CSCE 463/612
Networks and Distributed Processing
Spring 2024

Network Layer II

Dmitri Loguinov
Texas A&M University

April 5, 2024

Chapter 4: Roadmap

4.1 Introduction

4.2 Virtual circuit and datagram networks

4.3 What's inside a router

4.4 IP: Internet Protocol

4.5 Routing algorithms

4.6 Routing in the Internet

4.7 Broadcast and multicast routing

Network Layer = IP Layer

- Transports segments from sending to receiving host
- On the sending side, encapsulates segments into datagrams
- On the receiving side, delivers segments to transport layer
- Network layer protocols in *every* host and router
- Router examines header fields in all IP datagrams passing through it

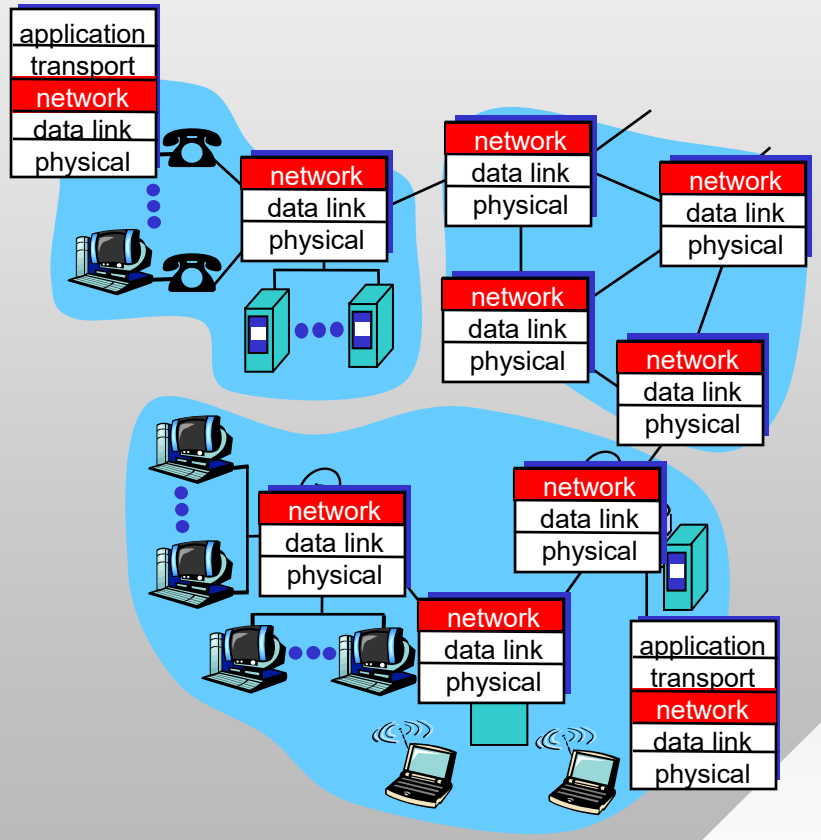

Key Network-Layer Functions

- 1) *Routing:* determine the path taken by packets from source to dest
  - Build a minimum-cost table at each router
  - Table has next-hop neighbor for each possible destination
  - Goal: send packet along the least-expensive path (e.g., in terms of hops, ISPs, or peering agreements)
- 2) *Forwarding:* move packets from a router's input port to appropriate router output port
  - Table lookup
  - Port-to-port transfer
  - Goal: efficiency

physical interface (NIC) inside router, not a TCP/UDP port!

