

abstract: Let G be a finite group acting on finite set X . This divides X into orbits and we ask ?how many orbits are there?, ?What is the typical size of an orbit?, ?do the orbits have 'nice names'?, ?do they fit together into some kind of 'moduli space'?. The Burnside process addresses the first two topics. It allows uniform sampling of a random orbit. It runs the following Markov chain on X : From x , choose g fixing x (uniformly) and then y fixed by this g (uniformly). The chain moves from x to y . This chain has stationary distribution $\pi(x)$ proportional to the size of the orbit containing x . Thus simply reporting the size of the current orbit gives a Markov chain with a uniform stationary distribution.

Problems abound: How do you actually carry out the two steps in problems of interest? How can you convert the output into a useful estimate of the number and size of orbits? What is the running time and typical behaviour of this Markov chain? There is progress in special cases, but in general, all problems are open.

Lecture 1: An introduction to Polya Theory, the Burnside process and auxiliary variables algorithms

Lecture 2: Three real examples: Polya trees, contingency tables with fixed row and column sums, The commuting graph process for generating random partitions of n (and counting the number of conjugacy classes in a finite group).

Lecture 3: Careful proof in a simple case: the binary Burnside process, from Tchebychev polynomials to Schur-Weyl duality.

I expect the audience to know something about Markov chains (at the level of Levin-Peres' book) and something about groups (a good undergraduate course should suffice). These topics are based on my current research and you can look at papers on my homepage for the past year or two.