Future Directions for On-Chip Interconnection Networks

William J. Dally
Computer Systems Laboratory
Stanford University

OCIN Workshop

December 7, 2006

State of Off-Chip Networks

Technology Trends...

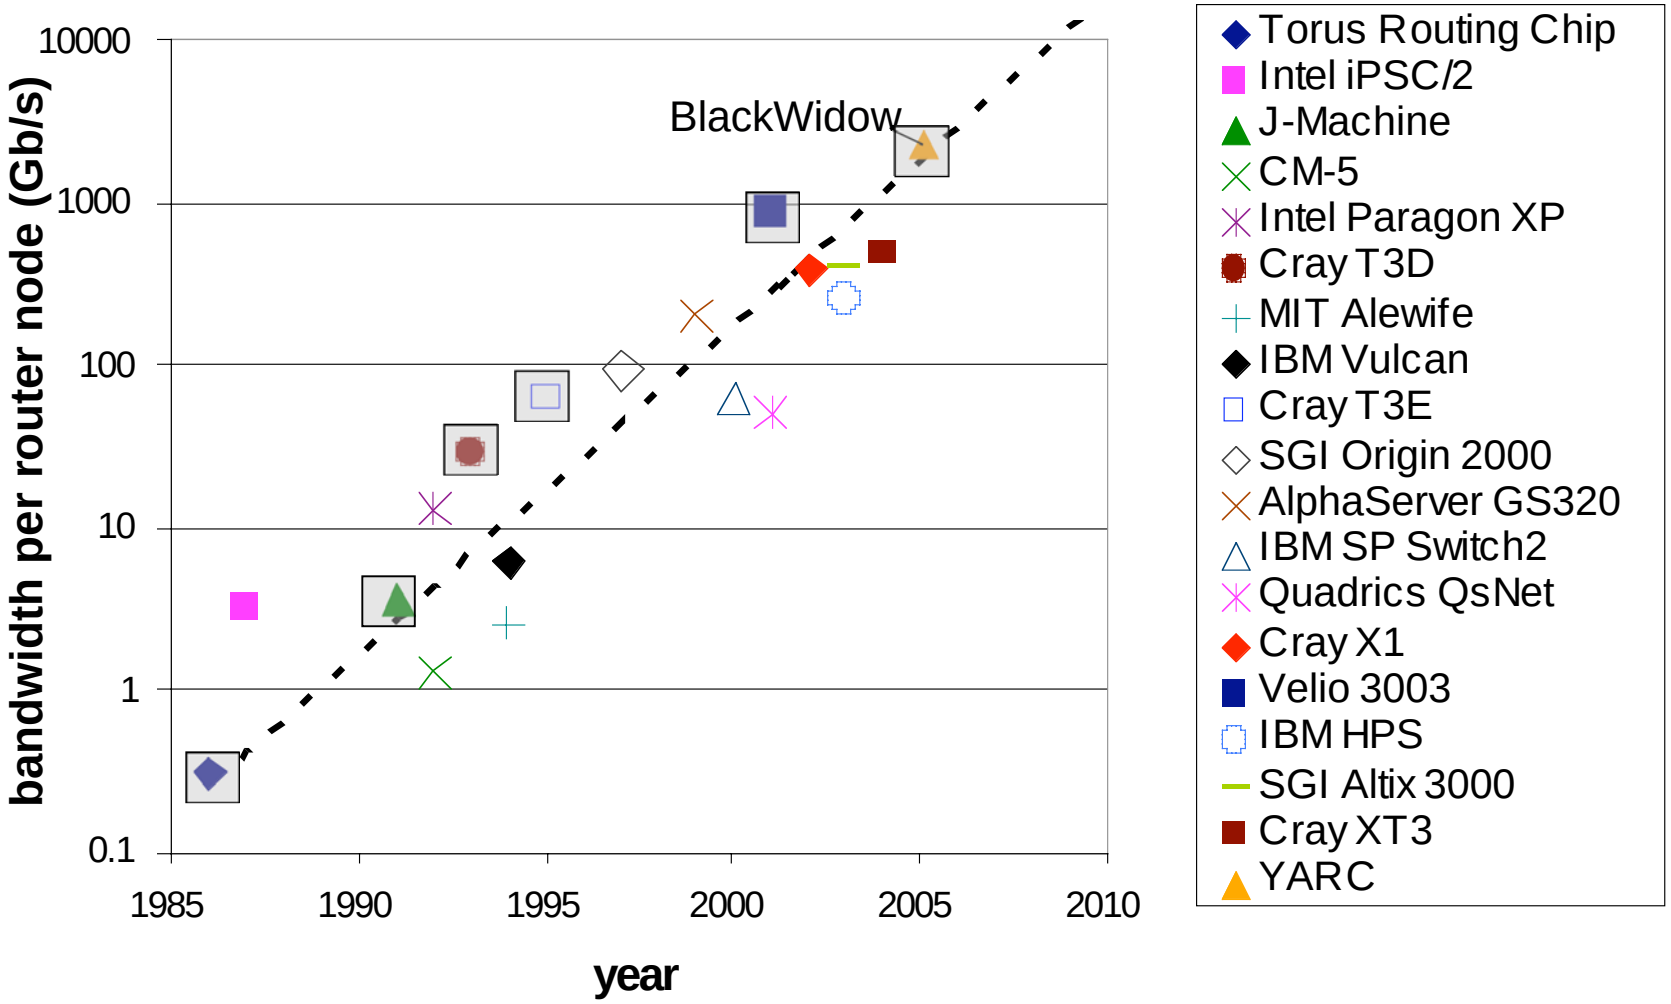

Some History

MARS Router
1984

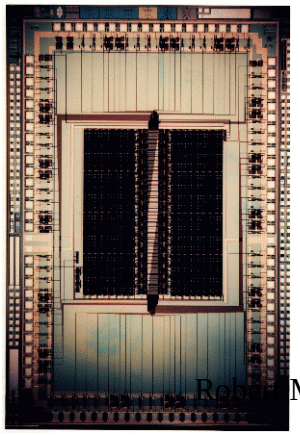

Torus Routing Chip
1985

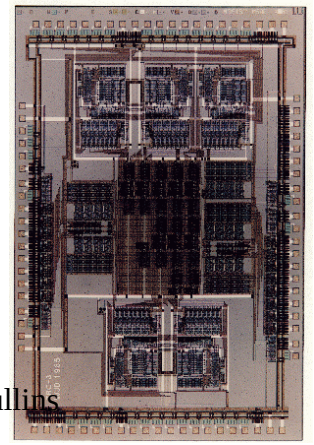

Network Design Frame
1988

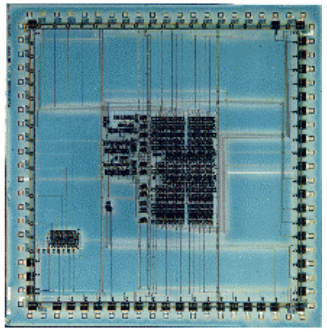

MDP 1991

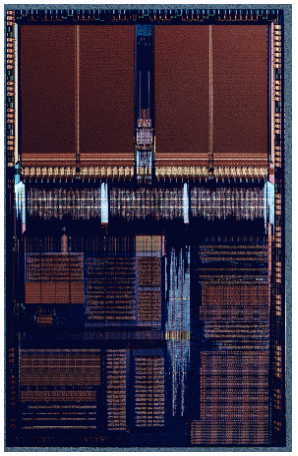

Reliable Router
1994

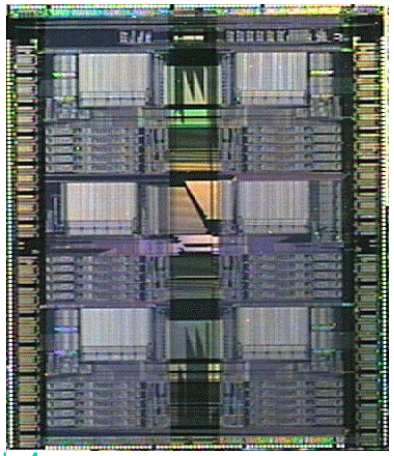

