TDDD95 Algorithmic Problem Solving 6hp, vt2022

Fredrik Heintz Dept of Computer and Information Science Linköping University

Outline



- What is algorithmic problem solving?
- Why is algorithmic problem solving important?
- What will be studied in this course?
- A method for algorithmic problem solving
- Common algorithmic problem solving approaches
- Common data structures and algorithms
- Pragmatic algorithmic problem solving using Kattis

What is Algorithmic Problem Solving?

- 3
- Algorithmic problem solving is about developing correct and working algorithms to solve classes of problems.
- The problems are normally very well defined and you know there is a solution, but they can still be very hard.





Those that really understand and take advantage of software technology owns the future!

TΜ

facebook.



University

. .

÷.

Warsaw

ICPC

100

0

ACM-ICPI

LINKÖPINGS UNIVERSITET



acm

IBM.

event sponsor international collegiate programming contest

Background

Problems in Computer Science are often classified as belonging to a certain class of problems (e.g., NP, Unsolvable, Recursive). In this problem you will be analyzing a property of an algorithm whose classification is not known for all possible inputs.

The Problem

Consider the following algorithm:

1.	inpu	t s		
2.	prin	t a		
3.	15 /	= 1 then SIOP		
4.		11	n is odd then	$n \longleftarrow 3n+1$
۶.		•13	$n \leftarrow n/2$	
6.	GOID	2		

Given the input 22, the following sequence of numbers will be printed 22 11 34 17 52 26 13 40 20 10 5 16 8 4 2 1

It is conjectured that the algorithm above will terminate (when a 1 is printed) for any integral input value. Despite the simplicity of the algorithm, it is unknown whether this conjecture is true. It has been verified, however, for all integers n such that 0 < n < 1,000,000 (and, in fact, for many more numbers than this.)

Given an input *n*, it is possible to determine the number of numbers printed (including the 1). For a given *n* this is called the *cycle-length* of *n*. In the example above, the cycle length of 22 is 16.

For any two numbers i and j you are to determine the maximum cycle length over all numbers between i and j.

The Input

The input will consist of a series of pairs of integers i and j, one pair of integers per line. All integers will be less than 1,000,000 and greater than 0.

You should process all pairs of integers and for each pair determine the maximum cycle length over all integers between and including i and j.

You can assume that no operation overflows a 32-bit integer.

The Output

For each pair of input integers *i* and *j* you should output *i*, *j*, and the maximum cycle length for integers between and including *i* and *j*. These three numbers should be separated by at least one space with all three numbers on one line and with one line of output for each line of input. The integers *i* and *j* must appear in the output in the same order in which they appeared in the input and should be followed by the maximum cycle length (on the same line).

Sample Input

1 10 100 200 201 210 900 1000

Sample Output

1 10 20 100 200 125 201 210 59 900 1000 174

Example: The 3n+1 problem

100 The 3n + 1 problem

Background

Problems in Computer Science are often classified as belonging to a certain class of problems (e.g., NP, Unsolvable, Recursive). In this problem you will be analyzing a property of an algorithm whose classification is not known for all possible inputs.

The Problem

Consider the following algorithm:

- 1. input n
- print n
- 3. if n = 1 then STOP
- 4. if n is odd then $n \leftarrow -3n + 1$
- 5. else $n \leftarrow n/2$
- 6. GOTO 2

Given the input 22, the following sequence of numbers will be printed

 $22\ 11\ 34\ 17\ 52\ 26\ 13\ 40\ 20\ 10\ 5\ 16\ 8\ 4\ 2\ 1$

It is conjectured that the algorithm above will terminate (when a 1 is printed) for any integral input value. Despite the simplicity of the algorithm, it is unknown whether this conjecture is true. It has been verified, however, for all integers n such that 0 < n < 1,000,000 (and, in fact, for many more numbers than this.)

Given an input n, it is possible to determine the number of numbers printed before and including the 1 is printed. For a given n this is called the *cycle-length* of n. In the example above, the cycle length of 22 is 16.

For any two numbers i and j you are to determine the maximum cycle length over all numbers between and including both i and j.

Example: The 3n+1 problem



The Input

The input will consist of a series of pairs of integers i and j, one pair of integers per line. All integers will be less than 10,000 and greater than 0.

You should process all pairs of integers and for each pair determine the maximum cycle length over all integers between and including i and j.

