ON DIFFERENTIAL POSETS

JOEL LEWIS
HARVARD UNIVERSITY

Abstract. We study differential posets, a family of partially ordered sets discovered by Richard Stanley. In the first half of this paper we present an introduction to poset theory as relevant to differential posets and theorems on the structure and combinatorial properties of differential posets, culminating in the explicit definition of a new differential poset. In the second half we focus on Young's lattice, the most well-studied differential poset, and the RSK algorithm.

0. Introduction

In this paper, we study the properties of differential posets, a family of structures introduced by Richard Stanley in [5] and since generalized in various different ways (see, for example, [6, 2]).[1] Differential posets are those which share certain structural properties of Young's lattice $\mathbf{Y}$, an extremely well-studied object of combinatorics and algebra.[2] Thus, they provide a route for generalizing results about $\mathbf{Y}$ to a much broader family of posets.

In Section 1 we present the terminology which will be used throughout the rest of the paper, most of which comes from [7]. For the sake of simplicity we have in some instances differed slightly from the conventional definition of a term (e.g. locally finite), but in no case does this change the content of the theorems which result.

In Section 2 we introduce differential posets and two linear transformations D, U which are extremely powerful tools in the study of differential posets. For much more on their uses, see [5, 1].

Section 3 deals with the construction of new differential posets, including the Fibonacci-Young lattice $\mathbf{FY}$, via recursive processes called reflections. This culminates in what is believed to be the first explicit

Date: April 2, 2007.

[1]In terms of Stanley's definition ([5]), we actually concern ourselves only with 1-differential posets rather than the broader family of r-differential posets. However, it seems likely that most of the interesting properties of r-differential posets are captured by the 1-differential case.

[2]See, for example, the note following Proposition 4.5.

definition of a differential poset other than Young's lattice and the Fibonacci-Young lattice.

Section 4 deals primarily with Young's lattice and the famous RSK algorithm. We define and prove the basic properties of the classical insertion version of RSK. We then outline an alternative approach based on Fomin's work in [2] that applies for all differential posets and show that this truly coincides with RSK on $\mathbf{Y}$. (A tableau-based insertion algorithm is also known for $\mathbf{FY}$ – for a proof that the generalized RSK presented here also generalizes that algorithm, see [3].)

Where possible, citations to the literature have been provided. Some results are extremely well-known but do not often appear in the form shown (e.g. Theorem 1.1, actually a corollary of Proposition 4.2; $\mathbf{Y}$ is often taken by definition to be a poset of order ideals and so automatically a distributive lattice). Other content is known but unpublished (e.g., the reflection and crown reflection operations, which were discovered by David Wagner – see [5], Section 6). What little remains (notably Theorem 3.5), as well as the exposition, we believe to be original.

We have tried to write this paper in an entirely self-contained manner, with references representing potential sources for further reading rather than presumed knowledge. Thus, we hope that this paper should be accessible to an advanced undergraduate with no previous knowledge of the subject.

1. Preliminaries

A *partially ordered set* P, or *poset* for short, is a pair $P = (S, \leq)$ of a set S together with an order relation $\leq$ which satisfies the following conditions for all $x, y, z \in S$:

(1) $x \leq x$ (Reflexivity)
(2) $x \leq y$ and $y \leq x$ imply $x = y$ (Anti-symmetry)
(3) $x \leq y$ and $y \leq z$ imply $x \leq z$ (Transitivity)

For brevity and because it does not harm comprehensibility, we will often omit specific mention of the order relation and will frequently refer to the "elements of a poset," rather than the elements of the underlying set. If x and y are elements of some poset P and $x \leq y$ or $y \leq x$ we say that x and y are *comparable*. Note, in particular, that we do not require that two different elements of P be comparable. We sometimes write $y \geq x$ for $x \leq y$ and $x < y$ or $y > x$ for $x \leq y$ and $x \neq y$. An element x is *minimal* if there is no y such that $y < x$ and x is *maximal* if there is no y such that $x < y$.

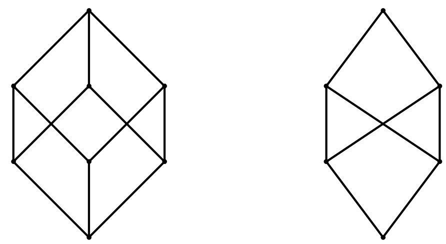

FIGURE 1. Two Hasse diagrams.

Given a poset P and elements $x, y \in P$, we say that y *covers* x if $x < y$ and there is no z such that $x < z < y$. Then y is *a cover* of x and x is *covered by* y, and we denote this relationship by $x \lessdot y$ or $y \gtrdot x$. We will also use the notations $C^-(x) = \{y \mid y \lessdot x\}$ and $C^+(x) = \{y \mid x \lessdot y\}$.

For many posets, including all of those we will consider, all order relations follow from the cover relations and transitivity. A poset with this property is called *locally finite*.

A wide variety of objects exhibit poset structure. We provide a few examples here:

- The non-negative integers $\mathbf{N}$ together with the usual order relation $\leq$ form a poset. This poset is *totally ordered*; that is, every two elements are comparable. Each element $n \in \mathbf{N}$ is covered uniquely by $n + 1$. This poset has a minimal element, 0.
- The non-negative integers $\mathbf{N}$ together with the order relation $\leq$ defined by $a \leq b$ if and only if there is some $k \in \mathbf{N}$ such that $b = a \cdot k$ form a poset. This poset has a minimal element, 1, and a maximal element, 0. An element $n \neq 0$ is covered by all integers of the form $p \cdot n$ for primes p.
- Given an arbitrary set S, we can form the *Boolean lattice* B_S on S. The elements of B_S are the subsets of S, ordered so that $X \leq Y$ if and only if $X \subseteq Y$ as sets. This poset has a minimal element, $\emptyset$, and a maximal element, S. An element $X \in B_S$ is covered by all those sets of the form $X \cup \{x\}$ for $x \in S \setminus X$.

Every locally finite poset has a naturally associated *Hasse diagram*, a graph whose vertices are elements of the poset and whose edges denote cover relations, where if $x \lessdot y$ we draw x below y. Figure 1 shows the Hasse diagram of two posets; the graph on the left is the Hasse diagram of the Boolean lattice on a three-element set.

A *chain* in a poset P is a set of elements of P which are pairwise comparable while an *anti-chain* is a set of points which are pairwise incomparable. A finite chain $x_0 < x_1 < \ldots < x_n$ is *saturated* if x_{i-1} is covered by x_i for $i \in \{1, \ldots, n\}$. The length of a chain is one less than the number of elements in the chain, so the trivial chain $\{x\}$ for any x has length 0.

We will call a locally finite poset P *graded* if $P = \amalg_n P_n$ is the disjoint union of antichains P_n, indexed by a set of consecutive integers, such that $x \in P_n$ and $y \gtrdot x$ imply $y \in P_{n+1}$ and all minimal elements (if any exist) belong to the same P_i. If P has a unique minimal element $\hat{0}$, the *rank* of $x \in P$ is the maximal length of a saturated chain with largest element x. Thus, $\hat{0}$ is rank 0, every element which covers $\hat{0}$ has rank 1, and so on.

Two posets are *isomorphic* if there exists a bijection between their elements which is order-preserving in both directions. Typically, we will be interested in studying posets up to isomorphism and also in finding "nice" members of particular isomorphism classes. Usually, this means finding a poset in which both the underlying set and the order relations, or at least cover relations, can be given in a compact, explicit form. For finite posets this goal is trivial, since the elements and order relations could in principle be listed out, but for infinite posets this may be more challenging.

We provide one more example of a poset. In order to define it, we need first to introduce some associated terminology.

Given a non-negative integer n, a *partition* of n is a finite, nonincreasing list of positive integers, $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$, such that $\lambda_1 + \ldots + \lambda_k = n$. We denote this by $\lambda \vdash n$. For example, there are five partitions of 4: $(4), (3, 1), (2, 2), (2, 1, 1)$ and $(1, 1, 1, 1)$. There is one partition of 0, the empty list. We define an order on the partitions as follows: given two partitions, $\lambda = (\lambda_1, \ldots, \lambda_k)$ and $\mu = (\mu_1, \ldots, \mu_j)$, $\lambda \leq \mu$ if and only if $k \leq j$ and $\lambda_i \leq \mu_i$ for $1 \leq i \leq k$. The poset composed of all integer partitions, ordered in this way, is known as *Young's lattice* and is denoted $\mathbf{Y}$.

