ECE 1749H: Interconnection Networks for Parallel Computer Architectures

Introduction

Prof. Natalie Enright Jerger

Interconnection Networks Introduction

- How to connect individual devices together into a community of communicating devices?

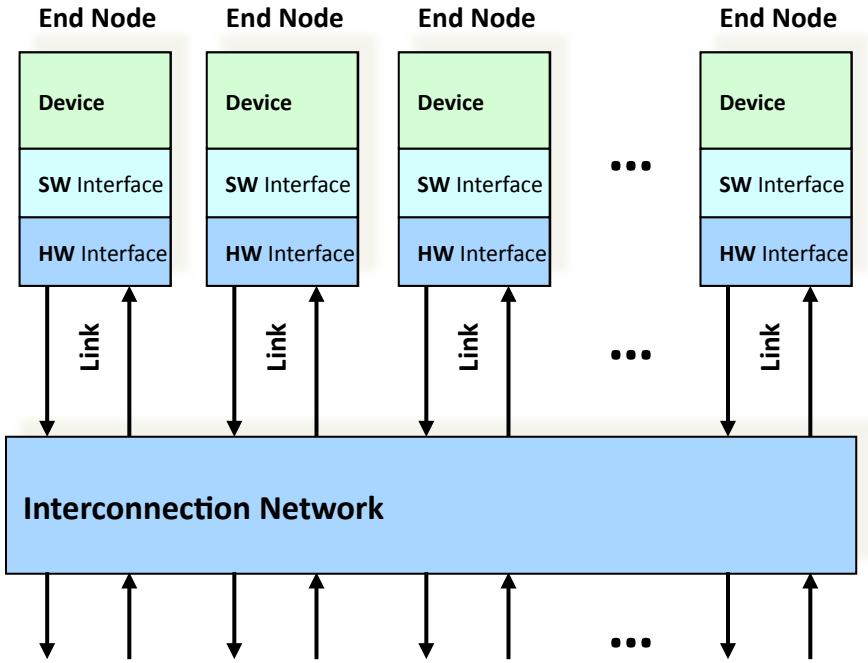

- Device:
  - Component within a computer
  - Single computer
  - System of computers
- Types of elements:
  - end nodes (device + interface)
  - links
  - interconnection network
- Internetworking: interconnection of multiple networks

Slide courtesy Timothy Mark Pinkston and José Duato

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*

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*

Slide courtesy Timothy Mark Pinkston and José Duato

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 designs: small number of devices
    - Ex: IBM Cell, Sun's Niagara
  - Projected systems: dozens, hundreds of devices
    - Ex: Intel TeraFLOPS research prototypes -- 80 cores
  - Proximity: millimeters

Slide courtesy Timothy Mark Pinkston and José Duato

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)
  - a few hundred meters (some)
    - InfiniBand: 120 Gbps over a distance of 300 m
- Examples (standards and proprietary)
  - InfiniBand, Myrinet, Quadrics, Advanced Switching Interconnect

Slide courtesy Timothy Mark Pinkston and José Duato

Local Area Network (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
  - few tens of kilometers (some)
- Example (most popular): Ethernet, with 10 Gbps over 40Km

Slide courtesy Timothy Mark Pinkston and José Duato

Wide Area Networks (WANs)

- Interconnect systems distributed across the globe
- Internetworking support is required
- Many millions of devices interconnected
- Maximum interconnect distance
  - many thousands of kilometers
- Example: ATM (asynchronous transfer mode)

