Ungarian Markov chains[*]

Colin Defant[†]   Rupert Li[‡]

Abstract

We introduce the *Ungarian Markov chain* $\mathbf{U}_L$ associated to a finite lattice L. The states of this Markov chain are the elements of L. When the chain is in a state $x \in L$, it transitions to the meet of $\{x\} \cup T$, where T is a random subset of the set of elements covered by x. We focus on estimating $\mathcal{E}(L)$, the expected number of steps of $\mathbf{U}_L$ needed to get from the top element of L to the bottom element of L. Using direct combinatorial arguments, we provide asymptotic estimates when L is the weak order on the symmetric group S_n and when L is the n-th Tamari lattice. When L is distributive, the Markov chain $\mathbf{U}_L$ is equivalent to an instance of the well-studied random process known as *last-passage percolation with geometric weights*. One of our main results states that if L is a trim lattice, then $\mathcal{E}(L) \leq \mathcal{E}(\mathrm{spine}(L))$, where $\mathrm{spine}(L)$ is a specific distributive sublattice of L called the *spine* of L. Combining this lattice-theoretic theorem with known results about last-passage percolation yields a powerful method for proving upper bounds for $\mathcal{E}(L)$ when L is trim. We apply this method to obtain uniform asymptotic upper bounds for the expected number of steps in the Ungarian Markov chains of Cambrian lattices of classical types and the Ungarian Markov chains of ν-Tamari lattices.

1 Introduction

1.1 Ungar moves

Given a set $\mathscr{X}$ of $n \geq 4$ points in the plane that do not all lie on a single line, one can consider the lines that pass through two or more of the points in $\mathscr{X}$. In 1970, Scott [51] asked for the minimum possible number of distinct slopes determined by such a
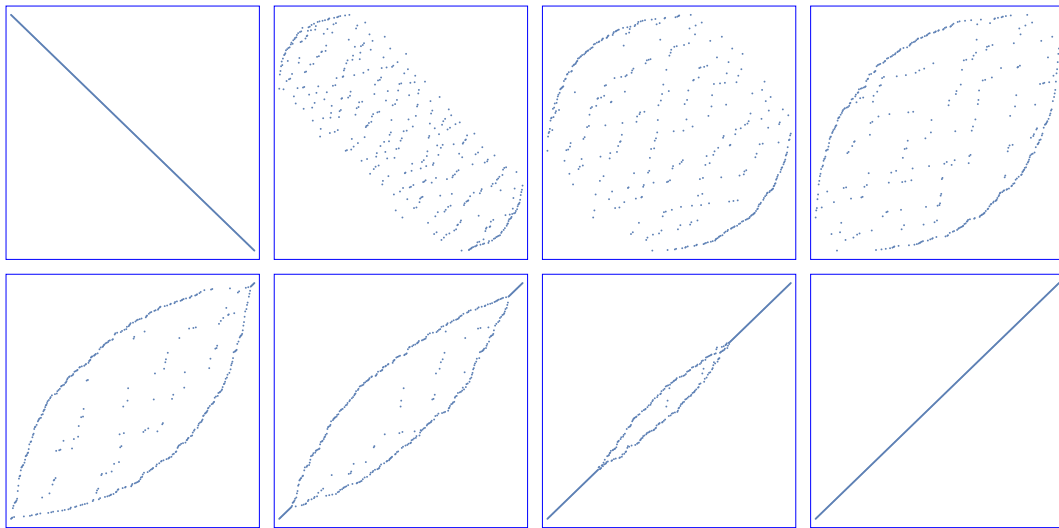
Figure 1: We ran the Ungarian Markov chain $\mathbf{U}_{S_{400}}$ with $p = 1/2$ starting with the decreasing permutation. Here, we show the plots of the permutations that the Markov chain reached at times $0, 200, 400, 600, 800, 1000, 1200, 1400$ (read from left to right, with the top row before the bottom).

collection of lines. He observed that the answer is at most n if n is even (take $\mathscr{X}$ to be the set of vertices of a regular n-gon) and is at most $n - 1$ if n is odd (take $\mathscr{X}$ to be the set consisting of the vertices and the center of a regular $(n-1)$-gon). In 1979, Burton and Purdy [18] proved that the answer is at least $\lfloor n/2 \rfloor$. Then, in 1982, Ungar [58] resolved the problem by showing that Scott's original upper bound is in fact the truth. His method, which followed a strategy proposed by Goodman and Pollack [34], involved projecting the points in $\mathscr{X}$ onto a generic line. The ordering of the projected points along the line can be interpreted as an element of S_n, the set of permutations of the set $[n] = \{1, \ldots, n\}$. As one slowly rotates the line, the projected points will sometimes swap positions in the ordering (see Figure 2). Thus, Ungar actually worked in a purely combinatorial setting in which he analyzed certain *moves* that can be performed on permutations.

Define an *Ungar move* to be an operation that reverses some disjoint collection of consecutive decreasing subsequences in a permutation. For example, we can perform an Ungar move on the permutation 763198542 by reversing the consecutive decreasing subsequences 63 and 542 to obtain the permutation 736198245. One can always perform a *trivial* Ungar move, which does nothing. At the other extreme, a *maximal* Ungar move reverses the maximal consecutive decreasing subsequences (also called *descending runs*) of the permutation. Applying a maximal Ungar move to 763198542 yields the permutation 136724589. Let $\mathrm{Pop} \colon S_n \to S_n$ be the operator that applies a maximal Ungar move. This operator is called the *pop-stack sorting map* because it coincides with the map that passes a permutation through a data structure called a *pop-stack* in a right-greedy manner. This map has already been studied in combinatorics and theoretical computer science [4, 5, 8, 24, 25, 39, 42]. If one starts with a permutation in S_n and repeatedly applies the pop-stack sorting map, then one will eventually reach the identity permutation $12 \cdots n$, which is fixed by Pop. Ungar proved that the maximum number of iterations of Pop needed to send a permutation in S_n to the identity is $n - 1$. He also proved the following theorem, which we record for later use (see [1, Chapter 12] for further exposition on this result).
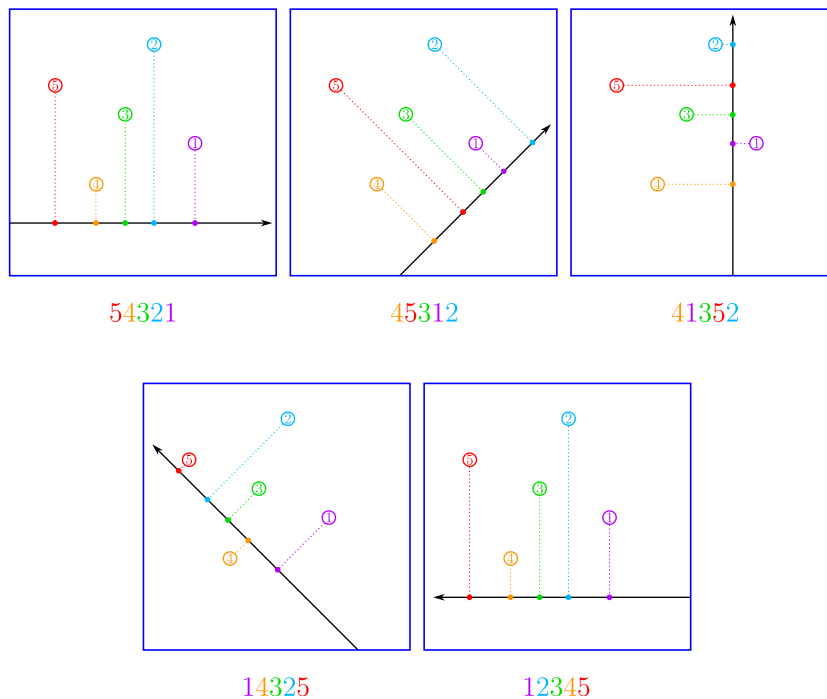
Figure 2: Five points in the plane are numbered $1, 2, 3, 4, 5$. After projecting the points onto a line, we can read the ordering of the projections along the line to obtain a permutation. As the line rotates, the associated permutation changes via an Ungar move.

Theorem 1.1 ([58])**.** *Let $n \geq 4$ be an integer. Suppose one starts with the decreasing permutation $n(n-1)\cdots 1$ and applies nontrivial Ungar moves until reaching the identity permutation $12\cdots n$. If the first Ungar move is not maximal, then the total number of Ungar moves used is at least n if n is even and is at least $n-1$ if n is odd.*

Remark 1.2. The hypothesis that the first Ungar move is not maximal in Theorem 1.1 corresponds to the condition in the original geometric problem that the points in $\mathscr{X}$ are not collinear.

All posets in this article are assumed to be finite. A *lattice* is a poset L such that any two elements $x, y \in L$ have a greatest lower bound—which is called their *meet* and denoted by $x \wedge y$—and a least upper bound—which is called their *join* and denoted by $x \vee y$. The meet and join operations are commutative and associative, so we can write $\bigwedge X$ and $\bigvee X$ for the meet and join, respectively, of any nonempty set $X \subseteq L$. Let $\mathrm{cov}_P(y) = \{x \in P : x \lessdot y\}$ denote the set of elements of a poset P that are covered by an element $y \in P$.

We can describe Ungar moves in a purely poset-theoretic manner if we consider the (right) weak order on S_n; this is a particular lattice that has been studied extensively in algebraic combinatorics. Applying an Ungar move to a permutation $x \in S_n$ is equivalent to taking the meet $\bigwedge(\{x\} \cup T)$ in the weak order for some set T of elements covered by x. This allows us to vastly generalize the definition of Ungar moves to the setting of lattices as follows.

Definition 1.3. *Let L be a lattice. An* Ungar move *is an operation that sends an element $x \in L$ to $\bigwedge(\{x\} \cup T)$ for some set $T \subseteq \mathrm{cov}_L(x)$. We say this Ungar move is* trivial *if $T = \emptyset$, and we say it is* maximal *if $T = \mathrm{cov}_L(x)$.*

Let L be a lattice. In [29, 30], the first author defined the *pop-stack sorting operator*

$\mathsf{Pop}_L \colon L \to L$ to be the map that acts on an element of L by applying a maximal Ungar move; that is,

$$\mathsf{Pop}_L(x) = \bigwedge(\{x\} \cup \mathrm{cov}_L(x)).$$

The articles [23, 29, 30, 33, 36, 50] study the dynamical properties and the image of Pop_L for various choices of interesting lattices L.

1.2 Random Ungar moves

The goal of this article is to investigate what happens when one applies Ungar moves randomly. Let us fix a probability $p \in (0, 1]$. Given an element x of a lattice L, we choose a random subset $T \subseteq \mathrm{cov}_L(x)$ by adding each element of $\mathrm{cov}_L(x)$ to T with probability p; we assume the choices for different elements are independent. We then apply the Ungar move that sends x to $\bigwedge(\{x\} \cup T)$; let us call this a *random Ungar move* (the probability p is implicit in this definition). This produces a Markov chain on L, a more symbol-heavy definition of which is as follows.

Definition 1.4. *Let L be a lattice, and fix a probability $p \in (0, 1]$. The* Ungarian Markov chain *of L is the Markov chain $\mathbf{U}_L$ with state space L such that for all $x, y \in L$, we have the transition probability*

$$\mathbb{P}(x \to y) = \sum_{\substack{T \subseteq \mathrm{cov}_L(x) \\ \bigwedge(\{x\} \cup T) = y}} p^{|T|}(1-p)^{|\mathrm{cov}_L(x)|-|T|}.$$

If we were to allow p to be 0, then $\mathbf{U}_L$ would simply fix each element of L; we assume $p > 0$ to avoid this boring scenario. If $p = 1$, then $\mathbf{U}_L$ is deterministic and agrees with the operator Pop_L.

Remark 1.5. One could define a more general version of an Ungarian Markov chain by assigning a probability p_e to each edge e in the Hasse diagram of L. However, for the sake of simplicity, we will only consider the case in which there is a single probability p.

The lattice L has a unique minimal element $\hat{0}$ and a unique maximal element $\hat{1}$. The Markov chain $\mathbf{U}_L$ is absorbing, and its unique absorbing state is $\hat{0}$. We will be primarily interested in estimating the expected number of steps required to reach the absorbing state $\hat{0}$ when we start at $\hat{1}$ in $\mathbf{U}_L$. We denote this expected number of steps by $\mathcal{E}(L)$. Note that $\mathcal{E}(L)$ implicitly depends on the fixed probability p.

Remark 1.6. The first author and Williams [33] recently introduced *semidistrim* lattices, which form a broad class of finite lattices that generalize distributive, semidistributive, and trim lattices. They showed how to define a natural invertible *rowmotion* operator on any semidistrim lattice, and they found that rowmotion is closely related to pop-stack sorting. In some sense, pop-stack sorting can be viewed as the "non-invertible cousin" of rowmotion. In a recent article [31], Nestoridi and the current authors introduced *rowmotion Markov chains* by inserting randomness into the definition of rowmotion. One can think of the Ungarian Markov chain of L as the "absorbing cousin" of the rowmotion Markov chain of L.

1.3 Symmetric groups under the weak order

We view the symmetric group S_n as a lattice under the weak order. A *descent* of a permutation $x \in S_n$ is an index $i \in [n-1]$ such that $x(i) > x(i+1)$; let $\mathrm{Des}(x)$ be the set of descents of x. To apply a random Ungar move to a permutation x, we choose a subset of $\mathrm{Des}(x)$ by including each descent with probability p, and we reverse the consecutive decreasing subsequences of x corresponding to those descents. For example, the set of descents of 731496852 is $\{1, 2, 5, 7, 8\}$; if we choose the subset $\{1, 7, 8\}$ (this happens with probability $p^3(1-p)^2$), then we obtain the new permutation 371496258.

Remark 1.7. Applying a random Ungar move to a permutation x is equivalent to sending x through a pop-stack (see [6, 52] for the definition) using the following two rules:

- the entries in the pop-stack must always be increasing from top to bottom;
- at each point in time when we have a choice whether to pop the entries out of the stack, we do so with probability $1 - p$.

There is also an algebraic way to interpret an Ungar move if we think of S_n as the Coxeter group generated by $\{s_1, \ldots, s_{n-1}\}$, where s_i is the transposition $(i\ i+1)$. A *parabolic subgroup* of S_n is a subgroup generated by a subset of $\{s_1, \ldots, s_{n-1}\}$. Suppose we apply an Ungar move to a permutation x by choosing a subset $T \subseteq \mathrm{Des}(x)$ and reversing the corresponding decreasing subsequences. The result is the permutation $xw_0(T)$, where $w_0(T)$ is the maximal element (in the weak order) of the parabolic subgroup of S_n generated by $\{s_i : i \in T\}$. The element $w_0(T)$ is an involution. Thus, if we start at a permutation x and run the Markov chain $\mathbf{U}_{S_n}$ until reaching the identity permutation $12 \cdots n$ (which is the absorbing state), we will generate a random factorization $x_1 \cdots x_m$, where each factor x_k is of the form $w_0(T)$ for some $T \subseteq [n-1]$. This factorization is *length-additive* in the sense that the Coxeter length (i.e., the number of inversions) of x is equal to the sum of the Coxeter lengths of the factors $x_1, \ldots, x_m$.

The *plot* of a permutation $x \in S_n$ is the graph showing the points $(i, x(i))$ for all $i \in [n]$. Figure 1 shows the plots of eight of the permutations that we obtained while running $\mathbf{U}_{S_{400}}$ with $p = 1/2$ starting with the decreasing permutation. It would be extremely interesting to determine the shapes of these permutations in a manner similar to what was done in [2, 3, 27, 28] for permutations obtained via random sorting networks.

It follows easily from Ungar's Theorem 1.1 that if we fix the probability $p \in (0, 1)$, then $\mathcal{E}(S_n)$, the expected number of steps in the Ungarian Markov chain $\mathbf{U}_{S_n}$ needed to go from the decreasing permutation to the identity permutation, is at least $n - 1 + o(1)$; see Lemma 3.1 below. Noting that S_n is a graded lattice of total rank $\binom{n}{2}$, one can easily show that $\mathcal{E}(S_n)$ grows at most quadratically in n. The following theorem greatly improves upon this naive upper bound.

Theorem 1.8. *For $p \in (0, 1)$ fixed, we have*

$$n - 1 + o(1) \leq \mathcal{E}(S_n) \leq \frac{8}{p} n \log n + O(n)$$

as $n \to \infty$.

1.4 Distributive lattices

An *order ideal* of a poset P is a set $Y \subseteq P$ such that if $x, y \in P$ satisfy $x \leq y$ and $y \in Y$, then $x \in Y$. When ordered by inclusion, the order ideals of P form a distributive lattice that we denote by $J(P)$. Birkhoff's Fundamental Theorem of Finite Distributive Lattices [12] states that every (finite) distributive lattice is isomorphic to the lattice of order ideals of some (finite) poset.

Let P be a poset, and consider an order ideal $I \in J(P)$. The set $\mathrm{cov}_{J(P)}(I)$ consists of the order ideals of P that can be obtained by removing a single maximal element from I. Therefore, applying a random Ungar move to I is equivalent to removing a random subset of the set of maximal elements of I, where each maximal element is removed with probability p. Associate to each $x \in P$ a geometric random variable G_x with parameter p (i.e., with expected value $1/p$), and assume that the random variables associated to different elements of P are all independent. If we start the Markov chain $\mathbf{U}_{J(P)}$ at the top element $\hat{1} = P$, then we can think of G_x as the number of steps throughout the process during which the state (i.e., order ideal) has x as a maximal element. Thus, the

number of steps needed to go from the top element $\hat{1} = P$ to the bottom element $\hat{0} = \emptyset$ is $\max\limits_{\mathcal{C} \in \mathrm{MC}(P)} \sum_{x \in \mathcal{C}} G_x$, where $\mathrm{MC}(P)$ is the set of maximal chains of P. Consequently,

$$\mathcal{E}(J(P)) = \mathbb{E}\left(\max_{\mathcal{C} \in \mathrm{MC}(P)} \sum_{x \in \mathcal{C}} G_x\right). \tag{1.1}$$

The above reformulation of the Markov chain $\mathbf{U}_{J(P)}$ in terms of geometric random variables is a well-studied random process known as *last-passage percolation with geometric weights*. Therefore, while we originally viewed Ungarian Markov chains as randomized versions of pop-stack sorting, one could equally well package them as generalizations of last-passage percolation with geometric weights.

Last-passage percolation has been investigated thoroughly when P is the poset associated to a Young diagram, where it has also been called the *multicorner growth process* [48, Chapters 4 & 5] and is very closely related to the *totally asymmetric simple exclusion process* (TASEP) [49]. In this setting, very precise results are known concerning the number of steps needed to reach the absorbing state $\hat{0}$ and even what the order ideals (i.e., Young diagrams) typically "look like" throughout the process. We refer the reader to [48, Chapters 4 & 5] for more details; here, we simply state a result that we will employ later. Let $R_{k \times \ell}$ denote the $k \times \ell$ rectangle poset, which is simply the product of a chain of length k and a chain of length ℓ. The following result is originally due to Cohn–Elkies–Propp [26] and Jockusch–Propp–Shor [38] (using different language); it appears as [48, Theorem 2.4].

Theorem 1.9 ([26, 38]). *Let $(k_n)_{n \geq 1}$ and $(\ell_n)_{n \geq 1}$ be sequences of positive integers such that the limit $(\overline{k}, \overline{\ell}) = \lim\limits_{n \to \infty} \frac{1}{n}(k_n, \ell_n)$ exists. As $n \to \infty$, we have*

$$\mathcal{E}(J(R_{k_n \times \ell_n})) = \frac{1}{p}\left(\overline{k} + \overline{\ell} + 2\sqrt{(1-p)\overline{k}\,\overline{\ell}}\right) n + o(n).$$

The next theorem follows from (1.1) and a Chernoff bound; we provide the details in Section 4.

Theorem 1.10. *Let $(P_n)_{n \geq 1}$ be a sequence of posets. Suppose there exist constants Γ and μ such that $|\mathrm{MC}(P_n)| \leq \Gamma^{(1+o(1))n}$ and such that the maximum size of a chain in P_n is at most μn. As $n \to \infty$, we have*

$$\mathcal{E}(J(P_n)) \leq \frac{1}{p}\left(\mu + \log\Gamma + \sqrt{2\mu\log\Gamma + (\log\Gamma)^2}\right) n + o(n).$$

We note that for rectangle posets, the upper bound in Theorem 1.10 is weaker than the bound one would obtain from Theorem 1.9.

1.5 Trim lattices

Thomas [56] introduced *trim lattices* as generalizations of distributive lattices that need not be graded. There are several notable examples of trim lattices such as Cambrian lattices and ν-Tamari lattices; see [57, Section 7] for a more extensive list. (Note, however, that the weak order on S_n is not trim when $n \geq 3$.) In [57], Thomas and Williams investigated fascinating dynamical properties of trim lattices and also related their results to quiver representation theory. They defined the *spine* of a trim lattice L to be the set $\mathrm{spine}(L)$ of elements that lie on at least one maximum-length chain of L, and they showed that $\mathrm{spine}(L)$ is a distributive sublattice of L. The following is one of our main results.

Theorem 1.11. *If L is a trim lattice, then*

$$\mathcal{E}(L) \leq \mathcal{E}(\mathrm{spine}(L)).$$

In fact, we have the following more general result, which we will obtain as a corollary of Theorem 1.11, (1.1), and known properties of trim lattices. Note that quotients and intervals of trim lattices are trim.

