CSE561 – Switches

David Wetherall
djw@cs.washington.edu

Switches

- Focus:
  - Devices that connect individual links

- Switch internals
- "Plug and play" LAN switches
- IP versus MPLS

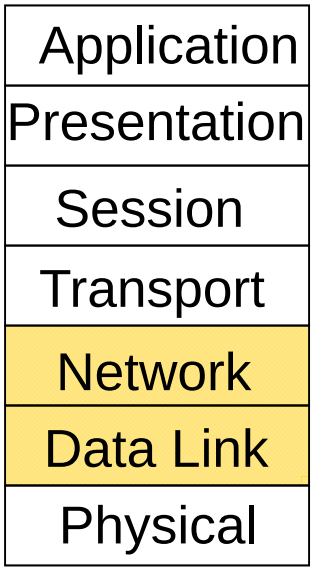

Terminology

- Bridge
  - Old fashioned name for a LAN switch, e.g., Ethernet switch
  - Works at the link (Ethernet) layer
- Router
  - Switch that works at the network (IP) layer
- Switch
  - Generic term for a low-level interconnection device
- Gateway
  - Generic term for a high-level interconnection device

Examples

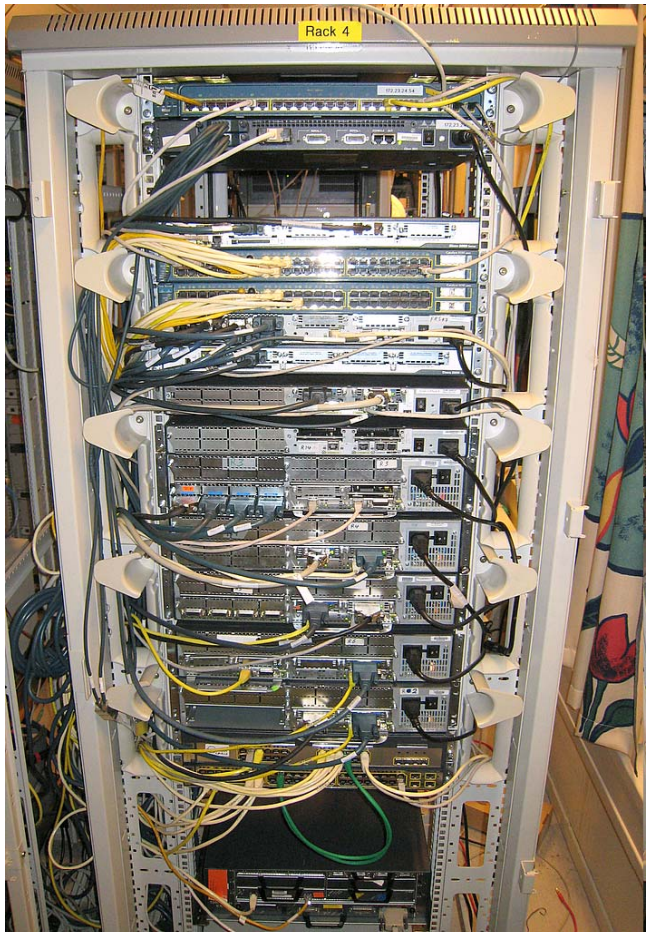

What's in a Router?

