Review[3] of

Probability and Computing: Randomized Algorithms and Probabilitic Analysis

Author of Book: Michael Mitzenmacher and Eli Upfal

Publisher: Cambridge University Press, 2005

Cost: $55.00, hardcover

Reviewer: Jonathan Katz, Dept of CS, Univ of MD.

Randomized algorithms and basic techniques of probabilistic analysis are essential components of the theoretical computer scientist's tool-kit; indeed, their applications have become so widespread that they should be (and, increasingly, are) part of *every* computer scientist's tool-kit. Randomization has of course proved useful in designing traditional algorithms; interestingly, here we don't actually know whether randomization truly "helps" (at least in the sense that we don't have a proof that $\mathcal{P} \neq \mathcal{BPP}$), but in practice randomized algorithms are often more efficient and/or simpler than their deterministic counterparts. Randomization is also essential (in a provable sense) for cryptography and for certain results in distributed computing. Finally, a good understanding of random processes is also needed any time one is trying to analyze a system that evolves probabilistically, or trying to bound the performance of an algorithm (even a deterministic one!) when the inputs are not chosen deterministically.

Because of the widespread interest in the subject, a textbook covering randomization in computer science must be many things to many different people: it should serve as an introduction to the area for an undergraduate or graduate student interested in randomized algorithms; a survey of applications and techniques for the general computer scientist; and also a solid reference for the advanced researcher. I am pleased to say that *Probability and Computing ...* succeeds on all these fronts. I found the book a joy to read: the writing is clear and concise, the proofs are well-structured, and the writing style invites the reader to explore the material. The book is also organized very well, and the selection of topics is excellent. I have already used the book multiple times as a reference, and have found it incredibly useful each time.

The chapters of the book can roughly be classified in three sections (essentially as suggested by the authors in the preface): introductory material and review; basic material of the sort that might constitute an undergraduate course on the subject; and more advanced material for a graduate course or independent study. The introductory material, in which I would include Chapters 1 and 2, begins with a review of basic probability theory along with some standard applications to motivate the study of the subject. Chapter 2 includes a discussion of random variables and their expectation, and also introduces the Bernoulli, binomial, and geometric distributions. It is fair to say that the treatment here is meant as a review for the student who has had some prior exposure to probability in a discrete mathematics course, and is not meant to teach this material from scratch.

The basic material, contained in Chapters 3–7, is actually quite comprehensive. Standard material such as Markov's inequality, Chebyshev's inequality, and Chernoff bounds are included; in the book's coverage of this material, I especially appreciated the applications that are stressed throughout and are covered at a good level of detail. Chapter 5 covers the "balls-into-bins" probabilistic model and the Poisson distribution, again giving a wealth of realistic examples to continually motivate the material. Chapter 6 details the probabilistic method as well as the Lovasz Local Lemma, and Chapter 7 deals with Markov chains and random walks.

Chapters 1–7 could form the basis for a solid introductory graduate course on randomized algorithms; I think these chapters (with some material cut so as to go at a slower pace) would also work well as a (tough) undergraduate course. (Indeed, the authors have used parts of this book in undergraduate courses taught at Harvard and Brown.) In particular, I think the writing is "friendly" enough to suit the dedicated undergraduate, and the focus on interesting applications of the material would catch the interest of an undergraduate even if they were not solely interested in theory.

The remaining seven chapters of the book could be used as the core of a second semester graduate course, for self-study, or for reference. I quite liked the selection of topics here; these include chapters covering continuous distributions; the basics of entropy and information theory; Monte Carlo methods; martingales and the Azuma-Hoeffding inequality; and pairwise independence/universal hash functions (and more). My only quibble here is that I would have liked to see the chapter on pairwise independence earlier; this material is both more central and less difficult than the other material included in these "advanced" chapters, and so I don't see why it comes so late in the book. This, however, is a rather minor point.

