2.3 High-Throughput Memories for Time-Switching Shared Buffers

- Time-Switching applied to Packet Switching:
  - multiplex all incoming traffic onto a high-throughput internal shared path that leads into a *shared buffer*
  - read from the shared buffer and demux onto output ports
- High-Throughput Memories $\Rightarrow$ Wide memories
  - how to best combine mux/demux and wide memory?
  - “pipelined memory”
  - interleaved memories & hidden crossbars

Time-Multiplexing the Switch Ports on a Wide-Memory Port:

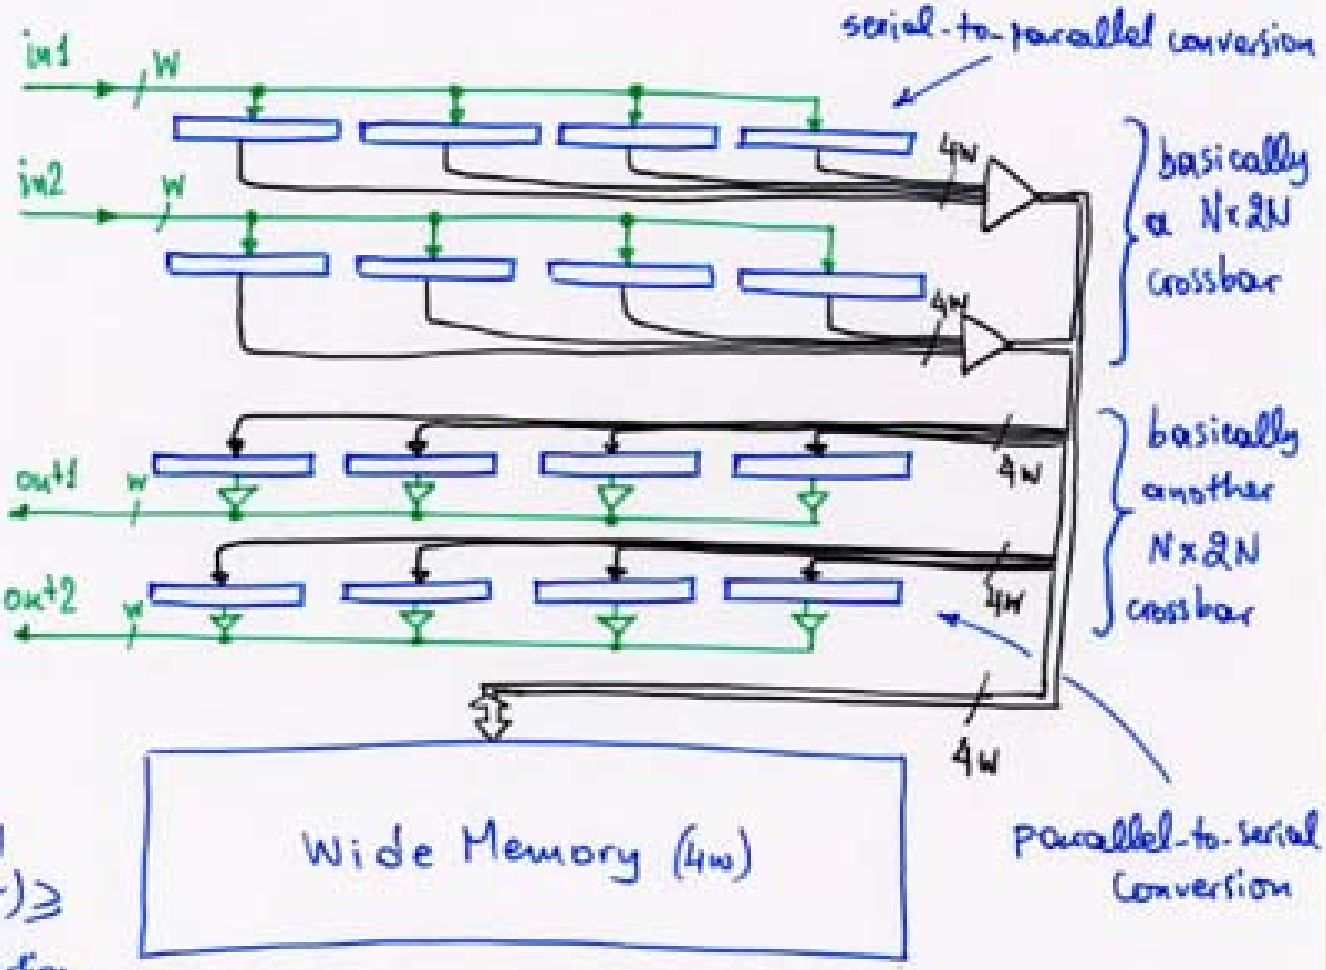

(this figure assumes that incoming packets are synchronized with each other and with the memory port "frame"; in reality, "double buffering" is needed at the serial-to-parallel conversion points).

