

More Network Flow

CSE 417 Winter 24
Lecture 19

Announcements

HW4 P5 (Force Dynamics---one of the programming questions) extended to deadline of this Monday

We're going to shift around office hours for the last few weeks.
Announcement with all the details coming over the weekend

Monday is a holiday!

A few hours shifted or cancelled. None in-person, only zoom.

Announcement coming tonight!


No lecture Monday.

Max Flow

We have a directed graph G , a source vertex s and a target vertex t .

We have some thing (water or data packets) we have to send from s to t .

Every edge has a capacity, it can only handle so many.




Flows

A **flow** moves units of water from s to t .

Water can only be created at s and only disappear at t .

And you cannot move more water than the capacity on any edge.





Flows

A **flow** moves units of water from s to t .

Water can only be created at s and only disappear at t .

And you cannot move more water than the capacity on any edge.





Finding Max Flows

Residual Graph

In general:

If the original graph has an edge (u, v) of capacity c , and the flow sends $f_{u,v}$ along (u, v) :

Include (u, v) in the residual with capacity $c - f_{u,v}$ as long as $c - f_{u,v} > 0$ (if equal to zero, don't include the edge)

Include (v, u) [the edge going in the reverse direction] with capacity $f_{u,v}$ as long as $f_{u,v} > 0$

Ford-Fulkerson Algorithm

While(flow is not maximum)

Run BFS in residual graph starting from s .

Record predecessors to find an s, t -path

Iterate through path, finding c minimum residual capacity on path.

Add c to every edge on path in flow

Update residual graph

Ford-Fulkerson Algorithm

While(true)

Run BFS in residual graph starting from s .

Record predecessors to find an s, t -path

If you don't reach t , break. //otherwise you can still augment

Iterate through path, finding c minimum residual capacity on path.

Add c to every edge on path in flow

Update residual graph

Assuming $E \geq V$ (isolated vertices won't affect the flow, so this is reasonable).

Running Time: $O(E)$ per iteration. Number of iterations?

Ford-Fulkerson Algorithm

If we have all integer capacities at the start...

The residual graph will always have integer capacities.

Why? The minimum capacity on the first path is an integer.

So we subtract or add integers to the residual graph.

And the result is more integers!

So in every iteration we add at least 1 unit of flow!

Ford-Fulkerson Algorithm

While(true)

Run BFS in residual graph starting from s .

Record predecessors to find an s, t -path

If you don't reach t , break. //otherwise you can still augment

Iterate through path, finding c minimum residual capacity on path.

Add c to every edge on path in flow

Update residual graph

Running Time:? $O(E)$ per iteration. Number of iterations? $O(f)$, where f is the value of the maximum flow. Total $O(Ef)$

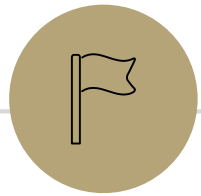
Wait... f ?

We haven't seen a running time before that depends on the **answer**

Normally, we want the running time directly in terms of the input.

There are tricks to speed up Ford-Fulkerson so you don't take too long if f is really big.

Some optional content about this at the end of Wednesday's slide deck.



Minimum Cut



What's a Cut?

For directed graphs (like we have here)


An (s, t) -cut, is a split of the vertices into two sets (S, T)

So that s is in S , t is in T ,
and every other vertex is in exactly one of S and T .


The capacity of a cut (or size of a cut) is the capacity of the edges going from s to t (don't count capacity from t to s).

An Example

Graph, with flow



"Residual Graph"




We can't get from s to t anymore in the residual graph. We're done! That's a maximum flow.


But...why?

An Example

Graph, with flow



"Residual Graph"



Cut is $(s, V - s)$ (i.e. s on one side, everything else on the other)

Edges from s to everything else? Capacity 30

We can't get more than 30 units of flow from s to t . Because it all (simultaneously) must cross from one side to the other.

How Do We Know?

How do we know the algorithm is done?

When we can't we get from s to t ?

We've **cut** the (residual) graph.

s and all the vertices you can reach from it on one side
and t and all the vertices s can't reach on the other.

Take a look at the edges spanning the cut in the original graph.

In our first graph, that capacity was equal to the value of the max flow.

Finding the min-cut

Maintain the residual graph.

