CSE 550: Systems for all

Au 2022

Ratul Mahajan

Transport protocol challenges

Flow control

- Do not send faster than the receiver’s consumption capacity

Congestion control

- Do not send faster than the networks delivery capacity

Congestion control challenges

Network capacity is unknown

- Individual links can vary from less then 100 Kbps to more than 100 Gbps

Network capacity varies with time

- E.g., distance from cellular base stations or WiFi access points

Network capacity is shared

- Don’t know how many other users are there and their needs

What if you send too slow?

What if you send too fast?

Congestion Collapse in the 1980s

- Early TCP used fixed size window (e.g., 8 packets)
  - Initially fine for reliability
- But something happened as the network grew
  - Links stayed busy but transfer rates fell by orders of magnitude!

Nature of Congestion

- Routers/switches have internal buffering

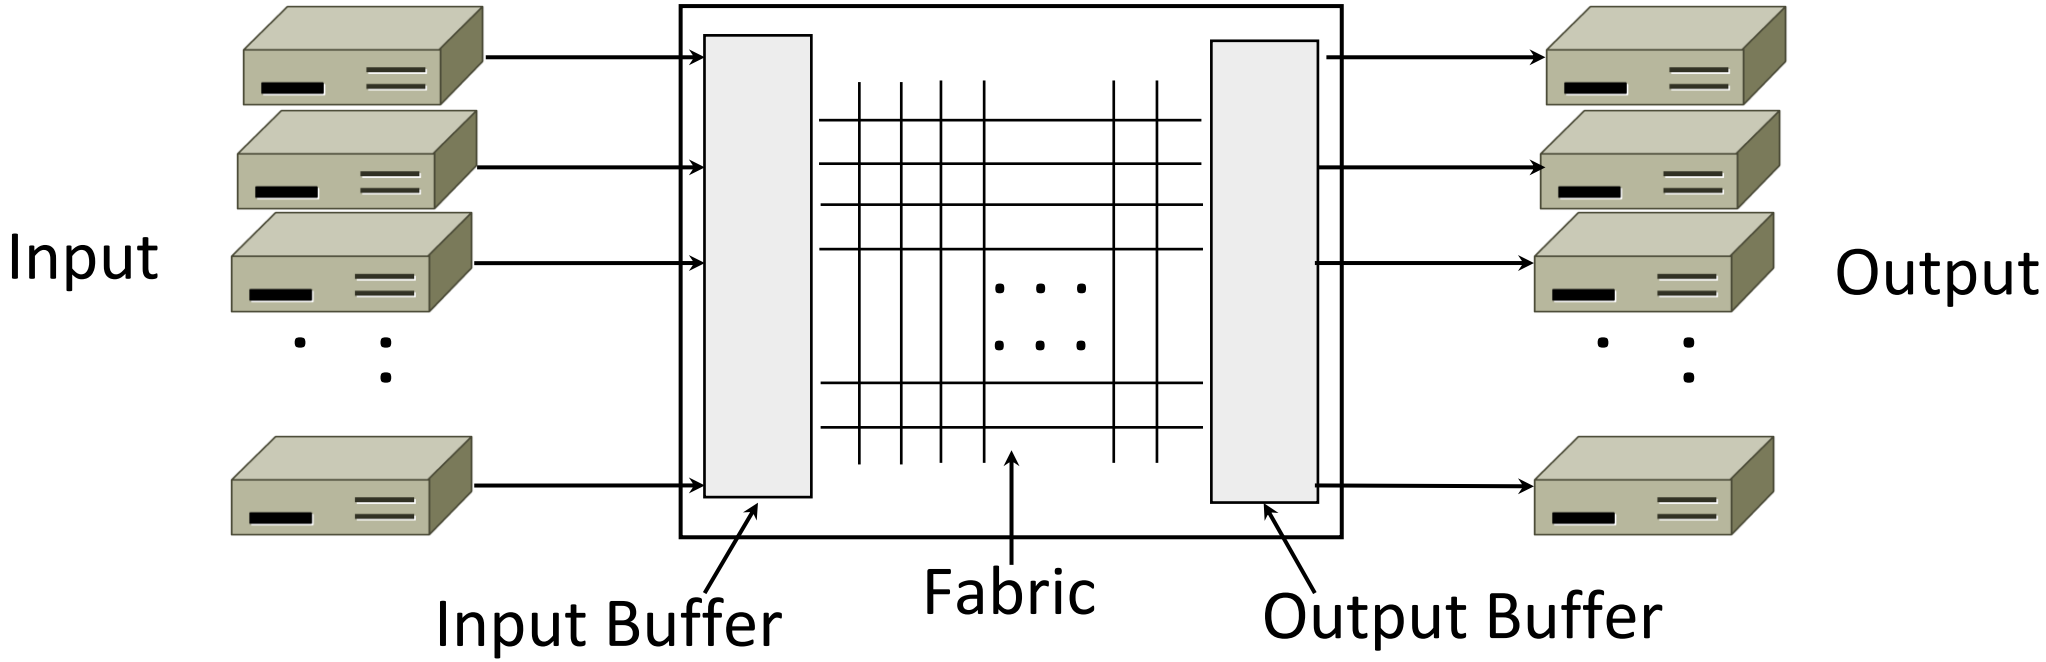

Nature of Congestion (2)

