



# Learn Algorithms through Programming and Puzzle Solving



**Alexander Kulikov and Pavel Pevzner**

## Welcome!

Thank you for joining us! This book powers our popular Data Structures and Algorithms online specialization on [Coursera](#) and online MicroMasters program at [edX](#) taken by over 700 000 students. See [the book website](#) for more information about these online courses. We encourage you to sign up for a session and learn this material while interacting with thousands of other talented students from around the world. As you explore this book, you will find a number of active learning components that help you study the material at your own pace.

**Stop and Think** questions invite you to slow down and check your knowledge before moving to the next topic.

**Exercise Breaks** offer “just in time” assessments testing your understanding of a topic before moving to the next one.

**Algorithmic puzzles** provide you with a fun way to “invent” the key algorithmic ideas on your own! Even if you fail to solve some puzzles, the time will not be lost as you will better appreciate the beauty and power of algorithms. We encourage you to try our puzzles before attempting to solve the programming challenges.

**Programming challenges** ask you to implement the algorithms that you will encounter in one of programming languages that we support: C, C++, Java, JavaScript, Python, Scala, C#, Haskell, Kotlin, Ruby, and Rust.

# Contents

<b>About This Book</b>	<b>7</b>
Programming Challenges . . . . .	10
Interactive Algorithmic Puzzles . . . . .	17
What Lies Ahead . . . . .	20
Meet the Authors . . . . .	21
Meet Our Online Co-Instructors . . . . .	22
Acknowledgments . . . . .	23
<b>1 Algorithms and Complexity</b>	<b>25</b>
1.1 What Is an Algorithm? . . . . .	25
1.2 Pseudocode . . . . .	25
1.3 Problem Versus Problem Instance . . . . .	25
1.4 Correct Versus Incorrect Algorithms . . . . .	27
1.5 Fast Versus Slow Algorithms . . . . .	28
1.6 Big-O Notation . . . . .	31
1.6.1 Advantages and Disadvantages . . . . .	35
1.6.2 Five Common Rules . . . . .	35
1.6.3 Visualizing Common Rules . . . . .	36
1.6.4 Frequently Arising Running Times . . . . .	40
<b>2 Algorithm Design Techniques</b>	<b>43</b>
2.1 Exhaustive Search Algorithms . . . . .	43
2.2 Branch-and-Bound Algorithms . . . . .	44
2.3 Greedy Algorithms . . . . .	44
2.4 Dynamic Programming Algorithms . . . . .	47
2.5 Recursive Algorithms . . . . .	50
2.6 Divide-and-Conquer Algorithms . . . . .	56
2.7 Randomized Algorithms . . . . .	58
<b>3 Programming Challenges</b>	<b>63</b>
3.1 Sum of Two Digits . . . . .	64
3.2 Maximum Pairwise Product . . . . .	67
3.2.1 Naive Algorithm . . . . .	68
3.2.2 Fast Algorithm . . . . .	72
3.2.3 Testing and Debugging . . . . .	73

3.2.4	Can You Tell Me What Error Have I Made? . . . . .	75
3.2.5	Stress Testing . . . . .	75
3.2.6	Even Faster Algorithm . . . . .	79
3.2.7	A More Compact Algorithm . . . . .	80
3.3	Solving a Programming Challenge in Five Easy Steps . . . . .	80
3.3.1	Reading Problem Statement . . . . .	81
3.3.2	Designing an Algorithm . . . . .	81
3.3.3	Implementing an Algorithm . . . . .	81
3.3.4	Testing and Debugging . . . . .	82
3.3.5	Submitting to the Grading System . . . . .	83
<b>4</b>	<b>Algorithmic Warm Up</b>	<b>85</b>
4.1	Programming Challenges . . . . .	87
4.1.1	Fibonacci Number . . . . .	87
4.1.2	Last Digit of Fibonacci Number . . . . .	93
4.1.3	Huge Fibonacci Number . . . . .	96
4.1.4	Last Digit of the Sum of Fibonacci Numbers . . . . .	99
4.1.5	Last Digit of the Partial Sum of Fibonacci Numbers . . . . .	107
4.1.6	Last Digit of the Sum of Squares of Fibonacci Numbers . . . . .	109
4.1.7	Greatest Common Divisor . . . . .	111
4.1.8	Least Common Multiple . . . . .	115
<b>5</b>	<b>Greedy Algorithms</b>	<b>117</b>
5.1	The Main Idea . . . . .	119
5.1.1	Examples . . . . .	119
5.1.2	Proving Correctness of Greedy Algorithms . . . . .	120
5.1.3	Implementation . . . . .	121
5.2	Programming Challenges . . . . .	123
5.2.1	Money Change . . . . .	123
5.2.2	Maximum Value of the Loot . . . . .	127
5.2.3	Car Fueling . . . . .	132
5.2.4	Maximum Advertisement Revenue . . . . .	137
5.2.5	Collecting Signatures . . . . .	142
5.2.6	Maximum Number of Prizes . . . . .	148
5.2.7	Maximum Salary . . . . .	151



<b>6</b>	<b>Divide-and-Conquer</b>	<b>155</b>
6.1	The Main Idea . . . . .	157
6.1.1	Guess a Number . . . . .	157
6.1.2	Searching Sorted Data . . . . .	161
6.1.3	Finding a White-Black Pair . . . . .	163
6.1.4	Finding a Peak . . . . .	165
6.1.5	Multiplying Integers . . . . .	166
6.1.6	The Master Theorem . . . . .	170
6.2	Programming Challenges . . . . .	177
6.2.1	Binary Search . . . . .	177
6.2.2	Binary Search with Duplicates . . . . .	181
6.2.3	Majority Element . . . . .	186
6.2.4	Improving QUICKSORT . . . . .	190
6.2.5	Number of Inversions . . . . .	192
6.2.6	Organizing a Lottery . . . . .	199
6.2.7	Closest Points . . . . .	205
<b>7</b>	<b>Dynamic Programming</b>	<b>211</b>
7.1	The Main Idea . . . . .	214
7.1.1	Number of Paths . . . . .	214
7.1.2	Dynamic Programming . . . . .	216
7.1.3	Shortest Path in Directed Acyclic Graph . . . . .	217
7.2	Programming Challenges . . . . .	220
7.2.1	Money Change Again . . . . .	220
7.2.2	Primitive Calculator . . . . .	225
7.2.3	Edit Distance . . . . .	230
7.2.4	Longest Common Subsequence of Two Sequences . . . . .	240
7.2.5	Longest Common Subsequence of Three Sequences . . . . .	244
7.2.6	Maximum Amount of Gold . . . . .	247
7.2.7	Splitting the Pirate Loot . . . . .	257
7.2.8	Maximum Value of an Arithmetic Expression . . . . .	261
7.3	Designing Dynamic Programming Algorithms . . . . .	268
<b>8</b>	<b>Best Programming Practices (Optional)</b>	<b>269</b>
8.1	Language Independent . . . . .	269
8.1.1	Code Format . . . . .	269
8.1.2	Code Structure . . . . .	269
8.1.3	Names and Comments . . . . .	271

8.1.4	Debugging . . . . .	272
8.1.5	Integers and Floating Point Numbers . . . . .	274
8.1.6	Strings . . . . .	275
8.1.7	Ranges . . . . .	275
8.2	C++ Specific . . . . .	278
8.2.1	Code Format . . . . .	278
8.2.2	Code Structure . . . . .	278
8.2.3	Types and Constants . . . . .	280
8.2.4	Classes . . . . .	282
8.2.5	Containers . . . . .	284
8.2.6	Integers and Floating Point Numbers . . . . .	285
8.3	Python Specific . . . . .	286
8.3.1	General . . . . .	286
8.3.2	Code Structure . . . . .	289
8.3.3	Functions . . . . .	290
8.3.4	Strings . . . . .	291
8.3.5	Classes . . . . .	293
8.3.6	Exceptions . . . . .	293
<b>Appendix</b>		<b>295</b>
	Compiler Flags . . . . .	295
	Frequently Asked Questions . . . . .	296

# About This Book

---

*I find that I don't understand things unless I try to program them.*

---



Donald E. Knuth, *The Art of Computer Programming*

There are many excellent books on Algorithms — why in the world would we write another one???

Because we feel that while these books excel in introducing algorithmic ideas, they have not yet succeeded in teaching you how to implement algorithms, the crucial computer science skill. As the famous quote by the legendary computer scientist Donald Knuth hints, learning algorithms without implementing them is not unlike learning surgery based solely on reading an anatomy book.

**Intelligent Tutoring System for learning algorithms.** Our goal is to develop an *Intelligent Tutoring System* for learning algorithms that can compete with the best professors in a traditional classroom. This *MOOC book* is the first step towards this goal written specifically for our Massive Open Online Courses (MOOCs) forming a specialization “[Algorithms and Data Structures](#)” on Coursera platform and [a microMasters program](#) on edX platform. Since the launch of our MOOCs in 2016, over 700 000 students enrolled in this specialization and tried to solve more than hundred algorithmic programming challenges to pass it. And some of them even got offers from small companies like Google after completing our specialization!

**Like Donald Knuth, learn algorithms by programming them!** Some professors have been concerned about the pedagogical quality of MOOCs and even called them the “junk food of education.” In contrast, we are among the growing group of professors who believe that traditional classes, that pack hundreds of students in a classroom, represent junk food of education. In a large classroom, once a student takes a wrong turn, there are limited

opportunities to ask a question, resulting in a *learning breakdown*, or the inability to progress further without individual guidance. Furthermore, the majority of time a student invests in an Algorithms course is spent completing assignments outside the classroom. That is why we stopped giving lectures in our offline classes (and we haven't got fired yet :-). Instead, we give *flipped classes* where students watch our recorded lectures, solve algorithmic puzzles, complete programming challenges using our automated homework checking system before the class, and come to class prepared to discuss their learning breakdowns with us.

**Repository of algorithmic coding challenges.** When a student experiences a learning breakdown, s/he needs immediate help in order to proceed. Traditional textbooks do not provide such help, but our automated grading system does! Algorithms is a unique discipline since students' ability to program provides the opportunity to automatically check their knowledge through coding challenges. These coding challenges are far superior to traditional quizzes that barely check whether a student fell asleep. Indeed, to implement a complex algorithm, the student must possess a deep understanding of its underlying algorithmic ideas.

We believe that a large portion of grading in thousands of Algorithms courses taught at various universities each year can be consolidated into a single automated system available at all universities. It did not escape our attention that many professors teaching algorithms have implemented their own custom-made systems for grading student programs, an illustration of academic inefficiency and lack of cooperation between various instructors. Our goal is to build a repository of algorithmic programming challenges, thus allowing professors to focus on teaching and helping students to prepare for the coding interviews. We have already invested thousands of hours into building such a system and thousands students in our MOOCs tested it. Below we briefly describe how it works.

**Solving programming challenges.** When you face a programming challenge, your goal is to implement an efficient algorithm for its solution. Solving programming challenges will help you better understand various algorithms and may even land you a job since many high-tech companies give our programming challenges during the interviews.

Your implementation will be automatically checked against many care-

fully selected tests to verify that it always produces a correct answer and fits into the time and memory constraints. Our system will teach you to write programs that work correctly on all of our test datasets rather than on some of them. This is an important skill since failing to thoroughly test your programs leads to undetected bugs that frustrate your boss, your colleagues, and, most importantly, users of your programs.

It took us thousands of hours to develop the automated homework checking system. First, we had to build a Compendium of Learning Breakdowns for each programming challenge, 10–15 most frequent errors that students make while solving it. Afterward, we had to develop test cases for each learning breakdown in each programming challenge, over 20 000 test cases for just 100+ programming challenges in this Specialization. This MOOC book discusses 30 of these programming challenges that are included in the first “Algorithmic Toolbox” MOOC.

We are grateful to many students for the feedback that helped us to improve this MOOC that is now ranked among the [top ten computer science courses on Coursera](#) and among [250 best MOOCs of all time](#) across all disciplines and all platforms.

Thank you for joining us!

## Programming Challenges

This book introduces basic algorithmic techniques using the following programming challenges.

$$2 + 3 = 5$$

### Sum of Two Digits

*Compute the sum of two single digit numbers.*

Input. Two single digit numbers.

Output. The sum of these numbers.

	5	6	2	7	4
5		30	10	35	20
6	30		12	42	24
2	10	12		14	8
7	35	42	14		28
4	20	24	8	28	

### Maximum Pairwise Product

*Find the maximum product of two distinct numbers in a sequence of non-negative integers.*

Input. An integer  $n$  and a sequence of  $n$  non-negative integers.

Output. The maximum value that can be obtained by multiplying two different elements from the sequence.

$$F_n = F_{n-1} + F_{n-2}$$

### Fibonacci Number

*Compute the  $n$ -th Fibonacci number.*

Input. An integer  $n$ .

Output.  $n$ -th Fibonacci number.


$$F_{100} = 354224848179 \\ 261915075$$

### Last Digit of Fibonacci Number

*Compute the last digit of the  $n$ -th Fibonacci number.*

Input. An integer  $n$ .

Output. The last digit of the  $n$ -th Fibonacci number.




### Huge Fibonacci Number

*Compute the  $n$ -th Fibonacci number modulo  $m$ .*

Input. Integers  $n$  and  $m$ .

Output.  $n$ -th Fibonacci number modulo  $m$ .






### Last Digit of the Sum of Fibonacci Numbers

Compute the last digit of  $F_0 + F_1 + \dots + F_n$ .

Input. An integer  $n$ .

Output. The last digit of  $F_0 + F_1 + \dots + F_n$ .




### Last Digit of the Partial Sum of Fibonacci Numbers

Compute the last digit of  $F_m + F_{m+1} + \dots + F_n$ .

Input. Integers  $m \leq n$ .

Output. The last digit of  $F_m + F_{m+1} + \dots + F_n$ .




### Last Digit of the Sum of Squares of Fibonacci Numbers

Compute the last digit of  $F_0^2 + F_1^2 + \dots + F_n^2$ .

Input. An integer  $n$ .

