

Analysis of Algorithms, I

CSOR W4231.002

Eleni Drinea
Computer Science Department

Columbia University

Tuesday, March 31, 2015

- 1 Recap
- 2 Minimum Spanning Trees (MSTs)
 - Prim's algorithm
 - Kruskal's algorithm
 - More MST algorithms
- 3 #spanning trees in K_n

- 1 Recap
- 2 Minimum Spanning Trees (MSTs)
 - Prim's algorithm
 - Kruskal's algorithm
 - More MST algorithms
- 3 #spanning trees in K_n

Recap

- ▶ Chain hashing: expected wasted space
- ▶ Hashing-based fingerprints
- ▶ Bloom filters

- 1 Recap
- 2 Minimum Spanning Trees (MSTs)
 - Prim's algorithm
 - Kruskal's algorithm
 - More MST algorithms
- 3 #spanning trees in K_n

The problem

Input: weighted undirected graph $G = (V, E, w)$

Output: a communication network over the locations in V such that

1. the network is **connected**
2. it is built as cheaply as possible

Equivalently, find subset $T \subseteq E$ such that

1. the graph (V, T) is connected and
2. $\sum_{e \in T} w(e)$ is minimal.

Minimum weight Spanning Trees (MST)

Remark 1.

The graph (V, T) is a tree: if there is a cycle, remove any edge from the cycle and obtain a connected graph with smaller cost.

- ▶ **Spanning tree:** a tree that spans the vertices of the graph
- ▶ **Goal:** find the minimum weight spanning tree

Remarks

- ▶ even simple graphs have many spanning trees
(*how many spanning trees in a simple cycle?*)
- ▶ #spanning trees in the **complete** graph on n vertices: n^{n-2}

The cut property

Definition 1 (Cut).

A cut $(S, V - S)$ is a bipartition of the vertices.

Claim 1 (Cut property).

Assume all edge weights are distinct. Let $S \subset V$ ($S \neq \emptyset$). Let e be the minimum-weight edge with one endpoint in S and the other in $V - S$. Then every MST contains e .

Remark 2.

The assumption serves the purposes of the analysis; we will get rid of it later.

Proof of the cut property

By an **exchange** argument.

- ▶ Let T' be a spanning tree that does not contain $e = (u, v)$.
- ▶ Then there must be some other path P in T' from u to v .
- ▶ Starting at u , follow the vertices of P : since (u, v) crosses from S to $V - S$, there must be some first vertex $v' \in V - S$ on P . Let u' be the last vertex before it in S .
- ▶ Then $e' = (u', v') \in T'$ **and** e' crosses between $S, V - S$.

Proof of the cut property (cont'd)

Exchange e with e' to obtain the set of edges

$$T = T' + \{e\} - \{e'\}.$$

T is a spanning tree:

- ▶ T is connected: any path in T' that used $e' = (u', v')$ is rerouted to follow P from u' to u , (u, v) and P from v to v' .
- ▶ T is acyclic (*why?*).

Since both e' and e cross between S and $V - S$ but e is the lightest edge with this property, $w(e) < w(e')$. Thus

$$w(T) < w(T').$$

Using the cut property for designing MST algorithms

The cut property says: construct MST **greedily** by taking the **lightest** edge across two regions not yet connected.

Generic-MST($G = (V, E, w)$)

$A = \emptyset$

while $|A| \leq n - 1$ **do**

 Pick $S \subseteq V$ s.t. no edge in A crosses between $S, V - S$

 Let $e \in E$ be a lightest edge that crosses between $S, V - S$

$A = A \cup \{e\}$

end while

Prim's algorithm

In Prim's algorithm, the edges in A always form a subtree and S is chosen to be the set of this subtree's vertices.

In other words,

1. start with a root node s
2. **greedily** grow a tree outward from s by adding the node that can be attached as cheaply as possible at every step.


Detailed description of Prim's algorithm

1. Maintain a set $S \subseteq V$ on which a spanning tree has been constructed so far.
2. In each iteration
 - 2.1 grow S by one vertex v , adding the vertex $v \in V - S$ that minimizes the attachment cost:


$$\min_{e=(u,v):u \in S} w(e)$$

- 2.2 include edge e to A

Example graph



Prim's MST for example graph (letters indicate the order in which edges were added)



Follows directly from the Cut property: at every iteration an edge (u, v) is added such that

- ▶ $u \in S, v \in V - S$; here S is the set of vertices on which a partial MST has been constructed;
- ▶ (u, v) is the lightest edge that crosses between S and $V - S$.

