

The rules of the exam are as follows:

- You are free to consult class notes, the textbook, homework solutions, and any other resources posted on the course website or in the “Resources” section of the course Piazza page. **You may not consult with other students or the internet.** You may also ask Prof. Goldner questions during office hours or virtually.
- Post clarifying questions as a **private note** on Piazza, or email them to Prof. Goldner. She will respond within 24 hours. Direct all questions to her and not to TAs.
- Solutions must be typeset in LaTeX.
- When asked for a runtime in big-Oh notation, give the tightest and simplest possible bound you can. If your algorithm is $O(n)$ and also $O(n^2)$, just say it is $O(n)$. If it is $O(2n)$, it is also $O(n)$, so omit the factor of 2. And so on.
- Submit the exam via gradescope by **3:00pm on Wednesday, May 10**. Late exams will not be accepted.
- Do not discuss the exam with your classmates until after the due date, even if you have finished.
- Whenever an algorithm is given, it must be accompanied with justification and a runtime.

Topics to be Covered

- Asymptotic runtime ($O/\Omega/\Theta$)
- Reading and writing pseudocode
- Sorting
- Induction
- Loop Invariants
- Data structures: stacks, queues, graphs
- Graph search
 - Algorithms: DFS, BFS (unweighted shortest path)
- Greedy algorithms
 - Algorithms: Dijkstra (non-negative weighted shortest path), caching, interval scheduling, scheduling to minimize lateness, Huffman codes

- Techniques: Greedy stays ahead proof, Greedy exchange proof
- Divide & Conquer
 - Algorithms: Mergesort, Closest pair of points, Integer Multiplication, Matrix Multiplication
 - Techniques: Prove the recurrence via induction, solve the runtime recurrence (recurrence trees, master theorem)
- Dynamic Programming
 - Algorithms: Bellman-Ford (all weighted shortest path), weighted interval scheduling, segmented least squares, knapsack
 - Techniques: Find the DP recurrence, proof via 7 part solution
- Linear Programming
 - Techniques: formulating problems as LPs, taking duals, using properties from duality

Covered in much less depth:

- NP-Completeness
 - Concepts: What is P, NP, NP-Hard, NP-Complete?
 - Techniques: How do we relate hard problems to each other?
- Remaining lectures:
 - Zero Sum Games
 - Multiplicative Weight Update
 - Randomized Algorithms
 - Online Algorithms

Types of Questions Asked

- a. Standard homework questions (give an algorithm, prove its correctness, analyze its runtime).
- b. True/False about properties of algorithms, theory, or proofs.
 - (a) And justify why.
 - (b) And correct the statement if wrong.
- c. Short answer, e.g., of counter examples, algorithms, proofs, and runtimes.

Final Exam Review

For **all** algorithms, *always* give:

- (1) a clear enough description that someone could code it up without knowing any specific language (even if it just an english description, it must be that clear to understand!),
- (2) a justification of why it gives the guarantees it does, and
- (3) an analysis of its running time.

Reiterating, “design and analyze” means given a word problem, introduce necessary notation, design an algorithm to solve the problem, analyze runtime, and analyze accuracy (along with any other problem specific requirements of the algorithm or solution).

Intuition and Reminders

Greedy:

- Appropriate when the next best choice (“myopic”) leads you to optimality.
- “Best” has to be by some metric—in interval scheduling, we scheduled by earliest finish time, *not* by shortest job time, so picking which metric to use as “best” is important.
- You should be thinking: What if we just take the next available thing that meets X criteria?
- Pitfalls to look out for: when you can’t just look at each piece individually/successively, when you need to be able to “look ahead” somehow.

Divide and Conquer:

- Appropriate when you have subproblems that can be solved independently.
- Usually for a D&C problem, brute force (e.g., check all pairs) should already be efficient (polynomial) for the problem; you just want to speed up over that.
- Runtime recurrences $T(n) = aT(n/b) + f(n)$ should remind you that you’re splitting the problem in a subproblems of size n/b , solving them (further recursively), and then combining the solutions, and at this level taking computing time $f(n)$.
- You should be thinking: If each subproblem was already solved, this would be easy. I just wish I could break it down smaller. . .
- Pitfalls to look out for: merging the solutions shouldn’t be too difficult, or the subproblems probably aren’t independent.