minimum in→out delay (even upon idle output) ≥

≥ serial-parallel conversion =

= memory-port frame duration

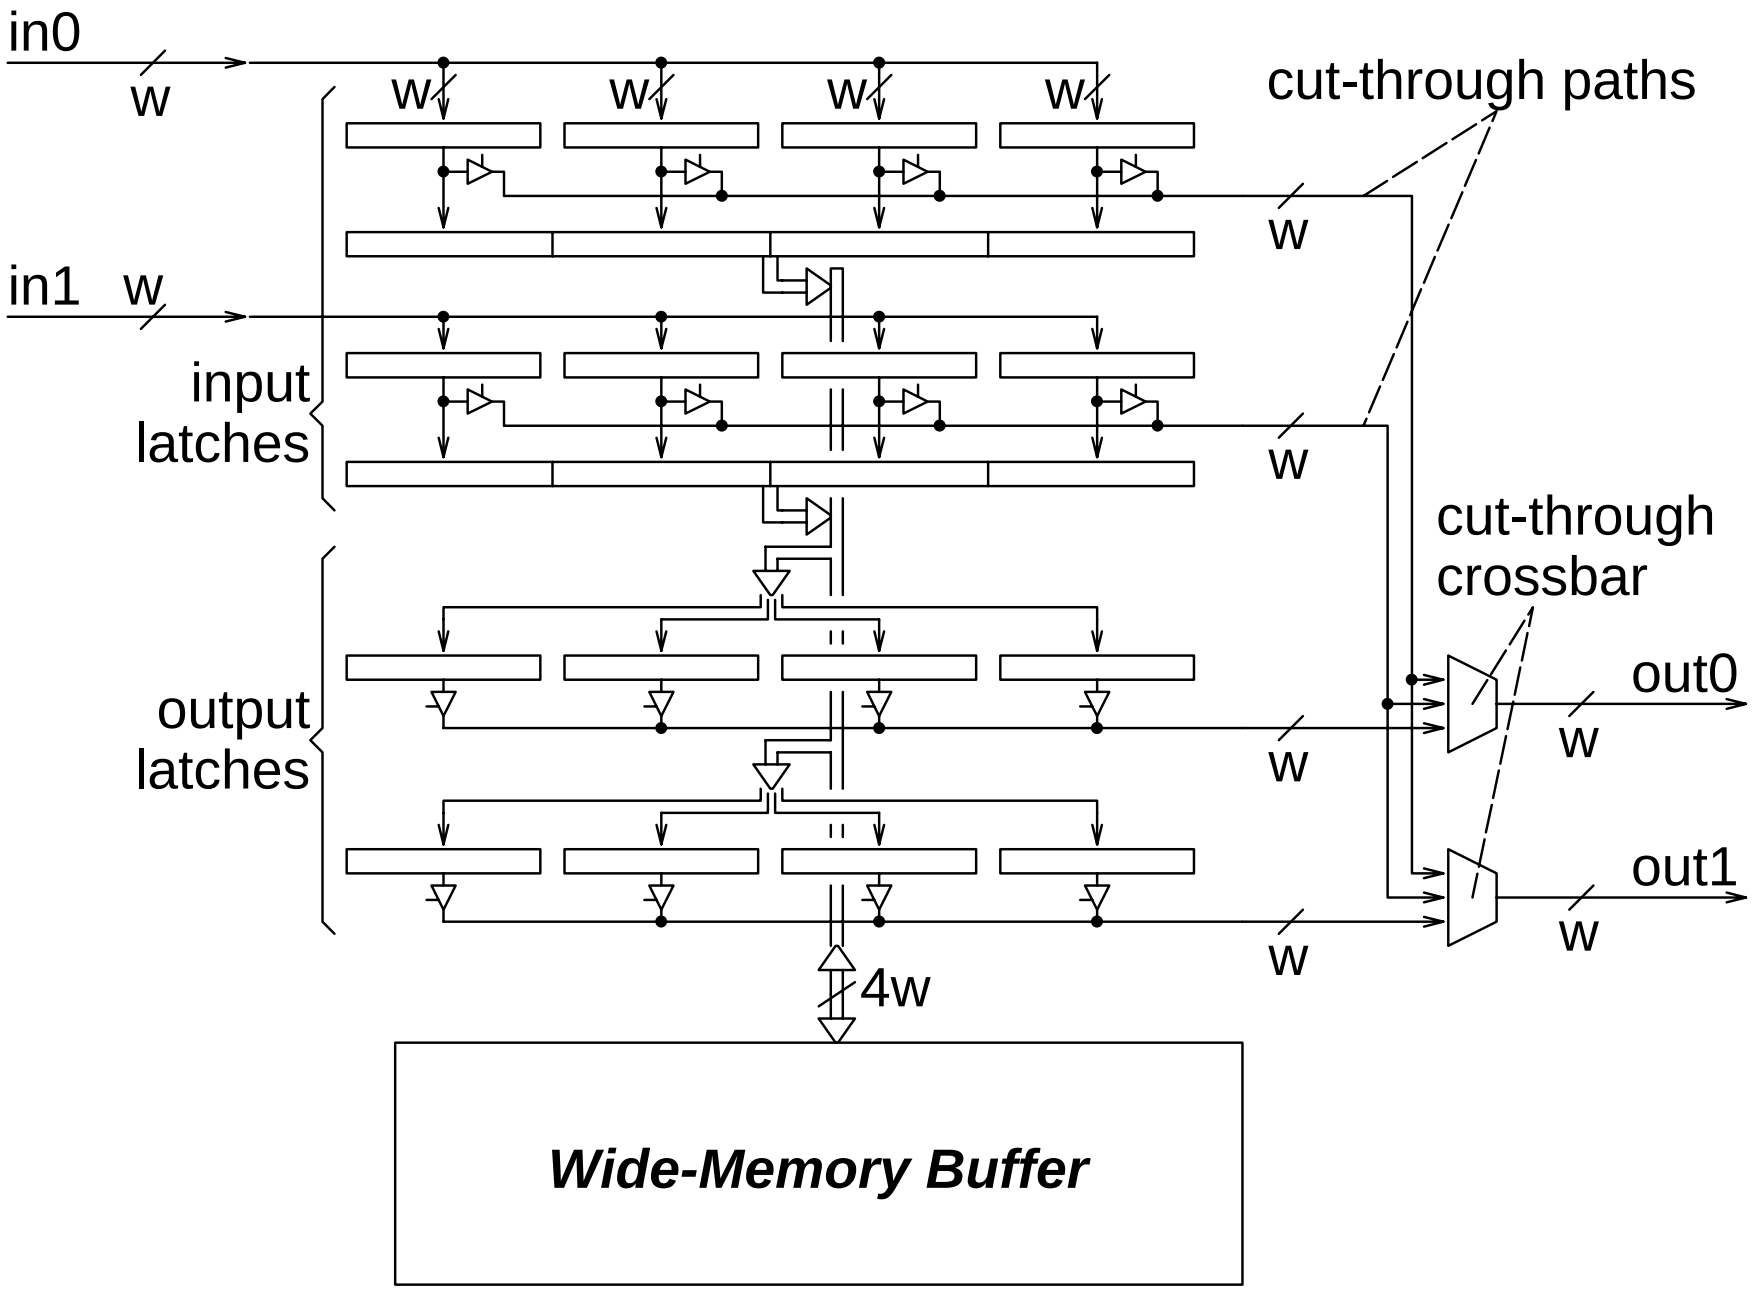
in0
w
in1
w
input
latches
output
latches
w
w
w
w
cut-through paths
w
w
w
w
cut-through
crossbar
out0
w
out1
w
4w
Wide-Memory Buffer