To the best of my knowledge, this book is the best available treatment of randomized algorithms in terms of its breadth, depth, and accessibility. It will serve handily as a reference or as a guide to self-study, and belongs on the bookshelf of every graduate student and computer scientist interested in the role of randomization in computing.

Review[4] of

Probability and Computing: Randomized Algorithms and Probabilitic Analysis

Authors of Book: Michael Mitzenmacher and Eli Upfal

Publisher: Cambridge University Press, 2005

Cost: $55.00, hardcover

Reviewer: Yannis C. Stamatiou

1 Introduction

The question of whether the evolution of our world obeys a still undiscovered complex set of rules, or to put it in modern terminology, behaves as a deterministic computer algorithm with each step following in a unique way from the previous step, is very old and touches the borders between philosophy and science. Isaac Newton's feat, to describe neatly and orderly the evolution of the movement of heavenly bodies, seemed to answer this question in the affirmative. Later, however, this venerable and reassuring deterministic point of view was shaken first, "mildly", by Poincaré's work about the inability to solve the equations precisely, even if the world behaves deterministically and then, more fundamentally, by quantum physics which completely defied the deterministic point of view.

The authors discuss the points above right from the very beginning of the book and conclude that randomness seems to be here to stay while, in addition, it can be put into very good use in algorithmics and complexity theory. This book, thus, is about how randomness can be harnesses in order to serve effective computation and how its effects can be assessed through the tools of probability theory.

2 Summary of the book's chapters

In what follows we will briefly present each chapter's contents along with comments that I hope will help the reader to assess the chapters as well as the book as a whole and form an idea as to whether this book would be of use to her/him.

Chapter 1. The first chapter lays the foundations of probability theory and nicely establishes links with computer science and algorithmics. Contrary to usual practices, the book's very first chapter starts with a just one simple problem, that of checking whether two polynomials are the same. This problem is simple and its solution intuitively clear. Using this problem, the authors link, in just two pages, algorithmics with randomness demonstrating, at the same time, the gains of using randomness instead of trying to formally prove (either by hand or an algebraic manipulation program) the polynomials identity. The author's clear and intuitive writing manages to reveal in a natural way the tools which are necessary in order to evaluate the gains of randomness, capturing the reader's awareness and interest about the sections that follow since in these chapters the reader will be presented with these tools in a most instructive way ("by example"). These tools are the three axioms of probability theory, along with the basic notion of a *sample space* and *probability events*, as well as some (fully derived) lemmas with properties about the union and intersection of probability events (and, of course, Bayes' law). The authors analyze with great detail the solution to the polynomial identity problem while presenting two more applications: the verification of matrix multiplication and a randomized min-cut algorithm. The analysis of these problems is again very detailed and instructive.

Chapter 2. Having successfully linked probability theory and randomness with algorithmics in a clear and intuitive way, the authors move a small step further by introducing the concepts of a random variable and its expectation (conditional expectation too). Also, linearity of expectation is presented and proved in a step-by-step manner (no "left to the reader" comments in the book!). Of course, there are plenty of examples to illustrate the use of the newly introduced concepts and tools. Also, the proof are very detailed, to a point never encountered by me in other books I have come across to! This is great help to people not introduced to probability theory before and help them concentrate more on proving things themselves in the exercises in the end of the chapter. Moreover, the authors introduce the reader to the Binomial and Geometric distributions and combine the material presented in the chapter in tackling the *coupon's collector problem* and the analysis of the expected time of quicksort. The authors follow the "introduce when needed" approach which is very natural and avoids distraction to the reader. Thus, while tackling the coupon's collector problem they introduce the reader to the method of approximating a discrete summation with an integral (using, simply, a figure!) a tool that will be of much use in later chapters. Also, the two examples used by the authors complete the introduction to the two focal point of the book: randomized algorithms and probabilistic analysis. The first point was taken care of in the first chapter while the two examples introduce the reader to the second.