Dynamic Programming:

- Appropriate when you have recursive subproblems that are not independent, and when there’s a clever order that allows us to build up the answers to avoid recursive computation.

- You should be thinking: I can't figure out whether this thing is in my optimal solution or not!! Wait, so there are multiple cases to maximize over... either this thing is in my optimal solution, or it's not (could be more than two cases)—leads to your recurrence!
- Pitfalls to look out for: Is your recurrence (1) well-defined (base-cases), (2) built in the right order (memo-table), and (3) correct (read it to yourself in English)?

Linear Programming:

- The *fractional* relaxation can be solved in polynomial-time—but not if we constrain to only integral solutions! (Because we can express NP-Hard problems this way.)
- The objective function only improves with fractional solutions.
- The dual problem of a maximization gives an *upper bound*, and its *objective function* is equal when they are both optimal. It is *not* an equivalent problem, however. It is a very different problem, asking a different question, with different types of solutions. The values are just equal when optimal (and only when optimal).
- Weak duality (upper bound), strong duality (equality), and complementary slackness (primal constraint vs. dual variable and vice versa) are the parts of duality theory we talked about.

Homework 7 Part 2: Dual Greed

Re: Maximum Spanning Tree (Kruskal) and Shortest Path (Dijkstra).

- a. Consider a connected undirected graph $G = (V, E)$ in which each edge e has a weight w_e . For a subset $F \subseteq E$, let $\kappa(F)$ denote the number of connected components in the subgraph (V, F) .

Prove that the spanning trees of G correspond to the integer solutions to the following linear program with the same objective function value (with decision variables $\{x_e\}_{e \in E}$):

$$\begin{aligned}
 & \max \quad \sum_{e \in E} w_e x_e \\
 & \text{subject to} \quad \sum_{e \in F} x_e \leq |V| - \kappa(F) && \forall F \\
 & \quad \quad \quad \sum_{e \in E} x_e = |V| - 1 \\
 & \quad \quad \quad x_e \geq 0 && \forall e \in E.
 \end{aligned}$$

(While this linear program has a huge number of constraints, we are using it purely for the analysis of Kruskal's algorithm.)

Let x_e be our decision variable indicating whether an edge e is in our spanning tree. Then, recall the properties of a spanning tree. It is some set of edges T such that it “spans”

V , or is a maximal set of edges of G that does not contain a cycle. We will show that a spanning tree is a valid solution to this LP.

Note that from this definition, we know $|T| = |V| - 1$, as if this were not the case then the set would contain a cycle, or we could add another edge (per the above definition of T). Thus the second constraint is satisfied. For the first constraint, we can imagine that we are in the process of building a spanning tree and currently have edges F , and can think of $\kappa(F) - 1$ as the number of edges left to pick in order to span the graph, to form a tree. If this constraint is not satisfied, $\sum_{e \in F} x_e > |V| - \kappa(F)$, in which case, when we add in the other edges, we will have $x_e + \kappa(F) > V$, which implies a cycle, since $\kappa(F)$ will then be 1 when we have formed a tree.

b. What is the dual of this linear program?

With y_F as our dual variable for the first constraints and z as our dual variable for the constraint, we get the following dual:

$$\begin{aligned} \min \quad & \sum_{F \subseteq E} (|V| - \kappa(F))y_F + (|V| - 1)z \\ \text{subject to} \quad & \sum_{F: e \in F} y_F + z \geq w_e & \forall e \in E \\ & y_F, z \geq 0 & \forall F \subseteq E. \end{aligned}$$

c. What are the complementary slackness conditions?

