

6.006
Study

5/21

Start of review of # theory

↳ Use 6.042 notes as well

- study of integers

- a divides b if integer k so $ak = b$

- gcd(a, b)

↳ greatest common divisor of a and b

$$\text{gcd}(18, 24) = 6$$

↳ everything has one

- gcd of a and b is the linear combo of a and b

$$\text{gcd}(a, b) = sa + tb$$

(I seem like I understand 6.042 much better now
↳ always like that 2nd time around)

2

Prime - divisible by itself and 1 only

Every \oplus # has a unique prime factorization
esp if sort by size (since can exist in order)
 $12 = 2 \cdot 2 \cdot 3$

Modular arithmetic

$$a \equiv b \pmod{n}$$

ie $29 \equiv 15 \pmod{7}$

Since $7 \mid (29 - 15)$

*preserved by addition and multiplication

So can do math by remainder

↳ add, multi, sub

1. replace every int by remainder mod n

2. immed: replace any result outside of range w/
remainder on div by n ~~over~~

(3)

\div_n = notation for add then imm. dividing by n
and taking the remainder

\times_n = same for mult.

So the set of int for which this applies = \mathbb{Z}_n
↳ "the ring of integers modulo n "
Ohhhh!!!

relative prime - integers that have no prime
factors in common

Same as $\gcd(a, b) = 1$

Canceling factors does not work

↳ normally $rs = st$

④

Ok back to 6006

\mathbb{Z}_n = the additive group of integers modulo n

↳ defn matches

Two integers rel prime $\gcd(a, b) = 1$

\mathbb{Z}_n^* = int $< n$ and rel prime to n

every els $a \in \mathbb{Z}_n^*$ has an inverse that is easy to compute

Multiplicative inverse

$$x \cdot x^{-1} = 1$$

$$7 \cdot \frac{1}{3} = 1$$

Every $\# \frac{n}{m}$ has an inverse $\rightarrow \frac{m}{n}$

Only $1, -1$ has inverse in \mathbb{Z}_n

In \mathbb{Z}_n $2 \cdot 8 = 1$

↑ are mult inverses

(5)

But 3 does not

since $3 \cdot j = 1$

(no j exists)

is an inverse if the #'s are rel prime
(aka coprime)

(How to find inverse)

Extended Euclidean algo - finds the GCD

"the table thing"

eg $\text{gcd}(120, 23)$

Dividend	Divisor	Quotient	Remainder
120	23	5	5
23	5	4	3
5	3	1	2
3	2	1	1
2	1	2	0

GCD = 1

$$a^{-1} \equiv x \pmod{m}$$

inverse of a is x

$$ax \equiv aa^{-1} \equiv 1 \pmod{m}$$

When found \rightarrow it's unique

6

But how to actually find

$$ax + by = \gcd(a, b)$$

\uparrow a, b are given

Sol to

$$ax = 1 \pmod{m}$$

So $ax - 1 = qm$

$$ax - qm = 1$$

$x = \text{inverse}$

$q = \text{int that discarded}$

$\gcd(a, m) = 1$ is pre det

So what do you get

(ans is $O(\log(m)^2)$)

Also: Euler's theorem

If a is rel prime to m

$$a^{\phi(m)} \equiv 1 \pmod{m}$$

8

$$\begin{array}{cccc} 15 & 8 & 1 & 7 \\ 8 & 7 & 1 & \textcircled{1} \rightarrow \text{gcd} = 1 \end{array}$$

~~$1 \cdot 8 + 0 \cdot 15 = 8$~~
 ~~$0 \cdot 8 + 1 \cdot 15 = 15$~~
 So return $1 \% 15 \in 1$

$$0 \cdot 8 + 1 \cdot 15 = 1$$

Need to do the table

$$7 = 15 - 1 \cdot 8$$

$$1 = 8 - 1 \cdot 7$$

$$= 8 - 1(15 - 1 \cdot 8)$$

$$= 2 \cdot 8 - 1 \cdot 15$$

↑ rem x quotient
 ↑ plug in for
 Combine

chk

So $1 = 2 \cdot 8 + (-1) \cdot 15$
 return $x \% m = 2 \pmod{15} = 2$

9

So after that long digression, I think I can find math. inverses,

Try it again inverse of 8 mod 15

a	b	ax	rem	
8	15	0	8	
15	8	1	7	$= 15 - 1 \cdot 8$
8	7	1	1	$\textcircled{1} = 8 - 1 \cdot 7$

$$= 8 - 1(15 - 1 \cdot 8)$$

$$= \underset{x}{2 \cdot 8} - 1 \cdot \underset{y}{15}$$

return $2 \% 15 = 2$ $\textcircled{2}$

SMX

So where were we?

\mathbb{Z}_n^* = multig. group of int $\leq n$ and rel prim to n
inverse easy (since $\text{gcd} = 1$)

\mathbb{Z}_p^* is a special case when $n = \text{prime}$

(10)

\mathbb{Z}_p^* is a cyclic group of order $p-1$

Cyclic group - a group that can be generated from a single element (g)

So that every element is a power of g

ie $G = \{g^0, g^1, g^2, g^3, g^4, g^5\}$ set defined as

then $g^6 = g^0$ and its cyclic

is isomorphic to $\{0, 1, 2, 3, 4, 5\} \text{ mod } 6$

For every \oplus int n , there is one group whose order is n

Order (group theory)

= Cardinality = # of items inside set

Sometimes also the period

(11)

So \mathbb{Z}_p^* is $\{g, g^2, \dots, g^{p-1}\}$

Example 2 is a generator of \mathbb{Z}_{11}^*

2^1	2	mod 11
2^2	4	
2^3	8	
2^4	5	
2^5	10	
2^6	9	
2^7	7	
2^8	3	
2^9	6	
2^{10}	1	

notice → unique!

How to find generators

↳ knew we talked about in class

②

Fermat's Little Theorem

$$a^{p-1} \equiv 1 \pmod{p} \text{ for all } a \in \mathbb{Z}_p^*$$

So for 2 which is in \mathbb{Z}_{11}^*

$$2^{11-1} \equiv 1 \pmod{11}$$

$$2^{10} = 1 \quad \checkmark$$

How about $2^7 \equiv 7$

$$7^{11-1} \pmod{11} = 1 \quad \checkmark$$

(I remember this from 6.042 - never understood...)

$\phi(n)$ is Euler totient function

$$\phi(n) = |\mathbb{Z}_n^*|$$

= The # of ints that are rel prime to n

How to find again?

(13)

Plus $k^{\phi(n)} \equiv 1 \pmod{n}$

$$\phi(mn) = \phi(m)\phi(n)$$

Gives the order of \mathbb{Z}_n^*

$$\phi(n) = n \prod_{p|n} \left(1 - \frac{1}{p}\right)$$

~~the~~ think that's all we need to know

When p is prime all int are rel prime to p

↳ kinda by def'n

$$\phi(p) = p - 1$$

If p^k is a prime power then all int b/w
1 and p^k are rel prime to p^k

except $p, 2p, 3p, \dots, p^{k-1}p$

$$\phi(p^k) = p^k - p^{k-1}$$

(14)

Mod exponentiation

↓ Ist actual problem

Given $a; x, n$ compute $a^x \text{ mod } n$

Naive

1. set $y = 1$
2. Repeat x times $y = y * a$
3. Return $y \text{ mod } n$

But ~~compute~~ that property cumulative:

1. set $y = 1$
2. Repeat x times $y = (y * a) \text{ mod } n$
3. Return y

But still has x multiplications

So Repeated Squaring

Uses $(a^k)^2 = a^{2k}$

So $(a^{2^i})^2 = a^{2 * 2^i} = a^{2^{i+1}}$

15

So we can obtain ~~a^{2ⁱ}~~ a^{2^i} in only i multiplications

How does this work exactly?

Example: Think saw one in lecture

Was this Fast exponentiation in 6.042?

Algorithmist example		
Say want	3^{13}	
Know	3^1	3
	3^2	$3 \cdot 3 = 9$
	3^4	$9 \cdot 9 = 81$
	3^8	$81 \cdot 81 = 6561$

So $3^{13} = 3^8 \times 3^4 \times 3^1$
 $[8+4+1=13]$
 $= 6561 \cdot 81 \cdot 3$
 $= 1594323$

But under exp by squaring on WP - lots of examples

(6)

So this works in $\log n$ multiplications

But only ~~for~~ if it's a power of 2!

If x is not a power of 2, can write

as

$$x = \sum_{i=1}^k b_i 2^i$$

where $b_k b_{k-1} \dots b_0$ is the binary representation of x

then $k = \lfloor \log_2(x) \rfloor$

$$\begin{aligned} a^{nx} &= a^{\sum_{i=1}^k b_i 2^i} \\ &= \prod_{i=1}^k a^{b_i 2^i} \\ &= \prod_{i=1}^k \left(a^{(2^i)} \right)^{b_i} \end{aligned}$$

Why are doing this? binary!

(17)

So this is prod of $a^{(2^i)}$ for i where $b_i = 1$

So $O(k) = O(\log x)$ multiply

i. Make a matrix $A[i] = a^{2^i}$

ii. Set $A[0] = a$

iii. For $i = 1, \dots, k$ $A[i] = (A[i-1])^2 \pmod n$

2. Obtain the binary representation

3. $y = 1$

4. For $i = 1, \dots, k$: If $b_i = 1$, set $y = (A[i] \cdot y) \pmod n$

5. Return y

So what is this binary rep?

↳ just the # in binary

I'm guessing k is # of digits

↳ yeah since that is $\lfloor \log_2(x) \rfloor$

(18)

Rehooked over lecture notes

The binary says where to square

$$b = 101 \dots$$
$$= 2^k + 2^{k+2} + \dots$$

Quadratic residues modulo p

We say that a is a quadratic residue (square) mod p if

exists $x \in \mathbb{Z}_p^*$ such that $a \equiv x^2 \pmod{p}$

Note that \mathbb{Z}_p^* has generator g

$x \in \mathbb{Z}_p^*$ as $x = g^y$

Set $\{g^1, \dots, g^{p-1}\}$ is $\{1, \dots, p-1\}$

the set of els of \mathbb{Z}_p^*

Deciding if a # is a quadratic residue mod p

Have a generator g

What els are quadratic residues?

Quadratic residue

if q is congruent to perfect square mod n

$$x^2 \equiv q \pmod{n}$$

if there exists an x s.t.

For a given n,

then a list of q.r. may be obtained by squaring 0, 1, ..., n-1

So for 5 q.r. are 1, 4

7 1, 4, 5, 9, 16

18 1, 4, 7, 9, 10, 13, 16

Since $9^2 \pmod{18} = 9$

think this is a coincidence

(2)

We know the set is $\{g^1, \dots, g^{p-1}\}$

Thus the set of squares is $g^2, g^4, \dots, g^{2p-2}$

But are counting els twice

So $g^2, g^4, \dots, g^{p-3}, g^{p-1} = g^0$

So are even powers of g

(Where is this leading?)

Eulers' criterion The el a is square mod p iff $a^{\frac{p-1}{2a}} \equiv 1 \pmod{p}$

Moreover a is not a square mod p iff $a^{\frac{p-1}{2}} \equiv -1 \pmod{p}$

(21)

Let me go through lecture notes again

Is "a" a square mod p^r

binary search

all generators distinct

Any el to power of group size $(p-1) = 1$

ex is g^{2i} square

g^i is in square


so $(g^i)^2$ is as well

$$\left(g^{\frac{p-1}{2}}\right)^2 = g^{p-1} = 1$$

oh thats where that comes from!

~~oh thats~~


(22)



Any quad has exactly 2 roots

g^{2i+1}
↑ is square

$g^{\text{odd \#}}$ could be 2



← it sq has 2 roots

1 10 p-1

can't be root of 2 diff groups

This is totally confusing!

(23)

Recitation

Given a prime, find a generator

$$\exists g \text{ s.t. } \{g, g^2, \dots, g^{p-1}\} = \mathbb{Z}_p^*$$

~~#####~~

ie a \mathbb{Z}_7^* is generated by 3

$$\begin{array}{l}
 3 = 3 \\
 3^2 = 2 \\
 3^3 = 6 \\
 3^4 = 4 \\
 3^5 = 5 \\
 3^6 = 1
 \end{array}$$

kinda proof

$$g^{p-1} = 1 \text{ for any } g$$

Then a bunch of iterations of code

↳ simplest try all $2 \rightarrow p$

$O(n^2)$ time

This has nothing to do w/ rest of G.O.O.F

(24)

Generators mod p are randomly evenly distributed

↳ One every $\log n$ #'s

$$\# \text{ generators} \approx \frac{p}{\log p}$$

$$= \Omega\left(\frac{p}{\log p}\right)$$

(I need to remember symbol)

Symbols

Θ tight bound



c_1 is a Θ constant

$$f(n) = \Theta(g(n))$$

O upper bound

Ω lower bound

o upper bound - not asy tight

ω lower bound - not asy tight

25

Lecture

Non square mod p iff $a \equiv g^{2i+1}$

How to decide if square or not?


1, Find generator

2, Find discrete log that gives exponent

3, Even or odd?

How to efficiently find a generator?

DLP gives $g, p, e \in \mathbb{Z}_p^*$



But very hard to actually solve

↳ that means we don't need to know?

26

(still really confused...!)

In recitation again - pick a random #
- test if generator

$p = \text{prime}$

$g = \text{gen mod } p$

Given a we can compute $g^a \text{ mod } p$

But given g^a we can't compute a

↳ I get this → st. crypto stuff

We can't easily find a g for a given p

But can compute a g and p at same time

↳ Don't need to know

Crypto uses this

27

A tells B a prime p, g

A picks a , announces g^a

B " b " " g^b

~~$g^{ab} = g^{ba}$~~

$$g^{ab} = (g^b)^a = (g^a)^b$$

[Diffie-Hellman]

What g^a will they ask here?

Lecture: NP completeness

tractable = NP complete

= solvable in poly. time

intractable can't solve in polynomial

28

Next lecture

This was this the "weird" lecture
Where he had a fn called weird
Testing if something is composite or not

Think going to start from top

(skipping basic stuff) -

- Vector space model
- Peak finding

Core concepts


- Linked data structures
- Divide + conquer
- Hashing
- Sorting
- Graphs
- DP
- number theory

29

1D Peak finding

never scan whole list

Or compare middle w/ neighbors
go that direction



$$T(n) = T(n/2) + \theta(1)$$

$$= \underbrace{\theta(1) + \dots + \theta(1)}_{\log n \text{ times}}$$

? I understand where this comes from much better!

2D Recycle 1D

By doing each col 'ind.

30

Or (exact details NOT important!)

$$\begin{aligned}T(n, m) &= T(n, m/2) + \Theta(n) \\ &= \underbrace{\Theta(n) + \dots}_{\log m \text{ times}} \\ &= \Theta(n \log m)\end{aligned}$$

Master Method

$$T(n) = a T(n/b) + f(n)$$

1. $f(n) = O(n^{\log_b a - \epsilon})$

then $T(n) = \Theta(n^{\log_b a})$

2. $f(n) = \Theta(n^{\log_b a})$

then $T(n) = \Theta(n^{\log_b a} \log n)$

3. $f(n) = \Omega(n^{\log_b a + \epsilon})$

(This is the hard way ...)

31


Continue lectures - then find when see it

Linked List



blocks

BST



\leq left

\geq right

? note can be =

insert, etc all all $\log n$ avg case

n worst case

h the height actually

32

insert pretty straight forward

just walk down tree


till get to end of branch

is it always at end? Yes - that's the code

delete = very complex

red-black - a form of balanced

- makes sure



are = $\pm n-1$

is correct

- forget exact balance mech, somewhat complicated

Goal for answering our binary problem

↳ everything \leq or \geq than a tree
w/ augment of # of subtrees

Back to master theorem

$$T(n) = a T\left(\frac{n}{b}\right) + f(n)$$

(same notation as before...)
thought it was diff

$a \geq 1$
branch factor

$b > 1$
how much to
shrink on recurrence

$$a f\left(\frac{n}{b}\right) \leq c f(n)$$

In this class balanced = ~~AVL~~ AVL
is ± 1 allowable diff invariant
height = $\log n$

34

there is the easy Master Theorem

$$n^c \lg_b a$$

$$n^c \cdot (n^d)$$

1. Compare "Cs"

If "c"s diff - bigger c wins
Polynomially diff

$$\text{eg } n(\lg n)^3 \gg \sqrt{n} (\lg n)^3$$
$$n^1 \gg n^{1/2}$$

2. Next compare d
only slightly diff

~~ans $n^1 < n^{10}$~~ to recurrence

$$n(\lg n)^3 < n (\lg n)^{10}$$

35

Case 1 if sig diff
ans is bigger one

Case 2 slight diff
ans to recurrence is bigger one $\circ \log n$

Oh so what are we comparing

$a T(\frac{n}{a})$
recursive routine

$O(f(n))$
top level routine

$n \log_b a$

$f(n)$

eg

$$T(n) = T(\frac{n}{2}) + O(1)$$

$$n \log_2 1$$

$$\log 1 = 0$$

$$n^0 = 1 \quad \text{vs} \quad 1$$

same, so $1 \circ \log n = \log n$

36

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

$$n \lg_2 a$$

$$\text{So } n \lg_2 1$$

$$n^0$$

$$1 \quad \text{vs } n \quad \tau_{\text{bigger}} = n \quad \textcircled{D}$$

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

$$n \lg_2 2$$

$$n^1$$

$$n \quad \text{vs } n \quad \text{in } \lg n \quad \textcircled{D}$$

$$T(n) = 25 T\left(\frac{n}{2}\right) + n^2 \lg n^3$$

$$n \lg_2 25$$

$$n^4 \text{, } \text{bigger}$$

τ_{bigger}

$$n^2$$

(37)

& there not done that in a while!