Interplay Between Routing and Forwarding

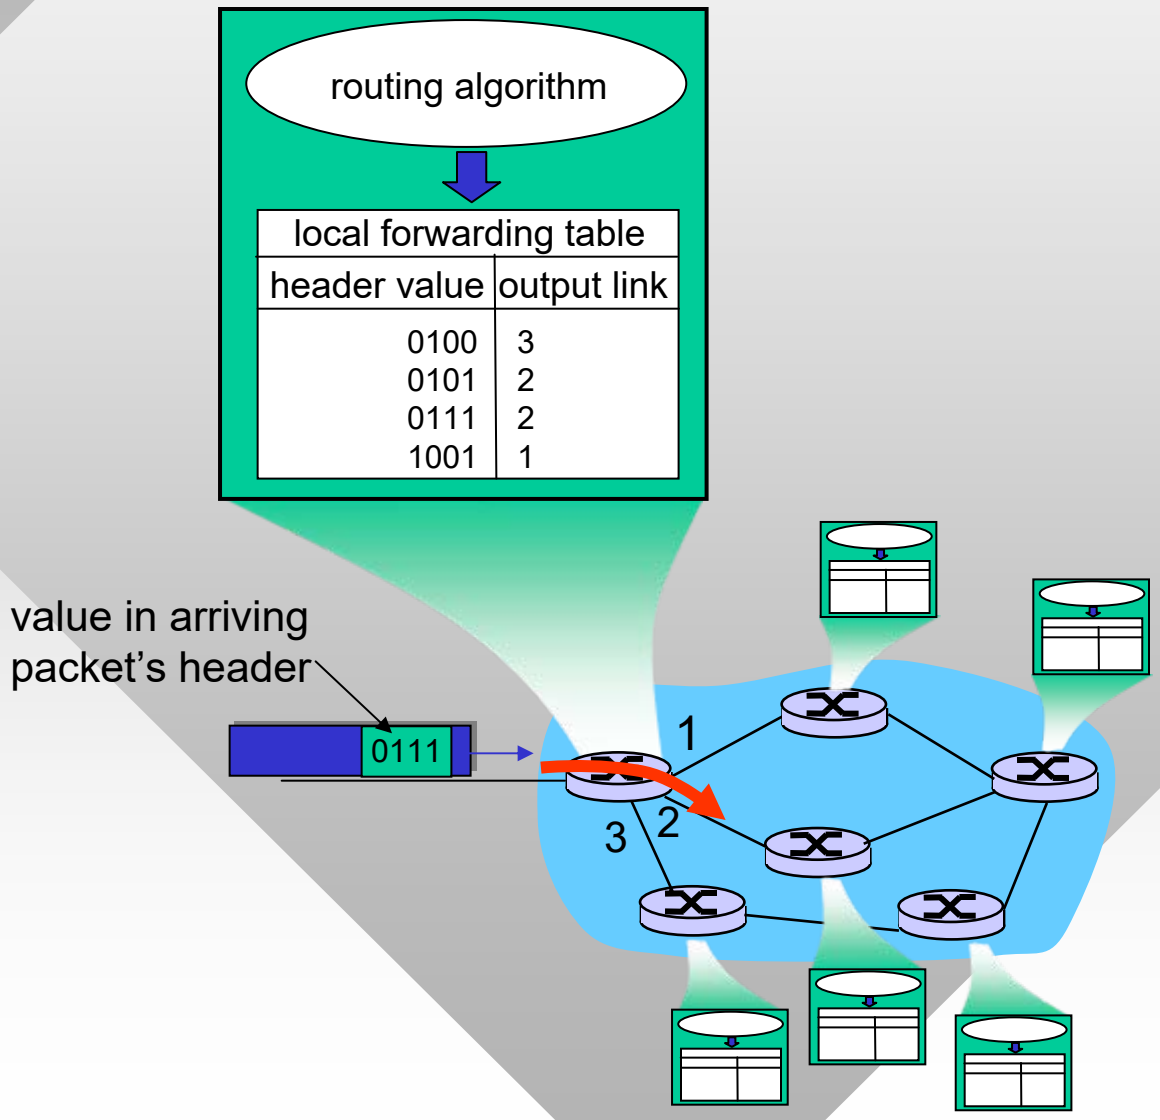

Datagram Forwarding Table

4 billion possible entries

Destination Address Range (32 bit)	Link Interface
11001000 00010111 00010000 00000000 through 11001000 00010111 00010111 11111111	0
11001000 00010111 00011000 00000000 through 11001000 00010111 00011000 11111111	1
11001000 00010111 00011000 00000000 through 11001000 00010111 00011111 11111111	2
otherwise	3

Longest Prefix Matching

Prefix Match	Link Interface
11001000 00010111 00010	0
11001000 00010111 00011000	1
11001000 00010111 00011	2
otherwise	3

Examples (DA = destination address)

DA: 11001000 00010111 00010110 10100001
DA: 11001000 00010111 00011001 10101010
DA: 11001000 00010111 00011000 10101010

Which interface?

