# CSE 373 <br> Data Structures \& Algorithms 

## Lectures 24

Final Review

## Third M idterm (a.k.a. Final)

- Friday, 12:30-1:30, here in class
- Logistics: Closed Book
- Comprehensive
- Everything up to and including Network Flow
- Not the material we will cover this Wednesday


## B+Trees

(book calls these B-trees)

- Each internal still has (up to) M-1 keys:
- Order property:
- subtree between two keys $x$ and $y$ contain leaves with values $v$ such that $x \leq v<y$
- Note the " $\leq$ "
- Leaf nodes contain up to L sorted keys.



## B+Tree Structure Properties

Root (special case)

- has between 2 and $\boldsymbol{M}$ children (or root could be a leaf)


## Internal nodes

- store up to M-1 keys

Nodes are at least $1 / 2$ full

- have between $\lceil M / 2\rceil$ and $M$ children

Leaf nodes

- where data is stored
- all at the same depth

Leaves are at least $1 / 2$

- contain between $\lceil L / 2\rceil$ and $L$ data iteinds


## B+Tree: Example $B+$ Tree with $M=4$ (\#pointers in internal hode)

 and $L=5 \quad$ (\#data items in leaf)

Défirificion for later: "neigh6oris is effie nextisibling to the left or righe.

## B+trees vs. AVL trees

Suppose again we have $\mathrm{n}=2^{30} \approx 10^{9}$ items:

- Depth of AVL Tree
- Depth of $B+$ Tree with $M=256, L=256$

$$
\text { Log_128 } 10^{\wedge 9}=4.3
$$

So let's see how we do this...

## Thinking about B+Trees

- B+Tree insertion can cause (expensive) splitting and propagation up the tree
- B+Tree deletion can cause (cheap) adoption or (expensive) merging and propagation up the tree
- Split/merge/propagation is rare if $\boldsymbol{M}$ and $\boldsymbol{L}$ are large (Why?)
- Pick branching factor $\boldsymbol{m}$ and data items/leaf L such that each node takes one full page/block of memory/disk.

Only $1 / \mathrm{L}$ inserts cause split, only $1 / \mathrm{M}$ of these go up!

## Hash Tables

- Find, insert, delete: constant time on average!
- A hash table is an array of some fixed size.
- General idea:

hash function: index $=\mathbf{h}(\mathbf{K})$

TableSize - 1

key space (e.g., integers, strings)

## Separate Chaining

| 0 | 10 |  |
| :---: | :---: | :---: |
| 1 |  | 22 |
|  |  | Thoughs about this? ${ }^{107}$ |
| 2 | 42, 12, 22 | Our goal is to teep it such that $\quad 12$ |
| 3 |  |  |
| 4 |  | Separate chaining: All |
| 5 |  | keys that map to |
| 6 |  | the same hash |
| 7 | 107 | value are kept in a |
| 8 |  | list (or "bucket"). |
| 9 |  |  |

## Open Addressing



## Linear Probing

$$
f(i)=i
$$

- Probe sequence:
$0^{\text {th }}$ probe $=h(K) \%$ TableSize
$1^{\text {th }}$ probe $=(\mathrm{h}(\mathrm{K})+1) \%$ TableSize
$2^{\text {th }}$ probe $=(\mathrm{h}(\mathrm{K})+2) \%$ TableSize
ith probe $=(\mathrm{h}(\mathrm{K})+\mathrm{i}) \%$ TableSize


## Linear Probing - Clustering

 no collision $\xrightarrow[\text { Lu Liflell }]{\longrightarrow}$ collision in
 Lu el Lu UL -
[R. Sedgewick]
collision in large cluster

## Quadratic Probing

Less likely
to encounter
Primary
Clustering

- Probe sequence:
$0^{\text {th }}$ probe $=h(K) \%$ TableSize
$1^{\text {th }}$ probe $=(\mathrm{h}(\mathrm{K})+1) \%$ TableSize
$2^{\text {th }}$ probe $=(h(K)+4) \%$ TableSize
$3^{\text {th }}$ probe $=(h(K)+9) \%$ TableSize


## $\mathrm{i}^{\text {th }}$ probe $=\left(\mathrm{h}(\mathrm{K})+\mathrm{i}^{2}\right) \%$ TableSize

## Double Hashing

Idea: given two different (good) hash functions h(K) and $g(K)$, it is unlikely two keys to collide with both.