Given any two elements in a poset, we can ask whether there are any elements smaller than both or larger than both. In particular, we can ask if they have a greatest lower bound and least upper bound. In general, the answer to this question will be "no." For example, consider in the poset represented on the right in Figure 1 the two elements of rank 2. If in a poset P, for every pair $x, y \in P$ we have that there exist $w, z \in P$ such that $w \leq x, y \leq z$ and for all u, v, $u \leq x, y \leq v$ implies $u \leq w$ and $z \leq v$, we say that P is a *lattice*. In this case, we write $w = \mathrm{glb}(x, y)$ and $z = \mathrm{lub}(x, y)$. Every poset we have seen so far is a

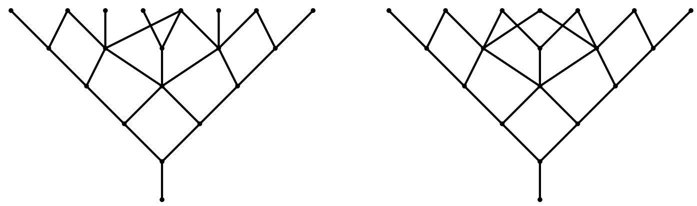

FIGURE 2. Both possibilities for the first six ranks of a differential poset

lattice: for example, in any totally ordered poset $\mathrm{glb}(x, y) = \min(x, y)$ and $\mathrm{lub}(x, y) = \max(x, y)$ while in a Boolean lattice, $\mathrm{glb}(x, y) = x \cap y$ and $\mathrm{lub}(x, y) = x \cup y$.

Theorem 1.1. *Young's lattice is a lattice.*

Proof. We prove only that $\mathbf{Y}$ is an *upper semi-lattice*, that is that least upper bounds exist. Given two partitions λ, μ, we can pad the shorter list with 0s so that both are of the same length k. Any partition ν such that $\nu \geq \lambda$ and $\nu \geq \mu$ must certainly satisfy $\nu_i \geq \max(\lambda_i, \mu_i)$ for $1 \leq i \leq k$. Thus, if

$$M(\lambda, \mu) = ((\max(\lambda_1, \mu_1), \ldots, \max(\lambda_k, \mu_k))$$

we have $\nu \geq \lambda$ and $\nu \geq \mu$ imply $\nu \geq M(\lambda, \mu)$. Surely $M(\lambda, \mu) \geq \lambda$ and $M(\lambda, \mu) \geq \mu$, so $M(\lambda, \mu)$ is the least upper bound of μ and λ. The proof that greatest lower bounds exist is nearly identical. □

2. Differential posets

We call a poset P *differential* if it satisfies the following three axioms:

(D1) P is locally finite and graded with a unique minimal element (often denoted $\hat{0}$).

(D2) If $x \neq y$ are two elements of P and there are k elements of P covered by both x and y, there are exactly k elements of P which cover both x and y.

(D3) If $x \in P$ covers k elements of P then x is covered by exactly $k + 1$ elements of P.

These conditions are enough to completely determine the number of elements and the cover relations among them in the zeroth through fourth ranks of every differential poset. Above the fourth rank, both the cover relations and the number of elements at each rank vary among the different posets. Figure 2 shows the two possibilities for the first six ranks of the Hasse diagram of a differential poset.

We may consolidate our definition in the following way. Given any poset P, we may define an abstract vector space $\mathbf{C}P = \oplus_{x \in P} \mathbf{C}x$ of finite linear combinations of elements of P with complex coefficients, where no additive relations hold among the elements of P. (In fact, any field of characteristic 0 would suffice in place of the complex numbers.) If in addition P is locally finite and each element of P is a member of only finitely many cover relations, we may define two linear transformations D and U on $\mathbf{C}P$ as follows: for $x \in P$, we define

$$Dx = \sum_{y \lessdot x} y \qquad \text{and} \qquad Ux = \sum_{x \lessdot y} y,$$

and we extend both to all of $\mathbf{C}P$ by linearity. In essence, Ux keeps track of all possible steps "up" the Hasse diagram from x and Dx keeps track of all the steps "down."

Now we investigate the behavior of certain combinations of D, U on an arbitrary $x \in P$. Applying our definitions directly gives

$$DUx = \sum_{\substack{y,z \\ z \lessdot y \text{ and } x \lessdot y}} z.$$

An element $z \in P$ appears in this sum exactly k times, where k is the number of elements of P which cover both x and z. In particular, x appears the same number of times as the number of elements that cover x.

Similarly, we have

$$UDx = \sum_{\substack{y,z \\ y \lessdot z \text{ and } y \lessdot x}} z.$$

An element $z \in P$ appears in this sum exactly k times, where k is the number of elements of P which are covered by both x and z. In particular, x appears the same number of times as the number of elements that x covers.

Thus, we see that $(DU - UD)x = x$ if and only if x is covered by exactly one more element than it covers and for each $z \neq x \in P$, the number of elements covering both x and z is equal to the number of elements covered by both x and z. It follows immediately that $DU - UD = I$ if and only if P is differential ([5], Theorem 2.2).

Proposition 2.1. *For a differential poset P, we have $DU^n = nU^{n-1} + U^n D$ for all $n \geq 1$.*

Proof. The case $n = 1$ follows immediately from the discussion above. In general, assuming $DU^k = kU^{k-1} + U^k D$ gives

$$\begin{aligned} DU^{k+1} &= (DU)U^k \\ &= (I + UD)U^k \\ &= U^k + U(DU^k) \\ &= U^k + kU^k + U^{k+1}D \\ &= (k+1)U^k + U^{k+1}D, \end{aligned}$$

so we have our claim by induction. □

Thus, the action of D on U has a distinct resemblance to that of a differential operator. The resemblance is even more clear if we apply the operator to our minimal element $\hat{0}$, since $D\hat{0} = 0$ so $DU^n\hat{0} = nU^{n-1}\hat{0}$. It is from a generalization of this result ([5], Corollary 2.4) that the name "differential poset" arises.

The transformations U and D have been studied in great detail. They provide an extremely useful tool for enumerating various sorts of paths (*Hasse walks*, in which each step consists of a move from the present element to an element which it covers or is covered by) in differential posets. For example, if $x, y \in P$ and the rank of y is k greater than the rank of x, then the number of minimal Hasse walks from x to y (also the number of saturated chains joining x and y) is just the coefficient of y in $U^k x$. The total number of paths from x which consist of k upward steps is the sum of all coefficients in $U^k x$. Similarly, paths with other structures are enumerated by sums of coefficients in other combinations of U and D. We provide one simple example here:

Proposition 2.2 ([5])**.** *In any differential poset P, the number of "loops" of length $2n$ beginning and ending at $\hat{0}$ (the number of Hasse walks consisting of n steps up the Hasse diagram followed by n steps down) is equal to $n!$.*

We will discuss this result further in the context of Young's lattice in Section 4.

Proof. The number of such loops is exactly equal to the coefficient of $\hat{0}$ in $D^nU^n\hat{0}$. We have trivially that $D^0U^0\hat{0} = \hat{0}$, so our proposition is correct for $n = 0$. Then

$$D^{k+1}U^{k+1}\hat{0} = D^k(DU^{k+1}\hat{0}) = (k+1)D^kU^k\hat{0},$$

and the result follows by induction. □

We now introduce a general method for building differential posets and with it one of the two most important differential posets.

3. Posets from reflections

A finite poset P is said to *look differential* if P is graded with a minimal element, every element of non-maximal rank is covered by one more element than it covers and each pair of elements of non-maximal rank both cover the same number of elements which cover both. That is, our poset satisfies (D1), (D2) and (D3) except for elements of maximal rank: it looks like the first n ranks of a differential poset.

Given a poset P of maximal rank n which looks differential, we construct a new poset P' as follows: each of the elements of P is an element of P' and all the order relations among the elements of P hold also as elements of P'. In addition, for each element x of rank $n-1$ in P, P' contains an element $2x$ which covers every element covering x and no others, and for every element y of rank n in P, P' contains an element $1y$ which covers y and no other elements. Geometrically, this operation is equivalent to reflecting the top level in the Hasse diagram of P, which guarantees that each pair of elements of rank n that cover elements in common are now covered by the same number of common elements, and adding one additional element covering each element of rank n of P, which ensures that every element of the n-th level is now covered by one more element than it covers.