Corollary 1.12. *If L is a trim lattice and L' is a quotient or an interval of L, then*

$$\mathcal{E}(L') \leq \mathcal{E}(\mathrm{spine}(L)).$$

The proof of Theorem 1.11 is lattice-theoretic and relies on several known properties of trim lattices. Because $\mathrm{spine}(L)$ is distributive, we can interpret its Ungarian Markov chain as last-passage percolation and use the results discussed in Section 1.4 to obtain immediate upper bounds for $\mathcal{E}(\mathrm{spine}(L))$. Thus, Theorem 1.11 and Corollary 1.12 combine lattice theory and probability theory to yield a powerful method for bounding $\mathcal{E}(L)$ for any trim lattice L. We illustrate this method with some applications for some special classes of trim lattices, whose definitions we postpone.

Associated to a Coxeter element c of a finite Coxeter group W is an important lattice Camb_c, called the *c-Cambrian lattice*, which Reading introduced in [43, 45]. Cambrian lattices have now been studied thoroughly because of their rich combinatorial and lattice-theoretic properties and their connections to representation theory, cluster algebras, and polyhedral geometry [7, 21, 37, 46, 47, 54, 57]. Because we are primarily interested in asymptotic results, it makes little sense for us to deal with Coxeter groups of exceptional types. Thus, we will focus on the classical types A, B, and D.

Theorem 1.13. *For each $n \geq 1$, choose a Coxeter element $c^{(n)}$ of the Coxeter group A_n. As $n \to \infty$, we have*

$$\mathcal{E}\left(\mathrm{Camb}_{c^{(n)}}\right) \leq \frac{1}{p}\left(2 + 2\sqrt{1-p}\right)n + o(n).$$

The Coxeter graph of the Coxeter group B_n is a path with vertices $s_0, \ldots, s_{n-1}$ and edges of the form $\{s_i, s_{i+1}\}$ for $0 \leq i \leq n-2$. A Coxeter element of B_n is uniquely determined by orienting the edges of this graph; let $r(c)$ denote the number of edges $\{s_i, s_{i+1}\}$ that are oriented from s_i to s_{i+1} in the orientation corresponding to a Coxeter element c.

Theorem 1.14. *For each $n \geq 2$, choose a Coxeter element $c^{(n)}$ of the Coxeter group B_n, and assume that the limit $\overline{r} = \lim_{n\to\infty} \frac{1}{n}r(c^{(n)})$ exists. As $n \to \infty$, we have*

$$\mathcal{E}\left(\mathrm{Camb}_{c^{(n)}}\right) \leq \frac{1}{p}\left(3 + 2\sqrt{(1-p)(2-\overline{r})(1+\overline{r})}\right)n + o(n).$$

The Coxeter graph of the Coxeter group D_n has vertices $s_0, \ldots, s_{n-1}$, edges of the form $\{s_i, s_{i+1}\}$ for $1 \leq i \leq n-2$, and an additional edge $\{s_0, s_2\}$. A Coxeter element of D_n is uniquely determined by orienting the edges of this graph. Let $r(c)$ denote the number of edges $\{s_i, s_{i+1}\}$ that are oriented from s_i to s_{i+1} in the orientation corresponding to a Coxeter element c (note that $r(c)$ does not depend on the orientation of $\{s_0, s_2\}$), and let $u(c)$ denote the number of edges in a maximum-length directed path in this orientation.

Theorem 1.15. *For each $n \geq 4$, choose a Coxeter element $c^{(n)}$ of the Coxeter group D_n, and assume that the limits $\overline{r} = \lim_{n\to\infty} \frac{1}{n}r(c^{(n)})$ and $\overline{u} = \lim_{n\to\infty} \frac{1}{n}u(c^{(n)})$ exist. As $n \to \infty$, we have*

$$\mathcal{E}\left(\mathrm{Camb}_{c^{(n)}}\right) \leq \frac{1}{p}\left(6 + 4\sqrt{(1-p)(2-\overline{r})(1+\overline{r})}\right)n + o(n)$$

and

$$\mathcal{E}(\mathrm{Camb}_{c^{(n)}}) \leq \frac{1}{p}\left(2 + \overline{u} + \log\left(5 \cdot 2^{\overline{u}}\right) + \sqrt{2(2+\overline{u})\log\left(5 \cdot 2^{\overline{u}}\right) + \left(\log\left(5 \cdot 2^{\overline{u}}\right)\right)^2}\right)n + o(n).$$

Remark 1.16. Let us consider numerical approximations of the upper bounds in Theorems 1.13 to 1.15 when $p = 1/2$. The bound in Theorem 1.13 is $(6.82843 + o(1))n$. The bound in Theorem 1.14 ranges from $(10 + o(1))n$ to $(10.24264 + o(1))n$ as $\bar{r}$ ranges from 0 to 1. The first bound in Theorem 1.15 ranges from $(20 + o(1))n$ to $(20.48528 + o(1))n$ as $\bar{r}$ ranges from 0 to 1, while the second bound ranges from $(13.22822 + o(1))n$ to $(19.34986 + o(1))n$ as $\bar{u}$ ranges from 0 to 1. (So the second bound in Theorem 1.15 is always better than the first when $p = 1/2$, but the first bound ends up being better when p is closer to 1.)

Remark 1.17. If we maximize the bounds in Theorems 1.13 to 1.15 over all values of $\bar{r}$ and $\bar{u}$, we obtain the upper bounds

$$\frac{1}{p}\left(2 + 2\sqrt{1-p}\right)n + o(n),$$
$$\frac{1}{p}\left(3 + 3\sqrt{1-p}\right)n + o(n),$$

and

$$\frac{1}{p}\min\left\{6 + 6\sqrt{1-p}, 3 + \log(10) + \sqrt{6\log(10) + (\log(10))^2}\right\}n + o(n),$$

respectively, which are linear in n and apply uniformly to all Cambrian lattices of the prescribed types. It is not clear how one could obtain such linear bounds (with any leading coefficients whatsoever) without the use of Theorem 1.11.

For the sake of completeness, we will also prove the following very easy result about Cambrian lattices of dihedral groups.

Theorem 1.18. *Let c be a Coxeter element of $I_2(m)$, the dihedral group of order $2m$. Then*

$$\mathcal{E}(\mathrm{Camb}_c) = \frac{1 + m(1-p)}{2p - p^2}.$$

Cambrian lattices are generalizations of Tamari lattices, which are significant objects in algebraic combinatorics that have been studied since the work of Tamari in 1962 [55]. Bergeron and Préville-Ratelle [11] introduced different generalizations of Tamari lattices called m-*Tamari lattices* in their study of trivariate diagonal harmonics. These lattices have now received a great deal of further attention [15, 16, 22, 32, 35]. Going further, Préville-Ratelle and Viennot [41] generalized m-Tamari lattices by defining the ν-*Tamari lattice* $\mathrm{Tam}(\nu)$, where ν is an arbitrary lattice path consisting of unit north and east steps; this more general class of lattices has been connected to diagonal coinvariant spaces [41], polyhedral geometry [10, 9, 19], and combinatorial dynamics [30, 32].

We will provide a simple explicit description of the spine of $\mathrm{Tam}(\nu)$ (see Proposition 7.1), which appears to be new. Note, however, that spines of m-Tamari lattices were studied in [35] (see also [40]). Combining this description with Theorem 1.9, we will deduce the following upper bound for $\mathcal{E}(\mathrm{Tam}(\nu))$.

Theorem 1.19. *For each $\ell \geq 1$, choose a lattice path $\nu^{(\ell)}$ of length ℓ consisting of north and east steps. Let n_ℓ be the number of north steps in $\nu^{(\ell)}$, and assume that the limit $\overline{n} = \lim_{\ell \to \infty} \frac{1}{\ell} n_\ell$ exists. As $\ell \to \infty$, we have*

$$\mathcal{E}(\mathrm{Tam}(\nu^{(\ell)})) \leq \frac{1}{p}\left(1 + 2\sqrt{(1-p)\overline{n}(1-\overline{n})}\right)\ell + o(\ell).$$

1.6 Tamari lattices

The original Tamari lattice Tam_n is a Cambrian lattice of type A_{n-1} and is also $\mathrm{Tam}(\nu)$ when $\nu = (\mathrm{NE})^n$. By specializing either Theorem 1.13 or Theorem 1.19, we find that

$\mathcal{E}(\mathrm{Tam}(\nu)) \leq \frac{1}{p}(2 + 2\sqrt{1-p})n + o(n)$. Because Tamari lattices are so fundamental, we will analyze them on their own and derive a better upper bound.

In [17], Bruss and O'Cinneide studied the asymptotic behavior of $\rho_p(n)$, which is defined to be the probability that the maximum of n independent geometric random variables, each with expected value $1/p$, is attained uniquely. Somewhat surprisingly, they found that $\lim_{n\to\infty} \rho_p(n)$ does not exist if $p < 1$ and that $\lim_{n\to\infty}(\rho_p(n) - \Upsilon_p(n)) = 0$, where

$$\Upsilon_p(x) = \begin{cases} px\sum_{k\in\mathbb{Z}}(1-p)^k e^{-(1-p)^k x} & \text{if } p < 1; \\ 0 & \text{if } p = 1. \end{cases}$$

Note that $\Upsilon_p((1-p)x) = \Upsilon_p(x)$ for all $x > 0$. Therefore, the quantity

$$\overline{\rho}_p = \max_{0 < x < 1} \Upsilon_p(x)$$

is equal to $\limsup_{n\to\infty} \rho_p(n)$. When $p = 1/2$, we have $\overline{\rho}_{1/2} \approx 0.72136$.

Theorem 1.20. *We have*

$$\mathcal{E}(\mathrm{Tam}_n) \leq \frac{2}{p}\left(\sqrt{\overline{\rho}_p(1+\overline{\rho}_p)} - \overline{\rho}_p\right)n + o(n)$$

as $n \to \infty$.

Remark 1.21. Because

$$\sqrt{\overline{\rho}_p(1+\overline{\rho}_p)} - \overline{\rho}_p < \sqrt{1+\overline{\rho}_p} - \overline{\rho}_p \leq 1,$$

the upper bound for $\mathcal{E}(\mathrm{Tam}_n)$ provided by either Theorem 1.13 or Theorem 1.19 is weaker than the upper bound in Theorem 1.20. When we set $p = 1/2$, the former upper bound is $(4 + 2\sqrt{2} + o(1))n \approx (6.82843 + o(1))n$, while the latter is approximately $(1.57186 + o(1))n$.

1.7 Outline

In Section 2, we recall some basic definitions concerning posets and lattices, and we prove a lemma that bounds the tails of a sum of i.i.d. geometric random variables. Section 3 is devoted to proving Theorem 1.8, which bounds $\mathcal{E}(S_n)$. In Section 4, we consider $\mathcal{E}(L)$ when L is distributive; we state some corollaries that follow from interpreting $\mathbf{U}_L$ as last-passage percolation, and we prove Theorem 1.10. Section 5 recalls necessary background about trim lattices and proves Theorem 1.11. In Section 6, we recall background about Cambrian lattices, describe the spines of Cambrian lattices, and combine Theorem 1.11 with Theorems 1.9 and 1.10 to prove Theorems 1.13 to 1.15; we also quickly prove Theorem 1.18. In Section 7, we discuss background about ν-Tamari lattices, derive a simple description of the spines of ν-Tamari lattices, and use this description to deduce Theorem 1.19 from Theorems 1.9 and 1.11. Section 3 provides a direct combinatorial proof of Theorem 1.20, which bounds $\mathcal{E}(\mathrm{Tam}_n)$ from above. Finally, Section 9 lists several enticing suggestions for further research.

2 Preliminaries

We assume basic familiarity with the theory of posets (partially ordered sets), as discussed in [53, Chapter 3]. As mentioned in the introduction, all posets considered in this article are assumed to be finite. Here, we recall some notions that we will need.

Let P be a poset. A *subposet* of P is a poset $P' \subseteq P$ such that if $x, y \in P'$ satisfy $x \leq y$ in P', then $x \leq y$ in P. The *dual* of P is the poset P^* with the same underlying set as P

such that $x \leq y$ in P if and only if $y \leq x$ in P^*. Given $x, y \in P$ with $x \leq y$, the *interval* from x to y is the set $[x, y] = \{z \in P : x \leq z \leq y\}$. If $x < y$ and $[x, y] = \{x, y\}$, then we say y *covers* x and write $x \lessdot y$. We write $\mathrm{cov}_P(y) = \{x \in P : x \lessdot y\}$ for the set of elements of P that are covered by y. A *chain* of P is a totally-ordered subset of P; we say a chain is *maximal* if it is not properly contained in another chain. Let $\mathrm{MC}(P)$ denote the set of maximal chains of P. The *length* of a chain $\mathcal{C}$ of P is $|\mathcal{C}| - 1$. The *length* of P is the maximum of the lengths of the chains of P. We say P is *graded* if all of its maximal chains have the same length.

As discussed in Section 1, a *lattice* is a poset L such that any two elements $x, y \in L$ have a meet $x \wedge y$ and a join $x \vee y$. Write $\bigwedge X$ and $\bigvee X$ for the meet and join, respectively, of a nonempty set $X \subseteq L$. The lattice L has a unique minimal element $\hat{0}$ and a unique maximal element $\hat{1}$; we use the convention that $\bigwedge \emptyset = \hat{1}$ and $\bigvee \emptyset = \hat{0}$. A *sublattice* of L is a lattice $L' \subseteq L$ whose meet and join operations agree with those of L. We implicitly view intervals of L as sublattices of L.

Given lattices L and L', a map $\varphi \colon L \to L'$ is a *lattice homomorphism* if $\varphi(x \wedge y) = \varphi(x) \wedge \varphi(y)$ and $\varphi(x \vee y) = \varphi(x) \vee \varphi(y)$ for all $x, y \in L$. We say L' is a *quotient* of L if there exists a surjective lattice homomorphism from L to L'.

Throughout this article, we omit floor and ceiling symbols when they do not affect the relevant asymptotics.

We will frequently need the following lemma, which allows us to bound the upper tails of a sum of i.i.d. geometric random variables.

Lemma 2.1. *Let $G_1, \ldots, G_k$ be independent geometric random variables with parameter p. For each real number $\gamma \geq 1$, we have*

$$\mathbb{P}\left(\sum_{i=1}^{k} G_i > \gamma k/p\right) \leq e^{-\frac{\gamma k}{2}\left(1 - \gamma^{-1}\right)^2}.$$

Proof. Extend $G_1, \ldots, G_k$ to an infinite sequence $G_1, G_2, \ldots$ of independent geometric random variables with parameter p. Let

$$Y_m = \begin{cases} 1 & \text{if } m = \sum_{i=1}^{j} G_i \text{ for some } j \geq 1; \\ 0 & \text{otherwise.} \end{cases}$$

Then $Y_1, Y_2, \ldots$ is a sequence of independent Bernoulli random variables, each with expected value p. Thus, $\sum_{1 \leq m \leq \gamma k/p} Y_m$ is a binomial random variable with parameters $\lfloor \gamma k/p \rfloor$ and p. Let $\mu = \gamma k$ and $\delta = 1 - \gamma^{-1}$ so that $k = (1 - \delta)\mu$. A standard Chernoff bound tells us that

$$\mathbb{P}\left(\sum_{i=1}^{k} G_i > \gamma k/p\right) = \mathbb{P}\left(\sum_{1 \leq m \leq \gamma k/p} Y_m < k\right) \leq e^{-\mu \delta^2/2} = e^{-\frac{\gamma k}{2}\left(1 - \gamma^{-1}\right)^2}. \qquad \square$$

3 Symmetric groups under the weak order

Let s_i denote the transposition $(i \; i+1)$ in the symmetric group S_n. The *weak order* on S_n is the partial order in which there is a cover relation $w \lessdot w'$ whenever i is a descent of w' and $w' = w s_i$ (more precisely, this is the *right* weak order). It is well known that the weak order on S_n is a lattice; we will simply write S_n for this lattice.

In this section, we analyze the Ungarian Markov chains $\mathbf{U}_{S_n}$. We will prove Theorem 1.8, which provides asymptotic estimates for the expected value $\mathcal{E}(S_n)$. We assume throughout this section that our fixed probability p is strictly between 0 and 1.

To begin, let us derive the lower bound in Theorem 1.8 from Ungar's Theorem 1.1.

Lemma 3.1. *We have $n - 1 + o(1) \leq \mathcal{E}(S_n)$.*

Proof. Let us run $\mathbf{U}_{S_n}$ starting at the top element $\hat{1} = n(n-1)\cdots 1$. This element $\hat{1}$ covers $n-1$ elements of S_n, so $\mathbf{U}_{S_n}$ has the transition probability $\mathbb{P}(\hat{1} \to \hat{1}) = (1-p)^{n-1}$. The only Ungar move that we can apply to $\hat{1}$ to obtain $\hat{0}$ is the maximal Ungar move, and this has transition probability $\mathbb{P}(\hat{1} \to \hat{0}) = p^{n-1}$. If the first random Ungar move that we apply is not trivial or maximal, then it follows immediately from Theorem 1.1 that the total number of steps we will need to reach $\hat{0}$ is at least $n-1$. Therefore,

$$
\begin{aligned}
\mathcal{E}(S_n) &\geq (1-p)^{n-1}(1 + \mathcal{E}(S_n)) + p^{n-1} \cdot 1 + \sum_{\emptyset \neq T \subsetneq [n-1]} p^{|T|}(1-p)^{n-1-|T|}(n-1) \\
&= (1-p)^{n-1}(1 + \mathcal{E}(S_n)) + p^{n-1} + (1 - (1-p)^{n-1} - p^{n-1})(n-1) \\
&= 1 + (1-p)^{n-1}\mathcal{E}(S_n) + (1 - (1-p)^{n-1} - p^{n-1})(n-2)
\end{aligned}
$$

It follows that

$$
\mathcal{E}(S_n) \geq \frac{1 + (1 - (1-p)^{n-1} - p^{n-1})(n-2)}{1 - (1-p)^{n-1}} = n - 2 + \frac{1 - p^{n-1}(n-2)}{1 - (1-p)^{n-1}} = n - 1 + o(1). \quad \square
$$

We now proceed to prove the upper bound in Theorem 1.8.

Given a permutation $w \in S_n$, we write $\mathrm{DB}(w)$ for the set $\{w(i+1) : i \in \mathrm{Des}(w)\}$ of *descent bottoms* of w. For integers $\beta \in [n-1]$ and $j \geq 0$, let $X_j^{(\beta)}$ be a Bernoulli random variable with expected value p; assume that all of these random variables for different choices of β and j are independent. We can simulate the Markov chain $\mathbf{U}_{S_n}$ using these random variables as follows. Starting with the decreasing permutation $\sigma_0 = n(n-1)\cdots 1$, we are going to create a sequence $\sigma_0, \sigma_1, \ldots$ of permutations. Suppose we have already generated the permutations $\sigma_0, \ldots, \sigma_t$. We will define σ_{t+1} to be the permutation obtained by choosing a random subset T_t of $\mathrm{Des}(\sigma_t)$ and then applying the corresponding random Ungar move to σ_t. In Coxeter-theoretic terminology, this means that $\sigma_{t+1} = \sigma_t w_0(T_t)$, where $w_0(T_t)$ is the maximal element of the parabolic subgroup of S_n generated by $\{s_i : i \in T_t\}$. For each $\beta \in \mathrm{DB}(\sigma_t)$, let $j_t(\beta) = |\{\ell \in \{0, \ldots, t-1\} : \beta \in \mathrm{DB}(\sigma_\ell)\}|$. The random variables from above now enter the picture: we define

$$
T_t = \left\{ i \in \mathrm{Des}(\sigma_t) : X_{j_t(\sigma_t(i+1))}^{(\sigma_t(i+1))} = 1 \right\}.
$$

Note that for permutations $w, w' \in S_n$, the conditional probability $\mathbb{P}(\sigma_{t+1} = w' \mid \sigma_t = w)$ is equal to the transition probability $\mathbb{P}(w \to w')$ in $\mathbf{U}_{S_n}$.

This construction is designed so that if β appears immediately to the right of some entry $\alpha > \beta$ in σ_t, then β will move to the left of α when we transition to σ_{t+1} if and only if $X_{j_t(\beta)}^{(\beta)} = 1$. Our definition of $j_t(\beta)$ is such that $X_0^{(\beta)}, \ldots, X_{j_t(\beta)-1}^{(\beta)}$ are precisely the random variables from the sequence $(X_j^{(\beta)})_{j \geq 0}$ that were already used when we generated the permutations $\sigma_0, \ldots, \sigma_t$ (so $X_{j_t(\beta)}^{(\beta)}$ is the next available random variable from this sequence).

Let N be the unique integer such that $\sigma_{N-1} \neq \sigma_N = 12 \cdots n$. Then $\sigma_t = 12 \cdots n$ for all $t \geq N$. For $\beta \in [n-1]$, let $\eta_0^{(\beta)} < \cdots < \eta_{q(\beta)}^{(\beta)}$ be the indices t such that $\beta \in \mathrm{DB}(\sigma_t)$. If $\beta < \beta'$ and β appears to the left of β' in a permutation σ_t, then β must also appear to the left of β' in all of the permutations $\sigma_{t+1}, \sigma_{t+2}, \ldots$. This implies that

$$
\sum_{j=0}^{q(\beta)} X_j^{(\beta)} \leq n - \beta. \tag{3.1}
$$

Observe that $N - 1 = \max_{1 \leq \beta \leq n-1} \eta_{q(\beta)}^{(\beta)}$.

Lemma 3.2. *Assume that*

$$\sum_{j=a}^{b-1} X_j^{(\beta)} \geq p(b-a) - \sqrt{8p \log n} \cdot \sqrt{b-a}$$

for all $1 \leq \beta \leq n-1$ and $0 \leq a < b \leq q(\beta) + 1$. Then $N \leq \frac{8}{p} n \log n + O(n)$.

