



On the Passage Time Geometry of the Last Passage Percolation Problem

Tom Alberts and Eric Cator

Department of Mathematics, University of Utah
155 S 1400 E, Rm 233,
Salt Lake City, Utah, USA.

E-mail address: alberts@math.utah.edu

URL: <http://www.math.utah.edu/~alberts/>

Department of Mathematics, Radboud University
PO Box 9010, 6500 GL
Nijmegen, the Netherlands.

E-mail address: e.cator@math.ru.nl

URL: <http://www.ru.nl/english/people/cator-e/>

Abstract. We analyze the geometrical structure of the passage times in the last passage percolation model. Viewing the passage time as a piecewise linear function of the weights we determine the domains of the various pieces, which are the subsets of the weight space that make a given path the longest one. We focus on the case when all weights are assumed to be positive, and as a result each domain is a pointed polyhedral cone. We determine the extreme rays, facets, and two-dimensional faces of each cone, and also review a well-known simplicial decomposition of the maximal cones via the so-called order cone. All geometric properties are derived using arguments phrased in terms of the last passage model itself. Our motivation is to understand path probabilities of the extremal corner paths on rectangles in \mathbb{Z}^2 , but all of our arguments apply to general, finite partially ordered sets.

1. Introduction

Last passage percolation is a well-studied model in probability theory that is simple to state but notoriously difficult to analyze. In recent years it has been shown to be related to many seemingly unrelated things: longest increasing subsequences in random permutations, eigenvalues of random matrices, long-time asymptotics of solutions to stochastic partial differential equations, and much more. All of

Received by the editors July 3rd, 2019; accepted October 28th, 2020.

2010 Mathematics Subject Classification. 60D05, 60K35, 06A11.

Key words and phrases. Last Passage Percolation, Convex Geometry, Partially Ordered Sets.

Research of Tom Alberts supported by supported by Simons Collaboration Grant 351687 and National Science Foundation grants DMS-1811087 and DMS-1715680.

these problems are of great interest due to the asymptotic behavior of various related statistics, neither of which are predicted by the classical strong law of large numbers or central limit theorem. The last passage model has been a particularly fertile ground for exploring this new frontier of probability theory due to its rich *solvability* structure. For certain choices of the random inputs the last passage model can be analyzed exactly, through various connections with representation theory of the symmetric group and rings of symmetric polynomials.

We briefly recall the setup of the last passage percolation model on \mathbb{Z}^2 . Consider the rectangle of integer points in \mathbb{Z}^2 with lower left corner at $(1, 1)$ and upper right corner at (m, n) , where $m, n > 1$. At each of the $m \cdot n$ integer points (i, j) we place a random variable $\omega(i, j)$ (a weight). The variables are typically assumed to be independent and identically distributed (iid) across points, and in this paper we will assume that they are always positive. We then consider the set $\mathcal{P}(m, n)$ of *up-right* paths from $(1, 1)$ to (m, n) , an up-right path being one whose steps are always either $(1, 0)$ or $(0, 1)$. To each $\gamma \in \mathcal{P}(m, n)$ we assign a random *length* $\ell(\gamma)$ that is the sum of the ω along the path, i.e.

$$\ell(\gamma) = \sum_{(i,j) \in \gamma} \omega(i, j).$$

In a lot of the literature on last passage percolation, either the start weight or the end weight of the path is left out in the length, so that concatenating paths is easier. It turns out that in our description, it is more convenient to consider all weights. Last passage percolation studies the maximal length over all paths, also known as the *passage time*:

$$G(m, n) := \max_{\gamma \in \mathcal{P}(m, n)} \ell(\gamma) = \max_{\gamma \in \mathcal{P}(m, n)} \sum_{(i,j) \in \gamma} \omega(i, j).$$

The passage time $G(m, n)$ is itself a random variable but its statistical distribution (or law) is very complicated. For any *fixed* path γ , the law of the length $\ell(\gamma)$ is well understood by the Strong Law of Large Numbers and the Central Limit Theorem. The maximum length, however, is determined by the joint law which describes the statistics of the entire collection of random lengths, and the complicating feature is that there is a very strong *correlation* between these different lengths. Whenever two paths share common vertices the random numbers at those vertices both contribute to their lengths, and so knowing the length of one path gives information about the length of the other. The more two paths intersect, the greater the correlation between their random lengths, and since there are $(m+n-2)!/((m-1)!(n-1)!)$ paths but only $m \cdot n$ vertices the correlation effects are significant.

Remarkably though, these correlation effects can be overcome when the choice of the input weights is assumed to be iid across vertices (i, j) and coming from either the geometric or exponential distribution. In these cases *exact* formulas can be computed for the distribution function of $G(m, n)$. The formulas are somewhat complicated, however, and typically involve determinants of linear operators on the sequence space ℓ^2 , with the operator determined by certain families of orthogonal polynomials. Nonetheless, the formulas are somewhat explicit and tractable enough to perform asymptotic analysis as $m, n \rightarrow \infty$. Three very important and well known such asymptotic results are the following:

- the *limit shape* (the almost sure, non-random limit of $G(\lfloor nx \rfloor, \lfloor ny \rfloor)/n$ as $n \rightarrow \infty$, as a function of x and y , whose existence follows from Kingman's subadditive ergodic theorem),
- the magnitude and distribution of the fluctuations of the passage time $G(\lfloor nx \rfloor, \lfloor ny \rfloor)$ as $n \rightarrow \infty$ (the growth of the fluctuations being $n^{1/3}$ and the convergence of the centered and appropriately normalized passage time to the Tracy-Widom law),
- and the magnitude of the transversal fluctuations of the maximal path away from the diagonal (the maximal path from $(1, 1)$ to (n, n) is thought to go distance $n^{2/3}$ away from the main diagonal, and is known to do so in certain solvable models).

The limit shape results are originally due to exact bijections between LPP and the TASEP process (Rost, 1981; Aldous and Diaconis, 1995; Seppäläinen, 1997), or for stationary models that exist for certain special weight distributions (O'Connell and Yor, 2001; Cator and Groeneboom, 2005; Balázs et al., 2006; Seppäläinen, 2012). More recent work (Georgiou et al., 2016, 2017) provides variational formulas for the limit shape for very general weight distributions in terms of infinite dimensional objects called cocycles, although obtaining explicit results for these formulas is generally difficult. Exact Tracy-Widom limits for fluctuations are originally based on connections with generalized permutations and the Robinson-Schensted-Knuth algorithm (Johansson, 1998), often based on ideas from random matrix theory (see also Priezzhev and Schütz, 2008; Georgiou and Ortmann, 2018 in the case of Bernoulli weights). In some cases these results have been re-understood through different means Johansson (2010), but in general all methods to date require a special choice of random input (see also Corwin, 2014, 2018 for further references).

Nonetheless, it is widely believed that there is a certain *universality* aspect to the last passage model. This specifically refers to the distribution of the fluctuations of the passage time $G(\lfloor nx \rfloor, \lfloor ny \rfloor)$ as $n \rightarrow \infty$, which is believed to be the same Tracy-Widom law for a wide class of random inputs, not just the special cases mentioned above. This is analogous to the Central Limit Theorem for sums of iid random variables, where the fluctuations of the sum follow the Gaussian distribution for a very broad class of input variables. While universality in the Central Limit Theorem is now understood via many different techniques and proofs, less progress has been made for universality results of the last passage model.

This paper explores a possible method for studying various aspects of the last passage model using tools from combinatorics and geometry. The main idea is to embed the model into a suitable high-dimensional space, determined by the random input weights, and in this space study the geometry of the last passage problem. The basic setup is relatively simple. For any fixed path γ its length $\ell(\gamma)$ is clearly a linear function of the weights ω , and therefore the passage time $G(m, n)$ is piecewise linear. The main purpose of this article is to determine the geometry of the domains of the pieces, each one of which corresponds to a different path. The main strength of this approach is that it is purely geometric, with no probabilistic input at all until a measure is put on the space of weights. This flexibility allows one to study many different types of random inputs with the same underlying geometric framework, and it is our hope that it will allow for a new conceptual framework for the last passage problem while at the same time shedding new light onto previously solvable models. Our interest in this approach was primarily driven by one simple question:

among all paths in $\mathcal{P}(n, n)$ (take $m = n$ for simplicity), which one is the most likely to be the maximizer? Even in the exactly solvable cases this does not seem to be an easy question to answer, as we explain later in Section 6. While much attention has been paid to the paths with transversal fluctuations $n^{2/3}$, in particular the recent work [Dauvergne et al. \(2018\)](#) proves the existence of scaling limit for these objects in terms of the so-called Airy sheet ([Corwin et al., 2015](#)), less attention has been paid to the more extreme paths. We are quite confident that the most likely maximal path is the extreme one that goes straight up from $(1, 1)$ to $(1, n)$ and then straight right from $(1, n)$ to (n, n) (or its symmetric version that goes right and then up). We do not have a proof but the intuition is straightforward: the weights that are picked up by the extreme path are shared by a relatively small number of other paths, and therefore the extremal path should have a much larger portion of the environment space in which it is longest. For example, the extremal path gets the weight at $(1, n)$ entirely to itself. In contrast, the paths going through the interior share the weights they pick up with *many* other paths, meaning each individual path has a hard time distinguishing itself as the longest. In fact, we expect that the probability of the extremal path being longest is substantially larger (in n) than the probability of the middle path being longest (the middle path being the one that alternates between up and right steps). This heuristic fits with the expectation that the transversal fluctuations are larger than the $n^{1/2}$ magnitude obtained by the uniform measure on paths. In fact some sort of behavior of this type seems necessary to obtain superdiffusive fluctuations, although on its own it does not explain why the magnitude of the fluctuations should be precisely $n^{2/3}$. We expect that the $n^{2/3}$ corresponds to the region where the low probabilities for the “middle paths” balances out the fact that the bulk of the paths are in the middle. In other words, even though we believe that the extremal path (which has transversal fluctuation of order n) is the mode of the path distribution, a typical sample from the path distribution has transversal fluctuations of much smaller magnitude because there are so many more paths there.

In the last section of the paper we discuss some other results that we believe follow from this intuition, such as a negative correlation result between the location of the maximizer and the path length itself. The present paper comes from a desire to solidify our intuition by understanding more about the structure of the parts of environment space that makes a given path the longest. We also hope it will help to separate out how much of the expected universal behavior is due to the geometry of the last passage time function and how much is due to the particular probability distribution on the weights. One advantage of our framework is that it extends beyond the traditional study of LPP on \mathbb{Z}^2 . In fact all that is required is a notion of directedness, which allows us to carry out the analysis on general finite posets.

General Setup and Main Results. Although we are largely motivated by the last passage problem on \mathbb{Z}^2 , our approach assumes nothing other than the paths being *directed*. On \mathbb{Z}^2 this is forced by the assumption that paths are up-right (and hence not allowed to go backwards), but in fullest generality we can force a direction by studying the problem on an arbitrary partially ordered set (poset). This has the advantage of allowing for different correlation structures among path lengths, which is determined by the structure of the underlying poset as follows.

Let (P, \leq) be a finite poset. The general last passage problem will be considered on P , which we often think of in terms of its Hasse diagram (definition below), and so we will commonly refer to the elements of P as vertices. We will often say that vertices $\mathbf{u}, \mathbf{v} \in P$ are *in order* if $\mathbf{u} \leq \mathbf{v}$ or $\mathbf{v} \leq \mathbf{u}$, and otherwise we say they are *out of order*. We will assume throughout that P is connected, meaning that its Hasse diagram is connected as a graph, since otherwise we may consider the problem individually on the different connected components. The cover relations of the poset will be denoted by \prec , where we recall that for $\mathbf{v}, \mathbf{w} \in P$, $\mathbf{v} \prec \mathbf{w}$ means that $\mathbf{v} < \mathbf{w}$ and there is no $\mathbf{u} \in P$ such that $\mathbf{v} < \mathbf{u} < \mathbf{w}$. The Hasse diagram of P is the directed graph with P as its vertex set and an edge from \mathbf{v} to \mathbf{u} iff $\mathbf{v} \prec \mathbf{u}$. On a general poset the paths of the last passage problem are the *maximal chains* of P , the set of which we denote by Π_P . Recall that a maximal chain is an ordered subset $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ of P such that $\mathbf{v}_1 \prec \mathbf{v}_2 \prec \dots \prec \mathbf{v}_n$ and there are no elements \mathbf{u} or \mathbf{w} such that $\mathbf{u} \prec \mathbf{v}_1$ or $\mathbf{v}_n \prec \mathbf{w}$. Intuitively we see that this corresponds to all nearest-neighbor paths in the Hasse diagram of P that are as “long” as possible.

For the (positive weight) last passage problem on P we place a weight $\omega(\mathbf{v}) \in \mathbb{R}_+ = [0, \infty)$ on each element $\mathbf{v} \in P$. The vector $\omega \in \mathbb{R}_+^P$ is collectively referred to as the weight, and the length of each element of $\pi \in \Pi_P$ is the the sum of the weights along the path:

$$\ell(\pi) := \sum_{\mathbf{v} \in \pi} \omega(\mathbf{v}).$$

Note that we can naturally associate each path $\pi \in \Pi_P$ to a vector in \mathbb{R}_+^P (which we also call π) via $\pi(\mathbf{v}) = \mathbf{1}\{\mathbf{v} \in \pi\}$, where $\mathbf{1}$ is the indicator function. Via this association we have that the length is simply the standard inner product between the path and the weight vector, i.e.

$$\ell(\pi) = \langle \omega, \pi \rangle.$$

The *passage time* of the poset P , under the weight vector ω , is the largest length of all possible paths, i.e.

$$G_P = G_P(\omega) := \max_{\pi \in \Pi_P} \langle \omega, \pi \rangle.$$

More generally we may consider the vector of passage times determined by the weight vector ω , which encodes the length of the longest path up to each given vertex and is defined as

$$G_P(\mathbf{v}) = G_P(\mathbf{v}; \omega) := \max_{\pi \in \Pi_P(\mathbf{v})} \langle \omega, \pi \rangle,$$

where $\Pi_P(\mathbf{v})$ is the set of all maximal chains in the subposet of elements below \mathbf{v} (the so-called lower set of \mathbf{v} , see below for a definition), extended to \mathbb{R}^P by adding zeros. Then the passage time can be written as

$$G_P = \max_{\mathbf{v} \in P} G_P(\mathbf{v}).$$

The collection of passage times $G_P(\mathbf{v})$ can also be built up from the weight vector ω via the recursion

$$G_P(\mathbf{v}) = \omega(\mathbf{v}) + \max_{\mathbf{u} \prec \mathbf{v}} G_P(\mathbf{u}) \tag{1.1}$$

with the “initial condition” $G_P(\mathbf{v}) = \omega(\mathbf{v})$ if \mathbf{v} is a minimal element of P . Conversely, given the vector of passage times $G_P(\mathbf{v})$ this recursion can be inverted to solve for the corresponding weight vector ω via

$$\omega(\mathbf{v}) = G_P(\mathbf{v}) - \max_{\mathbf{u}:\mathbf{u}<\mathbf{v}} G_P(\mathbf{u}),$$

again with $G_P(\mathbf{v}) = \omega(\mathbf{v})$ for \mathbf{v} minimal. Regardless of how G_P is constructed, for P fixed and ω allowed to vary, this definition implies that G_P is a piecewise linear function of ω , and the main purpose of this article is to determine the regions on which the function is equal to each of the various linear maps that define it. Since in this case the maps are defined by the paths π there is a natural region in \mathbb{R}_+^P associated to each path: the set of weight vectors ω that give path π the longest weight. More precisely, this is the set

$$\mathcal{C}(\pi) := \{\omega \in \mathbb{R}_+^P : G_P(\omega) = \langle \omega, \pi \rangle\} = \{\omega \in \mathbb{R}_+^P : \langle \omega, \pi \rangle \geq \langle \omega, \pi' \rangle \text{ for all } \pi' \in \Pi_P\}. \quad (1.2)$$

From this definition and especially the second equality we immediately see that each set $\mathcal{C}(\pi)$ is a *polyhedral cone*, namely a finite intersection of half-spaces of \mathbb{R}^P . The inequalities defining the half-spaces are those of the form $\langle \omega, \pi - \pi' \rangle \geq 0$, with π and π' regarded as vectors in \mathbb{R}_+^P , but also those inequalities implicitly given by the condition that the cone is a subset of \mathbb{R}_+^P . Explicitly, the cone being a subset of \mathbb{R}_+^P means that $\omega(\mathbf{v}) \geq 0$ for all $\mathbf{v} \in P$, and these inequalities are an important part of (1.2).

As with all polyhedral cones the sets $\mathcal{C}(\pi)$ are both convex and invariant under positive scaling, since the half-spaces that define them are also and these properties are preserved under intersection. Both properties also follow from their interpretation via the last passage model, since if two weight vectors make the same path maximal then so does their sum and any positive scalar multiple.

Beyond the fact that the maximal sets are polyhedral cones, a more detailed description of the structure of the sets is required to perform any meaningful analysis. There are two common descriptions of a polyhedral cone: via the set of half-spaces that bound it (the *H-decomposition*), or via the extreme rays that span it (the *V-decomposition*). For a polyhedral cone there are at most finitely many half-spaces and extreme rays that define it, and in this article we will determine both for each given path π . It turns out that both descriptions have a very beautiful structure, and moreover can be determined solely by working with their description in terms of the last passage model. The V-decomposition is already known in [Stanley \(1986\)](#) but our argument is different in that it is phrased in terms of the last passage model. To the best of our knowledge our determination of the H-decomposition is new, and we regard it as the most significant of our results. In both cases the idea is that the geometry of each maximal set $\mathcal{C}(\pi)$ is naturally encoded in the poset P , and our arguments are based on a comparison of π to the other paths in Π_P . In an intuitive sense we are analyzing the ability of the other paths to compete with π to be the maximal one. We will prove two main theorems on the *H* and *V* decompositions of the sets $\mathcal{C}(\pi)$. We begin with the *H*-decomposition, which requires the following definitions.

