Phase Transitions in Sampling Algorithms and the Underlying Random Structures

Dana Randall

College of Computing, Georgia Institute of Technology
Atlanta, GA 30332-0765
randall@cc.gatech.edu

Sampling algorithms based on Markov chains arise in many areas of computing, engineering and science. The idea is to perform a random walk among the elements of a large state space so that samples chosen from the stationary distribution are useful for the application. In order to get reliable results, we require the chain to be rapidly mixing, or quickly converging to equilibrium. For example, to sample independent sets in a given graph G, the so-called hard-core lattice gas model, we can start at any independent set and repeatedly add or remove a single vertex (if allowed). By defining the transition probabilities of these moves appropriately, we can ensure that the chain will converge to a useful distribution over the state space Ω. For instance, the Gibbs (or Boltzmann) distribution, parameterized by $\Lambda > 0$, is defined so that $p(\Lambda) = \pi(I) = \Lambda^{|I|}/Z$, where $Z = \sum_{J \in \Omega} \Lambda^{|J|}$ is the normalizing constant known as the partition function. An interesting phenomenon occurs as Λ is varied. For small values of Λ, local Markov chains converge quickly to stationarity, while for large values, they are prohibitively slow. To see why, imagine the underlying graph G is a region of the Cartesian lattice. Large independent sets will dominate the stationary distribution π when Λ is sufficiently large, and yet it will take a very long time to move from an independent set lying mostly on the odd sublattice to one that is mostly even. This phenomenon is well known in the statistical physics community, and characterizes by a phase transition in the underlying model.

In general, phase transitions occur in models where a small microscopic change to some parameter suddenly causes a macroscopic change to the system. This phenomenon is pervasive in physical systems, but often lacks rigorous analysis. Colloids, which are mixtures of two disparate substances in suspension, are an interesting example. At low enough density the two types of particles will be uniformly interspersed, but at sufficiently high density the particles of each substance will cluster together, effectively separating. Unlike seemingly related models where the clustering occurs because like particles are drawn together by enthalpic forces, such as the ferromagnetic Ising model, the behavior of colloids is purely entropic – the particles separate because because the overwhelming majority of configurations in the stationary distribution exhibit such a separation. In this talk I will give an overview of some techniques used to prove the existence of two distinct phases in various sampling algorithms and the underlying physical models. Finally, I will suggest some potential approaches for using our understanding of these two phases to inform the design of more efficient algorithms.