2x2 Shared-Buffer Switch using Pipelined Memory

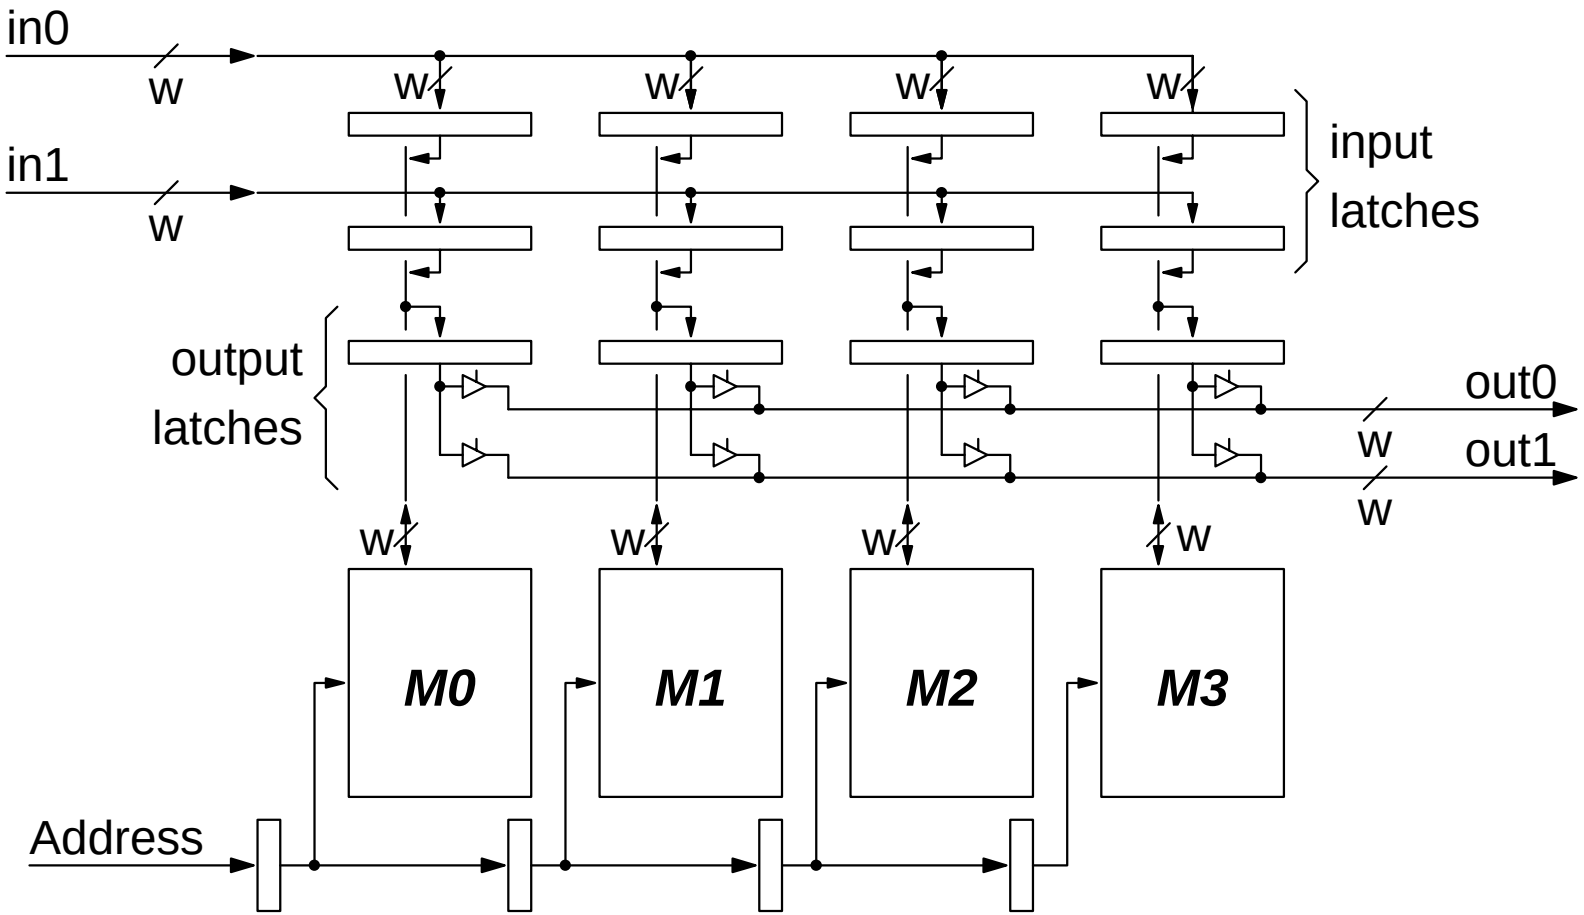

FORTH, Crete Greece (1994): USA Patent number 5,774,653 of 30 June 1998

Katevenis, Vatsolaki, Efthymiou: ACM SIGCOMM Conf., Sep. 1995. pp.39-48

http://archvlsi.ics.forth.gr/sw_arch/pipeMem.html

