

CSC 373 Week 4 Notes

Introduction to Network Flow:

- Input:

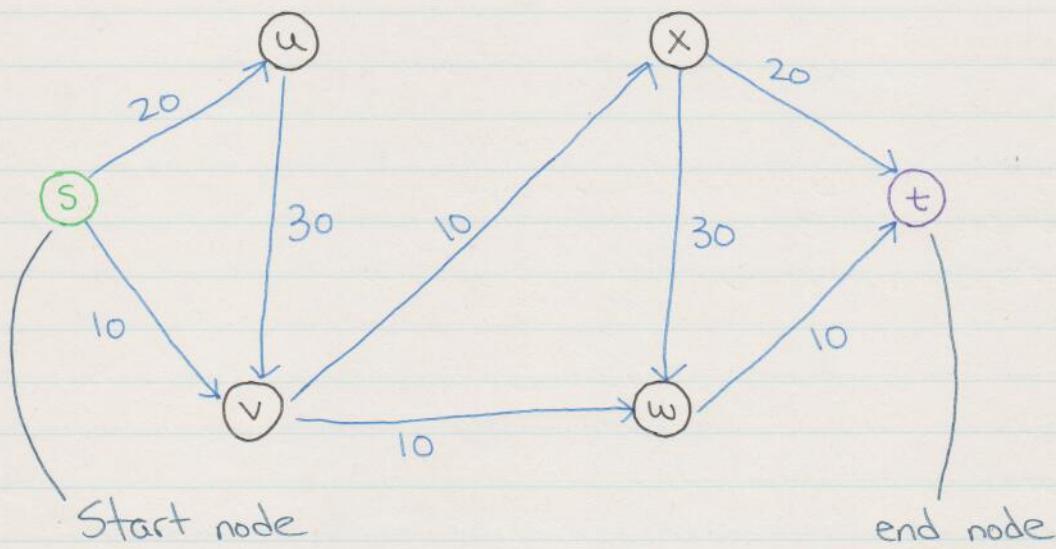
1. A directed graph $G = (V, E)$.
2. Edge capacities that are non-negative.
3. A source node, s , and a target node, t .

- Output: Maximum "flow" from s to t .

- Assumptions:

1. No edges enter s
2. No edges leave t
3. Each edge capacity is non-negative (≥ 0)

- E.g.



Note: Treat the edge capacities like stuff/weight, can go on the edge at a single time.

the max amount of

- The **flow** is the amount of stuff/weight carried over an edge.

- The flow is always between 0 and the capacity of the edge, inclusive.

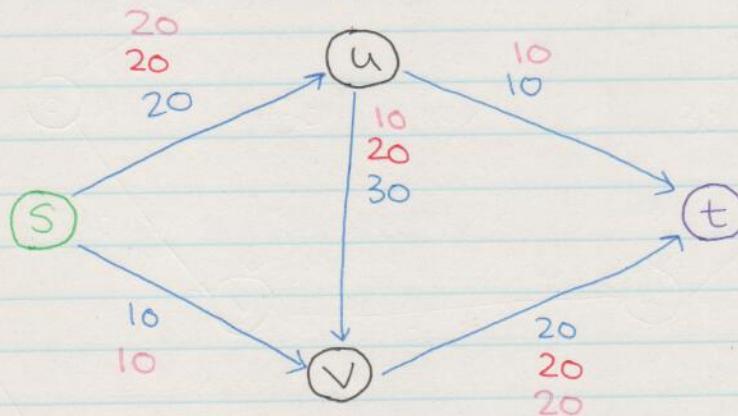
I.e. $0 \leq f(e) \leq c(e), \forall e \in E$

↑ ↑
flow edge capacity

- For all nodes except for s and t , the sum of the flow entering that node equals the sum of the flow leaving it.

I.e. $\forall v \in V \setminus \{s, t\}, \sum_{e \text{ entering } v} f(e) = \sum_{e \text{ leaving } v} f(e)$

- Consider the graph below:



I will use red to denote the numbers for the first example and pink for the second example.

Blue numbers mean the max capacity of that edge.

In the first example, we put 20 on the edge (S, u) , then 20 on the edge (u, v) and 20 on the edge (v, t) . However this is not the most optimal solution.

Note: The most optimal soln is the min between the total capacity of all the edges flowing from S or the total capacity of all the edges flowing into t .

In my second example, we do get an optimal solution. This is because the total amount flowing into t is 30, which is the same as both the total capacity of edges flowing out of S or the total capacity of edges going into t .

- We can use a **residual graph** to reverse bad decisions.
- In the example above, the pink numbers (2^{nd} soln) was the residual graph.