The Output

For each pair of input integers i and j you should output i, j, and the maximum cycle length for integers between and including i and j. These three numbers should be separated by at least one space with all three numbers on one line and with one line of output for each line of input. The integers i and j must appear in the output in the same order in which they appeared in the input and should be followed by the maximum cycle length (on the same line).

Sample Input

Sample Output

1 10 20 100 200 125 201 210 89 900 1000 174

Example: The 3n+1 problem

- Follow the instructions in the problem!
- Memoization to speed it up.
- Table lookup to solve it in constant time.
- Gotchas:
 - *j* can be smaller than *i*.
 - *j* can equal *i*.
 - The order of *i* and *j* in output must be the same as the input, even when *j* is smaller than *i*.

Course Goals



The goals of the course are you should be able to:

- analyze the efficiency of different approaches to solving a problem to determine which approaches will be reasonably efficient in a given situation,
- compare different problems in terms of their difficulty,
- use algorithm design techniques such as greedy algorithms, dynamic programming, divide and conquer, and combinatorial search to construct algorithms to solve given problems,
- strategies for testing and debugging algorithms and data structures,
- quickly and correctly implement a given specification of an algorithm or data structure,
- communicate and cooperate with other students during problem solving in groups.

Examination



- LAB1 4hp
 - individually solving the 4 lab assignments and
 - actively participating in at least 3 problem solving sessions.
- UPPG1 2hp,
 - individually solving the 13 weekly homework exercises, e.g.:
 - Data structures
 - Greedy Problems and Dynamic Programming
 - Graph Algorithms
 - Search
 - Math-related Problems
 - Computational Geometry.





UPG1

To pass the UPG1 part of the course, you must fulfill the following requirements:

- 1. Solve at least one problem from each of the 13 exercise sets, before or after the corresponding deadline.
- 2. Accumulate a sufficient number of UPG1 points. This also determines your UPG1 grade.

The problems come in three difficulty classes: A, B, and C. You are awarded 1 point in the corresponding difficulty class for each problem you solve before its deadline, and 0.5 points if you solve it after the deadline. The table below specifies how many points you must accumulate for the different UPG1 grades.

Α	В	С	Grade UPG1
13	0	0	3
13	7	0	4
13	7	6	5

You may of course replace easier problems with harder problems. For example, if you have 6 A points and 7 B points, you have 13 points at difficulty A or above, but not an *additional* 7 points at difficulty B or above. You would therefore get grade 3.

The Schedule – VT 1

13

- 19/1 How to solve algorithmic problems, intro seminar
- 20/1 Practice problem solving session: 13.30-17.00 with discussion
- 28/1 Deadline Ex 1 (Greedy and DP 1) Seminar Ex1 and Data structures
- 4/2 Deadline Ex2 (Data structures) Seminar Ex2 and Arithmetic
- 11/2 Deadline Ex3 (Arithmetic) Seminar Ex3 and Problem solving approaches
- 17/2 Deadline Lab Assignment 1 (Data structures, Greedy/Dynamic, Arithmetic)
- 17/2 Problem solving session (individual based on Lab 1)
- 18/2 Deadline Ex4 (Greedy and DP II) Seminar Ex 4 and Graphs I
- 25/2 Deadline Ex5 (Graphs I) Seminar Ex5 and Graphs II
- 4/3 Deadline Ex6 (Graphs II) Seminar Ex6 and Graphs III
- 10/3 Deadline Lab Assignment 2 (Graphs)
- 10/3 Problem solving session (individual based on Lab 2)

The Schedule – VT 2



- 28/3 Problem solving session (groups based on Lab 1-2)
- 1/4 Deadline Ex 7 (Graphs III) Seminar Ex7 and Strings I
- 8/4 Deadline Ex8 (Strings I) Seminar Ex8 and Strings II
- 13/4 Deadline Ex9 (Strings II) Seminar Ex9 and Number Theory
- 25/4 Deadline Lab Assignment 3 (Strings, String Matching and Number Theory)
- 25/4 Problem solving session (individual based on Lab 3)
- 22/4 Deadline Ex10 (Number Theory) Seminar Ex 11 and Combinatorial Search
- 29/4 Deadline Ex11 (Combinatorial Search) Seminar Ex12 and Computational Geometry
- 6/5 Deadline Ex12 (Computational Geometry)
- 9/5 Deadline Lab Assignment 4 (Computational Geometry)
- 9/5 Problem solving session (individual based on Lab 4)
- 13/5 Deadline Ex13 (Mixed)
- 16/5 Problem solving session (groups based on Lab 3-4)

Steps in solving algorithmic problems



- Estimate the difficulty
 - Theory (size of inputs, known algorithms, known theorems, ...)
 - Coding (size of program, many cases, complicated data structures, ...)
 - Have you seen this problem before? Have you solved it before? Do you have useful code in your code library?

