The Book Review Column[1]
by Frederic Green

Department of Mathematics and Computer Science
Clark University
Worcester, MA 02465
email: fgreen@clarku.edu

In this column, we review the following 5 books:

1. **Turing Computability: Theory and Applications**, by Robert Soare. A treatise on the theory of computability, by one of its leading expositors, reviewed by Bill Gasarch.

2. **Analysis of Boolean Functions**, by Ryan O'Donnell. A unified graduate-level approach to some deep foundations, including (for me) unusual and diverse applications (e.g., to social choice, statistics, noise stability, and hardness results). Reviewed by Daniel Apon.

3. **Distributed Systems: An Algorithmic Approach (2nd Edition)**, by Sukumar Ghosh. This book, which lays out both the theoretical groundwork and incorporates real-world applications, will be of value both to students and practitioners. Reviewed by Ramon de Vera Jr.

4. **The Golden Ratio and Fibonacci Numbers**, by Richard A. Dunlap. An introduction to diverse aspects of one of our favorite ratios and sequences, including such fascinating topics as quasiperiodic tilings of the plane. Reviewed by Michaël Cadilhac.

5. **The Fascinating World of Graph Theory**, by Arthur Benjamin, Gary Chartrand and Ping Zhang. A highly reader-friendly and fun introduction to graph theory. Reviewed by Frederic Green.

BOOKS THAT NEED REVIEWERS FOR THE SIGACT NEWS COLUMN

Algorithms

1. *Distributed Systems: An algorithmic approach (second edition)*, by Ghosh
2. *Tractability: Practical approach to Hard Problems*, Edited by Bordeaux, Hamadi, Kohli
3. *Recent progress in the Boolean Domain*, Edited by Bernd Steinbach
4. *A Guide to Graph Colouring Algorithms and Applications*, by R.M.R. Lewis
5. *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, Ariel D. Procaccia, Eds.

Programming Languages

1. *Selected Papers on Computer Languages* by Donald Knuth

Miscellaneous Computer Science

1. *Algebraic Geometry Modeling in Information Theory* Edited by Edgar Moro
2. *Communication Networks: An Optimization, Control, and Stochastic Networks Perspective* by Srikant and Ying
3. *CoCo: The colorful history of Tandy's Underdog Computer* by Boisy Pitre and Bill Loguidice
4. *Introduction to Reversible Computing*, by Kalyan S. Perumalla
5. *A Short Course in Computational Geometry and Topology*, by Herbert Edelsbrunner
6. *Network Science*, by Albert-László Barabási
7. *Actual Causality*, by Joseph Y. Halpern

Computability, Complexity, Logic

1. *The Foundations of Computability Theory,* by Borut Robič
2. *Models of Computation*, by Roberto Bruni and Ugo Montanari
3. *Proof Analysis: A Contribution to Hilbert's Last Problem* by Negri and Von Plato.

Cryptography and Security

1. *Cryptography in Constant Parallel Time,* by Benny Appelbaum
2. *Secure Multiparty Computation and Secret Sharing*, Ronald Cramer, Ivan Bjerre Damgård, and Jesper Buus Nielsen
3. *A Cryptography Primer: Secrets and Promises*, by Philip N. Klein

Combinatorics and Graph Theory

1. *Finite Geometry and Combinatorial Applications*, by Simeon Ball
2. *Introduction to Random Graphs*, by Alan Frieze and Michał Karoński
3. *Erdös–Ko–Rado Theorems: Algebraic Approaches*, by Christopher Godsil and Karen Meagher
4. *Words and Graphs*, by Sergey Kitaev and Vadim Lozin

Miscellaneous Mathematics and History

1. *The Magic of Math*, by Arthur Benjamin
2. *Professor Stewart's Casebook of Mathematical Mysteries* by Ian Stewart

Review of[2]
Turing Computability: Theory and Applications
by Robert Soare
Publisher: Springer, 2016
$54.00 hardcover, 253 pages

Reviewed by
William Gasarch `gasarch@cs.umd.edu`
University of Maryland, College Park, MD

1 Introduction

We discuss the basic objects of computability theory.

A function f is *computable* if there exists a Turing machine (a Java program also works) that will, on input x, output $f(x)$. A set is *computably enumerable (c.e.)* if there exists a Turing machine which, on input x, will (a) eventually halt if $x \in A$ and (b) never halt if $x \notin A$. If given B you can compute A (formally with an oracle Turing machine) then we write $A \leq_T B$. If $A \leq_T B$ and $B \leq_T A$ then we write $A \equiv_T B$. This is an equivalence relation. The equivalence classes are called *Turing degrees*.

Computability theory has its origins in logic; however, it has become a field unto itself. This book discusses some of the historical ties to logic as well as present the field on its own terms.

2 Is this a Second Edition of Soare's Prior Book?

NO! Let's get one thing out of the way first: This book is emphatically *not* a second edition of Soare's 1987 book *Recursively Enumerable Sets and Degrees: A Study of Computable Functions and Computably Generated Sets*. The new book does not have $0'''$-priority arguments, nor $0''$-priority arguments. It *does* have $0'$-priority arguments (also known as finite injury priority arguments). The old book was very focused on r.e. (now called c.e.) sets and degrees. The new book takes a much wider view of computability theory.

3 Summary of Contents

The book is in five parts each of which has several chapters.

Part I: Foundations of Computability

This part has seven chapters. The first chapter (*Defining Computability*) is a combination of history and math. It describes how the modern notions of computability came about. The second chapter (*Computably Enumerable Sets*) covers the basics of c.e. sets. There is one unusual section on Lachlan Games which seems to be a method that encompasses many constructions of c.e. sets. Later in the book he shows how to phrase finite injury priority arguments in terms of Lachlan games. I am curious if one can phrase $0''$ or $0'''$ arguments this way.