From what we have said, one can see that the resulting poset P' also looks differential, but in this case up to the $(n+1)$-st rank. Also, this operation does not alter the poset in any way below the n-th rank. Thus, given a poset P of rank n which looks differential, we can generate from it a differential poset $F(P)$ as follows: the underlying set of $F(P)$ is the union of the underlying sets of $P, P', P'', \ldots$. Given any two elements $x, y \in F(P)$, there exists some n such that $x, y \in P^{(n)}$, and we have $x \leq y$ in $F(P)$ if and only if $x \leq y$ in $P^{(n)}$. This is well-defined because the reflection operation does not alter the comparisons between elements of P and so the relation between x and y does not depend on the choice of n.

Since conditions (D2) and (D3) are local constraints on the poset, they are satisfied for any $x, y \in F(P)$ because they are satisfied for $x, y \in P^{(n)}$ for some sufficiently large n. (D1) is a global property, but it is possible to check that all three of its conditions are met without difficulty.

This tells us that any finite structure which looks differential can be extended to a complete differential poset in a canonical way. In particular, the poset P_0 with only one element looks differential. The poset $F(P_0)$ which results from applying the above construction to P_0

is known as the *Fibonacci-Young lattice* and denoted $\mathbf{FY}$; it is in some sense the canonical differential poset.

We can say more about $\mathbf{FY}$ than just its existence. If we take the lone element of P_0 to be the empty string $\emptyset$ then the elements of $\mathbf{FY}$ are $\{1,2\}^*$, the set of all finite strings composed entirely of 1s and 2s. We can also deduce the cover relations of $\mathbf{FY}$:

Theorem 3.1. *Given two such strings x and y, $y \gtrdot x$ in $\mathbf{FY}$ if and only if y can be formed either by adding a 1 to x so that the new 1 is the leftmost in x or by changing the leftmost 1 in x to a 2.*

Equivalently, $x \lessdot y$ if and only if x can be formed by removing the leftmost 1 from y or by changing a 2 in y into a 1 so that the new 1 is the leftmost 1 in x.

For example, Theorem 3.1 implies that 22121 is covered by 122121, 212121, 221121 and 22221, and no other elements of $\mathbf{FY}$.

Proof. The proof of this theorem is inductive. It can be verified by hand that the cover relation of $\mathbf{FY}$, formed according to the rules of the reflection, follows the pattern described for the first few ranks. Suppose it holds up to the k-th rank and consider an arbitrary element x of rank k. We work in several cases:

- If $x = 1w$ for any string w, then by assumption x covers only w. Thus by the properties of the reflection x is covered by $1x$ and $2w$, as needed.
- If $x = 2^n$, where exponentiation represents repetition, by assumption x covers those strings of the form $2^{m-1}12^{n-m}$ for $1 \leq m \leq n$. Then x is covered by $1x$ and those strings of the form 2^m12^{n-m}, as needed.
- If $x = 2^n1w$ for some string w with $n \geq 1$, then by assumption x covers 2^nw and $2^{m-1}12^{n-m}1w$ for $1 \leq m \leq n$, so x is covered by $1x$, $2^{n+1}w$ and $2^m12^{n-m}1w$, as needed.

Since every $x \in \{1,2\}^*$ falls into one of these categories, our explicit form corresponds exactly to the recursive construction of the Fibonacci poset. $\square$

Theorem 3.2 ([5], Proposition 5.4)**.** *The Fibonacci-Young lattice is a lattice.*

Proof. First, we prove the existence of greatest lower bounds in $\mathbf{FY}$. It is clear that $\mathrm{glb}(\emptyset, \emptyset) = \emptyset$. We proceed by induction on the sum of the ranks of the two elements x, y, whose lower bound we seek. If x and y are comparable, then their greatest lower bound is simply whichever of the two is smaller. Now, suppose x and y are incomparable. Since

$C^-(1w) = \{w\}$, any lower bound on x and y is a lower bound on $1x$ and y and vice versa, so we have $\mathrm{glb}(x, y) = \mathrm{glb}(1x, y)$, and so in particular the latter is defined if the former is. With the inductive hypothesis, we are only left with to consider the case that $x = 2x'$ and $y = 2y'$ for some x', y'.

Note that for any $u, v \in \mathbf{FY}$ we have $u \leq v$ if and only if $2u \leq 2v$. In addition, if $1u \leq 2v$ then we have either $2u \leq 2v$ or $11u \leq 2v$, so either $2u \leq 2v$ or $21u \leq 2v$ or $111u \leq 2v$, and so on. Thus $21^n u \leq 2v$ for some n, so $1^n u \leq v$ and $u \leq v$.

It follows that $2\mathrm{glb}(x', y') \leq 2x' = x$ and $2\mathrm{glb}(x', y') \leq 2y' = y$. In addition, the comments of the previous paragraph show that if w' is the result of removing the leftmost digit of w (so $w = 1w'$ or $w = 2w'$) then $w \leq x$ and $w \leq y$ imply $w' \leq x'$ and $w' \leq y'$, so $w' \leq \mathrm{glb}(x', y')$ and $2w' \leq 2\mathrm{glb}(x', y')$. Since $1w' \leq 2w'$, in either case we can say that $w \leq 2\mathrm{glb}(x', y')$ and so $2\mathrm{glb}(x', y')$ is the greatest lower bound of x and y. Then we have by induction that every pair of elements of $\mathbf{FY}$ have a greatest lower bound.

Note that being able to take a pairwise greatest lower bound automatically gives us the ability to take a setwise greatest lower bound over any finite set,

$$\mathrm{glb}(x_1, \ldots, x_{k-1}, x_k) = \mathrm{glb}(x_1, \mathrm{glb}(x_2, \ldots, \mathrm{glb}(x_{k-1}, x_k) \ldots)).$$

Now, we construct a least upper bound of arbitrary x and y. Let z be the string that is the same length as the longer of x and y but is composed entirely of 2s. It is easy to see that z is an upper bound on x, y. Consider the set S of all $z_i \in \mathbf{FY}$ such that $z_i \geq x, z_i \geq y$ and the rank of z_i is at most that of z, and take the greatest lower bound z' of all elements of this set. Since $x \leq z_i$ and $y \leq z_i$ for all i, we have $x, y \leq z'$.

Consider any w such that $x, y \leq w$. Then $x, y \leq \mathrm{glb}(w, z') \leq z'$, so $\mathrm{glb}(w, z')$ has rank at most that of z'. Then by construction $\mathrm{glb}(w, z')$ must be one of the z_i and so $z' \leq \mathrm{glb}(w, z')$, so $z' = \mathrm{glb}(w, z')$ and $z' \leq w$. Thus $z' = \mathrm{lub}(x, y)$ and $\mathbf{FY}$ is a lattice. □

We now give two theorems on the structure of differential posets:

Theorem 3.3 ([5], Proposition 1.2)**.** *Given any poset P which satisfies (D1) and (D2) and two elements $x \neq y \in P$, there is at most one element which covers both. In other words, the constant k in (D2) is always 0 or 1.*

Proof. Suppose the contrary. Then choose elements $x_1, x_2, y_1, y_2 \in P$ such that both x_i are covered by both y_j and the x_i are of minimal rank. But then by (D2), there must be two elements which are both covered by x_1 and x_2. These elements are of rank one less than the x_i, contradicting the assumption of minimal rank. $\square$

Equivalently, we may say that the Hasse diagram of any differential poset P does not admit an induced $K_{2,2}$ subgraph. There are also some substructures which must occur in every differential poset.

Theorem 3.4. *Every differential poset P contains saturated chains $\{a_n\}_{n\geq 0}$ and $\{b_n\}_{n\geq 1}$ which satisfy the following conditions:*

(1) *$a_0 = \hat{0}$, $b_1 = a_1$ and $b_n \neq a_n$ for $k \geq 2$.*
(2) *For $n \geq 0$, $C^-(a_{n+1}) = \{a_n\}$ and $C^-(b_{n+2}) = \{a_{n+1}, b_{n+1}\}$. (In the case $n = 0$, this says $C^-(b_2) = \{a_1\}$.)*

Proof. We will build the desired chains recursively. There is only one possible choice for a_0, a_1 and b_1. There are two choices for a_2 so b_2 will be the other element of rank 2. Our conditions certainly hold for these elements. Now, suppose we have constructed finite chains $\{a_n\}_{0\leq n\leq k}$ and $\{b_n\}_{0\leq n\leq k}$ which satisfy our conditions. We can find a_{k+1}, b_{k+1} as follows:

We know by assumption that $C^-(a_{k-1}) = \{a_{k-2}\}$, so by (D3), a_{k-1} is covered by exactly two elements. We also know by assumption that $a_k \gtrdot a_{k-1}$ and $b_k \gtrdot a_{k-1}$, so these are the unique two elements which cover a_{k-1}. Thus some element x must cover both a_k and b_k, and a_k must be covered by one additional element $y \neq x$. Then $C^-(y) = \{a_k\}$ and $C^-(x) = \{a_k, b_k\}$, so $b_{k+1} = x$ and $a_{k+1} = y$ satisfy conditions (i) and (ii). $\square$

Note that for a fixed poset P, the only choice we made in this process was which of the two elements of rank 2 to label a_2; after that choice, the a_i and b_i were entirely determined by condition (ii). Define $f_P(x)$ to be the number of minimal-length Hasse walks in a differential poset P from $\hat{0}$ to x, or equivalently the number of saturated chains in P with minimal element $\hat{0}$ and maximal element x. Then we can count these paths by summing over possible penultimate elements to get

$$f_P(x) = \sum_{y \lessdot x} f_P(y).$$

It follows by induction that $f_P(a_k) = 1$ for all k. Moreover, since $f_P(x) \geq 1$ for all x, if $x \neq \hat{0}$ and $f_P(x) = 1$ then $C^-(x)$ must have exactly one element y and $f_P(y) = 1$. Thus, the a_n of the Theorem 3.4 (over both possible choices for a_2) represent all elements of P for which $f_P(x) = 1$.