- (1) For all $F \subseteq E$, either $\sum_{e \in F} x_e = |V| - \kappa(F)$ or $y_F = 0$. Additionally, either $\sum_{e \in E} x_e = |V| - 1$ or $z = 0$.
- (2) For all $e \in E$, either $\sum_{F: e \in F} y_F + z = w_e$ or $x_e = 0$.

e. Now consider the problem of computing a shortest path from s to t in a directed graph $G = (V, E)$ with a nonnegative cost c_e on each edge $e \in E$. Prove that every simple s - t path of G corresponds to an integer solution of the following linear program with the same objective function value:¹

$$\begin{aligned} \min \quad & \sum_{e \in E} c_e x_e \\ \text{subject to} \quad & \sum_{e \in \delta^+(S)} x_e \geq 1 & \forall S \subseteq V \text{ with } s \in S, t \notin S \\ & x_e \geq 0 & \forall e \in E. \end{aligned}$$

¹Recall that $\delta^+(S)$ denotes the edges sticking out of S .

(Again, this huge linear program is for analysis only.)

For any simple s - t path P , let x_e be the decision variable indicating whether an edge e is in our path P . For every subset S such that $s \in S$ and $t \notin S$, the path P must eventually have some edge $e = (i, j)$ such that $i \in S$ but $j \notin S$ due to the fact that $t \notin S$. Then the constraint $\sum_{e \in \delta^+(S)} x_e \geq 1$ for this subset S will be satisfied by this edge. Hence every simple s - t path is a valid solution to this LP.

f. What is the dual of this linear program?

With y_S as our dual variable for the first constraints, we get the following dual:

$$\begin{aligned} & \max && \sum_{\{S \subseteq V \mid s \in S, t \notin S\}} y_S \\ \text{subject to} &&& \sum_{\{S \subseteq V \mid s \in S, t \notin S\}: e \in \delta^+(S)} y_S \geq c_e && \forall e \in E \\ &&& y_S \geq 0 && \forall \{S \subseteq V \mid s \in S, t \notin S\}. \end{aligned}$$

g. What are the complementary slackness conditions?

- (1) For all $S \subseteq V$ with $s \in S, t \notin S$, either $\sum_{e \in \delta^+(S)} x_e = 1$ or $y_S = 0$.
- (2) For all $e \in E$, either $\sum_{\{S \subseteq V \mid s \in S, t \notin S\}: e \in \delta^+(S)} y_S = c_e$ or $x_e = 0$.

Review Problems

Part 1

You're working at an investment company that asks you: given the opening price of the stock for n consecutive days in the past, days $i = 1, 2, \dots, n$, given the opening price of the stock $p(i)$ for each day i , on which day i should the company have bought and which later day j should they have sold shares in order to maximize their profits? If there was no way to make money during the n days, you should report this instead.

For example, suppose $n = 3, p(1) = 9, p(2) = 1, p(3) = 5$. Then you should return "buy on 2, sell on 3" (buying on day 2 and selling on day 3 means they would have made \$4 per share, the maximum possible for that period).

Clearly, there's a simple algorithm that takes time $O(n^2)$: try all possible pairs of buy/sell days and see which makes them the most money. Your investment friends were hoping for something a little better. Show how to find the correct numbers i and j in time $O(n \log n)$.

Whenever we see that brute-force is $O(n^2)$ and a speed-up is $O(n \log n)$, we should think of divide and conquer.

A natural approach would be to consider the first $n/2$ days and the final $n/2$ days separately, solving the problem recursively on each of these two sets, and then figure out how to get an overall solution from this in $O(n)$ time. This would give us $T(n) \leq 2T(\frac{n}{2}) + O(n)$, and hence $O(n \log n)$.

Our main observation is that there are three cases when we split the days into two sets:

- We buy then sell within the first $n/2$ days—this is the optimal solution on the days $1, \dots, n/2$.
- We buy then sell within the last $n/2$ days—this is the optimal solution on the days $n/2 + 1, \dots, n$.
- We buy in the first $n/2$ days and sell in the last $n/2$ days: then the day we buy i is the *minimum* price among days $1, \dots, n/2$ and the day we sell j is the *maximum* among days $n/2 + 1, \dots, n$.