MAP 1998

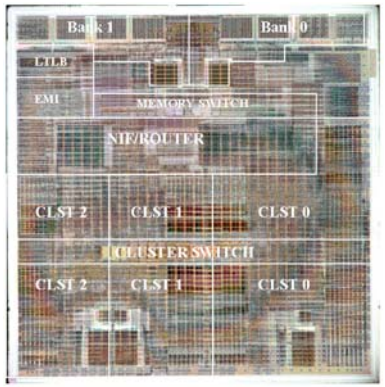

Imagine
2002

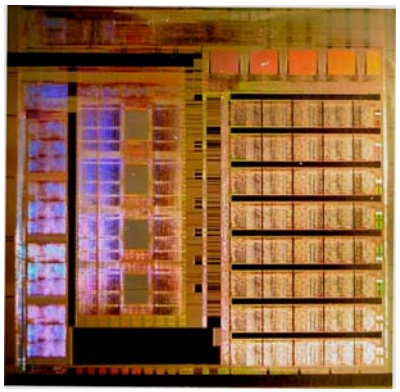

YARC
2006

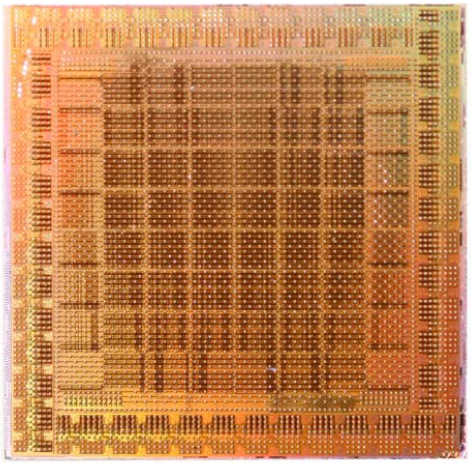

Some very good books

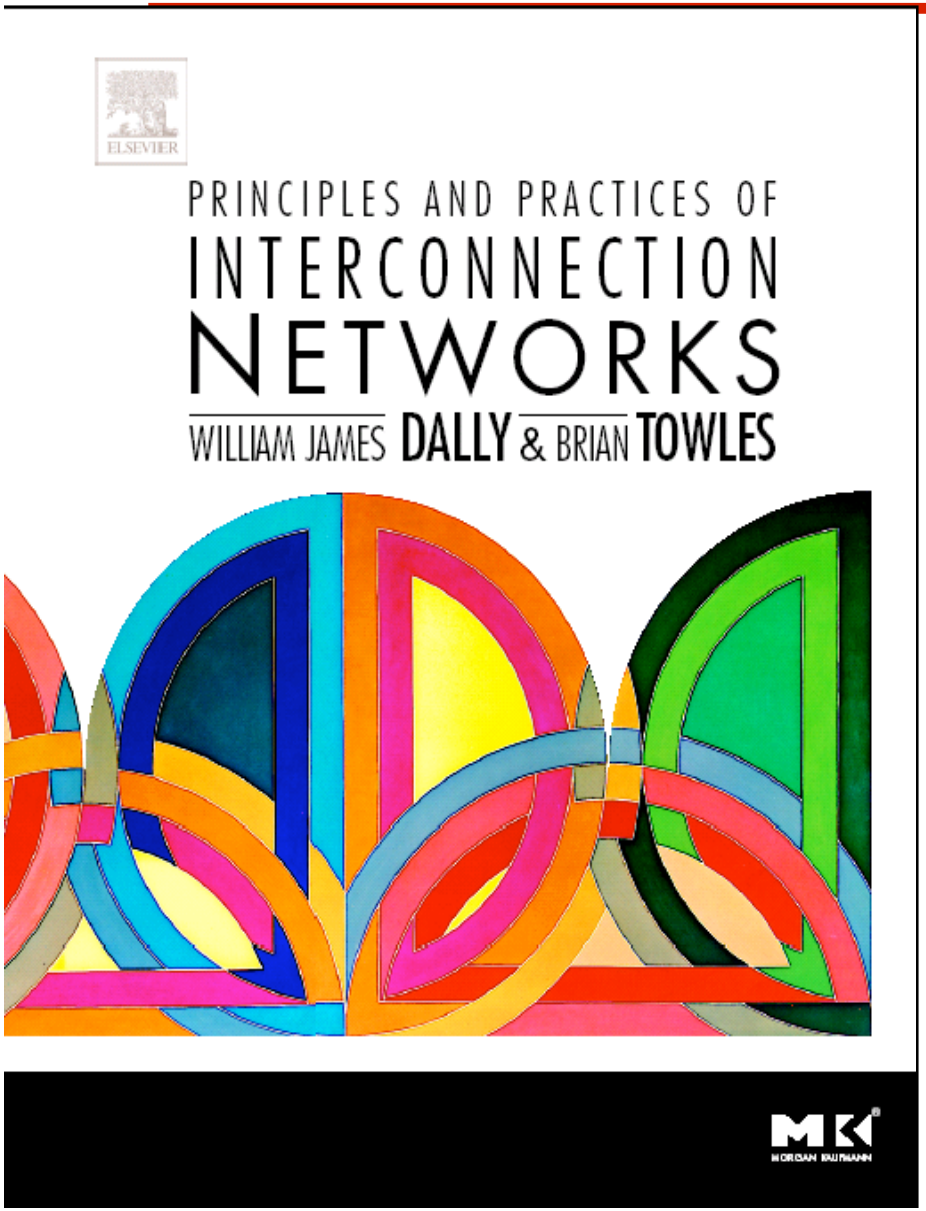

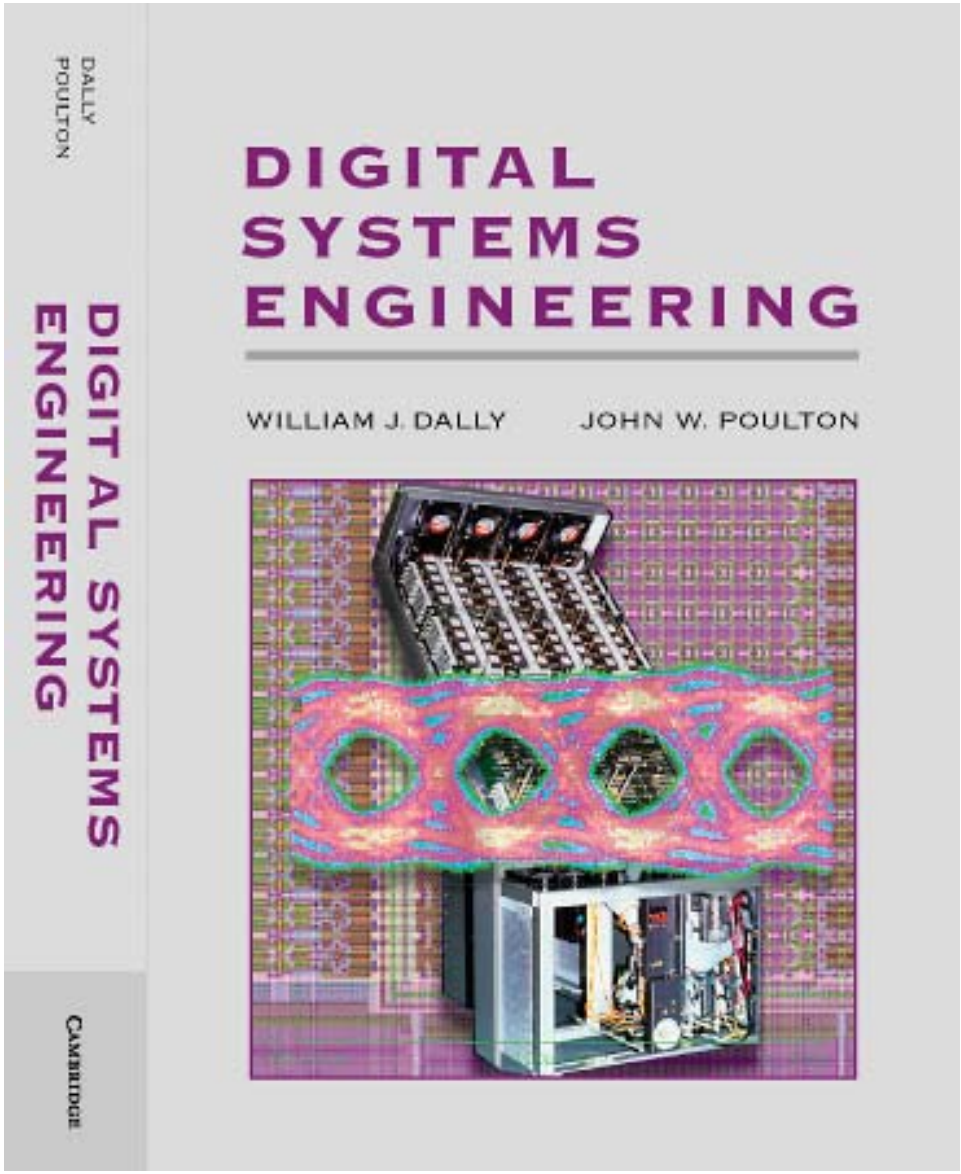

Summary of Off-Chip Networks*

- Topology
  - Fit to packaging and signaling technology
  - High-radix - Clos or FlatBfly gives lowest cost
- Routing
  - Global adaptive routing balances load w/o destroying locality
- Flow control
  - Virtual channels/virtual cut-through

*oversimplified

So, what's different about on-chip networks?

What's different about OCINs?

- Cost
  - Off: cost is channels - pins, connectors, cables, optics
  - On: cost is storage and switches, wires plentiful
  - Drives networks with many wide channels, few buffers
    - Low-radix augmented mesh
- Channel Characteristics
  - On-chip RC lines - need a repeater every 1mm
  - Short distance - low latency
  - Can put logic in repeaters, motivates low-latency routers
- Workload
  - CMP cache traffic
  - SoC isochronous flows
- Differences motivate some surprising differences from on-chip networks

