CS61B Lecture #34

earch, Minimum spanning trees, union-find.

Point-to-Point Shortest Path

- gorithm gives you shortest paths from a particular given others in a graph.
- you're only interested in getting to a particular vertex?
- algorithm finds paths in order of length, you *could* and stop when you get to the vertex you want.
- be really wasteful.
- to travel by road from Denver to a destination on lower in New York City is about 1750 miles (says Google).
- from Denver to Berkeley is about 1250 miles.
- lore much of California, Nevada, Arizona, etc. before r destination, even though these are all in the wrong
- en worse when graph is infinite, generated on the fly.

A* Search

- g for a path from vertex Denver to the desired NYC
- t we had a simple *heuristic estimate*, $h(V)$, of the length m any vertex V to NYC.
- that instead of visiting vertices in the fringe in order rtest known path to Denver, we order by the sum of plus this heuristic estimate of the remaining distance nver, $V) + h(V)$.
- rds, we look at places that are reachable from places ready know the shortest path to Denver and choose ook like they will result in the shortest trip to NYC, the remaining distance.
- ate is good, then we don't look at, say, Grand Junction est by road), because it's in the wrong direction.
- g algorithm is *A* search*.
- work, we must be careful about the heuristic.

Illustration

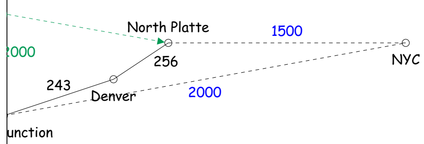

- ines indicate distances from Denver we have determined
- gorithm would have us look at Grand Junction next (smallest n Denver).
- d in the heuristic remaining distance to NYC (our goal), orth Platte instead.

missible Heuristics for A* Search

- stic estimate for the distance to NYC is too high (i.e., he actual path by road), then we may get to NYC without ng points along the shortest route.
- , if our heuristic decided that the midwest was literally f nowhere, and $h(C) = 2000$ for C any city in Michigan or l only find a path that detoured south through Kentucky.
- *missible*, $h(C)$ must never overestimate $d(C, \text{NYC})$, the h distance from C to NYC.
- hand, $h(C) = 0$ will work (what is the result?), but yield l algorithm.

missible Heuristics for A* Search

- stic estimate for the distance to NYC is too high (i.e., he actual path by road), then we may get to NYC without ng points along the shortest route.
- , if our heuristic decided that the midwest was literally f nowhere, and $h(C) = 2000$ for C any city in Michigan or l only find a path that detoured south through Kentucky.
- *missible*, $h(C)$ must never overestimate $d(C, \text{NYC})$, the h distance from C to NYC.
- hand, $h(C) = 0$ will work (what is the result?), but yield l algorithm. This is just Dijkstra's algorithm.

Summary of Shortest Paths

gorithm finds a *shortest-path tree* computing giving shortest paths in a weighted graph from a given starting ther nodes.

d =

emove V nodes from priority queue +

pdate all neighbors of each of these nodes and add or hem in queue ($E \lg E$)

$V + E \lg V) = \Theta((V+E)\lg V)$

arches for a shortest path to a *particular* target node.

kstra's algorithm, except:

n we take target from queue.

ue by estimated distance to start + heuristic guess of distance ($h(v) = d(v, \text{target})$)

must not overestimate distance and obey triangle inequality $d(b,c) \geq d(a,c)$).

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

, add the shortest edge connecting some node already o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

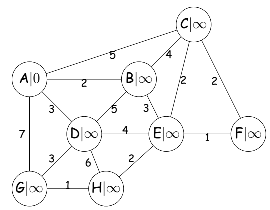

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

, add the shortest edge connecting some node already o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

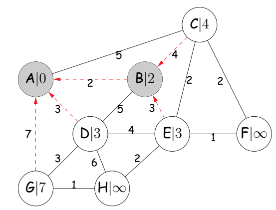

Consistency

t we estimate $h(\text{Chicago}) = 700$, and $h(\text{Springfield, IL}) =$ (Chicago, Springfield) $= 200$.

200 miles to Springfield, we guess that we are suddenly oser to NYC.

ssible, since both estimates are low, but it will mess up n.

will require that we put processed nodes back into the e our estimate was wrong.

ourse, anyway) we also require *consistent heuristics:* $+ d(A,B)$, as for the triangle inequality.

t heuristics are admissible (why?).

as the crow flies" is a reasonable $h(\cdot)$ in the trip-planning

search (and others) is in `cs61b-software` and on the machines as `graph-demo`.