Oops copy error

$$25 T\left(\frac{n}{3}\right) + n^2 \lg^3 n$$

$$n^2 \quad \text{vs} \quad n^2 \lg^3 n$$

So compare d
bigger one is $\lg n$

$$n^2 \lg^4 n$$

Oh didn't divide

$$\frac{n^2}{\lg^3 n}$$

So like

$$d=0 \quad d=-3$$

↑ here

$$n^2 \lg n \quad \text{⓪}$$

Shipping by wisdom

(38)

Dictionaries

insert

delete

switch

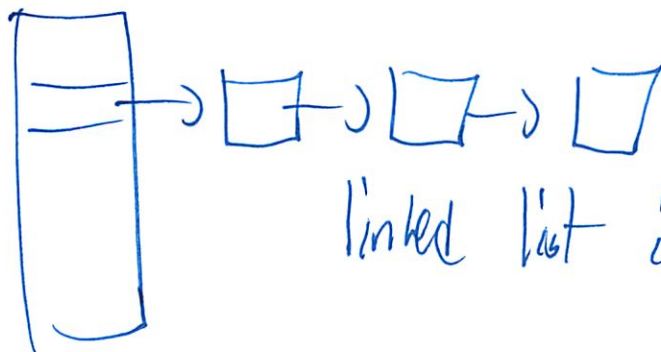
powered by hashing

Hashing

$h: U \rightarrow \{1, \dots, m\}$

puts stuff in buckets

→ so insert/find is \sim length of key
= time to hash



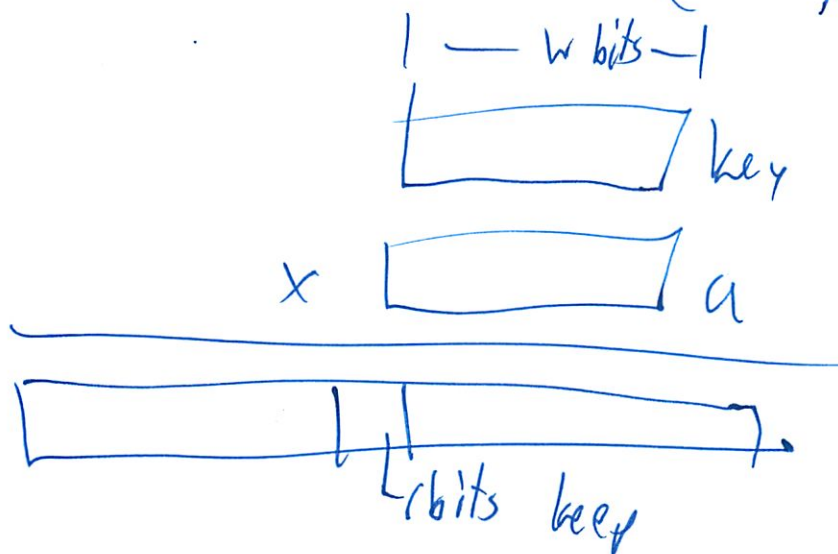
linked list if collision

39

Division Hash Fn

$$h(k) = k \bmod m$$

Collide if $k_1 = k_2 \pmod{m}$



(It seems that the actual hashing alg is a bit hardware)

p should be prime ...
for poly hash

$\frac{n}{p}$ collision prob

40

Rolling Hash (Rabin-Karp)

Best for seeing if P is a substring of S

P has len L

S has len n

1. Hash P $O(L)$

2. Iterate through S $O(nL)$

3. Compare to verify no collision

Initial

$$h(k) = (k[0] \cdot 10^4 + k[1] \cdot 10^3 + k[2] \cdot 10^2 + k[3] \cdot 10^1 + k[4] \cdot 10^0) \text{ mod } m$$

Then next

$$h(S_{i+1}) = [(h(S_i) - (10^5 \cdot \text{first digit } S_i)) \cdot 10 + \text{next digit after } S_i] \text{ mod } m$$

1. Subtract out last
2. Multiply all by 10 (mod of course)
3. Add on last digit

41

Resize when $n \geq m$

↳ Amortized analysis - avg over all cases

Open addressing - probe till find empty

STUA - that hashing result is uniform

$$\left[\right] = P(\text{collision})$$

~~Uniform~~

Universal hashing family - keyed hash?

↳ seems like it since pick a hash from a family of hashes

(42)

Sorting

want $A[1] \leq A[2] \leq \dots \leq A[n]$

~~Insertion~~


Insertion Sort

(the one at start of text book)

as go through deck of cards, insert
1 at a time to right place

$O(n^2)$

Merge Sort



split in half
+ half again

then reassemble in order

$O(n \log n)$

43

Data structure: Priority queue

insert (S, x)
↳ x into S

~~increase~~ max

implemented w/ Heap

Heap

Actually an array

But shows as a binary tree

~~Max~~

min-heap property $A[\text{parent}(i)] \leq A[i]$

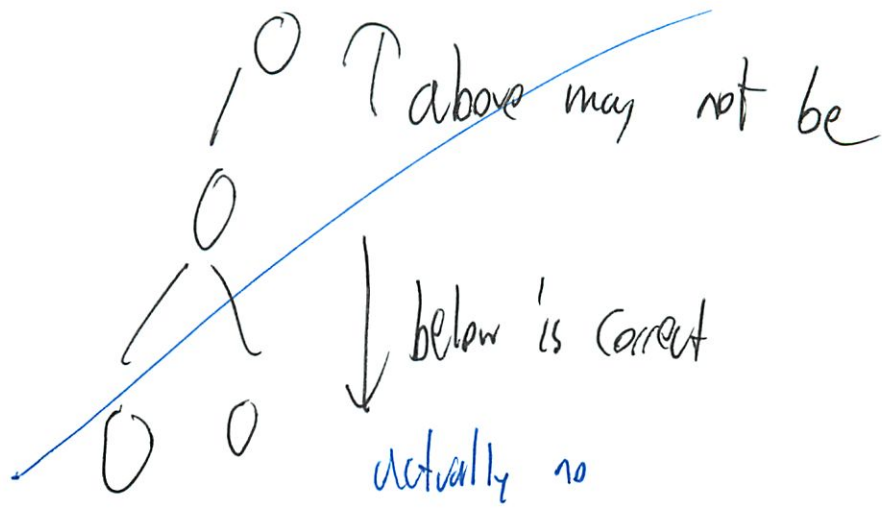
height = $\lg n$

Max heapify $O(\lg n)$ maintains property

Assumes $\text{left}(i)$ and $\text{right}(i)$ are max heaps
but $A[i]$ might be smaller than children

(114)

Basically



lets $A[i]$ float down

↓ moves it down

for a specific i

Build max heapify — does max heapify on all

$O(n \lg n)$

Heapsort

Sorts it in order by removing items (from end)

Remember heap sort is only

sorted max to the level

↳ within level arbitrary

45

So it removes max
then calls max heapify on spot 1
repeats till all items in array

$$O(n \lg n)$$

$$\text{parent} = \left\lfloor \frac{i}{2} \right\rfloor \quad \begin{array}{l} \text{left} = 2i \\ \text{right} = 2i + 1 \end{array}$$

Quicksort


Randomized $\rightarrow O(n \lg n)$

picks a pivot



divide + conquer

quicksort (A, p, r)



46



So slowly (l by l) sorts to < pivot or > pivot

Then recurse on that remaining group

Can be ~~done~~

Balanced \rightarrow merge sort $O(n \lg n)$

Unbal \rightarrow insertion sort $O(n^2)$



Randomize - to get good over all inputs

randomly picks a pivot

$O(n \lg n)$

47

Linear time sorts

Comparison sort

The algo we saw before

$O(n \lg n)$

decision tree model

go right or left based on comparisons

so no better than $O(n \lg n)$ worst case

Counting sort

Count the # of occurrences of int

$O(n+k)$

↑ # of possibilities

is stable w/ sublight data

↑ same order comes out (How does this work?)

(48)

Radix Sort

Card machines

most sig ~~digit~~ digit list

Graphs, Representations, Search

Adj list

- for each vertex, list neighbors

incidence list ^{- good for sparse graphs}

- for every vertex list edges

adj matrix

vertices $\left\{ \begin{array}{c} \text{vertices} \\ 1 \text{ if edge} \end{array} \right.$

good for ^{calculating} $A^2 = \# \text{ length } 2 \text{ paths}$

- good for dense graphs

(49)

Breadth First

0

Greedy

Coloring
queue

$O(n+m)$

Depth First



optimistic

like exploring a maze

~~Start~~ coloring as well

↳ start + finish times

Parentheses structure

Stack

$O(n+m)$

Types of Edges

1. Tree edges

- ↳ regular
- ↳ in DFS forest

2. Back edge

↳ reach back to an ancestor


3. Forward edge

Connect to a descendant (non tree)
(from an edge other than the one that was explored)

4. Cross edges

all others

directed only



nodes explored in

(51)

Topo sort

Call DFS

Record in order finished

For a DAG

Strongly Connected Components

all vertices are reachable from all other in the group

Spanning tree

build layers



can't skip

but can do back

Topo Revers

if top sort then reverse

from DFS

lets us find sinks

(52)

Dijkstra

Single source shortest path

Weighted graph

↳ why you have to vs DFS

* all edge non neg *

↳ why you can vs B-F

extract min from queue

↳ ~~greedy~~ greedy Ⓛ

Then relax all edges from that

↳ see if we can update shortest path estimates

Running time depends on how queue implemented

Remember having a lot of trouble w/ this

$O((V+E) \lg V)$ binary min-heap

$O(V \lg V + E)$ Fib heap

Should better understand how put together

~~42~~ 53

Bellman Ford

Weighted graph

Allow \ominus

for ~~every~~ $\#$ of vertices

for every edge \leftarrow in a specific order

relax

neg weight cycles \rightarrow return false

$O(mn)$ VE

DAG

Topo Sort

Go through once along the order

$\Theta(V+E)$

54

Heuristics

Google Maps is too slow w/ normal D_{ij}

So bi-directional - start from each side
pick shortest from both

Planner

use as cross fly distance

A*

has a potential $h()$

"aka" guess
* heuristic

use a consistent heuristic

Can bi-dir A* - hard

Dynamic Programming

Fib as example

Basically save results (memoization) to sub problems

And use ~~that~~ these results to make subproblems easier to solve

Is both a top-down and bottom-up way to visualize

All-pairs shortest path

output is usually a table

could run D_{ij} for each node

L's stupid $O(V^3)$

or Bellman Ford
 $O(V^4)$

Uses adj matrix

also compute a predecessor matrix

$$\Pi = \pi_{ij}$$

50

DP approach to this

- each step is very similar to matrix mult.

Remember the 4 steps

1. Char. the optimal structure
2. Recursively define the value of an optimal sol
3. Compute value of optimal sol bottom up
4. Construct an optimal sol

I never really understood this →

So (Recursive sol)

$l_{ij}^{(m)}$ is min weight of any path i to j w/ at most m edges

So start w/

$$l_{ij}^0 = \begin{cases} 0 & \text{if } i=j \\ \infty & \text{if } i \neq j \end{cases}$$

$$\text{Then } l_{ij}^m = \min_k l_{ik}^{(m-1)} \text{ and } l_{kj}^m$$

(from looking at all possible predecessors)

(lengthy notation!

57

$$\text{So } l_{ij}^m = \min \left(l_{ij}^{m-1}, \min_{1 \leq k \leq n} \{ l_{ik}^{(m-1)} + w_{kj} \} \right)$$

$$= \min_{1 \leq k \leq n} \{ l_{ik}^{(m-1)} + w_{kj} \}$$

(Since $w_{jj} = 0$ for all j)

Basically grow one at a time
look back and find the shortest

$J(i, j)$ = actually shortest path weights

Can't have more than $n-1$ int. nodes

So now compute some matrices L^1, L^2, L^3

$$L^m = l_{ij}^m$$

L^{n-1} is final, has shortest path weights

L^1 is the basic one = W

(58)

key is Extend - Shortest - Paths (L, W)

for $i \rightarrow 1$ to n

$j \rightarrow 1$ to n

$k \rightarrow 1$ to n

$$d_{ij} = \min(d_{ij}, d_{ik} + w_{kj})$$

$\Theta(n^3)$ Then do this for $m \rightarrow 2$ to $n-1$

→ Can see this is just matrix multiplication

And don't need each $m \rightarrow$ just $n-1$

So repeated squaring trick

$$\Theta(n^3 \lg n)$$

Floyd ~~W~~ Warshall

diff DP formulation

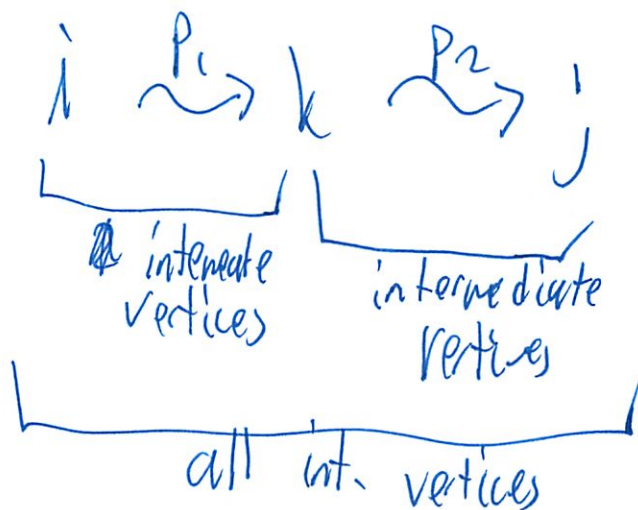
$$\Theta(V^3)$$

59

neg weights allowed - but not neg cycles

Considers the intermediate vertices of a shortest path

So containing set stuff



Recursive sol

same as before

$$d_{ij}^0 = w_{ij}$$

$$d_{ij}^k = \begin{cases} w_{ij} & k=0 \\ \min(d_{ij}^{k-1}, d_{ik}^{k-1} + d_{kj}^{k-1}) & k \geq 1 \end{cases}$$

code for bottom is

↳ seems the same

66
So the diff seems to be the introduction
of an "intermediate" vertex

Then ~~can~~ for actually printing shortest path
w/ Π the predecessor matrix

$$\Pi_{ij}^0 = \begin{cases} \text{nil} & \text{if } i=j \\ \lambda & \text{if } i \neq j \end{cases} \quad \begin{matrix} w_{ij} = \infty \\ w_{ij} < \infty \end{matrix}$$

$$\Pi_{ij}^k = \begin{cases} \Pi_{ij}^{k-1} & \text{if } d_{ij}^{k-1} \leq d_{ik}^{k-1} + d_{kj}^{k-1} \\ \Pi_{kj}^{k-1} & \text{if } d_{ij}^{k-1} > d_{ik}^{k-1} + d_{kj}^{k-1} \end{cases}$$

Transitive closure

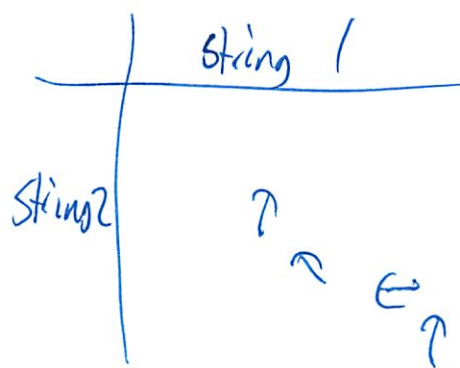
$F^* = \{(i,j) \mid \text{there is path from } i \text{ to } j \text{ in } G\}$

Proof or a field now

(61)

Notes have that top down vs bottom up perception thing
↳ save a lot by memoizing previous ans

LCS b/w 2 strings can use that table



knapsack

not possible for arbitrary size

~~answer~~

given a set of items each w/ weight + value
find best combo so weight < limit and
value largest

Linear programming?

WP

(0 or 1 of each item
or up to x_i of each
or unbounded

(62)

Or Text doc layout function

or vertex cover

What was the point of this?

Then back to Number Theory

Also done

5/12

Final Exam

- Do not open this exam booklet until directed to do so. Read all the instructions on this page.
- When the exam begins, write your name on every page of this exam booklet.
- You have 180 minutes to earn **180** points. Do not spend too much time on any one problem. Read them all through first, and attack them in the order that allows you to make the most progress.
- This exam booklet contains 20 pages, including this one. Two extra sheets of scratch paper are attached. Please detach them before turning in your exam at the end of the exam period.
- This exam is closed book. You may use **three** handwritten, $8\frac{1}{2}'' \times 11''$ or A4 crib sheets (both sides). No calculators or programmable devices are permitted. No cell phones or other communications devices are permitted.
- Write your solutions in the space provided. If you need more space, write on the back of the sheet containing the problem. Pages may be separated for grading.
- Do not waste time and paper rederiving facts that we have studied. It is sufficient to cite known results.
- Show your work, as partial credit will be given. You will be graded not only on the correctness of your answer, but also on the clarity with which you express it. Be neat.
- Good luck!

Problem	Parts	Points	Grade	Grader	Problem	Parts	Points	Grade	Grader
1	10	30			7	-	10		
2	8	40			8	-	10		
3	3	15			9	-	15		
4	-	10			10	-	10		
5	-	10			11	-	20		
6	-	10			Total		180		

Name: _____

Athena username: _____

Recitation: Nick Nick Tianren David Joe Joe Michael
 WF10 WF11 WF12 WF1 WF2 WF3a WF3b

(Handwritten notes in red and blue ink, including names like "Nick", "Tianren", "David", "Joe", "Michael" and some illegible scribbles)

Problem 1. True or false [30 points] (10 parts)

For each of the following questions, circle either T (True) or F (False). **Explain your choice.** (Your explanation is worth more than your choice of true or false.)

(a) T F For all positive $f(n)$, $f(n) + o(f(n)) = \Theta(f(n))$.

