


CMPS 6610/4610 – Fall 2016



Flow Networks

Carola Wenk

Slides adapted from slides by Charles Leiserson


Max flow and min cut

- Fundamental problems in combinatorial optimization
- Duality between max flow and min cut
- Many applications:
 - Bipartite matching
 - Image segmentation
 - Airline scheduling
 - Network reliability
 - Survey design
 - Baseball elimination
 - Gene function prediction
 - ...

Flow networks

Definition. A *flow network* is a directed graph $G = (V, E)$ with two distinguished vertices: a *source* s and a *sink* t . Each edge $(u, v) \in E$ has a nonnegative *capacity* $c(u, v)$. If $(u, v) \notin E$, then $c(u, v) = 0$. We require that if $(u, v) \in E$ then $(v, u) \notin E$.

Example:



Flow networks

Definition. A (positive) *flow* on G is a function $f: V \times V \rightarrow \mathbb{R}$ satisfying the following:

- **Capacity constraint:** For all $u, v \in V$,
 $0 \leq f(u, v) \leq c(u, v)$.

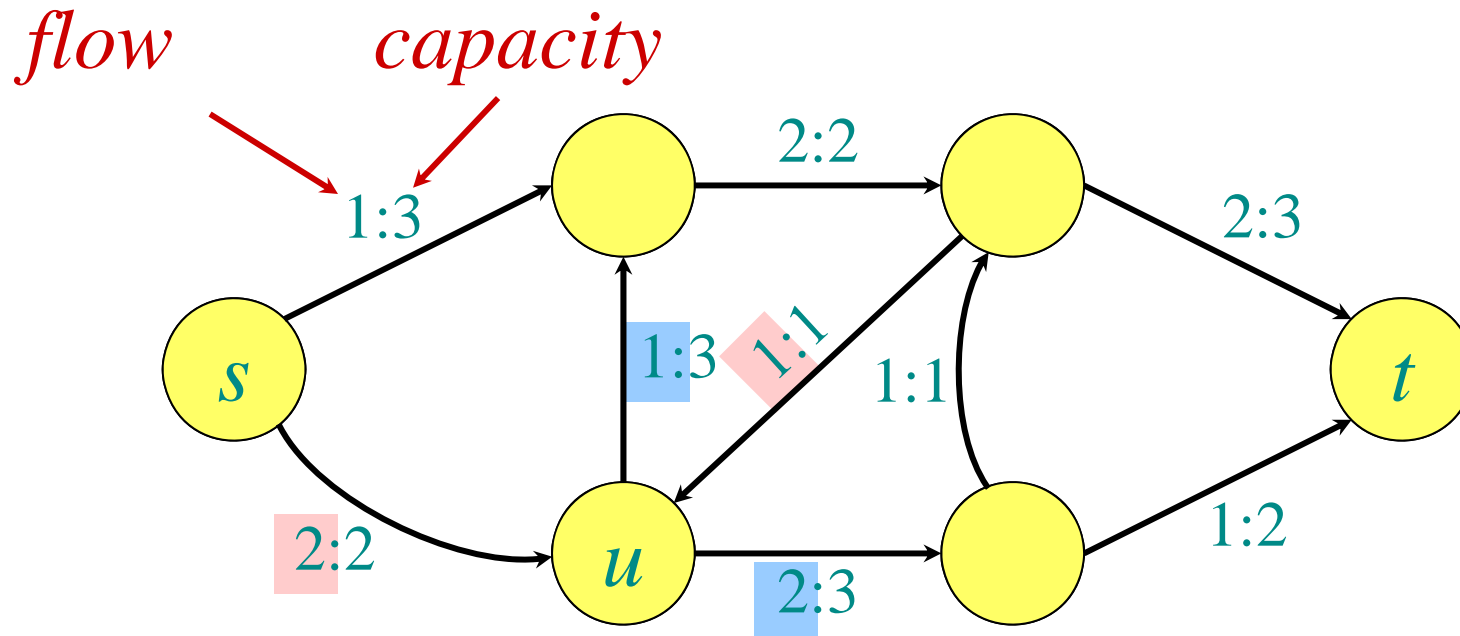
- **Flow conservation:** For all $u \in V \setminus \{s, t\}$,

$$\sum_{v \in V} f(u, v) - \sum_{v \in V} f(v, u) = 0$$

The *value* of a flow is the net flow out of the source:

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

A flow on a network




Flow conservation (like Kirchoff's current law):

- Flow into u is $2 + 1 = 3$.
- Flow out of u is $1 + 2 = 3$.

