



Admin

- Midterm 2
 - Due in exactly 1 week
 - Possible topics:
 - Greedy Algorithms
 - Dynamic Programming
 - Graph Algorithms
 - Earlier materials?
 - Review questions?
 - If you have questions you'd like to see covered Thursday, write them down on today's worksheet!

Where are we?

- Last time:
 - Discussed three different algorithms for computing APSP
- This week:
 - Network flow
- Ahead:
 - Linear Programming
 - NP-Completeness
 - Parallel Algorithms

All-Pairs Shortest Paths


- Run Bellman-Ford, once from each vertex
 - $O(V^2E)$
 - $O(V^4)$ if dense (E in $\Theta(V^2)$)
- Run Dijkstra's, once from each vertex
 - Limited to graphs with non-negative edges
 - $O(VE \log V)$ with binary heap
 - $O(V^3 \log V)$ if dense
 - $O(V^2 \log V + VE)$ with Fibonacci heap
 - $O(V^3)$ if dense
- Wouldn't it be nice to get the best of both worlds:
 - General algorithm that runs on any graph
 - $O(V^3)$ worst-case runtimes?
 - No funny complex data structures



Floyd-Warshall

- DP over the distance matrix (store values per edge, rather than per node)

- Intermediate vertex:



- Recursive formulation:

$$d_{ij}^{(k)} = \begin{cases} w_{ij} & \text{if } k = 0, \\ \min(d_{ij}^{(k-1)}, d_{ik}^{(k-1)} + d_{kj}^{(k-1)}) & \text{if } k \geq 1. \end{cases}$$

Floyd-Warshall

FLOYD-WARSHALL(W, n)

$D^{(0)} = W$

for $k = 1$ **to** n

let $D^{(k)} = (d_{ij}^{(k)})$ be a new $n \times n$ matrix

for $i = 1$ **to** n

for $j = 1$ **to** n

$d_{ij}^{(k)} = \min(d_{ij}^{(k-1)}, d_{ik}^{(k-1)} + d_{kj}^{(k-1)})$

return $D^{(n)}$

Time


$\Theta(n^3)$

Johnson's Algorithm


- The Problem:
 - If graph is sparse, running Dijkstra from each vertex would be faster
 - $O(V^2 \log V + VE)$ (Requires Fibonacci heap)
 - Idea:
 - Reweight edges so that they are non-negative
 - And still give you correct shortest paths (paths that are min-weight in the original graph)
- Add new node q , connect to each of the other nodes with zero-weight edges
 - Run Bellman-Ford from q to calculate min path to every node $h(v)$; return false if negative cycle is found
 - Reweight all edge weights w' using the rule $w'(u, v) = w(u, v) + h(u) - h(v)$
 - Remove q , run Dijkstra's from every node

Johnson's Algorithm


- The idea:
 - Reweight edges so that they are non-negative
 - And still give you correct shortest paths (paths that are min-weight in the original graph)



original graph
with negative edges



shortest path tree
found by Bellman-Ford



reweighted graph with
no negative edges

Worksheet: Johnson's ++?

- Proposal:

- Save steps 1-3 by using simpler reweighting function:


$$w^* = \min_{(u,v) \in E} w(u,v)$$

$$\hat{w}(u,v) = w(u,v) - w^* \quad \forall (u,v) \in E$$



- Why doesn't this work?

For Midterm 2: It's a good idea to study WHY graph algorithms work: e.g., What makes an edge "safe" to include in a solution? Why does a greedy choice work? When does a DP approach converge?,... etc.



All pairs shortest paths

Simple approach

- Call Bellman-Ford $|V|$ times
- $O(|V|^2 |E|)$


Floyd-Warshall – $O(|V|^3)$

Johnson's algorithm – $O(|V|^2 \log |V| + |V| |E|)$

Pettie's algorithm – $O(VE + V^2 \log \log V)$

Snowball [Planken et al. 2011]

$$\mathcal{O}(nm_c) \subseteq \mathcal{O}(n^2 w_d)$$



Max Credit Flow

On your worksheet:

- Select two integers A & B such that
 - A, B in $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$
 - $A+B \leq 10$

Rules of the game:

1. You can only pass forward or to your right (if there is someone in front or to the right of you)
2. You can only pass up to A units forward, B units to the right.
3. I will pass out 10 extra credit points in the back-left corner of the room.
4. Everyone gets extra credit participation points in proportion to the number of points that flow to the front, right-most student.