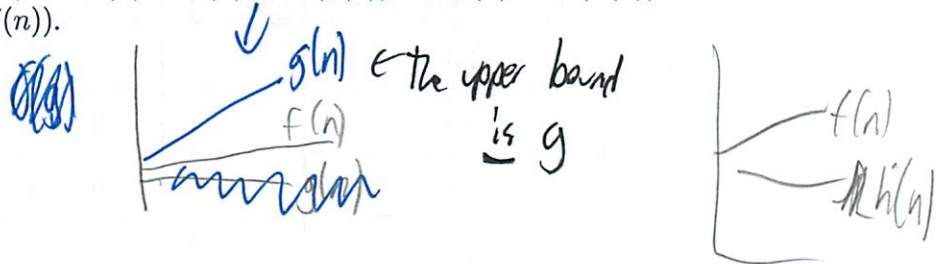
~~One bound~~

Θ tight
 O upper
 Ω lower
 o upper - not asy
 ω - lower not asy

No - doesn't make any sense

True
(but not explained)

(b) T F For all positive $f(n)$, $g(n)$ and $h(n)$, if $f(n) = O(g(n))$ and $f(n) = \Omega(h(n))$, then $g(n) + h(n) = \Omega(f(n))$.



So if add - since asy like $\frac{1}{2} + 2 = 1$
 basically < 1 and > 1 are lower
 don't think so

$$f(n) = O(g(n)) \rightarrow g(n) = \Omega(f(n))$$

still don't get ...

- (c) **T** Under the simple uniform hashing assumption, the probability that three specific data elements (say 1, 2 and 3) hash to the same slot (i.e., $h(1) = h(2) = h(3)$) is $1/m^3$, where m is a number of buckets.

So prob multiples $\frac{1}{m} \cdot \frac{1}{m} \cdot \frac{1}{m}$

Only for fixed bucket ~~and~~
(a particular bucket)

$\frac{1}{m^2}$ ← how did they get this?

- (d) **T** Given an array of n integers, each belonging to $\{-1, 0, 1\}$, we can sort the array in $O(n)$ time in the worst case.

? is -1 or 0 or 1?

Yes → Counting Sort

① $O(n+k)$ $k=3$

add 1 to all #s

Why? I think it would work

otherwise -- depends on specific Alg

- (e) **T** **F** The following array is a max heap: $[10, 3, 5, 1, 4, 2]$.

Every thing is smaller than parents

x No 3, 4

- (f) **T** **F** RADIX SORT does not work correctly (i.e., does not produce the correct output) if we sort each individual digit using INSERTION SORT instead of COUNTING SORT.

~~Does~~ Insertion sort is stable it is
- I guess it could be quite easily

- (g) T F Given a directed graph G , consider forming a graph G' as follows. Each vertex $u' \in G'$ represents a strongly connected component (SCC) of G . There is an edge (u', v') in G' if there is an edge in G from the SCC corresponding to u' to the SCC corresponding to v' .
Then G' is a directed acyclic graph.

So SCC is collapsed to a node

I believe you can do that



- (h) T F Consider two positively weighted graphs $G = (V, E, w)$ and $G' = (V, E, w')$ with the same vertices V and edges E such that, for any edge $e \in E$, we have $w'(e) = w(e)^2$. For any two vertices $u, v \in V$, any shortest path between u and v in G' is also a shortest path in G .

Yeah doesn't change anything

Since everything $\oplus > 0$

Oh paths not of = length

Duh - think!! I knew that!

- (i) **T** **F** An optimal solution to a knapsack problem will always contain the object i with the greatest value-to-cost ratio v_i/c_i .

Wait on WP was a height cutoff
could be above that

Was that in our knapsack problem?

✓

Diff reason

Greedy doesn't work
if max weight = 2

$\frac{2}{1}$ vs $\frac{3}{2}$ ← pick

(can you have multiple)

- (j) **T** **F** Every problem in NP can be solved in exponential time.

Forgot the def

I think NP is exp + ~~unsolvable~~


Those were kinda hit or miss thing

Problem 2. Short answer [40 points] (8 parts)

oh I dislike these

- (a) Rank the following functions by increasing order of growth; that is, find an arrangement g_1, g_2, g_3, g_4 of the functions satisfying $g_1 = O(g_2)$, $g_2 = O(g_3)$, $g_3 = O(g_4)$. (For example, the correct ordering of n^2, n^4, n, n^3 is n, n^2, n^3, n^4 .)

$f_1 = (n!)^{1/n}$ $f_2 = \log n^n$ $f_3 = n^{n^{1/2}}$ $f_4 = n \log n \log \log n$



- (b) Solve the following recurrences by giving tight Θ -notation bounds. You do not need to justify your answers, but any justification that you provide will help when assigning partial credit.

- i. $T(n) = 4T(n/2) + n^2 \log n$
- ii. $T(n) = 8T(n/2) + n \log n$
- iii. $T(n) = \sqrt{6006} \cdot T(n/2) + n^{\sqrt{6006}}$

Master theorem

$2^2 = 4$
 1. $n^{\log_2 4}$ vs $n^2 \log n \rightarrow n^2 \log^2 n$ ✓

$2^3 = 8$
 2. $n^{\log_2 8}$ vs $n \log n \rightarrow n^3$ ✓

3. $n^{\log_2 \sqrt{6006}}$ vs $n^{\sqrt{6006}}$
 $n^{6.27}$ vs n^{77} = $n^{6.27}$ ✓

How supposed to do w/o calc I don't get why!

(c) Give a recurrence $T(n) = \dots$ for the running time of each of the following algorithms, along with the asymptotic solution to that recurrence:

- i. Insertion sort
- ii. Merge sort
- iii. 2D peak finding (the fastest algorithm we've seen)

oh

forget to solve

Insertion sort $T(n) = T(n-1) + O(n) = O(n^2)$

Merge $T(n) = 2T(\frac{n}{2}) + O(n) = O(n \lg n)$

2D $T(n) = T(\frac{n}{2}) + O(n) = O(n)$

Review these

(that was in the notes!

the final one was never given!

(d) Could a binary search tree be built using $O(n \lg n)$ comparisons in the comparison model? Explain why or why not.

the bound not as tight

Yes - isn't this a red/black and AVL tree

No or can sort in $O(n \lg n)$ time by building a BST and then doing in-order tree walk in $O(n)$

↳ But a ^{AVL} BST is $O(\lg n)$ for insertion

I don't get ---

(g) In dynamic programming, we derive a recurrence relation for the solution to one subproblem in terms of solutions to other subproblems. To turn this relation into a bottom-up dynamic programming algorithm, we need an order to fill in the solution cells in a table, such that all needed subproblems are solved before solving a subproblem. For each of the following relations, give such a valid traversal order, or if no traversal order is possible for the given relation, briefly justify why.

- i. $A(i, j) = F(A(i, j - 1), A(i - 1, j - 1), A(i - 1, j + 1))$
- ii. $A(i, j) = F(A(\min\{i, j\} - 1, \min\{i, j\} - 1), A(\max\{i, j\} - 1, \max\{i, j\} - 1))$
- iii. $A(i, j) = F(A(i - 2, j - 2), A(i + 2, j + 2))$

Solve smallest last



or consider on format

Solve $A(i, j)$ for i from 0 to n
for j from 0 to n

Extra ①

(h) Consider an array $A[1 \dots n]$ of integers in the range $1 \dots n^2$. A number a is a **heavy hitter** in A if a occurs in A at least $n/2$ times.

Give an efficient algorithm that finds all heavy hitters in a given array A .

reminds me of parking wars

Counting sort style list of integers

Count the # of appearances

If $> n/2$, add to result

$$O(n+n^2) = O(n^2)$$

Radix sort + linear scan

Oh yeah that's better

I should have thought harder...

- (e) Given n integers in the range $0 \dots k$, describe how to preprocess these integers into a data structure that can answer the following query in $O(1)$ time: given two integers a and b , how many integers fall within the range $a \dots b$?

BST augmented

- how many #s $<$

- then find augment b - augment a ✓

Same as counting sort they say

I think mine works

- (f) Describe how any comparison-based sorting algorithm can be made stable, without affecting the running time by more than a constant factor.

When found the right place, append at the end of that list of #s

Or augment to be decimal 7.1

7.2

or 2 part #s ✓

Extra
①

ii)

$$\min\{i, j\} - 1 \quad \min\{i, j\} - 1$$

$$\max\{i, j\} - 1 \quad \max\{i, j\}$$

Solve $A(k, k)$ for k 0 to n

Solve rest in any order

I don't get this at all...

iii) $A(i, i) = F(A(i-2, j-2), A(i+2, j+2))$

Impossible
cycle

Problem 3. You are the computer [15 points] (3 parts)

- (a) Fill in the following grid with the correct subproblem solutions for this sequence alignment problem with these weights: 0 for mutation, 1 for insertion or deletion, and 3 for a match (the goal is to maximize the sum of the weights). Here "ATC" is the starting sequence and "TCAG" is the ending sequence.

What is the problem?

-	-	A	T	C
-	0	1	2	3
T	1	2	4	5
C	2	3	5	7
A	3	5	6	8
G	4	6	7	9

What is the optimal alignment?

- (b) Draw a max-heap on the following set of integers: {2, 3, 5, 7, 11, 13, 17}. (You do not need to use the Build-Heap algorithm.)

17
13 11
7 5 3 2

no ans - but said almost everyone got this

- (c) Compute $\sqrt[3]{6006}$ using two iterations of Newton's method, i.e., fill out the following table. Your entry for x_1 should be fully simplified. Your entry for x_2 can be left unsimplified.

i	x_i
0	1
1	
2	

this test
seems diff
from the typical
ones here

skipping

Problem 4. Rotated array [10 points]

Consider an array $A[1 \dots n]$ constructed by the following process: we start with n distinct elements, sort them, and then rotate the array k steps to the right. For example, we might start with the sorted array $[1, 4, 5, 9, 10]$, and rotate it right by $k = 3$ steps to get $[5, 9, 10, 1, 4]$. Give an $O(\log n)$ -time algorithm that finds and returns the position of a given element x in array A , or returns None if x is not in A . Your algorithm is given the array $A[1 \dots n]$ but does not know k .

here is
the tricky
q

Nieve Find min and find its pos $O(n)$
 Lbt not $\log n$
 essentially max heapify \uparrow ?
 call max heapify on 1
 see what pos the original # is
 then have k
 return $x + k$

Sol 1 - modified binary search
 2 - binary search on left and right
 find smallest

↑ seems very similar

fol

Problem 5. Taxachusetts [10 points]

Suppose you are given a weighted graph $G = (V, E, w)$ of highways, and the state government has implemented a new tax rule whereby the cost of a path gets doubled as penalty if the number of edges in the path is greater than 10. Explain how to reduce finding the shortest-path weight between every pair of vertices (under this penalty) to the usual all-pairs shortest paths problem (as solved by Floyd-Warshall).

$$d^{10} = d^9 \circ 2$$

↑
as normal

just change the d^{10} step to double the cost of the d^9 step path

or 10 is the # of nodes explored - we want actual - so not d^{10} but when $d^n = 10 \rightarrow$ double it

anything more complex

1. Must augment to track path lengths

↑ ah forgot that detail

d^n tracks weight sum not # edges totally missed

2. Must track weights for < 10 edges

> 10 edges Why? Oh in case a q is \leftarrow for higher weight w/o double - I see

Problem 6. Does this path make me look fat? [10 points]

Consider a connected weighted directed graph $G = (V, E, w)$. Define the fatness of a path P to be the maximum weight of any edge in P . Give an efficient algorithm that, given such a graph and two vertices $u, v \in V$, finds the minimum possible fatness of a path from u to v in G .

$$\text{fatness} = \max \text{ weight}$$

Can Dij - but track fatness

Can we do just by tracking fatness

No we need to track both weight
fatness

Select by min fatness

but add to weight

Re calc fatness



Don't forget correctness argument

Plus note neg weight edges normally break
Dij - but not here!

Problem 7. Indiana Jones and the Temple of Algorithms [10 points]

While exploring an ancient temple, Prof. Jones comes across a locked door. In front of this door are two pedestals, and n blocks each labeled with its positive integer weight. The sum of the weights of the blocks is W . In order to open the door, Prof. Jones needs to put every block on one of the two pedestals. However, if the difference in the sum of the weights of the blocks on each pedestal is too large, the door will not open.

Devise an algorithm that Prof. Jones can use to divide the blocks into two piles whose total weights are as close as possible. To avoid a boulder rolling quickly toward him, Prof. Jones needs your algorithm to run in pseudopolynomial time.

Can he pick ind blocks?
or are they ordered?

Like a balance tree
but w/ weights

Goal $\rightarrow w/2$

Add 1 then add closest to diff $\frac{w}{2}$

1. Random on left

2. $\frac{w}{2}$ -left on right \leftarrow closest

3. Then again try to get close

greedy - prob not best
but how to avoid?

Smallest pile is close to $\frac{w}{2}$ w/o going over
DP! - not comfortable w/ knapsack - style problems

extra
2

Extra 2

Subproblem $P(i, w)$

↳ true if pile $w \leq W/2$ using subset $0 \rightarrow i$
false otherwise

Initially $P(0, 0)$ true

$P(0, w)$ false for all $w > 0$

Recurrence $P(i, w) = P(i-1, w) \vee P(i-1, w - w_i)$

What is this notation?

i = blocks subset 0 to i

w is target weight

So $P(i, w)$ is can we (T or F) make a group of blocks that is $w \leq W$
~~0 to i~~
up to

So that fn is one less block or one less block and remove its weight

So this is work backwards (but really a time)

Where do we start from?

Extra
3

$P(i-1, w-w_i)$ is false if $w_i > W$
(so if that block is more than weight?)
or ~~if~~ if it goes 0 - makes sense

If $P(i, w)$ is true, record in table
Can d which is true

Find largest w for which $P(n, w)$ is true
all blocks

Then use B tables to find which blocks are in

 $O(nW)$ subproblems $O(1)$ each = $O(nW)$
Also for every weight \rightarrow every block

So ~~left = right~~

Still don't get where you start from?

So use that formula for top down procedure

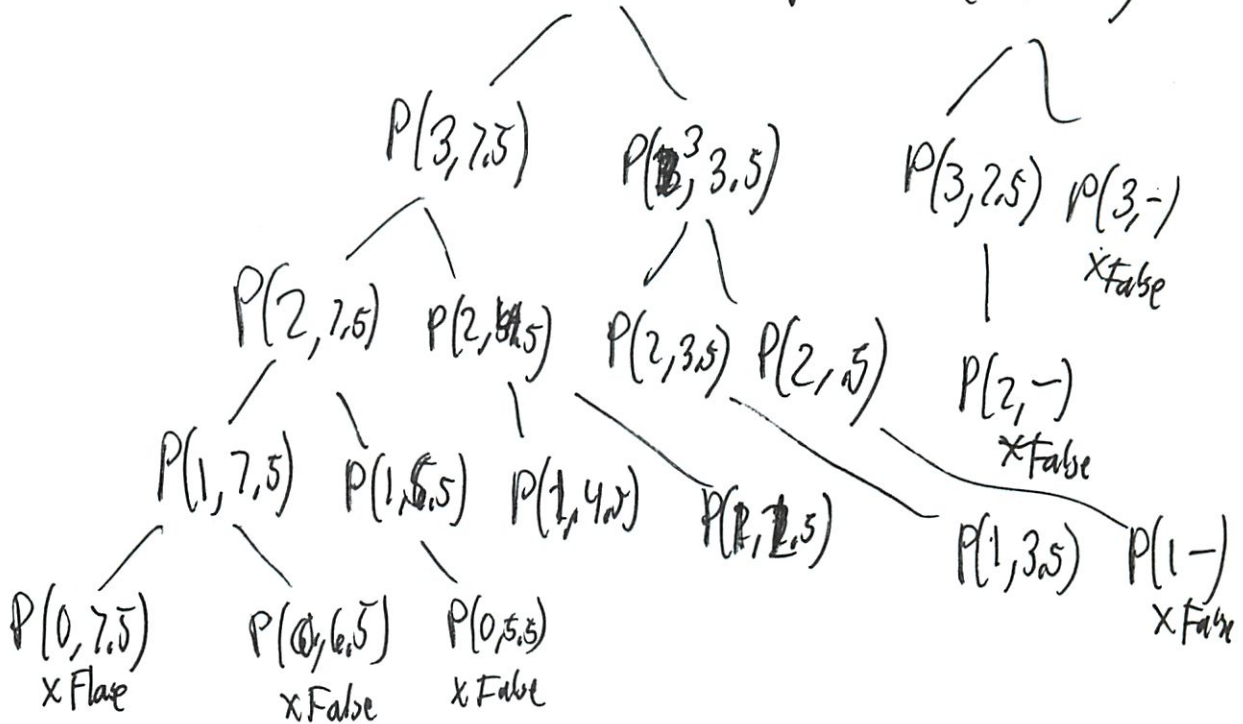
Extra
4

Say 5 blocks

i	1	2	3	4	5
w	1	2	3	4	5

~~weight~~ weight
 so $w = 15$
~~half~~ half = 7.5

Call $P(5, 7.5) = P(4, 7.5) \text{ or } P(4, 2.5)$



but only true at $P(0, 0)$ - so False

So need to try every weight, fill tree

Where is HW - just the tree - no that is $2^{(n+1)} + O(1)$
 $n \lg 2$
 I think

Still don't really get - but got FD vs BU in DP
 so 1 - but that is wrong

Extra
5

Yeah don't get recurrence / master theorem on DP
problems

— was on a hw

That was how many subproblems

Problem 8. Claustrophobic chickens [10 points]

Prof. Tyson has noticed that the chickens in his farm frequently get claustrophobic. He wants to build a monitoring system that will alert him when this happens, allowing him to manually rearrange the chickens. After thorough research, Prof. Tyson determines that a chicken becomes claustrophobic if it is within 2 feet of at least 8 other chickens.

Prof. Tyson has installed a tracker on each chicken, which reports the (x, y) coordinates (measured in feet) of each of his chickens. Suppose that there are n chickens whose locations are represented as a sequence $(x_1, y_1), \dots, (x_n, y_n)$. Prof. Tyson needs an efficient algorithm to determine whether there are any claustrophobic chickens (causing Prof. Tyson to go out and manually rearrange the chickens). Devise such an algorithm and save the chickens!

Oh boy!

