3. Time Switching, Multi-Queue Memories, Shared Buffers, Output Queueing Family

Manolis Katevenis

CS-534 – Univ. of Crete and FORTH, Greece

`http://archvlsi.ics.forth.gr/~kateveni/534`

3.5 Output Queueing & Shared Buffer Family

Table of Contents:

Output Queueing:

The "Reference" Architecture

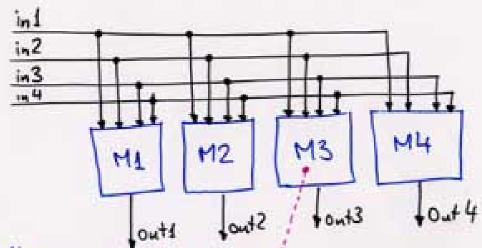

⇒ "Top" Performance:

- no head-of-line blocking
- full outgoing throughput utilization (no internal blocking)
- "minimum" delay
- adaptable to any QoS policy ⇒
  - inside each buf. memory:
    - one FIFO queue per output
    - or:
    - multiple queues + scheduler:
      - per connection, or
      - per flow, or
      - per group of flows, or
      - per priority, or...
- multicast traffic handled cleanly, at top performance

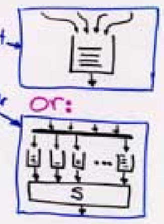

⇒ Unnecessarily High Cost:

- wasteful in memory throughput (but interesting for use with multicast packet pointers)
- partitioned buffer space is less efficient than shared

Shared Buffer:

Top Performance at Low Cost for small N

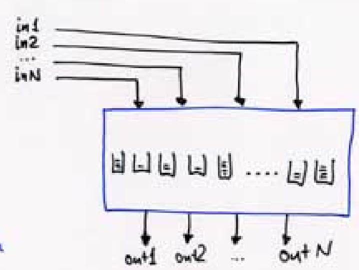

- total buffer memory throughput = 2N (versus N×(N+1) for output queueing)
- memory space is shared ⇒ better utilization
- same performance as output queueing for unicast traffic
  - multiple logical queues in a single memory, at least per output, possibly also per priority/flow)...
- for multicast packets: not enough throughput to enqueue each arriving packet into multiple (per output) queues. Hence, if fewer than 2^N multicast queues exist, some head-of-line blocking will occur in them. Interesting combination:
  - shared buffer for packet bodies
  - output queueing for queue pointers

3.5.2 Buffer Space Requirements: Analysis Results

- Analysis & simulation have yielded the results plotted below
  - *Reference:* M. Hluchyj, M. Karol: "Queueing in High-Performance Packet Switching", IEEE Journal on Sel. Areas in Communications (JSAC), vol. 6, no. 9, Dec. 1988, pp. 1587-1597
- **Assuming** that the input traffic consists of packets with/from:
  - uniformly-distributed destination (output) ports,
  - independent, identically distributed (i.i.d.) Bernoulli processes,
  - fixed-size packet (cell) traffic
- *Attention:* results derived for i.i.d. Bernoulli (non-bursty) arrivals, with uniformly-distributed destination (no overloaded hot-spot output ports), are only useful for gaining a first, rough insight into the behavior of systems, but are usually not representative of the real behavior of systems under real traffic patterns!...