So...let's try probing with a second hash function:

$$
f(i)=i * g(K)
$$

- Probe sequence:
$0^{\text {th }}$ probe $=h(\mathrm{~K}) \%$ TableSize
$1^{\text {th }}$ probe $=(h(K)+g(K)) \%$ TableSize
$2^{\text {th }}$ probe $=\left(h(K)+2^{*} g(K)\right) \%$ TableSize
$3^{\text {th }}$ probe $=\left(h(K)+3^{*} g(K)\right) \%$ TableSize
$i^{\text {th }}$ probe $=\left(\mathrm{h}(\mathrm{K})+\mathrm{i}^{*} \mathrm{~g}(\mathrm{~K})\right) \%$ TableSize


## Deletion in Separate Chaining

How do we delete an element with separate chaining?

Easy, just delete the item from the bucket

## Deletion in Open Addressing

Can we do something similar for open addressing?

- Delete
- Find
- Insert

|  |  |
| :---: | :---: |
|  |  |
|  | 16 |
| 3 | X |
|  | 59 |
| 5 |  |
| 6 | 76 |

$\mathrm{h}(\mathrm{k})=\mathrm{k} \% 7$
Linear probing
Delete(23)
Find(59)
Insert(30)

Need to leave a marker of a deletion

## Rehashing

Idea: When the table gets too full, create a bigger table (usually $2 x$ as large) and hash all the items from the original table into the new table.

- When to rehash?
- Separate chaining: full $(\lambda=1)$
- Open addressing: half full $(\lambda=0.5)$
- When an insertion fails
- Some other threshold
- Cost of a single rehashing?


## Why Sort?

- Allows binary search of an N -element array in $\mathrm{O}(\log \mathrm{N})$ time
- Allows O(1) time access to kth largest element in the array for any $k$
- People tend to like their output sorted
- Sorting algorithms are a frequently used and heavily studied family of algorithms in computer science


## Stability

## A sorting algorithm is stable if:

- Items in the input with the same value end up in the same order as when they began.

Input

| Adams | 1 | Adams | 1 |
| :--- | :--- | :--- | :--- |
| Black | 2 | Smith | 1 |
| Brown | 4 | Black | 2 |
| Jackson | 2 | Jackson | 2 |
| Jones | 4 | Washington | 2 |
| Smith | 1 | White | 3 |
| Thompson | 4 | Wilson | 3 |
| Washington | 2 | Brown | 4 |
| White | 3 | Jones | 4 |
| WWilson | 3 | CThempsen Dan situ |  |

Unstable sort
Adams 1
Smith $\quad 1$
Washington 2
Jackson 2
Black 2
White 3
Wilson 3
Thompson 4
Brown 4
Jones [Sedgewick]

## Sorting: The Big Picture

## Given n comparable elements in an array, sort them in an increasing order.

Simple
algorithms:
$\mathrm{O}\left(n^{2}\right)$

Insertion sort
Selection sort
Bubble sort

Fancier Comparison Specialized Handling algorithms: lower bound: algorithms: huge data
$\mathrm{O}(n \log n) \quad \Omega(n \log n)$


Heap sort
Binary tree sort
Merge sort
Quick sort (avg case)

Bucket sort External
Radix sort sorting
...

## Selection Sort: Idea

1. Find the smallest element, put it $1^{\text {st }}$
2. Find the next smallest element, put it $2^{\text {nd }}$
3. Find the next smallest, put it $3^{\text {rd }}$
4. And so on ...

## Bubble Sort Idea

- Take a pass through the array
- If a pair of neighboring elements are out of sort order, swap them.
- Take passes until no swaps are needed at any point in the pass.


## Insertion Sort: Idea

1. One element is by definition sorted
2. Sort first 2 elements.
3. Insert $3^{\text {rd }}$ element in order.

- (First 3 elements are now sorted.)

4. Insert $4^{\text {th }}$ element in order

- (First 4 elements are now sorted.)