Chapter 4: Roadmap

4.1 Introduction

4.2 Virtual circuit and datagram networks

4.3 What's inside a router

4.4 IP: Internet Protocol

4.5 Routing algorithms

4.6 Routing in the Internet

4.7 Broadcast and multicast routing

Router Architecture Overview

Two key router functions:

- Run routing algorithms/protocols
- Forward datagrams from incoming to outgoing link
  - Terminology: port = interface capable of sending/receiving

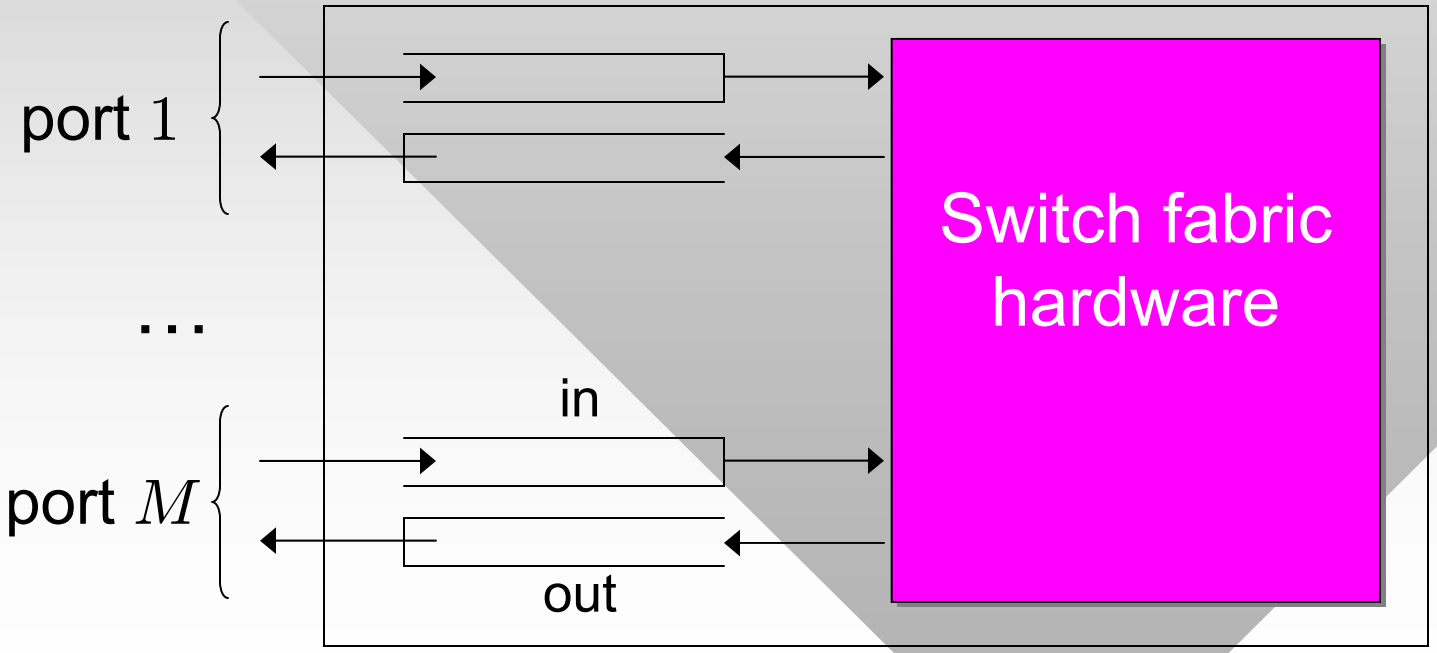

Input Port (Queue) Functions

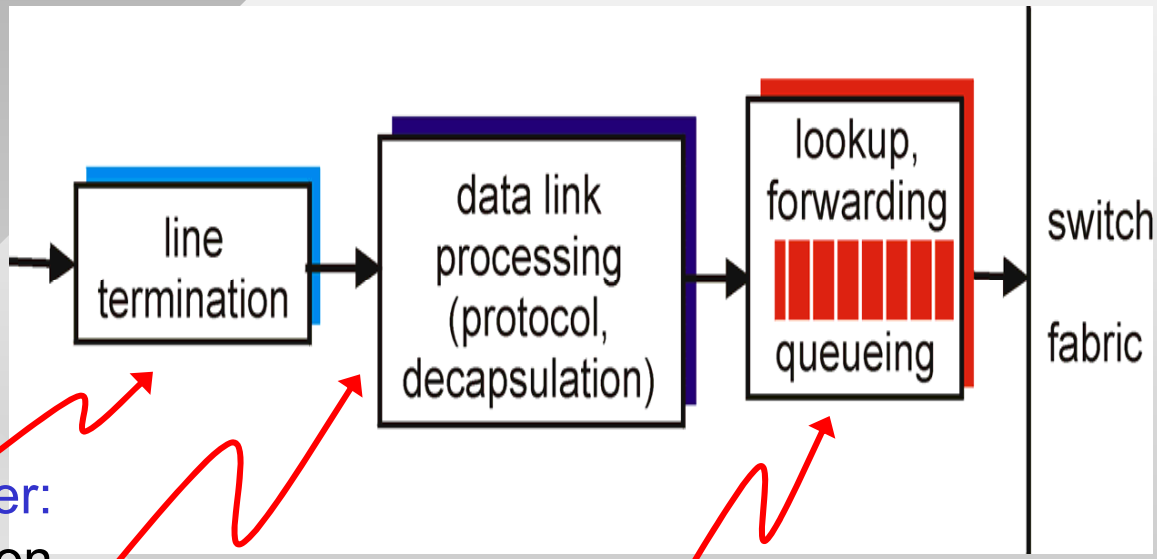

Physical layer:
bit-level reception

Data link layer
(e.g., Ethernet, ATM, Token Ring, 802.11b): see ch. 5

Decentralized switching:

- Given datagram destination, look up output port using forwarding table in input port memory