Proof. To ease notation, let $\chi(n) = \sqrt{8p \log n}$. Let $r_\beta(s) = n - \sigma_s^{-1}(\beta)$ be the number of entries that appear to the right of β in σ_s. We will prove by induction on β that

$$r_\beta(s) \geq \min\{n - \beta, ps - \chi(n)\sqrt{\beta s} - (p+1)(\beta-1)\} \tag{3.2}$$

for all $1 \leq \beta \leq n$ and $0 \leq s \leq N$.

We first consider the base case when $\beta = 1$. Note that $\eta_j^{(1)} = j$ for every $0 \leq j \leq q(1)$. Suppose $r_1(s) < n - 1$. Then $s \leq q(1)$, so

$$r_1(s) \geq \sum_{j=0}^{s-1} X_j^{(1)} \geq ps - \chi(n)\sqrt{s},$$

where we obtained the last inequality by setting $a = 0$ and $b = s$ in the hypothesis of the lemma. This proves the base case.

We may now assume $\beta \geq 2$ and proceed by induction on β. Suppose by way of contradiction that there exists $0 \leq s \leq N$ such that

$$r_\beta(s) < n - \beta \quad \text{and} \quad r_\beta(s) < ps - \chi(n)\sqrt{\beta s} - (p+1)(\beta-1).$$

Then it follows from our induction hypothesis that $r_\beta(s) \leq \min\{r_1(s), \ldots, r_{\beta-1}(s)\} - 2$. Let s^* be the largest integer such that $s^* \leq s - 1$ and $r_\beta(s^*) = \min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} - 1$ (so β appears immediately to the right of the rightmost entry from the list $1, \ldots, \beta - 1$ in σ_{s^*}); it is straightforward to see from the original definition of an Ungar move (for S_n) that such an integer s^* must exist. Because the entry immediately to the left of β in σ_{s^*} is smaller than β, we must have $r_\beta(s^* + 1) = r_\beta(s^*) = \min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} - 1$. It follows from our choice of s^* that β is a descent bottom of each of the permutations $\sigma_{s^*+1}, \ldots, \sigma_{s-1}$. Consequently, there exists k such that the numbers $\eta_k^{(\beta)}, \ldots, \eta_{k+s-s^*-2}^{(\beta)}$ are the same as the numbers $s^* + 1, \ldots, s - 1$, respectively. Applying the hypothesis of the lemma with $a = k$ and $b = k + s - s^* - 1$, we find that

$$r_\beta(s) \geq r_\beta(s^* + 1) + \sum_{j=k}^{k+s-s^*-2} X_j^{(\beta)}$$

$$= \min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} - 1 + \sum_{j=k}^{k+s-s^*-2} X_j^{(\beta)}$$

$$\geq \min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} - 1 + p(s - s^* - 1) - \chi(n)\sqrt{s - s^* - 1}. \tag{3.3}$$

Our induction hypothesis tells us that

$$\min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} \geq \min\{n - \beta + 1, ps^* - \chi(n)\sqrt{(\beta-1)s^*} - (p+1)(\beta-2)\}.$$

If $\min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} = n - \beta + 1$, then $r_\beta(s^*) = n - \beta$, contradicting the fact that $r_\beta(s) \leq \min\{r_1(s), \ldots, r_{\beta-1}(s)\} - 2$ (if β appears in position β in σ_{s^*}, then it cannot appear further to the right in σ_s). This shows that $\min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} < n - \beta + 1$, so $\min\{r_1(s^*), \ldots, r_{\beta-1}(s^*)\} \geq ps^* - \chi(n)\sqrt{(\beta-1)s^*} - (p+1)(\beta-2)$. Combining this with (3.3) yields

$$r_\beta(s) \geq ps^* - \chi(n)\sqrt{(\beta-1)s^*} - (p+1)(\beta-2) - 1 + p(s - s^* - 1) - \chi(n)\sqrt{s - s^* - 1}$$

$$= ps - \chi(n) \left[\sqrt{(\beta - 1)s^*} + \sqrt{s - s^* - 1} \right] - (p+1)(\beta - 1)$$
$$\geq ps - \chi(n)\sqrt{\beta s} - (p+1)(\beta - 1),$$

where the inequality $\sqrt{(\beta - 1)s^*} + \sqrt{s - s^* - 1} \leq \sqrt{\beta s}$ used to deduce the last line follows from basic calculus. This contradicts our original choice of s, so our inductive proof of (3.2) is complete.

For every $0 \leq s \leq N - 1$, there exists $\beta \in [n]$ such that $r_\beta(s) < n - \beta$. Therefore, if k is an integer such that

$$pk - \chi(n)\sqrt{\beta k} - (p+1)(\beta - 1) \geq n - \beta \tag{3.4}$$

for all $1 \leq \beta \leq n$, then (3.2) tells us that $N \leq k$. If k satisfies the inequality (3.4) when $\beta = n$, then it also satisfies (3.4) for all $\beta \in [n]$. Thus, to get an upper bound for N, we simply need to find an integer k such that $pk - \chi(n)\sqrt{n}\sqrt{k} - (p+1)(n-1) \geq 0$. The left-hand side of this inequality is quadratic in $\sqrt{k}$, so we can use the quadratic formula (and the definition of $\chi(n)$) to find that we just need k to satisfy

$$\sqrt{k} \geq \frac{1}{2p} \left[\sqrt{8pn \log n} + \sqrt{8pn \log n + 4p(p+1)(n-1)} \right].$$

Now, $\sqrt{8pn \log n + 4p(p+1)(n-1)} = \sqrt{8pn \log n} + O\left(\sqrt{n/\log n} \right)$, so

$$N \leq \left(\frac{1}{2p} \left(2\sqrt{8pn \log n} + O\left(\sqrt{n/\log n} \right) \right) \right)^2 = \frac{8}{p} n \log n + O(n). \qquad \square$$

Lemma 3.3. *We have*
$$\mathbb{P} \left(\max_{1 \leq \beta \leq n-1} q(\beta) > 2n/p \right) < ne^{-n/4}.$$

Proof. We know by (3.1) that $\sum_{j=0}^{q(\beta)} X_j^{(\beta)} < n$. This implies that $q(\beta)$ is bounded above by a sum of n independent geometric random variables, each with expected value $1/p$. Setting $k = n$ and $\gamma = 2$ in Lemma 2.1, we find that

$$\mathbb{P}(q(\beta) > 2n/p) \leq e^{-n/4}.$$

The desired result now follows from taking a union bound over all $1 \leq \beta \leq n - 1$. $\qquad \square$

We can now finish the proof of the main result of this section.

Proof of Theorem 1.8. The lower bound in Theorem 1.8 is Lemma 3.1, so we just need to prove the upper bound.

Consider simulating $\mathbf{U}_{S_n}$ by creating the sequence of permutations (i.e., states) $\sigma_0, \ldots, \sigma_N$ (where $\sigma_0 = n(n-1)\cdots 1$ and $\sigma_N = 12\cdots n \neq \sigma_{N-1}$) as above. Preserve the notation from earlier. Every maximal chain in the weak order on S_n has $\binom{n}{2} + 1 < n^2/2$ elements. At each step during the Markov chain before we reach the bottom state $12\cdots n$, the probability of moving to a strictly lower state is at least p. Therefore, $N \leq \sum_{\ell=1}^{\binom{n}{2}} G_\ell$, where $G_1, \ldots, G_{\binom{n}{2}}$ are independent geometric random variables with parameter p. Appealing to Lemma 2.1, we find that

$$\sum_{m=\lceil n^2/p \rceil + 1}^{\infty} m\mathbb{P}(N = m) \leq \sum_{m=\lceil n^2/p \rceil}^{\infty} (m+1)\mathbb{P}(N > m)$$

$$\leq \sum_{m=\lceil n^2/p \rceil}^{\infty} (m+1)\mathbb{P}\left(\sum_{\ell=1}^{\binom{n}{2}} G_\ell > m \right)$$

$$\leq \sum_{m=\lceil n^2/p \rceil}^{\infty} (m+1) \exp\left(-\frac{pm}{2}\left(1 - \frac{1}{pm}\binom{n}{2}\right)^2\right)$$

$$\leq \sum_{m=\lceil n^2/p \rceil}^{\infty} (m+1) \exp\left(-pm/8\right)$$

$$= O\left(n^2 e^{-n^2/8}\right). \tag{3.5}$$

Let A be the event that $\max_{1 \leq \beta \leq n-1} q(\beta) \leq 2n/p$, and let A' be the event that

$$\sum_{j=a}^{b-1} X_j^{(\beta)} \geq p(b-a) - \sqrt{8p\log n} \cdot \sqrt{b-a}$$

for all $1 \leq \beta \leq n-1$ and $0 \leq a < b \leq q(\beta)+1$. For any particular $1 \leq \beta \leq n-1$ and $0 \leq a < b \leq q(\beta)+1$, we can use a Chernoff bound to see that

$$\mathbb{P}\left(\sum_{j=a}^{b-1} X_j^{(\beta)} < p(b-a) - \sqrt{8p\log n} \cdot \sqrt{b-a}\right) \leq n^{-4}.$$

Therefore,

$$\mathbb{P}(A \setminus A') \leq \sum_{1 \leq \beta \leq n-1} \sum_{0 \leq a < b \leq 2n/p+1} n^{-4} = O(n^{-1}).$$

Combined with Lemma 3.3, this shows that $\mathbb{P}(\neg(A \cap A')) = O(n^{-1})$. According to Lemma 3.2, we have $N \leq \frac{8}{p} n \log n + O(n)$ if $A \cap A'$ occurs. Consequently,

$$\sum_{m=1}^{\lceil n^2/p \rceil} m\mathbb{P}(N=m) = \sum_{m=1}^{\lceil n^2/p \rceil} m\mathbb{P}(A \cap A' \cap (N=m)) + \sum_{m=1}^{\lceil n^2/p \rceil} m\mathbb{P}(\neg(A \cap A') \cap (N=m))$$

$$\leq \mathbb{E}(N \mid (A \cap A')) + \lceil n^2/p \rceil \, \mathbb{P}(\neg(A \cap A'))$$

$$\leq \frac{8}{p} n \log n + O(n).$$

Combining this with (3.5) shows that $\mathcal{E}(S_n) = \sum_{m=1}^{\infty} m\mathbb{P}(N=m) = \frac{8}{p} n \log n + O(n)$, as desired. $\qquad\square$

4 Distributive lattices

Because Ungarian Markov chains on distributive lattices can be reformulated in terms of last-passage percolation with geometric weights—a well-studied topic in probability—we do not have too many new things to say about them. This short section is devoted to expounding upon some of the discussion from Section 1.4.

We first mention two immediate corollaries of the equation (1.1) that are not obvious from the original definition of Ungarian Markov chains.

Corollary 4.1. *If P' is a subposet of a poset P, then $\mathcal{E}(J(P')) \leq \mathcal{E}(J(P))$.*

Proof. This follows from (1.1) because every maximal chain of P' is contained in a maximal chain of P. $\qquad\square$

Corollary 4.2. *Let L be a distributive lattice, and let L^* be the dual of L. We have $\mathcal{E}(L^*) = \mathcal{E}(L)$.*

Proof. There is a poset P such that $L \cong J(P)$ and $L^* \cong J(P^*)$. The result follows from (1.1) because every maximal chain of P is equal (as a set) to a maximal chain of P^* and vice versa. $\square$

Remark 4.3. The hypothesis that L is distributive in Corollary 4.2 is crucial. For example, suppose L is the c-Cambrian lattice of type A_3, where $c = s_1 s_3 s_2$ (see Section 6.2 for definitions). Then L is both trim and semidistributive, but it is not distributive. We have $\mathcal{E}(L) \neq \mathcal{E}(L^*)$ when the probability p is generic.

We now prove Theorem 1.10.

Proof of Theorem 1.10. As in Section 1.4, consider a collection $(G_x)_{x \in P_n}$ of independent geometric random variables with parameter p. Let

$$\delta = \frac{1}{p} \left(\mu + \log \Gamma + \sqrt{2\mu \log \Gamma + (\log \Gamma)^2} \right).$$

Fix $\varepsilon > 0$, and let $\eta = \frac{(\delta + \varepsilon)p}{\mu}$. One can check that this choice of η guarantees

$$\frac{\eta \mu n}{2} \left(1 - \eta^{-1} \right)^2 > \log \Gamma,$$

so

$$\Gamma^{(1+o(1))n} e^{-\frac{\eta \mu n}{2} \left(1 - \eta^{-1} \right)^2} = o(1).$$

Suppose $\mathcal{C}$ is a maximal chain of P_n of size k. For each integer $m \geq \eta \mu n / p$, we can use Lemma 2.1 and the fact that $k \leq \mu n$ to compute

$$\mathbb{P} \left(\sum_{x \in \mathcal{C}} G_x > m \right) \leq e^{-\frac{pm}{2} \left(1 - \frac{k}{pm} \right)^2} \leq e^{-\frac{pm}{2} \left(1 - \frac{\mu n}{pm} \right)^2} \leq e^{-\frac{pm}{2} \left(1 - \eta^{-1} \right)^2}.$$

Consequently,

$$\sum_{m \geq \eta \mu n / p} \mathbb{P} \left(\sum_{x \in \mathcal{C}} G_x > m \right) \leq \sum_{m \geq \eta \mu n / p} e^{-\frac{pm}{2} \left(1 - \eta^{-1} \right)^2} = O \left(e^{-\frac{p(\eta \mu n / p)}{2} \left(1 - \eta^{-1} \right)^2} \right)$$
$$= O \left(e^{-\frac{\eta \mu n}{2} \left(1 - \eta^{-1} \right)^2} \right).$$

By a union bound, this implies that

$$\sum_{m \geq \eta \mu n / p} \mathbb{P} \left(\max_{\mathcal{C} \in \mathrm{MC}(P_n)} \sum_{x \in \mathcal{C}} G_x > m \right) = O \left(|\mathrm{MC}(P_n)| \, e^{-\frac{\eta \mu n}{2} \left(1 - \eta^{-1} \right)^2} \right)$$
$$= O \left(\Gamma^{(1+o(1))n} e^{-\frac{\eta \mu n}{2} \left(1 - \eta^{-1} \right)^2} \right)$$
$$= o(1).$$

Finally, using (1.1), we find that

$$\mathcal{E}(J(P_n)) = \mathbb{E} \left(\max_{\mathcal{C} \in \mathrm{MC}(P_n)} \sum_{x \in \mathcal{C}} G_x \right)$$
$$= \sum_{m \geq 0} \mathbb{P} \left(\max_{\mathcal{C} \in \mathrm{MC}(P_n)} \sum_{x \in \mathcal{C}} G_x > m \right)$$
$$= \sum_{0 \leq m < \eta \mu n / p} \mathbb{P} \left(\max_{\mathcal{C} \in \mathrm{MC}(P_n)} \sum_{x \in \mathcal{C}} G_x > m \right) + \sum_{m \geq \eta \mu n / p} \mathbb{P} \left(\max_{\mathcal{C} \in \mathrm{MC}(P_n)} \sum_{x \in \mathcal{C}} G_x > m \right)$$

$$\leq \frac{\eta\mu n}{p} + 1 + o(1).$$

Since $\frac{\eta\mu n}{p} = (\delta + \varepsilon)n$ and ε was arbitrary, this proves that $\mathcal{E}(J(P_n)) \leq \delta n + o(n)$, as desired. $\qquad\square$

5 Trim lattices

5.1 Basics of trim lattices

Let L be a lattice. An element $j \in L$ is called *join-irreducible* if it covers exactly one element of L; in this case, we write j_* for the unique element of L covered by j. Dually, an element $m \in L$ is called *meet-irreducible* if it is covered by exactly one element of L; in this case, we write m^* for the unique element of L that covers m. Let $\mathcal{J}_L$ and $\mathcal{M}_L$ be the set of join-irreducible elements of L and the set of meet-irreducible elements of L, respectively. For each $x \in L$, we have $x = \bigvee\{j \in \mathcal{J}_L : j \leq x\} = \bigwedge\{m \in \mathcal{M}_L : m \geq x\}$. This implies that the length of L is at most $|\mathcal{J}_L|$ and also at most $|\mathcal{M}_L|$; we say L is *extremal* if its length is equal to both $|\mathcal{J}_L|$ and $|\mathcal{M}_L|$.

An element x of a lattice L is called *left modular* if for all $y, z \in L$ with $y \leq z$, we have

$$(y \vee x) \wedge z = y \vee (x \wedge z).$$

We say L is *left modular* if it has a maximal chain whose elements are all left modular.

A lattice is *trim* if it is both extremal and left modular. Let us recall some facts about trim lattices from [33, 56, 57].

Suppose L is trim. For each join-irreducible element $j \in \mathcal{J}_L$, there is a unique meet-irreducible element $\kappa_L(j) \in \mathcal{M}_L$ such that

$$j \wedge \kappa_L(j) = j_* \quad \text{and} \quad j \vee \kappa_L(j) = (\kappa_L(j))^*. \tag{5.1}$$

The resulting map $\kappa_L \colon \mathcal{J}_L \to \mathcal{M}_L$ is a bijection. The *Galois graph* of L is the loopless directed graph $\mathbf{G}(L)$ with vertex set $\mathcal{J}_L$ such that for all distinct $j, j' \in \mathcal{J}_L$, there is an arrow $j \to j'$ in $\mathbf{G}(L)$ if and only if $j \not\leq \kappa_L(j')$. This graph is acyclic (i.e., it has no directed cycles), so we can define a partial order $\preceq$ on $\mathcal{J}_L$ by declaring that $j \preceq j'$ if there exists a directed path in $\mathbf{G}(L)$ from j' to j. We call the resulting poset $\mathbf{P}(L) = (\mathcal{J}_L, \preceq)$ the *Galois poset* of L.

Proposition 5.1 ([56, Theorem 1], [57, Proposition 3.13], [33, Theorem 6.2 & Corollary 7.10]). *Let L be a trim lattice. If $u, v \in L$ are such that $u \leq v$, then the interval $[u, v]$ of L is also a trim lattice. Moreover, $\mathbf{G}([u, v])$ is isomorphic to the subgraph of $\mathbf{G}(L)$ induced by $\{j \in \mathcal{J}_L : j \leq v, \kappa_L(j) \geq u\}$.*

5.2 Spines

Not all trim lattices are graded; in fact, a trim lattice is graded if and only if it is distributive [56, Theorem 2]. The *spine* of a trim lattice L, denoted $\mathrm{spine}(L)$, is the set of elements that belong to a maximum-length chain of L. According to [57, Proposition 2.6], the spine of L is a distributive sublattice of L, and

$$\mathrm{spine}(L) \cong J(\mathbf{P}(L)). \tag{5.2}$$

In this subsection, we will prove Theorem 1.11, which states that $\mathcal{E}(L) \leq \mathcal{E}(\mathrm{spine}(L))$. First, we need the following lemma.

Lemma 5.2. *Let L be a trim lattice, and let $x \in \mathrm{spine}(L)$. The interval $[\hat{0}, x]$ of L is a trim lattice whose spine is $\mathrm{spine}(L) \cap [\hat{0}, x]$.*

Proof. We know by Proposition 5.1 that $[\hat{0}, x]$ is trim. Suppose $y \in \text{spine}(L) \cap [\hat{0}, x]$. Then $y \leq x$ in $\text{spine}(L)$, so there exists a maximal chain $\mathcal{C}$ of $\text{spine}(L)$ that contains x and y. Since $\text{spine}(L)$ is distributive (hence, graded), $\mathcal{C}$ is a maximum-length chain of L. It follows that $\mathcal{C} \cap [\hat{0}, x]$ is a maximum-length chain of $[\hat{0}, x]$, so $y \in \text{spine}([\hat{0}, x])$.

We have shown that $\text{spine}(L) \cap [\hat{0}, x] \subseteq \text{spine}([\hat{0}, x])$. To prove the reverse containment, suppose $z \in \text{spine}([\hat{0}, x])$. Let $\mathcal{C}'$ be a maximum-length chain of $[\hat{0}, x]$ that contains z, and let $\mathcal{C}''$ be a maximum-length chain of L that contains x. Then $\mathcal{C}' \cup (\mathcal{C}'' \setminus [\hat{0}, x])$ is a maximum-length chain of L that contains z, so $z \in \text{spine}(L)$. Since z is certainly in $[\hat{0}, x]$, this proves that $\text{spine}([\hat{0}, x]) \subseteq \text{spine}(L) \cap [\hat{0}, x]$. $\qquad\square$

Proof of Theorem 1.11. Let L be a trim lattice. We will prove the inequality $\mathcal{E}(L) \leq \mathcal{E}(\text{spine}(L))$ by induction on $|L|$. This is trivial if $|L| = 1$, so we may assume $|L| \geq 2$.