Example CMP OCIN

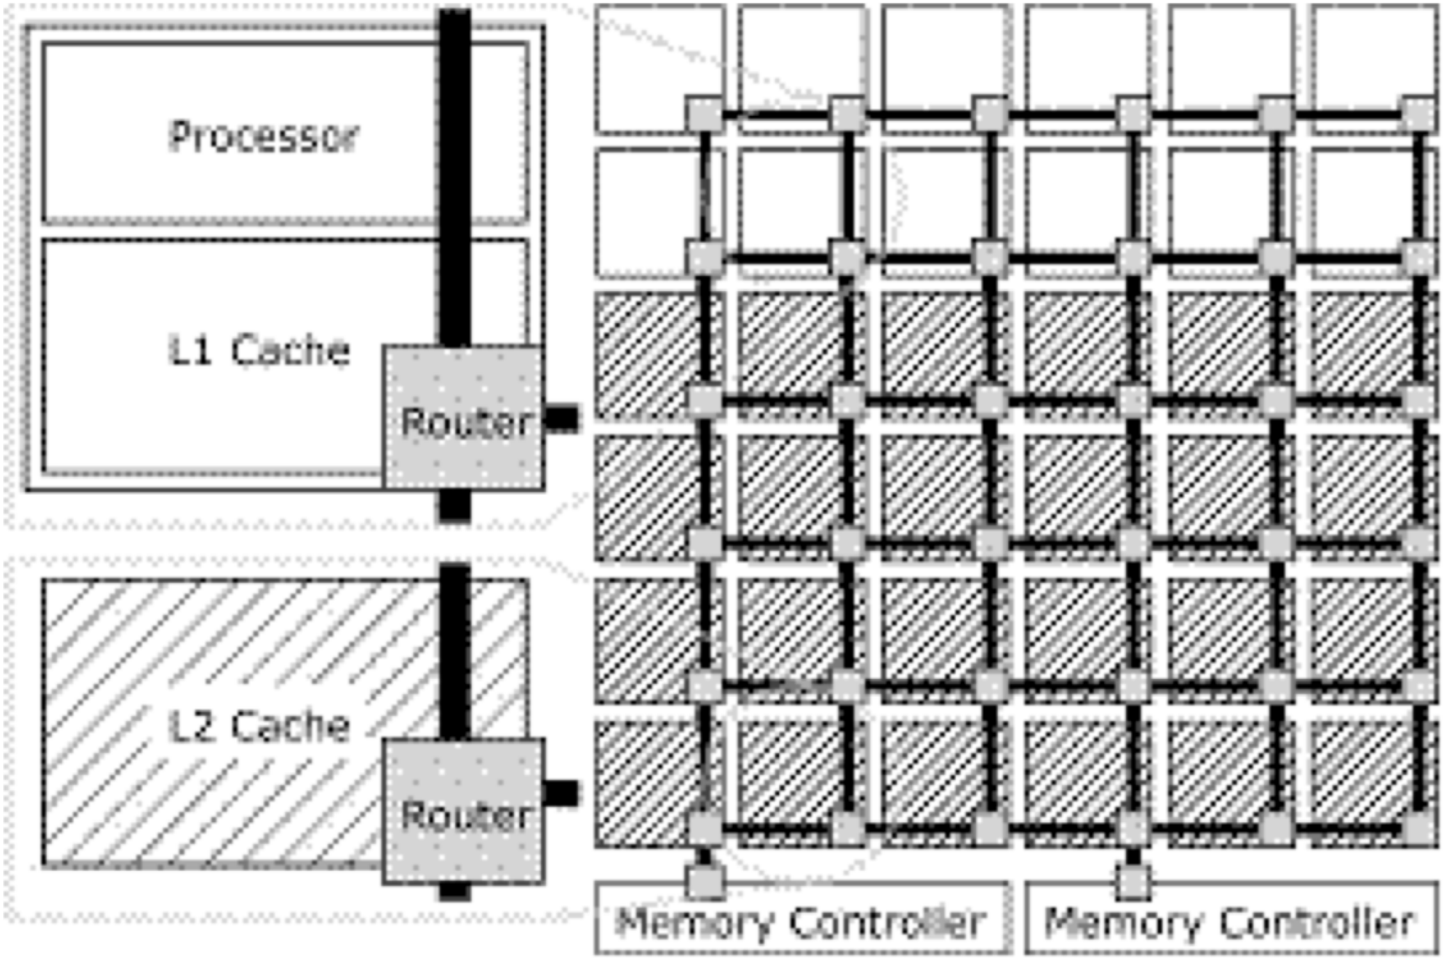

On-Chip Interconnection Network

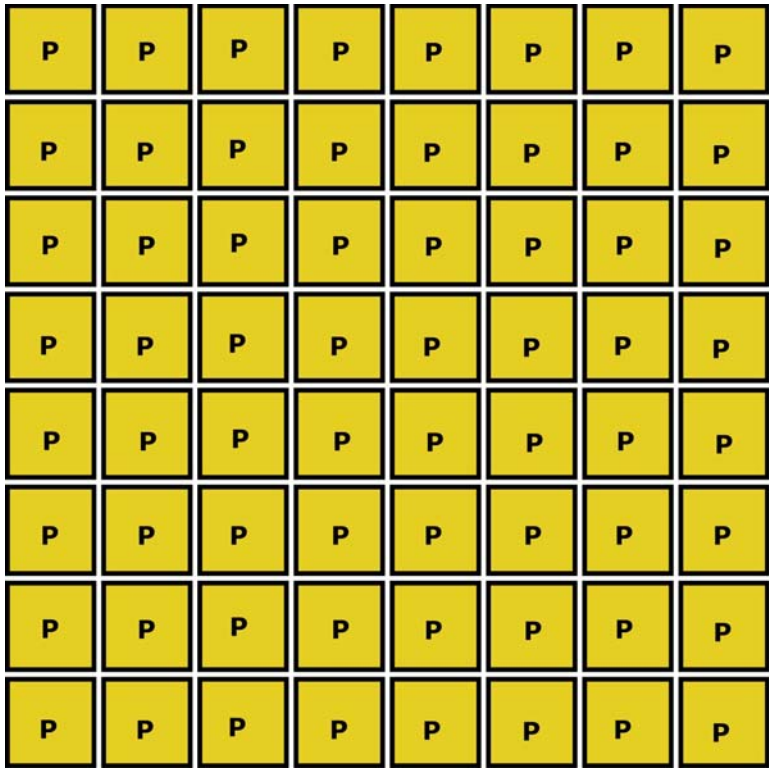

System = Processor Tiles

Source: Balfour and Dally, ICS 06

On-Chip Interconnection Network (2)

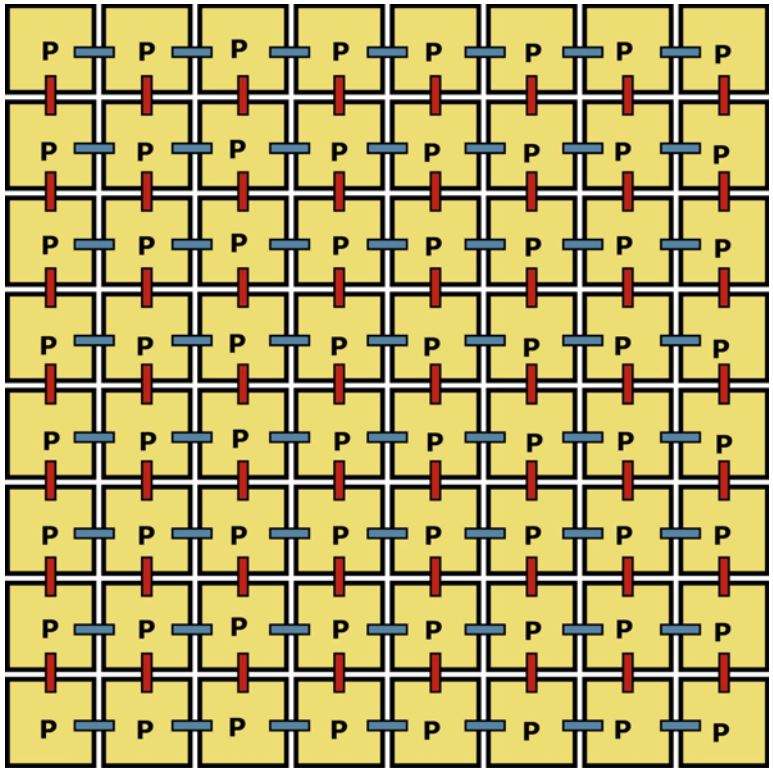

System = Processor Tiles + Channels

Source: Balfour and Dally, ICS 06

Interconnection Network (3)