Definition 1.1. A vertex $\mathbf{v} \in P$ is a corner of a path $\pi \in \Pi_P$ if $\mathbf{v} \in \pi$ and there exists another path $\pi' \in \Pi_P$ such that $\pi \setminus \mathbf{v} \subset \pi'$ and $\mathbf{v} \notin \pi'$.

The notion of a corner of a path can be easily visualized in the Hasse diagram of the poset.

Definition 1.2. Fix $\pi, \pi' \in \Pi_P$. The *disorder graph* of π and π' , denoted by $\Delta(\pi, \pi')$, has as its vertex set the symmetric difference $\pi \Delta \pi'$ (with π and π' seen as subsets of the poset), and an edge connecting $\mathbf{u}, \mathbf{v} \in \pi \Delta \pi'$ iff \mathbf{u} and \mathbf{v} are out of order.

An important property of $\Delta(\pi, \pi')$ is that it constitutes a bipartite graph, with one part consisting of the vertices belonging to π and the other part consisting of the vertices belonging to π' . The bipartiteness follows because two vertices belonging to the same path are always in order. See Figure 1.1 for examples of the disorder graphs.

With these definitions in hand we now state the H -decomposition.

Theorem 1.3 (*H*-decomposition of maximal sets). *For each path $\pi \in \Pi_P$ the minimal set of inequalities that define the cone $\mathcal{C}(\pi)$ are those of the form:*

- i) $\omega(\mathbf{v}) \geq 0$ for $\mathbf{v} \in P \setminus \pi$,
- ii) $\omega(\mathbf{v}) \geq 0$ for $\mathbf{v} \in \pi$ but not a corner of π ,
- iii) $\langle \omega, \pi - \pi' \rangle \geq 0$ for paths $\pi' \in \Pi_P \setminus \{\pi\}$, whose disorder graph with π is connected.

The second and third conditions are motivated by their meaning on \mathbb{Z}^2 , which for the second condition is geometrically intuitive and in the third condition means that π and π' form at most a single loop. See Figure 1.1 and the remark below for more.

Note that the statement of Theorem 1.3 is that this is the minimal set of inequalities needed to define the cone, so that removing any one of them would lead to a larger set than $\mathcal{C}(\pi)$. These inequalities define the *facets* of the cone, the co-dimension one boundary sets of $\mathcal{C}(\pi)$. Note that (1.2) already defines $\mathcal{C}(\pi)$ via these various inequalities, but what Theorem 1.3 amounts to showing is that many of these facets are *redundant*. Reducing the inequalities to only the irredundant ones allows for a fuller analysis of the cone, and is usually required for computational algorithms. In more probabilistic language Theorem 1.3 is describing the event that a given path π is longest in terms of the smallest possible number of intersections of events of the form $\{\pi$ is longer than $\pi'\}$.

Remark 1.4. For posets of the form $[1, m] \times [1, n]$ in \mathbb{Z}^2 the condition that the disorder graph $\Delta(\pi, \pi')$ is connected is equivalent to saying that π and π' form a single loop. That is, π and π' may start out the same, diverge for a while, and then recombine with each other, but after recombining cannot diverge again. Diverging more than once would mean that there are multiple loops between π and π' , which is equivalent to saying that $\pi - \pi'$ can be written as the sum of the individual loops. Since each individual loop is already a face of $\mathcal{C}(\pi)$ the sum is redundant. It is also easy to see that π and π' forming multiple loops is equivalent to the disorder graph being disconnected, since between any two consecutive loops there is a subpath in $\pi \cap \pi'$ that connects the loops together. This subpath prevents the two loops from being connected in $\Delta(\pi, \pi')$. See Figure 1.1.

We prove Theorem 1.3 for the H -representation in Section 3. In Section 4 we also describe a related object called the *order graph*, which forms connections based on ordering relations between the supported vertices of two extreme rays of $\mathcal{C}(\pi)$ (see the start of Section 2 for the definition of extreme rays). We use the order graph

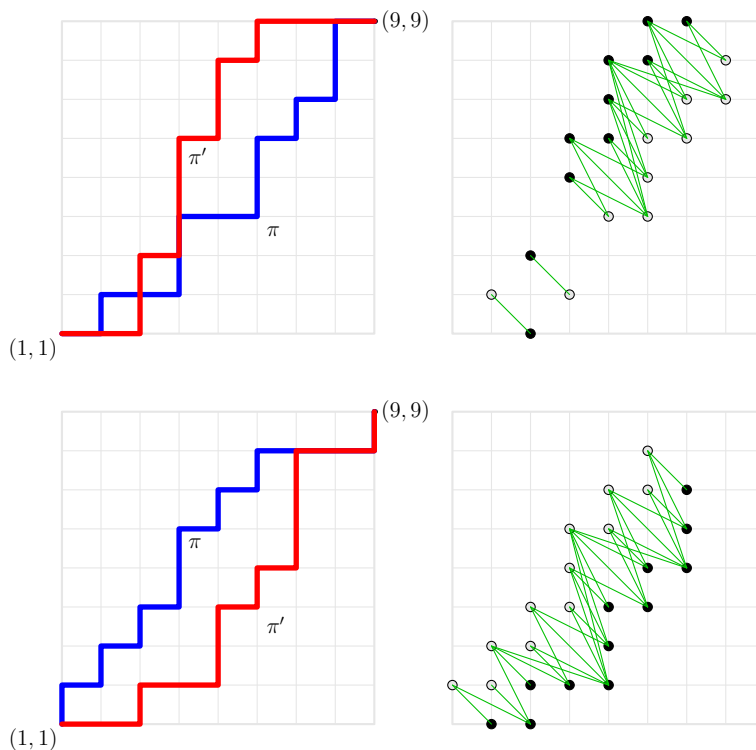


FIGURE 1.1. Two examples of paths π and π' on the subposet $[1, 9]^2$ of \mathbb{Z}^2 , with the associated disorder graphs on the right. In the disorder graph elements of $\pi \setminus \pi'$ are in grey while elements of $\pi' \setminus \pi$ are in black. In the first example the disorder graph is disconnected, while in the second graph it is connected. On rectangles in \mathbb{Z}^2 connectedness of the disorder graph is equivalent to π and π' forming a “single loop”.

to determine when two extreme rays of $\mathcal{C}(\pi)$ form a two-dimensional boundary face of the cone, see Theorem 4.2.

To describe the extreme rays of the cones requires the notion of an *antichain* of the poset P and a particular geometric embedding of it, which we define next.

Definition 1.5. An *antichain* of the poset P is a subset of P such that no two elements are in order. We will naturally embed an antichain $A \subset P$ into an element $a \in \mathbb{R}_+^P$ via $a(\mathbf{v}) = \mathbf{1}\{\mathbf{v} \in A\}$.

In particular, an antichain can contain at most one vertex from a given path in Π_P , since by definition the elements along a path are in complete order with each other. This leads to the following theorem.

Theorem 1.6 (*V-decomposition of maximal sets*). *For each path $\pi \in \Pi_P$ the extreme rays of the cone $\mathcal{C}(\pi)$ are precisely the geometric embeddings of the antichains which intersect π exactly once.*

Theorem 1.6 gives the following characterization of an extreme ray: a vector in \mathbb{R}_+^P is an extreme ray of $\mathcal{C}(\pi)$ if and only if there is one non-zero entry along the

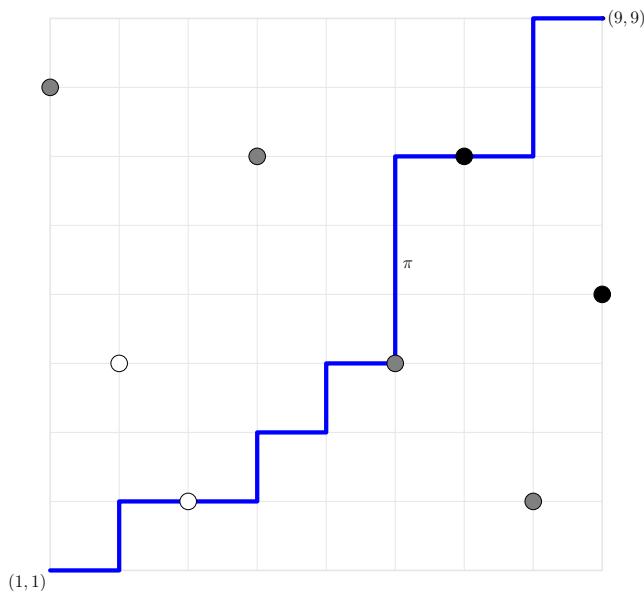


FIGURE 1.2. A single path π on the subposet $[1, 9]^2$ of \mathbb{Z}^2 , and three of the antichains of the poset that intersect π . The two vertices in white form one antichain, the four in gray form a second, and the two in black form the third. Note that the vertices within an antichain must be out of order with each other, the vertices in two distinct antichains do not need to be out of order. To convert these antichains into extreme rays of $\mathcal{C}(\pi)$ simply put an entry of 1 at each element of the antichain, and an entry of zero elsewhere.

path π , and all other non-zero entries have the same value and are out of order with each other in the poset. We will canonically take the non-zero entry to be 1, although by scaling invariance it could clearly be any positive value. Commonly we will use the notation:

Definition 1.7. For each fixed $\pi \in \Pi_P$ we let $\mathcal{ER}(\pi)$ denote the set of extreme rays of the polyhedral cone $\mathcal{C}(\pi)$.

The structure of extreme rays is essentially already stated by Stanley [Stanley \(1986\)](#) through what he calls the *chain polytope*. See also the earlier works referenced within [Stanley \(1986\)](#). The chain polytope can be formed by intersecting each maximal cone $\mathcal{C}(\pi)$ with the unit cube $[0, 1]^P$ and then taking the union of what remains over all paths π . Our description of the extreme rays for each individual cone $\mathcal{C}(\pi)$ is not a very extensive refinement of Stanley’s result, but our proof is different in that it is framed entirely in terms of the last passage model. The antichains turn out to be precisely the directions in which one can perturb the path lengths while keeping the longest path the longest, and this turns out to be the key argument in our proof. This is laid out in [Section 2](#).

The explicit structure of $\mathcal{ER}(\pi)$ also allows us to determine its size for certain types of posets, in particular for $[1, m] \times [1, n] \subset \mathbb{Z}^2$. See [Theorem 2.9](#). For all but pathological posets and paths the number of extreme rays in $\mathcal{ER}(\pi)$ is much

greater than the dimension $|P|$ of the ambient space \mathbb{R}^P , meaning that the maximal cones $\mathcal{C}(\pi)$ are far from simplicial. Nonetheless it is possible to use the extreme rays in $\mathcal{ER}(\pi)$ to give an explicit simplicial decomposition of each maximal cone $\mathcal{C}(\pi)$, without the need to introduce additional rays.

Theorem 1.8. *For each $\pi \in \Pi_P$ there is a decomposition of $\mathcal{C}(\pi)$ into disjoint simplicial cones (disjoint up to boundary intersections) such that the extreme rays of each simplicial cone only use elements from $\mathcal{ER}(\pi)$.*

This theorem can be found in Stanley (1986, 2012) so we only explain it briefly in Section 5. On Young diagrams (which we regard as subposets of \mathbb{Z}^2) it is equivalent to using Young tableaux to partition the space into simplices. From this partitioning we obtain the following result:

Corollary 1.9. *There exists functions $\Lambda_1, \dots, \Lambda_{|P|} : \mathbb{R}_+^P \rightarrow \mathbb{R}$ such that*

$$G_P(\omega) = \sum_{i=1}^{|P|} \Lambda_i(\omega).$$

The main purpose of this corollary is that it converts a complicated maximum of random variables into a sum of the same number of random variables. While sums are usually easier to handle, the mapping from ω to Λ is piecewise linear and induces a complicated correlation structure on the Λ_i random variables, even when the underlying ω distribution is nice. This representation of the passage time as a sum of random variables is equivalent to the corner growth representation of the last passage model Seppalainen (2009); Romik (2015), in which the elements of the poset are “filled in” at random times that obey the ordering of the poset. We briefly explain this connection towards the end of Section 5. In Section 6 we give some explanation of how the iid exponential distribution for the weight variables interacts nicely with the geometry of the last passage function; this gives some additional intuition into why the exponential distribution tends to produce the most precise results. In Section 7 we describe how the geometrical description of the maximal cones can be used to give an alternative description of the passage time for iid Uniform(0, 1) weights, in terms of Stanley’s order cone (Stanley, 1986). Finally, in Section 8 we list a series of open problems that this work has led us to.

Terminology and Notation: Throughout we write $\mathbb{R}_+ = [0, \infty)$. Functions $\omega : P \rightarrow \mathbb{R}$ and vectors $\omega \in \mathbb{R}^P$ are identified in the natural way. Subsets $A \subset P$ are often associated with their indicator functions, considered as elements of \mathbb{R}^P . If $\omega \in \mathcal{C}(\pi)$ we often say “ π is maximal under ω ”, or equivalently that “ π is a maximal path for ω ”. Subset inclusions \subset and \subseteq are the same throughout, for strict inclusion we write \subsetneq .

2. Extreme Rays

In this section we concentrate on proving Theorem 1.6. We recall that for a vector ω to be an extreme ray of a polyhedral cone means that it can only be written as a (positive) weighted sum of (positive) multiples of itself, i.e. if $\omega, \omega_1, \omega_2$ are all in the same polyhedral cone then

$$\omega = \alpha\omega_1 + \beta\omega_2 \text{ with } \alpha, \beta > 0 \implies \omega_1, \omega_2 \in \text{Span}_+\{\omega\}.$$

Another way of saying this is that the only linear subspace of directions in which one can move infinitesimally away from ω and still remain in the cone is $\text{Span}\{\omega\}$. We will use this type of argument throughout our analysis, which leads to the following definition:

Definition 2.1. For a path $\pi \in \Pi_P$ and a vector $\omega \in \mathcal{C}(\pi)$, we define the *perturbation space* of ω in the cone $\mathcal{C}(\pi)$ by

$$\mathcal{D}_\pi(\omega) = \{\sigma \in \mathbb{R}^P : \exists \epsilon > 0 \text{ such that } \omega \pm \epsilon\sigma \in \mathcal{C}(\pi)\}.$$

If $\omega \notin \mathcal{C}(\pi)$ we set $\mathcal{D}_\pi(\omega) = \emptyset$.

Note that as long as $\omega \in \mathcal{C}(\pi)$ then scale invariance of the cone implies that $\text{Span}\{\omega\} \subset \mathcal{D}_\pi(\omega)$. Furthermore, from this definition it is straightforward to verify:

Lemma 2.2. *The perturbation space $\mathcal{D}_\pi(\omega)$ is a linear subspace of \mathbb{R}^P . Moreover, if $\omega \in \mathcal{C}(\pi)$ then it is an extreme ray of $\mathcal{C}(\pi)$ if and only if $\mathcal{D}_\pi(\omega) = \text{Span}\{\omega\}$.*

The previous lemma will be our key tool for proving Theorem 1.6. First we will show that all antichains of P that intersect π have only their span in their perturbation space, and then conversely that all weight vectors that make a given path maximal and have only their span in their perturbation space must be maximal. To this end we first note the following simplification.

Remark 2.3. Fix a path π . Then to determine the perturbation space of a vector $\omega \in \mathcal{C}(\pi)$ it is enough to consider only the non-zero entries of the vector that can be perturbed. Indeed, the zero entries can never be perturbed since necessarily the perturbation in either the positive or negative direction will take them out of \mathbb{R}_+^P , which violates that $\mathcal{C}(\pi)$ is a subset of \mathbb{R}_+^P . We will use this simple fact repeatedly so we define:

Definition 2.4. For a vector $\omega \in \mathbb{R}^P$ we define the support of ω to be the subset of vertices on which ω is non-zero, i.e. $\text{supp}(\omega) := \{\mathbf{v} \in P : \omega(\mathbf{v}) \neq 0\}$. We let $\text{zero}(\omega) = P \setminus \text{supp}(\omega) = \{\mathbf{v} \in P : \omega(\mathbf{v}) = 0\}$.

Remark 2.5. In all proofs that follow we assume that $\text{supp}(\omega) \neq \emptyset$, since that only happens in the trivial case of ω equal to the zero vector.

Proof of Theorem 1.6: Antichains are extreme rays: Fix $\pi \in \Pi_P$ and suppose that a is an antichain of P . The first observation is that all paths have length either 0 or 1 under the weight vector $\omega = a$ (recall that we naturally associate the subset a with a vector through its indicator function; see Definition 1.5). This follows since the vertices along a path are in order and the elements of a are completely out of order, hence a given path can intersect a at most once. If a is non-zero at some vertex of π then clearly π has length 1 under a and therefore is maximal under a , i.e. $a \in \mathcal{C}(\pi)$.