Slide courtesy Timothy Mark Pinkston and José Duato

Interconnection Network Domains

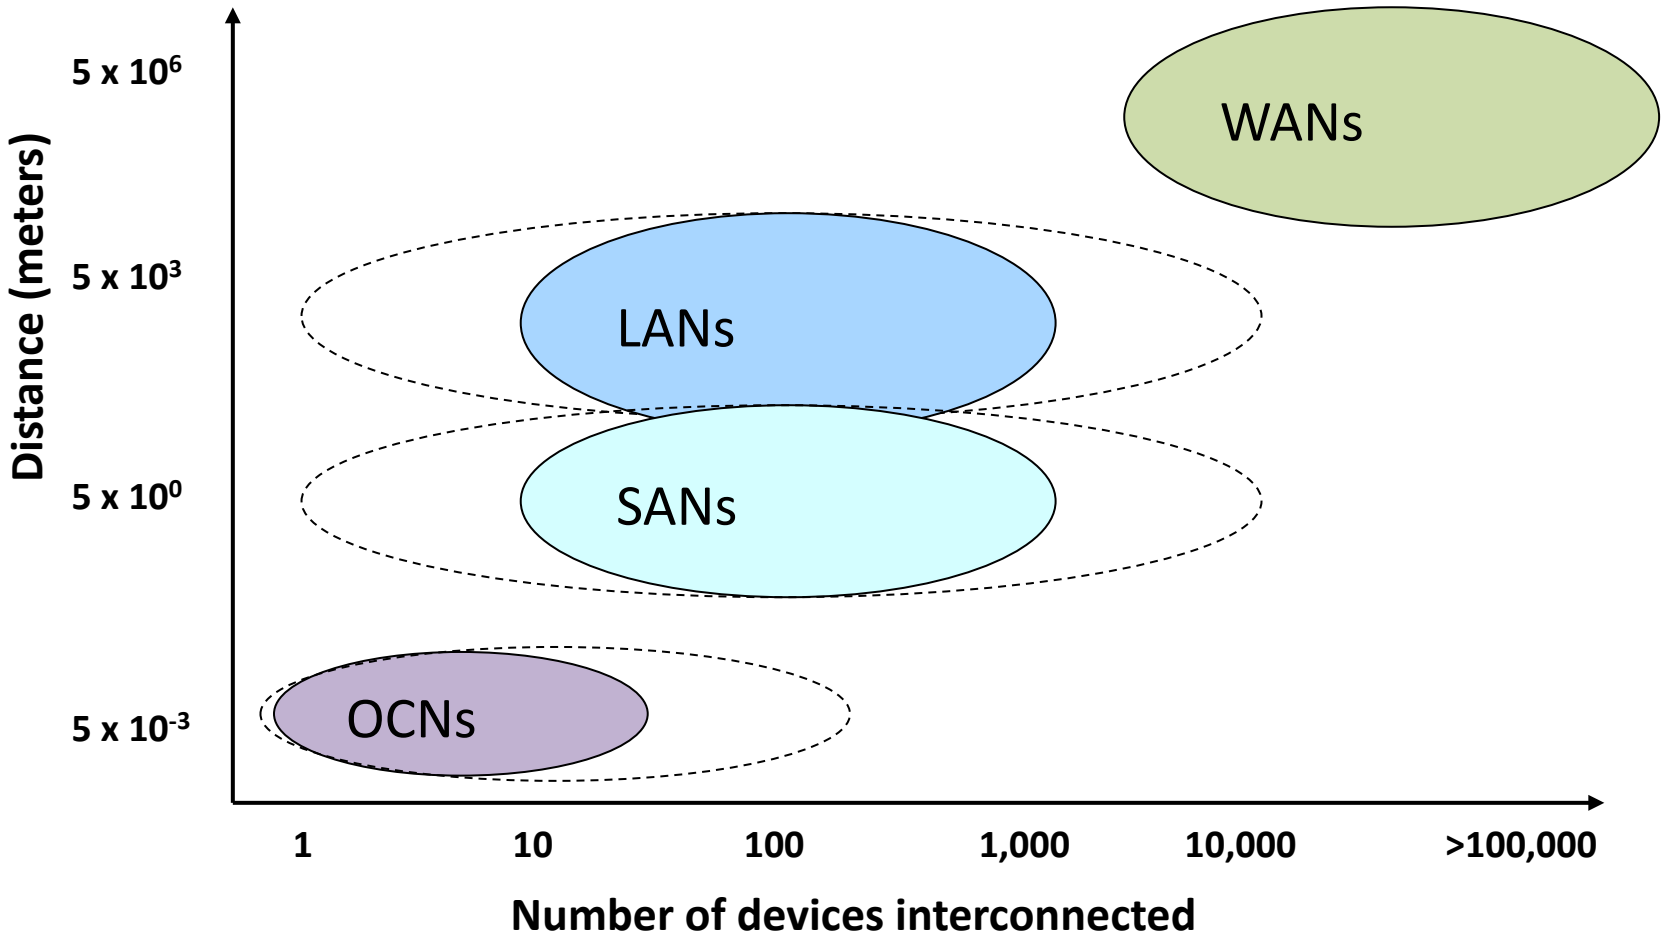

Slide courtesy Timothy Mark Pinkston and José Duato

Course Overview

- Focus on On-Chip Networks
  - Two aspects:
    - Interconnection network basics
      - Lectures covering: topology, routing, flow control and router microarchitecture
    - Second part:
      - Mix in current OCN research with each lecture topic

Moore's Law: Double the number of transistors on chip every 2 years

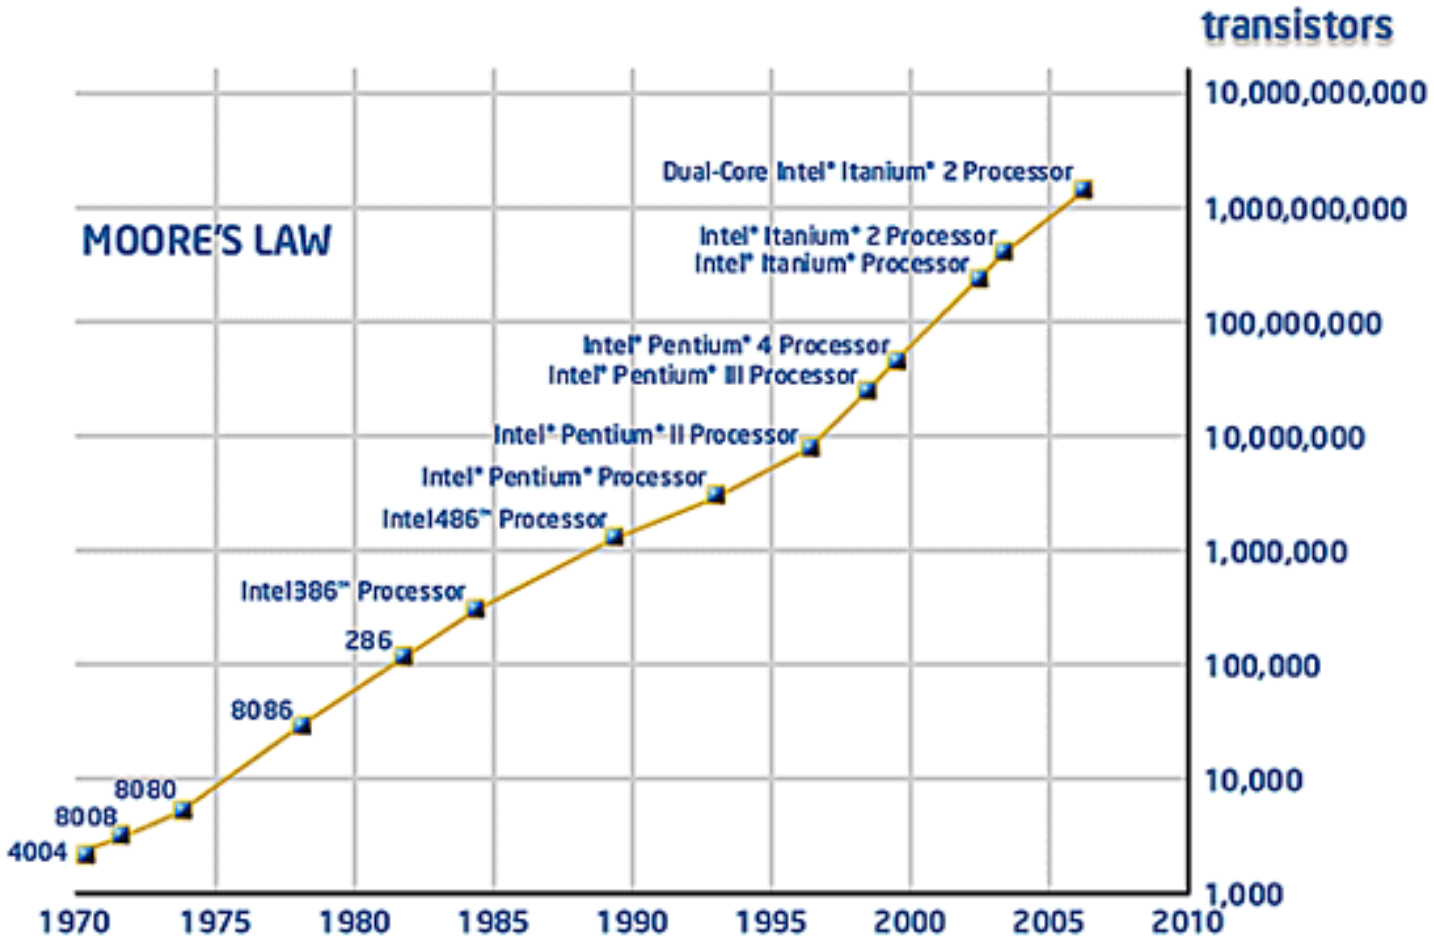

Chart courtesy Intel Corp.

What is Tera-scale?