• Understand the problem!

- What is being asked for? What is given? How large can instances be?
- Can you draw a diagram to help you understand the problem?
- Can you explain the problem in your own words?
- Can you come up with good examples to test your understanding?

Steps in solving algorithmic problems

- 16
- Determine the right algorithm or algorithmic approach
 - Can you solve the problem using brute force?
 - Can you solve the problem using a greedy approach?
 - Can you solve the problem using dynamic programming?
 - Can you solve the problem using search?
 - Can you solve the problem using a known algorithm in your code library?
 - Can you modify an existing algorithm? Can you modify the problem to suite an existing algorithm?
 - Do you have to come up with your own algorithm?
- Solve the problem! ③

Time Limits and Computational Complexity



- The normal time limit for a program is a few seconds.
- You may assume that your program can do about 100M operations within this time limit.

n	Worst AC Complexity	Comments	
≤ [1011]	$O(n!), O(n^6)$	Enumerating permutations	
≤ [1518]	$O(2^n \times n^2)$	DP TSP	
≤ [1822]	$O(2^n \times n)$	DP with bitmask technique	
≤ 100	<i>O</i> (<i>n</i> ⁴)	DP with 3 dimensions and O(n) loop	
≤ 450	$O(n^3)$	Floyd Warshall's (APSP)	
≤ 2K	$O(n^2 \log_2 n)$	2-nested loops + tree search	
≤ 10K	$O(n^2)$	Bubble/Selection/Insertion sort	
≤ 1M	$O(n \log_2 n)$	Merge Sort, Binary search	
≤ 100M	$O(n), O(\log_2), O(1)$	Simulation, find average	

Important Problem Solving Approaches

- Simulation/Ad hoc
 - Do what is stated in the problem
 - Example: Simulate a robot
- Greedy approaches
 - Find the optimal solution by extending a partial solution by making locally optimal decisions
 - Example: Minimal spanning trees, coin change in certain currencies
- Divide and conquer
 - Take a large problem and split it up in smaller parts that are solved individually
 - Example: Merge sort and Quick sort
- Dynamic programming
 - Find a recursive solution and compute it "backwards" or use memoization
 - Example: Finding the shortest path in a graph and coin change in all currencies
- Search
 - Create a search space and use a search algorithm to find a solution
 - Example: Exhaustive search (breadth or depth first search), binary search, heuristic search (A*, best first, branch and bound)

Complete Search a.k.a. Brute Force

- When a problem is small or (almost) all possibilities have to be tried *complete search* is a candidate approach.
- To determine the feasibility of complete search estimate the number of calculations that have to be made in the worst case.
- Iterative complete search uses nested loops to generate every possible complete solution and *filter* out the valid ones.
 - Iterating over all permutations using next_permutation
 - Iterating over all subsets using bit set technique
- Recursive complete search extends a partial solution with one element until a complete and valid solution is found.
 - This approach is often called *recursive backtracking*.
 - Pruning is used to significantly improve the efficiency by removing partial solutions that can not lead to a solution as soon as possible. In the best case only valid solutions are generated.

We have three different integers, *x*, *y* and *z*, which satisfy the following three relations:

- x + y + z = A
- xyz = B
- $x^2 + y^2 + z^2 = C$

You are asked to write a program that solves for *x*, *y* and *z* for given values of *A*, *B* and *C*.

Divide and Conquer

- 21
- Divide and conquer is very common and powerful technique which divides a problem into smaller parts, solves each part recursively and then puts together the answer from the pieces.
- Many well known algorithms are based on divide and conquer such as quick sort, merge sort and binary search.
- Binary search is a very versatile and useful technique which can be used to
 - find a particular value in a sorted range,
 - find the parameters of a (convex) function that gives a particular value,
 - find the minimum/maximum value of a function.
- Binary search can be implemented either using built in functions (lower_bound/upper_bound), iterating until the difference between the end points is small enough or iterate a constant but sufficiently large number of times.

