ECE 5970
Chip-Level Interconnection Networks
Lecture 1: Course Overview

Christopher Batten
Cornell University
January 26, 2010

`http://www.csl.cornell.edu/courses/ece5970`

Agenda

- Course Motivation
- Interconnection Network Basics
  - Topology
  - Routing
  - Flow Control
  - Router Microarchitecture
  - Examples
- Course Logistics

What is an Interconnection Network?

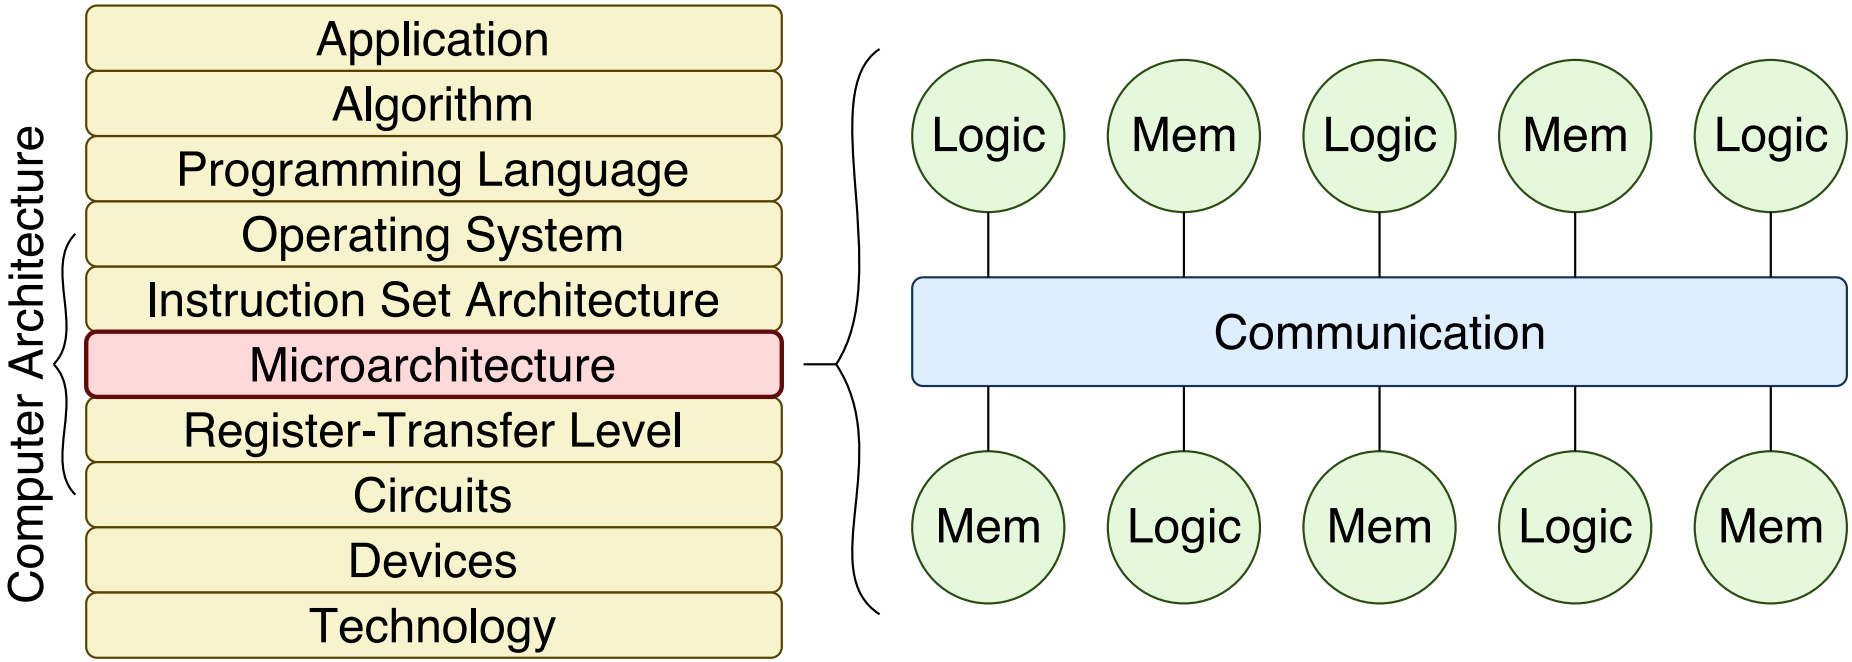

What is an Interconnection Network?