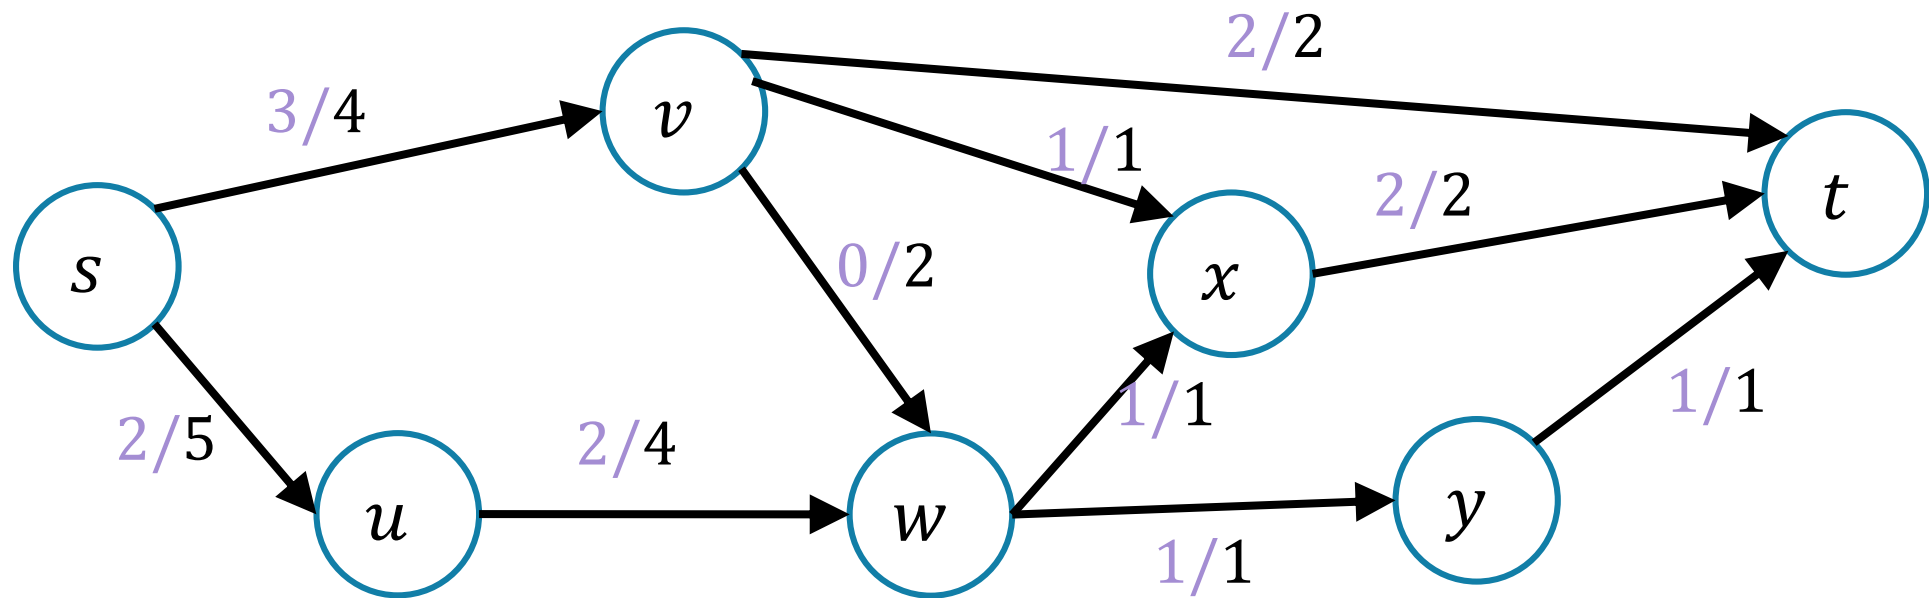
When you search from s and can't get to t :

s and everything you can reach from s is on one side of the cut
 t and everything you can't reach from s is on the other side.

Another Example

We started lecture with this flow.

What's the residual graph? What's the cut?

