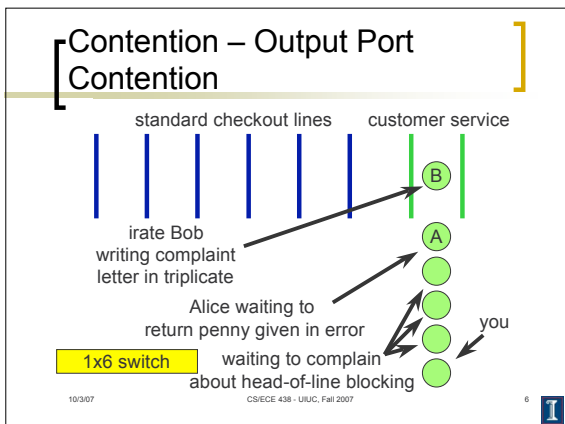
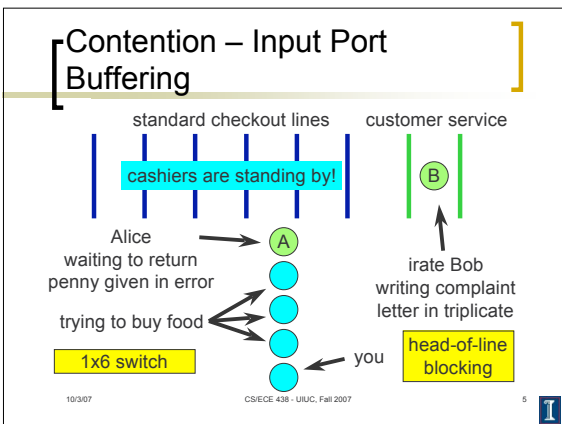
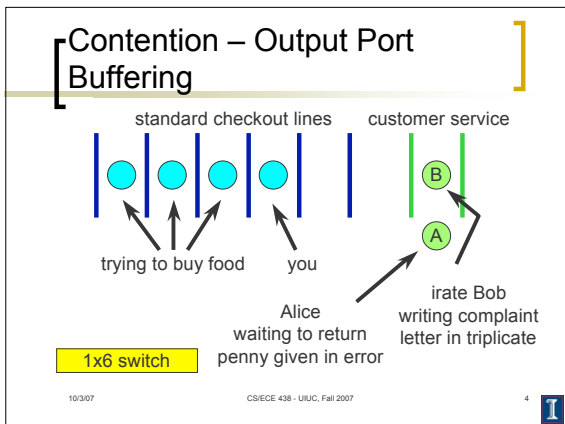
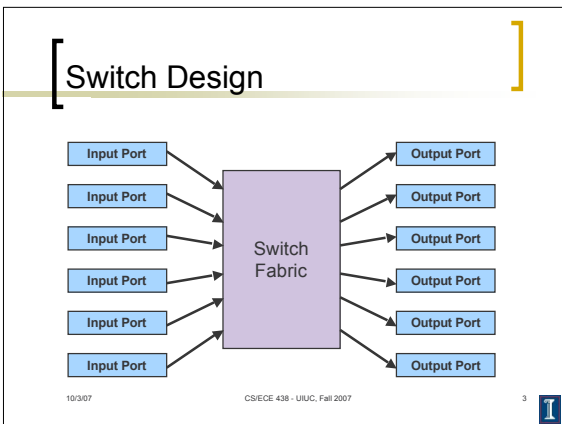


Switching Hardware

10/307 CS/EECE 438 - UIUC, Fall 2007 1

- ## Contention
- Bridges: same collision domain
 - If an output port is busy when forwarding packet from input port, cause collision
 - Switches: different collision domain
 - Use CSMA/CD before sending packet onward
 - Buffer packets
 - When output port is busy
 - When multiple packets are destined for same output port
- 10/307 CS/EECE 438 - UIUC, Fall 2007 2



Contention – Buffering Capacity

standard checkout lines customer service

buffering capacity per output is finite

1x6 switch

angry mob eagerly awaiting opportunity to address underpaid customer service representative

you

(others turned away at door)

10/3/07 7

Contention – Back Pressure

- Let the receiver tell the sender to slow down
 - Propagation delay requires that the receiver react before the buffer is full
 - Typically used in networks with small propagation delay

switch 1 switch 2

no more, please!

10/3/07 CS/ECE 438 - UIUC, Fall 2007 8

Contention – Back Pressure

- NOTE
 - Propagation delay requires that switch 2 exert backpressure at high-water mark rather than when buffer completely full. Backpressure is thus typically only used in networks with small propagation delays (e.g., switch fabrics).

Switch stop stop stop Switch

8 9 9 6 5 4 3 2 1

Discard: 9

10/3/07 CS/ECE 438 - UIUC, Fall 2007 9

Switch Design Goals

- Throughput
 - Number of packets a switch can forward per second
- Scalability
 - How many input/output ports can it connect
- Cost
 - Per port monetary costs

10/3/07 CS/ECE 438 - UIUC, Fall 2007 10

Special Purpose Switches

- Problem
 - Connect N inputs to M outputs
 - $N \times M$ ("N by M") switch
 - Often $N = M$
- Goals
 - High throughput
 - Best is $\text{MIN}(\text{sum of inputs}, \text{sum of outputs})$
 - Avoid contention
 - Good scalability
 - Linear size/cost growth

10/3/07 CS/ECE 438 - UIUC, Fall 2007 11

Switch Design

- Ports handle complexity
 - Forwarding decisions
 - Buffering
- Simple fabric
 - Move packets from inputs to outputs
 - May have a small amount of internal buffering

10/3/07 CS/ECE 438 - UIUC, Fall 2007 12

Switch Design Goals

- Throughput
 - Main problem is contention
 - Need a good traffic model
 - Arrival time
 - Destination port
 - Packet length
 - Telephony modeling is well understood
 - Until faxes and modems
 - Modeling of data traffic is new
 - Not well understood
 - Will good models help?