Sort in order - DP

- or graph

There no clue how to start

2x2 grid

Nieve $\rightarrow n^2$ - check every chicken w/ each other

Must be on int point

- could do hash grid

Square grid $\sqrt{2} \times \sqrt{2}$ square

Hash point into bucket

If ≥ 9 points \rightarrow yes

Otherwise \rightarrow calc distance to all pts in cell
Count those points


Extra 6

$O(n)$

Problem 9. Architects 'R Us [15 points]

You are assisting Prof. Gehry with designing the shape of a new room in the Stata Center. The professor has given you n columns, each of the same unit thickness, but with different heights: $A[1], A[2], \dots, A[n]$. He asks you to permute the columns in a line to define the shape of the room. To make matters difficult, MIT wants to be able to hang a large rectangular picture on the columns. If j consecutive columns in your order all have a height of at least k , then we can hang a rectangle of size $j \cdot k$.

The example below contains 3 consecutive columns with heights of at least 4, so we can hang a rectangle of area 12 on the first three columns.



- (a) Give an efficient algorithm to find the largest area of a hangable rectangle for the *initial* order $A[1], A[2], \dots, A[n]$ of columns.

This is the same as the football players problem,
Rewrite for practice

No wait is it - I think it is
Sort highest to lowest $O(n \log n)$

Then say $\text{Max}(\text{cols } 1, 2, 3, 4 \text{ etc})$
but those don't include or not
recurse $O(n)$

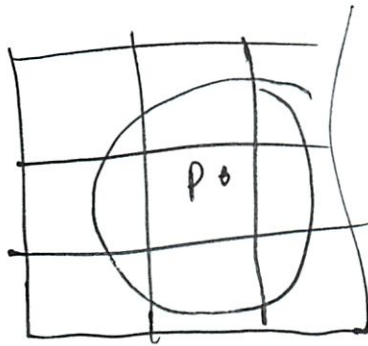
extra 7

Extra
6

So it does not have to be 'int'?

Since check all neighboring cells

- Calc distance to all neighboring cell



So this is not n^2 since
just checking 9 cells?

instead of each

I see how this works

Kinda seems like cheating - but

I guess it works ...

Extra 9

Simpler ~~Best~~

~~Biggest rectangle is banded by some col~~

Guess a col i $\leftarrow i$

Go left + right counting cols

How do we guess this?
Total DS

Other

$m[i, j]$ is min of interval $A[i], \dots, A[j]$


$$m[i, j] = \min \{ m[i, j-1], A[j] \}$$

Now sol is $\{ m[i, j] \cdot (j - i + 1) \cdot i \leq j$

So was what I had it?

Also review notes on 2D DP

it was



So is this a 2D one? - Don't think so

extra?

$O(n^2)$

worse than mine \rightarrow guess I was wrong

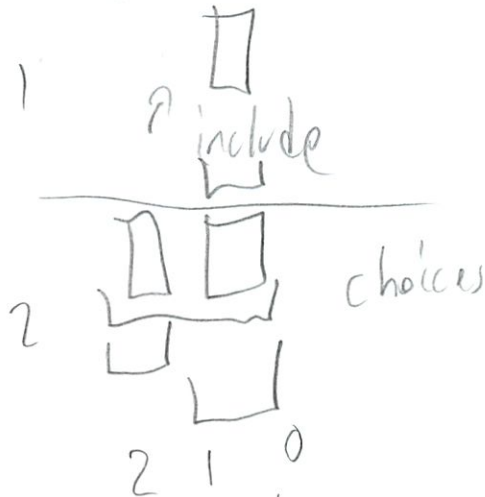
Ohhh can't rearrange cols!

Let me rethink this
now standard DP

$DP(n) = \max(\text{include or not})$
from i to n

= \max oh start can vary

Do from end



$$DP = \max(2 \rightarrow 1, 2 \rightarrow 0, 1 \rightarrow 0)$$

$$DP = \max$$

Extra 8



So 3 choices
 just that
 best of DP 2
 or all 3

$$DP(3) = \min(A[3], DP(2), A[3] + DP(2))$$

↑
from end

? what is this exactly
 (don't just add)
 So just return


~~ABA~~ square $[A[3] \cdot A[2], A[1]]$
 (find min
 multiply by 3)

? Can we DP that?
 Something about saving min
 ? so sep to save

↑
 but this
 is cost

- (b) Devise an efficient algorithm to permute the columns into an order that maximizes the area of a hangable rectangle.

So no we can move? Yes

Don't we want 

~~Just sort~~ ✓ sort $O(n \log n)$ ✓

wrote sol before

Compute hangable in $O(n)$

✓ So exactly as I had before


Problem 10. Guess Who? [10 points]

Woody the woodcutter will cut a given log of wood, at any place you choose, for a price equal to the length of the given log. Suppose you have a log of length L , marked to be cut in n different locations labeled $1, 2, \dots, n$. For simplicity, let indices 0 and $n + 1$ denote the left and right endpoints of the original log of length L . Let d_i denote the distance of mark i from the left end of the log, and assume that $0 = d_0 < d_1 < d_2 < \dots < d_n < d_{n+1} = L$. The **wood-cutting problem** is the problem of determining the sequence of cuts to the log that will cut the log at all the marked places and minimize your total payment. Give an efficient algorithm to solve this problem.

I don't get what it is asking...

But prob some sort of DP \circ

$$C(i, j) = \min_{i < k < j} \{ C(i, k) + C(k, j) + (d_j - d_i) \}$$



$O(n)$ each

$O(n^2)$ entries = $O(n^3)$

Greedy does not work

So it was the 2D - ~~pa~~ like the parenthesis q_v

Problem 11. Varying variance [20 points]

For a set S of numbers, define its **average** to be $\bar{S} = \frac{1}{|S|} \sum_{s \in S} s$, and its **variance** to be $V(S) = \frac{1}{|S|} \sum_{s \in S} (s - \bar{S})^2$.

A **segment variance data structure** supports the following operations on an array $A[1 \dots n]$:

- **Assignment:** given an index i and a value a , set $A[i] = a$.
- **Segment variance:** given a pair of indices i and j , compute $V(\{A[i], A[i+1], \dots, A[j]\})$.

Initially all entries in A are set to 0.

Design a data structure that supports both operations in $O(\log n)$ time.

Some sort of augmented tree

Or two

but you can't avg 2 avgs
or subtract

✓

but keep sum of all leaves

and sum of squares of all leaves

Oh can subtract that

I was close...

SCRATCH PAPER

could be to me and
and the to change to me
that - that's all
- and all

SCRATCH PAPER

Final Exam Solutions

Problem 1. True or false [30 points] (10 parts)

For each of the following questions, circle either T (True) or F (False). **Explain your choice.** (Your explanation is worth more than your choice of true or false.)

(a) **T F** For all positive $f(n)$, $f(n) + o(f(n)) = \Theta(f(n))$.

Solution: True.

(b) **T F** For all positive $f(n)$, $g(n)$ and $h(n)$, if $f(n) = O(g(n))$ and $f(n) = \Omega(h(n))$, then $g(n) + h(n) = \Omega(f(n))$.

Solution: True. This follows from $f(n) = O(g(n)) \Rightarrow g(n) = \Omega(f(n))$.

(c) **T F** Under the simple uniform hashing assumption, the probability that three specific data elements (say 1, 2 and 3) hash to the same slot (i.e., $h(1) = h(2) = h(3)$) is $1/m^3$, where m is a number of buckets.

Solution: False. The above formula only describes the probability of collision in a *fixed* bucket (say bucket number 1). The correct answer is $1/m^2$.

(d) **T F** Given an array of n integers, each belonging to $\{-1, 0, 1\}$, we can sort the array in $O(n)$ time in the worst case.

Solution: True. Use counting sort, e.g., after adding 1 to all numbers.

- (e) T F The following array is a max heap: [10, 3, 5, 1, 4, 2].

Solution: False. The element 3 is smaller than its child 4, violating the max-heap property.

- (f) T F RADIX SORT does not work correctly (i.e., does not produce the correct output) if we sort each individual digit using INSERTION SORT instead of COUNTING SORT.

Solution: False. INSERTION SORT (as presented in class) is a stable sort, so RADIX SORT remains correct. The change can worsen running time, though.

3 points for correct answer and explanation

1 point for incorrect answer, but mentioning that radix sort needs to use a stable sort

- (g) T F Given a directed graph G , consider forming a graph G' as follows. Each vertex $u' \in G'$ represents a strongly connected component (SCC) of G . There is an edge (u', v') in G' if there is an edge in G from the SCC corresponding to u' to the SCC corresponding to v' . Then G' is a directed acyclic graph.

Solution: True. If there were any cycles in the graph of strongly connected components, then all the components on the cycle would actually be one strongly connected component.

3 points for correct answer

- (h) T F Consider two positively weighted graphs $G = (V, E, w)$ and $G' = (V, E, w')$ with the same vertices V and edges E such that, for any edge $e \in E$, we have $w'(e) = w(e)^2$. For any two vertices $u, v \in V$, any shortest path between u and v in G' is also a shortest path in G .

Solution: False. Assume we have two paths in G , one with weights 2 and 2 and another one with weight 3. The first one is shorter in G' while the second one is shorter in G .

3 points for correct answer - most people got this right

- (i) T F An optimal solution to a knapsack problem will always contain the object i with the greatest value-to-cost ratio v_i/c_i .

Solution: False. Greedy choice doesn't work for the knapsack problem. For example, if the maximum cost is 2, and there are two items, the first with cost 1 and value 2, and the second with cost 2 and value 3, the optimal solution is to take just the second item.

3 points for correct answer

- (j) T F Every problem in NP can be solved in exponential time.

Solution: True.

3 points for correct answer

Problem 2. Short answer [40 points] (8 parts)

- (a) Rank the following functions by increasing order of growth; that is, find an arrangement g_1, g_2, g_3, g_4 of the functions satisfying $g_1 = O(g_2)$, $g_2 = O(g_3)$, $g_3 = O(g_4)$. (For example, the correct ordering of n^2, n^4, n, n^3 is n, n^2, n^3, n^4 .)

$$f_1 = (n!)^{1/n} \quad f_2 = \log n^n \quad f_3 = n^{n^{1/2}} \quad f_4 = n \log n \log \log n$$

Solution: The correct order is f_1, f_2, f_4, f_3

- (b) Solve the following recurrences by giving tight Θ -notation bounds. You do not need to justify your answers, but any justification that you provide will help when assigning partial credit.

i. $T(n) = 4T(n/2) + n^2 \log n$

ii. $T(n) = 8T(n/2) + n \log n$

iii. $T(n) = \sqrt{6006} \cdot T(n/2) + n^{\sqrt{6006}}$

Solution:

i. Case 2 of the Master Method - $T(n) = n^2 \log^2 n$.

ii. Case 1 of the Master Method - $T(n) = n^3$.

iii. Case 3 of the Master Method to obtain $T(n) = \Theta(n^{\sqrt{6006}})$, checking that the regularity condition $af(n/2) = \sqrt{6006} \frac{n^{\sqrt{6006}}}{2} \leq cn^{\sqrt{6006}}$ for some $c < 1$. Condition holds for any $c > \sqrt{6006}/2^{\sqrt{6006}}$.

- (e) Give a recurrence $T(n) = \dots$ for the running time of each of the following algorithms, along with the asymptotic solution to that recurrence:
- Insertion sort
 - Merge sort
 - 2D peak finding (the fastest algorithm we've seen)

Solution:

- $T(n) = T(n-1) + O(n) = O(n^2)$
- $T(n) = 2T(n/2) + \Theta(n) = \Theta(n \lg n)$
- $T(n) = T(n/2) + \Theta(n) = \Theta(n)$

- (d) Could a binary search tree be built using $o(n \lg n)$ comparisons in the comparison model? Explain why or why not.

Solution: No, or else we could sort in $o(n \lg n)$ time by building a BST in $o(n \lg n)$ time and then doing an in-order tree walk in $O(n)$ time.

- (e) Given n integers in the range $0 \dots k$, describe how to preprocess these integers into a data structure that can answer the following query in $O(1)$ time: given two integers a and b , how many integers fall within the range $a \dots b$?

Solution: same preprocessing as for counting sort, then numbers in range is $C'(b) - C'(a)$.

- (f) Describe how any comparison-based sorting algorithm can be made stable, without affecting the running time by more than a constant factor.

Solution: Tag elements with their original positions in the array, only increase by a factor of 2 at most

- (g) In dynamic programming, we derive a recurrence relation for the solution to one subproblem in terms of solutions to other subproblems. To turn this relation into a bottom-up dynamic programming algorithm, we need an order to fill in the solution cells in a table, such that all needed subproblems are solved before solving a subproblem. For each of the following relations, give such a valid traversal order, or if no traversal order is possible for the given relation, briefly justify why.

- $A(i, j) = F(A(i, j - 1), A(i - 1, j - 1), A(i - 1, j + 1))$
- $A(i, j) = F(A(\min\{i, j\} - 1, \min\{i, j\} - 1), A(\max\{i, j\} - 1, \max\{i, j\} - 1))$
- $A(i, j) = F(A(i - 2, j - 2), A(i + 2, j + 2))$

Solution:

- Solve $A(i, j)$ for i from 0 to n : for j from 0 to n)
- Solve $A(k, k)$ for k from 0 to n) then solve rest in any order
- Impossible: cyclic.

- (h) Consider an array $A[1 \dots n]$ of integers in the range $1 \dots n^2$. A number a is a *heavy hitter* in A if a occurs in A at least $n/2$ times.

Give an efficient algorithm that finds all heavy hitters in a given array A .

Solution: Radix-sort and linear scan.

Problem 3. You are the computer [15 points] (3 parts)

- (a) Fill in the following grid with the correct subproblem solutions for this sequence alignment problem with these weights: 0 for mutation, 1 for insertion or deletion, and 3 for a match (the goal is to maximize the sum of the weights). Here "ATC" is the starting sequence and "TCAG" is the ending sequence.

-	-	A	T	C
-	0			
T				
C				
A				
G				

What is the optimal alignment?

Solution: The optimal alignment is to match the "TC" in both sequences.

-	-	A	T	C
-	0	1	2	3
T	1	2	4	5
C	2	3	5	7
A	3	5	6	8
G	4	6	7	9

3 points for correctly filled grid, 2 points for correct alignment of sequences

- (b) Draw a max-heap on the following set of integers: {2, 3, 5, 7, 11, 13, 17}. (You do not need to use the Build-Heap algorithm.)

Solution: multiple possible solutions

5 points for a correct answer - (almost) everyone got this correct

- (c) Compute $\sqrt[3]{6006}$ using two iterations of Newton's method, i.e., fill out the following table. Your entry for x_1 should be fully simplified. Your entry for x_2 can be left unsimplified.

i	x_i
0	1
1	
2	

Solution:

i	x_i
0	1
1	$2002.\bar{6}$
2	$2/3x_1 + \frac{2002}{x_1^2}$

5 points for a correct answer

-2 for significant arithmetic/algebraic errors

-3 for errors that led to negative or very large solutions (should have been caught by simple sanity checks)

Problem 4. Rotated array [10 points]

Consider an array $A[1 \dots n]$ constructed by the following process: we start with n distinct elements, sort them, and then rotate the array k steps to the right. For example, we might start with the sorted array $[1, 4, 5, 9, 10]$, and rotate it right by $k = 3$ steps to get $[5, 9, 10, 1, 4]$. Give an $O(\log n)$ -time algorithm that finds and returns the position of a given element x in array A , or returns None if x is not in A . Your algorithm is given the array $A[1 \dots n]$ but does *not* know k .

Solution: Solution I (to be fixed): You can perform a modified binary search.

SEARCH(A, i, j, x)

$first \leftarrow A[i]$

$middle \leftarrow A[(i + j)/2]$

$last \leftarrow A[j]$

if $x \in \{first, middle, last\}$:

 then

 return the corresponding index

 else :

 if $(x < first \text{ and } x > middle)$ or $(x > first \text{ and } x > middle)$: [xxx Isn't it equivalent to just $x > m$

 then

 return SEARCH($A, (i + j)/2, j, x$)

 else :

 return SEARCH($A, i, (i + j)/2, x$)

Solution II: Let $k \geq 0$ be the number of elements A was rotated by. We show how to identify the value of k , after which one can simply perform binary search in the "left part" $A[1 \dots k]$ and the "right part" $A[k + 1 \dots n]$. To this end, observe that all elements in the left part are not smaller than $A[1]$, while all elements in the right part are smaller than $A[1]$. By binary search we find the smallest $j > 1$ such that $A[j]$ is smaller than $A[1]$ (or set $j = 1$ if no such j exists). We then report $k = j - 1$.

Problem 5. Taxachusetts [10 points]

Suppose you are given a weighted graph $G = (V, E, w)$ of highways, and the state government has implemented a new tax rule whereby the cost of a path gets doubled as penalty if the number of edges in the path is greater than 10. Explain how to reduce finding the shortest-path weight between every pair of vertices (under this penalty) to the usual all-pairs shortest paths problem (as solved by Floyd-Warshall).

Solution: Augment to keep track of path lengths between each pair. We also augment so we keep track of the shortest-path weight between each pair using fewer than 10 edges and using greater than 10 edges.

Problem 6. Does this path make me look fat? [10 points]

Consider a connected weighted directed graph $G = (V, E, w)$. Define the *fatness* of a path P to be the maximum weight of any edge in P . Give an efficient algorithm that, given such a graph and two vertices $u, v \in V$, finds the minimum possible fatness of a path from u to v in G .

Solution: There are two good solutions to this problem.

We can see that that fatness must be the weight of one of the edges, so we sort all edge weights and perform a binary search. To test whether or not a path with a fatness no more than x exists, we perform a breadth-first search that only walks edges with weight less than or equal to x . If we reach v , such a path exists.

If such a path exists, we recurse on the lower part of the range we are searching, to see if a tighter fatness bound also applies. If such a path does not exist, we recurse on the upper part of the range we are searching, to relax that bound. When we find two neighboring values, one of which works and one of which doesn't, we have our answer. This takes $O((V + E) \lg E)$ time.