We can use Theorem 3.4 together with the process of reflection from which we built the Fibonacci-Young lattice to construct a more varied family of differential posets. In Figure 2, the only difference between the two given structures occurs among those vertices covering the three vertices of rank four which cover a common vertex. In the image on the left, these three are all covered by a single vertex and each has in addition a vertex which covers it alone. This is exactly the structure we would get if we performed a reflection to the finite poset composed of the first five ranks pictured. In the image on the right, however, those three vertices have been covered pairwise by three vertices – there is a "crown" structure in the graph. This change preserves the number of covers of each vertex as well as the number of covers of each pair.

In fact, in a poset P which looks differential we can always switch between these two situations: if there are elements $x_1, x_2, x_3, y_0, y_1, y_2, y_3$, such that the y_i are of maximal rank,

$$C^-(y_0) = \{x_1, x_2, x_3\} \text{ and } C^-(y_i) = \{x_i\} \text{ for } i = 1, 2, 3$$

then there is a poset Q whose elements are the elements of P excluding the four y_i but including new elements z_1, z_2, z_3, and whose cover relations are those of P excluding those involving the y_i-s but including the six relations $x_i \lessdot z_j$ for $i \neq j$.

Note that the proof Theorem 3.4 applies equally well to posets which look differential as it does to differential posets. In addition, the b_i of Theorem 3.4 each cover exactly two elements for $i \geq 3$ and so they are each covered by exactly three elements. Thus, in every poset which looks differential of maximal rank greater than 3, we can find elements x_0, x_1, x_2, x_3 such that $C^+(x_0) = \{x_1, x_2, x_3\}$ and the x_1, x_2, x_3 are of maximal rank. Performing a reflection to this poset creates elements $y_0 = 2x_0, y_1 = 1x_1, y_2 = 1x_2$ and $y_3 = 1x_3$ which together with x_1, x_2, x_3 meet the conditions of the previous paragraph. Thus, we can replace these four elements with new elements z_1, z_2, z_3 as above, and the resulting poset will look differential. We will refer to this operation (reflection together with addition and deletion of elements) as a *crown reflection* or as *crowning above* x_0.

In the process of building a poset by repeated reflections we may substitute crown reflections as many times as we wish and in particular we have at each level the option to reflect or to crown reflect above the relevant b_i. Then to each infinite binary string we can associate a sequence of posets which look differential, where the $(n+1)$-st poset is formed from the n-th by a reflection if the corresponding digit of our sequence is 0, and by a crown reflection if the corresponding digit in the sequence is 1. Just as in the construction of $F(P)$, for a fixed

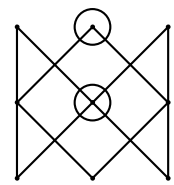

FIGURE 3. The two circled vertices have no greatest lower bound.

binary string we can take the union of all the posets in the associated sequence to give a differential poset.

Two differential posets are isomorphic only if they have the same number of elements of each rank. If we follow our construction from the same initial poset which looks differential (say, the unique such poset of maximal rank 4) with two different binary strings, we see that the first time the sequences differ, the poset produced by reflecting will have one more element of maximal rank then the poset produced by crown reflection. Thus, the two differential posets that eventually result will have differing numbers of elements at that rank and so will be non-isomorphic. It follows that we can assign to each infinite binary string a distinct isomorphism class of differential posets and so there are infinitely (in fact uncountably) many non-isomorphic differential posets.

Unfortunately, the crown reflection has a number of downsides. Although it produces a huge family of differential posets, it certainly does not contain them all. In particular, Young's lattice $\mathbf{Y}$ cannot be produced in this manner. (We will prove that Young's lattice is differential in Theorem 4.1.) In addition, the posets produced here will never be lattices: the Hasse diagram of every poset other than $\mathbf{FY}$ produced in this fashion will contain as an induced subgraph the structure shown in Figure 3, and the two circled vertices cannot have a greatest lower bound as elements of our poset. (In fact, no differential lattices are known other than $\mathbf{Y}$ and $\mathbf{FY}$.)

Finally, unlike the reflection we used to generate $\mathbf{FY}$, this process does not obviously produce a canonical labeling of the new elements. Thus, it is somewhat more difficult to produce "nice" labelings for these posets, a problem that is endemic to discussions of large, complicated structures. We provide here the explicit form for one poset produced in this manner; as far as we know, it is the first such expression for any differential poset other than $\mathbf{Y}$ or $\mathbf{FY}$.

Theorem 3.5. *Consider the poset P defined on the set of finite strings composed of* 1*s and* 2*s (including the empty string) except those strings* *which end in* $2211^n, n \geq 0$*, such that given two such strings* x *and* y*,* $y \gtrdot x$ *if and only if one of the following holds for some* $m, n \geq 0$*:*

(i) y *can be formed by inserting a* 1 *to the left of the leftmost* 1 *in* x*, except in the following subcases:*
  (a) $x = 2^m1^n22$ *and* $y = 2^m1^n122$
  (b) $x = 2^m1^n121$ *and* $y = 2^m1^n1121$
  (c) $x = 2111^n$ *and* $y = 12111^n$

(ii) y *can be formed by replacing the leftmost* 1 *in* x *with a* 2*, except in the following subcases:*
  (a) $x = 2^m1^n122$ *and* $y = 2^{m+1}1^n22$
  (b) $x = 2^m1^n1121$ *and* $y = 2^{m+1}1^n121$

(iii) $x = 2111^n$ *and* $y = 1^n122$ *or* $y = 1^n1121$
(iv) $x = 2^m1^n22$ *and* $y = 2^m1^n1121$
(v) $x = 2^m1^n1121$ *and* $y = 2^m21^n22$
(vi) $x = 2^m1^n22$ *and* $y = 2^m12111^n$
(vii) $x = 2^m12111^n$ *and* $y = 2^m21^n22$
(viii) $x = 2^m1^n121$ *and* $y = 2^m1^n122$
(ix) $x = 2^m1^n122$ *and* $y = 2^m21^n121$
(x) $x = 2^m1^n121$ *and* $y = 2^m12111^n$
(xi) $x = 2^m12111^n$ *and* $y = 2^m21^n121$

This poset is differential. (We make no claims about the elegance, relative or absolute, of this description.)

In fact, we will prove the following more detailed result: P arises from the process described above by applying standard reflections and crown reflections in a particular combination. Beginning with the one-element poset P_0 with element $\emptyset$, we apply the standard reflection four times. We then apply the crown reflection at every level above the element 21^{n+1}, which will always be covered by exactly three elements, x_1, x_2, x_3, such that the new element z_i has the same name as the element y_i which it replaces for $i = 1, 2, 3$. (That is, a reflection would give us the relations $x_i \lessdot y_i$, but for our crown reflection we take $x_i \lessdot y_j$ for $j \neq i$. Note that this is independent of which of the covers we happen to call x_1, etc.)

Figure 4 shows the first six ranks of this poset. Several edges are labeled, showing the first application of each of the rules.

Proof. It is clear from the definition that P has a minimum element, $\emptyset$, that it is locally finite and graded, and that the n-th rank consists of all elements of P whose entries sum to n.