Implementing Prim's algorithm

Similarly to Dijkstra's algorithm,

- ▶ store every node $v \in V - S$ in a priority queue Q , e.g., implemented as a binary min-heap (key= weight of the lightest edge from S to v). Initially, $S = \{s\}$.
- ▶ maintain two arrays
 - ▶ $dist[v]$: stores the weight of the lightest edge between v and any vertex in S (in Dijkstra, it stored a conservative overestimate of the distance of v from the source s)
 - ▶ $prev[v]$: the node responsible for adding v to S

Pseudocode

```
Prim( $G = (V, E, w), s$ )  
  for  $u \in V$  do  
     $dist[v] = \infty; prev[v] = NIL$   
  end for  
   $dist[s] = 0$   
   $Q = \{V; dist\}$   
   $S = \emptyset$   
  while  $Q \neq \emptyset$  do  
     $u = \text{ExtractMin}(Q)$   
     $S = S \cup \{u\}$   
    for  $(u, v) \in E$  and  $v \in V - S$  do  
      if  $dist[v] > w(u, v)$  then  
         $dist[v] = w(u, v)$   
         $prev[v] = u$   
        DecreaseKey( $Q, v$ )  
      end if  
    end for  
  end while
```

Further implementations of Prim's algorithm

Notation: $|V| = n$, $|E| = m$

Implementation	ExtractMin	Insert/ DecreaseKey	Time
Array	$O(n)$	$O(1)$	$O(n^2)$
Binary heap	$O(\log n)$	$O(\log n)$	$O((n + m) \log n)$
d -ary heap	$O(\log n)$	$O(\log n)$	$O((nd + m) \frac{\log n}{\log d})$
Fibonacci heap	$O(\log n)$	$O(1)$ amortized	$O(n \log n + m)$

- ▶ Optimal choice for $d \approx m/n$ (the *average* degree of the graph)
- ▶ d -ary heap works well for both sparse and dense graphs
 - ▶ If $m = n^{1+x}$, what is the running time of Prim's algorithm using a d -ary heap?
- ▶ **Amortized** analysis: coming up in the next lecture

Kruskal's algorithm

Short description: at every step, add to A the *lightest* edge that does not create a cycle with the edges already in A .

Thus, at all times, A is a subset of an MST.

Alternative view: merging partial trees

Initially, every vertex forms its own trivial tree (no edges).
Maintain a *forest* of trees at all times.


Let $T(v)$ be the tree where vertex v belongs.

1. Initialize $A = \emptyset$
2. **Sort** the edges by increasing weight
3. For every edge $e = (u, v)$ in **increasing order**:
 - ▶ If u and v belong to the same tree, discard e
 - ▶ Else
 - ▶ $A = A \cup \{e\}$;
 - ▶ **merge** $T(u), T(v)$ into a single tree


\triangle *Need a data structure that allows*

1. *to check if u, v belong to the same tree*
2. *for updates to reflect the merging of two trees into one*

Example graph



Kruskal's MST for example graph (letters indicate the order in which edges were added)



- ▶ Let (u, v) be the edge added at the current iteration.
 - ▶ Let S be the set of nodes that have a path to u by edges in A just before (u, v) is added; then $u \in S$ but $v \notin S$.
 - ▶ Also, (u, v) must be the **first** edge between S and $V - S$ encountered so far: otherwise, if such an edge was encountered before, it would have been added to A since its inclusion would not cause a cycle.
- ⇒ (u, v) is the **lightest** edge that crosses between S and $V - S$
- ▶ By the Cut Property, (u, v) belongs to the MST.

Implementing Kruskal's algorithm

Kruskal's algorithm maintains a forest of trees at all times, starting from n trivial trees (no edges).

Want a data structure that maintains a **collection of disjoint sets** and supports operations:

1. **MakeSet**(u): Given an element u , create a new tree containing only u . **Target worst-case time: $O(1)$**
2. **Find**(u): Given an element u , find which tree u belongs to. **Target worst-case time: $O(\log n)$**
3. **Union**(u, v): Merge the tree containing u and the tree containing v into a single tree. **Target worst-case time: $O(\log n)$**

Pseudocode

```
Kruskal( $G = (V, E, w)$ )  
   $A = \emptyset$   
  Sort( $E$ ) by  $w$   
  for  $u \in V$  do MakeSet( $u$ )  
  end for  
  for  $(u, v) \in E$  by increasing  $w$  do  
    if Find( $u$ )  $\neq$  Find( $v$ ) then  
       $A = A \cup \{(u, v)\}$   
      Union( $u, v$ )  
    end if  
  end for
```

Running time analysis

- ▶ Sorting: $O(m \log m) = O(m \log n)$
- ▶ n `Makeset()` operations: $O(n)$
- ▶ $2m$ `Find()` operations: $2mO(\log n)$
- ▶ $\leq n$ `Union()` operations: $nO(\log n)$

Running time: $O(m \log n)$

When is it safe to not include an edge to the MST?

Fact 2 (The Cycle Property).

Assume that all edge costs are distinct. Let C be any cycle in T , and let edge (u, v) be the heaviest edge in C . Then e does not belong to any MST of G .

Proof of the cycle property

- ▶ Let T be a spanning tree that contains e . We want to show that T is not optimal.
 - ▶ To this end, we will exchange e for some e' to get a spanning tree with less weight.
 - ▶ First, delete e from T ; T is now partitioned into two components: the set S containing u and the set $V - S$ containing v
- ⇒ We want an edge e' with one endpoint in S and another in $V - S$ so as to reconnect them.