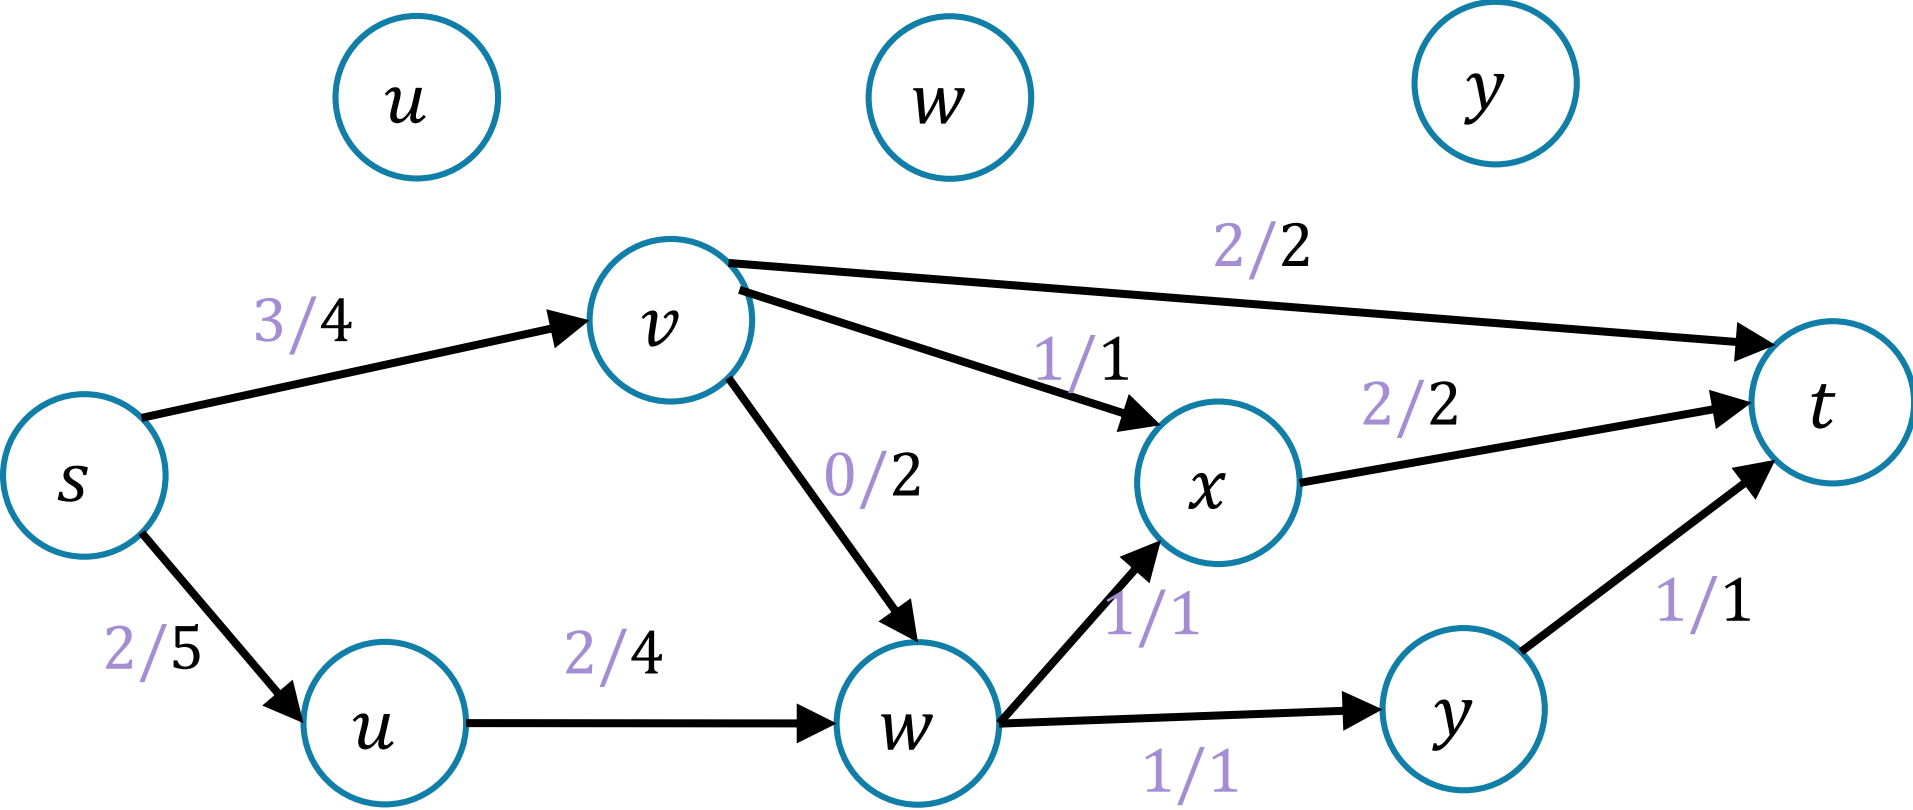


Another Example


Residual



Flow



Another Example



$(\{s, u, v, w\}, \{x, y, t\})$ is the cut in the residual graph.

What edges span?

$(v, x), (v, t), (w, x), (w, y)$

Total capacity? 5. What's the value of the flow? 5

Max Flow-Min Cut Theorem

Max-Flow-Min-Cut Theorem

The value of the maximum flow from s to t is equal to the value of the minimum cut separating s and t .

No proof, but here's some intuition:

Every cut is an upper-bound on the value of the flow (you can't have the total flow value higher than the cut).

Cut will be saturated (i.e. full flow going through all the edges) otherwise residual graph will have another edge. So flow and cut are equal values.

So What?

Max-flow and min-cut are each interesting algorithmic problems.

They were first studied in the 1950s

The U.S. military wanted to know how much the Soviets could ship on their rail network.

And also which rail lines they would target to *disrupt* the network.


So What?

Great quick check for if you've found the maximum flow (or min-cut).
Check the other and see if the value is the same!

We'll see examples next time of max-flow used for modeling. In those cases the min-cut can be interpreted as a "barrier" to a good assignment.

It's also a nice example of **duality**

If you know what a "dual linear program" is – flows and cuts are dual LPs.



Applications



Applications of Max-Flow-Min-Cut

Max-Flow and Min-Cut are useful if you work for the water company...
But they're also useful if you don't.

The most common application is assignment problems.

You have jobs and people who can do jobs – who is going to do which?

Big idea:

Let one unit of flow mean “assigning” one job to a person.

Hey Wait...

Isn't this what stable matching is for?

Stable matching is very versatile, and it lets you encode preferences.

Max-flow assignment is even more versatile on the types of assignments.

But there's not an easy way to encode preferences.

