

SOI Workshop 2015


Network Flows

Daniel Graf



ETH Zürich

November 8, 2015

Network Flow: Example




Network Flow: Example



Network Flow: Problem Statement

Input: A flow network consisting of

- directed graph $G = (V, E)$
- source and sink $s, t \in V$
- edge capacity $c : E \rightarrow \mathbb{N}$.



Output: A flow function $f : E \rightarrow \mathbb{N}$ such that:

- all capacity constraints are satisfied:
 $\forall u, v \in V : 0 \leq f(u, v) \leq c(u, v)$
(no pipe is overflowed)

- flow is conserved at every vertex:

$$\forall u \in V \setminus \{s, t\} :$$

$$\sum_{(v,u) \in E} f(v, u) = \sum_{(u,v) \in E} f(u, v)$$


(no vertex is leaking)

- the total flow is maximal:

$$|f| = \sum_{v \in V} f(s, v) - \sum_{v \in V} f(v, s) = \sum_{u \in V} f(u, t) - \sum_{u \in V} f(t, u)$$


Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Take any s - t -path and increase the flow along it.
- Update capacities and repeat as long as we can.
- Problem: We can get stuck at a local optimum.




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Take any s - t -path and increase the flow along it.
- Update capacities and repeat as long as we can.
- Problem: We can get stuck at a local optimum.




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Take any s - t -path and increase the flow along it.
- Update capacities and repeat as long as we can.
- Problem: We can get stuck at a local optimum.




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Take any s - t -path and increase the flow along it.
- Update capacities and repeat as long as we can.
- Problem: We can get stuck at a local optimum.




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Take any s - t -path and increase the flow along it.
- Update capacities and repeat as long as we can.
- Problem: We can get stuck at a local optimum.




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [BGL-Doc].




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [\[BGL-Doc\]](#).




Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [\[BGL-Doc\]](#).



Network Flow Algorithms: Ford-Fulkerson and Edmonds-Karp

- Solution: Keep track of the flow and allow paths that *reroute* units of flow. These are called *augmenting paths* in the *residual network*.
- Ford-Fulkerson: Repeatedly take any augmenting path: running time $\mathcal{O}(m|f|)$.
- Edmonds-Karp: Repeatedly take the shortest augmenting path: running time: best of $\mathcal{O}(m|f|)$, $\mathcal{O}(nm \max c)$ and $\mathcal{O}(nm^2)$. [\[BGL-Doc\]](#).



Edmonds-Karp Maximum Flow Algorithm

The headers and edge struct:

```
1 // Sample implementation of the Edmonds Karp Algorithm
2 // Daniel Graf, grafdan@ethz.ch, 7.11.2015
3 #include <iostream>
4 #include <vector>
5 #include <cassert>
6 #include <queue>
7
8 #define INF 1000000000
9 typedef long long int in;
10 using namespace std;
11
12 struct Edge {
13     in from, to, flow, cap, rev;
14     in residual_capacity() {
15         return cap-flow;
16     }
17 };
18 ...
```

Edmonds-Karp Maximum Flow Algorithm

Graph struct, add_edge function:

```
1 ...
2 struct Graph {
3     in s, t;
4     vector<vector<Edge> > E; // adjacency-list of edges
5     vector<Edge*> P; // predecessor map for the BFS
6
7     Graph(in N) {
8         E = vector<vector<Edge> >(N);
9     }
10
11     void add_edge(in from, in to, in cap) {
12         if(from==to) return;
13         E[from].push_back({from,to,0,cap,(in)E[to].size()});
14         E[to].push_back({to,from,0,0,(in)E[from].size()-1});
15     }
16 ...
```

Edmonds-Karp Maximum Flow Algorithm

Reset all flow values:

```
1 ...
2 void reset_flow() {
3     for(in v=0; v<E.size(); v++) {
4         for(in e=0; e<E[v].size(); e++) {
5             E[v][e].flow = 0;
6         }
7     }
8 }
9 ...
```

Edmonds-Karp Maximum Flow Algorithm

Can we find a path from s to t ?

```
1 bool find_flow() {
2     P = vector<Edge*>(E.size(), NULL);
3     // Breadth First Search through the edges with remaining capacity
4     queue<in> Q; Q.push(s);
5     while(!Q.empty() && P[t]==NULL) {
6         in v = Q.front(); Q.pop();
7         for(in e=0; e<E[v].size(); e++) {
8             if(E[v][e].residual_capacity()==0) {
9                 continue;
10            }
11            in w = E[v][e].to;
12            if(P[w]==NULL) {
13                P[w] = &(E[v][e]);
14                Q.push(w);
15            }
16        }
17    } ...
}
```