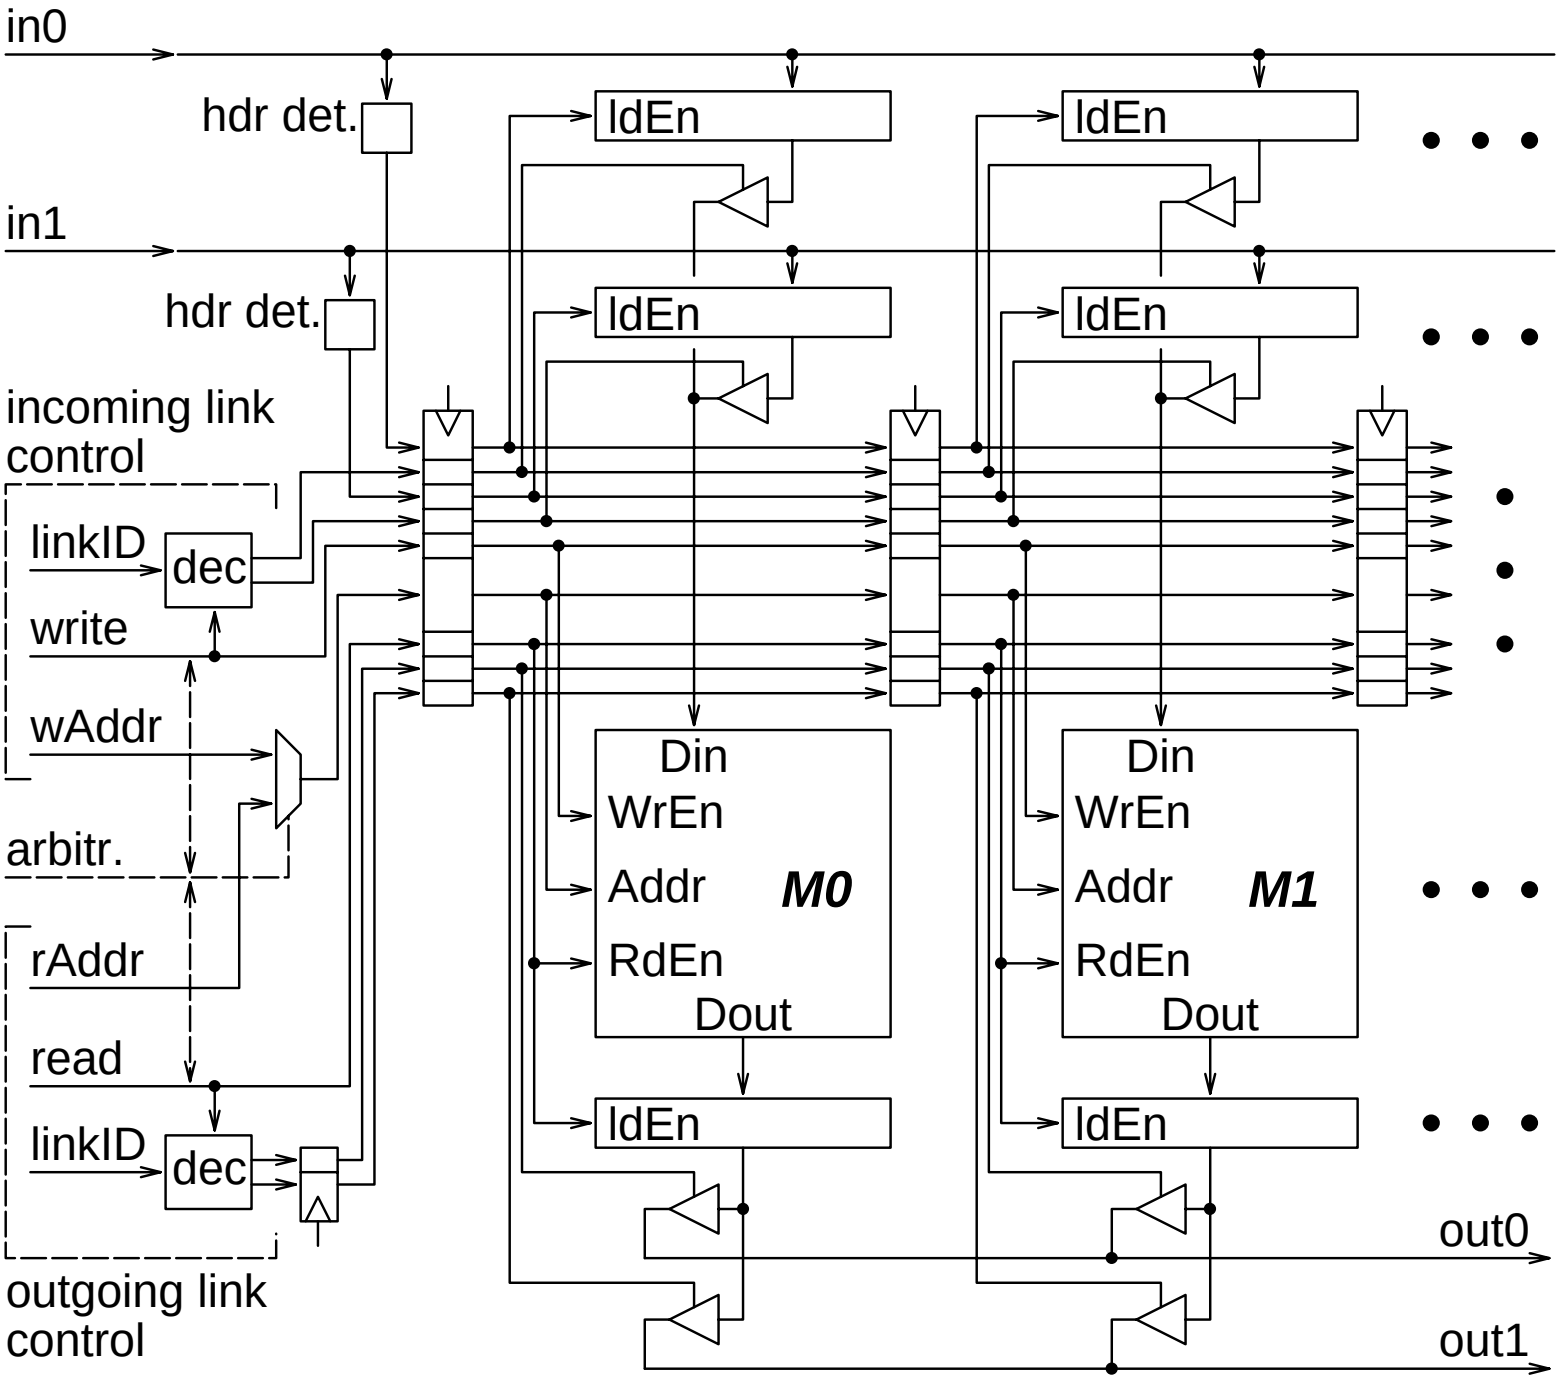
in0
hdr det.
in1
hdr det.
incoming link control
linkID
dec
write
wAddr
arbitr.
rAddr
read
linkID
dec
outgoing link control
ldEn
ldEn
ldEn
ldEn
Din
WrEn
Addr
M0
RdEn
Dout
Din
WrEn
Addr
M1
RdEn
Dout
ldEn
ldEn
out0
out1

