Fields Institute Communications
Volume **00**, 0000

Semigroup and Ring Theoretical Methods in Probability

Kenneth S. Brown
Department of Mathematics
Cornell University
Ithaca, NY 14853
kbrown@math.cornell.edu

Introduction

This is an expanded version of a series of four lectures designed to show algebraists how ring theoretical methods can be used to analyze an interesting family of finite Markov chains. The chains happen to be random walks on semigroups, and the analysis is based on a study of the associated semigroup algebras. The paper is divided into four sections, which correspond roughly to the four lectures:

1. Examples.
2. Semigroup interpretation.
3. Algebraic analysis.
4. Connections with Solomon's descent algebra.

Two appendices give a self-contained treatment of bands (also called idempotent semigroups) and their linear representations.

The work described here grew out of my joint work with Persi Diaconis [5]. It is a pleasure to thank him for many stimulating conversations.

1 Examples

A reference for all results stated in this section is Brown–Diaconis [5] and further references cited there. We begin with two concrete examples that will be used throughout the paper.

1.1 The Tsetlin library. Imagine a deck of n cards, labeled $1, 2, \dots, n$. Repeatedly shuffle the deck by picking a card at random and moving it to the top. "Random" here refers to an underlying probability distribution on the set of labels; thus one picks the card labeled i with some probability w_i, where $w_i > 0$ and $\sum_i w_i = 1$. This process, called the *Tsetlin library* or *random-to-top shuffle*, is a Markov chain with $n!$ states, corresponding to the possible orderings of the deck. One can identify the states with the elements of the symmetric group S_n, where a permutation $\sigma \in S_n$ corresponds to the ordering $\sigma(1), \sigma(2), \dots, \sigma(n)$ of the cards.

2000 *Mathematics Subject Classification.* Primary 60J10; Secondary 20M99, 60B99.
The author was supported in part by NSF Grant #0071428.

Warning Although this process is a Markov chain whose state space is a group, it is *not* a random walk on the group except in the special case where the weights w_i are uniform. [In a random walk on a group, one repeatedly multiplies by random group elements, chosen from some fixed probability measure on the group. See Saloff-Coste [14] for a recent survey of random walks on groups.] We will see later, however, that it can be interpreted as a random walk on a semigroup.

The random-to-top shuffle may seem foolish as a method of shuffling, but there are other interpretations that are more reasonable:

- Think of a pile of n file folders, where file i is used with frequency w_i. Each time a file is used, it is replaced on top of the pile. With this interpretation, the Markov chain describes a familiar procedure for keeping the most commonly used files near the top. It is a *self-organizing* scheme; one does not need to know the frequencies w_i in advance.
- Instead of files, one can visualize a row of books on a long shelf; each time a book is used, it is replaced at the left. (This interpretation explains the word "library" in the name "Tsetlin library".)
- Finally, one can think of data stored in a computer as a linked list. Each time an item is accessed, it is moved to the front of the list. The most commonly accessed items will then tend to be near the front of the list, thus reducing access time. This "move-to-front" list-management scheme has been extensively studied by computer scientists.

1.2 The inverse riffle shuffle. In the ordinary riffle shuffle, one divides a deck of cards roughly in half and then interleaves the top half with the bottom. There is a precise model for this, the "Gilbert–Shannon–Reeds" model, according to which the number of cards in the top half has a binomial distribution, and all interleavings are equally likely. The inverse riffle shuffle reverses the process: Pick a random set of cards and move them to the top, keeping them in the same relative order. All 2^n subsets are equally likely to be chosen. The inverse riffle shuffle has the same mixing properties as the riffle shuffle, and it happens to fit the framework of random walks on semigroups that I will be discussing.

There is a natural generalization of the inverse riffle shuffle, in which one assigns arbitrary weights to the subsets instead of making them equally likely. This includes the Tsetlin library as the special case where the weights are concentrated on the 1-element subsets.

1.3 Questions. The behavior of a Markov chain over time is governed by the powers of the transition matrix K. The latter is a matrix whose rows and columns are indexed by the states of the chain (permutations in our examples), where the (σ, τ)-entry $K(\sigma, \tau)$ is the chance of moving from σ to τ in one step. The l-th power K^l of this matrix gives the l-step transition probabilities. So the basic question is how K^l behaves as $l \to \infty$.

A fundamental theorem of Markov chain theory [9] implies that $K^l(\sigma, \tau)$ tends to a limit $\pi(\tau)$, independent of the starting state σ:

$$K^l(\sigma, \tau) \to \pi(\tau) \quad \text{as } l \to \infty.$$

(The theorem requires a mild regularity condition, which is satisfied by all examples treated in this paper.) Here π is a probability distribution on the set of states, and $\pi(\tau)$ is the long-term probability of being in the state τ. The distribution π

is called the *stationary distribution*; it can be characterized algebraically as the unique probability distribution satisfying the "equilibrium equation"

$$\sum_{\sigma} \pi(\sigma) K(\sigma, \tau) = \pi(\tau) \tag{1.1}$$

for all τ. This says that if we pick a state according to π and take a step in the chain, then the state is still described by π. If we view π as a row vector, Equation (1.1) can be written as $\pi K = \pi$, i.e., π is a left eigenvector for K with eigenvalue 1.

The following questions are of interest for any Markov chain:

(a) What is π?
(b) What are the eigenvalues of K?
(c) How fast does $K^l \to \pi$?

Question (c) is often phrased as, "How many times do you have to shuffle a deck of cards to mix it up?" The main interest in Question (b) is that one expects it to be useful in answering (c), though in practice this connection is not always as clear as one would hope.

1.4 Answers. I will give complete answers for the Tsetlin library, in order to give the flavor of the results; the answers for the inverse riffle shuffle (and its generalization where weights are assigned to the subsets) are similar.

(a) The stationary distribution π is the distribution of sampling without replacement from the weights w_i. In other words, for any ordering σ, the stationary probability $\pi(\sigma)$ is the same as the probability of getting σ by the following process: Pick a card according to the weights w_i and put it on top; now pick from the remaining cards, with probabilities proportional to their weights, and put it next; keep going in this way until all the cards have been picked.

(b) The transition matrix K is diagonalizable, and its eigenvalues are the partial sums of the weights w_i. More precisely, for each subset $X \subseteq \{1, 2, \ldots, n\}$, there is an eigenvalue

$$\lambda_X = \sum_{i \in X} w_i$$

of multiplicity

$$m_X = d_{n-|X|},$$

where d_k, the kth derangement number, is the number of fixed-point-free permutations of k elements. We can also describe m_X as the number of permutations in S_n with fixed-point set X. This description makes it obvious that $\sum_X m_X = n!$, so that we have the right number of eigenvalues. Note that $d_1 = 0$, so λ_X does not actually occur as an eigenvalue if $|X| = n - 1$. Note also that, for particular choices of weights, the eigenvalues λ_X for different X might coincide; one then has to add the corresponding numbers m_X to get the true multiplicity of the eigenvalue.

(c) Let K^l_σ be the distribution of the chain started at σ, after l steps; it is given by the σ-row of the lth power of K. Then it satisfies

$$\|K^l_\sigma - \pi\|_{\text{T.V.}} \le \sum_{|X|=n-2} \lambda^l_X, \tag{1.2}$$

where the left-hand side is the "total-variation" distance. (This is one of the standard measures of distance to stationarity; we do not need its precise definition here.) With uniform weights $w_i = 1/n$, for example, one can deduce that the distance to stationarity is small after about $l = n \log n$ steps. This is roughly 200 for an ordinary deck of cards with $n = 52$. The vague statement that the distance to stationarity is "small" can be made quite precise. There is in fact a striking cutoff phenomenon, in which the distance is close to 1 (the maximum possible value) until l reaches $n \log n$ and then quickly decreases to 0. See Diaconis [6] for a discussion of this.

There are a number of striking features of these results. First, there is no reason to expect, a priori, that K should be diagonalizable or have real eigenvalues. The standard hypothesis under which these properties hold is called "reversibility" in the Markov chain literature, and the Tsetlin library is not reversible. Secondly, one would not expect such simple formulas for the eigenvalues in terms of the data that determine the matrix K. Finally, although one does expect the "second-largest" eigenvalues λ_X to play a role in determining the convergence rate, there is no general theory that would predict an estimate as simple as (1.2).

The main point of these lectures is that many of the results stated above for the Tsetlin library hold for an interesting family of Markov chains, which turn out to be random walks on semigroups. We will focus primarily on the linear algebra results as in (b), since this is where the algebraic methods are the most useful.