5. And so on...

## Heap Sort: Sort with a Binary Heap



Use a max-heap, do it in-place

## Runtime: $\mathrm{O}(\mathrm{n} \log \mathrm{n})$

## "Divide and Conquer"

- Two divide and conquer sorting methods:
- Idea 1: Divide array into two halves, recursively sort left and right halves, then merge two halves $\rightarrow$ known as M ergesort
- Idea 2 : Partition array into small items and large items, then recursively sort the two smaller portions $\rightarrow$ known as Quicksort


## M ergesort

- Divide it in two at the midpoint
- Conquer each side in turn (by recursively sorting)
- Merge two halves together


## M ergesort Example



## Iterative M ergesort



## Quicksort

- Quicksort uses a divide and conquer strategy, but does not require the $\mathrm{O}(\mathrm{N})$ extra space that M ergeSort does
- Partition array into left and right sub-arrays
- the elements in left sub-array are all less than pivot
- elements in right sub-array are all greater than pivot
- Recursively sort left and right sub-arrays
- Concatenate left and right sub-arrays in O(1) time


## Quicksort Example



## So Which Is Best?

- It's a trick question, a naïve question
- M yth: "Quicksort is the best in-memory sorting algorithm."
- Mergesort and Quicksort make different tradeoffs regarding the cost of comparison and the cost of a swap
- M ergesort is also the basis for external sorting algorithms (large N sorting)


## Permutations

- How many possible orderings are there?
- Example: a, b, c

$$
\begin{array}{lll}
\mathrm{a}<\mathrm{b}<\mathrm{c} & \mathrm{~b}<\mathrm{a}<\mathrm{c} & \mathrm{c}<\mathrm{a}<\mathrm{b} \\
\mathrm{a}<\mathrm{c}<\mathrm{b} & \mathrm{~b}<\mathrm{c}<\mathrm{a} & \mathrm{c}<\mathrm{b}<\mathrm{a}
\end{array}
$$

## Decision Tree



The leaves contain all the possible orderings of $\mathrm{a}, \mathrm{b}, \mathrm{c}$

## Lower bound on Height

- The decision tree has how many leaves:

$$
L=N!
$$

- A binary tree with L leaves has height at least:

$$
h \geq \log _{2} L
$$

- So the decision tree has height:

$$
h \geq \log _{2}(N!)
$$

## $\log (N!)$

$$
\begin{aligned}
& \log (N!)=\log (N \cdot(N-1) \cdot(N-2) \cdots(2) \cdot(1)) \\
& =\log N+\log (N-1)+\log (N-2)+\cdots+\log 2+\log 1 \\
& \geq \log N+\log (N-1)+\log (N-2)+\cdots+\log \frac{N}{2} \\
& \geq \frac{N}{2} \log \frac{N}{2} \\
& \geq \frac{N}{2}(\log N-\log 2)=\frac{N}{2} \log N-\frac{N}{2} \log 2 \\
& =\Omega(N \log N)
\end{aligned}
$$

## $\Omega(\mathrm{N} \log \mathrm{N})$

- No matter how clever you are about which comparisons you perform, your sorting algorithm with always be $\Omega(N \log N)$
- Your worst case will be at least $\mathrm{N} \log \mathrm{N}$
- Proving this saves us the trouble of trying to do better than this, because we cannot
- Now that's some Computer Science


## Doing Better

- So how can we do better?
- Need to dodge one of the proof's assumptions
- What's our proof based in?
- Comparisons
- Can we sort without doing comparisons?


## BucketSort (aka BinSort)

If all values are known to be between 1 and K , create an array count of size $K$, increment counts while traversing the input, and finally output the result.

Example $\mathrm{K}=5$. Input $=(5,1,3,4,3,2,1,1,5,4,5)$

| count array |  |
| :--- | :--- |
| 1 | 3 |
| 2 | 1 |
| 3 | 2 |
| 4 | 2 |
| $5_{\text {2077200 }} 3$ |  |



Running time to sort $\mathbf{n}$ items?
CSE 373 Fall 2009 -- Dan Suciu $\mathbf{N}+\mathbf{K}$

## RadixSort

- Radix ="The base of a number system"
- We'll use 10 for convenience
- Use a larger number in any implementation
- ASCII Strings, for example, might use 128
- Idea:
- BucketSort on one digit at a time
- Requires stable sort!
- After sort k, the last k digits are sorted
- Set number of buckets: $B=$ radix.


