ECE 671 – Lecture 7

Network Systems and
Interconnects and Switching Fabrics

Network systems

- How to interconnect ports of the network system?

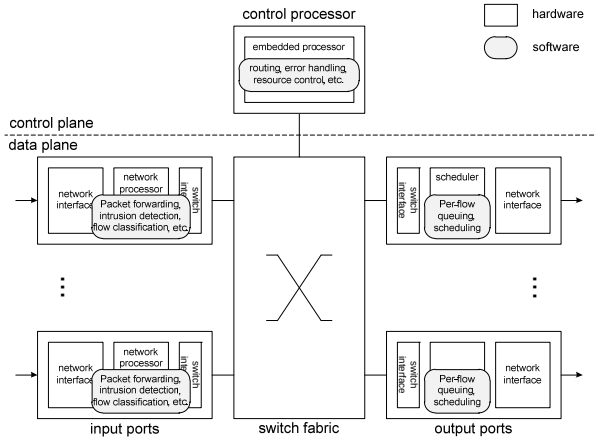

Interconnect designs

- What are possible designs for an interconnect?
  - Need to get data from any input to any output

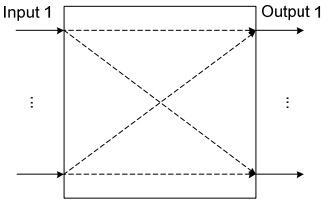

Interconnect designs

- Shared bus
- Shared memory
- Crossbar
- Multistage switching fabric

- Scalability is important
  - Assume N ports with C bandwidth each
  - What is the speed requirement for each component?
  - In practice: 1xC to 2xC is reasonable, more is not feasibles

Scalability evaluation

- Shared bus

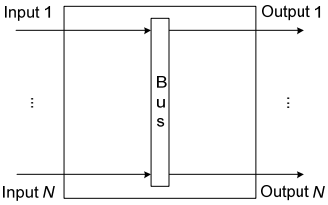

- Shared memory

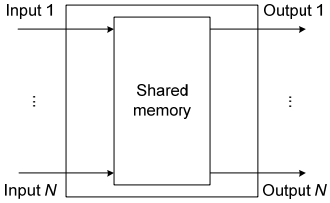

- Crossbar

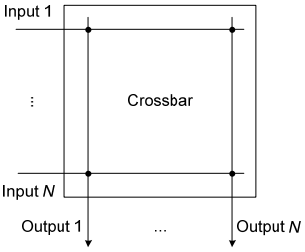

Multistage switching fabric

- Switching element
  - Two states: cross-over or straight
  - Very simple implementation

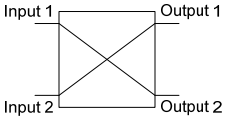

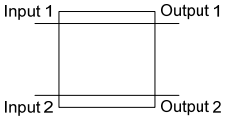

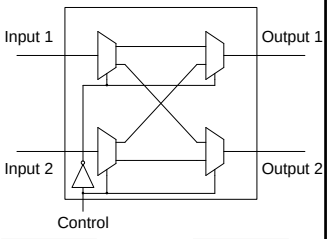

- Switch fabric assembled from multiple switching elements:

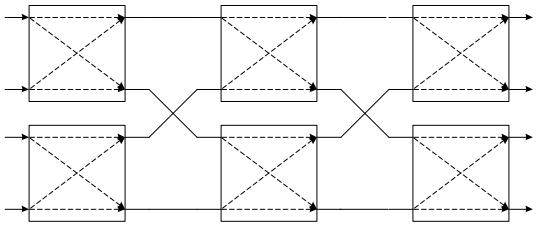

Switching fabric topologies

- Topologies determine properties
  - E.g., blocking probability, self-routing property, etc.

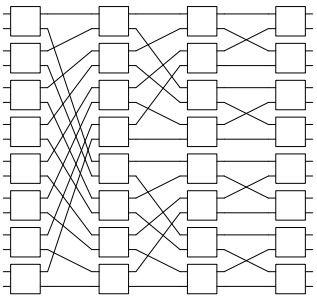

Delta network

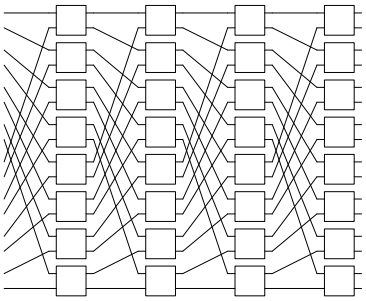

Omega network

Contention in interconnects

- Any traffic pattern should be handled
  - Focus on “permissible traffic load”
  - Over “long” term, each port receives and sends C
- Short-term variation in traffic can lead to contention
  - Switching system needs to deal with contention
- What are possible solutions