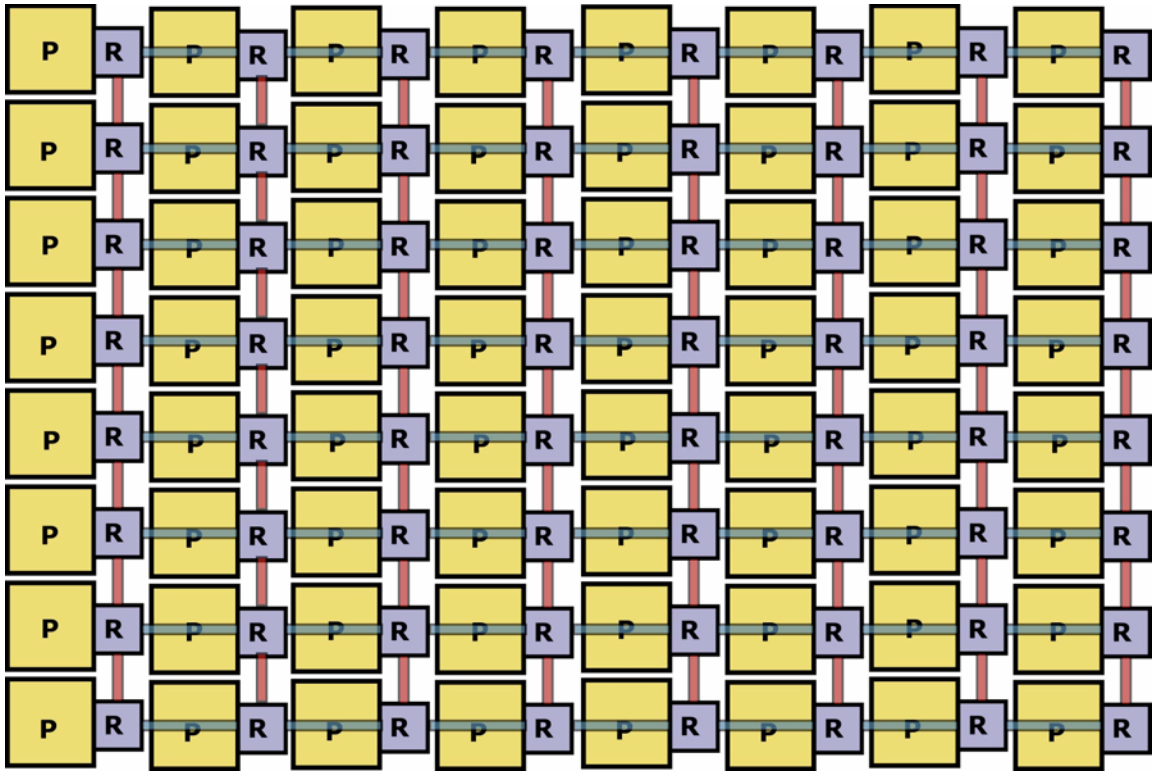

System = Processor Tiles + Channels + Routers

Source: Balfour and Dally, ICS 06

Router Architecture

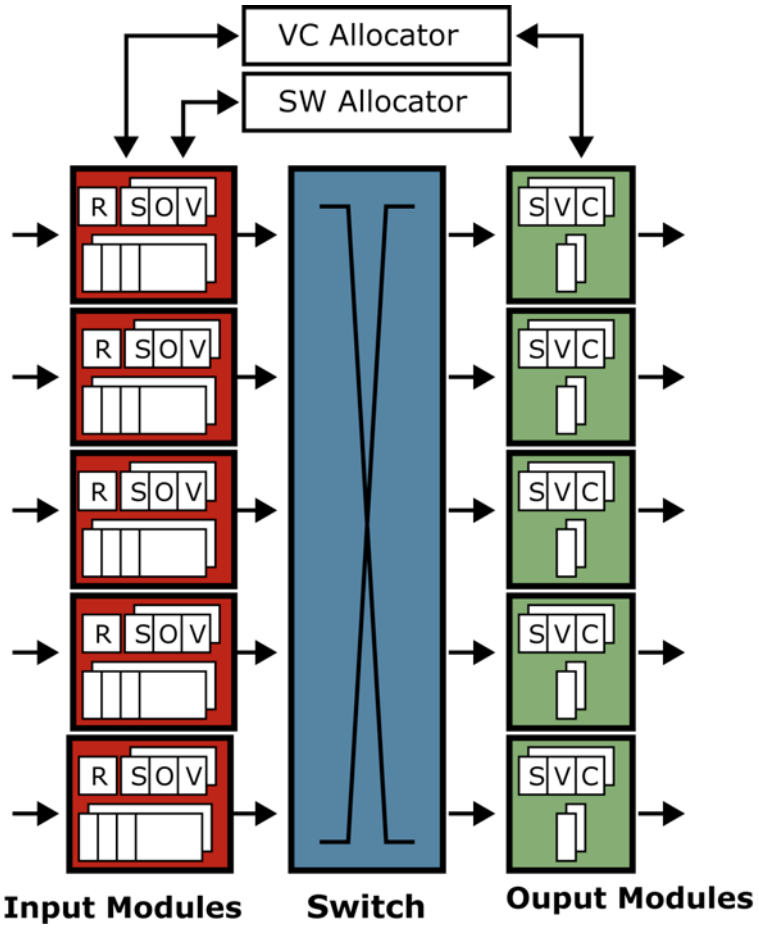

- Input-queued
- Virtual Channel
- Speculative Pipeline

Source: Balfour and Dally, ICS 06

Router Area

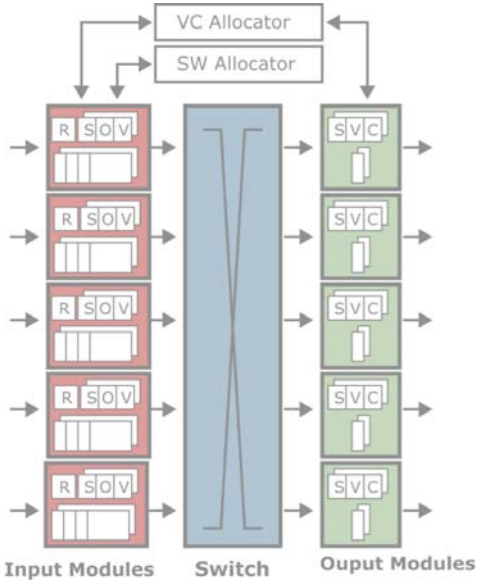

Accurate modeling requires floorplan

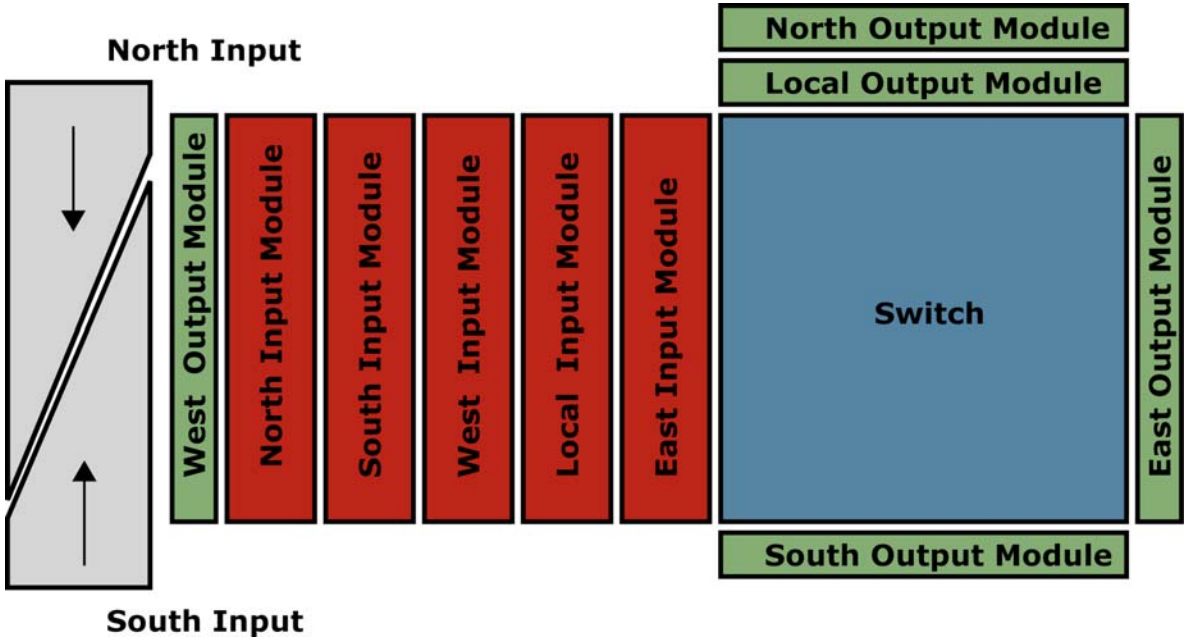

Source: Balfour and Dally, ICS 06

Architecture very sensitive to element properties

Enabling Technology is a Prerequisite

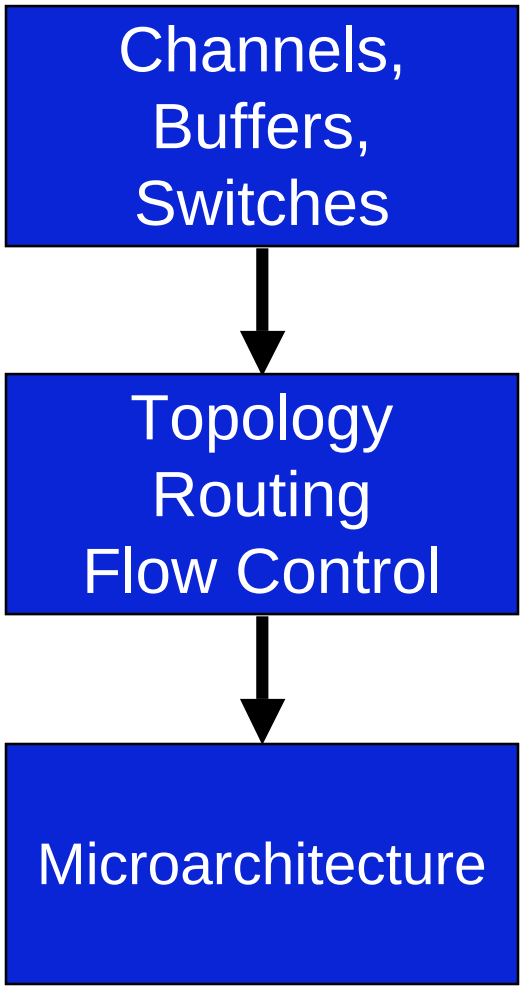

Channels