TIPs of compute power operating on Tera-bytes of data

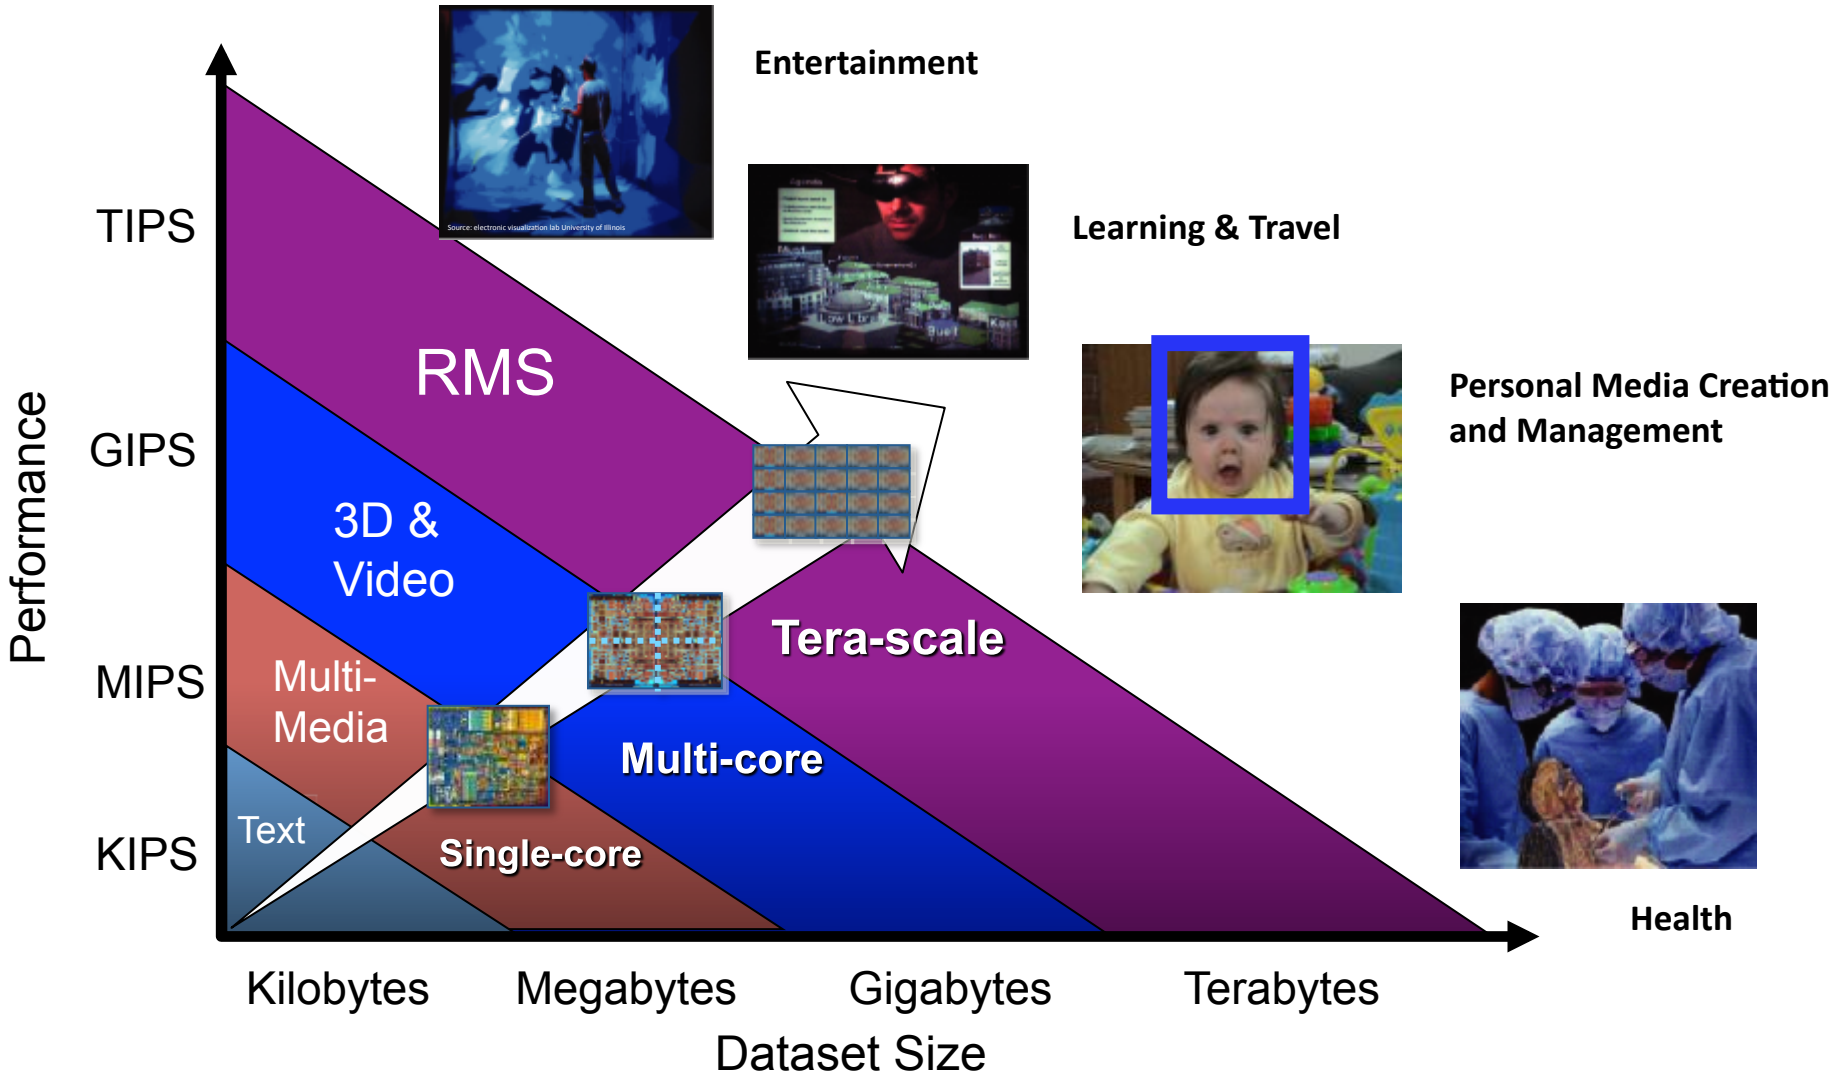

[Ack. Jim Held, Intel Fellow]

Performance Scaling Challenges

Energy Efficiency

Design Complexity

Programming Models

Emerging Applications

[Ack. Jim Held, Intel Fellow]