Another good solution is to modify Dijkstra's algorithm. We use "fatness" instead of the sum of edge weights to score paths, and the only change to Dijkstra itself that is necessary is to change the relaxation operation so that it compares the destination node's existing min-fatness with the max of the weight of the incoming edge (i, j) and the min-fatness of any path to i (the source of the incoming edge).

The correctness argument is almost precisely the same as that for Dijkstra's algorithm. A correct solution also had to note that negative-weight edges, which could be present here and normally break Dijkstra, don't do that here; adding negative numbers produces ever-more-negative path weights, but taking their max doesn't. This solution has the same time complexity as Dijkstra's algorithm.

Two less-efficient solutions were to use Bellman-Ford instead of Dijkstra (with the same modified relaxation step and an analogous correctness argument) and to perform the iterative search linearly instead of in a binary fashion. These received partial credit.

Problem 7. Indiana Jones and the Temple of Algorithms [10 points]

While exploring an ancient temple, Prof. Jones comes across a locked door. In front of this door are two pedestals, and n blocks each labeled with its positive integer weight. The sum of the weights of the blocks is W . In order to open the door, Prof. Jones needs to put every block on one of the two pedestals. However, if the difference in the sum of the weights of the blocks on each pedestal is too large, the door will not open.

Devise an algorithm that Prof. Jones can use to divide the blocks into two piles whose total weights are as close as possible. To avoid a boulder rolling quickly toward him, Prof. Jones needs your algorithm to run in pseudopolynomial time.

Solution: Just consider the smaller of the two piles. The goal is to make the weight of this pile as close to $W/2$ as possible, while not exceeding that weight. Our guess for each block is whether the block is included in this smaller pile or not.

Our subproblems are $P(i, w)$, which has the value “true” if it is possible to make a pile of weight $w \leq W/2$ with some subset of blocks 1 through i , and the value “false” otherwise. Initially, let $P(0, 0)$ be true, and $P(0, w)$ be false for all values of $w > 0$.

Let the weight of block i be w_i . Use the following recurrence:

$$P(i, w) = P(i - 1, w) \vee P(i - 1, w - w_i)$$

(where $P(i - 1, w - w_i)$ is considered to be false if $w_i > w$). If $P(i, w)$ is true, record in a separate table $B(i, w)$ which of $P(i - 1, w)$ or $P(i - 1, w - w_i)$ was true (choose arbitrarily if they are both true).

Find the largest value of w for which $P(n, w)$ is true. Then backtrack using the table B to determine which blocks were included in this pile to achieve this weight. This algorithm runs in $O(nW)$ time, because there are nW subproblems, each of which takes constant time.

Problem 8. Claustrophobic chickens [10 points]

Prof. Tyson has noticed that the chickens in his farm frequently get claustrophobic. He wants to build a monitoring system that will alert him when this happens, allowing him to manually rearrange the chickens. After thorough research, Prof. Tyson determines that a chicken becomes claustrophobic if it is within 2 feet of at least 8 other chickens.

Prof. Tyson has installed a tracker on each chicken, which reports the (x, y) coordinates (measured in feet) of each of his chickens. Suppose that there are n chickens whose locations are represented as a sequence $(x_1, y_1), \dots, (x_n, y_n)$. Prof. Tyson needs an efficient algorithm to determine whether there are any claustrophobic chickens (causing Prof. Tyson to go out and manually rearrange the chickens). Devise such an algorithm and save the chickens!


Solution: We shall solve this problem in a way similar to the close pair problem done in lecture, namely, partitioning the plane into a square grid and hashing points to corresponding buckets.

1. Impose a square grid onto the plane where each cell is a $\sqrt{2} \times \sqrt{2}$ square.
2. Hash each point into a bucket corresponding to the cell it belongs to.
3. If there is a bucket with ≥ 9 points in it, return YES.
4. Otherwise, for each point p , calculate distance to all points in the cell containing p as well as the neighboring cells. Return YES if the number of points within distance 2 is ≥ 8 . Clearly the algorithm will find a clumped chicken if one exists, either in step 3 or in step 4. By using a hash table of size $O(n)$, we can make the above algorithm run in expected linear time.

Problem 9. Architects 'R Us [15 points]

You are assisting Prof. Gehry with designing the shape of a new room in the Stata Center. The professor has given you n columns, each of the same unit thickness, but with different heights: $A[1], A[2], \dots, A[n]$. He asks you to permute the columns in a line to define the shape of the room. To make matters difficult, MIT wants to be able to hang a large rectangular picture on the columns. If j consecutive columns in your order all have a height of at least k , then we can hang a rectangle of size $j \cdot k$.

The example below contains 3 consecutive columns with heights of at least 4, so we can hang a rectangle of area 12 on the first three columns.



- (a) Give an efficient algorithm to find the largest area of a hangable rectangle for the *initial* order $A[1], A[2], \dots, A[n]$ of columns.

Solution: The best algorithms we know run in $O(n^2)$ time.

The simplest $O(n^2)$ algorithm is the following. The biggest rectangle is bounded on top by some column i . So the height of the rectangle is $A[i]$. Now walk left until reaching a column of height $< A[i]$, and similarly walk to the right. Count the number k of traversed columns, and multiply by $A[i]$.

Another $O(n^2)$ algorithm is the following. Define $m[i, j]$ to be the minimum of the interval $A[i], \dots, A[j]$. Then $m[i, j] = \min\{m[i, j-1], A[j]\}$, so by memoization, we can compute all $m[i, j]$'s in $O(n^2)$ time. Now the solution is $\max\{m[i, j] \cdot (j - i + 1) : i \leq j\}$, which takes $O(n^2)$ time to compute given the $m[i, j]$'s.

The easy brute-force algorithm already runs in $O(n^3)$ time (and was worth a base value of 4–5 out of 8 points for this part). Just use the computation above, but without memoizing the $m[i, j]$'s, so each takes $O(n)$ to compute, and we use $O(n^2)$ of them.

An $O(nh)$ algorithm, where $h = \max_i A[i]$, was worth a base value of 6 out of 8 points. We define one subproblem per (x, y) coordinate: $R(x, y)$ is the maximum possible area of a rectangle whose upper-right corner is at (x, y) . There are $O(nh)$

such subproblems. To solve the subproblem, we can use the $O(1)$ -time recurrence

$$R(x, y) = \begin{cases} R(x-1, y) + y & \text{if } A[x] \geq y \\ 0 & \text{otherwise.} \end{cases}$$

- (b) Devise an efficient algorithm to permute the columns into an order that maximizes the area of a hangable rectangle.

Solution: The intended algorithm is to sort the columns in decreasing order, e.g., using merge sort in $O(n \lg n)$ time. This works because, if the correct height of a rectangle is k , then at best it can involve all columns with height $\geq k$, and these are consecutive in the sorted order. In fact, increasing order works just as well, as does a strange order (suggested by several students) of putting the maximum in the middle, then repeatedly placing the next smaller column alternately between the left and right sides of the construction so far.

We can compute the hangable rectangle in $O(n)$ additional time, though this was not necessary to receive full credit. For each prefix $B[1 \dots i]$ of the sorted array, we'd like to compute $(\min B[1 \dots i]) \cdot i$, and take the maximum over all i . But $\min B[1 \dots i] = B[i]$, so this actually takes constant time per choice of i , for a total cost of $O(n)$ time.

2 out of 7 points were removed for lack of justification of sorting. 1 out of 7 points was removed for using counting (or radix) sort, which is not necessarily efficient given the setup of the problem.

Problem 10. Guess Who? [10 points]

Woody the woodcutter will cut a given log of wood, at any place you choose, for a price equal to the length of the given log. Suppose you have a log of length L , marked to be cut in n different locations labeled $1, 2, \dots, n$. For simplicity, let indices 0 and $n + 1$ denote the left and right endpoints of the original log of length L . Let d_i denote the distance of mark i from the left end of the log, and assume that $0 = d_0 < d_1 < d_2 < \dots < d_n < d_{n+1} = L$. The *wood-cutting problem* is the problem of determining the sequence of cuts to the log that will cut the log at all the marked places and minimize your total payment. Give an efficient algorithm to solve this problem.

Solution: Dynamic programming. $c(i, j) = \min_{i < k < j} \{c(i, k) + c(k, j) + (d_j - d_i)\}$ where $c(i, j)$ is the min cost of cutting a log with left endpoint i and right endpoint j at all its marked locations. Start with $c(i, i+1)$ (consecutive cuts) and move outwards to $c(i, i+2)$, $c(i, i+3)$ until the maximal distance $c(1, n)$, which gives the optimal score for the whole wood. Remember pointers to k that gave max score at each step, and trace back pointers to construct optimal solution. Each iteration takes $O(n)$ (linear search between i and j) and there are $O(n^2)$ entries to fill (a triangle really, not a square). Greedy solutions that pick the maximum cut each time do not work. Similarly, heuristics like picking the point closest to the center do not work.

Problem 11. Varying variance [20 points]

For a set S of numbers, define its *average* to be $\bar{S} = \frac{1}{|S|} \sum_{s \in S} s$, and its *variance* to be $V(S) = \frac{1}{|S|} \sum_{s \in S} (s - \bar{S})^2$.

A *segment variance data structure* supports the following operations on an array $A[1 \dots n]$:

- **Assignment:** given an index i and a value a , set $A[i] = a$.
- **Segment variance:** given a pair of indices i and j , compute $V(\{A[i], A[i+1], \dots, A[j]\})$.

Initially all entries in A are set to 0.

Design a data structure that supports both operations in $O(\log n)$ time.

Solution: We have $\frac{1}{|S|} \sum_{s \in S} (s - \bar{S})^2 = \frac{1}{|S|} \sum_{s \in S} s^2 - \bar{S}^2$. Both terms can be maintained using an augmented data structure that, in each internal node, keeps (i) the sum of all leaves and (ii) the sum of squares of all leaves.

5/22 Practice

Final Exam

- Do not open this quiz booklet until directed to do so. Read all the instructions on this page.
- When the quiz begins, write your name on every page of this quiz booklet.
- You have 180 minutes to earn 180 points. Do not spend too much time on any one problem. Read them all first, and attack them in the order that allows you to make the most progress.
- This exam is closed book. You may use **three** $8\frac{1}{2}'' \times 11''$ or A4 crib sheets (both sides). No calculators or programmable devices are permitted. No cell phones or other communications devices are permitted.
- Write your solutions in the space provided. If you need more space, write on the back of the sheet containing the problem. Pages may be separated for grading.
- Don't waste time and paper rederiving facts that we have studied. It is sufficient to cite known results.
- When writing an algorithm, a **clear** description will suffice. Pseudo-code isn't required.
- When asked for an algorithm, your algorithm should have the time complexity specified in the problem with a correct analysis. If you cannot find such an algorithm, you will generally receive partial credit for a slower algorithm **if you analyze your algorithm correctly**.
- Show your work, as partial credit will be given. You will be graded not only on the correctness of your answer, but also on the clarity with which you express it.

Problem	Parts	Points	Grade	Grader
1	2	2		
2	1	30		
3	2	30		
4	2	28		
5	3	30		
6	3	30		
7	2	30		
Total		180		

Name: _____

Friday
Recitation:

Aleksander
11 AM

Arnab
12 PM

Alina
3 PM

Matthew
4 PM

Problem 1. What is Your Name? [2 points] (2 parts)

(a) [1 point] Flip back to the cover page. Write your name there.



(b) [1 point] Flip back to the cover page. Circle your recitation section.



Problem 2. Storing Partial Maxima [30 points] (1 part)

6.006 student, Mike Velli, wants to build a website where the user can input a time interval in history, and the website will return the most exciting sports event that occurred during this interval. Formally, suppose that Mike has a chronologically sorted list of n sports events with associated integer "excitement factors" e_1, \dots, e_n . You can assume for simplicity that n is a power of 2. A user's query will consist of a pair (i, j) with $1 \leq i < j \leq n$, and the site is supposed to return $\max(e_i, e_{i+1}, \dots, e_j)$.

Mike wishes to minimize the amount of computation per query, since there will be a lot of traffic to the website. If he precomputes and stores $\max(e_i, \dots, e_j)$ for every possible input (i, j) , he can respond to user queries quickly, but he needs storage $\Omega(n^2)$ which is too much.

In order to reduce storage requirements, Mike is willing to allow a small amount of computation per query. He wants to store a cleverer selection of precomputed values than just $\max(e_i, \dots, e_j)$ for every (i, j) , so that for any user query, the server can retrieve two precomputed values and take the maximum of the two to return the final answer. Show that now only $O(n \log n)$ values need to be precomputed.

BST

- but interval

how deal w/ both

store max - but might be outside the range

Depends how large range is

- could store ~~max~~ sorted list of e max \rightarrow min

- search through till find in range

$O(n)$ worst case

that's actually pretty good

Extra

Extra

start, end

clue!

lower bound

Problem 3. Longest Simple Cycle [30 points] (2 parts)

Given an unweighted, directed graph $G = (V, E)$, a path $\langle v_1, v_2, \dots, v_n \rangle$ is a set of vertices such that for all $0 < i < n$, there is an edge from v_i to v_{i+1} . A cycle is a path such that there is also an edge from v_n to v_1 . A simple path is a path with no repeated vertices and, similarly, a simple cycle is a cycle with no repeated vertices. In this question we consider two problems:

- **LONGESTSIMPLEPATH**: Given a graph $G = (V, E)$ and two vertices $u, v \in V$, find a simple path of maximum length from u to v or output NONE if no path exists.
- **LONGESTSIMPLECYCLE**: Given a graph $G = (V, E)$, find a simple cycle of maximum length in G .

- (a) [20 points] Reduce the problem of finding the longest simple path to the problem of finding the longest simple cycle. Prove the correctness of your reduction and show that it runs in polynomial time in $|V|$ and $|E|$.

But longest path might not be longest cycle

Cycle is just special case ... longest path that folds back on itself

DFS $O(V+E)$

- (b) [10 points] As we discussed in class, finding a longest simple path is NP-Hard. Therefore, there is no known algorithm that, on input u, v , and G returns a longest simple path from u to v in polynomial time. Using this fact (which you do not need to prove) and Part (a), show that there is no known polynomial time algorithm that can find a longest simple cycle in a graph.

Note: If you were unable to solve Part (a), you may assume an algorithm **SIMPLEPATHFROMCYCLE** for finding a longest simple path from u to v that runs in time polynomial in $L, |V|$, and $|E|$ where L is the running time of a black-box algorithm for solving **LONGESTSIMPLECYCLE**.

This seems pretty simple

A longest simple cycle is just a special case of path ✓

But write us contradiction "proof format"

Path u v

↳ That's what I thought

extra
1

for each $1 \leq i \leq n/2$ store $\max(e_i, e_{i+1}, \dots, e_{n/2})$

so for every starting point store best up to $n/2$

for each $n/2 < j \leq n$ same

recurse sep on both lists
till list size 11

What are we recursing on?

if $i \leq n/2$ $j > n/2$ then return max of both

if $i, j \leq n/2$ or $i, j > n/2$ ans found recursively

Again how?

- Does that mean normal way?

~~$O(n) + 2$~~

$S(n) = 2S(n/2) + O(n)$

" $O(n \log n)$ master theorem

I still like my sol better

Extra
2

The sol is way longer than the space to ans

Create a new graph G'

Copy all vertices + edges

Add $|V|$ more vertices to G'

$w_1, \dots, w_{|V|}$

For $0 < i < |V|$ create a directed edge
from w_i to w_{i+1}

Also create directed edge $v \rightarrow w_1$
 $w_{|V|} \rightarrow u$

Run longest simple cycle

That makes it more difficult - ~~no~~
was not thinking that was possible

Problem 4. Closest pair [28 points] (2 parts)

We are interested in finding the closest pair of points in the plane, closest in the sense of the rectilinear distance (also called the Manhattan or L_1 distance). The rectilinear distance between two points p_1 and p_2 on the plane is defined as $d(p_1, p_2) = |x_1 - x_2| + |y_1 - y_2|$, where the x 's and y 's are the first and second coordinates, respectively. In the first part we consider a one-dimensional version of the problem as a warm-up. In both cases, coordinates of the points are real numbers with k significant digits beyond the decimal point, k constant.

- (a) [8 points] Warm up - Provide an efficient way to find a closest pair among n points in the interval $[0, 1]$ on the line. Full credit will be given to the most efficient algorithm with a correct analysis.



So what are the points?
How are they paired?
So smallest distance
None $O(n^2)$

On last exam this was grid + hash

Extra 3

- (b) [20 points] Case of the plane - A divide and conquer approach for n points in the square $[0, 1] \times [0, 1]$. Here is a possible strategy. Divide the set of points into two sets of about half size: those on the right of the x -coordinate median, and those on the left. Recursively, find the closest pair of points on the right, and the closest pair of points on the left and let δ_r and δ_l be the corresponding distances. The overall closest pair is either the minimum of these two options, or corresponds to a pair where one point is on the right of the median and the other is on the left. Given δ_r and δ_l , there should be an efficient way to find the latter. Explain how; then write the full recurrence for the running time of this approach; and conclude with its overall running time.

don't forget the special case →

very similar to what I had

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

\lg^2

$$O(n) \lg n$$

thought it was this

in worst cases all n pts on the strip

ah right dividing on space - not points

Problem 5. APSP Algorithm for Sparse Graphs [30 points] (3 parts)

Let $G = (V, E)$ be a weighted, directed graph that can have some of the weights negative. Let $n = |V|$ and $m = |E|$, and assume that G is strongly connected, i.e., for any vertex u and v there is a path from u to v in G .

We want to solve all-pairs-shortest-path problem (APSP) in G , i.e., we want to either find all the vertex-to-vertex distances $\{\delta(v, u)\}_{v, u \in V}$, or report existence of a negative-length cycle. We will design an algorithm for this task that runs in $O(mn + n^2 \log n)$ time. (Note that when G is not dense, i.e., when m is $o(n^2)$, the running time of this algorithm is asymptotically better than the one of the Floyd-Warshall algorithm.)