Now let $\sigma \in \mathcal{D}_\pi(a)$. By the last remark we can assume that σ has zero entries at all vertices where a has zero entries, i.e. $\text{supp}(\sigma) \subset \text{supp}(a)$. Now suppose that σ is not constant on $\text{supp}(a)$. Choose $\epsilon > 0$ such that $a - \epsilon\sigma \geq 0$ (this is possible since a is strictly positive on $\text{supp}(a)$). Then under the weight vector $a + \epsilon\sigma$ the maximal paths are those which were maximal under a and pass through vertices in $\text{supp}(a)$ at which σ achieves its maximal value, which by the non-constancy assumption is not all of $\text{supp}(a)$. If π is not one of these paths then it is no longer one of the longest, so by definition $a + \epsilon\sigma \notin \mathcal{C}(\pi)$, and this holds for all $\epsilon > 0$. If π is one of

these paths then the non-constancy assumption means it cannot be longest under $a - \epsilon\sigma$, again for all ϵ small enough. Thus if σ is not constant on $\text{supp}(a)$, it cannot be in $\mathcal{D}_\pi(a)$, which completes the proof. \square

To prove the opposite direction is relatively simple but slightly lengthier, so we break the proof into several smaller supporting results. First recall the following terminology:

Definition 2.6. For $A \subset P$ the lower set $L(A)$ of A is the set of elements below A in P , i.e. $L(A) = \{\mathbf{u} \in P : \mathbf{u} \leq \mathbf{v} \text{ for some } \mathbf{v} \in A\}$. Similarly the upper set is $U(A) = \{\mathbf{u} \in P : \mathbf{u} \geq \mathbf{v} \text{ for some } \mathbf{v} \in A\}$. We also define their boundaries $\partial L(A)$ and $\partial U(A)$ as the maximal and minimal elements of $L(A)$ and $U(A)$, respectively.

Note that both $\partial L(A)$ and $\partial U(A)$ are antichains of P , and from this one immediately has that:

Lemma 2.7. *A subset $A \subset P$ is an antichain of P iff $\partial L(A) = A = \partial U(A)$.*

Proof: See Stanley (2012, Section 3.1), where upper and lower sets are referred to as *order ideals*. \square

We will use this lemma for the subset $\text{supp}(\omega)$ determined by a weight vector $\omega \in \mathbb{R}^P$. In particular we use it to show that:

Lemma 2.8. *If ω is an extreme ray of $\mathcal{C}(\pi)$ then necessarily $\text{supp}(\omega)$ is an antichain of P .*

Proof: For shorthand write $\partial L = \partial L(\text{supp}(\omega))$ and $\partial U = \partial U(\text{supp}(\omega))$. We will show that $\partial L = \partial U$. First observe that any path π' can pass through at most one element from each of ∂L and ∂U , since both are antichains and the path is ordered. Suppose π' is one of the paths which is longest under ω . Define $S = \pi' \cap \text{supp}(\omega)$. Since $S \subset \pi'$, it has a unique minimal element \mathbf{v} . If $\mathbf{v} \notin \partial U$, there exists $\mathbf{w} < \mathbf{v}$ with $\omega(\mathbf{w}) > 0$ (since $\mathbf{v} \in U(\text{supp}(\omega))$). Then we can construct a path π'' such that $\mathbf{w} \cup S \subset \pi''$, which would mean that π'' is longer than π' under ω ; contradiction. This shows that $\pi' \cap \partial U \neq \emptyset$, and in a similar manner we can show that $\pi' \cap \partial L \neq \emptyset$. Hence for all such π' we have that $\langle \pi', \mathbf{1}_{\partial L} \rangle = 1 = \langle \pi', \mathbf{1}_{\partial U} \rangle$, where $\mathbf{1}_{\partial L}$ and $\mathbf{1}_{\partial U}$ are the indicator functions of ∂L and ∂U , respectively.

Now let $\epsilon_1 = \min\{\langle \pi', \omega \rangle - \langle \pi'', \omega \rangle : \langle \pi', \omega \rangle > \langle \pi'', \omega \rangle\}$, $\epsilon_2 = \min\{\omega(\mathbf{v}) : \mathbf{v} \in \text{supp}(\omega)\}$ and $\epsilon = \min(\epsilon_1, \epsilon_2)$. Note that $\epsilon > 0$. Define a vector $\sigma \in \mathbb{R}^P$ by

$$\sigma = \frac{\epsilon}{2}(\mathbf{1}_{\partial L} - \mathbf{1}_{\partial U}).$$

Because $\epsilon \leq \epsilon_2$, we have that $\omega - \sigma \geq 0$ (because $\text{supp}(\sigma) \subset \partial L \cup \partial U \subset \text{supp}(\omega)$, so that σ is zero at all vertices where ω is). We proceed by contradiction. Using from above that $\langle \pi', \mathbf{1}_{\partial L} \rangle = 1 = \langle \pi', \mathbf{1}_{\partial U} \rangle$ for *any* π' which is longest under ω (including π itself), we have

$$\langle \pi', \omega \pm \sigma \rangle = \langle \pi', \omega \rangle \pm \frac{\epsilon}{2}(\langle \pi', \mathbf{1}_{\partial L} \rangle - \langle \pi', \mathbf{1}_{\partial U} \rangle) = \langle \pi', \omega \rangle,$$

so that all paths which were longest under ω are still longest under $\omega \pm \sigma$ (here we use that $\epsilon \leq \epsilon_1$, so that second longest paths cannot overtake any of the longest paths). In particular π itself is still a longest path, so $\omega \pm \sigma \in \mathcal{C}(\pi)$, which implies that $\sigma \in \mathcal{D}_\pi(\omega)$. But the assumption $\partial L \neq \partial U$ also gives that $\sigma \notin \text{Span}\{\omega\}$, which is a contradiction to ω being an extreme ray. \square

To complete the proof of Theorem 1.6 it only remains to be shown that each extreme ray must be constant on its support.

Proof of Theorem 1.6: Extreme rays must be constant on antichains: Let ω be an extreme ray of $\mathcal{C}(\pi)$. Then by the last lemma its support is an antichain of P . Thus π can pass through at most one element of $\text{supp}(\omega)$, but it must pass through at least one since otherwise its length would be zero and it could not be maximal.

Let \mathbf{v} be the element of $\text{supp}(\omega)$ that π passes through. If $\omega(\mathbf{v}) < \omega(\mathbf{w})$ for some $\mathbf{w} \in \text{supp}(\omega)$ then π could not have been longest under ω since any path that goes through \mathbf{w} would be longer. If, on the other hand, ω achieves its maximal value at \mathbf{v} then let $A = \{\mathbf{w} \in \text{supp}(\omega) : \omega(\mathbf{v}) > \omega(\mathbf{w})\}$. Then under ω the maximal paths are those which pass through $\text{supp}(\omega) \setminus A$. Let $\epsilon = \min\{\omega(\mathbf{v}) - \omega(\mathbf{w}) : \mathbf{w} \in A\}$, which we note is strictly positive, and then define a vector $\sigma \in \mathbb{R}^P$ by

$$\sigma = \frac{\epsilon}{3}(1_{\text{supp}(\omega) \setminus A} - 1_A).$$

Then any path π' which was longest under ω is still longest under $\omega \pm \sigma$, and hence $\sigma \in \mathcal{D}_\pi(\omega)$. But if $A \neq \emptyset$ then $\sigma \notin \text{Span}\{\omega\}$, and this says $\mathcal{D}_\pi(\omega)$ is strictly larger than $\text{Span}\{\omega\}$. This contradicts that ω is an extreme ray of $\mathcal{C}(\pi)$. \square

Finally, we end this section by proving a formula for the number of extreme rays of a maximal cone $\mathcal{C}(\pi)$ on the subposet $[1, m] \times [1, n]$ of \mathbb{Z}^2 .

Theorem 2.9. *Let $m, n > 1$ and $P = [1, m] \times [1, n]$ as a subposet of \mathbb{Z}^2 with the componentwise ordering. Write a path $\pi \in \Pi_P$ as the ordered collection of vertices (u_i, v_i) with $(u_1, v_1) = (1, 1)$, $(u_{m+n-1}, v_{m+n-1}) = (m, n)$, and $(u_{i+1}, v_{i+1}) - (u_i, v_i) \in \{(0, 1), (1, 0)\}$. Then the number of extreme rays in $\mathcal{ER}(\pi)$ is*

$$\sum_{i=1}^{m+n-1} \binom{n + u_i - v_i - 1}{u_i - 1} \binom{m - u_i + v_i - 1}{v_i - 1}. \tag{2.1}$$

Proof: Consider a rectangle of the form $[1, a] \times [1, b]$ for integers $a, b \geq 1$, considered as a subposet of \mathbb{Z}^2 . Any antichain of $[1, a] \times [1, b]$ can have at most one non-zero entry in each row and column, thus at most $a \wedge b$ non-zero entries overall. To construct antichains with exactly $k \geq 1$ non-zero entries do the following: independently choose subsets $A \subset \{1, \dots, a\}$ and $B \subset \{1, \dots, b\}$ with $|A| = |B| = k$, and from them form k vertices by pairing the elements of A , sorted in increasing order, with the elements of B , sorted in decreasing order. By construction these k vertices are all out of order and hence form an antichain. Conversely, given any antichain of with exactly $k \geq 1$ vertices the corresponding subsets A and B are determined uniquely. Therefore the block $[1, a] \times [1, b]$ has exactly

$$J(a, b) = \sum_{k=1}^{a \wedge b} \binom{a}{k} \binom{b}{k}$$

antichains. By Vandermonde's identity

$$1 + \sum_{k=1}^{a \wedge b} \binom{a}{k} \binom{b}{k} = \sum_{k=0}^{a \wedge b} \binom{a}{k} \binom{b}{k} = \binom{a + b}{a}.$$

Now for the formula for $|\mathcal{ER}(\pi)|$, recall that by Theorem 1.6 every antichain must contain exactly one vertex along the path. The sum in (2.1) partitions the elements of $\mathcal{ER}(\pi)$ according to which vertex is included. Each such vertex (u_i, v_i) naturally

breaks the poset $[1, a] \times [1, b]$ into four quadrants, and any extreme ray containing (u_i, v_i) in its support must have the rest of its non-zero entries in the northwest and southeast quadrants. More precisely, the support must be in the complementary set of $L((u_i, v_i)) \cup U((u_i, v_i))$. The northwest quadrant is precisely $[1, u_i - 1] \times [v_i + 1, n]$ and the southwest one is $[u_i + 1, m] \times [1, v_i - 1]$. The total number of extreme rays containing (u_i, v_i) can then be broken into four distinct types: those with non-zero entries in both quadrants, those with non-zero entries in only one of the quadrants, and the single extreme ray supported only at (u_i, v_i) . Therefore the total number of extreme rays of $[1, m] \times [1, n]$ that have (u_i, v_i) in their support is

$$(1 + J(u_i - 1, n - v_i))(1 + J(m - u_i, v_i - 1))$$

Combining this with Vandermonde’s identity completes the proof. □

3. Facets

In this section we prove Theorem 1.3 on the facets of the maximal cones $\mathcal{C}(\pi)$, again using reasoning that is purely in terms of the last passage model. For each path $\pi \in \Pi_P$ we start with the definition (1.2) of $\mathcal{C}(\pi)$ and determine which inequalities that define it are redundant and which are necessary. The necessary ones are precisely the facets of the cone.

To accomplish this we let N_π be the set of normal vectors which describe the half-spaces defining $\mathcal{C}(\pi)$, i.e.

$$N_\pi = P \cup \{\pi - \pi' : \pi' \in \Pi_P \setminus \{\pi\}\}.$$

Note that we are considering the elements of P as the basis vectors $\delta_{\mathbf{v}}$, $\mathbf{v} \in P$, in this case. Then for each $\eta \in N_\pi$ we define $\mathcal{C}(\pi; \eta)$ to be the same polyhedral cone as $\mathcal{C}(\pi)$ but after removing the hyperplane with normal η , i.e.

$$\mathcal{C}(\pi; \eta) = \{\omega \in \mathbb{R}^P : \langle \omega, \eta' \rangle \geq 0 \text{ for all } \eta' \in N_\pi \setminus \{\eta\}\}.$$

Geometric considerations show that η is redundant if $\mathcal{C}(\pi; \eta) = \mathcal{C}(\pi)$ and necessary otherwise. Equivalently, $\eta \in N_\pi$ is necessary iff $\mathcal{C}(\pi)$ is a *proper* subset of $\mathcal{C}(\pi; \eta)$ i.e. $\mathcal{C}(\pi) \subsetneq \mathcal{C}(\pi; \eta)$. Our strategy is to go through the normal vectors in N_π and, for each one, try to find a weight vector that is in $\mathcal{C}(\pi; \eta)$ but not in $\mathcal{C}(\pi)$. That this strategy works can be seen by a duality argument, see the remark at the end of this section for more details. It can be used to quickly determine which of the inequalities of the form $\omega(\mathbf{v}) \geq 0$ are necessary and which are redundant.

Proof of Theorem 1.3: Weights off the path must be positive: Suppose $\mathbf{v} \in P \setminus \{\pi\}$. Take any $\omega \in \mathcal{C}(\pi)$ and make the entry at \mathbf{v} a negative value. That is, set $\omega'(\mathbf{u}) = \omega(\mathbf{u})$ for $\mathbf{u} \neq \mathbf{v}$, and $\omega'(\mathbf{v}) = -1$ (we use -1 for concreteness). Under ω' the length of π is unchanged, and the length of any other path does not increase, so π is still maximal under the new vector. This proves that $\omega' \in \mathcal{C}(\pi; \mathbf{v}) \setminus \mathcal{C}(\pi)$, hence $\mathcal{C}(\pi)$ is a proper subset of $\mathcal{C}(\pi; \mathbf{v})$ and $\omega(\mathbf{v}) \geq 0$ is a necessary inequality.

Proof of Theorem 1.3: $\omega(\mathbf{v}) \geq 0$ for \mathbf{v} a corner of π is redundant: If \mathbf{v} is a corner of π then there is another path π' such that $\pi \setminus \{\mathbf{v}\} \subset \pi'$ and $\mathbf{v} \notin \pi'$. Define $S = \pi' \setminus \pi$ (note that $S \neq \emptyset$) and $\omega \in \mathcal{C}(\pi; \mathbf{v})$. Since π is still the longest path for ω (but possibly ω is negative at \mathbf{v}), and therefore at least as long as π' , we get

$$\omega(\mathbf{v}) \geq \sum_{\mathbf{v}' \in S} \omega(\mathbf{v}').$$

But for $\omega \in \mathcal{C}(\pi; \mathbf{v})$ we still have that $\omega(\mathbf{v}') \geq 0$ for all $\mathbf{v}' \in S$, hence $\omega(\mathbf{v}) \geq 0$ also by the proof above (that weights off the path must be positive). Therefore inequality $\omega(\mathbf{v}) \geq 0$ is implied by the other inequalities, hence $\mathcal{C}(\pi; \mathbf{v}) = \mathcal{C}(\pi)$ and $\omega(\mathbf{v}) \geq 0$ is redundant.

Proof of Theorem 1.3: $\omega(\mathbf{v}) \geq 0$ for \mathbf{v} on the path but not a corner is necessary: Let $L > 0$, choose $0 < \epsilon < L$, and consider a weight vector ω defined by

$$\omega(\mathbf{u}) = \begin{cases} -\epsilon, & \mathbf{u} = \mathbf{v} \\ L, & \mathbf{u} \in \pi \setminus \{\mathbf{v}\} \\ 0, & \mathbf{u} \notin \pi \end{cases}$$

Then ω is negative at \mathbf{v} , so $\omega \notin \mathcal{C}(\pi)$. Now suppose π' is longer than π under ω , where π' is some other maximal chain π' . Since the only positive weights are in $\pi \setminus \{\mathbf{v}\}$, the only way this is possible is if $\pi \setminus \{\mathbf{v}\} \subset \pi'$. But this would imply that \mathbf{v} is a corner of π , and this is a contradiction. Since ω is non-negative on $P \setminus \{\mathbf{v}\}$, we see that $\omega \in \mathcal{C}(\pi; \mathbf{v})$, making this inequality necessary. \square

The necessity and redundancy of the normal vectors of the form $\pi - \pi'$ proof requires a better understanding of the properties of the disorder graph, which leads to the following proposition. We begin by reviewing the notion of saturated chains within a poset.

Definition 3.1. A saturated chain of a poset P is a totally ordered subset $A \subset P$ such that if $\mathbf{v}, \mathbf{w} \in A$ and $\mathbf{v} < \mathbf{w}$ then there is no $\mathbf{u} \in P \setminus A$ such that $\mathbf{v} < \mathbf{u} < \mathbf{w}$. On finite posets (which are all that we consider) this means that in every pair of successive elements the larger one covers the smaller one.

We will also use the following notation for the neighbor set of a vertex in $\Delta(\pi, \pi')$.

Definition 3.2. For $\mathbf{v} \in \Delta(\pi, \pi')$ let $\mathcal{N}_{\mathbf{v}}$ denote its set of neighbors in the disorder graph.

Part (i) of the next proposition shows that if $\mathbf{v} \in \pi \setminus \pi'$ then $\mathcal{N}_{\mathbf{v}}$ is a subset of π' , and vice versa.

Proposition 3.3. Suppose π and π' are two different paths (maximal chains) on P . The disorder graph $\Delta(\pi, \pi')$ has the following properties.

