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 $f_{u,v} > 0$

Another Example

Residual

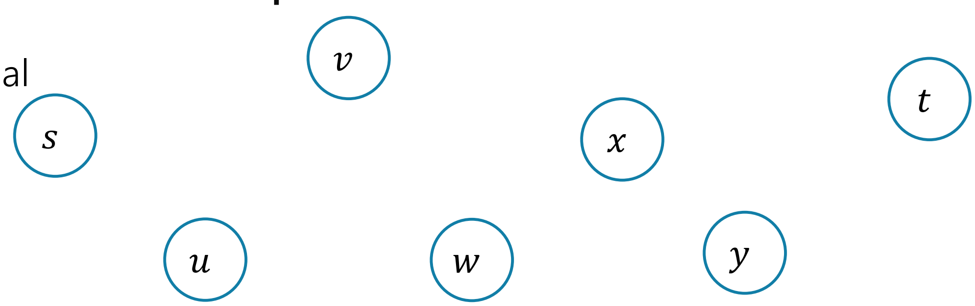

Flow

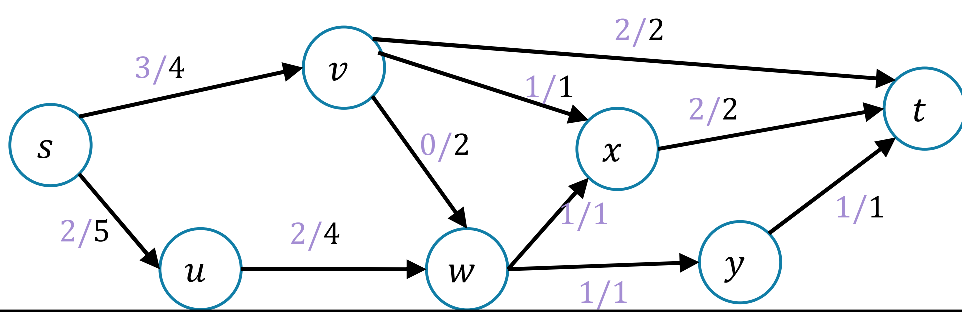

Some notation (more formally)

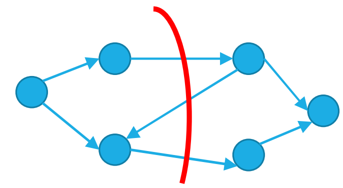

Let f be a flow.

For an edge e, $f(e)$ is the flow on e.

$\text{val}(f)$ is the sum of flow leaving s (equivalently entering t).

For a cut (A, B), $\text{cap}(A, B) = \sum_{e:e=(u,v),u\in A,v\in B} c(e)$

i.e., the sum of the capacities on edges going from A to B.
Direction matters!

Notice the capacity of a cut is independent of any particular flow. It's a property of the original graph, not the flow or the residual graph.

Step 1: The Flow Goes Somewhere

For every s-t cut, (A, B):

$$\text{val}(f) = f^{out}(A) - f^{in}(A) = \sum_{e=(u,v):u\in A,v\in B} f(e) - \sum_{e=(v,u):u\in A,v\in B} f(e)$$

Intuitively, the net-flow for *every* cut is the same as the net flow for the cut $(s, V \setminus \{s\})$.

Why? Well the flow has to go somewhere! It can only disappear at t.

Why care? It's a technical observation we'll need later.