2 Semigroup interpretation

2.1 Hyperplane face semigroups. The semigroups underlying the examples above come from an unexpected source: the theory of hyperplane arrangements. Let $\mathcal{A}$ be a finite set of hyperplanes in a real vector space V. For simplicity, we assume the hyperplanes are linear (i.e., they pass through the origin), though the theory works for affine hyperplanes also. We also assume, without loss of generality, that $\bigcap_{H \in \mathcal{A}} H = \{0\}$. (If this fails we can pass to the quotient of V by the intersection of the hyperplanes.) The dimension of V is called the *rank* of the arrangement. Figures 1 and 2 show two simple examples, the first of rank 2 and the second of rank 3. In the latter we have drawn the intersections of the hyperplanes with the unit sphere, viewed from the north pole; the dotted circle is the equator, which is not part of the arrangement. Only the northern hemisphere is visible in the picture.

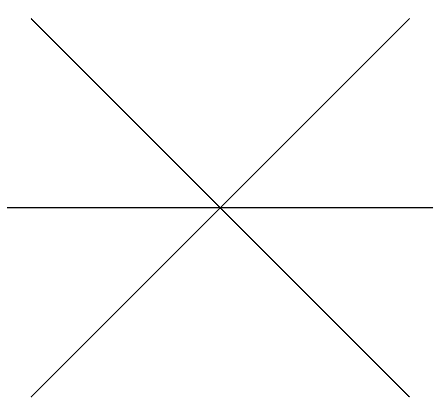

Figure 1 Three lines in $\mathbb{R}^2$

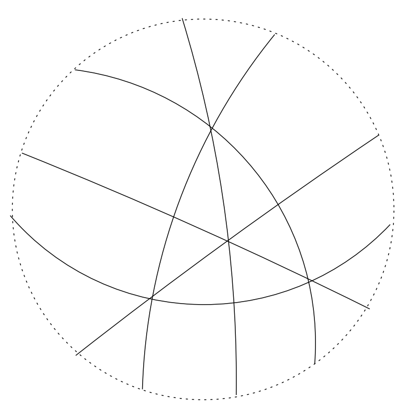

Figure 2 Six planes in $\mathbb{R}^3$

The hyperplanes in $\mathcal{A}$ divide V into regions called chambers. In Figure 1 there are 6 chambers, which are sectors. In Figure 2 there are 24 chambers, which are triangular cones that intersect the sphere in spherical triangles. Six of these are visible in the picture (they are in the northern hemisphere), six are opposite these and not shown in the picture, and 12 are partially visible.

The chambers are polyhedral sets and hence have faces. We denote by $\mathcal{F}$ the set of all (open) chambers and their (relatively open) faces of all dimensions. The elements of $\mathcal{F}$ will simply be called *faces*. In particular, we view a chamber as a face of itself. In Figure 1, for example, there are 13 faces: 6 open sectors, 6 open halflines, and the origin. As illustrated in Figure 2, the faces are in 1–1 correspondence with their intersections with the unit sphere in V; these form the cells of a cell decomposition of the sphere. (Here we need to include the empty cell, which corresponds to the face consisting of the origin.) If we draw the unit circle in Figure 1, we see that the cell complex is combinatorially equivalent to a hexagon. In Figure 2, the cell complex is the barycentric subdivision of the boundary of a tetrahedron.

Remark 2.1 It is often useful to characterize a face A by specifying, for each hyperplane in $\mathcal{A}$, which side of the hyperplane A lies on. There are three possibilities for a given hyperplane H: A can be strictly on one side, strictly on the other side, or contained in H.

All of this geometry is fairly transparent. What is less obvious is that there is a natural way to multiply faces, so that $\mathcal{F}$ becomes a semigroup. This product was introduced by Bland in the early 1970s in connection with linear programming, and it eventually led to one approach to the theory of oriented matroids; see [3]. Tits [19] discovered the product independently (in the setting of Coxeter complexes and buildings), although he phrased his version of the theory in terms of "projection operators" rather than products. Here are two descriptions of the product:

1. Given two faces $A, B \in \mathcal{F}$, start in A and move toward B; one immediately enters a face, which is defined to be the product AB.
2. For each hyperplane $H \in \mathcal{A}$, the product AB lies on the same side of H as A unless $A \subseteq H$, in which case AB lies on the same side as B.

See Figure 3 for a simple example, where A and B are halflines and AB turns out to be a chamber. For a second example, let B' be the halfline opposite A in the same figure; then $AB' = A$.

One can easily check from the second description that the associative law holds:

$$A(BC) = (AB)C. \tag{2.1}$$

In fact, the triple product, with either way of associating, can be characterized by the property that for each $H \in \mathcal{A}$ it lies on the same side of H as A unless $A \subseteq H$, in which case it lies on the same side of H as B unless $B \subseteq H$, in which case it lies on the same side of H as C. So $\mathcal{F}$ is indeed a semigroup. It has an identity, given by the trivial face $\{0\}$. We call $\mathcal{F}$ the *hyperplane face semigroup* associated with the arrangement $\mathcal{A}$.

Remark 2.2 In Tits's treatment cited above, the product AB is called the projection of B on A and is denoted $\operatorname{proj}_A B$. The associative law in the context of Coxeter complexes appears in Tits's appendix to Solomon's paper [17], where it

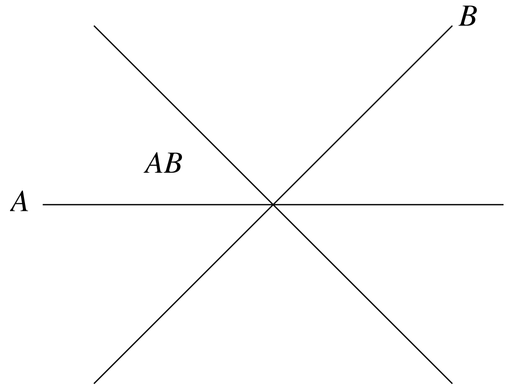

Figure 3 The product of two halflines

takes the form

$$\mathrm{proj}_A(\mathrm{proj}_B\, C) = \mathrm{proj}_{\mathrm{proj}_A\, B}\, C. \tag{2.2}$$

A comparison of Equations (2.1) and (2.2) shows the usefulness of thinking in terms of products rather than projection operators.

Remark 2.3 Note that hyperplane face semigroups are idempotent semigroups, i.e., $A^2 = A$ for all $A \in \mathcal{F}$. In particular, they are very far from being groups.

We are mainly interested in products where the second factor is a chamber. Given a face A and a chamber C, the product AC is always a chamber. It has A as a face and can be characterized as the nearest chamber to C having A as a face. Here "nearest" is defined in terms of the number of hyperplanes in $\mathcal{A}$ separating two chambers.

Remark 2.4 If, on the other hand, the first factor is a chamber, then the product is very boring. In fact, we have $CA = C$ for all A if C is a chamber.

2.2 Example: The braid arrangement. The *braid arrangement* in $\mathbb{R}^n$ consists of the $\binom{n}{2}$ hyperplanes $x_i = x_j$ $(i \neq j)$. Their intersection is the 1-dimensional subspace $x_1 = x_2 = \cdots = x_n$, which is not trivial. So, as explained at the beginning of Section 2.1, we should view the arrangement as living in an $(n-1)$-dimensional quotient of $\mathbb{R}^n$. It is also known as the *reflection arrangement of type* A_{n-1}. Figures 1 and 2 above show the cases $n = 3$ and $n = 4$, respectively.

There are $n!$ chambers, corresponding to the possible orderings of the coordinates. Thus for each permutation $\sigma \in S_n$ there is a chamber given by

$$x_{\sigma(1)} > x_{\sigma(2)} > \cdots > x_{\sigma(n)}.$$

Faces are gotten by changing zero or more inequalities to equalities. They correspond to ordered partitions $B = (B_1, B_2, \ldots, B_k)$ of the set $\{1, 2, \ldots, n\}$. Here the blocks B_i form a set partition in the usual sense, and their order matters. The partition B encodes the ordering of the coordinates and which coordinates are equal to one another. For example, the face

$$x_1 = x_3 > x_2 > x_4 = x_6 > x_5$$

corresponds to the ordered partition $(\{1,3\},\{2\},\{4,6\},\{5\})$. Notice that the chambers can be identified with the ordered partitions into singletons.

The product of faces has a simple interpretation in terms of ordered partitions: Take (nonempty) intersections of the blocks in lexicographic order; more precisely, if $B=(B_1,\dots,B_l)$ and $C=(C_1,\dots,C_m)$, then

$$BC=(B_1\cap C_1,\dots,B_1\cap C_m,\dots,B_l\cap C_1,\dots,B_l\cap C_m)\hat{},$$

where the hat means "delete empty intersections". More briefly, BC is obtained by using C to refine B. (To check that this description of the product is correct, one should use the second definition of the product in Section 2.1.) Observe that we do indeed have $B^2=B$ for all ordered partitions B, as claimed in Remark 2.3.

