Lecture contents

1. Graph contractions

2. Star contraction

3. Parallel connected components (random mates)

4. Intro to optimization

1 Minimum spanning tree algorithms

Now we'll shift our focus to parallel graph algorithms, beginning with minimum spanning trees. In the following sections, we'll denote our *connected* and *undirected* graph by $G = (V, E, w)$. The size of the vertex set $|V| = n$, the size of the edge set $|E| = m$, and we assume that the weights $w : E \to \mathbb{R}$ are distinct for convenience.[1]

A *tree* is an undirected graph G which satisfies any two of the three following statements, such a graph also then satisfies the other property as well.[2]

1. G connected

2. G has no cycles

3. G has exactly $n - 1$ edges

The *minimum spanning tree* (MST) of G is $T^* = (V, E_T, w)$ such that $\sum_{e \in E_{T^*}} w(e) \leq \sum_{e \in E_T} w(e)$ for any other spanning tree T. Under the assumptions on G, the MST T^* is unique and the previous inequality is strict.

---

[1] If the edge weights are not distinct, then we may simply perturb each edge weight by a small random factor $\epsilon > 0$, where $\epsilon < 1/n^3$. We no longer have to break ties in our algorithm and at the end, we may simply round back the edge weights to recover the weight of the minimum spanning tree.

[2] To see that $(1, 2) \implies 3$, use induction on the number of nodes n. The base case is trivial. For the inductive step, realize that any acyclic connected graph G must have a leaf v where $d(v) = 1$. Since $G - v$ also acyclic and connected, by the induction hypothesis, $e(G') = n - 2$. Since $d(v) = 1$, we have that $e(G) = n - 1$. To see that $(1, 3) \implies 2$, suppose toward contradiction G has a cycle. To see that $(2, 3) \implies 1$, consider the case that G split into k components. Since G has no cycles, each component G_i both connected and acyclic, hence each is a tree. So total number of edges in G given by $\sum_i (n(G_i) - 1) = n - k$. But since $|E| = n - 1$, $k = 1$.

1.1 Sequential approaches and Kruskal's algorithm

Sequential algorithms for finding the MST are nice and easy because greedy algorithms work well for this problem. The two clear approaches are Prim's algorithm and Kruskal's algorithm. Here, we'll focus on Kruskal's algorithm, given in Algorithm 1.

Algorithm 1: Kruskal's MST algorithm

Input: $G = (V, E, w)$, an undirected and connected graph
Result: T^*, the MST of G

1 Sort E in increasing order by edge weight $T \leftarrow \emptyset$
2 **while** T *not yet a spanning tree* **do**
3 $\quad$ $e \leftarrow$ next edge in the queue $\qquad$ // i.e. lightest edge yet to be considered
4 $\quad$ **if** $(T \cup \{e\})$ *contains no cycles* **then**
5 $\quad\quad$ $T \leftarrow T \cup \{e\}$ $\qquad\qquad$ // By red rule, e belongs in T^*.
6 $\quad$ **end**
7 **end**
8 **return** T

Kruskal's algorithm is based on the *cut property* (or *red rule*), which we now prove.

Theorem 1.1 (Cut-Property, Red-Rule) *Let $G(V, E, w)$ be an undirected and connected graph. Then the edge with minimum weight leaving any cut $S \subset V$ is in the MST of G.*

Proof: We are given a graph G and a cut $S \subset V$. We say that an edge is in a cut S if exactly one incident node is in S and the other incident node of the edge in $V \setminus S$. Let $e^* = (u, v)$ be the minimum weight edge in cut S. Assume toward contradiction that the edge with minimum weight leaving cut S is *not* in the MST T, i.e. that $e^* \notin T$ where T is the MST of G and

$$e^* = \arg\min_{e \in S} w(e).$$

Suppose $e^* = (u, v)$. Since T is a spanning tree of G, we know that u and v are connected in the edge set of T, i.e. there exists a u-v path in T. Of course, we have assumed that $(u, v) \notin T$. We construct a u-v path using depth first search. Let $(x, y) = e'$ denote the edge in T which crosses cut S in the u-v path. Replace (x, y) with (u, v). Call the result T'.