Residual Graph:

- Let $c(e)$ denote the capacity for edge e .
- Let $f(e)$ denote the flow of edge e .
- Suppose the current flow is f .
Let G_f be the residual graph of f .
 G_f has the same vertices as f
For each edge $e = (u, v)$ in f , G_f has at most 2 edges:

1. Forward Edge

- $e = (u, v)$
- Capacity = $c(e) - f(e)$
- How much more flow we can send on this path

2. Reverse Edge

$$- e^{\text{rev}} = (v, u)$$

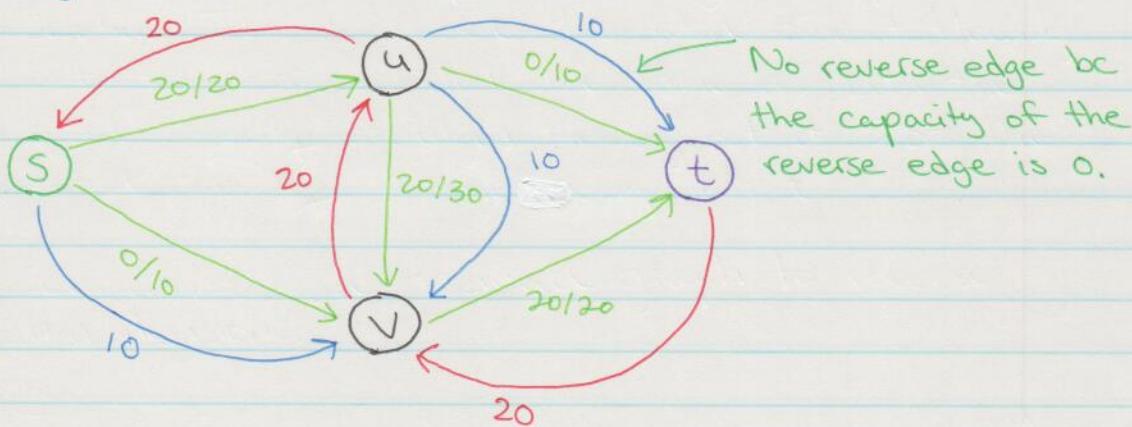
$$- \text{Capacity} = f(e)$$

- This is the max reverse flow. Since the original flow along edge e is $f(e)$, the max we can send back is $f(e)$.

I.e. Say that the flow from S to A is 3. Then, the reverse flow from A to S is 3.

Note: We only add edges where the capacity is greater than 0.

- E.g.



The red lines are the reverse edges and the blue lines are the forward edges.

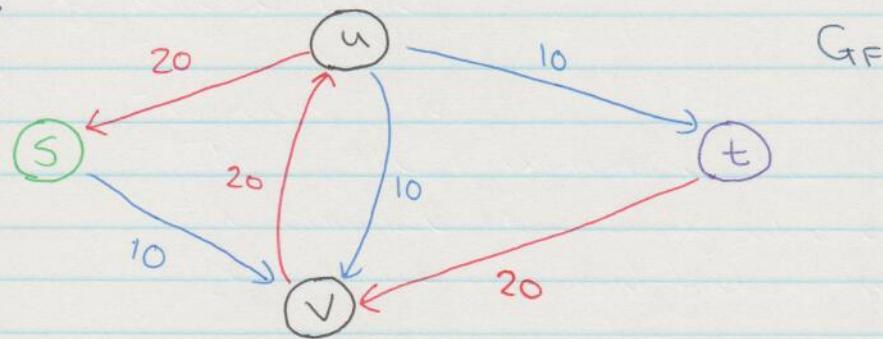
Notice that some edges only have a forward edge or a reverse edge. This is bc the other edge has a capacity of ≤ 0 , so we don't add it.

- Let P be an $S-T$ path in G_F .
- Let $\text{bottleneck}(P, F)$ be the smallest capacity across all edges in P .
- We can augment F by sending $\text{bottleneck}(P, F)$ units of flow along P .

Note: Sending x units of flow along P means:

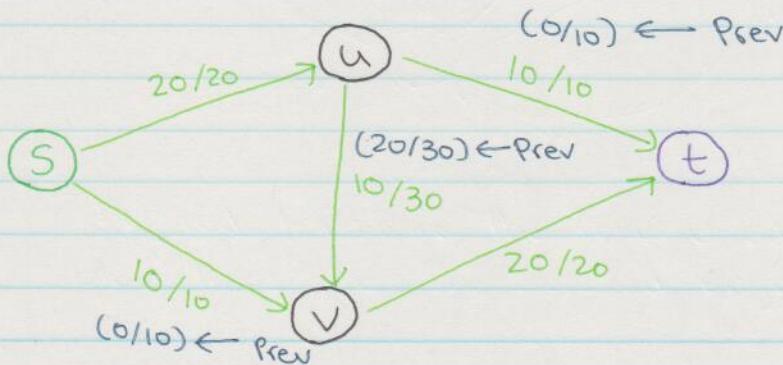
1. For each forward edge in P , inc the flow by x .
2. For each reverse edge in P , dec the flow by x .

- E.g.



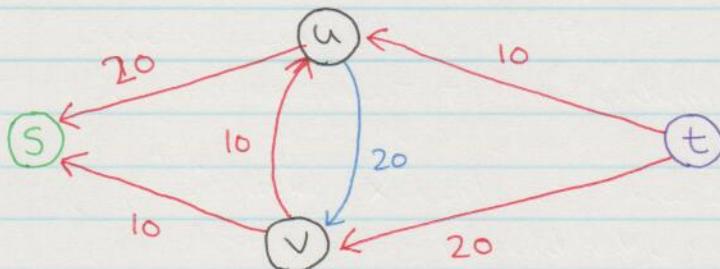
We see that (S, V) and (U, T) are bottlenecks as they have the smallest capacity (10). Hence, we will increase the flow by 10 along all forward edges and decrease the flow by 10 along all reverse edges.

