

# ECE/CS 757: Advanced Computer Architecture II Interconnects

Instructor: Mikko H Lipasti

Spring 2017

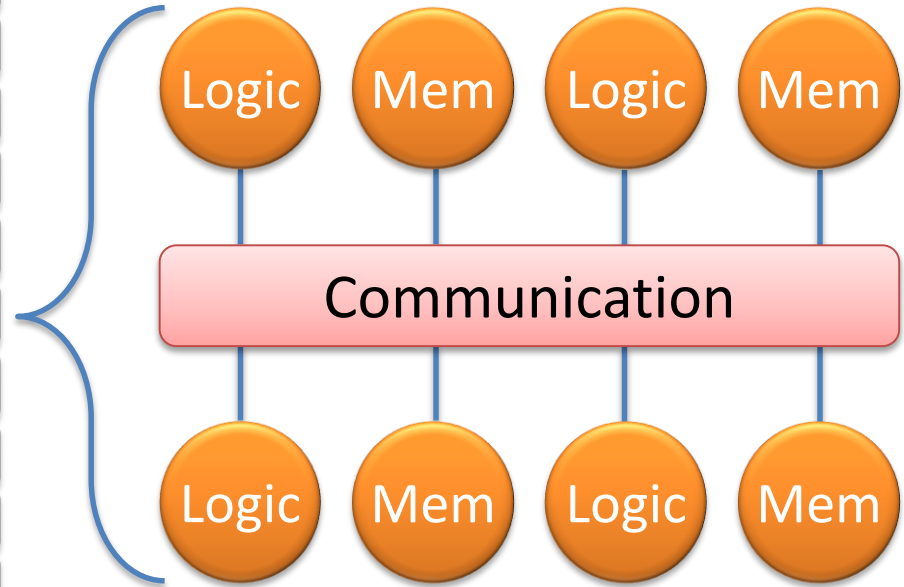
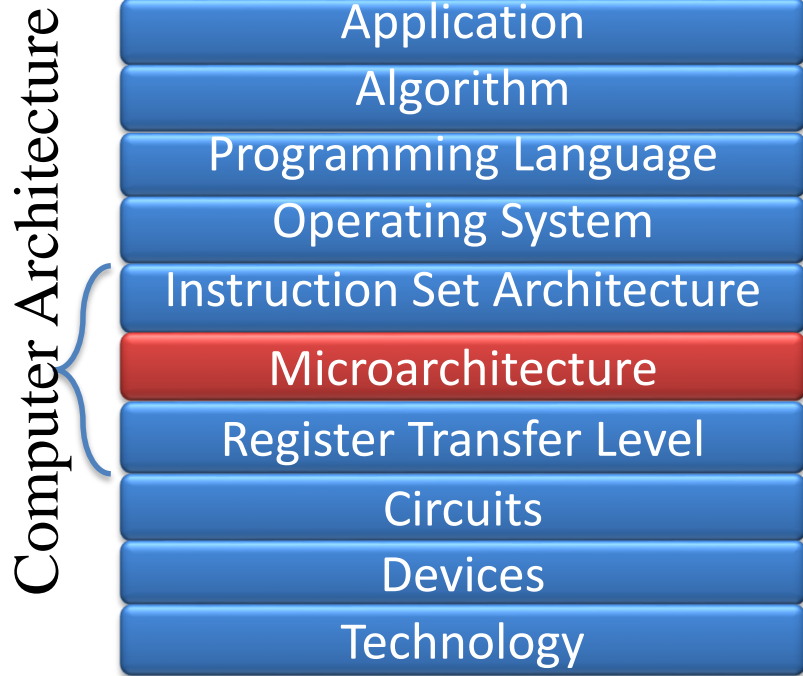
University of Wisconsin-Madison