The value of this flow is $1 + 2 = 3$.

The maximum-flow problem

Maximum-flow problem: Given a flow network G , find a flow of maximum value on G .



The value of the maximum flow is 4.

Cuts

Definition. A *cut* (S, T) of a flow network $G = (V, E)$ is a partition of V such that $s \in S$ and $t \in T$.

If f is a flow on G , then the *net flow across the cut* is

$$f(S, T) = \sum_{u \in S} \sum_{v \in T} f(u, v) - \sum_{u \in S} \sum_{v \in T} f(v, u)$$

The *capacity of the cut* is

$$c(S, T) = \sum_{u \in S} \sum_{v \in T} c(u, v)$$

Another characterization of flow value

Lemma. For any flow f and any cut (S, T) , we have $|f| = f(S, T)$.

$$\begin{aligned}
 \textit{Proof: } |f| &= \sum_{v \in V} f(s, v) - \sum_{v \in V} f(v, s) \\
 &= \sum_{v \in V} f(s, v) - \sum_{v \in V} f(v, s) + \overbrace{\sum_{u \in S \setminus \{s\}} \left(\sum_{v \in V} f(u, v) - \sum_{v \in V} f(v, u) \right)}^0 \\
 &= \sum_{v \in V} \sum_{u \in S} f(u, v) - \sum_{v \in V} \sum_{u \in S} f(v, u) \\
 &= \sum_{v \in S} \sum_{u \in S} \cancel{f(u, v)} + \sum_{v \in T} \sum_{u \in S} f(u, v) - \sum_{v \in S} \sum_{u \in S} \cancel{f(v, u)} - \sum_{v \in T} \sum_{u \in S} f(v, u) \\
 &= \sum_{v \in T} \sum_{u \in S} f(u, v) - \sum_{v \in T} \sum_{u \in S} f(v, u) = f(S, T)
 \end{aligned}$$

Upper bound on the maximum flow value

Theorem. The value of any flow is bounded from above by the capacity of any cut:


$$|f| \leq c(S, T).$$

Proof.

$$\begin{aligned} |f| &= f(S, T) \\ &= \sum_{u \in S} \sum_{v \in T} f(u, v) - \sum_{u \in S} \sum_{v \in T} f(v, u) \\ &\leq \sum_{u \in S} \sum_{v \in T} f(u, v) \\ &\leq \sum_{u \in S} \sum_{v \in T} c(u, v) \\ &= c(S, T) \end{aligned}$$



Flow into the sink



$$|f| = f(\{s\}, V \setminus \{s\}) = f(V \setminus \{t\}, t) = 4$$

Residual network

Definition. Let f be a flow on $G = (V, E)$. The *residual network* $G_f(V, E_f)$ is the graph with *residual capacities*

$$c_f(u, v) = \begin{cases} c(u, v) - f(u, v), & \text{if } (u, v) \in E \\ f(v, u) & , \text{if } (v, u) \in E \\ 0 & , \text{otherwise} \end{cases}$$

Edges in E_f admit more flow.

Lemma. $|E_f| \leq 2|E|$. □


Augmenting paths

$$c_f(u, v) = \begin{cases} c(u, v) - f(u, v), & \text{if } (u, v) \in E \\ f(v, u) & , \text{if } (v, u) \in E \\ 0 & , \text{otherwise} \end{cases}$$


Definition. Let p be a path from s to t in G_f . The **residual capacity** of p is $c_f(p) = \min_{(u,v) \in p} \{c_f(u, v)\}$.

If $c_f(p) > 0$ then p is called an **augmenting path** in G with respect to f . The flow value can be increased along an augmenting path p by $c_f(p)$.