Processor power skyrocketing

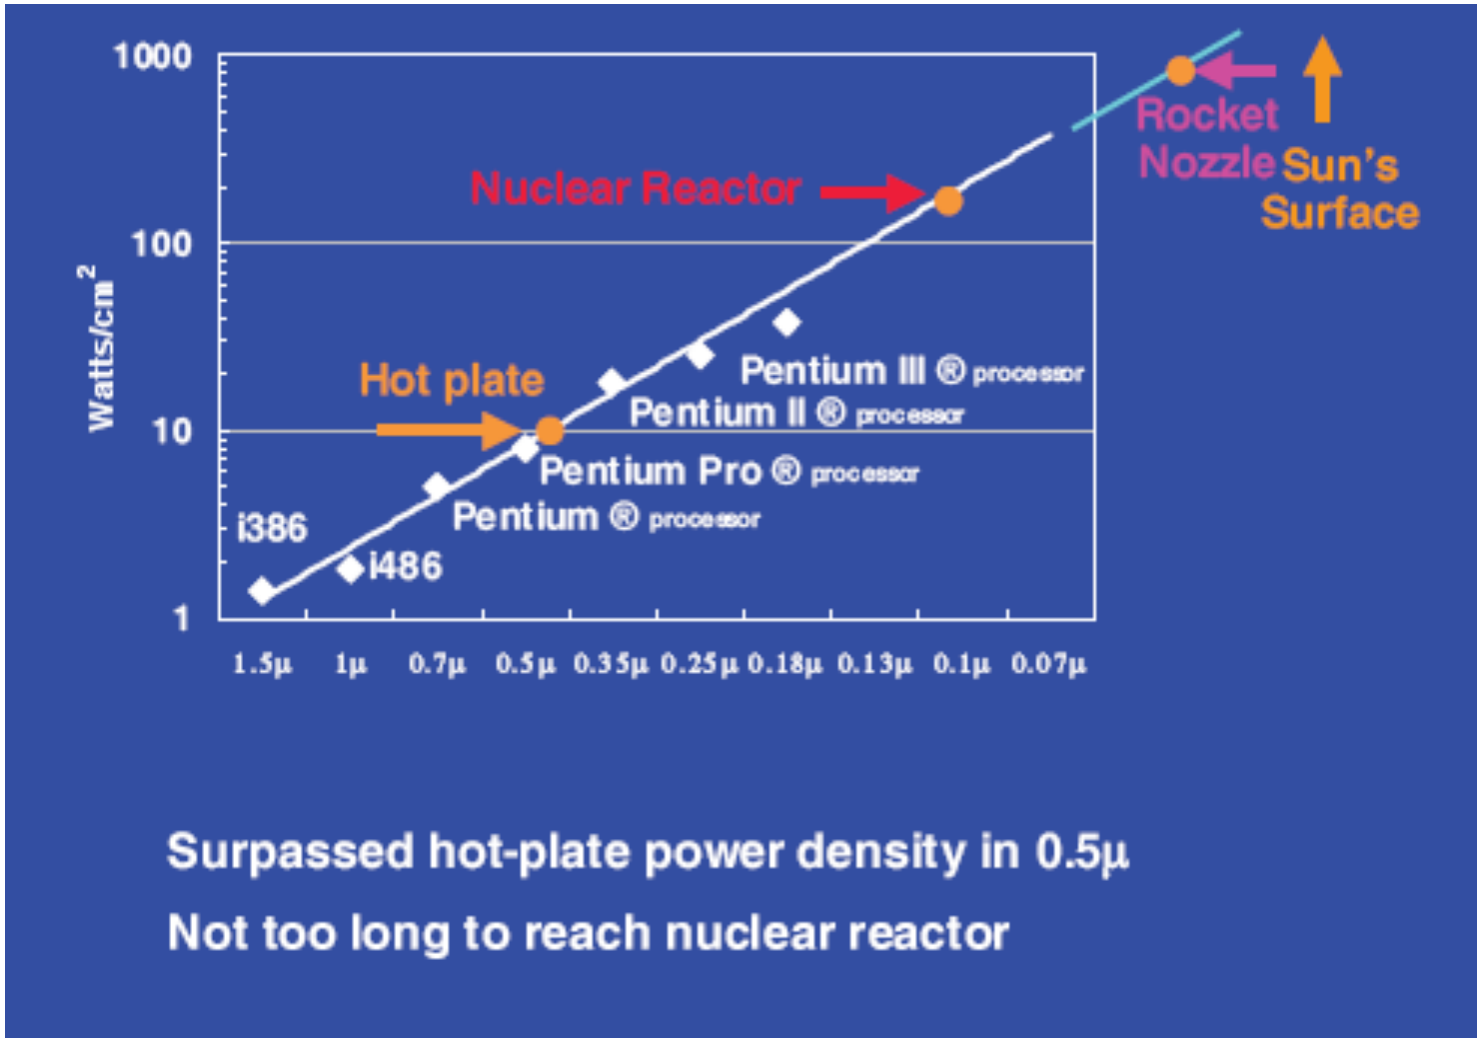

[Ack: Fred Pollack, Intel Fellow, 1996]