- (a) [5 points] Fix some vertex $t \in V$ and consider the vertex potential $\lambda_t(u) = \delta(u, t)$ where $\delta(u, t)$ is the shortest path from $u \in V$ to t . Give an algorithm for calculating $\lambda_t(u)$ for all $u \in V$ and analyze the running time.

- (b) [5 points] Show $\lambda_t(u)$, the potential from Part (a), is a feasible potential even if some of the original weights are negative.

ah whole thing strongly connected

clue

heuristic

Extra 3


But needs to be some way to do diff in 2D

I guess it can be hash groups

But need right group size

Largest group = best

Or divide out



$$\frac{1}{2} + \frac{1}{4} + \frac{1}{8} = 1$$

So $O(n)$ sol

Radix sort $O(kn)$

Calc distance b/w it and one before $O(k)$

So total $O(n)$ since k fixed

I think my sol would have worked

↳ almost like binary search

(c) [20 points] Show how you use the vertex potential from Parts (a) and (b) to solve APSP in G in $O(mn + n^2 \log n)$ time, including the time it takes to calculate the vertex potential.

[Faint handwritten notes at the bottom of the page, including the phrase "show on each edge" and some mathematical symbols.]

Problem 6. Airplane Scheduling [30 points] (3 parts) *lecture*

Consider the runway reservation system from PS2. Recall that we kept track of requested landing times from airplanes by storing them in a balanced binary search tree. In that problem set, we required, for safety, that no landing time be within three minutes of any other landing time in the tree. We showed we could insert into the tree, delete from the tree, and check the validity of a landing time in $O(\log n)$ time.

Sometimes severe weather hits, and the 3-minute window between flights just isn't safe, so a new window size is determined. In these cases, an extra runway might be opened up at a nearby airport to take flights which don't fit within the new window. For all parts, you may use data structure augmentations provided that you explain the augmentation. Its maintenance may not increase the asymptotic running time of other operations, but you are not required to prove this.

- (a) [5 points] Provide a very fast (constant-time) algorithm to determine if there are any flights which are not valid with the new window so that the extra runway can start opening immediately.

I don't fully remember the last aly
Basically used next largest
in order tree walk - checking every gap

isn't simple

(cool missing from paper)

- (b) [15 points] Also provide a slower algorithm which locates which specific flights are not valid within the new window so they can be rescheduled. For example, assuming a window of 3 minutes with flights at times 28, 31, 35, 40, 43, 48, 53, 57, 60, if the window size were expanded to 4 minutes, the algorithm should return that 28, 31, 40, 43, 57, 60 are invalid.

Why is 28 invalid?

So both are canceled?

isn't wouldn't an in-order tree walk work?

track min window size of subtree

look at root if any window smaller than new window
ahh very clever - wanted old time

Can traverse ~~extra~~ tree in $O(n)$ - so sol to a walk

- (c) [10 points] Can the running time of this slower algorithm be improved if we assume bounds on the maximum size the window could become?

Yes - can keep a side-structure

no extra details :)

Problem 7. Traveling on a Budget [30 points] (2 parts)

Arthur Dent has \$500 and 1000 ^{too long} hours to go from Cambridge, MA, to Berkeley, CA. He has a map of the States represented as a directed graph $G = (V, E)$. The vertices of the graph represent towns, and there is a directed edge $e = (A, B)$ from town A to town B if there is some means of public transportation connecting the two towns. Moreover, the edge is labeled with a pair (m_e, t_e) , representing the cost $m_e \in \{0, 1, \dots\}$ in dollars of transportation from A to B and the time $t_e \in \{0, 1, \dots\}$ in hours that it takes to go from A to B.

Arthur is interested in finding a path from Cambridge to Berkeley that does not cost more than \$500 and does not take more than 1000 hours, while also minimizing the objective $5M^2 + 2T^2$, where M is the cost of the trip in dollars and T is the duration of the trip in hours. He was looking for an algorithm that runs in time polynomial in $|V|$ and $|E| \dots$


- (a) [10 points] Due to his lack of knowledge in algorithms, he gave up on the idea of respecting the budget and time constraints. At least he thought he could efficiently find the path minimizing the objective $5M^2 + 2T^2$. He tried modifying Dijkstra's algorithm as follows: If an edge e was labeled (m_e, t_e) , he assigned it a weight $w_e = 5m_e^2 + 2t_e^2$, and ran Dijkstra's algorithm on the resulting weighted directed graph. Show that Arthur's algorithm may return incorrect results, i.e. return a path that does not minimize the objective $5M^2 + 2T^2$.

oh geeze →

inapsack
 don't be greedy
 DP
 7 basic rules
 (things I learned studying)

Something about 2 legs vs 1

Say



← this is 50050 ← when this is actually much better

17512 17518 = 29000

(pick smaller Hs on actual exam)

- (b) [20 points] Now provide an algorithm that solves Arthur's original problem in time polynomial in $|E|$ and $|V|$. Your algorithm should find the path that minimizes the objective $5M^2 + 2T^2$, while at the same time respecting the constraints $M \leq 500$ and $T \leq 1000$. Please describe your algorithm precisely, and justify its correctness and running time. More credit will be given to faster algorithms, provided that the analysis of the algorithm is correct.

[Hint 1: Use dynamic programming.] *Oh how to do this*

[Hint 2: For each town A , integer values $m \leq 500$ and $t \leq 1000$, either there is a path from Cambridge to A that requires cost m and time t , or there is not.]

$$\text{Cost} = \min_{\text{all edges}} ($$

The fn

but how to avoid the problem we saw

- save the cost + time and re calc the function each time

506 x 15⁻ matrix

- no further ans

but this is like the all pairs problem

- what I would have answered

SCRATCH PAPER

When I was in
the army we used to
write letters to our
loved ones and they
would come back to us
with their answers.

When I was in the
army we used to write
letters to our loved ones
and they would come back
to us with their answers.

SCRATCH PAPER

Final Exam Solutions

Problem 1. What is Your Name? [2 points] (2 parts)

(a) [1 point] Flip back to the cover page. Write your name there.

(b) [1 point] Flip back to the cover page. Circle your recitation section.

Problem 2. Storing Partial Maxima [30 points] (1 part)

6.006 student, Mike Velli, wants to build a website where the user can input a time interval in history, and the website will return the most exciting sports event that occurred during this interval. Formally, suppose that Mike has a chronologically sorted list of n sports events with associated integer "excitement factors" e_1, \dots, e_n . You can assume for simplicity that n is a power of 2. A user's query will consist of a pair (i, j) with $1 \leq i < j \leq n$, and the site is supposed to return $\max(e_i, e_{i+1}, \dots, e_j)$.

Mike wishes to minimize the amount of computation per query, since there will be a lot of traffic to the website. If he precomputes and stores $\max(e_i, \dots, e_j)$ for every possible input (i, j) , he can respond to user queries quickly, but he needs storage $\Omega(n^2)$ which is too much.

In order to reduce storage requirements, Mike is willing to allow a small amount of computation per query. He wants to store a cleverer selection of precomputed values than just $\max(e_i, \dots, e_j)$ for every (i, j) , so that for any user query, the server can retrieve two precomputed values and take the maximum of the two to return the final answer. Show that now only $O(n \log n)$ values need to be precomputed.

Solution: We are given the list e_1, \dots, e_n . For each $1 \leq i \leq n/2$, store $\max(e_i, e_{i+1}, \dots, e_{n/2})$, and for each $n/2 < j \leq n$, store $\max(e_{n/2+1}, \dots, e_j)$. Recurse separately on the two lists $e_1, \dots, e_{n/2}$ and $e_{n/2+1}, \dots, e_n$. Stop the recursion when the list size becomes 1.

If the user's query is (i, j) with $i \leq n/2$ and $j > n/2$, then we can return $\max(\max(e_i, e_{i+1}, \dots, e_{n/2}), \max(e_{n/2+1}, \dots, e_j))$. If both $i, j \leq n/2$ or $i, j > n/2$, then the answer is found recursively.

Let $S(n)$ be the number of values stored for a list of length n . By construction, $S(n) = O(n) + 2S(n/2)$, and therefore, $S(n) = O(n \log n)$.

Problem 3. Longest Simple Cycle [30 points] (2 parts)

Given an unweighted, directed graph $G = (V, E)$, a path $\langle v_1, v_2, \dots, v_n \rangle$ is a set of vertices such that for all $0 < i < n$, there is an edge from v_i to v_{i+1} . A cycle is a path such that there is also an edge from v_n to v_1 . A *simple path* is a path with no repeated vertices and, similarly, a *simple cycle* is a cycle with no repeated vertices. In this question we consider two problems:

- **LONGESTSIMPLEPATH:** Given a graph $G = (V, E)$ and two vertices $u, v \in V$, find a simple path of maximum length from u to v or output NONE if no path exists.
 - **LONGESTSIMPLECYCLE:** Given a graph $G = (V, E)$, find a simple cycle of maximum length in G .
- (a) [20 points] Reduce the problem of finding the longest simple path to the problem of finding the longest simple cycle. Prove the correctness of your reduction and show that it runs in polynomial time in $|V|$ and $|E|$.

Solution: Create a new graph G' . Copy all vertices and edges from G to G' . Then add $|V|$ more vertices to G' , $w_1, \dots, w_{|V|}$. For $0 < i < |V|$, create a directed edge from w_i to w_{i+1} . Also create a directed edge from v to w_1 and a directed edge from $w_{|V|}$ to u . Run LONGESTSIMPLECYCLE on G' . If the longest simple cycle involves vertices u and v , return the path from u to v in the cycle. Otherwise, output NONE.

Running Time: Creating the new graph takes $O(|V| + |E|)$. We run LONGESTSIMPLECYCLE once. Therefore, the running time of the algorithm is $O(|V| + |E| + L)$ where L is the running time of LONGESTSIMPLECYCLE.

Correctness: If a path from u to v exists then $c = \langle u, \dots, v, w_1, \dots, w_{|V|} \rangle$ is a cycle in G' . The cycle c contains at least $|V| + 2$ vertices and, by construction, is the only cycle containing the w vertices. Therefore, all other cycles must have length at most $|V|$. Hence LONGESTSIMPLECYCLE will return this cycle if it exists and our algorithm will return $p = \langle u, \dots, v \rangle$. Since the cycle is simple, p must contain only vertices in V and is thus the longest path from u to v in G . If the cycle does not exist, there is no path from u to v and the algorithm correctly outputs NONE.

Note: This is only one possible reduction. People came up with others and that's fine so long as they are correct and polynomial.

Grading:

- 20/20 For any polynomial time reduction. The reduction did not have to be linear to receive full credit.
- 19.5/20 if you added a weighted edge from v to u rather than a number of vertices (the graph was unweighted)
- 18/20 If you gave the reduction of adding an edge from v to u , finding the connected cluster containing v and u and running LSC on that. The problem is that

the connected cluster containing v and u might have a simple cycle that does not involve v and u longer than the one that does involve v and u . The whole connected cluster is a cycle, but not a simple one.

- 17/20 For removing back edges
- 17/20 For removing all nodes u can't reach
- 14/20 If you tried to remove edges one at a time, which does not work, but the algorithm was still polynomial.
- 10/20 If you returned the longest path in the graph, but not necessarily from u to v .
- 10/20 If you did the reduction correctly in the wrong direction
- From the above grades points were also subtracted for:
 - -1 for no runtime analysis
 - -2 for no correctness proof
 - -1 if you forgot to add an edge from v to u and your reduction required one.

- (b) [10 points] As we discussed in class, finding a longest simple path is NP-Hard. Therefore, there is no known algorithm that, on input u, v , and G returns a longest simple path from u to v in polynomial time. Using this fact (which you do not need to prove) and Part (a), show that there is no known polynomial time algorithm that can find a longest simple cycle in a graph.

Note: If you were unable to solve Part (a), you may assume an algorithm SIMPLEPATHFROMCYCLE for finding a longest simple path from u to v that runs in time polynomial in $L, |V|$, and $|E|$ where L is the running time of a black-box algorithm for solving LONGESTSIMPLECYCLE.

Solution: If L is polynomial then the algorithm outlined in Part (a) gives a polynomial time algorithm for finding the longest simple path in a graph. Contradiction.

Grading:

- 10/10 For anything like above proof. It was also fine to say that the reduction means LSC must be at least as hard as LSP since we stressed that formulation in class.
- 9/10 If you said you can use LSC to solve LSP and therefore by definition of NP-Hard, LSC is also NP-Hard. The definition of NP-Hard is that all NP optimization problems can be reduced to some problem in NP-Hard. So you need an extra step just tying all the ends together, but it's very close.

Problem 4. Closest pair [28 points] (2 parts)

We are interested in finding the closest pair of points in the plane, closest in the sense of the rectilinear distance (also called the Manhattan or L_1 distance). The rectilinear distance between two points p_1 and p_2 on the plane is defined as $d(p_1, p_2) = |x_1 - x_2| + |y_1 - y_2|$, where the x 's and y 's are the first and second coordinates, respectively. In the first part we consider a one-dimensional version of the problem as a warm-up. In both cases, coordinates of the points are real numbers with k significant digits beyond the decimal point, k constant.

- (a) [8 points] Warm up - Provide an efficient way to find a closest pair among n points in the interval $[0, 1]$ on the line. Full credit will be given to the most efficient algorithm with a correct analysis.

Solution: Use radix sort on the n numbers (k digits representation). Then go through the sorted list and calculate the distance between each point and the next one, keeping a running minimum along the way. Radix sort will take $O(kn)$, and distance calculations $n - 1$ time $O(k)$, so an overall linear time $O(n)$, since k is fixed. (Half points only given to an $O(n \log n)$ answer with otherwise correct analysis.)

- (b) [20 points] Case of the plane - A divide and conquer approach for n points in the square $[0, 1] \times [0, 1]$. Here is a possible strategy. Divide the set of points into two sets of about half size: those on the right of the x -coordinate median, and those on the left. Recursively, find the closest pair of points on the right, and the closest pair of points on the left and let δ_r and δ_l be the corresponding distances. The overall closest pair is either the minimum of these two options, or corresponds to a pair where one point is on the right of the median and the other is on the left. Given δ_r and δ_l , there should be an efficient way to find the latter. Explain how; then write the full recurrence for the running time of this approach; and conclude with its overall running time.

Solution: Searching for closest pair of points that cross the median x -coordinate can be limited to the vertical strip around the median line of width $\pm \delta$ where $\delta = \min\{\delta_r, \delta_l\}$. By having the points belonging to this strip sorted according to their y -coordinates, one can go through of all them bottom up, and for each check distances to points above them within a vertical bound of δ (one can show that only the 7 points above need to be checked for each point). In the worst case, all n points are in this strip, so we end up with a recurrence for the running time of $T(n) = 2T(n/2) + O(n)$ and so $T(n) = O(n \log n)$.

Problem 5. APSP Algorithm for Sparse Graphs [30 points] (3 parts)

Let $G = (V, E)$ be a weighted, directed graph that can have some of the weights negative. Let $n = |V|$ and $m = |E|$, and assume that G is strongly connected, i.e., for any vertex u and v there is a path from u to v in G .

We want to solve all-pairs-shortest-path problem (APSP) in G , i.e., we want to either find all the vertex-to-vertex distances $\{\delta(v, u)\}_{v, u \in V}$, or report existence of a negative-length cycle. We will design an algorithm for this task that runs in $O(mn + n^2 \log n)$ time. (Note that when G is not dense, i.e., when m is $o(n^2)$, the running time of this algorithm is asymptotically better than the one of the Floyd-Warshall algorithm.)

- (a) [5 points] Fix some vertex $t \in V$ and consider the vertex potential $\lambda_t(u) = \delta(u, t)$ where $\delta(u, t)$ is the shortest path from $u \in V$ to t . Give an algorithm for calculating $\lambda_t(u)$ for all $u \in V$ and analyze the running time.

Solution: Bellman-Ford in $O(|V||E|)$ time.

- (b) [5 points] Show $\lambda_t(u)$, the potential from Part (a), is a feasible potential even if some of the original weights are negative.

Proof of the feasibility of λ :

Consider $w^*(u, v)$ for any $u, v \in V$. By definition

$$w^*(u, v) := w(u, v) - \lambda(u) + \lambda(v) = w(u, v) - \delta(u, t) + \delta(v, t).$$

Since $\delta(u, v) \leq w(u, v)$, we have that

$$w^*(u, v) = w(u, v) - \delta(u, t) + \delta(v, t) \geq \delta(u, v) - \delta(u, t) + \delta(v, t).$$

But, by triangle inequality we have that $\delta(u, t) \leq \delta(u, v) + \delta(v, t)$, thus

$$w^*(u, v) \geq \delta(u, v) - \delta(u, t) + \delta(v, t) \geq 0,$$

as desired.

- (c) [20 points] Show how you use the vertex potential from Parts (a) and (b) to solve APSP in G in $O(mn + n^2 \log n)$ time, including the time it takes to calculate the vertex potential.

Solution: Choose any vertex t in G and run Bellman-Ford algorithm from it on the transposed graph G to either compute all the distances $\delta(u, t)$ for every $u \in V$, or detect a negative-length cycle. If such a cycle is detected in the transposed graph then it also exists in the original graph, so just report its existence and terminate.

Otherwise, define the vertex potential $\lambda(u) := \delta(u, t)$ and look at the resulting reduced weight $w^*(u, v) := w(u, v) - \lambda(u) + \lambda(v)$ in G . Since G is strongly connected (and we now know it does not have negative-length cycles), all distances $\delta(u, t)$ are finite and thus all $\lambda(u)$'s are finite as well. From the recitations (and pset 5B) we know that this implies that λ is a feasible potential, i.e., for every $u, v \in V$ $w^*(u, v)$ is non-negative (for the sake of completeness we include the proof of feasibility below).

