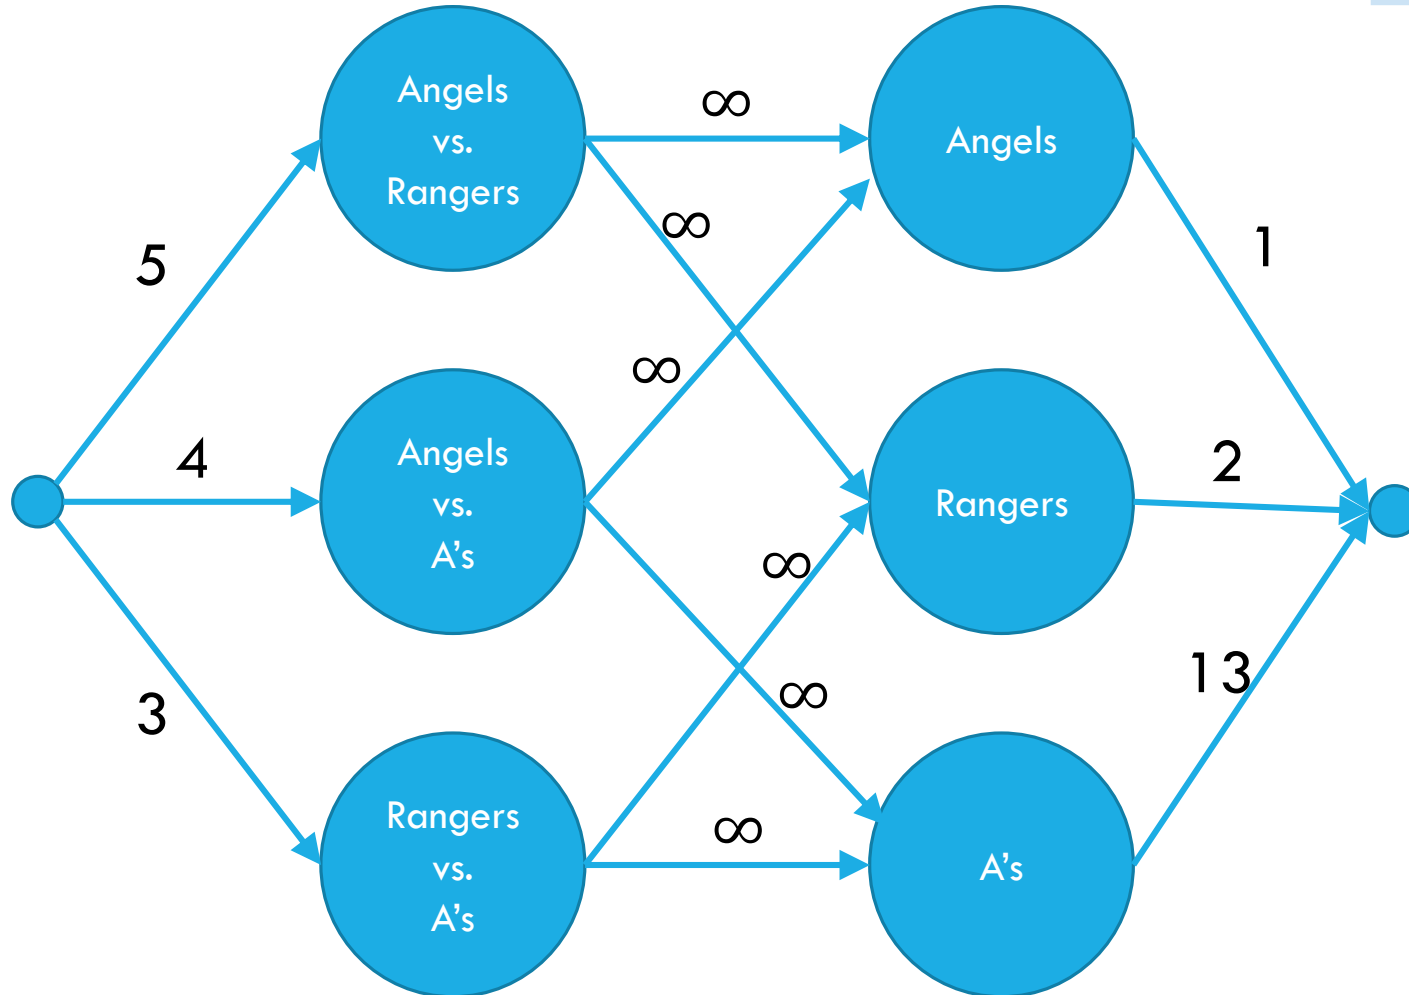


Somehow, Still More Flow

CSE 421 Winter 2023
Lecture 20

Making a Network

	Angels	Rangers	Mariners	A's
Angels	-	5	3	4
Rangers	5	-	4	3
Mariners	3	4	-	5
A's	4	3	5	-



Team	Wins (w)	Possible Wins (P)
Angels	81	93
Rangers	80	92
Mariners	70	82
A's	69	81

We're done!

