



CSE 332 Data Abstractions: Disjoint Set Union-Find and Minimum Spanning Trees

Kate Deibel
Summer 2012

August 13,
2012

CSE 332 Data Abstractions, Summer 2012

1

Making Connections

Answering these questions is much easier if we create disjoint sets of nodes that are connected:

Start: {1} {2} {3} {4} {5} {6} {7} {8} {9}
 3-5 {1} {2} {3, 5} {4} {6} {7} {8} {9}
 4-2 {1} {2, 4} {3, 5} {6} {7} {8} {9}
 1-6 {1, 6} {2, 4} {3, 5} {7} {8} {9}
 5-7 {1, 6} {2, 4} {3, 5, 7} {8} {9}
 4-8 {1, 6} {2, 4, 8} {3, 5, 7} {9}
 3-7 *no change*

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

3

Disjoint Set Union-Find ADT

Separate elements into *disjoint* sets

- If set $x \neq y$ then $x \cap y = \emptyset$ (i.e. no shared elements)

Each set has a *name* (usually an element in the set)

union(x,y): take the union of the sets x and y ($x \cup y$)

- Given sets: {3,5,7}, {4,2,8}, {9}, {1,6}
- $\text{union}(5,1) \rightarrow \{3,5,7,1,6\}, \{4,2,8\}, \{9\}$,

find(x): return the name of the set containing x .

- Given sets: {3,5,7,1,6}, {4,2,8}, {9},
- $\text{find}(1)$ returns 5
- $\text{find}(4)$ returns 8

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

5

Making Connections

You have a set of nodes (numbered 1-9) on a network. You are given a sequence of pairwise connections between them:

3-5 4-2 1-6 5-7 4-8 3-7

Q: Are nodes 2 and 4 connected? Indirectly?

Q: How about nodes 3 and 8?

Q: Are any of the paired connections redundant due to indirect connections?

Q: How many sub-networks do you have?

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

2

Making Connections

Let's ask the questions again.

3-5 4-2 1-6 5-7 4-8 3-7
 ↓
 {1, 6} {2, 4, 8} {3, 5, 7} {9}

Q: Are nodes 2 and 4 connected? Indirectly?

Q: How about nodes 3 and 8?

Q: Are any of the paired connections redundant due to indirect connections?

Q: How many sub-networks do you have?

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

4

Disjoint Set Union-Find Performance

Believe it or not:

- We can do Union in constant time.
- We can get Find to be **amortized** constant time with worst case $O(\log n)$ for an individual Find operation

Let's see how...

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

6

Beware of Minotaurs

FIRST, LET'S GET LOST

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

7

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

8


What Makes a Good Maze?

- We can get from any room to any other room (connected)
- There is just one simple path between any two rooms (no loops)
- The maze is not a simple pattern (random)

Making a Maze

A high-level algorithm for a random maze is easy:

- Start with a grid
- Pick Start and Finish
- Randomly erase edges



August 13, 2012


CSE 332 Data Abstractions, Summer 2012

9

The Middle of the Algorithm

So far, we've knocked down several walls while others still remain.

Consider the walls between A and B and C and D



- Which walls can we knock down and maintain both our **connectedness** and our **no cycles** properties?

How do we do this efficiently?

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

10


Maze Algorithm: Number the Cells

Number each cell and treat as disjoint sets:

- $S = \{ \{1\}, \{2\}, \{3\}, \{4\}, \dots, \{36\} \}$

Create a set of all edges between cells:

- $W = \{ (1,2), (1,7), (2,8), (2,3), \dots \}$ 60 walls total.



August 13, 2012

CSE 332 Data Abstractions, Summer 2012

11

Maze Algorithm: Building with DSUF

Algorithm sketch:

- Choose a wall at random.
- Erase wall if the neighbors are in disjoint sets (this avoids creating cycles)
- Take union of those cell's sets
- Repeat until there is only one set
 - Every cell is thus reachable from every other cell


August 13, 2012

CSE 332 Data Abstractions, Summer 2012

12

The Secret To Why This Works

Notice that a connected, acyclic maze is actually a Hidden Tree



This suggests how we should implement the Disjoint Set Union-Find ADT

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

13

I promise the first twenty minutes of this section will not be the saddest trees you have ever seen...

