Photonic On-Chip Networks for Performance-Energy Optimized Off-Chip Memory Access

GILBERT HENDRY

JOHNNIE CHAN, DANIEL BRUNINA,
LUCA CARLONI, KEREN BERGMAN

Lightwave Research Laboratory
Columbia University
New York, NY

Motivation

- The memory gap warrants a paradigm shift in how we move information to and from storage and computing elements

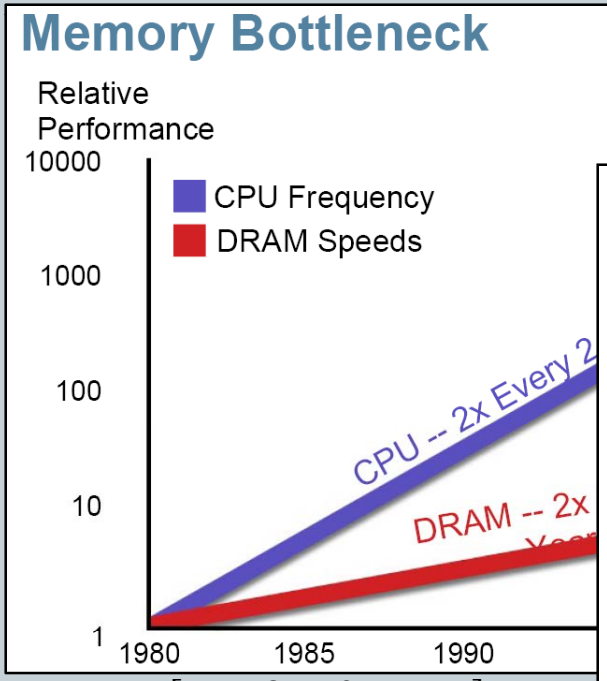

[www.OpenSparc.net]

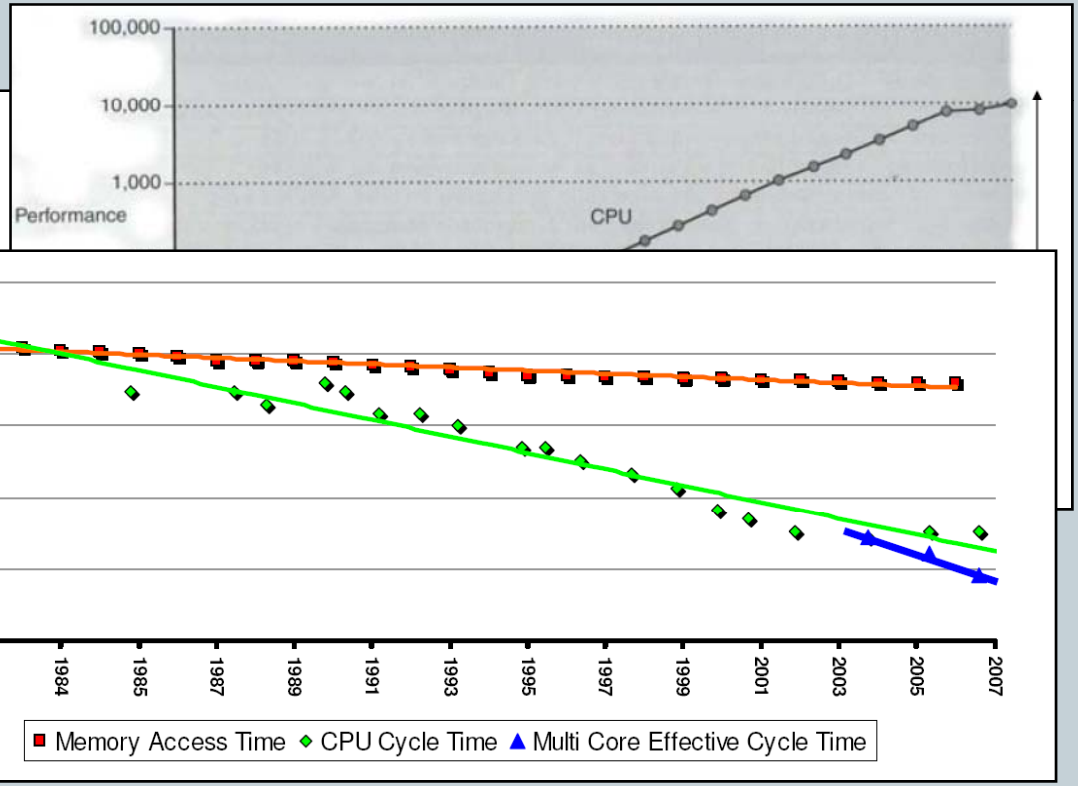

[Exascale Report, 2008]

Main Premise

- Current memory subsystem technology and packaging are not well-suited to future trends
  - Networks on chip
  - Growing cache sizes
  - Growing bandwidth requirements
  - Growing pin counts

SDRAM context

- DIMMs controlled fully in parallel, sharing access on data and address busses
- Many wires/pins
- Matched signal paths (for delay)
- DIMMs made for short, random accesses

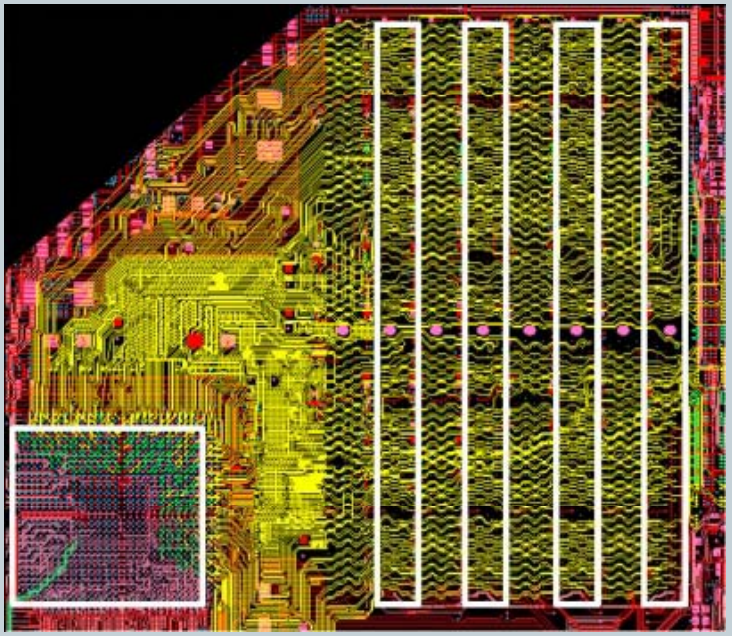

[Intel]

