CSCI 3104 Recitation- Network Flows

Michael Levet

1 Learning Objectives

The primary goal of recitation this week is to become familiar with key terminology and constructions related to network flows. In particular, we have the following learning outcomes.

- Students will correctly identify the amount of flow that can be pushed in both the forward and backward directions of edges.
- Students will correctly identify flow-augmenting paths, including the amount of flow that can be pushed along such paths.
- Students will correctly identify paths that are not flow-augmenting and provide correct justification as to why the paths are not flow-augmenting.
- Students will construct residual graphs.
- Students will identify matchings on graphs.
- Students will determine the sizes of maximum matchings.
- Students will use maximum flow procedures to determine maximum-size bipartite matchings.

2 Network Flows

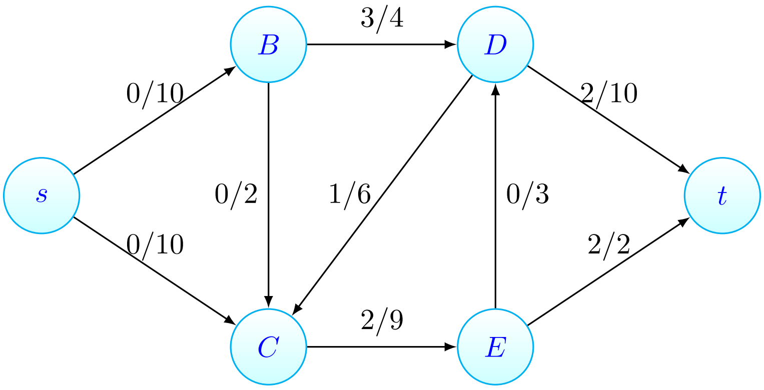

Figure 1. Consider the flow network $\mathcal{N}$.

(Recommended) Problem 1. Consider the flow network $\mathcal{N}$. How much flow can be pushed in the forward direction along the edge (B, D)? That is, how much flow can be pushed from $B \to D$?

Example 1. While we can push flow forward across an edge, we can also push flow backwards. The intuition is that we re-route flow along a flow-augmenting path. **However, only existing flow can be re-routed.**

(a) Consider the edge (D, C) in the flow network $\mathcal{N}$ above, which has a capacity of 6 units of flow and currently has 1 unit of flow being routed through it. So we can push 5 units of flow from $D \to C$, or we can push 1 unit of flow backwards from $C \to D$.

(b) Consider the edge (E, D). As there is no flow being pushed from $E \to D$, we cannot push any flow backwards from $D \to E$.

(Recommended) Problem 2. Consider the flow network $\mathcal{N}$ above.

(a) How much flow can be pushed in the backward direction along the edge (B, D)? That is, how much flow can be pushed from $D \to B$?

(b) How much flow can be pushed in the backward direction along the edge (B, C)? That is, how much flow can be pushed from $C \to B$?

The goal of the Network Flow problem is to maximize the amount flow that we can route from a source vertex s to a destination (or *sink*) vertex t. The approach is to find a path through which we can push positive flow and then push as much flow as possible. Such paths are called *flow-augmenting paths.*

(Recommended) Problem 3. Consider the flow networks $\mathcal{N}$ above. Determine whether each path is a flow-augmenting path. If the given path is flow-augmenting, then clearly indicate the maximum amount of flow that can be pushed through the path. If the given path is **not** flow-augmenting, then clearly explain **why** the path fails to be flow-augmenting.

For the purpose of this problem, do **not** update $\mathcal{N}$ after considering a flow-augmenting path. That is, you should work part (b) without considering part (a), and so on. The point of this problem is solely to identify flow-augmenting paths.

(a) $s \to B \to D \to t$.

(b) $s \to C \to B \to D \to t$.

(c) $s \to C \to E \to t$.

(d) $s \to C \to E \to D \to t$.

(e) $s \to C \to D \to t$.