We claim that $T' = (T \setminus (x, y)) \cup (u, v)$ is still a tree. We first claim T' retains connectivity among all nodes. To see this, realize first that since (x, y) and (u, v) each have exactly one incident node in S and one incident node in $V \setminus S$, then it is *not* possible that either (x, y) or (u, v), when removed from T, could result in S becoming disconnected or $V \setminus S$ becoming disconnected. Hence, without either of these edges, we can get from all nodes in S to all other nodes in S, and same goes for the set $V \setminus S$. It remains to show that we can still traverse from S to $V \setminus S$. Realize that all paths previously going from S to $V \setminus S$ using (x, y) can now simply utilize (u, v) instead.

Further, realize that we have not created any cycles. Specifically, when we add edge (u, v) to the tree T, we induce a cycle. But realize that the edge (x, y) is part of this cycle, and we immediately remove it. Hence in maintaining connectivity and keeping exactly $n - 1$ edges, by definition of a tree we know that the result is acyclic.

Since e^* is the minimum weight edge in S, and since it is not contained in T, we know that $w(e^*) < w(e')$. But then this implies that

$$w(T') = \sum_{e \in T'} w(e) = \sum_{e \in T} w(e) - \underbrace{[w(e') - w(e^*)]}_{>0} < \sum_{e \in T} w(e) = w(T)$$

But this is a contradiction, since T was defined as the MST of G. Thus, e^* is in the MST of G. ∎

1.2 Complexity Analysis for Kruskal's

A brief complexity analysis of Kruskal's algorithm is as follows.[3] Recall work is defined as $T_1 = W(n)$.

Sorting Edges Our QuickSort algorithm requires $\mathbb{E}[W(n)] = O(m \log m) = O(m \log n)$.

Checking for Cycles Specifically because we are considering adding a single edge to a tree, we can check if we are inducing a cycle in almost constant time; specifically the work required is $\alpha(m, n)$, where α denotes the Inverse Ackermann function.[**?**]

We show now quickly a way to check for cycles in $O(\log n)$ work that is a bit more accessible. We must ensure that if edge $e = (u, v)$ to be added, that the component of u is *not* the same component as v. To do this, we keep several data structures around: an array and a Binary Search Tree. We assume that with each component of our graph, we associate with it a binary search tree storing values of nodes in our graph G which are in a particular component k. In addition, we keep an array of length n where in each entry, it tells us which component node i currently in.

Hence, given a candidate (u, v), we go to our length n array and look up in constant time $a[v]$ which tells us that node v in component j. We then go to our BST corresponding to component j, and search in $O(\log n)$ time to check if node u in the same component.

Updating Data Structures if we Add an Edge If we find they share the same component already, we discard the edge and continue. Otherwise, the edge turns out to connect two components, hence we must update our data structures. We go to array A of length n, and write down that u now belongs to v's component.[4] This last step requires $O(1)$ work.

Now, we must merge the component of u with the component of v. Two merge two binary search trees, we first flatten each BST into a sorted linked list with $O(n)$ work using in order

---

[3]It's worthwhile to note that since $\binom{n}{2} \leq m \implies m = O(n^2)$, hence any $O(\log m)$ term may actually be replaced by $O(\log n)$, using the fact that $\log m = O(\log n^2) = O(2 \log n) = O(\log n)$. We leave $O(\log m)$ terms in place here for pedagogical purposes.

[4]We resolve the conflict of which node to join to which component by (arbitrarily) choose node index v as the parent since it is larger than node index u.

traversal. We then merge two sorted lists in $O(n)$ work to get a sorted *array*. We then need to form a BST with our sorted array. To do this, we fix attention to the element in the middle of our array. Realize that all preceding elements in the list are no larger than itself, and all elements following are no smaller. We may make the same observation for the left and right partitions, each time storing pointers from the median element of the list to the left and right children medians. In this way, we re-construct our BST with $O(n)$ work, since we only end up applying a constant amount of work to each element in our sorted array.

Total Work Sorting edges costs $O(m \log n)$ work. Our while loop iterates over at most m edges. Checking for cycles each iteration costs $O(\log n)$ work. Hence we have incurred $O(m \log n)$ work thus far without even considering how much it costs to update our data structures when we add an edge to our tree.

Realize that we only need to add $n - 1$ edges before we construct a tree. Further, each edge we add requires $O(n)$ work. Hence the entire tree construction process takes $O(n^2)$ work. Since $O(n^2) = O(m)$, the work required to sort our edges dominates, and total work required is $O(m \log n)$.