Network Flow

- The idea: Use a graph to model material that flows through conduits
 - Each edge represents one conduit and has a capacity, which is an upper bound on the flow rate = units/time
 - We want to compute max rate that we can ship material from a designated source to designated sink.



Network Flow (Key Ideas Preview)

- Terminology / Notation:
 - Capacity:
 - Source:
 - Sink:
 - Flow:
 - Capacity constraint:
 - Flow conservation
 - Value of flow:
- Maximum-flow problem
- Cuts:
 - Net flow:
 - Capacity:



Three Theorems about Flow

The Cut Theorem: For any cut S, T and flow f , the flow across the cut is equal to $|f|$. That is:

$$\sum_{x \in S, y \in T} f(x, y) = |f|$$


The Capacity Theorem: For any cut S, T and flow f , the flow is bounded by the capacity of the cut. That is: $|f| \leq \sum_{x \in S, y \in T} c(x, y)$

Max Flow/Min Cut Theorem: For any cut S, T and flow f , if $|f| = \sum_{x \in S, y \in T} c(x, y)$ then f is max flow and S, T is a min cut.

Student networking


You decide to create your own campus network:

- You get three of your friends and string some network cables
- Because of capacity (due to cable type, distance, computer, etc) you can only send a certain amount of data to each person
- If edges denote capacity, what is the maximum throughput you can send from S to T ?




Student networking

- You decide to create your own campus network:
 - You get three of your friends and string some network cables
 - Because of capacity (due to cable type, distance, computer, etc) you can only send a certain amount of data to each person
 - If edges denote capacity, what is the maximum throughput you can send from S to T?




Another flow problem



How much water flow can we continually send from s to t?


Another flow problem



Flow graph/networks


Flow network

- directed, weighted graph (V, E)
- positive edge weights indicating the “**capacity**” (generally, assume integers)
- contains a **single source** $s \in V$ with no incoming edges
- contains a **single sink/target** $t \in V$ with no outgoing edges
- every vertex is on a path from s to t




Flow

What are the constraints on flow in a network?



Max flow problem

Given a flow network: *what is the maximum flow we can send from s to t that meet the flow constraints?*



Flow

Flow Conservation:


in-flow = out-flow for every vertex (except s, t)

Capacity Constraint:

flow along an edge cannot exceed the edge capacity

Value of flow:

flows are positive



Applications?

network flow

- water, electricity, sewage, cellular...
- traffic/transportation capacity

bipartite matching

sports elimination

...

Max flow origins


Rail networks of the Soviet Union in the 1950's

The US wanted to know how quickly the Soviet Union could get supplies through its rail network to its satellite states in Eastern Europe.

In addition, the US wanted to know which rails it could destroy most easily to cut off the satellite states from the rest of the Soviet Union.

These two problems are closely related, and that solving the **max flow problem** also solves the **min cut problem** of figuring out the cheapest way to cut off the Soviet Union from its satellites.

Source: lbackstrom, The Importance of Algorithms, at www.topcoder.com



Max Credit Flow: Double down challenge

Double or nothing?!?

Using the same A&B that you chose earlier:

- A, B in {0,1,2,3,4,5,6,7,8,9}
- $A+B \leq 10$

You have 2:00 as a class to double your take!



Rules of the game:

1. You can only pass forward or to your right (if there is someone in front or to the right of you)
2. You can only pass up to A units forward, B units to the right.
3. I will pass out 10 extra credit points in the back-left corner of the room.
4. Everyone gets extra credit participation points in proportion to the number of points that flow to the front, right-most student.

Algorithm ideas?


graph algorithm?

- BFS, DFS, shortest paths...
- MST


divide and conquer?

greedy?

dynamic programming?




Algorithm idea




Algorithm idea

send some flow down a path



Algorithm idea


send some flow down a path



Now what?

Algorithm idea


reroute some of the flow




Total flow?

Algorithm idea

reroute some of the flow




Algorithm idea




Algorithm idea

send some flow down a path




Algorithm idea


send some flow down a path




Algorithm idea



Algorithm idea




Algorithm idea



Are we done?
Is this the best we can do?