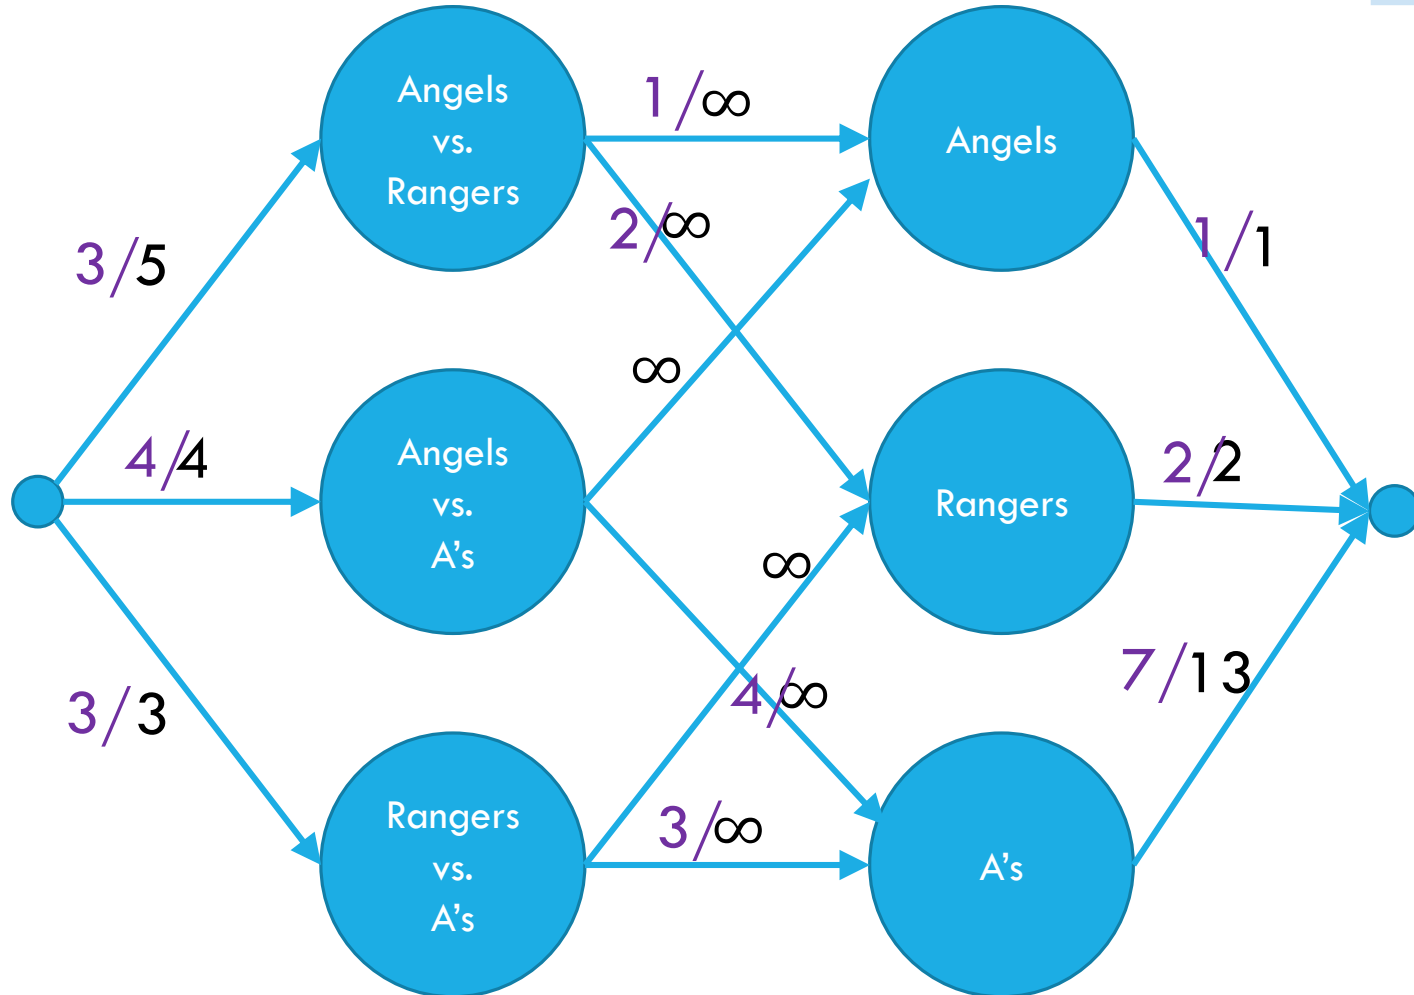
Interpreting the answer

If the max flow has value equal to number of games, we know how the Mariners can still win the division.

If the max flow is less than that, the Mariners can't win the division!

(if they could win the division, then there is a way that the remaining games could play out with the mariners having as many wins as anyone else, but then we could make a feasible flow by assigning a unit of flow for each winner).

Max Flow



	Angels	Rangers	Mariners	A's
Angels	-	5	3	4
Rangers	5	-	4	3
Mariners	3	4	-	5
A's	4	3	5	-

Team	Wins (w)	Possible Wins (P)
Angels	81	93
Rangers	80	92
Mariners	70	82
A's	69	81

This is the maximum flow. What's the min-cut?

$\{s, \text{Angels vs. Rangers}, \text{Angels}, \text{Rangers}\}$ is one side of the cut.

The Angels and Rangers were enough to prove that the Mariners couldn't win!

Generating Proof that you're eliminated

How do you describe to the general public that the Mariners are eliminated.

People are going to say "the Mariners can still win 82 games, no one has one 82, it's not over yet!"

Of the Angels and Rangers, they will win (combined) at least

81 + 80 + 5 games (Angels wins, Rangers wins, games to be played among these teams)

On average they win $\frac{166}{2} = 83$ games. That's more than 82. Someone is beating that average, and whoever that is the Mariners won't catch them.

In General