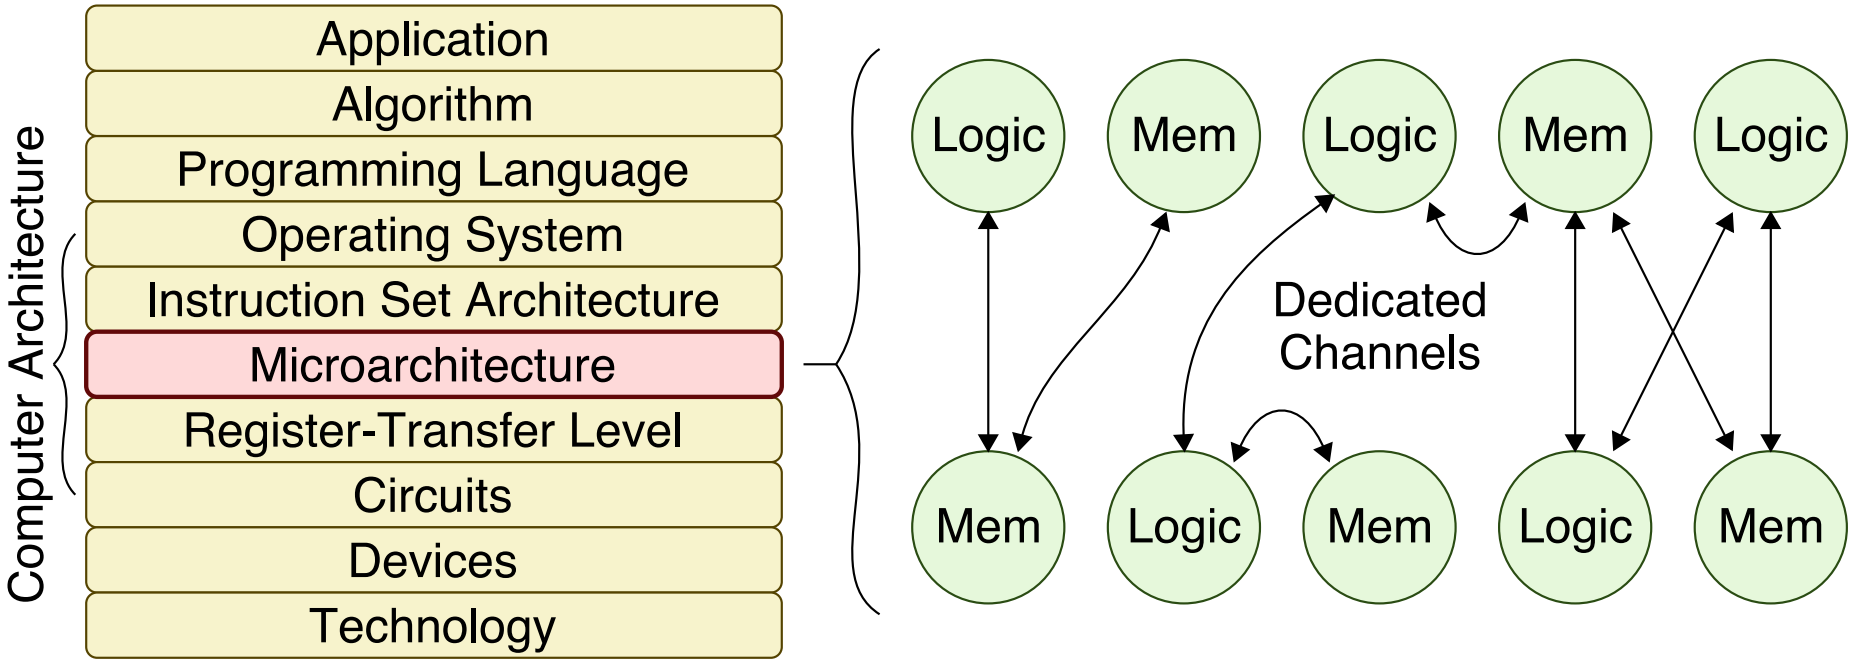

Application: Ideally wants low-latency, high-bandwidth, dedicated channels between logic and memory

Technology: Dedicated channels too expensive in terms of area and power

What is an Interconnection Network?