Edmonds-Karp Maximum Flow Algorithm

How much can we fit through that path from s to t?

```
1 ...
2 // Check if there is a path to t
3 if(P[t] == NULL) {
4     return 0;
5 }
6 // Check the minimum capacity
7 in flow = INF;
8 in pos = t;
9 while(pos != s) {
10     flow = min(flow, P[pos]->residual_capacity());
11     pos = P[pos]->from;
12 }
13 return flow;
14 }
```

Edmonds-Karp Maximum Flow Algorithm

Increase the flow along the path and reduce the reverse edges.

```
1 void update_flow(in flow) {
2     in pos = t;
3     while(pos != s) {
4         // cout << "update at vertex " << pos << endl;
5         P[pos]->flow += flow;
6         E[P[pos]->to][P[pos]->rev].flow -= flow;
7         pos = P[pos]->from;
8     }
9 }
```

Edmonds-Karp Maximum Flow Algorithm

Repeatedly search for the shortest augmenting path while one exists.

```
1      in edmonds_karp_max_flow(in _s, in _t) {
2          s = _s;
3          t = _t;
4          reset_flow();
5          in flow = 0;
6          in new_flow;
7          do {
8              new_flow = find_flow();
9              flow += new_flow;
10             if(new_flow > 0) {
11                 update_flow(new_flow);
12             }
13         } while(new_flow > 0);
14         return flow;
15     }
16 }; // end of the Graph struct
```


Edmonds-Karp Maximum Flow Algorithm

Read graph from input and call the algorithm.


```
1 void read_graph_from_stdin(Graph &G) {
2     in N,M;
3     cin >> N >> M;
4     G = Graph(N);
5     for(in m=0; m<M; m++) {
6         in a,b,c;
7         cin >> a >> b >> c;
8         G.add_edge(a,b,c);
9     }
10 }
11 int main() {
12     Graph G(0);
13     read_graph_from_stdin(G);
14     in s, t; cin >> s >> t;
15     in res = G.edmonds_karp_max_flow(s,t);
16     cout << "max_flow: " << res << endl;
17 }
```

Common tricks


Multiple sources/sinks:



Undirected Graphs



Vertex Capacities




Minimum Flow per Edge


[Exercise]

Common tricks


Multiple sources/sinks:



Undirected Graphs



Vertex Capacities




Minimum Flow per Edge


[Exercise]

Common tricks


Multiple sources/sinks:



Undirected Graphs



Vertex Capacities




Minimum Flow per Edge


[Exercise]

Common tricks


Multiple sources/sinks:



Undirected Graphs



Vertex Capacities




Minimum Flow per Edge


[Exercise]

Common tricks


Multiple sources/sinks:



Vertex Capacities



Undirected Graphs




Minimum Flow per Edge


[Exercise]

Common tricks


Multiple sources/sinks:



Undirected Graphs



Vertex Capacities




Minimum Flow per Edge


[Exercise]

Common tricks


Multiple sources/sinks:



Undirected Graphs



Vertex Capacities

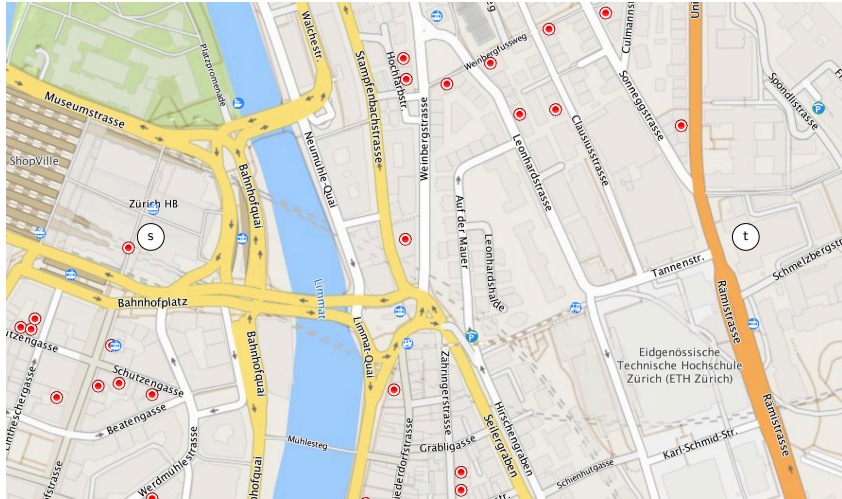


Minimum Flow per Edge

[Exercise]