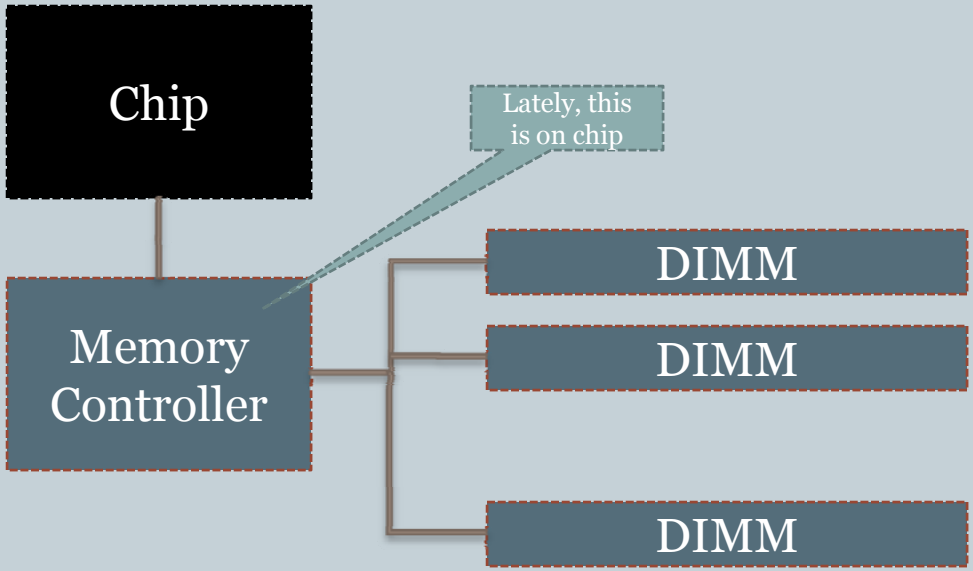

Future SDRAM context

- Example: Tilera TILE 64

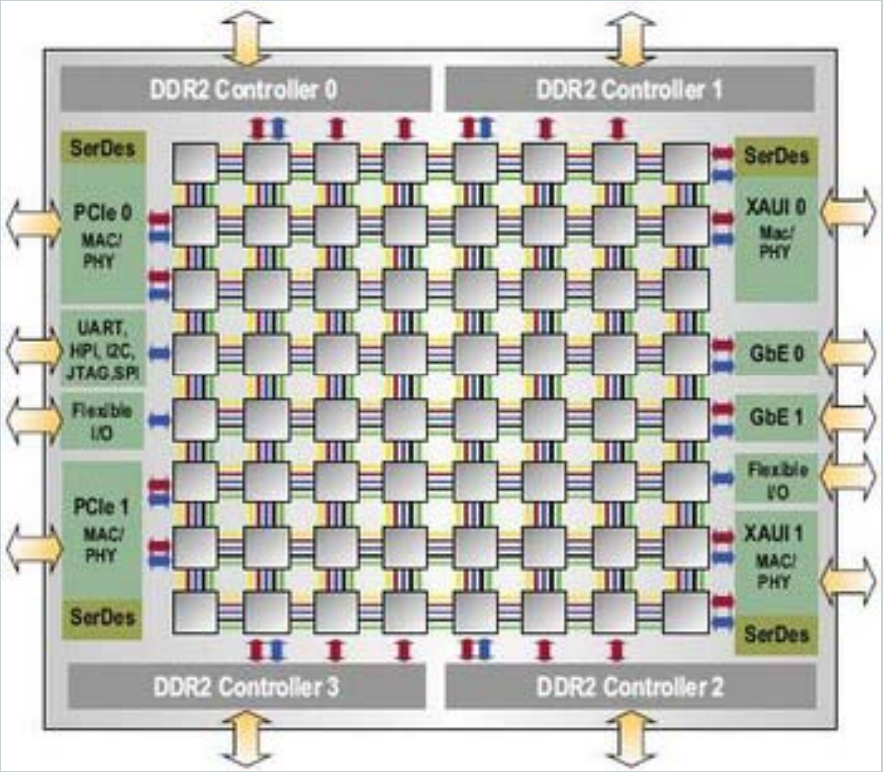

SDRAM DIMM Anatomy

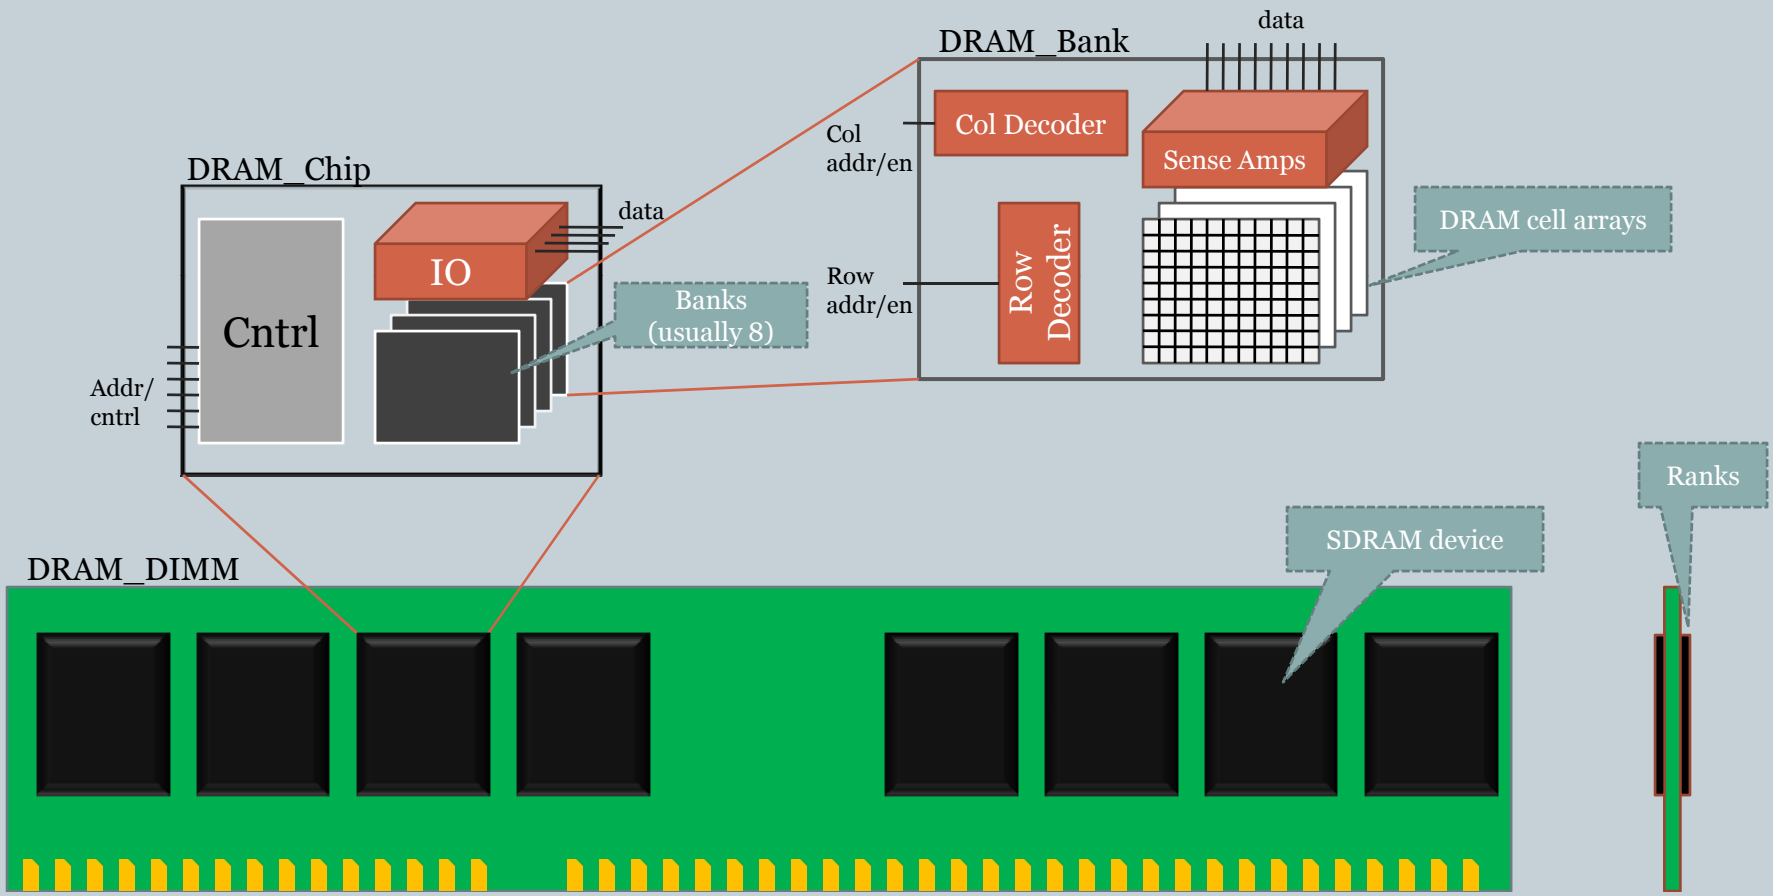

Memory Access in an Electronic NoC

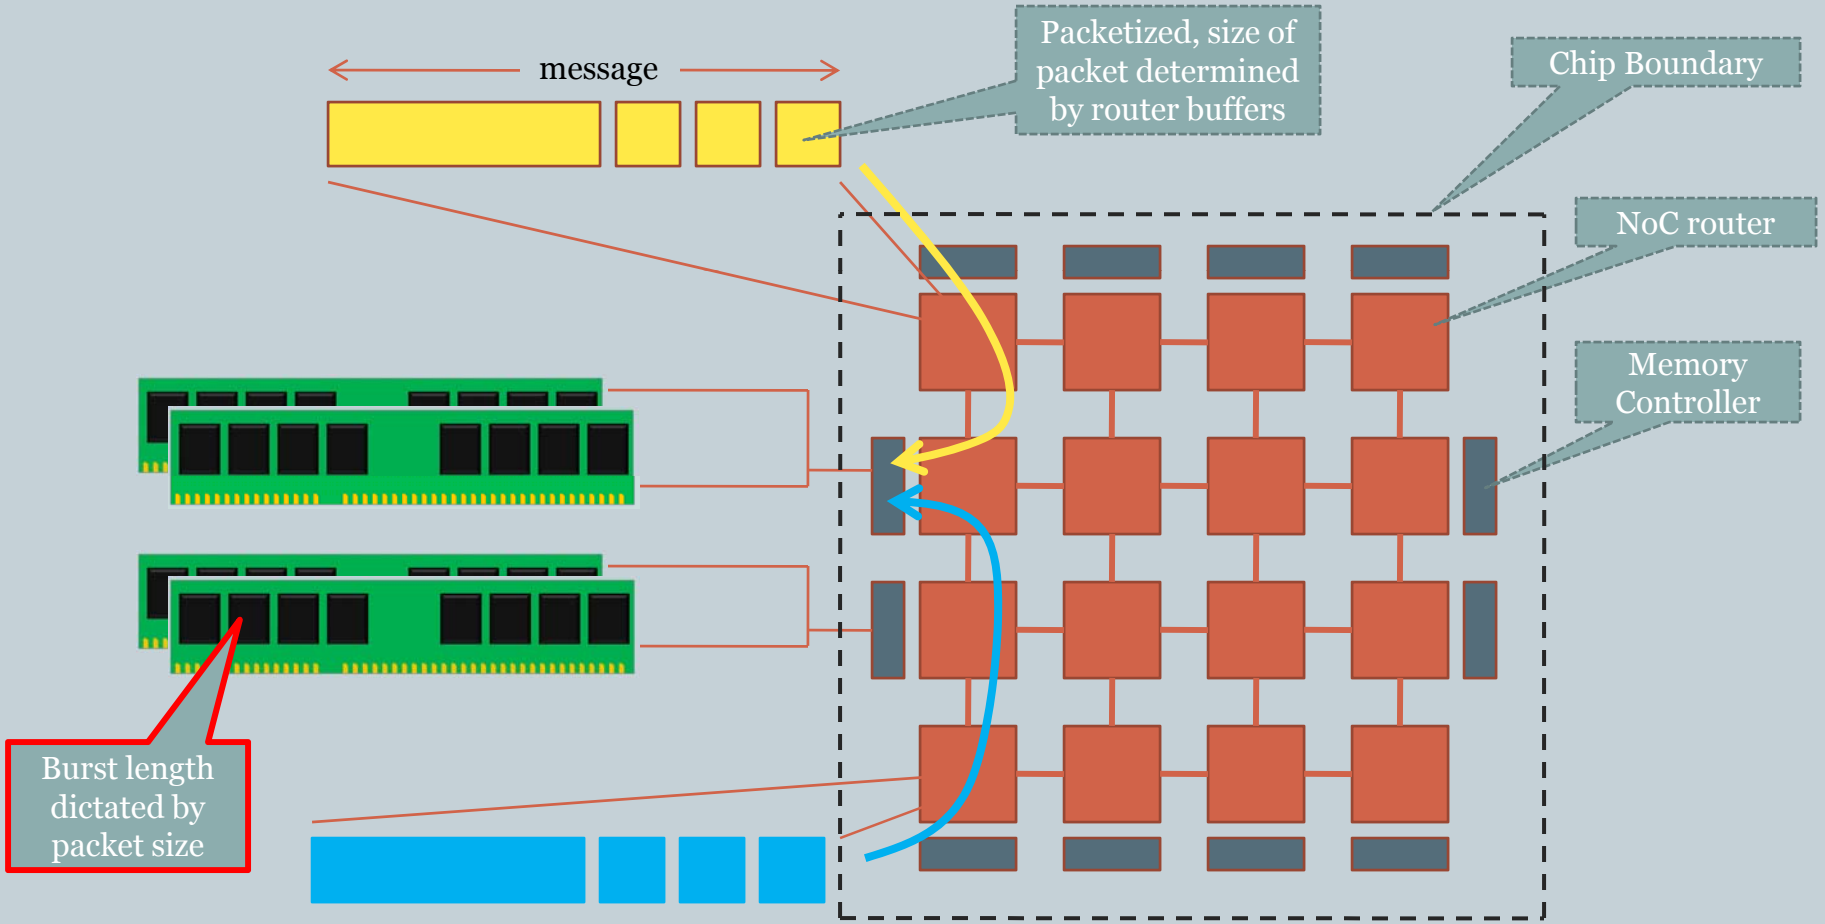

Memory Control

- Complex DRAM control
  - Scheduling accesses around:
    - Open/closed rows
    - Precharging
    - Refreshing
    - Data/Control bus usage

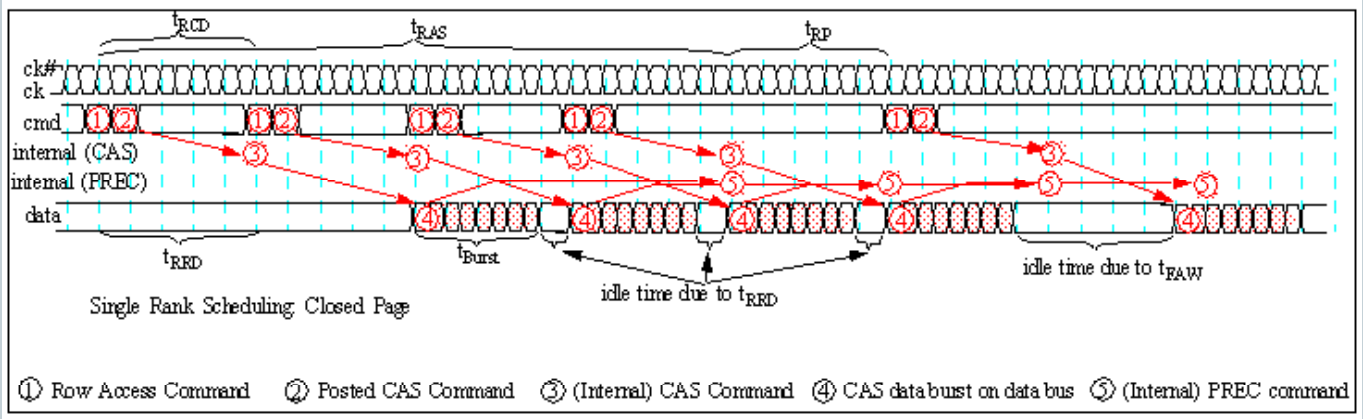

[DRAMsim, UMD]

Experimental Setup – Electronic NoC

System:

- 2cm×2cm chip
- 8×8 Electronic Mesh
  - 28 DRAM Access points (MCs)
  - 2 DIMMs per DRAM AP
- Routers:
  - 1 kb input buffers (per VC)
  - 4 virtual channels
  - 256 b packet size
  - 128 b channels
- 32 nm tech. point (ORION)
  - Normal V_t
  - V_{dd} = 1.0 V
  - Freq = 2.5 GHz

