A'_j \mid \tau_1 = j)\mathbb{P}\left(\bigcap_{i=2}^{h} A'_{\tau_i} \mid A'_j \cap \{\tau_1 = j\}\right).$$

Note that for $i \geq 2$ the restriction of τ_i on the set $A'_j \cap \{\tau_1 = j\}$ is again a stopping time. Indeed, by the definition, $j < \tau_i \leq N$ on $\{\tau_1 = j\}$, and for $k > j$ we have $\{\tau_i \leq k\} \cap A'_j \cap \{\tau_1 = j\} \in \mathcal{F}'_k$, since both $\{\tau_i \leq k\} \in \mathcal{F}'_k$ and $A'_j \in \mathcal{F}'_k$ and $\{\tau_1 = j\} \in \mathcal{F}'_k$. Therefore, using the induction assumption we conclude that if $\mathbb{P}(A'_j \cap \{\tau_1 = j\}) > 0$, then $\mathbb{P}(\bigcap_{i=2}^{h} A'_{\tau_i} \mid A'_j \cap \{\tau_1 = j\}) = p^{h-1}$. Hence,

$$\mathbb{P}\left(\bigcap_{i=1}^{h} A'_{\tau_i}\right) = \sum_{j=1}^{N} \mathbb{P}(\tau_1 = j)\mathbb{P}(A'_j \mid \tau_1 = j)p^{h-1} = \sum_{j=1}^{N} \mathbb{P}(\tau_1 = j)p^h = p^h.$$

Now, let

$$G'_a = \bigcap_{i=1}^{N} \big(\{X_i \neq a\} \cup A'_i\big) \subseteq G_a.$$

Applying the above claim to the r ordered stopping times τ_i defined by

$$\{\tau_1, \dots, \tau_r\} = \{k : X_k = a\}$$

we find $\mathbb{P}(G'_a) = p^r$. Moreover, for any set $S \subseteq [1, m]$, by taking the $r|S|$ ordered stopping times τ_i^S defined by

$$\{\tau_1, \dots, \tau_{r|S|}\} = \{k : X_k \in S\}$$

we find

$$\mathbb{P}\left(\bigcap_{a \in S} G'_a\right) = p^{r|S|}.$$

It follows that the events G'_a are independent, and so

$$\sum_{a=1}^{m} 1_{G_a} \geq \sum_{a=1}^{m} 1_{G'_a} \overset{d}{=} \mathrm{Bin}(m, p^r). \qquad \square$$

Lemma 9. *Let $X_1, \dots, X_N$ be random variables taking values in $\{1, \dots, m, \infty\}$ such that a.s. each $a \in [1, m]$ appears exactly r times. Denote $S_k(a) := \#\{i \le k : X_i = a\}$. Let $\widehat{\mathcal{F}}_k = \sigma(X_1, \dots, X_k)$, and suppose moreover that for some $c > 0$ and all a, k, on the event $S_k(a) < r$ (which lies in $\widehat{\mathcal{F}}_k$), we have*

$$\mathbb{P}\big(X_{k+1} = a \mid \widehat{\mathcal{F}}_k\big) > \frac{c}{m}.$$

Finally, let $D_k = \#\{a : S_k(a) = r\}$. Then for every $\varepsilon > 0$ there are constants c_1, c_2, depending on c, r but not on m or N, such that

$$\mathbb{P}\big(D_{c_1 m} \le (1-\varepsilon)m\big) < e^{-c_2 m}.$$

Proof. Let $T_k = \sum_{a=1}^m S_k(a)$, and note that $T_k > mr - \varepsilon m$ implies $D_k > (1-\varepsilon)m$.

On the event $D_k \le (1-\varepsilon)m$ we have $\mathbb{E}(T_{k+1} \mid \widehat{\mathcal{F}}_k) - T_k \ge c\varepsilon$. Let M_k be $c\varepsilon k - T_k$, stopped when D_k exceeds $(1-\varepsilon)m$, then we see that M_k is a supermartingale with bounded increments. By the Azuma-Hoeffding inequality for supermartingales (which follows from the martingale version by Doob decomposition; see e.g. [Az] or [W, E14.2 and 12.11]), for any $c_1 > 0$ there is a c_2 so that $\mathbb{P}(M_{c_1 m} > m) \le e^{-c_2 m}$.