Cuts

A cut is a partitioning of the vertices into two sets A and $B = V - A$



Flow across cuts


In flow graphs, we're interested in cuts that separate s from t , that is $s \in A$ and $t \in B$



Flow across cuts

The **net flow** “across” a cut is the total flow from nodes in A to nodes in B *minus* the total from from B to A


What is the flow across this cut?



Flow across cuts

The flow “across” a cut is the total flow from nodes in A to nodes in B *minus* the total from from B to A


$$10+10-6 = 14$$



Flow across cuts

Consider any cut where $s \in A$ and $t \in B$, i.e. the cut partitions the source from the sink


What do we know about the flow across the any such cut?



Flow across cuts

Consider any cut where $s \in A$ and $t \in B$, i.e. the cut partitions the source from the sink


The flow across ANY such cut is the same and is the current flow in the network



Flow across cuts

Consider any cut where $s \in A$ and $t \in B$, i.e. the cut partitions the source from the sink


$$4 + 10 = 14$$



Flow across cuts

Consider any cut where $s \in A$ and $t \in B$, i.e. the cut partitions the source from the sink


$$4 + 6 + 4 = 14$$



Flow across cuts

Consider any cut where $s \in A$ and $t \in B$, i.e. the cut partitions the source from the sink

$$10 + 10 - 6 = 14$$




Flow across cuts

Consider any cut where $s \in A$ and $t \in B$, i.e. the cut partitions the source from the sink

The flow across ANY such cut is the same and is the current flow in the network

Why? Can you prove it?



Flow across cuts

The flow across ANY such cut is the same and is the current flow in the network

Base case: $A = s$

- Flow is total from from s to t : therefore total flow out of s should be the flow
- All flow from s gets to t
 - every vertex is on a path from s to t
 - in-flow = out-flow

Flow across cuts

The flow across ANY such cut is the same and is the current flow in the network

Inductively?


- every vertex is on a path from s to t
- in-flow = out-flow for every vertex (except s, t)
- flow along an edge cannot exceed the edge capacity
- flows are positive

Flow across cuts

The flow across ANY such cut is the same and is the current flow in the network

Inductive case: Consider moving a node x from A to B

Is the flow across the different partitions the same?




Flow across cuts

The flow across ANY such cut is the same and is the current flow in the network

Inductive case: Consider moving a node x from A to B

$$\text{flow} = \text{left-inflow}(x) - \text{left-outflow}(x) \quad \text{flow} = \text{right-outflow}(x) - \text{right-inflow}(x)$$




$$\text{left-inflow}(x) + \text{right-inflow}(x) = \text{left-outflow}(x) + \text{right-outflow}(x) \quad \text{in-flow} = \text{out-flow}$$

$$\text{left-inflow}(x) - \text{left-outflow}(x) = \text{right-outflow}(x) - \text{right-inflow}(x)$$

Capacity of a cut

The “**capacity of a cut**” is the maximum flow that we *could* send from nodes in A to nodes in B (i.e. across the cut)

How do we calculate the capacity?




Capacity of a cut

The “**capacity of a cut**” is the maximum flow that we *could* send from nodes in A to nodes in B (i.e. across the cut)

Capacity is the sum of the edges from A to B

Why?



Capacity of a cut

The “**capacity of a cut**” is the maximum flow that we *could* send from nodes in A to nodes in B (i.e. across the cut)

Capacity is the sum of the edges from A to B


- Any more and we would violate the edge capacity constraint
- Any less and it would not be maximal, since we could simply increase the flow

Quick recap

A cut is a partitioning of the vertices into two sets A and $B = V - A$

For any cut where $s \in A$ and $t \in B$, i.e. the cut partitions the source from the sink

- the flow across any such cut is the same
- the maximum capacity (i.e. flow) across the cut is the sum of the capacities for edges from A to B



Worksheet: The kid's aren't alright


Professor Adam has two children who, unfortunately, dislike each other. The problem is so severe that not only do they refuse to walk to school together, but in fact each one refuses to walk on any block that the other child has stepped on that day. The children have no problem with their paths crossing at a corner. Fortunately both the professor's house and the school are on corners, but beyond that he is not sure if it is going to be possible to send both of his children to the same school. The professor has a map of his town. Show how to formulate the problem of determining whether both his children can go to the same school as a maximum-flow problem.