Let P_k be the differential-looking poset whose elements are those of the first k ranks of P and whose order relations are the same as the

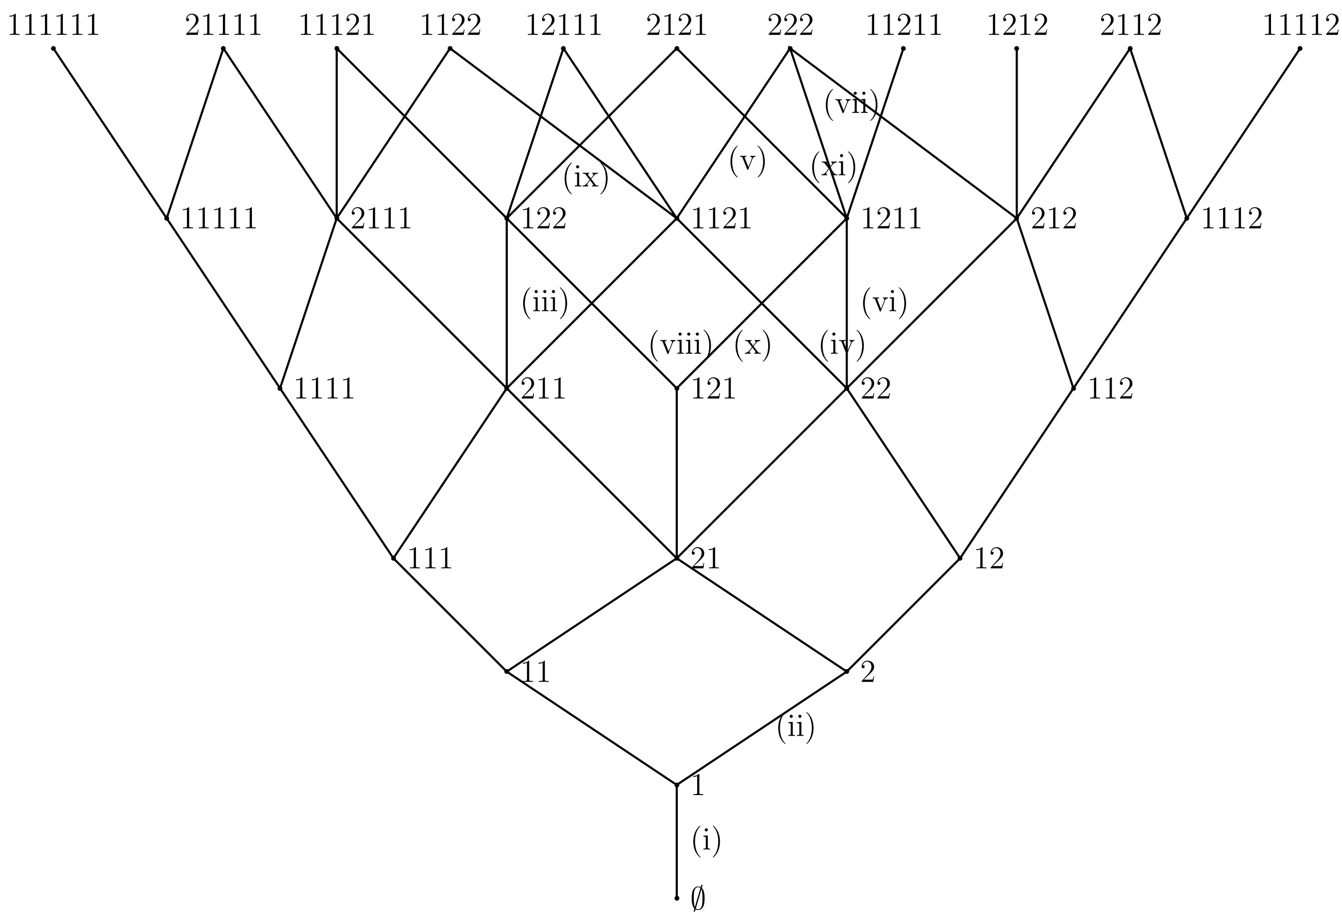

FIGURE 4. A new differential poset. Some edges are labeled with the rules that were used to generate them.

order relations between those elements as members of P. In order to prove our result, it is enough to show that P_{k+1} is identical to the poset Q_{k+1} which arises from applying the appropriate transformation (reflection or crown reflection) to P_k.

It is certainly true that $Q_k = P_k$ for $k \leq 4$ since only rules (i) and (ii) apply to elements of P of rank four or less and on these elements they describe exactly the same relations as we see in the Fibonacci-Young lattice and so exactly the relations that arise from reflection. (The higher-numbered rules and the exceptions to rules (i) and (ii) apply when one of the elements is of rank at least five.)

For all $k \geq 4$, taking covers in P_k we have

$$C^+(21^{k-3}) = \{21^{k-2}, 1^{k-4}22, 1^{k-2}21\}$$

and so in particular it is possible to apply a crown reflection above 21^{k-1}. Thus Q_{k+1} is indeed well-defined. Consider any element x of rank k in Q_{k+1}. Then by the nature of the crown reflection we have the following possibilities.

Case 1: $x = 21^{k-2}$. Then one can verify (by inverting and checking each possible rule; in this case, rules (i) and (ii) apply) that $C^-(x) = \{21^{k-3}, 1^{k-1}\}$ and so after crowning above 21^{k-3} we get

$$C^+(x) = \{1^{k-3}22, 1^{k-2}21, 21^{k-1}\}.$$

These are exactly the set of covers of x in P_{k+1}, where they are a result of rules (i) and (iii).

Case 2: $x = 1^{k-4}22$. If $k = 4$, $C^-(x) = \{21, 12\}$ (rule (i) twice) and so crown reflecting, $C^+(x) = \{1211, 1121, 212\}$, which match what we want in P_{k+1} (rules (i), (iv) and (vi)). If $k > 4$ then $C^-(x) = \{21^{k-3}, 1^{k-4}21\}$ (rules (iii), (viii)) and by crown reflection, $C^+(x) = \{121^{k-2}, 1^{k-2}21, 21^{k-4}21\}$. This matches the covers in P_{k+1} which follow from rules (iv), (vi) and (ix).

Case 3: $x = 1^{k-3}21$. If $k = 4$ then $C^-(x) = \{21\}$ and so $C^+(x) = \{1211, 122\}$, which matches P_{k+1} (rules (viii), (x)). If $k > 4$ then $C^-(x) = \{21^{k-1}, 1^{k-5}22\}$ (rules (iii), (iv)) and by reflection, $C^+(x) = \{21^{k-5}22, 1^{k-3}22, 121^{k-2}\}$ which matches the set of covers of x in P_{k+1} (rules (v), (viii), (x)).

Case 4: x is not one of these three elements. Then we have by rule (i) that $x \lessdot 1x$. In addition, rules (iv) through (xi) come in pairs for which it is evident that if $w \lessdot x$ as a consequence of one such rule then $x \lessdot 2w$ as a consequence of the paired rule. For example, (vi) and (vii) tell us that for any $m, n \geq 0$ we have

$$2^m 1^n 22 \lessdot 2^m 12111^n \lessdot 2^{m+1} 1^n 22,$$

and so on. In addition, the exceptions to rules (i) and (ii) correspond in such a way that if w does not fall into the exceptions and $w \lessdot x$ as a consequence of one of these rules then x will not fall into the exceptions either and $x \lessdot 2w$ just as in the case of **FY**. Thus, Q_{k+1} coincides with P_{k+1} around those elements not involved in the crown part of a crown reflection.

These cases are comprehensive. It follows that P can be formed through the sequence of reflections and crown reflections described and so in particular P is differential, as desired. $\square$

We now leave our discussion of the broader family of differential posets and turn our focus to one particular member, Young's lattice.

4. Young's lattice

We defined Young's lattice in terms of integer partitions, certain lists of positive integers. Rather than dealing with partitions directly, it is often useful to consider instead a certain picture associated with the given partition. The *Young diagram* of a partition $\lambda = (\lambda_1, \ldots, \lambda_k)$ is a figure composed of $|\lambda| = \lambda_1 + \ldots + \lambda_k$ boxes, upper-left justified, so that the number of boxes in the i-th row is λ_i. For example, the Young diagrams of the partitions $(5, 4, 3)$ and $(6, 2, 2, 2)$ are given below.