Flow Application: Edge Disjoint Paths

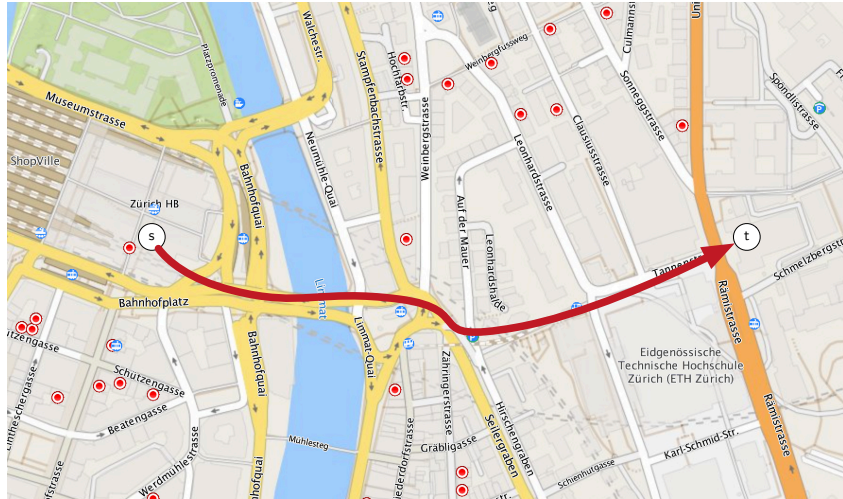
How many ways are there to get from HB to CAB without using the same street twice?



Map:
search.ch,
TomTom,
swisstopo,
OSM

Flow Application: Edge Disjoint Paths

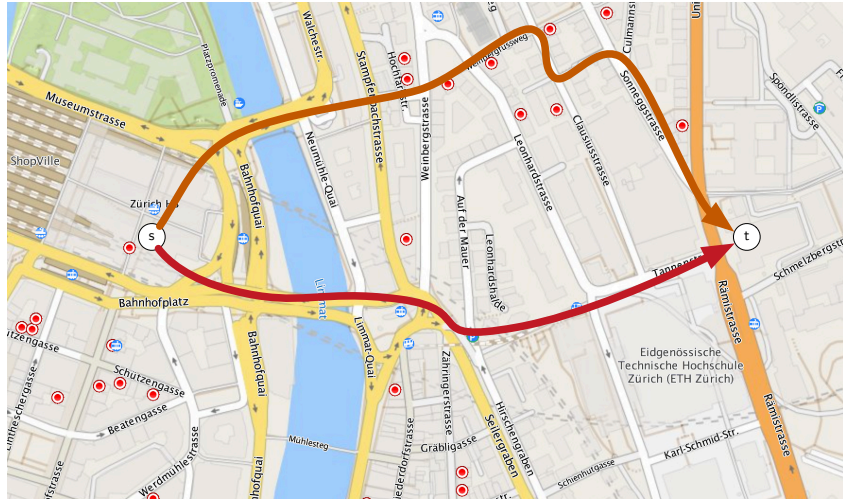
How many ways are there to get from HB to CAB without using the same street twice?



Map:
search.ch,
TomTom,
swisstopo,
OSM

Flow Application: Edge Disjoint Paths


How many ways are there to get from HB to CAB without using the same street twice?



Map:
search.ch,
TomTom,
swisstopo,
OSM

Flow Application: Edge Disjoint Paths

How many ways are there to get from HB to CAB without using the same street twice?



Map:
search.ch,
TomTom,
swisstopo,
OSM

Flow Application: Edge Disjoint Paths

How many ways are there to get from HB to CAB without using the same street twice?

- Is this a flow problem? No.
- Can it be turned into a flow problem? Maybe.
- Build directed street graph by adding edges in both directions.
- Set all capacities to 1.

Lemma

In a directed graph with unit capacities, the maximum number of edge-disjoint s - t -paths is equal to the maximum flow from s to t .

Flow Application: Edge Disjoint Paths

How many ways are there to get from HB to CAB without using the same street twice?

- Is this a flow problem? No.
- Can it be turned into a flow problem? Maybe.
- Build directed street graph by adding edges in both directions.
- Set all capacities to 1.

Lemma

In a directed graph with unit capacities, the maximum number of edge-disjoint s - t -paths is equal to the maximum flow from s to t .

Flow Application: Edge Disjoint Paths

How many ways are there to get from HB to CAB without using the same street twice?

- Is this a flow problem? No.
- Can it be turned into a flow problem? Maybe.
- Build directed street graph by adding edges in both directions.
- Set all capacities to 1.

Lemma

In a directed graph with unit capacities, the maximum number of edge-disjoint s - t -paths is equal to the maximum flow from s to t .

Flow Application: Edge Disjoint Paths


How many ways are there to get from HB to CAB without using the same street twice?

- Is this a flow problem? No.
- Can it be turned into a flow problem? Maybe.
- Build directed street graph by adding edges in both directions.
- Set all capacities to 1.

Lemma

In a directed graph with unit capacities, the maximum number of edge-disjoint s - t -paths is equal to the maximum flow from s to t .

Flow Application: Edge Disjoint Paths



Map:
search.ch,
TomTom,
swisstopo,
OSM