Quick-Sort



- Quick-sort is a randomized sorting algorithm based on the divide-and-conquer paradigm:
 - Divide: pick a random element *x* (called pivot) and partition *S* into
 - *L* elements less than *x*
 - *E* elements equal *x*
 - G elements greater than x
 - Recur: sort L and G
 - Conquer: join *L*, *E* and *G*



Pivot selection



23

Partition, recursive call, pivot selection



Partition, recursive call, base case



Recursive call, ..., base case, join



Recursive call, pivot selection



Partition, ..., recursive call, base case



28

Join, join



Greedy



- An algorithm is said to be greedy if it makes a locally optimal choice in each step towards the globally optimal solution.
- For a greedy algorithm to give a globally optimal result a problem must have two properties:
 - It has optimal sub-structures, i.e. an optimal solution contains the optimal solutions to sub problems.
 - It has the greedy choice property, i.e. if we extend a partial solution by making a locally optimal choice we will get the optimal complete solution without reconsidering previous choices.
- Classical examples: Coin change in some currencies, interval coverage and load balancing.
- Greedy algorithms can be very useful as heuristics for example in branch-and-bound search algorithms.
- In combinatorics matroids and the generalization greedoids characterize classes of problems with greedy solutions.



Consider an undirected, weight graph







Sort the edges by increasing edge weight

edge	d_v	
(D,E)	1	
(D,G)	2	
(E,G)	3	
(C,D)	3	
(G,H)	3	
(C,F)	3	
(B,C)	4	

edge	d_{v}	
(B,E)	4	
(B,F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	
(E,G)	3	
(C,D)	3	
(G,H)	3	
(C,F)	3	
(B,C)	4	

edge	d_{v}	
(B,E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	
(C,D)	3	
(G,H)	3	
(C,F)	3	
(B,C)	4	

edge	d_{v}	
(B,E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	



Select first |V|-1 edges which do not generate a cycle



edge	d_v		edge	d_v	
(D,E)	1	\checkmark	(B,E)	4	
(D,G)	2	\checkmark	(B,F)	4	
(E,G)	3	X	(B,H)	4	
(C,D)	3		(A,H)	5	
(G,H)	3		(D,F)	6	
(C,F)	3		(A,B)	8	
(B,C)	4		(A,F)	10	

Accepting edge (E,G) would create a cycle





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	X
(C,D)	3	\checkmark
(G,H)	3	
(C,F)	3	
(B,C)	4	

edge	d_{v}	
(B , E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	X
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	
(B,C)	4	

edge	d_{v}	
(B,E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	X
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	

edge	d_{v}	
(B,E)	4	
(B,F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	X
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B,E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	χ
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B , E)	4	X
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	χ
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B ,E)	4	χ
(B,F)	4	χ
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	X
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B , E)	4	χ
(B,F)	4	χ
(B,H)	4	χ
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	χ
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_v	
(B , E)	4	χ
(B , F)	4	χ
(B,H)	4	χ
(A,H)	5	\checkmark
(D,F)	6	
(A,B)	8	
(A,F)	10	



Select first |V|-1 edges which do not generate a cycle





Done

Total Cost = $\sum d_v = 21$

Dynamic Programming



- Dynamic Programming is a problem solving approach which computes the answer for every possible *state* exactly once.
- For DP to be suitable a problem must have two properties:
 - It has optimal sub-structures, i.e. an optimal solution contains the optimal solutions to sub problems.
 - Overlapping sub-problems, i.e. the same subproblem occurs many times.
- Top-down (memoization) vs Bottom-up
 - Top-down: no need to consider the order of computations, only compute states actually used, natural transition from complete search,
 - Bottom-up: no recursion, computes every state, table size can be reduced if only the previous row of states is used then only two rows are required.
- Displaying the optimal solution
 - Store the previous state for each solution
 - Use the DP table and the optimal sub-structures property to compute the path.

Classical DP Problems

- Max 1D sum
- Max 2D sum
- Longest increasing subsequence (LIS)
 - Longest decreasing subsequence (LDS)
- o-1 Knapsack (subset sum)
- Coin Change (general version)
- Travelling Salesman Problem (TSP)

	1D RSQ	2D RSQ	LIS	Knapsack	CoinChange	TSP
State	(i)	(i,j)	(i)	(id,remW)	(v)	(pos,mask)
Space	O(n)	$O(n^2)$	O(n)	O(nS)	O(V)	$O(n2^n)$
Transition	subarray	submatrix	all j <i< th=""><th>take/ignore</th><th>all <i>n</i> coins</th><th>all <i>n</i> cities</th></i<>	take/ignore	all <i>n</i> coins	all <i>n</i> cities
Time	<i>O(1)</i>	<i>O</i> (1)	$O(n^2)$	O(nS)	O(nV)	$O(2^n n^2)$