Proof of the cycle property (cont'd)

- ▶ We can find such an edge by following the cycle C .
- ▶ Consider the edges of C except for e : they form a path from u to v .
- ▶ So if we start at u , following this path, at some point there is an edge e' that crosses from S to $V - S$. Construct

$$T' = T - \{e\} + \{e'\}.$$

- ▶ Now T' is connected and has $n - 1$ edges. Moreover, since e is the heaviest edge in the cycle

$$w(T') < w(T).$$

More MST algorithms

Fact 2 yields yet another algorithm for MST

Reverse-Delete($G = (V, E, w)$)

- ▶ Start with the full graph
- ▶ Sort the edges in decreasing weight
- ▶ Repeatedly delete edges in order of decreasing weight, so long as the graph does not become disconnected

More MST algorithms: combine the Cut property (to add edges) and the Cycle property (to eliminate edges)

△ such algorithms may be subtle to implement

Eliminating the assumption on unequal edge weights

- ▶ Suppose some edges have equal weights.
 - ▶ Slightly perturb all edge weights by different, **tiny** amounts.
- ⇒ All edge weights are now distinct.
- ▶ Apply the algorithms discussed in the previous sections.

Remark 3.

*Perturbations serve as **tie-breakers**: edges whose weights differed before still have the same relative order.*

Today

- 1 Recap
- 2 Minimum Spanning Trees (MSTs)
 - Prim's algorithm
 - Kruskal's algorithm
 - More MST algorithms
- 3 #spanning trees in K_n

Cayley's formula

- ▶ Let K_n be the complete graph on n vertices.
- ▶ Let $T_n = \#$ spanning trees in K_n
- ▶ Cayley's formula: $T_n = n^{n-2}$
- ▶ Proof: via computing a quantity in **two different ways** to derive an expression for T_n .

Proof of Cayley's formula: counting in 2 different ways

Recall that a directed tree is a **rooted** graph that has a simple path from the root to every vertex in the graph

- ▶ A tree has $n - 1$ edges.

The quantity we will compute in two different ways is the number ν of different *sequences of directed edges* that can be added to an empty graph on n vertices to yield a rooted tree.

1. A formula for ν that directly involves T_n

1. Start with a spanning tree on the empty graph (T_n choices)
 2. Pick a root for the tree (n choices)
 3. Given the root, the direction of every edge is fully determined (*why?*)
- \Rightarrow There are $n - 1$ directed edges to insert in **any order** in our graph ($(n - 1)!$ ways to order them)

In total, there are $T_n \cdot n \cdot (n - 1)!$ different sequences of directed edges to add in a graph so as to form a directed rooted tree; so

$$\nu = T_n \cdot n \cdot (n - 1)!$$

2. Computing ν directly

Start with an empty graph on n nodes.

We will add $n - 1$ directed edges **one by one** so that we construct a rooted tree spanning the n nodes.

1. At every step $i = 1, 2, \dots, n - 1$, let n_i be the #possible directed edges from which to choose the edge to add.
2. Then the #different sequences of directed edges that yield a rooted tree is simply the product of all n_i .

Adding the first directed edge

#possible directed edges from which to choose at every step:

- ▶ An edge is completely defined when its tail and head are picked.
- ▶ Hence the #possible directed edges at every step is

$$(\text{\#ways to choose a tail}) \cdot (\text{\#ways to choose a head})$$

Initially, we have a forest of n empty rooted trees.

1. Adding the 1st edge:

- ▶ tail: pick **any** of the n vertices
 - ▶ head: direct the edge to **any** of the $n - 1$ *other* roots
- $\Rightarrow n_1 = n(n - 1)$ ways to choose the 1st edge

Adding the second directed edge

The graph is now a forest with $n - 1$ rooted trees.

2. Choosing the 2nd edge:

- ▶ tail: pick **any** of the n vertices
- ▶ direct the edge to the **root** of any tree (so that the resulting graph remains a rooted tree) **except** for the tree where the tail belongs (*why?*)

⇒ $n_2 = n(n - 2)$ ways to choose the 2nd edge

Adding the k -th edge

- k. k -th edge: the reasoning is entirely similar. After addition of the $(k - 1)$ -st edge, there are $n - (k - 1)$ rooted trees in the forest (by construction, every edge we add reduces the number of trees by 1).
- ▶ pick **any** of the n vertices as the tail of the edge
 - ▶ direct the edge to the **root** of any tree in the **except** for the tree where the tail belongs
- ⇒ $n_k = n(n - k)$ ways to choose the k -th edge
- n-1. $n - 1$ -st edge: $n_{n-1} = n \cdot 1$ ways to choose the $n - 1$ -st edge

Conclusion

- ▶ In total, there are $\prod_{i=1}^n n_i = n^{n-1}(n-1)!$ ways to add the edges. Hence $\nu = n^{n-1}(n-1)!$
- ▶ Equating the two expressions for ν , we obtain:

$$T_n = n^{n-2}$$

- ▶ Arbitrary graphs: #spanning trees computable in polynomial time