## RadixSort

- Input:126, 328, 636, 341, 416, 131, 328

BucketSort on 1sd:

|  | 341 <br> 131 |  |  |  |  | 126 <br> 636 <br> 416 |  | 328 <br> 328 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |

BucketSort on next-higher digit:

|  | 416 | 126 <br> 328 <br> 328 | 131 <br> 636 | 341 |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |

BucketSort on msd:

|  | 126 <br> 131 | 328 <br> 328 <br> 341 | 416 |  | 636 |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |

Output: $126,131,328,328,34$ CSE 373 Fall 2009 - - Dan Suciu

## Summary of sorting

O(N²) average, worst case:

- Selection Sort, Bubblesort, Insertion Sort

O(N $\log \mathrm{N}$ ) average case:

- Heapsort: In-place, not stable.
- Mergesort: O(N) extra space, stable, massive data.
- Quicksort: claimed fastest in practice, but $\mathrm{O}\left(\mathrm{N}^{2}\right)$ worst case. Recursion/ stack requirement. Not stable.
$\Omega(\mathrm{N} \log \mathrm{N})$ worst and average case:
- Any comparison-based sorting algorithm

O(N)

- Radix Sort: fast and stable. Not comparison based. Not inplace. Poor memory locality can undercut performance.


## Disjoint Set ADT

- Data: set of pairwise disjoint sets.
- Required operations
- Union - merge two sets to create their union
- Find - determine which set an item appears in
- A common operation sequence:
- Connect two elements if not already connected: if (Find $(x)!=$ Find $(y))$ then Union $(x, y)$


## Up-Tree for DU/F

Initial state

(7)

Intermediate state


## Find Operation

- Find $(x)$ : follow $x$ to the root and return the root


Find $(6)=7$


## Union Operation

- Union(i,j): assuming i and j roots, point i to $j$.



## Simple Implementation

- Array of indices

$\mathrm{Up}[\mathrm{x}]=0$ means x is a root.



## Weighted Union

- Weighted Union
- Instead of arbitrarily joining two roots, always point the smaller root to the larger root



## Elegant Array Implementation



## Elegant Array Implementation



Instead of a separate weight array, can re-use the empty parent reference

## Path Compression

- On a Find operation point all the nodes on the search path directly to the root.



## Graphs

Formalism representing relationships among objects
Graph $\mathbf{G}=(\mathbf{v}, \mathbf{E})$

- Set of vertices (aka nodes):
$\mathrm{v}=\left\{\mathrm{v}_{1}, \mathrm{v}_{\mathbf{2}}, \ldots, \mathrm{v}_{\mathrm{n}}\right\}$
- Set of edges:

$\mathbf{E}=\left\{\mathbf{e}_{1}, \mathbf{e}_{2}, \ldots, \mathbf{e}_{\mathrm{m}}\right\}$
where each $\mathbf{e}_{i}$ connects one
$\mathrm{E}=\underset{\text { (Han, Leia) }}{\text { (Heia) },}$

vertex to another ( $\mathbf{v}_{\mathbf{j}}, \mathbf{v}_{\mathbf{k}}$ )
 (Leia, Han) \}
Graphs can be directed or undirected


## Undirected Graphs

In undirected graphs, edges have no specific direction (edges are always two-way):


Thus, $(u, v) \in$ implies $(v, u) \in$. Only one of these edges needs to be in the set; the other is implicit.

Degree of a vertex: number of edges containing that vertex. (Same as number of adjacent vertices.)

## Directed Graphs

In directed graphs (aka digraphs), edges have a specific direction:


Thus, $(u, v) \in E$ does not imply $(v, u) \in E$.
In-degree of a vertex: number of inbound edges. Out-degree of a vertex : number of outbound edges.

## Paths and Cycles

- A path is a list of vertices $\left\{v_{1}, v_{2}, \ldots, v_{n}\right\}$ such that $\left(\mathbf{v}_{\mathrm{i}}, \mathbf{v}_{\mathrm{i}+1}\right) \in \mathrm{E}$ for all $\mathbf{0} \leq \mathbf{i}<\mathbf{n}$.
- A cycle is a path that begins and ends at the same node.

 Seattle,


## Directed Acyclic Graphs (DAGs)

DAGs are directed graphs with no (directed) cycles.