- (i) $\Delta(\pi, \pi')$ is a bipartite graph, where the two parts are $\pi \setminus \pi'$ and $\pi' \setminus \pi$.
- (ii) If $\Delta(\pi, \pi')$ is connected then $\pi \setminus \pi'$ and $\pi' \setminus \pi$ are saturated chains in P .
- (iii) If $\mathbf{v} \in \Delta(\pi, \pi')$ then $\mathcal{N}_{\mathbf{v}}$ is a non-empty saturated chain in P . Hence if $\mathbf{v} \in \pi \setminus \pi'$ (resp. $\pi' \setminus \pi$) then $\mathcal{N}_{\mathbf{v}}$ is a non-empty interval in π' (resp. π).
- (iv) If $\mathbf{v} < \mathbf{u}$ for $\mathbf{v}, \mathbf{u} \in \pi \setminus \pi'$ then

$$\min \mathcal{N}_{\mathbf{v}} \leq \min \mathcal{N}_{\mathbf{u}} \quad \text{and} \quad \max \mathcal{N}_{\mathbf{v}} \leq \max \mathcal{N}_{\mathbf{u}}.$$

The analogous statement holds for $\mathbf{v}, \mathbf{u} \in \pi' \setminus \pi$.

- (v) The minimal element of $\pi \setminus \pi'$ must be out of order with the minimal element of $\pi' \setminus \pi$, and hence are connected in the disorder graph. The same holds for the maximal elements.
- (vi) Suppose $\Delta(\pi, \pi')$ is not connected. Then there are $\mathbf{v} \in \pi, \mathbf{v}' \in \pi'$ (possibly $\mathbf{v} = \mathbf{v}'$) such that

$$\begin{aligned} & \{\mathbf{u} \in \pi \setminus \pi' : \mathbf{u} < \mathbf{v}\} \cup \{\mathbf{u}' \in \pi' \setminus \pi : \mathbf{u}' < \mathbf{v}'\} \\ & \text{and } \{\mathbf{u} \in \pi \setminus \pi' : \mathbf{u} \geq \mathbf{v}\} \cup \{\mathbf{u}' \in \pi' \setminus \pi : \mathbf{u}' \geq \mathbf{v}'\} \end{aligned}$$

are not connected in $\Delta(\pi, \pi')$, and these four sets are non-empty.

Proof:

- (i) Clearly, all vertices in π are ordered, so there is no edge between these vertices, and the same holds for vertices in π' .
- (ii) Let $\mathbf{u}, \mathbf{w} \in \pi \setminus \pi'$. Suppose $\mathbf{v} \in \pi$ and $\mathbf{u} < \mathbf{v} < \mathbf{w}$. We want to show that $\Delta(\pi, \pi')$ being connected implies that $\mathbf{v} \in \pi \setminus \pi'$. Suppose not. Then it must be that $\mathbf{v} \in \pi \cap \pi'$, from which it follows that every element in $U(\mathbf{v}) \cap \Delta(\pi, \pi')$ is in order with every element in $L(\mathbf{v}) \cap \Delta(\pi, \pi')$ (they are in order through \mathbf{v}) hence these two sets are disconnected subsets of $\Delta(\pi, \pi')$. This is a contradiction unless either $U(\mathbf{v}) \cap \Delta(\pi, \pi') = \emptyset$ or $L(\mathbf{v}) \cap \Delta(\pi, \pi') = \emptyset$, which is impossible because $\mathbf{u} \in L(\mathbf{v}) \cap \Delta(\pi, \pi')$ and $\mathbf{w} \in U(\mathbf{v}) \cap \Delta(\pi, \pi')$.
- (iii) We will prove a slightly more general statement: for any $\mathbf{v} \in P \setminus \pi$, the set $\mathcal{O}_{\mathbf{v}}$ of elements in π that are out of order with \mathbf{v} is a non-empty saturated chain. We do this by contradiction, so first suppose $\mathcal{O}_{\mathbf{v}}$ is empty. This means that \mathbf{v} is in order with every element of π , contradicting the fact that π is a maximal chain. Hence $\mathcal{O}_{\mathbf{v}}$ is non-empty, and by definition it is a subset of the maximal chain π , thus it must have a unique minimal element denoted by \mathbf{u}_0 and a unique maximal element denoted by \mathbf{u}_1 . It is possible that $\mathbf{u}_0 = \mathbf{u}_1$. Now suppose $\mathbf{u} \geq \mathbf{u}_0$. Then it cannot happen that $\mathbf{u} < \mathbf{v}$, since this would imply that $\mathbf{u}_0 < \mathbf{v}$, and we know that \mathbf{u}_0 and \mathbf{v} are out of order. A completely analogous argument shows that if $\mathbf{u} \leq \mathbf{u}_1$, then it cannot happen that $\mathbf{u} > \mathbf{v}$. Therefore if $\mathbf{u}_0 \leq \mathbf{u} \leq \mathbf{u}_1$ then \mathbf{u} must be out of order with \mathbf{v} , so $\mathcal{O}_{\mathbf{v}}$ is indeed a saturated chain.
- (iv) We proceed by contradiction. If $\min \mathcal{N}_{\mathbf{u}} < \min \mathcal{N}_{\mathbf{v}}$ then there must be a $\mathbf{u}' \in \mathcal{N}_{\mathbf{u}}$ such that $\mathbf{u}' < \min \mathcal{N}_{\mathbf{v}}$. The latter means that $\mathbf{u}' \notin \mathcal{N}_{\mathbf{v}}$, which implies that \mathbf{u}' must be in order with \mathbf{v} , i.e. either $\mathbf{u}' < \mathbf{v}$ or $\mathbf{v} < \mathbf{u}'$. The latter is impossible because it would imply $\mathbf{v} < \min \mathcal{N}_{\mathbf{v}}$ and we know that these two vertices are out of order. Thus $\mathbf{u}' < \mathbf{v}$. But also $\mathbf{v} < \mathbf{u}$ by assumption, so therefore $\mathbf{u}' < \mathbf{u}$, which contradicts the fact that $\mathbf{u}' \in \mathcal{N}_{\mathbf{u}}$. Thus $\min \mathcal{N}_{\mathbf{u}} < \min \mathcal{N}_{\mathbf{v}}$ is impossible, but because $\mathcal{N}_{\mathbf{u}}$ and $\mathcal{N}_{\mathbf{v}}$ are both subsets of the ordered chain π' the only remaining option is that $\min \mathcal{N}_{\mathbf{v}} \leq \min \mathcal{N}_{\mathbf{u}}$. The statements for the max follow from completely similar arguments.
- (v) Suppose \mathbf{v} is the minimal element of $\pi \setminus \pi'$ and \mathbf{v}' is the minimal element of $\pi' \setminus \pi$. If $\mathbf{v} < \mathbf{v}'$ was true, then \mathbf{v} could be “inserted” into the path π' to form a longer chain, i.e.

$$(\pi \cap \pi' \cap L(\mathbf{v})) \cup (\pi' \cap U(\mathbf{v}))$$

would be a chain in P that contains \mathbf{v} and π' (recall that $L(\mathbf{v})$ and $U(\mathbf{v})$ are the lower and upper sets of \mathbf{v} , see Definition 2.6). But since $\mathbf{v} \notin \pi'$ this contradicts that π' is a maximal chain. For the analogous reason we cannot have $\mathbf{v}' < \mathbf{v}$. Therefore \mathbf{v} and \mathbf{v}' must be out of order.

- (vi) By part (v) the minimal element of $\pi \setminus \pi'$ is connected to the minimal element of $\pi' \setminus \pi$ in the disorder graph. Let C be the connected component of $\Delta(\pi, \pi')$ containing them both. Then $C \neq \Delta(\pi, \pi')$ by assumption. Thus we may suppose that there exists a minimal $\mathbf{v} \in \pi \setminus \pi'$ that is not in C (the case of a minimal element of $\pi' \setminus \pi$ not in C is handled similarly). For this \mathbf{v} it automatically follows that $\{\mathbf{u} \in \pi \setminus \pi' : \mathbf{u} < \mathbf{v}\}$ is non-empty (it includes the minimal element of $\pi \setminus \pi'$) as is $\{\mathbf{u} \in \pi \setminus \pi' : \mathbf{u} \geq \mathbf{v}\}$ (it contains \mathbf{v}). By

part (iii), there is a minimal $\mathbf{v}' \in \pi' \setminus \pi$ such that \mathbf{v} and \mathbf{v}' are out of order (i.e. \mathbf{v}' is the minimal neighbor of \mathbf{v} in the disorder graph). Then we must have $\mathbf{v}' \notin C$, because if $\mathbf{v}' \in C$ then \mathbf{v} being connected to \mathbf{v}' in the disorder graph would imply $\mathbf{v} \in C$ also, which contradicts the definition of \mathbf{v} . Thus $\{\mathbf{u} \in \pi' \setminus \pi : \mathbf{u} < \mathbf{v}'\}$ is non-empty (it contains the minimal element of $\pi' \setminus \pi$ which is in C and therefore different from \mathbf{v}'), as is $\{\mathbf{u} \in \pi' \setminus \pi : \mathbf{u} \geq \mathbf{v}'\}$. This shows that each of the four sets is non-empty.

Finally, we show that the two sets are not connected in the disorder graph. Let $\mathbf{u}^* \in \pi \setminus \pi'$ with $\mathbf{u}^* < \mathbf{v}$. Then by definition of \mathbf{v} this means $\mathbf{u}^* \in C$. Therefore \mathbf{u}^* and \mathbf{v}' cannot be connected in the disorder graph (else it would imply $\mathbf{v} \in C$) so therefore $\mathbf{u}^* < \mathbf{v}'$. Thus \mathbf{u}^* cannot be connected in the disorder graph to $\{\mathbf{u}' \in \pi \setminus \pi' : \mathbf{u}' \geq \mathbf{v}'\}$. Further \mathbf{u}^* already cannot be connected to $\{\mathbf{u} \in \pi \setminus \pi' : \mathbf{u} \geq \mathbf{v}\}$. Thus $\{\mathbf{u} \in \pi \setminus \pi' : \mathbf{u} < \mathbf{v}\}$ is not connected to $\{\mathbf{u} \in \pi \setminus \pi' : \mathbf{u} \geq \mathbf{v}\} \cup \{\mathbf{u}' \in \pi' \setminus \pi : \mathbf{u}' \geq \mathbf{v}'\}$, and by an analogous argument the latter set is also not connected to $\{\mathbf{u} \in \pi' \setminus \pi : \mathbf{u} < \mathbf{v}'\}$. This completes the proof. \square

Now we return to the redundancy and necessity of inequalities of the form $\langle \pi - \pi', \omega \rangle \geq 0$.

Proof of Theorem 1.3: If $\Delta(\pi, \pi')$ is disconnected then $\pi - \pi'$ is redundant.

By part (vi) of Proposition 3.3, there exists $\mathbf{v} \in \pi$ and $\mathbf{v}' \in \pi'$ such that \mathbf{v} is above every vertex $\mathbf{u}' < \mathbf{v}'$, and \mathbf{v}' is above every vertex $\mathbf{u} < \mathbf{v}$, and these two sets are non-empty. Let π_1 be the path which follows π' until just before \mathbf{v}' and then switches to \mathbf{v} and follows π afterwards. Then $\langle \pi - \pi_1, \omega \rangle \geq 0$ on $\mathcal{C}(\pi; \pi - \pi')$. Similarly, let π_2 be the path which follows π until just before \mathbf{v} , then switches to \mathbf{v}' and follows π' afterwards. Then $\langle \pi - \pi_2, \omega \rangle \geq 0$ on $\mathcal{C}(\pi; \pi - \pi')$. But since $\pi - \pi_1 + \pi - \pi_2 = \pi - \pi'$, this implies that $\langle \pi - \pi', \omega \rangle \geq 0$ automatically on $\mathcal{C}(\pi; \pi - \pi')$, which means that $\pi - \pi'$ is redundant (or $\mathcal{C}(\pi; \pi - \pi') = \mathcal{C}(\pi)$). \square

Proof of Theorem 1.3: If $\Delta(\pi, \pi')$ is connected then $\pi - \pi'$ is necessary.

We need to show that the cone $\mathcal{C}(\pi; \pi - \pi')$ is strictly larger than $\mathcal{C}(\pi)$, meaning that there is a weight $\omega \in \mathbb{R}_+^P$ under which π' is the unique longest path and π is the second longest path (although π may be tied for second longest with several other paths). We will give an explicit such weight vector. To do so we write $\pi \setminus \pi'$ as $\mathbf{u}_0 \leq \mathbf{u}_1 \leq \dots \leq \mathbf{u}_\ell$, recalling that $\pi \setminus \pi'$ is a saturated chain in P because the disorder graph is connected (Proposition 3.3, part ii). Recall that we denote the upper set of \mathbf{u} by $U(\mathbf{u})$. The claimed weight vector ω is

- $\omega(\mathbf{v}) = 1$ for $\mathbf{v} \in \pi'$,
- $\omega(\mathbf{u}_0) = |\pi' \setminus \pi| - |(\pi' \setminus \pi) \cap U(\mathbf{u}_0)| - 1$, and
- $\omega(\mathbf{u}_i) = |(\pi' \setminus \pi) \cap U(\mathbf{u}_{i-1})| - |(\pi' \setminus \pi) \cap U(\mathbf{u}_i)|$ for $i = 1, \dots, \ell$,
- $\omega(\mathbf{v}) = 0$ otherwise.

Note that all weights are non-negative, since $U(\mathbf{u}_i) \subset U(\mathbf{u}_{i-1})$ and there exists at least one vertex in $\pi' \setminus \pi$ that is not an element of $U(\mathbf{u}_0)$ (the lowest element of $\pi' \setminus \pi$ is out of order with \mathbf{u}_0 , by part (v) of Proposition 3.3). Under this weight vector the length of π' is $\langle \pi', \omega \rangle = |\pi'|$, since there is weight one on each element

of π' . Furthermore, the length of π under ω is

$$\begin{aligned}
\langle \pi, \omega \rangle &= \langle \pi \cap \pi', \omega \rangle + \langle \pi \setminus \pi', \omega \rangle \\
&= |\pi \cap \pi'| + \omega(\mathbf{u}_0) + \omega(\mathbf{u}_1) + \dots + \omega(\mathbf{u}_\ell) \\
&= |\pi \cap \pi'| + |\pi' \setminus \pi| - 1 - |(\pi' \setminus \pi) \cap U(\mathbf{u}_\ell)| \\
&= |\pi'| - 1 \\
&= \langle \pi', \omega \rangle - 1.
\end{aligned}$$

The fourth equality follows because \mathbf{u}_ℓ , the highest element of $\pi \setminus \pi'$, is out of order with the highest element of $\pi' \setminus \pi$ and hence *no* element of $\pi' \setminus \pi$ can be above \mathbf{u}_ℓ . Thus $|(\pi' \setminus \pi) \cap U(\mathbf{u}_\ell)| = 0$. Now since ω has only integer entries, any paths with distinct weights must have lengths that differ by at least one, hence $\langle \pi, \omega \rangle = \langle \pi', \omega \rangle - 1$ shows that π is necessarily a second longest path. Thus ω satisfies the required conditions so long as there is no other path (distinct from π and π') with the same length as π' under ω . Since the non-zero weights in ω are only distributed along $\pi \cup \pi'$, any such path would have to switch between π and π' , while also containing the common part $\pi \cap \pi'$.

We start by proving that switching from π to π' will not increase the length of a path. To do this we show that wherever we switch along π , the remaining weight available along π' is the same as the remaining weight available along π . Suppose that \mathbf{u}_i for $i \geq 1$ is the first weight along π that we do not use, meaning that \mathbf{u}_{i-1} is the last element along π that we still go through. Then the weight available along π' by making this switch is $|U(\mathbf{u}_{i-1}) \cap (\pi' \setminus \pi)|$. On the other hand, had we stayed on π the remaining weight we would have picked up on $\pi \setminus \pi'$ is

$$\sum_{j=i}^{\ell} |(\pi' \setminus \pi) \cap U(\mathbf{u}_{j-1})| - |(\pi' \setminus \pi) \cap U(\mathbf{u}_j)| = |(\pi' \setminus \pi) \cap U(\mathbf{u}_{i-1})|.$$

Again the last equality uses that $|(\pi' \setminus \pi) \cap U(\mathbf{u}_\ell)| = 0$. This proves that making a switch from π to π' and staying on π' is never profitable.

Now consider a switch from π' to π . Let \mathbf{v}' be the element of $\pi' \setminus \pi$ that we switch from and \mathbf{u}_i be the first element of $\pi \setminus \pi'$ that it is possible to switch to from \mathbf{v}' . Then it must be that $\mathbf{v}' < \mathbf{u}_i$. Let \mathbf{w}' be the unique element in π' that covers \mathbf{v}' (so it is the next point on π' after \mathbf{v}'). It cannot be that $\mathbf{u}_i < \mathbf{w}'$, since otherwise the path π' could be extended by going from \mathbf{v} to \mathbf{u}_i to \mathbf{w}' ; this would contradict that π' is a maximal chain. Thus $\mathbf{w}' \in \pi' \setminus \pi$ also. Now suppose $\mathbf{u}_{i-1} < \mathbf{w}'$. Then $\{\mathbf{u}_0, \dots, \mathbf{u}_{i-1}\}$ is not connected in $\Delta(\pi, \pi')$ to the set $\{\mathbf{u}' \in \pi' : \mathbf{u}' \geq \mathbf{w}'\}$, while $\{\mathbf{u}_i, \dots, \mathbf{u}_\ell\}$ and $\{\mathbf{u}' \in \pi' : \mathbf{u}' < \mathbf{w}'\}$ are not connected either: this would imply that $\Delta(\pi, \pi')$ is not connected. Thus $\mathbf{u}_{i-1} < \mathbf{w}$ is impossible, and therefore $\mathbf{w}' \notin U(\mathbf{u}_{i-1})$. This implies that

$$|(\pi' \setminus \pi) \cap U(\mathbf{w}')| > |(\pi' \setminus \pi) \cap U(\mathbf{u}_{i-1})|.$$

But the left hand side is how much weight would be picked up along $\pi' \setminus \pi$ by staying on π' after \mathbf{v}' , while the right hand side is the weight that would be picked up along $\pi \setminus \pi'$ by switching from \mathbf{v}' to \mathbf{u}_i , since

$$|(\pi' \setminus \pi) \cap U(\mathbf{u}_{i-1})| = \sum_{j=i}^{\ell} |(\pi' \setminus \pi) \cap U(\mathbf{u}_{j-1})| - |(\pi' \setminus \pi) \cap U(\mathbf{u}_j)| = \sum_{j=i}^{\ell} \omega(\mathbf{u}_j).$$

Thus it is more profitable to stay on π' than to switch to π .