Since the reduced weight w^* is non-negative, one can run Dijkstra's algorithm in G from every node $u \in V$ to compute all the vertex-to-vertex distances $\{\delta^*(u, v)\}_{u, v \in V}$ with respect to w^* . Next, extract the true vertex-to-vertex distances in G by computing $\delta(u, v) = \delta^*(u, v) - \lambda(v) + \lambda(u)$ for all $u, v \in V$, and output them.

The running time of the above algorithm is dominated by one execution of Bellman-Ford algorithm – which is $O(mn)$ time – and n executions of Dijkstra's algorithm – which is $O(n(m + n \log n)) = O(mn + n^2 \log n)$ time using the implementation of Dijkstra with Fibonacci heaps. Thus the total running time is $O(mn + n^2 \log n)$, as desired.

Problem 6. Airplane Scheduling [30 points] (3 parts)

Consider the runway reservation system from PS2. Recall that we kept track of requested landing times from airplanes by storing them in a balanced binary search tree. In that problem set, we required, for safety, that no landing time be within three minutes of any other landing time in the tree. We showed we could insert into the tree, delete from the tree, and check the validity of a landing time in $O(\log n)$ time.

Sometimes severe weather hits, and the 3-minute window between flights just isn't safe, so a new window size is determined. In these cases, an extra runway might be opened up at a nearby airport to take flights which don't fit within the new window. For all parts, you may use data structure augmentations provided that you explain the augmentation. Its maintenance may not increase the asymptotic running time of other operations, but you are not required to prove this.

- (a) [5 points] Provide a very fast (constant-time) algorithm to determine if there are any flights which are not valid with the new window so that the extra runway can start opening immediately.

missing sol!

- (b) [15 points] Also provide a slower algorithm which locates which specific flights are not valid within the new window so they can be rescheduled. For example, assuming a window of 3 minutes with flights at times 28, 31, 35, 40, 43, 48, 53, 57, 60, if the window size were expanded to 4 minutes, the algorithm should return that 28, 31, 40, 43, 57, 60 are invalid.

Solution: Like question 2d on the problem set, but instead of tracking for the presence of a particular window size, track the minimum window of the subtree. Constant-time lookup at the root to see if any window is smaller than the new window value. Traversing the tree is both $O(n)$ and $O(k \log n)$ to find all newly invalid times. With bounds on the window size, we could keep a side-structure with all flights belonging to a given window.

- (c) [10 points] Can the running time of this slower algorithm be improved if we assume bounds on the maximum size the window could become?

No sol here either

Problem 7. Traveling on a Budget [30 points] (2 parts)

Arthur Dent has \$500 and 1000 hours to go from Cambridge, MA, to Berkeley, CA. He has a map of the States represented as a directed graph $G = (V, E)$. The vertices of the graph represent towns, and there is a directed edge $e = (A, B)$ from town A to town B if there is some means of public transportation connecting the two towns. Moreover, the edge is labeled with a pair (m_e, t_e) , representing the cost $m_e \in \{0, 1, \dots\}$ in dollars of transportation from A to B and the time $t_e \in \{0, 1, \dots\}$ in hours that it takes to go from A to B.

Arthur is interested in finding a path from Cambridge to Berkeley that does not cost more than \$500 and does not take more than 1000 hours, while also minimizing the objective $5M^2 + 2T^2$, where M is the cost of the trip in dollars and T is the duration of the trip in hours. He was looking for an algorithm that runs in time polynomial in $|V|$ and $|E| \dots$

- (a) [10 points] Due to his lack of knowledge in algorithms, he gave up on the idea of respecting the budget and time constraints. At least he thought he could efficiently find the path minimizing the objective $5M^2 + 2T^2$. He tried modifying Dijkstra's algorithm as follows: If an edge e was labeled (m_e, t_e) , he assigned it a weight $w_e = 5m_e^2 + 2t_e^2$, and ran Dijkstra's algorithm on the resulting weighted directed graph. Show that Arthur's algorithm may return incorrect results, i.e. return a path that does not minimize the objective $5M^2 + 2T^2$.

- (b) [20 points] Now provide an algorithm that solves Arthur's original problem in time polynomial in $|E|$ and $|V|$. Your algorithm should find the path that minimizes the objective $5M^2 + 2T^2$, while at the same time respecting the constraints $M \leq 500$ and $T \leq 1000$. Please describe your algorithm precisely, and justify its correctness and running time. More credit will be given to faster algorithms, provided that the analysis of the algorithm is correct.

[Hint 1: Use dynamic programming.]

[Hint 2: For each town A , integer values $m \leq 500$ and $t \leq 1000$, either there is a path from Cambridge to A that requires cost m and time t , or there is not.]

Solution sketch: We construct a 501-by-15 sized matrix M_v at each vertex v so that the (i, j) 'th entry of M_v is 1 if there's a path from the node for Cambridge to v of total cost i and total duration j and is 0 otherwise. $M_{\text{Cambridge}}$ is set to be 1 at entry $(0, 0)$ and 0 everywhere else.

Final Exam

- Do not open this quiz booklet until directed to do so. Read all the instructions on this page.
- When the quiz begins, write your name on every page of this quiz booklet.
- You have 180 minutes to earn **200** points. Do not spend too much time on any one problem. Read them all through first, and attack them in the order that allows you to make the most progress.
- This quiz booklet contains 12 pages, including this one. Two extra sheets of scratch paper are attached. Please detach them before turning in your quiz at the end of the exam period.
- This quiz is closed book. You may use **three** $8\frac{1}{2}'' \times 11''$ or A4 crib sheets (both sides). No calculators or programmable devices are permitted. No cell phones or other communications devices are permitted.
- Write your solutions in the space provided. If you need more space, write on the back of the sheet containing the problem. Pages may be separated for grading.
- Do not waste time and paper rederiving facts that we have studied. It is sufficient to cite known results.
- Show your work, as partial credit will be given. You will be graded not only on the correctness of your answer, but also on the clarity with which you express it. Be neat.
- Good luck!

Problem	Parts	Points	Grade	Grader	Problem	Parts	Points	Grade	Grader
1	8	24			5	1	24		
2	4	32			6	1	24		
3	3	24			7	1	24		
4	3	24			8	1	24		
					Total		200		

Name: _____

Problem 1. True or False [24 points] (8 parts)

For each of the following questions, circle either T (True) or F (False). **Explain your choice.** (No credit if no explanation given.)

(a) **T F** There exists an algorithm to build a binary search tree from an unsorted list in $O(n)$ time.
Explain:

(b) **T F** There exists an algorithm to build a binary heap from an unsorted list in $O(n)$ time.
Explain:

(c) **T F** To solve the SSSP problem for a graph with no negative-weight edges, it is necessary that some edge be relaxed at least twice.
Explain:

(d) **T F** On a connected, directed graph with only positive edge weights, Bellman-Ford runs asymptotically as fast as Dijkstra.
Explain:

(e) **T F** A Givens rotation requires $O(1)$ time.

Explain:

(f) **T F** In the worst case, merge sort runs in $O(n^2)$ time.

Explain:

(g) **T F** There exists a stable implementation of merge sort.

Explain:

(h) **T F** An AVL tree T contains n integers, all distinct. For a given integer k , there exists a $\Theta(\lg n)$ algorithm to find the element x in T such that $|k - x|$ is minimized.

Explain:

Problem 2. Short Answer [32 points] (4 parts)

- (a) The *eccentricity* $\epsilon(u)$ of a vertex u in a connected, undirected, *unweighted* graph G is the maximum distance from u to any other vertex in the graph. That is, if $\delta(u, v)$ is the shortest path from u to v , then $\epsilon(u) = \max_{v \in V} \delta(u, v)$.

Give an efficient algorithm to find the eccentricity of a given vertex s . Analyze its running time.

- (b) What is the asymptotic cost of solving a linear system of equations with $n-1$ equations of the form

$$a_{i,i}x_i + a_{i,i+1}x_{i+1} = b_i \quad i = 1, \dots, n-1$$

and one equation of the form

$$a_{n,1}x_1 + a_{n,2}x_2 + \dots + a_{n,n}x_n = b_n$$

(none of the a 's in this equation are zero). The a 's and b 's are given numbers and the x 's are unknowns. Assume that we use Givens rotations to reduce the coefficient matrix to a triangular form. Justify your answer.

- (c) Suppose a hash function h maps arbitrary keys to values between 0 and $m - 1$, inclusive. We hash n keys, k_1, k_2, \dots, k_n . If $h(k_i) = h(k_j)$, we say that k_i and k_j collide. Assuming simple uniform hashing, how should we choose m (in terms of n) such that the expected number of collisions is $O(1)$? Justify your answer.

- (d) Suppose you have a directed acyclic graph with n vertices and $O(n)$ edges, all having nonnegative weights. Propose an efficient method of finding the shortest path to each vertex from a single source, and give its running time in terms of n .

Problem 3. Judge Jill [24 points] (3 parts)

Judge Jill has created a web site that allows people to file complaints about one another. Each complaint contains exactly two names: that of the person who filed it and that of the person he/she is complaining about.

Jill had hoped to resolve each complaint personally, but the site has received so many complaints that she has realized she wants an automated approach.

She decides to try to label each person as either good or evil. She only needs the labeling to be consistent, not necessarily correct. A labeling is consistent if every complaint labels one person as good and the other person as evil, and no person gets labeled both as good and evil in different complaints.

- (a) [8 points] Propose a way to model the consistent labeling problem as a graph problem.

(b) [10 points] Propose an efficient algorithm to consistently label all the names as good or evil, or to decide that no such classification exists. Use the graph model you proposed in the previous part of the problem. Analyze the running time of the algorithm.

(c) [6 points] Later, Judge Jill wants to be more thorough. She will interview some people to figure out who is good and who is evil. She can always determine whether a person is good or evil by interviewing him or her. Assuming that one person in every complaint is good and the other is evil, what is the minimum number of people she needs to interview to correctly classify all the people named in the complaints?

Problem 4. Bitdiddle Bins [24 points] (3 parts)

Ben Bitdiddle has devised a new data structure called a Bitdiddle Bin. Much like an array or a set, you can INSERT values into it, and you can LOOKUP values to see if they are contained in the structure. (He'll figure out DELETE later.)

A Bitdiddle Bin is implemented as a pair of lists (arrays), designated the *neat list* and the *messy list*, with these properties:

- The *neat list* is always in sorted order. (The messy list may or may not be sorted.)
- The *messy list* has a size of at most \sqrt{n} , where n is the total number of values in the entire Bitdiddle Bin.

The LOOKUP algorithm for a Bitdiddle Bin is as follows:

1. Use binary search to look for the value in the neat list.
2. If it wasn't in the neat list, iterate over the entire messy list to look for the value.

The INSERT algorithm is as follows:

1. Append the value to the messy list.
2. If the messy list is now too big, CLEANUP.

This CLEANUP subroutine is run whenever the messy list grows beyond \sqrt{n} items:

1. Sort the messy list.
2. Merge the messy list with the neat list.
3. The merge result is the new neat list. The new messy list is empty.

(a) [4 points] What is the worst-case asymptotic runtime of LOOKUP on a Bitdiddle Bin?

(b) [8 points] What is the worst-case asymptotic runtime of INSERT on a Bitdiddle Bin? Explain.

(c) [12 points] What is the *amortized* asymptotic runtime of each INSERT operation, when inserting n values into an empty Bitdiddle Bin? Explain.

Problem 5. Local Minimum [24 points]

Consider an array A containing n distinct integers. We define a *local minimum* of A to be an x such that $x = A[i]$, for some $0 \leq i < n$, with $A[i - 1] > A[i]$ and $A[i] < A[i + 1]$. In other words, a local minimum x is less than its neighbors in A (for boundary elements, there is only one neighbor). Note that A might have multiple local minima.

As an example, suppose $A = [10, 6, 4, 3, 12, 19, 18]$. Then A has two local minima: 3 and 18.

Of course, the absolute minimum of A is always a local minimum, but it requires $\Omega(n)$ time to compute.

Propose an efficient algorithm to find some local minimum of A , and analyze the running time of your algorithm.

Problem 6. Straits [24 points]


Let G be a connected, weighted, undirected graph. All the edge weights are positive. An edge e is a *strait* between vertices u and v if it has weight ℓ , is on a path from u to v whose other edges weigh ℓ or more, and every path between u and v contains an edge of weight ℓ or less. In other words, to go from u to v you must cross an edge of weight ℓ , but you do not need to cross edges lighter than ℓ .

Describe an efficient algorithm for finding the weight of the strait between every pair of vertices in G . Analyze the running time of the algorithm.

Hint: Use a $|V|$ -by- $|V|$ table to keep track of the weights of all the straits. Initialize it so that it is correct for a graph with no edges. Then add the edges one by one.

Problem 7. The Cake Is a Lie! [24 points]

At Aperture Bakeries, every cake comes with a binary boolean-valued tree indicating whether or not it is available. Each leaf in the tree has either a *true* or a *false* value. Each of the remaining nodes has exactly two children and is labeled either *and* or *or*; the value is the result of recursively applying the operator to the values of the children. One example is the following tree:



If the root of a tree evaluates to *false*, like the one above, the cake is a lie and you cannot have it. Any *true* cake is free for the taking. You may modify a tree to make it *true*; the only thing you can do to change a tree is to turn a *false* leaf into a *true* leaf, or vice versa. This costs \$1 for each leaf you change. You can't alter the operators or the structure of the tree.

Cake is good. Cheap cake is even better. Describe an efficient algorithm to determine the minimum cost of a cake whose tree has n nodes, and analyze its running time.

Problem 8. I Am Locutus of Borg, You Will Respond To My Questions [24 points]

Upon arrival at the planet Vertex T, you and Ensign Treaps are captured by the Borg. They promptly throw the ensign out of the airlock. You had better solve their problem, lest you share his fate.

The Vertex T parking lot is an $n \times n$ matrix A . There are already a number of spaceships parked in the lot. For $0 \leq i, j < n$, let $A[i][j] = 0$ if there is a ship occupying position (i, j) , and 1 otherwise.

The Borg want to find the *largest square parking space* in which to park the Borg Cube. That is, find the largest k such that there exists a $k \times k$ square in A containing all ones and no zeros. In the example figure, the solution is 3, as illustrated by the 3×3 highlighted box.

	0	1	2	3	4
0	0	1	1	1	0
1	1	1	1	1	1
2	1	1	1	1	1
3	1	1	0	0	0
4	1	1	1	0	1

Describe an efficient algorithm that finds the size of the largest square parking space in A . Analyze the running time of your algorithm.

Hint: Call $A[0][0]$ be the *top-left* of the parking lot, and call $A[n - 1][n - 1]$ the *bottom-right*. Use dynamic programming, with the subproblem $S[i, j]$ being the side length of the largest square parking space whose bottom-right corner is at (i, j) .

SCRATCH PAPER

SCRATCH PAPER

Final Exam Solutions

Problem 1. True or False [24 points] (8 parts)

For each of the following questions, circle either T (True) or F (False). Explain your choice. (No credit if no explanation given.)

- (a) T F There exists an algorithm to build a binary search tree from an unsorted list in $O(n)$ time.
Explain:

Solution: False. Because an in-order walk can create a sorted list from a BST in $O(n)$ time, the existence of such an algorithm would violate the $\Omega(n \lg n)$ lower bound on comparison-based sorts.

- (b) T F There exists an algorithm to build a binary heap from an unsorted list in $O(n)$ time.
Explain:

Solution: True. The standard BUILD-HEAP algorithm runs in $O(n)$ time.

- (c) T F To solve the SSSP problem for a graph with no negative-weight edges, it is necessary that some edge be relaxed at least twice.
Explain:

Solution: False. Dijkstra's algorithm solves the SSSP problem relaxing each edge at most once.

- (d) T F On a connected, directed graph with only positive edge weights, Bellman-Ford runs asymptotically as fast as Dijkstra.
Explain:

Solution: False. Bellman-Ford requires $\Theta(VE)$, regardless of the edge weights. Dijkstra runs in $\Theta(E + V \lg V)$. Because the graph is connected, $E = \Omega(V)$, so $\Theta(VE) = \Omega(V^2)$, which is clearly worse than Dijkstra.

- (e) T F A Givens rotation requires $O(1)$ time.
Explain:

Solution: False. It may take $O(1)$ time if the matrix is sparse, but in general a Givens rotation requires $O(n)$ for a matrix with n columns.

- (f) T F In the worst case, merge sort runs in $O(n^2)$ time.
Explain:

Solution: True. Merge sort runs in $O(n \lg n)$ time which is $O(n^2)$.

- (g) T F There exists a stable implementation of merge sort.
Explain:

Solution: True. If the items in the two lists being merged are equal, choose the item in the "left" half (the half that came first in the original array).

- (h) T F An AVL tree T contains n integers, all distinct. For a given integer k , there exists a $\Theta(\lg n)$ algorithm to find the element x in T such that $|k - x|$ is minimized.
Explain:

Solution: True. INSERT k , then find the PREDECESSOR and SUCCESSOR of k . Return the one whose difference with k is smaller. All three methods take $\Theta(\lg n)$ time.

Problem 2. Short Answer [32 points] (4 parts)

- (a) The *eccentricity* $\epsilon(u)$ of a vertex u in a connected, undirected, *unweighted* graph G is the maximum distance from u to any other vertex in the graph. That is, if $\delta(u, v)$ is the shortest path from u to v , then $\epsilon(u) = \max_{v \in V} \delta(u, v)$.

Give an efficient algorithm to find the eccentricity of a given vertex s . Analyze its running time.

Solution: Run BFS starting at s . Keep track of how many "levels" have been explored. The number of levels required to explore the entire graph is equal to $\epsilon(s)$. This requires $O(V + E)$ time.

- (b) What is the asymptotic cost of solving a linear system of equations with $n-1$ equations of the form

$$a_{i,i}x_i + a_{i,i+1}x_{i+1} = b_i \quad i = 1, \dots, n-1$$

and one equation of the form

$$a_{n,1}x_1 + a_{n,2}x_2 + \dots + a_{n,n}x_n = b_n$$

(none of the a 's in this equation are zero). The a 's and b 's are given numbers and the x 's are unknowns. Assume that we use Givens rotations to reduce the coefficient matrix to a triangular form. Justify your answer.