IMPLEMENTING DSUF WITH UP TREES

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

14

Up Trees for Disjoin Set Union-Find

Up trees

- Notes point to parent, not children
- Thus only one pointer per node


In a DSUF

- Each disjoint set is its own up tree
- The root of the tree is the name for the disjoint set

Initial State



After Unions



August 13, 2012


CSE 332 Data Abstractions, Summer 2012

15

Find Operation

find(x): follow x to the root and return the root (the name of the disjoint set)

- find(1) = 1
- find(3) = 3
- find(4) = 1
- find(6) = 7



August 13, 2012


CSE 332 Data Abstractions, Summer 2012

16

Find Operation

union(i,j): assuming i and j are roots, point root i to root j

union(1,7)



What if i or j is not a root?

- Run a find on i and j first and use the returned values for the joining

Why do we join roots and not just the nodes?

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

17

Simple Implementation

Once again, it is better to implement a tree using an array than with node objects

- Leave up[0] empty (or # of disjoint sets)
- up[x] = i means node x's parent is node i
- up[x] = 0 means x is a root



August 13, 2012

CSE 332 Data Abstractions, Summer 2012

18

Performance

Using array-based up trees, what is the cost for

- union(i,j)?
- find(x)?

union(i,j) is O(1) if i and j are roots
 ▪ Otherwise depends on cost of find

find(x) is O(n) in worst-case
 ▪ What does the worst-case look like?



August 13, 2012

CSE 332 Data Abstractions, Summer 2012

19

Performance – Doing Better

The problem is that up trees get too tall

In order to make DSUF perform as we promised, we need to improve both our union and find algorithms:

- Weighted Union
- Path Compression

Only with BOTH of these will we get find to average-case O(log n) and amortized O(1)

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

20


Weighted Union

Instead of arbitrarily joining two roots, always point the smaller tree to the root of the larger tree

- Each up tree has a weight (number of nodes)
- The idea is to limit the height of each up tree
- Trees with more nodes tend to be deeper

Union by rank or height are similar ideas but more complicated to implement

union(1,7)




August 13, 2012

CSE 332 Data Abstractions, Summer 2012

21

Weighted Union Implementation

We can just use an additional array to store weights of the roots...



August 13, 2012


CSE 332 Data Abstractions, Summer 2012

22

Weighted Union Implementation

... or we use negative numbers to represent roots and their weights

But generally, saving O(n) space is not critical



August 13, 2012

CSE 332 Data Abstractions, Summer 2012

23

Weighted Union Performance

Weighted union gives us guaranteed worst-case O(log n) for find

- The union rule prevents linear up trees
- Convince yourself that it will produce at worst a fairly balanced binary tree

However, we promised ourselves O(1) amortized time for find

- Weighted union does not give us enough
- Average-case is still O(log n)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

24

General Problem: Spanning a Graph

A simple problem: Given a *connected* graph $G=(V,E)$, find a minimal subset of the edges such that the graph is still connected

- A graph $G_2=(V,E_2)$ such that G_2 is connected and removing any edge from E_2 makes G_2 disconnected




Observations

1. Any solution to this problem is a tree
 - Recall a tree does not need a root; just means acyclic
 - For any cycle, could remove an edge and still be connected
 - We usually just call the solutions spanning trees
2. Solution not **unique** unless original graph was already a tree
3. Problem ill-defined if original graph not connected
 - We can find a spanning tree per connected component of the graph
 - This is often called a *spanning forest*
4. A tree with $|V|$ nodes has $|V|-1$ edges
 - This every spanning tree solution has $|V|-1$ edges

We Saw This Earlier

Our acyclic maze consisted of a tree that touched ever square of the grid



Motivation

A **spanning tree** connects all the nodes with as few edges as possible

Example: A "phone tree" so everybody gets the message and no unnecessary calls get made