- Goal: complete input port processing at "line speed"
- Queuing: if datagrams arrive faster than forwarding rate into switch fabric

Switching Via Memory

First generation routers (1960s-mid 1980s):

- Traditional computers with switching under direct control of CPU
- Packet copied to system memory
- Speed limited by CPU, memory latency/bandwidth, and bus bandwidth (two bus crossings per datagram)
- Honeywell 316 (1969) →

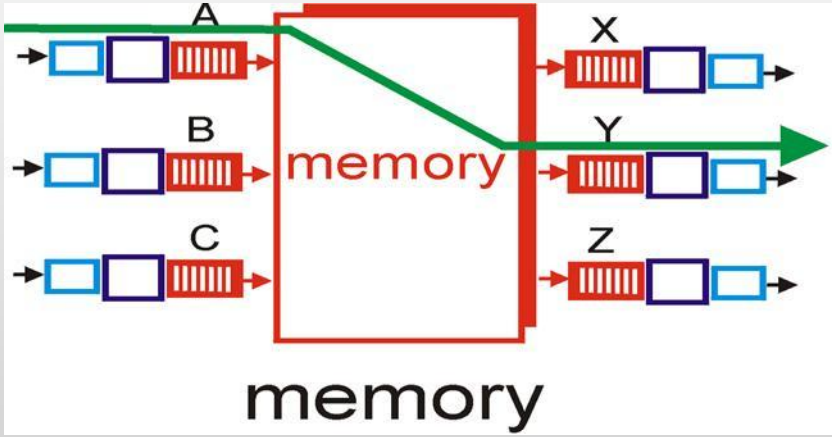

Switching Via a Bus

- Datagram from input port memory to output port memory via a shared bus
- Bus contention: switching speed limited by bus bandwidth
- 1 Gbps bus in Cisco 1900: sufficient speed for access and small enterprise networks (not ISPs)

