

Instructions (that will appear on the real final exam)

- **DO NOT OPEN THE EXAM UNTIL YOU ARE INSTRUCTED TO.**
- Answer all of the questions as well as you can. You have **180 minutes**.
- The exam is **non-collaborative**; you must complete it on your own. If you have any clarification questions, please ask the course staff. We cannot provide any hints or help.
- This exam is **closed-book**, except for **up to three double-sided sheets of paper** that you have prepared ahead of time. You can have anything you want written on these sheets of paper.
- **Please DO NOT separate pages of your exam.** The course staff is not responsible for finding lost pages, and you may not get credit for a problem if it goes missing.
- There are a few pages of extra paper at the back of the exam in case you run out of room on any problem. If you use them, please clearly indicate on the relevant problem page that you have used them, and please clearly label any work on the extra pages.
- Please make sure to sign out of the roster when handing in your completed exam to the teaching team.
- **Please do not discuss the exam until after solutions are posted!** Of course for this practice exam, feel free to discuss the solutions with the course staff and anyone else!

General Advice

- If you get stuck on a question or a part, move on and come back to it later. The questions on this exam have a wide range of difficulty, and you can do well on the exam even if you don't get a few questions.
- Pay attention to the point values. Don't spend too much time on questions that are not worth a lot of points.
- There are **110** total points on this exam. There are **seven problems** across **22 pages**.

Name and SUNet ID (please print clearly):

This page intentionally blank. Please do not write anything you want graded here.

Honor Code

The Honor Code is an undertaking of the Stanford academic community, individually and collectively. Its purpose is to uphold a culture of academic honesty. Students will support this culture of academic honesty by neither giving nor accepting unpermitted academic aid on this examination.

This course is participating in the proctoring pilot overseen by the Academic Integrity Working Group (AIWG), therefore proctors will be present in the exam room. The purpose of this pilot is to determine the efficacy of proctoring and develop effective practices for proctoring in-person exams at Stanford.

Unpermitted Aid on this exam includes but is not limited to the following: collaboration with anyone else; reference materials other than your cheat-sheet (see below); and internet access.

Permitted aid on this exam includes a “cheat-sheet:” two double-sided sheets of paper with anything written on them, which you have prepared yourself ahead of time.

I acknowledge and attest that I will abide by the Honor Code:

[signed] -----

Exam Break Sign-out

I pledge that during my exam break:

- I will not bring any paper, electronic devices (phone, smart watch, smart glasses, etc), or aid (permitted or unpermitted) *out of or into* the exam room.
- I will not communicate with anyone other than the course instructional staff about the content of the exam.

Signature Confirming Honor Code Pledge	Exit Time	Return Time	Proctor Initial	Length (min)

If you are feeling unwell and are not able to complete the exam, please connect with the proctor to discuss options.

Good Luck!

This page intentionally blank. Please do not write anything you want graded here.

Multiple Choice and Short Answer

1. (29 pt.) [Multiple Choice!] For each of the parts below, clearly fill in your answers. Ambiguously filled-in answers will be marked incorrect.



Filled in means:
“This is a
correct choice”

Empty means:
“This is **not** a
correct choice”

Anything else
may be marked
as incorrect!

[We are expecting: *For each part, just clearly filled-in-answers. No explanation is required or will be considered while grading.*]

(a) (5 pt.) Which of the following algorithms can find shortest paths in weighted directed graphs, that may have negative edge weights (but no negative cycles). Select “Yes” for all that apply.

A. BFS

Yes

No

B. Dijkstra’s Algorithm

Yes

No

C. Bellman-Ford Algorithm

Yes

No

D. Floyd-Warshall Algorithm

Yes

No

E. Ford-Fulkerson Algorithm

Yes

No

(b) **(3 pt.)** In which of the following scenarios would Prim vs. Kruskal be asymptotically faster for finding a valid minimum spanning tree? Select all that apply.