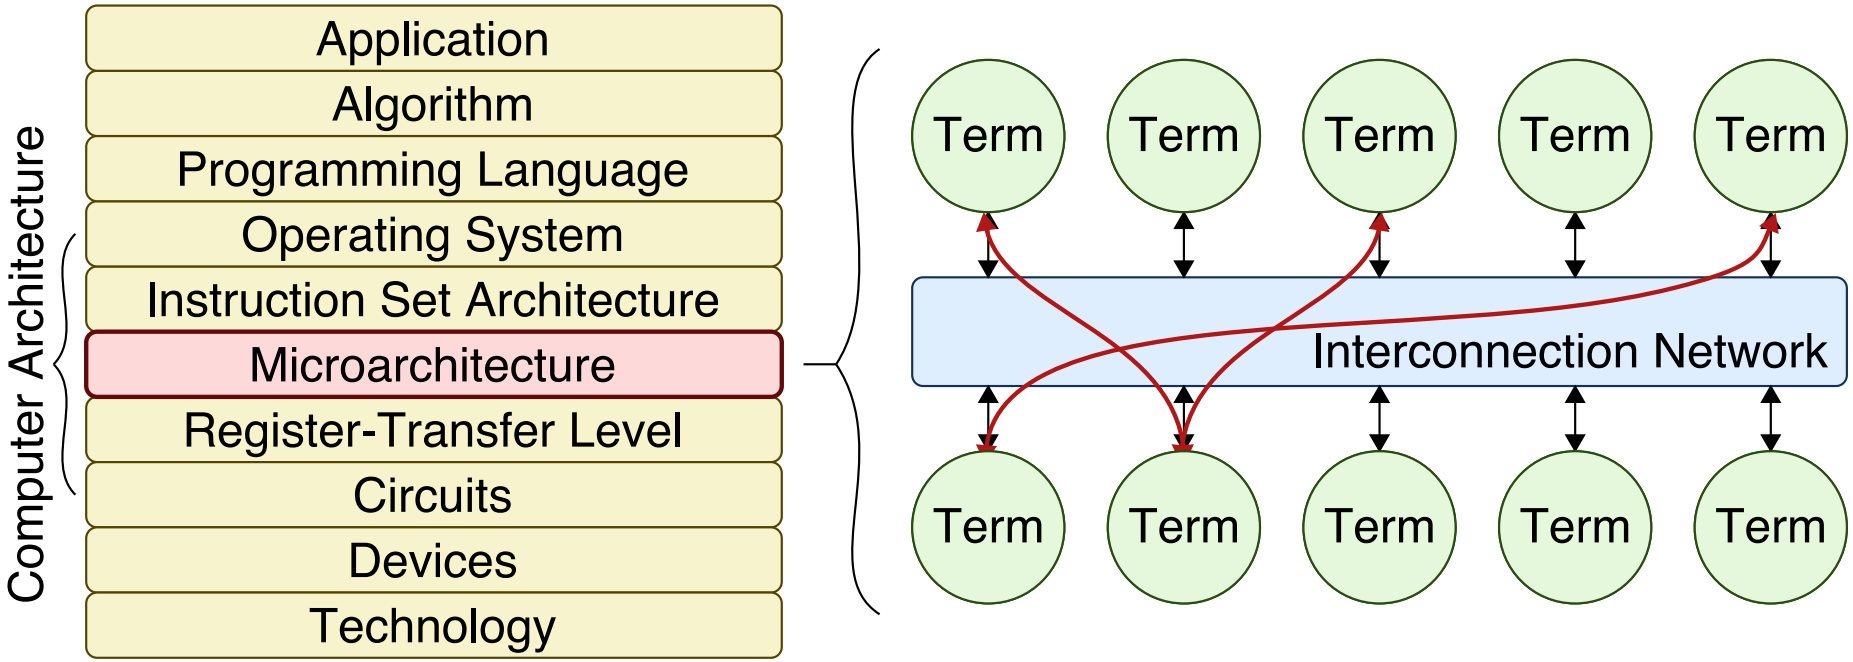

An Interconnection Network is a programmable system that transports data between terminals

Technology: Interconnection network helps efficiently utilize scare resources such as area and power

Application: Managing interconnection network can be critical to achieving good performance

What is an Interconnection Network?