Output. The last digit of  $F_0^2 + F_1^2 + \dots + F_n^2$ .



### Greatest Common Divisor

Compute the greatest common divisor of two positive integers.

Input. Two positive integers  $a$  and  $b$ .

Output. The greatest common divisor of  $a$  and  $b$ .




### Least Common Multiple

Compute the least common multiple of two positive integers.

Input. Two positive integers  $a$  and  $b$ .

Output. The least common multiple of  $a$  and  $b$ .



### Money Change

Compute the minimum number of coins needed to change the given value into coins with denominations 1, 5, and 10.

Input. An integer  $money$ .

Output. The minimum number of coins with denominations 1, 5, and 10 that changes  $money$ .



### Maximizing the Value of the Loot

*Find the maximal value of items that fit into the backpack.*

Input. The capacity of a backpack  $W$  as well as the weights ( $w_1, \dots, w_n$ ) and costs ( $c_1, \dots, c_n$ ) of  $n$  different compounds.

Output. The maximum total value of items that fit into the backpack of the given capacity: i.e., the maximum value of  $c_1 \cdot f_1 + \dots + c_n \cdot f_n$  such that  $w_1 \cdot f_1 + \dots + w_n \cdot f_n \leq W$  and  $0 \leq f_i \leq 1$  for all  $i$  ( $f_i$  is the fraction of the  $i$ -th item taken to the backpack).




### Car Fueling

*Compute the minimum number of gas tank refills to get from one city to another.*

Input. Integers  $d$  and  $m$ , as well as a sequence of integers  $stop_1 < stop_2 < \dots < stop_n$ .

Output. The minimum number of refills to get from one city to another if a car can travel at most  $m$  miles on a full tank. The distance between the cities is  $d$  miles and there are gas stations at distances  $stop_1, stop_2, \dots, stop_n$  along the way. We assume that a car starts with a full tank.




### Maximum Product of Two Sequences

*Find the maximum dot product of two sequences of numbers.*

Input. Two sequences of  $n$  positive integers:  $price_1, \dots, price_n$  and  $clicks_1, \dots, clicks_n$ .

Output. The maximum value of  $price_1 \cdot c_1 + \dots + price_n \cdot c_n$ , where  $c_1, \dots, c_n$  is a permutation of  $clicks_1, \dots, clicks_n$ .




### Covering Segments by Points

*Find the minimum number of points needed to cover all given segments on a line.*

Input. A sequence of  $n$  segments  $[l_1, r_1], \dots, [l_n, r_n]$  on a line.

Output. A set of points of minimum size such that each segment  $[l_i, r_i]$  contains a point, i.e., there exists a point  $x$  from this set such that  $l_i \leq x \leq r_i$ .




### Distinct Summands

*Represent a positive integer as the sum of the maximum number of pairwise distinct positive integers.*

Input. An integer  $n$ .

Output. The maximum  $k$  such that  $n$  can be represented as the sum  $a_1 + \dots + a_k$  of  $k$  distinct integers.




### Largest Concatenate

*Compile the largest number by concatenating the given numbers.*

Input. A sequence of positive integers.

Output. The largest number that can be obtained by concatenating the given integers in some order.




### Sorted Array Multiple Search

*Search multiple keys in a sorted sequence of keys.*

Input. A sorted array  $K$  of distinct integers and an array  $Q = \{q_0, \dots, q_{m-1}\}$  of integers.

Output. For each  $q_i$ , check whether it occurs in  $K$ .




### Binary Search with Duplicates

*Find the index of the first occurrence of a key in a sorted array.*

Input. A sorted array of integers (possibly with duplicates) and an integer  $q$ .

Output. Index of the first occurrence of  $q$  in the array or “-1” if  $q$  does not appear in the array.




### Majority Element

*Check whether a given sequence of numbers contains an element that appears more than half of the times.*

Input. A sequence of  $n$  integers.

Output. 1, if there is an element that is repeated more than  $n/2$  times, and 0 otherwise.




### Improving QUICKSORT

*Modify the QUICKSORT algorithm so that it works fast even on sequences containing many identical elements.*

Input. A sequence of  $n$  integers.

Output. Sequence sorted in non-decreasing order.




### Number of Inversions

*Compute the number of inversions in a sequence of integers.*

Input. A sequence of  $n$  integers  $a_1, \dots, a_n$ .

Output. The number of inversions in the sequence, i.e., the number of indices  $i < j$  such that  $a_i > a_j$ .




### Points and Segments

*Given a set of points and a set of segments on a line, compute, for each point, the number of segments it is contained in.*

Input. A list of  $n$  segments and a list of  $m$  points.

Output. The number of segments containing each point.



### Closest Points

*Find the closest pair of points in a set of points on a plane.*

Input. A list of  $n$  points on a plane.

Output. The minimum distance between a pair of these points.




### Money Change Again

*Compute the minimum number of coins needed to change the given value into coins with denominations 1, 3, and 4.*

Input. An integer money.

Output. The minimum number of coins with denominations 1, 3, and 4 that changes money.




### Primitive Calculator

*Find the minimum number of operations needed to get a positive integer  $n$  from 1 by using only three operations: add 1, multiply by 2, and multiply by 3.*

Input. An integer  $n$ .

Output. The minimum number of operations “+1”, “x2”, and “x3” needed to get  $n$  from 1.




### Edit Distance

*Compute the edit distance between two strings.*

Input. Two strings.

Output. The minimum number of single-symbol insertions, deletions, and substitutions to transform one string into the other one.




### Longest Common Subsequence of Two Sequences

*Compute the maximum length of a common subsequence of two sequences.*

Input. Two sequences.

Output. The maximum length of a common subsequence.




### Longest Common Subsequence of Three Sequences

*Compute the maximum length of a common subsequence of three sequences.*

Input. Three sequences.

Output. The maximum length of a common subsequence.




### Maximum Amount of Gold

Given a set of gold bars of various weights and a backpack that can hold at most  $W$  pounds, place as much gold as possible into the backpack.

Input. A set of  $n$  gold bars of integer weights  $w_1, \dots, w_n$  and a backpack that can hold at most  $W$  pounds.

Output. A subset of gold bars of maximum total weight not exceeding  $W$ .




### 3-Partition

Partition a set of integers into three subsets with equal sums.

Input. A sequence of integers  $v_1, v_2, \dots, v_n$ .

Output. Check whether it is possible to partition them into three subsets with equal sums, i.e., check whether there exist three disjoint sets  $S_1, S_2, S_3 \subseteq \{1, 2, \dots, n\}$  such that  $S_1 \cup S_2 \cup S_3 = \{1, 2, \dots, n\}$  and

$$\sum_{i \in S_1} v_i = \sum_{j \in S_2} v_j = \sum_{k \in S_3} v_k.$$



### Maximum Value of an Arithmetic Expression

Parenthesize an arithmetic expression to maximize its value.


Input. An arithmetic expression consisting of digits as well as plus, minus, and multiplication signs.

Output. Add parentheses to the expression in order to maximize its value.



## Interactive Algorithmic Puzzles


You are also welcome to solve the following interactive algorithmic puzzles available at <https://discrete-math-puzzles.github.io/everything.html>:



**Eight Queens.** Place eight queens on the chessboard such that no two queens attack each other (a queen can move horizontally, vertically, or diagonally).




**Hanoi Towers.** Move all disks from one peg to another using a minimum number of moves. In a single move, you can move a top disk from one peg to any other peg provided that you don't place a larger disk on the top of a smaller disk.




**Covering Segments by Points.** Find the minimum number of points that cover all given segments on a line.



**Activity Selection.** Select as many non-overlapping segments as possible.



**Largest Concatenate Problem.** Compile the largest number by concatenating the given numbers.




**Black and White Squares.** Use the minimum number of questions “What is the color of this square?” to find two neighboring squares of different colors. The leftmost square is white, the rightmost square is black, but the colors of all other squares are unknown.




**Twenty One Questions Game.** Find an unknown integer  $1 \leq x \leq N$  by asking the minimum number of questions “Is  $x = y$ ?” (for any  $1 \leq y \leq N$ ). Your opponent will reply either “Yes”, or “ $x < y$ ”, or “ $x > y$ .”




**Book Sorting.** Rearrange books on the shelf (in the increasing order of heights) using minimum number of swaps.




**Number of Paths.** Find out how many paths are there to get from the bottom left circle to any other circle and place this number inside the corresponding circle.



**Antique Calculator.** Find the minimum number of operations needed to get a positive integer  $n$  from the integer 1 using only three operations: add 1, multiply by 2, or multiply by 3.




**Two Rocks Game.** There are two piles of ten rocks. In each turn, you and your opponent may either take one rock from a single pile, or one rock from both piles. Your opponent moves first and the player that takes the last rock wins the game. Design a winning strategy.




**Three Rocks Game.** There are two piles of ten rocks. In each turn, you and your opponent may take up to three rocks. Your opponent moves first and the player that takes the last rock wins the game. Design a winning strategy.




**Map Coloring.** Use minimum number of colors such that neighboring countries are assigned different colors and each country is assigned a single color.




**Clique Finding.** Find the largest group of mutual friends (each pair of friends is represented by an edge).




**Icosian Game.** Find a cycle visiting each node exactly once.




**Guarini Puzzle.** Exchange the places of the white knights and the black knights. Two knights are not allowed to occupy the same cell of the chess board.




**Room Assignment.** Place each student in one of her/his preferable rooms in a dormitory so that each room is occupied by a single student (preferable rooms are shown by edges).



**Tree Construction.** Remove the minimum number of edges from the graph to make it acyclic.



**Subway Lines.** You are planning a subway system where the subway lines should not cross. Can you connect each pair of the five stations except for a single pair?



**Plow Truck.** A plow truck needs to clean every street and get back to its original location. Help it find the shortest route for this.



**Delivery Van.** A delivery van needs to visit all the points and get back to the initial location. Help it find the shortest route for this.

# What Lies Ahead

Watch for our future editions that will cover the following topics.

## Data Structures

- Arrays and Lists

- Priority Queues

- Disjoint Sets

- Hash Tables

- Binary Search Trees

## Algorithms on Graphs

- Graphs Decomposition

- Shortest Paths in Graphs

- Minimum Spanning Trees

- Shortest Paths in Real Life

## Algorithms on Strings

- Pattern Matching

- Suffix Trees

- Suffix Arrays

- Burrows–Wheeler Transform

## Advanced Algorithms and Complexity

- Flows in Networks

- Linear Programmings

- NP-complete Problems

- Coping with NP-completeness

- Streaming Algorithms

## Meet the Authors



**Alexander S. Kulikov** (left) is a senior research fellow at Steklov Mathematical Institute of the Russian Academy of Sciences, Saint Petersburg, Russia and a lecturer at the Department of Computer Science and Engineering at University of California, San Diego, USA. He also is the head of B.Sc. program “Modern Software Engineering” at St. Petersburg State University and a board member of Computer Science Center in St. Petersburg that provides free advanced computer science courses complementing the standard university curricula. Alexander holds a Ph. D. from Steklov Mathematical Institute. His research interests include algorithms and complexity theory. He co-authored online courses “Data Structures and Algorithms” and “Introduction to Discrete Mathematics for Computer Science” that are available at Coursera and edX.

**Pavel Pevzner** (right) is Ronald R. Taylor Professor of Computer Science at the University of California, San Diego. He holds a Ph. D. from Moscow Institute of Physics and Technology, Russia and an Honorary Degree from Simon Fraser University. He is a Howard Hughes Medical Institute Professor (2006), an Association for Computing Machinery (ACM) Fellow (2010), an International Society for Computational Biology (ISCB) Fellow (2012), and a Member of the the Academia Europaea (2016). He was awarded the Senior Scientist Award from ISCB (2017) and the Kanellakis Theory and Practice Award from ACM (2019). He has authored the textbooks Computational Molecular Biology: An Algorithmic Approach (2000), An Introduction to Bioinformatics Algorithms (2004) (jointly with Neil Jones), and Bioinformatics Algorithms: An Active Learning Approach (2014) (jointly with Phillip Compeau). He co-authored online courses “Data Structures and Algorithms”, “Bioinformatics”, “Analyze Your Genome!”, and “Hacking COVID-19” that are available at Coursera and edX.

## Meet Our Online Co-Instructors



**Daniel Kane** (left) is an associate professor at the University of California, San Diego with a joint appointment between the Department of Computer Science and Engineering and the Department of Mathematics. He has diverse interests in mathematics and theoretical computer science, though most of his work fits into the broad categories of number theory, complexity theory, or combinatorics.

**Michael Levin** (middle) is an Associate Professor at the Computer Science Department of Higher School of Economics, Moscow, Russia and the Chief Data Scientist at the Yandex.Market, Moscow, Russia. He also teaches Algorithms and Data Structures at the Yandex School of Data Analysis.

**Neil Rhodes** (right) is a lecturer in the Computer Science and Engineering department at the University of California, San Diego and formerly a staff software engineer at Google. Neil holds a B.A. and M.S. in Computer Science from UCSD. He left the Ph.D. program at UCSD to start a company, Palomar Software, and spent fifteen years writing software, books on software development, and designing and teaching programming courses for Apple and Palm. He's taught Algorithms, Machine Learning, Operating Systems, Discrete Mathematics, Automata and Computability Theory, and Software Engineering at UCSD and Harvey Mudd College in Claremont, California.



## Acknowledgments

This book was greatly improved by the efforts of a large number of individuals, to whom we owe a debt of gratitude.

Our co-instructors and partners in crime Daniel Kane, Michael Levin, and Neil Rhodes invested countless hours in the development of our online courses at Coursera and edX platforms.

Hundreds of thousands of our online students provided valuable feedback that led to many improvements in our MOOCs and this MOOC book. In particular, we are grateful to the mentors of the Algorithmic Toolbox class at Coursera: Ayoub Falah, Denys Diachenko, Kishaan Jeeveswaran, Irina Pinjaeva, Fernando Gonzales Vigil Richter, and Gabrio Secco.