Minimum Spanning Trees

en a set of places and distances between them (assume ive), find a set of connecting roads of minimum total llows travel between any two.

ou get will not necessarily be shortest paths.

that such a set of connecting roads and places must because removing one road in a cycle still allows all to

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

, add the shortest edge connecting some node already o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

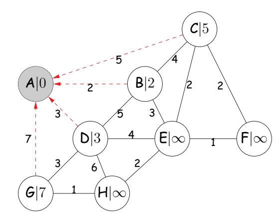

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

), add the shortest edge connecting some node already
o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
 fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

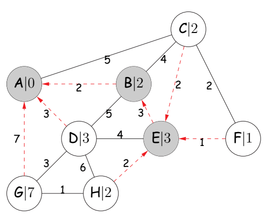

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

), add the shortest edge connecting some node already
o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
 fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

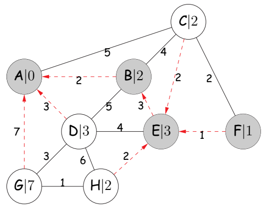

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

), add the shortest edge connecting some node already
o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
 fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

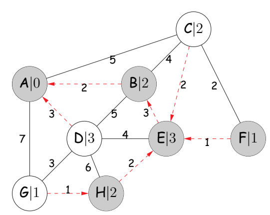

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

), add the shortest edge connecting some node already
o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
 fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

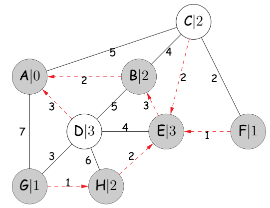

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

), add the shortest edge connecting some node already
o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
 fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

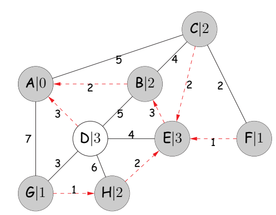

m Spanning Trees by Prim's Algorithm

ow a tree starting from an arbitrary node.

), add the shortest edge connecting some node already
o one that isn't yet.

is work?

```
inge;
{ v.dist() = ∞; v.parent() = null; }
y starting node, s;

queue ordered by smallest .dist();
 fringe;
sEmpty()) {
nge.removeFirst();

v,w) {
ge && weight(v,w) < w.dist())
 = weight(v, w); w.parent() = v; }
```

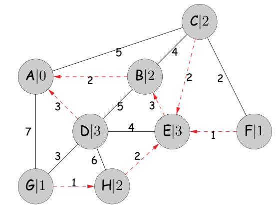

Example of Kruskal's Algorithm

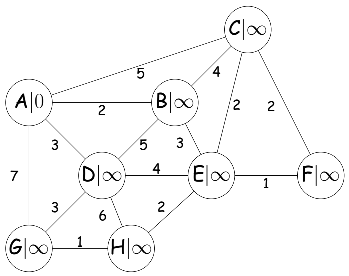

- ugh the edges in increasing order of weight (here, I phabetically).
- connect two unconnected groups get added to the tree).
- oin already-joined groups are discarded ('X'ed out here).
- a minimal spanning tree (a free tree).

Example of Kruskal's Algorithm

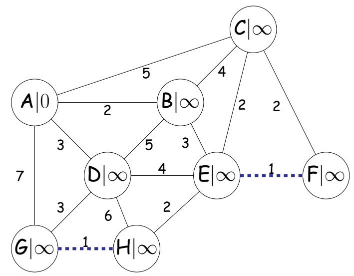

- ugh the edges in increasing order of weight (here, I phabetically).
- connect two unconnected groups get added to the tree).
- oin already-joined groups are discarded ('X'ed out here).
- a minimal spanning tree (a free tree).

Example of Kruskal's Algorithm

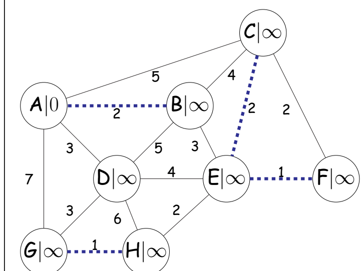

- ugh the edges in increasing order of weight (here, I phabetically).
- connect two unconnected groups get added to the tree).
- oin already-joined groups are discarded ('X'ed out here).
- a minimal spanning tree (a free tree).

Spanning Trees by Kruskal's Algorithm

- the shortest edge in a graph can always be part of a nning tree.
- have a bunch of subtrees of a MST, then the shortest nnects two of them can be part of a MST, combining rees into a bigger one.

(trivial) subtree for each node in the graph;

```
dge(v,w), in increasing order of weight {
,w) connects two different subtrees ) {
(v,w) to MST;
ine the two subtrees into one;
```


Example of Kruskal's Algorithm

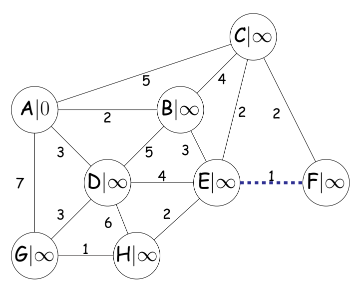

- ugh the edges in increasing order of weight (here, I phabetically).
- connect two unconnected groups get added to the tree).
- oin already-joined groups are discarded ('X'ed out here).
- a minimal spanning tree (a free tree).