Lecture notes created by Natalie Enright Jerger

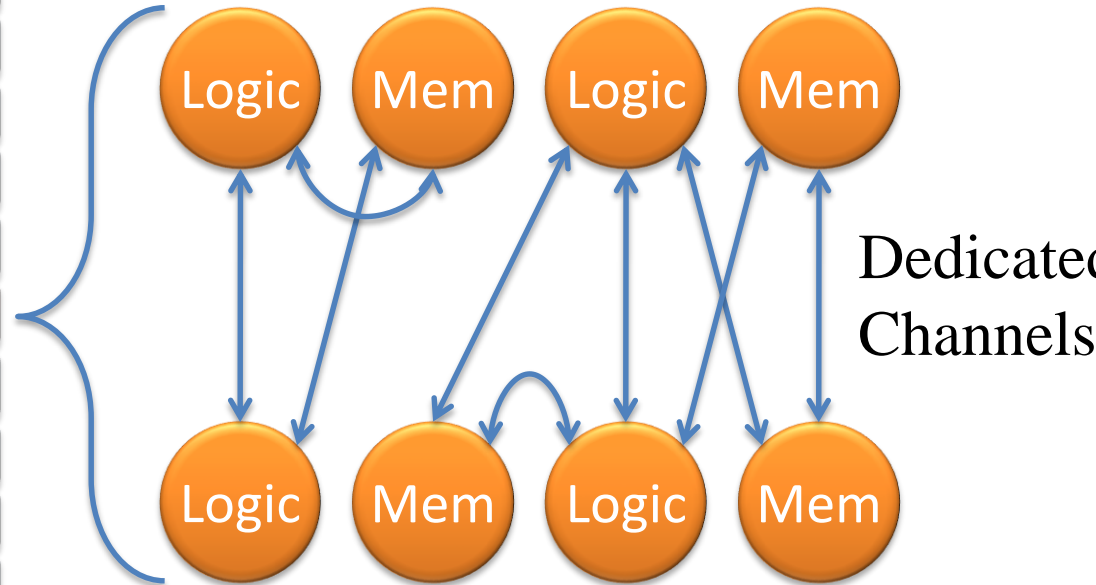

# Lecture Outline

- Introduction to Networks
- Network Topologies
- Network Routing
- Network Flow Control
- Router Microarchitecture

# What is an Interconnection Network?

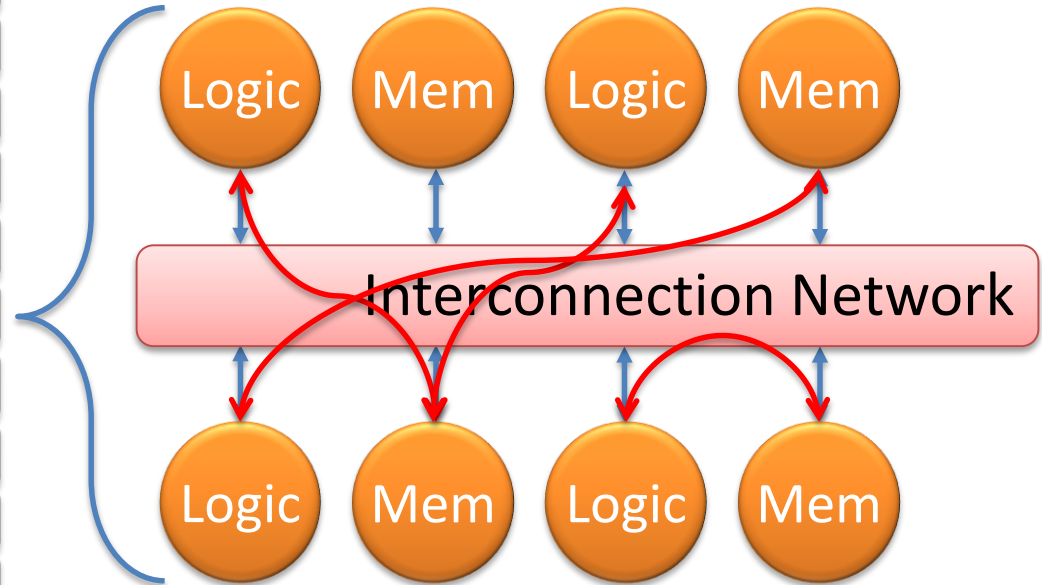
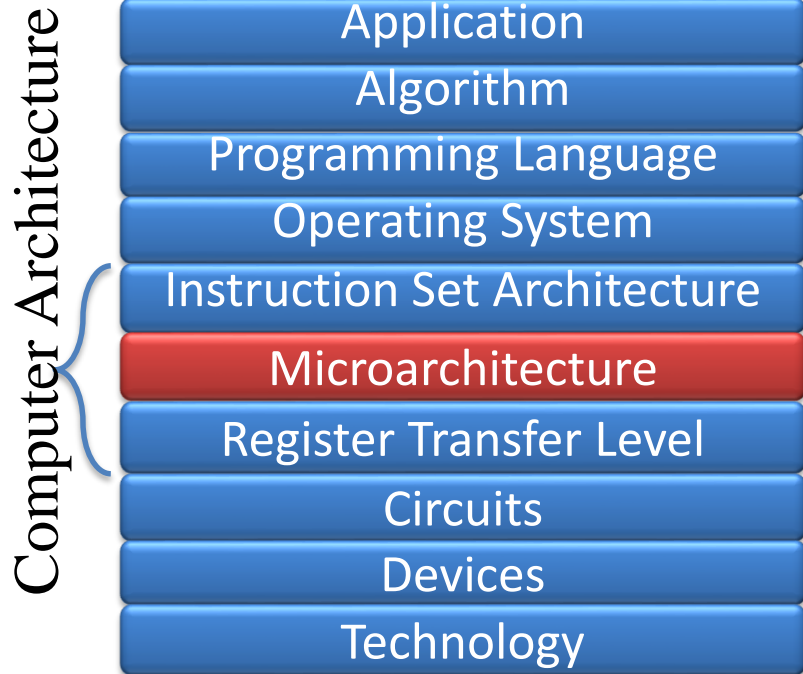