Aside: If program's call-graph is a DAG, then all procedure calls
can be in-lined

$\{$ Rooted, directed tree $\} \subset\{$ DAG $\} \subset\{$ Graph $\}$

## |E| and |V|

How many edges | $\mathrm{E} \mid$ in a graph with | $\mathrm{V} \mid$ vertices?

$$
0 \leq|\mathrm{E}| \leq|\mathrm{V}|^{2}
$$

What if the graph is directed?

$$
0 \leq|\mathrm{E}| \leq 2|\mathrm{~V}|^{2}
$$

What if it is undirected and connected?

$$
|\mathrm{V}|-1 \leq|\mathrm{E}| \leq|\mathrm{V}|^{2}
$$

Can the following bounds be simplified?

- Arbitrary graph: $\mathrm{O}\left(|\mathrm{E}|+|\mathrm{V}|^{2}\right)$
$\mathrm{O}\left(|\mathrm{V}|^{2}\right)$
- Undirected, connected: O(|E| log| V| +|V| log| V| )
$\mathrm{O}(|\mathrm{E}| \log |\mathrm{V}|)$
Some (semi-standard) terminology:
- A graph is sparse if it has $\mathrm{O}(|\mathrm{V}|)$ ) edges (upper bound).
- A graph is dense if it has $\Theta\left(|\mathrm{V}|^{2}\right)$ edges.


## Representation 1: Adjacency M atrix

A $|\mathrm{v}| \mathbf{x}|\mathrm{v}|$ matrix $\mathbf{m}$ in which an element $\mathbf{M}_{[\mathbf{u}, \mathrm{v}]}$ is true if and only if there is an edge from uto v


Runtimes:
Iterate over vertices?
Iterate over edges?
$\mathrm{O}(|\mathrm{V}|)$
$\mathrm{O}\left(|\mathrm{V}|^{2}\right)$ Space requirements? $\mathrm{O}\left(|\mathrm{V}|^{2}\right)$
Iterate edges adj. to vertex? $\mathrm{O}(|\mathrm{V}|)$ Best for what kinds of graphs?
Existence of edge?
CSE ЗО(1) 1009 - - Dan Suciu
dense

## Representation 2: Adjacency List

A list (array) of length $|\mathrm{v}|$ in which each entry stores a list (linked list) of all adjacent vertices


Runtimes:
Iterate over vertices?
Iterate over edges?
$\mathrm{O}(|\mathrm{V}|)$
Iterate edges adj. to vertex? $\mathrm{O}(\mathrm{d})$ Best for what kinds of graphs?
Existence of edge?
CSE 3 $\bigcirc$ (d)
sparse

## Application: Topological Sort

Given a graph, $\mathbf{g}=(\mathbf{v}, \mathbf{e})$, output all the vertices in $v$ sorted so that no vertex is output before any other vertex with an edge to it.


What kind of input graph is ${ }_{9}$ allowed?

Is the output unique?
CSE 373 FalNo, offericicalled a partial ordering ${ }^{59}$

## Topological Sort: Take Two

1. Label each vertex with its in-degree
2. Initialize a queue Q to contain all in-degree zero vertices
3. While Q not empty
a. $v=Q$. dequeue; output $v$
b. Reduce the in-degree of all vertices adjacent to v
c. If new in-degree of any such vertex $u$ is zero Q.enqueue(u)

> Is the use of a queue here important?

## Runtime:

$\mathrm{O}\left(\mathrm{V} \mathrm{V}_{20 \mathrm{O} 9}+|\mathrm{E}|\right)$
No, can use a stack, list, set, box, etc.
Changes behaviorpobutnresult is still topological sort

## Comparison: DFS versus BFS

- Breadth-first search
-Always finds shortest paths - optimal solutions
-M arking visited nodes can improve efficiency, but even without this search guaranteed to terminate
- Depth-first search
-Does not always find shortest paths
-M ust be careful to mark visited vertices, or you could go into an infinite loop if there is a cycle
- Is BFS always preferable?


## Single Source Shortest Paths (SSSP)

- Given a graph $G$, edge costs $c_{i, j}$, and vertex s, find the shortest paths from s to all vertices in $G$.
- Is finding paths to all the vertices harder or easier than the previous problem?
- The same difficulty
(imagine the one we want is the last one we reach)
- But we still haven't dealt with edge costs...