- Simplified view of per port output queues
  - Typically FIFO (First In First Out), discard when full

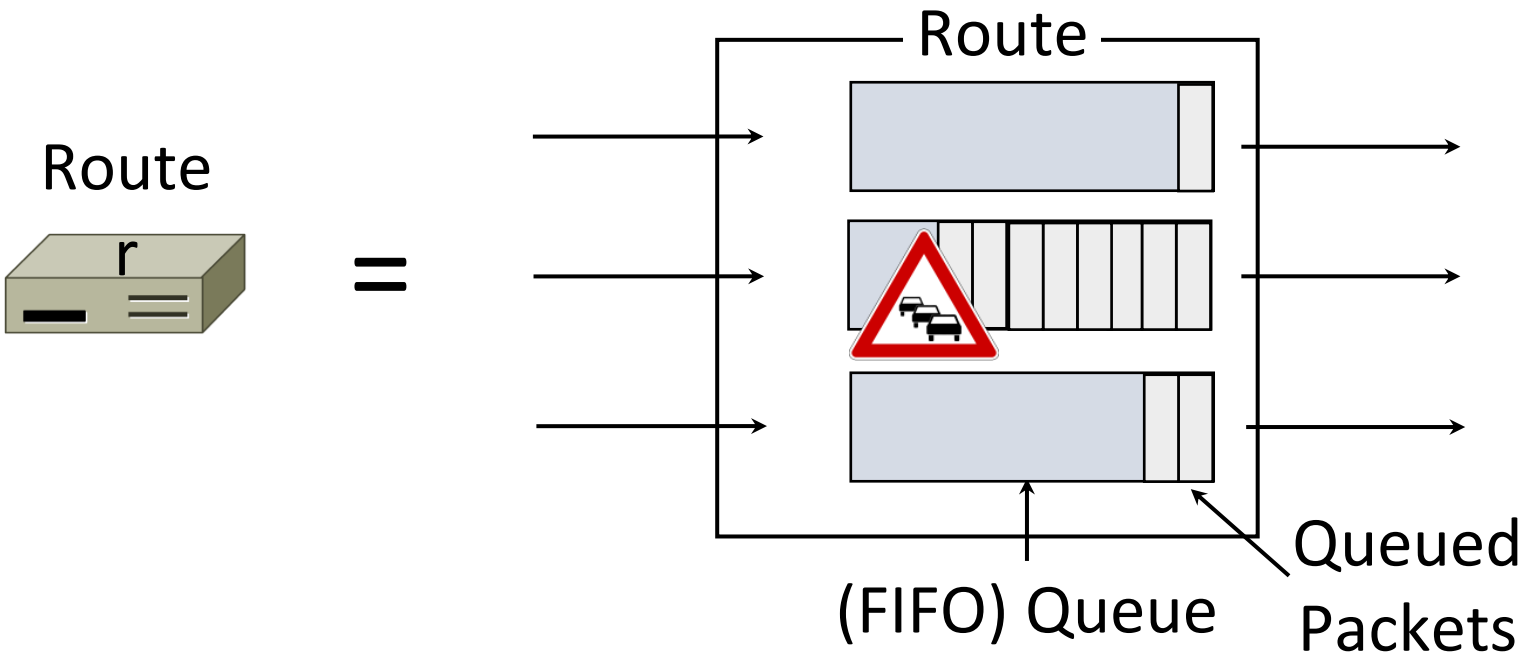

Nature of Congestion (3)

Queues help by absorbing bursts when input > output rate

But if input > output rate persistently, queue will overflow

- This is congestion

Congestion is a function of the traffic patterns – can occur even if every link has the same capacity