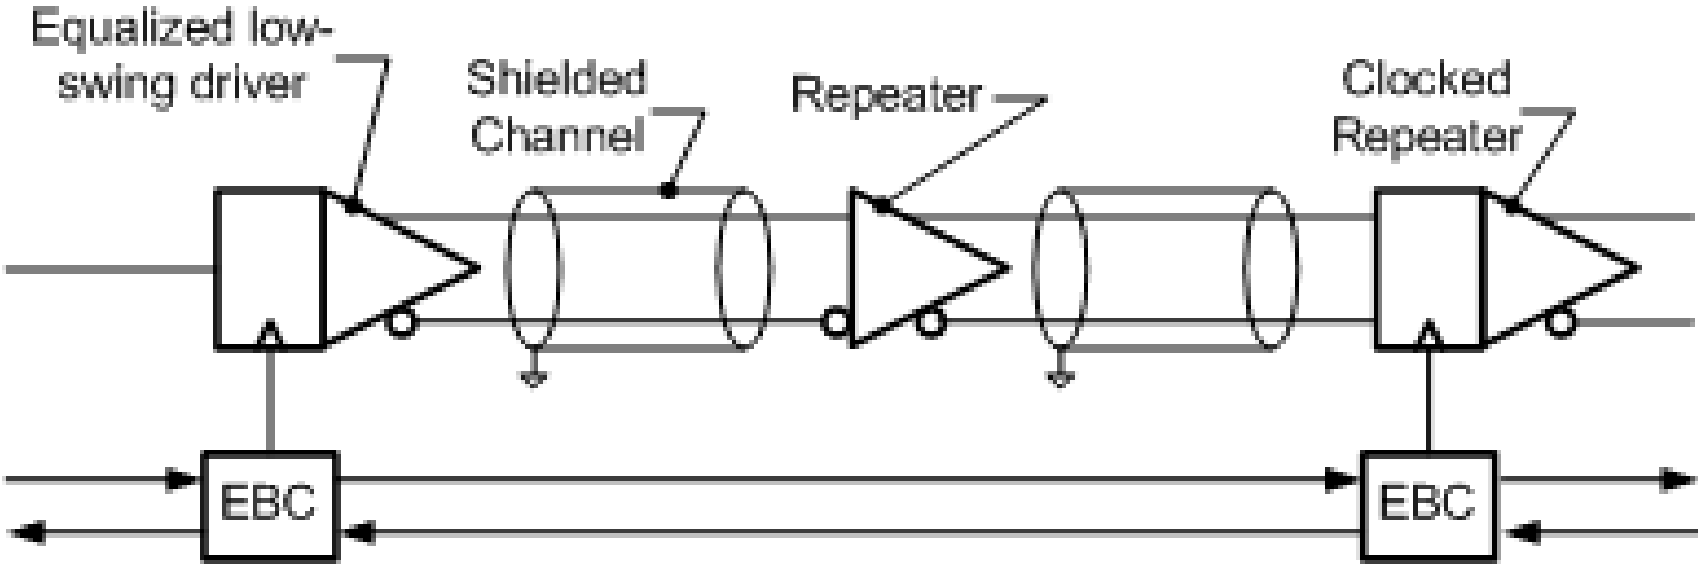

- 10x to 100x power reduction
- Eq signaling for faster propagation and increased repeater distance (D & P Chapter 8, Heaton 01)
- Elastic channels provide “free” buffers (Mizuno 01)
- Send 4-8 bits per cycle per wire (assuming 20FO4 cycle)

Buffers

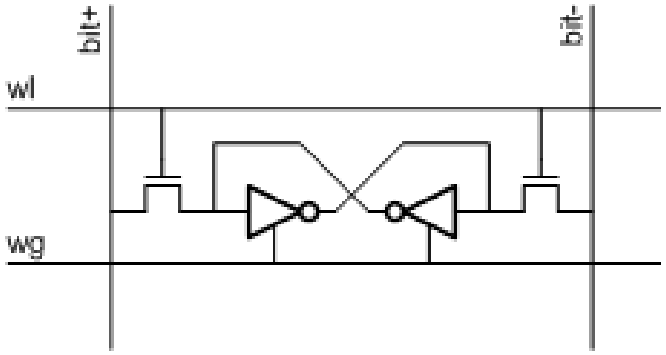

- Dense arrays (vs. Flip-Flops or Latches) 10x area/bit
- Low-swing write

Switches

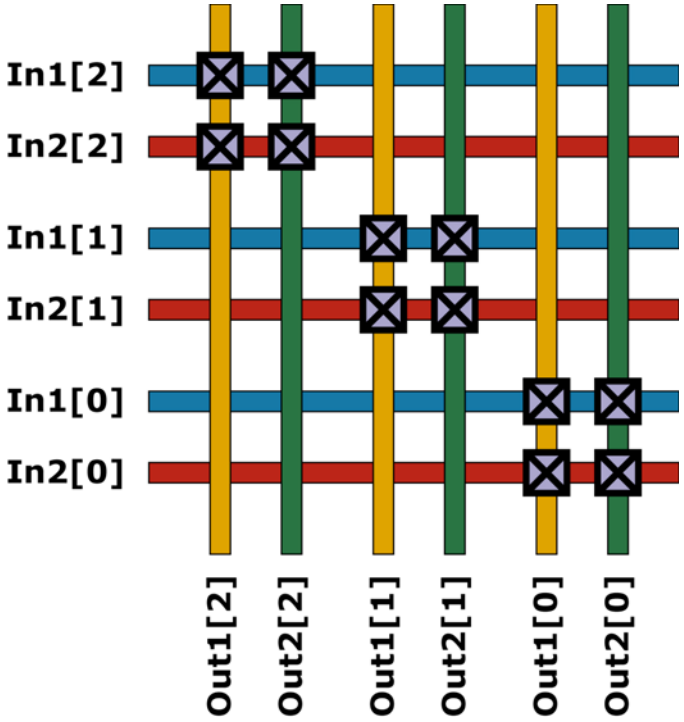

- Low-swing bit lines
- Operate at channel rate
  - Reduces area and hence power
- Equalized drive
- Buffered crosspoints
- Integral allocation