## Dijkstra’s Algorithm: Idea



At each step:

1) Pick closest unknown vertex
2) Add it to known vertices
3) Update distances

## Dijkstra's Algorithm: Pseudocode

Initialize the cost of each node to $\infty$
Initialize the cost of the source to 0
While there are unknown nodes left in the graph
Select an unknown node $b$ with the lowest cost
Mark b as known
For each node a adjacent to b
if b's cost +cost of ( $b, a$ ) <a's old cost
a's cost =b's cost +cost of (b, a)
a's prev path node $=b$

## Dijkstra's Algorithm in action



| Vertex | Visited? | Cost | Found by |
| :---: | :---: | :---: | :---: |
| A | Y | 0 |  |
| B | Y | 2 | A |
| C | Y | 1 | A |
| D | Y | 4 | A |
| E | Y | 11 | G |
| F | Y | 4 | B |
| G | Y | 8 | H |
| H | CSES23Eall |  | F |

```
void Graph::dijkstra(Vertex s) {
```

    Vertex v,w;
    Initialize s.dist $=0$ and set dist of all other vertices to infinity

```
while (there exist unknown vertices, find the
``` one \(b\) with the smallest distance)
b.known = true;
for each a adjacent to b
Sounds Lif (!a.known)

Sounds like deleteMin on a heap.
like if (b.dist + weight (b,a) < a.dist) \(\{\) a.dist \(=\) (b.dist + weight (b,a)); a.path \(=\) b;
\}
\}
\}
Running time: \(\mathrm{O}(|\mathrm{E}| \log |\mathrm{V}|)\) - there are \(|\mathrm{E}|\) edges to examine,


\section*{Correctness: The Cloud Proof}


How does Dijkstra's decide which vertex to add to the Known set next?
- If path to v is shortest, path to \(\mathbf{w}\) must be at least as long
(or else we would have picked w as the next vertex)
-20750cthe path through w to veannet-be any shorter!

\section*{Follow-On Question}
- What if I had multiple potential start points, and need to know the minimum cost of reaching each node from any start point?
- Can do this by changing the algorithm
- Add each start point to initial queue with cost 0
- If the algorithm is encapsulated (and highly tuned for efficiency), this seems bad
- You need to re-implement the whole thing
- Your implementation probably isn't as good

\section*{Thinking About Graph Structure}
- Working with graphs is often a problem of setting up the right graph so that you can apply the unmodified standard algorithm
- Change the graph, apply the encapsulated and optimized SSSP implementation
- Add a meta-start node
- Include 0 cost edges from it to the start nodes

\section*{Floyd-Warshall}
for (int \(k=1 ; k=<\mathrm{V}\); \(\mathrm{k}+\) +)
for (int \(i=1 ; i=<v ; i++)\)
```

for (int j = 1; j =< V; j++)

```
    if ( ( M[i][k]+ M[k][j] ) < M[i][j] )
    M[i][j] = M[i][k]+ M[k][j]

Invariant: After the kth iteration, the matrix includes the shortest paths for all pairs of vertices ( \(\mathrm{i}, \mathrm{j}\) ) containing only vertices \(1 . . \mathrm{k}\) as intermediate vertices
Simple for loop implementation intended to be fast (especially with the help of a modern compiler). Does not bother with if statements to filter out comparisons that will never result in a change.

\section*{Problem: Large Graphs}

It is expensive to find optimal paths in large graphs, using BFS or Dijkstra's algorithm (for weighted graphs)

How can we search large graphs efficiently by using "commonsense" about which direction looks most promising?

\section*{M inimum Spanning Trees}

Given an undirected graph \(\mathbf{G}=(\mathbf{V}, \mathrm{E})\), find a graph \(\mathbf{G}^{\prime}=\left(V, E^{\prime}\right)\) such that:
\(-E^{\prime}\) is a subset of \(E\)
\(-\left|\mathrm{E}^{\prime}\right|=|\mathrm{V}|-1\)
\(-G^{\prime}\) is connected
\(\mathrm{G}^{\prime}\) is a minimum
spanning tree.
\(-\sum_{(u, v) \in E^{\prime}} \mathrm{c}_{u v}\) is minimal

\section*{Reducing Best to Minimum}

Let \(\mathrm{P}(\mathrm{e})\) be the probability that an edge doesn't fail. Define:
\[
C(e)=-\log _{10}(P(e))
\]

Minimizing \(\quad \sum_{e \in T} C(e)\)
is equivalent to maximizing \(\quad \prod_{e \in T} P(e)\)
because \(\prod_{e \in T} P(e)=\prod_{e \in T} 10^{-C(e)}=10^{-\sum_{e \in T} C(e)}\)

\section*{Example of Reduction}


Best Spanning Tree Problem


Minimum Spanning Tree Problem

\section*{Two Different Approaches}


Prim's Algorithm
Looks familiar!