The third and fourth chapters (*Turing Reducibility* and *The Arithmetic hierarchy*) contain standard material on these topics as well as an assortment of other topics: (1) n-c.e. sets, (2) the low and high hierarchy, and (3) the low basis theorem (later there is an entire part on trees).

The fifth, sixth, and seventh chapters (*Classifying C.E. sets*, *Oracle Constructions and Forcing*, and *The Finite Injury Method*) cover material that is more advanced – so much so that I wonder why it is in a Part called *Foundations of Computability*. This is not an objection, it is an observation and may also reflect a difference in taste between me and the author. Some topics are (1) simple, hypersimple, and hyperhypersimple sets, (2) forcing arguments to construct a set of Turing degree strictly in between decidable and Halting, and (3) a finite injury priority arguments to construct a c.e. set of Turing degree strictly in between decidable and Halting.

The entire Part is well written with a sense that the author has a broad prospective on the material and has thought deeply about it. This comes through the most when he talks about the history of the field, which is more the subject of Part V.

Part II: Trees and Π^0_1 Classes

A *tree* is a subset of $\{0,1\}^*$ that is closed under prefix. This really is a tree in the way you usually think of it. A set of infinite strings is Π^0_1 if it consists of the branches of an infinite tree.

Why are trees important in logic? Consider the following tree T:

1. Every number codes a statement using (say) the usual logic symbols and also $[+, \times, <, 0, 1]$. The quantifiers are over the natural numbers.

2. A string $\sigma = b_0 \cdots b_n$ codes the statement (1) all i such that $b_i = 1$ are true, and (2) all i such that $b_i = 0$ are false.

3. Fix an axiom system (e.g., Peano Arithmetic). A string σ is in T if none of the assertions of σ (as defined by point 2) have been proven false by a $|\sigma|$-step derivation.

Studying models of Peano Arithmetic is equivalent to studying branches of T. This tree is decidable. Hence studying decidable trees will yield theorems about models of Peano Arithmetic. For example, the low basis theorem states that every c.e. tree has a low branch. Hence there is a low model of Peano Arithmetic.

This chapter first studies trees and then applies them to PA. This is very nice since it's good for any field of math, once it proves something, to go back to the original motivation. Soare proves far more about trees then one needs for PA, but what is proved may be useful at some later time.

This chapter then discusses how trees, randomness, and measure interact. This is an immense topic that the author can only touch on briefly. We give one example without defining terms formally: if a Π^0_1 class has measure 0 then it has no Martin-Löf Random sets.

Part III: Minimal Degrees

This Part has two chapters. A degree is *minimal* if for any A in that degree (1) A is not decidable, and (2) all $B <_T A$ are decidable. In this chapter the author shows that there is a minimal degree that is $\leq_T 0''$ and then that there is a minimal degree $\leq_T 0'$. The presentation is very educational in that the author says, at every step of the way, why he is doing what he is doing. Frankly, that is true for most of the book, but here it stands out more since he spends two chapters doing what other texts do in 5 pages.

Part IV: Games in Computability Theory

My wife has commented that math games are not fun games. This part will not change her mind.

Here is an example: Let A be a subset of $\{0,1\}^\omega$. The players are Alice and Bob. Alice goes first. Alice picks a bit a_0. Bob picks a bit b_0. Alice picks a bit a_1. Bob picks a bit b_1. etc. If the string $a_0b_0a_1b_1\cdots$ is in A then Alice wins. If not then Bob wins. The question is: for which sets A does Alice have a winning strategy. Martin proved that all Borel sets have winning strategies. The Axiom of Determinacy ("AD") declares that all sets have a winning strategy. It turns out that AD is an interesting alternative to the Axiom of Choice. However, that is not the direction that the book takes. Instead, it talks about how computable a strategy has to be in certain cases.

This chapter also talks about Lachlan games. In particular, finite injury priority arguments are rewritten as Lachlan games.

Part V: History of Computability

This is a history of computability from Hilbert (1900) until Post's problem was solved (1950's). It is masterful with enough details to be interesting, but not so much as to be boring. The author is interested in correcting some misconceptions such as (1) Church's Thesis is not merely a "thesis" it's... more than that, and (2) Recursive Function Theory is a terrible name for the field and historically inaccurate.

With regard to (1): while people still call it *Church's Thesis*, they really do believe it's true. With regard to (2): he has won the battle – people now call the field *computability*. But for the author this is not just a name change, nor even just a clarification to others of what we do (though that is part of it). It is a correct terminology based on history.

4 Opinion

Who should read this book? This question breaks down into several questions:

What level is this book at? Someone who knows no computability (but is good at math) could pick up this book, read it, and learn something. It would be tough going but they could. This could certainly be a classroom text on the topic, although one would have pick which chapters to cover carefully.

Is the book well written? Absolutely yes. Robert Soare knows the topic, knows the origins, knows not just how the proofs go, but why they go that way, and shares that with us.

People in field X should read this book. Fill in the X. In 1980 I would have said that logicians and computer scientists (especially theorists) should learn computability theory (I had a course from Michael Stob out of a preprint of Soare's old book in 1980.) What about in the year 2016?

1. *Computer Scientists:* Since 1980 computer science theory has shifted away from logic and towards combinatorics. Most theorists today have never seen a priority argument! I am *not* going to say *in my day we all did priority arguments while walking to school, uphill both ways, without shoes in the snow*. In fact, I think it's healthy for a field to evolve and change and not be stuck where it was 30 years ago. And the oracle results prove that the shift was needed. But see next note.