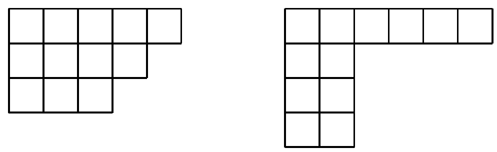

The ordering of Young's lattice is such that for two partitions λ, μ, we have $\lambda \leq \mu$ if and only if $\lambda_i \leq \mu_i$ for all i. In terms of Young diagrams, this says that $\lambda \leq \mu$ if and only if the Young diagram of λ fits entirely inside the Young diagram of μ when the two diagrams are arranged so their upper-left corners coincide. Then it is clear, for example, that $\lambda \lessdot \mu$ if and only if the Young diagram of μ can be formed by adding exactly one square to the Young diagram of λ so that the result is still a valid Young diagram. In terms of integer sequences, $\lambda \lessdot \mu$ if and only if μ arises by adding 1 to exactly one of the λ_i so that the result is still properly ordered, or by adding a new entry of 1 to λ.

Theorem 4.1 ([5], Corollary 1.4)**.** *Young's lattice is differential.*

Proof. Suppose for some partition λ that

$$\lambda_1 = \ldots = \lambda_{a_1} > \lambda_{a_1+1} = \ldots = \lambda_{a_1+a_2} > \lambda_{a_1+a_2+1} = \ldots$$

for some positive integers $a_1, a_2, \ldots$. If we let $i_1 = 1$, $i_{n+1} = i_n + a_n$, then we can write

$$\lambda = (\lambda_{i_1}^{a_1}, \lambda_{i_2}^{a_2}, \ldots, \lambda_{i_k}^{a_k}), a_i \geq 1, \lambda_{i_1} > \lambda_{i_2} > \ldots > \lambda_{i_k} > 0$$

where exponentiation represents multiplicities. Then λ is covered by exactly the $k + 1$ partitions

$$(\lambda_{i_1} + 1, \lambda_{i_1}^{a_1-1}, \lambda_{i_2}^{a_2}, \ldots, \lambda_{i_k}^{a_k}), \quad (\lambda_{i_1}^{a_1}, \lambda_{i_2} + 1, \lambda_{i_2}^{a_2-1}, \ldots, \lambda_{i_k}^{a_k}), \ldots$$
$$\ldots, (\lambda_{i_1}^{a_1}, \lambda_{i_2}^{a_2}, \ldots, \lambda_{i_k}^{a_k}, 1).$$

Conversely, λ covers the k partitions

$$(\lambda_{i_1}^{a_1-1}, \lambda_{i_1} - 1, \lambda_{i_2}^{a_2}, \ldots, \lambda_{i_k}^{a_k}), \quad (\lambda_{i_1}^{a_1}, \lambda_{i_2}^{a_2-1}, \lambda_{i_2} - 1, \ldots, \lambda_{i_k}^{a_k}), \ldots$$
$$\ldots, (\lambda_{i_1}^{a_1}, \lambda_{i_2}^{a_2}, \ldots, \lambda_{i_k}^{a_k-1}, \lambda_{i_k} - 1)$$

where in the last case if $\lambda_{i_k} = 1$ we simply omit the last entry. This establishes (D3) for Young's lattice.

(D1) follows from the nature of the cover relation: the empty partition $\emptyset$ is the minimal element of $\mathbf{Y}$, $\mathbf{Y}$ is locally finite because we can move from any partition to any partition larger than it by repeatedly adding 1s, and $\mathbf{Y}$ is graded, with the n-th rank consisting of all partitions of n.

Finally, if two partitions λ and μ have common cover, there is some partition ν such that λ and μ arise from adding 1 at two different locations in ν. Then adding both 1s to ν gives a partition which is easily seen to cover both λ and μ, so (D2) holds as well and $\mathbf{Y}$ is differential. $\square$

A *lower order ideal* in a locally finite poset P with a minimal element is a finite set I of elements of P such that $y \in I$, $x \in P$ and $x < y$ imply that $x \in I$. These order ideals have a natural poset structure through inclusion as sets, and the poset of order ideals of P so ordered is denoted $J(P)$. A poset is a *distributive lattice* if it is isomorphic to $J(P)$ for some poset P. That distributive lattices are lattices is straightforward: least upper bounds are set unions and greatest lower bounds are set intersections. The term distributive lattice arises because these lattices are exactly those for which the lattice operations glb and lub distribute over each other, a result we will not prove. For example, the Boolean lattice B_S discussed earlier is the poset of lower order ideals of the poset whose underlying set is S and which has no order relations among its elements.

We mentioned earlier that Young's lattice and the Fibonacci-Young lattice are the only two known differential lattices, and others may exist. In the case of differential distributive lattices, though, an actual result is known.

Proposition 4.2 (Folklore). *Young's lattice is a distributive lattice.*

Proof. Consider the poset P on $\mathbf{N} \times \mathbf{N}$ with $(a, b) \leq (c, d)$ in P if and only if $a \leq c$ and $b \leq d$ as members of $\mathbf{N}$ with the conventional ordering. If we arrange the Hasse diagram of P so that $(0, 0)$ is in the upper-left corner, it becomes immediately clear that order ideals correspond to Young diagrams, so that $J(P) \cong \mathbf{Y}$ and $\mathbf{Y}$ is a distributive lattice. $\square$

Theorem 4.3 ([7], p.180). *Any differential distributive lattice is isomorphic to Young's lattice.*

Proof. Suppose L is a differential distributive lattice. Thus, in particular, L is a graded, locally finite distributive lattice with a least element

$\hat{0}$ such that every element is covered by one more element than it covers. Call an element $x \in L$ *irreducible* if there do not exist $y, z \in L$ such that $y, z < x$ and $x = \mathrm{lub}(y, z)$, and let S_k be the set (in fact, poset) of irreducible elements of L of rank k or less. We can extend this set to a sublattice L_k of L by taking the least upper bounds of every subset of S_k. Every element of L of rank k or less is either in S_k or can be written as the least upper bound of smaller elements of L. Thus, every element of L of rank k or less is also an element of L_k. It follows that L is the union of the L_k; if we can show that the L_k are unique (up to isomorphism), it will follow that L is the unique differential distributive lattice and so $L \cong \mathbf{Y}$.

$S_0 = \{\hat{0}\}$ so L_0 is the single-element poset, and thus unique. Suppose that L_k can be constructed in a unique manner. Consider any element $x \in L_k$ of rank k, and suppose x covers n elements and is covered by m elements in L_k. The elements x covers in L_k are the same as the elements x covers in L, so x has $n - m + 1$ additional covers in L. These new covers are of rank $k + 1$ and so must be irreducible, since otherwise we would be able to express them as the least upper bound of elements of rank k or less and so as the least upper bound of elements of L_k, which would also make them elements of L_k, a contradiction. These new elements, together with S_k, are therefore included in S_{k+1}. In addition, this must exhaust all of S_{k+1}, since any other irreducibles would by necessity be of rank $k+1$ and so would cover some element of L of rank k, which is impossible. Thus S_{k+1} can be uniquely constructed. The structure of L_{k+1} follows uniquely from the structure of S_{k+1}, so by induction L is uniquely determined and $\mathbf{Y}$ is the only differential distributive lattice. $\square$

It is interesting to note that we could substitute for (D3) in this proof the fact that whenever x covers n vertices, it is covered by $f(n)$ vertices for any function f, and the result (that there exists at most one distributive lattice with this property, although it would certainly not be Young's lattice) would still be valid.

A *semi-standard Young tableau*, or SSYT, is a filling of the boxes of a Young diagram with positive integers so that rows weakly increase from left to right while columns strictly increase from top to bottom. The *shape* $\mathrm{sh}(T)$ of a SSYT T is the shape of its associated Young diagram and the *type* $\mathrm{type}(T)$ is the sequence of multiplicities of its entries. Thus, if

$$T = \begin{array}{|c|c|c|c|c|}\hline 1 & 1 & 1 & 3 & 4 \\ \hline 2 & 2 & 4 & 5 & 5 \\ \hline 3 & 4 & 5 \\ \cline{1-3}\end{array}$$

we have $\mathrm{sh}(T) = (5, 5, 3)$ and $\mathrm{type}(T) = (3, 2, 2, 3, 3)$.

The basic unit of the RSK algorithm is an insertion algorithm on tableaux. Given a SSYT T and an integer a_1, the tableau $T \leftarrow a_1$ is constructed as follows:

(0) To begin, $n = 1$ and a_1 is "active."
(1) If the active number is at least as large as all elements of the n-th row (and especially if the n-th row is empty), insert it as the rightmost element in this row and stop. Otherwise, proceed to step (2).
(2) Let a_2 be the smallest member of the n-th row which is larger than a_1. Replace a_2 with a_1 in the n-th row, make a_2 active, increment n and go back to step (1).

