Admin

- Assignment 8 back soon
- Final review out tonight

Skip List Details

- Search(x):
 - Start at top list, go right just before value gets > target
 - Go down and repeat until element is found or hit bottom



– **Example:** Search for 72

- * Level 1: 14 too small, 79 too big; go down 14 * Level 2: 14 too small, 50 too small, 79 too big; go down 50 * Level 3: 50 too small, 66 too small, 79 too big; go down 66

- * Level 4: 66 too small, 72 spot on

Skip List Analysis

- Let us first define the height of a skip list formally.
- Let L_k be the set of all items in level $k \ge 1$.
- Height of an element. $\ell(x) = \max\{k \mid x \in L_k\}$
- Height of a skip list. $h(L) = \max\{\ell(x) \mid x \in L_0\}$





Skip List Expected Analysis

- Expected height of a node:
 - $E[\ell(x)] = 1 + \frac{1}{2} \cdot 0 + \frac{1}{2}(1 + E[\ell(x)])$
 - $E[\ell(x)] = 2$
- Worst-case height? $h(L) = \max\{\ell(x) \mid x \in L\}$



Skip List Analysis

- high probability.
- is, $\leq 1/n^c$ where the constant $c \geq 1$
- to level 1?
 - · 1/2

• **Claim.** A skip list with n elements has height $O(\log n)$ levels with

• Recall. (Informally) An event happens with high probability if the probability that it does not happen is polynomially small in n, that

• (More formally) Skip list of size n has $O(\log n)$ levels with high probability if the probability that it has more than $d \log n$ levels is at most $1/n^c$ where the constants c, d usually depend on each other

• **Proof idea.** What is the probability that an element gets promoted

Skip List Analysis

- high probability.
- . $\Pr[\ell(x) = k] = \frac{1}{2^k}$ $\Pr[\ell(x) > k] = \sum_{k=1}^{\infty} \Pr[\ell(x) = k]$ *k*+1
- $\Pr[h(L) > k] = \Pr[\bigcup_{x \in L} \ell(x)]$
- $\Pr[h(L) > c \log n] \le \frac{1}{n^{c-1}}$
- Thus, height of skip is $O(\log n)$ with high probability

• **Claim.** A skip list with *n* elements has height $O(\log n)$ levels with

• **Proof.** For any $x \in L$, $k \ge 1$, the probability that height of x is k

$$= i] = \sum_{i=k+1}^{\infty} \frac{1}{2^{i}} = \frac{1}{2^{k}}$$
$$> k] \le \sum_{x \in L} \Pr[\ell(x) > k] = \frac{n}{2^{k}}$$

Union bound

[pick any c > 2 for w.h.p.]



- **Claim.** Search cost in a skip list is $O(\log n)$ with high probability
- Proof.
- Idea think of the search path "backwards"
- Starting at the target element
- Going left or up until you reach top-left element



Skip List Search Cost

- Backwards search path, when do go up versus left?
- If node wasn't promoted (got tails here), then we go [came from] left
- If node was promoted (got heads here), then we go [came from] top
- How many consecutive tails in a row? (left moves on a level)
 - Same analysis as the height! $O(\log n)$
 - $O(\log^2 n)$ length overall—but I claimed $O(\log n)$ earlier



Skip List Search Cost

- Thus, number of "up" moves is at most $c \log n$ with high probability
- Search path is a sequence of *HHHTTTHHTT*...
- Search cost:
 - heads with high probability?



Skip List Search Cost

• We know height is $O(\log n)$ with high probability; say it is $c \log n$

• How many times do we need to flip a coin to get $c \log n$

- Claim. Number of flips until $c \log n$ heads is $\Theta(\log n)$ with high probability, that is, with probability $1 - 1/n^c$
- Note. Constant in $\Theta(\log n)$ will depend on c
- Proof. •
 - Say we flip $10c \log n$ coins
 - When are there at least $c \log n$ heads?
 - Pr[exactly $c \log n$ heads] $= \left(\frac{10c\log n}{c\log n}\right) \cdot \left(\frac{1}{2}\right)^{c\log n} \cdot \left(\frac{1}{2}\right)^{9c}$

Coin Flipping

```
• Pr[at most c \log n heads] \leq \left( \frac{10c \log n}{c \log n} \right) \cdot \left( \frac{1}{2} \right)^9
```

- Claim. Number of flips until $c \log n$ heads is $\Theta(\log n)$ with high probability, that is, with probability $1 - 1/n^c$
- Proof.
 - Pr[at most $c \log n$ heads] ≤

 $d = 9 - \log(10e) \rightarrow \infty$, independent of c

Coin Flipping

$$\leq \left(\frac{e \cdot 10c \log n}{c \log n}\right)^{c \log n} \cdot \left(\frac{1}{2}\right)^{9c \log n}$$
$$= (10e)^{c \log n} \cdot \left(\frac{1}{2}\right)^{9c \log n}$$
$$= 2^{\log(10e) \cdot c \log n} \cdot \left(\frac{1}{2}\right)^{9c \log n}$$
$$= 2^{(\log(10e) - 9) \cdot c \log n} = 2^{-d \log n}$$
$$= 1/n^d$$

• If we instead look at probability of at most $d'c \log n$ heads, as $d' \to \infty$,

- Using $O(\log n)$ linked lists, achieve same performance as binary search tree
- No stored information about balance, no tricky balancing rules
- Just flip coins while inserting each new element to decide what lists it goes in

Skip Lists

Approximation Algorithms

Aside: Online Algorithms

- For algorithms we've seen, we have all data up front
- Not always true in practice!
- commit to a solution before you see all of it?

