

Undergraduate Topics in Computer Science

Undergraduate Topics in Computer Science (UTiCS) delivers high-quality instructional content for undergraduates studying in all areas of computing and information science. From core foundational and theoretical material to final-year topics and applications, UTiCS books take a fresh, concise, and modern approach and are ideal for self-study or for a one- or two-semester course. The texts are all authored by established experts in their fields, reviewed by an international advisory board, and contain numerous examples and problems. Many include fully worked solutions.

More information about this series at <http://www.springer.com/series/7592>

Bernhard Reus

Limits of Computation

From a Programming Perspective



Springer

Bernhard Reus
Department of Informatics
School of Engineering and Informatics
University of Sussex
Brighton
UK

Series editor
Ian Mackie

Advisory Board

Samson Abramsky, University of Oxford, Oxford, UK
Karin Breitman, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, Brazil
Chris Hankin, Imperial College London, London, UK
Dexter Kozen, Cornell University, Ithaca, USA
Andrew Pitts, University of Cambridge, Cambridge, UK
Hanne Riis Nielson, Technical University of Denmark, Kongens Lyngby, Denmark
Steven Skiena, Stony Brook University, Stony Brook, USA
Iain Stewart, University of Durham, Durham, UK

ISSN 1863-7310 ISSN 2197-1781 (electronic)
Undergraduate Topics in Computer Science
ISBN 978-3-319-27887-2 ISBN 978-3-319-27889-6 (eBook)
DOI 10.1007/978-3-319-27889-6

Library of Congress Control Number: 2015960818

© Springer International Publishing Switzerland 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by SpringerNature
The registered company is Springer International Publishing AG Switzerland

Foreword

Computer Science centers on questions about computational problems, and computer *programs* to solve problems:

- *What* is a computational problem, what is an algorithm, what does it mean to have solved a problem, are some problems unsolvable by any computing device, are there problems intrinsically difficult or impossible to solve automatically, are there problems intrinsically easier to solve than others, ...?
- *How* can a machine solve a problem, how to build such a machine, how to specify an algorithm for computer execution, how to design good algorithms, which “language” can a human use to direct the computer, how to design and “debug” programs for computer execution, how to build good programs, ...?

Good news: rapid progress has been made in both areas; we stand on the shoulders of giants in both theory and practice. The questions above have many and various answers. The first questions led in mathematical directions to foundational studies: the theories of computability, recursive functions, automata theory and more. The second led in engineering directions: computer architectures, the architectures of programming languages, the art or discipline of programming, software engineering and more.

Since the 1930s our understanding of both areas has developed hand in hand, led by theoreticians such as Kleene, Church, and Gödel; by hardware and software inventors such as Babbage, Von Neumann, and McCarthy; and by Alan Turing’s genius at the borderline between the two areas.

This book focuses on the first question area by an approach near the borderline: the theory of computability and complexity (C&C for short) is presented by using a simple programming language. In this language one is able to perform the (many) program constructions needed for the theory. This is done abstractly enough to reveal the great breadth and depth of C&C. Further, it is done concretely and precisely enough to satisfy practice-oriented readers about the constructions’ feasibility. Effect: a reader can see the relative efficiency of the constructions and understand the efficiency of what is constructed.

My 1997 C&C book was a step in this direction, but suffers from several problems: a scope too great for a single-semester university course; sections that (without warning) require more mathematical maturity than others; too few exercises; and too few historical and current references to research contexts.

Bernhard Reus has succeeded very well in overcoming these problems, and writing a deep, interesting, up-to-date, and even entertaining book on computability and complexity. I recommend it highly, both for systematic study and for spot reading.

Neil D. Jones
DIKU, University of Copenhagen

Preface

About 12 years ago a student¹ asked me after one of my lectures in *Computability and Complexity* why he had to write tedious Turing machine programs, given that everyone programmed in languages like Java. He had a point. What is the best way to teach a *Computability and Complexity* module in the twenty-first century to a cohort of students who are used to programming in high level languages with modern tools and libraries; to students who live in a world of smart phones and 24h connectivity, and, even more importantly maybe, who have not been exposed to very much formal reasoning and mathematics?