We thank our colleagues who helped us with preparing programming challenges: Maxim Akhmedov, Roman Andreev, Gleb Evstropov, Nikolai Karpov, Sergey Poromov, Sergey Kopeliovich, Ilya Kornakov, Gennady Korotkevich, Paul Melnichuk, and Alexander Tiunov.

We are grateful to Anton Konev and Darya Borisyak for leading the development of interactive puzzles.

We thank Alexey Kladov, Sergey Lebedev, Alexei Levin, Sergey Shulman, and Alexander Smal for help with the “Best Programming Practices” section of the book.

Randall Christopher brought to life our idea for the textbook cover.

Finally, our families helped us preserve our sanity when we were working on this MOOC book.



# Chapter 1: Algorithms and Complexity

This book presents algorithmic programming challenges and puzzles that will teach you how to design and implement algorithms. Solving a programming challenge is one of the best ways to understand an algorithm's design as well as to identify its potential weaknesses and fix them.

## 1.1 What Is an Algorithm?

Roughly speaking, an algorithm is a sequence of instructions that one must perform in order to solve a well-formulated problem. We will specify problems in terms of their *inputs* and their *outputs*, and the algorithm will be the method of translating the inputs into the outputs. A well-formulated problem is unambiguous and precise, leaving no room for misinterpretation.

After you designed an algorithm, two important questions to ask are: “Does it work correctly?” and “How much time will it take?” Certainly you would not be satisfied with an algorithm that only returned correct results half the time, or took 1 000 years to arrive at an answer.

## 1.2 Pseudocode

To understand how an algorithm works, we need some way of listing the steps that the algorithm takes, while being neither too vague nor too formal. We will use *pseudocode*, a language computer scientists often use to describe algorithms. Pseudocode ignores many of the details that are required in a programming language, yet it is more precise and less ambiguous than, say, a recipe in a cookbook.

## 1.3 Problem Versus Problem Instance

A problem describes a class of computational tasks. A problem instance is one particular input from that class. To illustrate the difference between a problem and an instance of a problem, consider the following example. You find yourself in a bookstore buying a book for \$4.23 which you pay

for with a \$5 bill. You would be due 77 cents in change, and the cashier now makes a decision as to exactly how you get it. You would be annoyed at a fistful of 77 pennies or 15 nickels and 2 pennies, which raises the question of how to make change in the least annoying way. Most cashiers try to minimize the number of coins returned for a particular quantity of change.

The example of 77 cents represents an instance of the Change Problem that assumes that there are  $d$  denominations represented by an array  $c = (c_1, c_2, \dots, c_d)$ . For simplicity, we assume that the denominations are given in decreasing order of value. For example,  $c = (25, 10, 5, 1)$  for United States denominations.

**Stop and Think.** What is the minimum number of coins of denominations (25, 10, 5, 1) needed to change 77 cents?

---

### Change Problem

*Convert some amount of money into given denominations, using the smallest possible number of coins.*

**Input:** An integer *money* and an array of  $d$  denominations  $c = (c_1, c_2, \dots, c_d)$ , in decreasing order of value ( $c_1 > c_2 > \dots > c_d$ ).

**Output:** A list of  $d$  integers  $i_1, i_2, \dots, i_d$  such that  $c_1 \cdot i_1 + c_2 \cdot i_2 + \dots + c_d \cdot i_d = \text{money}$ , and  $i_1 + i_2 + \dots + i_d$  is as small as possible.

---

The algorithm that is used by cashiers all over the world to solve this problem is simple:

```
CHANGE(money, c, d):  
while money > 0:  
    coin ← coin with the largest denomination that does not exceed money  
    give coin with denomination coin to customer  
    money ← money − coin
```

Here is a faster version of CHANGE:

```

CHANGE(money, c, d):
  r ← money
  for k from 1 to d:
     $i_k \leftarrow \lfloor \frac{r}{c_k} \rfloor$ 
    r ← r − ck · ik
  return (i1, i2, ..., id)

```

## 1.4 Correct Versus Incorrect Algorithms

We say that an algorithm is correct when it translates every input instance into the correct output. An algorithm is incorrect when there is at least one input instance for which the algorithm gives an incorrect output.

**Stop and Think.** What is the minimum number of coins of denominations (25, 20, 10, 5, 1) needed to change 40 cents?

CHANGE is an incorrect algorithm! Suppose you were changing 40 cents into coins with denominations of  $c_1 = 25$ ,  $c_2 = 20$ ,  $c_3 = 10$ ,  $c_4 = 5$ , and  $c_5 = 1$ . CHANGE would incorrectly return 1 quarter, 1 dime, and 1 nickel, instead of 2 twenty-cent pieces. As contrived as this may seem, in 1875 a twenty-cent coin existed in the United States. How sure can we be that CHANGE returns the minimal number of coins for the modern US denominations or for denominations in any other country?

To correct the CHANGE algorithm, we could consider every possible combination of coins with denominations  $c_1, c_2, \dots, c_d$  that adds to *money*, and return the combination with the fewest. We only need to consider combinations with  $i_1 \leq \text{money}/c_1$  and  $i_2 \leq \text{money}/c_2$  (in general,  $i_k$  should not exceed  $\text{money}/c_k$ ), because we would otherwise be returning an amount of money larger than *money*. The pseudocode below uses the symbol  $\sum$  that stands for summation:  $\sum_{i=1}^m a_i = a_1 + a_2 + \dots + a_m$ . The pseudocode also uses the notion of “infinity” (denoted as  $\infty$ ) as an initial value for *smallestNumberOfCoins*; there are a number of ways to carry this out in a real computer, but the details are not important here.

```

BRUTEFORCECHANGE(money, c, d):
  smallestNumberOfCoins ← ∞
  for each (i1, ..., id) from (0, ..., 0) to ( $\frac{\textit{money}}{c_1}, \dots, \frac{\textit{money}}{c_d}$ ):
    valueOfCoins ←  $\sum_{k=1}^d i_k \cdot c_k$ 
    if valueOfCoins = M:
      numberOfCoins =  $\sum_{k=1}^d i_k$ 
      if numberOfCoins < smallestNumberOfCoins:
        smallestNumberOfCoins ← numberOfCoins
        change ← (i1, i2, ..., id)
  return change

```

The second line iterates over every feasible combination (*i*<sub>1</sub>, ..., *i*<sub>*d*</sub>) of the *d* indices, and stops when it has reached

$$\left( \frac{\textit{money}}{c_1}, \dots, \frac{\textit{money}}{c_d} \right).$$

How do we know that BRUTEFORCECHANGE does not suffer from the same problem as CHANGE did, namely that it generates incorrect result for some input instance? Since BRUTEFORCECHANGE explores all feasible combinations of denominations, it will eventually come across an optimal solution and record it as such in the *change* array. Any combination of coins that adds to *M* must have at least as many coins as the optimal combination, so BRUTEFORCECHANGE will never overwrite *change* with a suboptimal solution.

So far we have answered only one of the two important algorithmic questions (“Does it work?”, but not “How much time will it take?”).

**Stop and Think.** How fast is BRUTEFORCECHANGE?

## 1.5 Fast Versus Slow Algorithms

Real computers require a certain amount of time to perform an operation such as addition, subtraction, or testing the conditions in a while loop. A supercomputer might take  $10^{-10}$  second to perform an addition, while a calculator might take  $10^{-5}$  second. Suppose that you had a computer that took  $10^{-10}$  second to perform an elementary operation such as addition,



and that you knew how many operations a particular algorithm would perform. You could estimate the running time of the algorithm simply by taking the product of the number of operations and the time per operation. However, computers are constantly improving, leading to a decreasing time per operation, so your notion of the running time would soon be outdated. Rather than computing an algorithm's running time on every computer, we rely on the total number of operations that the algorithm performs to describe its running time, since this is an attribute of the algorithm, and not an attribute of the computer you happen to be using.

Unfortunately, determining how many operations an algorithm will perform is not always easy. If we know how to compute the number of basic operations that an algorithm performs, then we have a basis to compare it against a different algorithm that solves the same problem. Rather than tediously count every multiplication and addition, we can perform this comparison by gaining a high-level understanding of the growth of each algorithm's operation count as the size of the input increases.

Suppose an algorithm  $A$  performs  $n^2$  operations on an input of size  $n$ , and an algorithm  $B$  solves the same problem in  $3n + 2$  operations. Which algorithm,  $A$  or  $B$ , is faster? Although  $A$  may be faster than  $B$  for some small  $n$  (e.g., for  $n$  between 1 and 3),  $B$  will become faster for large  $n$  (e.g., for all  $n > 4$ ). See Figure 1.1. Since  $f(n) = n^2$  is, in some sense, a “faster-growing” function than  $g(n) = n$  with respect to  $n$ , the constants 3 and 2 in  $3n + 2$  do not affect the competition between the two algorithms for large  $n$ . We refer to  $A$  as a *quadratic* algorithm and to  $B$  as a *linear* algorithm, and say that  $A$  is less efficient than  $B$  because it performs more operations to solve the same problem when  $n$  is large. Thus, we will often be somewhat imprecise when we count operations of an algorithm—the behavior of algorithms on small inputs does not matter.

Let's estimate the number of operations `BRUTEFORCECHANGE` will take on an input instance of  $M$  cents, and denominations  $(c_1, c_2, \dots, c_d)$ . To calculate the total number of operations in the for loop, we can take the approximate number of operations performed in each iteration and multiply this by the total number of iterations. Since there are roughly

$$\frac{\text{money}}{c_1} \times \frac{\text{money}}{c_2} \times \dots \times \frac{\text{money}}{c_d}$$

iterations, the for loop performs on the order of  $d \times \frac{\text{money}^d}{c_1 c_2 \dots c_d}$  operations, which dwarfs the other operations of the algorithm.




Figure 1.1: Plots of  $n^2$  and  $3n + 2$ .

This type of algorithm is often referred to as an *exponential* algorithm in contrast to quadratic, cubic, or other *polynomial* algorithms. The expression for the running time of exponential algorithms includes a term like  $n^d$ , where  $n$  and  $d$  are parameters of the problem (i.e.,  $n$  and  $d$  may deliberately be made arbitrarily large by changing the input to the algorithm), while the running time of a polynomial algorithm is bounded by a term like  $n^k$  where  $k$  is a constant not related to the size of any parameters.

For example, an algorithm with running time  $n^1$  (linear),  $n^2$  (quadratic),  $n^3$  (cubic), or even  $n^{2018}$  is polynomial. Of course, an algorithm with running time  $n^{2018}$  is not very practical, perhaps less so than some exponential algorithms, and much effort in computer science goes into designing faster and faster polynomial algorithms. Since  $d$  may be large when the algorithm is called with a long list of denominations (e.g.,  $c = (1, 2, 3, 4, 5, \dots, 100)$ ), we see that BRUTEFORCECHANGE can take a very long time to execute.

## 1.6 Big-O Notation

To figure out how long a program would take to run on a real computer, we would need to know things like: speed of the computer, the system architecture, the compiler being used, details of the memory hierarchy, etc. Hence, carefully estimating the running time is a rather difficult task. Moreover, in practice, you might not even know some of these details. That is why computer scientists use the *big-O* notation to estimate the running time of an algorithm without knowing anything about all these details!

If we say that the running time of an algorithm is quadratic, or  $O(n^2)$ , it means that the running time of the algorithm on an input of size  $n$  is limited by a quadratic function of  $n$ . That limit may be  $99.7n^2$  or  $0.001n^2$  or  $5n^2 + 3.2n + 99993$ ; the main factor that describes the growth rate of the running time is the term that grows the fastest with respect to  $n$ , for example  $n^2$  when compared to terms like  $3.2n$ , or  $99993$ . All functions with a leading term of  $n^2$  have more or less the same rate of growth, so we lump them into one class which we call  $O(n^2)$ . The difference in behavior between two quadratic functions in that class, say  $99.7n^2$  and  $5n^2 + 3.2n + 99993$ , is negligible when compared to the difference in behavior between two functions in different classes, say  $5n^2 + 3.2n$  and  $1.2n^3$ . Of course,  $99.7n^2$  and  $5n^2$  are different functions and we would prefer an algorithm that takes  $5n^2$  operations to an algorithm that takes  $99.7n^2$ . However, computer scientists typically ignore the leading constant and pay attention only to the fastest growing term.

When we write  $f(n) = O(n^2)$ , we mean that the function  $f(n)$  does not grow faster than a function with a leading term of  $cn^2$ , for a suitable choice of the constant  $c$ . In keeping with the healthy dose of pessimism toward an algorithm's performance, we measure an algorithm's efficiency as its worst case efficiency, which is the largest amount of time an algorithm can take given the worst possible input of a given size. The advantage to considering the worst case efficiency of an algorithm is that we are guaranteed that our algorithm will never behave worse than our worst case estimate, so we are never surprised or disappointed. Thus, when we derive a big-O bound, it is a bound on the worst case efficiency.


Consider two algorithms and denote by  $f(n)$  and  $g(n)$  their running times on an input of size  $n$ . We say that  $f$  grows no faster than  $g$ , if there exists a constant  $c$  such that for every positive integer  $n$ ,  $f(n) \leq c \cdot g(n)$

(equivalently,  $\frac{f(n)}{g(n)} \leq c$  for all  $n$ ). In this case, we write  $f = O(g)$  or  $f \leq g$ . The notation  $f = O(g)$  is the standard one, whereas some learners find the notation  $f \leq g$  to be more intuitive.

To give an example, let's visualize functions  $f(n) = 4n \log_2 n$  and  $g(n) = n^2$  by plotting them for  $1 \leq n \leq 30$  using a simple Python code. You can use this code as a template for plotting any other functions.

```
import matplotlib.pyplot as plt
import numpy as np

n = np.linspace(1, 30)
plt.plot(n, 4 * n * np.log2(n), label='$4n \log_2(n)$')
plt.plot(n, n ** 2, label='$n^2$')
plt.legend()
plt.savefig('plot_nlogn_and_n2.png')
```



As the picture reveals,  $4n \log_2 n \geq n^2$  for  $n \leq 16$ , but then the two functions switch. Indeed,  $4 \log_2 n \leq n$  for  $n \geq 16$ . If we set  $c = 16$ , then

$$4n \log_2 n \leq cn^2 \text{ for all } n \geq 1.$$

We conclude that  $4n \log_2 n = O(n^2)$ .


