

Cosc 201
Algorithms and Data Structures
Lecture 18 (4/5/2026)
Two proofs

Brendan McCane
brendan.mccane@otago.ac.nz
And
Michael Albert



Proving connection and union-find

Connection and union-find

A graph, G , is connected if and only if the union-find instance formed by starting with the vertices of G and taking the union between the endpoints of each edge has only one group.

Actually we'll prove something stronger.

Connection and union-find

In a graph, G there is a walk from v to w if and only if the union-find instance formed by starting with the vertices of G and taking the union between the endpoints of each edge v and w belong to the same group, i.e., $find(v) = find(w)$.

The proof (I)

- ▶ Since we do a union along every edge, the two endpoints of any edge belong to the same group.
- ▶ So, if there is a walk from v to w , say $v = v_0, v_1, \dots, v_k = w$ then v_0 and v_1 belong to the same group but so do v_1 and v_2 and v_2 and v_3 . Eventually, we conclude that v (as v_0) and w (as v_k) belong to the same group.
- ▶ The remaining problem is to show that if v and w belong to the same group then there must be a walk between them.
- ▶ To do this, we'll show that at any point during the process of doing the union operations and for any vertex v there is always a walk from v to any other vertex u such that $find(v) = find(u)$, i.e., to any other vertex in the same group as v at that point.
- ▶ This will be by induction over the sequence of union operations performed.
- ▶ This is certainly true before the first union operation since $find(v) = find(u)$ means $v = u$ before the first *union* operation.

The proof (II)

- ▶ Suppose the property holds and we now do our next union $union(x, y)$ for some edge xy of G . Does it still hold?
- ▶ If, previously, $find(x) = find(y)$ then (in terms of groups) nothing has changed, so the property we want still persists.
- ▶ Otherwise, the group containing x and that containing y have merged. Suppose that v belonged to one of them (the group containing x say).
- ▶ Can we now walk from v to any vertex in the new group?
- ▶ Yes! Because we could already walk to x (and any vertex in its original group), and now we can follow the edge from there to y and then from y to any vertex in *its* original group.
- ▶ So, by induction, the property holds after each and every *union* operation (including the last one, which is when we need it).

Is union-find the correct way to determine if a graph is connected?

- ▶ Probably not!
- ▶ Union-find deals with the *dynamic* addition of edges and analyses the evolution of the groups that they form.
- ▶ If we have a *static* (in the sense of unchanging) graph **we just don't care** what the connectivity status is after we've looked at half the edges. So union-find is probably doing too much work.
- ▶ To find the set of vertices, c we can reach from v by walks we just modify the BFT a bit (see next slide).
- ▶ In common graphics contexts (where we are working in a graph which is a grid) variations on this idea include the **flood fill algorithm**.

BFT to find the vertices reachable from v

q : an initially empty queue.

c : an initially empty set of vertices.

```
1:  $q.addAll(v.neighbours)$ ,  $c.add(v)$ 
2: while  $q$  is not empty do
3:    $w \leftarrow q.remove()$ ,  $c.add(w)$ 
4:   for  $n$  in  $w.neighbours$  do
5:     if  $n$  is not in  $c$  then
6:        $q.add(n)$ ,  $c.add(n)$ 
7:     end if
8:   end for
9: end while
```

Reminder: Dijkstra's shortest path algorithm

```
1: procedure DIJKSTRA( $G, v$ )
2:    $dist[v] \leftarrow 0$ 
3:   for each edge  $e = (v, u)$  from  $v$  do
4:      $q.add(e, e.weight)$ 
5:   end for
6:   while  $q$  is not empty do
7:      $e = (x, y) \leftarrow q.remove()$ 
8:     if  $y.dist$  is not known then
9:        $dist[y] \leftarrow dist[x] + e.weight$ 
10:       $y.parent \leftarrow x$ 
11:      for each edge  $e = (y, z)$  from  $y$  do
12:        if  $dist[z]$  is not known then
13:           $q.add(e, dist[y] + e.weight)$ 
14:        end if
15:      end for
16:    end if
17:  end while
18: end procedure
```

The weighted shortest-path algorithm works?

The data structures associated with the weighted single-source (*base*) shortest-path algorithm are:

- dist* A map from vertices, w , to integers that represents the known distance of the least-distance path from *base* to w . Vertices are supposed to be added to this map in increasing order of distance.
- parent* A map from vertices, w , to vertices that represents the parent of w along the least-distance path from *base* to w .
- pq* A min-priority queue of edges. The source of each edge in *pq* is a vertex of known distance, and the priority of the edge is the sum of that distance and the weight of the edge.

Our job is to analyse the process of the algorithm and see that at all times these conditions are satisfied.

Initialisation

The initialisations are:

dist The dist of *base* is set to 0.

parent The parent of *base* is set to some “end of path” marker

pq Each edge with source *base* is added.

These are all appropriate for the algorithm (i.e., satisfy the conditions).

Polling the queue

When we poll the queue, we consider the target, t , of the edge removed. There are two cases:

- ▶ Its distance is already known, in which case we continue.
- ▶ Its distance is not yet known, in which case we assign it.

We need to check that both these are appropriate actions. To do this we proceed by induction on the number of keys in $dist$ and add the condition that, if there are k keys in $dist$ then these are the k least-distance neighbours of $base$.

Known distance case

The only thing that could go wrong is if we had somehow found a better path to t than the *dist* (and *parent*) maps recognise.

For this to happen the source, s of that edge has to have lesser distance than the distance of t . So, that edge was added to pq before we knew the distance of t . But then, the edge was already in pq at the time we determined the distance of t and would have been preferred to the choice we allegedly made.

So that can't happen and it's safe to just continue.

Unknown distance case

We now need to convince ourselves that t is the next-closest vertex to $base$.

Whatever that vertex is, the path to it must consist of a path to a vertex of already known distance plus a single edge. But, all these edges are already in pq .

The edge we've just removed is the one of least priority, i.e., representing the path of least distance – which is exactly what we want!

Huzzah!