Turn the clock back only two or three decades. Then, a first year in a computer science Bachelor degree mainly consisted of mathematics (analysis and linear algebra, later discrete maths, numerical analysis, and basic probability theory). When computability and complexity was taught in the second or third year, students were already acquainted with a formal and very mathematical language of discourse, maybe because computer science lecturers in the 80s were usually mathematicians by trade. Things have changed significantly. Curriculum designers for Bachelor degrees are under pressure to push more and more new exciting material into a three-year degree program that should prepare students for their lives as working IT professionals. Any new module moved into the curriculum necessarily forces another one out. Often, allegedly “unpopular” modules, including formal theory and mathematics, are the victims. As a consequence, computer science students have to a degree lost the skills to digest material presented in an extremely formal and symbolic fashion while, at the same time, they are prolific programmers and quite knowledgeable in the use of tools.

So, *Computability and Complexity*, do they really have to be taught using Turing machines or μ -recursive functions? Do they have to be presented in the style a logician or mathematician would prefer? Seeking for alternatives, I eventually stumbled across Neil Jones’ fantastic book *Computability and Complexity—From a Programming Perspective*. The subtitle already gives away the book’s philosophy.

¹Alexis Petrounias.

The leitmotif of Neil’s textbook is to present the most important results in computability and complexity theory “using programming techniques and motivated by programming language theory” as well as “using a novel model of computation, differing from traditional ones in crucial aspects”.² The latter, WHILE, is a simple imperative language with one datatype of lists (s-expressions) à la LISP. Admittedly, this language is not Java, but it has the hallmarks of a modern high level language and is infinitely more comfortable to program in than Turing machines or Gödel numbers. Java, or any similar powerful language, would be impractical for our purposes “since proofs about them would be too complex to be easily understood”.³

So when rebranding the module under the name *Limits of Computation* in 2008, I adopted Neil’s book as course textbook. Delivering an introductory, one semester final-year module, I picked the most important and appealing chapters. This was easy as the design of the book was exactly made to mix and match.⁴ Soon, however, it turned out that students found Neil’s book tough going. In fact, this became more apparent as the years went past. There were several factors. First of all, Neil’s students would have had ML, a functional language with built in list type, as a first programming language, whereas our students were raised on Java. Yet, the datatype of WHILE is a functional one, and this caused more problems to the students than anticipated. Second, and more importantly, I had not put enough attention to the prerequisites. Neil expected readers of his book to be senior undergraduate students “with good mathematical maturity.”⁵ It turned out that not all the third-year students had this maturity (given the heterogeneity of backgrounds and reduction of maths teaching in the undergraduate years one and two).

As a response to mitigate the issues above, I started writing explicit notes to accompany my slides, intended as additional comments and explanations for the selected book chapters. I ended up adding more and more new material and rearranging it. The results of this effort are the 23 individual lectures of this book. A (British) semester is 12 weeks long, which usually requires 24 lectures to be delivered. The shortfall of one lecture is intentional, it acts as a buffer (in case things take more time) and also allows for extra events like invited talks or in-class tests.

This book was heavily influenced by Neil Jones’s textbook, which is clearly visible in some chapters. To pay homage to his book, its telling subtitle “From a Programming Perspective” has been adopted.

Brighton
November 2015

Bernhard Reus

²Preface of Neil’s book, page x.

³The book “Understanding Computation From Simple Machines to Impossible Programs” by Tom Stuart, published by O’Reilly in 2013, appears to follow the same idea and philosophy, using Ruby as programming language. It does not deal with complexity however.

⁴Neil’s Preface, page xii.

⁵Neil’s Preface, page xiii.