2. *People interested in Randomness:* This interest runs through logic, computer science, physics, and probably other fields. There has been a huge increase of interest in randomness. There is an 800 page book on the subject: *Algorithmic Randomness and Complexity* by Downey and Hirschfeldt. Computability theory is at the heart of the study of randomness. Hence people who read (or want to read) Downey/Hirschfeldt should also read Soare.

3. *Logicians:* Given a theorem whose proof is nonconstructive (defined in a variety of ways) is there a more constructive proof? How to measure this? The Reverse Math program deals with these issues and computability theory is an important part of that.

Review of
Analysis of Boolean Functions
by Ryan O'Donnell
Cambridge University Press, 2014
423 pages, Hardcover (Amazon: \$45 - \$50)
Request (Free!) Textbook Download via Author's Webpage:
`http://get.analysisofbooleanfunctions.net`

Review by
Daniel Apon, `dapon@cs.umd.edu`
University of Maryland, College Park, MD

1 Introduction

Boolean functions are perhaps the most basic objects of study in theoretical computer science. They arise naturally in complexity theory, cryptography, learning theory, combinatorics, statistics, and social choice theory – among many other areas of study. The field of analysis of Boolean functions seeks to understand these functions via *analytic* methods, such as their Fourier transform. Novel (and uniquely rewarding) concepts from this field include studies of the *noise stability* and *(partial) derivatives* of Boolean functions as well as the *hypercontractivity* of random variables.

At a high level, the textbook *Analysis Of Boolean Functions* is a collection of tools to – lo and behold! – analyze Boolean functions. Since Boolean functions are a virtually *universal* concept across different areas of computer science and mathematics, the concepts collected in this book originate from many different places as well. A particular strength of this book is its ability to "tie together" ideas from *all* of these areas into a single, step-by-step story with techniques that naturally evolve from one step to the next. (Indeed, this guarantees that at least *some part* of this book is "new" to most people!)

O'Donnell's book introduces each concept in the area, beginning with simple (though rigorous) definitions and ending with advanced topics such as isoperimetry and analysis of Gaussian functions. Each chapter includes a nice "highlight application." Typically, the "highlight" is a major, landmark theorem in the theme of the chapter: for example, the Goldreich-Levin theorem from cryptography is the highlight of the chapter on Learning Theory.

Each of the 11 chapters concludes with a highlight section, followed by exercises, additional notes, and numerous references to the original literature relevant to the chapter. The book has more than 500 exercises in total, making it suitable for a one-semester, advanced graduate course.

2 Summary

In the following, I give a chapter-by-chapter abbreviated overview of *Analysis of Boolean Functions.* I also *state* the "highlight theorem" of a few chapters. The goal is to provide a *taste* of the material, but to understand the chapters or their highlights in detail, you'll have to check out the book!

– Chapter 01: Boolean functions and the Fourier expansion. The book begins by describing the basics of Boolean function analysis. Specifically, the view presented of the Fourier expansion of a Boolean function is that of its representation as a real multilinear polynomial. (Insights from harmonic analysis are deferred to Chapter 3.)

◦ **Highlight #1: Almost Linear Functions and the BLR Test.** Imagine you have "black-box" access to a function f, and you want to verify that it is indeed linear.
The **BLR Test** is:

- *Choose* $\mathbf{x} \leftarrow \mathbb{F}_2^n$ *and* $\mathbf{t} \leftarrow \mathbb{F}_2^n$ *independently.*
- *Query* f *at* $\mathbf{x}$, $\mathbf{y}$, *and* $\mathbf{x} + \mathbf{y}$.
- *"Accept" if* $f(\mathbf{x}) + f(\mathbf{y}) = f(\mathbf{x} + \mathbf{y})$.

If the BLR Test accepts f with high probability, then f is "close" to being linear.

– Chapter 02: Basic concepts and social choice. In this chapter, O'Donnell introduces a number of important basic concepts including influences and noise stability. The motivations for these topics come from social choice theory and election design. The chapter concludes with Kalai's Fourier-based proof of Arrow's Theorem.

◦ **Highlight #2: Arrow's Theorem.** *Consider a 3-candidate Condorcet election using* $f : \{-1, 1\}^n \to \{-1, 1\}$. *Under the impartial culture assumption, the probability of a Condorcet winner is precisely* $\frac{3}{4} - \frac{3}{4}\mathbf{Stab}_{-1/3}[f]$.

– Chapter 03: Spectral structure and learning. This chapter examines the complexity of Boolean functions by exploring the "complexity" of the function's Fourier spectrum. For example, functions with sufficiently simple Fourier spectra can be efficiently learned from examples. Vital concepts include understanding the location, magnitude, and structure of a Boolean function's Fourier spectrum.

◦ **Highlight #3: The Goldreich-Levin Theorem.** *Given query access to a target* $f : \{-1, 1\}^n \to \{-1, 1\}$ *as well as input* $0 < \tau \leq 1$, *there is a* $\mathrm{poly}(n, 1/\tau)$*-time algorithm that with high probability outputs a list* $L = \{U_1, ..., U_\ell\}$ *of subsets of* $[n]$ *such that:*

- $|\widehat{f}(U)| \geq \tau \implies U \in L$;
- $U \in L \implies |\widehat{f}(U)| \geq \tau/2$.

– Chapter 04: DNF Formulas and small-depth circuits. This chapter explores Boolean functions written as small DNF formulas and constant-depth circuits. Besides being a natural model of computation, these representation classes are close to the limit of what complexity theorists currently "understand" (i.e., can prove explicit lower bounds for). One reason for this is that functions in these classes have strong Fourier concentration properties.