• What happens when data comes in gradually, and you need to

Challenges: Approximation Algorithms

- Approximating problems that are NP hard
 - optimal when the optimal solution is not known/NP hard
 - solution for minimization (maximization) problems
- Approximation for online algorithms
 - element
 - algorithm has cost $k \cdot OPT$

Main challenge is showing that the algorithm performs close to

• Usually done by lower (upper) bounding the cost of the optimal

• High benchmark. Comparison against an optimal that knows the entire future, while the algorithm does not even know the next

• k-competitive: if the optimal offline algorithm has cost OPT, our

Online: Ski Rental Problem

- Assume that you are taking ski lessons
- any information about when that will be.
- Question: rent or buy the skis?
- Cost of renting \$1
- Cost of buying \$*B*
- (you know t) what is the best strategy?
 - If $t \ge B$ times, then buy, else rent
 - In other words, optimal offline cost is $min\{t, B\}$

• At some point (after t days) the ski season ends. But you don't have

• Offline strategy. If you knew in advance when the ski season ends



Online: Ski Rental Problem

- such you buy skis on the kth visit (renting before then)
- buying on the Bth ski visit is 2-competitive
- Offline cost is $\min\{t, B\} = B$
- Online strategy's cost?
 - are 2-competitive (in fact we are 1-competitive)
 - are 2-competitive

• Online strategy. We need to figure out a decision point, a number k

• Claim. If we set k = B (the cost of buying skis), we are gauranteed to never pay more than twice of the best offline optimal strategy. That is,

• If $t \leq B$, then our cost is t, and OPT has cost t. In that case, we

• If t > B, then our cost is 2B, and OPT has cost B. In that case, we



- **Input.** *m* identical machines
- *n* jobs with processing times t_1, \ldots, t_m , where job *j* has processing time t_i (on any machine)
- Job j must run contiguously on one machine
- A machine can process at most one job at a time.
- Let S[i] be the subset of jobs assigned to machine i.



- $L = \max L[i]$
- makespan.
- Claim. Load balancing is NP hard even with m = 2 machines
- Proof. Reduction from PARTITION problem.
- We will design an approximation algorithm for this problem
- [Greedy returns!] Consider the following greedy strategy:
 - Fix some order on the jobs
 - Assign job *j* to machine *i* whose load is smallest so far

• The makespan of an algorithm is the maximum load on any machine

• Load balancing Problem. Assign jobs to machines so as to minimize

Load Balancing: Greedy

- Go through the jobs one by one
- Assign each job to the machine with the smallest load so far

- How can we keep track of this efficiently?
 - Priority queue

Load Balancing: Greedy



- Running time?
 - $O(n \log m)$ using a priority queue for loads L[k]

Load Balancing: Greedy Analysis

- **Claim.** Greedy algorithm is a 2-approximation.
- two worse than the optimal
- optimal is NP hard.
- We want to:
 - Lower bound the cost of optimal solution
 - cannot be too much worse than the optimal
- an optimal algorithm?

• To show this, we need to show greedy solution never more than a factor

• **Challenge.** We don't know the optimal solution. In fact, finding the

A good enough lower bound can help show that our algorithm

• In our problem, what are some lower bounds on the makespan of even

Load Balancing: Greedy Analysis

- Let OPT be the optimal makespan.
- Lemma. OPT $\geq \max_{j} t_{j}$.
- Any other lower bounds?

Lemma. OPT
$$\geq \frac{1}{m} \sum_{j} t_{j}$$

Proof. •

The total processing time is $\sum t_j$

• Some machine must do a 1/m fraction of the total work.

Proof. Some machine must process the most time-consuming job.