For Tutors

The students using this book are expected to be senior undergraduates who can master at least one imperative programming language. They are supposed to know arithmetic, Boolean algebra, graphs and some basic graph algorithms. Similarly, knowledge of basic set theory, function, and relations is needed, but Chap. 2 contains a short summary of basic definitions used in the book. Some exposition to formal reasoning would be helpful as well, but is not strictly required. To understand the probabilistic complexity classes in Chap. 21 some basic knowledge of probability theory will be needed.

This book is divided into 23 chapters with the intention that each chapter corresponds to one lecture. If not all chapters can be delivered, the following chapters can be omitted without interrupting the natural narrative:

- Chapter 10, *Self-referencing Programs*, but the self-producing program is a brain teaser which turns out to be very popular with students;
- Chapter 21, *How to Solve NP-complete Problems*, which may, however, be the chapter that has the most impact on students' future projects;
- The last two chapters *Molecular Computing* (Chap. 22) and *Quantum Computing* (Chap. 23), but students find this material particularly exciting.

A few chapters are significantly longer than others. These are Chap. 11, *The Church-Turing Thesis*, Chap. 16, *Famous Problems in P*, Chap. 20 *Complete Problems*, Chap. 21 *How Solve NP-complete Problems?*, and Chap. 22 *Molecular Computing*. The longer chapters allow the tutor to pick some of the sections and present them in the lecture, leaving the remaining ones for self-study or exercises.

Acknowledgements

This book would not have been possible without the help and influence of so many people. They all deserve a big thank you.

First of all, I wish to thank Neil Jones for his book on *Computability and Complexity*, which has been inspirational. The results of many brilliant researchers have been reported in this textbook. Many of them have been mentioned and referenced but, this being an introductory textbook, some of them may have been left out. My sincerest apologies to those who were omitted.

I am grateful for the feedback I have received over the years from my teaching assistants Matthew Wall, Jan Schwinghammer, Cristiano Solarino, Billie Joe Charlton, Ben Horsfall, and Shinya Sato.

I enjoyed and benefited from talking to many students on the *Limits of Computation* course I taught at Sussex. Thanks go in particular to Alexis Petrounias, Marek Es, Thomas Weedon Hume, Alex Jeffery, Susan Coleman, Sarah Aspery, Benjamin Hayward, Jordan Hobday, and Lucas Rijllart. The latter five also gave feedback on early versions of various chapters.

I also benefited from discussions with Thomas Streicher and my colleagues at Sussex: thanks go to Martin Berger, Matthew Hennessy, Julian Rathke, Des Watson, and particularly, George Parisis. Des and George also provided most welcome proof-reading services.

Luca Cardelli was kind enough to discuss current topics in molecular computing and gave valuable pointers regarding Chap. 22. Neil Jones and Matthew Hennessy provided comments on an entire draft of this book which helped improve the presentation. All remaining errors are of course solely mine.

I would also like to thank series editor Ian Mackie for his encouragement, Helen Desmond from Springer for her continuous support and Divya Meiyazhagan from the production team for all the last minute edits. The wonderful *LaTeX* typesetting and the *TikZ* (PGF) vector graph drawing systems have been used. Thanks to all those people who contributed to their development.

My sister's family and my friends deserve acknowledgment for moral support, and for being there for me when it counted.

Finally, I would like to dedicate this book to the memory of my parents and my brother. His interest in the sciences and in computing aroused my curiosity already at a young age.

Contents

1	Limits? What Limits?	1
1.1	Physical Limits of Computation	2
1.1.1	Fundamental Engineering Constraints to Semiconductor Manufacturing and Scaling	2
1.1.2	Fundamental Limits to Energy Efficiency	3
1.1.3	Fundamental Physical Constraints on Computing in General	3
1.2	The Limits Addressed	4
1.2.1	Computability Overview	4
1.2.2	Complexity Overview	6
	References	9