The first two alternatives are computed in time $T(n/2)$, each by recursion, and the third alternative is computed by finding the minimum in the first half and the maximum in the second half, which takes time $O(n)$. Thus the running time $T(n)$ satisfies

$$T(n) \leq 2T\left(\frac{n}{2}\right) + O(n),$$

as desired.

This is actually not the best running time achievable for this problem. In fact, one can find the optimal pair of days in $O(n)$ time using dynamic programming—do you see how? (But, this question was definitely set up to get you to use D&C. Either answer would be correct.)

Algorithm 1 investing($p(1), \dots, p(n)$).

Input: Prices $p(i)$ for days $i = 1, \dots, n$.

if $n = 1$ **then**

return (Null, Null)

else

 Let Buy1, Sell1 = investing($p(1), \dots, p(n/2)$) and Buy2, Sell2 = investing($p(n/2+1), \dots, p(n)$)

 Let Buy3 = argmin{ $p(1), \dots, p(n/2)$ } and Sell3 = argmax{ $p(n/2 + 1), \dots, p(n)$ }

if $p(\text{Sell3}) - p(\text{Buy3}) < 0$, **then** (Buy3, Sell3) = (Null, Null)

return (Buy, Sell) \in argmax{ $p(\text{Sell1}) - p(\text{Buy1}), p(\text{Sell2}) - p(\text{Buy2}), p(\text{Sell3}) - p(\text{Buy3})$ }

end if

Claim 1. For any natural number n days, given prices $p(1), \dots, p(n)$, the above algorithm returns the optimal day Buy and day Sell to maximize profits $p(\text{Sell}) - p(\text{Buy})$.

Proof by strong induction on n .

Base Case: $k = 1$. When there is only 1 day, we cannot both buy and sell, so we return Null for the days to buy and sell on—do not trade.

Inductive Hypothesis: Assume the algorithm correctly finds the best Buy and Sell days in order if there are some, and otherwise returns (Null, Null) on $k < n$ days for some n .

Inductive Case: Given n days of prices, the algorithm considers 3 cases: when we buy and sell in the first $n/2$ days, when we buy and sell in the second $n/2$ days, and when we buy in the first $n/2$ days and sell in the second $n/2$ days. In the third instance, the revenue will be maximized by choosing the minimum price from the first $n/2$ days and the maximum from the second $n/2$ days, as the algorithm does. The first instance requires the algorithm's solution the first $n/2$ days, which is correct by the inductive hypothesis, and the second instance the algorithm's solution on the second $n/2$ days, again correct by the inductive hypothesis. We then take the maximum of these these profits and return the days that give this. If none of these three options give profits, then we return Null—no days should be bought/sold on—which is the same as the solutions on the first and second $n/2$ days. This implies that the algorithm is correct. \square

Runtime Analysis: Let $T(n)$ denote the algorithm's worst-case runtime on two lists of size n . There are two recursive subproblems of size $n/2$ and $O(n)$ computational steps outside the recursive call (finding min/max). Hence, we have the recurrence: $T(n) = 2T(n/2) + O(n)$. This can be solved to obtain $T(n) = O(n \log n)$: each layer of recursion has $O(n)$ work, and there are $\log n$ layers.

Part 2

Consider the problem of making change for n cents using the fewest number of coins. Assume that each coin's value is an integer.

- a. Describe a greedy algorithm to make change consisting of quarters, dimes, nickels, and pennies. Prove that your algorithm yields an optimal solution.

Algorithm: Starting with no coins, take the largest denomination coin given that it is worth at most n minus how much the coins you've taken so far are worth.