To ease notation, let $K = \text{spine}(L)$ and $Q = \text{cov}_L(\hat{1})$. It is immediate from the definition of the spine that $\text{cov}_K(\hat{1}) = Q \cap K$. For $x \in L$ and $x' \in K$, we consider

$$\Delta_L(x) = \{y \in L : y \leq x\} \quad \text{and} \quad \Delta_K(x') = \{y \in K : y \leq x'\},$$

which are sublattices of L and K, respectively. It is immediate from the definition of the Ungarian Markov chain $\mathbf{U}_L$ that

$$\mathcal{E}(L)\left(1 - (1-p)^{|Q|}\right) - 1 = \sum_{\emptyset \neq T \subseteq Q} p^{|T|}(1-p)^{|Q|-|T|}\,\mathcal{E}(\Delta_L(\textstyle\bigwedge T)) = \mathcal{A}_1 + \mathcal{A}_2, \qquad (5.3)$$

where

$$\mathcal{A}_1 = \sum_{T' \subseteq Q \setminus K} p^{|T'|}(1-p)^{|Q \setminus K|-|T'|} \sum_{\emptyset \neq T'' \subseteq Q \cap K} p^{|T''|}(1-p)^{|Q \cap K|-|T''|}\,\mathcal{E}(\Delta_L(\textstyle\bigwedge(T' \cup T'')))$$

and

$$\mathcal{A}_2 = \sum_{\emptyset \neq T' \subseteq Q \setminus K} p^{|T'|}(1-p)^{|Q \setminus K|-|T'|}(1-p)^{|Q \cap K|}\,\mathcal{E}(\Delta_L(\textstyle\bigwedge T')).$$

We will prove that

$$\mathcal{E}(\Delta_L(\textstyle\bigwedge(T' \cup T''))) \leq \mathcal{E}(\Delta_K(\textstyle\bigwedge T'')) \qquad (5.4)$$

for all $T' \subseteq Q \setminus K$ and $T'' \subseteq Q \cap K$ such that $T' \cup T'' \neq \emptyset$. Because $\Delta_K(\bigwedge \emptyset) = \Delta_K(\hat{1}) = K$, this will imply that

$$\begin{aligned}
\mathcal{A}_1 &\leq \sum_{T' \subseteq Q \setminus K} p^{|T'|}(1-p)^{|Q \setminus K|-|T'|} \sum_{\emptyset \neq T'' \subseteq Q \cap K} p^{|T''|}(1-p)^{|Q \cap K|-|T''|}\,\mathcal{E}(\Delta_K(\textstyle\bigwedge T'')) \\
&= \sum_{\emptyset \neq T'' \subseteq Q \cap K} p^{|T''|}(1-p)^{|Q \cap K|-|T''|}\,\mathcal{E}(\Delta_K(\textstyle\bigwedge T'')) \\
&= \sum_{T'' \subseteq Q \cap K} p^{|T''|}(1-p)^{|Q \cap K|-|T''|}\,\mathcal{E}(\Delta_K(\textstyle\bigwedge T'')) - (1-p)^{|Q \cap K|}\,\mathcal{E}(K) \\
&= \mathcal{E}(K) - 1 - (1-p)^{|Q \cap K|}\mathcal{E}(K)
\end{aligned}$$

and

$$\begin{aligned}
\mathcal{A}_2 &\leq \sum_{\emptyset \neq T' \subseteq Q \setminus K} p^{|T'|}(1-p)^{|Q \setminus K|-|T'|}(1-p)^{|Q \cap K|}\mathcal{E}(K) \\
&= \sum_{T' \subseteq Q \setminus K} p^{|T'|}(1-p)^{|Q \setminus K|-|T'|}(1-p)^{|Q \cap K|}\mathcal{E}(K) - (1-p)^{|Q \setminus K|}(1-p)^{|Q \cap K|}\mathcal{E}(K) \\
&= (1-p)^{|Q \cap K|}\mathcal{E}(K) - (1-p)^{|Q|}\mathcal{E}(K),
\end{aligned}$$

so it will follow from (5.3) that

$$\mathcal{E}(L)\left(1 - (1-p)^{|Q|}\right) - 1 \le \mathcal{E}(K)\left(1 - (1-p)^{|Q|}\right) - 1,$$

which is equivalent to our desired inequality.

Fix $T' \subseteq Q \setminus K$ and $T'' \subseteq Q \cap K$ with $T' \cup T'' \ne \emptyset$. Let $L' = \Delta_L\left(\bigwedge(T' \cup T'')\right)$. Because L' is an interval in the trim lattice L, Proposition 5.1 tells us that L' is also trim. Let $K' = \mathrm{spine}(L')$. The assumption that $T' \cup T''$ is nonempty guarantees that $|L'| < |L|$, so we know by induction that $\mathcal{E}(L') \le \mathcal{E}(K')$. Hence, in order to prove (5.4), it suffices to show that $\mathcal{E}(K') \le \mathcal{E}(\Delta_K(\bigwedge T''))$.

Because $T'' \subseteq K$ and K is a sublattice of L, we have $\bigwedge T'' \in K$. According to Lemma 5.2, $\Delta_K(\bigwedge T'')$ is the spine of $\Delta_L(\bigwedge T'')$. Therefore, we know by (5.2) that

$$K' \cong J(\mathbf{P}(L')) \quad \text{and} \quad \Delta_K(\textstyle\bigwedge T'') \cong J(\mathbf{P}(\Delta_L(\textstyle\bigwedge T''))).$$

Now, L' is an interval of $\Delta_L(\bigwedge T'')$, so Proposition 5.1 tells us that $\mathbf{G}(L')$ is isomorphic to an induced subgraph of $\mathbf{G}(\Delta_L(\bigwedge T''))$. This implies that $\mathbf{P}(L')$ is a subposet of $\mathbf{P}(\Delta_L(\bigwedge T''))$, so the desired inequality $\mathcal{E}(K') \le \mathcal{E}(\Delta_K(\bigwedge T''))$ follows from Corollary 4.1. $\qquad\square$

Proof of Corollary 1.12. Let L be a trim lattice, and suppose L' is an interval or a quotient of L. Then L' is a trim lattice whose Galois graph $\mathbf{G}(L')$ is isomorphic to an induced subgraph of $\mathbf{G}(L)$; this follows from Proposition 5.1 if L' is an interval of L, and it follows from [57, Lemma 3.10 & Remark 3.11] if L' is a quotient of L. This implies that $\mathbf{P}(L')$ is a subposet of $\mathbf{P}(L)$, so we can use (5.2) and Corollary 4.1 to see that $\mathcal{E}(\mathrm{spine}(L')) \le \mathcal{E}(\mathrm{spine}(L))$. On the other hand, we can apply Theorem 1.11 to L' to obtain the inequality $\mathcal{E}(L') \le \mathcal{E}(\mathrm{spine}(L'))$. $\qquad\square$

6 Cambrian lattices

In this section, we review relevant notions related to Reading's Cambrian lattices, and we analyze the spines of Cambrian lattices in order to prove Theorems 1.13 to 1.15. We also give a simple proof of Theorem 1.18.

6.1 Background

Let W be a finite *Coxeter group*, and let S be the set of *simple reflections* of W; this means that W has a presentation of the form $\langle S : (ss')^{m(s,s')} = e \rangle$, where e is the identity element of W, $m(s,s) = 1$ for all $s \in S$, and $m(s,s') = m(s',s) \in \{2, 3, \ldots\} \cup \{\infty\}$ for all distinct $s, s' \in S$. Note that each simple reflection is an involution. The *Coxeter graph* of W is the graph with vertex set S in which two simple reflections s and s' are connected by an edge whenever $m(s,s') \ge 3$; this edge is labeled with the number $m(s,s')$ if $m(s,s') \ge 4$. We will assume that W is *irreducible*, which means that its Coxeter graph is connected (equivalently, W cannot be expressed as a direct product of smaller Coxeter groups).

When we refer to a *word over S*, we mean a (possibly infinite) word whose letters are in S, where we view any two of the letters of the word as distinct from one another (even if they are copies of the same simple reflection). We can apply a *commutation move* to a word over S by swapping two consecutive letters s and s' if $m(s,s') = 2$ (we do not allow such a commutation move if $m(s,s') = 1$). The *commutation class* of a word Q over S is the set of words that can be obtained from Q via a sequence of commutation moves. Finite words in the same commutation class represent the same element of W. Following Viennot [59], we define a certain poset $\mathrm{Heap}(Q)$ called the *heap* of Q. The elements of this poset are the letters in Q (which are seen as distinct from one another).

The order relation is defined so that if s and s' are two letters (which could represent the same simple reflection), then $s < s'$ if and only if s appears to the left of s' in every word in the commutation class of Q. The Hasse diagram of $\mathrm{Heap}(Q)$, which we will also denote by $\mathrm{Heap}(Q)$, is called the *combinatorial AR quiver* of Q [54, Chapter 9] and is typically drawn sideways so that each cover relation $s \lessdot s'$ is depicted with s to the left of s'; see subsequent subsections for several such situations.

A *reduced word* for an element $w \in W$ is a word over S that represents w and has minimum length among all such words. The (right) *weak order* on W is the partial order on W defined by declaring $u \leq v$ if there is a reduced word for v that contains a reduced word for u as a prefix. We have assumed that W is finite, so a seminal result due to Björner [14] states that the weak order on W is a lattice; the maximal element of this lattice is called the *long element* of W and is typically denoted w_0.

Let $s_{i_1}, \ldots, s_{i_n}$ be an ordering of the simple reflections of W. Let c be the word $s_{i_1} \cdots s_{i_n}$. The element c of W represented by c is called a *Coxeter element* of W. Two words c and c′ represent the same Coxeter element if and only if $\mathrm{Heap}(\mathsf{c}) = \mathrm{Heap}(\mathsf{c'})$. Let us orient each edge $\{s, s'\}$ in the Coxeter graph from s to s' if s appears to the left of s' in c (equivalently, $s \leq s'$ in $\mathrm{Heap}(\mathsf{c})$), and let us orient it from s' to s if s appears to the right of s' in c (equivalently, $s \geq s'$ in $\mathrm{Heap}(\mathsf{c})$). This orientation depends only on c. In fact, this establishes a one-to-one correspondence between Coxeter elements of W and acyclic orientations of the Coxeter graph of W.

Let $\mathsf{c}^\infty = \mathsf{c}^{\langle 1 \rangle} \mathsf{c}^{\langle 2 \rangle} \cdots$, where each word $\mathsf{c}^{\langle i \rangle}$ is a copy of c. Following Reading [44], we define the *c-sorting word* of an element $w \in W$, denoted $\mathrm{sort}_\mathsf{c}(w)$, to be the reduced word for w that is lexicographically first as a subword of c^∞. We can write $\mathrm{sort}_\mathsf{c}(w) = \mathsf{w}^{\langle 1 \rangle} \mathsf{w}^{\langle 2 \rangle} \cdots$, where $\mathsf{w}^{\langle i \rangle}$ is the subword of $\mathrm{sort}_\mathsf{c}(w)$ that came from $\mathsf{c}^{\langle i \rangle}$ when we found $\mathrm{sort}_\mathsf{c}(w)$ as a lexicographically minimum subword of c^∞. Let $\mathrm{supp}(\mathsf{w}^{\langle i \rangle})$ be the set of simple reflections appearing in $\mathsf{w}^{\langle i \rangle}$. We say w is *c-sortable* if we have the chain of containments $\mathrm{supp}(\mathsf{w}^{\langle 1 \rangle}) \supseteq \mathrm{supp}(\mathsf{w}^{\langle 2 \rangle}) \supseteq \mathrm{supp}(\mathsf{w}^{\langle 3 \rangle}) \supseteq \cdots$. Whether or not w is c-sortable depends only on the Coxeter element c and not on the word c [44].

The set of c-sortable elements of W forms a sublattice (and a quotient lattice) Camb_c of the weak order on W called the *c-Cambrian lattice*. Moreover, Camb_c is trim [56, 37]. Thomas and Williams [57] provided a description of the Galois graph $\mathbf{G}(\mathrm{Camb}_c)$, from which one can deduce that the Galois poset of the c-Cambrian lattice is isomorphic to the heap of the c-sorting word of the long element of W:

$$\mathbf{P}(\mathrm{Camb}_c) \cong \mathrm{Heap}(\mathrm{sort}_\mathsf{c}(w_0)). \tag{6.1}$$

All Coxeter elements of W are conjugate to each other, so they all have the same group-theoretic order h, which is called the *Coxeter number* of W. Let c^h be the word obtained by concatenating c with itself h times. The long element w_0 is an involution in W, so it gives rise to an involution $\psi \colon S \to S$ defined by $\psi(s) = w_0 s w_0$. We can extend ψ to a map defined on words over S by letting $\psi(s_{i_1} \cdots s_{i_k}) = \psi(s_{i_1}) \cdots \psi(s_{i_k})$ (and similarly for infinite words). It follows from [54, Lemma 2.6.5] that

$$\mathrm{sort}_\mathsf{c}(w_0) \psi(\mathrm{sort}_\mathsf{c}(w_0)) = \mathsf{c}^h. \tag{6.2}$$

6.2 Cambrian lattices of type A

The purpose of this subsection is to prove Theorem 1.13, which provides an asymptotic upper bound for $\mathcal{E}(L)$ when L is a (large) Cambrian lattice of type A.

The Coxeter group A_n is the symmetric group S_{n+1}; its simple reflections are $s_1, \ldots, s_n$, where s_i is the transposition $(i \ i+1)$. The Coxeter graph of A_n is a path with edges $\{s_i, s_{i+1}\}$ for $i \in [n-1]$, and each of these edges is unlabeled (i.e., $m(s_i, s_{i+1}) = 3$). The Coxeter number of A_n is $n+1$.

Fix a reduced word c of a Coxeter element c of A_n. The isomorphisms (5.2) and (6.1) tell us that the spine of Camb_c is isomorphic to $J(\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0)))$, so we wish to describe the combinatorial AR quiver of $\mathsf{sort}_\mathsf{c}(w_0)$. We refer the reader to Example 6.1 and Figures 3 and 4 for concrete illustrations the following discussion.

Begin by drawing $\mathrm{Heap}(\mathsf{c}^{n+1})$ by drawing $n + 1$ copies of $\mathrm{Heap}(\mathsf{c})$ in a row and adding in edges as appropriate so that the result has the shape of a chain-link fence. We can coordinatize this drawing in the xy-plane so that the leftmost point has x-coordinate 0, each edge extends 1 unit horizontally and 1 unit vertically, and all of the vertices that are copies of s_i lie on the line $y = i$. Draw a path that has the same shape as an upside-down version of $\mathrm{Heap}(\mathsf{c})$, and use this path to cut some of the edges in $\mathrm{Heap}(\mathsf{c}^{n+1})$ so that exactly half of the $n(n + 1)$ vertices are to the left of the cut (the remainder of this paragraph will imply that this is indeed possible). The involution $\psi\colon S \to S$ is given by $\psi(s_i) = s_{n+1-i}$. This implies that $\mathrm{Heap}(\psi(\mathsf{sort}_\mathsf{c}(w_0)))$ has the same "shape" as $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$, except that it is flipped upside-down, and (6.2) tells us that if we place $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$ immediately to the left of $\mathrm{Heap}(\psi(\mathsf{sort}_\mathsf{c}(w_0)))$ and add edges as appropriate, we will obtain $\mathrm{Heap}(\mathsf{c}^{n+1})$. Hence, the path that we drew to make the cut has the same shape as $\mathrm{Heap}(\psi(\mathsf{c}))$, and it follows from (6.2) that the combinatorial AR quiver $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$ is the part of the graph on the left of the cut.

In our drawing of $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$, if $(a, 1)$ and (b, n) are the leftmost points in the bottom and top rows, respectively, then $(n - 1 + b, 1)$ and $(n - 1 + a, n)$ are the rightmost points in the bottom and top rows, respectively. It follows that this drawing of $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$ lies within the square $\mathscr{R}_\mathsf{c}$ whose sides lie on the lines $y = x + n - b$, $y = x - n + 2 - b$, $y = -x + a + 1$, and $y = -x + 2n - 1 + a$.

Example 6.1. Suppose $n = 9$. Let c be the word $s_3 s_2 s_1 s_4 s_5 s_7 s_6 s_8 s_9$, and let $c \in A_9$ be the element represented by c. Figure 3 shows how to arrange $n + 1 = 10$ copies of $\mathrm{Heap}(\mathsf{c})$ (drawn in different colors) and add edges (drawn in black) to obtain $\mathrm{Heap}(\mathsf{c}^{10})$. The thick black-and-blue path has the same shape as $\mathrm{Heap}(\psi(\mathsf{c}))$ and cuts $\mathrm{Heap}(\mathsf{c}^{10})$ into two halves; the left half is $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$. Indeed, the c-sorting word for w_0 is

$$s_3 s_2 s_1 s_4 s_5 s_7 s_6 s_8 s_9 s_3 s_2 s_1 s_4 s_5 s_7 s_6 s_8 s_9 s_3 s_2 s_1 s_4 s_5 s_7 s_6 s_8 s_9 s_3 s_2 s_1 s_4 s_5 s_7 s_6 s_8 s_9 s_3 s_2 s_1 s_4 s_5$$
$$s_3 s_2 s_1 s_4.$$

In the notation from above, we have $a = 2$ and $b = 4$. Figure 4 shows that the square $\mathscr{R}_\mathsf{c}$ (drawn dotted) whose sides lie on the lines $y = x + 5$, $y = x - 11$, $y = -x + 3$, and $y = -x + 19$ contains our drawing of $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$.

The fact that our drawing of $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$ fits inside the square $\mathscr{R}_\mathsf{c}$ readily implies that $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$ is a subposet of the $n \times n$ rectangle poset $R_{n \times n}$. It follows from Theorem 1.9 and Corollary 4.1 that

$$\mathcal{E}(J(\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0)))) \leq \mathcal{E}(J(R_{n \times n})) = \frac{1}{p}\left(2 + 2\sqrt{1 - p}\right) n + o(n).$$

On the other hand, (5.2) and (6.1) tell us that $J(\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0)))$ is isomorphic to the spine of Camb_c. Invoking Theorem 1.11, we find that

$$\mathcal{E}(\mathrm{Camb}_c) \leq \frac{1}{p}\left(2 + 2\sqrt{1 - p}\right) n + o(n),$$

and this proves Theorem 1.13.

6.3 Cambrian lattices of type B

This subsection is devoted to proving Theorem 1.14, which provides an asymptotic upper bound for $\mathcal{E}(L)$ when L is a (large) Cambrian lattice of type B.
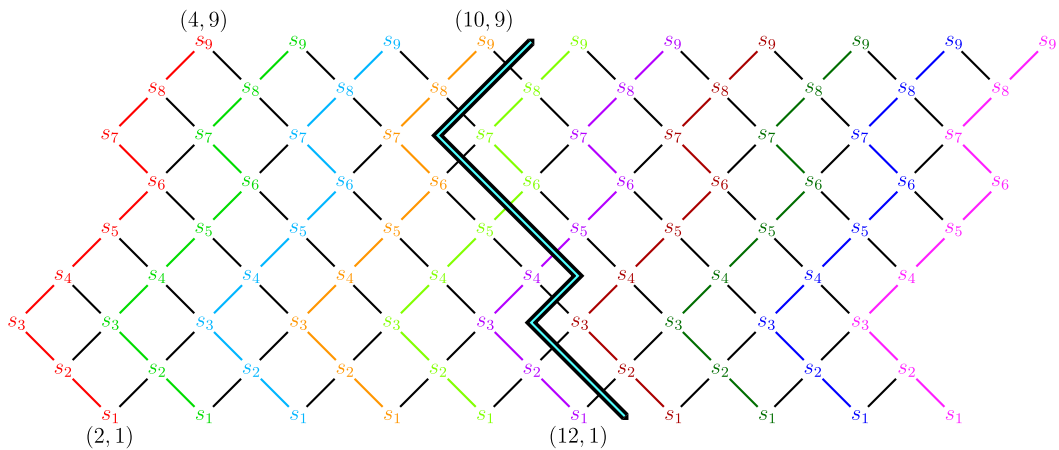
Figure 3: A drawing of the (type A) combinatorial AR quiver $\mathrm{Heap}(\mathsf{c}^{10})$, where $\mathsf{c} = s_3 s_2 s_1 s_4 s_5 s_7 s_6 s_8 s_9$ is the word from Example 6.1. The thick black-and-blue path cuts $\mathrm{Heap}(\mathsf{c}^{10})$ into a left half, which is $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$, and a right half, which is $\mathrm{Heap}(\psi(\mathsf{sort}_\mathsf{c}(w_0)))$.
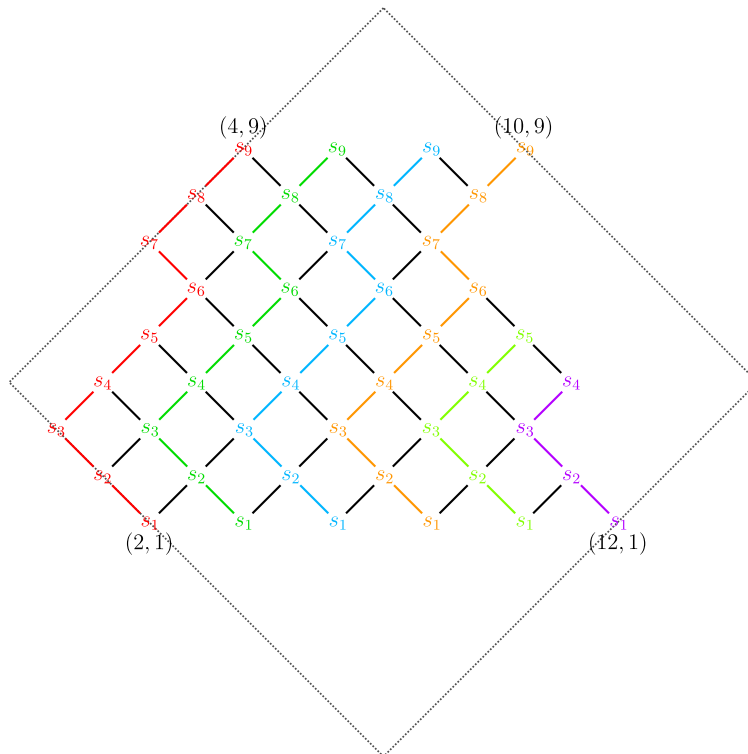
Figure 4: The (type A) combinatorial AR quiver $\mathrm{Heap}(\mathsf{sort}_\mathsf{c}(w_0))$ from Example 6.1 fits inside the square $\mathscr{R}_\mathsf{c}$.