- Bad example since would prefer a balanced tree

In most compelling uses, we have a *weighted* undirected graph and want a tree of least total cost

- Minimize electrical wiring for a house or wires on a chip
- Minimize road network if you cared about asphalt cost


This is the **minimum spanning tree** problem

- Will do that next, after intuition from the simpler case

Finding Unweighted Spanning Trees

Different algorithmic approaches to the spanning-tree problem:

1. Do a graph traversal (e.g., depth-first search, but any traversal will do) and keep track of edges that form a tree
2. or, iterate through edges and add to output any edge that doesn't create a cycle




Spanning Tree via DFS

```
spanning_tree(Graph G) {
  for each node i: i.marked = false
  for some node i: f(i)
}
f(Node i) {
  i.marked = true
  for each j adjacent to i:
    if(!j.marked) {
      add(i,j) to output
      f(j) // DFS
    }
}
```

Correctness: DFS reaches each node. We add one edge to connect it to the already visited nodes. Order affects result, not correctness.


Time: $O(|E|)$

DFS Spanning Tree Example




August 13, 2012 CSE 332 Data Abstractions, Summer 2012 37

DFS Spanning Tree Example




August 13, 2012 CSE 332 Data Abstractions, Summer 2012 38

DFS Spanning Tree Example




August 13, 2012 CSE 332 Data Abstractions, Summer 2012 39

DFS Spanning Tree Example




August 13, 2012 CSE 332 Data Abstractions, Summer 2012 40

DFS Spanning Tree Example



August 13, 2012 CSE 332 Data Abstractions, Summer 2012 41


DFS Spanning Tree Example



August 13, 2012 CSE 332 Data Abstractions, Summer 2012 42

DFS Spanning Tree Example

Stack
 f(1)
 f(2)
 f(7)
 f(5)
 f(4)
 f(3)
 f(6)



Output: (1,2), (2,7), (7,5), (5,4),
 (4,3), (5,6)


August 13, 2012

CSE 332 Data Abstractions, Summer 2012

43

DFS Spanning Tree Example

Stack



Output: (1,2), (2,7), (7,5), (5,4),
 (4,3), (5,6)

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

44

Second Approach

Iterate through edges; output any edge that does not create a cycle

Correctness (hand-wavy):

- Goal is to build an acyclic connected graph
- When we add an edge, it adds a vertex to the tree (or else it would have created a cycle)
- The graph is connected, we consider all edges

Efficiency:

- Depends on how quickly you can detect cycles
- Reconsider after the example

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

45

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7),
 (2,3), (4,5), (4,7)



Output:

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

46

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7),
 (2,3), (4,5), (4,7)



Output: (1,2)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

47

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7),
 (2,3), (4,5), (4,7)



Output: (1,2), (3,4)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

48

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7), (2,3), (4,5), (4,7)



Output: (1,2), (3,4), (5,6)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

49

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7), (2,3), (4,5), (4,7)



Output: (1,2), (3,4), (5,6), (5,7)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

50

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7), (2,3), (4,5), (4,7)



Output: (1,2), (3,4), (5,6), (5,7), (1,5)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

51

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7), (2,3), (4,5), (4,7)



Output: (1,2), (3,4), (5,6), (5,7), (1,5)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

52

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7), (2,3), (4,5), (4,7)



Output: (1,2), (3,4), (5,6), (5,7), (1,5)

August 13, 2012


CSE 332 Data Abstractions, Summer 2012

53

Example

Edges in some arbitrary order:

(1,2), (3,4), (5,6), (5,7), (1,5), (1,6), (2,7), (2,3), (4,5), (4,7)



Output: (1,2), (3,4), (5,6), (5,7), (1,5), (2,3)

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

54

Cycle Detection

To decide if an edge could form a cycle is $O(|V|)$ because we may need to traverse all edges already in the output