(Recommended) Problem 4. Find a flow augmenting path in $\mathcal{N}$ that was not listed in Problem 3 . Clearly indicate the maximum amount of flow that can be pushed through this path.

Remark 1. Finding flow-augmenting paths can be challenging, particularly as we can push flow in both the forward and backwards directions. To this end, we introduce the residual graph. Informally, the residual graph replaces each directed edge in a flow network with at most weighted two edges. The weights on the edges indicate how much flow can be pushed along a given direction. Should we not be able to push any flow along an edge in the original flow network, then the corresponding directed edge is omitted in the residual graph. We make this precise with the following definition.

Definition 1. Let $\mathcal{N}(V, E, \text{cap}, f)$ be a flow network, where V is our vertex set, E is our edge set, $\text{cap} : E \to \mathbb{R}_{\geq 0}$ identifies our edge capacities, and $f : E \to \mathbb{R}_{\geq 0}$ identifies the amount of flow being pushed across each edge. The *residual graph* $\mathcal{R}(V, E, w)$ is a directed, weighted graph with the same vertex set as $\mathcal{N}$. For each directed edge $(u, v) \in E$, we include at most two directed edges in $\mathcal{R}$, as follows:

- If $\text{cap}((u, v)) - f((u, v)) > 0$, we include the edge (u, v) in $E(\mathcal{R})$ with weight $\text{w}((u, v)) = \text{cap}((u, v)) - f((u, v))$. That is, the edge (u, v) in $\mathcal{R}$ indicates that we can push $\text{cap}((u, v)) - f((u, v))$ units of flow from $u \to v$.
- If $f((u, v)) > 0$, we include the edge (v, u) in $E(\mathcal{R})$ with weight $w((u, v)) = f((u, v))$. That is, the edge (v, u) in $\mathcal{R}$ indicates that we can push $f((u, v))$ units of flow in the backwards direction along the edge (u, v) in $\mathcal{N}$. That is, in $\mathcal{N}$, we can push $f((u, v))$ units of flow from $v \to u$.

Note that if cannot push flow in a given direction, then we do not include a directed edge in the residual graph $\mathcal{R}$.

Example 2. Recall the flow graph $\mathcal{N}$ at the top of Page 1. We construct the residual graph $\mathcal{R}$ corresponding to $\mathcal{N}$, which is pictured below.

- In $\mathcal{N}$, we have the directed edge (B, D) which has a capacity of 4 and 3 units of flow being pushed through it.
  - As we can push at most 1 more unit of flow in the $B \to D$ direction, we include the directed edge (B, D) in the residual graph $\mathcal{R}$, with the weight of (B, D) as 1.
  - As we have 3 units of flow being pushed from $B \to D$, we can push at most 3 units of flow in the $D \to B$ direction. So we also include the directed edge (D, B) in the residual graph $\mathcal{R}$. The weight of the (D, B) edge in $\mathcal{R}$ is 3, indicating we can push 3 units of flow from $D \to B$ in the original flow network $\mathcal{N}$.
- In $\mathcal{N}$, we hvae the directed edge (s, C), which has a capacity of 10 an no flow being pushed through it. For this reason, we only include the directed edge (s, C) with weight 10 in the residual graph $\mathcal{R}$. As there is no flow being pushed through the (s, C) edge, we do not include the edge (C, s) in $\mathcal{R}$.

The rationale for including the remaining edges in the residual graph $\mathcal{R}$ is similar.

