

Final Review

CSE 461 – Autumn 2015

Panji Wisesa

With thanks to the TAs of CSE 461 Autumn 2014

Routing

- More complex than spanning tree used by the link layer
- Source/Sink Trees
- Dijkstra
- Distance Vector
- Link State

BGP

- Hosts grouped into Autonomous Systems (AS)
- Within an AS, shortest path routing
- Between ASes, routes governed by real world considerations such as costs so shortest path is not always taken.

BGP Rules for Traversing

- Go through your customers first
- Go through paths announced by your peers second
- Go through your provider last
- Take the path with the lowest number of hops

BGP Rules for Announcement

- Always announce to your customers all of the nodes you can reach
- Don't announce any paths through your peers to your provider
- Don't announce provider links to any other provider
- Announce only customer links to your providers
- Also Announce only customer links to your peers
- If there are 2 paths where one is through a provider and one through peer then only announce peer paths to your customers
- Paths: Customer > Peer > Provider, then by how many hops

Transport Layer

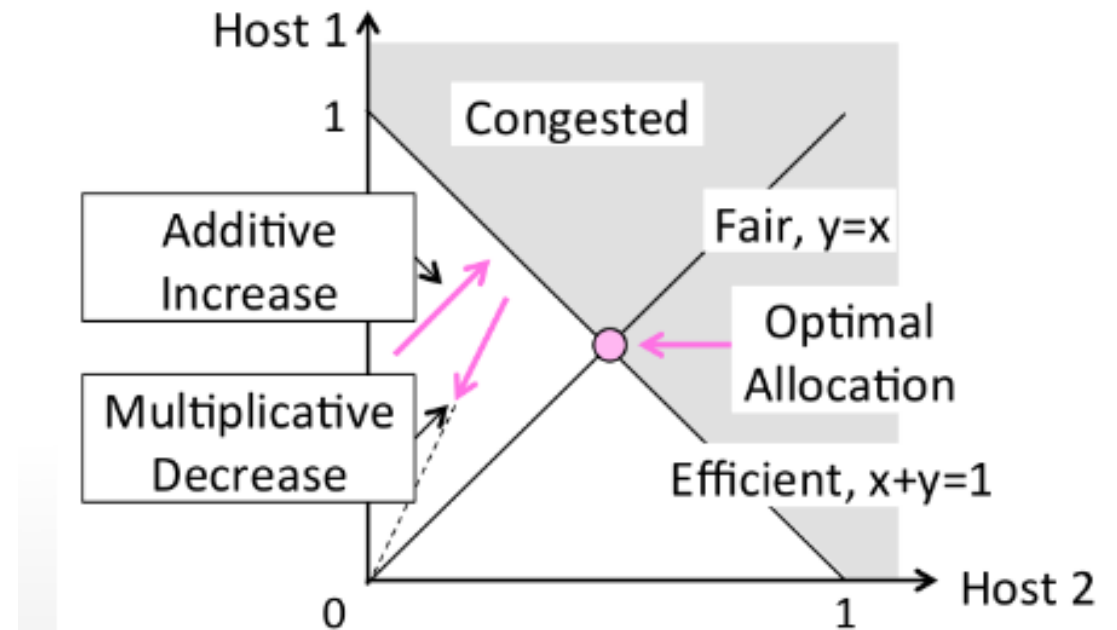

TCP (Streams)	UDP (Datagrams)
Connections	Datagrams
Bytes are delivered once, reliably, and in order	Messages may be lost, reordered, duplicated
Arbitrary length content	Limited message size
Flow control matches sender to receiver	Can send regardless of receiver state
Congestion control matches sender to network	Can send regardless of network state

TCP Sliding Window

- Allow W packets to be outstanding
- Can send W packets per RTT
 - The first packet in W must go to the receiver and the ack must come back before the window slides and another packet may be sent
- Variants:
 - Go Back N
 - Selective Repeat
- Fast retransmit