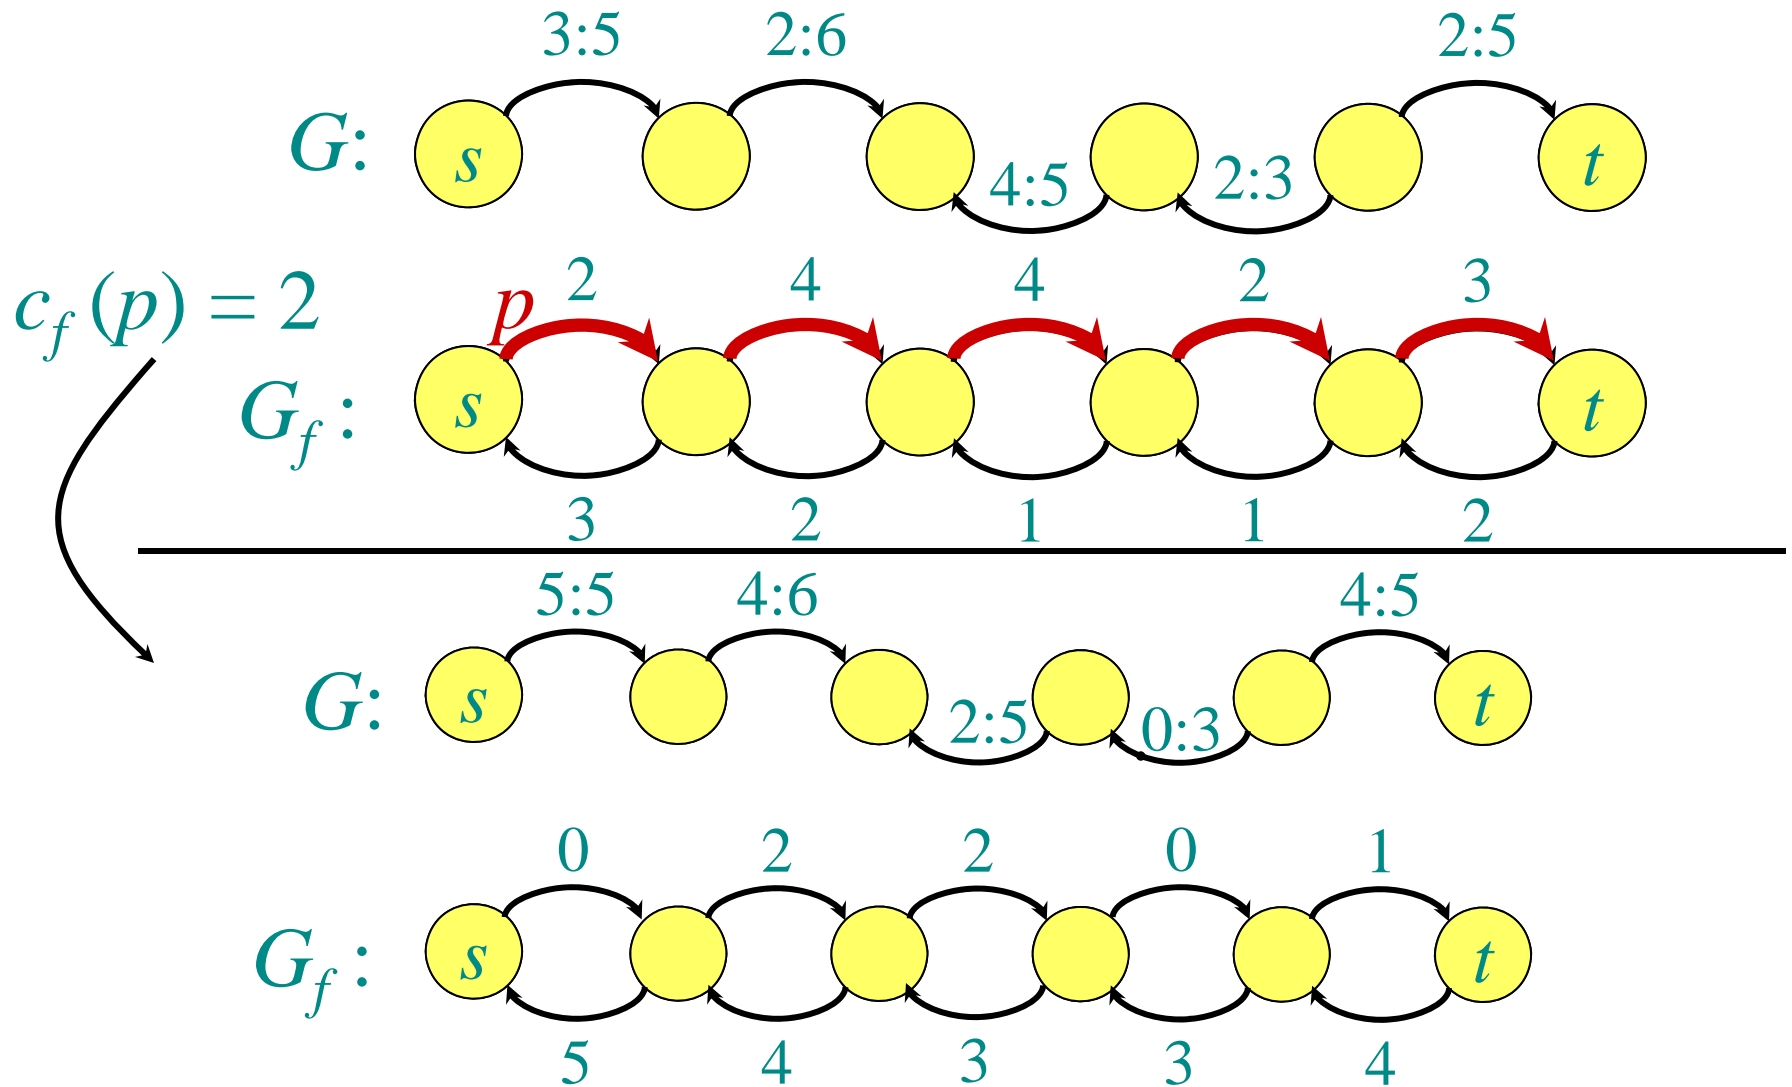
Ex.:



$c_f(p) = 2$



Augmenting paths (cont.)



Max-flow, min-cut theorem

Theorem. The following are equivalent:

1. $|f| = c(S, T)$ for some cut (S, T) . ← min-cut
2. f is a maximum flow.
3. f admits no augmenting paths.


Proof.

(1) \Rightarrow (2): Since $|f| \leq c(S, T)$ for any cut (S, T) , the assumption that $|f| = c(S, T)$ implies that f is a maximum flow.

(2) \Rightarrow (3): If there was an augmenting path, the flow value could be increased, contradicting the maximality of f .

Proof (continued)

(3) \Rightarrow (1): Define $S = \{v \in V : \text{there exists an augmenting path in } G_f \text{ from } s \text{ to } v\}$, and let $T = V \setminus S$. Since f admits no augmenting paths, there is no path from s to t in G_f . Hence, $s \in S$ and $t \in T$, and thus (S, T) is a cut. Consider any vertices $u \in S$ and $v \in T$.



We must have $c_f(u, v) = 0$, since if $c_f(u, v) > 0$, then $v \in S$, not $v \in T$ as assumed. Thus, $f(u, v) = c(u, v)$ if $(u, v) \in E$ since $c_f(u, v) = c(u, v) - f(u, v)$. And otherwise $f(u, v) = 0$. Summing over all $u \in S$ and $v \in T$ yields $f(S, T) = c(S, T)$, and since $|f| = f(S, T)$, the theorem follows. □

Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Can be slow:



Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Can be slow:



Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Can be slow:



Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Can be slow:



Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

 augment f by $c_f(p)$

Can be slow:



Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Can be slow:



Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Can be slow:



Ford-Fulkerson max-flow algorithm


Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Can be slow:



2 billion iterations on a graph with 4 vertices!

Ford-Fulkerson max-flow algorithm

Algorithm:

$f[u, v] \leftarrow 0$ for all $(u, v) \in E$

while an augmenting path p in G wrt f exists:

augment f by $c_f(p)$

Runtime:

- Let $|f^*|$ be the value of a maximum flow, and assume it is an integral value.
 - The initialization takes $O(|E|)$ time
 - There are at most $|f^*|$ iterations of the loop
 - Find an augmenting path with DFS in $O(|V| + |E|)$ time
 - Each augmentation takes $O(|V|)$ time
- $\Rightarrow O(|E| \cdot |f^*|)$ time in total

Edmonds-Karp algorithm

Edmonds and Karp noticed that many people's implementations of Ford-Fulkerson augment along a *breadth-first augmenting path*: a shortest path in G_f from s to t where each edge with positive capacity has weight 1. These implementations would always run relatively fast.