Find the max flow. If its value is the number of games remaining, great! Mariners can still win.

If its value is less than that, find the min cut. The set of all teams reachable from s in the residual graph will show you **why** the Mariners are eliminated.

Takeaways


If you want to “assign” things, max-flow might be a good option.

If you say “at most” you can probably just make a capacity constraint

Once you can do an “exactly equal” or “at least” by checking the value of the max-flow.

Sometimes you want an extra layer or two if you have a multiple types of assignments.

Sometimes you can convert an “at least” in one group into an “at most” on another group.



Optional: Why is there always an explanation?

An Explanation Always Exists

g_{ij} is games to be played between i and j
 P is number of wins possible for Mariners
 w_i is current number of wins for team i .

Let (S, \bar{S}) be a min-cut.

There's a lot of structure in the min-cut.


Let R be the set of teams whose vertices are reachable from s after the edges have been cut.

The capacity of the cut is

$$\sum_{i \notin R \text{ or } j \notin R} g_{ij} + \sum_{i \in R} P - w_i$$

And the capacity of the cut is less than $\sum_{i,j} g_{ij}$ (because that is a cut, and we can't have a flow of that value).

If R is a set of teams, let $a(R) = \frac{\sum_{i \in R} w_i + \sum_{i,j \in R} g_{i,j}}{|R|}$ the average number of games won by a team in R .