Multi-core chips the answer to the power problem

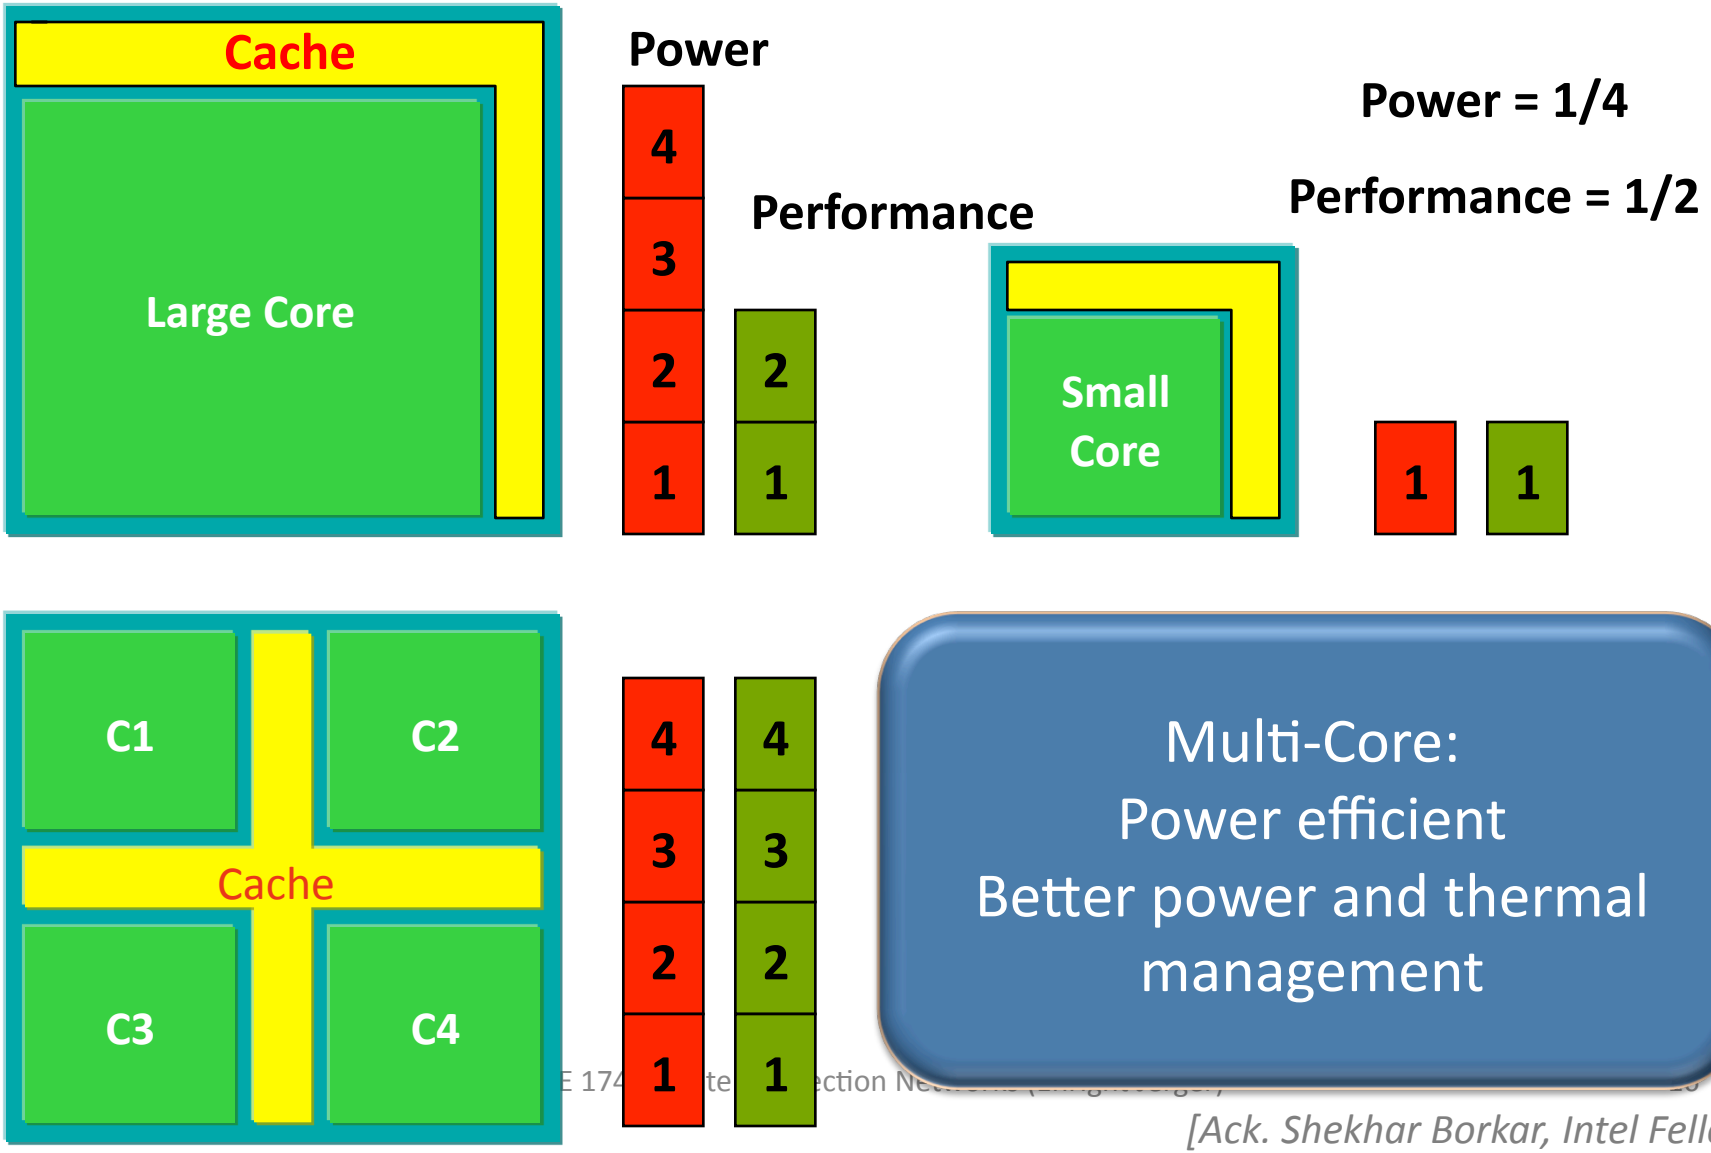

[Ack. Shekhar Borkar, Intel Fellow]

Tera-scale Research

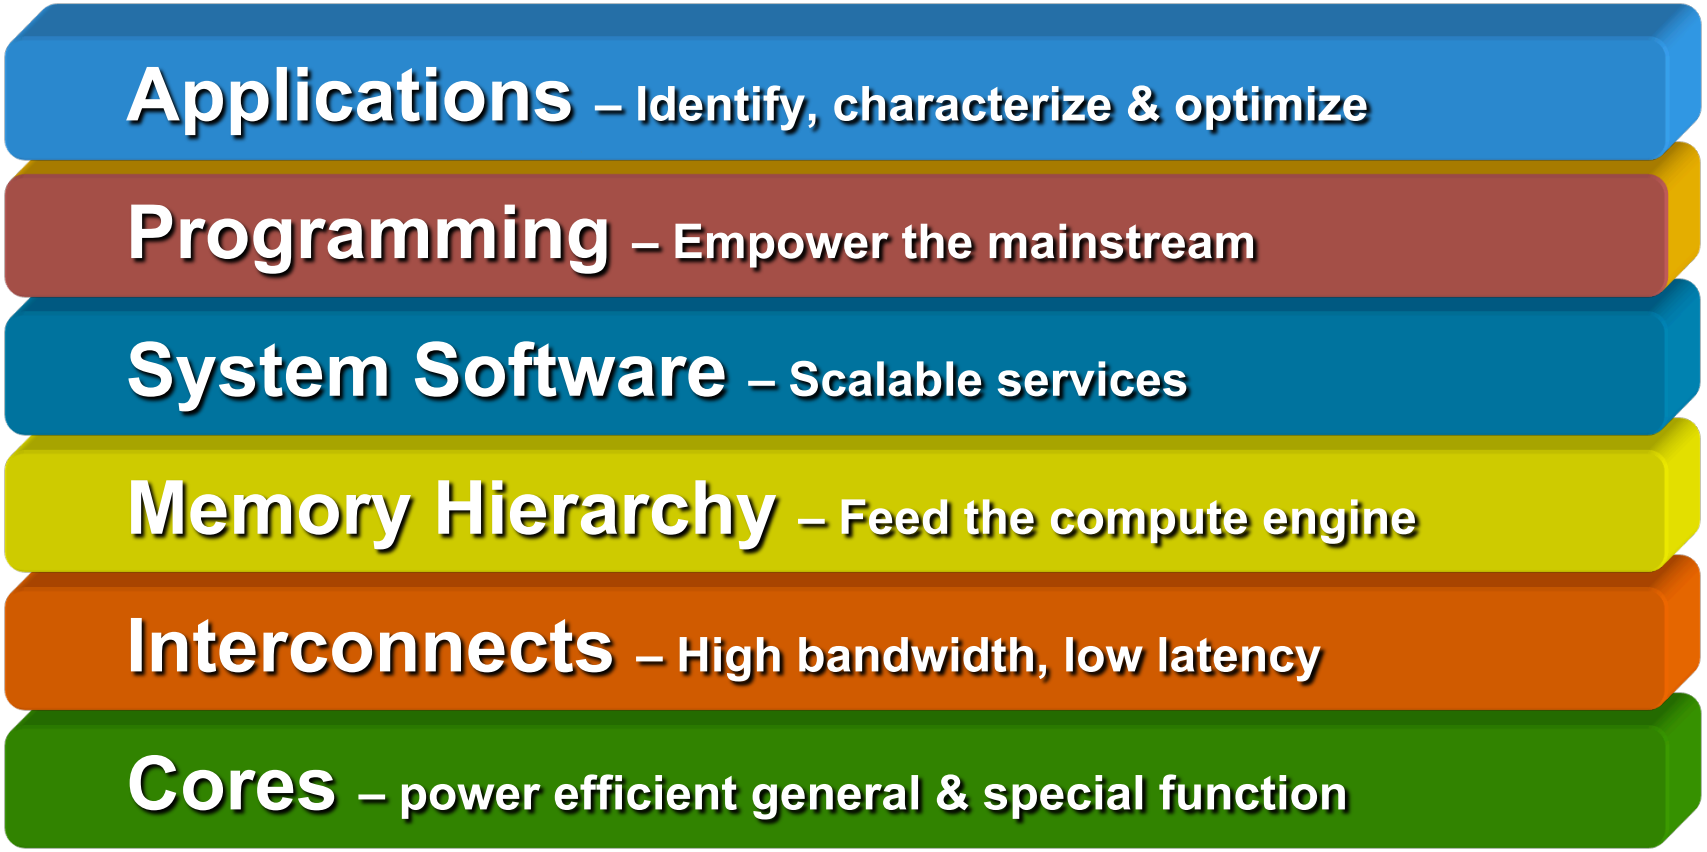

[Ack. Jim Held, Intel Fellow]