- So overall algorithm would be $O(|V||E|)$

But it is faster way to use the **DSUF ADT**

- Initially, each vertex is in its own 1-element set
- `find(u)`: what set contains `u`?
- `union(u,v)`: combine the sets containing `u` and `v`

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

55

Using Disjoint-Set to Detect Cycles

Invariant:

- `u` and `v` are connected in output-so-far if and only if `u` and `v` in the same set

Algorithm:

- Initially, each node is in its own set
- When processing edge `(u,v)`:
 - If `find(u) == find(v)`, then do not add the edge
 - Else add the edge and `union(u,v)`

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

56

Summary so Far

The **spanning-tree problem**

- Add nodes to partial tree approach is $O(|E|)$
- Add acyclic edges approach is $O(|E| \log |V|)$

But what we really want to solve is the **minimum-spanning-tree problem**

- Given a weighted undirected graph, find a spanning tree of minimum weight
- The above approaches suffice with minor changes
- Both will be $O(|E| \log |V|)$

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

57

Like vi versus emacs except people do not typically fight over which one is better (emacs and Kruskal are best!)

PRIM AND KRUSKAL'S ALGORITHMS

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

58

One Problem, Two Algorithms

Algorithm #1: Prim's Algorithm

- Shortest-path is to Dijkstra's Algorithm as Minimum Spanning Tree is to Prim's Algorithm
- Both based on expanding cloud of known vertices, basically using a priority queue

Algorithm #2: Kruskal's Algorithm

- Exactly our forest-merging approach to spanning tree but process edges in cost order

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

59

Idea: Prim's Algorithm

Central Idea:

- Grow a tree by adding an edge from the "known" vertices to the "unknown" vertices.
- Pick the edge with the smallest weight that connects "known" to "unknown."

Recall Dijkstra picked "edge with closest known distance to source."

- But that is not what we want here
- Otherwise identical
- Feel free to look back and compare

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

60

Pseudocode: Prim's Algorithm

1. For each node v , set $v.cost = \infty$ and $v.known = false$
2. Choose any node v .
 - a) Mark v as known
 - b) For each edge (v,u) with weight w , set $u.cost = w$ and $u.prev = v$
3. While there are unknown nodes in the graph
 - a) Select the unknown node v with lowest cost
 - b) Mark v as known and add $(v, v.prev)$ to output
 - c) For each edge (v,u) with weight w ,


```


                    if (w < u.cost) {
                        u.cost = w;
                        u.prev = v;
                    }
                    
```

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

61

Example: Prim's Algorithm




vertex	known?	cost	prev
A			
B			
C			
D			
E			
F			
G			

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

62

Example: Prim's Algorithm




vertex	known?	cost	prev
A	Y	0	-
B		2	A
C		2	A
D		1	A
E			
F			
G			

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

63

Example: Prim's Algorithm




vertex	known?	cost	prev
A	Y	0	-
B		2	A
C		2 1	A D
D	Y	1	A
E		1	D
F		6	D
G		5	D

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

64

Example: Prim's Algorithm




vertex	known?	cost	prev
A	Y	0	-
B		2	A
C	Y	2 1	A D
D	Y	1	A
E		1	D
F		6 2	D C
G		5	D

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

65

Example: Prim's Algorithm




vertex	known?	cost	prev
A	Y	0	-
B		2 1	A E
C	Y	2 1	A D
D	Y	1	A
E	Y	1	D
F		6 2	D C
G		5 3	D E

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

66


Example: Prim's Algorithm



vertex	known?	cost	prev
A	Y	0	-
B	Y	2 1	A E
C	Y	2 1	A D
D	Y	1	A
E	Y	1	D
F		6 2	D C
G		5 3	D E

August 13, 2012 CSE 332 Data Abstractions, Summer 2012 67


Example: Prim's Algorithm



vertex	known?	cost	prev
A	Y	0	-
B	Y	2 1	A E
C	Y	2 1	A D
D	Y	1	A
E	Y	1	D
F	Y	6 2	D C
G		5 3	D E