Example Problem

You and your housemates need to decide who is going to do each of the chores this week.

Some of your housemates are unable to do some chores.

Housemates: 1,2,3

Chores:

Arrange furniture, clean the **B**athroom, **C**ook dinner, do the **D**ishes


Housemate 1 is unable to arrange furniture, 2 is unable to cook.

Example Problem

Housemate 1 is unable to arrange furniture, 2 is unable to cook.

Vertex for each housemate and chore.

Edge if the housemate **could** do the chore




Example Problem

Housemate 1 is unable to arrange furniture, 2 is unable to cook.

Vertex for each housemate and chore.

Edge if the housemate **could** do the chore




Example Problem

Housemate 1 is unable to arrange furniture, 2 is unable to cook.

Vertex for each housemate and chore.

Edge if the housemate **could** do the chore




Example Problem

Housemate 1 is unable to arrange furniture, 2 is unable to cook.

Vertex for each housemate and chore.

Edge if the housemate **could** do the chore




Example Problem

Idea: Flow from 1 to B means "make housemate 1 do chore B."

Every chore needs to be done (by one person).

Every person needs to do at most two chores.




Example Problem

Idea: Flow from 1 to B means "make housemate 1 do chore B."

Every chore needs to be done (by one person).

Every person needs to do at most two chores.




Example Problem

Idea: Flow from 1 to B means "make housemate 1 do chore B."

Every chore needs to be done (by one person).

Every person needs to do at most two chores.




Example Problem

Idea: Flow from 1 to B means "make housemate 1 do chore B."

Every chore needs to be done (by one person).

Every person needs to do at most two chores.




Example Problem

What are the capacities for the middle edges?

Could make them 1 (make sure you don't get "two units of cooking")

All our requirements are already (implicitly) encoded. So could make them ∞ instead.




Example Problem

Find a max flow...And read off the assignment!

Full color: 1 unit of flow, faded: 0 units of flow

1 cleans the bathroom and does the dishes, 2 arranges furniture, 3 cooks.



Why are all of our constraints met?

Every chore gets done

No one does more than 2 chores

People only do chores they're capable of

Why are all of our constraints met?

Every chore gets done

A flow of value 4 sends one unit of flow through each of A,B,C,D (because the edges to t are all capacity 1), so a max-flow ensures if possible we'll find an assignment.

No one does more than 2 chores

Only 2 units of flow can go through any person vertex (because edges from s to people are all capacity 2).

People only do chores they're capable of

There is only an edge from a person to a chore if they can do that chore.

One More Requirement...

There's another requirement we haven't mentioned:

People only get "whole units" of chores
i.e. you don't have two people each doing half of the cooking.

The max-flow approach guarantees this! As long as our requirements are integers (or ∞) as well.

Another Problem

pollev.com/robbie

You run two coffee shops. You have to decide who will work at which of your shops today:

A, B, C are all capable of managing a shop.

D, E, F, G are all regular employees (can't be a manager)

You need at least one manager at each shop, at least 3 people (total) at shop 1 and at least 4 people (total) at shop 2.

Hint: think of assigning managers and non-managers as separate...



One More Example

A classic example

We'll also be able to use the min-cut in addition to the flow!

Question: Can the Mariners still win* the division?

*or at least tie for first place.

And if they can't, can you explain why.

Can The Mariners Win The Division?

It's late at night September 14, 1998.

You're working for the Seattle Times.

The Mariners won! But the Angels did too.

How do you frame the Mariners current situation in your postgame article?

Team	Wins (<i>w</i>)	Games Left
Angels	81	12
Rangers	80	12
Mariners	70	12
A's	69	12

MLB rules say all games will be played (if they end up mattering) so you can assume those will happen.

Can The Mariners Win The Division?

Team	Wins (w)	Games Left	Possible Wins (P)
Angels	81	12	93
Rangers	80	12	92
Mariners	70	12	82
A's	69	12	81

$P_{MARINERS} \geq w_i$ for all i , so the Mariners can win the division, right?

Well...No

The teams will play each other, here are the number of games to be played against each other.

g_{ij}	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)	Games Left	Possible Wins (P)
Angels	81	12	93
Rangers	80	12	92
Mariners	70	12	82
A's	69	12	81

Well...No

At least one of the Angels and Rangers is going to win at least 83 games

someone wins at least three of the five they play against each other.

The Mariners can only win 82 games.

Lessons

Comparing P_i to w_j is insufficient to tell if a team is eliminated.

The teams are interconnected by the games they play against each other.

Let's find a way to do this calculation...not by hand.

What do we need to assign?

Assignment


We need to assign who wins each of the remaining games.

Safe to assume the Mariners will win them all.

Just need to figure out the others.

One unit of flow represents one win.