- A. On a graph $G = (V, E)$, where $|E| = \Theta(|V|)$ and the edge weights are all integers from 1 to 10. Assume Prim's algorithm uses a Fibonacci heap.
 - Prim's is asymptotically faster
 - Kruskal's is asymptotically faster
 - They are asymptotically the same
- B. On a graph $G = (V, E)$ where $|E| = \Theta(|V|^2)$, and the edge weights are arbitrary comparable items. Assume Prim's algorithm uses a Fibonacci heap.
 - Prim's is asymptotically faster
 - Kruskal's is asymptotically faster
 - They are asymptotically the same
- C. You have no guarantees on the graph (other than that it is connected), and you are using an RBTree in Prim's algorithm.
 - Prim's is asymptotically faster
 - Kruskal's is asymptotically faster
 - They are asymptotically the same

(c) **(3 pt.)** Suppose that $s \rightarrow v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_\ell \rightarrow t$ is a shortest path from s to t in a connected graph with (possibly negative) edge weights, but no negative cycles.

Which of the following are true?

- A. The path $s \rightarrow v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_i$ is a shortest path for all $i = 1, \dots, \ell$.
 - True
 - False
- B. If all the edge weights in G are distinct, then this shortest path must be unique.
 - True
 - False
- C. If you were to double all the edge weights in G , then $s \rightarrow v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_\ell \rightarrow t$ would still be a shortest path in the transformed graph.
 - True
 - False

(d) (5 pt.) Suppose you want to count the number of ways to tile a $2 \times n$ grid with tiles. You have (an infinite number of) tiles that are 2×2 , 1×2 , and 2×1 . For example, if $n = 6$, here are two legitimate tilings:



Let $D[j]$ be the number of legitimate tilings of a $2 \times j$ board. Which of the following is a correct recursive relationship between these sub-problems?

- A. $D[j] = D[j - 2] + D[j - 1]$
- B. $D[j] = D[j - 2] \times D[j - 1]$
- C. $D[j] = D[j - 2] + 2 \cdot D[j - 1]$
- D. $D[j] = 2 \cdot D[j - 2] + D[j - 1]$
- E. None of the above.

(e) (5 pt.) Suppose you have M identical items that you want to allocate to n people. Each person wants as many items as possible, up to some limit B , but they have diminishing returns. If person i already has $s - 1$ items, then the value they get from the s 'th item is $\Delta_{i,s}$, where


$$\Delta_{i,1} \geq \Delta_{i,2} \geq \cdots \geq \Delta_{i,B}.$$

There is no benefit to anyone getting more than B items.

You want to decide how to allocate items to maximize total value among all n people. You will give out the items one at a time, following a greedy strategy. Which greedy rule is correct for this problem? Choose the best answer. Assume that ties are broken arbitrarily.

- A. Give the next item to the person with the smallest *total value so far*. (If person i has $s - 1$ things so far, their *total value so far* is $\sum_{j=1}^{s-1} \Delta_{i,j}$).
- B. Give the next item to the person who has the most *marginal benefit* from the next item. (If person i has $s - 1$ things so far, their *marginal benefit* from the next item is $\Delta_{i,s}$).
- C. Give the next item to the person who has the most *left to gain*. (If person i has $s - 1$ things so far, the *amount they have left to gain* is $\sum_{j=s}^B \Delta_{i,j}$).

(f) (5 pt.) Consider the following graph, which shows an s - t flow. Below, the notation $x : y$ means that the edge has flow x out of total capacity y .



A. Which of the following are augmenting paths in G ? Fill in one bubble for each path.

Path	Augmenting Path	Not an augmenting path
$s \rightarrow a \rightarrow t$	<input type="radio"/>	<input type="radio"/>
$s \rightarrow a \rightarrow b \rightarrow t$	<input type="radio"/>	<input type="radio"/>
$s \rightarrow b \rightarrow a \rightarrow t$	<input type="radio"/>	<input type="radio"/>

B. What is the maximum flow in this graph? Choose exactly one.

- 2
- 3
- 4
- 5

(g) (3 pt.) Suppose Alice is designing an algorithm to decide which public school students should get need-based access to a free tutoring service. Alice's algorithm uses a student's GPA as an input. Barbara raises a concern that the algorithm does not consider whether the student is taking advanced, on-level, or remedial math courses, and a low grade in an advanced course may show less need than a low grade in a remedial course. What *best* describes the issue Barbara is raising?