An Explanation Always Exists

g_{ij} is games to be played between i and j
 P is number of wins possible for Mariners
 w_i is current number of wins for team i .

$$\sum_{i \notin R \text{ or } j \notin R} g_{ij} + \sum_{i \in R} P - w_i < \sum_{i,j} g_{ij}$$

$$\sum_{i \in R} P - w_i < \sum_{i \in R, j \in R} g_{ij}$$

After subtracting pairs where at least one of i, j are not in R all that remains are pairs where both i, j are in R .

$$|R|P < \sum_{i \in R, j \in R} g_{ij} + \sum_{i \in R} w_i$$

Move w_i to the other side. P is a constant, so we just add $|R|$ copies of P .


$$P < \frac{\sum_{i \in R, j \in R} g_{ij} + \sum_{i \in R} w_i}{|R|}$$

That is, the average number of wins for a team in R (after all games are played) is strictly more than the possible number of wins for the Mariners.

Summary

To tell whether your favorite team is eliminated, you can run a max-flow computation on a graph with $O(n^2)$ vertices and $O(n^2)$ edges.

If your team is eliminated, there is a witness set of teams that must average more wins than is possible for your team.



Bipartite Applications



Today

One last day of max-flow.

Two new problems we can solve with max-flow/min-cut.

Today's proofs are short but subtle; I'm intentionally taking us through slowly. Please ask questions.


A (possibly) simple problem

Given a bipartite graph, find a maximum matching.

A *matching* is a set of edges that do not share an endpoint.

Think of it as *matching* the endpoints of the edges to each other.

A *stable matching* is a particular matching on a complete bipartite graph.




A (possibly) simple problem

Given a bipartite graph, find a maximum matching.

A **matching** is a set of edges that do not share an endpoint.


For example, on this graph the red edges are a maximum matching. There is no way to get 4 edges in a matching.



A (possibly) simple problem


Design an algorithm to find a maximum matching on a bipartite graph.

(hint: what if the vertices on one side are chores and the other are housemates).




Algorithm for Bipartite Maximum Matching

Use the “tuple selection/assignment” we did last week!



Algorithm for Bipartite Maximum Matching


Add a dummy source and sink. Attach each to one side.



Algorithm for Bipartite Maximum Matching

Add a dummy source and sink. Attach each to one side.

Direct (previously undirected) edges toward sink




Algorithm for Bipartite Maximum Matching

Add a dummy source and sink. Attach each to one side.

Direct (previously undirected) edges toward sink.

Capacity 1 entering and leaving each (real) vertex ensures matching




Algorithm for Bipartite Maximum Matching

Add a dummy source and sink. Attach each to one side.

Direct (previously undirected) edges toward sink.

Capacity 1 entering and leaving each (real) vertex ensures matching

Capacity ∞ in the middle (it'll help later)




Algorithm for Bipartite Maximum Matching

Add a dummy source and sink. Attach each to one side.

Direct (previously undirected) edges toward sink.

Capacity 1 entering and leaving each (real) vertex ensures matching

Capacity ∞ in the middle (it'll help later)



Algorithm for Bipartite Maximum Matching


Add a dummy source and sink. Attach each to one side.

Direct (previously undirected) edges toward sink.

Capacity 1 entering and leaving each (real) vertex

Capacity ∞ in the middle (it'll help later)

Red edges have flow.
Take edges with flow
for the matching.



Algorithm for Bipartite Matching

Modify the (undirected) graph G into the network flow graph.

Find a maximum flow, taking all edges of G which have flow.

Is it correct?

The set of edges found should be a matching.

There should be no larger matching.

Algorithm for Bipartite Matching

Modify the (undirected) graph G into the network flow graph.

Find a maximum flow, taking all edges of G which have flow.

Is it correct?

The set of edges found should be a matching.

Capacities ensure no edge has more than one unit of flow through it. By integrality of flow, we have only one edge.

There should be no larger matching.

Suppose, for contradiction, that M is a larger matching, then put a unit of flow on M , entering each vertex with an endpoint in M on the left and leaving on the right. This is a valid flow! But then its value is $|M|$, which would be larger than the max-flow, a contradiction.

Solving a new problem (with matchings)

Let G be an undirected graph.

A minimum vertex cover is the smallest set of vertices such that every edge has at least one of its endpoints in the set.