The Coxeter group B_n is the n-th *hyperoctahedral group*; it has simple reflections $s_0, s_1, \ldots, s_{n-1}$. The Coxeter graph of B_n is a path with edges $\{s_i, s_{i+1}\}$ for $i \in \{0, \ldots, n-2\}$. The label of the edge $\{s_0, s_1\}$ is $m(s_0, s_1) = 4$, while each edge $\{s_i, s_{i+1}\}$ with $1 \leq i \leq n-2$ is unlabeled (i.e., $m(s_i, s_{i+1}) = 3$). The Coxeter number of B_n is $2n$.

Fix a reduced word c of a Coxeter element c of B_n. As discussed in Section 6.1, there is a unique orientation of the Coxeter graph of B_n corresponding to c; let $r(c)$ be the number of edges $\{s_i, s_{i+1}\}$ that are oriented from s_i to s_{i+1} in this orientation.

Finding the c-sorting word for w_0 is actually quite simple because the involution $\psi \colon S \to S$ is the identity map. Therefore, it follows from (6.2) that $\mathsf{sort}_c(w_0) = c^{h/2} = c^n$. To draw $\mathrm{Heap}(c^n)$, simply draw n copies of $\mathrm{Heap}(c)$ in a row and add in edges as appropriate so that the result has the shape of a chain-link fence. Let us coordinatize this drawing in the xy-plane so that the leftmost point has x-coordinate 0, each edge extends 1 unit horizontally and 1 unit vertically, and all of the vertices that are copies of s_i lie on the line $y = i$. It is a straightforward exercise to show that there exist integers q_1 and q_2 such that this drawing of $\mathrm{Heap}(c^n)$ lies in the rectangle $\mathscr{R}_c$ whose sides lie on the lines $y = x + q_1$, $y = x + q_1 - 4n + 4 + 2r(c)$, $y = -x + q_2$, and $y = -x + q_2 - 2n + 2 - 2r(c)$.

Example 6.2. Suppose $n = 7$. Let $c = s_1 s_0 s_2 s_3 s_5 s_4 s_6$, and let $c \in B_7$ be the element represented by c. The orientation of the Coxeter graph of B_7 corresponding to c is

$$s_0 \overset{4}{\twoheadleftarrow} s_1 \twoheadrightarrow s_2 \twoheadrightarrow s_3 \twoheadrightarrow s_4 \twoheadleftarrow s_5 \twoheadrightarrow s_6 \ .$$

Figure 5 shows how to arrange 7 copies of $\mathrm{Heap}(c)$ (drawn in different colors) and add edges (drawn in black) to obtain $\mathrm{Heap}(c^7) = \mathrm{Heap}(\mathsf{sort}_c(w_0))$. In the notation from above, we have $r(c) = 4$, $q_1 = 4$ and $q_2 = 22$. Figure 5 shows that the square $\mathscr{R}_c$ (drawn dotted) whose sides lie on the lines $y = x + 4$, $y = x - 12$, $y = -x + 22$, and $y = -x + 2$ contains our drawing of $\mathrm{Heap}(\mathsf{sort}_c(w_0))$.

As in the statement of Theorem 1.14, let us now choose a sequence $(c^{(n)})_{n \geq 2}$, where each $c^{(n)}$ is a Coxeter element of B_n, and let us assume that the limit $\overline{r} = \lim\limits_{n \to \infty} \frac{1}{n} r(c^{(n)})$ exists. For each n, let $c^{(n)}$ be a reduced word for $c^{(n)}$. Because our drawing of $\mathrm{Heap}(\mathsf{sort}_{c^{(n)}}(w_0))$ fits inside the rectangle $\mathscr{R}_{c^{(n)}}$, $\mathrm{Heap}(\mathsf{sort}_{c^{(n)}}(w_0))$ is a subposet of the $(2n - 1 - r(c^{(n)})) \times (n + r(c^{(n)}))$ rectangle poset $R_{(2n-1-r(c^{(n)})) \times (n+r(c^{(n)}))}$. It follows from Theorem 1.9 and Corollary 4.1 that

$$\mathcal{E}(J(\mathrm{Heap}(\mathsf{sort}_{c^{(n)}}(w_0)))) \leq \mathcal{E}\big(J\big(R_{(2n-1-r(c^{(n)})) \times (n+r(c^{(n)}))}\big)\big)$$
$$= \frac{1}{p} \left(3 + 2\sqrt{(1-p)(2-\overline{r})(1+\overline{r})} \right) n + o(n).$$

On the other hand, (5.2) and (6.1) tell us that $J(\mathrm{Heap}(\mathsf{sort}_{c^{(n)}}(w_0)))$ is isomorphic to the spine of $\mathrm{Camb}_{c^{(n)}}$. Invoking Theorem 1.11, we find that

$$\mathcal{E}(\mathrm{Camb}_{c^{(n)}}) \leq \frac{1}{p} \left(3 + 2\sqrt{(1-p)(2-\overline{r})(1+\overline{r})} \right) n + o(n),$$

and this proves Theorem 1.14.

6.4 Cambrian lattices of type D

In this subsection, we prove Theorem 1.15, which provides an asymptotic upper bound for $\mathcal{E}(L)$ when L is a (large) Cambrian lattice of type D.

The Coxeter group D_n has simple reflections $s_0, s_1, \ldots, s_{n-1}$. The Coxeter graph of D_n has unlabeled edges $\{s_i, s_{i+1}\}$ for $i \in [n-2]$ together with the additional unlabeled edge $\{s_0, s_2\}$. The Coxeter number of D_n is $2n - 2$.

Consider a reduced word c of a Coxeter element c of D_n. As discussed in Section 6.1, there is a unique orientation of the Coxeter graph of D_n corresponding to c; let $r(c)$ be the number of edges of the form $\{s_i, s_{i+1}\}$ with $1 \leq i \leq n-2$ that are oriented from s_i to s_{i+1} in this orientation (so the definition of $r(c)$ ignores the edge $\{s_0, s_2\}$).

If n is even, the involution $\psi \colon S \to S$ is the identity map, but if n is odd, this involution is the transposition that fixes $s_2, s_3, \ldots, s_{n-1}$ and swaps s_0 and s_1. It is certainly possible
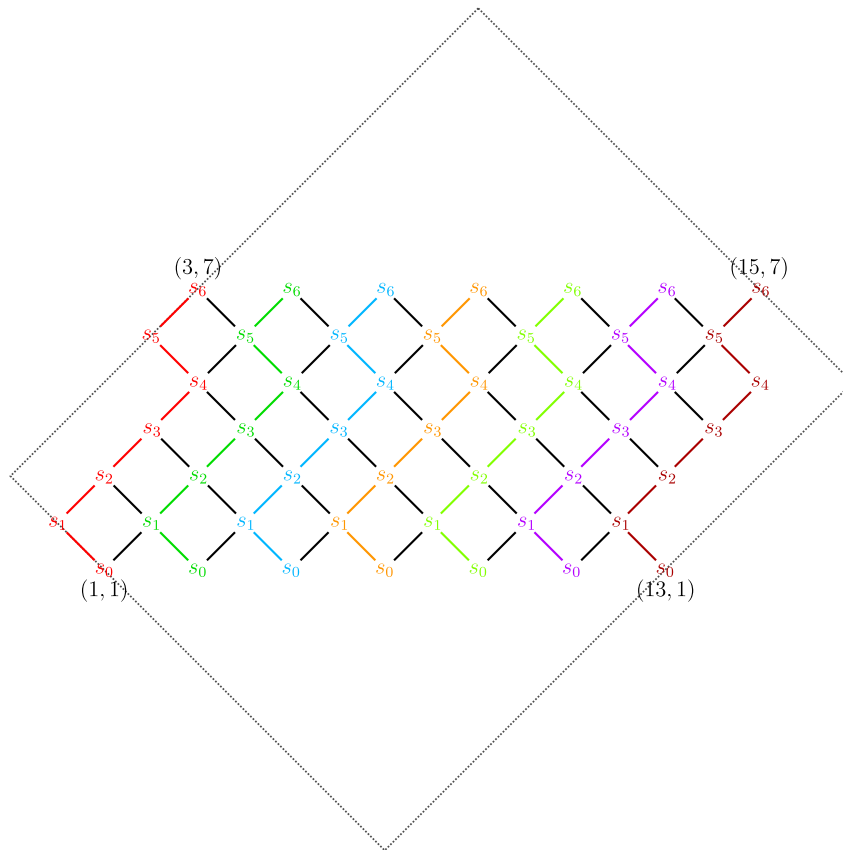
Figure 5: The (type B) combinatorial AR quiver $\mathrm{Heap}(\mathrm{sort}_{\mathsf{c}}(w_0)) = \mathrm{Heap}(\mathsf{c}^7)$, where $\mathsf{c} = s_1 s_0 s_2 s_3 s_5 s_4 s_6$ is the word from Example 6.2. This drawing fits inside the rectangle $\mathscr{R}_{\mathsf{c}}$.

to prove Theorem 1.15 by describing $\mathrm{sort}_{\mathsf{c}}(w_0)$ in a manner that depends on the parity of n. However, because Theorem 1.15 is only concerned with asymptotics, we can take a shortcut by reducing our analysis to the case when n is even. We do so via the following lemma, which makes tacit use of the standard identification of D_n with the subgroup of D_{n+1} generated by $s_0, s_1, \ldots, s_{n-1}$.

Lemma 6.3. *Suppose $n \geq 4$. Let c be a Coxeter element of D_n, and let c' be the Coxeter element $s_n c$ of D_{n+1}. Then $r(c) = r(c')$, and*

$$\mathcal{E}(\mathrm{Camb}_c) \leq \mathcal{E}(\mathrm{spine}(\mathrm{Camb}_{c'})).$$

Proof. The identity $r(c) = r(c')$ is immediate from the relevant definitions. Denote by $w_0(\{s_0, \ldots, s_{n-1}\})$ the long element of D_n, seen as an element of D_{n+1}; it is known that this element is c'-sortable. It is straightforward to show that an element $v \in D_{n+1}$ is a c-sortable element of D_n if and only if v is c'-sortable and $v \leq w_0(\{s_0, \ldots, s_{n-1}\})$ in the weak order on D_{n+1}. Consequently, Camb_c is the interval $[e, w_0(\{s_0, \ldots, s_{n-1}\})]$ of $\mathrm{Camb}_{c'}$. The desired inequality now follows from Corollary 1.12. $\qquad\square$

Let us now assume n is even so that ψ is the identity map. Fix a reduced word c of a Coxeter element $c \in D_n$. It follows from (6.2) that $\mathrm{sort}_{\mathsf{c}}(w_0) = \mathsf{c}^{h/2} = \mathsf{c}^{n-1}$. We will draw $\mathrm{Heap}(\mathsf{c}^{n-1})$ in the xy-plane as follows (see Example 6.4 and Figure 6 for an illustration). Draw $n - 1$ copies of $\mathrm{Heap}(\mathsf{c})$ in a row and add in edges as appropriate. For each $i \in [n-1]$, draw all of the vertices that are copies of s_i on the line $y = i$. Also,

draw all of the vertices that are copies of s_0 on the line $y = 2$. Assume the leftmost point in the drawing has x-coordinate 0. Draw each edge that does not include s_0 as a line segment that extends 1 unit horizontally and 1 unit vertically, and draw each edge that does include s_0 as a horizontal line segment of length 1.

As in Sections 6.2 and 6.3, we wish to find a rectangle $\mathscr{R}_c$ that encloses our drawing of the combinatorial AR quiver $\mathrm{Heap}(\mathsf{sort}_c(w_0)) = \mathrm{Heap}(\mathsf{c}^{n-1})$. Describing this rectangle is a bit cumbersome because of the presence of vertices that are copies of s_0. However, these vertices can only increase the length and width of the rectangle by a small amount that will be absorbed by the error term in our asymptotic results anyway. To be more precise, consider the orientation of the Coxeter graph of D_n corresponding to c. Let $\epsilon_1 = 1$ if one of the following conditions holds:

- $\{s_0, s_2\}$ is oriented from s_2 to s_0, and $\{s_1, s_2\}$ is oriented from s_1 to s_2;
- $\{s_0, s_2\}$ is oriented from s_0 to s_2, and $\{s_i, s_{i+1}\}$ is oriented from s_i to s_{i+1} for all $2 \leq i \leq n-2$;

otherwise, let $\epsilon_1 = 0$. Let $\epsilon_2 = 1$ if one of the following conditions holds:

- $\{s_0, s_2\}$ is oriented from s_0 to s_2, and $\{s_1, s_2\}$ is oriented from s_2 to s_1;
- $\{s_0, s_2\}$ is oriented from s_2 to s_0, and $\{s_i, s_{i+1}\}$ is oriented from s_{i+1} to s_i for all $2 \leq i \leq n-2$;

otherwise, let $\epsilon_2 = 0$. Then there are integers q_1 and q_2 such that the sides of $\mathscr{R}_c$ lie on the lines $y = x + q_1$, $y = x + q_1 - 4n + 8 + 2r(c) - \epsilon_1$, $y = -x + q_2$, and $y = -x + q_2 - 2n + 4 - 2r(c) - \epsilon_2$.

Example 6.4. Suppose $n = 6$. Let $\mathsf{c} = s_0 s_3 s_2 s_1 s_5 s_4$, and let $c \in D_6$ be the element represented by c. The orientation of the Coxeter graph of D_6 corresponding to c is
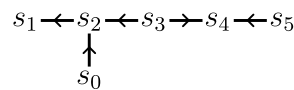
$$s_1 \leftarrow s_2 \leftarrow s_3 \rightarrow s_4 \leftarrow s_5$$
$$\uparrow$$
$$s_0$$

.

Figure 6 shows how to arrange 5 copies of $\mathrm{Heap}(\mathsf{c})$ (drawn in different colors) and add edges (drawn in black) to obtain $\mathrm{Heap}(\mathsf{c}^5) = \mathrm{Heap}(\mathsf{sort}_c(w_0))$. In the notation from above, we have $\epsilon_1 = 0$, $\epsilon_2 = 1$, $r(c) = 1$, $q_1 = 5$, and $q_2 = 13$. Figure 6 shows that the rectangle $\mathscr{R}_c$ (drawn dotted) whose sides lie on the lines $y = x + 5$, $y = x - 9$, $y = -x + 13$, and $y = -x + 2$ contains our drawing of $\mathrm{Heap}(\mathsf{sort}_c(w_0))$. Figure 6 also uses additional diagonal grid lines to illustrate how $\mathrm{Heap}(\mathsf{sort}_c(w_0))$ is a subposet of the rectangle poset $R_{(4n-7-2r(c)+\epsilon_1) \times (2n-3+2r(c)+\epsilon_2)} = R_{15 \times 12}$.

As illustrated in Figure 6, the fact that our drawing of $\mathrm{Heap}(\mathsf{sort}_c(w_0))$ fits inside $\mathscr{R}_c$ allows us to see that $\mathrm{Heap}(\mathsf{sort}_c(w_0))$ is a subposet of the rectangle poset

$$R_{(4n-7-2r(c)+\epsilon_1) \times (2n-3+2r(c)+\epsilon_2)}.$$

Hence, $\mathrm{Heap}(\mathsf{sort}_c(w_0))$ is a subposet of $R_{(4n-2r(c)) \times (2n+2r(c))}$.

As in the statement of Theorem 1.15, let us now choose a sequence $(c^{(n)})_{n \geq 4}$, where each $c^{(n)}$ is a Coxeter element of D_n. For each n, let $\mathsf{c}^{(n)}$ be a reduced word for $c^{(n)}$. Let us assume that the limits $\overline{r} = \lim_{n \to \infty} \frac{1}{n} r(c^{(n)})$ and $\overline{u} = \lim_{n \to \infty} \frac{1}{n} u(c^{(n)})$ exist, where $u(c^{(n)})$ is the maximum number of edges in a directed path in the orientation of the Coxeter graph of D_n corresponding to $c^{(n)}$ (equivalently, $u(c^{(n)})$ is the length of the poset $\mathrm{Heap}(\mathsf{c}^{(n)})$). When n is even, we have seen that $\mathrm{Heap}(\mathsf{sort}_{c^{(n)}}(w_0))$ is a subposet of $R_{(4n-2r(c^{(n)})) \times (2n+2r(c^{(n)}))}$; it follows from Theorem 1.9 and Corollary 4.1 that

$$\mathcal{E}(J(\mathrm{Heap}(\mathsf{sort}_{c^{(n)}}(w_0)))) \leq \mathcal{E}(J(R_{(4n-2r(c^{(n)})) \times (2n+2r(c^{(n)}))}))$$
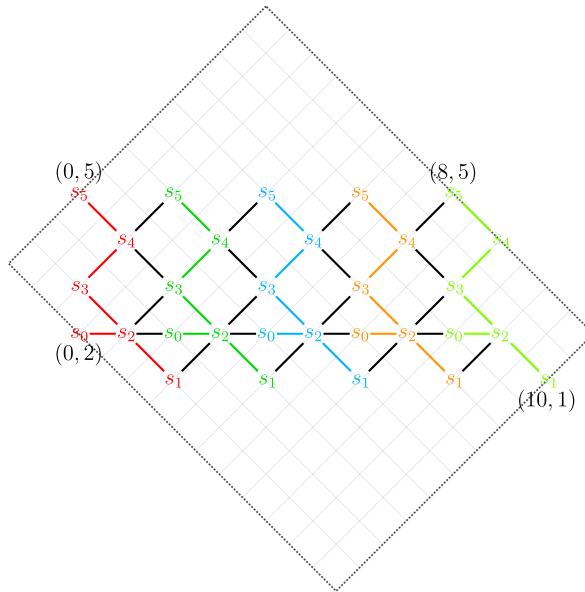
Figure 6: The (type D) combinatorial AR quiver $\text{Heap}(\text{sort}_\mathsf{c}(w_0)) = \text{Heap}(\mathsf{c}^5)$, where $\mathsf{c} = s_0 s_3 s_2 s_1 s_5 s_4$ is the word from Example 6.4. This drawing fits inside the rectangle $\mathscr{R}_\mathsf{c}$. The extra diagonal grid lines serve to reveal how $\text{Heap}(\text{sort}_\mathsf{c}(w_0))$ is a subposet of the rectangle poset $R_{15 \times 12}$.

$$= \frac{1}{p} \left(6 + 4\sqrt{(1-p)(2-\overline{r})(1+\overline{r})} \right) n + o(n).$$

We also know by (5.2) and (6.1) that $J(\text{Heap}(\text{sort}_{c^{(n)}}(w_0)))$ is isomorphic to the spine of $\text{Camb}_{c^{(n)}}$. Invoking Theorem 1.11, we find that

$$\mathcal{E}(\text{Camb}_{c^{(n)}}) \leq \frac{1}{p} \left(6 + 4\sqrt{(1-p)(2-\overline{r})(1+\overline{r})} \right) n + o(n)$$

for even n. Using Lemma 6.3, we find that the same inequality holds for odd n.

To complete the proof of Theorem 1.15, we must show that

$$\mathcal{E}(\text{Camb}_{c^{(n)}}) \leq \frac{1}{p} \left(2 + \overline{u} + \log\left(5 \cdot 2^{\overline{u}}\right) + \sqrt{2(2+\overline{u})\log\left(5 \cdot 2^{\overline{u}}\right) + \left(\log\left(5 \cdot 2^{\overline{u}}\right)\right)^2} \right) n + o(n).$$

Because $J(\text{Heap}(\text{sort}_\mathsf{c}(w_0))) \cong \text{spine}(\text{Camb}_{c^{(n)}})$, we can do so by appealing to Theorems 1.10 and 1.11. For even n, we need to prove that the maximum size of a chain in $\text{Heap}(\text{sort}_{c^{(n)}}(w_0))$ is at most $(2 + \overline{u} + o(1))n$ and that

$$|\text{MC}(\text{Heap}(\text{sort}_{c^{(n)}}(w_0)))| \leq \left(5 \cdot 2^{\overline{u}}\right)^{(1+o(1))n};$$

the result for odd n will then follow from Lemma 6.3. Suppose n is even. The statement about the maximum size of a chain is immediate; indeed, it follows from the above discussion that the maximum size of a chain in $\text{Heap}(\text{sort}_{c^{(n)}}(w_0))$ is $2n - 3 + u(c^{(n)})$. Now imagine constructing a maximal chain $\mathcal{C}$ in $\text{Heap}(\text{sort}_{c^{(n)}}(w_0))$ that uses exactly k copies of s_0. Specifying these k copies will also determine $2k - 1$ or $2k$ of the edges used in $\mathcal{C}$; we must then choose the remaining edges. The number of remaining edges is at most $2n - 3 + u(c^{(n)}) - 2k$. If we choose these remaining edges one-by-one, starting at the left of the combinatorial AR quiver $\text{Heap}(\text{sort}_{c^{(n)}}(w_0))$, then we will have at most 2

choices for each such edge. This shows that the total number of choices for $\mathcal{C}$ is at most $\binom{n-1}{k}2^{2n-3+u(c^{(n)})-2k}$. Consequently,

$$
\begin{aligned}
|\mathrm{MC}(\mathrm{Heap}(\mathsf{sort}_{\mathsf{c}^{(n)}}(w_0)))| &\leq \sum_{k=0}^{n-1}\binom{n-1}{k}2^{2n-3+u(c^{(n)})-2k} \\
&= 2^{(2+\overline{u})n+o(n)}\sum_{k=0}^{n-1}\binom{n-1}{k}2^{-2k} \\
&= 2^{(2+\overline{u})n+o(n)}(5/4)^{n-1} \\
&= \left(5\cdot 2^{\overline{u}}\right)^{(1+o(1))n}.
\end{aligned}
$$

6.5 Cambrian lattices of dihedral type

The dihedral group of order $2m$, denoted $I_2(m)$, is a Coxeter group with two simple generators s and t such that $m(s,t)=m$. This group has two Coxeter elements: st and ts. The Cambrian lattices for these Coxeter elements are isomorphic, so we may assume $c=st$. Let $[s|t]_k$ denote the word of length k that starts with s and alternates between s and t. Let α_k be the element of $I_2(m)$ represented by $[s|t]_k$. For example, $\alpha_3=sts$, and $\alpha_4=stst$. The c-Cambrian lattice consists of the two chains $e\lessdot\alpha_1\lessdot\alpha_2\lessdot\cdots\lessdot\alpha_m$ and $e\lessdot t\lessdot\alpha_m$. If $\mathcal{C}$ is a chain of length r, then $\mathcal{E}(\mathcal{C})=r/p$. It follows immediately from the definition of the Ungarian Markov chain $\mathbf{U}_{\mathrm{Camb}_c}$ that

$$
\mathcal{E}(\mathrm{Camb}_c)=p^2\cdot 1+p(1-p)\cdot(1+(m-1)/p)+p(1-p)\cdot(1+1/p)+(1-p)^2\cdot(1+\mathcal{E}(\mathrm{Camb}_c)).
$$

Solving this equation yields

$$
\mathcal{E}(\mathrm{Camb}_c)=\frac{1+m(1-p)}{2p-p^2},
$$

which proves Theorem 1.18.