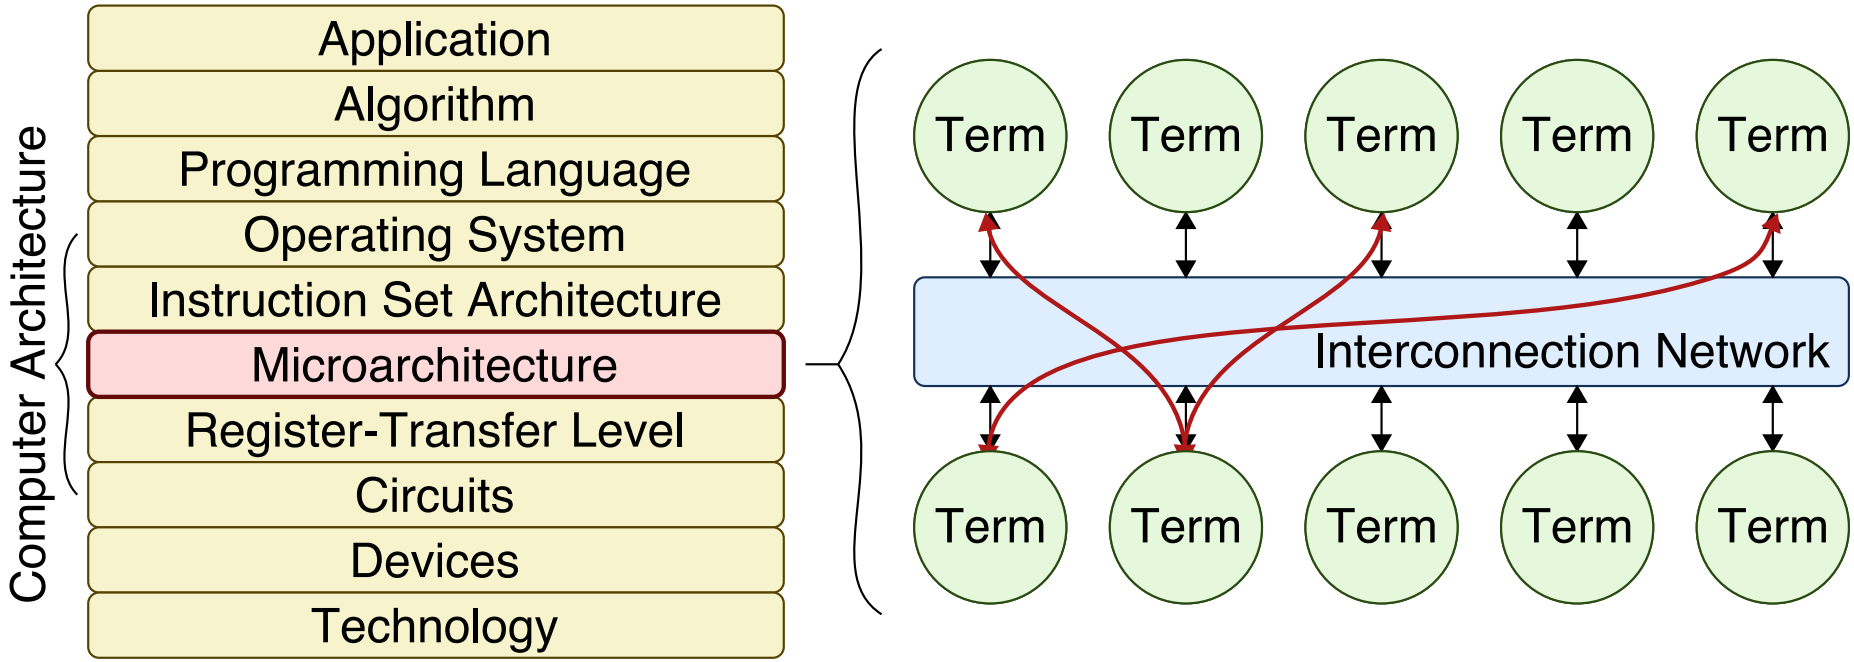

One Theme for Course:
Interplay between application requirements, technology constraints, and interconnection networks