– Chapter 05: Majority and threshold functions. Here, O'Donnell explores linear threshold functions, their generalization to higher degrees, and – their best example – the *majority* function. This leads to the Central Limit Theorem and the introduction of Gaussian random variables. Using these tools, it is shown that the Fourier spectrum of the $\mathbf{Maj}_n$ function "converges" as $n \to \infty$. Other concepts include: degree-1 Fourier weight, noise stability, and total influence of general linear threshold functions.

– Chapter 06: Pseudorandomness and $\mathbb{F}_2$-polynomials. This chapter explores various definitions for *pseudorandomness for Boolean functions* – that is, properties of a particular function that are characteristic of a randomly sampled function. Several of the results in the chapter involve interplay between the representation of $f : \{0,1\}^n \to \{0,1\}$ as a polynomial over the reals and its representation as a polynomial over $\mathbb{F}_2$.

– Chapter 07: Property testing, PCPPs, and CSPs. In this chapter, several interesting and intertwined concepts are studied: (i) property testing, (ii) probabilistically checkable proofs of proximity (PCPPs), and (iii) constraint satisfaction problems (CSPs). These notions coincide in the problem of determining whether some unknown Boolean function is a *dictator.*

The study proceeds as follows: First, the BLR Test is extended to give a 3-query property testing algorithm for the class of dictator functions. Then, we get a 3-query testing algorithm for any property, under some constraints. Next, CSPs – or, "string-testing algorithms" – are introduced. Finally, dictator tests are translated into *computational complexity results for CSPs;* for example:

◦ **Highlight #7: Håstad's Hardness Theorems.** *For any constant $\delta > 0$, it is NP-hard to $(\frac{7}{8} + \delta, 1)$-approximate Max-E3-Sat or to $(\frac{1}{2} + \delta, 1 - \delta)$-approximate Max-E3-Lin.*

– Chapter 08: Generalized domains. Up to this point in the book, the functions in question have been $f : \{0,1\}^n \to \mathbb{R}$. What about $f : \{0,1,2\}^n \to \mathbb{R}$? This chapter observes that, in fact, very few of the ideas so far depend on the domain being $\{0,1\}^n$. Rather, it mostly depends on the domain being a product probability distribution. Almost all prior analysis carries over to the case $f : \Omega_1 \times ... \times \Omega_n \to \mathbb{R}$ where the domain has a product probability distribution $\pi_1 \otimes ... \otimes \pi_n$.

There are two main exceptions: the "derivative" operator D_i does not generalize to the case when $|\Omega_i| > 2$ (though the Laplacian operator L_i does); and, the important notion of hypercontractivity (introduced in the next chapter) depends strongly on the probability distributions π_i.

– Chapter 09: Basics of hypercontractivity. In 1970, Bonami proved a central result: *The Hypercontractivity Theorem* (formal statement omitted). The theorem, however, is quite dense to break into and understand. This chapter provides some special cases of the theorem, which are easier to understand, easier to prove, and which encompass almost all of the greater theorem's uses. The full theorem is deferred to Chapter 10.

◦ **Highlight #9: The Kahn-Kalai-Linial Theorem.** Consider a 2-candidate, n-voter election using a monotone voting rule $f : \{-1,1\}^n \to \{-1,1\}$. Say that one of the candidates is able to secretly bribe k voters. We want to design unbiased voting rules f that minimize the effect of the bribed k-coalition.

The KKL Theorem states (as a special case) that every unbiased function f has maximum influence $\Omega(\frac{\log n}{n})$. That is: *For any* $f : \{-1,1\}^n \to \{-1,1\}$,

$$\mathbf{MaxInf}[f] \stackrel{\text{def}}{=} \max\{\mathbf{Inf}_i[f] : i \in [n]\} \geq \mathbf{Var}[f] \cdot \Omega\left(\frac{\log n}{n}\right).$$

– Chapter 10: Advanced hypercontractivity. In this chapter, the proof of the Hypercontractivity Theorem is completed for uniform ± 1 bits. Then it is generalized to statements about arbitrary product probability spaces; that is, functions $f \in L^2(\Omega^n, \pi^{\otimes n})$.

– Chapter 11: Gaussian space and Invariance Principles. The final destination of the textbook is the proof of Mossel, O'Donnell, and Oleszkiewics's seminal Majority is Stablest Theorem. In general, for *any* fixed volume, the theorem states that "Hamming ball indicators have maximum noise stability among all small-influence functions." More formally, we arrive at the final highlight of the book:

◦ **Highlight #11: The Majority is Stablest Theorem.** *Let* $f : \{-1,1\}^n \to [0,1]$. *Suppose that* $\mathbf{MaxInf}[f] \leq \epsilon$, *or more generally, that* f *has no* $\left(\epsilon, \frac{1}{\log(1/\epsilon)}\right)$*-notable coordinates. Then for any* $0 \leq \rho < 1$,

$$\mathbf{Stab}_\rho[f] \leq \Lambda_\rho(\mathbf{E}[f]) + O\left(\frac{\log\log(1/\epsilon)}{\log(1/\epsilon)}\right) \cdot \frac{1}{1-\rho}.$$

3 Opinion

Reading through the material in this book, and understanding it reasonably well, took a long-term commitment from me. This book undoubtedly contains a challenging sequence of ideas to internalize. On the other hand, learning this material is *worth the investment* of time and effort. Indeed, some of the perspectives and intuitions in this book clearly helped me in recent research of my own (read a book, write a review, get a research insight – not a bad bargain!).