In the important special case where the second factor C is a chamber (an ordering of $1,2,\dots,n$), the product is a new ordering that maintains the ordering of the blocks $B_1,B_2,\dots,B_k$ and, within each block, uses the ordering given by C. We can think of this concretely in terms of card shuffling. Imagine a deck of cards numbered $1,2,\dots,n$. A chamber corresponds to an ordering of the deck, and multiplying by $B=(B_1,B_2,\dots,B_k)$ performs the following "B-shuffle": Remove the cards with labels in B_1 and place them on top (keeping them in their original order), remove the cards with labels in B_2 and place them next, and so on.

Remark 2.5 Although chambers correspond to permutations, their product in the hyperplane face semigroup has nothing to do with the usual product of permutations; see Remark 2.4.

2.3 Hyperplane chamber walks. Since $\mathcal{F}$ is a semigroup, we can run a random walk on it: Start at some element of $\mathcal{F}$ and repeatedly left-multiply by randomly chosen elements of $\mathcal{F}$. Here "randomly chosen" refers to some given probability distribution $\{w_A\}_{A\in\mathcal{F}}$ on $\mathcal{F}$. The fruitful idea of studying random walks on hyperplane face semigroups is due to Bidigare, Hanlon, and Rockmore [2]. For readers familiar with random walks on groups but not semigroups, we mention one difference: Semigroups can have proper ideals (nonempty proper subsets closed under multiplication by arbitrary semigroup elements). If we start the random walk in an ideal, it stays there, so we get a Markov chain on the ideal. In the present setting, the natural ideal to use is the ideal $\mathcal{C}$ consisting of chambers. Taking $\mathcal{A}$ to be the braid arrangement, for example, we see that any choice of probability distribution on the faces gives a Markov chain on S_n (identified with the set of chambers).

We recover the Tsetlin library by concentrating the probabilities on the 2-block ordered partitions in which the first block is a singleton. And we recover the inverse riffle shuffle by assigning probability $1/2^n$ to each 2-block ordered partition and probability $2/2^n$ to the 1-block partition. [Intuitively, the 1-block partition should be thought of as occurring twice, once from each of the "weak ordered partitions" $(\{1,2,\dots,n\},\emptyset)$ and $(\emptyset,\{1,2,\dots,n\})$.] The generalized inverse riffle shuffle mentioned at the end of Section 1.2 is obtained by putting arbitrary weights on the 1-block and 2-block partitions.

3 Algebraic analysis

3.1 Random walks and semigroup algebras. Our analysis of random walks on semigroups will be based on the structure of the associated semigroup algebras. Let S be a finite semigroup, C a left ideal, and $\{w_x\}_{x\in S}$ a probability distribution on S. As above, we then get a (left) random walk on C: If we are at $c \in C$, we choose $x \in S$ with probability w_x and move to xc. The transition matrix is given by

$$K(c,d) = \sum_{xc=d} w_x. \tag{3.1}$$

Note, again, how this differs from the case where the semigroup is a group; in that case, the sum degenerates to a single term, corresponding to $x = dc^{-1}$.

Form the semigroup algebra $\mathbb{R}S$ consisting of real linear combinations of elements of S, and encode the probability distribution in the element

$$w = \sum_{x \in S} w_x x.$$

Then the transition matrix K arises algebraically when one considers the operator "left-multiplication by w" acting on the ideal $\mathbb{R}C \subseteq \mathbb{R}S$. Indeed, for any $a = \sum_c a_c c$ in $\mathbb{R}C$ we have

$$wa = \sum_x w_x x \sum_c a_c c = \sum_d \left(\sum_{\substack{x,c \\ xc=d}} w_x a_c \right) d = \sum_d \left(\sum_c a_c K(c,d) \right) d,$$

where the last equality follows from (3.1). Thus left multiplication by w acting on $\mathbb{R}C$ corresponds to right multiplication by K acting on row vectors $(a_c)_{c \in C}$.

If we are interested in algebraic properties of K, then, such as eigenvalues or diagonalizability, we should study multiplication by w acting on $\mathbb{R}C$. Similarly, the powers of K correspond to the powers of w, so a key object to analyze is $\mathbb{R}[w]$, the algebra generated by w, i.e., the linear span of the powers w^l. The semigroups for which I can give the most complete analysis are called *left-regular bands*, although some of the results are valid in greater generality.

3.2 The semigroup algebra of a band. A *band*, also called an *idempotent semigroup*, is a semigroup in which $x^2 = x$ for all x. A band is called *left-regular* if $xyx = xy$ for all x, y. It is easy to check that hyperplane face semigroups are always left-regular bands (see the discussion of the triple product following Equation (2.1) in Section 2.1). The first fact we need is that the representation theory of a finite band (whether left-regular or not) is particularly simple:

Theorem 3.1 *For any field k and any finite band S, every irreducible kS-module is 1-dimensional. In other words, the quotient of kS by its radical is isomorphic to a product of copies of k.*

More briefly, the theorem says that kS is an *elementary algebra*. This result is well-known to experts, but we give a self-contained proof in Appendix B motivated by the geometry of hyperplane arrangements. This approach is due to Bidigare [1] for hyperplane face semigroups and was generalized to left-regular bands in [4]; the further generalization to arbitrary bands given in Appendix B is not much more difficult. The remainder of this section consists of a sketch of the proof for the case of hyperplane face semigroups.

Let $\mathcal{A}$ be a hyperplane arrangement and let $\mathcal{F}$ be its face semigroup as in Section 2.1. The *intersection lattice* associated with $\mathcal{A}$ is the poset $\mathcal{L}$ consisting of all intersections $\bigcap_{H \in \mathcal{A}'} H$, where $\mathcal{A}' \subseteq \mathcal{A}$. We order $\mathcal{L}$ by inclusion. It is a lattice, with greatest lower bound given by intersection. The least upper bound $U \vee V$ of two elements $U, V \in \mathcal{L}$ is the intersection of all hyperplanes in $\mathcal{A}$ containing both

U and V. [Warning: In much of the literature on hyperplane arrangements, $\mathcal{L}$ is ordered by the opposite of the inclusion relation.]

There is a *support* map $\operatorname{supp}\colon \mathcal{F} \twoheadrightarrow \mathcal{L}$, where $\operatorname{supp} A$ for $A \in \mathcal{F}$ is the linear span of A. The support of A can also be described as the intersection of all hyperplanes in $\mathcal{A}$ containing A. It satisfies

$$\operatorname{supp}(AB) = \operatorname{supp} A \vee \operatorname{supp} B \tag{3.2}$$

and

$$\operatorname{supp} A \supseteq \operatorname{supp} B \iff A = AB. \tag{3.3}$$

Note that Equation (3.2) says that the support map is a semigroup homomorphism, where $\mathcal{L}$ is viewed as a (commutative) semigroup under the $\vee$-operation.

Example 3.2 Let $\mathcal{A}$ be the braid arrangement, and identify $\mathcal{F}$ with the set of ordered partitions of $\{1, 2, \ldots, n\}$ as in Section 2.2. Then $\mathcal{L}$ can be identified with the set of unordered set partitions, and the support map simply forgets the ordering.

In view of Equation (3.2), the support map induces a homomorphism

$$\operatorname{supp}\colon k\mathcal{F} \twoheadrightarrow k\mathcal{L}$$

of semigroup algebras. Theorem 3.1 for $\mathcal{F}$ now follows from:

Theorem 3.3

1. *The algebra $k\mathcal{L}$ is isomorphic to a product of copies of k. In particular, it is semisimple.*
2. *The kernel of* supp: $k\mathcal{F} \twoheadrightarrow k\mathcal{L}$ *is nilpotent and hence is the radical of $k\mathcal{F}$.*

The first assertion is well-known, and its proof is recalled in Section B.1 below. The second assertion, due to Bidigare [1], is proved by a straightforward calculation based on (3.2) and (3.3). The calculation is repeated in Brown [4] in the more general context of left-regular bands. See also Section B.1 below for a further generalization to arbitrary bands, which takes slightly more work.

Remark 3.4 The proof yields explicit formulas for the nonzero characters of the semigroup $\mathcal{F}$ (see [4] or Theorem B.2 in Section B.2 below.) There is a character $\chi_X\colon \mathcal{F} \to k$ for each $X \in \mathcal{L}$, which takes the value 1 on all faces contained in X and 0 on all other faces. More briefly,

$$\chi_X(A) = 1_{A \subseteq X} \tag{3.4}$$

for $A \in \mathcal{F}$.

Remark 3.5 Algebraically-oriented readers might not find the theory of hyperplane arrangements to be convincing motivation for the approach to Theorem 3.1 sketched here. Such readers might prefer the motivation provided by the theory of free bands (Section A.1), where the support map again arises in a very natural way.

3.3 Eigenvalues. Recall that the transition matrix of the Tsetlin library turns out, surprisingly, to have real eigenvalues, which are simply the partial sums of the weights. Using Theorem 3.1, we can now give a clear explanation of this and a generalization of it to random walks on arbitrary finite bands. For simplicity, we confine ourselves here to the hyperplane chamber walk, where the result was first proved by Bidigare, Hanlon, and Rockmore [2]. The generalization is treated in Section B.3 below.