Making a Network



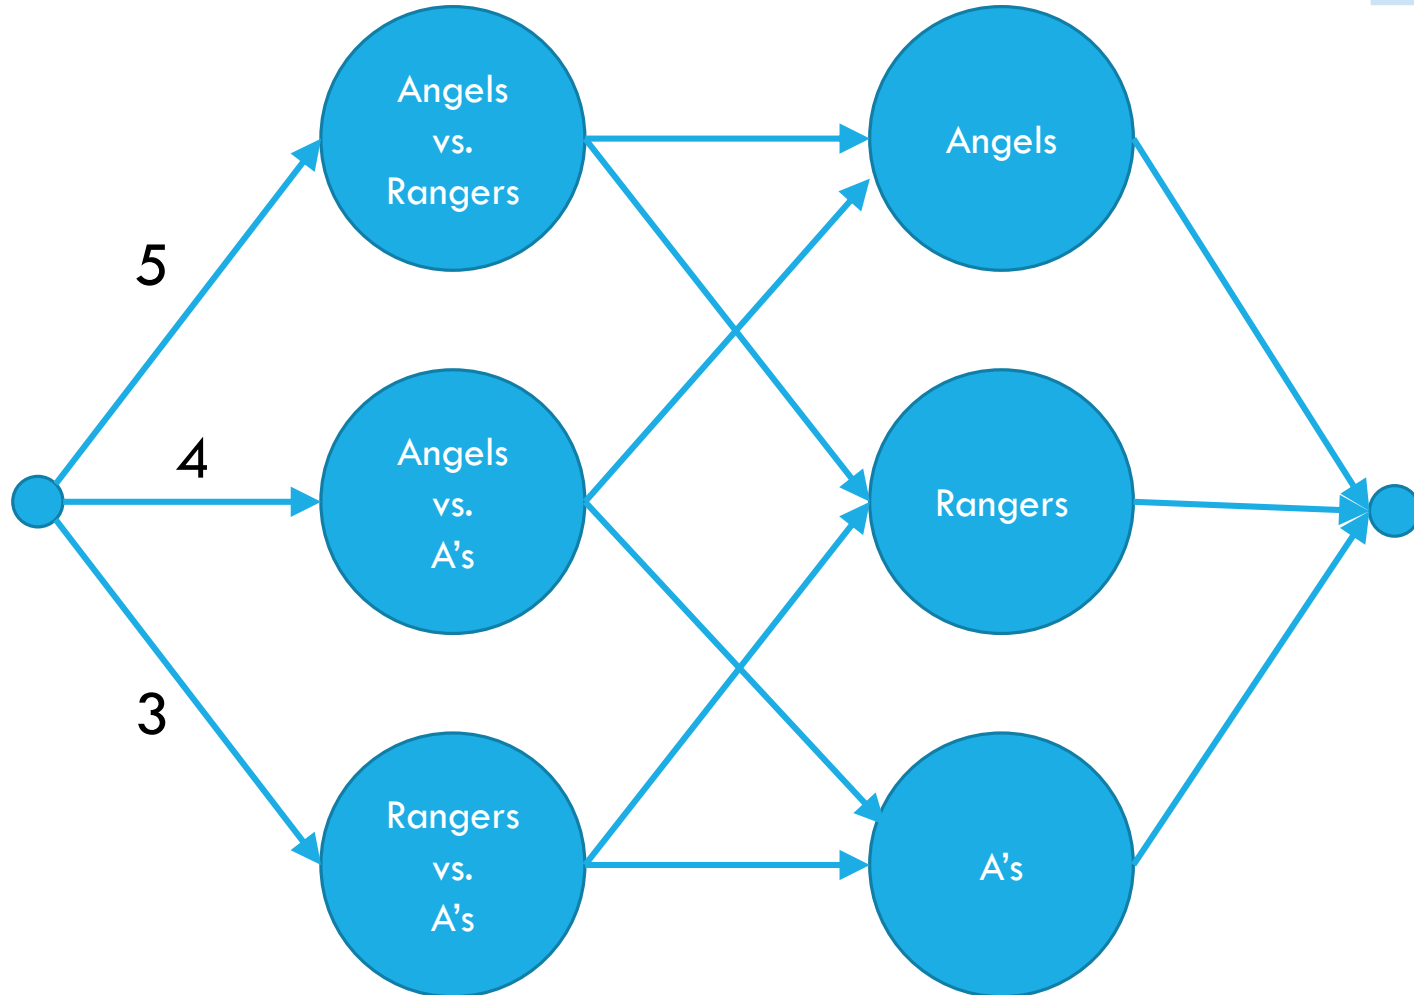
	Angels	Rangers	Mariners	A's
Angels	-	5	3	4
Rangers	5	-	4	3
Mariners	3	4	-	5
A's	4	3	5	-

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

s, t on the ends

First layer is pairs of opponents
(i.e. what game is being played)
Second layer is individual teams.