# What is an Interconnection Network?



- **Application:** Ideally wants low-latency, high-bandwidth, dedicated channels between logic and memory
- **Technology:** Dedicated channels too expensive in terms of area and power

# What is an Interconnection Network?



- An **Interconnection Network** is a programmable system that transports data between terminals
- **Technology**: Interconnection network helps efficiently utilize scarce resources
- **Application**: Managing communication can be critical to performance

# Interconnection Networks

## Introduction

- *Interconnection networks should be designed*
  - *to transfer the maximum amount of information*
  - *within the least amount of time (and cost, power constraints)*
  - *so as not to bottleneck the system*

# Types of Interconnection Networks

- Interconnection networks can be grouped into four domains
  - Depending on number and proximity of devices to be connected
- **On-Chip Networks (OCNs or NoCs)**
  - Devices include microarchitectural elements (functional units, register files), caches, directories, processors
  - Current/Future systems: dozens, hundreds of devices
    - Ex: Intel TeraFLOPS research prototypes – 80 cores
    - Intel Single-chip Cloud Computer – 48 cores
  - Proximity: millimeters

# Types of Interconnection Networks (2)

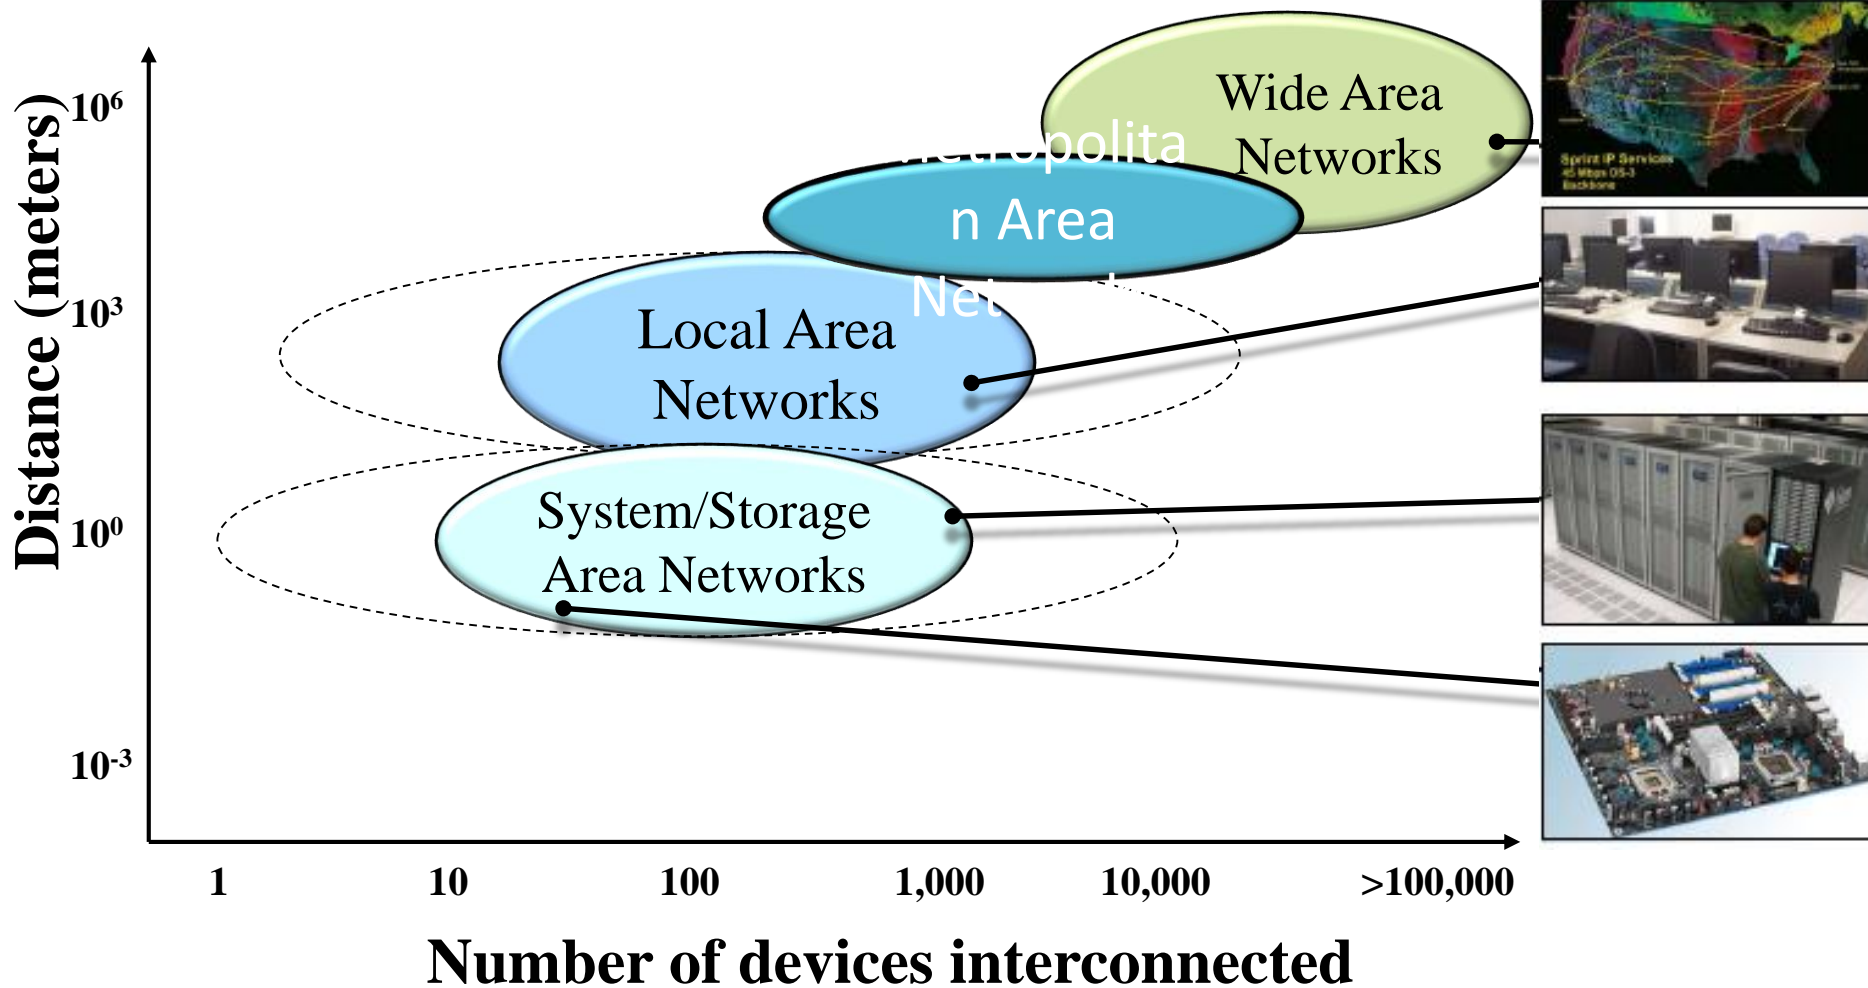
- System/Storage Area Networks (SANs)
  - Multiprocessor and multicomputer systems
    - Interprocessor and processor-memory interconnections
  - Server and data center environments
    - Storage and I/O components
  - Hundreds to thousands of devices interconnected
    - IBM Blue Gene/L supercomputer (64K nodes, each with 2 processors)
  - Maximum interconnect distance: tens of meters (typical) to a few hundred meters
  - Examples (standards and proprietary)
    - InfiniBand, Myrinet, Quadrics, Advanced Switching Interconnect