It is clearly the case that the resulting diagram is a SSYT. In this SSYT, we may consider those entries which were at some point active; they are called the *insertion path* of a_1 in T. For example, the image below shows one example of an insertion with the insertion path in bold.

$$T = \begin{array}{|c|c|c|c|c|} \hline 1 & 1 & 1 & 3 & 4 \\ \hline 2 & 2 & 4 & 5 & 5 \\ \hline 3 & 4 & 5 \\ \cline{1-3} \end{array} \qquad T \leftarrow 2 = \begin{array}{|c|c|c|c|c|} \hline 1 & 1 & 1 & \mathbf{2} & 4 \\ \hline 2 & 2 & \mathbf{3} & 5 & 5 \\ \hline 3 & 4 & \mathbf{4} \\ \cline{1-3} \mathbf{5} \\ \cline{1-1} \end{array}$$

It is straightforward to show that insertion paths move only down and to the left, never to the right, and that if $a \leq b$ then the insertion path of b in $(T \leftarrow a) \leftarrow b$ lies strictly to the right and not below the insertion path of a.

The RSK algorithm consists of a particular sequence of repeated insertions. Consider a finite matrix $A = (a_{ij})_{i,j \geq 1}$ of natural numbers and let $n = \sum_{i,j} a_{ij}$. We associate to A its $2 \times n$ *scheme matrix* as follows: the columns of the scheme are $\binom{i}{j}$, appearing a_{ij} times, sorted first along the top row and then along the bottom row. For example, the scheme of the matrix

$$\begin{pmatrix} 1 & 0 & 1 & 2 \\ 0 & 0 & 2 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}$$

is

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 2 & 2 & 3 \\ 1 & 3 & 4 & 4 & 3 & 3 & 2 \end{pmatrix}.$$

Since we may reconstruct the matrix A given its scheme we consider the two to be interchangeable.

Given a matrix A with a $2 \times n$ scheme matrix whose t-th column is $\binom{i_t}{j_t}$, let $P(0) = Q(0) = \emptyset$ be two empty SSYTs and define recursively

$P(t+1) = P(t) \leftarrow j_{t+1}$. $Q(t+1)$ is formed by augmenting $Q(t)$ by one box so that $Q(t)$ and $P(t)$ are of the same shape and giving this box the entry i_{t+1}. The output of the algorithm is the final pair $(P, Q) = (P(n), Q(n))$, and we write $A \xrightarrow{\text{RSK}} (P, Q)$. P is called the *insertion tableau* and Q is the *recording tableau*.

Theorem 4.4 ([8], Theorem 7.11.5)**.** *This algorithm is a bijection between scheme matrices of the desired form, or equivalently finite matrices over* $\mathbf{N}$*, and pairs of SSYT of the same shape.*

Proof. We have two things to show: first, that Q is in fact an SSYT, and second, that this algorithm is invertible and so a bijection.

Note that the elements of Q are added in weakly increasing order, and each new entry is along the outermost edge of the diagram at the time it is inserted. Thus the rows and columns of Q must be weakly increasing. That the columns of Q are strictly increasing follows from our observation above that if $b \geq a$ then the insertion path of b in $(T \leftarrow a) \leftarrow b$ is to the right and cannot pass below the insertion path of a: the copies of i in Q were added consecutively, and later additions are associated with larger insertions into P. Thus, a copy of i added to Q must occur in a column to the right of the preceding copy of i, so the copies of i in Q are in different columns. Then the columns of Q strictly increase and Q is indeed an SSYT.

Now, we must show how to invert the algorithm. Take any two SSYTs (P, Q) of the same shape $\lambda \vdash n$ and consider the right-most instance of the largest entry in Q. The order of insertion together with the observations on why Q is a SSYT show that this must have been the last entry added to $Q = Q(n)$ and so the corresponding entry in $P = P(n)$ was the last to be inserted. Reversing the insertion algorithm is straightforward: an active number must have been bumped from the previous row by the right-most entry smaller than it, so we can trace the insertion path upwards until we discover exactly which number was inserted into $P(n-1)$ to get P. The numbers removed from P and Q in this fashion give us the rightmost column of the scheme matrix, and we can repeat the process with the new matrices $P(n-1)$ and $Q(n-1)$. After n iterations we will have completely constructed the scheme matrix, thus inverting RSK. $\square$

One particular interesting class of SSYTs is the set of *standard Young tableaux*, or SYTs. An SYT of shape $\lambda \vdash n$ is an SSYT in which the diagram of λ is filled with the integers from 1 to n, each appearing exactly once. There is a particularly nice relationship between SYTs and Young's lattice. Suppose we have an SYT T_n of shape $\lambda \vdash n$ and we remove the box labeled n. This gives a new SYT T_{n-1} of shape $\mu \vdash n-1$. Then $\mu \lessdot \lambda$ as elements of $\mathbf{Y}$. We can perform the same operation to μ to find a new element of $\mathbf{Y}$ covered by μ, and so on. Thus, an SYT of shape λ has a natural interpretation as a Hasse walk in $\mathbf{Y}$ joining λ to $\emptyset$.

It is natural to ask at this point what RSK does when applied to a pair of SYTs.

Proposition 4.5 ([8], Theorem 7.11.5)**.** *The RSK algorithm gives a bijection between pairs (P, Q) of SYT of the same shape and permutation matrices.*

Proof. By definition, P and Q are SYTs if and only if $\text{type}(P) = \text{type}(Q) = (1, 1, \ldots, 1)$. The entries of P and Q are the elements of the bottom and top rows of the scheme matrix, respectively, and from the construction of the scheme we see that i appears in the top row exactly $\sum_j a_{ij}$ times while j appears in the bottom row $\sum_i a_{ij}$ times. Thus, if A is the matrix associated via RSK with (P, Q) then we must have $\sum_i a_{ik} = \sum_j a_{kj} = 1$ for all $k = 1, \ldots, n$, so A must have exactly one 1 in each row and column and 0s elsewhere. Conversely, if A is a permutation matrix then P and Q will have type $(1, \ldots, 1)$ and so be SYTs. Since RSK is invertible, this is a bijection. $\square$

Denote by f^λ the number of SYT of shape λ. From the comments above, this is equal to $f_{\mathbf{Y}}(\lambda)$, the number of minimal Hasse walks from $\emptyset$ to λ in $\mathbf{Y}$. Then the number of pairs of SYTs of shape λ is $(f^\lambda)^2$. Proposition 4.5 gives a bijection between $n \times n$ permutations matrices and pairs of SYT of the same shape with n boxes, so as a direct consequence

$$\sum_{\lambda \vdash n} (f^\lambda)^2 = n!.$$

In other words, we now have a bijective proof of Proposition 2.2 in the case $P = \mathbf{Y}$.[3]

Surprisingly, this algorithm, which appears to rely quite heavily on a particular representation of the elements of $\mathbf{Y}$, can be extended to

[3]This result has an entirely different significance to algebraists. The conjugacy classes of the symmetric group S_n are classified by their cycle structure, so in fact they are indexed by integer partitions. Thus, the irreducible representations of S_n can be similarly indexed, and the sum of the squares of their dimensions is the size of the group, $n!$ ([4], Prop. 1.10.1). These dimensions turn out to be exactly the f^λ, so our theorem is in fact a special case of a more general theorem on group representations. For much more on this topic, including a proof that the f^λ are the dimensions of the irreducible representations, see [8], chapter 7, or [4].

all differential posets. Here we describe the general algorithm due to Fomin ([2]). We will not prove that this alorithm applies in the general setting, but we will use it to establish two results on Hasse walks in differential posets and will show that it coincides with RSK when we take $\mathbf{Y}$ to be our differential poset.

First, fix a differential poset P and choose $n \in \mathbf{N}$. At each element $y \in P$, choose a bijection $\phi_y : C^-(y) \cup \{y\} \to C^+(y)$. (There is in general no canonical way to make this choice, and the resulting bijection will depend on the choices of the maps ϕ_y.) Now, we construct a *growth diagram* for P: for a given $\pi \in S_n$, draw an $n \times n$ table and place an X in the i, j-th box of the table if and only if $\pi(i) = j$. (Here we are using Cartesian coordinates: the box in the lower-left corner is $(1, 1)$ and the box in the lower-right is $(n, 1)$.) Next, label each vertex along the left and bottom edges with our minimal element $\hat{0}$, as in Figure 5. Given a box whose three corners are labeled as shown below, we fill it in according to the following rules:

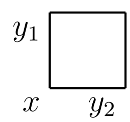

(1) If there is no X in the box and $x = y_1$ (or $x = y_2$) then we label the upper-right corner with y_2 (respectively y_1).
(2) If there is no X in the box and $x \lessdot y_1 = y_2$ then we label the upper-right corner with $\phi_{y_1}(x)$
(3) If there is no X in the box and $x \lessdot y_1 \neq y_2 \gtrdot x$ then we label the upper-right corner with the unique z covering y_1 and y_2.
(4) If there is an X in the box then $x = y_1 = y_2$ and we label the upper-right corner with $\phi_x(x)$.