Part I Computability

2	Problems and Effective Procedures	13
2.1	On Computability	14
2.1.1	Historical Remarks	14
2.1.2	Effective Procedures	16
2.2	Sets, Relations and Functions	17
2.2.1	Sets	17
2.2.2	Relations	21
2.2.3	Functions	22
2.2.4	Partial Functions	22
2.2.5	Total Functions	23
2.3	Problems	24
2.3.1	Computing Solutions to Problems	26
	References	27

3 The WHILE-Language	29
3.1 The Data Type of Binary Trees	31
3.2 WHILE-Syntax	32
3.2.1 Expressions	32
3.2.2 Commands	32
3.2.3 Programs	33
3.2.4 A Grammar for WHILE	33
3.2.5 Layout Conventions and Brackets	34
3.3 Encoding Data Types as Trees	35
3.3.1 Boolean Values	35
3.3.2 Lists and Pairs	36
3.3.3 Natural Numbers	37
3.3.4 Finite Words	40
3.4 Sample Programs	40
3.4.1 Addition	40
3.4.2 List Reversal	41
3.4.3 Tail Recursion	42
3.4.4 Analysis of Algorithms	43
References	45
4 Semantics of WHILE	47
4.1 Stores	48
4.2 Semantics of Programs	49
4.3 Semantics of Commands	50
4.4 Semantics of Expressions	52
References	54
5 Extensions of WHILE	55
5.1 Equality	55
5.2 Literals	56
5.2.1 Number Literals	56
5.2.2 Boolean Literals	57
5.3 Adding Atoms	57
5.4 List Constructor	58
5.5 Macro Calls	59
5.6 Switch Statement	60
References	63
6 Programs as Data Objects	65
6.1 Interpreters Formally	66
6.2 Abstract Syntax Trees	67
6.3 Encoding of WHILE-ASTs in \mathbb{D}	67
Reference	70

7 A Self-interpreter for WHILE	71
7.1 A Self-interpreter for WHILE -Programs with One Variable	72
7.1.1 General Tree Traversal for ASTs	72
7.1.2 The STEP Macro	73
7.2 A Self-interpreter for WHILE	81
7.2.1 Store Manipulation Macros	83
References	86
8 An Undecidable (Non-computable) Problem	87
8.1 WHILE-Computability and Decidability	87
8.2 The Halting Problem for WHILE	89
8.3 Diagonalisation and the Barber “Paradox”	90
8.4 Proof of the Undecidability of the Halting Problem	92
References	95
9 More Undecidable Problems	97
9.1 Semi-decidability of the Halting Problem	97
9.2 Rice’s Theorem	99
9.3 The Tiling Problem	101
9.4 Problem Reduction	103
9.5 Other (Famous) Undecidable Problems	105
9.6 Dealing with Undecidable Problems	106
9.7 A Fast-Growing Non-computable Function	107
References	111
10 Self-referencing Programs	113
10.1 The S-m-n Theorem	114
10.2 Kleene’s Recursion Theorem	116
10.3 Recursion Elimination	118
References	121
11 The Church-Turing Thesis	123
11.1 The Thesis	124
11.2 Semantic Framework for Machine-Like Models	125
11.3 Turing Machines TM	126
11.4 GOTO-Language	129
11.5 Register Machines RAM and SRAM	131
11.6 Counter Machines CM	134
11.7 Cellular Automata	135
11.7.1 2D: Game of Life	138
11.7.2 1D: Rule 110	140
11.8 Robustness of Computability	141
11.8.1 The Crucial Role of Compilers	141
11.8.2 Equivalence of Models	142
References	147