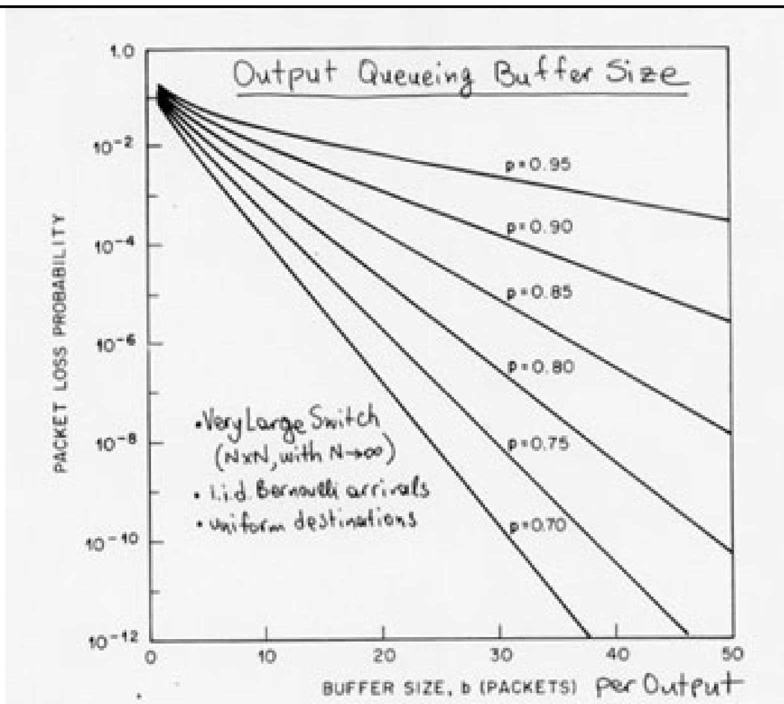

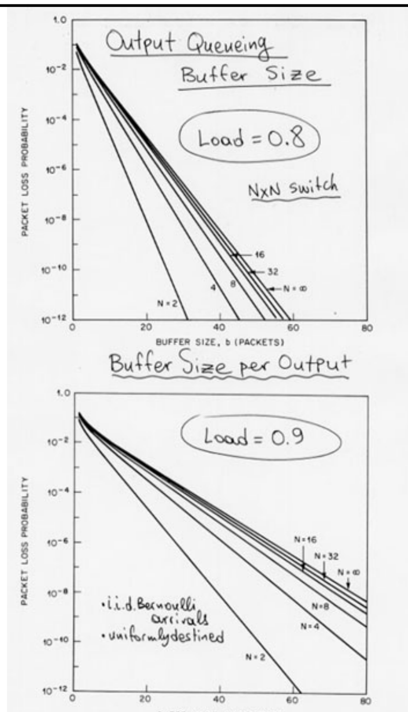
Output Queueing
Buffer Size
Load = 0.8
NxN switch
PACKET LOSS PROBABILITY
BUFFER SIZE, b (PACKETS)
Buffer Size per Output
Load = 0.9
• i.i.d. Bernoulli arrivals
• uniformly destined

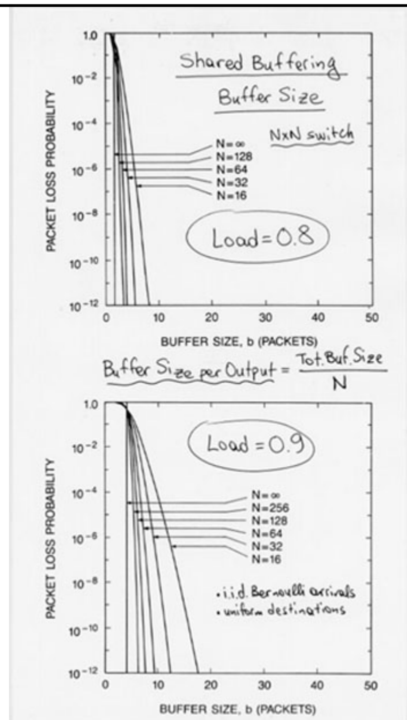
Shared Buffering
Buffer Size
NxN switch
N=∞
N=128
N=64
N=32
N=16
Load = 0.8
PACKET LOSS PROBABILITY
BUFFER SIZE, b (PACKETS)
Buffer Size per Output = Tot. Buf. Size / N
Load = 0.9
N=∞
N=256
N=128
N=64
N=32
N=16
• i.i.d. Bernoulli arrivals
• uniform destinations