Chapter 3. In this chapter the authors complete the description and illustration of the basic tools of probability theory by introducing the concept of variance and generalized moments of a random variable. In addition, they derive Chebyshev's inequality and use it to bound the tails of the distribution of a random variable. The authors use many examples too in order to demonstrate the usefulness of these techniques in practice. They prove, for instance, the tightness of the average case analysis to the coupon collector's problem tackled in Chapter 2. In addition, the authors define the problem of selecting the median of a set of values using the sampling technique. The

explanation is very detailed and the derivations complete.
Chapter 4. The focus of this chapter is one of the most important tools for bounding the deviations from the expectation of the number of successes in Poisson distributed random variables. The authors derive the relevant basic inequalities and apply a stronger version of them to the problem of parameter estimation using an example involving mutations of DNA. Then the authors apply Chernoff bounds to a problem related to design of experiments in statistics as well as (advanced material) to the analysis of packet routing in Hypercube and Butterfly computer interconnection networks. Especially in the latter application (i.e. routing) the authors present so nicely the domain as well as the problem that one need not have any previous exposure to computer networks or parallel computing in order to understand the problem and its solution. Of course the analysis is likely to present difficulties to newcomers but the reader may go back later or simply skip the section without missing the main points of the chapter.
Chapter 5. One of the most important models in combinatorics is the *balls into bins*. It models a wide variety of situations of randomly pairing members of two classes of objects (balls and bins) in pure analogy with randomly placing a number of balls into a number of bins (the pairing). This is exactly the focus of this chapter where the balls into bins model is introduced along with various statistics of the model (e.g. maximum number of balls among all bins). As usual, the authors introduce the model through an example, the *birthday paradox* which fits perfectly the general ball and bins model. Then a variety of applications follows (bucket sort, leader election in distributed systems, sharper analysis of coupon's collector problem etc.) with (as usual again!) very detailed and illuminating description of solutions. The chapter ends in a very interesting and illuminating way: the introduction of random graphs within the framework of the balls and bins model! This is something that I have not come across with in the books I happen to have a look into. Usually, random graphs are introduced as a separate chapter with no reference to the rich theory of the balls and bins model.
Chapter 6. Here is something not usually found on probability theory or algorithmics textbooks: the *probabilistic method* and applications of it. Probability textbooks usually treat this subject as advanced (though its basics require minimal probability background) and algorithmics books treat it as a theoretical tool alone for proving the existence of a "needle in a haystack" which is not possible to find in reasonable time (not always the case, as the authors stress). To introduce the reader to the subject, the authors avoid the classical *Ramsey numbers* problem, which may seem a little confusing to a student, and use, instead, the problem of edge coloring of a clique using two colors so as to avoid a monochromatic sub-clique. The method gives readily conditions under which this is possible. Then the authors introduce the *expectation argument* and apply it to finding large cuts in graphs and maximal satisfiability of Boolean formulas. Then the derandomization technique is applied to the former problem, leading to a deterministic algorithm. The next subject is the sample-and-modify method while then the authors turn to the important *second moment method*. Then the authors prove the celebrated *Lovász Local Lemma* and give a very intuitive example of proof of existence of edge-disjoint paths in graphs when the paths share edges with a limited number of other paths. What gives a great value to this chapter, however, in the context of the book is the section that explains a general methodology (using an example from the problem of satisfiability of Boolean formulas) which can sometimes be used in order to locate the object guaranteed to exist by the probabilistic method. Thus, the probabilistic method chapter of the authors is a concise and very intuitive explanation of a methodology where one proves the existence of an object with the desired properties while, in some cases, the design of an efficient randomized algorithm to locate it