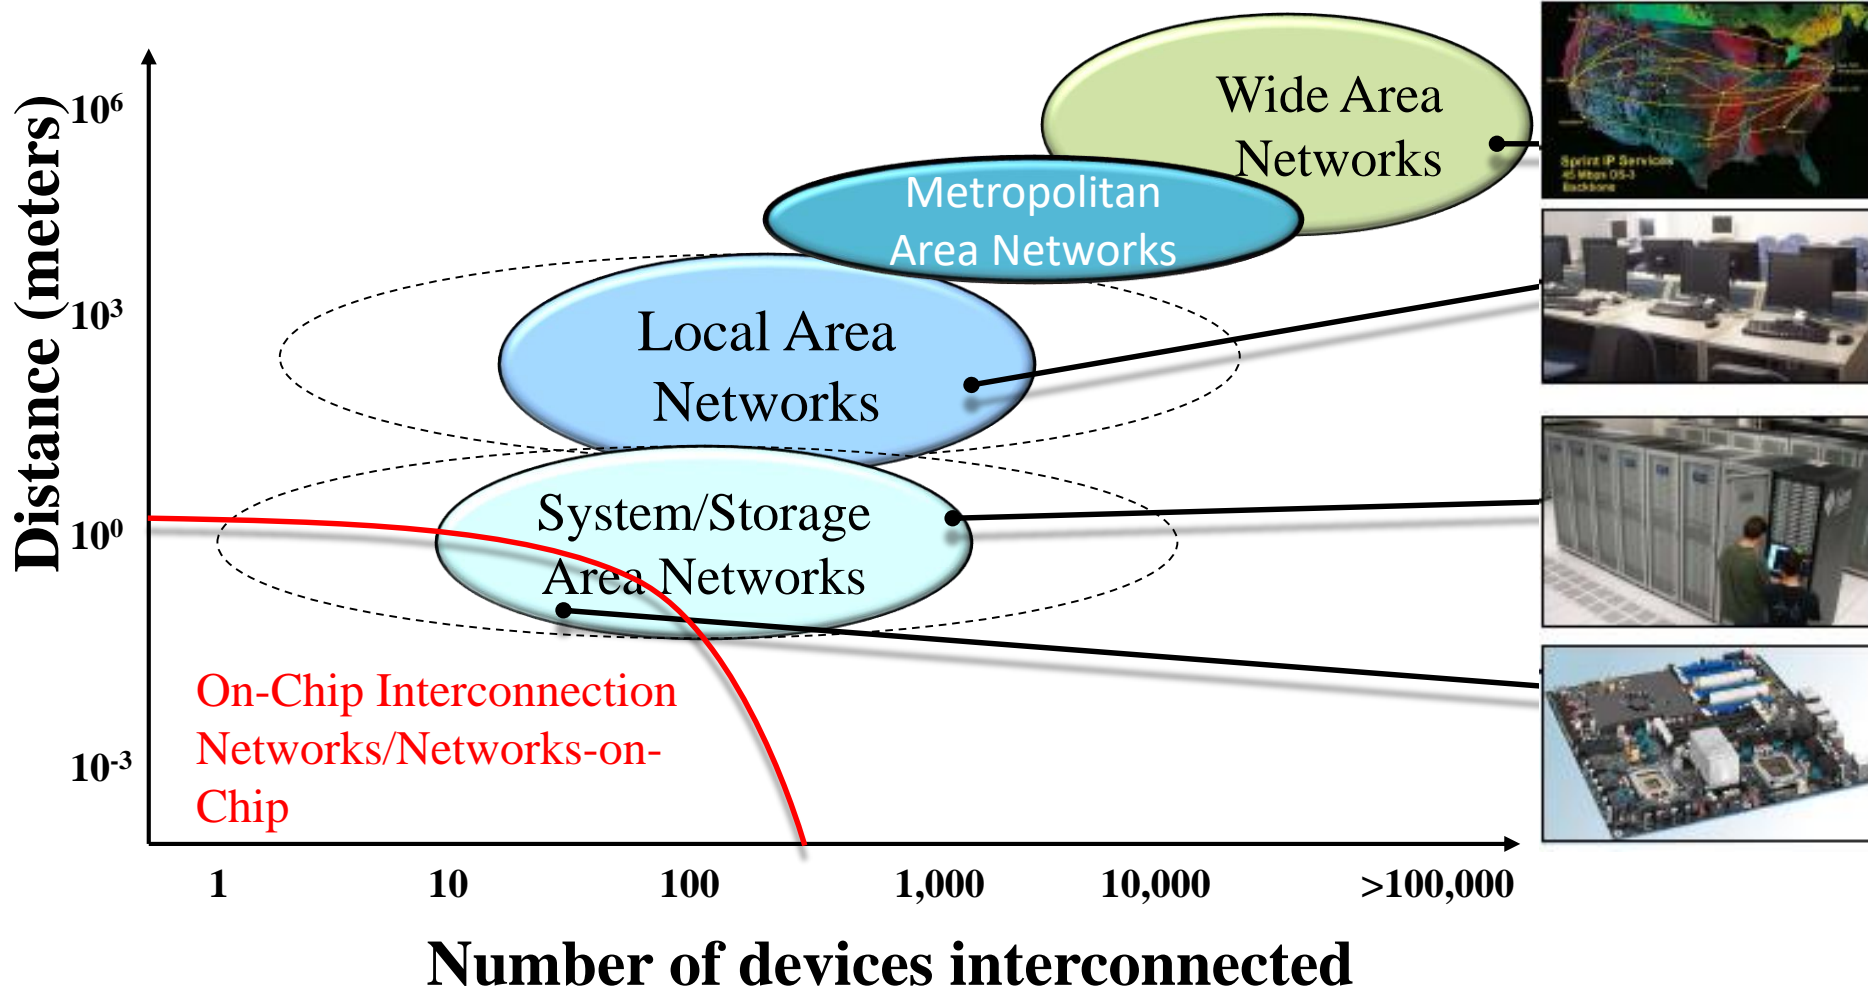
# Types of Interconnection Networks (3)