7 ν-Tamari lattices

In this section, we define the ν-Tamari lattice $\mathrm{Tam}(\nu)$ of Préville-Ratelle and Viennot, and we give a simple description of its spine. This will allow us to use Theorem 1.11 to prove the upper bound for $\mathcal{E}(\mathrm{Tam}(\nu))$ stated in Theorem 1.19.

A *lattice path* is a finite path in $\mathbb{R}^2$ that starts at a point in $\mathbb{Z}^2$ and uses unit north (i.e., $(0,1)$) steps and unit east (i.e., $(1,0)$) steps; we consider lattice paths that are translations of each other to be the same. We denote north steps by N and east steps by E, and we identify lattice paths with finite words over the alphabet $\{\mathrm{N},\mathrm{E}\}$. We frequently use superscripts to denote concatenation of words; for instance, $(\mathrm{NE}^2)^3=\mathrm{NEENEENEE}$. Fix nonnegative integers $n\leq\ell$ and a lattice path ν that uses n north steps and $\ell-n$ east steps, and let $\mathrm{Tam}(\nu)$ denote the collection of lattice paths that lie weakly above ν and have the same endpoints as ν.

Given a lattice point $v=(x,y)$ that lies weakly above ν and satisfies $0\leq x\leq\ell-n$ and $0\leq y\leq n$, we define the *horizontal distance* of v to be the largest integer d such that $(x+d,y)$ is on ν. Now suppose $\mu\in\mathrm{Tam}(\nu)$ and v is a lattice point on μ that is immediately preceded by an east step and immediately followed by a north step in μ. Let v' be the first lattice point on μ that appears after v and has the same horizontal distance as v. Let $\mathrm{D}_{[v,v']}$ be the subpath of μ whose endpoints are v and v'. There are lattice paths X and Y such that $\mu=\mathrm{XED}_{[v,v']}\mathrm{Y}$. Let $\mu'=\mathrm{XD}_{[v,v']}\mathrm{EY}$. Then $\mu\lessdot\mu'$ is a cover relation in $\mathrm{Tam}(\nu)$. (See Figure 7.)

The cover relations defined in the previous paragraph make $\mathrm{Tam}(\nu)$ into a poset. Préville-Ratelle and Viennot [41] showed that $\mathrm{Tam}(\nu)$ is a lattice called the *ν-Tamari*
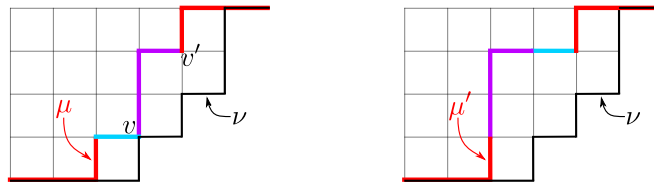
Figure 7: The lattice paths $\mu = $ EENENNENEE (left) and $\mu' = $ EENNNEENEE (right) form the cover relation $\mu \lessdot \mu'$ in $\mathrm{Tam}(\nu)$, where $\nu = $ EEENENENNE.

lattice. The n-th *m-Tamari lattice* $\mathrm{Tam}_n(m)$, which was originally introduced by Bergeron and Préville-Ratelle [11], is $\mathrm{Tam}((\mathrm{NE}^m)^n)$. Finally, $\mathrm{Tam}_n(1) = \mathrm{Tam}((\mathrm{NE})^n)$ is the n-th *Tamari lattice*, which was originally defined by Tamari [55]; we denote it simply by Tam_n. Préville-Ratelle and Viennot actually proved that the ν-Tamari lattice is an interval of some Tamari lattice; since Tamari lattices are known to be trim, it follows from Proposition 5.1 that $\mathrm{Tam}(\nu)$ is trim. Figure 8 portrays $\mathrm{Tam}_3(2)$.

If we draw the lattice path ν in the square grid, then we can consider the set $\mathrm{Cells}(\nu)$ of all unit grid cells that lie directly north of one of the east steps in ν and lie directly west of one of the north steps in ν. Define a partial order $\leq$ on $\mathrm{Cells}(\nu)$ so that for all $\square_1, \square_2 \in \mathrm{Cells}(\nu)$, we have $\square_1 \leq \square_2$ if and only if $\square_1$ is weakly southwest of $\square_2$. For example, Figure 9 shows $\mathrm{Cells}(\mathrm{ENEEENNEENNE})$ with each cover relation $\square_1 \lessdot \square_2$ represented by an arrow from $\square_1$ to $\square_2$.

Proposition 7.1. *For any lattice path ν, the Galois poset $\mathbf{P}(\mathrm{Tam}(\nu))$ is isomorphic to $\mathrm{Cells}(\nu)$.*

The previous proposition provides a simple description of the spine of $\mathrm{Tam}(\nu)$ because, according to (5.2), $\mathrm{spine}(\mathrm{Tam}(\nu)) \cong J(\mathbf{P}(\mathrm{Tam}(\nu)))$. Before proving it, let us quickly see how it implies Theorem 1.19. For each $\ell \geq 1$, choose a lattice path $\nu^{(\ell)}$ with n_ℓ north steps and $\ell - n_\ell$ east steps, and assume that the limit $\overline{n} = \lim\limits_{\ell \to \infty} \frac{1}{\ell} n_\ell$ exists. Because $\mathrm{Cells}(\nu^{(\ell)})$ is clearly a subposet of the rectangle $R_{n_\ell, \ell - n_\ell}$, it follows from Theorem 1.9 and Corollary 4.1 that

$$\mathcal{E}(J(\mathrm{Cells}(\nu^{(\ell)}))) \leq \frac{1}{p}\left(1 + 2\sqrt{(1-p)\overline{n}(1-\overline{n})}\right)\ell + o(\ell).$$

If we assume Proposition 7.1, then we can deduce from (5.2) and Theorem 1.11 that

$$\mathcal{E}(\mathrm{Tam}(\nu^{(\ell)})) \leq \mathcal{E}(\mathrm{spine}(\mathrm{Tam}(\nu^{(\ell)}))) = \mathcal{E}(J(\mathrm{Cells}(\nu^{(\ell)}))),$$

thereby proving Theorem 1.19.

Let ν be a lattice path that starts at $(0,0)$ and ends at $(\ell - n, n)$. In order to prove Proposition 7.1, we will need an alternative way of thinking about the ν-Tamari lattice. Let $\mathbf{h}(\nu) = (h_0(\nu), \ldots, h_\ell(\nu))$ be the vector obtained by reading the heights (i.e., y-coordinates) of the lattice points on ν in the order they appear in ν. For $0 \leq k \leq n$, let f_k be the maximum index such that $h_{f_k}(\nu) = k$. A *ν-bracket vector* is an integer vector $\mathbf{b} = (\mathbf{b}_0, \ldots, \mathbf{b}_\ell)$ satisfying the following conditions:

(I) $\mathbf{b}_{f_k} = k$ for all $0 \leq k \leq n$;

(II) $h_i(\nu) \leq \mathbf{b}_i \leq n$ for all $0 \leq i \leq \ell$;

(III) $\mathbf{b}$ avoids the pattern 121.

The third condition means that there do not exist indices $i_1 < i_2 < i_3$ such that $\mathbf{b}_{i_1} = \mathbf{b}_{i_3} < \mathbf{b}_{i_2}$. Ceballos, Padrol, and Sarmiento [20] proved that for each lattice path
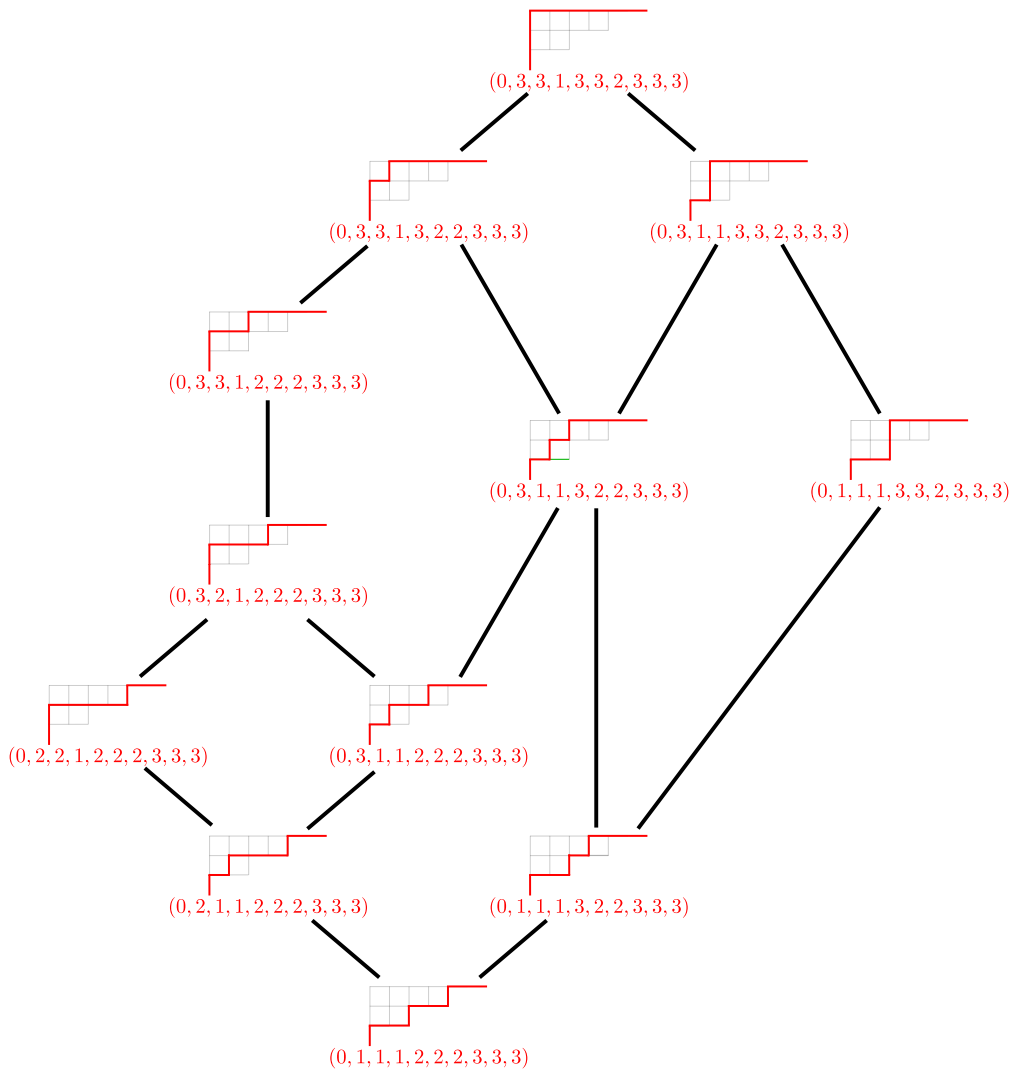
Figure 8: The Hasse diagram of $\mathrm{Tam}_3(2) = \mathrm{Tam}(\nu)$, where $\nu = (\mathrm{NE}^2)^3$. Each lattice path appears above its associated ν-bracket vector.
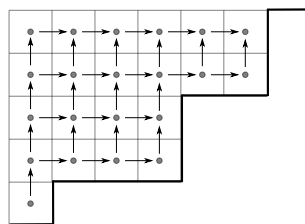

Figure 9: The set $\mathrm{Cells}(\nu)$, where $\nu = \mathrm{ENEEENNEEINNE}$. Each cover relation $\square_1 \lessdot \square_2$ is represented by an arrow from $\square_1$ to $\square_2$.

$\mu \in \mathrm{Tam}(\nu)$, there is a unique ν-bracket vector $\mathbf{v}_\nu(\mu)$ such that for every $k \in \{0, \dots, n\}$, the number of lattice points on μ with height k is equal to the number of occurrences of k in $\mathbf{v}_\nu(\mu)$. (In particular, $\mathbf{v}_\nu(\nu) = \mathbf{h}(\nu)$.) This correspondence is illustrated in Figure 8.

There is an obvious partial order on the set of ν-bracket vectors given by componentwise comparison: $\mathsf{b} \leq \mathsf{b}'$ if and only if $\mathsf{b}_i \leq \mathsf{b}'_i$ for all $0 \leq i \leq \ell$. This poset has a meet operation given by the componentwise minimum:

$$\mathsf{b} \wedge \mathsf{b}' = (\min\{\mathsf{b}_0, \mathsf{b}'_0\}, \ldots, \min\{\mathsf{b}_\ell, \mathsf{b}'_\ell\}).$$

Indeed, it is straightforward to check that the componentwise minimum of two ν-bracket vectors is a ν-bracket vector. It is well known that a finite poset with a meet operation and a unique maximal element is a lattice. This poset on ν-bracket vectors has $\mathbf{v}_\nu(\mathrm{N}^n \mathrm{E}^{\ell-n})$ as its unique maximal element, so it is a lattice. In fact, Ceballos, Padrol, and Sarmiento [20] proved that the map $\mu \mapsto \mathbf{v}_\nu(\mu)$ is a lattice isomorphism from $\mathrm{Tam}(\nu)$ to this lattice on the set of ν-bracket vectors. Therefore, in the remainder of this subsection, we will abuse notation and identify $\mathrm{Tam}(\nu)$ (via this isomorphism) with the lattice of ν-bracket vectors under the componentwise order.

Note that we 0-index our vectors. For instance, we would say that 5 appears in positions 0 and 3 in the vector $(5, 2, 1, 5, 1, 3)$. We always write b_i for the entry in position i of a vector b.

Define a *descent* of a ν-bracket vector b to be an index $i \in \{0, \ldots, \ell - 1\}$ such that $\mathsf{b}_i > \mathsf{b}_{i+1}$. Let $\mathrm{Des}(\mathsf{b})$ denote the set of descents of b. According to [30, Proposition 4.4], the number of descents of b is equal to the number of elements covered by b in $\mathrm{Tam}(\nu)$. In particular, b is join-irreducible if and only if $|\mathrm{Des}(\mathsf{b})| = 1$.

Suppose $1 \leq k \leq n$, $f_{k-1} < i < f_k$, and $h_i(\nu) + 1 \leq m \leq n$ (note that $h_i(\nu) = k$). Let $\mathfrak{b}^{i,m}$ and $\mathfrak{c}^{i,m}$ be the ν-bracket vectors such that for all $0 \leq r \leq \ell$, the entries in position r of $\mathfrak{b}^{i,m}$ and $\mathfrak{c}^{i,m}$ are

$$\mathfrak{b}_r^{i,m} = \begin{cases} m & \text{if } f_{k-1} + 1 \leq r \leq i; \\ h_r(\nu) & \text{otherwise.} \end{cases} \tag{7.1}$$

and

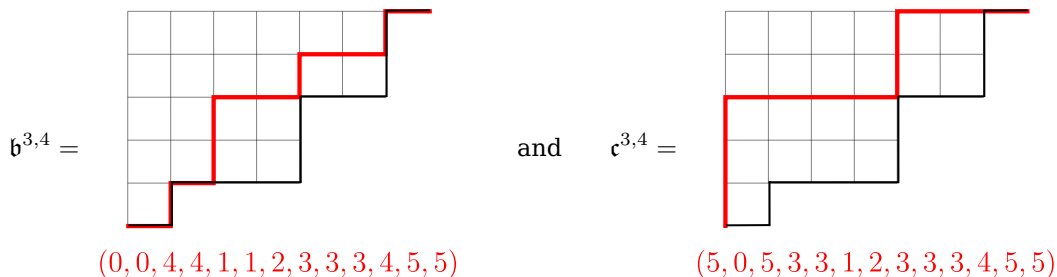$$\mathfrak{c}_r^{i,m} = \begin{cases} s & \text{if } r = f_s \text{ for some } s \in \{0, \ldots, n\}; \\ m - 1 & \text{if } i \leq r < f_{m-1} \text{ and } r \notin \{f_0, \ldots, f_n\}; \\ n & \text{otherwise.} \end{cases} \tag{7.2}$$

Example 7.2. Suppose $\nu = \mathrm{ENEEENNEENNE}$. Then

$$\mathbf{h}(\nu) = (0, 0, 1, 1, 1, 1, 2, 3, 3, 3, 4, 5, 5),$$

so $f_0 = 1$, $f_1 = 5$, $f_2 = 6$, $f_3 = 9$, $f_4 = 10$, and $f_5 = 12$. Setting $i = 3$ and $m = 4$ in (7.1) and (7.2), we find that

$$\mathfrak{b}^{3,4} = \qquad \text{and} \qquad \mathfrak{c}^{3,4} =$$

$$(0, 0, 4, 4, 1, 1, 2, 3, 3, 3, 4, 5, 5) \qquad\qquad (5, 0, 5, 3, 3, 1, 2, 3, 3, 3, 4, 5, 5)$$

(where we are representing elements of $\mathrm{Tam}(\nu)$ as both lattices paths and ν-bracket vectors).

We will need the notation and terminology from Section 5.1. The lattice path corresponding to $\mathfrak{c}^{i,m}$ only has a single east step that is immediately followed by a north step, so it follows from the above description of cover relations in $\mathrm{Tam}(\nu)$ (in terms of lattice

paths) that $\mathfrak{c}^{i,m}$ is meet-irreducible in $\mathrm{Tam}(\nu)$. Moreover, it is straightforward to see that every meet-irreducible element of $\mathrm{Tam}(\nu)$ arises in this way. In other words,

$$\mathcal{M}_{\mathrm{Tam}(\nu)} = \{\mathfrak{c}^{i,m} : i \in \{0, \dots, \ell\} \setminus \{f_0, \dots, f_n\}, h_i(\nu) + 1 \le m \le n\}. \qquad (7.3)$$

The vector $\mathfrak{b}^{i,m}$ is join-irreducible because its only descent is i. Because $\mathrm{Tam}(\nu)$ is trim, we know from Section 5.1 that there is a bijection $\kappa = \kappa_{\mathrm{Tam}(\nu)} \colon \mathcal{J}_{\mathrm{Tam}(\nu)} \to \mathcal{M}_{\mathrm{Tam}(\nu)}$. In light of (7.3), this implies that every join-irreducible element of $\mathrm{Tam}(\nu)$ is of the form $\mathfrak{b}^{i,m}$. In other words,

$$\mathcal{J}_{\mathrm{Tam}(\nu)} = \{\mathfrak{b}^{i,m} : i \in \{0, \dots, \ell\} \setminus \{f_0, \dots, f_n\}, h_i(\nu) + 1 \le m \le n\}. \qquad (7.4)$$

In order to understand the Galois poset $\mathbf{P}(\mathrm{Tam}(\nu))$, we must understand the Galois graph $\mathbf{G}(\mathrm{Tam}(\nu))$. This, in turn, requires us to understand the bijection $\kappa \colon \mathcal{J}_{\mathrm{Tam}(\nu)} \to \mathcal{M}_{\mathrm{Tam}(\nu)}$.

Lemma 7.3. *Suppose* $1 \le k \le n$, $f_{k-1} < i < f_k$, *and* $h_i(\nu) + 1 \le m \le n$. *Then*

$$\kappa(\mathfrak{b}^{i,m}) = \mathfrak{c}^{i,m}.$$

Proof. By the definition of κ in (5.1), we just need to demonstrate that $\mathfrak{b}^{i,m} \wedge \mathfrak{c}^{i,m}$ is covered by $\mathfrak{b}^{i,m}$ and that $\mathfrak{b}^{i,m} \vee \mathfrak{c}^{i,m}$ covers $\mathfrak{c}^{i,m}$. The meet operation is given by the componentwise minimum, so it follows from (7.1) and (7.2) that

$$(\mathfrak{b}^{i,m} \wedge \mathfrak{c}^{i,m})_r = \begin{cases} m & \text{if } f_{k-1} + 1 \le r \le i - 1; \\ m - 1 & \text{if } r = i; \\ h_r(\nu) & \text{otherwise.} \end{cases}$$

It is immediate from this description that $\mathfrak{b}^{i,m} \wedge \mathfrak{c}^{i,m}$ is covered by $\mathfrak{b}^{i,m}$.

Now let $\mathsf{b} = (\mathfrak{c}^{i,m})^*$ be the unique element that covers $\mathfrak{c}^{i,m}$. Note that $\mathfrak{b}^{i,m} \not\le \mathfrak{c}^{i,m}$ because $\mathfrak{b}_i^{i,m} = m > m - 1 = \mathfrak{c}_i^{i,m}$. Hence, to prove that $\mathfrak{b}^{i,m} \vee \mathfrak{c}^{i,m} = \mathsf{b}$, we just need to show that $\mathfrak{b}^{i,m} \le \mathsf{b}$. For every index $q \in \{0, \dots, \ell\} \setminus \{i\}$, we have $\mathfrak{b}_q^{i,m} \le \mathfrak{c}_q^{i,m} \le \mathsf{b}_q$ by (7.1) and (7.2). Thus, we are left to show that $\mathfrak{b}_i^{i,m} \le \mathsf{b}_i$.