Let $\mathcal{A}$, $\mathcal{F}$, and $\mathcal{L}$ be as above. For each $X \in \mathcal{L}$, let

$$\mathcal{A}_X = \{H \in \mathcal{A} : H \supseteq X\}.$$

Then $\mathcal{A}_X$ is a hyperplane arrangement in its own right, and we denote by c_X the number of chambers of $\mathcal{A}_X$. Define a family of integers m_X ($X \in \mathcal{L}$) by the system of equations

$$c_X = \sum_{Y \supseteq X} m_Y \tag{3.5}$$

for each $X \in \mathcal{L}$. Zaslavsky [20] gave the explicit formula $m_X = |\mu(X, V)|$, where μ is the Möbius function of $\mathcal{L}$ and V is the ambient vector space; but this does not hold for general bands, so I prefer the implicit definition in Equation (3.5). In practice one can often find the solution to (3.5) by inspection.

Theorem 3.6 *Let $\{w_A\}$ be a probability distribution on $\mathcal{F}$ and let K be the transition matrix of the corresponding chamber walk, i.e., K is the $\mathcal{C} \times \mathcal{C}$ matrix given by*

$$K(C, D) = \sum_{AC=D} w_A.$$

Then K has an eigenvalue

$$\lambda_X = \sum_{A \subseteq X} w_A$$

for each $X \in \mathcal{L}$, with multiplicity m_X.

Sketch of proof Recall from Section 3.1 that we are interested in the eigenvalues of multiplication by $w = \sum_{A \in \mathcal{F}} w_A A$ acting on $\mathbb{R}\mathcal{C}$. Choose a composition series for $\mathbb{R}\mathcal{C}$ as an $\mathbb{R}\mathcal{F}$-module. Theorem 3.1 and Remark 3.4 yield a triangular matrix representation for multiplication by w, with diagonal entries of the form $\chi_X(w)$; the latter occurs m'_X times, where m'_X is the number of composition factors given by χ_X. Now

$$\begin{aligned} \chi_X(w) &= \chi_X\left(\sum_A w_A A\right) \\ &= \sum_A w_A \chi_X(A) \\ &= \sum_{A \subseteq X} w_A \\ &= \lambda_X, \end{aligned}$$

where the third equality follows from (3.4). To complete the proof, one gives a counting argument to show that the multiplicities m'_X satisfy the system (3.5), so $m'_X = m_X$. See [4] or the proof of Theorem B.3 below for details. □

3.4 Diagonalizability. Brown and Diaconis [5] showed that the transition matrix of the hyperplane chamber walk is diagonalizable. The proof used topology. The result was generalized to left-regular bands by Brown [4], using purely algebraic methods. Here I just want to give the general principles. These seem to be well-known in some circles (in connection with perturbation theory or linear recurrence relations, for instance), but I think they deserve to be more widely known. I begin with the simplest version.

Fix an $n \times n$ matrix A with complex entries, and form the generating function of the powers of A:

$$f(t) = \sum_{m \geq 0} A^m t^m = \frac{1}{I - tA}\,,$$

where the right-hand side denotes the inverse of $I - tA$. The series converges for small complex t and defines a matrix-valued holomorphic function of t, initially defined in a neighborhood of the origin. It is convenient to make a change of variable and introduce

$$g(z) = (1/z)f(1/z) = \frac{1}{zI - A}\,,$$

initially defined in a neighborhood of infinity.

Proposition 3.7 *The function $g(z)$ is a rational function, with poles precisely at the eigenvalues of A. The matrix A is diagonalizable if and only if the poles are all simple. In this case, the partial-fractions decomposition of g is given by*

$$g(z) = \sum_{\lambda} \frac{E_\lambda}{z - \lambda}\,,$$

where λ ranges over the eigenvalues and E_λ is the projection onto the λ-eigenspace.

"Projection" here refers to the eigenspace decomposition. The proof of the proposition is an exercise based on the Jordan decomposition of A; it is written out in [4, Section 8.1].

Now the fancier version: Let k be a field and let R be a finite-dimensional k-algebra with identity, generated by a single element a. Form

$$f(t) = \sum_{m \geq 0} a^m t^m = \frac{1}{1 - at}$$

and

$$g(z) = (1/z)f(1/z) = \frac{1}{z1_R - a}\,.$$

We can view these as holomorphic functions if k is a subfield of the complex numbers, or we can just work with them as formal power series.

Proposition 3.8 *R is split semisimple (isomorphic to a product of copies of k) if and only if $g(z)$ has the form*

$$g(z) = \sum_i \frac{e_i}{z - \lambda_i} \tag{3.6}$$

for distinct $\lambda_i \in k$, where the e_i are nonzero elements of R. In this case the e_i are the primitive idempotents of R, and

$$a = \sum_i \lambda_i e_i.$$

The proof is similar to the proof of the previous proposition and is written out in [4, Section 8.2].

Using this I was able to prove:

Theorem 3.9 *Let S be a finite left-regular band with identity, and let $w = \sum_{x \in S} w_x x$ be an element of $\mathbb{R}S$ with $w_x \geq 0$ for all x. Then the subalgebra $\mathbb{R}[w]$ generated by w is split semisimple. In particular, multiplication by w is diagonalizable, so the transition matrix of any random walk on S is diagonalizable.*

The proof is based on a mindless computation of the generating function $f(t) = \sum_{m\geq 0} w^m t^m$. This is complicated, but left-regularity makes it possible to reduce $g(z)$ to something that visibly has the form (3.6). See [4, Sections 8.3 and 8.4]. The proof yields explicit formulas for the primitive idempotents of $\mathbb{R}[w]$; for generic w, these lift the primitive idempotents of $\mathbb{R}S$ mod its radical and may be of purely algebraic interest in connection with Solomon's descent algebra (Section 4 below).

3.5 Stationary distribution and convergence rate. I mentioned in Section 1.3 three questions that one can ask about a finite Markov chain, but I have mostly discussed only one of them (eigenvalues). The other two, involving stationary distribution and convergence rate, also have nice answers for random walks on bands. But the proofs are not algebraic, so they do not fit into these lectures. Suffice it to say that there are general results in the spirit of those stated in Section 1.4 for the Tsetlin library. A proof for hyperplane face semigroups is given in Brown–Diaconis [5], and virtually the same proof goes through for arbitrary finite bands. There is an explicit statement in [4] for the left-regular case.

One hope I have is that the explicit diagonalization mentioned in Section 3.4 will lead to more precise results about convergence rate. So far, however, I have not been able to get past the complexity of the formulas for the primitive idempotents.

4 Connections with Solomon's descent algebra

The goal of this final section is to apply some of the ideas of this paper to algebra rather than probability. Specifically, I would like to show the usefulness of hyperplane face semigroups for understanding Solomon's descent algebra. The inspiration for this again comes from the work of Bidigare [1]. I will be brief and concise, but I hope that the examples will make the discussion comprehensible. Full details can be found in [4, Section 9].

Solomon [17] introduced, for any finite Coxeter group W, an interesting subalgebra of the group algebra kW. Here k is an arbitrary commutative ring. For $W = S_n$, the definition of this subalgebra can be phrased in terms of descent sets of permutations, so Solomon's subalgebra has come to be known as the "descent algebra".

4.1 Notation. Throughout this section W denotes a finite Coxeter group or, equivalently, a finite (real) reflection group. Thus W is a finite group of orthogonal transformations of a real inner-product space V, and W is generated by reflections s_H with respect to hyperplanes. The set $\mathcal{A}$ of hyperplanes H such that $s_H \in W$ is called the *reflection arrangement* associated with W. It is a W-invariant set of hyperplanes. As before, we will assume (without loss of generality) that $\bigcap_{H\in\mathcal{A}} H = \{0\}$. This is equivalent to requiring that the fixed-point set of W in V be trivial.

Let Σ be the face semigroup. It turns out that all the faces are simplicial cones, so the decomposition of the sphere induced by $\mathcal{A}$ is simplicial. Thus Σ can be identified with the set of simplices of a simplicial complex that triangulates the sphere in V. By abuse of language we will often say that Σ "is" a simplicial complex and we will speak of the vertices of Σ and so on. The chambers of the arrangement are the top-dimensional simplices of Σ. When we represent a hyperplane arrangement by a spherical picture as in Figure 2 above, what we are really drawing is a picture of the simplicial complex Σ.

Since $\mathcal{A}$ is W-invariant, there is an action of W on Σ that is simplicial and preserves the semigroup structure. The action is simply-transitive on the set $\mathcal{C}$ of chambers, so if we choose a "fundamental chamber" C, we get a bijection between $\mathcal{C}$ and W; a group element w corresponds to the chamber wC.