New flow f :



We added/increased the flow from (S, V) and (U, T) by 10 and decreased the flow from (U, V) by 10.

New Residual Graph G_F :



Notice there's no $S-T$ path in G_F anymore.

- Now, we'll prove that the new flow is valid.

1. Capacity Constraints:

- If we increase flow on e , we can do so by at most the capacity of the forward edge e in G_F which is $c(e) - f(e)$.

$$f(e) + (c(e) - f(e)) = c(e) \leftarrow \text{Max new flow}$$

- If we decrease flow on e , we can do so by at most the capacity of the reverse edge e^{rev} in G_F , which is $f(e)$. Hence, the new flow is at least $f(e) - f(e)$ or 0.

2. Flow Conservation:

- Each node except for s and t has 2 incident edges.
- If they are both forward/both reverse, that means one is incoming, one is outgoing. The flow is increased on both or decreased on both. The net flow is 0.
- If one is forward and the other reverse then we get both incoming or both outgoing. Flow is increased on one and decreased on the other. The net flow is 0.

Ford - Fulkerson Algorithm:

Max Flow (G):

// Initialize

Set $f(e) = 0$ for all e in G

// while there's an $s-t$ path in GF :

While $P = \text{FindPath}(s, t, \text{Residual}(G, f)) \neq \text{None}$:

$f = \text{augment}(f, P)$

Update Residual (G, f)

return f

- Running time: $O((m+n) \cdot C)$

Has $2m$ edges

Has n vertices

The max flow / max num of augmentations is at most

$$C = \sum_{e \text{ leaving } s} c(e)$$

For path P in GF , because $\text{bottleneck}(P, f) \geq 1$, each augmentation increases flow by at least 1. That's why the max flow is at most $\sum_{e \text{ leaving } s} c(e)$.

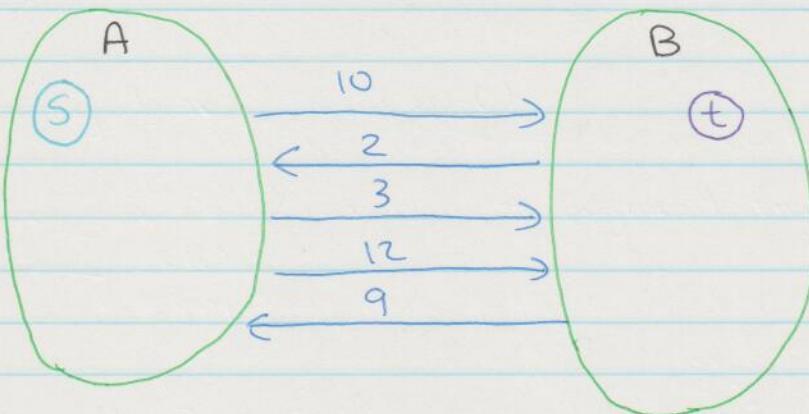
The time it takes to perform an augmentation is $O(mn)$. This is bc GF has n vertices and at most $2m$ edges. So finding P , computing $\text{bottleneck}(P, f)$ and updating G_f takes $O(mn)$ time.

Proof of Optimality:

- Let (A, B) be an $s-t$ cut. This means we partition the nodes into 2 groups, A and B . Furthermore, $s \in A$ and $t \in B$.
- $A \cup B = V$
- $A \cap B = \emptyset$

- Let $\text{cap}(A, B)$ be the sum of the capacities of the edges leaving A .

- E.g.



$$\text{Cap}(A, B) = 10 + 3 + 12 \\ = 25$$

- Thm: For any flow f and any $s-t$ cut (A, B) ,

$$v(f) = f^{\text{out}}(A) - f^{\text{in}}(A)$$

Note: $v(f) = f^{\text{out}}(s) = f^{\text{in}}(t)$

Proof:

Only edges that connect nodes in both partitions can deliver the flow to t .

Edges that connect nodes in the same partition cancel each other out.

I.e. Consider nodes x and y s.t. they're in the same partition. Suppose the flow from x to y is e . That means y has to "release" e amount of flow. If it's to other nodes in the same partition, it doesn't directly lead to t , so it cancels out.

Assume only x goes into y .

So, we're only interested in the edges that span both partitions.

Therefore, we need to sum the capacity of the edges that go from A to B and subtract the capacity of the edges that go from B to A.

We need to do the subtraction because the edges that go from B to A "return" some of the flow.

- Thm: For any flow f and any s-t cut (A, B) , $v(f) \leq \text{cap}(A, B)$

Proof:

$$\begin{aligned} v(f) &= f^{\text{out}}(A) - f^{\text{in}}(B) \\ &\leq f^{\text{out}}(A) \\ &= \sum_{e \text{ leaving } A} f(e) \\ &\leq \sum_{e \text{ leaving } A} c(e) \\ &= \text{cap}(A, B) \end{aligned}$$

Hence, $\max_f v(f) \leq \min_{(A, B)} \text{cap}(A, B)$

I.e. The max value of any flow \leq min capacity of any s-t cut

Implications:

1. Max flow = Min cut
2. Ford-Fulkerson generates max flow

- Thm: Ford-Fulkerson returns / finds the max flow

Proof:

f = flow returned by F-F

A^* = nodes reachable from S in G_F

B^* = $V \setminus A^*$

Note: We look at the residual graph G_F but define the cut in G

Note: The F-F algo terminates when there are no paths from S to t . (See pg 7)

Claim: (A^*, B^*) is a valid cut.