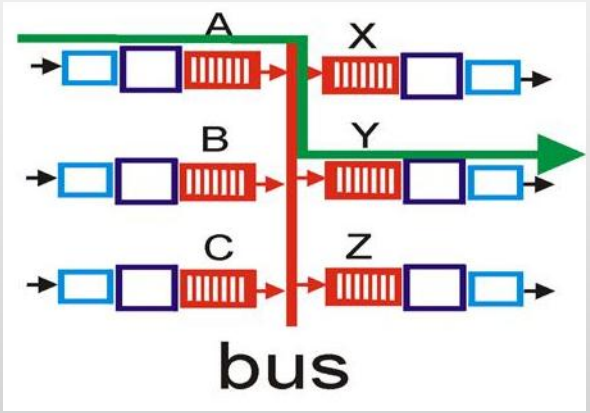

Switching Via An Interconnection Network

- Overcomes bus bandwidth limitations
  - Crossbar: packets transmitted in parallel as long as they do not occupy the same horizontal or vertical bus
- Cisco 12000 (1996): uses an interconnection network
  - CRS-X (2013): 1600 lbs, 84” rack, 7.6 KWatt, 800 Gbps/slot
  - 16 slots/rack = 12.8 Tbps
  - Up to 72 racks (922 Tbps)