4.2 Example. Let W be the symmetric group S_n acting by permuting the coordinates on $V = \mathbb{R}^n/\{x_1 = x_2 = \cdots = x_n\}$. The corresponding reflection arrangement of rank $n-1$ is the braid arrangement (Section 2.2). There is one vertex of Σ for each proper nonempty subset of $\{1, 2, \ldots, n\}$, and the simplices are the chains $E_1 < E_2 < \cdots < E_l$ of such subsets. (Combinatorially, Σ is the barycentric subdivision of the boundary of an $(n-1)$-simplex.) See Figure 4. To

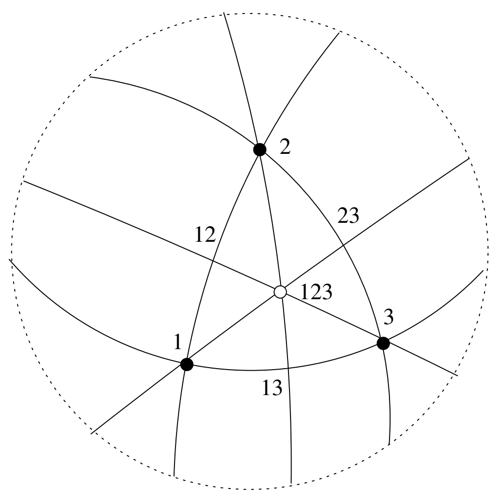

Figure 4 $W = S_4$: Vertices are subsets of $\{1, 2, 3, 4\}$

relate the present description of Σ to our earlier description of the faces of the braid arrangement, note that an ordered partition $(B_1, B_2, \ldots, B_k)$ yields a chain $B_1 < B_1 \cup B_2 < \cdots < B_1 \cup B_2 \cup \cdots \cup B_{k-1}$.

Finally, if we take the fundamental chamber to be the chain

$$\{1\} < \{1, 2\} < \cdots < \{1, 2, \ldots, n-1\},$$

then we get a bijection between chambers and permutations that makes a permutation w correspond to the chamber

$$\{w(1)\} < \{w(1), w(2)\} < \cdots < \{w(1), w(2), \ldots, w(n-1)\}.$$

This is the same bijection we described in Section 2.2 in different language. It is illustrated in Figure 5, in which have followed the convention of identifying a permutation w with the list of numbers $w(1)\, w(2) \cdots w(n)$.

4.3 Types of simplices. A basic fact about reflection arrangements is that the vertices of Σ can be assigned *types.* The set I of types has size $r = \operatorname{rank} \mathcal{A}$, and every chamber has exactly one vertex of each type. In the example above, for instance, a vertex is a nonempty proper subset of $\{1, 2, \ldots, n\}$, and the type of a vertex is its cardinality; thus $I = \{1, 2, \ldots, n-1\}$. The vertices in Figures 4 and 5 have been drawn in three different "colors" to illustrate the types.

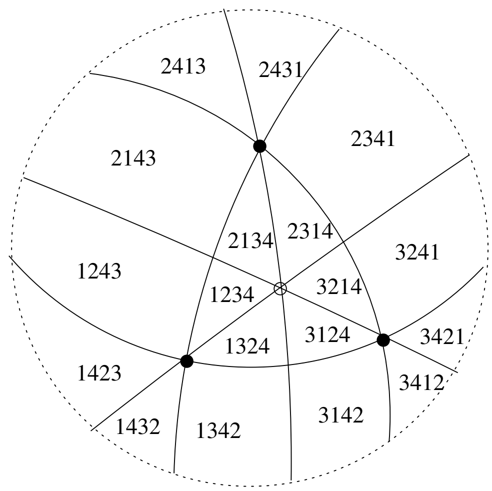

Figure 5 $W = S_4$: Chambers correspond to permutations

Since vertices have types, so do arbitrary simplices; the type of a simplex is the subset of I consisting of the types of its vertices. For example, every chamber has type I, while every *panel* (codimension 1 simplex) has type $I - \{i\}$ for some i.

4.4 Adjacency and galleries. Two chambers are *adjacent* if they have a common panel. More precisely, two chambers are i-adjacent ($i \in I$) if they have the same panel of type $I - \{i\}$. See Figure 6. We include the degenerate case: Any chamber is i-adjacent to itself. A *gallery* is a sequence of chambers $C_0, C_1, \ldots, C_l$ such that any two consecutive chambers are adjacent. It is a *geodesic* gallery if there is no shorter gallery joining C_0 and C_l. The length l is then the *distance* between C_0 and C_l; it is equal to the number of hyperplanes in $\mathcal{A}$ that separate C_0 from C_l.

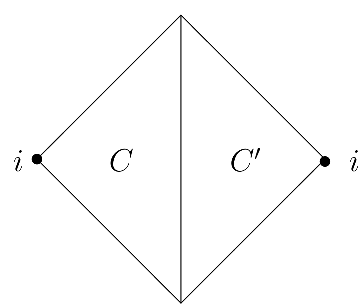

Figure 6 C and C' are i-adjacent chambers

Since chambers correspond to elements of W (once a fundamental chamber has been chosen), we can apply this terminology to elements of W. For $W = S_n$, let us identify a permutation w with a list of numbers as above. Then w is adjacent to w' if and only if w' is obtained from w by swapping two consecutive elements of the list. The type of the adjacency, which is an element of $I = \{1, 2, \ldots, n-1\}$, indicates the position of the swap. For example, 31425 is 2-adjacent to 34125 and is 4-adjacent to 31452. The distance from a permutation w to the identity is the number of inversions of w. To construct a geodesic gallery from w to the identity, we can start with any $i \in \{1, 2, \ldots, n-1\}$ at which w has a *descent* (i.e., $w(i) > w(i+1)$) and swap $w(i)$ and $w(i+1)$. Repeating the process leads to the identity in the required number of steps. The reader should try a few examples in Figure 5.

4.5 Descent sets. Returning to the general case, the discussion above motivates the following definition. Given chambers C, C', we define the *descent set* of C' with respect to C, denoted $\operatorname{des}(C, C')$, to be the set of types $i \in I$ such that there is a geodesic gallery from C' to C starting with an i-adjacency. Equivalently,

$$i \in \operatorname{des}(C, C') \iff C \text{ and } C' \text{ lie on opposite sides of } \operatorname{supp} A,$$

where A is the panel of C' of type $I - \{i\}$. See Figure 7. If C is the fundamental chamber and C' is the chamber wC corresponding to w, we will also write $\operatorname{des}(w)$ for $\operatorname{des}(C, C')$, and we will call this the descent set of w. In view of the discussion at the end of the previous subsection, this agrees with the usual descent set of a permutation if $W = S_n$.

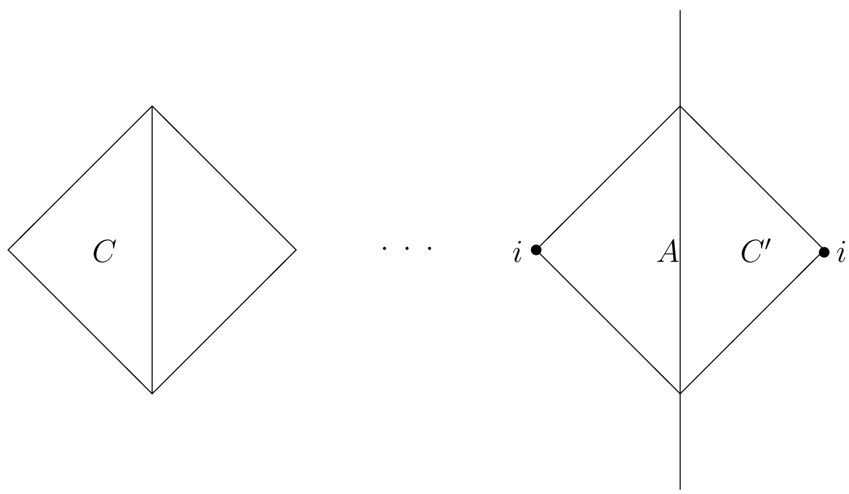

Figure 7 The type i is in $\operatorname{des}(C, C')$

4.6 Semigroup interpretation. Recall that Σ is a semigroup. We can use the product to characterize descent sets. For motivation, note that in Figure 7 $AC \neq C'$ and the type of A does not contain $\operatorname{des}(C, C')$. With a little careful checking, one can turn this observation into a proof of the following:

Proposition 4.1 *Given chambers C, C', there is a smallest face A of C' such that $AC = C'$. The descent set* $\operatorname{des}(C, C')$ *is the type of that face A.*

4.7 Connection with the h-vector. The h-vector is a standard and important object in algebraic combinatorics. We show here that the h-vector of Σ encodes the number of chambers (or elements of W) with a given descent set.

