Combinatorial Ergodicity

James Propp
(UMass Lowell; visiting UCB and MSRI)

April 23, 2012

Slides for this talk are on-line at

http://jamespropp.org/ucbcomb12.pdf

Acknowledgments

This talk describes on-going work with Tom Roby.

Thanks also to Drew Armstrong, Karen Edwards, Bob Edwards. Svante Linusson, Richard Stanley, and Ben Young.

Overview

For many cyclic actions τ on a finite set S of combinatorial objects, and for many natural statistics ϕ on S, one finds that the average of ϕ over each τ-orbit in S is the same as the average of ϕ over the whole set S.

We say that (S, τ, ϕ) exhibits **combinatorial ergodicity**.

I'll illustrate this with two actions on $J([a] \times [b])$ — rowmotion and promotion, in Striker and Williams' terminology — that can also be viewed as actions on associated antichains and lattice-paths.

An invertible operation on antichains

Let $\mathcal{A}(P)$ be the set of antichains of a finite poset P.

Given $A \in \mathcal{A}(P)$, let $\tau(A)$ be the set of minimal elements of the complement of the downward-saturation of A.

τ is invertible since it is a composition of three invertible operations:

$$\text{antichains} \longleftrightarrow \text{downsets} \longleftrightarrow \text{upsets} \longleftrightarrow \text{antichains}$$

This map and its inverse have been considered with varying degrees of generality, by many people more or less independently (using a variety of nomenclatures and notations): Duchet, Brouwer and Schrijver, Cameron and Fon Der Flaass, Fukuda, and Panyushev.

Panyushev's conjecture

Most of the work on τ has focussed on its orbit structure, with the notable exception of a conjecture of Panyushev (Conjecture 2.1(iii) in his 2009 article "On orbits of antichains of positive roots"), proved by Armstrong, Stump, and Thomas in their 2011 article "A uniform bijection between nonnesting and noncrossing partitions":

Panyushev's conjecture: Let Δ be a reduced irreducible root system in $\mathbf{R}^n$.
Choose a system of positive roots and make it a poset by decreeing that y covers x iff $y - x$ is a simple root.
Let $\mathcal{O}$ be an arbitrary τ-orbit. Then

$$\frac{1}{\#\mathcal{O}} \sum_{A \in \mathcal{O}} \#(A) = \frac{n}{2}.$$

(A more general assertion of this kind, Panyushev's Conjecture 2.3(iii), remains open. See also Panyushev's Conjecture 2.4(ii).)

Products of two chains

A simpler phenomenon of this kind concerns the poset $[a] \times [b]$ (where $[n]$ denotes the linear ordering of $\{1, 2, \ldots, n\}$):

Let $\mathcal{O}$ be an arbitrary τ-orbit in $\mathcal{A}([a] \times [b])$. Then

$$\frac{1}{\#\mathcal{O}} \sum_{A \in \mathcal{O}} \#(A) = \frac{ab}{a+b}.$$

Indeed, for $i \in [a]$, let $f_i(A)$ be 1 or 0 according to whether or not the projection of $A \subset [a] \times [b]$ onto $[a]$ contains i; then
Claim 1:

$$\frac{1}{\#\mathcal{O}} \sum_{A \in \mathcal{O}} f_i(A) = \frac{b}{a+b}.$$

Note that $\sum_{1 \leq i \leq a} f_i(A) = \#(A)$, so the second equation refines the first.

(A similar situation holds when $[a]$ and $[b]$ switch roles.)

Armstrong's clarification of Stanley's observation

In his 2009 article "Promotion and evacuation", Stanley observed (in the last paragraph of section 2) that there is a connection between the action of τ on $\mathcal{A}([a] \times [b])$ and the action of promotion on linear extensions of the disjoint union of $[a]$ and $[b]$.

In unpublished work, Armstrong has presented a very clear picture of this.

Given an antichain A in $[a] \times [b]$, define the associated "Armstrong word" $w_1, \dots, w_a, w_{a+1}, \dots, w_{a+b}$ as follows:

For $1 \leq i \leq a$, w_i is 1 if the projection of A onto $[a]$ contains i and 0 otherwise.
For $1 \leq j \leq b$, w_{a+j} is $\underline{0}$ if the projection of A onto $[b]$ contains j and $\underline{1}$ otherwise.
That is, the Armstrong word of A says which rows of $[a] \times [b]$ **do** contain an element of A, and which columns of $[a] \times [b]$ **don't**.

The Armstrong map is a bijection

The Armstrong word contains b 1's and a 0's.
(Because: If $\#(A) = k$, then $w_1, \ldots, w_a$ contains k 1's and $w_{a+1}, \ldots, w_{a+b}$ contains $b - k$ 1's.)