Making a Network

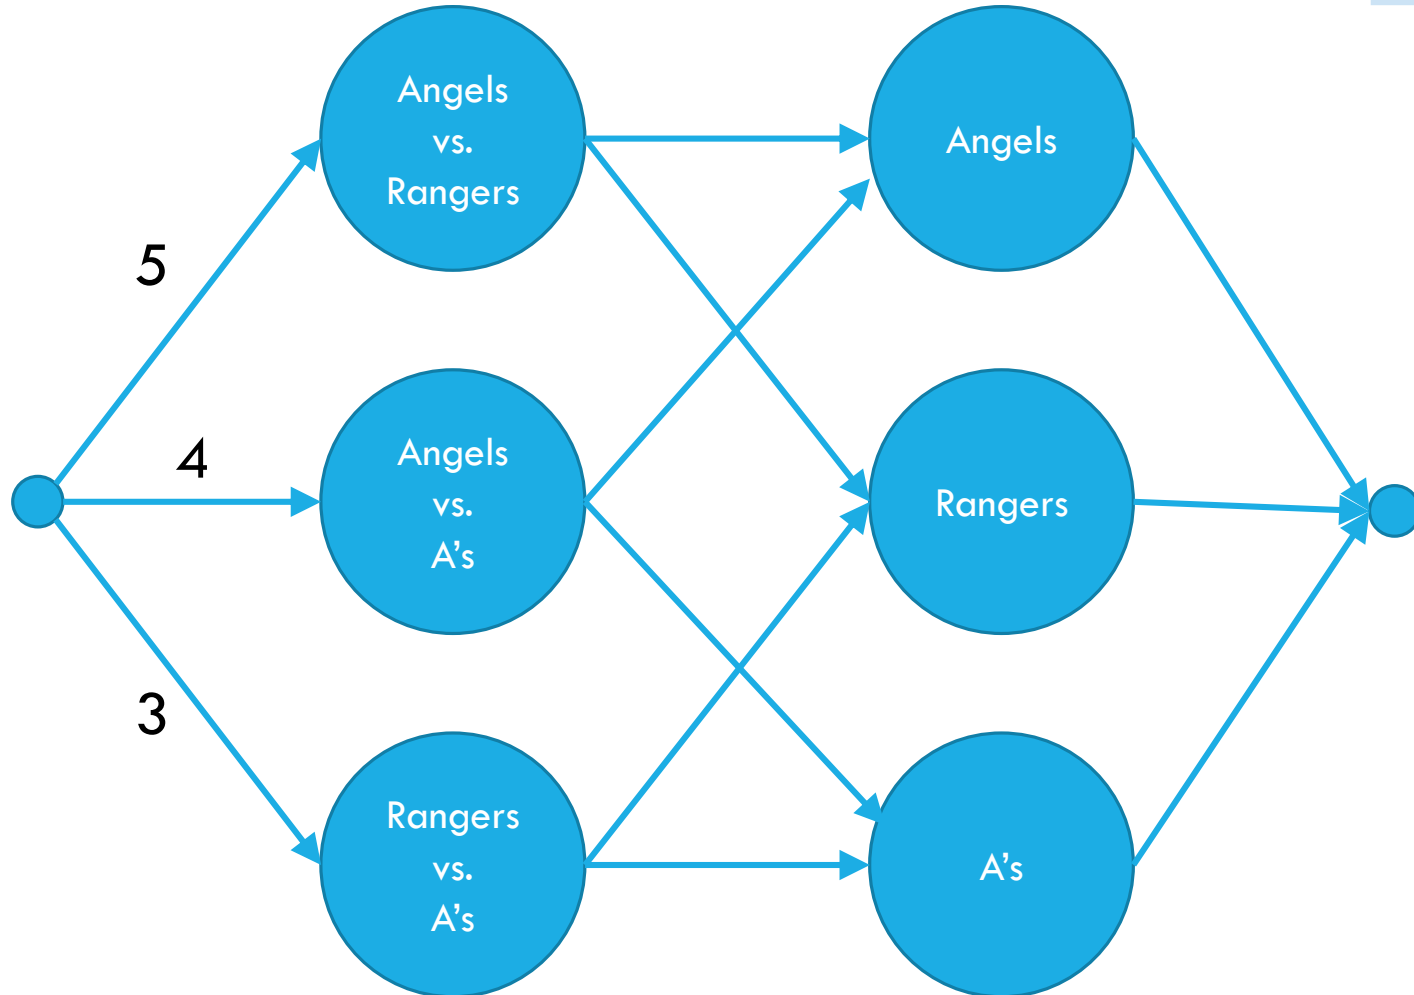


	Angels	Rangers	Mariners	A's
Angels	-	5	3	4
Rangers	5	-	4	3
Mariners	3	4	-	5
A's	4	3	5	-

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

Put number of games to be played from s to pairs

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

How do we make sure Mariners win? They'll end the season with 82 wins (current + games left). How many more can each team win?
 Mariners poss total – team current

Making a Network

Angels have 81 wins, 1 more is ok (total matches Mariners possible) 2 is not. Capacity is 1.