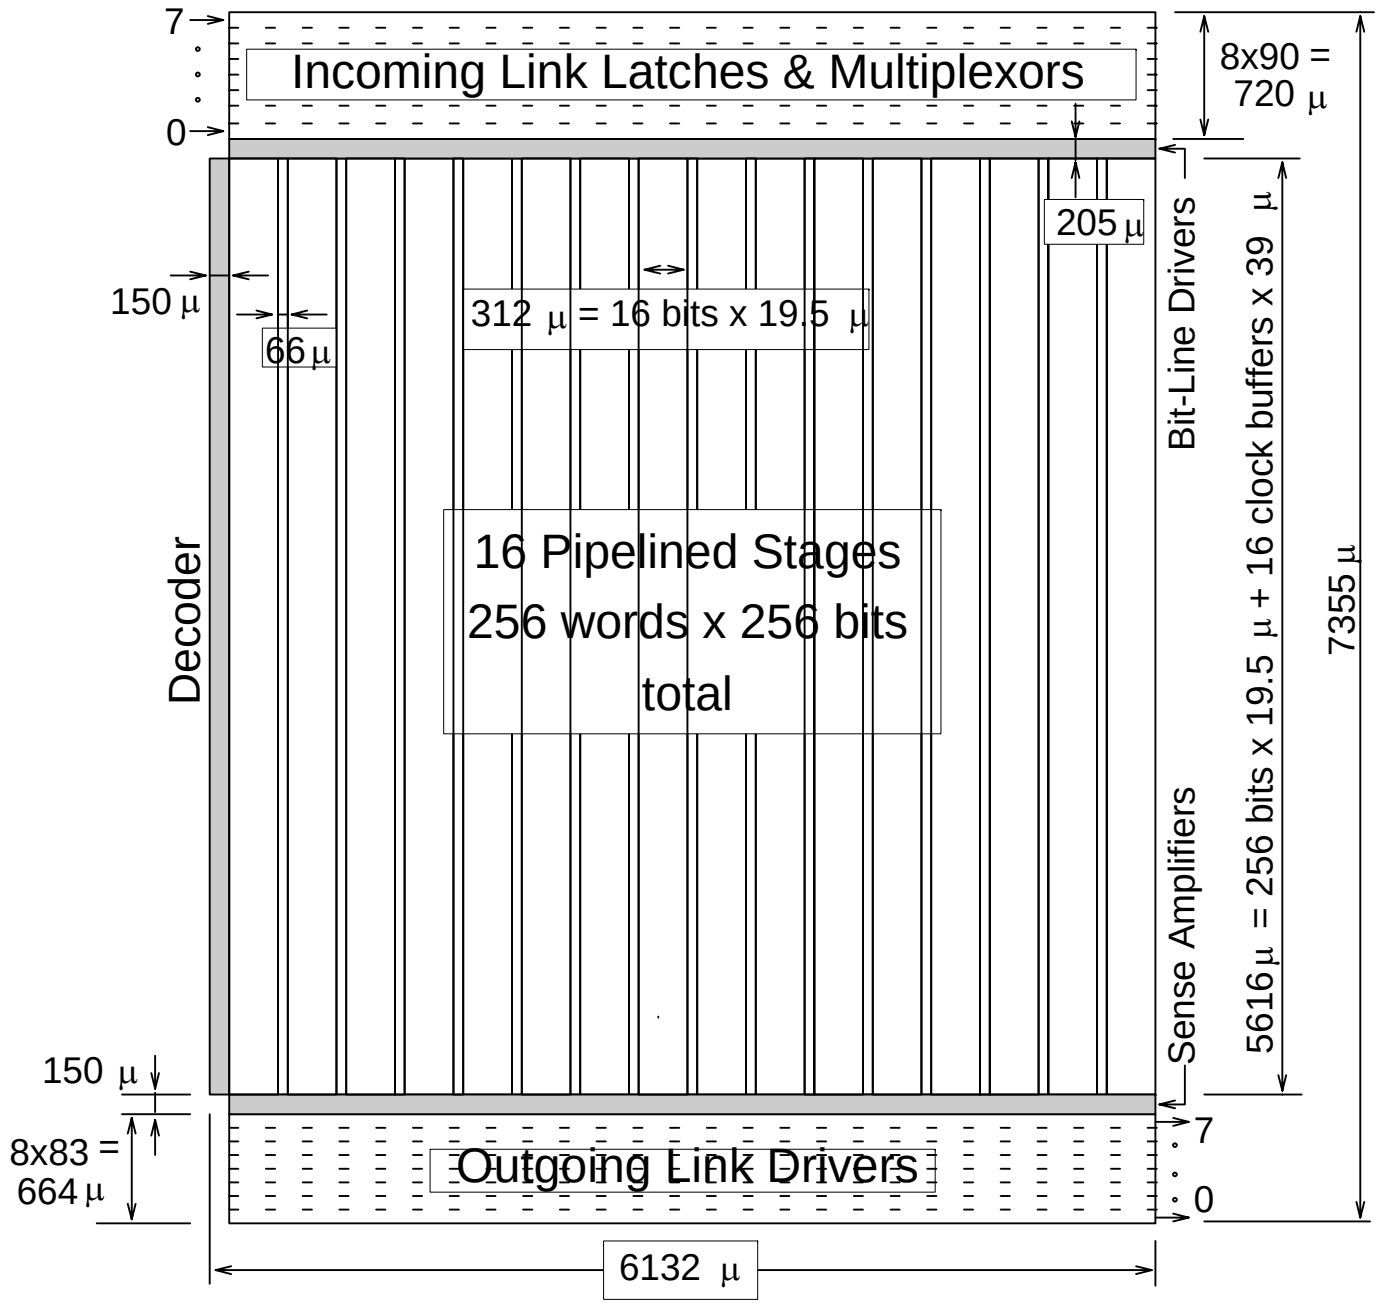

64 Kbit
Pipelined Memory
Shared Buffer
for an 8x8 Switch

1-micron
full-custom CMOS

Link throughput:
1.6 Gb/s (typical),
1.0 Gb/s (worst case)

