Communication Networks (0368-3030) / Spring 2011

The Blavatnik School of Computer Science, Tel-Aviv University

Allon Wagner

Max-Min Fairness

Recommended references:

- Computer Networks - Performance and Quality of Service \ Ivan Marsic (available online)
- An Engineering Approach to Computer Networking \ S.Keshav

Fair allocation

- Unfair allocation is the result of uncontrolled competition between users for the same resource.
- Fairness and user prioritization are complementary network demands.
    - Maintain fair allocation while supporting “premium users”
- There’s often a tradeoff between fair and optimality

Fairness vs. optimality

- A simple network with 3 nodes and 2 links with capacity 1 each.
- 3 flows with demand 1: A→B, A→C, B→C
- Optimal throughput:

$$\gamma_{AB} = \gamma_{BC} = 1$$
$$\gamma_{AC} = 0$$

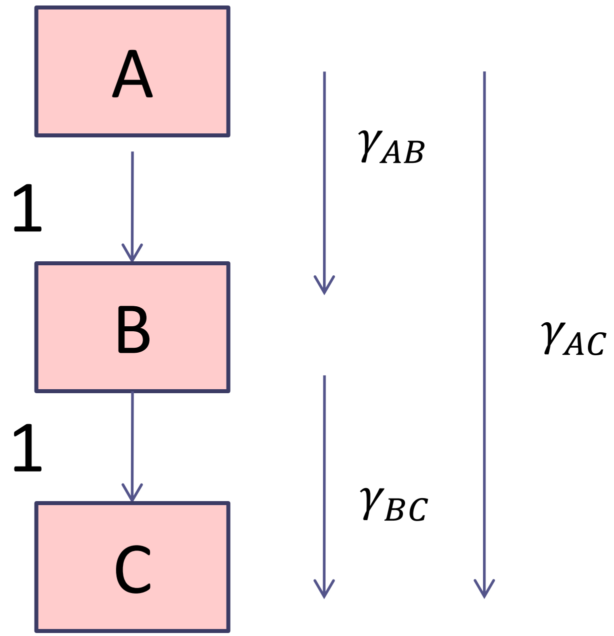

But what is "fair allocation"?

- All users get the same throughput?
  - $\gamma_{AD} = \gamma_{BC} = \gamma_{CD} = 0.5$
  - But then γ_{CD} actually gets more resources
- All users get the same amount of network resources?
  - $\gamma_{BC} = \gamma_{CD} = 0.75$
  - $\gamma_{AD} = 0.25$
    (total on 3 links is 0.75)

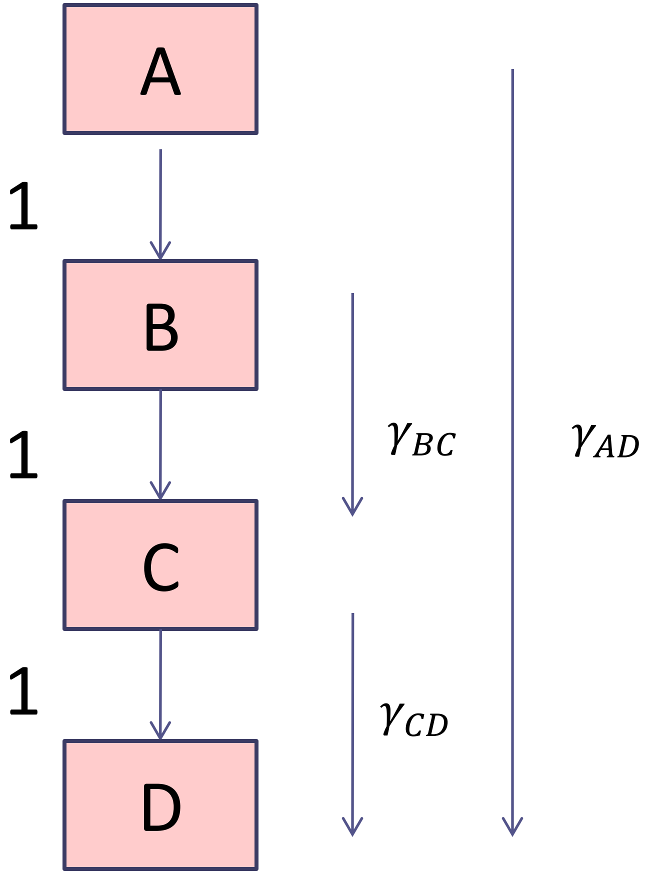

But what is “fair allocation”?

- Each user gets a throughput proportional to how much it damages other users?
  - $\gamma_{BC} = \gamma_{CD} = \frac{2}{3}$
  - $\gamma_{AD} = \frac{1}{3}$ (total interference on the two links is $\frac{2}{3}$)