Here's an algorithm to find a minimum vertex cover in a bipartite graph...

Vertex Cover

How to find a vertex cover?


Well a max-flow says "you can't use this edge; (at least) one of its endpoints is already used by the matching."

Here's a vertex cover (in gold). Notice something about it?




Vertex Cover

Here's a vertex cover (in gold). Find the min-cut (write residual graph)




Vertex Cover

Flow graph



Residual;
Cut in blue X's




Vertex Cover

How to find a vertex cover?

Well a max-flow says "you can't use this unsaturated edge; (at least) one of its endpoints is already used by the matching.

Here's a vertex cover (in gold). Notice something about it? It looks a lot like the min cut.




Vertex Cover

Let A be the left bipartition, B the right side. Let (S, T) be the min-cut.

$A_S = A \cap S$; A_T, B_S, B_T defined correspondingly.

Here, A_T and B_S form a vertex cover



Vertex Cover

Is $A_T \cup B_S$ always a vertex cover? If so, how big is it?


There are 4 potential kinds of edges. Which kind is a problem for the vertex cover? Can they all exist?

A_S to B_S

A_S to B_T

A_T to B_S

A_T to B_T



Vertex Cover

Is $A_T \cup B_S$ always a vertex cover? If so, how big is it?

There are 4 potential kinds of edges. Which kind is a problem for the vertex cover? Can they all exist?


A_S to B_S

A_S to B_T Only kind we have to worry about!!

A_T to B_S

A_T to B_T

But you can't have an edge from S to T ! We have a vertex cover.




Vertex Cover

How big is $A_T \cup B_S$?

Well, it looks a lot like the min-cut. (s, A_T) and (B_S, t) are cut.

Each of those edges is capacity 1. Each has an endpoint we put in the vertex cover! $|A_T \cup B_S| = \text{cap}(S, T) = |f|$.



But is it a minimum VC?


So it's a vertex cover, of size equal to the min-cut...what if there's a smaller vertex cover?

Let X be any vertex cover, define A_X, B_X as before.

No smaller VC

Let X be any vertex cover, define A_X, B_X as before.

Claim: $(\{s\} \cup B_X \cup (A \setminus A_X), \{t\} \cup (B \setminus B_X) \cup A_X)$ is a cut with cut edges are $(s, A_X), (B_X, t)$. [this is the corresponding cut to before]



But is it a minimum VC?

So it's a vertex cover, of size equal to the min-cut...what if there's a smaller vertex cover?

Let X be any vertex cover, define A_X, B_X as before.


Claim: $(\{s\} \cup B_X \cup (A \setminus A_X), \{t\} \cup (B \setminus B_X) \cup A_X)$ is a cut with cut edges are $(s, A_X), (B_X, t)$.

No $(A \setminus A_X), (B \setminus B_X)$ edges because we're a cover!

No smaller VC

Every vertex cover lets you discover a cut of the same size.

Our algorithm found a VC of size equal to the minimum cut! A smaller VC would give a smaller cut! But there isn't one. So we have found the min VC.



Wrapping it up

You can find the following of a bipartite graph (using only flow)

1. Maximum matching
2. Minimum Vertex cover

And their value is equal to the size of the max-flow (and the min-cut) in the modified graph.

“König’s Theorem” (aka König-Egevary Theorem)

In a bipartite graph, the size of the maximum matching is equal to the size of the minimum vertex cover.

Wrapping It Up

We've found maximum matchings and minimum vertex covers in bipartite graphs using flow.

If your graph isn't bipartite:

Efficiently finding a maximum matching is still possible.
but more complicated. Look up "Edmond's Blossom Algorithm"

Efficiently finding a minimum vertex cover, is a taller task.

It's an NP-complete problem. We'll more clearly define what that means next week.

But Wait, There's More Applications!

Finding Edge disjoint paths in a directed graph.

Maximum number of paths from u to v that don't repeat an edge.

I want to send a packet from u to v but edges are unreliable. Want completely independent copies.

Finding internally-vertex-disjoint paths in a directed graph

Send packet, but I don't trust the vertices.

Image Segmentation

A surprisingly useful tool for matching and separating things.

Also for proving graph theory results (Konig-Egevary, Hall's Theorem)