Maximum flow

For any cut where $s \in A$ and $t \in B$

- the flow across the cut is the same
- the maximum capacity (i.e. flow) across the cut is the sum of the capacities for edges from A to B




Are we done?
Is this the best we can do?

Maximum flow

For any cut where $s \in A$ and $t \in B$

- the flow across the cut is the same
- the maximum capacity (i.e. flow) across the cut is the sum of the capacities for edges from A to B




We can do no better than the minimum capacity cut!

Maximum flow

What is the minimum capacity cut for this graph?

Capacity = 10 + 4




Is this the best we can do?

Maximum flow

What is the minimum capacity cut for this graph?


Capacity = 10 + 4



flow = minimum capacity, so we can do no better

Algorithm idea


send some flow down a path



How do we determine the path to send flow down?

Algorithm idea


send some flow down a path



Search for a path with remaining capacity from s to t


Algorithm idea

reroute some of the flow



How do we handle "rerouting" flow?


Algorithm idea



During the search, if an edge has some flow, we consider "reversing" some of that flow

Algorithm idea

reroute some of the flow



During the search, if an edge has some flow, we consider "reversing" some of that flow


The residual graph

The **residual graph** G_f is constructed from G


For each edge e in the original graph (G):

- if $flow(e) < capacity(e)$
 - introduce an edge in G_f with capacity = $capacity(e) - flow(e)$
 - this represents the remaining flow we can still push
- if $flow(e) > 0$
 - introduce an edge in G_f in the **opposite direction** with capacity = $flow(e)$
 - this represents the flow that we can reroute/reverse


Algorithm idea





Algorithm idea



Algorithm idea





Algorithm idea





None exist... done!