If $M_{c_1 m} \le m$ and M is not yet stopped at time $c_1 m$, then $T_{c_1 m} \ge (c\varepsilon c_1 m - 1)m$. If c_1 is such that $c\varepsilon c_1 m - 1 > r$, this cannot hold, so M is stopped by time $c_1 m$ with probability at least $1 - e^{-c_2 m}$. □

Corollary 10. *Let X_i, A_i for $i = 1, \dots, N$ be two random sequences satisfying the assumptions of both Lemmas 8 and 9. Let $\widehat{G}(a, i)$ be the intersection of the events G_a and $\{S_i(a) = r\}$, i.e.*

$$\widehat{G}(a,i) = \{S_i(a) = r\} \cap \bigcap_{j=1}^{N} \big(\{X_j \ne a\} \cup A_j\big).$$

Set $\widehat{Q}(i) = \sum_a 1_{\widehat{G}(a,i)}$. There exist positive constants c_1, c_2, c_3 (which depend on r, p, c, but not on m, N) such that $\mathbb{P}\big(\widehat{Q}(c_1 m) > c_2 m\big) > 1 - e^{-c_3 m}$.

If we again think about m counters, then the corollary means simply that after time $c_1 m$, with probability at least $1 - e^{-c_3 m}$, at least $c_2 m$ counters will have advanced r times.

Proof of Corollary 10. Denote $Q = \sum_a 1_{G_a}$. Lemma 8 implies that Q stochastically dominates a binomial random variable. Thus, by a standard large deviation estimate (see e.g. [K, Chapter 27]), for some positive constants c_4, c_5 we have

$$\mathbb{P}(Q > c_4 m) > 1 - e^{-c_5 m}.$$

Take $\varepsilon = c_4/2$ in Lemma 9. It follows that for some c_1 with probability at least $1 - e^{-c_6 m}$ random variable $\widehat{Q}(c_1 m)$ differs from Q by not more than $c_4 m/2$. Thus,

$$\mathbb{P}\Big(\widehat{Q}(c_1 m) > c_4 m/2\Big) > 1 - e^{-c_3 m}. \qquad \square$$

5. Proofs of the main results

We are now ready to prove Theorems 1 and 2. We denote by $\widehat{T}$ a standard staircase-shape Young tableau of size k and by T a uniformly random standard staircase-shape Young tableau of size n. In what follows k and $\widehat{T}$ are fixed while n tends to infinity. Given $\widehat{T}$, the idea is to consider cn specific disjointly supported subtableaux of T in columns $\lfloor n/3 \rfloor, \ldots, \lfloor 2n/3 \rfloor$ and show that linearly many (in n) of them are identically ordered with $\widehat{T}$.

Proof of Theorem 2. Within the staircase Young diagram λ of size n we fix $m := \lfloor n/(3k-3) \rfloor$ disjoint subdiagrams $K_1, \ldots, K_m$ of λ in columns $\lfloor n/3 \rfloor, \ldots, \lfloor 2n/3 \rfloor$, each a translation of a staircase Young diagram of size k. Let θ_i be the translation mapping K_i to the staircase Young diagram.

Let $N := \binom{n}{2}$ and $r := \binom{k}{2}$. We now construct sequences X_t and A_t to which we shall apply Lemmas 8 and 9, as random variables on the probability space of standard staircase-shape Young tableaux T of size n with uniform measure. Set $X_t = a$ if $T^{-1}(N+1-t)$ belongs to K_a and set $X_t = \infty$ if $T^{-1}(N+1-t)$ does not belong to $\bigcup_a K_a$. Note that each $a \in \{1, \ldots, m\}$ appears exactly r times among $X_1, \ldots, X_N$.