Common Data Structures





Important Data Structures and Algorithms

Data structures

- Standard library data structures
 - Vector, stack, queue, heap, priority queue, sets, maps
- Other data structures
 - Graph (adjacency list and adjacency matrix), Union/find, Segment tree, Fenwick tree, Trie
- Sorting
 - Quick sort, Merge sort, Radix sort, Bucket sort
- Strings
 - String matching (Knuth Morris Pratt, Aho-Corasick), pattern matching, trie, suffix trees, suffix arrays, recursive decent parsing



Dynamic programming

- Longest common subsequence, Longest increasing subsequence, o/1 Knapsack, Coin Change, Matrix Chain Multiplication, Subset sum, Partitioning
- Graphs
 - Traversal (pre-, in- and post-order), finding cycles, finding connected components, finding articulation points, topological sort, flood fill, Euler cycle/Euler path, SSSP Single source shortest path (Dijkstra, Bellman-Ford), APSP All pairs shortest path (Floyd Warshall), transitive closure (Floyd Warshall), MST Minimum spanning tree (Prim, Kruskal (using Union/find)), Maximal Bipartite Matching, Maximum flow, Maximum flow minimal cost, Minimal cut

Search

 Exhaustive search (depth-first, breadth-first search, backtracking), binary search (divide and conquer), greedy search (hill climbing), heuristic search (A*, branch and bound), search trees



Mathematics

- Number theory (prime numbers, greatest common divisor (GCD), modulus), big integers, combinatorics (count permutations), number series (Fibonacci numbers, Catalan numbers, binomial coefficients), probabilities, linear algebra (matrix inversion, linear equations systems), finding roots to polynomial equations, diofantic equations, optimization (simplex)
- Computational geometry
 - Representations of points, lines, line segments, polygons, finding intersections, point localization, triangulation, Voronoi diagrams, area and volume calculations, convex hull (Graham scan), sweep line algorithms

Kattis (https://liu.kattis.com)



How Kattis checks a program





UVA Online Judge



Programming languages



- Allowed languages are C, C++, Java, and Python.
- C++ or Java is strongly recommended, use the language that you are most familiar with and want to learn more about.
- Get to know their standard libraries.
- Get to know input and output. Remember that I/O in Java is very slow, use Kattio. Remember that cout/cerr also is relatively slow, learn how to use scanf/printf if you use C++.
- Learn to use an appropriate IDE such as eclipse, emacs, or vim
- Create a problem template to speed up problem solving and to create a common format for your problems.

Pragmatic Algorithmic Problem Solving

59

```
/* -*- Mode: C++ -*- */
/**
 * XXX NAME C++ "Approach"
 *
    Started:
 *
    Finished:
 * Total time:
 * Submission 1:
 ж.
 * Comments:
 * Lessons learned:
 *7
#include <algorithm>
#include <cassert>
#include <cmath>
#include <cstdio>
#include <cstdlib>
#include <cstring>
#include <functional>
#include <iomanip>
#include <iostream>
#include <sstream>
#include <map>
#include <set>
#include <queue>
#include <stack>
#include <string>
#include <utility>
#include <vector>
using namespace std;
typedef vector<int> vi;
int
main(int argc, char* argv[])
return 0;
```

Testing and debugging



- Always create an example input (.in) and example output (.out) file with verbatim copies of the example input and output from the problem statement!
- For most problems it is enough to diff your output with the example output: ./prog < prog.in | diff - prog.out
- Create additional tests, such as:
 - Extreme inputs, i.e. smallest and largest values (0, 1, "", empty line, 2³¹⁻¹)
 - Small inputs that you can compute by hand
 - Potentially tricky cases such as when all inputs are equal, in the case of floating points numbers when you have to round both up and down
 - Very large cases, randomly generated to test that your program computes an answer fast enough (even though you might not know the correct answer).
- Use a correct but slow algorithm to compute answers.
- Print intermediate information, such as values of relevant variables. cout << "a=" << a << "; b=" << b << endl; Remember to remove the debug output before submitting! (or use cerr)

Summary



- What is algorithmic problem solving?
- Why is algorithmic problem solving important?
- What will be studied in this course?
- A method for algorithmic problem solving
- Common algorithmic problem solving approaches
- Common data structures and algorithms
- Pragmatic algorithmic problem solving using Kattis