- Local Area Networks (LANs)
  - Interconnect autonomous computer systems
    - Machine room or throughout a building or campus
  - Hundreds of devices interconnected (1,000s with bridging)
  - Maximum interconnect distance
    - few kilometers to few tens of kilometers
  - Example (most popular): Ethernet, with 10 Gbps over 40Km
- Wide Area Networks (WANs)
  - Interconnect systems distributed across globe
    - Internetworking support required
  - Many millions of devices interconnected
  - Max distance: many thousands of kilometers
  - Example: ATM (asynchronous transfer mode)


# Interconnection Network Domains



# Interconnection Network Domains




# Why Study Networks on Chip?



Data partially collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond

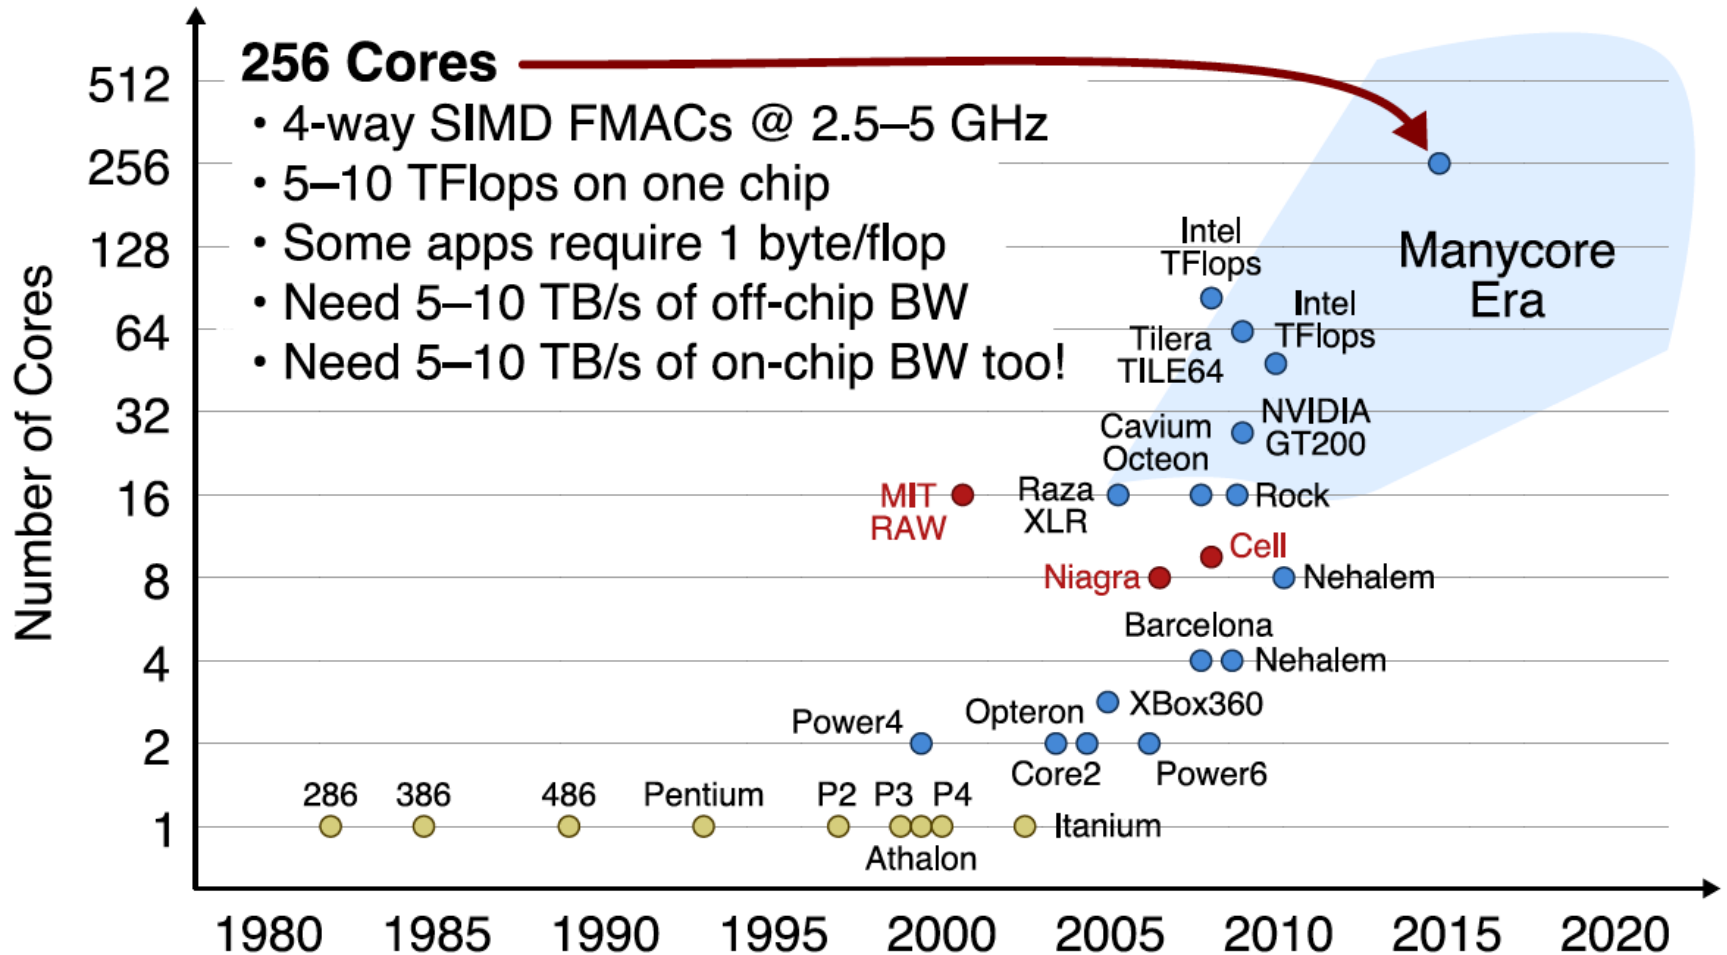
# Why Study On-Chip Networks?