7

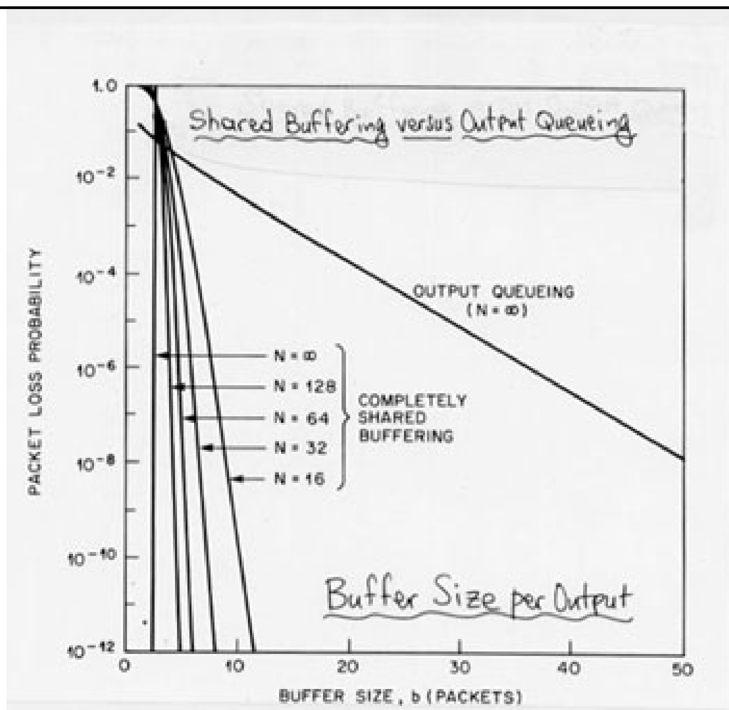
Shared Buffering versus Output Queueing
PACKET LOSS PROBABILITY
OUTPUT QUEUEING (N = ∞)
N = ∞
N = 128
N = 64
N = 32
N = 16
COMPLETELY SHARED BUFFERING
Buffer Size per Output
BUFFER SIZE, b (PACKETS)