$S \in A^*$ by definition

$t \in B^*$ because there's no path from S to t in G_F . Hence, $t \notin A^*$

Consider the following:

1. Each edge (u, v) going out of A^* must be ^{in G} $\overset{\text{in } G}{\underset{\text{Saturated}}{(c(e) = f(e))}}$. Otherwise, G_F would have its own forward edge (u, v) and $v \in A^*$.

Recall: If edge (v, u) in G has a flow less than its capacity, then in G_F , there will be a forward edge with capacity $c(e) - f(e)$.

2. Each edge $\overset{(v, u)}{\longrightarrow}$ coming into A^* in G must have 0 flow. Otherwise G_F would have its reverse edge (u, v) and then $v \in A^*$.

E.g. Suppose $A \in A^*$ and $B \in B^*$

$A \xrightarrow{5/10} B$

G

, then in G_F :

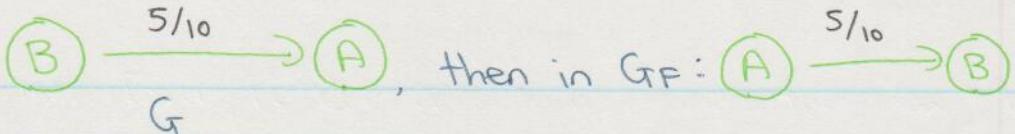
Forward Edge

$A \xrightarrow{5/10} B$

This means we can reach B from S , which means $B \in A^*$, contradicting our

Reverse Edge

11



This means that we can reach B from S, contradicting our assumption.

Based on 1 and 2, we conclude:

$$f^{\text{out}}(A^*) = \text{Cap}(A^*, B^*)$$

$$f^{\text{in}}(A^*) = 0$$

$$\begin{aligned} \therefore v(f) &= f^{\text{out}}(A^*) - f^{\text{in}}(A^*) \\ &= \text{Cap}(A^*, B^*) \end{aligned}$$

Edmonds - Karp Algo:

- A modification on Ford - Fulkerson algo.
- Instead of finding any s-t path in GF, like the Ford - Fulkerson algo does, it finds the shortest s-t path.
- Algo:

MaxFlow(G):

// Initialize

Set $f(e) = 0 \quad \forall e \in G$

// Find the shortest s-t path in GF and augment

while $P = \text{BFS}(S, t, \text{Residual}(G, f)) \neq \text{None}$:

$f = \text{augment}(f, P)$

$\text{UpdateResidual}(G, f)$

return f

- Let $d(v)$ be the shortest dist to node v from s in the residual graph G_f .
- Lemma 1: During the execution of the algo, $d(v)$ does not decrease for any v .

Proof:

Suppose the augmentation $f \rightarrow f'$ dec $d(v)$ for some v .

Choose the v with the smallest $d(v)$ in G_f' . — G_f is the original graph.
Say $d(v) = k$ in G_f' and so, $d(v) \geq k+1$ in G_f .

Note: G_f is the residual graph for f and
 G_f' is the residual graph for f' .

G_f' is its augmented graph.

Recall that f' , the augmented graph of f , is created by finding the bottleneck in the $s-t$ path in G_f and then sending that many units along that path. (See pg 4)

Let node w be the node just before v on the shortest path $s \rightarrow v$ in G_f' .

Since $d(v) = k$ in G_f' , $d(w) = k-1$ in G_f' .

Furthermore, since $d(w)$ didn't decrease, $d(w) \leq k-1$ in G_f .

Note: Since we start with G_f and get G_f' later, by the lemma, $d(w) \geq d(v)$. However, because of our assumption, $d(w) < d(v)$.

Here's a table to summarize the values:

	$d(w)$	$d(v)$	
G_f	$\leq k-1$	$\geq k+1$	Before
G_f'	$k-1$	k	After

in G_f , (u, v) must be missing.

This is because in G_f , $d(u) \leq k-1$ and $d(v) \geq k+1$.

If there is an edge (u, v) , then

$$d(v) = d(u) + 1, \text{ but it doesn't.}$$

Since (u, v) doesn't exist in G_f , we must have added it by selecting (v, u) in the augmented path P . However, P is a shortest path so it cannot have an edge (v, u) with $d(v) > d(u)$.

Note: Consider $d(t)$. We know that it will always increase as the execution continues and that eventually, $d(t) = \infty$ because there is no path from s to t .

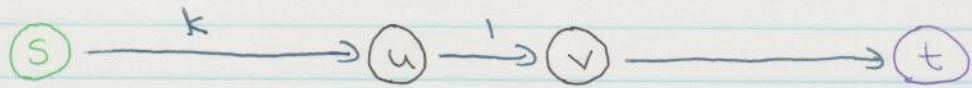
- Let edge (u, v) be **critical** in an augmentation step if it's part of the augmented path P and its capacity is equal to bottleneck(P, f) and the augmentation step removes e and adds e^{rev} , if it's missing.
- Lemma 2: Between any 2 steps in which (u, v) is critical, $d(u)$ increases by at least 2.