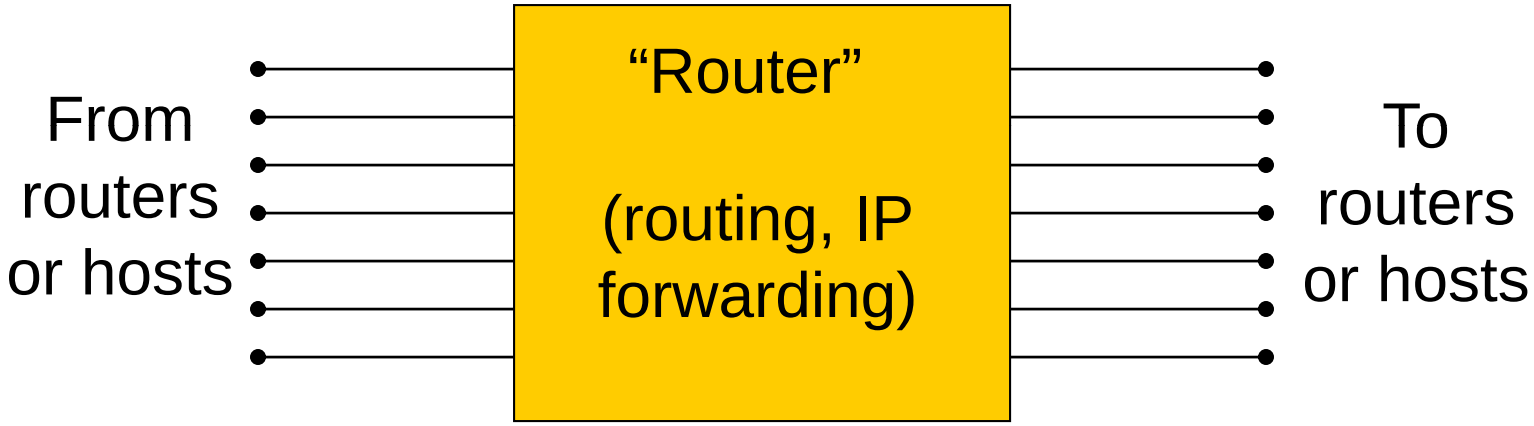

- By convention, draw input ports on left, output on right. (But in reality a single physical port handles both directions.)

Router Model: “FIFO with Tail Drop”