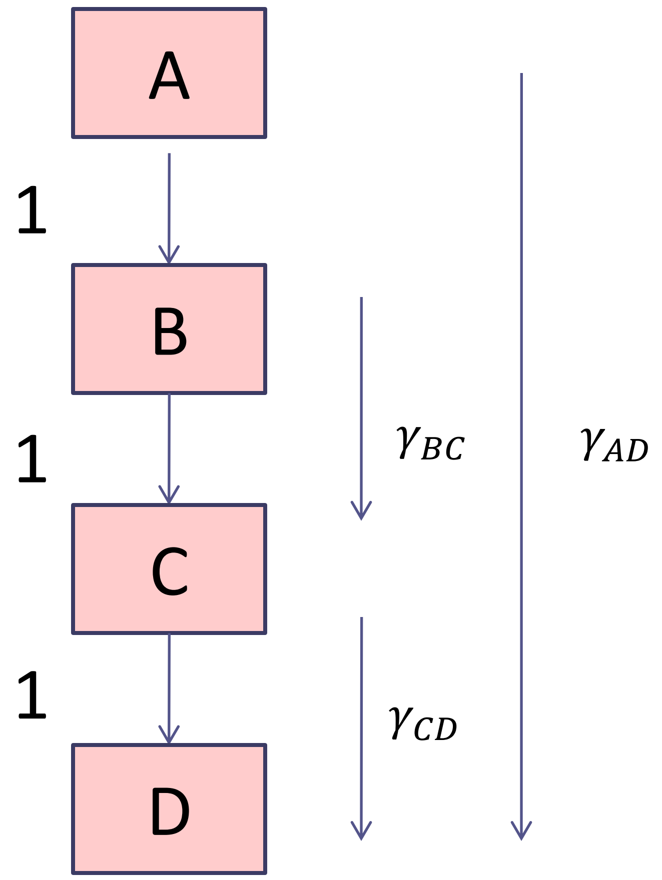

An example

- Assume we have a simple resource with capacity 30
  - No network considerations for the time being
  - For instance: a single link, a queue in a router etc.
- User demands are $A = 20, B = 20, C = 20$
  - Fair allocation is trivial
- But what if the demands are:

$$A = 4, B = 20, C = 20$$

Max-Min Fairness

- A sharing technique commonly used in practice
- The idea: Users who need less than what they are entitled to get their full demand. The excess is evenly distributed among the “heavy” users.
- Formally:
  - Resources are allocated in order of increasing demand
  - No source obtains a resource share larger than its demand
  - Sources with unsatisfied demands obtain an equal share of the resource

An example

- Compute the max-min fair allocation for a set of four sources with demands 2, 2.6, 4, 5 when the resource has capacity 10.

1. Give 2 to every user. 1’s demand is fulfilled, we have 2 excess capacity to distribute.
2. Give 2.6 to users 2,3,4. 2’s demand is fulfilled, we have excess 0.2 to distribute.
3. Final allocation is: 2, 2.6, 2.7, 2.7

Generalization to a graph

- We have a directed graph $G = (V, E)$, with capacy c_e for each edge e.
- We have a set of ongoing calls (flows). Each call i has demand r_i and a (fixed) path p_i.
- Algorithm:
  - Increase all flows equally until one link fills.
  - Fix the rate of the bottleneck flows.
  - Continue with the unfixed flows.

An example

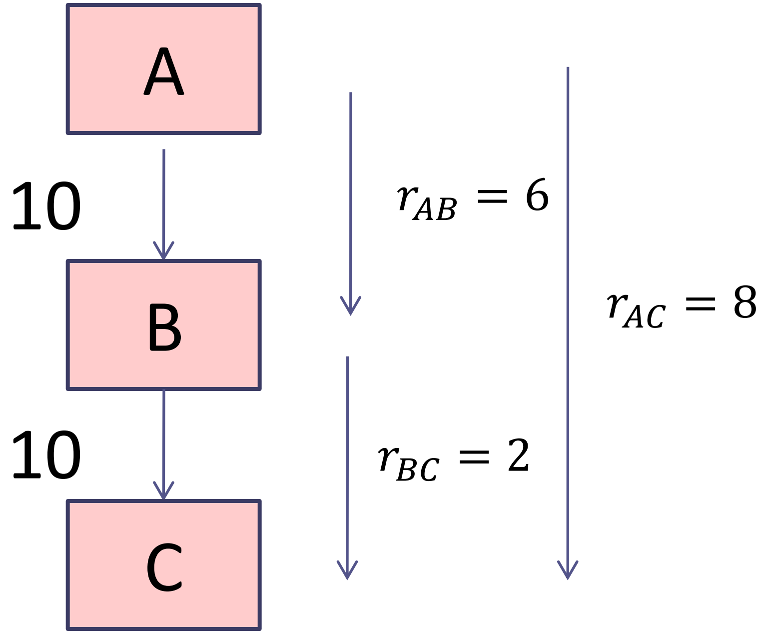

An example

- Give 2 to all flows
- Demand BC is fulfilled.

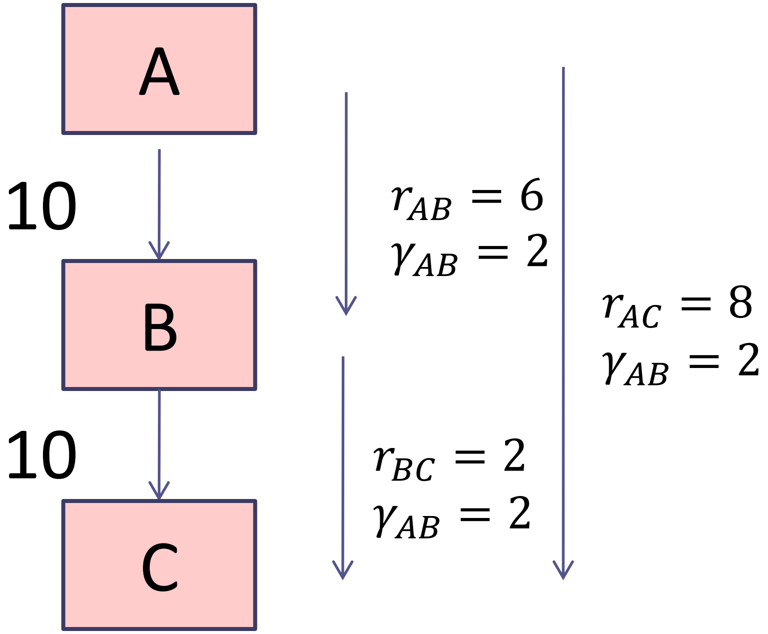

An example

- Continue increasing flows equally. When AB and AC get each 5, edge (A,B) is saturated.