Proof:

If (u, v) is critical, then that means its capacity is the lowest along path P . Furthermore, during the augmentation step, we send that many units of flow down path P and therefore, (u, v) is saturated. Hence, the residual graph will remove (u, v) and replace it with (v, u) .

Suppose that $d(u) = k$ in G_f . Since (u, v) is a shortest path, $d(v) = k+1$ in G_f .

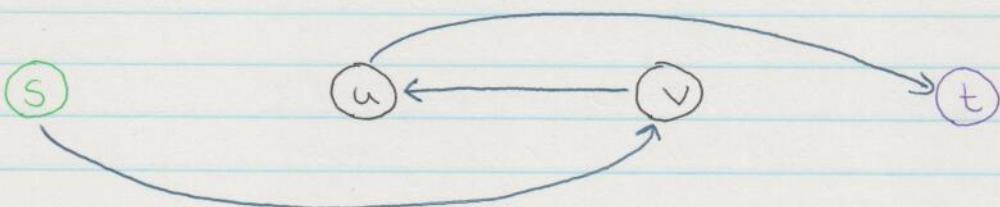
I.e.



For (u, v) to be critical a second time, it must have been added back at some point.

Since the residual graph removed (u, v) and added (v, u) in its place, to get (u, v) back, the augmented path must have selected (v, u) .

I.e.



This path is still a shortest path.

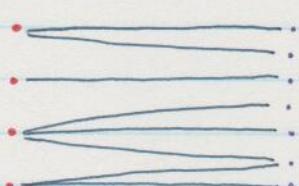
Hence, in G_f' , $d(u) = d(v)+1 \leftarrow$ The +1 comes from
 $\geq (k+1)+1 \quad (v, u).$
 $\geq k+2$

- Proof of Running Time:
- Each $d(u)$ can go from 0 to n . (Lemma 1)
- So, each edge (u, v) can be critical at most $\frac{n}{2}$ times. (Lemma 2)
- So, there can be at most $\frac{mn}{2}$ augmentation steps.
- Since each augmentation takes $O(m)$ time, there are $O(m^2n)$ operations in total.

- Applications of Network Flow:
- Integral Theorem: If edge capacities are integers, then the max-flow computed by Ford-Fulkerson and its variants are also integral, meaning that the flow on each edge is an integer.
- Bipartite Matching:
 - Problem: Given a bipartite graph $G = (U \cup V, E)$, find a max cardinality matching.

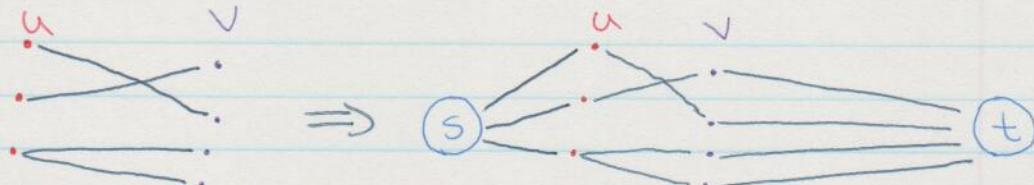
Note: A bipartite graph is a set of vertices decomposed into 2 disjoint sets s.t. no 2 graph vertices within the same set are adjacent.

E.g.



Note: A matching is a set of pairwise non-adjacent edges s.t. no 2 edges share common vertices.

- Consider this:



We can "reduce" this problem to max flow by adding a starting node that connects to each node in U and an end node that connects to each node from V .

- Furthermore, let the capacity of each edge be 1. This way, we can "mimic" choosing an edge as 1 and not choosing an edge as 0.
- Note: There is a 1-1 correspondence btwn matchings of size k in the original graph and flows with value k in the corresponding flow networks.

Proof:

Note: Similar to proving an iff, we must prove both directions.

Matching \Rightarrow Flow:

Take a matching $M, \{(u_1, v_1), \dots, (u_k, v_k)\}$ of size k .

Construct the corresponding flow f_M where:

- Edges $s \rightarrow u_i \rightarrow v_i \rightarrow t$ have flow 1, $\forall i \in 1 \dots k$
- All other edges have flow 0.

Hence, the total flow is k .

Flow \Rightarrow Matching:

Take any flow f with value k .

The corresponding unique matching M_f equals to the set of edges from U to V with a flow of 1.

Since k units of flow come from S , it must go into k vertices in U and that will go to k distinct vertices in V and finally, all k flows will go to t .

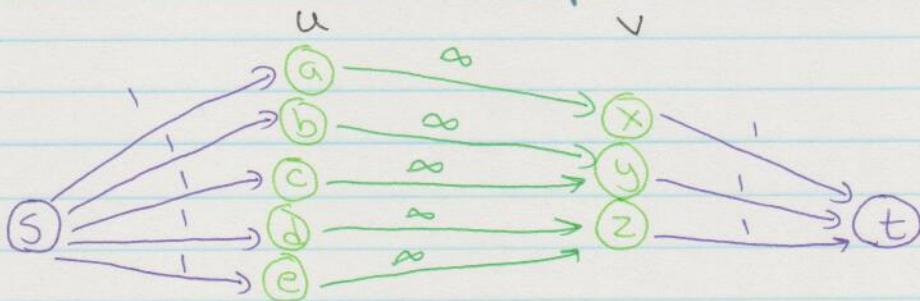
- Note: A perfect matching occurs when we have a flow with value n where $n = |U| = |V|$

A perfect matching is a matching that covers every vertex of the graph.