Now, let  $f(n) = 2n^2 + 5n + 3$  and  $g(n) = n^2$ . On the one hand,  $f(n)$  is larger than  $g(n)$  for all positive  $n$ . On the other hand,

$$f(n) = 2n^2 + 5n + 3 \leq 2n^2 + 5n^2 + 3n^2 = 10n^2 = 10g(n)$$

for all positive integers  $n$ . That is,  $f(n)$  is at most ten times larger than  $g(n)$ . We conclude that  $f$  grows no faster than  $g$  (and write  $f = O(g)$  or  $f \leq g$ ).

```
import matplotlib.pyplot as plt
import numpy as np


n = np.linspace(1, 100)
plt.plot(n, 2 * n ** 2 + 5 * n + 3, label='$2n^2+5n+3$')
plt.plot(n, n ** 2, label='$n^2$')
plt.plot(n, 10 * n ** 2, label='$10n^2$')
plt.legend(loc='upper left')
plt.savefig('plot_two_squares.png')
```



One can also plot the fraction  $\frac{f(n)}{g(n)}$ .

```
import matplotlib.pyplot as plt
import numpy as np

n = np.linspace(1, 100)
plt.plot(n, (2 * n ** 2 + 5 * n + 3) / (n ** 2))
plt.savefig('plot_two_squares_fraction.png')
```




This plot shows that (at least, in the considered range  $1 \leq n \leq 100$ ), the fraction  $\frac{f(n)}{g(n)}$  is at most 10 and, in fact, approaches 2 as  $n$  grows.

Now, let's compare  $f(n) = 11n$  and  $g(n) = 2n^2 + 5n + 3$ . First, let's look at their plots.

```
import matplotlib.pyplot as plt
import numpy as np

n = np.linspace(1, 100)
plt.plot(n, 11 * n, label='$11n$')
plt.plot(n, 2 * n ** 2 + 5 * n + 3, label='$2n^2+5n+3$')
plt.legend()
plt.savefig('plot_square_and_linear.png')
```



This plot reveals that both functions grow (as  $n$  grows) but  $11n$  grows “slower”. This can be formally stated as follows.

For two functions  $f, g$  we say that  $f$  grows slower than  $g$  and write  $f = o(g)$  or  $f < g$ , if the fraction  $\frac{f(n)}{g(n)}$  goes to zero as  $n$  grows.

**Exercise Break.** Plot the fraction  $\frac{11n}{2n^2+5n+3}$  to ensure that it goes to zero as  $n$  grows.

Of course, if  $f < g$  (equivalently,  $f = o(g)$ ), then also  $f \leq g$  (equivalently,  $f = O(g)$ ). In plain English: if  $f$  grows slower than  $g$ , then certainly  $f$  grows no faster than  $g$ .

### 1.6.1 Advantages and Disadvantages

Using the big- $O$  notation to evaluate running times has several advantages:

1. In many cases, computer scientists mainly care about how the running time grows with the input size — the big- $O$  notation clarifies the growth rate.
2. The big- $O$  notation simplifies the formulas for the running time:
  - $O(n^2)$  vs.  $3n^2 + 5n + 2$ .
  - $O(n)$  vs.  $n + \log_2 n + 7$ .
  - $O(n \log n)$  vs  $4n \log_2 n + 5$ . In particular,  $\log_2 n$ ,  $\log_3 n$ , and  $\log_a n$  differ by constant multipliers, so we don't need to specify the base of the logarithm in the big- $O$  notation.
3. With the big- $O$  notation, we no longer need to worry about things like how fast the computer is, or what the memory hierarchy looks like, or what compiler we used. Although these things have a big impact on the final running time, that impact will generally only be a constant multiple.

These advantages come with some disadvantages. Indeed, the big- $O$  notation “loses” some information since it ignores constant multipliers. If you have two algorithms, and one of them is a hundred times faster, they still have the same estimate of the running time in the big- $O$  notation. But, in practice, if you want to make things fast, a factor of 100 is a big deal.

Nevertheless, the big- $O$  notation is very useful and we will use it throughout this book.

### 1.6.2 Five Common Rules

Let's review the common rules of comparing the order of growth of functions arising in algorithm analysis.

1. Multiplicative constants can be omitted:

$$c \cdot f \leq f.$$

Examples:  $5n^2 \leq n^2$ ,  $\frac{n^2}{3} \leq n^2$ ,  $7n \leq n$ .

2. Out of two polynomials, the one with larger degree grows faster:

$$n^a < n^b \text{ for } 0 \leq a < b.$$

Examples:  $n < n^2$ ,  $\sqrt{n} < n^{2/3}$ ,  $n^2 < n^3$ ,  $n^0 < \sqrt{n}$ .

3. Any polynomial grows slower than any exponential:

$$n^a < b^n \text{ for } a \geq 0, b > 1.$$

Examples:  $n^3 < 2^n$ ,  $n^{10} < 1.1^n$ .

4. Any polylogarithm (that is, a function of the form  $(\log n)^a$ ) grows slower than any polynomial:

$$(\log n)^a < n^b \text{ for } a, b > 0.$$

Examples:  $(\log n)^3 < \sqrt{n}$ ,  $n \log n < n^2$ .

5. Smaller terms can be omitted:


$$\text{if } f \leq g, \text{ then } f + g \leq g.$$

Examples:  $n + n^2 \leq n^2$ ,  $n^9 + 2^n \leq 2^n$ .

## 1.6.3 Visualizing Common Rules

### 1.6.3.1 Multiplicative Constants Can Be Omitted


The following plot illustrates that that  $n^3$  and  $n^3/2$  have the same growth rate.






### 1.6.3.2 Out of Two Polynomials, The One With Larger Degree Grows Faster

The following plot illustrates that  $n^3$  grows faster than  $n^2$  and  $n^2$  grows faster than  $n$ .




### 1.6.3.3 Any Polynomial Grows Slower Than Any Exponential

The following plot illustrates that  $n^4$  grows faster than  $2^n$  in the range  $1 \leq n \leq 10$ .




The plot reveals that in this range  $n^4$  is always greater than  $2^n$ . This however does *not* mean that  $n^4$  grows faster than  $2^n$ . Take a look at a larger range  $1 \leq n \leq 20$ .




### 1.6.3.4 Any Polylogarithm Grows Slower Than Any Polynomial


The following plot illustrates that  $\log n$  grows slower than  $n$ .



Now, let's compare  $(\log n)^3$  versus  $\sqrt{n}$ . The following plot illustrates that the polylogarithm  $((\log n)^3)$  grows faster than the polynomial  $n^{1/2}$  in the range  $[1, 100]$ .




After extending the range from  $[1, 100]$  to  $[1, 1\,000\,000]$ , the polylogarithmic function is still above the polynomial one.




The following plot illustrates that it is in fact dangerous to decide which function grows faster just by looking at how they behave for some values of  $n$ . The rule “any polynomial grows faster than any polylogarithm” means that *eventually* any polynomial will become larger than any polylogarithm. But the rule does not specify for what value of  $n$  this happens for the first time.

To show that  $\sqrt{n}$  grows faster than  $(\log n)^3$ , let's increase the range to  $10^8$ .




Moreover, let's consider an even larger interval to make sure that these two functions don't switch back.



**Exercise Break.** Find the value of  $n$ , where  $n^{0.1}$  becomes larger than  $(\log n)^5$ .

#### 1.6.3.5 Smaller Terms Can Be Omitted

The plot below shows that the contribution of  $n$  into  $n^2 + n$  is negligible for large  $n$ :  $n + n^2$  and  $n^2$  have the same growth rate.



#### 1.6.4 Frequently Arising Running Times

The table below shows an approximate running time of four algorithms on various input sizes for a hypothetical computer that performs  $10^9$  operations per second. The last row shows the maximum value of  $n$  for which the running time of the corresponding algorithm fits into one second.

	$n$	$n \log n$	$n^2$	$2^n$
$n = 20$	< 1 sec	< 1 sec	< 1 sec	< 1 sec
$n = 50$	< 1 sec	< 1 sec	< 1 sec	13 days
$n = 10^2$	< 1 sec	< 1 sec	< 1 sec	$4 \cdot 10^{13}$ year
$n = 10^6$	< 1 sec	< 1 sec	17 min	
$n = 10^9$	1 sec	30 sec	30 years	
max $n$	$10^9$	$10^{7.5}$	$10^{4.5}$	30

In all programming challenges in this book, your goal is to implement a program that works in at most a few seconds for all allowed input sizes  $n$ . If  $n$  can be as large as  $10^6$ , then you want to design an algorithm with running time  $O(n)$  or  $O(n \log n)$ . If  $n$  is at most  $10^4$ , then you probably want your algorithm to work in time  $O(n^2)$ . And if  $n \leq 20$ , then even an algorithm with the running time  $O(2^n)$  should be OK.



# Chapter 2: Algorithm Design Techniques

Over the last half a century, computer scientists have discovered that many algorithms share similar ideas, even though they solve very different problems. There appear to be relatively few basic techniques that can be applied when designing an algorithm, and we cover some of them later in various programming challenges in this book. For now we will mention the most common algorithm design techniques, so that future examples can be categorized in terms of the algorithm's design methodology.

To illustrate the design techniques, we will consider a very simple problem that plagued nearly everyone before the era of mobile phones when people used cordless phones. Suppose your cordless phone rings, but you have misplaced the handset somewhere in your home. How do you find it? To complicate matters, you have just walked into your home with an armful of groceries, and it is dark out, so you cannot rely solely on eyesight.

## 2.1 Exhaustive Search Algorithms

An *exhaustive search*, or *brute force*, algorithm examines every possible alternative to find one particular solution. For example, if you used the brute force algorithm to find the ringing telephone, you would ignore the ringing of the phone, as if you could not hear it, and simply walk over every square inch of your home checking to see if the phone was present. You probably would not be able to answer the phone before it stopped ringing, unless you were very lucky, but you would be guaranteed to eventually find the phone no matter where it was.

BRUTEFORCECHANGE is a brute force algorithm, and our programming challenges include some additional examples of such algorithms—these are the easiest algorithms to design, and sometimes they work for certain practical problems. In general, though, brute force algorithms are too slow to be practical for anything but the smallest instances and you should always think how to avoid the brute force algorithms or how to finesse them into faster versions.

## 2.2 Branch-and-Bound Algorithms

In certain cases, as we explore the various alternatives in a brute force algorithm, we discover that we can omit a large number of alternatives, a technique that is often called *branch-and-bound*.

Suppose you were exhaustively searching the first floor and heard the phone ringing above your head. You could immediately rule out the need to search the basement or the first floor. What may have taken three hours may now only take one, depending on the amount of space that you can rule out.

## 2.3 Greedy Algorithms

Many algorithms are iterative procedures that choose among a number of alternatives at each iteration. For example, a cashier can view the Change Problem as a series of decisions he or she has to make: which coin (among  $d$  denominations) to return first, which to return second, and so on. Some of these alternatives may lead to correct solutions while others may not.

Greedy algorithms choose the “most attractive” alternative at each iteration, for example, the largest denomination possible. In the case of the US denominations, CHANGE used quarters, then dimes, then nickels, and finally pennies (in that order) to make change. Of course, we showed that this greedy strategy produced incorrect results when certain new denominations were included.

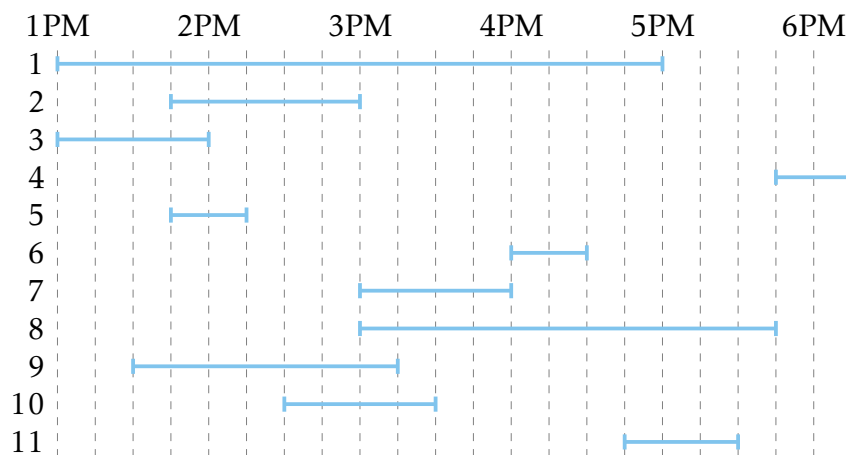
In the telephone example, the corresponding greedy strategy would be to walk in the direction of the telephone’s ringing until you found it. The problem here is that there may be a wall (or a fragile vase) between you and the phone, preventing you from finding it. Unfortunately, these sorts of difficulties frequently occur in most realistic problems. In many cases, a greedy approach will seem “obvious” and natural, but will be subtly wrong.

In the Activity Selection Problem ([try it online!](#)), one is given a set of intervals and is asked to select as many of them as possible so that no two of them overlap (two intervals are overlapping if there is a point that belongs to both of them). The name of the problem comes from the following scenario. Imagine a meeting room and the following requests to hold a meeting in this room from eleven groups of people.



1	1:00PM–5:00PM
2	1:45PM–3:00PM
3	1:00PM–2:00PM
4	5:45PM–6:15PM
5	1:45PM–2:15PM
6	4:00PM–4:30PM
7	3:00PM–4:00PM
8	3:00PM–5:45PM
9	1:30PM–3:15PM
10	2:30PM–3:30PM
11	4:45PM–5:30PM


We cannot satisfy all the requests (as some of them are overlapping), but we want to satisfy as many of them as possible. To do this, let us start by visualizing the input data, in a more convenient format.



As we are talking about greedy strategies, let's experiment with various "most profitable" moves here. One common-sense approach is to select the shortest interval, remove all intervals that overlap it, and iterate.