Algorithm idea





Algorithm idea



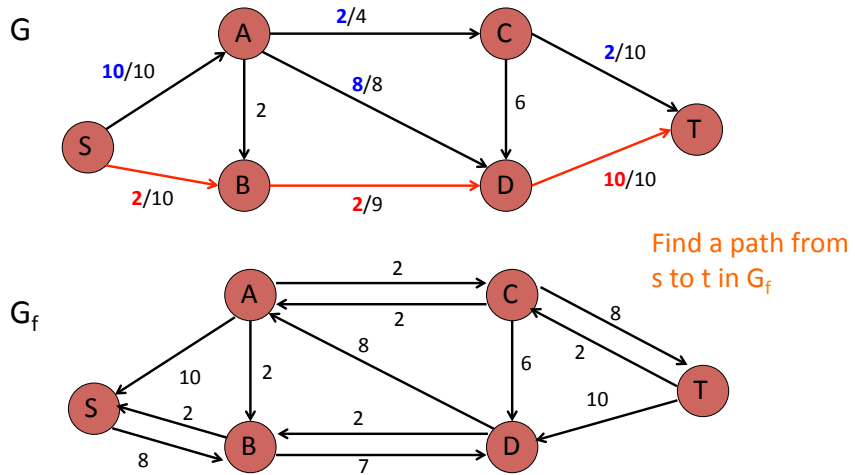
Find a path from
s to t in G_f

Algorithm idea

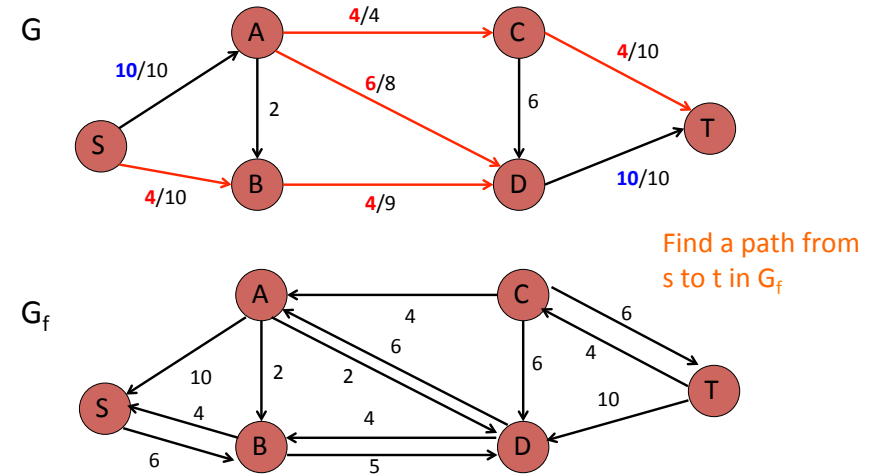


Find a path from
s to t in G_f


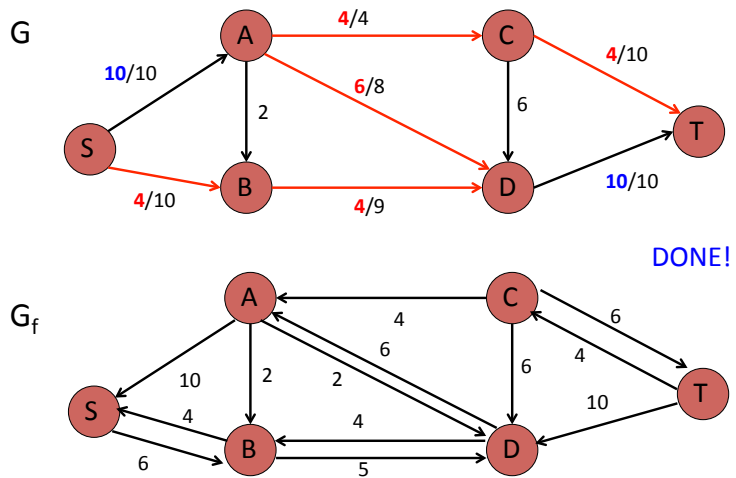
Algorithm idea



Algorithm idea



Algorithm idea



Ford-Fulkerson

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
  a simple path contains no repeated vertices
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from s to t in  $G_f$ 
    send as much flow along the path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

Ford-Fulkerson: is it correct?

Does the function terminate?

Every iteration increases the flow from s to t

- Every path must start with s
- The path has positive flow (or it wouldn't exist)
- The path is a simple path (so it cannot revisit s)
- conservation of flow

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from s to t in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

Ford-Fulkerson: is it correct?

Does the function terminate?

- Every iteration increases the flow from s to t
- the flow is bounded by the min-cut

Ford-Fulkerson: is it correct?

When it terminates is it the maximum flow?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from s to t in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from s to t in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

Ford-Fulkerson: is it correct?

When it terminates is it the maximum flow?

Assume it didn't

- We know then that the flow < min-cut
- therefore, the flow < capacity across EVERY cut
- therefore, across each cut there must be a forward edge in G_f
- thus, there must exist a path from s to t in G_f
 - start at s (and $A = s$)
 - repeat until t is found
 - pick one node across the cut with a forward edge
 - add this to the path
 - add the node to A (for argument sake)
- However, the algorithm would not have terminated... a contradiction

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from  $s$  to  $t$  in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from  $s$  to  $t$  in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

Can we simplify this expression?

- traverse the graph
- at most add 2 edges for original edge
- $O(V + E)$

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from  $s$  to  $t$  in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

- traverse the graph
- at most add 2 edges for original edge
- $O(V + E) = O(E)$
- (all nodes exists on paths exist from s to t)

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from s to t in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

- BFS or DFS
- $O(V + E) = O(E)$

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from s to t in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```

- max-flow!
- increases ever iteration
- integer capacities, so integer increases

Can we bound the number of times the loop will execute?

Ford-Fulkerson: runtime?

```
Ford-Fulkerson(G, s, t)
  flow = 0 for all edges
   $G_f = \text{residualGraph}(G)$ 
  while a simple path exists from s to t in  $G_f$ 
    send as much flow along path as possible
     $G_f = \text{residualGraph}(G)$ 
  return flow
```


- max-flow!
- increases ever iteration
- integer capacities, so integer increases

Overall runtime? $O(\text{max-flow} * E)$

$O(\text{max-flow} * E)$


Can you construct a graph that *could* get this running time?

Hint:




$O(\text{max-flow} * E)$

Can you construct a graph that *could get* this running time?