10/3/07

CSE/ECE 438 - UIUC, Fall 2007

13



Switch Design Goals

- Contention
 - Avoid contention through intelligent buffering
 - Use output buffering when possible
 - Apply back pressure through switch fabric
 - Improve input buffering through non-FIFO buffers
 - Reduces head-of-line blocking
 - Drop packets if input buffers overflow

10/3/07

CSE/ECE 438 - UIUC, Fall 2007

14



Switch Design Goals

- Scalability
 - $O(N)$ ports
 - Port design complexity $O(N)$ gives $O(N^2)$ for entire switch
 - Port design complexity of $O(1)$ gives $O(N)$ for entire switch

10/3/07

CSE/ECE 438 - UIUC, Fall 2007

15



Switch Design

- Crossbar switches
- Banyan Networks
- Batcher Networks
- Sunshine Switch

10/3/07

CSE/ECE 438 - UIUC, Fall 2007

16



Crossbar Switch

- Every input port is connected to every output port
 - $N \times N$
- Output ports
 - Complexity scales as $O(N^2)$


10/3/07

CSE/ECE 438 - UIUC, Fall 2007

17



Crossbar Switch



10/3/07

CSE/ECE 438 - UIUC, Fall 2007

18




Knockout Switch

- Assumption:
 - It is unlikely that N inputs will have packets destined for the same output port
- Pick L from N packets at a port
 - Output port maintains L cyclic buffers
 - Shifter places up to L packets in one cycle
 - Each buffer gets only one packet
 - Output port uses round-robin between buffers
 - Arrival order is maintained
- Problem
 - Hot spots
- Output ports scale as $O(N)$

Knockout Switch

- Output port design
 - Packet filters
 - Recognize packets destined for a specific port
 - Concentrator
 - Selects up to L packets from those destined for this port
 - Discards excess packets
 - Queue
 - Length L

Knockout Switch




Self-Routing Fabrics


- Idea
 - Use source routing on "network" in switch
 - Input port attaches output port number as header
 - Fabric routes packet based on output port
- Types
 - Banyan Network
 - Batchner-Banyan Network
 - Sunshine Switch

Banyan Network

- A network of 2x2 switches
 - Each element routes to output 0 or 1 based on packet header
 - A switch at stage i looks at bit i in the header



Banyan Network



Banyan Network

- Perfect Shuffle
 - N inputs requires $\log_2 N$ stages of $N/2$ switching elements
 - Complexity on order of $N \log_2 N$
- Collisions
 - If two packets arrive at the same switch destined for the same output port, a collision will occur
 - If all packets are sorted in ascending order upon arrival to a banyan network, no collisions will occur!

10/3/07

CS/ECE 438 - UIUC, Fall 2007

25



Batcher Network

- Performs merge sort
- A network of 2×2 switches
 - Each element routes to output 0 or 1 based on packet header
 - A switch at stage i looks at the whole header
 - Two types of switches
 - Up switch
 - Sends higher number to top output (0)
 - Down switch
 - Sends higher number to bottom output (1)


10/3/07

CS/ECE 438 - UIUC, Fall 2007

26



Batcher Network




10/3/07

CS/ECE 438 - UIUC, Fall 2007

27



Batcher Network



10/3/07

CS/ECE 438 - UIUC, Fall 2007

28



Batcher Network

- How it really works
 - Merger is presented with a pair of sorted lists, one in ascending order, one in descending order
 - First stage of merger sends packets to the correct half of the network
 - Second stage sends them to the correct quarter
- Size
 - $N/2$ switches per stage
 - $\log_2 N \times (1 + \log_2 N)/2$ stages
 - Complexity = $N \log_2^2 N$

10/3/07

CS/ECE 438 - UIUC, Fall 2007

29



Batcher-Banyan Network

- Idea
 - Attach a batcher network back-to-back with a banyan network
 - Arbitrary unique permutations can be routed without contention
 - Two packets destined for same output port still collide!
- Sunshine Switch
 - Like a knockout switch
 - Can handle up to L packets per output port
 - Recirculates overflow packets
 - If more than L packets arrive for any output port in one cycle

10/3/07

CS/ECE 438 - UIUC, Fall 2007

30



Sunshine Switch

- Elements
 - Multiple Banyan networks
 - Enables multiple packets per output port
 - Delay Box
 - Excess (K) packets are recirculated and resubmitted to the switch
 - Batcher network
 - N new packets
 - K delayed packets
 - Trap
 - Identifies packets destined for banyan
 - Identifies excess packets
 - Selector
 - Routes multiple packets for same output on separate banyans


10/307

CS/ECE 438 - UIUC, Fall 2007

31



Sunshine Switch



10/307

CS/ECE 438 - UIUC, Fall 2007

32



Invariants

- Batcher
 - All packets are sorted by destination
- Trap
 - Compares line i with line $i+L$, marks $i+L$ for recirculation if address equal
- Selector
 - Another batcher network, sorting based on the recirculation bit
- Banyan
 - Alternate every L lines to each Banyan

10/307

CS/ECE 438 - UIUC, Fall 2007

33



Sunshine Switch

- Can packets circulate for ever?
 - Priority bit is used to favor older packets
 - Priority bit also ensure packet order is preserved through the switch

10/307

CS/ECE 438 - UIUC, Fall 2007

34