Properties of these elements drives optimal network organization

Torus

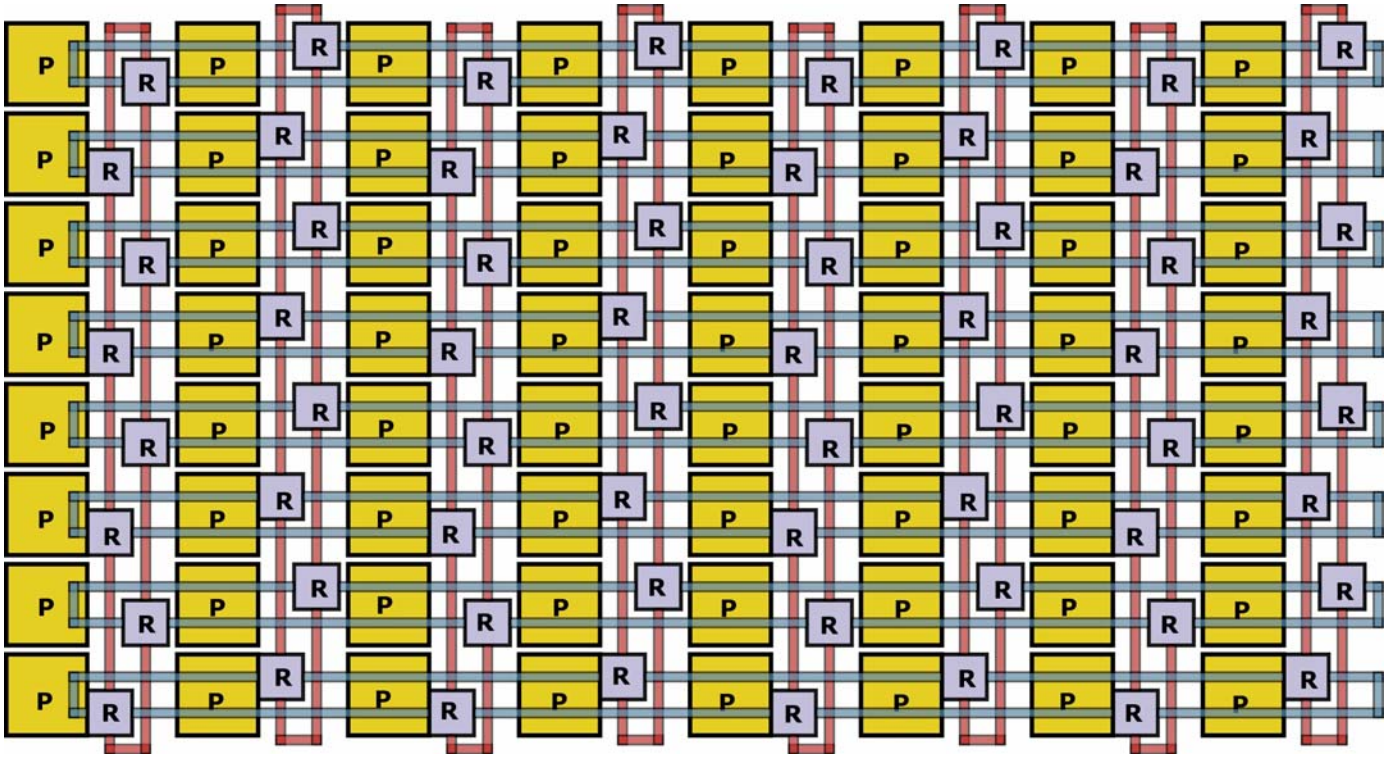

Source: Balfour and Dally, ICS 06

Concentrated Mesh

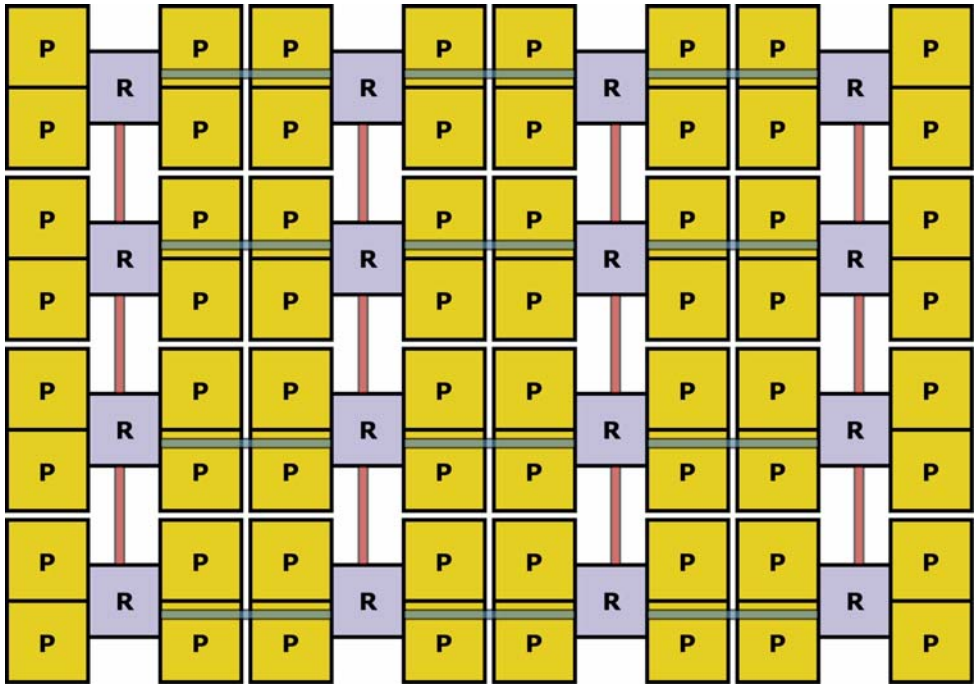

Source: Balfour and Dally, ICS 06

Express Links

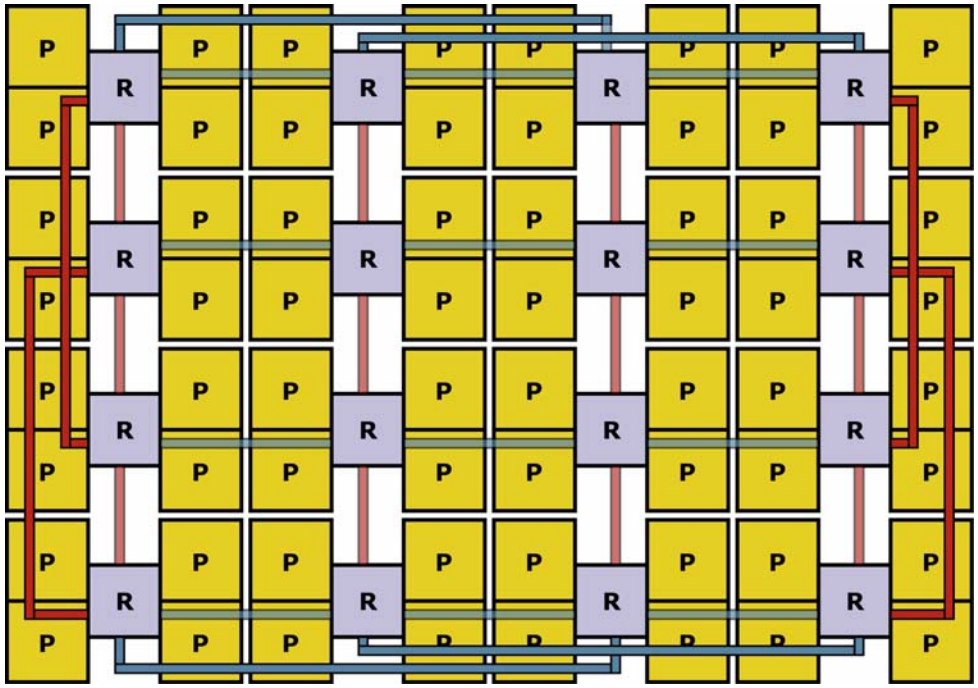

Source: Balfour and Dally, ICS 06

Network Replication

- Abundant wire resources ➡ build second network
  - Resource allocation tradeoff

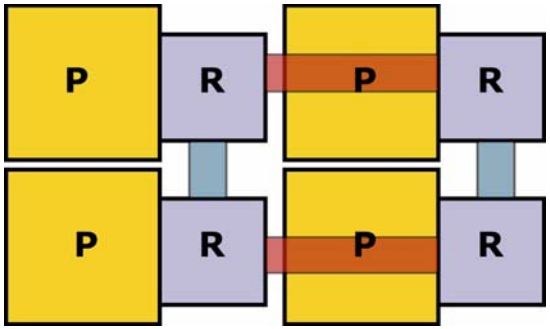

Wide:
[+] Serialization Latency
[+] Router Energy Efficiency
[–] Router Area

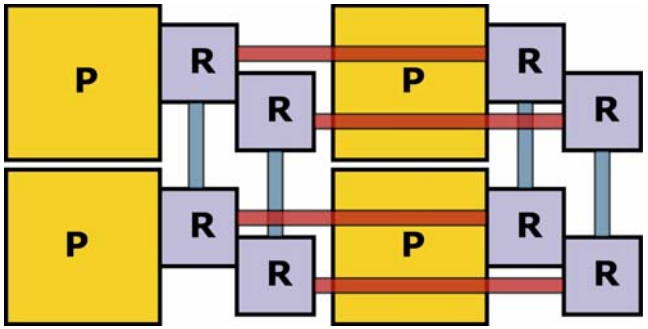

Replicated:
[+] Decoupled Resources
[+] Area Efficiency
[?] Energy Efficiency
[–] Serialization Latency

[+] SCALABLE

Source: Balfour and Dally, ICS 06

Energy Efficiency

Network Energy × Completion Time
(normalized to Torus network)