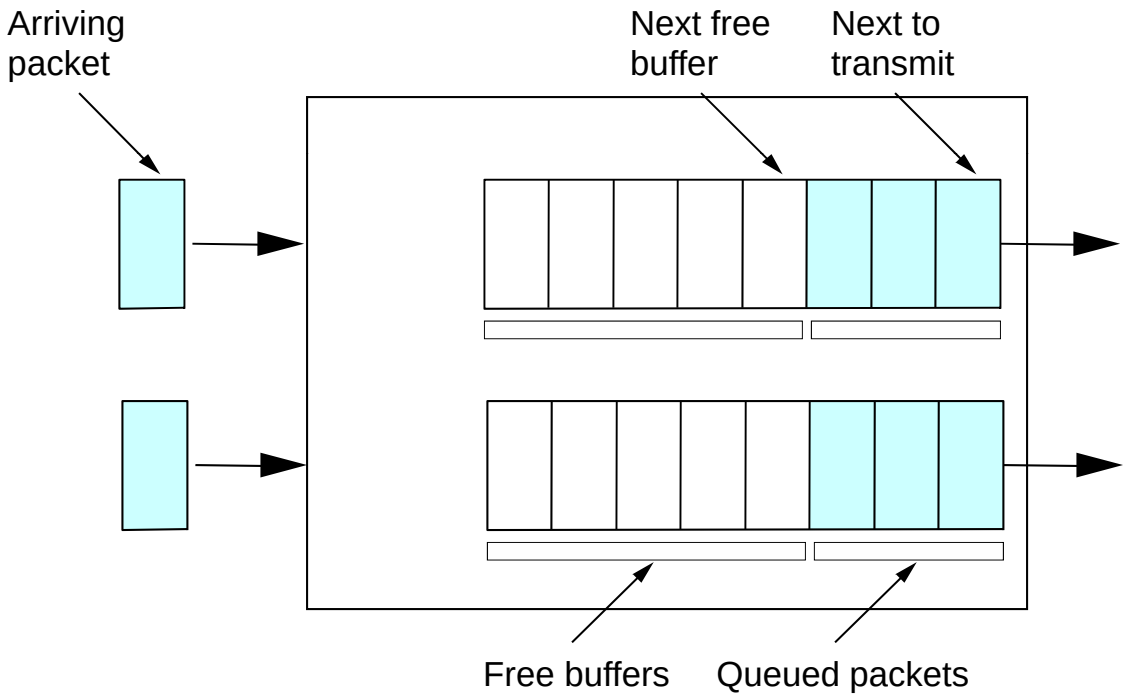

Example for 2 inputs/outputs

Model of a Router

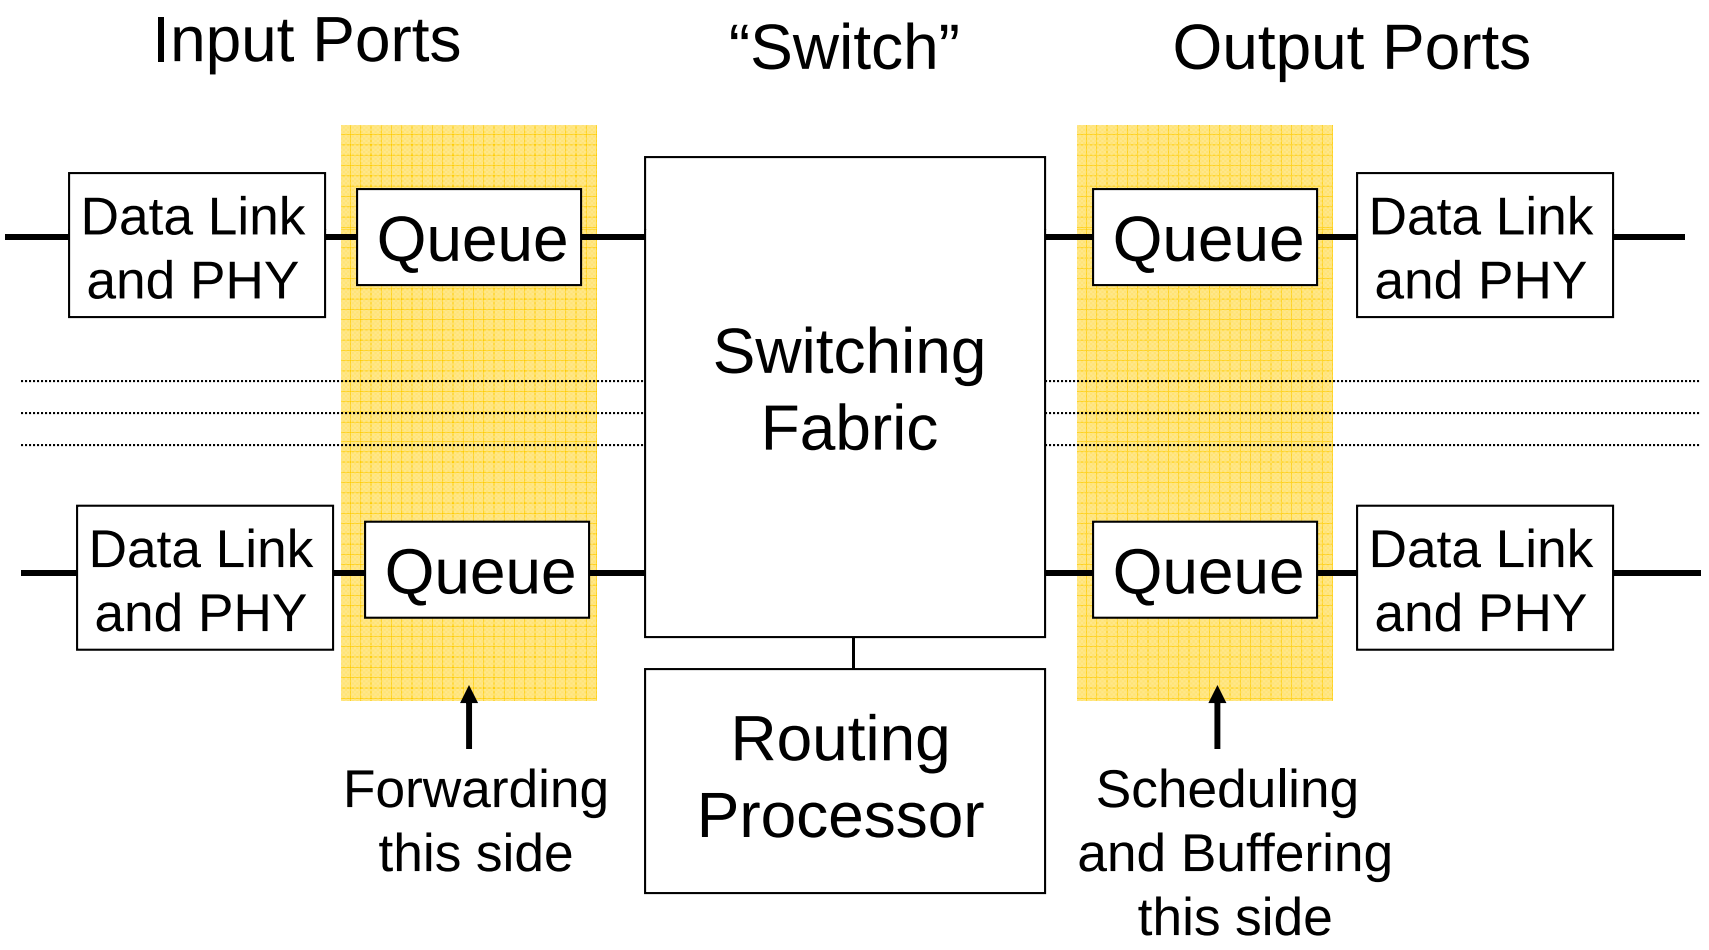

Crossbar switch

- On/off setting of intersection points control connections from inputs to outputs