Next, we define the events A_t. If $X_t = \infty$ then A_t occurs. Otherwise, let $a = X_t$ and suppose X_t is the ith occurrence of a among $X_1, \ldots, X_t$ (or equivalently, $N-t+1$ is the ith largest entry in K_a). If there is any $s < t$ with $X_s = a$ for which A_s does not occur, then A_t does occur. Finally, if the box $T^{-1}(N-t+1)$ is in the same position within K_a as $\widehat{T}^{-1}(r-i+1)$ is within the staircase Young diagram of size k (in other words, if $\theta_a(T^{-1}(N-t+1)) = \widehat{T}^{-1}(r-i+1)$), then A_t occurs. If it is not in the same position, then A_t does not occur. In other words, A_t fails to occur precisely if for some a, number t is the minimal number such that the locations of entries $\{N-t+1, \ldots, N\}$ imply that the subtableau supported by K_a and $\widehat{T}$ are not identically ordered.

Rephrasing in terms of counters, we do the following. Recall that a uniformly random standard staircase-shape Young tableau T is associated with a Markov chain of decreasing Young diagrams λ^t. Each step of this Markov chain is a removal of a box from a Young diagram. If the box x removed at step t belongs to K_h, then we choose the hth counter at this step. This counter advances if either the position of x is the correct one (as dictated by the order in $\widehat{T}$), or if the correct order of the entries of T inside K_h was already broken before tth step. Clearly, if the hth counter advances r times, then the subtableau of T with support K_h is identically ordered with $\widehat{T}$.

Let us check that the sequences X_t and A_t, and the numbers r, m, N, satisfy the conditions of Lemma 8 with

$$p = \frac{1}{2k^3}.$$

As noted, every $a \in \{1, \dots, m\}$ appears among $X_1, \dots, X_N$ exactly r times. Thus, it remains to bound from below the conditional probabilities of A_t. Define $\mathcal{F}_t$ as in Lemma 8 and let W be an elementary event of $\mathcal{F}_t$. We must prove that $\mathbb{P}(A_t \mid W) \geq p$. If $X_t = \infty$ on W, then $\mathbb{P}(A_t \mid W) = 1 \geq p$. If some previous A_s with $s < t$ and $X_s = X_t$ did not occur (on W) then again $\mathbb{P}(A_t \mid W) = 1 \geq p$.

In the remaining case, whether or not T belongs to A_t depends on the position of the box $T^{-1}(N - t + 1)$; specifically, T belongs to A_t if this box is the correct one according to $\widehat{T}$ of the possible boxes in the subdiagram K_{X_t}. Let a denote the value of X_t on W. The Markov property of the sequence λ^t implies that

$$\begin{aligned} \mathbb{P}(A_t \mid W) &= \sum_{\mu} \mathbb{P}(\lambda^{t-1} = \mu \mid W) \cdot \mathbb{P}(A_t \mid \lambda^{t-1} = \mu, W) \\ &= \sum_{\mu} \mathbb{P}(\lambda^{t-1} = \mu \mid W) \cdot \mathbb{P}(A_t \mid X_t = a, \lambda^{t-1} = \mu), \end{aligned} \tag{2}$$

where the sum is taken over all Young diagrams μ with $|\mu| = N - t + 1$ boxes that are contained in the staircase Young diagram of size n.

Let us bound $\mathbb{P}(A_t \mid X_t = a, \lambda^{t-1} = \mu)$ from below. The condition $X_t = a$ means that the box $T^{-1}(N - t + 1)$ is situated in the subdiagram K_a. Thus, given $X_t = a$ and $\lambda^{t-1} = \mu$, there are at most k possible positions for the box $T^{-1}(N + t - 1)$. Lemma 7 implies that the conditional probabilities of different positions differ at most by a factor of $2k^2$ (since the parameter ℓ in that lemma is at most $k - 2$). Consequently, the conditional probability of each position is at least $1/(2k^3)$. Exactly one of the positions corresponds to the event A_t. We conclude that

$$\mathbb{P}(A_t \mid X_t = b,\, \lambda^{t-1} = \mu) \geq \frac{1}{2k^3}$$

Hence, (2) gives

$$\mathbb{P}(A_t \mid W) \geq \frac{1}{2k^3} \sum_{\mu} \mathbb{P}(\lambda^{t-1} = \mu \mid W) = \frac{1}{2k^3}.$$