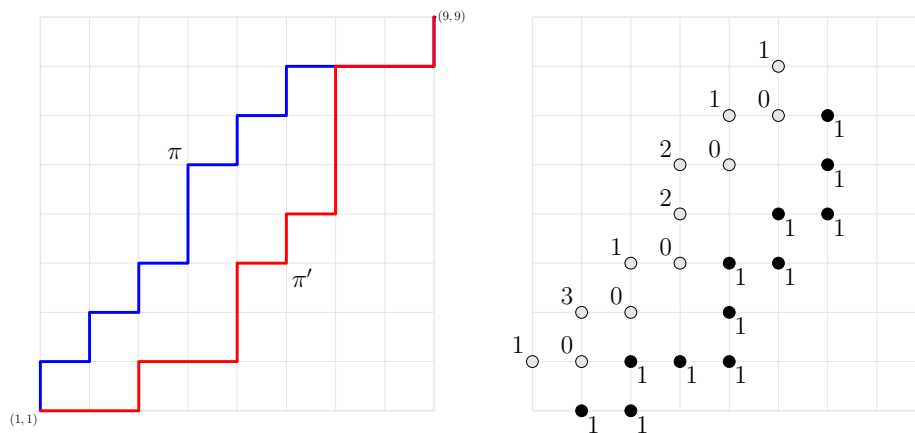


FIGURE 3.3. Two paths π, π' on the subposet $[1, 9]^2$ of \mathbb{Z}^2 , and the vertices of the disorder graph on the right. The vertices of $\pi \setminus \pi'$ are in grey, those of $\pi' \setminus \pi$ are in black. On the right is the weight vector in $\mathcal{C}(\pi; \pi - \pi')$ described in the proof of Theorem 1.3 (all unlabeled vertices have zero weight). Note that π' has length 12 and π has length 11, and there are no other maximal chains in the poset with length greater than 11.

Together, these two facts imply that π' is the unique longest path, and therefore that $\omega \in \mathcal{C}(\pi; \pi - \pi') \setminus \mathcal{C}(\pi)$. \square

Remark 3.4. A more standard way of proving Theorem 1.3 would be to use the dual cone

$$\mathcal{C}(\pi)^* = \{\xi \in \mathbb{R}^P : \langle \xi, \omega \rangle \geq 0 \text{ for all } \omega \in \mathcal{C}(\pi)\},$$

which defines the set of half-spaces containing $\mathcal{C}(\pi)$ (through their normal vectors). The extreme rays of $\mathcal{C}(\pi)^*$ are exactly the normals to the facets of $\mathcal{C}(\pi)$, and arguments similar to those in Section 2 can be used to verify which vectors are extreme rays of the dual cone. We chose the exposition above since it is more in the spirit of the last passage model, but the duality argument is also useful in its own right. For example, it makes clear the assertion that η is necessary iff $\mathcal{C}(\pi) \subsetneq \mathcal{C}(\pi; \eta)$, since on the dual side it corresponds to the statement that a cone becomes smaller when an extreme ray is removed from it. Moreover, the duality argument in this case is made much simpler by the nature of the extreme rays to $\mathcal{C}(\pi)$, as proved in Theorem 1.6. Indeed, using that theorem the dual cone can be rewritten as

$$\mathcal{C}(\pi)^* = \{\xi \in \mathbb{R}^P : \langle \xi, a \rangle \geq 0 \text{ for every antichain } a \text{ with } \text{supp}(a) \cap \text{supp}(\pi) \neq \emptyset\}. \tag{3.1}$$

The only candidates for extreme rays of the dual cone are vectors of the form $\xi = \delta_v$ or $\xi = \pi - \pi'$, and in that case it is clear that $\langle \xi, a \rangle$ is always either 0 or 1, for any antichain a that intersects π exactly once. Using this fact, one can show that a vector ξ of the type either δ_v or $\pi - \pi'$ is an extreme ray of $\mathcal{C}(\pi)^*$ if and only if

$$\bigcap_{\substack{a \in \mathcal{ER}(\pi): \\ \langle \xi, a \rangle = 0}} \text{Span}(a)^\perp = \text{Span}(\xi), \tag{3.2}$$

where $\mathcal{ER}(\pi)$ is the set of extreme rays of $\mathcal{C}(\pi)$. This simple fact can be used to make a relatively quick determination of the extreme rays of the dual cone. Specifically, for the ξ of the form $\pi - \pi'$ the relevant antichains are those which are supported on one vertex of $\pi \setminus \pi'$ and one vertex of $\pi' \setminus \pi$, with the two vertices being out of order to maintain the antichain condition. This forces that any vector in the left hand side of (3.2) must have entries with opposite values at each such pair of vertices, and thus on each connected component of $\Delta(\pi, \pi')$ the entries on each partite set within the component must be the same, and the negative of the common entry on the other partite set.

4. Two-Dimensional Faces of Maximal Cones

Fix a path $\pi \in \Pi_P$. In this section we will describe when two extreme rays $a_1, a_2 \in \mathcal{ER}(\pi)$ form a two-dimensional face of $\mathcal{C}(\pi)$. There are several equivalent definitions of what this means, and we will use one that has a description in terms of the last passage problem. Clearly since a_1 and a_2 are in $\mathcal{C}(\pi)$ so too is their sum $a_1 + a_2$. Then a_1 and a_2 form a two-dimensional face precisely when there are *only* two linearly independent directions from which one can perturb away from $a_1 + a_2$ and remain in the cone $\mathcal{C}(\pi)$. Formally this means

$$a_1, a_2 \text{ form a two-dimensional face of } \mathcal{C}(\pi) \iff \mathcal{D}_\pi(a_1 + a_2) = \text{Span}\{a_1, a_2\}. \quad (4.1)$$

We will assume throughout that a_1 and a_2 are distinct, otherwise they certainly do not form a two-dimensional face. Since they are extreme rays this means there must be at least one vertex where a_1 or a_2 takes on the value one and the other is zero. It is possible that there are vertices at which both a_1 and a_2 take the value one, but excluding the case $a_1 = a_2$ means that $\text{supp}(a_1)$ and $\text{supp}(a_2)$ are necessarily distinct. With this in mind we state the following result about edges. It relies on an object which we define next, called the **order graph**.

Definition 4.1. For two extreme rays a_1 and a_2 of a maximal cone $\mathcal{C}(\pi)$, their *order graph* is the bipartite graph $\mathcal{G}(a_1, a_2)$ with one part being the vertices in $\text{supp}(a_1) \setminus \text{supp}(a_2)$, the other part being the vertices in $\text{supp}(a_2) \setminus \text{supp}(a_1)$, and an edge placed between two vertices iff they are in order.

With this definition in hand the result is:

Theorem 4.2. *Two extreme vectors a_1, a_2 of $\mathcal{C}(\pi)$ form a two-dimensional face in the cone $\mathcal{C}(\pi)$ iff their order graph $\mathcal{G}(a_1, a_2)$ is connected.*

Observe that the order graph and the theorem don't make any reference to the vertices in $\text{supp}(a_1) \cap \text{supp}(a_2)$. Those vertices would only make trivial changes to the order graph. Since they appear in both extreme rays the natural choice would be to include the vertex in both partite sets with an edge between them. However, such vertices cannot be in order with any of the other vertices from either a_1 or a_2 since they are extreme rays. Therefore they would only appear in the graph as isolated components, and we will see in the proof that their presence would only make trivial changes to the statement of the theorem.

Also note that because $a_1 \neq a_2$ the order graph is always non-empty, but it is possible that one of the parts $\text{supp}(a_1) \setminus \text{supp}(a_2)$ or $\text{supp}(a_2) \setminus \text{supp}(a_1)$ is empty. In that case the order graph is connected iff the non-empty part consists of exactly

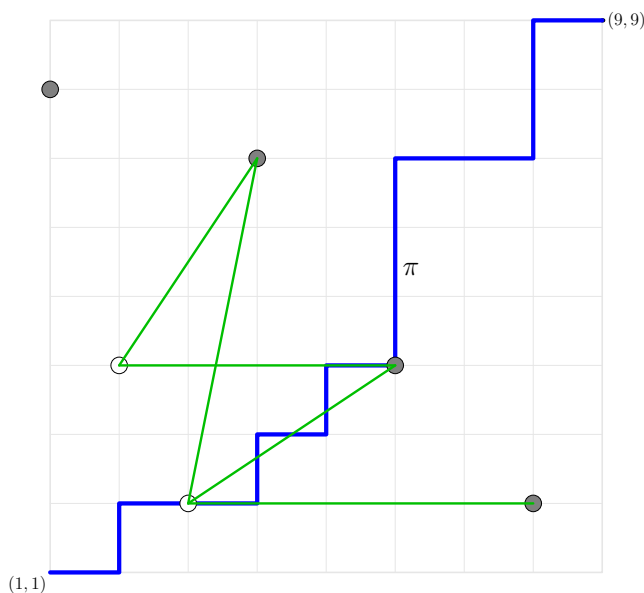


FIGURE 4.4. A path π in the subposet $[1, 9]^2$ of \mathbb{Z}^2 , along with two extreme rays of the cone $\mathcal{C}(\pi)$ (in white and grey). Every vertex of the antichains/extreme rays is in the order graph, but the order graph is not connected because of the isolated grey vertex in the top left. Hence these two extreme rays do not form a two-dimensional face of $\mathcal{C}(\pi)$.

one element. Thus it is always possible to build two-dimensional faces of $\mathcal{C}(\pi)$ by picking an arbitrary $a_1 \in \mathcal{ER}(\pi)$ and then forming a_2 by adding one element of P that is out of order with $\text{supp}(a_1)$, if such an element exists. Adding more than one out of order element will *not* form a two-dimensional face. As the theorem shows, however, not all two-dimensional faces come about from this type of construction.

In the rest of this section we prove Theorem 4.2. We begin with simplifications of (4.1).

Lemma 4.3. *For $\omega \in \mathcal{C}(\pi)$ the perturbation space $\mathcal{D}_\pi(\omega)$ of Definition 2.1 is also equal to*

$$\mathcal{D}_\pi(\omega) = \text{Span} \{a \in \mathcal{ER}(\pi) : \exists \epsilon > 0 \text{ s.t. } \omega \pm \epsilon a \in \mathcal{C}(\pi)\}.$$

Proof: Define the set

$$E = \{a \in \mathcal{ER}(\pi) : \exists \epsilon > 0 \text{ s.t. } \omega \pm \epsilon a \in \mathcal{C}(\pi)\}.$$

Comparing with the definition of $\mathcal{D}_\pi(\omega)$ we see that $E = \mathcal{ER}(\pi) \cap \mathcal{D}_\pi(\omega) \subset \mathcal{D}_\pi(\omega)$, and hence $\text{Span}(E) \subset \mathcal{D}_\pi(\omega)$. For the reverse inclusion, suppose $\sigma \in \mathcal{D}_\pi(\omega)$ so that there exists $\epsilon > 0$ such that $\omega \pm \epsilon \sigma \in \mathcal{C}(\pi)$. Since $\omega \pm \epsilon \sigma$ are in the cone $\mathcal{C}(\pi)$ they are positive linear combinations of at least some elements of $\mathcal{ER}(\pi)$. That is, there exist finite sets $F, G \subset \mathcal{ER}(\pi)$ such that $\omega + \epsilon \sigma$ and $\omega - \epsilon \sigma$ are strictly positive linear combinations of all extreme rays in F and G respectively. But then ω , as the mean of these two vectors, will be a strictly positive combination of all elements of $F \cup G$. Since it is a strictly positive combination one can perturb away from

ω in the direction of any extreme ray in $F \cup G$ and remain in $\mathcal{C}(\pi)$, so therefore $F \cup G \subset E$. Using the latter fact, and writing $\epsilon\sigma = \omega + \epsilon\sigma - \omega$, we see that σ is a linear combination of extreme rays in E , hence $\sigma \in \text{Span}(E)$. \square

Remark 4.4. Suppose $a_1, a_2 \in \mathcal{ER}(\pi)$. Clearly a_1 and a_2 are always in $\mathcal{C}(\pi)$. Therefore, using Lemma 4.3, a_1 and a_2 forming a face is equivalent to the fact that for all $a \in \mathcal{ER}(\pi)$

$$a_1 + a_2 \pm \epsilon a \in \mathcal{C}(\pi) \iff a = a_1 \text{ or } a_2,$$

or, more succinctly, that $\mathcal{ER}(\pi) \cap \mathcal{D}_\pi(a_1 + a_2) = \{a_1, a_2\}$.

Using the last remark and the properties of extreme rays it is possible to greatly reduce the number of extreme rays $a \in \mathcal{ER}(\pi)$ that are possibly in $\mathcal{D}_\pi(a_1 + a_2)$.

Lemma 4.5. *Let a be an extreme ray of $\mathcal{C}(\pi)$. Then $a \in \mathcal{D}_\pi(a_1 + a_2)$ iff a is zero at all vertices where both a_1 and a_2 are (i.e. $\text{supp}(a) \subset \text{supp}(a_1) \cup \text{supp}(a_2) = \text{supp}(a_1 + a_2)$) and any path that is maximal under $a_1 + a_2$, is also maximal under a .*

Proof: First suppose $a \in \mathcal{D}_\pi(a_1 + a_2)$. Then the vector $a_1 + a_2 - \epsilon a$ must be non-negative at every vertex. Thus if a_1 and a_2 are both zero at some vertex then so must be a .

For the second part, suppose that π' is a maximal path for $a_1 + a_2$. Note that π is a maximal path for all the vectors $a_1 + a_2, a$ and $a_1 + a_2 \pm \epsilon a$. So using that $\langle a_1 + a_2, \pi - \pi' \rangle = 0$ and $\langle a_1 + a_2 \pm \epsilon a, \pi - \pi' \rangle \geq 0$, we see that $\langle a, \pi' \rangle = \langle a, \pi \rangle$. This shows that π' is also a maximal path for a .

For the reversed statement, suppose that $\text{supp}(a) \subset \text{supp}(a_1) \cup \text{supp}(a_2)$ and that all paths maximal under $a_1 + a_2$ are also maximal under a . For ϵ small enough, $a_1 + a_2 - \epsilon a$ is still non-negative at every vertex. Now consider a path π' . If $\langle a_1 + a_2, \pi \rangle > \langle a_1 + a_2, \pi' \rangle$, then again for ϵ small enough and for all such π' we would have

$$\langle a_1 + a_2 \pm \epsilon a, \pi - \pi' \rangle > 0.$$

Now suppose that $\langle a_1 + a_2, \pi' \rangle = \langle a_1 + a_2, \pi \rangle$, so π' is a maximal path for $a_1 + a_2$. Then it is also maximal under a , so we get

$$\langle a_1 + a_2 \pm \epsilon a, \pi - \pi' \rangle = 0.$$

This shows that indeed $a \in \mathcal{D}_\pi(a_1 + a_2)$. \square

In geometric terms the previous lemma can be recast as saying that $a \in \mathcal{ER}(\pi) \cap \mathcal{D}_\pi(a_1 + a_2)$ iff $\text{supp}(a) \subset \text{supp}(a_1 + a_2)$ and

$$a \in \bigcap_{\substack{\pi' \in \Pi_P: \\ \langle a_1 + a_2, \pi - \pi' \rangle = 0}} \text{Span}(\pi - \pi')^\perp.$$

However, it is simpler to recast it in terms of the order graph.

Lemma 4.6. *Let a be an extreme ray of $\mathcal{C}(\pi)$. Then $a \in \mathcal{D}_\pi(a_1 + a_2)$ iff the following three conditions simultaneously hold:*

- (1) $\text{supp}(a) \subset \text{supp}(a_1) \cup \text{supp}(a_2)$.
- (2) For every edge $(\mathbf{v}, \mathbf{w}) \in \mathcal{G}(a_1, a_2)$ either $a(\mathbf{v}) = 1$ or $a(\mathbf{w}) = 1$ (but not both).
- (3) If $\mathbf{v} \in \text{supp}(a_1) \cap \text{supp}(a_2)$ then $a(\mathbf{v}) = 1$.

Proof: Suppose $a \in \mathcal{D}_\pi(a_1 + a_2)$. The first condition follows immediately from Lemma 4.5. Now suppose $(\mathbf{v}, \mathbf{w}) \in \mathcal{G}(a_1, a_2)$. This means that \mathbf{v} and \mathbf{w} are in order, and hence there exists a path π' passing through both \mathbf{v} and \mathbf{w} . Therefore π' is maximal under $a_1 + a_2$ (a path can pick up at most one vertex of an extreme ray, since all its vertices are out of order). But then by Lemma 4.5 π' is also maximal under any $a \in \mathcal{ER}(\pi) \cap \mathcal{D}_\pi(a_1 + a_2)$. Since π' can't pass through any other vertices from $\text{supp}(a_1) \cup \text{supp}(a_2)$ other than \mathbf{v} or \mathbf{w} , a must be non-zero at either \mathbf{v} or \mathbf{w} in order for π' to be maximal under a (here we use Condition 1). However a can't be non-zero at both, otherwise a would be not be extreme.