August 13, 2012 CSE 332 Data Abstractions, Summer 2012 68


Example: Prim's Algorithm



vertex	known?	cost	prev
A	Y	0	-
B	Y	2 1	A E
C	Y	2 1	A D
D	Y	1	A
E	Y	1	D
F	Y	6 2	D C
G	Y	5 3	D E

August 13, 2012 CSE 332 Data Abstractions, Summer 2012 69

Example: Prim's Algorithm



Output:
 (A, D) (C, F)
 (B, E) (D, E)
 (C, D) (E, G)

Total Cost: 9

vertex	known?	cost	prev
A	Y	0	-
B	Y	2 1	A E
C	Y	2 1	A D
D	Y	1	A
E	Y	1	D
F	Y	6 2	D C
G	Y	5 3	D E

August 13, 2012 CSE 332 Data Abstractions, Summer 2012 70

Analysis: Prim's Algorithm

Correctness

- Intuitively similar to Dijkstra's algorithm

Run-time

- Same as Dijkstra's algorithm
- $O(|E| \log |V|)$ using a priority queue

August 13, 2012 CSE 332 Data Abstractions, Summer 2012 71

Idea: Kruskal's Algorithm

Central Idea:

- Grow a forest out of edges that do not grow a cycle, just like for the spanning tree problem.
- But now consider the edges in order by weight

Basic implementation:

- Sort edges by weight $\rightarrow O(|E| \log |E|) = O(|E| \log |V|)$
- Iterate through edges using DSUF for cycle detection $\rightarrow O(|E| \log |V|)$

Somewhat better implementation:

- Floyd's algorithm to build min-heap with edges $\rightarrow O(|E|)$
- Iterate through edges using DSUF for cycle detection and deleteMin to get next edge $\rightarrow O(|E| \log |V|)$
- Not better worst-case asymptotically, but often stop long before considering all edges

August 13, 2012 CSE 332 Data Abstractions, Summer 2012 72

Pseudocode: Kruskal's Algorithm

1. Put edges in min-heap using edge weights
2. Create DSUF with each vertex in its own set
3. While output size < $|V|-1$
 - a) Consider next smallest edge (u,v)
 - b) if $\text{find}(u,v)$ indicates u and v are in different sets
 - output (u,v)
 - union (u,v)

Recall invariant:


u and v in same set if and only if connected in output-so-far

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

73

Example: Kruskal's Algorithm



Edges in sorted order:

- 1: (A,D) (C,D) (B,E) (D,E)
- 2: (A,B) (C,F) (A,C)
- 3: (E,G)
- 5: (D,G) (B,D)
- 6: (D,F)
- 10: (F,G)

Sets: (A) (B) (C) (D) (E) (F) (G)

Output:


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

74

Example: Kruskal's Algorithm



Edges in sorted order:

- 1: ~~(A,D)~~ (C,D) (B,E) (D,E)
- 2: (A,B) (C,F) (A,C)
- 3: (E,G)
- 5: (D,G) (B,D)
- 6: (D,F)
- 10: (F,G)

Sets: (A,D) (B) (C) (E) (F) (G)

Output: (A,D)


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

75

Example: Kruskal's Algorithm



Edges in sorted order:

- 1: ~~(A,D)~~ ~~(C,D)~~ (B,E) (D,E)
- 2: (A,B) (C,F) (A,C)
- 3: (E,G)
- 5: (D,G) (B,D)
- 6: (D,F)
- 10: (F,G)

Sets: (A,C,D) (B) (E) (F) (G)

Output: (A,D) (C,D)


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

76

Example: Kruskal's Algorithm



Edges in sorted order:

- 1: ~~(A,D)~~ ~~(C,D)~~ ~~(B,E)~~ (D,E)
- 2: (A,B) (C,F) (A,C)
- 3: (E,G)
- 5: (D,G) (B,D)
- 6: (D,F)
- 10: (F,G)

Sets: (A,C,D) (B,E) (F) (G)

Output: (A,D) (C,D) (B,E)