**Stop and Think.** Does this always lead to an optimum solution?


It turns out that this natural greedy strategy can give a suboptimal solution. In the example below, it returns a solution consisting of a single segment (the middle one) where as there is a solution consisting of two non-overlapping segments.



Another common-sense approach is to select the leftmost interval with the earliest start time, remove all intervals that overlap it, and iterate.

**Stop and Think.** Does this lead to an optimum solution?

The following example illustrates that this strategy may also fail.



**Stop and Think.** Do you see any other natural greedy strategy for this problem?

It turns out that the following greedy algorithm maximizes the number of non-overlapping segments: repeatedly select an interval with the smallest endpoint (referred to as the *champion* interval), remove all intervals that start before this endpoint, and iterate.

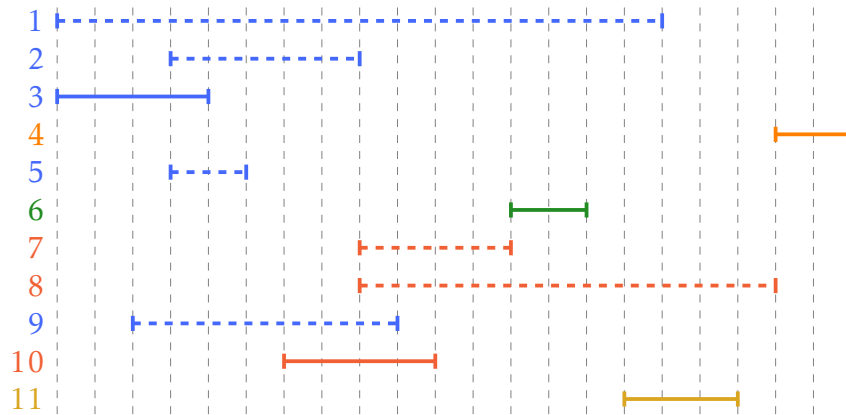
**Exercise Break.** Prove that if a set of non-overlapping intervals does not contain the champion interval then substituting the first interval in this set by the champion interval results in a set of non-overlapping intervals.

This Exercise Break explains why the described greedy strategy is optimal! Indeed, since there exists a solution of the Activity Selection Problem that contains the champion interval, you can safely select the champion interval at the first step and iterate!

For our working example, the algorithm proceeds as follows.

1. Take segment 3 and discard segments 1, 2, 5, and 9.
2. Take segment 10 and discard segments 7 and 8.
3. Take segment 6.
4. Take segment 11.

5. Take segment 4.



## 2.4 Dynamic Programming Algorithms

Some algorithms break a problem into smaller subproblems and use the solutions of the subproblems to construct the solution of the larger one. During this process, the number of subproblems may become very large, and some algorithms solve the same subproblem repeatedly, needlessly increasing the running time. Dynamic programming organizes computations to avoid recomputing values that you already know, which can often save a great deal of time.

The Ringing Telephone Problem does not lend itself to a dynamic programming solution, so we consider a different problem to illustrate the technique. Suppose that instead of answering the phone you decide to play the “Rocks” game ([try it online!](#)) with two piles of rocks, say ten in each. In each turn, one player may take either one rock (from either pile) or two rocks (one from each pile). Once the rocks are taken, they are removed from play. The player that takes the last rock wins the game. You make the first move. We encourage you to play this game using our interactive puzzle.

To find the winning strategy for the  $10 + 10$  game, we can construct a table, which we can call  $R$ , shown in Figure 2.1. Instead of solving a problem with 10 rocks in each pile, we will solve a more general problem with  $n$  rocks in one pile and  $m$  rocks in the other pile (the  $n + m$  game) where  $n$  and  $m$  are arbitrary non-negative integers.

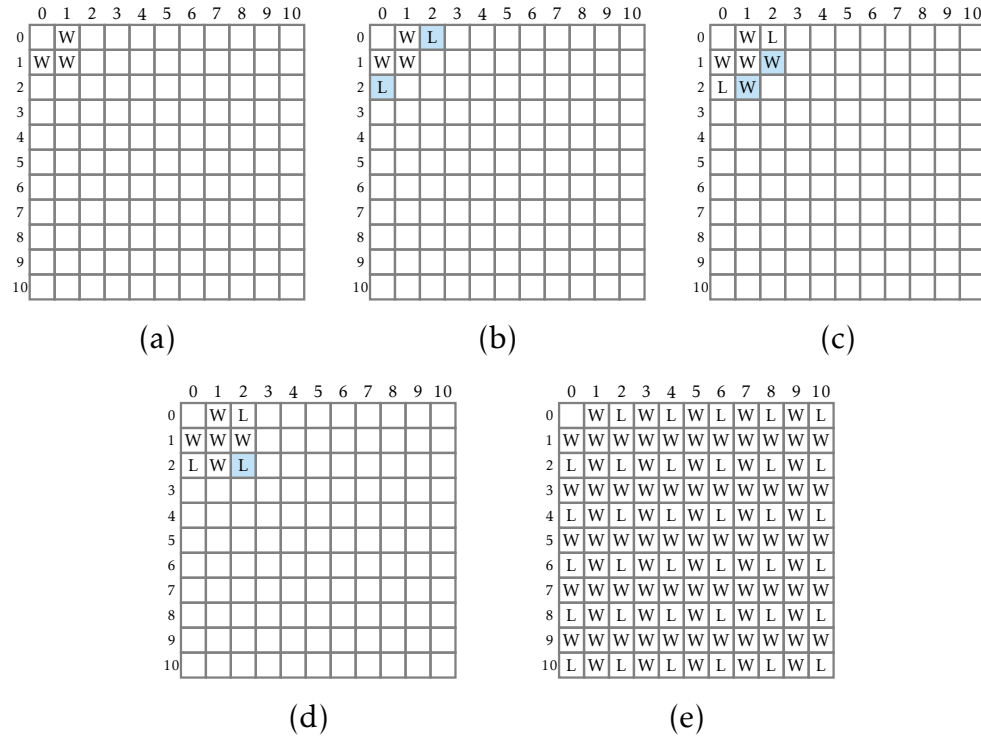


Figure 2.1: Table  $R$  for the  $10 + 10$  Rocks game.

If Player 1 can always win the  $n + m$  game, then we would say  $R(n, m) = W$ , but if Player 1 has no winning strategy against a player that always makes the right moves, we would write  $R(n, m) = L$ . Computing  $R(n, m)$  for arbitrary  $n$  and  $m$  seems difficult, but we can build on smaller values. Some games, notably  $R(0, 1)$ ,  $R(1, 0)$ , and  $R(1, 1)$ , are clearly winning propositions for Player 1 since in the first move, Player 1 can win. Thus, we fill in entries  $(1, 1)$ ,  $(0, 1)$ , and  $(1, 0)$  as  $W$ . See Figure 2.1(a).

After the entries  $(0, 1)$ ,  $(1, 0)$ , and  $(1, 1)$  are filled, one can try to fill other entries. For example, in the  $(2, 0)$  case, the only move that Player 1 can make leads to the  $(1, 0)$  case that, as we already know, is a winning position for his opponent. A similar analysis applies to the  $(0, 2)$  case, leading to the table in Figure 2.1(b).

In the  $(2, 1)$  case, Player 1 can make three different moves that lead respectively to the games of  $(1, 1)$ ,  $(2, 0)$ , or  $(1, 0)$ . One of these cases,  $(2, 0)$ ,

leads to a losing position for his opponent and therefore  $(2, 1)$  is a winning position. The case  $(1, 2)$  is symmetric to  $(2, 1)$ , so we have the table shown at Figure 2.1(c).

Now we can fill in  $R(2, 2)$ . In the  $(2, 2)$  case, Player 1 can make three different moves that lead to entries  $(2, 1)$ ,  $(1, 2)$ , and  $(1, 1)$ . All of these entries are winning positions for his opponent and therefore  $R(2, 2) = L$ , see Figure 2.1(d).

We can proceed filling in  $R$  in this way by noticing that for the entry  $(i, j)$  to be  $L$ , the entries above, diagonally to the left, and directly to the left, must be  $W$ . These entries  $((i - 1, j)$ ,  $(i - 1, j - 1)$ , and  $(i, j - 1))$  correspond to the three possible moves that Player 1 can make. See Figure 2.1(e).

The Rocks algorithm determines if Player 1 wins or loses. If Player 1 wins in an  $n + m$  game, Rocks returns  $W$ . If Player 1 loses, Rocks returns  $L$ . We introduced an artificial initial condition,  $R(0, 0) = L$  to simplify the pseudocode.

```

Rocks( $n, m$ ):
 $R(0, 0) \leftarrow L$ 
for  $i$  from 1 to  $n$ :
    if  $R(i - 1, 0) = W$ :
         $R(i, 0) \leftarrow L$ 
    else:
         $R(i, 0) \leftarrow W$ 
for  $j$  from 1 to  $m$ :
    if  $R(0, j - 1) = W$ :
         $R(0, j) \leftarrow L$ 
    else:
         $R(0, j) \leftarrow W$ 
for  $i$  from 1 to  $n$ :
    for  $j$  from 1 to  $m$ :
        if  $R(i - 1, j - 1) = W$  and  $R(i, j - 1) = W$  and  $R(i - 1, j) = W$ :
             $R(i, j) \leftarrow L$ 
        else:
             $R(i, j) \leftarrow W$ 
return  $R(n, m)$ 

```

A faster algorithm to solve the Rocks puzzle relies on the simple pattern

in  $R$ , and checks if  $n$  and  $m$  are both even, in which case the player loses (see table above).

```
FASTROCKS( $n, m$ ):  
if  $n$  and  $m$  are both even:  
    return  $L$   
else:  
    return  $W$ 
```

However, though FASTROCKS is more efficient than ROCKS, it may be difficult to modify it for similar games, for example, a game in which each player can move up to three rocks at a time from the piles. This is one example where the slower algorithm is more instructive than a faster one.

**Exercise Break.** Play the Three Rocks game using our interactive puzzle ([try it online!](#)) and construct the dynamic programming table similar to the table above for this game.

## 2.5 Recursive Algorithms

Recursion is one of the most ubiquitous algorithmic concepts. Simply, an algorithm is recursive if it calls itself.

The *Towers of Hanoi puzzle* ([try it online!](#)) consists of three pegs, which we label from left to right as 1, 2, and 3, and a number of disks of decreasing radius, each with a hole in the center. The disks are initially stacked on the left peg (peg 1) so that smaller disks are on top of larger ones. The game is played by moving one disk at a time between pegs. You are only allowed to place smaller disks on top of larger ones, and any disk may go onto an empty peg. The puzzle is solved when all of the disks have been moved from peg 1 to peg 3. Try our interactive puzzle Hanoi Towers to figure out how to move all disks from one peg to another.

---

### Towers of Hanoi Problem

*Output a list of moves that solves the Towers of Hanoi.*

**Input:** An integer  $n$ .


**Output:** A sequence of moves that solve the  $n$ -disk Towers of Hanoi puzzle.

---


Solving the puzzle with one disk is easy: move the disk to the right peg. The two-disk puzzle is not much harder: move the small disk to the middle peg, then the large disk to the right peg, then the small disk to the right peg to rest on top of the large disk.

The three-disk puzzle is somewhat harder, but the following sequence of seven moves solves it:


1. Move disk from peg 1 to peg 3




2. Move disk from peg 1 to peg 2




3. Move disk from peg 3 to peg 2




4. Move disk from peg 1 to peg 3




5. Move disk from peg 2 to peg 1



6. Move disk from peg 2 to peg 3



7. Move disk from peg 1 to peg 3



Now we will figure out how many steps are required to solve a four-disk puzzle. You cannot complete this game without moving the largest disk. However, in order to move the largest disk, we first had to move all the smaller disks to an empty peg. If we had four disks instead of three, then we would first have to move the top three to an empty peg (7 moves), then move the largest disk (1 move), then again move the three disks from their temporary peg to rest on top of the largest disk (another 7 moves). The whole procedure will take  $7 + 1 + 7 = 15$  moves.

More generally, to move a stack of size  $n$  from the left to the right peg, you first need to move a stack of size  $n - 1$  from the left to the middle peg, and then from the middle peg to the right peg once you have moved the  $n$ -th disk to the right peg. To move a stack of size  $n - 1$  from the middle to the right, you first need to move a stack of size  $n - 2$  from the middle to the left, then move the  $(n - 1)$ -th disk to the right, and then move the stack of size  $n - 2$  from the left to the right peg, and so on.



At first glance, the Towers of Hanoi Problem looks difficult. However, the following *recursive algorithm* solves the Towers of Hanoi Problem with just 9 lines!

```
HANOITOWERS( $n$ ,  $fromPeg$ ,  $toPeg$ )
if  $n = 1$ :
    output "Move disk from peg  $fromPeg$  to peg  $toPeg$ "
    return
 $unusedPeg \leftarrow 6 - fromPeg - toPeg$ 
HANOITOWERS( $n - 1$ ,  $fromPeg$ ,  $unusedPeg$ )
output "Move disk from peg  $fromPeg$  to peg  $toPeg$ "
HANOITOWERS( $n - 1$ ,  $unusedPeg$ ,  $toPeg$ )
```

The variables  $fromPeg$ ,  $toPeg$ , and  $unusedPeg$  refer to the three different pegs so that  $HANOITOWERS(n, 1, 3)$  moves  $n$  disks from the first peg to the third peg. The variable  $unusedPeg$  represents which of the three pegs can serve as a temporary destination for the first  $n - 1$  disks. Note that  $fromPeg + toPeg + unusedPeg$  is always equal to  $1 + 2 + 3 = 6$ , so the value of the variable  $unusedPeg$  can be computed as  $6 - fromPeg - toPeg$ . Table below shows the result of  $6 - fromPeg - toPeg$  for all possible values of  $fromPeg$  and  $toPeg$ .