	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

How do we make sure Mariners win? They'll end the season with 82 wins (current + games left).
 How many more can each team win?
 Mariners poss total – team current

Making a Network

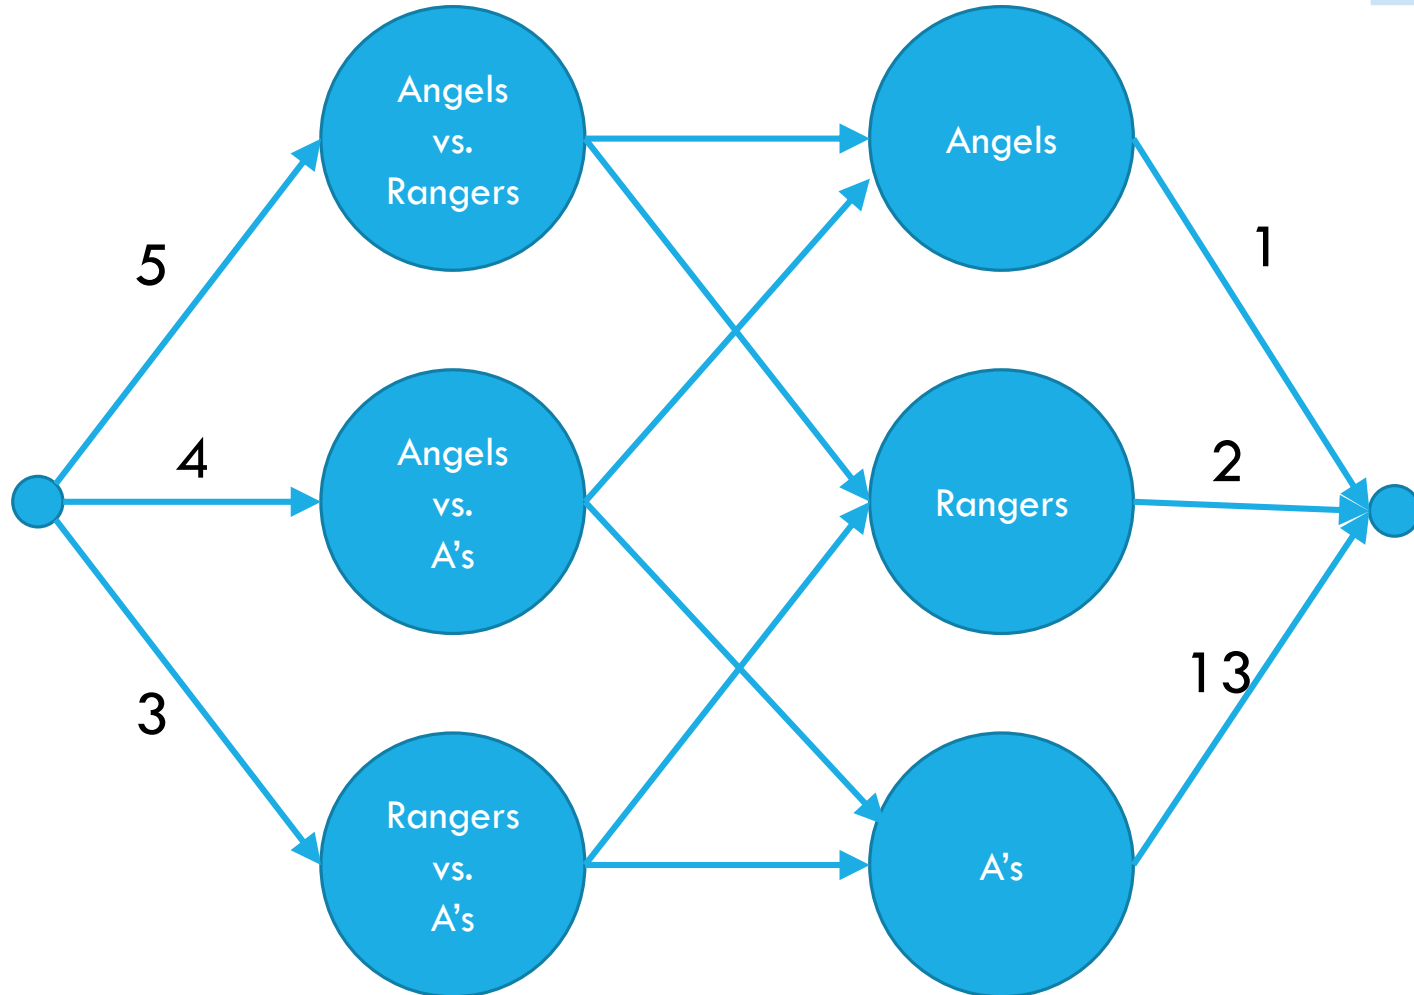


	Angels	Rangers	Mariners	A's
Angels	-	5	3	4
Rangers	5	-	4	3
Mariners	3	4	-	5
A's	4	3	5	-

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

How do we make sure Mariners win? They'll end the season with 82 wins (current + games left).
 How many more can each team win?
 Mariners poss total – team current

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

Edges in the middle?
Only to the two teams playing.

We've handled are constraints, can leave capacities at ∞ .