Example of Kruskal's Algorithm

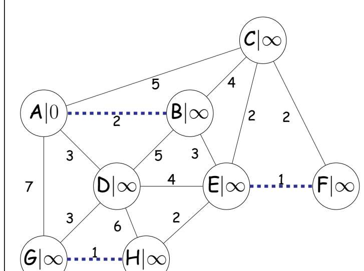

- ugh the edges in increasing order of weight (here, I phabetically).
- connect two unconnected groups get added to the tree).
- oin already-joined groups are discarded ('X'ed out here).
- a minimal spanning tree (a free tree).

Example of Kruskal's Algorithm

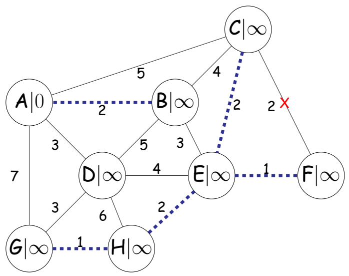

ugh the edges in increasing order of weight (here, I phabetically).

:onnect two unconnected groups get added to the tree ;).

oin already-joined groups are discarded ('X'ed out here).

a minimal spanning tree (a free tree).

Example of Kruskal's Algorithm

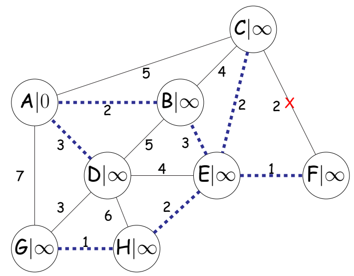

ugh the edges in increasing order of weight (here, I phabetically).

:onnect two unconnected groups get added to the tree ;).

oin already-joined groups are discarded ('X'ed out here).

a minimal spanning tree (a free tree).

Union Find

orithm required that we have a set of sets of nodes with ns:

h of the sets a given node belongs to.

vo sets with their *union*, reassigning all the nodes in the al sets to this union.

g to do is to store a set number in each node, making

es changing the set number in one of the two sets being smaller is better choice.

n individual union can take $\Theta(N)$ time.

fast?

Example of Kruskal's Algorithm

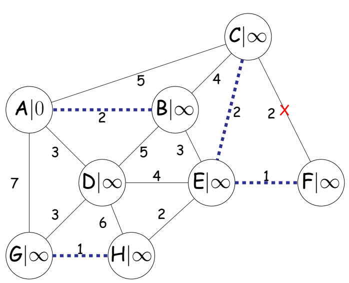

ugh the edges in increasing order of weight (here, I phabetically).

:onnect two unconnected groups get added to the tree ;).

oin already-joined groups are discarded ('X'ed out here).

a minimal spanning tree (a free tree).

Example of Kruskal's Algorithm

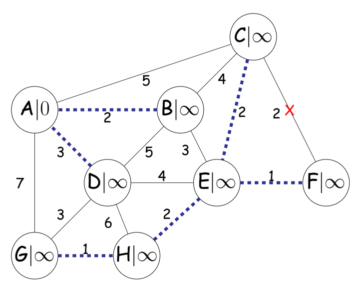

ugh the edges in increasing order of weight (here, I phabetically).

:onnect two unconnected groups get added to the tree ;).

oin already-joined groups are discarded ('X'ed out here).

a minimal spanning tree (a free tree).

Example of Kruskal's Algorithm

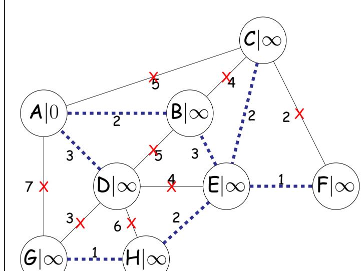

ugh the edges in increasing order of weight (here, I phabetically).

:onnect two unconnected groups get added to the tree ;).

oin already-joined groups are discarded ('X'ed out here).

a minimal spanning tree (a free tree).

Path Compression

unioning really fast, but the find operation potentially
).

ollowing trick: whenever we do a *find* operation, *compress*
the root, so that subsequent finds will be faster.

e each of the nodes in the path point directly to the

very fast, and sequence of unions and finds each have
early constant amortized time.

d 'g' in last tree (result of compression on right):

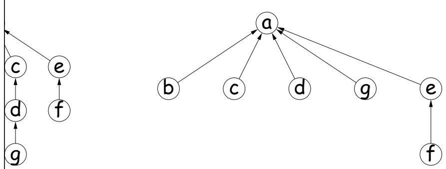

A Clever Trick

to represent a set of nodes by *one* arbitrary representative
set.

de contain a pointer to another node in the same set.

each pointer to represent the *parent* of a node in a tree
representative node as its root.

set a node is in, follow parent pointers.

such trees, make one root point to the other (choose
the larger tree as the union representative).

Two Sets

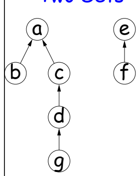

Their Union