- A. Barbara is concerned about an abstraction
- B. Barbara is concerned about an idealization
- C. Barbara is concerned about an NP-hard problem

2. (9 pt.) [Recurrence Relations] For each part below, give the best big-Oh bound you can for each $T(n)$.

[We are expecting: For each part, an expression “ $O(\text{something in terms of } n)$.” No explanation is required or will be considered. Your answer should be as simplified as possible.]

(a) (3 pt.) $T_1(n) = 4T_1(\lfloor n/2 \rfloor) + n^3$ for $n \geq 3$, $T_1(n) = 1$ for $n < 3$.

$$T_1(n) = \text{-----}$$

(b) (3 pt.) $T_2(n) = T_2(\lfloor n/2 \rfloor) + \log_2 n$ for $n \geq 2$, $T_2(n) = 1$ for $n < 2$.

[HINT: Recall that $\sum_{j=0}^T j = \frac{T(T+1)}{2}$, and $\log_2(a/b) = \log_2(a) - \log_2(b)$.]

$$T_2(n) = \text{-----}$$

(c) (3 pt.) $T_3(n) = T_3(\lfloor n/2 \rfloor) + T_3(\lfloor n/4 \rfloor) + n$ for $n \geq 4$, $T_3(n) = 1$ for $n < 4$.

$$T_3(n) = \text{-----}$$

3. (20 pt.) [True or False and Why?] For each of the parts below, say whether the statement is true or false, and give a *short* explanation saying why. You don’t have to give a formal proof, but your explanation should be convincing to the grader. You can use any statements/algorithms we have seen in class.

(-) (Example:) Let G be an undirected unweighted connected graph. True or False: Given s and t , it is possible to find a shortest path from s to t in G in time $O(n + m)$.

Answer: True. We can run BFS starting at s to find the distance from s to any other vertex in an unweighted graph, in time $O(n + m)$.

(a) (5 pt.) Suppose that $f(n)$ and $g(n)$ are positive functions of n , and that $f(n) = O(g(n))$. True or False: $f(n) + g(n) = O(g(n))$.

Explain your answer, using the definition of big-O.

(b) **(5 pt.)** Your friend thinks they have a way to get Dijkstra's algorithm to work with negative edge weights. Suppose that G is a weighted directed graph, so that the weight of $(u, v) \in E$ is $w(u, v)$. Suppose that the smallest (most negative) edge weight is $-w^* < 0$, so $w^* = -\min_{(u,v) \in E} w(u, v)$. Let G' be the same as G , except with edge weights $w'(u, v) = w(u, v) + w^*$. By construction, G' has no negative edge weights, so we can run Dijkstra's algorithm on G' and return the resulting shortest paths.

True or false: this algorithm (constructing G' and running Dijkstra's algorithm on it) will correctly return shortest *paths* in G . (It's okay if it doesn't return the correct path lengths). If your answer is "True," explain why. If your answer is "False," give a counter-example.

(c) **(5 pt.)** Your friend has developed a new hash family \mathcal{H} that maps the integers $1, \dots, M$ to n buckets, for $n \geq 2$. Every function $h \in \mathcal{H}$ has the property that it maps any pair $x \neq y$ of integers in $\{1, \dots, M\}$ to distinct buckets (that is, $h(x) \neq h(y)$ for all $h \in \mathcal{H}$), *except* for 3 and 5: for these two values, we will have $h(3) = h(5)$ for all $h \in \mathcal{H}$. However, your friend tells you that for different $h \in \mathcal{H}$, which bucket 3 and 5 land in is uniformly distributed over the buckets. (That is, for any $i \in \{1, \dots, n\}$ $\Pr_{h \in \mathcal{H}}[h(3) = i] = 1/n$, and similarly for $h(5)$).

True or False: \mathcal{H} is a universal hash family. Explain your answer, using the definition of universal hash family.

(d) **(5 pt.)** In class, we saw a *minimum* spanning tree, which is a spanning tree of minimum cost. A *maximum* spanning tree is a spanning tree of maximum cost. Consider modifying Kruskal's algorithm by, instead of greedily taking the lowest-cost edge that doesn't form a cycle, you take the *highest*-cost edge that doesn't form a cycle. Call this new algorithm max-Kruskal.