2   Parallel MST via Boruvka's Algorithm

Since we need to evaluate the edges in order of weight, Kruskal's algorithm is inherently sequential. We need a new approach if we want a parallel MST algorithm. Kruskal's algorithm is instructive in that it uses the Cut Property which we proved last lecture. Boruvka's algorithm, will also use this property. [1]

2.1   Observation - Smallest Weight Incident Edge on each Node belongs in MST

A single node is a subset of the nodes, i.e. each node can be considered individually to define a cut. And so if you look at a single node and its incident edges, the smallest one must be in the MST by applying the cut property. Realize that if we determined the smallest weight edge incident to *each* node in our graph, we *must* find at least $n/2$ unique edges belonging to the MST. Why $n/2$? When we search for a lowest weight incident edge on each node, realize that it's possible the node to which such an edge connects also selects the same edge for inclusion in the MST. At worst, this could happen for each node. Hence we must find at least $n/2$ edges in our MST.

2.2   Edge Contraction

Before we can recursively apply our algorithm, we must first take each of the (at least) $n/2$ edges we found belonging to the MST and "merge" the two incident nodes into one.

Deleting Duplicate Edges Suppose we add edge (u, v) to our tree. The neighborhood of u and the neighborhood of v may overlap, hence we have two ways to get from super-node u-v to such a neighbor t. When we contract edge (u, v), we would yield a multi-graph, in which two vertices

may be connected by more than one edge. We wish to avoid this, and since our algorithm uses the Red-Rule (and is interested in the lowest weight edge in a cut), we can ignore the edge with larger weight. That is, if $\Gamma(u) \cap \Gamma(v) \neq \emptyset$ in our original graph G, then for each node t in the intersection of these two neighborhoods, we compare $w(u,t)$ with $w(v,t)$ and delete the edge with larger weight.

Recursive formulation But notice that when we contract the edges, the resulting "super" node (consisting of multiple nodes) itself defines a cut. Hence we can again apply the Red-Rule and find the smallest weight edge incident to each of the nodes in the *contracted* graph, and again continue to find more edges belonging to the MST. Hence we repeat the algorithm until we realize an MST. Since each time we find at least $n/2$ edges, we only must repeat our algorithm at most $\log_2 n$ times before finding an MST.

2.3 Boruvka's Algorithm

We assume that the graph is stored in an Adjacency-List, i.e. for each node we store its neighbors in a linked list. Since each node can have at most $n - 1$ neighbors, each adjacency list can have at most $n - 1$ entries. There are n adjacency lists (one for each node).

Algorithm 2: Boruvka's Algorithm

   Input : Graph G, represented by adjacency list
1 for *each node, in parallel* **do**
2   Compute smallest weight incident edge   `// `$O(m)$` work, `$O(\log n)$
             `depth`
3 end
4 Contract edges
5 Merge Adjacency Lists   `// `$O(n)$` work and `$O(1)$` depth`
6 Recurse

2.4 Analysis

2.4.1 Finding Smallest Weight Incident Edges

Examine the initial `for` loop which is executed in parallel. We know that we're looking for smallest weight edges, hence we know that we need to examine each edge at least once, i.e. we require at least $O(m)$ work. By the handshake lemma, $\sum_i \deg(v_i) = 2|E|$. Since for each node $v \in V$, we need to find the minimal weight incident edge, we must scan through $2|E|$ neighbors in total. We may use our parallel quickSelect algorithm on each adjacency list which requires $O(\deg(v_i))$ work and $O(\log \deg(v_i))$ depth per node. To bound total work, we must sum across all nodes, hence total work is $O(m)$. For depth, we are concerned with $\max_i \deg(v_i)$, which is $O(n)$. Hence depth is $O(\log n)$.

2.4.2 Contracting Edges

Now that we've shown that we remove $n/2$ isolated vertices in each round, we need to analyze how many edges we remove in each iteration. This will give a full picture of how fast are our subproblems decreasing in size. Notice that the number of edges removed in a contraction is *at least* equal to the number of vertices (since each vertex selects a lowest weight incident edge). Consider a best case possible sequence of events.

round	vertices	edges
1	n	m
2	$n/2$	$m - n/2$
3	$n/4$	$m - n/2 - n/4 = m - 3n/4$
$\vdots$	$\vdots$	$\vdots$
k	$n/2^{k-1}$	$m - (2^{k-1} - 1)n/(2^{k-1})$

Notice that for all $k \in \mathbb{N}$,

$$n \frac{2^{k-1} - 1}{2^{k-1}} = n \underbrace{\left(1 - \frac{1}{2^{k-1}}\right)}_{<1}$$

hence the number of edges never quite drops below $m - n$. So as long as there are $m > 2n$ edges to start with, work is $O(m \log n)$.