- **Proof.** Consider load L[i] of bottleneck machine imachine that ends up with highest load
 - Let j be the last scheduled job on machine i
 - When job *j* was assigned to machine *i*, *i* must have had the smallest load
 - That is, $L[i] t_i \le L[k] \quad \forall 1 \le k \le m$



- **Proof.** Consider load L[i] of bottleneck machine i
 - Let *j* be the last scheduled job on machine *i*
 - When job *j* was assigned to machine *i*, *i* must have had the smallest load
 - That is, $L[i] t_j \le L[k] \quad \forall 1 \le k \le m$

• Summing over all
$$k$$
 and d

$$L[i] - t_j \leq \frac{1}{m} \sum_{k} L[k]$$

$$\leq \frac{1}{m} \sum_{j'} t_{j'}$$

$$\leq OPT$$

- diving by m

- · Proof.
- Consider load L(i) of bottleneck machine i
- $L[i] t_j \leq OPT$
- We know that $t_j \leq \text{OPT}$
- Thus, $L = L[i] \leq OPT + t_j$

≤ 20pt ■



- Is our analysis tight?
- Close to it.
- Consider m(m-1) jobs of length 1 and 1 job of length m
- How would greedy schedule these jobs?
 - Greedy will evenly divide the first m(m-1) jobs among m machines, will place the final long job on any one machine
 - Makespan: m 1 + m = 2m 1
- How would optimal schedule it?
 - Give the long job to one machine, the rest split the other small jobs with a makespan *m*
- Ratio: $(2m-1)/m \approx 2$

Greedy is Online

- Notice that our greedy algorithm is an online algorithm
- Assigns jobs to machines in the order they arrive
 - Does not depend on future jobs
- Can we do better, if we assume all jobs are available at start time?
- **Offline.** Slight modification of greedy gets better approximation!

Improving on Online Greedy

- giant job at the end messed things up
- What can we do to avoid this?
 - Idea: deal with larger jobs first
 - Small jobs can only hurt so much
- Turns out this improves our approximation factor
- Longest-processing-time (LPT) first. Sort n jobs in decreasing order of processing times; then run the greedy algorithm on them
- Claim. LPT has a makespan at most $1.5 \cdot \text{OPT}$
- **Observation.** If we have fewer than *m* jobs, then the greedy solution is clearly optimal (as it puts each job on its own machine)

Worst case of our greedy algorithm: spreading jobs out evenly when a

LPT-first is a 1.5-Approximation

- **Lemma.** LPT-first has a makespan at most $1.5 \cdot \text{OPT}$
- **Observation.**
 - If we have fewer than *m* jobs, then the greedy solution is clearly optimal (as it puts each job on its own machine)
- Claim. If more than *m* jobs then, $OPT \ge 2 \cdot t_{m+1}$
- **Proof.** Consider the first m + 1 jobs in sorted order.
 - They each take at least t_{m+1} time
 - m + 1 jobs and m machines, there must be a machine with at least two jobs
 - Thus the optimal makespan OPT $\geq 2 \cdot t_{m+1}$

LPT-first is a 1.5-Approximation

- Lemma. LPT-first has a makespan at most $1.5\cdot \text{OPT}$
- **Proof.** Similar to our original proof. Consider the machine M_i that has the maximum load
- If M_i has a single job, then our algorithm is optimal
- Suppose M_i has at least two jobs and let t_j be the last job assigned to the machine, note that $j \geq m+1$ (why?)

Thus,
$$t_j \leq t_{m+1} \leq \frac{1}{2}$$
OPT



LPT-first is a 1.5-Approximation

- Lemma. LPT-first has a makespan at most $1.5\cdot \text{OPT}$
- **Proof.** Similar to our original proof. Consider the machine M_i that has the maximum load
- If M_i has a single job, then our algorithm is optimal
- Suppose M_i has at least two jobs and let t_j be the last job assigned to the machine, note that $j \geq m+1$ (why?)
- Thus, $t_j \le t_{m+1} \le \frac{1}{2}$ OPT

•
$$L[i] - t_j \leq OPT$$

• $L[i] \leq \frac{3}{2}OPT$



Is our 1.5-Approximation tight?

- Question. Is out 3/2-approximation analysis tight?
 - Turns out, no
- Theorem [Graham 1969]. LPT-first is a 4/3-approximation.
 - Proof via a more sophisticated analysis of the same algorithm
- Question. Is the 4/3-approximation analysis tight?
 - Pretty much.
- Example \bullet
 - m machines, n = 2m + 1 jobs
 - 2 jobs each of length $m, m + 1, \dots, 2m 1$ + one job of length m• Approximation ratio = $(4m - 1)/3m \approx 4/3$

Can we do better than 4/3?

- Long series of improvements
- Shmoys 87]
- Specifically: $(1 + \epsilon)$ approximation
- **PTAS:** Polynomial time approximation scheme
- For any desired constant-factor approximation, there exists a polynomial-time algorithm

• Polynomial time algorithm for *any* constant approximation [Hochbaum]

tion in
$$O\left((n/\epsilon)^{1/\epsilon^2}\right)$$
 time

Acknowledgments

- Some of the material in these slides are taken from
 - Kleinberg Tardos Slides by Kevin Wayne (<u>https://www.cs.princeton.edu/~wayne/kleinberg-tardos/pdf/04GreedyAlgorithmsl.pdf</u>)
 - Jeff Erickson's Algorithms Book (<u>http://jeffe.cs.illinois.edu/teaching/</u> <u>algorithms/book/Algorithms-JeffE.pdf</u>)
 - Lecture slides: <u>https://web.stanford.edu/class/archive/cs/cs161/</u> <u>cs161.1138/</u>