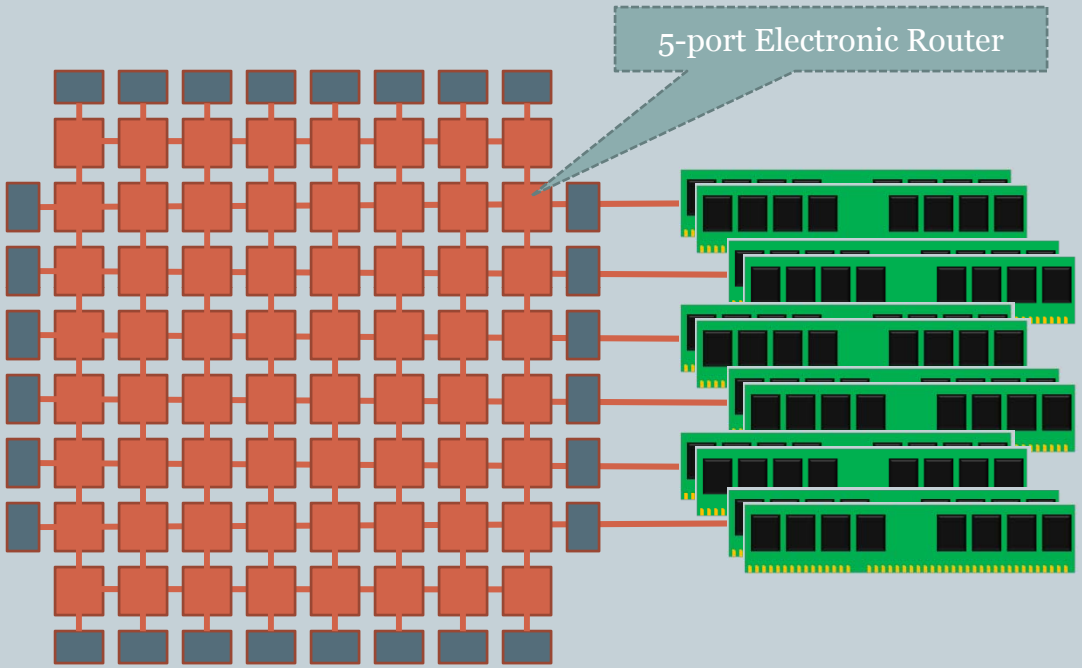

Traffic:

- Random core-DRAM access point pairs
- Random read/write
- Uniform message sizes
- Poisson arrival at 1µs

DRAM:

- Modeled cycle-accurately with DRAMsim [Univ. MD]
- DDR3 (10-10-10) @ 1333 MT/s
- 8 chips per DIMM, 8 banks per Chip, 2 ranks