ICS-FORTH,
Crete, Greece,
1995

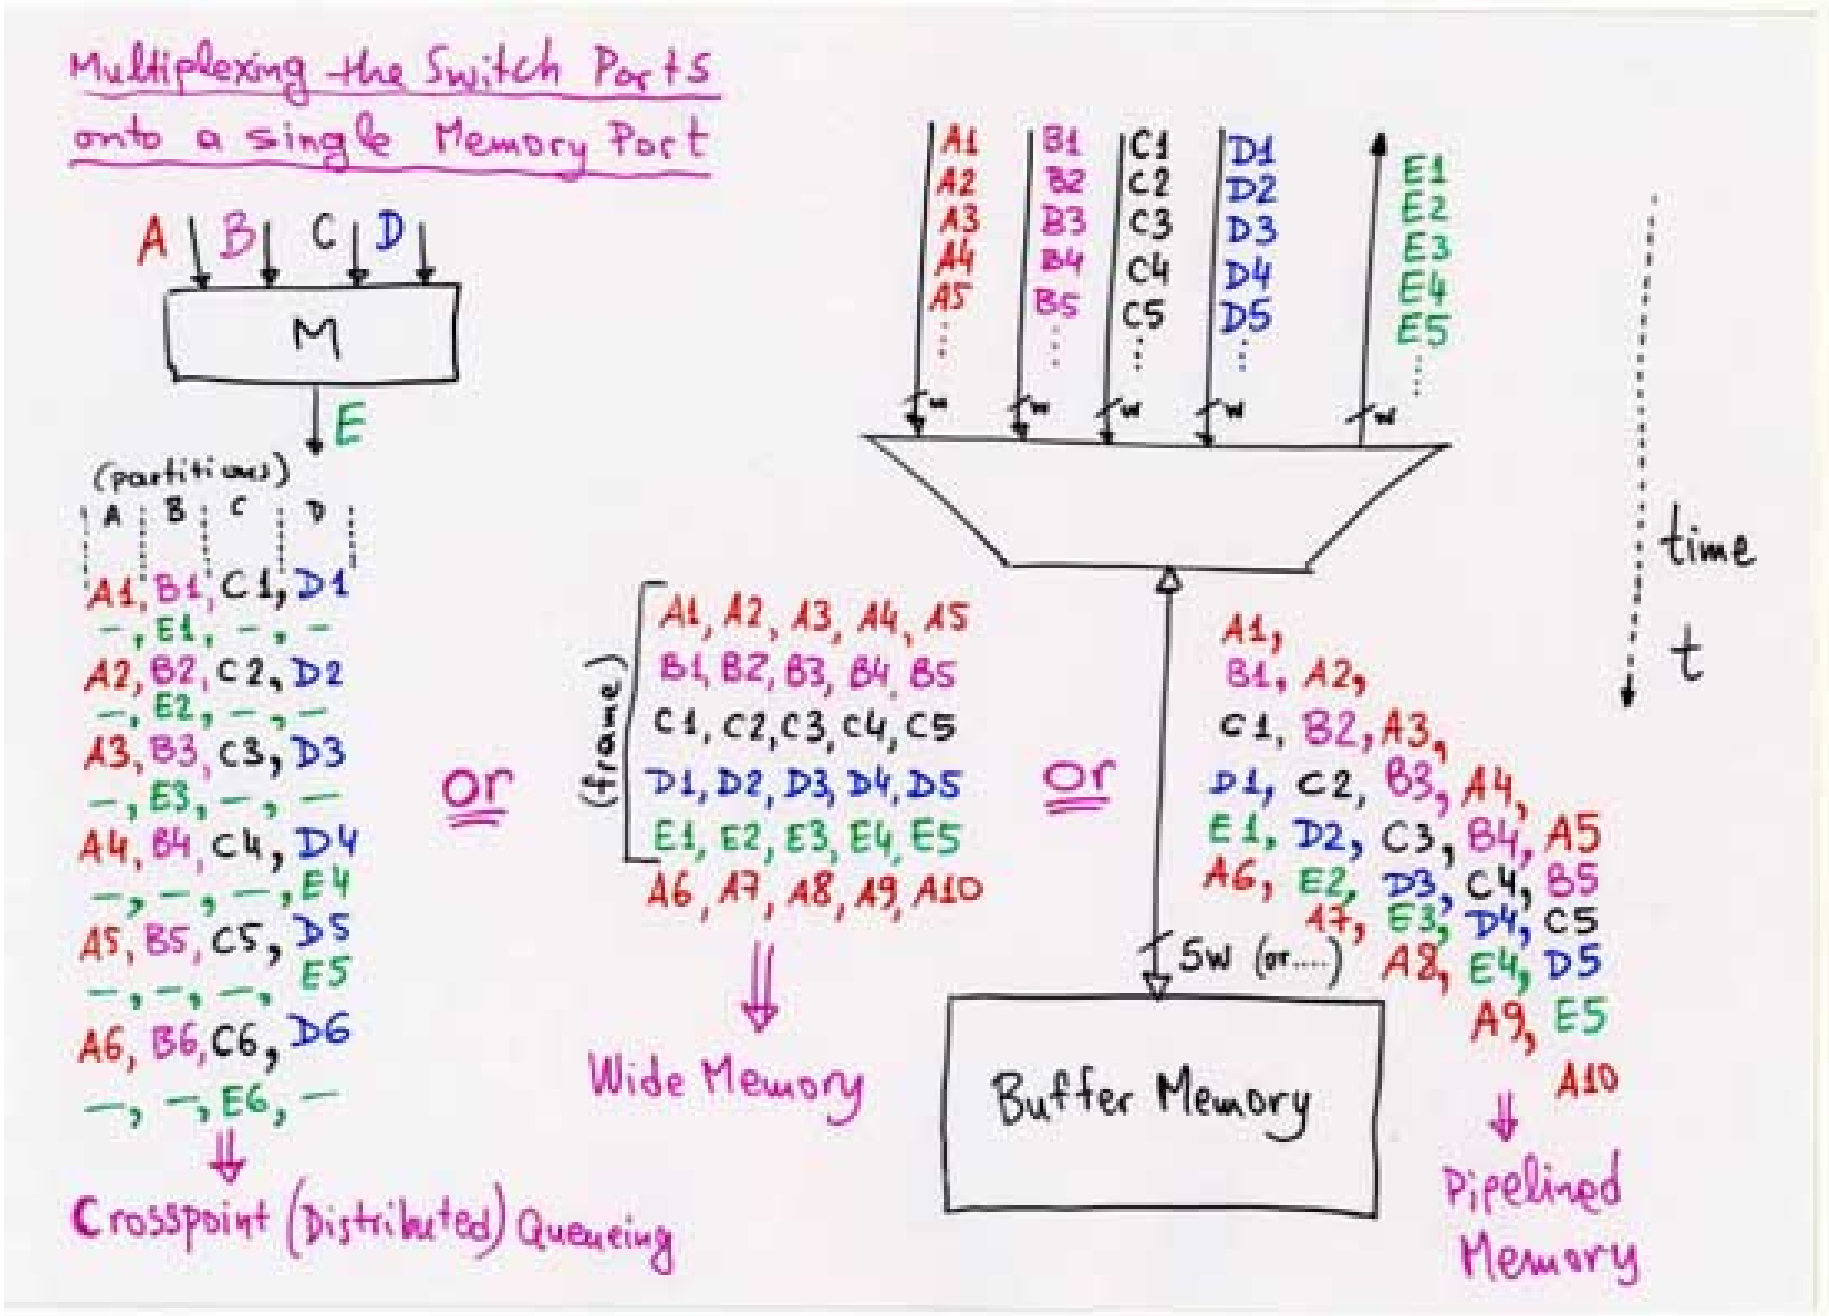
Multiplexing the Switch Ports onto a single Memory Port
A
B
C
D
M
E
(partitions)
A1, B1, C1, D1
–, E1, –, –
A2, B2, C2, D2
–, E2, –, –
A3, B3, C3, D3
–, E3, –, –
A4, B4, C4, D4
–, –, –, E4
A5, B5, C5, D5
–, –, –, E5
A6, B6, C6, D6
–, –, E6, –
Crosspoint (Distributed) Queueing
or
(frame)
A1, A2, A3, A4, A5
B1, B2, B3, B4, B5
C1, C2, C3, C4, C5
D1, D2, D3, D4, D5
E1, E2, E3, E4, E5
A6, A7, A8, A9, A10
Wide Memory
or
A1, B1, C1, D1, E1
A2, B2, C2, D2, E2
A3, B3, C3, D3, E3
A4, B4, C4, D4, E4
A5, B5, C5, D5, E5
A6, A7, A8, A9, A10
w
5w (or ....)
Buffer Memory
Pipelined Memory
time
t