Since a breadth-first augmenting path can be found in $O(|V|+|E|)$ time, their analysis, which provided the first polynomial-time bound on maximum flow, focuses on bounding the number of flow augmentations.

(In independent work, Dinic also gave polynomial-time bounds.)

Running time of Edmonds-Karp

- One can show that the number of flow augmentations (i.e., the number of iterations of the while loop) is $O(|V|/|E|)$.

- Breadth-first search runs in $O(|V|+|E|)$ time

- All other bookkeeping is $O(|V|)$ per augmentation.

⇒ The Edmonds-Karp maximum-flow algorithm runs in $O(|V|/|E|^2)$ time.

Monotonicity lemma

Lemma. Let $\delta(v) = \delta_f(s, v)$ be the breadth-first distance from s to v in G_f . During the Edmonds-Karp algorithm, $\delta(v)$ increases monotonically.

Proof. Suppose that f is a flow on G , and augmentation produces a new flow f' . Let $\delta'(v) = \delta_{f'}(s, v)$. We'll show that $\delta'(v) \geq \delta(v)$ by induction on $\delta(v)$. For the base case, $\delta'(s) = \delta(s) = 0$.

For the inductive case, consider a breadth-first path $s \rightarrow \dots \rightarrow u \rightarrow v$ in $G_{f'}$. We must have $\delta'(v) = \delta'(u) + 1$, since subpaths of shortest paths are shortest paths. Certainly, $(u, v) \in E_{f'}$, and now consider two cases depending on whether $(u, v) \in E_f$.

Case 1

Case: $(u, v) \in E_f$.

We have

$$\begin{aligned}\delta(v) &\leq \delta(u) + 1 && \text{(triangle inequality)} \\ &\leq \delta'(u) + 1 && \text{(induction)} \\ &= \delta'(v) && \text{(breadth-first path),}\end{aligned}$$

and thus monotonicity of $\delta(v)$ is established.

Case 2

Case: $(u, v) \notin E_f$.

Since $(u, v) \in E_{f'}$, the augmenting path p that produced f' from f must have included (v, u) . Moreover, p is a breadth-first path in G_f :

$$p = s \rightarrow \cdots \rightarrow v \rightarrow u \rightarrow \cdots \rightarrow t.$$

Thus, we have

$$\begin{aligned} \delta(v) &= \delta(u) - 1 && \text{(breadth-first path)} \\ &\leq \delta'(u) - 1 && \text{(induction)} \\ &= \delta'(v) - 2 && \text{(breadth-first path)} \\ &< \delta'(v), \end{aligned}$$


thereby establishing monotonicity for this case, too. □

Counting flow augmentations

Theorem. The number of flow augmentations in the Edmonds-Karp algorithm (Ford-Fulkerson with breadth-first augmenting paths) is $O(|V|/|E|)$.

Proof. Let p be an augmenting path, and suppose that we have $c_f(u, v) = c_f(p)$ for edge $(u, v) \in p$. Then, we say that (u, v) is **critical**, and it disappears from the residual graph after flow augmentation.

Example:




Counting flow augmentations

Theorem. The number of flow augmentations in the Edmonds-Karp algorithm (Ford-Fulkerson with breadth-first augmenting paths) is $O(|V||E|)$.

Proof. Let p be an augmenting path, and suppose that we have $c_f(u, v) = c_f(p)$ for edge $(u, v) \in p$. Then, we say that (u, v) is **critical**, and it disappears from the residual graph after flow augmentation.