For $\mathbf{v} \in \text{supp}(a_1) \cap \text{supp}(a_2)$ any path that passes through \mathbf{v} is longest under $a_1 + a_2$, and hence must also be longest under a . But this implies that $a(\mathbf{v}) = 1$.

Now we wish to prove the reverse implication. Suppose π' is maximal under $a_1 + a_2$. Since π is maximal under both a_1 and a_2 , either there exists an edge in $\mathcal{G}(a_1, a_2)$ or $\text{supp}(a_1) \cap \text{supp}(a_2) \neq \emptyset$. In both cases we would have that $\langle \pi', a_1 + a_2 \rangle = 2$ (otherwise π' would not be maximal). If π' picks up two vertices $\text{supp}(a_1) \cup \text{supp}(a_2)$, then Condition 2 implies that $\langle \pi', a \rangle = 1$, so π' is maximal under a . If π' picks up a vertex \mathbf{v} in $\text{supp}(a_1) \cap \text{supp}(a_2)$, then Condition 3 implies that $a(\mathbf{v}) = 1$, so also in that case π' is maximal under a . Now Lemma 4.5 implies that $a \in \mathcal{D}_\pi(a_1 + a_2)$. \square

This allows us to finish the proof of the theorem.

Proof of Theorem 4.2: First assume that a_1 and a_2 form a two-dimensional face. Then $a \in \mathcal{D}_\pi(a_1 + a_2) \cap \mathcal{ER}(\pi)$ implies that $a = a_1$ or $a = a_2$. Suppose $\mathcal{C} \subset \mathcal{G}(a_1, a_2)$ is a connected component, strictly smaller than $\mathcal{G}(a_1, a_2)$. Define $a \in \mathbb{R}^P$ in the following way: $a(\mathbf{v}) = 1$ for all $\mathbf{v} \in \text{supp}(a_1) \cap \text{supp}(a_2)$, all $\mathbf{v} \in \mathcal{C} \cap \text{supp}(a_1)$ and all $\mathbf{v} \in (\mathcal{G}(a_1, a_2) \setminus \mathcal{C}) \cap \text{supp}(a_2)$. For all other $\mathbf{v} \in P$ $a(\mathbf{v}) = 0$. Note that $a \in \mathcal{ER}(\pi)$: two vertices in $\text{supp}(a)$ cannot be in order, since all elements of \mathcal{C} are out of order with all elements of $\mathcal{G}(a_1, a_2) \setminus \mathcal{C}$. Furthermore, define \mathbf{v}_i (for $i = 1, 2$) as the unique element of $\text{supp}(\pi) \cap \text{supp}(a_i)$. If $\mathbf{v}_1 = \mathbf{v}_2$, then $\mathbf{v}_1 \in \text{supp}(a_1) \cap \text{supp}(a_2)$, so $a(\mathbf{v}_1) = 1$. If $\mathbf{v}_1 \neq \mathbf{v}_2$, then they are in order, so they are either both in \mathcal{C} (so $a(\mathbf{v}_1) = 1$), or both in $\mathcal{G}(a_1, a_2) \setminus \mathcal{C}$ (so $a(\mathbf{v}_2) = 1$). In all of these cases we have that $\langle \pi, a \rangle = 1$. It is also clear that a satisfies the three conditions of Lemma 4.6, so we conclude that $a \in \mathcal{D}_\pi(a_1 + a_2) \cap \mathcal{ER}(\pi)$. However, a cannot be equal to a_1 or a_2 : there exists at least one vertex $\mathbf{v} \in \mathcal{C}$. If $\mathbf{v} \in \text{supp}(a_1) \setminus \text{supp}(a_2)$, then $a(\mathbf{v}) = 1$. If $\mathbf{v} \in \text{supp}(a_2) \setminus \text{supp}(a_1)$, then $a(\mathbf{v}) = 0$. In both cases we have $a \neq a_2$. Since there exists also at least one vertex in $\mathcal{G}(a_1, a_2) \setminus \mathcal{C}$, we also see that $a \neq a_1$. We now contradict the assumption we started with, so $\mathcal{G}(a_1, a_2)$ must be connected.

For the reverse statement, we use a similar approach. Suppose $a \in \mathcal{D}_\pi(a_1 + a_2) \cap \mathcal{ER}(\pi)$. Then a must satisfy the three conditions of Lemma 4.6. There also must be at least one vertex $\mathbf{v} \in \mathcal{G}(a_1, a_2)$. Suppose $\mathbf{v} \in \text{supp}(a_1) \setminus \text{supp}(a_2)$ (the other case follows the same arguments). Suppose $a(\mathbf{v}) = 1$. Since $\mathcal{G}(a_1, a_2)$ is a connected bipartite graph and a satisfies Condition 2 of Lemma 4.6, it follows that a must be 0 on $\text{supp}(a_2) \setminus \text{supp}(a_1)$ and 1 on $\text{supp}(a_1) \setminus \text{supp}(a_2)$, and therefore (together with Condition 3 of Lemma 4.6) $a = a_1$. If, on the other hand, $a(\mathbf{v}) = 0$, then by the same reasoning we would have that $a = a_2$. This proves that a_1 and a_2 form a two-dimensional face. \square

Remark 4.7. We note that the definition of the order graph seems to somehow be “dual” to the definition of the disorder graph, especially one if considers paths and antichains as being in duality. However the order graph describes dimension 2 faces

of the maximal cones, while the disorder graph describes the codimension 1 facets. We are unsure of the exact reason for this discrepancy.

5. A Simplicial Decomposition of Maximal Sets

The analysis of polyhedral cones in a high-dimensional space is much simpler when the cones are *simplicial*, which we recall means the following:

Definition 5.1. A polyhedral cone is simplicial iff its extreme rays form a basis for the ambient space.

The major consequence of this fact is that every point in the cone can be uniquely written as a positive, linear combination of the extreme rays, whereas the uniqueness fails for a non-simplicial cone. Put another way, a simplicial cone is the image of a one-to-one linear transformation of the positive orthant of the ambient Euclidean space (the map that sends the standard basis vectors to the extreme rays of the cone), and as a result integrals over the simplicial cone can be transformed into integrals over the positive orthant. The results of Section 2 show that the number of extreme rays of a maximal cone $\mathcal{C}(\pi)$ is typically much larger than the dimension $|P|$ of the Euclidean space that the problem is embedded into, and hence the cones are far from simplicial. However, computations can be made tractable by partitioning the polyhedral cone into a disjoint union of simplices (disjoint up to measure zero boundary intersections), and general theory ensures that such a partition always exists. In fact, it is always possible to find a decomposition such that the extreme rays of every simplicial cone are also extreme rays of the original cone. In this section we describe a general scheme for finding such a decomposition for the last passage model. Although we found this scheme independently, it already appears in Stanley (1986).

The key is to consider the set of upper sets of the poset, sometimes also called the *order ideals*. Recall that $U \subset P$ is an upper set if $\mathbf{v} \in U$ and $\mathbf{v} \leq \mathbf{u}$ implies $\mathbf{u} \in U$. Further recall that ∂U is the set of minimal elements of U , which clearly forms an antichain, and this establishes a bijection between upper sets and antichains. In the context of the last passage model this bijection is very natural: if an antichain is the weight vector then its corresponding upper set is the vector of passage times, which encodes the maximal length up to each given vertex.

Definition 5.2. For a poset P the set of all upper sets of P , ordered by inclusion, is itself a poset called $J(P)$.

It is well known and straightforward to see that $J(P)$ is in fact a *distributive lattice*, meaning that any two elements of $J(P)$ have a unique least upper bound and greatest lower bound that are, in this case, given by the union and intersection of the elements, respectively. The unique minimal element of $J(P)$ is the empty set and the unique maximal element is P itself. Furthermore, $J(P)$ is graded of rank $|P|$ and in this case the rank function $\rho(U)$ of $U \in J(P)$ is simply the cardinality of U . In particular the maximal chains of $J(P)$ all contain $|P| + 1$ elements and can all be written in the form $U_0 \triangleleft U_1 \triangleleft \dots \triangleleft U_{|P|}$ where each lower set U_j contains exactly j vertices in P . For a review of these concepts see Stanley (2012, Chapter 3).

The maximal chains of $J(P)$ can be used to produce a simplicial decomposition of the maximal sets $\mathcal{C}(\pi)$. The construction is best explained through an object

called the *order cone* of P , which is a mild generalization of the order polytope of Stanley (1986).

Definition 5.3. Define the *order cone* $\mathcal{OC}(P)$ of P to be the subset of vectors in \mathbb{R}_+^P that obey the ordering of the poset, i.e.

$$\mathcal{OC}(P) := \{\eta \in \mathbb{R}_+^P : \mathbf{v} \leq \mathbf{w} \implies \eta(\mathbf{v}) \leq \eta(\mathbf{w})\}.$$

It is straightforward to verify that $\mathcal{OC}(P)$ is a polyhedral cone, but also that it is a proper subset of \mathbb{R}_+^P .

Theorem 5.4. Let $U_0 \triangleleft U_1 \triangleleft \dots \triangleleft U_{|P|}$ be a maximal chain in $J(P)$. The conical combinations of the antichains $\partial U_1, \partial U_2, \dots, \partial U_{|P|}$, embedded as vectors in \mathbb{R}_+^P , form a simplicial cone in \mathbb{R}_+^P . Moreover the set of all such simplicial cones forms a partition of \mathbb{R}_+^P (up to measure zero boundaries).

Proof: $J(P)$ forms a simplicial decomposition for $\mathcal{OC}(P)$ in the following way: for any $\eta \in \mathcal{OC}(P)$ and for $k = 1, 2, \dots, |P|$ let \mathbf{v}_k be the vertex at which η achieves its k^{th} smallest value. These \mathbf{v}_k are well-defined so long as η doesn't take on the same value at multiple vertices, which is Lebesgue almost all of $\mathcal{OC}(P)$. Let $U_k^* = \{\mathbf{v}_i : i \geq k\}$ for $k = 1, \dots, |P|$ and $U_0^* = \emptyset$. Then each U_k^* is an upper set of P , due to the η obeying the order relation, and since $U_k^* \subset U_{k+1}^*$ with $U_{k+1}^* \setminus U_k^* = \{\mathbf{v}_k\}$ it follows that $U_0^* \triangleleft U_1^* \triangleleft \dots \triangleleft U_{|P|}^*$ is a maximal chain in $J(P)$. Thus each $\eta \in \mathcal{OC}(P)$ produces a maximal chain in $J(P)$, and furthermore the set of vectors η which produce any particular maximal chain forms a simplicial cone in $\mathcal{OC}(P)$. It is in fact a canonical simplicial cone, since it is determined by an ordering of the η variables, and as such the indicator functions of the upper sets $U_1^*, U_2^*, \dots, U_{|P|}^*$ are its extreme rays. Since each vector in $\mathcal{OC}(P)$ uniquely determines one of these simplicial cones (Lebesgue almost surely), it follows that we have a simplicial decomposition. Moreover this construction is reversible: every maximal chain $\{U_k\}_{k=0, \dots, |P|}$ in $J(P)$ determines a sequence $\mathbf{v}_{k+1} := U_{k+1} \setminus U_k$ and the simplicial cone $\{\eta \in \mathbb{R}_+^P : 0 \leq \eta(\mathbf{v}_1) \leq \eta(\mathbf{v}_2) \leq \dots \leq \eta(\mathbf{v}_{|P|})\}$ is in $\mathcal{OC}(P)$.

To complete the proof we now associate each simplicial cone in $\mathcal{OC}(P)$ to one in \mathbb{R}_+^P . Use the linear mapping that takes $U_i \rightarrow \partial U_i$ for $i = 1, \dots, |P|$, which is invertible as a transformation of \mathbb{R}^P due to the bijection between upper sets and antichains. Hence it takes the simplex with extreme rays U_i to a simplex with extreme rays ∂U_i , which is in \mathbb{R}_+^P since the ∂U_i are. Note that it is a *different* linear map for each maximal chain $\{U_k\}_{k=0, \dots, |P|}$, and hence a different linear map applied to each simplex defined by the maximal chain. Each mapping is simply the map from the passage time vector back to the weight vector. Since each weight vector uniquely determines an ordering of the passage times (Lebesgue almost surely) one sees that disjoint simplices in $\mathcal{OC}(P)$ are mapped to disjoint simplices in \mathbb{R}_+^P (up to measure zero boundaries). Moreover, since there is a bijection between weight vectors and the vector of passage times all of \mathbb{R}_+^P is covered by these mappings from $\mathcal{OC}(P)$ to \mathbb{R}_+^P . This completes the proof. \square

Taking the above as an algorithm for producing simplicial cones, the next step is to associate them to the maximal sets $\mathcal{C}(\pi)$ for $\pi \in \Pi_P$. By Theorem 1.6 it is clear that the simplicial cone of Theorem 5.4 is a subset of $\mathcal{C}(\pi)$ iff the support of each of its extreme rays ∂U_i intersects the support of π . The next result shows that there is exactly one such path π , and that we can generate it directly from the

sequence of extreme rays. The basic algorithm is to start at the element of P with the longest passage time and then moving backwards to the element below it with the next longest passage time, repeating until arrival at a minimal element.

Theorem 5.5. *Let $U_0 \triangleleft U_1 \triangleleft \dots \triangleleft U_{|P|}$ be a maximal chain in $J(P)$ and \mathcal{C} be its associated simplicial cone. For $k = 1, 2, \dots, |P|$ let $\mathbf{v}_k = U_k \setminus U_{k-1}$. Define an integer-valued sequence g_i by $g_1 = 1$ and*

$$g_{j+1} = \min \{k : \mathbf{v}_k \triangleleft \mathbf{v}_{g_j}\},$$

until reaching the first integer m such that \mathbf{v}_{g_m} is a minimal element of P . Then the reversed subsequence $\pi_i = \mathbf{v}_{g_{m-i+1}}$ is a path in Π_P , and moreover $\mathcal{C} \subset \mathcal{C}(\pi)$.

Proof: That π is a maximal chain in Π_P is immediate from the way it is constructed as a sequence of vertices one below the next. The construction is well-defined since each vertex in P appears exactly once in the sequence \mathbf{v}_k . That π is defined as a reversed subsequence is only so that the elements along it are in increasing rather than decreasing order. Finally, to see that $\mathcal{C} \subset \mathcal{C}(\pi)$ simply observe that every extreme ray ∂U_i contains (exactly) one vertex in π . This is again by construction: \mathbf{v}_{g_1} is in ∂U_i for $i < g_2$, since U_i does not contain a vertex below v_{g_1} . For the same reason we have that $\mathbf{v}_{g_j} \in \partial U_i$ for $g_j \leq i < g_{j+1}$. The last element \mathbf{v}_{g_m} is in ∂U_i for $g_m \leq i \leq |P|$. This shows that every antichain ∂U_i is in $\mathcal{C}(\pi)$, by Theorem 1.6, and since the maximal sets $\mathcal{C}(\pi)$ are disjoint for different π (up to measure zero boundaries) this shows that $\mathcal{C} \subset \mathcal{C}(\pi)$. \square

Remark 5.6. There is also a simplicial decomposition of $\mathcal{OC}(P)$, and hence of \mathbb{R}_+^P , by lower sets. This can be seen via the standard bijection between upper and lower sets, which relates the two through the bijection from upper sets to antichains and then antichains to lower sets. The decomposition works in the same way as the above, with each maximal chain in the poset of lower sets (ordered by inclusion) determining a simplex in $\mathcal{OC}(P)$ which is then linearly mapped to a simplex in \mathbb{R}_+^P . In fact, the bijection between upper and lower sets also shows that there is a bijection between maximal chains of each, and it follows that a maximal chain of upper sets produces the same simplicial cone (in \mathbb{R}_+^P , not in $\mathcal{OC}(P)$) as the corresponding maximal chain of lower sets. Since upper sets have a natural interpretation as passage time vectors of antichains we prefer to make the description in terms of upper sets.

Remark 5.7. The maximal chains in P are also in bijection with the linear extensions of P : the set of bijections $\sigma : P \rightarrow \{1, \dots, |P|\}$ such that $\mathbf{v} \leq \mathbf{w} \implies \sigma(\mathbf{v}) \leq \sigma(\mathbf{w})$. The bijection is defined by letting $U_k = \{\mathbf{v} \in P : \sigma(\mathbf{v}) \leq k\}$, with $U_0 = \emptyset$. As such the linear extensions simply correspond to the ordering of the passage time vector $(G_P(\mathbf{v}) : \mathbf{v} \in P)$, with the ordering being well-defined for Lebesgue almost all ω . This implies the connection to the order growth model that we explain below.

Remark 5.8. The path produced by each maximal chain of $J(P)$ is often referred to as the Schützenberger or jeu-de-taquin path (see [Fulton, 1997](#) for review and [Romik and Śniady, 2015](#) for related results on infinite Young tableaux). While each maximal chain in $J(P)$ determines the maximal path for a simplicial cone of weight vectors, each given path $\pi \in \Pi_P$ is typically produced by many maximal chains of $J(P)$. Intuitively one expects that the probability that a given path is the longest one should be larger for those paths produced by more maximal chains, although

this is not entirely precise because it does not take into account the probability of each simplicial cone/maximal chain under a given weight distribution. Enumerating the number of a maximal chains which produce a given path π also appears to be difficult, even on posets of the form $[1, m] \times [1, n]$.