The Armstrong word of A determines A.
(Because: There is a unique order-reversing pairing between $\{i \in [a] : w_i = 1\}$ and $\{j \in [b] : w_{a+j} = 0\}$.)

Every sequence consisting of a 0's and b 1's is the Armstrong word of a unique antichain in $[a] \times [b]$.

The Armstrong map is an equivariant bijection

The Armstrong word of $\tau(A)$ is the rightward cyclic shift of the Armstrong word of A.

That is: the action of τ on the "Armstrong set" of A (the set of indices at which the Armstrong word takes the value 1) increments each index by 1 mod $a + b$.

It follows that the order of every τ-orbit $\#(\mathcal{O})$ is a divisor of $a + b$ (and indeed the entire orbit structure can be quickly derived, including the cyclic sieving phenomenon).

Proof of Claim 1

Look at the antichains A, $\tau(A)$, $\tau^2(A), \ldots, \tau^{a+b-1}(A)$ and their associated Armstrong words $w^0, w^1, w^2, \ldots, w^{a+b-1}$.

We have $f_i(\tau^k(A))$ if and only if the Armstrong word w^k has a 1 in the ith position.

Since the w^k's are just the successive cyclic shifts of w^0, each of the b 1's in w^0 appears in the ith position exactly once as k runs through $0, 1, 2, \ldots, a+b-1$.

So

$$\sum_{k=0}^{a+b-1} f_i(\tau^k)(A) = b.$$

Proof of Claim 1 (concluded)

Since

$$\sum_{k=0}^{a+b-1} f_i(\tau^k)(A) = b,$$

the average of $f_i(\tau^k)(A)$ over $0 \leq k \leq a+b-1$ is $b/(a+b)$.

But each antichain in the τ-orbit of A occurs equally often (namely, $n/\#(\mathcal{O})$ times) in the list $\tau^0(A), \ldots, \tau^{a+b-1}(A)$.

So the average of f_i over the orbit $\mathcal{O}$ is also $b/(a+b)$. $\square$

Note that the variation in orbit-sizes (of great concern in the theory of cyclic sieving) is irrelevant here.

From antichains to order ideals

Given a poset P and an antichain A in P, let $\mathcal{I}(A)$ be the order ideal $I = \{y \in P \ : y \leq x \text{ for some } x \in A\}$ associated with A, so that for any order ideal I in P, $\mathcal{I}^{-1}(I)$ is the antichain of maximal elements of I.

As usual, we let $J(P)$ denote the set of order ideals of P.

We define $\overline{\tau} : J(P) \to J(P)$ by $\overline{\tau}(I) = \mathcal{I}(\tau(\mathcal{I}^{-1}(I)))$.

We will sometimes write $\overline{\tau}$ as just τ, since it is usually clear from context which map we mean (even though technically some sets are both antichains and order ideals).

An order-ideal counterpart of Claim 1

Let $\mathcal{O}$ be an arbitrary $\overline{\tau}$-orbit in $J([a] \times [b])$. Then

$$\frac{1}{\#\mathcal{O}} \sum_{I \in \mathcal{O}} \#(I) = \frac{ab}{2}.$$

Indeed, for $1 - b \leq i \leq a - 1$, let $f_i(I)$ be the number of elements $(x, y) \in I$ with $x - y = i$; then

Claim 2 (Propp and Roby):

$$\frac{1}{\#\mathcal{O}} \sum_{I \in \mathcal{O}} f_i(I) = \begin{cases} \frac{(a-i)b}{a+b} & \text{if } i \geq 0 \\ \frac{a(b+i)}{a+b} & \text{if } i \leq 0. \end{cases}$$

(Summing over i yields the first equation.)

From order ideals to lattice paths

The asymmetry between the cases $i \geq 0$ and $i \leq 0$ disappears if we change our point of view and replace antichains by lattice paths.

Associate each order ideal I with a Young diagram in the usual way. Then flip each diagram across the line $y = x$, rotate by 45 degrees counterclockwise, and scale up by $\sqrt{2}$.

The boundary between the Young diagram and its complement becomes a lattice path from $(-a, a)$ to (b, b) consisting of a steps of type $(1, -1)$ ("down-steps") and b steps of type $(1, 1)$ ("up-steps").

Restating Claim 2

Think of the lattice path associated with I as a function $F_I : \{-a, -a+1, \ldots, b\} \to \{0, 1, \ldots, a+b\}$ such that $F(-a) = a$, $F(b) = b$, and $F(-a+i+1) = F(-a+i) \pm 1$ for $0 \leq i < a+b$.