$O(\text{max-flow} * E)$

Can you construct a graph that *could get* this running time?




$O(\text{max-flow} * E)$

Can you construct a graph that *could get* this running time?




$O(\text{max-flow} * E)$

Can you construct a graph that *could get* this running time?




$O(\text{max-flow} * E)$

Can you construct a graph that *could get* this running time?




$O(\text{max-flow} * E)$

Can you construct a graph that *could get* this running time?



$O(\text{max-flow} * E)$

Can you construct a graph that *could get* this running time?



What is the problem here?
Could we do better?

Faster variants

Edmunds-Karp

- Select the *shortest path* (in number of edges) from s to t in G_f
 - How can we do this?
 - use BFS for search
- Running time: $O(V E^2)$
 - avoids issues like the one we just saw
 - see the book for the proof
 - or <http://www.cs.cornell.edu/courses/CS4820/2011sp/handouts/edmondskarp.pdf>

preflow-push (aka push-relabel) algorithms

- $O(V^3)$

Other variations...

Method	Complexity
Linear programming	
Ford-Fulkerson algorithm	$O(E \max f)$
Edmonds-Karp algorithm	$O(VE^3)$
Dinitz blocking flow algorithm	$O(V^2E)$
General push-relabel maximum flow algorithm	$O(V^2E)$
Push-relabel algorithm with FIFO vertex selection rule	$O(V^3)$
Dinitz blocking flow algorithm with dynamic trees	$O(VE \log(V))$
Push-relabel algorithm with dynamic trees	$O(VE \log(V^2/E))$
Binary blocking flow algorithm [1]	$O(E \min(V^{2/3}, \sqrt{E}) \log(V^2/E) \log U)$
MPM (Malhotra, Pramodh-Kumar and Maheshwari) algorithm	$O(V^3)$


http://en.wikipedia.org/wiki/Maximum_flow

TABLE I. POLYNOMIAL-TIME ALGORITHMS FOR THE MAXIMUM FLOW PROBLEM^a

Algorithm no.	Date	Discoverer	Running time	References
1	1969	Edmonds and Karp	$O(nm^2)$	[5]
2	1970	Dinic	$O(n^3m)$	[4]
3	1974	Karzanov	$O(n^3)$	[18]
4	1977	Cherkasky	$O(n^3m^{1/2})$	[3]
5	1978	Malhotra, Pramodh Kumar, and Maheshwari	$O(n^3)$	[21]
6	1978	Galil	$O(n^{3/2}m^{2/3})$	[11]
7	1978	Galil and Naamad; Shiloach	$O(nm(\log n)^2)$	[12, 25]
8	1980	Sleator and Tarjan	$O(nm \log n)$	[27, 28]
9	1982	Shiloach and Vishkin	$O(n^3)$	[26]
10	1983	Gabow	$O(nm \log U)$	[10]
11	1984	Tarjan	$O(n^3)$	[31]
12	1985	Goldberg	$O(n^3)$	[14]
13	1986	Goldberg and Tarjan	$O(nm \log(n^2/m))$	[16, 15]
14	1986	Aluja and Orlin	$O(nm + n^3 \log U)$	[1]

^a Algorithm 13 is presented in this paper.

http://akira.ruc.dk/~keld/teaching/algorithmdesign_f03/Artikler/08/Goldberg88.pdf



Network Flow (Key Ideas)

- Terminology / Notation:
 - Capacity:
 - Source:
 - Sink:
 - Flow:
 - Capacity constraint:
 - Flow conservation
 - Value of flow:
- Maximum-flow problem
- Cuts:
 - Net flow:
 - Capacity:



Three Theorems about Flow

The Cut Theorem: For any cut S, T and flow f , the flow across the cut is equal to $|f|$. That is:

$$\sum_{x \in S, y \in T} f(x, y) = |f|$$

The Capacity Theorem: For any cut S, T and flow f , the flow is bounded by the capacity of the cut. That is: $|f| \leq \sum_{x \in S, y \in T} c(x, y)$

Max Flow/Min Cut Theorem: For any cut S, T and flow f , if $|f| = \sum_{x \in S, y \in T} c(x, y)$ then f is max flow and S, T is a min cut.

Proof / Applications next time!