Kruskals's Algorithm Completely different!

\section*{Prim's algorithm}

Idea: Grow a tree by adding an edge from the "known" vertices to the "unknown" vertices. Pick the edge with the smallest weight.


\section*{Prim's Algorithm for M ST}

\section*{A node-based greedy algorithm Builds MST by greedily adding nodes}
1. Select a node to be the "root"
- mark it as known
- Update cost of all its neighbors
2. While there are unknown nodes left in the graph
a. Select an unknown node b with the smallest cost to reach from some known node a
b. Mark b as known
c. Add \((a, b)\) to M ST
d. Update cost of all nodes adjacent to b

\section*{Find M ST using Prim's}

Start with \(\mathbf{V}_{1}\)
\begin{tabular}{|l|l|l|l|}
\hline V & Kwn & Distance & path \\
\hline v1 & Y & - & - \\
\hline v2 & Y & 2 & V1 \\
\hline v3 & Y & 2 & V4 \\
\hline v4 & Y & 1 & V1 \\
\hline v5 & Y & 6 & V7 \\
\hline v6 & Y & 1 & V7 \\
\hline v7 & Y & 4 & V4 \\
\hline
\end{tabular}


> Order Declared Known: V1, V4, V2, V3, V7, V6, V5

Selected Edges:
\{V2, V1\}, \{V3, V4\}, \{V4, V1\}, \{V5, V7\}, \{V6, V7\}, \{V7, V4\}

\section*{Kruskal's M ST Algorithm}

Idea: Grow a forest out of edges that do not create a cycle. Pick an edge with the smallest weight.


\section*{Kruskal's Algorithm for M ST}

\section*{An edge-based greedy algorithm Builds MST by greedily adding edges}
1. Initialize with
- empty MST
- all vertices marked unconnected
- all edges unmarked
2. While there are still unmarked edges
a. Pick the lowest cost edge ( \(u, v\) ) and mark it
b. If \(\mathbf{u}\) and \(\mathbf{v}\) are not already connected, add ( \(u, v\) ) to the M ST and mark u and v as connected

\section*{Optimized Kruskal code}


\section*{Find M ST using Kruskal's}

- Is this MST unique?
- Under what condition is an MST unique?
- Unique edge weights guarantee uniqueness

\section*{Best-First}


\section*{Network Flow}
- So, how do we want to go about this?


\section*{Ford-Fulkerson M ethod}
- Our greedy algorithm makes choices about how to route flow, and we never reconsider those choices
- Can we develop a way to efficiently reconsider the choices we already made?
- Can we do it by just modifying the graph?

\section*{Residual Graph}
- Constructing a residual graph:
- Use the same vertices
- Edge weights are the remaining capacity on the edges, given the existing augmenting paths
- Add additional edges for backward capacity
- If there is a path from s to \(t\) in the residual graph, then there is available capacity there

\section*{Example}

Augment along AEBCD (which saturates AE and EB, and empties I


\section*{Min Cut - Example}


Capacity of cut \(=5 \quad\) CSE 373 Fall \(2009-\) - Dan Suciu

\section*{Coincidence?}
- No, M ax-flow always equals M in-cut
- If there is a cut with capacity equal to the flow, we have a maxflow:
- We can't have a flow that's bigger than the capacity cutting the graph! So any cut puts a bound on the maxflow, and if we have an equality, then we must have a maximum flow.
- If we have a maxflow, then there are no augmenting paths left
- Or else we could augment the flow along that path, which would yield a higher total flow.
- If there are no augmenting paths, we have a cut of capacity equal to the maxflow
- Pick a cut (ST) where S contains all vertices reachable in the residual graph from \(s\), and \(T\) is everything else. Then every edge from \(S\) to \(T\) must be saturated (or else there would be a path in the residual graph). So \(c(S, T)=\) \(f(S, T)=f(s, t)=|f|\) and we're done.```

