Minimum Spanning Trees Prim's Algorithm Priority Queues

CSCI 3100

Review & Overview

Last time

- Spanning trees and spanning tree properties
- Minimum spanning tree of a Graph
- Kruskal's algorithm for MST uses greedy approach: select next smallest edge, if one of the end-points of that edge is not yet in the tree, add the edge and both end-points to the tree.

This time

- Prim's algorithm for MST
- Priority Queues
- Proof of Fluff

Prim's Algorithm

In Kruskal's algorithm we grow a *spanning forest* rather than a spanning tree • Only at the end is it guaranteed to be connected, hence a *spanning tree*

In Prim's algorithm we grow a spanning tree

T = the set of nodes currently forming the tree

Prim's algorithm: always add the cheapest edge crossing the cut



















Operations of a priority queue

Possible operations of a priority queue

- Build a priority queue
- Add an element to the priority queue
- Remove smallest element from the priority queue
- Get smallest element (without removing)
- Check if priority queue is empty
- Get the size of the priority queue

Java has a PriorityQueue class

Evaluating implementations

When we choose a data structure, it is important to look at usage patterns

- If we load an array once and do thousands of searches on it, we want to make searching fast—so we would probably sort the array
- If we load a huge array and expect to do only a few searches, we probably *don't* want to spend time sorting the array

For almost all uses of a queue (including a priority queue), we eventually remove everything that we add

Hence, when we analyze a priority queue, neither "add" nor "remove" is more important—we need to look at the timing for "add + remove"

Unsorted array implementation

A priority queue could be implemented as an *unsorted* array. If n is the number of elements in the queue, what is the complexity of adding and removing an element

- $\,\circ\,$ A. Adding an element $O(1),\,$ removing an element O(n) time
- $\,\circ\,$ B. Adding an element O(n), removing an element O(1)
- $\,\circ\,$ C. Adding an element $O(1),\,$ removing an element O(1)
- D. Adding an element O(n), removing an element $O(n \lg(n))$
- E. None of the above



Unsorted linked list implementation

A priority queue could be implemented as an *unsorted* linked list. If n is the number of elements in the queue, what is the complexity of adding and removing an element?

- A. Adding an element O(1), removing an element O(n) time
- $\,\circ\,$ B. Adding an element $O(n),\,$ removing an element O(1)
- $\,\circ\,$ C. Adding an element $O(1),\,$ removing an element O(1)
- $\,\circ\,$ D. Adding an element O(n), removing an element $O(n \; lg(n)$)
- E. None of the above

Unsorted linked list implementation

A priority queue could be implemented as an *sorted* linked list. If n is the number of elements in the queue, what is the complexity of adding and removing an element?

A priority queue could be implemented as a sorted linked list

- A. Adding an element O(1), removing an element O(n) time
- B. Adding an element O(n), removing an element O(1)
- C. Adding an element O(1), removing an element O(1)
- D. Adding an element O(n), removing an element $O(n \lg(n))$
- E. None of the above

Unbalanced binary tree implementation

A priority queue could be represented as a (not necessarily balanced) binary search tree

- A. Adding an element $O(\log n)$ removing an element O(n)
- $\circ\,$ B. Adding an element ranges from $O(log\,\,n)$ to $\,O(n),$ removing an element ranges $O(log\,\,n)$ to $\,O(n)$
- $\circ\,$ C. Adding an element $O(log\,\,n),$ removing an element $O(log\,\,n)$
- D. Adding an element $O(n \log n)$, removing an element O(1)
- E. Adding an element O (1), removing an element O(log n)

Binary tree implementations

A priority queue could be represented as a balanced binary search tree

- Insertion and removal could destroy the balance
- We need an algorithm to *rebalance* the binary tree
- Good rebalancing algorithms require only O(log n) time, and are complicated





Using the heap

To add an element:

- Increase lastIndex and put the new value there
- Reheap the newly added node by swapping with parent node, until heap property is restored
 - This is called up-heap bubbling or percolating up
 - Up-heap bubbling requires O(log n) time

To remove an element:

- Remove the element at location 0
- Move the element at location lastIndex to location 0, and decrement lastIndex
- Reheap the new root node (the one now at location 0) by swapping with smallest child element until heap property is restored
 - $\circ~$ This is called down-heap bubbling or percolating down
 - Down-heap bubbling requires O(log n) time

Thus, it requires O(log n) time to add and remove an element









MST-PRIM(G, w, r)

1 for each $u \in G$. V 2 $u.key = \infty$ 3 $u.\pi = NIL$ 4 r.key = 05 Q = G.V6 while $Q \neq \emptyset$ 7 u = EXTRACT-MIN(Q)8 for each $v \in G$. Adj[u]9 if $v \in Q$ and w(u, v) < v. key 10 $v.\pi = u$ 11 v.key = w(u, v)

Proof or Fluff

Let T=(V, E) be a tree. Prove that |E|=|V|-1

Consider the following proof by induction on V:

Base case: Clearly, this is true for |V|=1 and |V|=2.

Inductive hypothesis: suppose true for trees with |V|-1 vertices. Then this tree has |V|-1-1 edges. We can construct another tree by adding one new vertex and connecting it to one of vertices in the tree with one edge. Thus, we have a tree with |V| vertices and |V|-1-1+1=|V|-1 edges. This proves that |E|=|V|-1.

A. This proof is valid

B. This proof is flawed because it proves that there exists a tree with |E|=|V|-1 and not the general case

C. This proof is flawed because it doesn't prove that the newly constructed tree is in fact a tree.

D. This proof is flawed for some other reason

Let T=(V, E) be a tree. Prove that |E|=|V|-1

Consider the following proof by induction on V:

Base case: Clearly, this is true for |V|=1 and |V|=2.

Inductive hypothesis: suppose true for trees with n < |V| vertices.

Let T be a tree with |V| vertices. Let e be an edge connecting vertices u and v in T. Since T is a tree, there is a unique path from u to v and it has to be via edge e. If we remove e, T will become disconnected. Now T-{e} consists of two components T_1 and T_2 and those components are trees (since there were no cycles in T to begin with).

Let n_1 be the number of vertices in T_1 and n_2 be the number of vertices in T_2 , so $n_1+n_2=|V|$.

Also $0 < n_1 < |V|$ and $0 < n_2 < |V|$. By inductive hypothesis the number of edges in T₁ is n₁-1 and the number of edges in T₂ is n₂-1. Thus, the number of edges in T is n₁-1+n₂-1+1=n₁+n₂-1=|V|-1.

- A. This proof is valid
- B. This proof is flawed because $n_1+n_2=|V|$ is false

C. This proof is flawed because T₁ and T₂ are not guaranteed to be trees.

D. This proof is flawed for some other reason