Crucial Observation: Minimum Weight Edges form a Forest   Crucially, in the first stage of our algorithm, each node its own connected component. It's also a tree, since a component of size 1 with 0 edges is connected and acyclic hence it's a tree. Realize that when we expand via minimum weight edges, since these edges are guaranteed to be in the MST, they *cannot* form a cycle, hence the result we get is a *forest*. Hence, our job is now to actually contract an entire tree (such that we may apply this to each tree in the forest).

To contract a tree, have each node (in parallel) flip a coin. If a vertex flips heads, it becomes the *center* of a star. If it flips tails, then the vertex attempts to become a *satellite* of (some) star by finding a neighbor which is a center.[5] If no such neighbor exists, i.e. if all other neighbors flipped tails or the vertex is isolated, then the vertex becomes a center.

When we perform a contraction, we need to select one node as the center, and the rest are satellites. Edges between satellites and centers must be deleted. Edges which cross partitions must be kept.[6]

---

[5]If multiple a vertex has multiple neighbors which are centers, it may select one arbitrarily. For example, we select the one with lowest node index.

[6]We now turn to Lemma 16.17 of Blelloch and Maggs.

Claim 2.1 *For a graph G with n non-isolated vertices, let X_n be the random variable indicating the number of satellites in a call to our tree contraction process described above. Then,*

$$\mathbb{E}[X_n] \geq n/4.$$

Proof: Consider any non-isolated vertex $v \in V(G)$. By definition, a non-isolated vertex has at least one neighbor, hence the probability that v becomes a satellite is *at least* $\Pr(\text{node } v \text{ flips Tails}) \times \Pr(\text{neighbor flips Heads}) \geq \frac{1}{4}$. Hence by linearity of expectation,

$$\mathbb{E}[\# \text{ Isolated Vertices}] = \sum_{\text{non-isolated vertices}} \mathbb{E}[\text{vertex } i \text{ becomes isolated}] \geq n/4.$$

■

Contracting a tree yields another Tree Since when doing start contraction on a tree, it remains a tree on each step, the number of edges does down with the number of vertices. Further, since a tree on n nodes has $n-1$ edges, the number of edges in a forest is never more than n. Hence the number of edges in our graph decrease geometrically in expectation. Hence total expected work given by

$$\mathbb{E}[W(n,m)] = \sum_{i=0}^{\infty} \left(\frac{3}{4}\right)^i kn = O(n).$$

The *Depth* of the contraction is given by $\mathbb{E}[D(n,m)] = O(\log n)$ since when we accumulate *neighborhood lists*, adding an edge to it requires $O(\log n)$ depth since the height of the neighbor hood tree is at most $O(\log n)$.

2.4.3 Merge Adjacency Lists

In terms of our data structures, we have n adjacency lists, and if we are contracting edge (u, v), then we must merge their corresponding adjacency lists together. Suppose that our `list` data structure is singly linked with a head to both the head and tail. Since $u < v$ via lexicographical comparison, we choose to make u the "parent" node, i.e. v's neighbors are absorbed into the neighborhood of u. Hence, we may simply take the tail of u's adjacency list and point it to the head of v's. This takes $O(1)$ work.

Since we must contract at least $n/2$ edges each iteration, this requires $O(n)$ work total. Notice that we may contract each edge independently of the rest, hence we have $O(1)$ depth.

2.4.4 Total Work and Depth

Total Work Notice that our algorithm employs *unary tail recursion*. We have a chain of $O(\log n)$ recursive calls to our algorithm. At each recursive call, we require $O(m)$ work to be performed, since our bottleneck is in finding the smallest weight edge incident each component. Hence total work simply $W(n, m) = O(m \log n)$.

Total Depth   Note that at each of the $O(\log n)$ recursive calls, we require $O(\log n)$ depth in order to perform a min-scan. Hence total depth given by $D(n, m) = O(\log^2 n)$.

References

[1] J. NESETRIL, *On minimum spanning tree problem (translation)*, Elsevier, (2001).