TCP Congestion

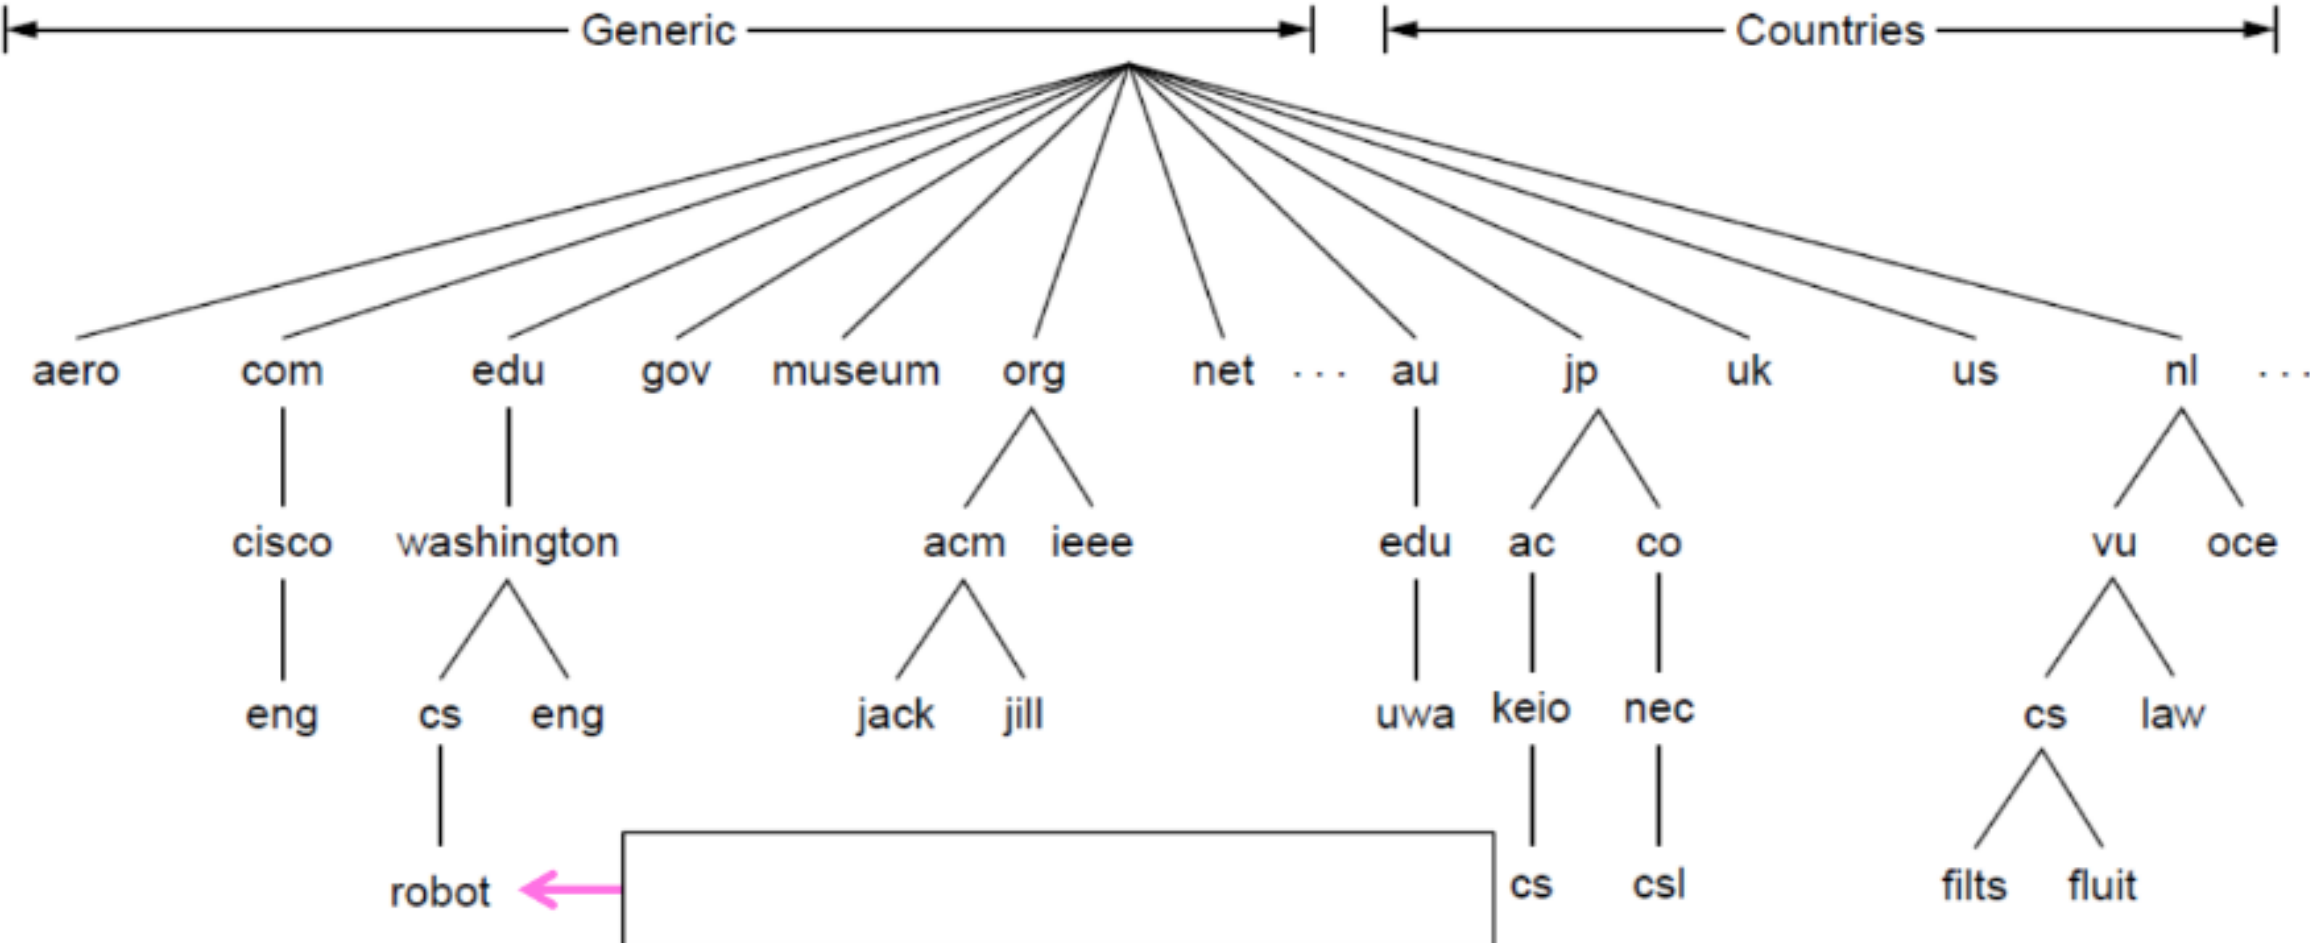
- Max-Min Fairness
 - Flows bottlenecked on a link get an equal share of it.
- Hosts cooperate using AIMD
 - Additive Increase, Multiplicative Decrease



TCP Slow Start

- Exponentially increase
 - $W *= 2$ each RTT
 - misnamed?
- When packet loss occurs, start over, but switch to AI before packet loss occurs
 - $W++$ each RTT

DNS



HTTP

- “simple” text-based protocol
- Allows fetching of content, and much more
- Returns status codes with each request
- Short-lived TCP connections give poor performance
 - So, re-use TCP connections (persistent connections)
 - Or even multiplex (SPDY)

Thanks for coming!

Good luck on the finals!