Experiment Results

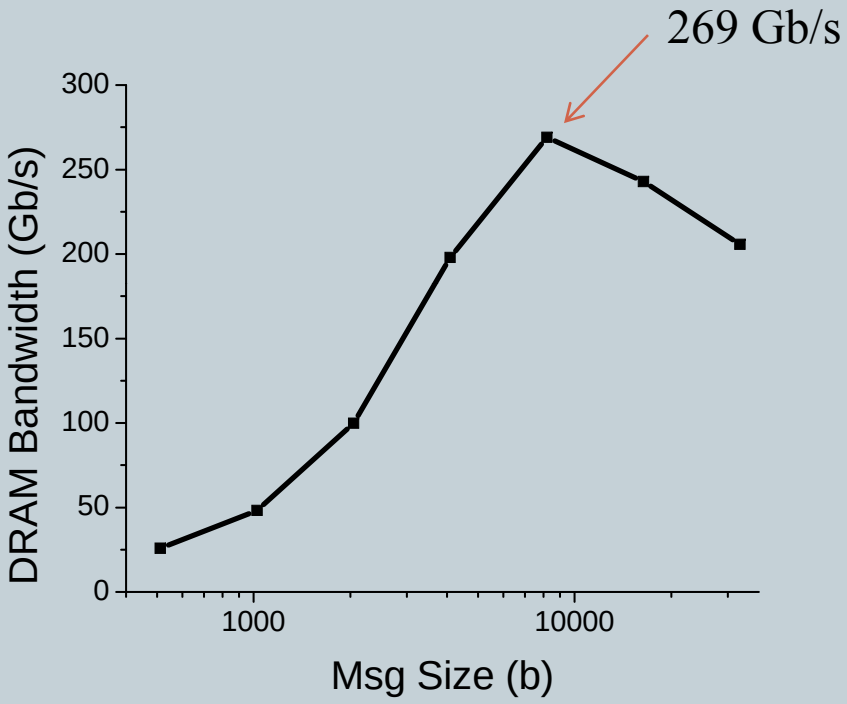

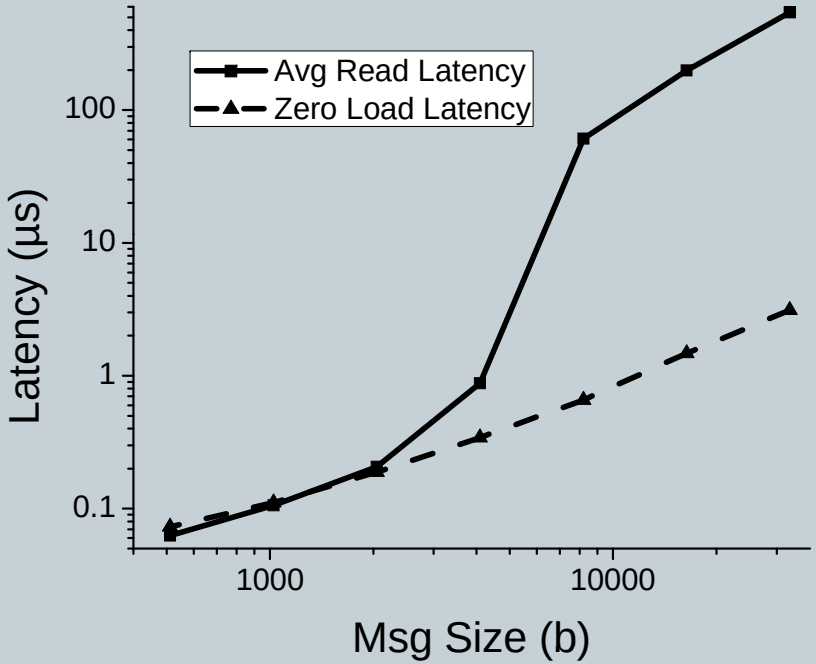

Current

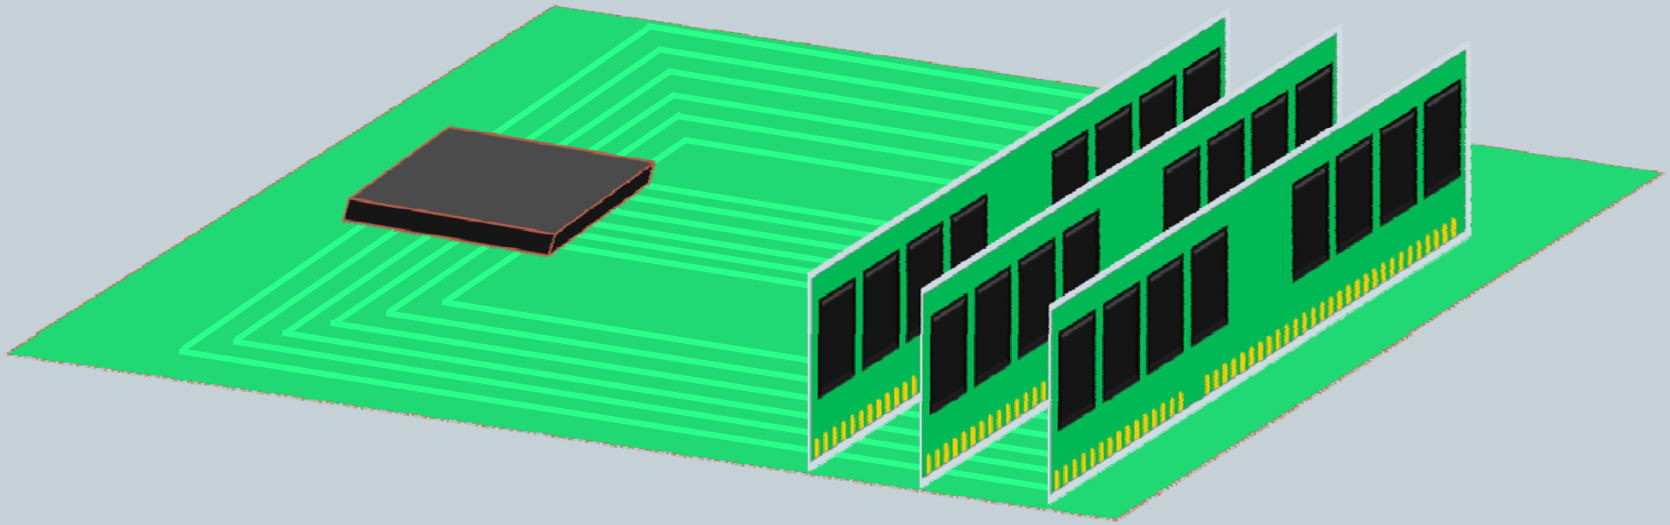

Goal: Optically Integrated Memory

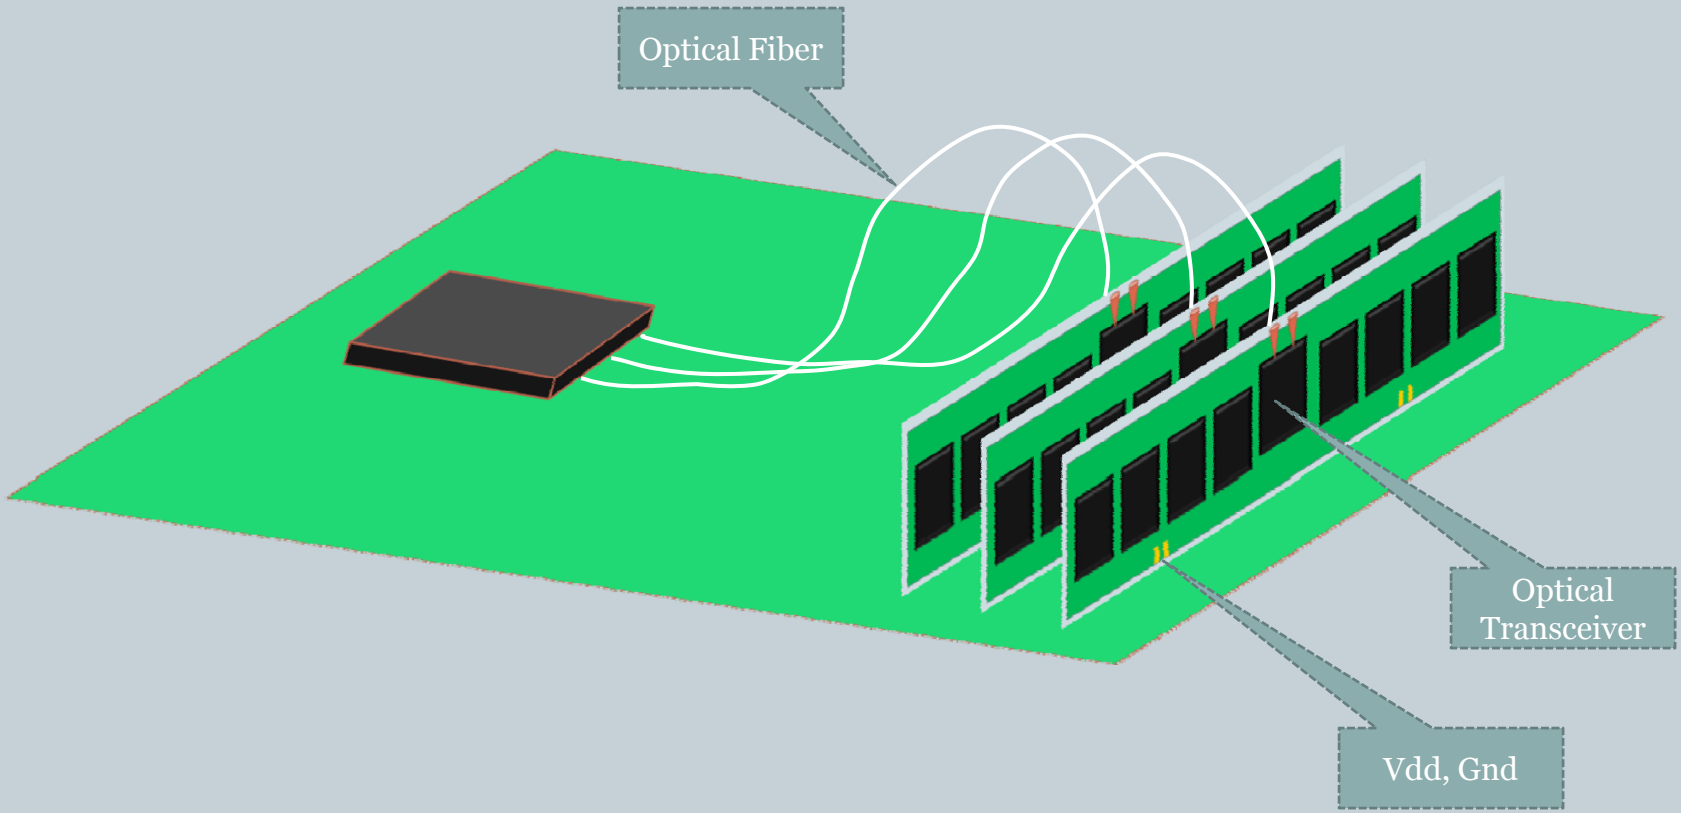

Advantages of Photonics

- Decoupled energy-distance relationship
- No long traces to drive and synch with clock
  - DRAM chips can run faster
  - Less power
- Less pins on DIMM module and going into chip
  - Eventually required by packaging constraints
  - Waveguides can achieve dramatically higher density due to WDM
- DRAM can be arbitrarily distant – fiber is low loss

Hybrid Circuit-Switched Photonic Network

Broadband
1×2 Switch

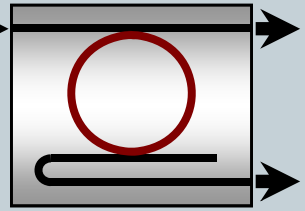

Broadband
2×2 Switch

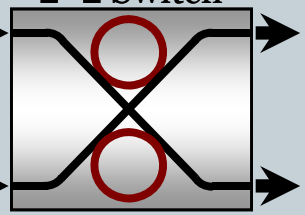

[Cornell, 2008]

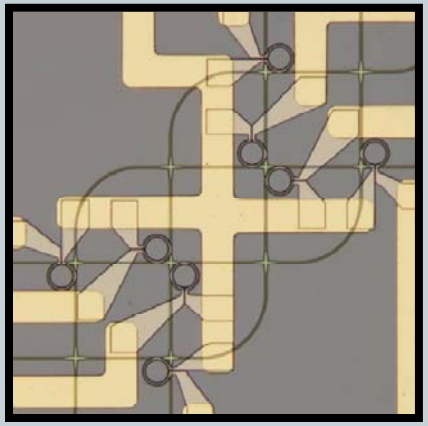

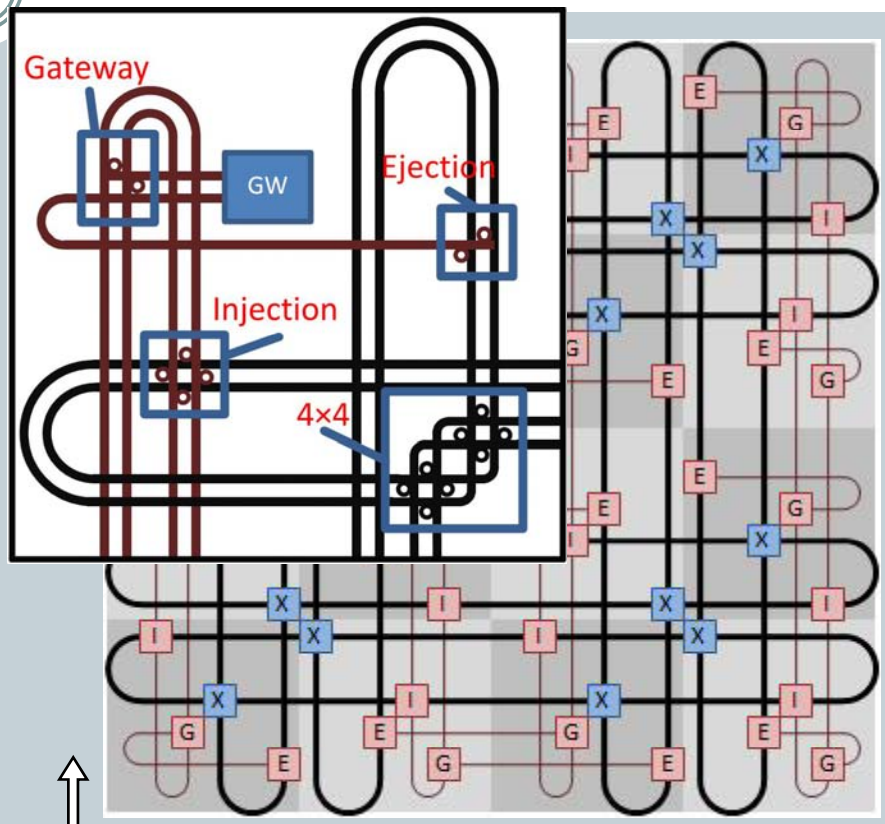

[Shacham, NOCS '07]