Types of Interconnection Networks

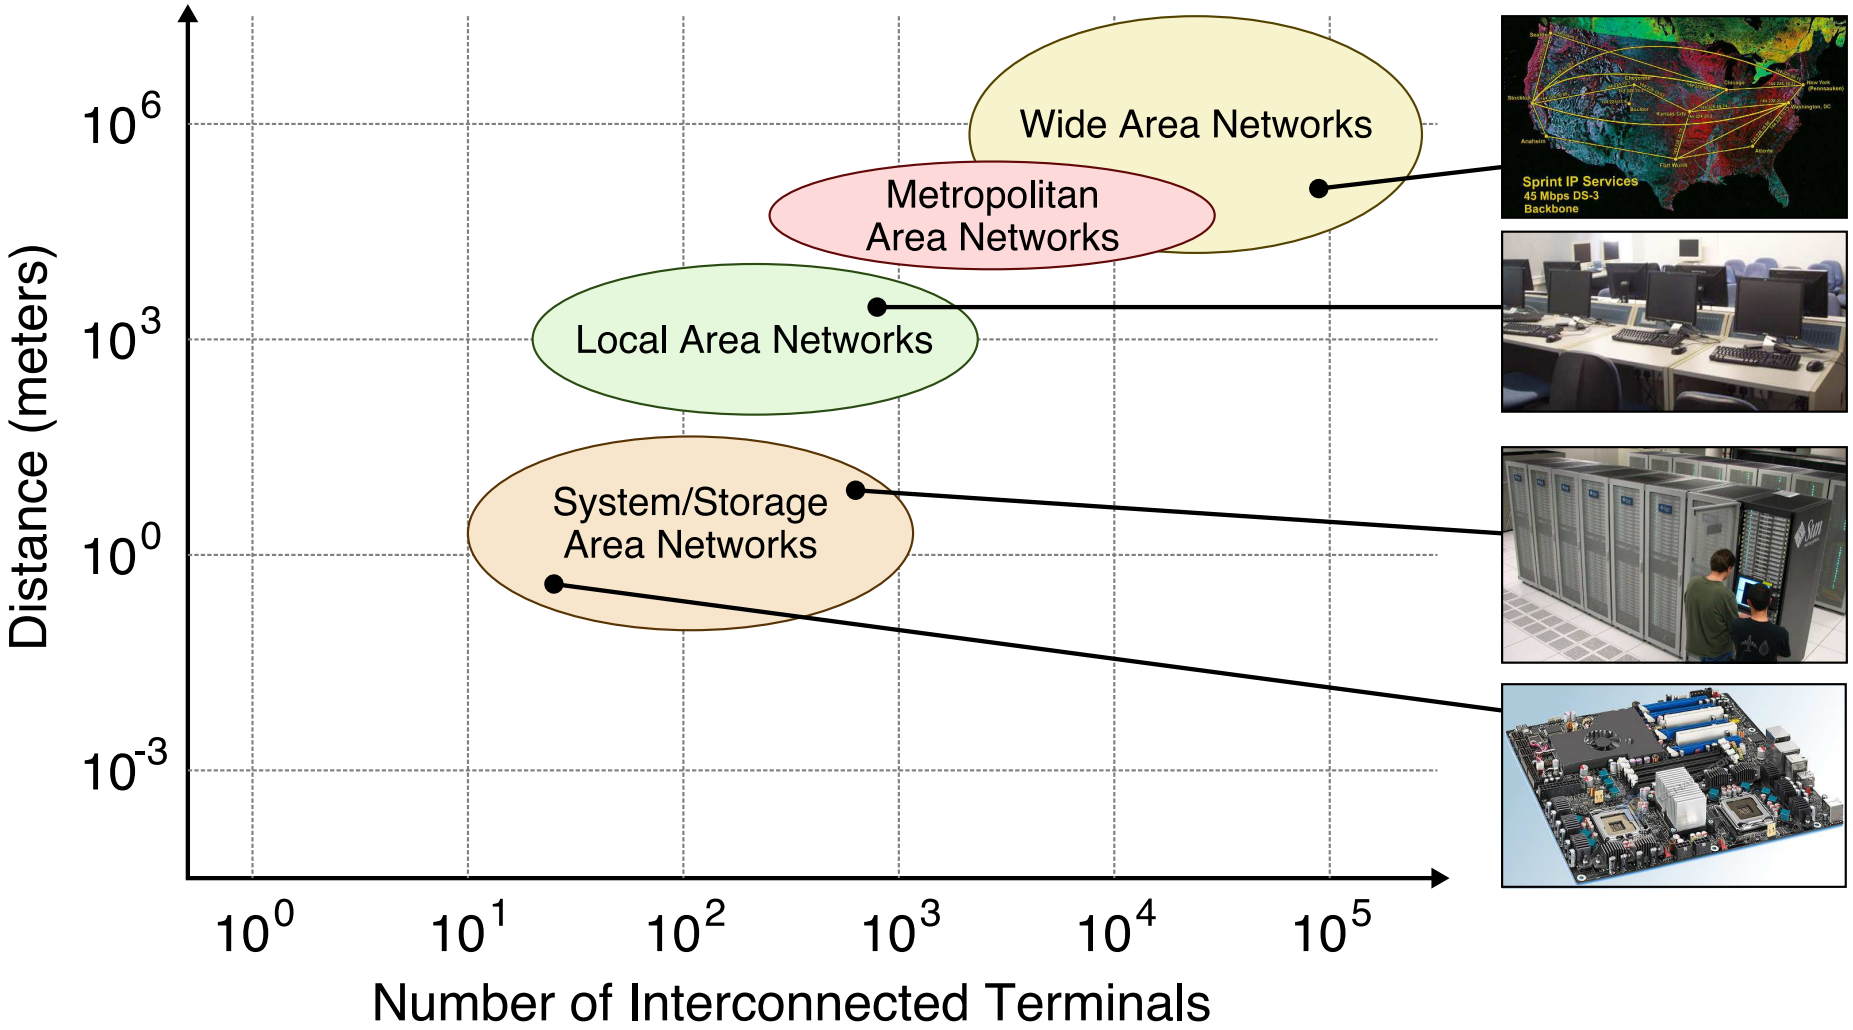

Types of Interconnection Networks

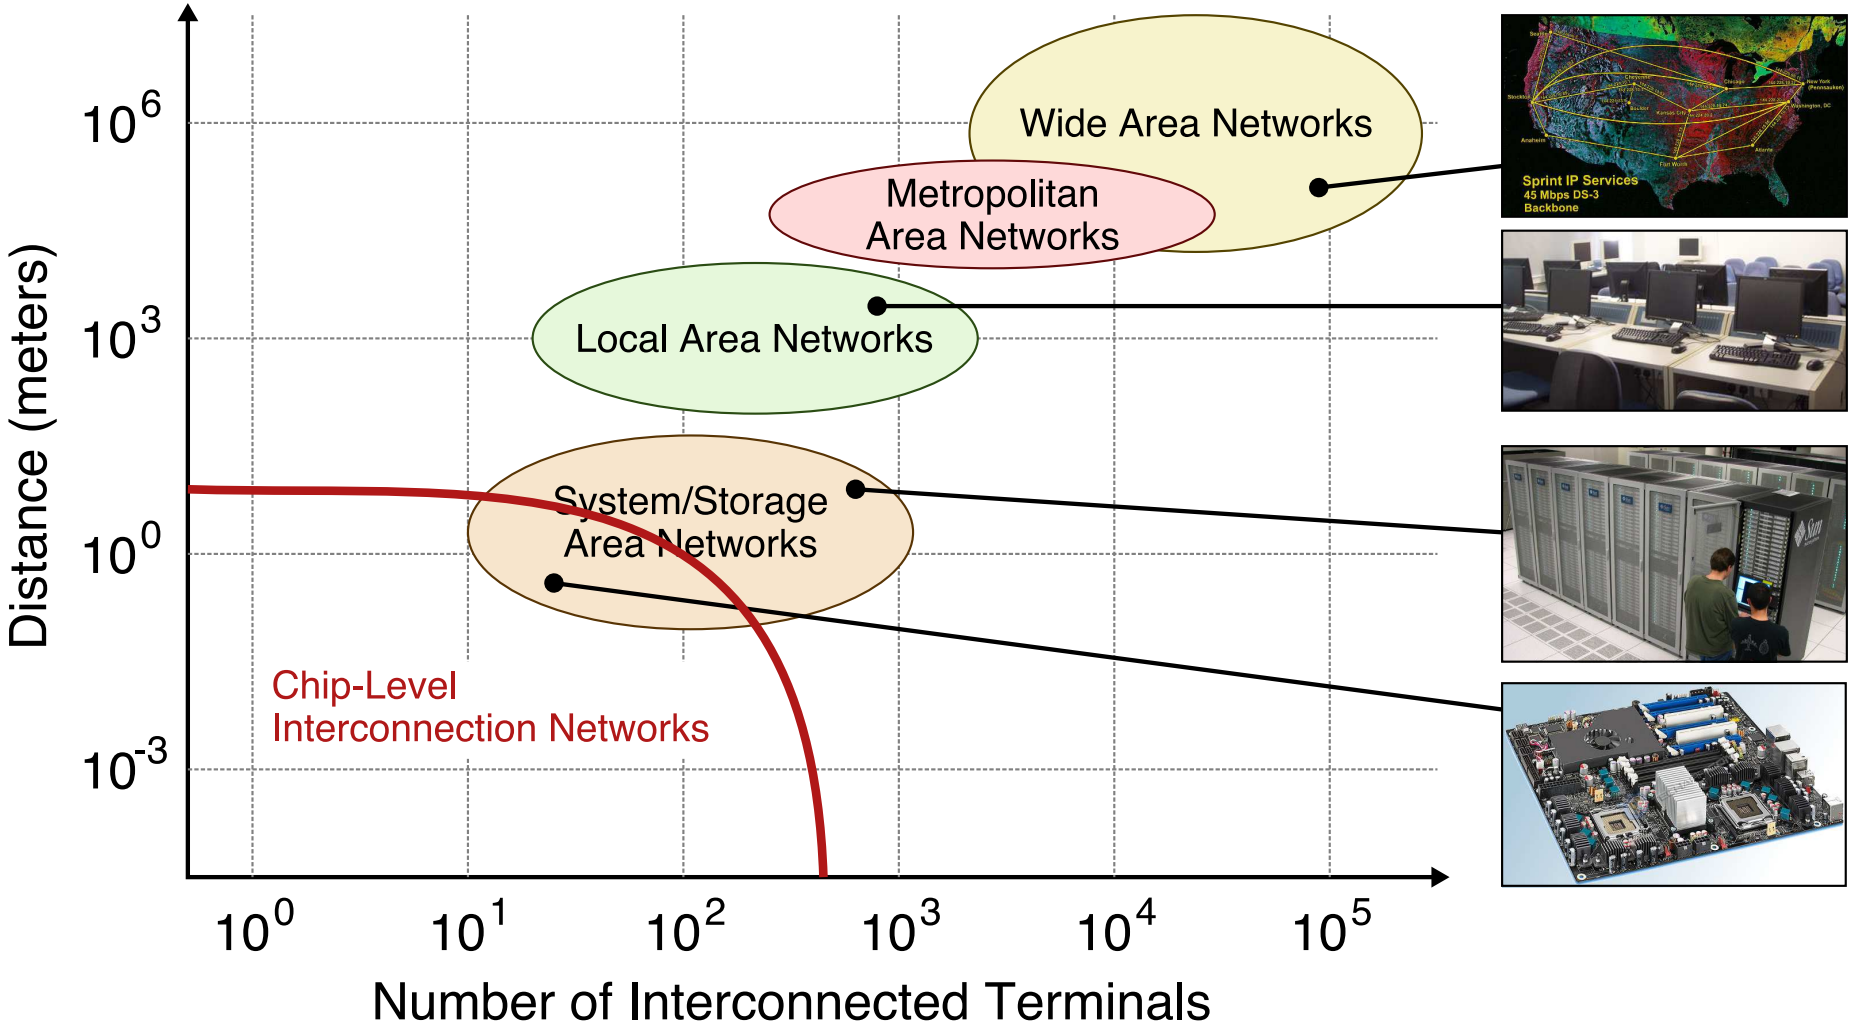