- Hall's Marriage Thm:
- We know that a bipartite graph has a perfect matching when the corresponding flow network has a value of n .
- However, can we interpret this condition in terms of edges of the original bipartite graph? \rightarrow Yes we can
- For $S \subseteq U$, let $N(S) \subseteq V$ be the set of nodes that is in V and is adjacent to some node in S .
- **Observation:** If G has a perfect matching, then $|N(S)| \geq |S|$ for each $S \subseteq U$. This is bc each node in S must be matched to a distinct node in $N(S)$.
- We'll consider a slightly different network flow, which is still equivalent to bipartite matching. Now, all $U \rightarrow V$ edges have ∞ capacity. All $S \rightarrow V$ and $V \rightarrow t$ edges still have a capacity of 1.
- **Hall's Thm:** G has a perfect matching iff $|N(S)| \geq |S|$ for each $S \subseteq U$.

Proof (Reverse Direction):

First, I'll draw an example of our network.



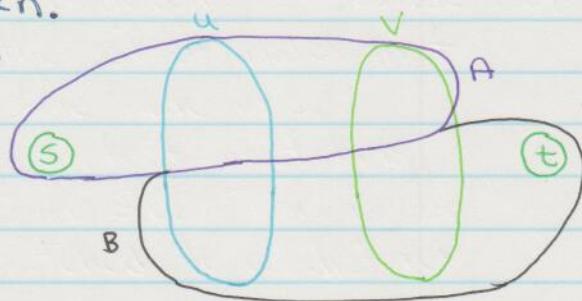
Note: We want to prove "if $|N(S)| \geq |S|$ for each $S \subseteq U$, then G has a perfect matching."

We will prove the contrapositive: IF G does not have a perfect matching then for some $S \subseteq V$, $|N(S)| < |S|$.

Suppose that G doesn't have a perfect matching.
Then $\text{max-flow} = \text{min-cut} < n$.

Let (A, B) be a min-cut \rightarrow

Note: There can't be any edges going from $V \cap A$ to $V \cap B$. This is because, since the capacity of all (u, v) edges are infinity, if such an edge exists, then $\text{cap}(A, B) = \infty$, but we know that the maxflow $< n$.



$$\begin{aligned}\text{Cap}(A, B) &= |V \cap B| + |V \cap A| \\ &< n \\ &= |V|\end{aligned}$$

$$\left. \begin{aligned}|V \cap B| + |V \cap A| &< |V| \\ |V \cap A| &< |V| - |V \cap B| \\ &= |V \cap B|\end{aligned}\right\} |V \cap A| < |V \cap B|$$

Furthermore, since we know that no edges between $(V \cap A)$ and $(V \cap B)$ exist (see above), we know that:
 $N(V \cap A) \subseteq V \cap B$ and that $|N(V \cap A)| \leq |V \cap B|$.

$$\begin{aligned}\text{Hence, } |N(V \cap A)| &\leq |V \cap B| \\ &< |V \cap A|\end{aligned}$$

- Edge - Disjoint Paths:
- Problem: Given a directed graph $G = (V, E)$, 2 nodes s and t , find the max number of edge-disjoint $s \rightarrow t$ paths.
- Note:** 2 $s \rightarrow t$ paths P and P' are edge-disjoint if they don't share an edge.
- Soln: Assign 1 as the capacity for all edges
- Thm: There is a 1-1 correspondence btwn sets of k edge-disjoint $s-t$ paths and integral flows of value k .

Proof:

1. Path \rightarrow Flow:

Let $\{P_1, \dots, P_k\}$ be k edge-disjoint $s-t$ paths.

Define flow f where $f(e) = 1$ whenever $e \in P_i$ for some i and 0 otherwise.

Since paths are edge-disjoint, flow conservation and capacity constraints are fulfilled/satisfied.

There is a unique integral flow of value k .

2. Flow \rightarrow Path:

Let f be an integral flow of value k .

That means that there are k outgoing edges, all of which have unit flow.

Pick 1 edge (s, u_i) :

- By flow conservation, u_i must have an outgoing flow of value 1, that we haven't used yet.

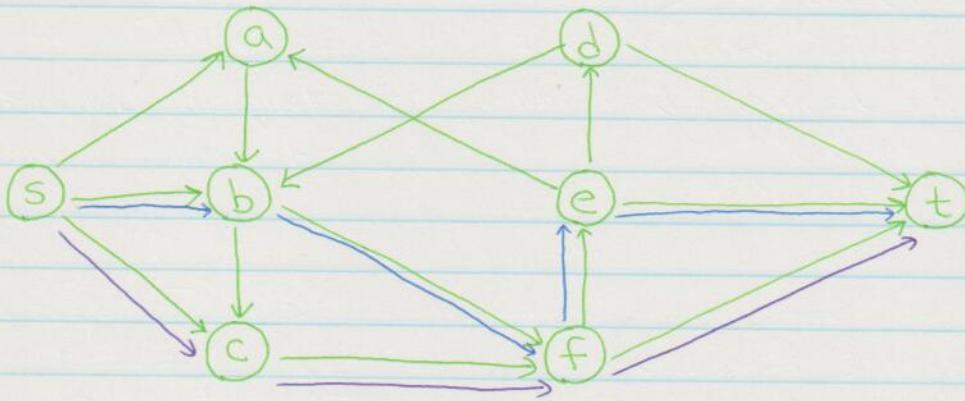
- Pick such an edge and build a path until we reach t .