True or false: max-Kruskal correctly finds a *maximum* spanning tree in a connected, undirected, weighted graph. Explain your answer.

Note: This next problem may be trickier, and is only worth 10 points out of 110 points total. (There are two independent 5-point parts). You may want to come back to this after you have finished the rest of the exam.

4. **(10 pt.) How to do it?** For each of the tasks below, explain briefly how you would do it. You may use any algorithm we have seen in class as a black box.

[We are expecting: *For each, a clear English explanation of your algorithms. You do not need to include pseudocode, but you can if you think it makes your answer clearer. You do not need to justify correctness or running time.]*

(a) **(5 pt.)** *[May be trickier.]* Let $G = (V, E)$ be an unweighted directed graph. Recall that G is *strongly connected* if for every $u, v \in V$, there is a path from u to v in G , and there is a path from v to u in G . Say that G is “*not-so-strongly-connected*” if, for every $u, v \in V$, either there is a path from u to v in G ; or there is a path from v to u in G . Give an algorithm to decide if a graph is not-so-strongly-connected in time $O(n + m)$.

Another part on next page

(b) **(5 pt.)** *[May be trickier.]* Let $G = (V, E)$ be a connected, unweighted, undirected graph. Let $s, t \in V$ be distinct vertices. Call a vertex $v \in V$ *essential* if every shortest path from s to t in G passes through v . Given G , s , and t , find a list of *all* of the essential vertices, in time $O(n + m)$.

[HINT: Suppose that v lies on a shortest path from s to t . What can you say about the distance from s to v and the distance of t to v , relative to the distance from s to t ?]


Algorithm Design

5. (16 pt.) [Keeping the dream alive] Suppose that you follow a sports league with $n+1$ teams, teams $0, 1, \dots, n$. Team 0 (let's call them Team Cardinal) is your favorite team. Each team plays each other some number of times, and at the end of the season, the team who has won the most games is declared the champion. (If it's a tie, there are multiple champions).

Right now, we are part-way through the season. Suppose that each team i has won w_i games so far. Team Cardinal (aka Team 0), has G games left to play in the season. For each other pair of teams $i, j \in \{1, 2, \dots, n\}$ (*not* including Team Cardinal), there are $g_{i,j}$ games left in the season between Team i and Team j .

As an avid fan, you want to know if it's still possible for Team Cardinal to become champion. You decide to solve this problem using the Ford-Fulkerson Algorithm.

In more detail, consider the following graph:



This graph has:

- **Vertices:**
 - Special vertices s and t
 - One vertex for each (unordered) pair $\{i, j\}$ for $i, j \in \{1, \dots, n\}$ of distinct teams (*not* including Team Cardinal)
 - One vertex for each team $i \in \{1, \dots, n\}$ (*not* including Team Cardinal)
- **Edges:**

- A directed edge from s to each “ $\{i, j\}$ ” vertex. The weight of these edges is TBD.
- A directed edge from each “ i ” vertex to t . The weight of these edges is TBD.
- For each distinct pair $i, j \in \{1, \dots, n\}$, there is a directed edge from vertex “ $\{i, j\}$ ” to vertex “ i ”; and a directed edge from vertex “ $\{i, j\}$ ” to vertex “ j ”. All these edges have weight ∞ .

(a) **(7 pt.)** Explain how to set the weights on the “?” edges (the ones with weight TBD) so that a maximum flow in this graph tells you whether or not Team Cardinal can still be champions.

That is, your goal is to decide if it is still possible for Team Cardinal to win at least as many games as any team $1, \dots, n$, assuming they win all their remaining G games, as well as the w_0 games they have already won.

[**HINT:** You'll need to use the parameters $g_{i,j}$, w_i , and G defined above.]

[**HINT:** Imagine that the flow on the edge $(\{i, j\}, i)$ represents how many games Team i will win against Team j going forward...]

[**We are expecting:** Just a description of the edge weights. No justification is required (yet)]

(b) **(3 pt.)** Given your edge weights in the previous part, explain how to use the Ford-Fulkerson algorithm to tell if Team Cardinal can still be champions.

[**We are expecting:** A clear English description of your algorithm.]

(c) **(6 pt.)** Explain why your algorithm is correct.