One thing that greatly assisted me personally (and must be mentioned) is the author's extensive website for the book, which can be found at `http://analysisofbooleanfunctions.net`. This not only contains a copy of the book's content broken into sections/lectures, but it has links to some great video resources. Specifically, one can find videos for 23 lectures from Carnegie Mellon's Fall 2012 course using the book. Also available on the book's website, there is a "live blog" style recording of a Simons Symposium from 2014 on analysis of Boolean functions (with great links to current, active topics in this area).

The applications of the ideas in this book are plentiful and diverse, and O'Donnell does an excellent job of leading the reader from one viewpoint to the next. I found it especially enjoyable to see theorems that I'm personally familiar with as a cryptographer, such as the Goldreich-Levin theorem, placed alongside other things I didn't know as well, like Arrow's theorem from social choice – with everything woven into a single, consistent story. I suspect other "fresh readers" will similarly find parts of this book that they recognize, and others they don't. The relationships exposed between these ideas should be of interest to everyone.

Altogether, I highly recommend that you take a glance at *Analysis of Boolean Functions.* As mentioned previously, the majority of the textbook's content is already online, and a full copy of the book can be requested via the author's website. Given the unique advantages (and enjoyment!) of the ideas here, it's worth setting aside some time to read through and learn for yourself.

Review of
Distributed Systems: An Algorithmic Approach (2nd Edition)
by Sukumar Ghosh
Published by CRC Press, 2014
Hardcover, 500 pages

Review by
Ramon de Vera Jr. (rdevera@live.com)
Product Engineering Software Group
Micron Technology Inc.

1 Overview

For distributed systems, barring academic work on the subject, the exposure of people to the concept is quite often through examples or implementations. People use various services on social, peer-to-peer, and mobile networks ubiquitously on a day-to-day basis. There is no mind to how these distributed systems do what they do - other than the occasional recourse to question the sanity of how well these distributed systems are running when their services get interrupted.

The book delves into distributed systems from the definition of what a distributed system is and its relevance and use, to concepts relating to its components (topology, communication, failure handling, and so on) up to coverage of prevalent distributed systems to date – with a twist. The author focuses on the analysis and discussion of the algorithmic/theoretical models that make up distributed systems and at the end of book relates them to real-world applications or implementations.

The author states that the intended audience of the book was for senior undergraduates to first-year graduate students. But, upon review of the book, that specification might be a little narrower than it should be. Developers that are out in the 'real world' working on distributed systems will benefit from books such as this as they serve as handy references and guides for ideas and concepts they should be considering when performing their jobs. There is much to be said for applying more informed decisions to the development of existing distributed systems in use in the industry right now.

2 Summary of Content

The book is well organized and follows a very logical progression from section to section. People who are less familiar with the concepts can build their foundation of the topics at a good pace. The breadth of the topics in discussion does not lend itself to a very in-depth coverage. But what is covered is sufficient for its purpose as a foundational work. More experienced practitioners can skip some of the initial sections or use them as references.

The style that was used in the book foregoes some formalisms, but instead goes for more direct and straightforward discussions of the algorithms and their corresponding characteristics backed by proofs and theorems. The evidence presented to reinforce the discussion are appropriate to the topics.

There are five sections to the book: "Background Materials," "Foundational Topics," "Important Paradigms," "Faults and Fault Tolerant Systems," and "Real World Issues."

The first section, "Background Materials," as is the case with such things, covers basic definitions and introductory materials regarding Distributed Systems in order to ground the reader before foundational concepts are discussed. The survey and subsequent overview of inter-process communication techniques usable for distributed systems is a good way to frame the perspective of readers about what is the architecture that supports distributed systems.

The second section, "Foundational Topics," covers models for communication, notation for distributed algorithm and correctness criteria and proofs. The section ends with a discussion of the notion of time in a distributed system and its impact.

The third section of the book, named "Important Paradigms," includes discussions of mutual exclusion, distributed snapshots, global state collection, graph algorithms, and coordination algorithms. As part of the discussion of these foundational concepts for distributed systems, there are discussions of theorems, lemmas, and proofs that underline the behavior or characteristics of the algorithms discussed.

As a representation of how the author goes through a discussion topic, let's look at the chapter covering Graph Algorithms (Chapter 10). The discussion starts with indicating how graph theoretic solutions have a number of applications for distributed systems – mainly in the aspects of communication and networking. Examples of issues with routing messages across nodes and maintaining routing tables, and even issues with transmitting messages in a possibly changing topology, are given to initiate the discussion of where graph theory concepts will be applicable. The section regarding Routing Algorithms, for example, goes through different routing algorithm options. The section starts with a discussion of the computation of the shortest path – a pseudo-code of the algorithm (Bellman-Ford). Behavior characteristics are identified via a lemma and a corresponding proof, and analysis of the complexity of the algorithm. And the discussion moves forward with the Chandy-Misra Modification for the Shortest Path Algorithm. And this new section continues with a presentation of the pseudocode and its benefits and so on, into the other sections.

The fourth section, "Faults and Fault-Tolerant Systems," focuses on how handling different fault classifications can influence the algorithms that we use to do distributed consensus, distributed transactions, group communication, replicated data management. At the end of the section, self-stabilizing systems are covered as well.

The last section of the book, "Real-World Issues," covered what are considerations that should be foremost in a distributed system developer's mind when a distributed system is implemented for a particular problem domain in 'the real world.' The author focuses on examples of distributed systems for discussion that are relevant to current times and which pose interesting problems to resolve – also shedding light on how challenging distributed systems implementations can be.