The proof of Theorem 5.4 contains the useful fact that each maximal chain in $J(P)$ induces a linear map from \mathbb{R}^P into itself defined by $U_i \mapsto \partial U_i$ for $i = 1, \dots, |P|$. Coordinatewise the mapping works out to be of the form $\eta \mapsto \omega$ where

$$\omega(\mathbf{v}) = \eta(\mathbf{v}) - \max_{\mathbf{u} \triangleleft \mathbf{v}} \eta(\mathbf{u}),$$

if \mathbf{v} is not a minimal element of P , and $\omega(\mathbf{v}) = \eta(\mathbf{v})$ if it is. It is worth recording the following important but well known observation (see Stanley, 2012, Chapter 3 for example).

Lemma 5.9. *For each maximal chain $U_0 \triangleleft U_1 \triangleleft \dots \triangleleft U_{|P|}$ of $J(P)$ the linear mapping of \mathbb{R}^P to itself defined by $U_i \mapsto \partial U_i$, $i = 1, \dots, |P|$, is volume preserving.*

Proof: This follows because the matrix representing the mapping can be put into an upper triangular form with all ones on the diagonal, in the following way. Define for $k = 1, \dots, |P|$ $\mathbf{v}_k = U_k \setminus U_{k-1}$. We take $\mathbf{v}_1, \dots, \mathbf{v}_{|P|}$ as the standard basis vectors of $\mathbb{R}^{|P|}$. Let A be the matrix defining our map in terms of our basis. Furthermore, consider the sets U_i and ∂U_i as vectors in \mathbb{R}^P . Clearly, $A\mathbf{v}_1 = \mathbf{v}_1$ and for $k > 1$,

$$A\mathbf{v}_k = AU_k - AU_{k-1} = \mathbf{v}_k + \sum_{i < k} 1_{\{\mathbf{v}_i \in \partial U_k\}} - 1_{\{\mathbf{v}_i \in \partial U_{k-1}\}}.$$

This shows that A has 1's on the diagonal and 0's below the diagonal, proving that it is volume preserving. \square

Converting the Passage Time into a Sum. The simplicial decomposition of \mathbb{R}_+^P by maximal chains in $J(P)$ provides a useful way of converting the passage time into a sum, in the following way. For $\omega \in \mathbb{R}_+^P$ let $\pi^*(\omega)$ be the (Lebesgue almost surely unique) longest path in Π_P corresponding to ω , so that

$$G_P = \langle \pi^*(\omega), \omega \rangle.$$

Now for a given maximal chain $U_0 \triangleleft U_1 \triangleleft \dots \triangleleft U_{|P|}$ in $J(P)$ let $E = E_U$ be the $|P| \times |P|$ matrix with columns $\partial U_1, \partial U_2, \dots, \partial U_{|P|}$, as vectors in \mathbb{R}_+^P . Then E is an invertible linear map of \mathbb{R}^P into itself (by Lemma 5.9), hence there exists $\Lambda_U = \Lambda \in \mathbb{R}^P$ such that

$$\omega = E_U \Lambda_U. \tag{5.1}$$

This leads to the expression

$$G_P = \langle \pi^*(\omega), E\Lambda \rangle = \langle E'\pi^*(\omega), \Lambda \rangle,$$

where E' denotes the transpose of E , so that the rows of E' are antichains of P . Since each antichain intersects any given path at most once, this implies that $E'\pi^*(\omega)$ is a vector whose entries are either zero or one:

$$(E'\pi^*(\omega))_i = \mathbf{1} \{ \text{supp}(\pi^*(\omega)) \cap \text{supp}(\partial U_i) \neq \emptyset \}.$$

Combining these together leads to the formula

$$G_P = \sum_{i=1}^{|P|} \Lambda_i \mathbf{1} \{ \text{supp}(\pi^*(\omega)) \cap \text{supp}(\partial U_i) \neq \emptyset \}. \quad (5.2)$$

Note that this formula holds for *any* choice of maximal chain U of $J(P)$, although the value of the Λ changes with different choices of U . If the ω are random but have density $f(\mathbf{x})$ with respect to Lebesgue measure on \mathbb{R}_+^P then the Λ have law $f(E\mathbf{x})$.

The choice of U can be made depending on ω . For each $\omega \in \mathbb{R}_+^P$ let $U(\omega)$ be the maximal chain corresponding to the passage time vector $(G_P(\mathbf{v}; \omega) : \mathbf{v} \in P)$. Then the longest path intersects each antichain $\partial U_i(\omega)$ exactly once, leading to the identity

$$G_P = \sum_{i=1}^{|P|} \Lambda_i. \quad (5.3)$$

In this case all Λ_i are positive but the density is more complicated. Now it becomes the mixture

$$\mathbb{P}((\Lambda_1, \dots, \Lambda_{|P|}) \in d\mathbf{x}) = \sum_{\substack{U \text{ maximal} \\ \text{chains of } J(P)}} f(E_U \mathbf{x}) \mathbb{P}(\omega \in \text{Span}_+ \{ \partial U_1, \dots, \partial U_{|P|} \}).$$

Corner Growth Model. The formulas above express the last passage percolation problem in terms of the corner growth model, which is a well known equivalent description (see [Seppalainen, 2009](#); [Romik, 2015](#) for reviews). In the continuous time version of corner growth the elements of P are "filled in" at random times, subject to the constraint that an element cannot be filled in until all elements in its lower set have also been filled in. At any given time the "corners" are the elements of P which are admissible to be filled in; this nomenclature is motivated by the process on \mathbb{Z}^2 . The process starts at time zero and $\omega(\mathbf{v})$ is the additional time it takes for $\mathbf{v} \in P$ to be filled in *after* all elements of $L(\mathbf{v}) \setminus \{\mathbf{v}\}$ have been filled in. In the case that $L(\mathbf{v}) = \{\mathbf{v}\}$, meaning that \mathbf{v} is a minimal element of P , then there is no waiting rule and $\omega(\mathbf{v})$ is the time at which \mathbf{v} is filled in. Now $G_P(\mathbf{v})$ is exactly the time at which element \mathbf{v} is filled in, and if we take this as a definition of G_P then it implies the recursion (1.1). In fact this recursion shows that the corner growth description is equivalent to the last passage one. The longest path is the maximal chain π of P that takes the longest amount of time to be filled in, together with the additional requirement that for every $\mathbf{v} \in \pi$ this same condition holds on $L(\mathbf{v})$.

The corner growth representation also makes clear the basic idea behind (5.3). The vector $G_P(\mathbf{v}; \omega)_{\mathbf{v} \in P}$ is in the order cone $\mathcal{OC}(P)$ of P , due to the positivity of $\omega \in \mathbb{R}_+^P$. Then $G_P(\mathbf{v}; \omega)$ belongs to a unique (Lebesgue almost surely) simplex in $\mathcal{OC}(P)$ that corresponds to a maximal chain in $J(P)$. The simplex describes the ordering according to which the elements of P are filled in for this particular ω , and then each Λ_i is the time between the filling in of the i th and $(i-1)$ st elements of P . More precisely, letting $U_0 < U_1 < \dots < U_{|P|}$ be the maximal chain of $J(P)$ determined uniquely by ω (Lebesgue almost surely), we let $\mathbf{v}_k = U_k \setminus U_{k-1} \in P$ for $k = 1, 2, \dots, |P|$ (the vertices ordered according to the time at which they appear) and then it follows that

$$\Lambda_i = G_P(\mathbf{v}_i) - G_P(\mathbf{v}_{i-1})$$

with $G_P(\mathbf{v}_0) = 0$. This clearly implies (5.3).

Remark 5.10. When P is a Young diagram (including the poset $[1, m] \times [1, n]$) the maximal chains of $J(P)$ are in bijection with the Young tableaux for the particular diagram. The Young tableaux describes a linear map from a simplex in $\mathcal{OC}(P)$ to a simplex in \mathbb{R}_+^P , with the outputted simplex being precisely the set of weights that produce that particular ordering for the passage times.

6. Independent Exponential Weights

The last passage model with iid exponential weights is solvable, meaning that exact formulas can be derived for various statistics such as the passage time, at least on the poset $[1, m] \times [1, n]$. One basic reason for this is that the memoryless property of the exponential distribution makes the corner growth process a continuous time Markov chain. It is a straightforward calculation to show that, when the ω are iid exponential(1) random variables, at any fixed time the random amount of time until the next corner is filled has an exponential distribution with parameter equal to the number of available corners. This is made clear by the following simple fact about the exponential distribution in several variables.

Lemma 6.1. *Let $X = (X_1, \dots, X_n)$ with $X_i \sim \text{exponential}(\lambda_i)$ independent, and let $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}_+^n$ be linearly independent. Let $\mathcal{C} = \text{Span}_+\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$. Then*

$$\mathcal{P}(X \in \mathcal{C}) = |\det V| \prod_{i=1}^n \frac{\lambda_i}{\langle \mathbf{v}_i, \lambda \rangle}, \tag{6.1}$$

where $\lambda = (\lambda_1, \dots, \lambda_n)$ and V is the $n \times n$ matrix with columns $\mathbf{v}_1, \dots, \mathbf{v}_n$. Moreover

$$\mathcal{L}(X | X \in \mathcal{C}) = \mathcal{L}\left(\sum_{i=1}^n \Lambda_i \mathbf{v}_i\right) = \mathcal{L}(V\Lambda),$$

where $\mathcal{L}(X | X \in \mathcal{C})$ denotes the conditional law of the random variable and $\Lambda_i \sim \text{exponential}(\lambda_i \|\mathbf{v}_i\|_1)$ are independent and $\Lambda = (\Lambda_1, \dots, \Lambda_{|P|})$.

The proof of this lemma comes by mapping \mathbb{R}_+^n into \mathcal{C} using the matrix V , but this type of mapping is one-to-one iff \mathcal{C} is a simplex. Note that in this representation the chosen length of the extreme rays is irrelevant since it always cancels out, as in (6.1). Since we canonically take the extreme rays of the cones in Section 5 to have entries either zero or one, it follows that the ℓ^1 norm of any extreme ray of a maximal cone $\mathcal{C}(\pi)$ is equal to the number of corners available in the corner growth process.

As in the last section the weight vector ω can always be uniquely rewritten as a linear combination of the extreme rays of the simplex that it belongs to (Lebesgue almost surely), and the simplices are in bijection with the maximal chains of $J(P)$. Lemma 6.1 implies that in the exponential last passage model, conditionally on the choice of cone, the coefficients in the linear combination are again independent exponentials. Alternatively, by forgetting about the conditioning one can think of the Λ_i in (5.3) as exponential random variables with random parameters. The parameters are independent of the exponentials and their joint law is determined by the probabilities (6.1) and the ℓ_1 lengths of the extreme rays, which as mentioned are simply counting the number of corners available at a given time. Equivalently, the joint law on parameters is determined by the directed random walk on $J(P)$

started from the minimal element \emptyset and with transitions proportional to the parameters λ_i of the vertices available at each time. This random walk produces a random maximal chain of $J(P)$ and the vector of parameters is simply the number of corners available at each time of the walk. The difficulty in using this approach is that the sheer number of maximal chains of $J(P)$ makes it difficult to average out over the random parameters. Even on the poset $[1, m] \times [1, n]$ the distribution of the random parameters appears to be complicated.

The geometric point-of-view sheds additional some light on the unique properties of the exponential distribution. Equation (6.1) can also be equated to the (suitably normalized) volume of the cone \mathcal{C} intersected with the hyperplane $\{\mathbf{x} : \langle \lambda, \mathbf{x} \rangle = 1\}$, since on that part of the hyperplane the exponential density is constant. The volume description of the probability does not require that \mathcal{C} be a simplex, and in particular implies that the probability of any given path $\pi \in \Pi_P$ being the maximal path is the (normalized) volume of the maximal cone $\mathcal{C}(\pi)$ intersected with the same hyperplane. This intersection is a polyhedron with codimension 1, and while algorithms for computing its volume exist they are in general $\#P$ hard [Brightwell and Winkler \(1991\)](#). On certain posets it may be that the structure of $\mathcal{C}(\pi)$ allows for more efficient computation, but in full generality it appears to be intractable.

Since Lemma 6.1 only works for simplices, to compute path probabilities via this formula would require using a simplicial decomposition of $\mathcal{C}(\pi)$ and summing over all of the simplices. The simplicial decomposition of Section 5 is an obvious choice but the sheer number of simplices involved makes it impractical. On the poset $[1, m] \times [1, n]$ the decomposition involves finding all Young tableaux which produce the π as its Schützenberger path. This appears to be difficult, but we expect that among all paths the extreme corner paths have the largest number of associated Young tableaux.

7. IID Uniform Weights

Independent and identically distributed uniform weights correspond to Lebesgue measure on $[0, 1]^P$, and in this case the *chain polytope* of [Stanley \(1986\)](#) is a useful tool in the analysis.

Definition 7.1. The *chain polytope* $\mathcal{C}(P)$ of a poset P is the subset of $[0, 1]^P$ defined by

$$\mathcal{C}(P) = \{\omega \in [0, 1]^P : \mathbf{v}_1 < \mathbf{v}_2 < \dots < \mathbf{v}_k \implies \omega(\mathbf{v}_1) + \omega(\mathbf{v}_2) + \dots + \omega(\mathbf{v}_k) \leq 1\},$$

where k is an arbitrary positive integer.

Clearly $\mathcal{C}(P)$ is a bounded polytope, and by the positivity assumption on the ω it is enough to restrict the chains in the definition to just the maximal chains. Therefore $\mathcal{C}(P)$ is the same as the event that $\{G_P \leq 1\}$ for iid uniform weights. This left-tail probability of the passage time distribution is shown in [Stanley \(1986, Corollary 4.2\)](#) to be equal to

$$\mathbb{P}(G_P \leq 1) = \text{Vol}(\mathcal{C}(P)) = \frac{e(P)}{|P|!}, \quad (7.1)$$

where $e(P)$ is the number of linear extensions of P (equivalently the number of maximal chains in $J(P)$). For example, on the subposet $[1, m] \times [1, n]$ of \mathbb{Z}^2 this

probability is

$$\mathbb{P}(G(m, n) \leq 1) = \prod_{i=1}^m \prod_{j=1}^n \frac{1}{m - i + n - j + 1},$$

thanks to the hook length formula. This particular probability is of limited use since the event $\{G_P \leq 1\}$ is so far from the typical behavior $G(m, n) \sim c(m + n)$ as $m, n \rightarrow \infty$, but we can still use the chain polytope to give a characterization of G_P . In particular, we can represent the passage time G_P in terms of the ℓ^1 norm of a uniformly chosen point from a random chain polytope $\mathcal{C}(P^*)$. Here P^* is a random poset whose distribution is determined by P . It is constructed in the following way: begin with (5.3) and let $U = U(\omega)$ be the (Lebesgue almost surely unique) maximal chain corresponding to ω . Then $\omega \in \text{Span}_+ \{\partial U_1(\omega), \dots, \partial U_{|P|}(\omega)\}$ implies that there are Lebesgue almost surely positive $\lambda_i(\omega)$ such that $\omega = E_U \lambda_U$, or equivalently

$$\omega = \sum_{i=1}^{|P|} \lambda_i(\omega) \partial U_i(\omega).$$

But also, since $\|\omega\|_\infty \leq 1$ almost surely, this means that the λ_i must satisfy

$$\left\| \sum_{i=1}^{|P|} \lambda_i \partial U_i \right\|_\infty \leq 1.$$

Since the entries of each antichain ∂U_i are either one or zero, this means that

$$\sum_{i=1}^{|P|} \lambda_i \mathbf{1} \{ \mathbf{v} \in \text{supp}(\partial U_i) \} \leq 1, \quad \text{for all } \mathbf{v} \in P. \tag{7.2}$$

Combined with the positivity condition, the set of λ_i satisfying these inequalities is the chain polytope of some poset $P^* = P^*(U)$ that is determined by the particular ordering U . In fact the choice of P^* is not unique, but this will not concern us since we will only be concerned with the number of linear extensions of P^* which turns out to be an invariant. For now let P^* be any poset such that

$$\mathcal{C}(P^*) = \left\{ \lambda_i \geq 0 : \sum_{i=1}^{|P|} \lambda_i \mathbf{1} \{ \mathbf{v} \in \text{supp}(\partial U_i) \} \leq 1 \text{ for all } \mathbf{v} \in P \right\}.$$

Now ω being uniformly distributed on $[0, 1]^P$ means also that its density on $\mathcal{C}_U \cap [0, 1]^P := \text{Span}_+ \{\partial U_1, \dots, \partial U_{|P|}\} \cap [0, 1]^P$ is uniform. On the latter set there is a bijection between $\omega \in \mathcal{C}_U \cap [0, 1]^P$ and $\lambda \in \mathcal{C}(P^*)$ given by $\omega = E_U \lambda_U$, and since E_U is volume preserving (by Lemma 5.9) it follows that λ is uniformly distributed on $\mathcal{C}(P^*)$, and that

$$\text{Vol}(\mathcal{C}_U \cap [0, 1]^P) = \frac{e(P^*)}{|P^*|!} = \frac{e(P^*)}{|P|!}.$$

Now via the formula

$$G_P = \sum_{i=1}^{|P|} \lambda_i$$

for the passage time, this leads to a method for sampling G_P when ω is iid Uniform(0, 1):

- sample a maximal chain $U_0 \triangleleft U_1 \triangleleft \dots \triangleleft U_{|P|}$ according to the probabilities $e(P^*(U))/|P|!$,
- sample λ as a uniform point in the associated chain polytope $\mathcal{C}(P^*(U))$,
- return the ℓ^1 norm $\sum_{i=1}^{|P|} \lambda_i$ as the passage time.