[We are expecting: *A clear explanation; it doesn't have to be a formal proof.*]

6. (14 pt.) [Procrastination Optimization] Over the upcoming break, Lucky the Lackadaisical Lemur is excited to kick back and binge-watch some shows. Lucky wants to watch k episodes over break; he's not very discerning, so he doesn't care what he watches.

Lucky could access to m different platforms (LemFlix, Madagascar Prime-ate, Ring-tail+, etc). Suppose that the i 'th platform has a fee of s_i to join. (Note: Lucky is planning on subscribing for one month during the break and then canceling, so it's a one-time fee for him). Once he's joined Platform i , each piece of content on Platform i costs some additional fee: Say that there are Q things on each platform, and that the j 'th one costs $f_{i,j}$. You can assume that the pieces of content are ordered by price, so

$$0 \leq f_{i,1} \leq f_{i,2} \leq \cdots \leq f_{i,Q}.$$

Design a DP algorithm to help Lucky find the minimum cost to watch k things across all these m platforms. Your algorithm should use the sub-problems:

$D(i, j)$ = The least money Lucky can spend to watch j things across the first i platforms.

If there aren't j distinct things to watch on the first i platforms, then $D(i, j) = \infty$.

(a) (2 pt.) What is $D(i, 0)$ for each $i \in \{1, \dots, m\}$?

[We are expecting: Your answer, no explanation required.]

(b) (2 pt.) What is $D(1, j)$ for each $j \in \{1, \dots, k\}$?

[We are expecting: Your answer, and a sentence or two of explanation.]

(c) **(5 pt.)** Write a recurrence relation for the problems $D(i, j)$.

[We are expecting: *Your answer, and a sentence or two of explanation.*]

(d) **(5 pt.)** Design an algorithm that runs in time $O(mQk)$ to compute the minimum cost that Lucky can spend. Your algorithm should take as input k , and array F so that $F[i, j] = f_{i,j}$.

[We are expecting: *Clear pseudocode, and a justification of the running time.*]

Algorithm Analysis

7. (12 pt.) [FIF Caching]

Consider the following caching problem. Suppose we are given:

- A sequence p_1, \dots, p_n of pages in memory that an algorithm is going to ask for, in order. (We know this up-front). Some of the pages might repeat, so it's possible, say, that $p_3 = p_{17}$.
- A small fast cache, which can store at most $k < n$ pages.

At each timestep t , the page p_t arrives. If p_t is not in the cache at time t , then:

- We have a *cache miss*, and have to go find p_t in memory, which is costly.
- We add p_t to our cache.
- We *evict* another page from the cache (of our choice) to make room for p_t .

If p_t is in the cache at time t , then:

- We get a *cache hit*, and can retrieve the page quickly.
- If we like, we have the option to evict a page from cache and retrieve another one from memory, although this usually wouldn't be a good idea (since it would just be extra memory calls).

Our goal is to design an *eviction policy* (that is, a schedule for which page gets evicted at each time step, if any) to minimize the number of memory accesses over all n timesteps.

In this problem, we will analyze the FIF (farthest-in-the-future) policy. This policy works as follows. At time t , if p_t is not in the cache, we find the page x in the cache that next occurs furthest in the future. That is, for each x , we compute

$$\text{NEXTTIME}(x) = \min\{i > t : x = p_i\},$$

and choose the x with the largest NEXTTIME value. By convention, if x never appears again, $\text{NEXTTIME}(x) = \infty$. The policy breaks ties arbitrarily.

(a) **(2 pt.)** To make sure that you understand the algorithm, suppose that the stream of items was a, b, c, a, b, a, c ; $k = 2$; and that the original cache was $\{a, b\}$. At each time-step, what is the cache after FIF has evicted an element? Was there a cache hit or a cache miss? We have filled in the first three for you.

[We are expecting: *The rest of the table filled in.*]

Time t	Item p_t	Cache	Hit or Miss?
1	a	$\{a, b\}$	hit
2	b	$\{a, b\}$	hit
3	c	$\{a, c\}$	miss
4	a		
5	b		
6	a		
7	c		

(b) **(10 pt.)** Our goal is to prove that the FIF policy is optimal (that is, it minimizes the number of memory accesses). To do this, we will follow our usual recipe of showing that our greedy choices don't rule out success. More precisely, we will prove the following claim.