Note in particular that the last rule describes the only situation which can possibly arise: if a given box has an X in it then since π is a permutation none of the boxes directly to its left or below it have an X. This ensures that only rules (1) and (2) were called upon in the labeling of the edges of our box and so the three corners of the box will have the same label.

We use these rules to label every vertex of the grid, starting in the lower-left corner and working our way up and right. Reading along the right and top edges of the resulting diagram gives us two saturated chains connecting $\hat{0}$ to the element in the upper-right corner of the diagram. This sends permutations to pairs of Hasse walks in P. Given an element $x \in P$ and two paths between $\hat{0}$ and x, we can invert the algorithm by labeling the top and right edges with the paths and moving down and left. Given the functions ϕ_y, the labels on the three

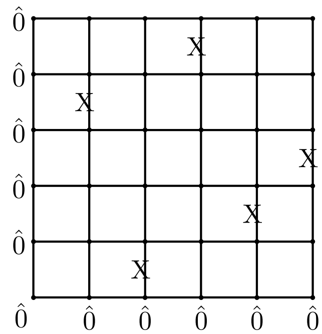

FIGURE 5. The initial setup of a geometric RSK generalization for the permutation $\pi = 41523$.

right and upper corners of a box uniquely determine the label in the lower-left and establish whether or not the box contains an X, allowing us to reconstruct the permutation which corresponds to the paths. That this actually represents a bijection is not obvious, and we do not provide the proof.

This algorithm immediately extends the bijective proof of Proposition 2.2 to all differential posets. It also has another important consequence: given two minimal Hasse walks connecting $\hat{0}$ to some element x in a differential poset P, swtiching the order of the walks reflects the growth diagram across its diagonal. Then if the original output of the algorithm was the permutation π, the new output will be the inverse permutation π^{-1}. In terms of the insertion version of RSK for $\mathbf{Y}$, this says that for permutation matrices, if $A \xrightarrow{\text{RSK}} (P, Q)$ then $A^\top \xrightarrow{\text{RSK}} (Q, P)$. This fact (which extends to all matrices) is a far-from-obvious property of RSK.

As an additional consequence in the general case, we have that the two paths are the same if and only if the resulting permutation satisfies $\pi = \pi^{-1}$, i.e. is an involution. Since the number of pairs of paths such that the two paths are the same is simply the number of paths, we have that

$$\sum_x f_P(x) = \#\{\pi \in S_n \mid \pi^2 = 1\}$$

where the sum is over all x of rank n in P.

In order to show that the two algorithms correspond in the case $P = \mathbf{Y}$, it is necessary to choose the functions ϕ_λ for each partition λ. The following seems to be a natural choice:

- $\phi_\lambda(\lambda)$ is the partition that arises from adding 1 to the largest part of λ.

- If λ arises from adding 1 to part μ_i of μ then $\phi_\lambda(\mu)$ is the partition which arises from adding 1 to λ_{i+1}.

Theorem 4.6 ([8], Theorem 7.13.5)**.** *In case $P = \mathbf{Y}$ with ϕ_λ as above, the growth diagram version of RSK corresponds to the standard version after identification of Hasse walks with SYTs.*

Proof. Coordinatize the vertices of our diagram so that each box has the same coordinates as its upper-right corner. Then at each vertex (i, j) of our complete growth diagram is a partition $\lambda(i, j)$, with $\lambda(0, j) = \lambda(i, 0) = \emptyset$. By the nature of the construction rules for the growth diagram, we have for any j that

$$\lambda(i, 0) \leq \lambda(i, 1) \leq \ldots \leq \lambda(i, n)$$

and for each j either

$$\lambda(i, j) = \lambda(i, j + 1) \text{ or } \lambda(i, j) \lessdot \lambda(i, j + 1).$$

Then replace each partition with a SYT $T(i, j)$ such that $\mathrm{sh}(T(i, j)) = \lambda(i, j)$ and if $\lambda(i, j) \lessdot \lambda(i, j + 1)$ then $T(i, j + 1)$ arises from $T(i, j)$ by filling the new box with $j + 1$.

For fixed i, j, let $(i_1, j_1), \ldots, (i_k, j_k)$ be the coordinates of the X-ed boxes below and to the left of (i, j), so $i_l \leq i$ and $j_l \leq j$, ordered so that $i_1 < \ldots < i_k$. We claim that

$$T(i, j) = (\ldots((\emptyset \leftarrow j_1) \leftarrow j_2) \leftarrow \ldots) \leftarrow j_k.$$

If $i = 0$ or $j = 0$ this claim is obvious. Suppose that the statement is true of $T(i-1, j-1), T(i-1, j)$ and $T(i, j-1)$. There are several cases to verify for a full inductive proof, each corresponding to a local rules by which we built the growth diagram, and we provide proofs for three of them.

(1) If box (i, j) does not contain an X and $\lambda(i-1, j-1) = \lambda(i-1, j)$ then $\lambda(i, j - 1) = \lambda(i, j)$ and there are no Xs directly to the left of the (i, j)-th box. Since there are no such Xs, our claim predicts that $T(i, j)$ is formed by the same sequence of insertions as $T(i, j - 1)$. But $\lambda(i, j - 1) = \lambda(i, j)$ implies that $T(i, j) = T(i, j - 1)$, so this case is done.
(2) If box (i, j) does not contain an X and $\lambda(i-1, j-1) = \lambda(i, j-1)$ then $\lambda(i-1, j) = \lambda(i, j)$ and there are no Xs directly below the (i, j)th box. Since there are no such Xs, our claim predicts that $T(i, j)$ should be formed by the same sequence of insertions as $T(i - 1, j)$. We also have by hypothesis that $T(i - 1, j - 1) = T(i, j-1)$, so $T(i-1, j)$ and $T(i, j)$ arise from $T(i-1, j-1)$ and $T(i, j - 1)$ by identical additions of a box. Thus $T(i - 1, j) = T(i, j)$ and this case is done as well.

(3) If box (i, j) contains an X then no box in the same row or column does so $T(i-1, j-1) = T(i-1, j) = T(i, j-1)$ arise according to hypothesis from insertions of a collection of $j_l < j$. Since $\lambda(i, j)$ is formed from $\lambda(i, j-1)$ by addition of a box to the first row, $T(i, j)$ arises from $T(i, j-1)$ by placing j at the end of the first row. Since $j > j_l$, this is equivalent to $T(i, j-1) \leftarrow j$, as desired.

From this claim it is immediately clear that

$$T(i, n) = (\ldots(\emptyset \leftarrow \pi(1)) \leftarrow \ldots) \leftarrow \pi(i)$$

and so in particular the chain $\lambda(0, n) \lessdot \lambda(1, n) \lessdot \ldots \lessdot \lambda(n, n)$ corresponds to the recording tableau while $T(n, n)$ is the insertion tableau associated with the permutation matrix for π. □

5. Acknowledgements

We would like to thank Professor Thomas Lam, under whose guidance this thesis was written, for all his help from start to finish. We would also like to thank Yan Zhang for his edits and suggestions.

References

[1] S. Fomin: Duality of graded graphs, *J. Algebraic Combin.* **3** (1994), 357–404.
[2] S. Fomin: Schensted algorithms for dual graded graphs, *J. Algebraic Combin.* **4** (1995), 5–45.
[3] K. Killpatrick: Evacuation and a geometric construction for Fibonacci tableaux, *Journal of Combinatorial Theory, Series A*, **110** (2005), 337–351
[4] B. Sagan, *The Symmetric Group*, Wadsworth and Brooks, 1991.
[5] R. Stanley: Differential posets, *J. Amer. Math. Soc.* **1** (1988), 919–961.
[6] R. Stanley: "Variations on differential posets," *Invariant Theory and Tableaux*, D. Stanton, Ed., Springer- Verlag, 1990, 145–165.
[7] R. Stanley, *Enumerative Combinatorics, Vol 1*, Cambridge, 1997.
[8] R. Stanley, *Enumerative Combinatorics, Vol 2*, Cambridge, 1999.