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

77

Example: Kruskal's Algorithm



Edges in sorted order:

- 1: ~~(A,D)~~ ~~(C,D)~~ ~~(B,E)~~ ~~(D,E)~~
- 2: (A,B) (C,F) (A,C)
- 3: (E,G)
- 5: (D,G) (B,D)
- 6: (D,F)
- 10: (F,G)

Sets: (A,B,C,D,E) (F) (G)

Output: (A,D) (C,D) (B,E) (D,E)


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

78

Example: Kruskal's Algorithm



- Edges in sorted order:
- 1: (A,D) (C,D) (B,E) (D,E)
 - 2: (A,B) (C,F) (A,C)
 - 3: (E,G)
 - 5: (D,G) (B,D)
 - 6: (D,F)
 - 10: (F,G)

Sets: (A,B,C,D,E) (F) (G)
 Output: (A,D) (C,D) (B,E) (D,E)


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

79

Example: Kruskal's Algorithm



- Edges in sorted order:
- 1: (A,D) (C,D) (B,E) (D,E)
 - 2: (A,B) (C,F) (A,C)
 - 3: (E,G)
 - 5: (D,G) (B,D)
 - 6: (D,F)
 - 10: (F,G)

Sets: (A,B,C,D,E,F) (G)
 Output: (A,D) (C,D) (B,E) (D,E) (C,F)


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

80

Example: Kruskal's Algorithm



- Edges in sorted order:
- 1: (A,D) (C,D) (B,E) (D,E)
 - 2: (A,B) (C,F) (A,C)
 - 3: (E,G)
 - 5: (D,G) (B,D)
 - 6: (D,F)
 - 10: (F,G)

Sets: (A,B,C,D,E,F) (G)
 Output: (A,D) (C,D) (B,E) (D,E) (C,F)


At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

81

Example: Kruskal's Algorithm



- Edges in sorted order:
- 1: (A,D) (C,D) (B,E) (D,E)
 - 2: (A,B) (C,F) (A,C)
 - 3: (E,G)
 - 5: (D,G) (B,D)
 - 6: (D,F)
 - 10: (F,G)

Sets: (A,B,C,D,E,F,G)
 Output: (A,D) (C,D) (B,E) (D,E) (C,F) (E,G)

At each step, the union/find sets are the trees in the forest

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

82

Analysis: Kruskal's Algorithm

Correctness: It is a spanning tree

- When we add an edge, it adds a vertex to the tree (or else it would have created a cycle)
- The graph is connected, we consider all edges

Correctness: That it is minimum weight

- Can be shown by induction
- At every step, the output is a subset of a minimum tree

Run-time

- $O(|E| \log |V|)$

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

83

So Which Is Better?

Time/space complexities essentially the same

Both are fairly simple to implement

Still, Kruskal's is slightly better

- If the graph is not connected, Kruskal's will find a forest of minimum spanning trees

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

84

sniff

WRAPPING UP DATA ABSTRACTIONS

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

85

Your Programming Mind has Changed

Before, you often thought first about code

- Declare a variable, a for-loop here, an if-else statement there, etc.

Now, you will see a problem and also think of the data structure

- Lots of lookups... use a hashtable
- Is this a graph and shortest path problem?
- Etc.

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

87

Cheers, Thanks, Whee!

Take care

Fill out the evaluations... I read these!!

Good luck on the final

Remember: Optional Section on Thursday

- Get your final back
- Free doughnuts!
- And maybe another cool data structure

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

89

That's All Folks

Disjoint Set Union-Find and minimum spanning trees are the last topics we will get to cover

Still, there are plenty more data structures, algorithms and applications out there to learn

You have the basics now

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

86

Most Important Lesson

There is rarely a best programming solution

Every solution has strengths and weaknesses

The key is to be able to argue in favor of your approach over others

Just remember:

Even though QuickSort's name says it is fast, it is not always the best sort every time

August 13, 2012

CSE 332 Data Abstractions, Summer 2012

88