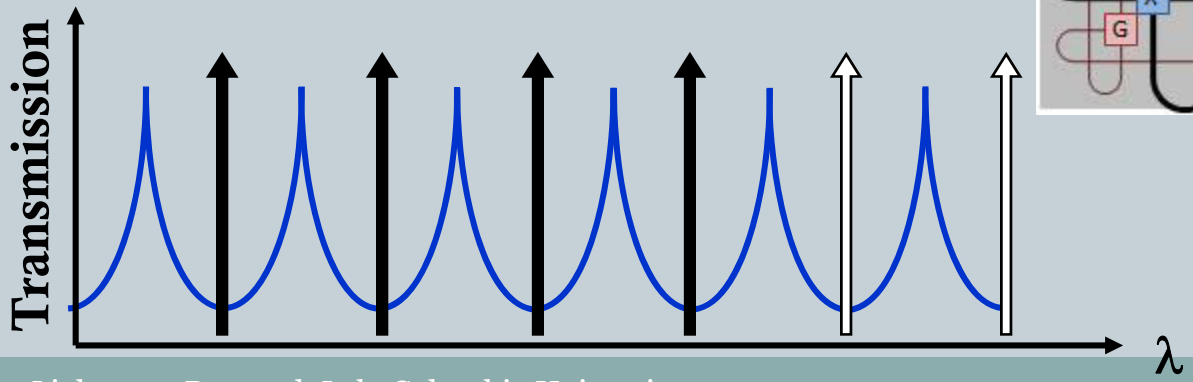

Hybrid Circuit-Switched Photonic Network

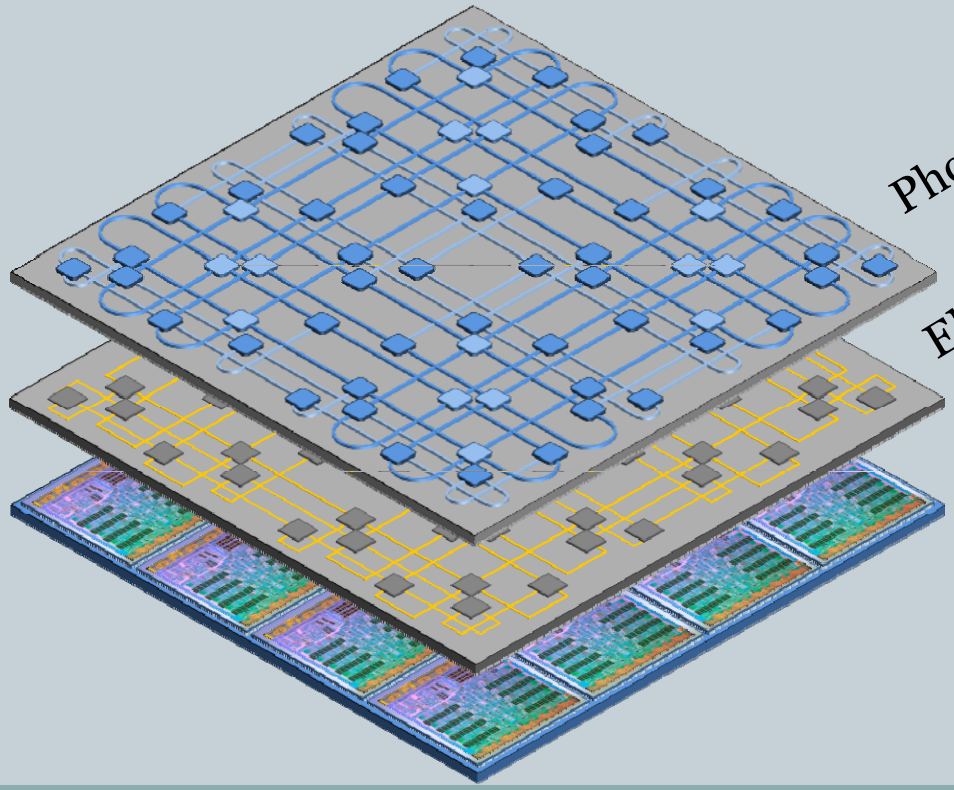

Hybrid Circuit-Switched Photonic Network

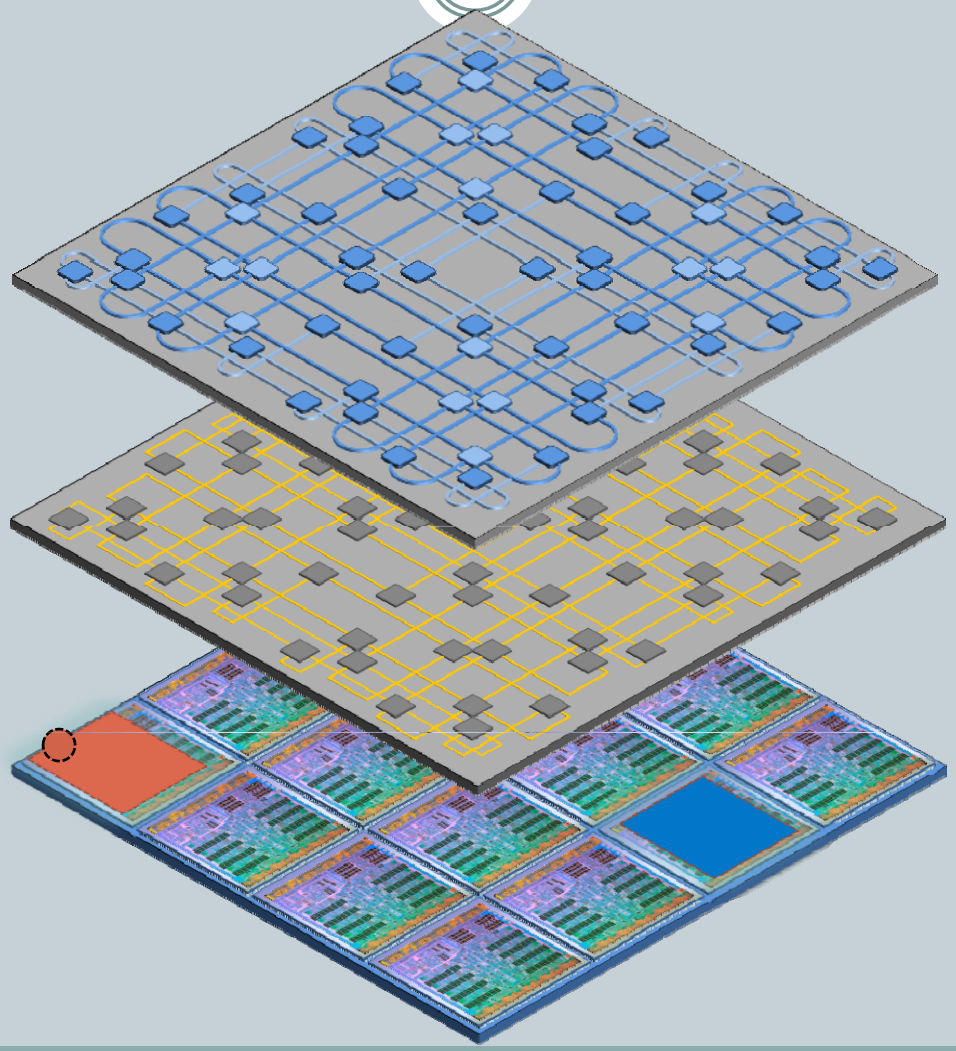

Hybrid Circuit-Switched Photonic Network

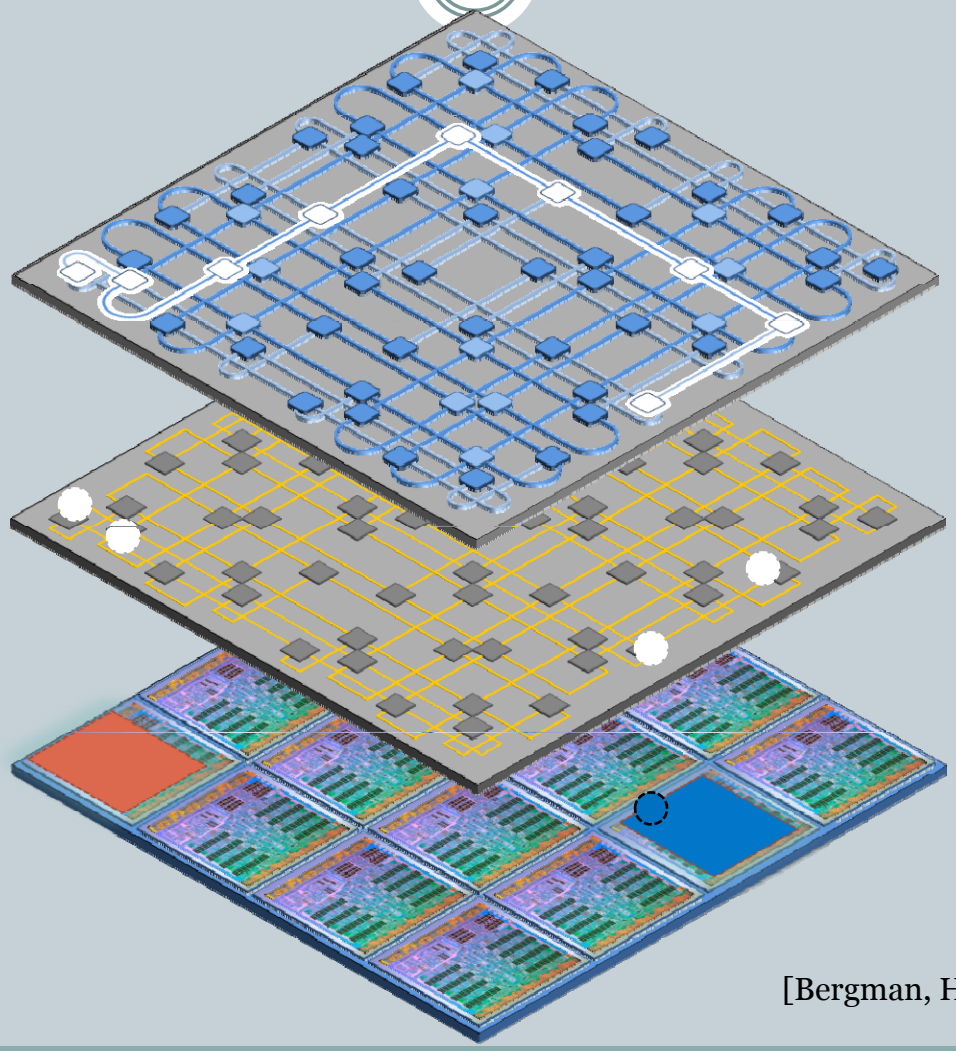

[Bergman, HPEC '07]