This characterizes G_P as the ℓ^1 norm of a point chosen uniformly from a random chain polytope, although the distribution of the polytope seems to be complicated. Evidently the final answer does not depend on the choice of P^* but for the sake of concreteness we give one possible way of constructing it. By (7.2) the inequalities defining $\mathcal{C}(P^*)$ can be written as

$$\sum_{i=\eta_{\mathbf{v}}}^{\tau_{\mathbf{v}}} \lambda_i \leq 1,$$

where $\eta_{\mathbf{v}} = \min\{i : \mathbf{v} \in \text{supp}(\partial U_i)\}$ and $\tau_{\mathbf{v}} = \max\{i : \mathbf{v} \in \text{supp}(\partial U_i)\}$ are the first and last times that \mathbf{v} is in $\text{supp}(\partial U_i)$, respectively. Now define a graph with the times $\{1, 2, \dots, |P|\}$ as its vertices and an edge connecting i, j iff there exists a $\mathbf{v} \in P$ such that $i, j \in [\eta_{\mathbf{v}}, \tau_{\mathbf{v}}]$. Thus an edge between two times i and j means that there is an element of P which is in both ∂U_i and ∂U_j , or in other words $\partial U_i \cap \partial U_j \neq \emptyset$. This graph may have multiple connected components, and each component corresponds to a component of $P^*(U)$. It can quickly be seen that a new component is born every time that there is a ∂U_i with exactly one non-zero entry, so that

$$\text{number of components of } P^*(U) = 1 + |\{i : \text{supp}(\partial U_i) \text{ is a singleton}\}|.$$

The constructed graph is the *comparability graph* of the poset P^* . The comparability graph of any poset is defined with the elements of P as its vertices and an edge connecting two vertices iff the corresponding poset elements are comparable to each other (i.e. one is in order with the other). See Golumbic (2004, Chapter 5) for more. We have already encountered this object implicitly: the order graph of Definition 4.1 is the comparability graph of the poset restricted to $\text{supp}(a_1)\Delta\text{supp}(a_2)$. Similarly, the disorder graph of Definition 1.2 is the incomparability graph of the poset restricted to $\pi\Delta\pi'$, where the incomparability graph is defined similarly but with edges between elements that are incomparable to each other. The number of linear extensions of the poset is determined by its comparability graph, showing that the choice of P^* above is irrelevant. Constructing a candidate P^* is the problem of choosing a transitive orientation for the graph: an assignment of directions to the edges such that the adjacency relation of the resulting directed graph is transitive. Algorithms for finding transitive orientations are found in Golumbic (2004, Chapter 5), as are formulas for the number of transitive orientations of the graph.

8. Open Questions

Most Likely Paths. Our interest in this problem was primarily motivated by the following question: on the subposet $P = [1, n]^2$ of \mathbb{Z}^2 , which up-right path has the highest probability of being the largest? Or in other words, which path $\pi \in \Pi_P$ maximizes $\mathbb{P}(\mathcal{C}(\pi))$ for a fixed probability measure \mathbb{P} on \mathbb{R}_+^P ? Even in the exactly solvable case of iid exponential weights this question appears to be difficult, for the reasons described in Section 6. We are confident that the answer is the extreme corner path that goes straight up to $(1, n)$ and then straight over to (n, n)

(or its obvious symmetric copy across the main diagonal) but we have been unable to prove this. Certainly since the fluctuations of the path away from the main diagonal are known to be of order $n^{2/3}$ in the exactly solvable cases, and thought to be of the same order in most iid cases, the paths that are outside of the window of size $n^{1/2}$ should have more weight than the paths inside this window, which of course supports the bulk of the probability under the uniform measure on paths.

Negative Correlations Between Transversal Fluctuations of the Maximizer and its Length. Again we consider the poset $[1, n]^2$ in \mathbb{Z}^2 . Our belief that the outside paths have the largest individual probability of being longest is motivated by the idea that their maximal cones $\mathcal{C}(\pi)$ take up more of the environment space than the other maximal cones. Since the passage time G_P is the inner product between the weight and the corresponding longest path, if a maximal cone $\mathcal{C}(\pi)$ is relatively large, then conditionally on ω being in that cone there is more room for it to point away from the path vector π . Conversely, if $\mathcal{C}(\pi)$ is relatively small then there is little room for an ω inside of it to point away from π . Thus we expect that there should be a *negative* correlation between the length of a path and its location in the rectangle. It would be interesting to see if this relationship effectively cancels in terms of expected values, so that each path contributes close to the same amount to the overall expected value of the passage time.

Structure of the Path Measure. For Lebesgue almost all $\omega \in \mathbb{R}_+^P$ the path lengths are all distinct, this leads to a total of $|\Pi_P|$ possible path lengths. Yet they are created from only $|P|$ random variables, leading to a strong linear dependency between the different path lengths. For example, in the case $P = [1, m] \times [1, n]$ there are a total of

$$\binom{m+n}{n}$$

different path lengths, yet it can be shown that the vector of these path lengths lives in a subspace of dimension $(m-1)(n-1)+1$ (it is actually slightly smaller than mn). Consequently, every path length can be expressed as a linear combination of $(m-1)(n-1)+1$ well chosen path lengths. We are curious to know if a similar relationship holds for the path probabilities $\mathbb{P}(\mathcal{C}(\pi))$. Does knowing $\mathbb{P}(\mathcal{C}(\pi))$ for a relatively small subset of $\pi \in \Pi_P$ determine $\mathbb{P}(\mathcal{C}(\pi))$ for all $\pi \in \Pi_P$? A result of this type would be useful as it is generally difficult to describe the measure $\mathbf{P}(\pi) := \mathbb{P}(\mathcal{C}(\pi))$ on Π_P . If one thinks of LPP as a random walk in a random environment (although really it is just a walk in a random environment), then \mathbf{P} is the *averaged* path measure. This \mathbf{P} is not Markov, which immediately takes away one of the nicest descriptions for path measures, and is generally difficult to compute explicitly. On the other hand there should be more structure than it simply being a point in the probability simplex of dimension $|\Pi_P|$. It does not seem unreasonable to expect \mathbf{P} to have some “low-dimensional” structure, although we are uncertain of what precisely it would be.

Monotonicity of the Path Measure. If, for squares $[1, n]^2$ in \mathbb{Z}^2 , the outside paths have the highest probability of being the longest and the middle paths have the lowest probability, then we expect that there should be some sort of monotonicity in the path probabilities as the paths move from the middle to the outside. It

is not entirely clear to us what the proper ordering on paths should be, or even if a total ordering exists. We expect that there should be a natural partial ordering on the paths, based solely on their relative locations in space, such that their respective probabilities obey the partial ordering. It would be interesting to derive a rate of growth of these path probabilities along a chain in this partial ordering.

Path Measure Proportional to the Number of Extreme Rays. On $[1, m] \times [1, n]$ in \mathbb{Z}^2 it would be interesting to study the path measure that is proportional to the number of extreme rays per path, i.e. the measure on $\Pi_{(m,n)}$ with probabilities

$$\mathbf{P}(\pi) = \frac{|\mathcal{ER}(\pi)|}{\sum_{\pi} |\mathcal{ER}(\pi)|}.$$

This measure is purely combinatorial but may still exhibit many of the features of the annealed last passage measure. Given that the extreme rays of $\mathcal{ER}(\pi)$ have relatively large angles between them, one might expect that the number of extreme rays of $\mathcal{C}(\pi)$ is a reasonable proxy its volume, under certain measures on \mathbb{R}_+^P . If so it would give some understanding of the expected universality behavior of the paths. In particular it would be interesting to know if the transversal fluctuations of the path are superdiffusive under the annealed measure. One might also hope that the asymptotics of probabilities of certain special paths, such as the extremal ones, could be computed under this measure.

Face Lattice of $\mathcal{C}(\pi)$. The faces of a polyhedral cone are any of its intersections with half-spaces with the property that no interior point of the cone lies on the boundary of the half-space. The set of faces can be made into a lattice (in fact an Eulerian lattice), where the partial ordering is determined by set containment of faces. We have not been able to fully determine the full structure of these inclusions for $\mathcal{C}(\pi)$, beyond Theorem 4.2 which explains the inclusion of the one dimensional faces (the extreme rays) in the two-dimensional faces. Being simplicial, each of the simplices $\partial U_1, \dots, \partial U_{|P|}$ described in Section 5 has the well known binomial poset of the appropriate size as its face lattice, but many of the faces will be interior to $\mathcal{C}(\pi)$. A more useful description of the face lattice would be in terms of the extreme rays of $\mathcal{C}(\pi)$ or the normal vectors $\pi - \pi'$ that define its facets, with the inclusions being expressed in terms of relations between these vectors. Part of our decision to explain the last passage model on general posets is motivated by a desire to explain the face lattice. Our hope is that the description of the facets can be iterated in some way to provide a description of the face lattice. Since the lower dimensional faces can be seen as “facets of facets”, if the facets can be described as last passage percolation problems on a smaller poset, then Theorem 1.3 can be applied again without modification. This indeed works for the facets of the form $\omega(\mathbf{v}) = 0$, but is more complicated for those of the form $\langle \pi - \pi', \omega \rangle = 0$. We do not know of a description of these facets as a last passage model on a smaller poset, but perhaps there is such a description in terms of a matroid or a related object.

Number of Young Tableaux that Produce a Given Longest Path. The simplicial decomposition of Section 5 suffers from the curse of dimensionality, in that the number of simplices used to partition each cone $\mathcal{C}(\pi)$ is exponentially larger than the ambient space. The number of such simplices could perhaps be used as a rough proxy for the probability of each $\mathcal{C}(\pi)$. On $[1, m] \times [1, n]$ we expect

that among all paths $\pi \in \Pi_P$ the extreme corner paths have the largest number of simplices in their decomposition, which would lead credence to our belief that the corner paths are the modes of the distribution. This assertion is equivalent to stating that the corner path is the Schützenberger path for the largest number of Young tableaux. We would be interested in asymptotics of the number of Young tableaux that produce a given longest path.

Polymerization. Closely related to the last passage percolation problem is the notion *directed polymers*. Here the environment variables do not on their own determine a path, instead there is some extra randomness involved. For each inverse temperature $\beta \geq 0$ and $\omega \in \mathbb{R}^P$ the directed polymer measure on Π_P is defined by the Gibbs measure

$$\mathbf{P}_\beta^\omega(\pi) = e^{\beta\langle\pi,\omega\rangle} / Z_\beta^\omega$$

where $Z_\beta^\omega = \sum_{\pi \in \Pi_P} e^{\beta\langle\pi,\omega\rangle}$ is the partition function. Note that in this case we remove the restriction that the ω have positive coordinates. With ω fixed this is often referred to as the quenched measure. As $\beta \rightarrow \infty$ it is clear that the quenched measure concentrates on the longest path (or splits uniformly amongst all paths that achieve the longest length, if there are several). In the finite temperature setting of directed polymers the analogue of the maximal cones $\mathcal{C}(\Pi)$ are the sets of ω which produce a given value of $\mathbf{P}_\beta^\omega(\pi)$, i.e. for each \mathbf{Q} a probability measure on Π_P one considers the set

$$\mathcal{V}_\beta(\mathbf{Q}) = \{\omega \in \mathbb{R}^P : \mathbf{P}_\beta^\omega = \mathbf{Q}\}.$$

Note that for many \mathbf{Q} this measure is likely empty, and it is an interesting problem to determine useful conditions on \mathbf{Q} for which this is not the case. If it is not then in the variables $e^{\beta\omega(\mathbf{v})}$ the set is an algebraic variety, and it would be interesting if any meaningful description of it can be made. In particular, one might hope that properties of the maximal cones $\mathcal{C}(\pi)$ can be transferred to properties of these varieties.

Acknowledgements

Tom Alberts gratefully acknowledges the contributions of Bryant Lin, who performed computer work that laid the groundwork for this project as part of a Summer Undergraduate Research Fellowships (SURF) program at Caltech. Alberts and Lin thank the SageMath project (www.sagemath.org) for their excellent software which enabled us to explore and verify properties of the maximal cones. Alberts and Cator thank Leonid Petrov for helpful discussions, and an anonymous referee who made many helpful suggestions that lead to a much improved presentation.

References

- Aldous, D. and Diaconis, P. Hammersley’s interacting particle process and longest increasing subsequences. *Probab. Theory Related Fields*, **103** (2), 199–213 (1995). [MR1355056](#).
- Balázs, M., Cator, E., and Seppäläinen, T. Cube root fluctuations for the corner growth model associated to the exclusion process. *Electron. J. Probab.*, **11**, no. 42, 1094–1132 (2006). [MR2268539](#).

- Brightwell, G. and Winkler, P. Counting linear extensions. *Order*, **8** (3), 225–242 (1991). [MR1154926](#).
- Cator, E. and Groeneboom, P. Hammersley’s process with sources and sinks. *Ann. Probab.*, **33** (3), 879–903 (2005). [MR2135307](#).
- Corwin, I. Macdonald processes, quantum integrable systems and the Kardar-Parisi-Zhang universality class. In *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. III*, pp. 1007–1034. Kyung Moon Sa, Seoul (2014). [MR3729062](#).
- Corwin, I. Exactly solving the KPZ equation. In *Random growth models*, volume 75 of *Proc. Sympos. Appl. Math.*, pp. 203–254. Amer. Math. Soc., Providence, RI (2018). [MR3838899](#).
- Corwin, I., Quastel, J., and Remenik, D. Renormalization fixed point of the KPZ universality class. *J. Stat. Phys.*, **160** (4), 815–834 (2015). [MR3373642](#).
- Dauvergne, D., Ortman, J., and Virág, B. The directed landscape. *ArXiv Mathematics e-prints* (2018). [arXiv: 1812.00309](#).
- Fulton, W. *Young tableaux. With applications to representation theory and geometry*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge (1997). ISBN 0-521-56144-2; 0-521-56724-6. [MR1464693](#).
- Georgiou, N. and Ortman, J. Optimality regions and fluctuations for Bernoulli last passage models. *Math. Phys. Anal. Geom.*, **21** (3), Paper No. 22, 29 (2018). [MR3835628](#).
- Georgiou, N., Rassoul-Agha, F., and Seppäläinen, T. Variational formulas and cocycle solutions for directed polymer and percolation models. *Comm. Math. Phys.*, **346** (2), 741–779 (2016). [MR3535900](#).
- Georgiou, N., Rassoul-Agha, F., and Seppäläinen, T. Stationary cocycles and Busemann functions for the corner growth model. *Probab. Theory Related Fields*, **169** (1-2), 177–222 (2017). [MR3704768](#).
- Golumbic, M. C. *Algorithmic graph theory and perfect graphs*, volume 57 of *Annals of Discrete Mathematics*. Elsevier Science B.V., Amsterdam, second edition (2004). ISBN 0-444-51530-5. [MR2063679](#).
- Johansson, K. The longest increasing subsequence in a random permutation and a unitary random matrix model. *Math. Res. Lett.*, **5** (1-2), 63–82 (1998). [MR1618351](#).
- Johansson, K. A multi-dimensional Markov chain and the Meixner ensemble. *Ark. Mat.*, **48** (1), 79–95 (2010). [MR2594587](#).
- O’Connell, N. and Yor, M. Brownian analogues of Burke’s theorem. *Stochastic Process. Appl.*, **96** (2), 285–304 (2001). [MR1865759](#).
- Priezzhev, V. B. and Schütz, G. M. Exact solution of the Bernoulli matching model of sequence alignment. *Journal of Statistical Mechanics: Theory and Experiment*, **2008** (09), P09007 (2008). DOI: [10.1088/1742-5468/2008/09/p09007](#).
- Romik, D. *The surprising mathematics of longest increasing subsequences*, volume 4 of *Institute of Mathematical Statistics Textbooks*. Cambridge University Press, New York (2015). ISBN 978-1-107-42882-9; 978-1-107-07583-2. [MR3468738](#).
- Romik, D. and Śniady, P. Jeu de taquin dynamics on infinite Young tableaux and second class particles. *Ann. Probab.*, **43** (2), 682–737 (2015). [MR3306003](#).
- Rost, H. Nonequilibrium behaviour of a many particle process: density profile and local equilibria. *Z. Wahrsch. Verw. Gebiete*, **58** (1), 41–53 (1981). [MR635270](#).

- Seppäläinen, T. Increasing sequences of independent points on the planar lattice. *Ann. Appl. Probab.*, **7** (4), 886–898 (1997). [MR1484789](#).
- Seppäläinen, T. Unpublished Lecture Notes on the Corner Growth Model (2009). URL: <https://www.math.wisc.edu/~seppalai/cornergrowth-book/ajo.pdf>. Last visited on 2019/05/20.
- Seppäläinen, T. Scaling for a one-dimensional directed polymer with boundary conditions. *Ann. Probab.*, **40** (1), 19–73 (2012). [MR2917766](#).
- Stanley, R. P. Two poset polytopes. *Discrete Comput. Geom.*, **1** (1), 9–23 (1986). [MR824105](#).
- Stanley, R. P. *Enumerative combinatorics. Volume 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition (2012). ISBN 978-1-107-60262-5. [MR2868112](#).