Data partially collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond

ECE 1749H: Interconnection  
Networks (Enright Jerger)


# Example of Multi- and Many-Core Architectures



# Why study interconnects?

- They provide external connectivity from system to outside world
  - Also, connectivity within a single computer system at many levels
    - I/O units, boards, chips, modules and blocks inside chips
- *Trends*: high demand on communication bandwidth
  - increased computing power and storage capacity
  - switched networks are replacing buses
- *Computer architects/engineers must understand interconnect problems and solutions in order to more effectively design and evaluate systems*

# On-Chip Networks (OCN or NoCs)



- Why On-Chip Network?
  - Ad-hoc wiring does not scale beyond a small number of cores
    - Prohibitive area
    - Long latency
- OCN offers
  - scalability
  - efficient multiplexing of communication
  - often modular in nature (ease verification)



# Differences between on-chip and off-chip networks

- Significant research in multi-chassis interconnection networks (off-chip)
  - Supercomputers
  - Clusters of workstations
  - Internet routers
- Leverage research and insight but...
  - Constraints are different
  - New opportunities

# Off-chip vs. on-chip

- Off-chip: I/O bottlenecks
  - Pin-limited bandwidth
  - Inherent overheads of off-chip I/O transmission
- On-chip
  - Wiring constraints
    - Metal layer limitations
    - Horizontal and vertical layout
    - Short, fixed length
    - Repeater insertion limits routing of wires
      - Avoid routing over dense logic
      - Impact wiring density

# Off-Chip vs On-Chip

- On-Chip
  - Power
    - Consume 10-15% or more of die power budget
  - Latency
    - Different order of magnitude
    - Routers consume significant fraction of latency


# New opportunities

- Abundant wiring
  - Change in relative cost of wires and buffers
  - Many short flits
- Tightly integrated into system
  - Not commodity – fully customized design
  - Allows for optimization with uncore
    - Cache coherence
- Emerging technology
  - Optics
  - 3D


# On-Chip Network Evolution

- Ad hoc wiring
  - Small number of nodes
- Buses and Crossbars
  - Simplest variant of on-chip networks
  - Low core counts
  - Like traditional multiprocessors
    - Bus traffic quickly saturates with a modest number of cores
  - Crossbars: higher bandwidth
    - Poor area and power scaling

# Multicore Examples (1)



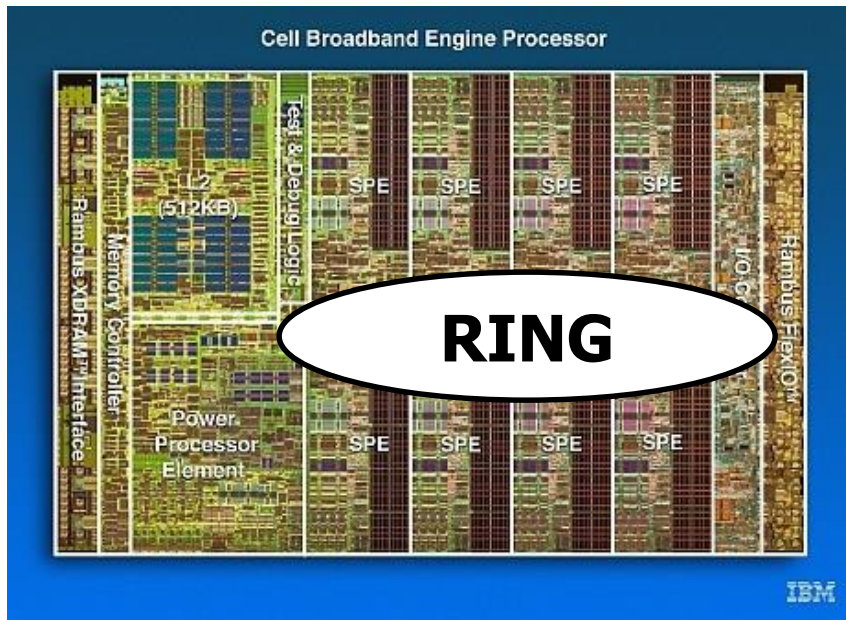
Sun Niagara



- Niagara 2: 8x9 crossbar (area  $\approx$  core)
- Rock: Hierarchical crossbar (5x5 crossbar connecting clusters of 4 cores)

