"Counting, Sampling and Integrating: Algorithms and Complexity" by Mark Jerrum (Birkhauser, 2003, softcover, xi+112pp., $29.95, ISBN 3-7643-6946-9).

The phrase *Lecture Notes* is too often used to make a dry monograph sound more appealing, so it is a pleasure to find this thin volume of well-written notes based on real lectures. The topic is at the intersection of two large subjects, which I should first mention.

At the theoretical end of Computer Science lies the subject of *combinatorial algorithms and complexity* (say CA&C), the design and analysis of algorithms whose input is a combinatorial structure, often a graph, and the associated structural theory of computational complexity. A quite different subject, originating 50 years ago in statistical physics, is *Markov Chain Monte Carlo* (MCMC). Given a space of configurations $\mathbf{x}$ and a function $H(\mathbf{x})$, suppose we want to study numerically the probability distribution $\Pr(\mathbf{x}) \propto \exp(-H(\mathbf{x}))$; for instance we may wish to find the mean

$$\bar{f} = \sum_{\mathbf{x}} f(\mathbf{x}) \Pr(\mathbf{x})$$

$$= \sum_{\mathbf{x}} f(\mathbf{x}) \exp(-H(\mathbf{x})) \Big/ \sum_{\mathbf{x}} \exp(-H(\mathbf{x}))$$

of a function f on configurations. The point of MCMC is that there are general methods, not requiring any theoretical knowledge of $\Pr(\cdot)$, to define and simulate a Markov chain whose stationary distribution is $\Pr(\cdot)$, so that long-run averages of f computed along the simulation will approach the true mean $\bar{f}$.

How these two subjects can intersect is best illustrated by example. Suppose we want an algorithm which *approximately* counts the number of matchings in a finite graph G. By induction on number of edges, it is enough to find the proportion of matchings in G which contain a specified edge e_0. This is equivalently the probability that a uniform random matching of G contains e_0, and MCMC allows us to find this probability approximately by running a Markov chain on matchings.

To make this argument rigorous, in the sense of relating the error of approximation to the number of steps of an algorithm, the key issue is the *mixing time* of the Markov chain, which indicates the number of steps required to approach the stationary distribution. This same issue of bounding the mixing time arises in any application of MCMC: in most scientific applications such bounds are out of reach of useful theory, but the topic of this book (along with Ising-type models of statistical physics [5] and random walks on finite groups [1]) provided impetus for development, over the last 15 years, of a suite of technical tools.

Jerrum's book starts by illustrating exact counting algorithms for spanning trees and for perfect matching in a planar graph, and describing the complexity-theoretic notion of #P completeness. It then describes Markov chains and mixing times, the classical coupling method for bounding mixing times, and the reduction (mentioned above) of approximate counting to approximate sampling. Then comes the core of the book, the development of three main tools, each illustrated by a paradigm example. In Chapter 4, coupling is applied to approximate counting of colorings of a finite graph, and the related notion of path-coupling is applied to the problem of approximately counting the number of extensions of a partial order. In Chapter 5 the method of canonical paths is applied to the problem of approximately counting matchings. In Chapter 6, the problem of finding approximately the volume of a convex set in high dimensions is studied via more geometric *Poincare inequality* methods.

The book does a wonderful job of going in 112 pages from modest prerequisites to "close to the coal face of current research", in the Introduction's words. The focus on a specific topic and the slightly sophisticated style – more survey paper than textbook – may not give the novice much perspective on the big picture of the surrounding two subjects. In the case of MCMC, the equally thin lecture notes [2, 4] provide a more standard introduction to the basic mathematics of MCMC, while thicker textbooks such as [6, 3] provide broad global perspective. But this book will be invaluable to anyone, from beginning graduate student onwards, wishing to learn its particular topic. Moreover, for an established researcher in one of the two subjects (CA&C or MCMC) wishing a friendly introduction to the other, this book is impossible to beat.

David Aldous
University of California, Berkeley

References

[1] P. Diaconis. *Group Representations in Probability and Statistics*. Institute of Mathematical Statistics, Hayward CA, 1988.

[2] O. Häggstrom. *Finite Markov Chains and Algorithmic Applications*. Cambridge University Press, 2002.

[3] J.S. Liu. *Monte Carlo Techniques in Scientific Computing*. Springer–Verlag, 2001.

[4] N. Madras. *Lectures on Monte Carlo Methods*. American Math. Society, 2002.

[5] F. Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability Theory and Statistics (St. Flour 1997)*, volume 1117 of *Lecture Notes in Math.*, pages 93–191. Springer–Verlag, 1999.

[6] C.P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer–Verlag, 2000.