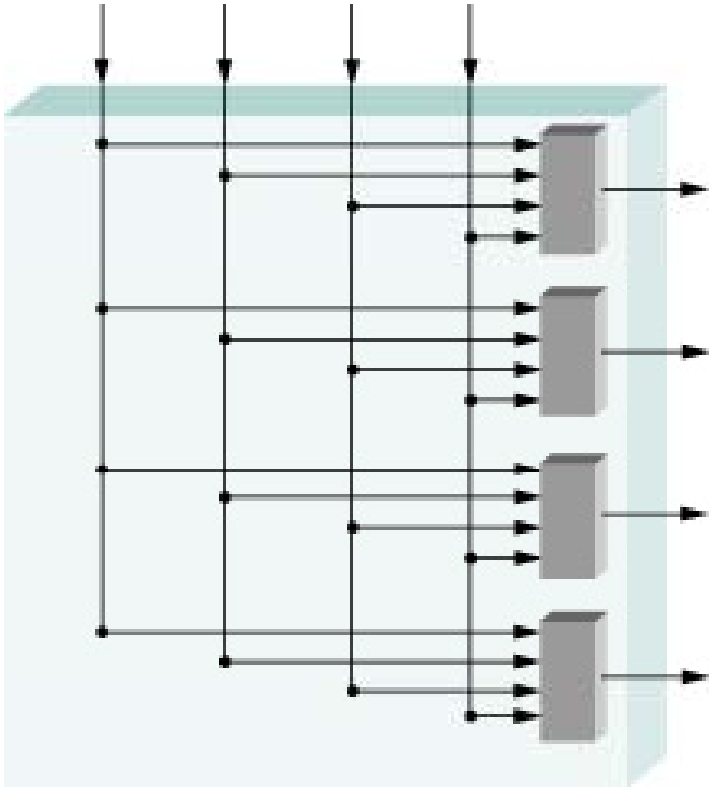

LAN Switches / Bridges

- Plug and play using two algorithms
- 1. Backward learning
- 2. Spanning tree computation
- Frames are forwarded using destination MAC address

- Book uses “classic Ethernet” but today bridges work with switched Ethernet – they are just multi-port, high-performance bridges.

Radia Perlman says …

Algorhyme

I think that I shall never see
A graph more lovely than a tree.

A tree whose crucial property
Is loop-free connectivity.

A tree which must be sure to span
So packets can reach every LAN.

First the Root must be selected.
By ID it is elected.

Least cost paths from Root are traced.
In the tree these paths are placed.

A mesh is made by folks like me
Then bridges find a spanning tree.

From:
"An Algorithm for Distributed Computation of a Spanning Tree in an Extended LAN",
R. Perlman, SIGCOMM 1985.

Routers and MPLS switches

- Routers
  - Packets are forwarded using destination IP address
  - Datagram model
  - We will get to how routes are set up next lecture …
- MPLS switches
  - Packets forwarded using labels
  - Virtual circuit model
  - Circuits need to be set up just like routes

- Q: Which is the more general mechanism?

Fish topology

- Consider traffic from R1/R2 to R7.
- What are the options with IP, and with MPLS?

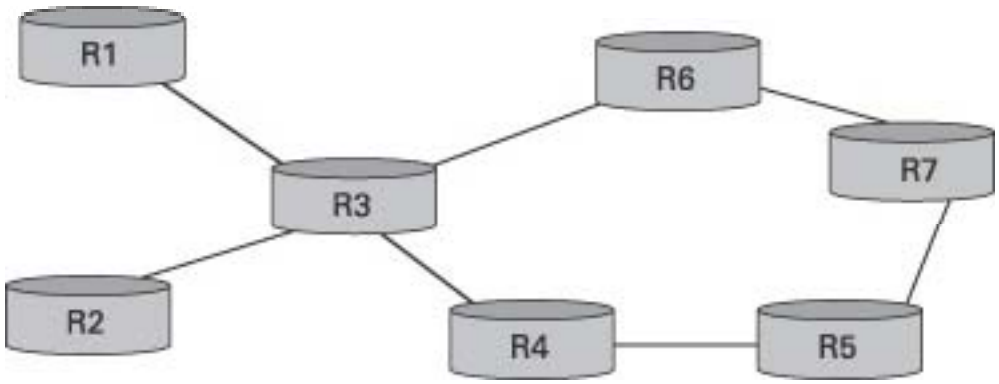

iSLIP discussion

- What is the problem and approach?
- What is another (simple) solution?
- What is head of line (HOL) blocking?
- What are virtual output (VOQ) queues?
- What properties do we want of a scheduler?
- What is an ideal performance curve?
- How is this curve affected by traffic patterns?
- What is the iSLIP algorithm?
- What is the key innovation?
- How much does it matter?