Load = 85 %

8

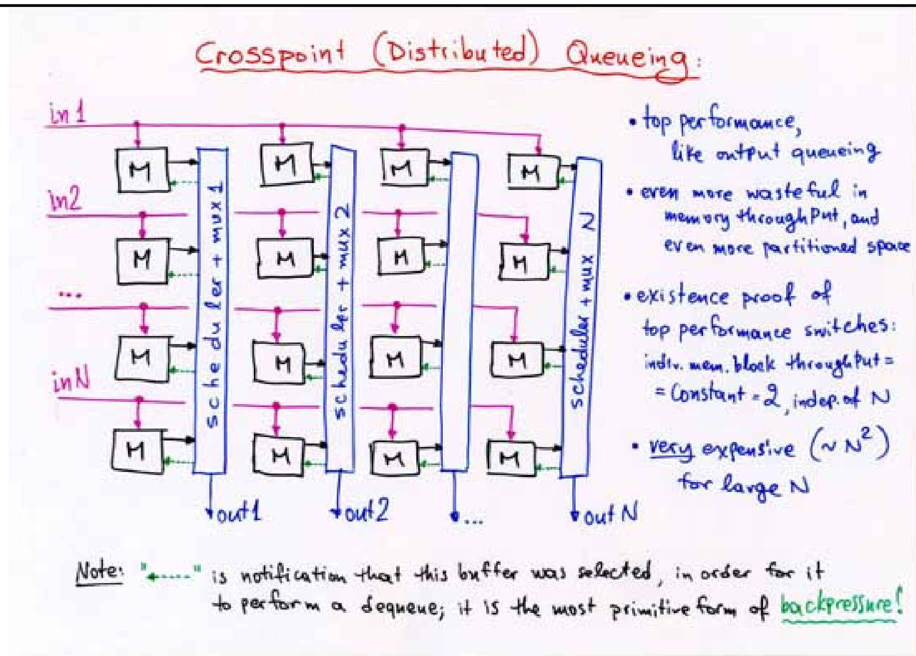

3.5 - U. Crete - M. Katevenis - CS-534

9

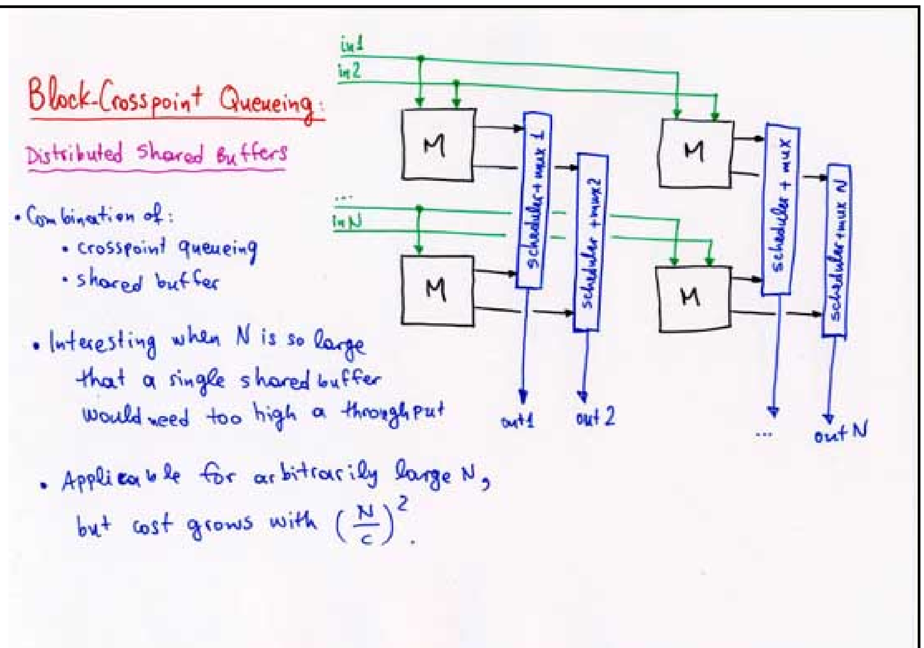

3.5 - U. Crete - M. Katevenis - CS-534

10

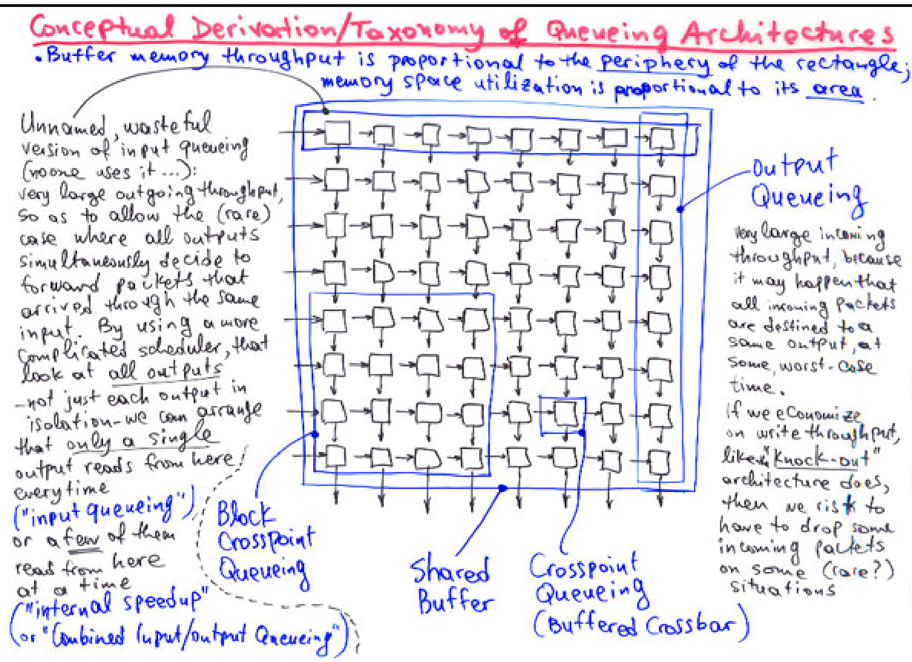

Other Variations of Output Queueing: **Knock-Out Switch**

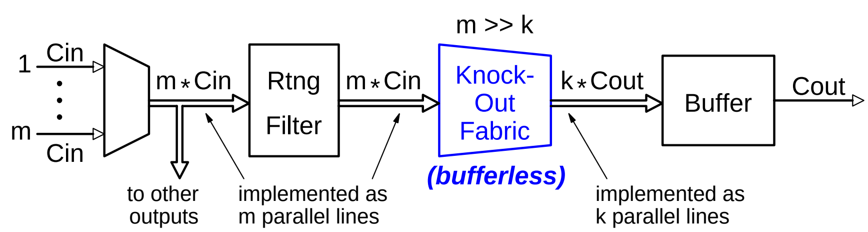

Knock-Out Fabric:

- has m inputs and k outputs, k << m
- passes on up to k non-idle packets to its outputs
- when more than k packets arrive in the same time slot, all destined to the same output, k of them are passed and the rest are dropped
- if the traffic is uniformly destined, and k is on the order of 8 to 12, packets will rarely be dropped

See: Yeh, Hluchyj, Acampora: IEEE JSAC, October 1987, pp. 1274-1283.