$fromPeg$	$toPeg$	$unusedPeg$
1	2	3
1	3	2
2	1	3
2	3	1
3	1	2
3	2	1

After computing  $unusedPeg$  as  $6 - fromPeg - toPeg$ , the statements

```
HANOITOWERS( $n - 1$ ,  $fromPeg$ ,  $unusedPeg$ )
output "Move disk from peg  $fromPeg$  to peg  $toPeg$ "
HANOITOWERS( $n - 1$ ,  $unusedPeg$ ,  $toPeg$ )
```

solve the smaller problem of moving the stack of size  $n - 1$  first to the temporary space, moving the largest disk, and then moving the  $n - 1$

remaining disks to the final destination. Note that we do not have to specify which disk the player should move from *fromPeg* to *toPeg*: it is always the top disk currently residing on *fromPeg* that gets moved.

**Stop and Think.** How many moves are needed to move 6 disks?

Although the Hanoi Tower solution can be expressed in just 9 lines of pseudocode, it requires a surprisingly long time to run. To solve a five-disk tower requires 31 moves, but to solve a hundred-disk tower would require more moves than there are atoms on Earth. The fast growth of the number of moves that HANOITOWERS requires is easy to see by noticing that every time HANOITOWERS( $n, 1, 3$ ) is called, it calls itself twice for  $n - 1$ , which in turn triggers four calls for  $n - 2$ , and so on.

We can illustrate this situation in a *recursion tree*, which is shown in Figure 2.2. A call to HANOITOWERS(4, 1, 3) results in calls HANOITOWERS(3, 1, 2) and HANOITOWERS(3, 2, 3); each of these results in calls to HANOITOWERS(2, 1, 3), HANOITOWERS(2, 3, 2) and HANOITOWERS(2, 2, 1), HANOITOWERS(2, 1, 3), and so on. Each call to the subroutine HANOITOWERS requires some amount of time, so we would like to know how much time the algorithm will take.

To calculate the running time of HANOITOWERS of size  $n$ , we denote the number of disk moves that HANOITOWERS( $n$ ) performs as  $T(n)$  and notice that the following equation holds:

$$T(n) = 2 \cdot T(n - 1) + 1.$$

Starting from  $T(1) = 1$ , this recurrence relation produces the sequence:

$$1, 3, 7, 15, 31, 63,$$

and so on. We can compute  $T(n)$  by adding 1 to both sides and noticing

$$T(n) + 1 = 2 \cdot T(n - 1) + 1 + 1 = 2 \cdot (T(n - 1) + 1).$$

If we introduce a new variable,  $U(n) = T(n) + 1$ , then  $U(n) = 2 \cdot U(n - 1)$ . Thus, we have changed the problem to the following recurrence relation.

$$U(n) = 2 \cdot U(n - 1).$$

Starting from  $U(1) = 2$ , this gives rise to the sequence

$$2, 4, 8, 16, 32, 64, \dots$$




Figure 2.2: The recursion tree for a call to `HANOITOWERS(4,1,3)`, which solves the Towers of Hanoi problem of size 4. At each point in the tree,  $(i, j, k)$  stands for `HANOITOWERS(i, j, k)`.

implying that at  $U(n) = 2^n$  and  $T(n) = U(n) - 1 = 2^n - 1$ . Thus, `HANOITOWERS`( $n$ ) is an exponential algorithm.

## 2.6 Divide-and-Conquer Algorithms

One big problem may be hard to solve, but two problems that are half the size may be significantly easier. In these cases, divide-and-conquer algorithms fare well by doing just that: splitting the problem into smaller subproblems, solving the subproblems independently, and combining the solutions of subproblems into a solution of the original problem. The situation is usually more complicated than this and after splitting one problem into subproblems, a divide-and-conquer algorithm usually splits these subproblems into even smaller sub-subproblems, and so on, until it reaches a point at which it no longer needs to recurse. A critical step in many divide-and-conquer algorithms is the recombining of solutions to subproblems into a solution for a larger problem.

To give an example of a divide-and conquer algorithm, we will consider the sorting problem:

---

### Sorting Problem

*Sort a list of integers.*

**Input:** A list of  $n$  distinct integers  $a = (a_1, a_2, \dots, a_n)$ .

**Output:** Sorted list of integers, that is, a reordering  $(b_1, b_2, \dots, b_n)$  of integers from  $a$  such that  $b_1 < b_2 < \dots < b_n$ .

---

`SELECTIONSORT` is a simple iterative method to solve the Sorting Problem. It first finds the smallest element in  $a$ , and moves it to the first position by swapping it with whatever happens to be in the first position (i.e.,  $a_1$ ). Next, it finds the second smallest element in  $a$ , and moves it to the second position, again by swapping with  $a_2$ . At the  $i$ -th iteration, `SELECTIONSORT` finds the  $i$ -th smallest element in  $a$ , and moves it to the  $i$ -th position. This is an intuitive approach at sorting, but is not the fastest one.

If  $a = (7, 92, 87, 1, 4, 3, 2, 6)$ , `SELECTIONSORT`( $a, 8$ ) takes the following seven steps:

(7, 92, 87, 1, 4, 3, 2, 6)

(1, 92, 87, 7, 4, 3, 2, 6)  
 (1, 2, 87, 7, 4, 3, 92, 6)  
 (1, 2, 3, 7, 4, 87, 92, 6)  
 (1, 2, 3, 4, 7, 87, 92, 6)  
 (1, 2, 3, 4, 6, 87, 92, 7)  
 (1, 2, 3, 4, 6, 7, 92, 87)  
 (1, 2, 3, 4, 6, 7, 87, 92)

MERGESORT is a canonical example of divide-and-conquer sorting algorithm that is much faster than SELECTIONSORT. We begin from the problem of *merging*, in which we want to combine two sorted lists  $List_1$  and  $List_2$  into a single sorted list.

$List_1$	2578		2578		2578		2578		2578		2578
$List_2$	346		346		346		346		346		346
sortedList	2		3		4		5		6		78

The MERGE algorithm combines two sorted lists into a single sorted list in  $O(|List_1| + |List_2|)$  time by iteratively choosing the smallest remaining element in  $List_1$  and  $List_2$  and moving it to the growing sorted list.

```

MERGE( $List_1, List_2$ ):
  SortedList  $\leftarrow$  empty list
  while both  $List_1$  and  $List_2$  are non-empty:
    if the smallest element in  $List_1$  is smaller than the smallest element in  $List_2$ :
      move the smallest element from  $List_1$  to the end of SortedList
    else:
      move the smallest element from  $List_2$  to the end of SortedList
  move any remaining elements from either  $List_1$  or  $List_2$  to the end of SortedList
  return SortedList
  
```

MERGE would be useful for sorting an arbitrary list if we knew how to divide an arbitrary (unsorted) list into two already sorted half-sized lists. However, it may seem that we are back to where we started, except now we have to sort two smaller lists instead of one big one. Yet sorting two smaller lists is a preferable algorithmic problem. To see why, let's consider the MERGESORT algorithm, which divides an unsorted list into two parts and

then recursively conquers each smaller sorting problem before merging the sorted lists.

```
MERGESORT(List):  
if List consists of a single element:  
    return List  
FirstHalf ← first half of List  
SecondHalf ← second half of List  
SortedFirstHalf ← MERGESORT(FirstHalf)  
SortedSecondHalf ← MERGESORT(SecondHalf)  
SortedList ← MERGE(SortedFirstHalf, SortedSecondHalf)  
return SortedList
```

Figure 2.3 shows the recursion tree of MERGESORT, consisting of  $\log_2 n$  levels, where  $n$  is the size of the original unsorted list. At the bottom level, we must merge two sorted lists of approximately  $n/2$  elements each, requiring  $O(n/2 + n/2) = O(n)$  time. At the next highest level, we must merge four lists of  $n/4$  elements, requiring  $O(n/4 + n/4 + n/4 + n/4) = O(n)$  time. This pattern can be generalized: the  $i$ -th level contains  $2^i$  lists, each having approximately  $n/2^i$  elements, and requires  $O(n)$  time to merge. Since there are  $\log_2 n$  levels in the recursion tree, MERGESORT requires  $O(n \log_2 n)$  runtime overall, which offers a speedup over a naive  $O(n^2)$  sorting algorithm.

**Exercise Break.** Use the divide-and-conquer strategy to solve our [21 questions](#) interactive puzzle.

## 2.7 Randomized Algorithms

If you happen to have a coin, then before even starting to search for the phone, you could toss it to decide whether you want to start your search on the first floor if the coin comes up heads, or on the second floor if the coin comes up tails. If you also happen to have a die, then after deciding on the second floor of your mansion, you could roll it to decide in which of the six rooms on the second floor to start your search. Although tossing coins and rolling dice may be a fun way to search for the phone, it is certainly not the intuitive thing to do, nor is it at all clear whether it gives you any




Figure 2.3: The recursion tree for sorting an 8-element array with MERGESORT. The divide (upper) steps consist of  $\log_2 8 = 3$  levels, where the input array is split into smaller and smaller subarrays. The conquer (lower) steps consist of the same number of levels, as the sorted subarrays are merged back together.

algorithmic advantage over a deterministic algorithm. Our programming challenges will help you to learn why randomized algorithms are useful and why some of them have a competitive advantage over deterministic algorithms.

To give an example of a randomized algorithm, we will first discuss a fast sorting technique called QUICKSORT. It selects an element  $m$  (typically, the first) from an array  $c$  and simply partitions the array into two subarrays:  $c_{small}$ , containing all elements from  $c$  that are smaller than  $m$ ; and  $c_{large}$  containing all elements larger than  $m$ .

This partitioning can be done in linear time, and by following a divide-and-conquer strategy, QUICKSORT recursively sorts each subarray in the same way. The sorted list is easily created by simply concatenating the sorted  $c_{small}$ , element  $m$ , and the sorted  $c_{large}$ .

```
QUICKSORT( $c$ ):  
if  $c$  consists of a single element:  
    return  $c$   
 $m \leftarrow c[1]$   
determine the set of elements  $c_{small}$  smaller than  $m$   
determine the set of elements  $c_{large}$  larger than  $m$   
QUICKSORT( $c_{small}$ )  
QUICKSORT( $c_{large}$ )  
combine  $c_{small}$ ,  $m$ , and  $c_{large}$  into a sorted array  $c_{sorted}$   
return  $c_{sorted}$ 
```

It turns out that the running time of QUICKSORT depends on how lucky we are with our selection of the element  $m$ . If we happen to choose  $m$  in such a way that the array  $c$  is split into even halves (i.e.,  $|c_{small}| = |c_{large}|$ ), then

$$T(n) = 2T\left(\frac{n}{2}\right) + a \cdot n,$$

where  $T(n)$  represents the time taken by QUICKSORT to sort an array of  $n$  numbers, and  $a \cdot n$  represents the time required to split the array of size  $n$  into two parts;  $a$  is a positive constant. This is exactly the same recurrence as in MERGESORT that leads to  $O(n \log n)$  running time.

However, if we choose  $m$  in such a way that it splits  $c$  unevenly (e.g., an extreme case occurs when  $c_{small}$  is empty and  $c_{large}$  has  $n - 1$  elements),



then the recurrence looks like

$$T(n) = T(n-1) + a \cdot n.$$

This is the recurrence that leads to  $O(n^2)$  running time, something we want to avoid. Indeed, QUICKSORT takes quadratic time to sort the array  $(n, n-1, \dots, 2, 1)$ . Worse yet, it requires  $O(n^2)$  time to process  $(1, 2, \dots, n-1, n)$ , which seems unnecessary since the array is already sorted.

The QUICKSORT algorithm so far seems like a bad imitation of MERGESORT. However, if we can choose a good “splitter”  $m$  that breaks an array into two equal parts, we might improve the running time. To achieve  $O(n \log n)$  running time, it is not actually necessary to find a perfectly equal (50/50) split. For example, a split into approximately equal parts of size, say, 51/49 will also work. In fact, one can prove that the algorithm will achieve  $O(n \log n)$  running time as long as the sets  $c_{small}$  and  $c_{large}$  are both larger in size than  $n/4$ .

It implies that, of  $n$  possible choices for  $m$  as elements of the array  $c$ , at least  $\frac{3n}{4} - \frac{n}{4} = \frac{n}{2}$  of them make good splitters! In other words, if we randomly choose  $m$  (i.e., every element of the array  $c$  has the same probability to be chosen), there is at least a 50% chance that it will be a good splitter. This observation motivates the following randomized algorithm:

```
RANDOMIZEDQUICKSORT( $c$ ):  
if  $c$  consists of a single element:  
    return  $c$   
randomly select an element  $m$  from  $c$   
determine the set of elements  $c_{small}$  smaller than  $m$   
determine the set of elements  $c_{large}$  larger than  $m$   
RANDOMIZEDQUICKSORT( $c_{small}$ )  
RANDOMIZEDQUICKSORT( $c_{large}$ )  
combine  $c_{small}$ ,  $m$ , and  $c_{large}$  into a sorted array  $c_{sorted}$   
return  $c_{sorted}$ 
```

RANDOMIZEDQUICKSORT is a fast algorithm in practice, but its worst case running time remains  $O(n^2)$  since there is still a possibility that it selects bad splitters. Although the behavior of a randomized algorithm varies on the same input from one execution to the next, one can prove that its *expected* running time is  $O(n \log n)$ . The running time of a randomized algorithm is a *random variable*, and computer scientists are often interested

in the mean value of this random variable which is referred to as the expected running time.

The key advantage of randomized algorithms is performance: for many practical problems randomized algorithms are faster (in the sense of expected running time) than the best known deterministic algorithms. Another attractive feature of randomized algorithms, as illustrated by RANDOMIZEDQUICKSORT, is their simplicity.

We emphasize that RANDOMIZEDQUICKSORT, despite making random decisions, always returns the correct solution of the sorting problem. The only variable from one run to another is its running time, not the result. In contrast, other randomized algorithms usually produce incorrect (or, more gently, *approximate*) solutions. Randomized algorithms that always return correct answers are called *Las Vegas algorithms*, while algorithms that do not are called *Monte Carlo algorithms*. Of course, computer scientists prefer Las Vegas algorithms to Monte Carlo algorithms, but the former are often difficult to come by.