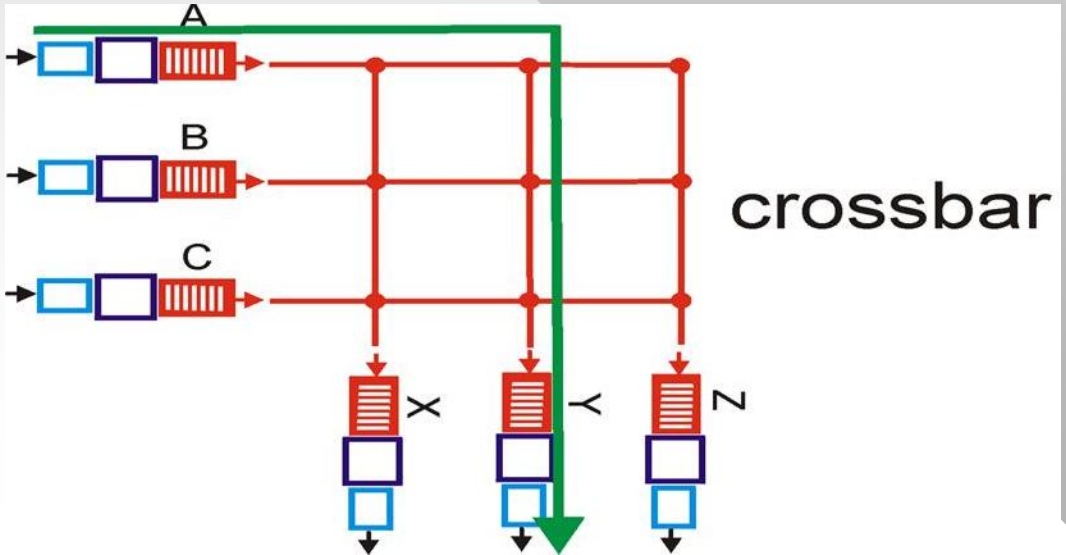

Output Ports

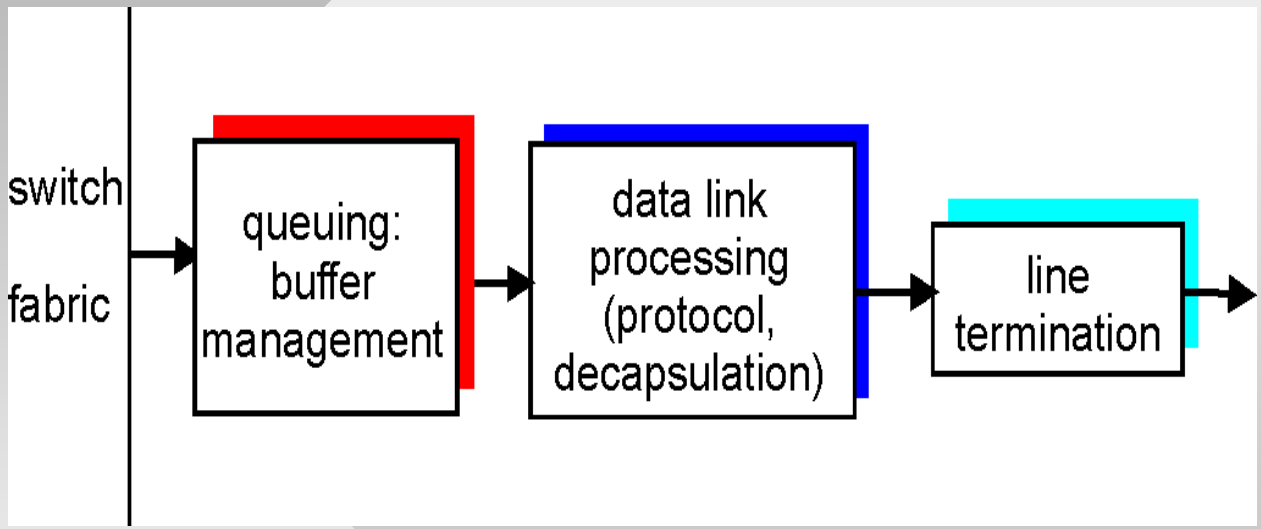

- *Buffering/queuing* required when datagrams arrive from fabric faster than the transmission rate
- *Scheduling discipline* chooses among queued datagrams for transmission
  - Customer traffic: single FIFO drop-tail queue
  - ISP traffic: multiple queues with WRR or priority queuing