Since $\mathsf{b} > \mathfrak{c}^{i,m}$, we know that $\mathsf{b}_i \ge \mathfrak{c}_i^{i,m} = m - 1$ and that there is an index r such that $\mathsf{b}_r > \mathfrak{c}_r^{i,m}$. Appealing to (7.2) and the definition of a ν-bracket vector, we find that $i \le r < f_{m-1}$, $r \notin \{f_0, \dots, f_n\}$, and $\mathfrak{c}_r^{i,m} = m - 1$. If $\mathsf{b}_i = m - 1$, then $\mathsf{b}_i = \mathsf{b}_{f_{m-1}} = m - 1 < \mathsf{b}_r$, so the entries of b in positions i, r, f_{m-1} form a 121-pattern, contradicting the definition of a ν-bracket vector. It follows that $\mathsf{b}_i \ge m = \mathfrak{b}_i^{i,m}$, as desired. $\qquad \square$

Proof of Proposition 7.1. Assume that ν starts at $(0,0)$ and ends at $(\ell - n, n)$. Suppose $1 \le k \le n$, $f_{k-1} < i < f_k$, and $h_i(\nu) + 1 \le m \le n$ (note that $h_i(\nu) = k$). The i-th step in ν is an east step at height k; let $\square^{i,m}$ be the grid cell that lies above that east step and has its center at height $m - \frac{1}{2}$. Every grid cell in $\mathrm{Cells}(\nu)$ arises uniquely in this way; that is,

$$\mathrm{Cells}(\nu) = \{\square^{i,m} : i \in \{0, \dots, \ell\} \setminus \{f_0, \dots, f_n\}, h_i(\nu) + 1 \le m \le n\}.$$

Recall from Section 5.1 that the underlying set of the Galois poset $\mathbf{P}(\mathrm{Tam}(\nu))$ is $\mathcal{J}_{\mathrm{Tam}(\nu)}$. It is now immediate from (7.4) that we have a bijection $\mathbf{P}(\mathrm{Tam}(\nu)) \to \mathrm{Cells}(\nu)$ given by $\mathfrak{b}^{i,m} \mapsto \square^{i,m}$; we just need to prove that this map is a poset isomorphism.

We first show that if $\square^{i,m} \le \square^{i',m'}$ in $\mathrm{Cells}(\nu)$ (so $i \le i'$ and $m \le m'$), then $\mathfrak{b}^{i,m} \preceq \mathfrak{b}^{i',m'}$ in $\mathbf{P}(\mathrm{Tam}(\nu))$; it suffices to do so when either $i = i'$ or $m = m'$. If $i = i'$, then $m - 1 \ge h_i(\nu) = h_{i'}(\nu)$. If $m = m'$, then $m - 1 = m' - 1 \ge h_{i'}(\nu)$. In either case, $m - 1 \ge h_{i'}(\nu)$. Recall that $f_{h_{i'}(\nu)}$ is the largest position in which the vector $\mathbf{h}(\nu)$ has the entry $h_{i'}(\nu)$. Since $h_{i'}(\nu)$ is, by definition, the entry in position i' of $\mathbf{h}(\nu)$, it follows that $i' \le f_{h_{i'}(\nu)}$.

Hence, $i' \le f_{m-1}$. This shows that $i \le i' \le f_{m-1}$, so we know by (7.2) and Lemma 7.3 that the entry in position i' of $\kappa(\mathfrak{b}^{i,m})$ is $m-1$. On the other hand, (7.1) tells us that the entry in position i' of $\mathfrak{b}^{i',m'}$ is m', which is strictly greater than $m-1$. This proves that $\mathfrak{b}^{i',m'} \not\preceq \kappa(\mathfrak{b}^{i,m})$ in $\mathrm{Tam}(\nu)$. Hence, there is an arrow $\mathfrak{b}^{i',m'} \to \mathfrak{b}^{i,m}$ in the Galois graph $\mathbf{G}(\mathrm{Tam}(\nu))$, which implies that $\mathfrak{b}^{i,m} \preceq \mathfrak{b}^{i',m'}$ in $\mathbf{P}(\mathrm{Tam}(\nu))$, as desired.

We now show that if $\mathfrak{b}^{i,m} \preceq \mathfrak{b}^{i',m'}$ in $\mathbf{P}(\mathrm{Tam}(\nu))$, then $\square^{i,m} \le \square^{i',m'}$ in $\mathrm{Cells}(\nu)$ (i.e., $i \le i'$ and $m \le m'$); it suffices to do so when there is an arrow $\mathfrak{b}^{i',m'} \to \mathfrak{b}^{i,m}$ in $\mathbf{G}(\mathrm{Tam}(\nu))$. Thus, $\mathfrak{b}^{i',m'} \not\preceq \kappa(\mathfrak{b}^{i,m})$ in $\mathrm{Tam}(\nu)$. This means that there is some index r such that $(\kappa(\mathfrak{b}^{i,m}))_r < \mathfrak{b}_r^{i',m'}$. Let k' be the index such that $f_{k'-1} < i' < f_{k'}$. We know from the definition of a ν-bracket vector that $(\kappa(\mathfrak{b}^{i,m}))_r \ge h_r(\nu)$, so $\mathfrak{b}_r^{i',m'} > h_r(\nu)$. According to (7.1), we have $\mathfrak{b}_r^{i',m'} = m'$ and $f_{k'-1} + 1 \le r \le i'$. Also, we have $\mathfrak{b}_r^{i',m'} \le n$, so $(\kappa(\mathfrak{b}^{i,m}))_r \ne n$. Since $r \notin \{f_0, \ldots, f_n\}$ (because $f_{k'-1} < r \le i' < f_{k'}$), it follows from (7.2) and Lemma 7.3 that $(\kappa(\mathfrak{b}^{i,m}))_r = m-1$ and $i \le r < f_{m-1}$. We have shown that $i \le r \le i'$ and $m - 1 = (\kappa(\mathfrak{b}^{i,m}))_r < \mathfrak{b}_r^{i',m'} = m'$. Hence, $i \le i'$ and $m \le m'$, so $\square^{i,m} \le \square^{i',m'}$. $\qquad\square$

8 Tamari lattices

A permutation $x \in S_n$ is called [312-avoiding]{.underline} if there do not exist indices $i_1 < i_2 < i_3$ such that $x(i_2) < x(i_3) < x(i_1)$. Let $\mathrm{Av}_n(312)$ denote the set of 312-avoiding permutations in S_n. Then $\mathrm{Av}_n(312)$ is precisely the set of c-sortable elements of S_n, where $c = s_1 s_2 \cdots s_{n-1} \in A_{n-1} = S_n$, so Camb_c is the sublattice of the weak order on S_n consisting of 312-avoiding permutations. The Tamari lattice Tam_n is isomorphic to this Cambrian lattice Camb_c [43], so it follows from Theorem 1.13 that $\mathcal{E}(\mathrm{Tam}_n) \le \frac{1}{p}\left(2 + 2\sqrt{1-p}\right) n + o(n)$. This bound is also a special case of Theorem 1.19. Because Tamari lattices are such important objects, we wish to improve upon this bound by proving Theorem 1.20. To do so, we will analyze Tamari lattices directly (without appealing to Theorem 1.11), using the interpretation of the n-th Tamari lattice as $\mathrm{Av}_n(312)$ under the weak order.

Figure 10 shows the plots of twelve of the permutations that we obtained while running $\mathbf{U}_{\mathrm{Av}_{400}(312)}$ with $p = 1/2$ starting with the decreasing permutation. These plots illustrate behavior that appears to be typical: the permutations become "mostly sorted" after a small number of steps, but there are a few entries that straggle behind and take a long time to reach their correct positions. Experiments also suggest that the number of steps needed to go from the decreasing permutation to the identity permutation in $\mathbf{U}_{\mathrm{Av}_n(312)}$ has a relatively high variance, and this makes the analysis difficult.

In order to analyze $\mathbf{U}_{\mathrm{Av}_n(312)}$, we first need an explicit combinatorial description of how to apply a random Ungar move to a permutation in $\mathrm{Av}_n(312)$. Consider $x \in S_n$. If there exist indices i and i' such that $i + 1 < i'$ and $x(i+1) < x(i') < x(i)$, then we can perform an [allowable swap]{.underline} by swapping the entries $x(i)$ and $x(i+1)$ (i.e., replacing x by xs_i). Let $\pi_\downarrow(x)$ denote the permutation obtained by starting with x and repeatedly applying allowable swaps until no more allowable swaps can be performed. The element $\pi_\downarrow(x)$ is well defined (i.e., does not depend on the sequence of allowable swaps) and is in $\mathrm{Av}_n(312)$ [43]. Thus, we have a *projection* $\pi_\downarrow \colon S_n \to \mathrm{Av}_n(312)$. Note that $\pi_\downarrow(x) = x$ if and only if x is 312-avoiding.

The first author showed [30, Equation (1)] that $\mathsf{Pop}_{\mathrm{Av}_n(312)}(x) = \pi_\downarrow(\mathsf{Pop}_{S_n}(x))$ for every $x \in \mathrm{Av}_n(312)$. In other words, if $x \in \mathrm{Av}_n(312)$, then applying a maximal Ungar move to x within the Tamari lattice $\mathrm{Av}_n(312)$ is equivalent to applying a maximal Ungar move to x within the weak order lattice S_n and then applying the projection $\pi_\downarrow$. The exact same argument (which we omit) shows that applying a random Ungar move to x within $\mathrm{Av}_n(312)$ is equivalent to applying a random Ungar move to x within S_n and then applying $\pi_\downarrow$. Thus, a random Ungar move within $\mathrm{Av}_n(312)$ will send x to $\pi_\downarrow(xw_0(T))$, where T is a random subset of $\mathrm{Des}(x)$ (with each descent of x added to T with probability
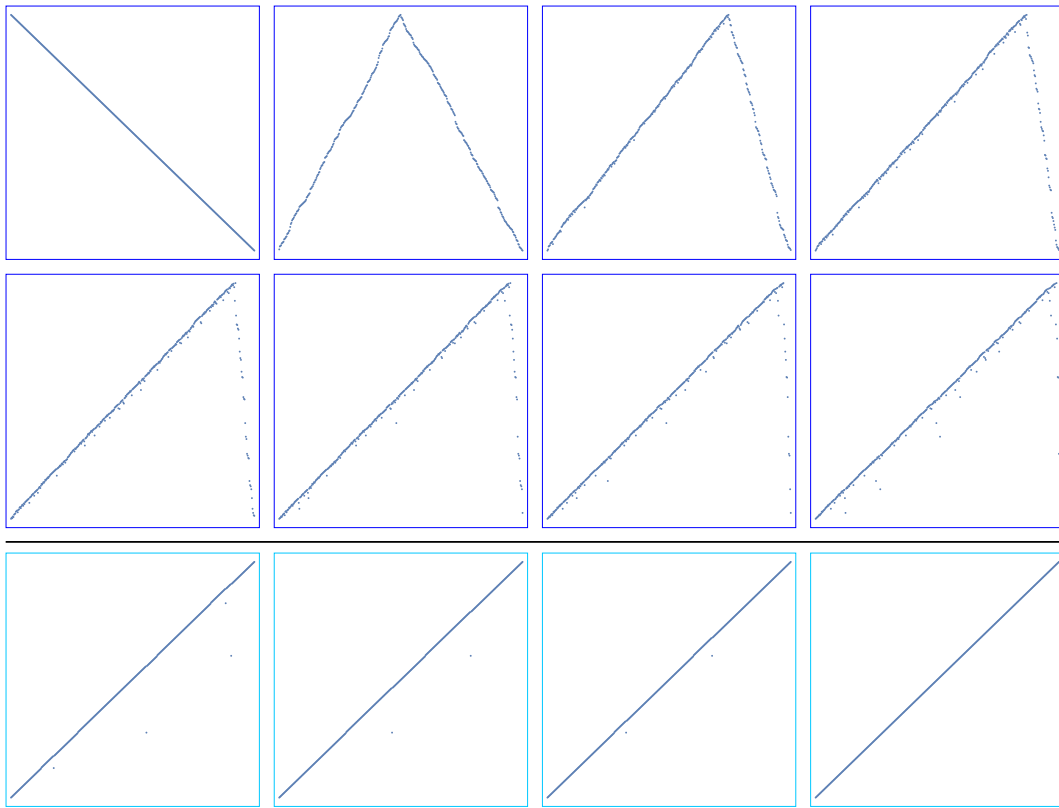
Figure 10: We ran the Ungarian Markov chain $\mathbf{U}_{\mathrm{Av}_{400}(312)}$ with $p = 1/2$ starting with the decreasing permutation. Here, we show the plots of the permutations that the Markov chain reached at times $0, 1, 2, 3, 4, 5, 6, 7, 90, 180, 270, 360$ (read from left to right, row by row). Note that the permutations get "mostly sorted" within the first few steps. To indicate that the permutations in the bottom row follow a time skip, we have colored their bounding boxes with a slightly different hue.

p) and $w_0(T)$ is the maximal element of the parabolic subgroup of S_n generated by $\{s_i : i \in T\}$.

As in Section 3, we write $\mathrm{DB}(w)$ for the set $\{w(i+1) : i \in \mathrm{Des}(w)\}$ of descent bottoms of a permutation $w \in S_n$. Let $\mathrm{LRMax}(w)$ denote the set

$$\{w(i) : w(j) < w(i) \text{ for all } 1 \leq j \leq i-1\}$$

of *left-to-right maxima* of w. If $w' \leq w$ in the weak order, then $\mathrm{LRMax}(w) \subseteq \mathrm{LRMax}(w')$. A simple but useful fact that we will exploit is that if w is 312-avoiding, then $\mathrm{DB}(w)$ and $\mathrm{LRMax}(w)$ form a partition of $[n]$. This implies that if $w' \leq w$ in $\mathrm{Av}_n(312)$ (for instance, if w' is obtained by applying an Ungar move to w), then $\mathrm{DB}(w') \subseteq \mathrm{DB}(w)$.

For integers $\beta \in [n-1]$ and $j \geq 0$, let $X_j^{(\beta)}$ be a Bernoulli random variable with expected value p; assume that all of these random variables for different choices of β and j are independent. We can simulate $\mathbf{U}_{\mathrm{Av}_n(312)}$ using these random variables as follows. Starting with the decreasing permutation $\sigma_0 = n(n-1)\cdots 1$, we will create a sequence $\sigma_0, \sigma_1, \ldots$ of 312-avoiding permutations. Suppose we have already generated the permutations $\sigma_0, \ldots, \sigma_t$. Let

$$T_t = \left\{ i \in \mathrm{Des}(\sigma_t) : X_t^{(\sigma_t(i+1))} = 1 \right\}.$$

We define σ_{t+1} to be the permutation obtained by applying the random Ungar move (within $\mathrm{Av}_n(312)$) corresponding to the set T_t to σ_t; that is, $\sigma_{t+1} = \pi_{\downarrow}(\sigma_t w_0(T_t))$. Note that for permutations $w, w' \in \mathrm{Av}_n(312)$, the conditional probability $\mathbb{P}(\sigma_{t+1} = w' \mid \sigma_t = w)$ is equal to the transition probability $\mathbb{P}(w \to w')$ in $\mathbf{U}_{\mathrm{Av}_n(312)}$.

Let N be the unique integer such that $\sigma_{N-1} \neq \sigma_N = 12 \cdots n$. Then $\sigma_t = 12 \cdots n$ for all $t \geq N$. For $\beta \in [n-1]$, let $q(\beta)$ be the largest integer t such that $\beta \in \mathrm{DB}(\sigma_t)$. As the sequence $\sigma_0, \sigma_1, \ldots$ is weakly decreasing in the weak order, we have the chain of containments $\mathrm{DB}(\sigma_0) \supseteq \mathrm{DB}(\sigma_1) \supseteq \cdots$. Thus, $\beta \in \mathrm{DB}(\sigma_t)$ for all $t \leq q(\beta)$. This tells us that whenever $0 \leq t \leq q(\beta)$ and $X_t^{(\beta)} = 1$, the number β *must* move to the left when we transition from σ_t to σ_{t+1}. If $\beta < \beta'$ and β appears to the left of β' in σ_t, then β must also appear to the left of β' in each of the permutations $\sigma_{t+1}, \sigma_{t+2}, \ldots$. Consequently,

$$\sum_{j=0}^{q(\beta)} X_j^{(\beta)} \leq n - \beta. \tag{8.1}$$

Observe that $N - 1 = \max_{1 \leq \beta \leq n-1} q(\beta)$. This already shows that $\mathcal{E}(\mathrm{Tam}_n) \leq \frac{1}{p} n + o(n)$, which improves upon the estimate from Theorem 1.13, but we will further improve upon this bound in what follows.

For each $\beta \in [n-1]$, let F_β be the smallest nonnegative integer j such that $X_j^{(\beta)} = 1$. Then $F_1, \ldots, F_{n-1}$ are independent geometric random variables, each with expected value $1/p$. For each real number $\zeta \in [1/n, 1]$, let A_ζ be the event that $F_\beta \leq \max\{F_i : \beta + 1 \leq i \leq \beta + \zeta n\}$ for all integers β satisfying $1 \leq \beta < (1 - \zeta)n$.

Lemma 8.1. *Suppose B is an event that only depends on the random variables $X_j^{(\beta)}$ with $1 \leq \beta \leq n - 1$ and $0 \leq j \leq \sqrt{n}$. If $\xi \in (0, 1]$, $n \geq 1/\xi$, and $B \subseteq A_\xi$, then $\mathbb{E}(N \mid B) \leq \xi n/p + o(n)$.*

Proof. Suppose the event B occurs. In particular, this implies that A_ξ occurs. Consider some $\beta \in [n-1]$. Let β' be the entry immediately to the left of β in σ_{F_β}. If $1 \leq \beta < (1-\xi)n$, then $F_\beta \leq \max\{F_i : \beta + 1 \leq i \leq \beta + \xi n\}$, so $\beta' \leq \beta + \xi n$. On the other hand, if $\beta \geq (1 - \xi)n$, then the inequality $\beta' \leq \beta + \xi n$ holds automatically. Since $X_{F_\beta}^{(\beta)} = 1$, the entry β appears to the left of β' in $\sigma_{F_\beta + 1}$. Since $\sigma_{F_\beta + 1}$ is 312-avoiding, all entries that are greater than β and appear to the left of β in $\sigma_{F_\beta + 1}$ must be at most $\beta + \xi n$; hence, there are at most $\xi n - 1$ such entries (the -1 term comes from the fact that β' is not one of these entries). Consequently, $\sum_{j = F_\beta + 1}^{q(\beta)} X_j^{(\beta)} \leq \xi n - 1$. The definition of F_β now guarantees that $\sum_{j=0}^{q(\beta)} X_j^{(\beta)} \leq \xi n$. Hence, $\sum_{\sqrt{n} < j \leq q(\beta)} X_j^{(\beta)} \leq \xi n$. Because B is independent of the variables $X_j^{(\beta)}$ with $j > \sqrt{n}$, this implies that $q(\beta) - \sqrt{n}$ is bounded above by a sum of ξn independent geometric random variables, each with expected value $1/p$. For $m \geq \frac{\xi n/p}{1 - n^{-1/4}}$, we can set $k = \xi n$ and $\gamma = \frac{pm}{\xi n}$ in Lemma 2.1 to find that

$$\mathbb{P}(q(\beta) - \sqrt{n} > m \mid B) \leq \exp\left(-(p/2)m\left(1 - \frac{\xi n}{pm}\right)^2\right) \leq \exp\left(-(p/2)mn^{-1/2}\right).$$

As β was arbitrary, we can use the identity $N - 1 = \max_{1 \leq \beta \leq n-1} q(\beta)$ and a union bound to see that

$$\mathbb{P}(N - 1 - \sqrt{n} > m \mid B) \leq n \exp\left(-(p/2)mn^{-1/2}\right).$$

Hence,

$$\sum_{m \geq \frac{\xi n/p}{1 - n^{-1/4}}} \mathbb{P}(N - 1 - \sqrt{n} > m \mid B) \leq \sum_{m \geq \frac{\xi n/p}{1 - n^{-1/4}}} n \exp\left(-(p/2)mn^{-1/2}\right)$$

$$\leq n \frac{\exp\left(-(p/2)n^{-1/2}\frac{\xi n/p}{1-n^{-1/4}}\right)}{1-\exp\left(-(p/2)n^{-1/2}\right)}$$

$$\leq n \frac{\exp\left(-\xi n^{1/2}/2\right)}{1-\exp\left(-(p/2)n^{-1/2}\right)}$$

$$= O\left(n^{3/2}e^{-\xi n^{1/2}/2}\right)$$

$$= o(1).$$

We deduce that

$$\mathbb{E}(N \mid B) = \sum_{0 \leq m < \frac{\xi n/p}{1-n^{-1/4}}+\sqrt{n}+1} \mathbb{P}(N > m \mid B) + \sum_{m \geq \frac{\xi n/p}{1-n^{-1/4}}+\sqrt{n}+1} \mathbb{P}(N > m \mid B)$$

$$\leq \frac{\xi n/p}{1-n^{-1/4}} + \sqrt{n} + 2 + \sum_{m \geq \frac{\xi n/p}{1-n^{-1/4}}} \mathbb{P}(N - 1 - \sqrt{n} > m \mid B)$$

$$= \xi n/p + o(n). \qquad \square$$

Now let A' be the event that $\max\{F_1, \ldots, F_{n-1}\} \leq \sqrt{n}$. Lemma 8.1 yields the following corollaries.