On-Chip Networks (OCN or NoCs)

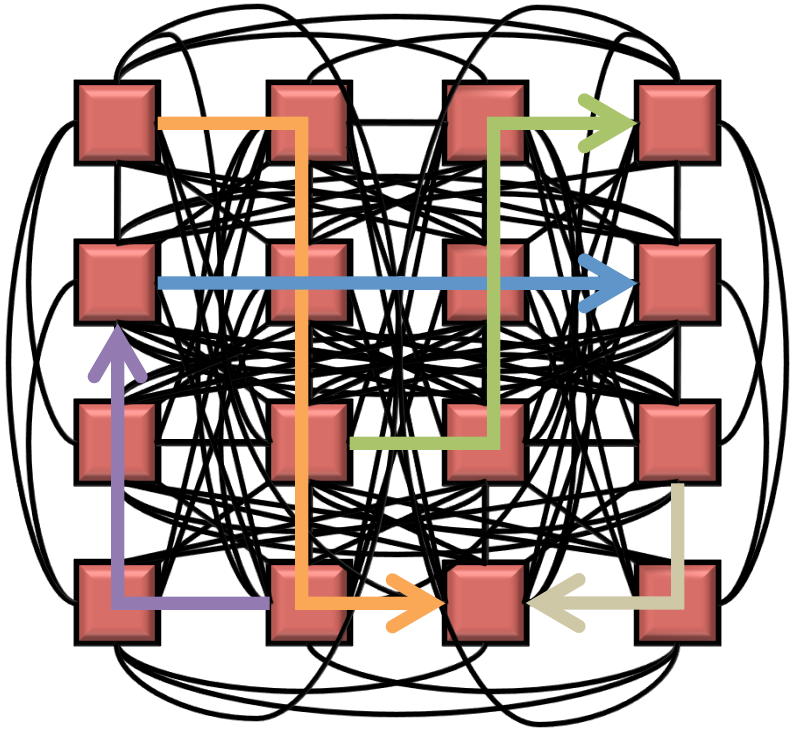

- Why On-Chip Network?
  - Ad-hoc wiring does not scale beyond a small number of cores
    - Prohibitive area
    - Long latency
- OCN offers
  - scalability
  - efficient multiplexing of communication
  - often modular in nature (ease verification)

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

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

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)

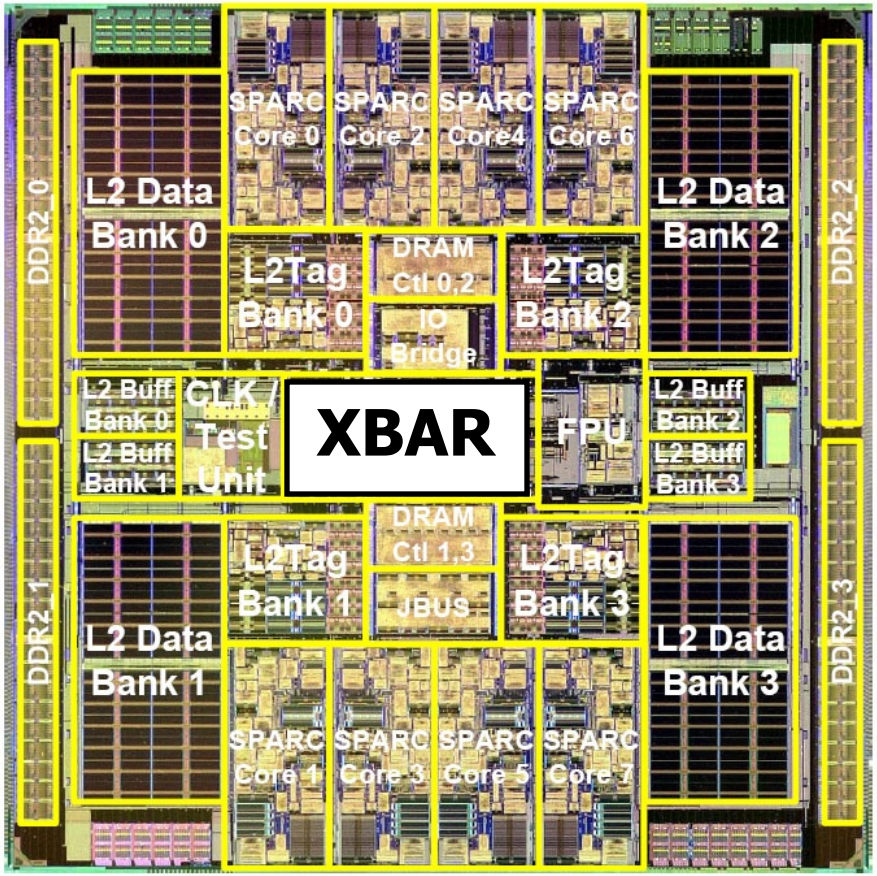

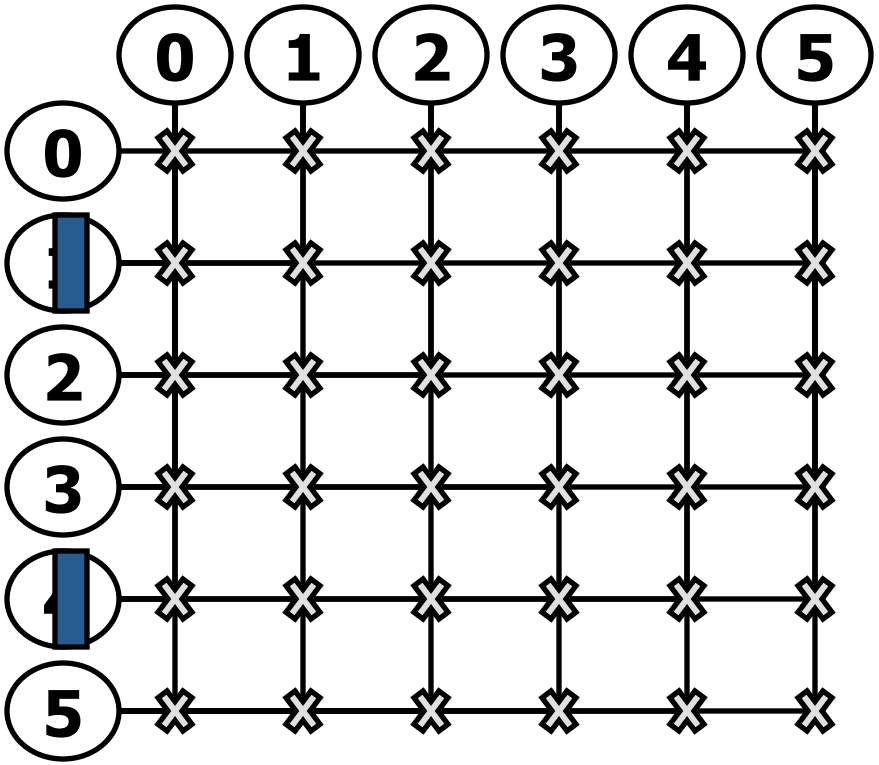

Sun Niagara

Multicore Examples (2)