Finally, let us check that the sequence X_t satisfies the conditions of Lemma 9. Note that the condition $S_t(a) < r$ means that the subdiagram K_a is not filled by boxes $T^{-1}(N - s + 1)$ with $s \leq t$. Thus, $S_t(a) < r$ if and only if $\lambda^t \cap K_a \neq \emptyset$. Note that the event $\{S_t(a) < r\}$ belongs to $\widehat{\mathcal{F}}_t$. Let V be an elementary event from $\widehat{\mathcal{F}}_t$ such $\{S_t(a) < r\}$ on V. Using the Markov property of the sequence λ^i we obtain:

$$\begin{aligned} \mathbb{P}(X_{t+1} = a \mid V) &= \sum_{\mu} \mathbb{P}(\lambda^t = \mu \mid V)\, \mathbb{P}(X_{t+1} = a \mid \lambda^t = \mu, V) \\ &= \sum_{\mu} \mathbb{P}(\lambda^t = \mu \mid V)\, \mathbb{P}(X_{t+1} = a \mid \lambda^t = \mu), \end{aligned}$$

where the sum is taken over the set of all Young diagrams μ with $|\mu| = N - t$ boxes that are subsets of the staircase Young diagram of size n. Since $V \subseteq \{S_t(a) < r\}$, we have $\mathbb{P}(\lambda^t = \mu \mid V) \neq 0$ only for μ such that $\mu \cap K_a$ is non-empty.

Consequently, in order to prove that

$$\mathbb{P}(X_{t+1} = a \mid V) > \frac{c}{m}$$

for some positive constant c, it suffices to show that

$$\mathbb{P}(X_{t+1} = a \mid \lambda^t = \mu) > \frac{c}{m}$$

for any Young diagram μ contained in the staircase Young diagram of size n and such that $\mu \cap K_a$ is non-empty. Any such diagram μ has at least one corner inside K_a. Applying Lemma 6 for μ and this corner yields the required bound.

Applying Corollary 10 to the sequences X_t and A_t we get the statement of Theorem 2. □

We now deduce Theorem 1 using the Edelman-Greene bijection.

Proposition 11. *Fix any pattern γ of size k. There exist constants c_3, c_4 and c_5 (depending on γ) such that for every $n \geq k$, the pattern γ occurs at least $c_3 n$ times within the time interval $[1, c_4 n]$ of a uniformly random sorting network of size n with probability at least $1 - e^{-c_5 n}$.*

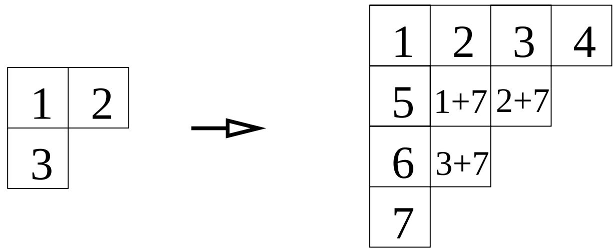

FIGURE 8. "Padding" a tableau T_γ to get $\widehat{T}$. Here $k = 3$.

Note that Proposition 11 differs from Theorem 1 in that we consider only the beginning of the network and hence only find a linear number of occurrences of γ.

Proof of Proposition 11. Clearly, it suffices to prove Proposition 11 for patterns of the maximum length $k(k-1)/2$, in other words a sorting network of size k. Such a pattern $\gamma = (\gamma_1, \dots, \gamma_{k(k-1)/2})$ corresponds via the Edelman-Greene bijection to some standard staircase-shape Young tableau T_γ of size k. Consider a larger standard staircase-shape Young tableau $\widehat{T}$ of size $k+2$, where entries of the hook of $(1,1)$ are the numbers $1, \dots, 2k+1$ (in an arbitrary admissible order) and the remaining staircase-shaped Young tableau of size $k-1$ contains $2k+2, \dots, (k+2)(k+1)/2$ and is identically ordered with T_γ. An example of this construction is shown in Figure 8.