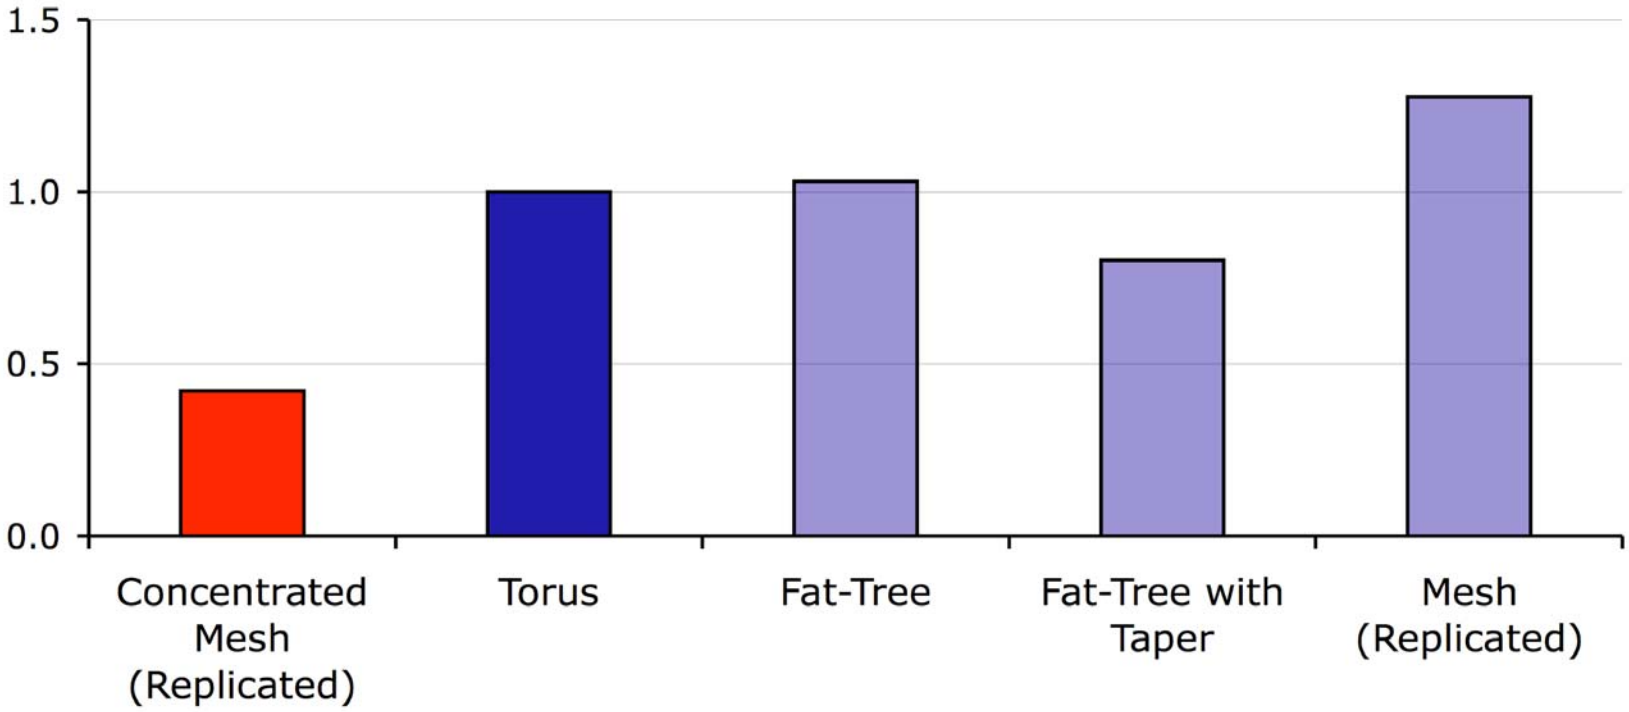

Source: Balfour and Dally, ICS 06

Large differences in efficiency.

Optimal topology not obvious, not regular and very sensitive to properties of network elements

Where is Energy Expended?

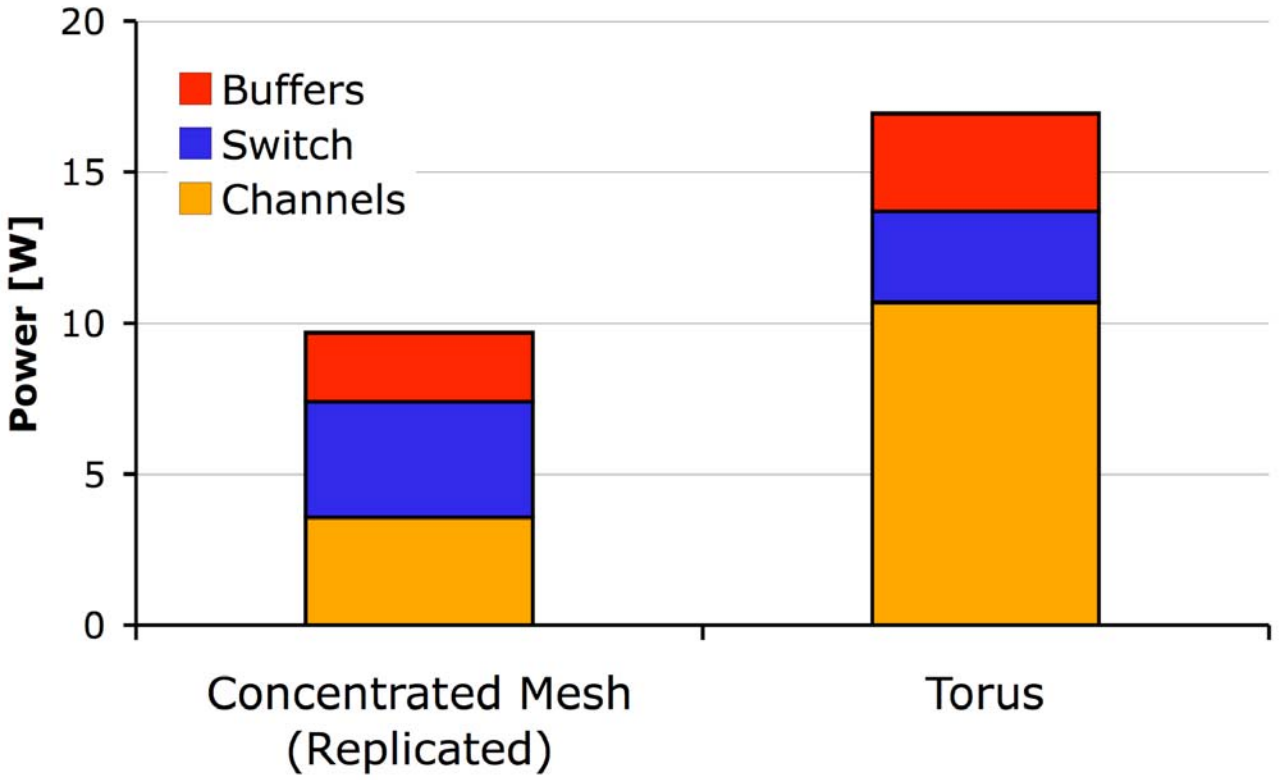

Source: Balfour and Dally, ICS 06

A Research Agenda

1. Develop efficient network elements
   - Channels, buffers, switches, allocators
   - Opportunities for >10x improvements in efficiency
   - Enabling technology
2. Capture workloads representative of CMPs and SoCs
3. Develop optimal topologies for 1 and 2
4. Develop efficient routing and flow-control methods
   - Load-balanced routing
   - Buffer-efficient flow control
5. Develop efficient router microarchitectures
   - Single cycle, area efficient
6. Prototype to test assumptions
7. Iterate

Summary

- OCNs critically important
  - Emerging CMPs, SoCs
- Very different than off-chip networks
  - Cost - largely area
  - Channel properties
- OCIN design very sensitive to implementation
  - Need floorplans, accurate estimates
- Efficient network elements are enabling technology
  - Change the equation for network design
- Optimal design
  - Concentration, replicated networks
- Many research opportunities