For $J \subseteq I$, let f_J be the number of simplices of Σ of type J. Define a family of integers h_J by the system of equations

$$f_J = \sum_{K \subseteq J} h_K. \tag{4.1}$$

Proposition 4.2 *Fix a fundamental chamber C. For any $J \subseteq I$, h_J is the number of chambers C' such that* $\mathrm{des}(C, C') = J$*. Equivalently, h_J is the number of $w \in W$ such that* $\mathrm{des}(w) = J$*.*

Sketch of proof Let Σ_J be the set of simplices of type J. There is a 1–1 correspondence between Σ_J and the set of chambers C' such that $\mathrm{des}(C, C') \subseteq J$; it is given by $A \mapsto AC$ for $A \in \Sigma_J$. Counting such chambers C' according to their descent sets, we obtain

$$f_J = \sum_{K \subseteq J} h'_K,$$

where h'_K is the number of C' with $\mathrm{des}(C, C') = K$. Comparing this system of equations with (4.1), we conclude that $h'_J = h_J$ for all J. $\square$

4.8 Invariants in the semigroup algebra. Fix a commutative ring k. Then W acts on the semigroup algebra $k\Sigma$, and we can form the ring of invariants $R = (k\Sigma)^W$. As a k-module it is free with one basis element

$$\sigma_J = \sum_{A \in \Sigma_J} A$$

for each subset $J \subseteq I$. Motivated by the discussion of the h-vector, we introduce a new basis $(\tau_J)_{J \subseteq I}$ by

$$\sigma_J = \sum_{K \subseteq J} \tau_K.$$

Theorem 4.3 *There is a 1–1 anti-homomorphism of k-algebras*

$$R = (k\Sigma)^W \hookrightarrow kW$$

such that $\tau_J \mapsto \sum_{\mathrm{des}(w) = J} w$.

The image of this map is Solomon's descent algebra, which is therefore anti-isomorphic to R. Solomon described it in a different but equivalent way, and he had to work hard to prove that it is an algebra. The semigroup point of view gives this immediately.

Corollary 4.4 (Bidigare) *Solomon's descent algebra is anti-isomorphic to $(k\Sigma)^W$.* $\square$

Sketch of proof of Theorem 4.3 Let $\mathcal{C}$ be the set of chambers. Then $k\mathcal{C}$ is a $k\Sigma$-module, hence an R-module, and the R-action on $k\mathcal{C}$ commutes with the W-action on $k\mathcal{C}$. This yields

$$R \to \mathrm{End}_{kW}(k\mathcal{C}) \cong \mathrm{End}_{kW}(kW) \cong (kW)^{\mathrm{op}}.$$

One checks that $\sigma_J \mapsto \sum_{\mathrm{des}(w) \subseteq J} w$, which implies that $\tau_J \mapsto \sum_{\mathrm{des}(w) = J} w$. $\square$

4.9 Connection with random walks. We close by coming back to the random walks that we started with. Recall that a probability distribution $\{p_A\}_{A\in\Sigma}$ gives rise to a Markov chain (left random walk) on the ideal $\mathcal{C}$ in the semigroup Σ. Having chosen a fundamental chamber, we can identify $\mathcal{C}$ with W, so we have a Markov chain on W. As I remarked in Section 1.1 in connection with the Tsetlin library, this is not generally a random walk on the group W. It *is* a random walk on W, however, if the system of weights p_A is W-invariant. In fact, one can use the ideas of the previous subsection to prove:

Theorem 4.5 *Suppose* $\{p_A\}$ *is* W*-invariant and let*

$$\mu = \sum_{w\in W} \mu_w w$$

be the element of Solomon's descent algebra corresponding to

$$p = \sum_{A\in\Sigma} p_A A$$

under the anti-isomorphism of Corollary 4.4. Then $\{\mu_w\}$ *is a probability distribution on* W*, and the left random walk on* $\mathcal{C}$ *driven by* $\{p_A\}$ *is the same as the right random walk on* W *driven by* $\{\mu_w\}$*.*

The significance of the theorem is that, in view of Section 3, we now have an interesting family of random walks on Coxeter groups that we know how to analyze.

Appendix A Bands and their support semilattices

A *band* is a semigroup in which $x^2 = x$ for all x. Bands are also called *idempotent semigroups.* The results of this appendix are not new, but the point of view is; it is motivated by the theory of hyperplane arrangements. For other treatments of some of this material, see [7, 10, 13, 15].

For motivation, we begin with an example that is easy to understand, independent of the theory of hyperplane arrangements.

A.1 The free band. For any set A, the *free band* $S = F(A)$ on A is the semigroup generated by A, subject to the relations $w^2 = w$ for every word w. It consists of equivalence classes of nonempty A-words, where the equivalence relation is generated by the rewriting rules $w^2 \to w$ and $w \to w^2$. [These rules may be applied to any subword w of a given word.] This description implies:

1. If $B \subseteq A$, then $F(B)$ can be identified with a subsemigroup of $F(A)$.
2. Every element $x \in F(A)$ has a well-defined *support*, denoted $\operatorname{supp} x$, consisting of the letters that occur in some (every) word representing x.
3. Every $x \in F(A)$ has a well-defined first letter and last letter.

Example A.1 (One generator) If A consists of a single generator a, then $F(A)$ is the one-element semigroup $\{a\}$.

Example A.2 (Two generators) If A has two elements a, b, then $F(A)$ has order 6; its elements are

$$a, b, ab, ba, aba, bab.$$

Indeed, it is clear that any element can be represented by an alternating word, which can be taken to have length ≤ 3 since an alternating word of length ≥ 4 contains *abab* or *baba*. And the alternating words of length ≤ 3 listed above represent distinct elements, since any two can be distinguished by support, first letter, or last letter.

For example, aba is the unique element whose support is $\{a, b\}$ and whose first and last letters are both a.

Example A.3 (Three generators) Let $A = \{a, b, c\}$. Elements whose support is a proper subset of A can be listed as in the previous example. There are three of length 1, six of length 2, and six of length 3. It is much less obvious that there are only finitely many elements whose support is the whole set A. But we will see in Example A.9 below that there are exactly 144 of these, all of length ≤ 8. Thus $F(a, b, c)$ is finite and of order $3 + 6 + 6 + 144 = 159$.

Let $S = F(A)$ be a free band.

Lemma A.4 *Given* $x, y \in S$, *if* $\operatorname{supp} x \supseteq \operatorname{supp} y$ *then* $x \in xyS$.

Proof We work with fixed words ξ and η representing x and y. The hypothesis implies that the last letter a of η occurs in ξ. So the word $\xi\eta$ has the form $\sigma a \tau a$ for some words σ, τ (which might be empty). Right-multiplying by τ gives $\sigma a \tau a \tau$, which is equivalent to $\sigma a \tau$, i.e., to $\xi\eta$ with the last letter deleted. Repeating this argument, we can delete the letters of η in $\xi\eta$ one by one by right-multiplication, until we are left with ξ. Thus there is a word w, possibly empty, such that $\xi\eta w$ is equivalent to ξ. If w is nonempty, this shows $x \in xyS$. Otherwise, $x = xy = (xy)^2 \in xyS$. □

We can now prove that S has two surprising properties, which we call "swallowing" and "deletion".

Proposition A.5 (Swallowing) *If* $\operatorname{supp} x \supseteq \operatorname{supp} y$, *then* $xyx = x$.

More generally:

Proposition A.6 (Deletion) *If* $\operatorname{supp} x = \operatorname{supp} z \supseteq \operatorname{supp} y$, *then* $xyz = xz$.

We say that x *swallows* y in the situation of Proposition A.5.

Proof of Proposition A.5 We have $x \in xyS$ by Lemma A.4, so x is fixed by left-multiplication by xy. □

Proof of Proposition A.6 The strategy is to introduce a second x into xyz to swallow the y:

$$\begin{aligned} xyz &= xy(zxz) && \text{because } z \text{ swallows } x \\ &= (xyzx)z \\ &= xz && \text{because } x \text{ swallows } yz. \end{aligned}$$

□

As an example of deletion, we have $abacbabcb = abacabcb$ [delete the middle letter on the left]. This illustrates the assertion in Example A.3 above that every element of $F(a, b, c)$ can be represented by a word of length ≤ 8. Note, however, that there is no rewriting rule of the form $w^2 \to w$ that can be used to shorten the word on the left; one has to first make it longer before one can shorten it.

The detailed analysis of free bands is based on the concepts of "prefix" and "suffix" of an element. Let w be a word with support B. The *prefix* of w is the smallest initial subword that uses all the letters in B. For example, the prefix of $abacbabcb$ is $abac$. It is easy to check that equivalent words have equivalent prefixes. So an element $x \in F(A)$ has a well-defined prefix $u \in F(A)$. Similarly, we can speak of the suffix of x. For example, *abacbabcb* has suffix *abcb*. One can also define the *strict prefix* and *strict suffix* of an element; these are obtained by deleting the last letter of the prefix and the first letter of the suffix. Again, it is easy to check that these are well-defined as elements of $F(A)$. For example, *abacbabcb* has strict prefix *aba* and strict suffix *bcb*.