# Chapter 3: Programming Challenges

To introduce you to our automated grading system, we will discuss two simple programming challenges and walk you through a step-by-step process of solving them. We will encounter several common pitfalls and will show you how to fix them.

Below is a brief overview of what it takes to solve a programming challenge in five steps:

**Reading problem statement.** Problem statement specifies the input-output format, the constraints for the input data as well as time and memory limits. Your goal is to implement a fast program that solves the problem and works within the time and memory limits.

**Designing an algorithm.** When the problem statement is clear, start designing an algorithm and don't forget to prove that it works correctly.

**Implementing an algorithm.** After you developed an algorithm, start implementing it in a programming language of your choice.

**Testing and debugging your program.** Testing is the art of revealing bugs. Debugging is the art of exterminating the bugs. When your program is ready, start testing it! If a bug is found, fix it and test again.

**Submitting your program to the grading system.** After testing and debugging your program, submit it to the grading system and wait for the message "Good job!". In the case you see a different message, return back to the previous stage.

## 3.1 Sum of Two Digits

---

### Sum of Two Digits Problem

*Compute the sum of two single digit numbers.*

**Input:** Two single digit numbers.

**Output:** The sum of these numbers.

$$2 + 3 = 5$$

---

We start from this ridiculously simple problem to show you the pipeline of reading the problem statement, designing an algorithm, implementing it, testing and debugging your program, and submitting it to the grading system.

**Input format.** Integers  $a$  and  $b$  on the same line (separated by a space).

**Output format.** The sum of  $a$  and  $b$ .

**Constraints.**  $0 \leq a, b \leq 9$ .

**Sample.**

Input:

9 7

Output:

16

**Time limits (sec.):**

---

C++	Java	Python	C	C#	Go	Haskell	JavaScript	Kotlin	Ruby	Rust	Scala
1	1.5	5	1	1.5	1.5	2	5	1.5	5	1	3

---

**Memory limit.** 512 Mb.

For this trivial problem, we will skip “Designing an algorithm” step and will move right to the pseudocode.

```
SUMOfTwoDIGITS( $a$ ,  $b$ ):  
return  $a + b$ 
```

Since the pseudocode does not specify how we input  $a$  and  $b$ , below we provide solutions in C++, Java, and Python3 programming languages as well as recommendations on compiling and running them. You can copy-and-paste the code to a file, compile/run it, test it on a few datasets, and then submit (the source file, not the compiled executable) to the grading system. Needless to say, we assume that you know the basics of one of programming languages that we use in our grading system.

## C++

```
#include <iostream>  
  
int sum_of_digits(int first, int second) {  
    return first + second;  
}  
  
int main() {  
    int a = 0;  
    int b = 0;  
    std::cin >> a;  
    std::cin >> b;  
    std::cout << sum_of_digits(a, b);  
    return 0;  
}
```

Save this to a file (say, `aplusb.cpp`), compile it, run the resulting executable, and enter two numbers (on the same line).

## Java

```
import java.util.Scanner;  
  
class SumOfDigits {  
    static int sumOfDigits(int first_digit,
```

```

        int second_digit) {
            return first_digit + second_digit;
        }

        public static void main(String[] args) {
            Scanner s = new Scanner(System.in);
            int a = s.nextInt();
            int b = s.nextInt();
            System.out.println(sumOfDigits(a, b));
        }
    }
}

```

Save this to a file `SumOfDigits.java`, compile it, run the resulting executable, and enter two numbers (on the same line).

## Python

```

def sum_of_digits(first_digit, second_digit):
    return first_digit + second_digit

if __name__ == '__main__':
    a, b = map(int, input().split())
    print(sum_of_digits(a, b))

```

Save this to a file (say, `aplusb.py`), run it, and enter two numbers on the same line.

Your goal is to implement an algorithm that produces a correct result under the given time and memory limits for any input satisfying the given constraints. You do not need to check that the input data satisfies the constraints, e.g., for the Sum of Two Digits Problem you do not need to check that the given integers  $a$  and  $b$  are indeed single digit integers (this is guaranteed).

## 3.2 Maximum Pairwise Product

---

### Maximum Pairwise Product Problem

*Find the maximum product of two distinct numbers in a sequence of non-negative integers.*

**Input:** An integer  $2 \leq n \leq 2 \cdot 10^5$  and a sequence of  $n$  non-negative integers.

**Output:** The maximum value that can be obtained by multiplying two different elements from the sequence.

	5	6	2	7	4
5		30	10	35	20
6	30		12	42	24
2	10	12		14	8
7	35	42	14		28
4	20	24	8	28	

---

Given a sequence of non-negative integers  $a_1, \dots, a_n$ , compute

$$\max_{1 \leq i \neq j \leq n} a_i \cdot a_j.$$

Note that  $i$  and  $j$  should be different, though it may be the case that  $a_i = a_j$ .

**Input format.** The first line contains an integer  $n$ . The next line contains  $n$  non-negative integers  $a_1, \dots, a_n$  (separated by spaces).

**Output format.** The maximum pairwise product.

**Constraints.**  $2 \leq n \leq 2 \cdot 10^5$ ;  $0 \leq a_1, \dots, a_n \leq 2 \cdot 10^5$ .

#### Sample 1.

Input:

```
3
1 2 3
```

Output:

```
6
```

### Sample 2.

Input:

```
10
7 5 14 2 8 8 10 1 2 3
```

Output:

```
140
```

**Time and memory limits.** The same as for the previous problem.

### 3.2.1 Naive Algorithm

A naive way to solve the Maximum Pairwise Product Problem is to go through all possible pairs of the input elements  $A[1 \dots n] = [a_1, \dots, a_n]$  and to find a pair of distinct elements with the largest product:

```
MAXPAIRWISEPRODUCTNAIVE( $A[1 \dots n]$ ):
    product  $\leftarrow$  0
    for  $i$  from 1 to  $n$ :
        for  $j$  from 1 to  $n$ :
            if  $i \neq j$ :
                if product  $< A[i] \cdot A[j]$ :
                    product  $\leftarrow A[i] \cdot A[j]$ 
    return product
```

This code can be optimized and made more compact as follows.

```
MAXPAIRWISEPRODUCTNAIVE( $A[1 \dots n]$ ):
    product  $\leftarrow$  0
    for  $i$  from 1 to  $n$ :
        for  $j$  from  $i + 1$  to  $n$ :
            product  $\leftarrow$  max(product,  $A[i] \cdot A[j]$ )
    return product
```

Implement this algorithm in your favorite programming language. If you are using C++, Java, or Python3, you may want to download the starter files (we provide starter solutions in these three languages for all the problems in the book). For other languages, you need to implement your solution from scratch.



Starter solutions for C++, Java, and Python3 are shown below.

## C++

```
#include <iostream>
#include <vector>
#include <algorithm>

int MaxPairwiseProduct(const std::vector<int>& numbers) {
    int max_product = 0;
    int n = numbers.size();

    for (int first = 0; first < n; ++first) {
        for (int second = first + 1; second < n; ++second) {
            max_product = std::max(max_product,
                                    numbers[first] * numbers[second]);
        }
    }

    return max_product;
}

int main() {
    int n;
    std::cin >> n;
    std::vector<int> numbers(n);
    for (int i = 0; i < n; ++i) {
        std::cin >> numbers[i];
    }

    std::cout << MaxPairwiseProduct(numbers); << "\n";
    return 0;
}
```

## Java

```
import java.util.*;
import java.io.*;

public class MaxPairwiseProduct {
    static int getMaxPairwiseProduct(int[] numbers) {
        int max_product = 0;
        int n = numbers.length;

        for (int first = 0; first < n; ++first) {
            for (int second = first + 1; second < n; ++second) {
                max_product = Math.max(max_product,
                    numbers[first] * numbers[second]);
            }
        }

        return max_product;
    }

    public static void main(String[] args) {
        FastScanner scanner = new FastScanner(System.in);
        int n = scanner.nextInt();
        int[] numbers = new int[n];
        for (int i = 0; i < n; i++) {
            numbers[i] = scanner.nextInt();
        }
        System.out.println(getMaxPairwiseProduct(numbers));
    }

    static class FastScanner {
        BufferedReader br;
        StringTokenizer st;

        FastScanner(InputStream stream) {
            try {
                br = new BufferedReader(new
```

```

        InputStreamReader(stream));
    } catch (Exception e) {
        e.printStackTrace();
    }
}

String next() {
    while (st == null || !st.hasMoreTokens()) {
        try {
            st = new StringTokenizer(br.readLine());
        } catch (IOException e) {
            e.printStackTrace();
        }
    }
    return st.nextToken();
}

int nextInt() {
    return Integer.parseInt(next());
}
}
}

```

## Python

```

def max_pairwise_product(numbers):
    n = len(numbers)
    max_product = 0
    for first in range(n):
        for second in range(first + 1, n):
            max_product = max(max_product,
                              numbers[first] * numbers[second])

    return max_product

```

```
if __name__ == '__main__':  
    input_n = int(input())  
    input_numbers = [int(x) for x in input().split()]  
    print(max_pairwise_product(input_numbers))
```

After submitting this solution to the grading system, many students are surprised when they see the following message:

Failed case #4/17: time limit exceeded

After you submit your program, we test it on dozens of carefully designed test cases to make sure the program is fast and error proof. As the result, we usually know what kind of errors you made. The message above tells that the submitted program exceeds the time limit on the 4th out of 17 test cases.

**Stop and Think.** Why the solution does not fit into the time limit?

MAXPAIRWISEPRODUCTNAIVE performs of the order of  $n^2$  steps on a sequence of length  $n$ . For the maximal possible value  $n = 2 \cdot 10^5$ , the number of steps is of the order  $4 \cdot 10^{10}$ . Since many modern computers perform roughly  $10^8$ – $10^9$  basic operations per second (this depends on a machine, of course), it may take tens of seconds to execute MAXPAIRWISEPRODUCTNAIVE, exceeding the time limit for the Maximum Pairwise Product Problem.

We need a faster algorithm!

### 3.2.2 Fast Algorithm

In search of a faster algorithm, you play with small examples like [5, 6, 2, 7, 4]. Eureka—it suffices to multiply the two largest elements of the array—7 and 6!

Since we need to find the largest and the second largest elements, we need only two scans of the sequence. During the first scan, we find the largest element. During the second scan, we find the largest element among the remaining ones by skipping the element found at the previous scan.

```

MAXPAIRWISEPRODUCTFAST( $A[1 \dots n]$ ):
 $index_1 \leftarrow 1$ 
for  $i$  from 2 to  $n$ :
    if  $A[i] > A[index_1]$ :
         $index_1 \leftarrow i$ 
 $index_2 \leftarrow 1$ 
for  $i$  from 2 to  $n$ :
    if  $A[i] \neq A[index_1]$  and  $A[i] > A[index_2]$ :
         $index_2 \leftarrow i$ 
return  $A[index_1] \cdot A[index_2]$ 

```

### 3.2.3 Testing and Debugging

Implement this algorithm and test it using an input  $A = [1, 2]$ . It will output 2, as expected. Then, check the input  $A = [2, 1]$ . Surprisingly, it outputs 4. By inspecting the code, you find out that after the first loop,  $index_1 = 1$ . The algorithm then initializes  $index_2$  to 1 and  $index_2$  is never updated by the second for loop. As a result,  $index_1 = index_2$  before the return statement. To ensure that this does not happen, you modify the pseudocode as follows:

```

MAXPAIRWISEPRODUCTFAST( $A[1 \dots n]$ ):
 $index_1 \leftarrow 1$ 
for  $i$  from 2 to  $n$ :
    if  $A[i] > A[index_1]$ :
         $index_1 \leftarrow i$ 
if  $index_1 = 1$ :
     $index_2 \leftarrow 2$ 
else:
     $index_2 \leftarrow 1$ 
for  $i$  from 1 to  $n$ :
    if  $A[i] \neq A[index_1]$  and  $A[i] > A[index_2]$ :
         $index_2 \leftarrow i$ 
return  $A[index_1] \cdot A[index_2]$ 

```

Check this code on a small datasets [7,4,5,6] to ensure that it produces correct results. Then try an input

```
2
100000 90000
```

You may find out that the program outputs something like 410 065 408 or even a negative number instead of the correct result 9 000 000 000. If it does, this is most probably caused by an *integer overflow*. For example, in C++ programming language a large number like 9 000 000 000 does not fit into the standard `int` type that on most modern machines occupies 4 bytes and ranges from  $-2^{31}$  to  $2^{31} - 1$ , where

$$2^{31} = 2\,147\,483\,648.$$

Hence, instead of using the C++ `int` type you need to use the `int64_t` type when computing the product and storing the result. This will prevent integer overflow as the `int64_t` type occupies 8 bytes and ranges from  $-2^{63}$  to  $2^{63} - 1$ , where

$$2^{63} = 9\,223\,372\,036\,854\,775\,808.$$

You then proceed to testing your program on large data sets, e.g., an array  $A[1 \dots 2 \cdot 10^5]$ , where  $A[i] = i$  for all  $1 \leq i \leq 2 \cdot 10^5$ . There are two ways of doing this.

1. Create this array in your program and pass it to `MAXPAIRWISEPRODUCTFAST` (instead of reading it from the standard input).
2. Create a separate program that writes such an array to a file `dataset.txt`. Then pass this dataset to your program from console as follows:

```
yourprogram < dataset.txt
```

Check that your program processes this dataset within time limit and returns the correct result: 39 999 800 000. You are now confident that the program finally works!

However, after submitting it to the testing system, it fails again...

```
Failed case #5/17: wrong answer
```

But how would you generate a test case that make your program fail and help you to figure out what went wrong?

### 3.2.4 Can You Tell Me What Error Have I Made?

You are probably wondering why we did not provide you with the 5th out of 17 test datasets that brought down your program. The reason is that nobody will provide you with the test cases in real life!

Since even experienced programmers often make subtle mistakes solving algorithmic problems, it is important to learn how to catch bugs as early as possible. When the authors of this book started to program, they naively thought that nearly all their programs are correct. By now, we know that our programs are *almost never* correct when we first run them.