Then Claim 2 is equivalent to the assertion that the average of $F_I(-a+i)$ along a $\overline{\tau}$-orbit is

$$a - \frac{a-b}{a+b}\, i.$$

(I'll prove this about ten slides from now.)

Binary words

We can represent each lattice path by a string of a 0's and b 1's (not to be confused with the Armstrong word) where a 0 signifies a down-step and a 1 signifies an up-step.

Then the map $\overline{\tau}$, viewed as a map from the set of such "path-words" to itself, can be described as a "block-reversal map": Divide the word into "primary blocks" (subwords of the form 01) and "secondary blocks" (the blocks that remain when the primary blocks are removed).
Then just reverse each block in place.

Example:
$0111000010110 = 01.11000.01.01.10$
maps to $10.00011.10.10.01 = 1000011101001$.

Inversion in binary words

The cardinality of an order ideal I is equal to $\mathrm{inv}(w)$, where w is the associated binary word and $\mathrm{inv}(w) = \#\{1 \le i < j \le a+b : w_i > w_j\} = \#\{1 \le i < j \le a+b : w_i = 1,\ w_j = 0\}$.

Example: The binary word 10100 has 5 inversions, so the associated order ideal in $[3] \times [2]$ has cardinality 5.

The formula $\frac{1}{\#\mathcal{O}} \sum_{I \in \mathcal{O}} \#(I) = \frac{ab}{2}$ can be rewritten as

$$\frac{1}{\#\mathcal{O}} \sum_{w \in \mathcal{O}} \mathrm{inv}(w) = \frac{ab}{2}$$

where now $\mathcal{O}$ is an orbit of binary words under the action of block-reversal.

Toggling

In their 1995 article "Orbits of antichains revisited", Cameron and Fon-der-Flaass give an alternative description of τ in terms of toggle-operations applied to order ideals.

Given $I \in J(P)$ and $x \in P$, let $\tau_x(I) = I \triangle \{x\}$ provided that $I \triangle \{x\}$ is an order ideal of P; otherwise, let $\tau_x(I) = I$.

We call the involution τ_x "toggling at x".

The involutions τ_x and τ_y commute unless x covers y or y covers x.

Toggling from top to bottom

Theorem (Cameron and Fon-der-Flaass): Let $x_1, x_2, \ldots, x_n$ be any order-preserving enumeration of the elements of the poset P. Then the action on $J(P)$ given by the composition $\tau_{x_1} \circ \tau_{x_2} \circ \cdots \circ \tau_{x_n}$ coincides with the action of $\overline{\tau}$.

In the particular case $P = [a] \times [b]$, we can enumerate P rank-by-rank; that is, we can list the (i,j)'s in order by $i + j$.

Note that all the involutions coming from a given rank of P commute with one another, since no two of them are in a covering relation.

Toggling from side to side

Define a **file** of $P = [a] \times [b]$ as the set of all $(i, j) \in P$ with $i - j$ fixed.

Note that all the involutions coming from a given file commute with one another, since no two of them are in a covering relation.

Theorem (Striker and Williams): Let $x_1, x_2, \ldots, x_n$ be any enumeration of the elements of the poset $[a] \times [b]$ arranged in order of increasing $i - j$.
Then the action on $J(P)$ given by $\tau_{x_1} \circ \tau_{x_2} \circ \cdots \circ \tau_{x_n}$ (Striker and Williams call this well-defined composition promotion since it is closely related to Schützenberger's notion of promotion on linear extensions of posets), conjugated to an action on the set of path-words, coincides with the action of the cyclic shift.

A Claim about promotion

Claim 3 (Propp and Roby): Let $\mathcal{O}$ be an arbitrary orbit in $J([a] \times [b])$ under the action of promotion ∂. Then

$$\frac{1}{\#\mathcal{O}} \sum_{I \in \mathcal{O}} \#(I) = \frac{ab}{2}.$$

Equivalently: Let $\mathcal{O}$ be an orbit in the set of words w composed of a 0's and b 1's under the action of rotation. Then

$$\frac{1}{\#\mathcal{O}} \sum_{w \in \mathcal{O}} \mathrm{inv}(w) = \frac{ab}{2}.$$

I know two simple ways to prove this: one can show pictorially that the value of the sum doesn't change when you mutate w (replacing a 01 somewhere in w by 10 or vice versa), or one can write the number of inversions in w as $\sum_{i<j} w_i(1-w_j)$ and then perform algebraic manipulations.

Refining the claim

As we did earlier, we can replace each order ideal I by a function $F_I : \{-a, -a+1, \ldots, b\} \to \{0, 1, \ldots, a+b\}$.