3 Conclusion

The book lives up to its plan of providing a comprehensive reference for the fundamental concepts behind distributed systems. The progression of, and the coverage itself, was well done. The theoretical underpinnings were given their due discussion, emphasizing the point of the differences in the algorithms and when to use them – a point that at times escapes us in 'the real world.' The notations are not too far removed from other usual notations for computational models so they are familiar and not too daunting. The examples used were relevant and current. The examples regarding social networks and centrality, clustering coefficients, modeling, and community detection are particularly interesting.

This book now lies at this reviewer's desk at his office, where it should lie in wait so that as the need arises we can whip it out to counter dubious claims of 'excellent real world design decisions.'

Review of
The Golden Ratio and Fibonacci Numbers
by Richard A. Dunlap
World Scientific, 1997
172 pages, Hardcover $58, ebook $46

Review by
Michaël Cadilhac <michael@cadilhac.name>
WSI, Universität Tübingen

1 Overview

Professor Dunlap, a physicist specialized in the study of materials, presents a few tokens of the pervasiveness of the golden ratio and the Fibonacci sequence in our mathematical, physical, and biological world. Mainly focused on geometrical facts and their applications, the book contains a fantasy mix of the author's favorite curiosities around this topic. If the book is by no means exhaustive—and it has no pretension to be—it is a nice potpourri that anyone can enjoy, while claiming to not be addressed to any particular profile.

2 Summary of Contents

Chapter 1: Introduction. Using simple mathematical facts as well as historical and artistic considerations, the author introduces the golden ratio as a natural and compelling value.

Chapter 2: Basic properties of the golden ratio. Some links between the golden ratio and sequences, in particular those of Lucas and Fibonacci, are presented. More importantly, the name "ratio" is explained, using geometry. Indeed, if a segment AB of length 1 is divided into two segments AC and CA such that $AB/AC = AC/CA$, then this ratio is the golden ratio.

Chapter 3: Geometric problems in two dimensions. Regular polygons are shown to have some dimensions related to the golden ratio. The golden gnomons[3] are introduced.

Chapter 4: Geometric problems in three dimensions. A similar study is carried in three dimensions.

Chapter 5: Fibonacci numbers. The classical rabbit problem is introduced, and some natural occurrences of the Fibonacci numbers are explored. Binet's formula is presented:

$$F_n = \frac{1}{\sqrt{5}} \times \left(\tau^n - (-\tau)^{-n}\right) \ ,$$

with F_n the n-th Fibonacci number, and τ the golden ratio. Necessary and sufficient conditions are given to "recognize" a Fibonacci sequence of rabbits (given by the rules "Adult $\mapsto$ Adult, Baby" and "Baby $\mapsto$ Adult").

[3] According to the Oxford dictionary, a gnomon is the part of a parallelogram left when a similar parallelogram has been taken from its corner.

Chapter 6: Lucas numbers and generalized Fibonacci numbers. Some generalizations of the Fibonacci recurrence relation are studied, leading to Lucas numbers, the Tribonacci sequence, and the like.

Chapter 7: Continued fractions and rational approximants. Continued fractions are introduced, in particular that of the golden ratio.

Chapter 8: Generalized Fibonacci representation theorems. This chapter focuses on representing numbers in "bases" that are extracted from Fibonacci-like sequences. For instance, using $F_2 = 1, F_3 = 2, F_4 = 3, F_5 = 5, \ldots$ as a base, the number 10 can be written 10010 or 1110, as $10 = F_6 + F_3 = F_5 + F_4 + F_3$. Using some extra rules, this representation can be chosen to be unique.

Chapter 9: Optimal spacing and search algorithms. Golden ratio search, an algorithm to search for an extremum of unimodal functions, is presented. This relies in part on the properties of iterating rotations around the unit circle by transcendental multiples of π.

Chapter 10: Commensurate and incommensurate projections. Herein, two-dimensional latices are projected onto lines, and the cut angle determines the nature of the segments defined by the projections on the line. In particular, when the cut angle is an irrational tangent ($\tan\theta = 1/\tau$), this is an *incommensurate projection*, and the grid is projected down the line in such a way that the line is cut in long (L) and short (s) segments that reproduce the sequence of Adult and Baby rabbits.

Chapter 11: Penrose tilings. These tilings are created by incommensurate projections of higher dimensions. Generalizations of the concepts of the previous chapter are presented. This gives rise to *quasiperiodic* tilings, i.e., tilings that can fill the space without exhibiting translational symmetry.

Chapter 12: Quasicrystallography. Quasiperiodic tilings are used to arrange atoms to create crystal-like structures. The main difference with crystals is the lack of translational symmetry. Their refraction properties are studied.

Chapter 13: Biological applications. In this chapter, the classical links between the Fibonacci sequence and flowers, marine animals, viruses, pine cones, pineapples, sunflowers, and mollusks are presented.

3 Opinion

On the intended audience. I have a hard time deciphering to whom this book is written. On the one hand, I learned a few things reading it, both in geometry and physics. I feel that the writing style is very dry, if not technical, and the book lacks the usual entertaining side of scientific popularization. The lack of a compelling driving force makes it hard to stay interested, save for being a scientist at heart. The reader is expected to be versed in some terminology that were way beyond my (arguably limited) knowledge.

On the other hand, the lack of proofs is blatant, appealing to the nonscientist. There is, altogether, no prior knowledge required, and the mathematical facts are usually within grasp (provided the reader knows about basic geometry, trigonometry, transcendental numbers, ...). The sheer number of tables giving numerical values may help the reader to understand how values go, although it is in my opinion generally overdone (e.g., Chapter 8 is 6.5 pages, 3 of which are just numerical values). The succinctness of the chapters allows the reader to pick up the book, read a chapter, and let it sit for a few days, to digest the content.