Corollary 8.2. *For $\zeta \in (0,1]$ and $n \geq 1/\zeta$, we have $\mathbb{E}(N \mid A_\zeta \cap A') \leq \zeta n/p + o(n)$.*

Proof. The event $A_\zeta \cap A'$ only depends on the random variables $X_j^{(\beta)}$ with $\beta \in [n-1]$ and $0 \leq j \leq \sqrt{n}$, so we can set $\xi = \zeta$ in Lemma 8.1 to deduce the desired result. $\square$

Corollary 8.3. *We have $\mathbb{P}(\neg A')\mathbb{E}(N \mid \neg A') = O(n^2 e^{-p\sqrt{n}/3})$.*

Proof. For each $\beta \in [n-1]$, we can set $k = 1$ and $\gamma = p\sqrt{n}$ in Lemma 2.1 to see that

$$\mathbb{P}(F_\beta > \sqrt{n}) \leq e^{-\frac{p\sqrt{n}}{2}(1-(p\sqrt{n})^{-1})^2} = O(e^{-p\sqrt{n}/3}).$$

Taking a union bound over all $\beta \in [n-1]$ shows that $\mathbb{P}(\neg A') = O(ne^{-p\sqrt{n}/3})$.

To complete the proof, we just need to show that $\mathbb{E}(N \mid \neg A') = O(n)$. The event $\neg A'$ only depends on the random variables $X_j^{(\beta)}$ with $\beta \in [n-1]$ and $0 \leq j \leq \sqrt{n}$. Also, the event A_1 always occurs vacuously, so $(\neg A') \subseteq A_1$. Setting $\xi = 1$ in Lemma 8.1, we derive the estimate $\mathbb{E}(N \mid \neg A') \leq n/p + o(n) = O(n)$. $\square$

Corollary 8.4. *For $\zeta \in (0,1]$ and $n \geq 1/\zeta$, we have $\mathbb{E}(N \mid A' \setminus A_\zeta) \leq n/p + o(n)$.*

Proof. The event $A' \setminus A_\zeta$ only depends on the random variables $X_j^{(\beta)}$ with $\beta \in [n-1]$ and $0 \leq j \leq \sqrt{n}$. The event A_1 always occurs vacuously, so $(A' \setminus A_\zeta) \subseteq A_1$. Setting $\xi = 1$ in Lemma 8.1 yields the estimate $\mathbb{E}(N \mid A' \setminus A_\zeta) \leq n/p + o(n)$. $\square$

Let $G_1, G_2, \ldots$ be independent geometric random variables, each with expected value $1/p$. Recall from Section 1 that Bruss and O'Cinneide [17] proved that $\lim_{n \to \infty}(\rho_p(n) - \Upsilon_p(n)) = 0$, where

$$\Upsilon_p(x) = \begin{cases} px \displaystyle\sum_{k \in \mathbb{Z}} (1-p)^k e^{-(1-p)^k x} & \text{if } p < 1; \\ 0 & \text{if } p = 1 \end{cases}$$

and $\rho_p(n)$ is the probability that there is a unique integer $i \in [n]$ such that $G_i = \max\{G_1, \ldots, G_n\}$. Moreover, $\Upsilon_p(x)$ is logarithmically periodic, so $\max\limits_{0 < x < 1} \Upsilon_p(x)$ is equal to $\limsup\limits_{n \to \infty} \rho_p(n)$. Let

$$\overline{\rho}_p = \max_{0 < x < 1} \Upsilon_p(x) = \limsup_{n \to \infty} \rho_p(n).$$

We can now complete the proof of Theorem 1.20.

Proof of Theorem 1.20. If $p = 1$, then Theorem 1.20 states that $\mathcal{E}(\mathrm{Tam}_n) = o(n)$, which is certainly true because $\hat{1}$ transitions to $\hat{0}$ in a single step in $\mathbf{U}_{\mathrm{Tam}_n}$. Now assume $p < 1$. Let $\zeta = \sqrt{\overline{\rho}_p/(1 + \overline{\rho}_p)} \in (0, 1]$, and assume $n \geq 1/\zeta$. If β is an integer satisfying $1 \leq \beta < (1 - \zeta)n$, then $F_\beta, F_{\beta+1}, \ldots, F_{\beta + \lfloor \zeta n \rfloor}$ are independent geometric random variables, each with expected value $1/p$, so

$$\mathbb{P}\left(F_\beta > \max\{F_i : \beta + 1 \leq i \leq \beta + \zeta n\}\right) = \frac{1}{\lfloor \zeta n \rfloor + 1}\rho_p(n) \leq \frac{1}{\zeta n}(\overline{\rho}_p + o(1)).$$

Consequently,

$$1 - \mathbb{P}(A_\zeta) \leq \frac{1 - \zeta}{\zeta}(\overline{\rho}_p + o(1)).$$

Appealing to Corollaries 8.2 to 8.4, we compute

$$\begin{aligned}
\mathcal{E}(\mathrm{Tam}_n) &= \mathbb{E}(N) \\
&= \mathbb{P}(A_\zeta \cap A')\mathbb{E}(N \mid A_\zeta \cap A') + \mathbb{P}(\neg A')\mathbb{E}(N \mid \neg A') + \mathbb{P}(A' \setminus A_\zeta)\mathbb{E}(N \mid A' \setminus A_\zeta) \\
&\leq \mathbb{P}(A_\zeta \cap A')(\zeta n/p + o(n)) + O(n^2 e^{-p\sqrt{n}/3}) + \mathbb{P}(A' \setminus A_\zeta)(n/p + o(n)) \\
&= \mathbb{P}(A_\zeta \cap A')\zeta n/p + \mathbb{P}(A' \setminus A_\zeta)n/p + o(n) \\
&\leq \mathbb{P}(A_\zeta)\zeta n/p + (1 - \mathbb{P}(A_\zeta))n/p + o(n) \\
&= \zeta n/p + (1 - \mathbb{P}(A_\zeta))(1 - \zeta)n/p + o(n) \\
&\leq \zeta n/p + \left(\frac{1 - \zeta}{\zeta}(\overline{\rho}_p + o(1))\right)(1 - \zeta)n/p + o(n) \\
&= \frac{1}{p}\left(\zeta + \frac{(1 - \zeta)^2}{\zeta}\overline{\rho}_p\right)n + o(n).
\end{aligned}$$

Substituting $\zeta = \sqrt{\overline{\rho}_p/(1 + \overline{\rho}_p)}$ and simplifying yields

$$\mathcal{E}(\mathrm{Tam}_n) \leq \frac{2}{p}\left(\sqrt{\overline{\rho}_p(1 + \overline{\rho}_p)} - \overline{\rho}_p\right)n + o(n),$$

as desired. $\square$

9 Future directions

9.1 Better estimates

It would be interesting to improve upon the asymptotic upper bounds that we proved for $\mathcal{E}(L)$ for various lattices L. With regard to the weak order on S_n, we conjecture that the upper bound in Theorem 1.8, which is on the order of $n \log n$, can be improved to a linear bound.

Conjecture 9.1. *We have $\mathcal{E}(S_n) = O(n)$.*

In Theorem 1.20, we focused on Tamari lattices and improved upon the upper bound that one would obtain from Theorem 1.13 by viewing Tam_n as a Cambrian lattice of type A. There are other specific interesting families of Cambrian lattices where one could

try to make similar improvements. For example, Reading's *Tamari lattices of type B_n* [43] are the Cambrian lattices of type B_n given by the Coxeter elements $s_0 s_1 \cdots s_{n-1}$ and $s_{n-1} \cdots s_1 s_0$. Another interesting class of Cambrian lattices worth exploring in more detail consists of the *bipartite Cambrian lattices*. To construct these, choose a bipartition $X \sqcup Y$ of the Coxeter graph of a finite irreducible Coxeter group W (this is possible because the Coxeter graph is a tree), and let $c_X = \prod_{s \in X} s$ and $c_Y = \prod_{s \in Y} s$. The two bipartite Cambrian lattices of type W are the Cambrian lattices corresponding to the Coxeter elements $c_X c_Y$ and $c_Y c_X$.

Most of our results focus on upper bounds for $\mathcal{E}(L)$. It would be very interesting to have nontrivial lower bounds, especially for specific choices of L. Appealing to Ungar's work, we did find a lower bound for $\mathcal{E}(S_n)$ in Theorem 1.8; can one find an improvement? Can one obtain any nontrivial asymptotic lower bound for $\mathcal{E}(\mathrm{Tam}_n)$?

9.2 The elements we reach along the way

Figure 1 shows the plots of some permutations obtained by applying the Ungarian Markov chain $\mathbf{U}_{S_{400}}$. There are very clear patterns in these plots that we do not know how to explain. Describing these shapes rigorously would be highly substantial. This problem seems similar, at least superficially, to the monumental works [2, 3, 27, 28] on random sorting networks.

More generally, we have the following.

Problem 9.2. Given a lattice L, describe the elements that we are likely to reach when we run the Ungarian Markov chain $\mathbf{U}_L$ starting at $\hat{1}$.

There are multiple ways to interpret the previous problem. On the one hand, one could study the probability distribution on L that results from running $\mathbf{U}_L$ for a prescribed number of steps (starting at $\hat{1}$). In a similar but slightly different vein, one could study the hitting probabilities of the elements of L.

9.3 Variance

This article has focused exclusively on $\mathcal{E}(L)$, the expected value of the number of steps needed to go from the top element $\hat{1}$ to the bottom element $\hat{0}$ in $\mathbf{U}_L$ (i.e., the expected hitting time of the unique absorbing state). It would be very interesting to say something about the *variance* of this number of steps. Experiments suggest that this variance is relatively small when L is the weak order on S_n but is relatively large when L is the Tamari lattice Tam_n.

9.4 Other lattices

There are several diverse families of lattices, and we have certainly not explored all of them through the lens of Ungarian Markov chains in this paper. Considering other families would likely be a rich source of future research directions.

One particular family that stands out is that of *geometric* lattices, which are fundamental because of their numerous examples and their strong connection with matroid theory. A lattice is *atomic* if each of its elements can be expressed as a join of atoms. If L is a graded lattice, then it has a unique *rank function* $\mathrm{rk} \colon L \to \mathbb{Z}_{\geq 0}$ such that $\mathrm{rk}(\hat{0}) = 0$ and $\mathrm{rk}(y) = \mathrm{rk}(x) + 1$ whenever $x \lessdot y$. We say a lattice L is *semimodular* if it is graded and its rank function satisfies

$$\mathrm{rk}(x \wedge y) + \mathrm{rk}(x \vee y) \leq \mathrm{rk}(x) + \mathrm{rk}(y)$$

for all $x, y \in L$. A lattice is *geometric* if it is atomic and semimodular. It is well known that a lattice is geometric if and only if it is isomorphic to the lattice of flats of a matroid.

Let us say a geometric lattice L is *primitive* if it cannot be written as a product of smaller geometric lattices. Then L is primitive if and only if it is isomorphic to the lattice of flats of a connected matroid.

Conjecture 9.3. *If $(L^{(n)})_{n \geq 1}$ is a sequence of primitive geometric lattices such that $|L^{(n)}| \to \infty$ as $n \to \infty$, then*

$$\mathcal{E}(L^{(n)}) = 1 + o(1).$$

References

[1] M. Aigner and G. M. Ziegler, Proofs from the Book, fifth edition. Springer-Verlag, 2014. MR3288091

[2] O. Angel, D. Dauvergne, A. Holroyd, and B. Virág, The local limit of random sorting networks. *Ann. Inst. H. Poincaré Probab. Statist.*, **55** (2019), 412–440. MR3901651

[3] O. Angel, A. Holroyd, D. Romik, and B. Virág, Random sorting networks. *Adv. Math.*, **215** (2007), 839–868. MR2355610

[4] A. Asinowski, C. Banderier, and B. Hackl, Flip-sort and combinatorial aspects of pop-stack sorting. *Discrete Math. Theor. Comput. Sci.*, **22** (2021). MR4252323

[5] A. Asinowski, C. Banderier, S. Billey, B. Hackl, and S. Linusson, Pop-stack sorting and its image: permutations with overlapping runs. *Acta. Math. Univ. Comenian.*, **88** (2019), 395–402. MR4012843

[6] D. Avis and M. Newborn. On pop-stacks in series. *Util. Math.*, **19** (1981), 129–140. MR0624050

[7] E. Barnard, E. Gunawan, E. Meehan, and R. Schiffler, Cambrian combinatorics on quiver representations (type A). *Adv. Appl. Math.*, **143** (2023). MR4492558

[8] M. Bauer and K. Copenhaver, A new lower bound for deterministic pop-stack-sorting. arXiv:2307.08188v1.

[9] M. von Bell, R. S. González D'León, F. A. Mayorga Cetina, and M. Yip, A unifying framework for the ν-Tamari lattice and principal order ideals in Young's lattice. To appear in *Combinatorica*. MR4630465

[10] M. von Bell and M. Yip, Schröder combinatorics and ν-associahedra. *European J. Combin.*, **98** (2021). MR4299738

[11] F. Bergeron and L.-F. Préville-Ratelle, Higher trivariate diagonal harmonics via generalized Tamari posets. *J. Comb.*, **3** (2012), 317–341. MR3029440

[12] G. Birkhoff, Rings of sets. *Duke Math. J.*, **3** (1937), 443–454. MR1546000

[13] A. Björner, Essential chains and homotopy type of posets. *Proc. Amer. Math. Soc.*, **116** (1992), 1179–1181. MR1140664

[14] A. Björner, Orderings of Coxeter groups. In: Combinatorics and Algebra (Boulder, Colo., 1983), vol. 34 of Contemp. Math. American Mathematical Society, 1984. MR0777701

[15] M. Bousquet-Mélou, G. Chapuy, and L.-F. Préville-Ratelle, The representation of the symmetric group on m-Tamari intervals. *Adv. Math.*, **247** (2013), 309–342. MR3096799

[16] M. Bousquet-Mélou, E. Fusy, and L.-F. Préville-Ratelle, The number of intervals in the m-Tamari lattices. *Electron. J. Combin.*, **18** (2011). MR2880681

[17] F. T. Bruss and C. A. O'Cinneide, On the maximum and its uniqueness for geometric random samples. *J. Appl. Probab.*, **27** (1990), 598–610. MR1067025

[18] G. R. Burton and G. B. Purdy, The directions determined by n points in the plane. *J. Lond. Math. Soc.*, **20** (1979), 109–114. MR0545207

[19] C. Ceballos, A. Padrol, and C. Sarmiento, Geometry of ν-Tamari lattices in types A and B. *Trans. Amer. Math. Soc.*, **371** (2019), 2575–2622. MR3896090

[20] C. Ceballos, A. Padrol, and C. Sarmiento, The ν-Tamari lattice via ν-trees, ν-bracket vectors, and subword complexes. *Electron. J. Combin.*, **27** (2020). MR4057170

[21] F. Chapoton and B. Rognerud, On the wildness of Cambrian lattices. *Algebr. Represent. Theory*, **22** (2019), 603–614. MR3951763

[22] G. Châtel and V. Pons, Counting smaller elements in the Tamari and m-Tamari lattices. *J. Combin. Theory Ser. A*, **134** (2015), 58–97. MR3345297

[23] Y. Choi and N. Sun, The image of the Pop operator on various lattices. arXiv:2209.13695v1.

[24] A. Claesson and B. Á. Guðmundsson, Enumerating permutations sortable by k passes through a pop-stack. *Adv. Appl. Math.*, **108** (2019), 79–96. MR3940979

[25] A. Claesson, B. Á. Guðmundsson, and J. Pantone, Counting pop-stacked permutations in polynomial time. *Experiment. Math.* (2021). MR4574423

[26] H. Cohn, N. Elkies, and J. Propp, Local statistics for random domino tilings of the Aztec diamond. *Duke Math. J.*, **85** (1996), 117–166. MR1412441

[27] D. Dauvergne, The Archimedian limit of random sorting networks. *J. Amer. Math. Soc.*, **35** (2022), 1215–1267. MR4467309

[28] D. Dauvergne and B. Virág, Circular support in random sorting networks. *Trans. Amer. Math. Soc.*, **373** (2020), 1529–1553. MR4068272

[29] C. Defant, Pop-stack-sorting for Coxeter groups. *Comb. Theory*, **2** (2022). MR4498592

[30] C. Defant, Meeting covered elements in ν-Tamari lattices. *Adv. Appl. Math.*, **134** (2022). MR4348238

[31] C. Defant, R. Li, and E. Nestoridi, Rowmotion Markov chains. arXiv:2212.14005v1.

[32] C. Defant and J. Lin, Rowmotion on m-Tamari and biCambrian lattices. arXiv:2208.10464v1.

[33] C. Defant and N. Williams, Semidistrim lattices. *Forum Math. Sigma*, **11** (2023). MR4603109

[34] J. E. Goodman and R. Pollack, A combinatorial perspective on some problems in geometry. *Congr. Numer.*, **32** (1981), 383–394. MR0681897

[35] S. Fishel and L. Nelson, Chains of maximum length in the Tamari lattice. *Proc. Amer. Math. Soc.*, **142** (2014), 3343–3353. MR3238412

[36] L. Hong, The pop-stack-sorting operator on Tamari lattices. *Adv. Appl. Math.*, **139** (2022). MR4432015

[37] C. Ingalls and H. Thomas, Noncrossing partitions and representations of quivers. *Compos. Math.*, **145** (2009), 1533–1562. MR2575093

[38] W. Jockusch, J. Propp, and P. Shor, Random domino tilings and the arctic circle theorem. arXiv:math/9801068v1.

[39] L. Lichev, Lower bound on the running time of pop-stack sorting on a random permutation. arXiv:2212.09316v1.

[40] L. Nelson, A recursion on maximal chains in the Tamari lattices. *Discrete Math.*, **340** (2017), 661–677. MR3603545

[41] L.-F. Préville-Ratelle and X. Viennot, An extension of Tamari lattices. *Trans. Amer. Math. Soc.*, **369** (2017), 5219–5239. (Assigned the incorrect title "The enumeration of generalized Tamari intervals" by the journal.) MR3632566

[42] L. Pudwell and R. Smith, Two-stack-sorting with pop stacks. *Australas. J. Combin.*, **74** (2019), 179–195. MR3939483

[43] N. Reading, Cambrian lattices. *Adv. Math.*, **205** (2006), 313–353. MR2258260

[44] N. Reading, Clusters, Coxeter-sortable elements and noncrossing partitions. *Trans. Amer. Math. Soc.*, **359** (2007), 5931–5958 MR2336311

[45] N. Reading, Sortable elements and Cambrian lattices. *Algebra Universalis*, **56** (2007), 411–437. MR2318219

[46] N. Reading and D. E. Speyer, Cambrian fans. *J. Eur. Math. Soc.*, **11** (2009), 407–447. MR2486939

[47] N. Reading and D. E. Speyer, Combinatorial frameworks for cluster algebras. *Int. Math. Res. Not. IMRN*, **2016** (2016), 109–173. MR3514060

[48] D. Romik, The Surprising Mathematics of Longest Increasing Subsequences. Cambridge University Press, 2015. MR3468738

[49] H. Rost, Non-equilibrium behaviour of a many particle process: Density profile and local equilibria. *Z. Wahrsch. Verw. Gebiete*, **58** (1981), 41–53. MR0635270

[50] A. Sapounakis, I. Tasoulas, and P. Tsikouras, On the dominance partial ordering on Dyck paths. *J. Integer Seq.*, **9** (2006). MR2217231

[51] P. R. Scott, On the sets of directions determined by n points. *Amer. Math. Monthly*, **77** (1970), 502–505. MR0262933

[52] R. Smith and V. Vatter, A stack and a pop stack in series. *Australas. J. Combin.*, **58** (2014), 157–171. MR3211776

[53] R. P. Stanley, Enumerative Combinatorics, vol. 1, second edition. Cambridge University Press, 2012. MR2868112

[54] C. Stump, H. Thomas, and N. Williams, Cataland: Why the Fuß? To appear in *Mem. Amer. Math. Soc.* MR4111284

[55] D. Tamari, The algebra of bracketings and their enumeration. *Nieuw Archief voor Wiskunde*, **10** (1962), 131–146. MR0146227

[56] H. Thomas, An analogue of distributivity for ungraded lattices. *Order*, **23** (2006), 249–269. MR2308910

[57] H. Thomas and N. Williams, Rowmotion in slow motion. *Proc. Lond. Math. Soc.*, **119** (2019), 1149–1178. MR3968720

[58] P. Ungar, $2N$ noncollinear points determine at least $2N$ directions. *J. Combin. Theory Ser. A*, **33** (1982), 343–347. MR0676751

[59] X. G. Viennot, Heaps of Pieces, I. Basic Definitions and Combinatorial Lemmas. Lecture Notes in Math., vol. 1234. Springer, Berlin, 1986, 321–350. MR0927773

Electronic Journal of Probability
Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards

- Free for authors, free for readers

- Quick publication (no backlog)

- Secure publication (LOCKSS[1])

- Easy interface (EJMS[2])

Economical model of EJP-ECP

- Non profit, sponsored by IMS[3], BS[4], ProjectEuclid[5]

- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund[6] (click here to donate!)

- Submit your best articles to EJP-ECP

- Choose EJP-ECP over for-profit journals

---

[1] LOCKSS: Lots of Copies Keep Stuff Safe http://www.lockss.org/
[2] EJMS: Electronic Journal Management System: https://vtex.lt/services/ejms-peer-review/
[3] IMS: Institute of Mathematical Statistics http://www.imstat.org/
[4] BS: Bernoulli Society http://www.bernoulli-society.org/
[5] Project Euclid: https://projecteuclid.org/
[6] IMS Open Access Fund: https://imstat.org/shop/donation/