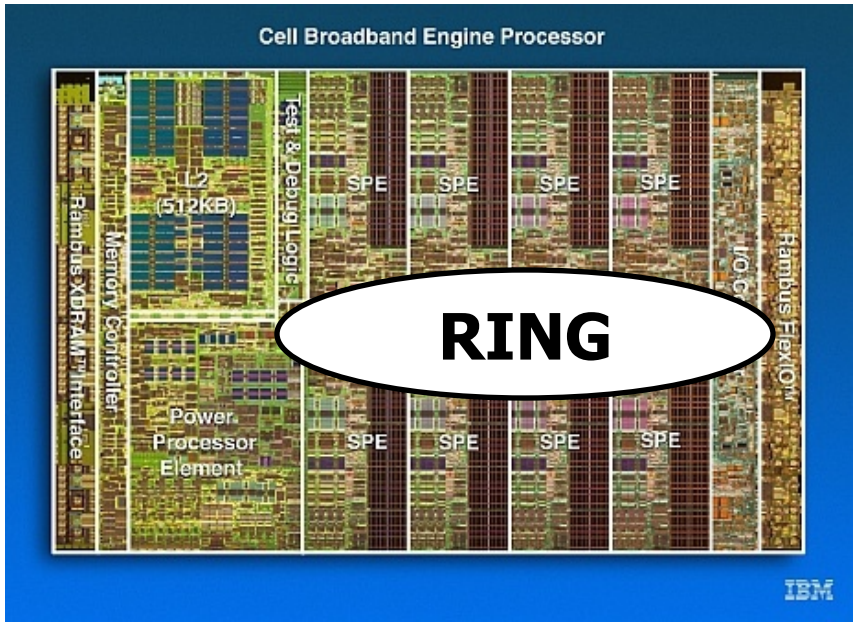

IBM Cell

- Element Interconnect Bus
  - 12 elements
  - 4 unidirectional rings
    - 16 Bytes wide
    - Operate at 1.6 GHz

Many Core Example

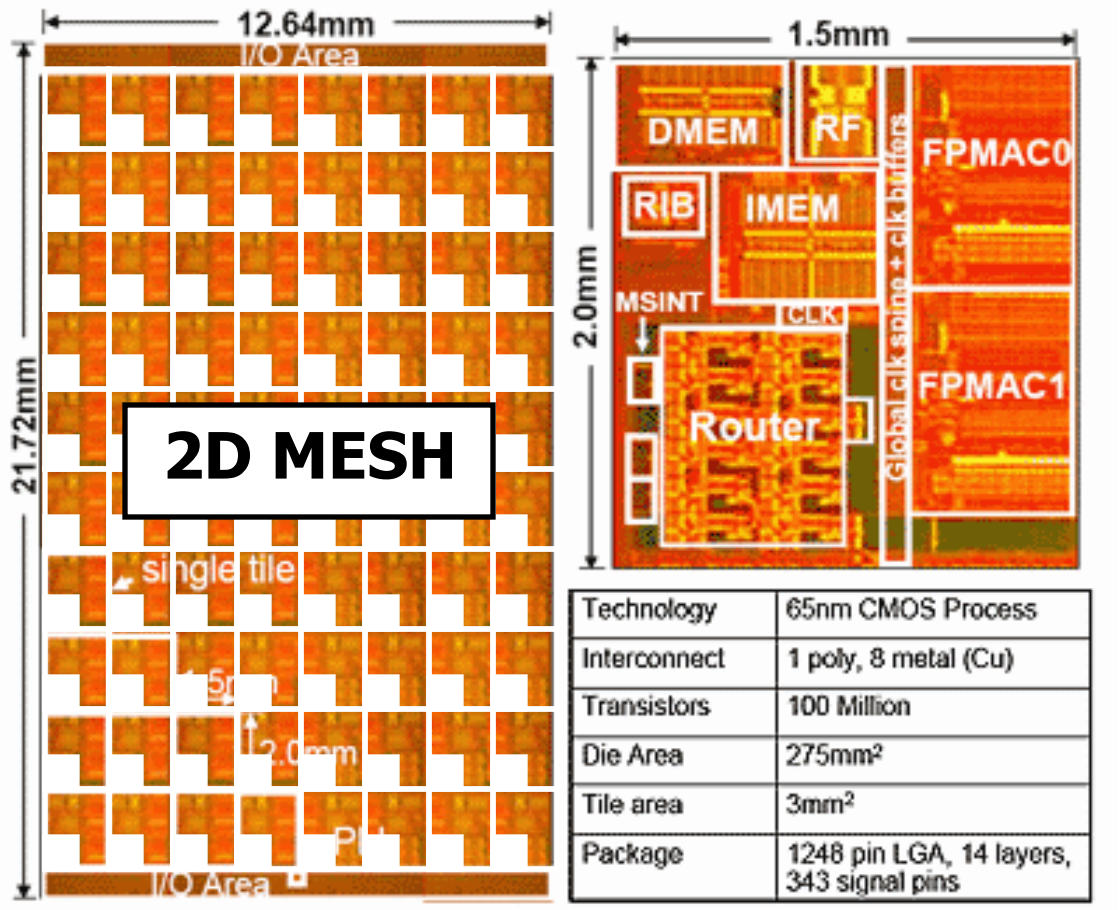

Technology	65nm CMOS Process
Interconnect	1 poly, 8 metal (Cu)
Transistors	100 Million
Die Area	275mm²
Tile area	3mm²
Package	1248 pin LGA, 14 layers, 343 signal pins