On the content. The choice of topics is the author's own, and I believe, not to be argued against. It is certainly a diverse and healthy mix, that may have aged a tad bit since its publication some twenty years ago. In a way, the content can be seen as being steered toward quasicrystals, stopping along the way to discover some nuggets of mathematical knowledge.

On the form. The book contains a large number of drawings that do come as a relief in the driest parts. Their position in the text, on the physical pages, is often amiss, forcing the reader to go back and forth, and search for the correct figure number. This is a more general problem in fact, since equations are cross referenced throughout the book without a single aid to find them (e.g., when reaching page 56, referencing "Eq. (2.1)," appearing page 7, is quite a stretch). More than once was I tempted not to look at a figure or an equation as I felt it was "too far."

Final opinion. I would recommend this book to physics students, not as an object of study, but as mathematical entertainment. I do not think laypeople would enjoy it, because of the writing style and the choice of topics. I feel that the lack of proofs and the over abundance of numerical values would discourage a mathematician. All in all, the text end up occupying a small niche, despite a (overly?) general title.

Review of[4]
The Fascinating World of Graph Theory
by Arthur Benjamin, Gary Chartrand and Ping Zhang
Princeton University Press, 2015
322 pages, Hardcover, $29.95

Review by
Frederic Green `fgreen@clarku.edu`
Department of Mathematics and Computer Science
Clark University, Worcester, MA

1 Introduction

In central Spain in mainly rain
Three houses stood upon the plain.

The houses of our mystery
To which, from realms of industry

Flowed pipes and wires to light and heat
And other pipes with water sweet.

The owners said "What these things cross
Burn, leak or short, we'll suffer loss,

So let a graphman living near
Plan each from each to keep them clear."

Tell them, graphman, come in vain,
They'll bear one cross that must remain.
Explain the planeness of the plain.

– Blanche Descartes, Waterloo, 1997

"The Fascinating World of Graph Theory" looks, feels (dust jacket and all), and often reads like a novel, or an entertaining history. But it's actually a thinly-disguised textbook on graph theory. Graph theory is not an easy subject, and it's too easily presented so as to appear dull and uninteresting. But its main ideas and problems are not only deep and fascinating, but also fun and easily described to just about anyone. What a great idea! An especially accessible book about an especially exciting and accessible subject!

Everyone (depending on how you define "everyone") knows the famous story of the bridges of Königsberg and how, in solving it, Euler invented graph theory and topology in one fell swoop. But graph theory has a rich history underlying all of its most prominent problems. Some of the mathematicians involved were famous, some virtually unknown, but the stories are all quite engaging.

This branch of mathematics seems to have attracted quite a colorful cast of characters. One (or more) of them took up the pen-name Blanche Descartes to write serious mathematics as well as whimsical mathematical verses such as the one quoted above, attributed to this same "Descartes." Anyone acquainted enough with graph theory probably knows who Blanche Descartes really was. Whether or not you manage to resist the temptation to immediately google the answer, I urge you to read this book by Benjamin, Chartrand and Zhang, and find out about not only "Descartes" but many others in this cast of characters and, most importantly, their mathematics.

2 Summary of Contents

Here, in some detail, is what the individual chapters cover.

- Chapter 1: We encounter problems like five princes, three houses/utilities (cf. the verse opening this review), three friends, matching, Königsberg, four-color, and the polyhedron problem. This is followed at the end by a formal definition of graphs.

- Chapter 2: "Classifying Graphs" first discusses beauty in math generally, listing things like the Euler identity, infinitude of primes etc. Graph problems are listed among them. Then basic results are presented on irregular and regular graphs, ending with the graph reconstruction problem.

- Chapter 3: "Analyzing Distance," about properties related to connectivity, and bipartite graphs. Locating sets are discussed on pgs 54-55: One example of many where the description of a real-life application (sensors in rooms that only detect distance, used to detect the location of a node) serves to "anthropomorphize" the problem. Also the "Lights Out Puzzle." The story of the Erdös number and other tidbits like the papers of G. W. Peck and Blanche Descartes acquaint the reader with the human and humorous side of mathematics.

- Chapter 4: "Constructing Trees". Basic definitions, numbers of vertices and edges. Counting trees (saturated hydrocarbons, and labeled trees, following Cayley). Cayley's tree formula, Prüfer's proof[5] of the latter via Prüfer codes, and then decision trees, minimum spanning trees (motivated, at the beginning of the chapter, by a nice example about economically paving a system of roads), and Kruskal's algorithm.

- Chapter 5: "Traversing Graphs": This begins with Eulerian paths. I really enjoyed the history behind the bridges of Königsberg problem, which is much sketchier in the other accounts I know of. And how the importance of the "geometry of position" (now known to us as topology), originally envisioned by Leibniz, emerged out of Euler's consideration of this curious problem about walking around in an ancient city. Variants are discussed, like the Chinese Postman Problem. One very minor quibble, which is more about formatting than content: It is a little disconcerting to have three sub-sections have exactly the same name ("The Chinese Postman Problem") as each other as well as the name of the main section. This causes some moments of déja-vu. They are about different versions of the problem, and the section headings should say so.

- Chapter 6: "Encircling Graphs": Essentially Hamiltonian circuits and paths. Along the way, we learn about the origin of the Icosian Game, to say nothing of the introduction of croquet and ping pong to Great Britain.

- Chapter 7: "Factoring Graphs": Here things get quite interesting, possibly because I'm less familiar with some of the nuances covered here. The authors discuss matching, Hall's theorem, Tutte's

[5] or is it Proofer's prüf?

Theorem, Petersen's Theorem, the Petersen graph, and a few nice details about the lives of Tutte and Petersen. As throughout the book, many results are stated and not proved. But then they are well-motivated, which enables the reader to appreciate them.