Solution: If we form a matrix and apply Givens rotations, eliminating the first column will fill in all the elements below the main diagonal, so we will need to perform $\Theta(n^2)$ rotations at a total cost of $\Theta(n^3)$. If we move the last equation (the one without zero coefficients) to the top, we only need to perform one Givens rotation per column, so the total cost reduces to $\Theta(n^2)$.

- (c) Suppose a hash function h maps arbitrary keys to values between 0 and $m-1$, inclusive. We hash n keys, k_1, k_2, \dots, k_n . If $h(k_i) = h(k_j)$, we say that k_i and k_j collide. Assuming simple uniform hashing, how should we choose m (in terms of n) such that the expected number of collisions is $O(1)$? Justify your answer.

Solution: Any two keys collide with probability $1/m$. There are $\binom{n}{2} = O(n^2)$ pairs of elements, so the expected number of collisions is $O(n^2/m)$. Thus, let $m = \Omega(n^2)$.

- (d) Suppose you have a directed acyclic graph with n vertices and $O(n)$ edges, all having nonnegative weights. Propose an efficient method of finding the shortest path to each vertex from a single source, and give its running time in terms of n .

Solution: The given conditions allow us to use either the DAG shortest paths algorithm or Dijkstra's. DAG shortest paths runs in $\Theta(V + E) = \Theta(n)$, and Dijkstra's runs in $O((V + E) \lg V)$, or $O(V \lg V + E)$ if using a Fibonacci heap. In either case, Dijkstra's runs in $O(n \lg n)$. DAG shortest paths is faster in this case.

Problem 3. Judge Jill [24 points] (3 parts)

Judge Jill has created a web site that allows people to file complaints about one another. Each complaint contains exactly two names: that of the person who filed it and that of the person he/she is complaining about.

Jill had hoped to resolve each complaint personally, but the site has received so many complaints that she has realized she wants an automated approach.

She decides to try to label each person as either good or evil. She only needs the labeling to be consistent, not necessarily correct. A labeling is consistent if every complaint labels one person as good and the other person as evil, and no person gets labeled both as good and evil in different complaints.

- (a) [8 points] Propose a way to model the consistent labeling problem as a graph problem.

Solution: We can model the problem using an undirected graph in which every edge represents one complaint and every vertex represents one name (one person).

- (b) [10 points] Propose an efficient algorithm to consistently label all the names as good or evil, or to decide that no such classification exists. Use the graph model you proposed in the previous part of the problem. Analyze the running time of the algorithm.

Solution: The problem is equivalent to testing whether we can color the vertices of graph with two colors, or equivalently, whether the graph is bipartite. This can be done using DFS in $\Theta(V + E)$ time. We label the root of DFS traversals arbitrarily good or bad. When we first visit a vertex, we label it differently than its parent. When we encounter back or cross edges, we check for consistency; if we find an inconsistency, the graph cannot be labeled consistently.

- (c) [6 points] Later, Judge Jill wants to be more thorough. She will interview some people to figure out who is good and who is evil. She can always determine whether a person is good or evil by interviewing him or her. Assuming that one person in every complaint is good and the other is evil, what is the minimum number of people she needs to interview to correctly classify all the people named in the complaints?

Solution: She needs to interview one person in each connected component of the complaint graph.

Problem 4. Bitdiddle Bins [24 points] (3 parts)

Ben Bitdiddle has devised a new data structure called a Bitdiddle Bin. Much like an array or a set, you can INSERT values into it, and you can LOOKUP values to see if they are contained in the structure. (He'll figure out DELETE later.)

A Bitdiddle Bin is implemented as a pair of lists (arrays), designated the *neat list* and the *messy list*, with these properties:

- The *neat list* is always in sorted order. (The messy list may or may not be sorted.)
- The *messy list* has a size of at most \sqrt{n} , where n is the total number of values in the entire Bitdiddle Bin.

The LOOKUP algorithm for a Bitdiddle Bin is as follows:

1. Use binary search to look for the value in the neat list.
2. If it wasn't in the neat list, iterate over the entire messy list to look for the value.

The INSERT algorithm is as follows:

1. Append the value to the messy list.
2. If the messy list is now too big, CLEANUP.

This CLEANUP subroutine is run whenever the messy list grows beyond \sqrt{n} items:

1. Sort the messy list.
2. Merge the messy list with the neat list.
3. The merge result is the new neat list. The new messy list is empty.

(a) [4 points] What is the worst-case asymptotic runtime of LOOKUP on a Bitdiddle Bin?

Solution: Binary search on the neat list takes $\Theta(\log n)$, and iterating over the messy list takes $\Theta(\sqrt{n})$, for a total time of $\Theta(\sqrt{n})$.

(b) [8 points] What is the worst-case asymptotic runtime of INSERT on a Bitdiddle Bin? Explain.

Solution: Appending is $\Theta(1)$. Sorting the messy list is $\Theta(\sqrt{n} \log \sqrt{n})$. Merging is $\Theta(n)$. The $\Theta(n)$ term dominates, so the asymptotic runtime is $\Theta(n)$.

(c) [12 points] What is the *amortized* asymptotic runtime of each INSERT operation, when inserting n values into an empty Bitdiddle Bin? Explain.

Solution: Use the aggregate method. The cost is $\Theta(1)$ to append to the messy list, plus $\Theta(n)$ per CLEANUP. Because CLEANUP is called $\Theta(\sqrt{n})$ times, the total cost for the n insertions is $O(n\sqrt{n})$, so the amortized cost per operation is $O(n\sqrt{n}/n) = O(\sqrt{n})$.

Problem 5. Local Minimum [24 points]

Consider an array A containing n distinct integers. We define a *local minimum* of A to be an x such that $x = A[i]$, for some $0 \leq i < n$, with $A[i-1] > A[i]$ and $A[i] < A[i+1]$. In other words, a local minimum x is less than its neighbors in A (for boundary elements, there is only one neighbor). Note that A might have multiple local minima.

As an example, suppose $A = [10, 6, 4, 3, 12, 19, 18]$. Then A has two local minima: 3 and 18.

Of course, the absolute minimum of A is always a local minimum, but it requires $\Omega(n)$ time to compute.

Propose an efficient algorithm to find some local minimum of A , and analyze the running time of your algorithm.

Solution: Use a divide-and-conquer algorithm. Let $m = n/2$, and examine the value $A[m]$ (that is, the element in the middle of the array).

Case 1: $A[m-1] < A[m]$. Then the left half of the array must contain a local minimum, so recurse on the left half. We can show this by contradiction: assume that $A[i]$ is not a local minimum for each $0 \leq i < m$. Then $A[m-1]$ is not a local minimum, which implies that $A[m-2] < A[m-1]$. Similarly, $A[m-3] < A[m-2]$. Continuing in this fashion, we obtain $A[0] < A[1]$. But then $A[0]$ is a local minimum, contrary to our initial assumption.

Case 2: $A[m+1] > A[m]$. Then the right half of the array must contain a local minimum, so recurse on the right half. This is symmetrical to Case 1.

Case 3: $A[m-1] > A[m]$ and $A[m+1] < A[m]$. Then $A[m]$ is a local minimum, so return it.

The running time recurrence is $T(n) = T(n/2) + \Theta(1)$, which yields $T(n) = \Theta(\log n)$.

Problem 6. Straits [24 points]

Let G be a connected, weighted, undirected graph. All the edge weights are positive. An edge e is a *strait* between vertices u and v if it has weight ℓ , is on a path from u to v whose other edges weigh ℓ or more, and every path between u and v contains an edge of weight ℓ or less. In other words, to go from u to v you must cross an edge of weight ℓ , but you do not need to cross edges lighter than ℓ .

Describe an efficient algorithm for finding the weight of the strait between every pair of vertices in G . Analyze the running time of the algorithm.

Hint: Use a $|V|$ -by- $|V|$ table to keep track of the weights of all the straits. Initialize it so that it is correct for a graph with no edges. Then add the edges one by one.


Solution: The initial table has infinity on the diagonal and zero everywhere else. Suppose the table T is correct for $E \setminus \{u, v\}$. We define the table T' for E using the rule

$$T'[x, y] = \max [T(x, y), \min (T(x, u), w(u, v), T(v, y)), \min (T(x, v), w(v, u), T(u, y))].$$

The cost of this update is constant. We perform $V(V-1)$ such updates for every edge, so the total cost is $\Theta(EV^2)$.

Problem 7. The Cake Is a Lie! [24 points]

At Aperture Bakeries, every cake comes with a binary boolean-valued tree indicating whether or not it is available. Each leaf in the tree has either a *true* or a *false* value. Each of the remaining nodes has exactly two children and is labeled either *and* or *or*; the value is the result of recursively applying the operator to the values of the children. One example is the following tree:



If the root of a tree evaluates to *false*, like the one above, the cake is a lie and you cannot have it. Any *true* cake is free for the taking. You may modify a tree to make it *true*; the only thing you can do to change a tree is to turn a *false* leaf into a *true* leaf, or vice versa. This costs \$1 for each leaf you change. You can't alter the operators or the structure of the tree.

Cake is good. Cheap cake is even better. Describe an efficient algorithm to determine the minimum cost of a cake whose tree has n nodes, and analyze its running time.

Solution: One can compute the truth value of each node in $O(n)$ time by starting with the leaves and going up. The precise order can be computed in $O(n)$ time using either a topological sort or BFS. If the root is *true*, the cake is free, so return \$0.

Otherwise, we can use dynamic programming to determine this cost. Augment each node with a field C , the minimum cost to make the node *true*. Each node corresponds to the subproblem of determining C at that node.

The base case consists of the leaves. For each *true* leaf, set C to \$0. For each *false* leaf, set C to \$1.

Recursive cases:

For all *and* nodes: Set C to the sum of the children's C values.

For all *or* nodes: Set C to the smaller of the children's C values.

After computing C for all nodes, in the same order in which the truth values were computed initially, return the C value of the root.

Since there are $O(n)$ subproblems, and the cost of each subproblem is constant, the total running time is $O(n)$.

Problem 8. I Am Locutus of Borg, You Will Respond To My Questions [24 points]

Upon arrival at the planet Vertex T, you and Ensign Treaps are captured by the Borg. They promptly throw the ensign out of the airlock. You had better solve their problem, lest you share his fate.

The Vertex T parking lot is an $n \times n$ matrix A . There are already a number of spaceships parked in the lot. For $0 \leq i, j < n$, let $A[i][j] = 0$ if there is a ship occupying position (i, j) , and 1 otherwise.

	0	1	2	3	4
0	0	1	1	1	0
1	1	1	1	1	1
2	1	1	1	1	1
3	1	1	0	0	0
4	1	1	1	0	1

The Borg want to find the *largest square parking space* in which to park the Borg Cube. That is, find the largest k such that there exists a $k \times k$ square in A containing all ones and no zeros. In the example figure, the solution is 3, as illustrated by the 3×3 highlighted box.

Describe an efficient algorithm that finds the size of the largest square parking space in A . Analyze the running time of your algorithm.

Hint: Call $A[0][0]$ be the *top-left* of the parking lot, and call $A[n-1][n-1]$ the *bottom-right*. Use dynamic programming, with the subproblem $S[i, j]$ being the side length of the largest square parking space whose bottom-right corner is at (i, j) .

Solution: The base cases are the positions along the top and left sides of the parking lot. In each of these positions, you can fit a 1×1 square parking space if and only if the space is unoccupied. Thus, for the base cases ($i = 0$ or $j = 0$), we have $S[i, j] = A[i][j]$.

Now for the general case. Again, if $A[i][j] = 0$, then $S[i, j] = 0$, so we'll only consider the case where $A[i][j] = 1$. There is a parking space of size x with bottom-right corner at (i, j) if and only if there are three (overlapping) spaces of size $x - 1$ at each of the locations $(i - 1, j)$, $(i, j - 1)$, and $(i - 1, j - 1)$. Thus, the largest possible parking space is

$$S[i, j] = 1 + \min(S[i - 1, j], S[i, j - 1], S[i - 1, j - 1])$$

The desired answer is $\max_{i, j} S[i, j]$, which requires $O(n^2)$ to compute.

There are n^2 subproblems, and each takes $O(1)$ time to solve, for a time of $O(n^2)$. This, added to the $O(n^2)$ to extract the desired answer, gives a total $O(n^2)$ running time, which is clearly optimal because all the data must be examined.

Exam Thoughts

5/23

not as bad as I thought
time goes fast once exams start

I think I did well on this
Can I make a B?

Break after every problem

Don't put on book too fast!

You just need consent.

Thought more deeply about the problems than any time in semester

I feel really confident about - even though prob got a few wrong
But did really deliberate

Afterwards

I don't know why I felt so confident
or feel

felt like I knew the stuff

but I was not able to figure 2 of 3 of
the long problems
and I short

but felt like I got everything else

Perhaps I just liked my focus + concentration

#5

Topo ordering

k edges inserted in

may add cycles

Arbitrary weights

So what matters about the previous topo ordering?

$O(k \cdot (V + E + k))$ algo for shortest path

di is $O((V + E) \log V)$ 'don't fully remember

So have topo best in $V + E$

So check every path through each of the k possibilities if shorter?

↳ I'm going off the $O()$ - that's usually right.

This seems likely

Visit each

- can see it in path

Find shortest path from k

- is better b/c cycles

(Not sure - but time looks correct)

check indiv

↳ no if 2 nodes.

What does checking mean?

Hand waving it now...

(5.7)

Checking is exploding out from it in \forall Exh time
in ~~DFS~~ DFS

Oh weights!

Dij

- but what is that?

- the queue is no good

I have a feeling this fails...

↳ Leave it...

No longer think about t_0

Now go to Gb and do that!

#6

 V_i and p_i

$$\text{So score} = \prod p_i \cdot \sum v_i \quad \text{but only adds score}$$

Each score multiplies prob

$$1) p_1 \cdot v_1$$

$$2) p_1 \cdot p_2 \cdot (v_1 + v_2)$$

$$p_1(v_1 + v_2) \cdot p_2(v_1 + v_2)$$

$$3) p_1 \cdot p_2 \cdot p_3 \cdot (v_1 + v_2 + v_3)$$

$$p_1(v_1 + v_2 + v_3) \cdot p_2(v_1 + v_2 + v_3) \cdot p_3(v_1 + v_2 + v_3)$$

Easy problem $v_i = 1$

$$\text{So } p_1(2) \cdot p_2(2) = 2(p_1 \cdot p_2)$$

So we add

$$p_1 \cdot 1$$

$$p_1 \cdot \underbrace{p_1 \cdot p_2 \cdot p_2}_{}$$

$$\text{or } 3(p_1 \cdot p_2 \cdot p_3)$$

Want optimal subset

DP - include or not!

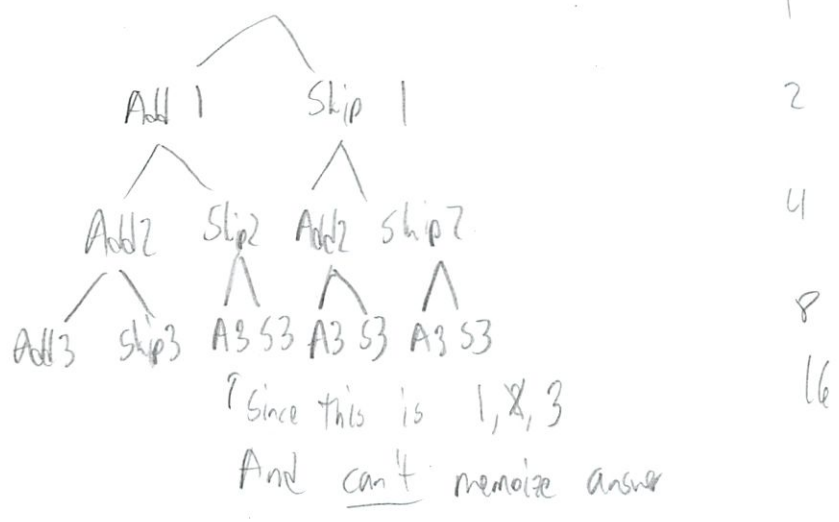
(2)

but doesn't need to be contiguous

DP(i) = max (p_i added, p_i not added)
Best score

This is O(1) to compute
O(n) to check
w/ memo

No its not just a line



Recursion $2T(n) + O(1)$

	2^{n-1}	n^2
1	2	1
2	4	4
3	8	9
4	16	16
5	32	25

? its exponential

Only choose best col from each

↳ greedy
but thats not best here

DFS
↳ not nlog n

③ I guess the key is - is one always better?

I think so

Either the EV of the question is good or its not

Plus use the fact its \perp here

keep p sep

just multiply on p

and know whats included, multiply score by that


but still branching exponential

how to limit branching

by reusing

Can we say 1 branch is better than the others?

Yeah I think that is the advantage - pick the highest so far
at each level - confine that



$\log n$ is usually binary search

Or! order the p s from lowest to highest!

Then include or not

I like this

(often I think of the part to sol lot)

Remember ~~an~~ explain, analysis, correctness

9)

b) Now v_i can be an int b/w 1 and C

- So limited upper value
- So can bullet etc
- Want a $O(n^2)$
- What I had before was exponential
- So C means we can try each value
- EV each individually
 - time won't allow
- What's n^2
 - each problem w/ every other problem
 - that gets you best of 2
- But at each cost level?

At each cost level, compare each problem w/ every other problem. See what is larger and use that;

↑
? Start w/ local best + work down?

But how compare b/w V_i levels

? Start from high points to low points
or back to local EV

5

I have some ideas - but don't see how they fit in $O(n^2)$

And I can't see how to get something in $O(n^2)$ to work

Wierd \rightarrow no credit since not in time

Perhaps write

At every possible cost level...

Or does that come from the calc of it

which is n

Start w/ max points

- try each include or not

- $O(n)$ to calc

Go through each pts level

Skip it, out of time