Let c_3, c_4 and c_5 be the constants c_1', c_2' and c_3' of Theorem 2, respectively. Let T be a standard staircase-shape Young tableau of size n having at least $c_3 n$ disjointly supported subtableaux identically ordered with $\widehat{T}$, furthermore, all the entries of these subtableaux are greater than $N - c_4 n$. (Theorem 2 implies that a uniformly random standard staircase-shape Young tableau of size $n \geq k$ is of this kind with probability at least $1 - e^{-c_5 n}$.) Suppose that the support of the ℓth such subtableau ($\ell = 1, 2, \dots, c_3 n$) is a subdiagram K_ℓ with top-left corner $(n - j_\ell - k, j_\ell)$. Let K_ℓ' denote the subdiagram with top-left corner $(n - j_\ell - k + 1, j_\ell + 1)$ and note that the subtableau with support K_ℓ' is identically ordered with T_γ.

Let ω be the sorting network corresponding to T via the Edelman-Greene bijection. Note that in the Edelman-Greene bijection, every tableau entry moves towards the boundary of the staircase Young diagram until it becomes the maximal entry in the tableau, and then it disappears and adds to the sorting network a swap in position j, where j is the column of the entry just before it disappeared. It follows that all the entries starting in K_ℓ disappear in the columns $j_\ell, \dots, j_\ell+k$ and, thus, add to the sorting network swaps s_i satisfying $j_\ell \le s_i \le j_\ell + k$. Furthermore, observe that all the entries starting in K'_ℓ disappear (in columns s_i satisfying $j_\ell < s_i < j_\ell + k$) before the entries in $K_\ell \setminus K'_\ell$. Finally, note that until all entries starting in K'_ℓ disappeared no other entry can disappear in columns $j_\ell, \dots, j_\ell + k$.

We conclude that for every ℓ, the pattern γ occurs in ω at $[1, t_\ell] \times [j_\ell+1, j_\ell+k-1]$. Thus, pattern γ occurs in ω at least $c_3 n$ times within the time interval $[1, c_4 n]$. □

Proof of Theorem 1. Let c_3, c_4, c_5 be the constants from Proposition 11, and let $m := \lceil c_4 n \rceil$. For $t = 1, \dots, \lfloor N/m \rfloor$ let I_t be the set of all sorting networks ω of size n such that γ occurs in ω at least $c_3 n$ times within the time interval $[(t-1)m+1, tm]$. Proposition 11 yields that $\mathbb{P}(I_1) \ge 1 - e^{-c_5 n}$.

A uniformly random sorting network $(s_1, s_2 \dots, s_N)$ is stationary in the sense that $(s_1, \dots, s_{N-1})$ and $(s_2, \dots, s_N)$ have the same distributions (see [AHRV, Theorem 1]). Thus $\mathbb{P}(I_t)$ does not depend on t.

There exist constants $c_6 > 0$ and n_0 such that if $n > n_0$, then $\lfloor N/m \rfloor e^{-c_5 n} \le e^{-c_6 n}$. Let $c_1 = \min(\frac{c_3}{4c_4}, \frac{c_3}{n_0})$ and $c_2 = \min(c_5, c_6)$. Let I denote the set of all sorting networks ω of size n such that γ occurs $c_1 n^2$ times in ω. If $n > n_0$ then we have

$$\mathbb{P}(I) \ge \mathbb{P}\Big(\bigcap_t I_t\Big) \ge 1 - \sum_t \big(1 - \mathbb{P}(I_t)\big) \ge 1 - \lfloor N/m \rfloor e^{-c_5 n} \ge 1 - e^{-c_2 n}.$$

And if $k \le n \le n_0$, then $I_1 \subseteq I$ and

$$\mathbb{P}(I) \ge \mathbb{P}(I_1) \ge 1 - e^{-c_5 n} \ge 1 - e^{-c_2 n}.$$ □

6. Uniform sorting networks are not geometrically realizable

Proof of Theorem 3. Goodman and Pollack proved in the paper [GP] that there exists a sorting network γ of size 5 that is not geometrically realizable. This sorting network is shown in Figure 9. (This is the smallest possible size of such a network.)