When you are confident that your program works, you often test it on just a few test cases, and if the answers look reasonable, you consider your work done. However, this is a recipe for a disaster. To make your program *always* work, you should test it on a set of carefully designed test cases. Learning how to implement algorithms as well as test and debug your programs will be invaluable in your future work as a programmer.

### 3.2.5 Stress Testing

We will now introduce *stress testing*—a technique for generating thousands of tests with the goal of finding a test case for which your solution fails.

A stress test consists of four parts:

1. Your implementation of an algorithm.
2. An alternative, trivial and slow, but correct implementation of an algorithm for the same problem.
3. A random test generator.
4. An infinite loop in which a new test is generated and fed into both implementations to compare the results. If their results differ, the test and both answers are output, and the program stops, otherwise the loop repeats.

The idea behind stress testing is that two correct implementations should give the same answer for each test (provided the answer to the problem is unique). If, however, one of the implementations is incorrect, then there exists a test on which their answers differ. The only case when it is not so is when there is the same mistake in both implementations,

but that is unlikely (unless the mistake is somewhere in the input/output routines which are common to both solutions). Indeed, if one solution is correct and the other is wrong, then there exists a test case on which they differ. If both are wrong, but the bugs are different, then most likely there exists a test on which two solutions give different results.

Here is the stress test for `MAXPAIRWISEPRODUCTFAST` using `MAXPAIRWISEPRODUCTNAIVE` as a trivial implementation:

```
STRESSTEST( $N, M$ ):  
while true:  
     $n \leftarrow$  random integer between 2 and  $N$   
    allocate array  $A[1 \dots n]$   
    for  $i$  from 1 to  $n$ :  
         $A[i] \leftarrow$  random integer between 0 and  $M$   
    print( $A[1 \dots n]$ )  
     $result_1 \leftarrow$  MAXPAIRWISEPRODUCTNAIVE( $A$ )  
     $result_2 \leftarrow$  MAXPAIRWISEPRODUCTFAST( $A$ )  
    if  $result_1 = result_2$ :  
        print("OK")  
    else:  
        print("Wrong answer:",  $result_1$ ,  $result_2$ )  
    return
```

The while loop above starts with generating the length of the input sequence  $n$ , a random number between 2 and  $N$ . It is at least 2, because the problem statement specifies that  $n \geq 2$ . The parameter  $N$  should be small enough to allow us to explore many tests despite the fact that one of our solutions is slow.

After generating  $n$ , we generate an array  $A$  with  $n$  random numbers from 0 to  $M$  and output it so that in the process of the infinite loop we always know what is the current test; this will make it easier to catch an error in the test generation code. We then call two algorithms on  $A$  and compare the results. If the results are different, we print them and halt. Otherwise, we continue the while loop.

Let's run `STRESSTEST(10, 100 000)` and keep our fingers crossed in a hope that it outputs "Wrong answer." We see something like this (the result can be different on your computer because of a different random number generator).



```

...
OK
67232 68874 69499
OK
6132 56210 45236 95361 68380 16906 80495 95298
OK
62180 1856 89047 14251 8362 34171 93584 87362 83341 8784
OK
21468 16859 82178 70496 82939 44491
OK
68165 87637 74297 2904 32873 86010 87637 66131 82858 82935
Wrong answer: 7680243769 7537658370

```

Hurrah! We've found a test case where `MAXPAIRWISEPRODUCTNAIVE` and `MAXPAIRWISEPRODUCTFAST` produce different results, so now we can check what went wrong. Then we can debug this solution on this test case, find a bug, fix it, and repeat the stress test again.

**Stop and Think.** Do you see anything suspicious in the found dataset?

Note that generating tests automatically and running stress test is easy, but debugging is hard. Before diving into debugging, let's try to generate a smaller test case to simplify it. To do that, we change  $N$  from 10 to 5 and  $M$  from 100 000 to 9.

**Stop and Think.** Why did we first run `STRESSTEST` with large parameters  $N$  and  $M$  and now intend to run it with small  $N$  and  $M$ ?

We then run the stress test again and it produces the following.

```

...
7 3 6
OK
2 9 3 1 9
Wrong answer: 81 27

```

The slow `MAXPAIRWISEPRODUCTNAIVE` gives the correct answer 81 ( $9 \cdot 9 = 81$ ), but the fast `MAXPAIRWISEPRODUCTFAST` gives an incorrect answer 27.

**Stop and Think.** How `MAXPAIRWISEPRODUCTFAST` can possibly return 27?

To debug our fast solution, let's check which two numbers it identifies

as two largest ones. For this, we add the following line before the return statement of the `MAXPAIRWISEPRODUCTFAST` function:

```
print(index1, index2)
```

After running the stress test again, we see the following.

```
...
7 3 6
1 3
OK
5
2 9 3 1 9
2 3
Wrong answer: 81 27
```

Note that our solutions worked and then failed on exactly the same test cases as on the previous run of the stress test, because we didn't change anything in the test generator. The numbers it uses to generate tests are pseudorandom rather than random—it means that the sequence looks random, but it is the same each time we run this program. It is a convenient and important property, and you should try to have your programs exhibit such behavior, because deterministic programs (that always give the same result for the same input) are easier to debug than non-deterministic ones.

Now let's examine  $\text{index}_1 = 2$  and  $\text{index}_2 = 3$ . If we look at the code for determining the second maximum, we will notice a subtle bug. When we implemented a condition on  $i$  (such that it is not the same as the previous maximum) instead of comparing  $i$  and  $\text{index}_1$ , we compared  $A[i]$  with  $A[\text{index}_1]$ . This ensures that the second maximum differs from the first maximum by the value rather than by the index of the element that we select for solving the Maximum Pairwise Product Problem. So, our solution fails on any test case where the largest number is equal to the second largest number. We now change the condition from

```
 $A[i] \neq A[\text{index}_1]$ 
```

to

```
 $i \neq \text{index}_1$ 
```

After running the stress test again, we see a barrage of “OK” messages

on the screen. We wait for a minute until we get bored and then decide that `MAXPAIRWISEPRODUCTFAST` is finally correct!

However, you shouldn't stop here, since you have only generated very small tests with  $N = 5$  and  $M = 10$ . We should check whether our program works for larger  $n$  and larger elements of the array. So, we change  $N$  to 1 000 (for larger  $N$ , the naive solution will be pretty slow, because its running time is quadratic). We also change  $M$  to 200 000 and run. We again see the screen filling with words "OK", wait for a minute, and then decide that (finally!) `MAXPAIRWISEPRODUCTFAST` is correct. Afterward, we submit the resulting solution to the grading system and pass the Maximum Pairwise Product Problem test!

As you see, even for such a simple problem like Maximum Pairwise Product, it is easy to make subtle mistakes when designing and implementing an algorithm. The pseudocode below presents a more "reliable" way of implementing the algorithm.

```
MAXPAIRWISEPRODUCTFAST( $A[1 \dots n]$ ):  
   $index \leftarrow 1$   
  for  $i$  from 2 to  $n$ :  
    if  $A[i] > A[index]$ :  
       $index \leftarrow i$   
  swap  $A[index]$  and  $A[n]$   
   $index \leftarrow 1$   
  for  $i$  from 2 to  $n - 1$ :  
    if  $A[i] > A[index]$ :  
       $index \leftarrow i$   
  swap  $A[index]$  and  $A[n - 1]$   
  return  $A[n - 1] \cdot A[n]$ 
```

In this book, besides learning how to design and analyze algorithms, you will learn how to implement algorithms in a way that minimizes the chances of making a mistake, and how to test your implementations.

### 3.2.6 Even Faster Algorithm

The `MAXPAIRWISEPRODUCTFAST` algorithm finds the largest and the second largest elements in about  $2n$  comparisons.

**Exercise Break.** Find two largest elements in an array in  $1.5n$  comparisons.

After solving this problem, try the next, even more challenging Exercise Break.

**Exercise Break.** Find two largest elements in an array in  $n + \lceil \log_2 n \rceil - 2$  comparisons.

And if you feel that the previous Exercise Break was easy, here are the next two challenges that you may face at your next interview!

**Exercise Break.** Prove that no algorithm for finding two largest elements in an array can do this in less than  $n + \lceil \log_2 n \rceil - 2$  comparisons.

**Exercise Break.** What is the fastest algorithm for finding three largest elements?

### 3.2.7 A More Compact Algorithm

The Maximum Pairwise Product Problem can be solved by the following compact algorithm that uses sorting (in non-decreasing order).

```
MAXPAIRWISEPRODUCTBYSORTING( $A[1 \dots n]$ ):  
  SORT( $A$ )  
  return  $A[n - 1] \cdot A[n]$ 
```

This algorithm does more than we actually need: instead of finding two largest elements, it sorts the entire array. For this reason, its running time is  $O(n \log n)$ , but not  $O(n)$ . Still, for the given constraints ( $2 \leq n \leq 2 \cdot 10^5$ ) this is usually sufficiently fast to fit into a second and pass our grader.

## 3.3 Solving a Programming Challenge in Five Easy Steps

Below we summarize what we've learned in this chapter.

### 3.3.1 Reading Problem Statement

Start by reading the problem statement that contains the description of a computational task, time and memory limits, and a few sample tests. Make sure you understand how an output matches an input in each sample case.

If time and memory limits are not specified explicitly in the problem statement, the following default values are used.

**Time limits (sec.):**

C++	Java	Python	C	C#	Go	Haskell	JavaScript	Kotlin	Ruby	Rust	Scala
1	1.5	5	1	1.5	1.5	2	5	1.5	5	1	3

**Memory limit:** 512 Mb.

### 3.3.2 Designing an Algorithm

After designing an algorithm, prove that it is correct and try to estimate its expected running time on the most complex inputs specified in the constraints section. If your laptop performs roughly  $10^8$ – $10^9$  operations per second, and the maximum size of a dataset in the problem description is  $n = 10^5$ , then an algorithm with quadratic running time is unlikely to fit into the time limit (since  $n^2 = 10^{10}$ ), while a solution with running time  $O(n \log n)$  will. However, an  $O(n^2)$  solution will fit if  $n = 1\,000$ , and if  $n = 100$ , even an  $O(n^3)$  solution will fit. Although polynomial algorithms remain unknown for some hard problems in this book, a solution with  $O(2^n n^2)$  running time will probably fit into the time limit as long as  $n$  is smaller than 20.

### 3.3.3 Implementing an Algorithm

Start implementing your algorithm in one of the following programming languages supported by our automated grading system: C, C++, C#, Haskell, Java, JavaScript, Kotlin, Python2, Python3, Ruby, Rust, or Scala. For all problems, we provide starter solutions for C++, Java, and Python3. For other programming languages, you need to implement a solution from scratch. The grading system detects the programming language of your submission automatically, based on the extension of the submission file.

We have reference solutions in C++, Java, and Python3 (that we don't share with you) which solve the problem correctly under the given constraints, and spend at most 1/3 of the time limit and at most 1/2 of the memory limit. You can also use other languages, and we've estimated the time limit multipliers for them. However, we have no guarantee that a correct solution for a particular problem running under the given time and memory constraints exists in any of those other languages.

In the Appendix, we list compiler versions and flags used by the grading system. We recommend using the same compiler flags when you test your solution locally. This will increase the chances that your program behaves in the same way on your machine and on the testing machine (note that a buggy program may behave differently when compiled by different compilers, or even by the same compiler with different flags).

### 3.3.4 Testing and Debugging

Submitting your implementation to the grading system without testing it first is a bad idea! Start with small datasets and make sure that your program produces correct results on all sample datasets. Then proceed to checking how long it takes to process a large dataset. To estimate the running time, it makes sense to implement your algorithm as a function like `solve(dataset)` and then implement an additional procedure `generate()` that produces a large dataset. For example, if an input to a problem is a sequence of integers of length  $1 \leq n \leq 10^5$ , then generate a sequence of length  $10^5$ , pass it to your `solve()` function, and ensure that the program outputs the result quickly.

Check the boundary values to ensure that your program processes correctly both short sequences (e.g., with 2 elements) and long sequences (e.g., with  $10^5$  elements). If a sequence of integers from 0 to, let's say,  $10^6$  is given as an input, check how your program behaves when it is given a sequence  $0, 0, \dots, 0$  or a sequence  $10^6, 10^6, \dots, 10^6$ . Afterward, check it also on randomly generated data. Check degenerate cases like an empty set, three points on a single line, a tree which consists of a single path of nodes, etc.

After it appears that your program works on all these tests, proceed to stress testing. Implement a slow, but simple and correct algorithm and check that two programs produce the same result (note however that this is not applicable to problems where the output is not unique). Generate

random test cases as well as biased tests cases such as those with only small numbers or a small range of large numbers, strings containing a single letter “a” or only two different letters (as opposed to strings composed of all possible Latin letters), and so on. Think about other possible tests which could be peculiar in some sense. For example, if you are generating graphs, try generating trees, disconnected graphs, complete graphs, bipartite graphs, etc. If you generate trees, try generating paths, binary trees, stars, etc. If you are generating integers, try generating both prime and composite numbers.

### 3.3.5 Submitting to the Grading System

When you are done with testing, submit your program to the grading system! Go to the submission page, create a new submission, and upload a file with your program (make sure to upload a source file rather than an executable). The grading system then compiles your program and runs it on a set of carefully constructed tests to check that it outputs a correct result for all tests and that it fits into the time and memory limits. The grading usually takes less than a minute, but in rare cases, when the servers are overloaded, it might take longer. Please be patient. You can safely leave the page when your solution is uploaded.

As a result, you get a feedback message from the grading system. You want to see the “Good job!” message indicating that your program passed all the tests. The messages “Wrong answer”, “Time limit exceeded”, “Memory limit exceeded” notify you that your program failed due to one of these reasons. If you program fails on one of the first two test cases, the grader will report this to you and will show you the test case and the output of your program. This is done to help you to get the input/output format right. In all other cases, the grader will *not* show you the test case where your program fails.