is also possible.
Chapter 7. By this time the authors have covered the necessary concepts in order to introduce the student to one of the most useful types of *stochastic processes*, the *Markov chains*, as well as the concept of a random walk. As it has always been the case in the previous chapters, the authors present the definitions in a clear intuitive way, along with examples. One thing that impressed my in the authors' presentation (which I don't recall to have seen it stressed in some other book) is the remark that in a Markov chain it is *not* true that the random variable corresponding to a time instance is independent from all all random variables corresponding to previous time instances except the one immediately preceding it. The subtlety is, simply, that this dependency is "lumped" into the dependency on the value of the immediately preceding random variable. The authors then introduce the transition probability matrix and draw on examples from the satisfiability problem to demonstrate the concepts. Afterwards the authors cover the usual things one expects in the study of Markov chains: classification of states and stationary distributions. The application they choose to demonstrate the concepts is related to a simple queuing system. The authors conclude with random walks on graphs and an interesting subtlety that arises in the study of Markov as exemplified by *Parrondo's Paradox* (lengthy discussion but, by all means, clear in presentation and instructive to read).
Chapter 8. In this chapter the authors introduce the reader to the concept of a *continuous* random variable and probability distribution functions. They give the relevant definitions as they naturally emerge from the example of a continuous roulette wheel. They also cover joint and conditional distribution functions and derive the relevant properties. Then they consider properties of the uniform and the exponential distribution functions as well as the ubiquitous in applications Poisson process. The chapter concludes with an extensive discussion of continuous time Markov processes and examples from queuing theory with different queuing policies.
Chapter 9. A book about randomness could not, of course, avoid studying randomness itself and its characterization! Thus, the authors aptly discuss in Chapter 9 the *entropy function* and its relationship to randomness, using the coin-flipping example. Then the authors discuss how one can extract random bits from high-entropy random variables and touch on how randomness of a sequence can be related to its compressibility. Finally, the authors introduce the reader to the famous Shannon's Coding Theorem and, generally, to coding theory (the exercise section amplifies in a very instructive way on this aspect).
Chapter 10. This chapter is devoted to the *Monte Carlo* and its numerous, as well as widely varying in nature, applications. Following the nice pedagogical methodology they have used so far, the authors start with an example in order to introduce the method: the estimation of π using random throws of points on a circle inscribed in a unit square. In a step-by-step manner, the authors introduce the reader to the concept of a good probabilistic approximation algorithm. In order to clarify the point that lies in the heart of the method, i.e. that the "positive" samples should form a size with considerable size in relation to the sample space, the authors solve the problem of counting the solution of a Boolean formula given in DNF form. They consider a bad naive approach and then an efficient one providing an excellent introduction to the issue of appropriately designing the sampling process in the application of the Monte Carlo method. The authors then move on to approximate counting and the Monte Carlo method relying on Markov Chains (a famous example of the latter one being the Metropolis algorithm discussed in the end of the chapter).
Chapter 11. One area in Markov Chain theory, usually not discussed in introductory probability theory books, is that of *coupling* and its appropriate design in order to prove fast convergence to

the stationary distribution. The authors define the *variational distance* and construct a coupling able to show that a Markov Chain converges fast to the stationary distribution (three examples are discussed: shuffling of a deck of cards, random walks on hypercubes, and finding fixed-size independent sets). The chapter concludes with two nice applications: the approximate sampling of proper colorings of graphs and the *path sampling* method.

Chapter 12. Apart from Markov Chains, *Martingales* are another class of stochastic processes usually covered at an advanced undergraduate level. The authors give the necessary definitions and very early in the chapter show the connection of margingales with the evaluation of functions on random graphs. Then they discuss stopping times of martingales and prove the *Martingale Stopping Theorem* and *Wald's Equation*. One of the most useful characteristics of martingales is that quantities associated with them seem to be highly concentrated. The authors state and derive the *Azuma-Hoeffding* tail inequalities that implies this concentration and provide numerous examples of its uses.