- Chapter 8: "Decomposing Graphs": This chapter is about decomposing graphs into triangles and cycles, beginning with Kirkman's results on Steiner triple systems (including why they are not called Kirkman triple systems), and the "Schoolgirl Problem." Continuing with cycle decompositions, cyclic decompositions, graceful labelings and graceful graphs. And how to win at Instant Insanity via decomposition into regular subgraphs. Another small quibble about format: The Graceful Tree Conjecture looks like the name of a subsection, but the statement of the conjecture leads directly into an unrelated discussion of cyclic decomposition.

- Chapter 9: "Orienting Graphs": Begins with a brief history of Herbert Robbins' career (that, in turn, commencing with the 1931 game between the Harvard and Army football teams, which led to the Putnam competition), which includes some great quotes. Robbins' theorem on strong orientations of graphs. Then tournaments, pecking orders, and voting procedures. Few proofs here, other than some nice tidbits such as Redei's theorem that every tournament has a directed Hamiltonian path, and the King Chicken Theorem.

- Chapter 10: "Drawing Graphs": Concentrates on planar graphs. We see the proof of the Euler identity (generalizing the polyhedron identity to arbitrary planar graphs), Kuratowski's Theorem (enough to give a substantial appreciation of the result), crossing numbers, embedding graphs on surface, graph minors, Wagner's conjecture, and the Graph Minor Theorem of Robertson and Seymour. The chapter ends with the following as yet unresolved cliff-hanger:

  > [For] the set of graphs that can be embedded on the torus, there is a finite set M of forbidden minors. That is, a graph G can be embedded on the torus if and only if no graph in M is a minor of G. No one knows what such a set M is, except M must contain more than 80 graphs.

  Wow.

- Chapter 11: "Coloring Graphs": A great history of the four color problem, beginning in the 19th century (in which de Morgan played a role) tracing the development through Cayley, who publicized the problem and whose student Kempe's "proof" (found incorrect after 10 years, but nevertheless containing important and useful ideas) eventually resulted in Heawood's five color theorem (which is proved in the book), all of which laid the groundwork for the formulation as a graph problem and its resolution (or not?...still controversial in some circles) in 1976 by Appel and Haken. Some basic results and applications of vertex colorings are proved and given.

- Chapter 12: "Synchronizing Graphs": A chapter about edge coloring. It begins with the definition of chromatic index and the statement of Vizing's Theorem. Applications of edge coloring to scheduling problems, and then, of course, an introduction to Ramsey's Theorem and Ramsey numbers. Ramsey's Theorem is not proved, but extensive illuminating examples give the reader an appreciation for its content. The chapter, and the book, concludes with a careful explanation of synchronizing graphs (which in turn entails an encounter with finite automata, so this is one of the more elaborate problems described in the book), culminating in the Road Coloring Problem and the relevant theorem proved by Trahtman in 2008.

3 Opinion

I really like the fact that the subject is presented in narrative form, and that not all formal results are proved, nor are all problems solved. When not engaged in the serious business of proving theorems, the book has a breezy, open-ended feel to it. It also conveys quite well how mathematical research works (how mathematicians come up with proofs, write papers, get them published...).

I am compelled to make one substantive critical point. In Chapter 6, which deals with Hamiltonian circuits and related ideas, I found it surprising that there is nothing about NP-hardness in the context of, e.g., the Traveling Salesman Problem. Near the end of the chapter there is a remark that the TSP is "an extraordinarily complicated problem," but no hint as to how "extraordinarily complicated" it is widely believed to be (or what the phrase might mean). More generally, although algorithms like Kruskal's algorithm are discussed, there seems to be no mention of algorithmic efficiency or intractability anywhere in the book. Many problems and properties discussed here border on, or are even central to, algorithms and complexity. And, I would argue, *vice versa*: algorithms and complexity are a central concern of modern graph theory. I very much understand that authors may have reasons not to cover such topics. However, these topics strike me as sufficiently central to the discipline that at least something ought to be said about them, even if only to acknowledge their importance and mention that they would be omitted.

Even with that (in my view) shortcoming, this is a wonderful book. It ties together diverse concepts via recurring themes: extremal arguments (in proofs), complete graphs, planarity, connectivity, Eulerianicity and Hamiltonicity, and on it goes. In Chapter 8, for example, it becomes very clear how unified one can understand graph theory to be. Seemingly disparate themes are tied together and arise naturally in different contexts. E.g., Petersen's theorem is generalized at the same moment as Alspach's Conjecture is proved.

Proofs in this book are informal, often "just" giving the idea of the proof. This is more illuminating in most cases than formal proofs, and is much closer to how mathematicians explain mathematics to each other, rather than how it is expounded (not really explained) in conventional textbooks. Speaking of "conventional textbooks," this one is eminently suitable for a course. There is an extensive set of exercises of varying difficulty, set off from the text at the end of the book, in part to facilitate its narrative structure.

While it assumes very little mathematical background on the part of the reader, beyond a willingness to think abstractly, it nevertheless gets into some serious mathematics. One might question whether knowing something about the lives or careers of graph theorists has anything to do with graph theory, and of course one usually divorces the subject matter from the creators. But learning how the problems were formulated, solved, often "unsolved," and finally *re*solved gives one a real sense of what mathematical research, and the life it brings, is about. The twists and turns, the blind alleys, the circuitous paths (real life, Eulerian, Hamiltonian, or otherwise). . . these are all worth knowing about. Furthermore, the historical narrative enlivens the mathematical narrative, and gives one a deeper than usual appreciation of the significance of the subject.