Queuing techniques

- Input queuing

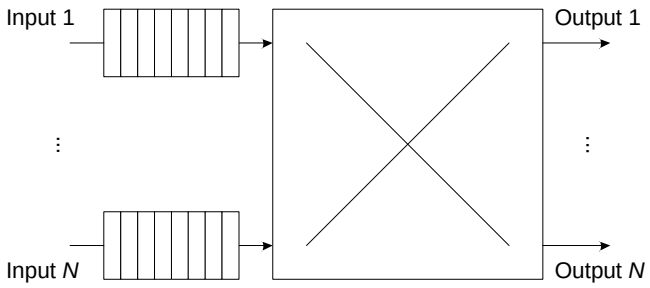

Queuing techniques

- Output queuing

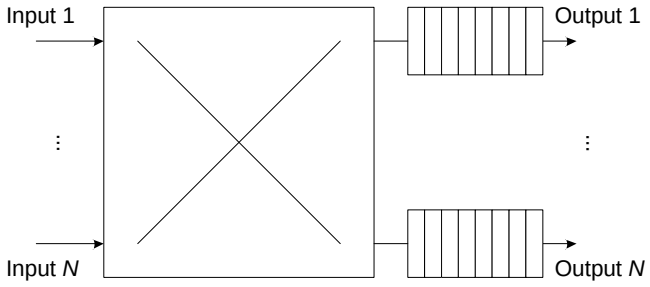

Queuing techniques

- Virtual output queuing

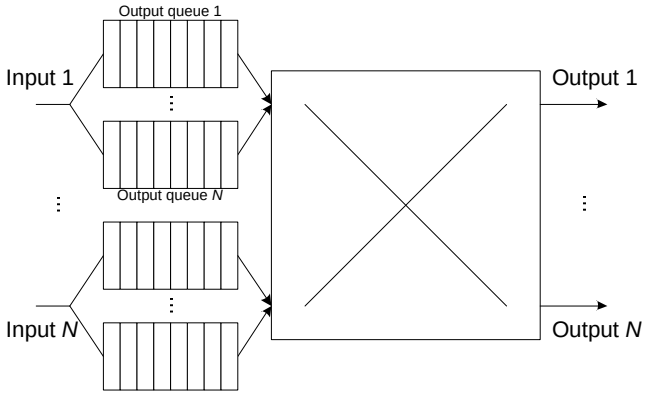

Queuing techniques

- Combined input-output queuing

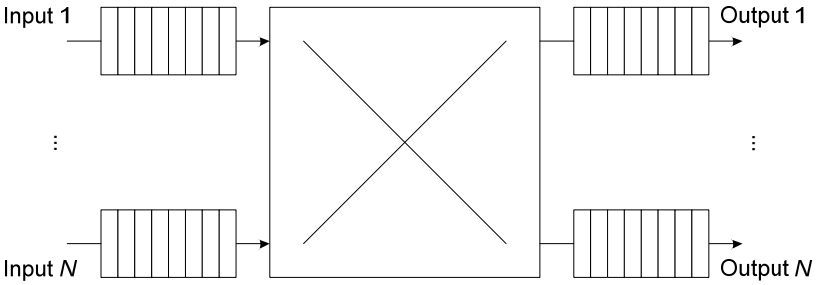

Queuing techniques

- Input queuing
  - Head-of-line blocking
  - Theoretical throughput limit 58.6%
- Output queuing
  - Requires speedup of N
- Virtual output queuing
  - N^2 queues
- VOQ and CIOQ require scheduling algorithm

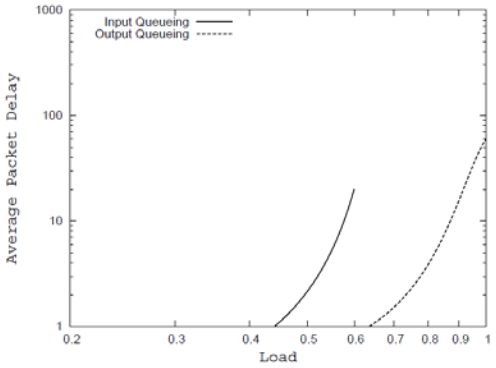

Switch scheduling

- Model of scheduling problem
  - Request matrix or bipartite graph
  - Goal: find match in graph

Inputs \ Outputs	0	1	2	3
0	1	0	0	1
1	0	1	1	0
2	0	1	0	1
3	1	1	1	0

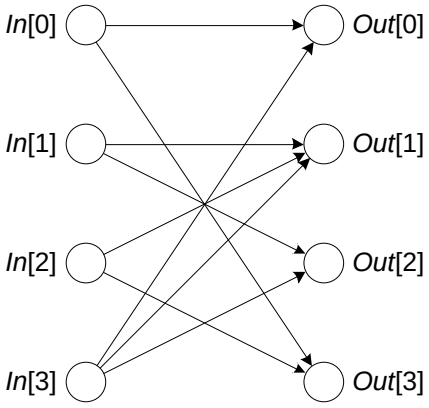

Switch scheduling

- Goals for scheduler
  - High throughput
  - Low latency
  - Fairness
  - Low-cost implementation
- Comparison
  - iSlip, FIRM, MP
  - Output queuing is optimum