Then we can show that the average of $F_I(-a+i)$ along a promotion orbit is

$$a - \frac{a-b}{a+b}\, i$$

(just as we claimed for rowmotion orbits).

Proving the claim

Create a rectangular array with $a+b$ rows and $a+b$ columns (indexed 0 through $a+b+1$).

The k, i entry is $F_{\partial^k(I)}(-a+i+1) - F_{\partial^k(I)}(-a+i) = \pm 1$ (where ∂ is the cyclic shift) so that $F_{\partial^k(I)}(-a+i)$ equals a plus the sum of the first i entries in row k, and the orbit-average of $F_{\partial^k(I)}(-a+i)$ equals a plus the average of the sum of the first $i-1$ entries in row k, or equivalently a plus the sum of the first $i-1$ column-averages.

But each column is a cyclic shift of the previous column, and each row and column contains a -1's and b $+1$'s, so each column-average is $((a)(-1)+(b)(+1)/(a+b) = (b-a)/(a+b)$, so the orbit-average of $F_{\partial^k(I)}(-a+i)$ is $a - i(a-b)/(a+b)$ as claimed. $\square$

Proving Claim 2

Create a $(a+b)$-by-$(a+b)$ square array whose k, i entry is $F_{\tau^k(I)}(-a+i+1) - F_{\tau^k(I)}(-a+i) = \pm 1$ where τ denotes rowmotion.

A row contains $-1, +1$ in the ith and $i+1$st positions respectively if and only if the next row (in cyclic order) contains $+1, -1$ in the ith and $i+1$st positions respectively.

Hence the ith column equals the $i+1$st column sum.

Since all the column-averages are equal, and since their grand average is $(b-a)/(a+b)$, each column-average is $(b-a)/(a+b)$.

The rest of the argument is the same. $\square$

The story thus far

We've looked at 2 different actions (rowmotion and promotion) and 2 different notions of cardinality for the objects they act on (antichains and order ideals).

In 3 of the 2×2 cases, the average cardinality along an orbit equals the average cardinality over the whole space; equivalently, the average cardinality along an orbit doesn't depend on the orbit.

We know many other equalities like this (some proved, most still conjectural).

We call this the constant-averages-along-orbits property, or combinatorial ergodicity.

"Ergodicity"?

This may seem like a misnomer: A measurable action is ergodic iff the only invariant sets have measure zero or full measure, so in the combinatorial setting, an action is ergodic iff it is transitive.

However, the coinage makes more sense if you think back to Boltzmann's original notion of the equality between space-averages and long-term time-averages.

Note that if x is a periodic point for the invertible map τ (and there is no other kind of point if τ is a permutation!) we have

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \phi(\tau^k(x)) = \frac{1}{\#(\mathcal{O})} \sum_{y\in\mathcal{O}} \phi(y)$$

where $\mathcal{O}$ is the orbit of x.

Research strategy

Given an action on a combinatorial set S, identify features of interest in S (e.g., if S is a set of binary words a feature might be an inversion between two specific positions).

Create Boolean functions on S, each indicating the presence or absence of a feature, and take the linear span of these functions.

Within this space, look for functions that are combinatorially ergodic.

Also look for functions that are invariant along orbits.

We expect that these two (ortho)complementary kinds of functions can coordinatize S in useful ways.

Example

Go back to Claim 1 (τ is rowmotion acting on antichains).

The functions

$$1_{i,j}(A) = \begin{cases} 1 & \text{if } (i,j) \in A, \\ 0 & \text{if } (i,j) \notin A, \end{cases}$$

do not themselves exhibit combinatorial ergodicity, but the functions $f_i = \sum_j 1_{i,j}(A)$ in their span do.

Moreover, in the span of the function $1_{i,j}(A)$ and $1_{i,j}(\tau(A))$, one finds that the functions $f_i(A) - f_{i+1}(\tau(A))$ are invariant.

From such relations (easy to conjecture with the aid of computer linear algebra) the whole Armstrong picture can readily be guessed.

Example, continued

Computation (by hand or computer) also leads us to conjecture that $\sum_{i,j}(i-j)1_{i,j}$ exhibits combinatorial ergodicity.

In terms of binary words, this is saying that the major index has the constant-averages-along-orbits property for rowmotion.

Roby and I are working on a proof of this.

The last slide of this talk

I'm happy to talk about this stuff further with anyone who's interested.

My office hour is at 11 am on Tuesdays and Thursdays in 1063 Evans (at least until the term ends!).

I divide my out-of-the-house time between Evans, MSRI, and nearby cafes.

Slides for this talk are on-line at

http://jamespropp.org/ucbcomb12.pdf