Photonic DRAM Access

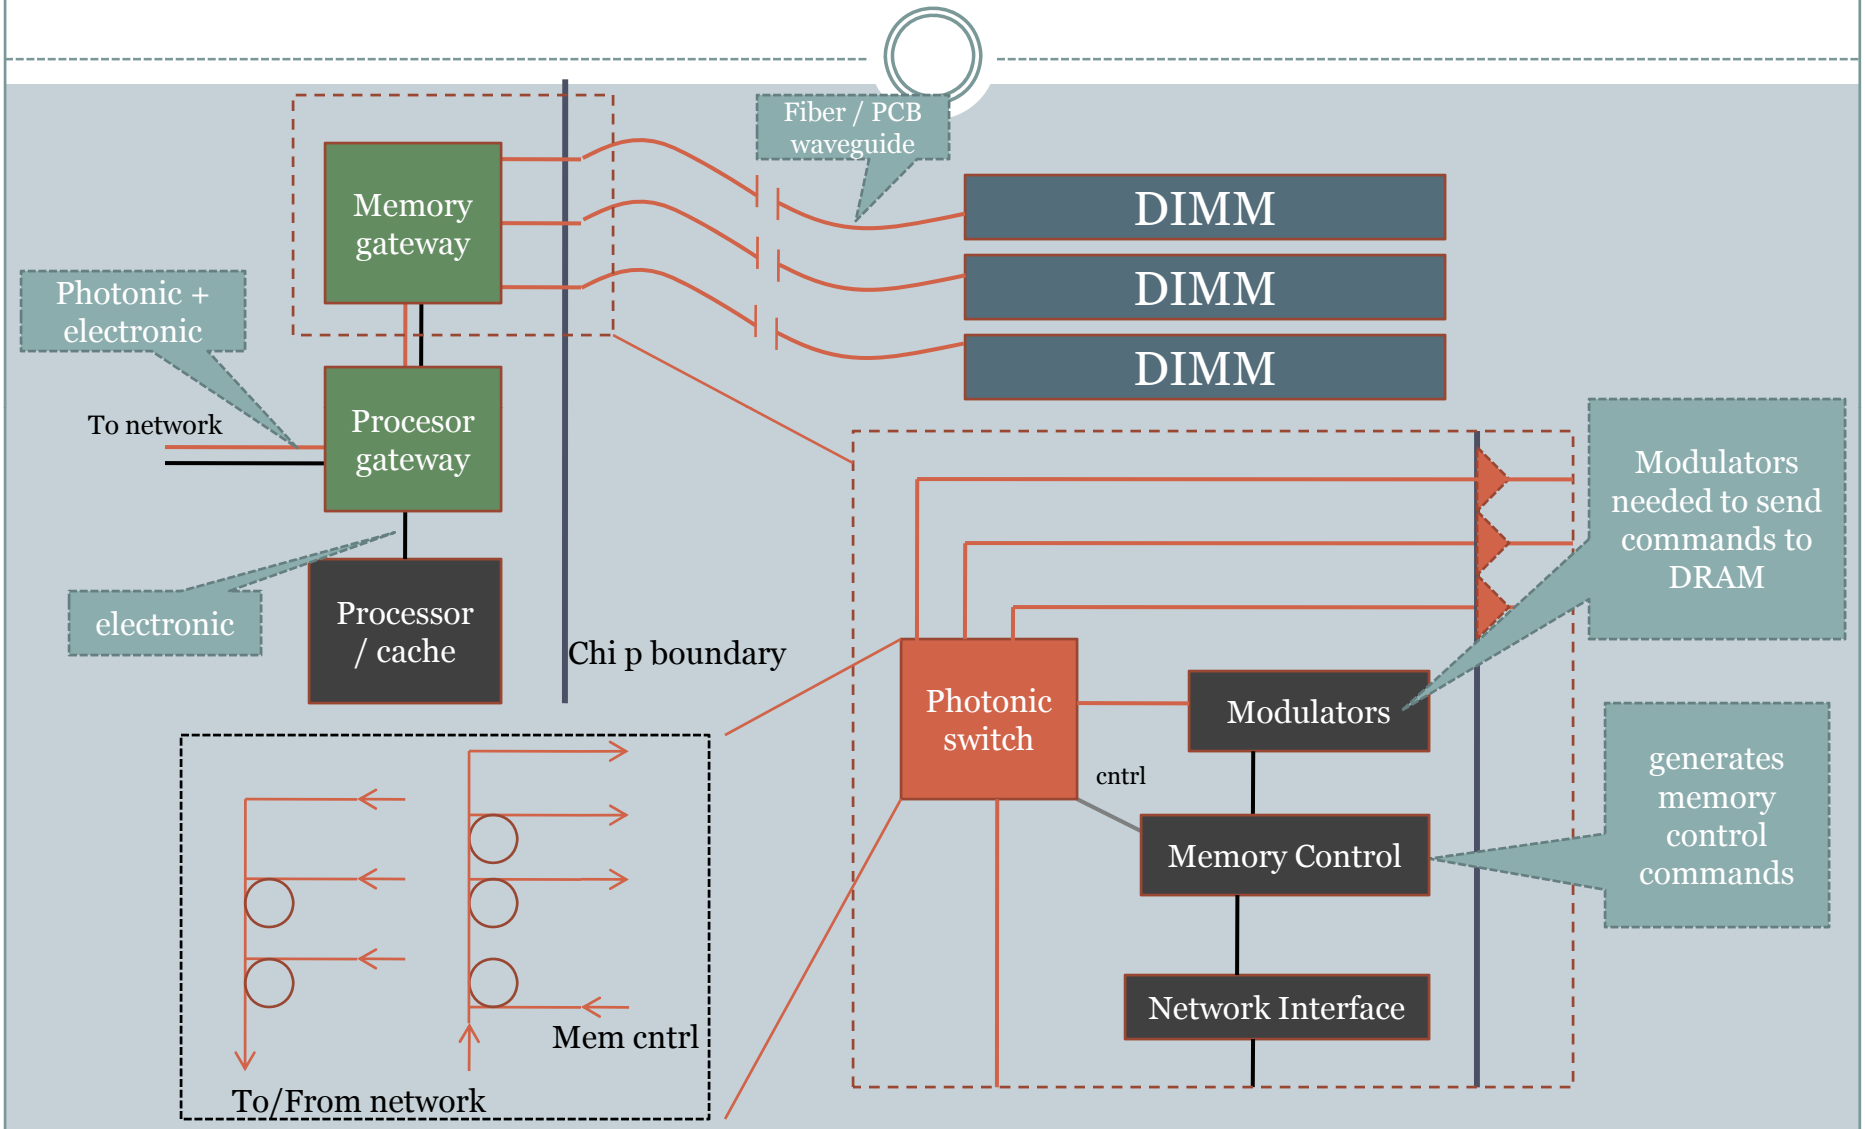

Memory Transaction

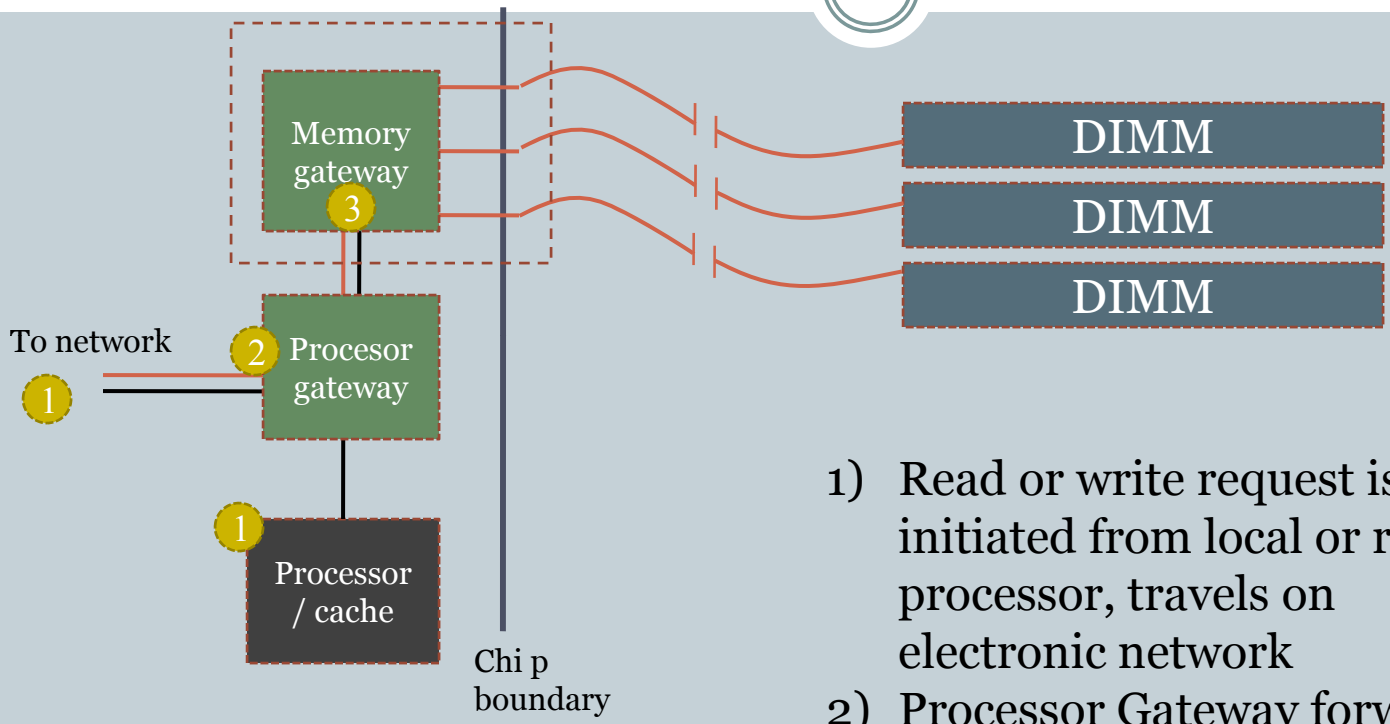

1) Read or write request is initiated from local or remote processor, travels on electronic network
2) Processor Gateway forwards it to Memory gateway
3) Memory gateway receives request

Memory READ Transaction

4) MC receives READ command
5) Switch is setup from modulators to DIMM, and from DIMM to network
6) Path setup travels back to receiving Processor. Path ACK returns when path is set up
7) Row/Col addresses sent to DIMM optically
8) Read data returned optically
9) Path torn down, MC knows how long it will take

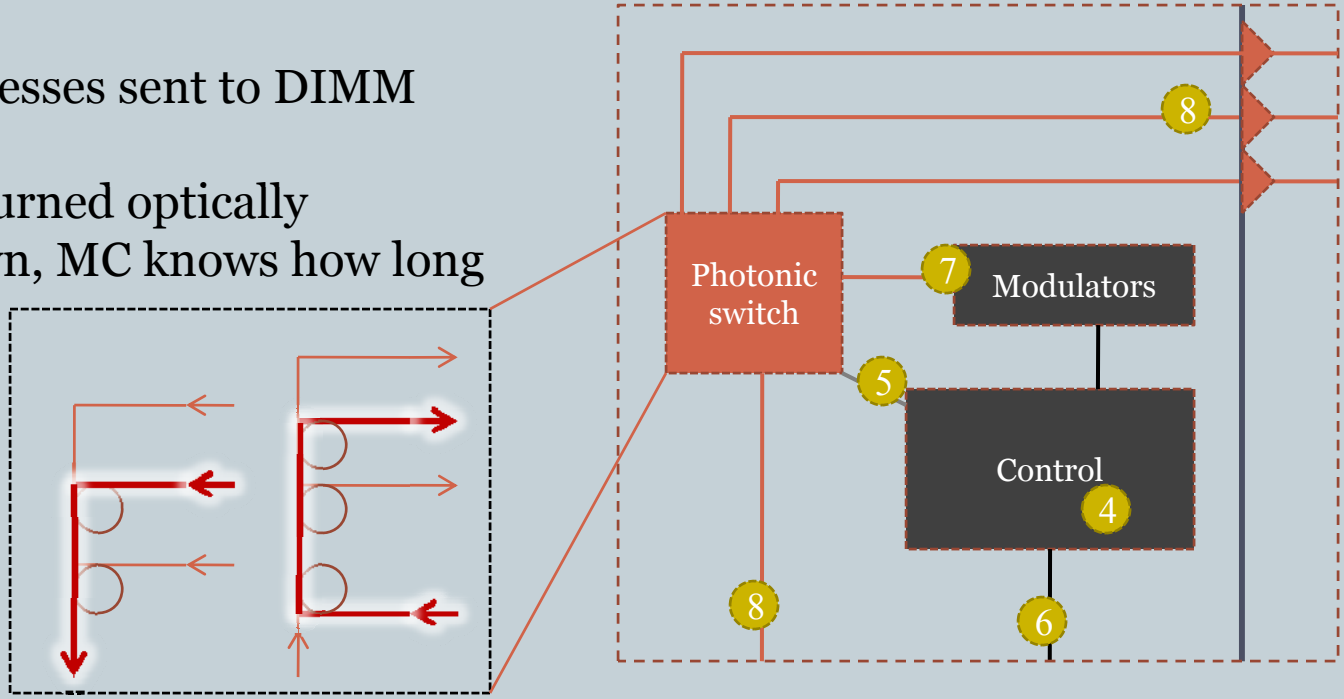

Memory WRITE Transaction

4) MC receives WRITE command, which is also a path setup from the processor to memory gateway
5) Switch is setup from modulators to DIMM
6) Row/Col addresses sent to DIMM
7) Switch is setup from network to DIMM
8) Path ACK sent back to Processor
9) Data transmitted optically to DIMM
10) Path torn down from Processor after data transmitted