Claim: Suppose that, at time t , there is an optimal eviction policy (for p_1, \dots, p_n and cache size k) consistent with the choices that FIF has made so far. Then at time $t + 1$, there is still an optimal eviction policy consistent with the choices FIF has made so far.

We will walk you through the proof. Suppose that P is an optimal policy that is consistent with the choices made by FIF at time t . We will construct an optimal policy P' that is consistent with the choices made by FIF at time $t + 1$.

To begin, define P' to do exactly the same thing as P (and hence FIF) up to time t . Then define P' to do the same thing that FIF would have done at time $t + 1$.

i. **(1 pt.)** Suppose that there is no cache miss at time $t + 1$. Explain why we are done proving the claim in this case.

[We are expecting: *A sentence or two.*]

Given the previous part, suppose that there is a cache miss at time $t + 1$, so FIF evicts the element x that appears furthest in the future. Since we defined P' to do the same thing as FIF at time $t + 1$, P' evicts x at time $t + 1$. Suppose that P evicts y at time $t + 1$.

ii. **(1 pt.)** Suppose that $y = x$. Explain why we are done proving the claim in this case.

[We are expecting: *A sentence or two.*]

Thus, for the rest of the problem, assume that $y \neq x$.

Let t^* be the first time after $t + 1$ when we can't define P' to do the same thing that P does. For example, if P would evict some w at time t^* , but w is not in the cache of P' at time t^* , then we can't define P' to do the same thing that P does. Define P' to do the same thing as P up until $t^* - 1$. Notice that, up until time $t^* - 1$, the caches only differ by one element: P has x (having evicted y), and P' has y (having evicted x). All other cached items are the same, since P and P' made all the same decisions up until time $t^* - 1$.

We need to figure out how to define P' at time t^* . There are three cases for what can happen. You will address them in the following sub-parts.

The goal is for P and P' to have the same cache after time t^* , and for P' to make no more trips to memory than P .

- iii. **(2 pt.)** Suppose that $p_{t^*} = z$ is the request at time t^* , and that z is not in the cache for *either* P or P' . Suppose that P chooses to evict x and replace it z . How would you define P' 's behavior in this case?

[We are expecting: *A clear description of what P' would do, and why it satisfies the goal above.*]

- iv. **(3 pt.)** Suppose that $p_{t^*} = y$. Then P has a cache miss. Suppose that P chooses to evict an element w . How would you define P' 's behavior in this case?

[HINT: *Recall that it's okay for a schedule to evict something even if there's not a cache miss.*]

[We are expecting: *A clear description of what P' would do, and why it satisfies the goal above.*]

- v. **(3 pt.)** The third case is that $p_{t^*} = x$. Explain why this can't happen.

[We are expecting: *A clear explanation about why this can't happen.*]

Altogether, these cases show us how we can define a schedule P' at time t^* so that P and P' have the same cache at time $t^* + 1$. After that, we can define P' to do the same thing that P would have done, since P' will face all the same situations going forward.

Now that we have fully defined P' , we see that P' is a policy with no more memory accesses than P , so it's still optimal. Moreover, at time $t + 1$, P' made the same decision as FIF. This proves the statement. So FIF is optimal – hooray!

This is the end of the exam!

This is the end of the exam! You can use this page for extra work on any problem.
Keep this page attached to the exam packet (whether or not you use it), and if you want extra work on this page to be graded, clearly label which question your extra work is for, and make a note on the problem page itself.

This page is for extra work on any problem. **Keep this page attached** to the exam packet (whether or not you use it), and if you want extra work on this page to be graded, clearly label which question your extra work is for, and make a note on the problem page itself.

This page is for extra work on any problem. **Keep this page attached** to the exam packet (whether or not you use it), and if you want extra work on this page to be graded, clearly label which question your extra work is for, and make a note on the problem page itself.

This page is for extra work on any problem. **Keep this page attached** to the exam packet (whether or not you use it), and if you want extra work on this page to be graded, clearly label which question your extra work is for, and make a note on the problem page itself.