Proof by Greedy Exchange. Let $A = (c_1, \dots, c_k)$ be the coins generated by the greedy algorithm and let $O = (o_1, \dots, o_m)$ be the coins generated by some other algorithm. Let q_A, q_O be the number of quarters generated by each, and dimes d_A, d_O , and so on. We can write

$$n = 25q_A + 10d_A + 5n_A + p_A \quad \text{and} \quad n = 25q_O + 10d_O + 5n_O + p_O.$$

By definition of the greedy algorithm, $q_A \geq q_O$ and $n - 25q_A < 25$. If $q_O < q_A$ and $n - 25q_O \geq 25$, we exchange multiple smaller coins for one larger quarter, improving its solution. We do this until $q_O = q_A$ (or it already does), then we set aside quarters. We then continue process on dimes, then nickels, until $O = A$, and we have only reduced the number of coins, hence A is optimal. \square

- b. Give a set of coin denominations (so *not* penny, nickel, dime, quarter) for which the greedy algorithm does not yield an optimal solution. Your set should include a penny so that there is a solution for every value of n .

Let the coin denominations be 1, 5, 7. Suppose you want to make change for $n = 10$. The greedy algorithm will use one 7 coin and three 1 coins, for a total of four coins. It is optimal to use two 5 coins.

Part 3

A palindrome is a nonempty string over some alphabet that reads the same forward and backward. Examples of palindromes are: (1) all strings of length 1, (2) civic, (3) racecar, and (4) aibohphobia (fear of palindromes). Give an efficient algorithm to find the longest palindrome that is a subsequence of a given input string. For example, given the input “character,” your algorithm should return “carac.” What is the running time of your algorithm?

Subproblem: Let $\text{OPT}(i, j)$ be the length of the longest palindrome of the input string from character i through character j .

Recurrence:
$$\text{OPT}(i, j) = \begin{cases} \text{OPT}(i + 1, j - 1) + 2 & \text{if } i = j \\ \max \{ \text{OPT}(i + 1, j), \text{OPT}(i, j - 1) \} & \text{otherwise.} \end{cases}$$

Proof of Recurrence: Consider any subsequence from i to j —either the first character is equal to the last, or it is not. If they are, the longest palindrome subsequence from i to j contains these characters, so we count them (+2), remove them from the ends, and then count the longest palindrome subsequence on the remaining $i + 1$ to $j - 1$. If not, then the longest palindrome subsequence on i to j is either the longest palindrome subsequence on i to $j - 1$ or the longest palindrome subsequence on $i + 1$ to j , as we know the ends are not the same.

Base Cases: $\text{OPT}(i, i) = 0$ and $\text{OPT}(i, i + 1) = 1$ for all i —all sequences of length 0 count 0 and all sequence of length 1 are length 1 palindromes.

Solution to Original Problem: $\text{OPT}(1, n)$

Algorithm: See below.

Runtime: The algorithm fills a table of size $O(n^2)$, doing $O(1)$ table lookups to fill each entry. Base cases and returns are constant. The total runtime is therefore $O(n^2)$.

Computing a Solution: See below, return `computePalindromeSol(String, 1, n, memo)`.

Algorithm 2 `longestPalindromeSubsequence(String)`.

Input: A string `String` of length n .

`Memo[][] = new int[n][n]`

for ℓ from 0 to $n - 1$ **do**

for i from 1 to $n - \ell$ **do**

$j = i + \ell$

if $\ell = 0$ **then**

`Memo[i][i] = 0`

else if $\ell = 1$ **then**

`Memo[i][i + 1] = 1`

else if `String[i] = String[j]` **then**

`Memo[i][j] = Memo[i + 1][j - 1] + 2`

else

`Memo[i][j] = max{Memo[i + 1][j], Memo[i][j - 1]}`

end if

end for

end for

return `Memo[1][n]`

Algorithm 3 `computePalindromeSol(String, i, j, memo)`.

Input: `String`, indices $i \leq j$, and filled out memo table.

if $i == j$ **then**

return `Null`

else if $j == i + 1$ **then**

return `String[i]`

else

if `String[i] = String[j]` **then**

return `String[i] + computePalindromeSol(String, i + 1, j - 1, memo) + String[i]`

else if `Memo[i + 1][j] > Memo[i][j - 1]` **then**

return `computePalindromeSol(String, i + 1, j, memo)`

else

return `computePalindromeSol(String, i, j - 1, memo)`

end if

end if