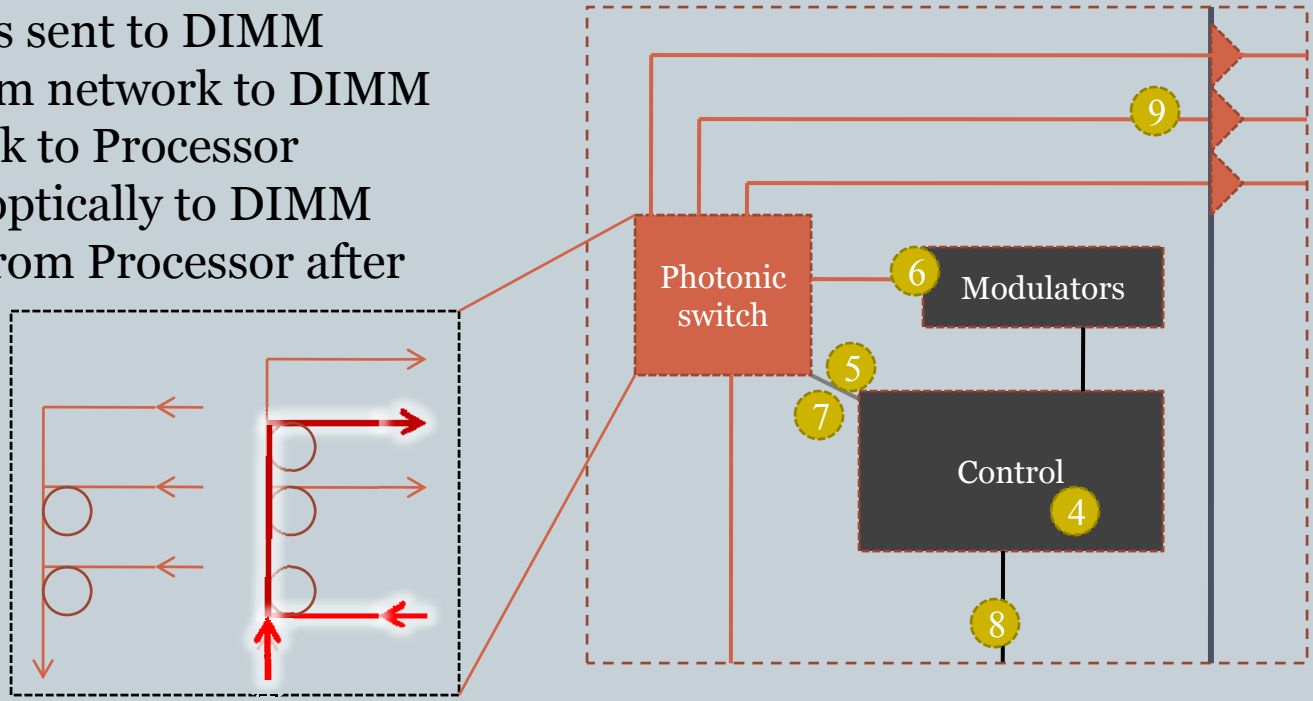

Optical Circuit Memory (OCM) Anatomy

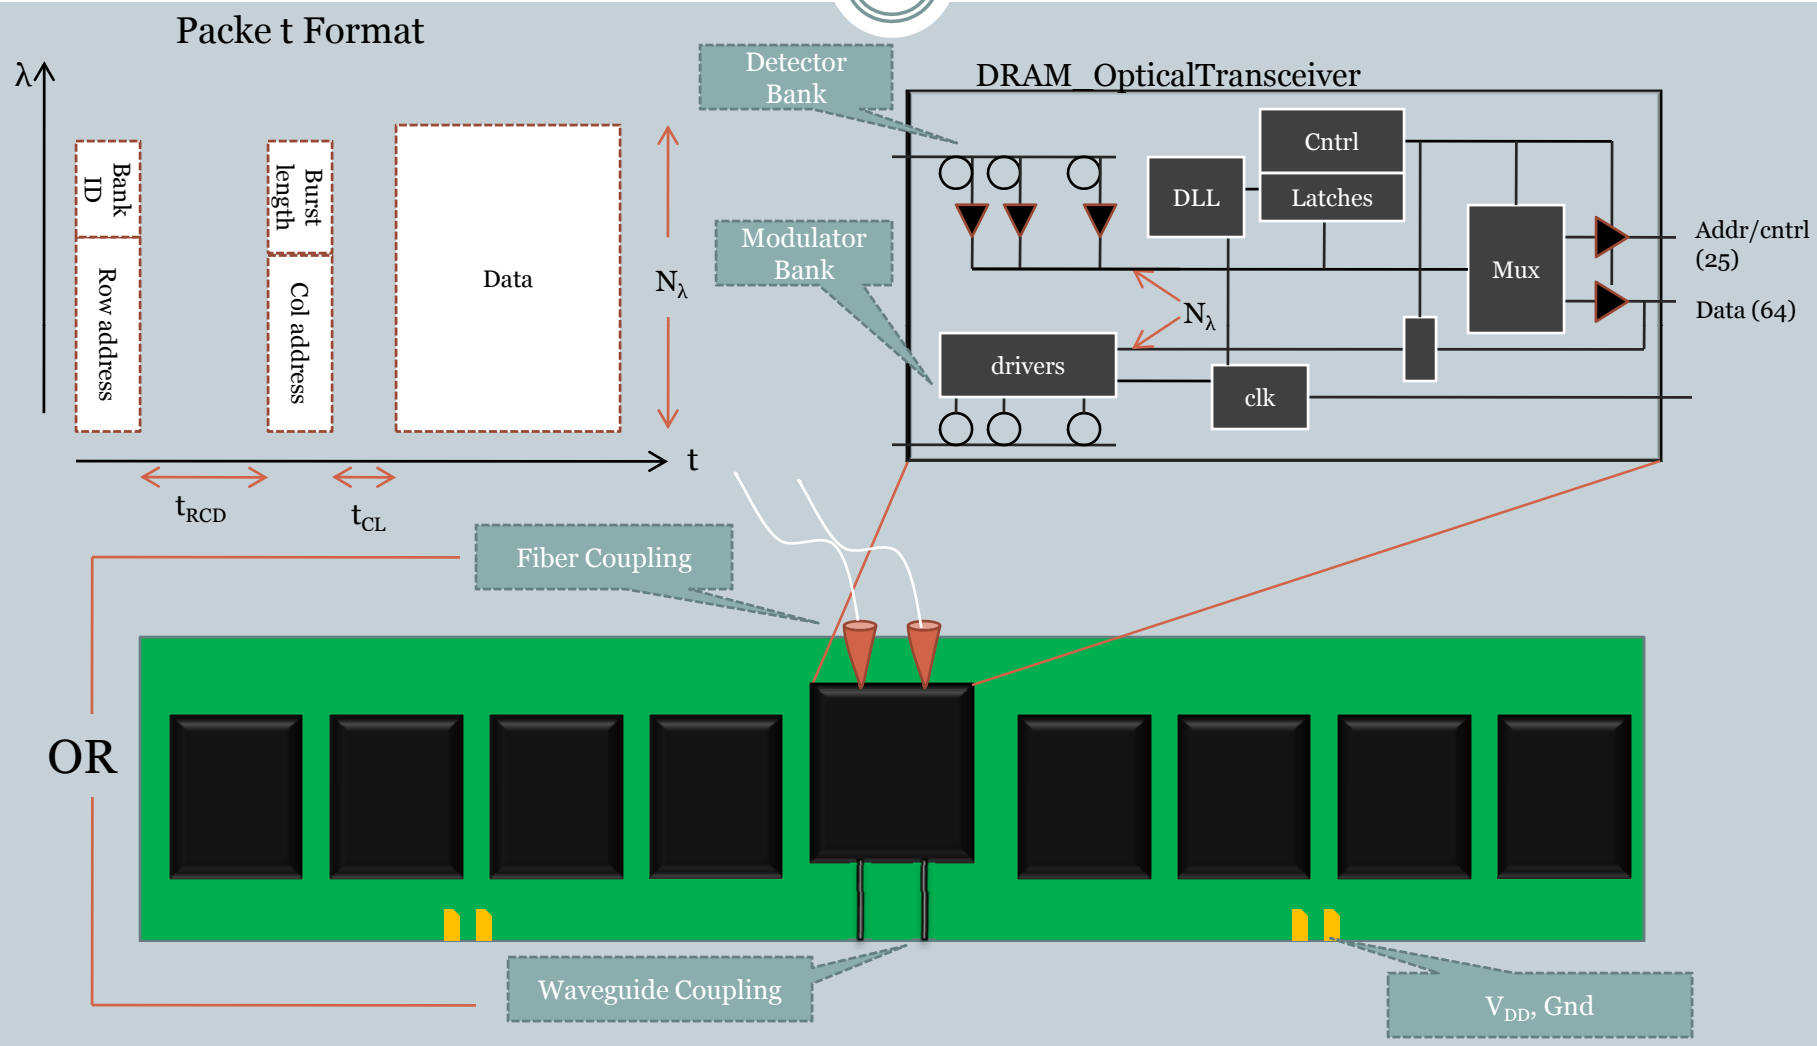

Advantages of Photonics

- Decoupled energy-distance relationship
- No long traces to drive and synch with clock
  - DRAM chips can run faster
  - Less power
- Less pins on DIMM module and going into chip
  - Eventually required by packaging constraints
  - Waveguides can achieve dramatically higher density due to WDM
- DRAM can be arbitrarily distant – fiber is low loss
- Simplified memory control logic – no contending accesses, contention handled by path setup
  - Accesses are optimized for large streams of data

Experimental Setup - Photonic

System:

- 2cm×2cm chip
- 8×8 Photonic Torus
  - 28 DRAM Access points (MCs)
  - 2 DIMMs per DRAM AP
- Routers:
  - 256 b buffers
  - 32 b packet size
  - 32 b channels
- 32 nm tech. point (ORION)
  - High V_t
  - V_{dd} = 0.8 V
  - Freq = 1 GHz
- Photonics - 13λ

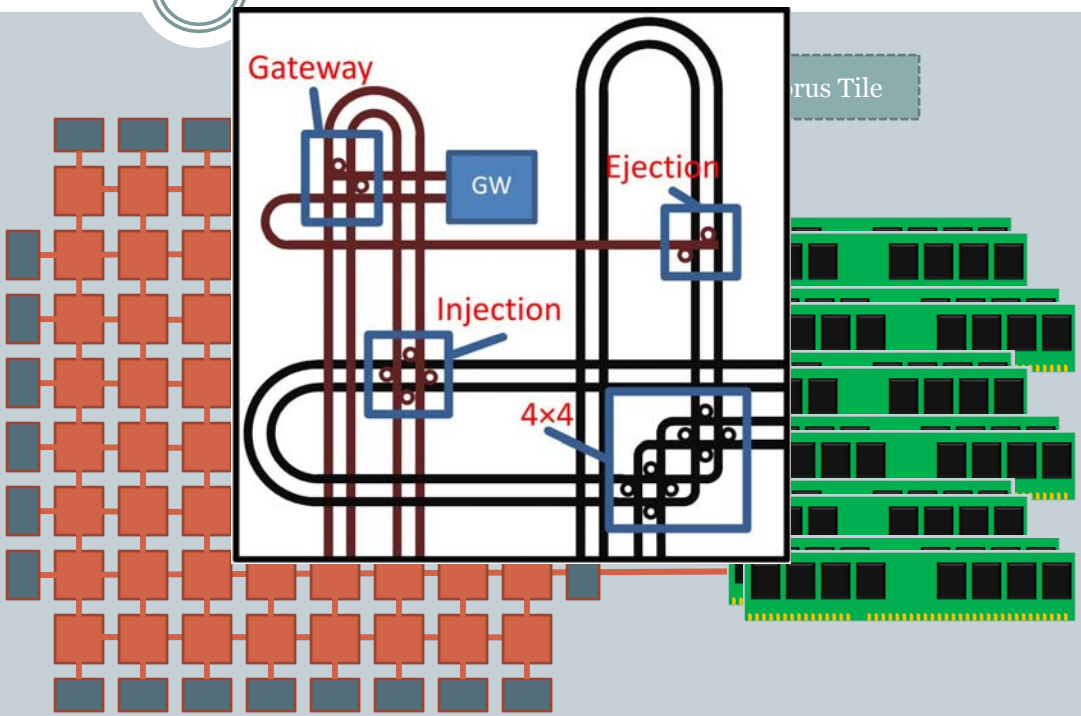

Traffic:

- Random core-DRAM access point pairs
- Random read/write
- Uniform message sizes
- Poisson arrival at 1μs

DRAM:

- Modeled with our event-driven DRAM model
- DDR3 (10-10-10) @ 1600 MT/s
- 8 chips per DIMM, 8 banks per Chip