Output Port Queuing

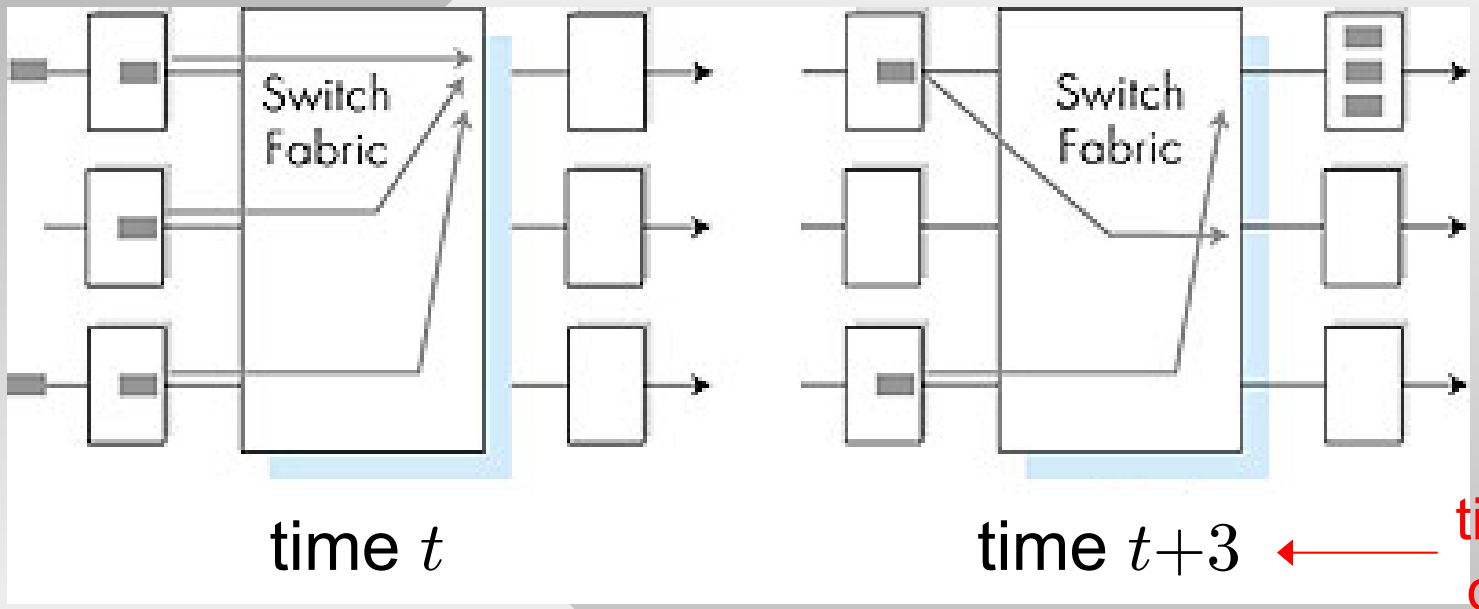

- Buffering when arrival rate via switch fabric exceeds output line speed
  - Queuing delay and loss due to output buffer overflow
- Switch fabric often faster than individual ports
  - Produces large bursts of arrivals into output queues

Input Port Queuing

- Reasons for input-port queuing:
  - Head-of-Line (HOL) blocking: queued datagram at front of queue prevents others in queue from moving forward

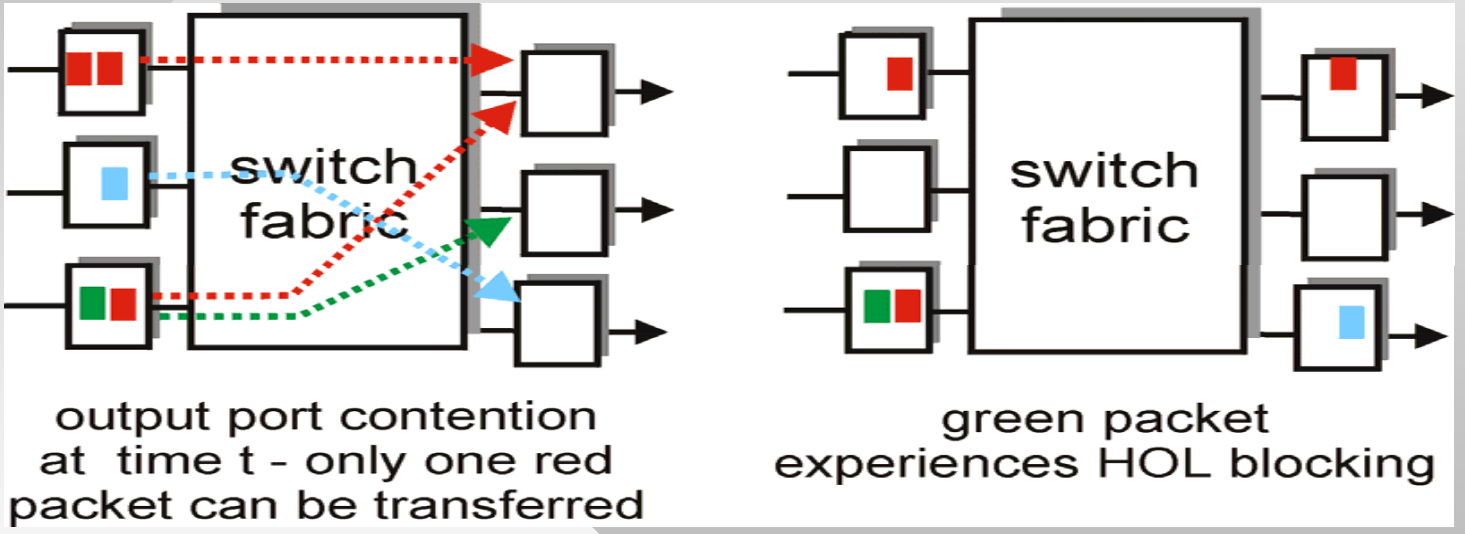

time t time $t+1$

- Queuing delay and loss due to input buffer overflow!
  - How likely is this compared to output port queuing/loss?