Abstract

Graded posets frequently arise throughout combinatorics, where it is natural to try to count the number of elements of a fixed rank. These counting problems are often #**P**-complete, so we consider approximation algorithms for counting and uniform sampling. We show that for certain classes of posets, biased Markov chains that walk along edges of their Hasse diagrams allow us to approximately generate samples with any fixed rank in expected polynomial time. Our arguments do not rely on the typical proofs of log-concavity, which are used to construct a stationary distribution with a specific mode in order to give a lower bound on the probability of outputting an element of the desired rank. Instead, we infer this directly from bounds on the mixing time of the chains through a method we call *balanced bias*. A noteworthy application of our method is sampling restricted classes of integer partitions of n. We give the first provably efficient Markov chain algorithm to uniformly sample integer partitions of n from general restricted classes. Several observations allow us to improve the efficiency of this chain to require $O(n^{1/2}\log(n))$ space, and for unrestricted integer partitions, expected $O(n^{9/4})$ time. Related applications include sampling permutations with a fixed number of inversions and lozenge tilings on the triangular lattice with a fixed average height.