Example:




Counting flow augmentations (continued)

The first time an edge (u, v) is critical, we have $\delta(v) = \delta(u) + 1$, since p is a breadth-first path. We must wait until (v, u) is on an augmenting path before (u, v) can be critical again. Let δ' be the distance function when (v, u) is on an augmenting path. Then, we have

$$\begin{aligned}\delta'(u) &= \delta'(v) + 1 && \text{(breadth-first path)} \\ &\geq \delta(v) + 1 && \text{(monotonicity)} \\ &= \delta(u) + 2 && \text{(breadth-first path).}\end{aligned}$$

Example:




Counting flow augmentations (continued)

The first time an edge (u, v) is critical, we have $\delta(v) = \delta(u) + 1$, since p is a breadth-first path. We must wait until (v, u) is on an augmenting path before (u, v) can be critical again. Let δ' be the distance function when (v, u) is on an augmenting path. Then, we have

$$\begin{aligned}\delta'(u) &= \delta'(v) + 1 && \text{(breadth-first path)} \\ &\geq \delta(v) + 1 && \text{(monotonicity)} \\ &= \delta(u) + 2 && \text{(breadth-first path).}\end{aligned}$$

Example:




Counting flow augmentations (continued)

The first time an edge (u, v) is critical, we have $\delta(v) = \delta(u) + 1$, since p is a breadth-first path. We must wait until (v, u) is on an augmenting path before (u, v) can be critical again. Let δ' be the distance function when (v, u) is on an augmenting path. Then, we have

$$\begin{aligned}\delta'(u) &= \delta'(v) + 1 && \text{(breadth-first path)} \\ &\geq \delta(v) + 1 && \text{(monotonicity)} \\ &= \delta(u) + 2 && \text{(breadth-first path).}\end{aligned}$$

Example:




Counting flow augmentations (continued)

The first time an edge (u, v) is critical, we have $\delta(v) = \delta(u) + 1$, since p is a breadth-first path. We must wait until (v, u) is on an augmenting path before (u, v) can be critical again. Let δ' be the distance function when (v, u) is on an augmenting path. Then, we have

$$\begin{aligned}\delta'(u) &= \delta'(v) + 1 && \text{(breadth-first path)} \\ &\geq \delta(v) + 1 && \text{(monotonicity)} \\ &= \delta(u) + 2 && \text{(breadth-first path).}\end{aligned}$$

Example:




Counting flow augmentations (continued)

The first time an edge (u, v) is critical, we have $\delta(v) = \delta(u) + 1$, since p is a breadth-first path. We must wait until (v, u) is on an augmenting path before (u, v) can be critical again. Let δ' be the distance function when (v, u) is on an augmenting path. Then, we have

$$\begin{aligned}\delta'(u) &= \delta'(v) + 1 && \text{(breadth-first path)} \\ &\geq \delta(v) + 1 && \text{(monotonicity)} \\ &= \delta(u) + 2 && \text{(breadth-first path).}\end{aligned}$$

Example:




Counting flow augmentations (continued)


The first time an edge (u, v) is critical, we have $\delta(v) = \delta(u) + 1$, since p is a breadth-first path. We must wait until (v, u) is on an augmenting path before (u, v) can be critical again. Let δ' be the distance function when (v, u) is on an augmenting path. Then, we have

$$\begin{aligned}\delta'(u) &= \delta'(v) + 1 && \text{(breadth-first path)} \\ &\geq \delta(v) + 1 && \text{(monotonicity)} \\ &= \delta(u) + 2 && \text{(breadth-first path).}\end{aligned}$$

Example:



Running time of Edmonds-Karp

Distances start out nonnegative, never decrease, and are at most $|V| - 1$ until the vertex becomes unreachable. Thus, (u, v) occurs as a critical edge $O(|V|)$ times, because $\delta(v)$ increases by at least 2 between occurrences. Since the residual graph contains $O(|E|)$ edges, the number of flow augmentations is $O(|V|/|E|)$. 

Corollary. The Edmonds-Karp maximum-flow algorithm runs in $O(|V|/|E|^2)$ time.

Proof. Breadth-first search runs in $O(|E|)$ time, and all other bookkeeping is $O(|V|)$ per augmentation. 

Best to date

- The asymptotically fastest algorithm to date for maximum flow, due to King, Rao, and Tarjan, runs in $O(|V||E| \log_{|E|/(|V| \log |V|)} |V|)$ time.
- If we allow running times as a function of edge weights, the fastest algorithm for maximum flow, due to Goldberg and Rao, runs in time $O(\min\{|V|^{2/3}, |E|^{1/2}\} \cdot |E| \log(|V|^2/|E| + 2) \cdot \log C)$, where C is the maximum capacity of any edge in the graph.