Repeat for all other edges.

$k-1$

- Menger's Thm: In any graph, regardless if it's directed or undirected, the max number of edge-disjoint $s-t$ paths equal the min number of edges whose removal disconnects s and t .

E.g. Consider the graph below



The blue and purple lines each represent a $s-t$ path and they're edge disjoint.

If we remove the edges (b, f) and (c, f) , we disconnect s and t .

Notice that the number of edge-disjoint $s-t$ paths equals to the min number of edges removed to disconnect s and t .

- Multiple Source / Sinks:
- Problem: Given a directed graph $G = (V, E)$ with edge capacities $c: E \rightarrow N$, sources s_1, \dots, s_k and sinks t_1, \dots, t_l , find the max total flow from sources to sinks.
- Soln:
 - Add a new source s and connect s with each node s_1, \dots, s_k s.t. the edges have infinite capacity.
 - Add a new sink t s.t. t is connected with each node t_1, \dots, t_l and the edges all have infinite capacity.
 - Find the max flow from s to t .

- Circulation:

- Input:
 - Directed graph $G = (V, E)$
 - Edge capacities $c: E \rightarrow N$
 - Node demands $d: V \rightarrow Z$

- Output:
 - Some circulation $f: E \rightarrow N$ satisfying:
 1. For each $e \in E$: $0 \leq f(e) \leq c(e)$
 2. For each $V \in V$: $\sum_{e \text{ entering } v} f(v) - \sum_{e \text{ leaving } v} f(v) = d(v)$

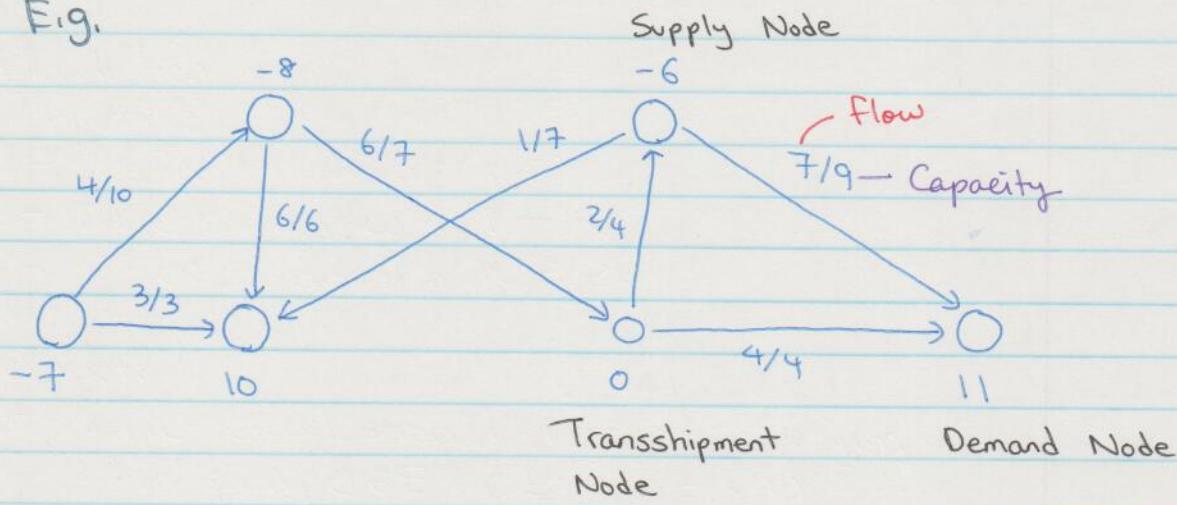
Note: We need $\sum_{v: d(v) > 0} d(v) = -\sum_{v: d(v) < 0} d(v)$

$d(v)$ equals to the sum
of the incoming flows -
the sum of the outgoing
flows.

- Demand at v , $d(v)$, is the amount of flow you need to take out at node v .

- If $d(v) > 0$:
 - There should be more incoming flow than outgoing flow.
 - This is a demand node.
- If $d(v) < 0$:
 - There should be more outgoing flow than incoming flow.
 - This is a supply node.
- If $d(v) = 0$:
 - Equal outgoing and incoming flows.
 - Transshipment node

- E.g.