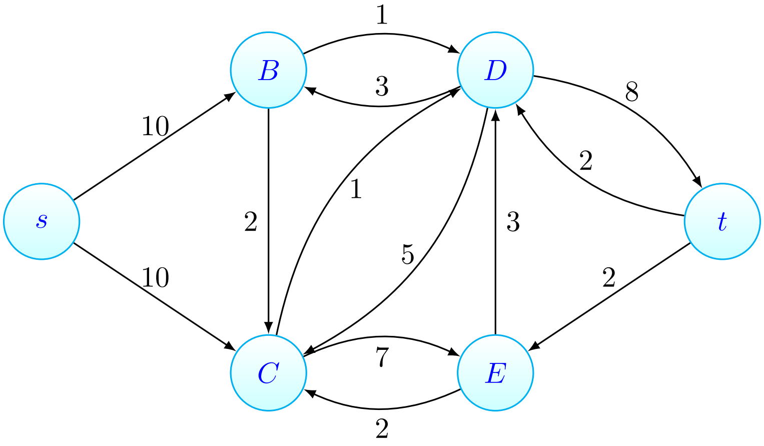

(Recommended) Problem 5. Draw the residual graph corresponding to the flow network $\mathcal{N}$ pictured below.

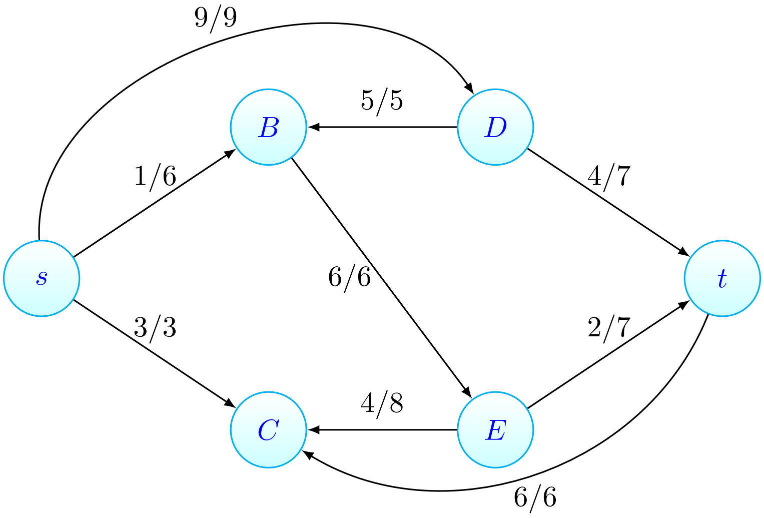

3 Applications of Network Flows: Bipartite Matching

A key application of finding maximum flows is in constructing maximum-size matchings in bipartite graphs. We begin with the notion of a matching before examining how network flows are useful in constructing matchings in bipartite graphs.

3.1 Preliminaries- Bipartite Matchings

We begin with the definition of a matching.

Definition 2. Let $G(V, E)$ be a simple, undirected graph. A *matching* $\mathcal{M}$ is a subset of $E(G)$ such that if $\{i, j\}, \{u, v\} \in \mathcal{M}$ are distinct edges, then $\{i, j\} \cap \{u, v\} = \emptyset$. That is, no two distinct edges in $\mathcal{M}$ touch or share any endpoints.

Remark 2. The above definition of a matching makes sense, even if the graph is not bipartite. We illustrate with the next example.

Example 3. Consider the cycle graph C_3 pictured below.

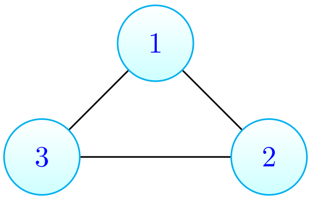

The graph C_3 has four matchings.

- Let $\mathcal{M}_1 = \{12\}$. As $\mathcal{M}_1$ only has one edge, no two distinct edges in $\mathcal{M}_1$ share a common endpoint. So $\mathcal{M}_1$ is a matching. By similar argument, $\mathcal{M}_2 = \{23\}$ and $\mathcal{M}_3 = \{13\}$ are also matchings.
- Let $\mathcal{M}_4 = \emptyset$. As $\mathcal{M}_4$ does not contain any edges, it vacuously satisfies the definition of a matching.

It is natural to ask whether we can add any edges to a matching. To this end, we again consider $\mathcal{M}_1 = \{12\}$. Observe that as 23 shares the endpoint 2 with 12, the set $\{12, 23\}$ is **not** a matching. By similar argument, $\{12, 13\}$ and $\{13, 23\}$ are also **not** matchings. Clearly, $\{12, 23, 13\}$ is also not a matching. So the four matchings listed above are the only matchings of C_3.