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

Many-Core: Intel SCC

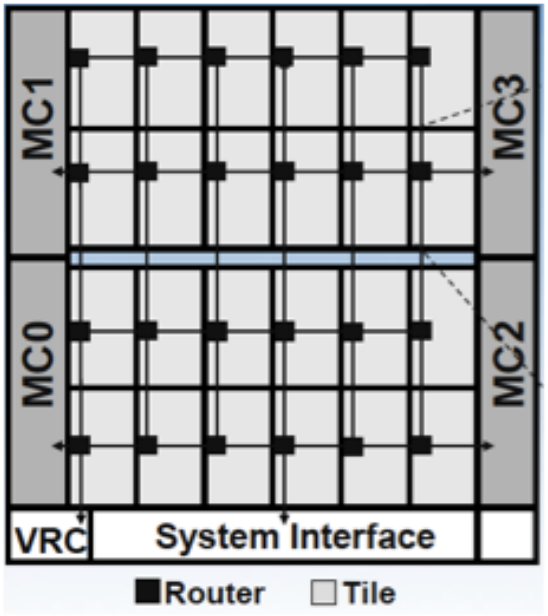

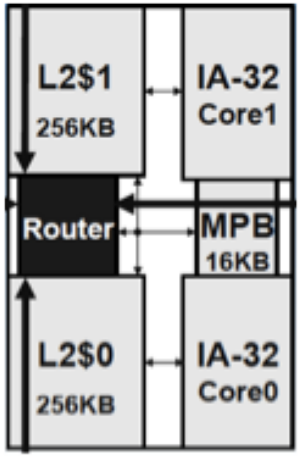

Courtesy: Jason Howard, Intel

- Intel's Single-chip Cloud Computer (SCC) uses a 2D mesh with state of the art routers
  - Will discuss in detail later in lecture
  - Potential platform for project

Performance and Cost

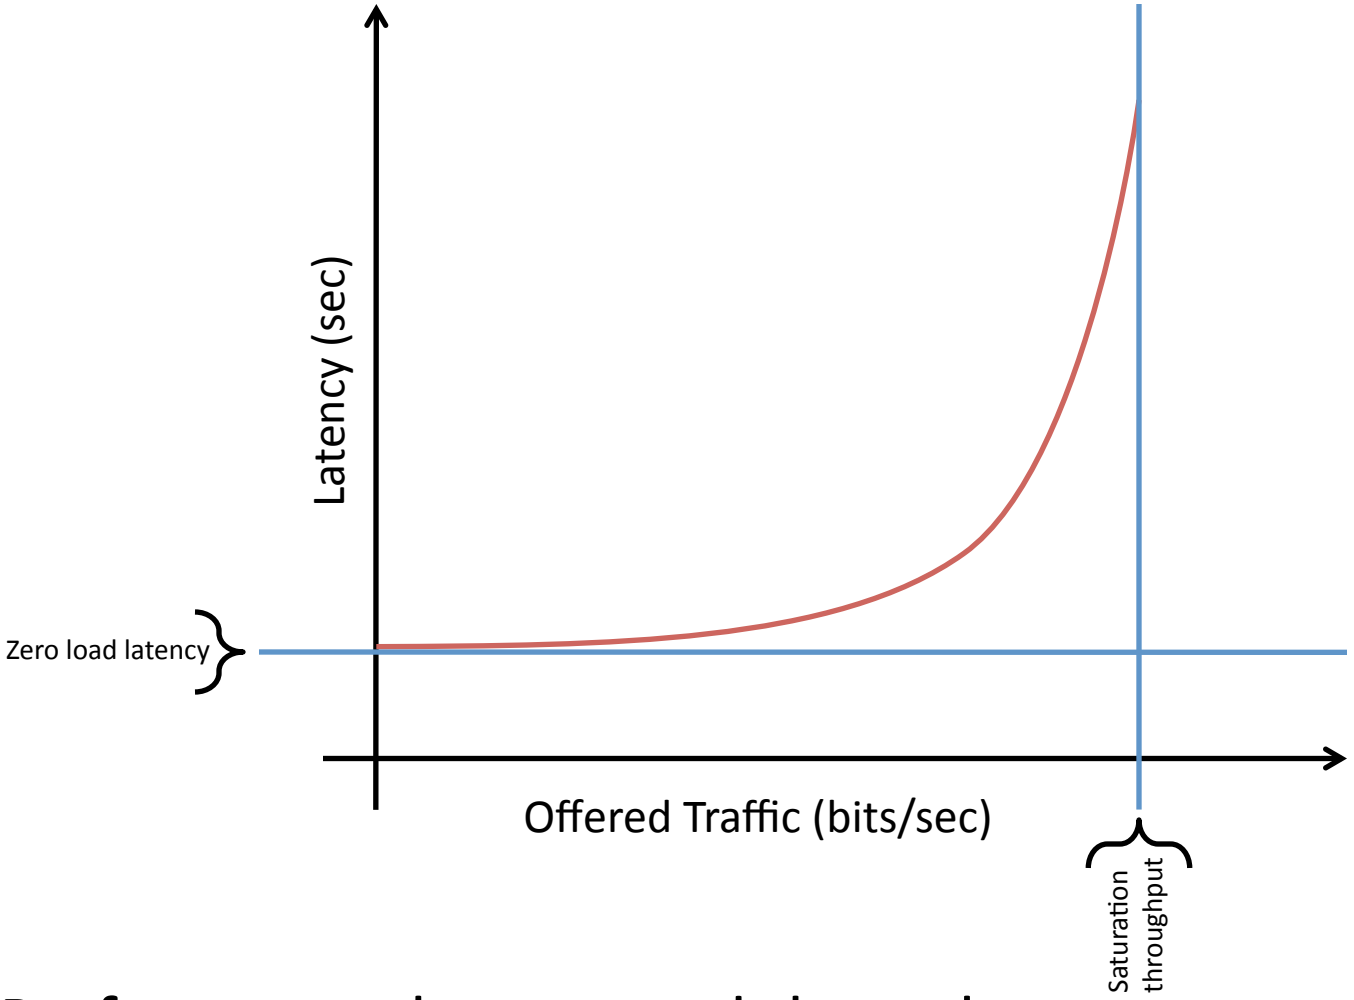

- Performance: latency and throughput
- Cost: area and power

Course Information

- Website: www.eecg.toronto.edu/~enright/1749h/index.html
- Lecture: Wed 10-12, BA 4164
- Contact Information:
  - enright@eecg.toronto.edu
  - Office: 374A Pratt
  - Stop by or send me email for appointment

Course Format

- ~5 weeks of lecture
  - Covering material from book: On-Chip Networks
    - Available for free download (within UofT)

- Remaining weeks:
  - Presentation of research papers by you
    - Should be interactive/foster discussion
  - Critiques (1 page)
    - identify one idea in paper that is major contribution or major limitation
    - describe new idea (of yours) that builds on paper

Course Marking Scheme

- Evaluation Scheme:
  - Presentation: 20%
  - Critiques: 35%
  - Project: 45%

Course Project

- Propose your own topic
  - Topics do not need to be limited to OCN
  - Look for ways to relate OCN to your research
- Re-evaluate recent paper (or compare 2 papers)
  - Can use a paper from class
  - Or see bibliographic notes in book for ideas
  - Or check out recent conferences: ISCA, MICRO, HPCA, NOCS etc.
- Can work in groups of 2
  - Project must be sized accordingly
- Enrolled in 1755? – Parallel Computer Architecture and Programming
  - Do one larger project that spans both classes
    - Must get approval from Prof. Steffan as well
- See website for details

Next Time

- Look at Topology
- Discuss various potential projects
  - Not too early to start thinking about your project and bring questions
    - Send me e-mail/stop by if you want to discuss your project early
  - Project should
    - Reinforce/enhance class material
    - Add value to your research – publish at conference/ workshop?