Performance Comparison

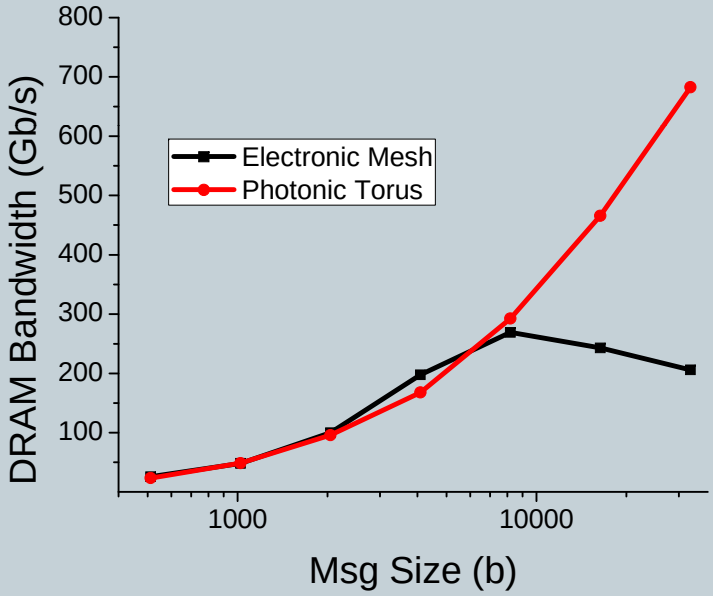

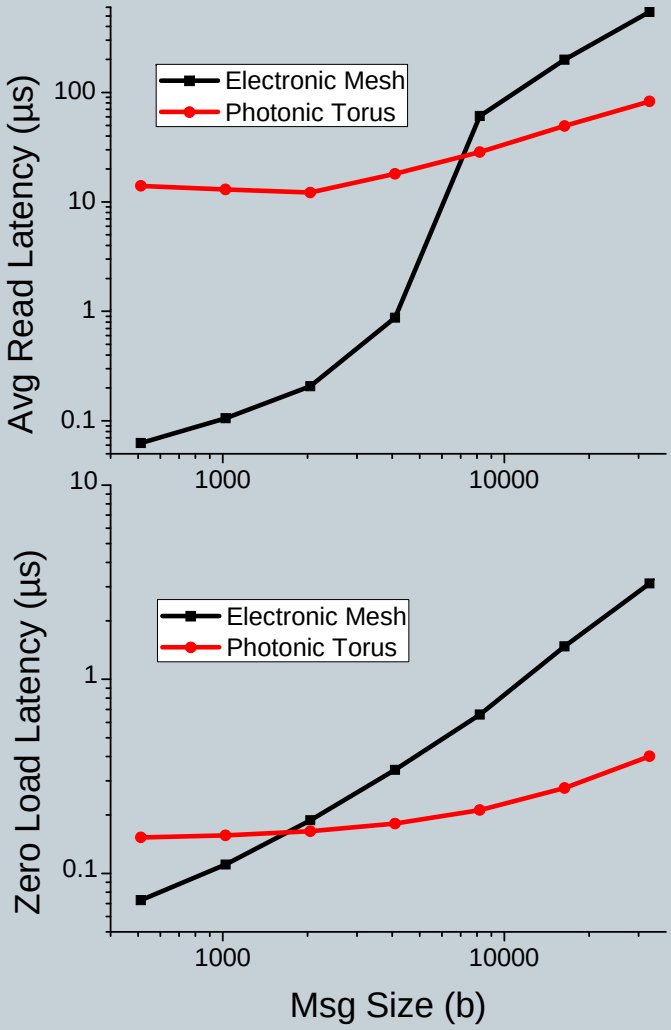

Experiment #2

Random

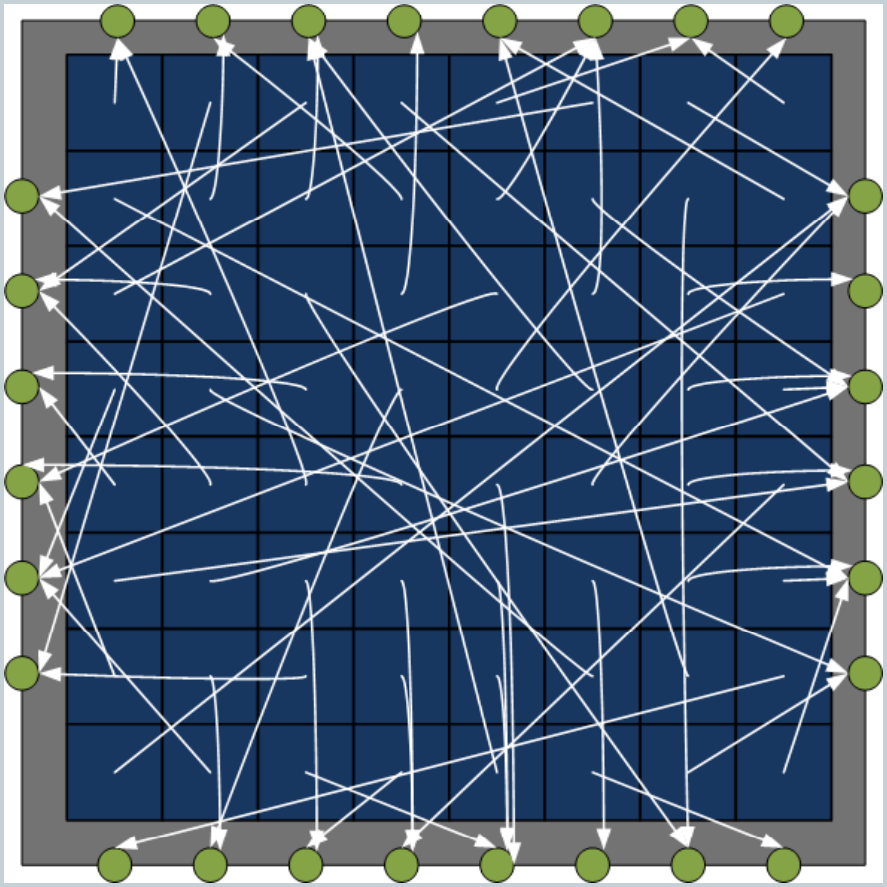

Statically Mapped Address Space

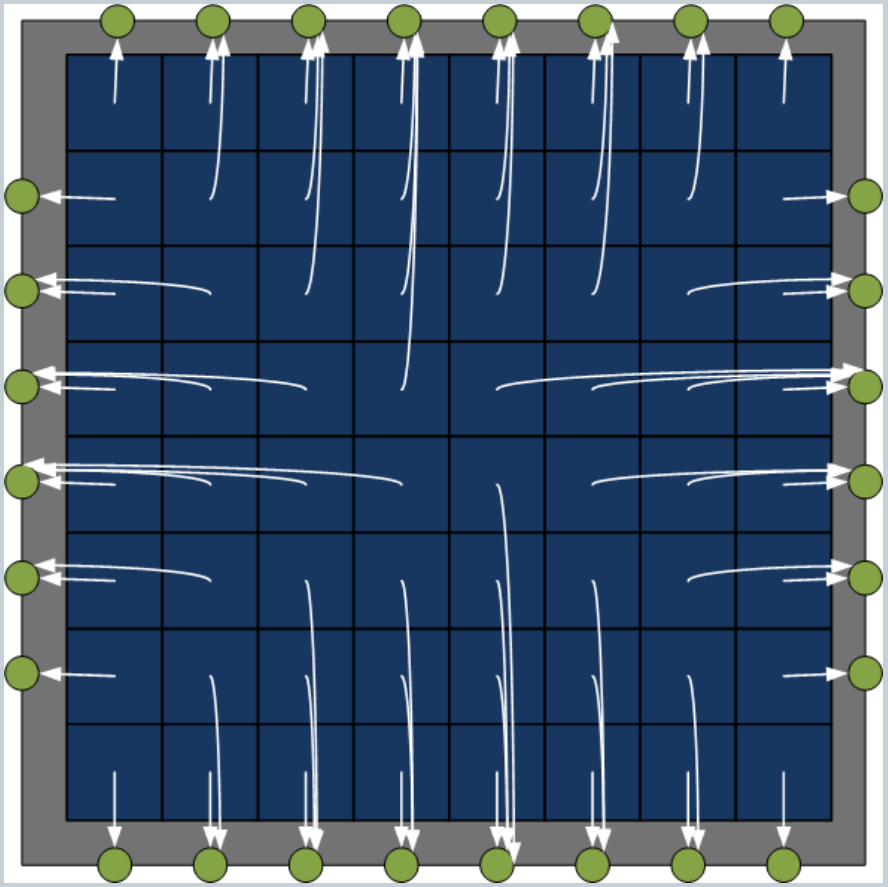

Results

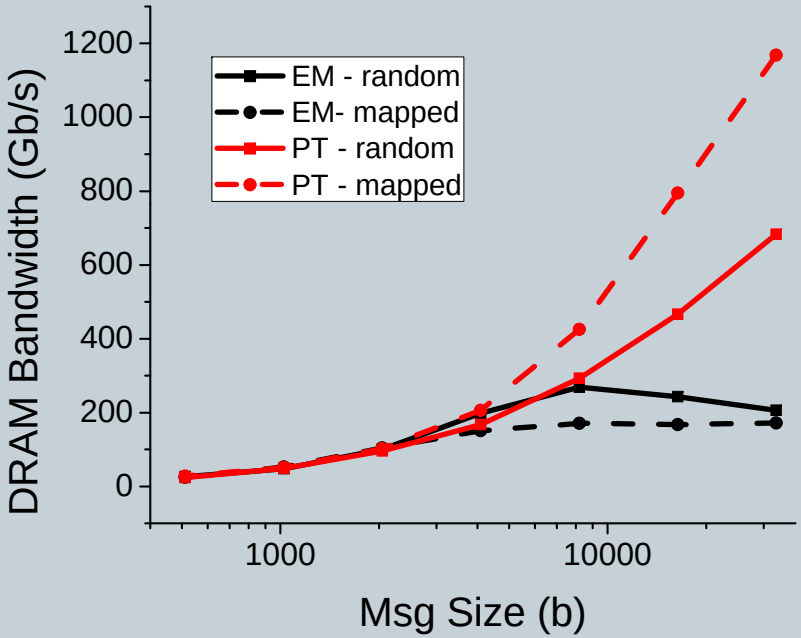
DRAM Bandwidth (Gb/s)
1200
1000
800
600
400
200
0
EM - random
EM- mapped
PT - random
PT - mapped
1000
10000
Msg Size (b)

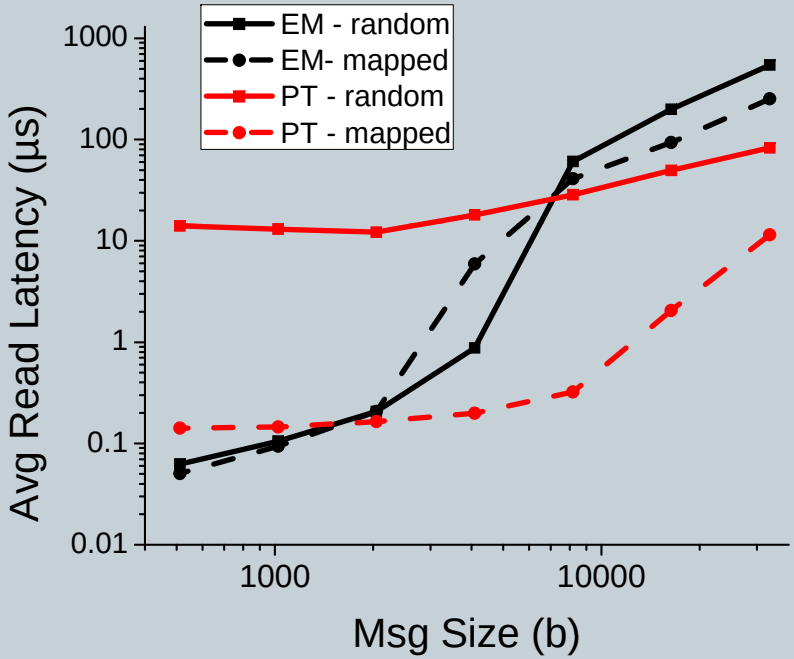
Avg Read Latency (µs)
1000
100
10
1
0.1
0.01
EM - random
EM- mapped
PT - random
PT - mapped
1000
10000
Msg Size (b)

Network Energy Comparison

Electronic Mesh

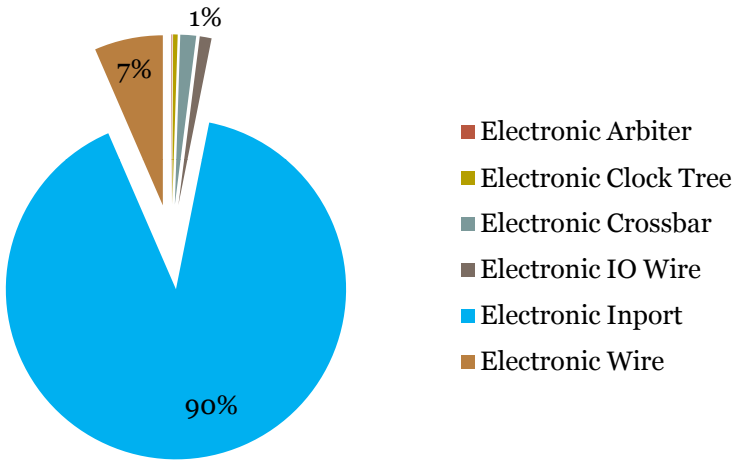

Power = 13.3 W

Photonic Torus

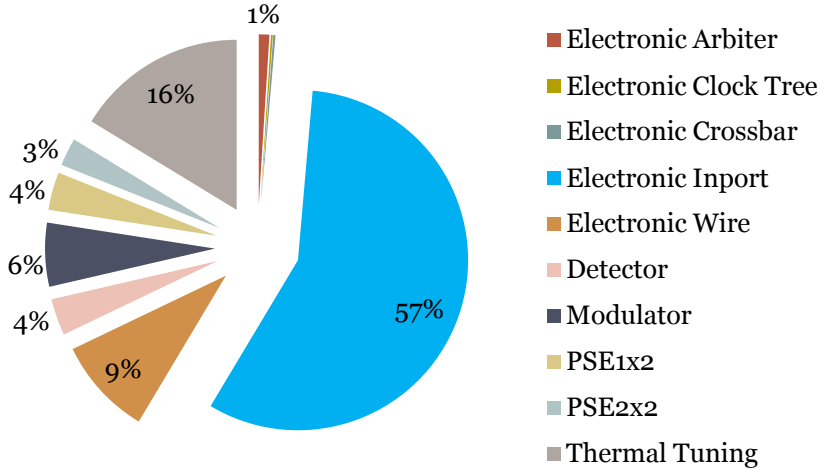

Power = 0.42 W

Total Power = 2.53 W

(Including laser power)

Summary

- Extending a photonic network to include access to DRAM looks good for many reasons:
  - Circuit-switching allows large burst lengths and simplified memory control, for increased bandwidth.
  - Energy efficient end-to-end transmission
  - Alleviates pin count constraints with high-density waveguides

PhotoMAN