Irregular Threaded Application Running on Processor-to-Memory Network

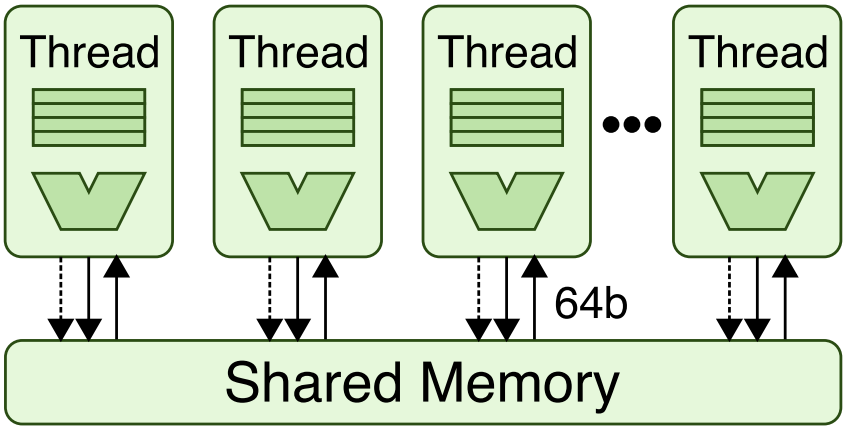

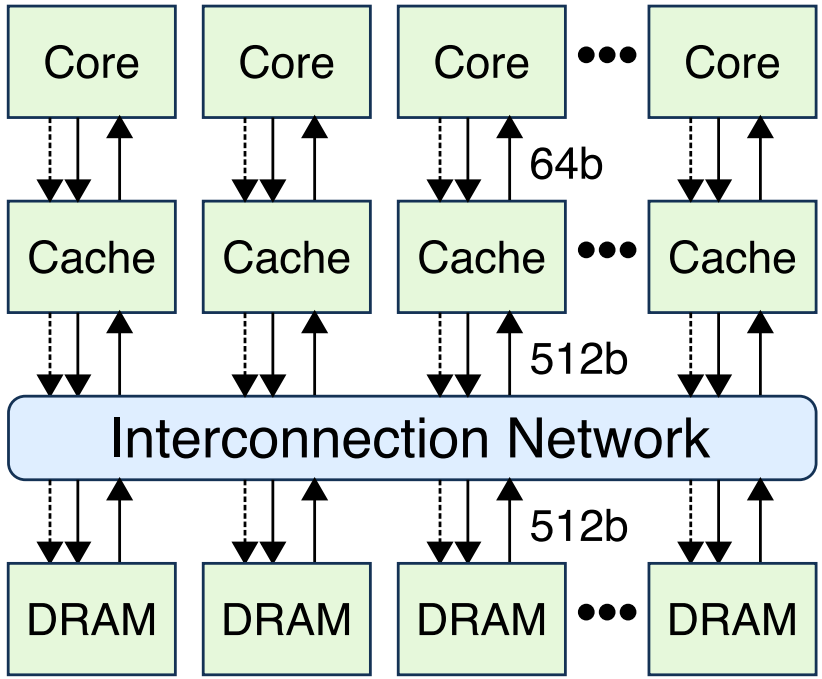

Application Requirements