Let us view γ as a pattern. Suppose that γ occurs in a sorting network ω at time interval $[1, t]$ and position $[a, b]$. We claim that w is not geometrically realizable. Indeed, if ω were a geometrically realizable sorting networks associated with points $x_1, \dots, x_n \in \mathbb{R}^2$ (labeled from left to right), then γ would be a geometrically realizable sorting network associated with the points $x_a, \dots, x_b$.

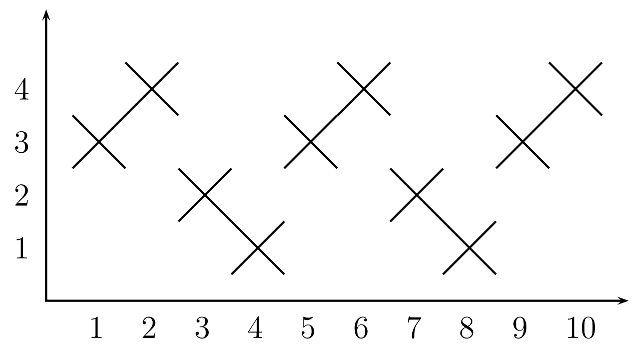

FIGURE 9. A sorting network that is not geometrically realizable.

Proposition 11 yields that with tending to 1 probability γ occurs within the time interval $[1, c_4 n]$ of a uniformly random sorting network ω of size n and thus ω is not geometrically realizable. □

Acknowledgements

O.A. has been supported by the University of Toronto, NSERC and the Sloan Foundation. V.G. has been supported by Microsoft Research, Moebius Foundation for Young Scientists, "Dynasty" foundation, RFBR-CNRS grant 10-01-93114, the program "Development of the scientific potential of the higher school" and by IUM-Simons foundation scholarship.

References

[AHRV] O. Angel, A. E. Holroyd, D. Romik and B. Virag, Random Sorting Networks. Adv. in Math., 215(2007), no. 2, 839-868. arXiv: math/0609538.

[AH] O. Angel, A. E. Holroyd, Random Subnetworks of Random Sorting Networks. Elec. J. Combinatorics, 17 (2010), no. 1, paper. 23.

[Az] K. Azuma, Weighted Sums of Certain Dependent Random Variables. Tôhoku Math. Journ., 19 (1967), 357–367.

[EG] P. Edelman and C. Greene, Balanced tableaux, Adv. in Math., 63 (1987), no.1, 42–99.

[FRT] J. S. Frame, G. de B. Robinson, and R. M. Thrall. The hook graphs of the symmetric groups. Canadian J. Math., 6 (1954), 316–324.

[GP] J. E. Goodman, R. Pollack, On the combinatorial classification of non-degenerate configurations in the plane, Journal of Combinatorial Theory, Series A, 29 (1980), no. 2, 220–235.

[GR] J. E. Goodman and J. O'Rourke, editors. Handbook of discrete and computational geometry. Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL, second edition, 2004.

[K] O. Kallenberg, Foundations of Modern Probability, Second Edition, Springer, 2002.

[M] I. G. Macdonald, Symmetric functions and Hall polynomials. Oxford Mathematical Monographs, Clarendon Press, Oxford, 1979.

[St] R. P. Stanley, On the number of reduced decompositions of elements of Coxeter groups, European J. Combin. 5 (1984), 359–372.

[W] D. Williams, Probability with martingales. Cambridge University Press, 1991.

(O. Angel) Department of Mathematics, University of British Columbia, Vancouver, BC V6T 1Z2, Canada

E-mail address: angel@math.ubc.ca

URL: http://math.ubc.ca/~angel/

(V. Gorin) Institute for Information Transmission Problems, Bolshoy Karetny 19, Moscow 127994, Russia

E-mail address: vadicgor@gmail.com

URL: http://www.mccme.ru/~vadicgor/

(A. E. Holroyd) Microsoft Research, 1 Microsoft Way, Redmond, WA 98052, USA

E-mail address: holroyd@microsoft.com

URL: http://research.microsoft.com/~holroyd/