Note:

1. Each $d(v)$ equals to the sum of the incoming flow(s) minus the sum of the outgoing flow(s).
2. The sum of $d(v)$'s > 0 equal the sum of the $d(v)$'s < 0 times -1 .

$$\sum_{d(v)>0} d(v) = 10 + 11 = 21 \quad -\sum_{d(v)<0} d(v) = -(-7 - 8 - 6) = -(-21) = 21$$

- Soln:

- Create a start node S and a sink node T .
- Connect S with each supply node v , such that the edge (S, v) has capacity $-d(v)$.
- Connect T with each demand node v , s.t. the edge (v, T) has capacity $d(v)$.
- Claim: G has a circulation iff G' has a max flow of value

$$\sum_{v: d(v) > 0} d(v) = - \sum_{v: d(v) < 0} d(v)$$

- Survey Design:

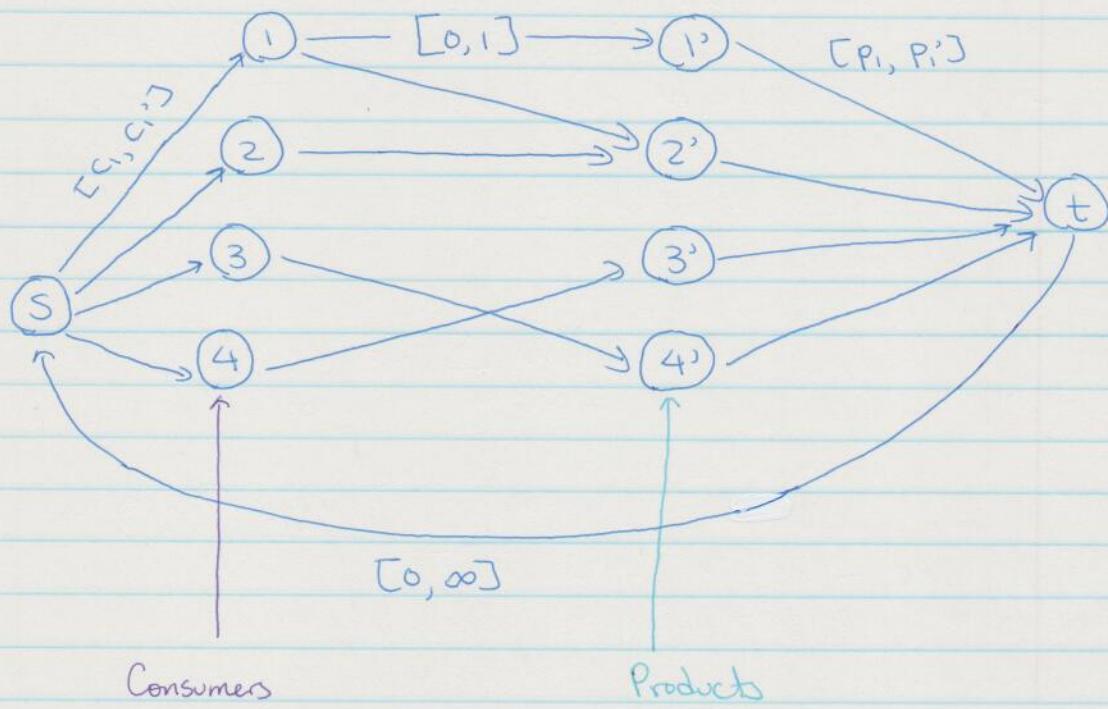
- Problem: We want to design a survey about m products. We have 1 question per product and we need to ask product j 's question to btwn p_j and p'_j consumers.

There are a total of n consumers and ^{each} consumer i owns a subset of products O_i . We can only ask consumer i questions about these products and we want to ask each consumer i between c_i and c'_i questions.

- Soln:

- Create a network with special nodes s and t .
- Create an edge from s to each consumer i with flow $\in [c_i, c'_i]$.
- Create an edge from each consumer i to each product $j \in O_i$ with flow $\in [0, 1]$.
- Create an edge from each product j to t with flow $\in [p_j, p'_j]$.
- Create an edge from t to s with flow $\in [0, \infty]$. This is to ensure that the flow coming into t equals the flow leaving s .

- Fig.



- Profit Maximization:
- Problem: Given n tasks, each of which generates some profit (and the profit could be 0 or negative) and a set E of precedence relations, we want to find a subset of tasks, S , that is subject to the precedence constraints s.t. $\text{profit}(S)$ is maximized.

Note: If $(i, j) \in E$, that means we have to perform job/task i before we can perform task j . This is the precedence relation/predence constraint.

- Soln:
- We can represent this problem as a network flow problem.

We will represent the input as a graph:

1. Tasks are nodes
2. Profits are node weights
3. Precedence constraints are edges.

Note: The capacity of all edges s.t. the edges are not connected to s or t will be infinite. This is because when we take the min-cut, we don't want node i to be on one side and node j to be on the other side if we have the precedence constraint (i, j) .

Furthermore, we will have a start node, s , and a sink node, t , and connect s and t with the other nodes like this:

1. If P_i , the profit of node i , is positive, we connect node i to s and the capacity of $(s, i) = P_i$.
2. If P_i is negative, then we connect node i to t and the capacity of edge (i, t) will be $-P_i$.

- If $\text{cap}(A, B)$ is finite, then $A \setminus \{s\}$ is a valid soln.
- $\text{Min Cap}(A, B) \geq \text{max profit}(A \setminus \{s\})$
- $\text{Cap}(A, B) = \sum_{i \in \text{positive nodes}} P_i - \text{profit}(A \setminus \{s\})$.