Effects of Congestion

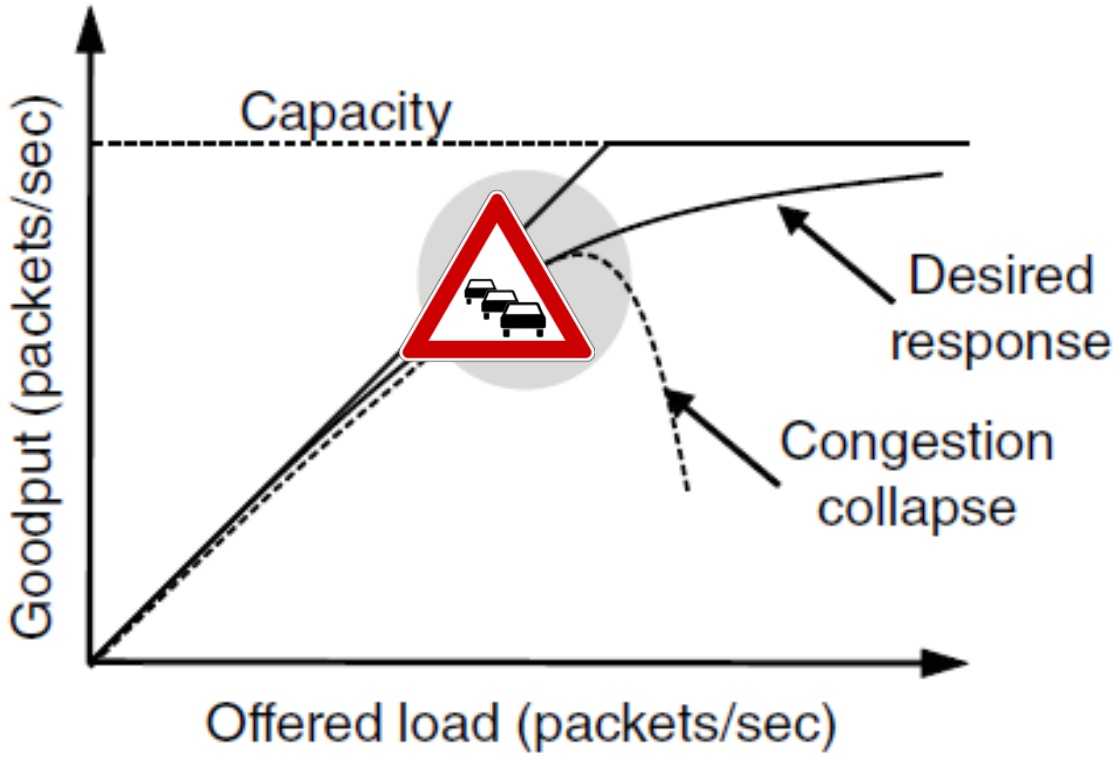

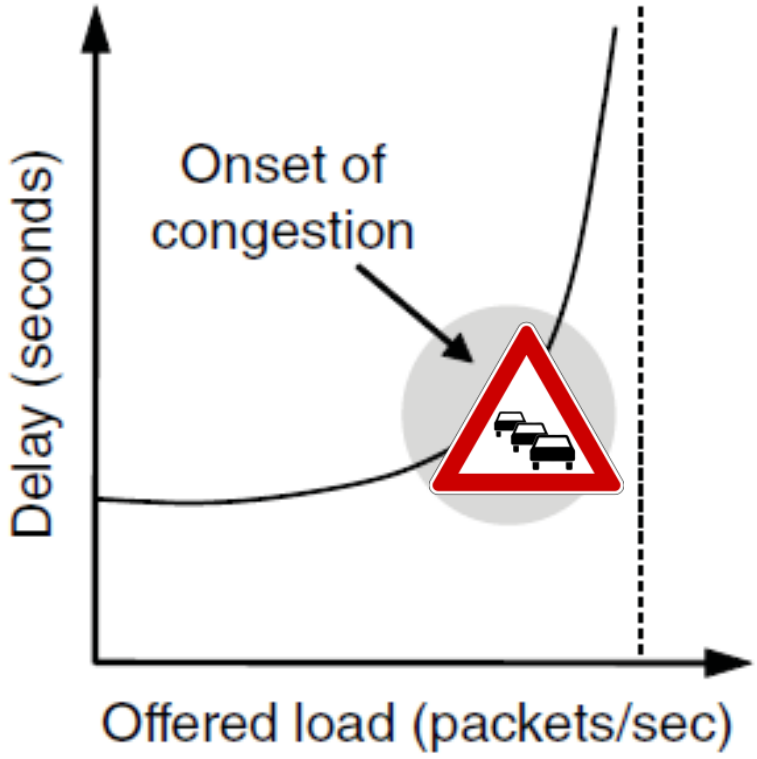

Effects of Congestion (2)

As offered load rises, congestion occurs as queues begin to fill:

- Delay and loss rise sharply with load
- Throughput < load (due to loss)
- Goodput << throughput (due to spurious retransmissions)

None of the above is good!

- Want network performance just before congestion

Congestion control design space: Congestion indicator

Several possible signals, with different pros/cons

Signal	Example Protocol	Pros / Cons
Packet loss	TCP Cubic TCP (Linux)	Hard to get wrong Hear about congestion late Other events can cause loss
Packet delay	BBR (Google)	Hear about congestion early Need to infer congestion
Router indication	ECN, DCTCP	Hear about congestion early Requires router support

Congestion control design space: Control mechanism

Two main controls, but hybrids also possible

Control	Example Protocol	Pros/Cons
Window based	TCP Cubic TCP (Linux) DCTCP	Self-clocking Can be Jittery
Rate based	BBR (Google)	Smoother Tricky estimation

Over to Oliver and Wenxuan