Chapter 13. This chapter is focused on the concept of *pairwise independent* random variables and its application to the construction of universal hash functions. The authors discuss the general concept of k-independent random variables, of which special case are pairwise independent random variables. This limited independence provides a powerful tool in algorithmics. The authors describe the construction of pairwise independent bits and how these can be put into use in derandomizing algorithms (an very nice and easy to understand example is given for the large cuts problem). Then the Chebyshev inequality for pairwise independent variables is proved and used for the construction of samplings of a solution space that requires less random bits in order to be effective (in the target approximation problem). Finally the authors introduce the concept of a universal hash function as well as perfect hashing and apply the chapter's material to a very interesting network problem, that of locating pairs of sender and receiver nodes that have exchanged packets over a predetermined packet limit.

Chapter 14. The last chapter (marked as advanced by the authors) is about a variation of the balls and bins framework. In this variation, each ball may fall into one of d randomly chosen bins. The ball is finally placed into the least full of the d bins, with ties broken arbitrarily. The main theorem about this model concerns the maximum load among all bins (a long and technical proof but very well explained). Then the authors show that the estimate given by the main theorem is actually tight, by giving a matching lower bound. Finally, the authors apply the model of balanced ball allocations to two important problems: hashing and dynamic resource allocation.

3 Opinion

This book is about the algorithmic side of randomness and how the tools of probability theory can assess randomness' effectiveness in the field of algorithmics. The book really embraces two kinds of audience: mathematical oriented and computer science oriented. The mathematical oriented audience will be introduced to the field of probability through a very carefully written and easy to understand text and will be benefited from exposure to the very interesting applications studied by the authors. The computer science oriented audience will acquire a firm background in probability theory and see how it can be applied to the design and analysis of probabilistic algorithms.

The text also includes numerous exercises. These exercises are placed at the end of each chapter and have the extremely important (in my opinion) feature of extending the material presented in the corresponding chapter with the student actually doing the extension herself/himself. The exercises cover both theoretical issues as well as a variety of computer science applications. In addition, the authors include at the appropriate places carefully designed and well explained programming assignments that will be of great help to the readers in improving their understanding about how randomness interacts with computation.

One minor negative point I would like to make is the non-existence of citations. However, to do justice to the authors, such a detailed bibliography would, possibly, confuse or frighten newcomers to the field, and would act as a distraction. Nevertheless, the authors cite a number of excellent textbooks in the end of the book for the reader who wants to pursue the subject further. I think, however, that there should be annotations to the books cited in order to help the reader find out what each book is about and how she/he could use it to further her/his understanding of the field. Some of them, although looking similar from a look on their titles are, nevertheless diverse in the way they approach the field of randomized algorithms and probability as well as in the material they cover.

In summary, this book certainly fills a gap in the relevant bibliography by providing an excellent, *easy* to follow introductory text on probability theory and its applications to algorithmics.

Review[5] of

Computational Techniques of the Simplex Method

Author of Book: István Maros

Kluwer Academic Publishers, 2003, 325 pages

Review by Brian Borchers

In the period after the second world war, one of the most important early applications of digital computers was in the area of linear programming. Problems that would now be recognized as linear programming problems had arisen during the war in logistics planning, but there were no generally effective algorithms for solving these problems until a young mathematician named George B. Dantzig developed the simplex algorithm. Dantzig, who recently passed away, is now widely recognized as the one person most influential in starting the new field of linear programming.

In its simplest form, a linear programming problem can be written as

$$\begin{array}{rrcl} \max & cx & & \\ & Ax & = & b \\ & x & \geq & 0 \end{array}$$

where x is a column vector with n elements, b is a column vector with m elements, c is a row vector with n elements, and A is a matrix with m rows and n columns.

A basis is a collection of m of the variables such that the constraints can be rewritten with the basic variables isolated on the left hand side. By dividing the vector of variables into m basic variables, x_B, and $n-m$ non-basic variables x_N, and similarly dividing the columns of A into A_B and A_N, the constraints can be rewritten as

$$A_B x_B + A_N x_N = b$$

or

$$x_B = B^{-1}b - B^{-1}A_N x_N.$$