High-Throughput Buffers: Interleaved Memories

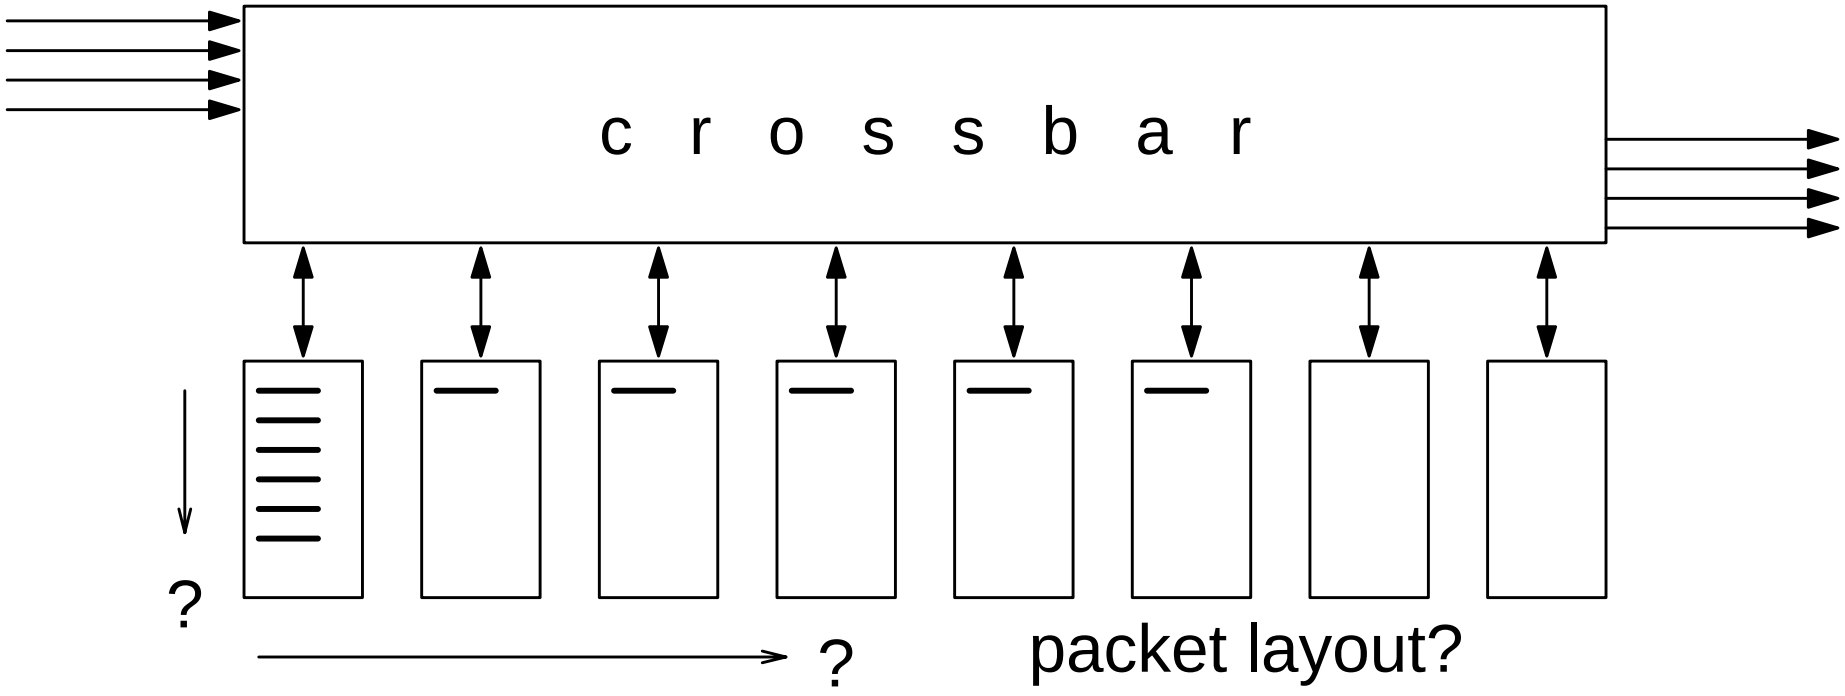

- How to layout the packets so as to avoid bank conflicts?
- Is the memory so wide (number of banks so large) that some packets do not extend across all banks?

All Switches need a Crossbar somewhere inside them....

Crosspoint Queueing:

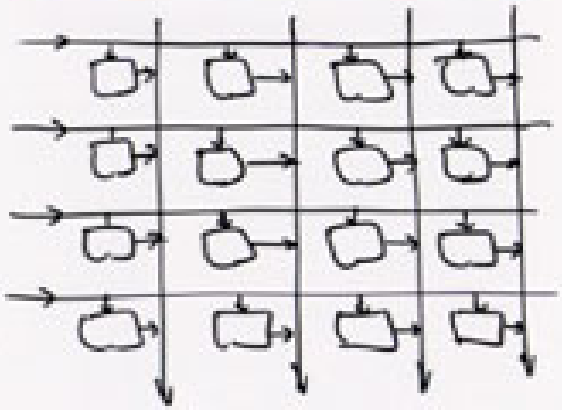

Output Queueing:

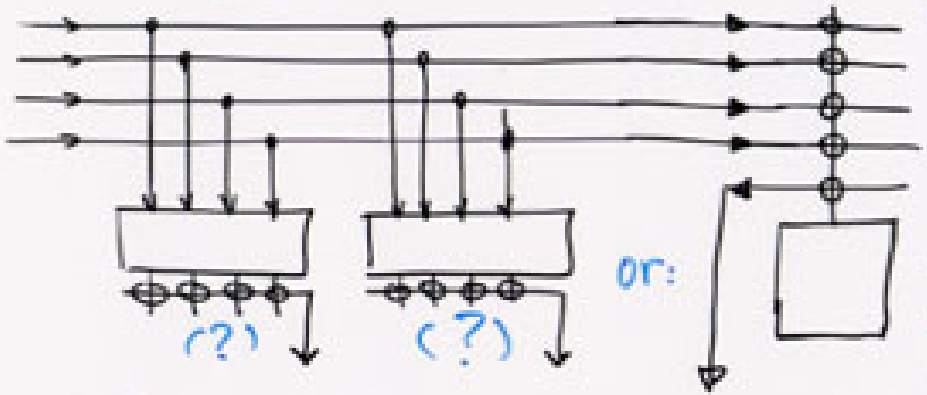

If the memory clock can be run faster than the line clock

Input Queueing:

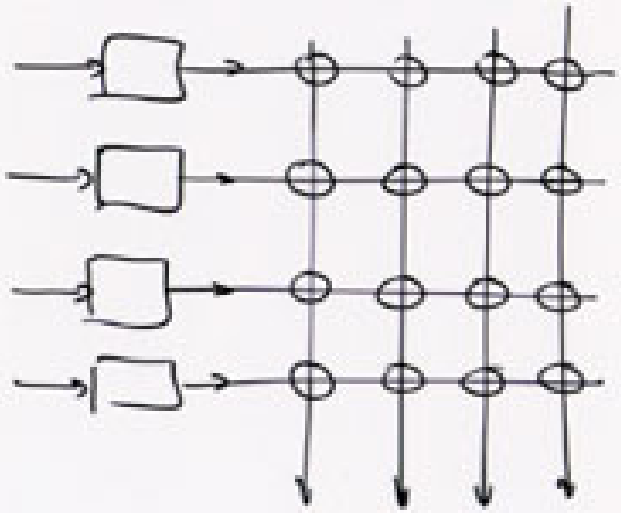

Shared Buffer:

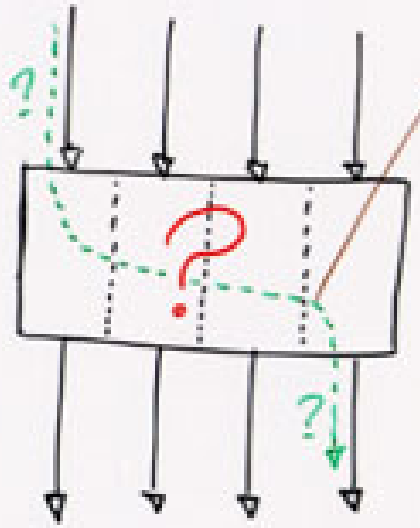

How can a packet/segment/word go this way?

How does "wide" (high-throughput) memory work?