Theorem A.7 *Two elements of $F(A)$ are equal if and only if they have the same prefix and suffix.*

Proof Let x have prefix y and suffix z. We will show that $x = yz$. Choose a fixed word representing x. If the prefix and suffix overlap in that word, then we have $x = y'uz'$, with $y'u = y$ and $uz' = z$, so $x = y'uuz' = yz$. Otherwise, x has the form $x = yuz$ with u possibly empty. Proposition A.6 allows us to delete u. □

Corollary A.8 *If A is finite, then $F(A)$ is finite. Consequently, every finitely generated band is finite.*

Proof The strict prefix and strict suffix of an element $x \in F(A)$ have smaller support than x. Arguing by induction on the size of A, we conclude that there are only finitely many possibilities for the prefix and suffix. □

Example A.9 There are 12 possible prefixes and 12 possible suffixes using three letters a, b, c, all of length ≤ 4. This explains the number 144 in Example A.3 above, as well as the assertion that every element can be represented by a word of length ≤ 8.

Remark A.10 For general A it is not hard to enumerate the possible prefixes and suffixes inductively. This leads to normal forms and a precise count of the number of elements in $F(A)$. If A has n elements, then the order of $F(A)$ is

$$\sum_{i=1}^{n} \binom{n}{i} \prod_{j=1}^{i} (i-j+1)^{2^j}.$$

The normal forms are gotten by multiplying the prefix by the suffix and deleting the repeated overlap, if any. For example, the normal form for the element with prefix *abac* and suffix *acbc* is *abacbc*.

Finally, we spell out the formal properties of the support map, in order to motivate its generalization to arbitrary bands.

Let L be the (upper) semilattice consisting of the nonempty finite subsets of A, ordered by inclusion. (An upper semilattice is a poset in which any two elements have a least upper bound.) The support map is then a surjection

$$\operatorname{supp}: S \twoheadrightarrow L$$

satisfying

$$\operatorname{supp}(xy) = \operatorname{supp} x \vee \operatorname{supp} y \tag{A.1}$$

and

$$\operatorname{supp} x \geq \operatorname{supp} y \iff x = xyx. \tag{A.2}$$

Here $\vee$ denotes the join (least upper bound) in L, given by union in this case.

Note that L is an abelian semigroup under the join operation, and Equation (A.1) simply says that supp is a semigroup homomorphism. In fact, this is the abelianization homomorphism. For suppose ϕ is a homomorphism from S to an abelian semigroup; we must show that ϕ factors through the support map. If

$\operatorname{supp} x = \operatorname{supp} y$, then $x = xyx$ and $y = yxy$, so $\phi(x) = \phi(x)\phi(y)\phi(x) = \phi(x)\phi(y)$, and similarly $\phi(y) = \phi(x)\phi(y)$. Thus $\phi(x) = \phi(y)$, as required.

A.2 The general case. Let S be an arbitrary band. We will construct a semilattice having the properties (A.1) and (A.2) above. This was already done for left-regular bands in [4]. Note that, in the left-regular case, the right side of (A.2) takes the simpler form $x = xy$, as in Section 3.2, Equation (3.3).

Theorem A.11 *For any band S, there is a semilattice L together with a surjection*

$$\operatorname{supp}\colon S \twoheadrightarrow L$$

satisfying (A.1) *and* (A.2).

Proof Define $x \succeq y \iff x = xyx$. I claim that this is a weak partial order (reflexive and transitive). Reflexivity ($x \succeq x$) is obvious. For transitivity, suppose $x = xyx$ and $y = yzy$. We wish to show $x = xzx$, which says x is fixed by left-multiplication by xz, i.e., $x \in xzS$. We have

$$x = xyx = x(yzy)x,$$

which implies $x \in Szyx$, so x is fixed by right-multiplication by zyx. Thus $x = xzyx \in xzS$, as required.

We may now pass to the quotient poset L by setting

$$x \sim y \iff x \succeq y \text{ and } y \succeq x.$$

Define

$$\operatorname{supp}\colon S \twoheadrightarrow L$$

to be the quotient map. Then (A.2) holds by definition. To prove that L has least upper bounds and that Equation (A.1) holds, we need to show

$$z \succeq x,\ z \succeq y \iff z \succeq xy. \tag{A.3}$$

It is immediate that $xy \succeq x$ and $xy \succeq y$, so we need only prove the forward implication. It is also immediate that $xy \sim yx$ for any x, y; this will be needed in the proof of (A.3).

We are given that $z = zxz$ and $z = zyz$, and we wish to show $z = zxyz$. Multiply the two given equations to get

$$z = (zyz)(zxz) = zyzxz.$$

This implies

$$z \succeq yzx = (yz)(zx) \sim (zx)(yz) = zxyz,$$

hence $z = z(zxyz)z = zxyz$, as required. □

We call L the *support semilattice* of S. Note that (A.2) completely characterizes L up to canonical isomorphism. Consequently:

Proposition A.12 *For any subsemigroup $S' \subseteq S$, the image of S' in L is the support semilattice of S'.* □

Example A.13 For any $x \in S$, the right ideal xS is a subsemigroup of S whose support semilattice is $L_{\geq X}$, where $X = \operatorname{supp} x$.

Next, as in the proof of Proposition A.6, one can generalize (A.2) to a deletion property:

Proposition A.14 *If* $\operatorname{supp} x = \operatorname{supp} z \geq \operatorname{supp} y$*, then* $xyz = xz$. □

Remark A.15 As in the free case, L is the abelianization of S.

Appendix B Representation theory of finite bands

The representation theory of finite bands can be developed as a special case of general results in semigroup theory. See, for instance, [11, 12, 13]. Here we give instead a self-contained treatment based on the support semilattice. This was already done for left-regular bands in [4], based on the work of Bidigare [1]. It is not much harder to treat the general case now that we have the support semilattice.

Throughout this appendix S denotes a finite band, L is its associated semilattice of supports, and k is an arbitrary field.

B.1 The radical of the semigroup algebra. Let kS and kL be the semigroup algebras of S and L. (Recall that L is a semigroup under the join operation.) Our starting point is the algebra homomorphism

$$\mathrm{supp}\colon kS \twoheadrightarrow kL \tag{B.1}$$

induced by the semigroup homomorphism $\mathrm{supp}\colon S \twoheadrightarrow L$. Recall next that kL is isomorphic to the algebra k^L of functions from L to k, which is a product of copies of k indexed by L. This was first proved by Solomon [16]; see also [8] and [18, Section 3.9]. Let's make the isomorphism explicit:

The algebra kL has a k-basis consisting of the elements $X \in L$, with product $(X, Y) \mapsto X \vee Y$. The algebra k^L has the standard k-basis $\{\delta_X\}_{X\in L}$, where δ_X is the function $Y \mapsto 1_{Y=X}$ whose value at Y is 1 if $Y = X$ and 0 otherwise. The product is given by $\delta_X^2 = \delta_X$ and $\delta_X\delta_Y = 0$ if $X \neq Y$. We get an isomorphism

$$\phi\colon kL \xrightarrow{\cong} k^L$$

by sending X to the function $Y \mapsto 1_{Y\geq X}$, i.e.,

$$\phi(X) = \sum_{Y\geq X} \delta_Y.$$

Note that ϕ preserves products because $1_{Y\geq X}1_{Y\geq X'} = 1_{Y\geq X\vee X'}$. And it is an isomorphism because of the triangular nature of ϕ.

We can use Möbius inversion to write down the inverse. Indeed, we have $X = \sum_{Y\geq X} \phi^{-1}(\delta_Y)$, so Möbius inversion gives

$$\phi^{-1}(\delta_X) = \sum_{Y\geq X} \mu(X, Y)Y,$$

where μ is the Möbius function of the poset L. The elements

$$e_X := \sum_{Y\geq X} \mu(X, Y)Y$$

are therefore the primitive idempotents of kL.

We now compose the support map in (B.1) with the isomorphism ϕ to obtain an algebra homomorphism

$$\psi\colon kS \twoheadrightarrow k^L.$$

Theorem B.1 *The kernel of ψ is nilpotent and is therefore the radical of kS.*

Proof The second assertion follows from the first because k^L is semisimple. To prove the first assertion, note that the kernel J of ψ is the kernel of $\operatorname{supp}\colon kS \twoheadrightarrow kL$. It consists of linear combinations of elements of S such that if we lump the terms according to supports, the coefficient sum of each lump is zero. Thus $J = \sum_{X \in L} J_X$, where J_X consists of linear combinations $\sum_{\operatorname{supp} x = X} a_x x$ with $\sum_x a_x = 0$. Let $\hat{1}$ denote the largest element of the poset L; it exists because L is finite and any two elements have an upper bound. For any $Y < \hat{1}$ in L, let $S_{\le Y}$ be the subsemigroup $\{\operatorname{supp} x \le Y\}$. We may assume, inductively, that the kernel of $\operatorname{supp}\colon kS_{\le Y} \to kL$ is nilpotent. Let n be an integer such that the nth power of this kernel is 0 for all such Y. I claim that $J^{2n+1} = 0$.