Part II Complexity

12 Measuring Time Usage	151
12.1 Unit-Cost Time Measure	152
12.2 Time Measure for WHILE	154
12.3 Comparing Programming Languages Considering Time	157
References	160
13 Complexity Classes	161
13.1 Runtime Bounds	162
13.2 Time Complexity Classes	163
13.3 Lifting Simulation Properties to Complexity Classes	165
13.4 Big-O and Little-o	166
References	171
14 Robustness of P	173
14.1 Extended Church–Turing Thesis	174
14.2 Invariance or Cook’s Thesis	174
14.2.1 Non-sequential Models	175
14.2.2 Evidence for Cook’s Thesis	176
14.2.3 Linear Time	178
14.3 Cobham–Edmonds Thesis	179
References	181
15 Hierarchy Theorems	183
15.1 Linear Time Hierarchy Theorems	184
15.2 Beyond Linear Time	189
15.3 Gaps in the Hierarchy	192
References	193
16 Famous Problems in P	195
16.1 Decision Versus Optimisation Problems	197
16.2 Predecessor Problem	198
16.3 Membership Test for a Context Free Language	201
16.4 Primality Test	202
16.5 Graph Problems	203
16.5.1 Reachability in a Graph	203
16.5.2 Shortest Paths in a Graph	204
16.5.3 Maximal Matchings	206
16.5.4 Min-Cut and Max-Flow	207
16.5.5 The Seven Bridges of Königsberg	208
16.6 Linear Programming	210
References	215
17 Common Problems Not Known to Be in P	217
17.1 The Travelling Salesman Problem (TSP)	218
17.2 The Graph Colouring Problem	220

17.3	Max-Cut Problem	221
17.4	The 0-1 Knapsack Problem	222
17.5	Integer Programming Problem	223
17.6	Does Not Being in P Matter?	224
	References	226
18	The One-Million-Dollar Question	227
18.1	The Complexity Class NP	228
18.2	Nondeterministic Programs	229
18.2.1	Time Measure of Nondeterministic Programs	231
18.2.2	Some Basic Facts About NP	233
18.3	Robustness of NP	234
18.4	Problems in NP	235
18.5	The Biggest Open Problem in (Theoretical) Computer Science	237
	References	239
19	How Hard Is a Problem?	241
19.1	Reminder: Effective Reductions	242
19.2	Polynomial Time Reduction	242
19.3	Hard Problems	245
	References	249
20	Complete Problems	251
20.1	A First NP -complete Problem	252
20.2	More NP -complete Problems	255
20.3	Puzzles and Games	256
20.3.1	Chess	258
20.3.2	Sudoku	259
20.3.3	Tile-Matching Games	260
20.4	Database Queries	261
20.5	Policy Based Routing	264
20.6	“Limbo” Problems	266
20.7	Complete Problems in Other Classes	268
20.7.1	P -complete	268
20.7.2	RE -complete	269
	References	273
21	How to Solve NP-Complete Problems	275
21.1	Exact Algorithms	276
21.2	Approximation Algorithms	276
21.3	Parallelism	281
21.4	Randomization	282
21.4.1	The Class RP	282
21.4.2	Probabilistic Algorithms	284

21.5	Solving the Travelling Salesman Problem	285
21.5.1	Exact Solutions	285
21.5.2	Approximative Solutions	286
21.6	When Bad Complexity is Good News.	289
	References	295
22	Molecular Computing.	299
22.1	The Beginnings of DNA Computing.	300
22.2	DNA Computing Potential.	301
22.3	DNA Computing Challenges	302
22.4	Abstract Models of Molecular Computation.	302
22.4.1	Chemical Reaction Networks (CRN)	303
22.4.2	CRNs as Effective Procedures.	305
22.4.3	Are CRNs Equivalent to Other Notions of Computation?	308
22.4.4	Time Complexity for CRNs	309
22.4.5	Implementing CRNs	309
	References	315
23	Quantum Computing	317
23.1	Molecular Electronics	318
23.2	The Mathematics of Quantum Mechanics	319
23.3	Quantum Computability and Complexity.	321
23.4	Quantum Algorithms	323
23.4.1	Shor’s Algorithm	323
23.4.2	Grover’s Algorithm	324
23.5	Building Quantum Computers	325
23.6	Quantum Computing Challenges	325
23.7	To Boldly Go	326
	References	328
	Further Reading—Computability and Complexity Textbooks	331
	Glossary	335
	Index	341