(Recommended) Problem 6. Consider the cycle graph C_4 pictured below.

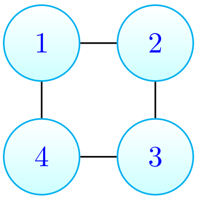

Do the following.

(a) How many matchings of size 0 does C_4 have? List them.

(b) How many matchings of size 1 does C_4 have? List them.

(c) How many matchings of size 2 does C_4 have? List them.

(d) How many matchings of size 3 does C_4 have? List them.

(e) How many matchings of size 4 does C_4 have? List them.

Remark 3. It is of particular interest in algorithm design (as well as in graph theory, and application areas such as algorithmic economics) to find a matching in the graph G with largest cardinality. We call such matchings *maximum*, and the size of such a matching is called the *matching number* of G. We denote the matching number of G as $\nu(G)$.

(Recommended) Problem 7. Determine $\nu(G)$ for the following graphs.

(a) $G = C_3$.

(b) $G = C_{2n+1}$ $(n \geq 1)$.

(c) $G = C_4$.

(d) $G = C_{2n}$ $(n \geq 2)$.

(e) $G = K_{2,3}$. [Refer back to Page 3, Example 6 of my Graph Theory notes on Canvas for the definition of the Complete Bipartite Graph.]

(f) $G = K_{m,n}$.

3.2 Using Network Flows to Find Maximum Matchings in Bipartite Graphs

We now describe how to use maximum flow algorithms to find maximum-size matchings in bipartite graphs. Suppose we are given a bipartite graph $G(L\dot{\cup}R, E)$. We proceed as follows.

- First construct a flow network $\mathcal{N}(V, E, \text{cap})$ in the following manner:
  - $V(\mathcal{N}) = V(G) \cup \{s, t\}$. That is, $\mathcal{N}$ has precisely the same vertices as G, and then we add a source vertex s and a sink vertex t.
  - For each edge $\{i, j\} \in E(G)$ where $i \in L$ and $j \in R$, we add a directed edge (i, j) with capacity 1 in $\mathcal{N}$.
  - For each $u \in L$, we add a directed edge (s, u) with capacity 1 in $\mathcal{N}$.
  - For each $v \in R$, we add a directed edge (v, t) with capacity 1 in $\mathcal{N}$.
- Next, we compute the maximum flow f^* that we can push from $s \to t$ in $\mathcal{N}$.
- We construct the matching $\mathcal{M} \subset E(G)$ corresponding to f^* by including the edges $\{i, j\} \in E(G)$ if and only if $f^*((i, j)) = 1$. That is, the saturated edges from $L \to R$ in $\mathcal{N}$ are precisely the edges in the corresponding maximum matching of G.

Example 4. We now work through an example of this construction. Let G be the graph pictured below.

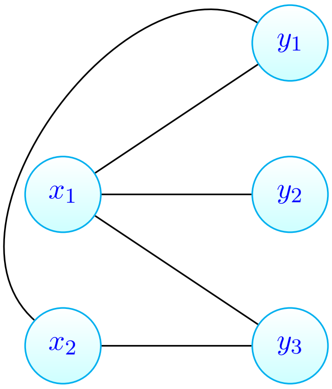

- **Step 1:** We begin by constructing the corresponding flow network $\mathcal{N}$:

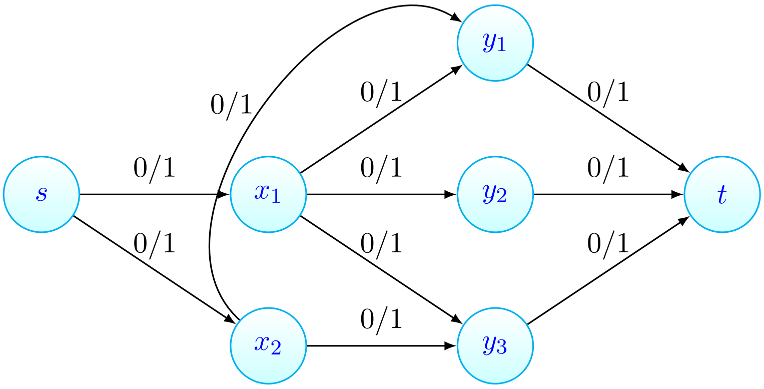

- **Step 2:** We now find a maximum flow. **We stress that the selection of our flow-augmenting paths below is a choice. There are other valid choices; that is, our selection is in a sense arbitrary.**
  - Observe that we have a flow-augmenting path $s \to x_1 \to y_2 \to t$, allowing us to push 1 unit of flow from $s \to t$. This yields the updated flow network:

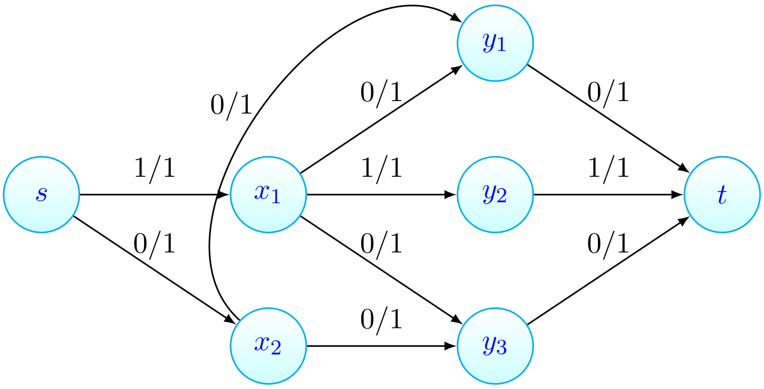

  - Observe that we have a flow-augmenting path $s \to x_2 \to y_1 \to t$, allowing us to push 1 unit of flow from $s \to t$. This yields the updated flow network:

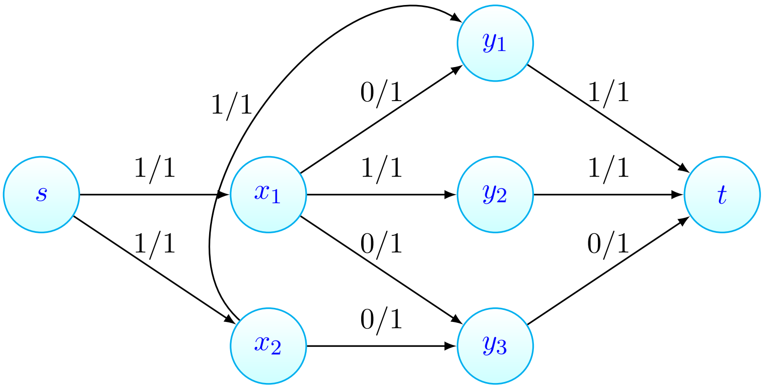

  There are no more flow-augmenting paths. So we now have a maximum flow on $\mathcal{N}$.

- **Step 3:** We now construct the matching on G based on our maximum flow on $\mathcal{N}$. Since the edges (x_1, y_2) and (x_1, y_2) are saturated on $\mathcal{N}$, our corresponding maximum matching $\mathcal{M} = \{\{x_1, y_2\}, \{x_2, y_1\}\}$.

Remark 4. We note that other maximum flows on $\mathcal{N}$ would yield different maximum matchings. We leave it as an exercise for the reader to find the other maximum matchings of G.

(Recommended) Problem 8. Let $G = K_{3,4}$. The goal of this problem is to use the network flow reduction to find a maximum matching of G.

(a) Construct the corresponding flow network $\mathcal{N}$.

(b) Find a maximum flow on $\mathcal{N}$. Clearly identify each flow-augmenting path and the amount of flow you can push through said path.

(c) Use the maximum flow from part (b) to determine the corresponding maximum matching on G.