Consider a product of $2n+1$ elements from various J_X. If there are n consecutive factors such that the join of the corresponding X's is less than $\hat{1}$, then the product is 0 by the induction hypothesis. Otherwise, we may group the first n factors and the last n factors to write our product as uvw with $u, w \in J_{\hat{1}}$ and v in some J_X. Now uvw is a sum of terms of the form xvy with $x, y \in S$ and $\operatorname{supp} x = \operatorname{supp} y = \hat{1}$. Since v has coefficient sum 0, we have $xvy = 0$ by the deletion property (Proposition A.14). □

B.2 Representations. For any $X \in L$ the X-component of the homomorphism $\psi\colon kS \twoheadrightarrow k^L$ is the homomorphism $\chi_X\colon kS \to k$ given by

$$\chi_X(y) = 1_{\operatorname{supp} y \le X}$$

for $y \in S$. These are characters of S, which correspond to 1-dimensional representations.

Theorem B.2 *Let S be a finite band with support semilattice L, and let k be a field.*

1. *Every irreducible representation of kS is 1-dimensional.*
2. *There is one nonzero irreducible representation for each $X \in L$, given by the character χ_X.*
3. *The elements of S are simultaneously triangularizable in every finite dimensional representation.*

This is an immediate consequence of Theorem B.1 and standard ring theory. Direct arguments can be found in [4, Section 7.2] for the convenience of readers not familiar with the necessary ring theory.

B.3 Eigenvalues. Motivated by the theory of hyperplane arrangements, we introduce the 2-sided ideal $C \subseteq S$ consisting of the elements whose support is the largest element $\hat{1}$, and we call the elements of C *chambers*. In the free case, C consists of the *complete* elements of S (those using all the letters).

For applications to random walks, we are interested in the eigenvalues of an arbitrary element

$$w = \sum_{x \in S} w_x x$$

of kS acting by left multiplication on kC. For simplicity, we will assume that S has at least one left identity e. Then $xex = x$ for all $x \in S$, so $\operatorname{supp} e$ is the smallest element $\hat{0}$ of L, which is therefore a lattice. The assumption of a left identity is harmless, since we can always adjoin a 2-sided identity to S and a smallest element $\hat{0}$ to L. [I do *not* want to assume S has a 2-sided identity in general, however, since I want to consider subsemigroups xS, in which x is a left identity but not necessarily a right identity. The choice of left rather than right here is dictated by the convention that random walks are *left* random walks in this paper.]

For any $X \in L$, choose $x \in S$ with $\operatorname{supp} x = X$ and let c_X be the number of chambers in the band xS. This is independent of the choice of x; for if we also have $\operatorname{supp} y = X$, then left multiplication by x and y define mutually inverse chamber-preserving bijections between xS and yS. Define a family of integers m_X $(X \in L)$ by the system of equations

$$c_X = \sum_{Y \geq X} m_Y \tag{B.2}$$

for each $X \in L$. The proof of the following theorem will show that $m_X \geq 0$ for all X.

Theorem B.3 *Let S be a finite band with at least one left identity, let L be its support lattice, let C be the ideal of chambers, and let k be a field. For any element $w = \sum_{x \in S} w_x x \in kS$, let T_w be the operator on kC given by left multiplication by w. Then T_w has an eigenvalue*

$$\lambda_X = \sum_{\operatorname{supp} y \leq X} w_y$$

for each $X \in L$, with multiplicity m_X.

Note, as a check, that the sum of the multiplicities is indeed the dimension of kC, i.e., the total number of chambers, by Equation (B.2) with $X = \hat{0}$. In view of the theorem, we can interpret the general case of Equation (B.2) as saying the same thing with S replaced by xS for any $x \in S$ with $\operatorname{supp} x = X$.

Remark B.4 For particular choices of w, some of the λ_X might coincide, so one has to add the corresponding numbers m_X to get the true multiplicity (which might be 0, in which case λ_X does not actually occur as an eigenvalue). The precise meaning of the theorem is that the characteristic polynomial of T_w is $f(\lambda) = \prod_{X \in L} (\lambda - \lambda_X)^{m_X}$.

Proof of Theorem B.3 Choose a composition series for kC as a left kS-module. By Theorem B.2, each composition factor is 1-dimensional and is given by one of the characters χ_X. (The zero-representation cannot occur here because of our assumption that S has at least one left identity.) For each $X \in L$, let m'_X be the number of times χ_X occurs. Then T_w has a triangular matrix representation whose diagonal entries are the values $\chi_X(w)$, the latter occurring m'_X times. Now

$$\chi_X(w) = \sum_{y \in S} w_y \chi_X(y) = \sum_{y \in S} w_y 1_{\operatorname{supp} y \leq X} = \sum_{\operatorname{supp} y \leq X} w_y = \lambda_X,$$

so the proof will be complete if we show that $\sum_{Y \geq X} m'_Y = c_X$ for all $X \in L$.

Consider an arbitrary $x \in S$. It acts on kC as an idempotent operator, projecting kC onto the linear span of the chambers in xS. The rank of this projection is therefore c_X, where $X = \operatorname{supp} x$. On the other hand, the rank is also the multiplicity of 1 as an eigenvalue, i.e., the number of composition factors whose character takes the value 1 at x. So

$$c_X = \sum_{\substack{Y \in L \\ \chi_Y(x) = 1}} m'_Y = \sum_{Y \geq X} m'_Y.$$

□

References

[1] T. Patrick Bidigare, *Hyperplane arrangement face algebras and their associated Markov chains*, Ph.D. thesis, University of Michigan, 1997.

[2] T. Patrick Bidigare, Phil Hanlon, and Daniel N. Rockmore, *A combinatorial description of the spectrum for the Tsetlin library and its generalization to hyperplane arrangements*, Duke Math. J. **99** (1999), no. 1, 135–174.

[3] Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler, *Oriented matroids*, Encyclopedia of Mathematics and its Applications, vol. 46, Cambridge University Press, Cambridge, 1993.

[4] Kenneth S. Brown, *Semigroups, rings, and Markov chains*, J. Theoret. Probab. **13** (2000), no. 3, 871–938.

[5] Kenneth S. Brown and Persi Diaconis, *Random walks and hyperplane arrangements*, Ann. Probab. **26** (1998), no. 4, 1813–1854.

[6] Persi Diaconis, *The cutoff phenomenon in finite Markov chains*, Proc. Nat. Acad. Sci. U.S.A. **93** (1996), no. 4, 1659–1664.

[7] J. A. Green and D. Rees, *On semi-groups in which $x^r = x$*, Proc. Cambridge Philos. Soc. **48** (1952), 35–40.

[8] Curtis Greene, *On the Möbius algebra of a partially ordered set*, Advances in Math. **10** (1973), 177–187.

[9] John G. Kemeny and J. Laurie Snell, *Finite Markov chains*, The University Series in Undergraduate Mathematics, D. Van Nostrand Co., Inc., Princeton, N.J.-Toronto-London-New York, 1960.

[10] David McLean, *Idempotent semigroups*, Amer. Math. Monthly **61** (1954), 110–113.

[11] Jan Okniński, *Semigroup algebras*, Monographs and Textbooks in Pure and Applied Mathematics, vol. 138, Marcel Dekker Inc., New York, 1991.

[12] ———, *Semigroups of matrices*, Series in Algebra, vol. 6, World Scientific Publishing Co. Inc., River Edge, NJ, 1998.

[13] Heydar Radjavi and Peter Rosenthal, *Simultaneous triangularization*, Springer-Verlag, New York, 2000.

[14] Laurent Saloff-Coste, *Random walks on finite groups*, Combinatorics and probability on discrete structures (Harry Kesten, ed.), Encyclopedia of mathematical sciences, Springer-Verlag, 2003, to appear.

[15] J. Siekmann and P. Szabó, *A Noetherian and confluent rewrite system for idempotent semigroups*, Semigroup Forum **25** (1982), no. 1-2, 83–110.

[16] Louis Solomon, *The Burnside algebra of a finite group*, J. Combinatorial Theory **2** (1967), 603–615.

[17] ———, *A Mackey formula in the group ring of a Coxeter group*, J. Algebra **41** (1976), no. 2, 255–264.

[18] Richard P. Stanley, *Enumerative combinatorics. Vol. 1*, Cambridge Studies in Advanced Mathematics, vol. 49, Cambridge University Press, Cambridge, 1997, with a foreword by Gian-Carlo Rota, corrected reprint of the 1986 original.

[19] Jacques Tits, *Buildings of spherical type and finite BN-pairs*, Springer-Verlag, Berlin, 1974, Lecture Notes in Mathematics, Vol. 386.

[20] Thomas Zaslavsky, *Facing up to arrangements: face-count formulas for partitions of space by hyperplanes*, Mem. Amer. Math. Soc. **1** (1975), no. 154, vii+102 pp.