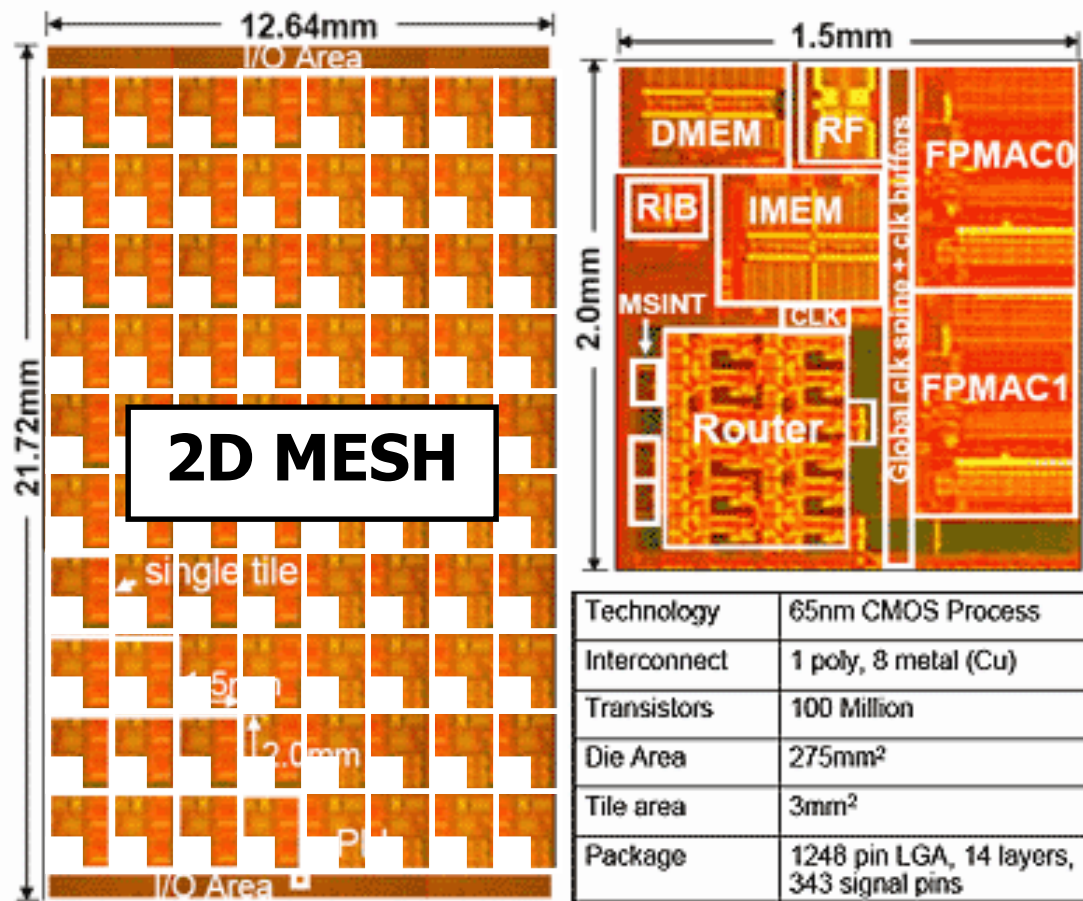
# Multicore Examples (2)

- IBM Cell
- Element Interconnect Bus
  - 12 elements
  - 4 unidirectional rings
    - 16 Bytes wide
    - Operates at 1.6 GHz



IBM Cell



# Many Core Example



- Intel TeraFLOPS
  - 80 core prototype
  - 5 GHz
  - Each tile:
    - Processing engine + on-chip network router




# Many-Core Example (2): Intel SCC



Courtesy: Jason Howard, Intel

- Intel's Single-chip Cloud Computer (SCC) uses a 2D mesh with state of the art routers

# Performance and Cost



- Performance: latency and throughput
- Cost: area and power

# Lecture Outline

- Introduction to Networks
- Network Topologies
- Network Routing
- Network Flow Control
- Router Microarchitecture

# Readings

- Read: [18] N. Enright Jerger, L.-S. Pei, “On-Chip Networks,” Synthesis Lectures on Computer Architecture,  
<http://www.morganclaypool.com/doi/abs/10.2200/S00209ED1V01Y200907CAC008>
- Review: [19] D. Wentzlaff, P. Griffin, H. Hoffmann, L. Bao, B. Edwards, C. Ramey, M. Mattina, C.-C. Miao, J. F. B. III, and A. Agarwal. On-Chip Interconnection Architecture of the Tile Processor. IEEE Micro, vol. 27, no. 5, pp. 15-31, 2007