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	-



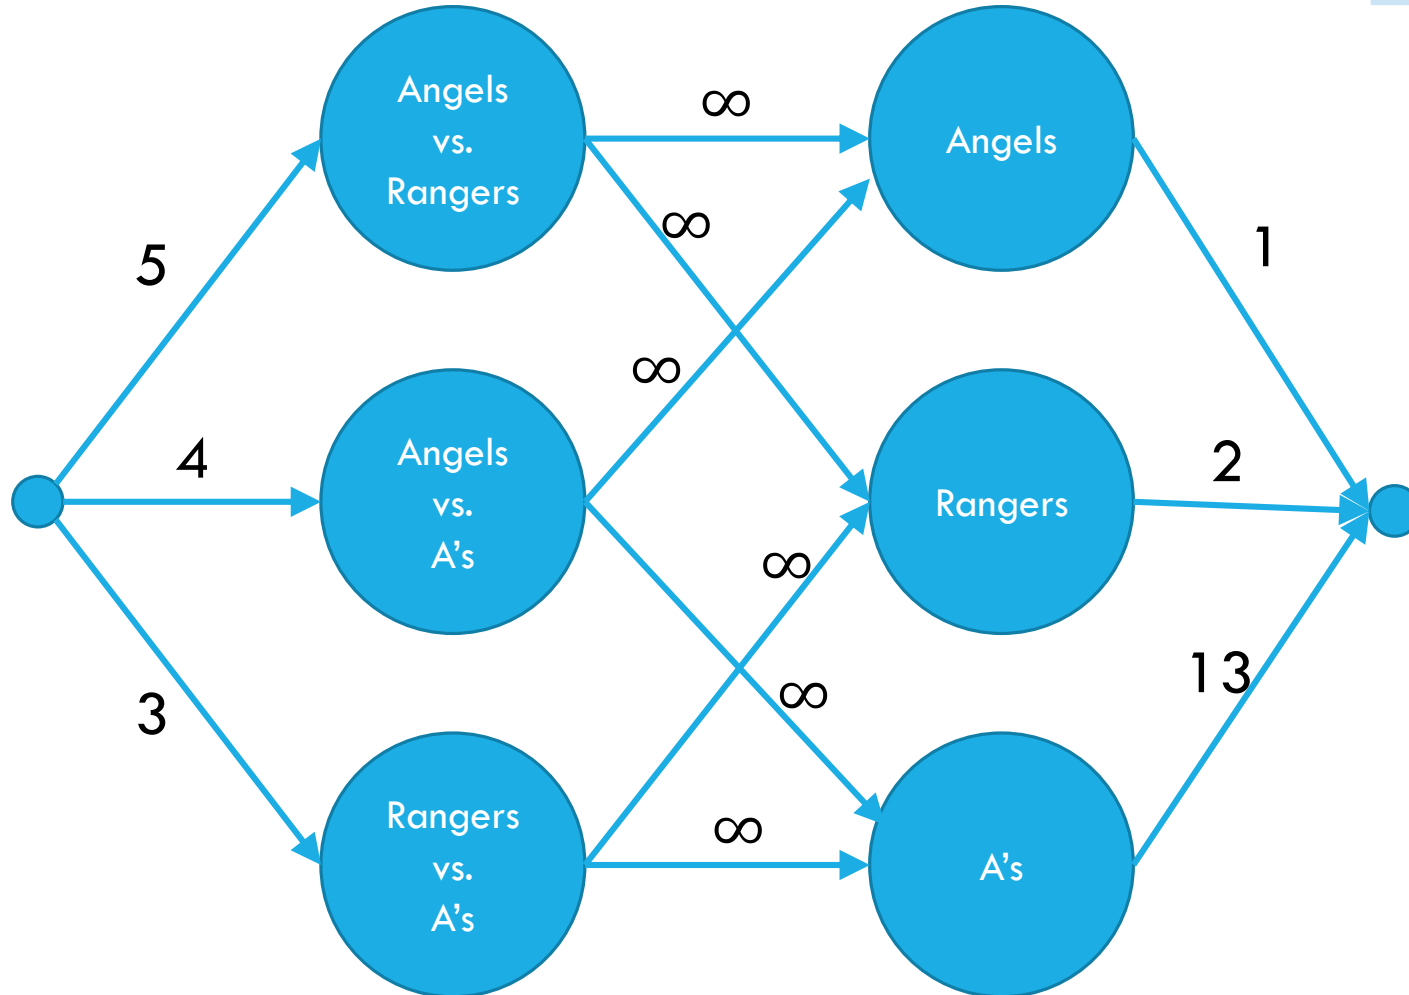
Team	Wins (w)	Possible Wins (P)
Angels	81	93
Rangers	80	92
Mariners	70	82
A's	69	81

Edges in the middle?
Only to the two teams playing.

We've handled are constraints, can leave capacities at ∞ .

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!

Why are all the constraints met?

How many games are there to play? Equal to the capacities leaving s .

So if we have a flow of at least that value, we'll assign winners to all the games.

Why will the Mariners win with this assignment?

The capacity from team A to t ensures A will not end with more wins.

No "half-wins" or anything weird?

All capacities are integers, so we'll get an integer solution!


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.