Message Size	512b
Average Bandwidth	400 MB/s
Peak Bandwidth	8 GB/s
Latency	Minimum
Traffic Pattern	Arbitrary

What network design meets these requirements within the technology constraints and with the least area, power, and lowest latency?

Streaming Application Running on Processor-to-Processor Network

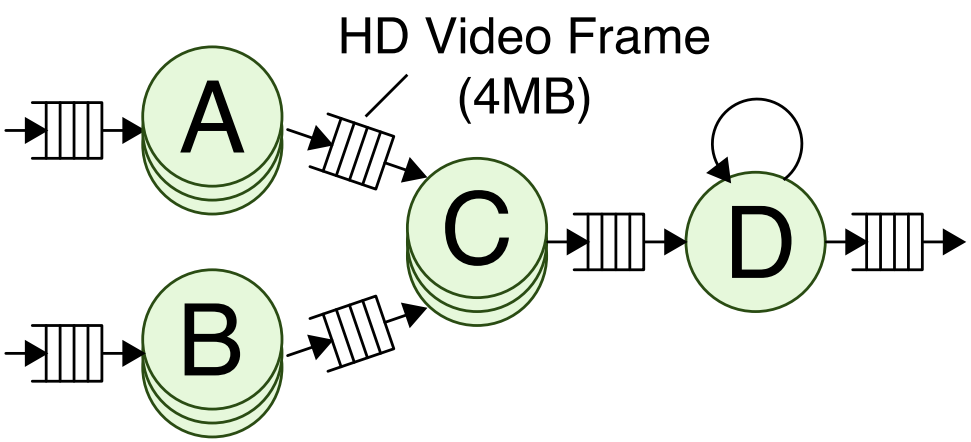

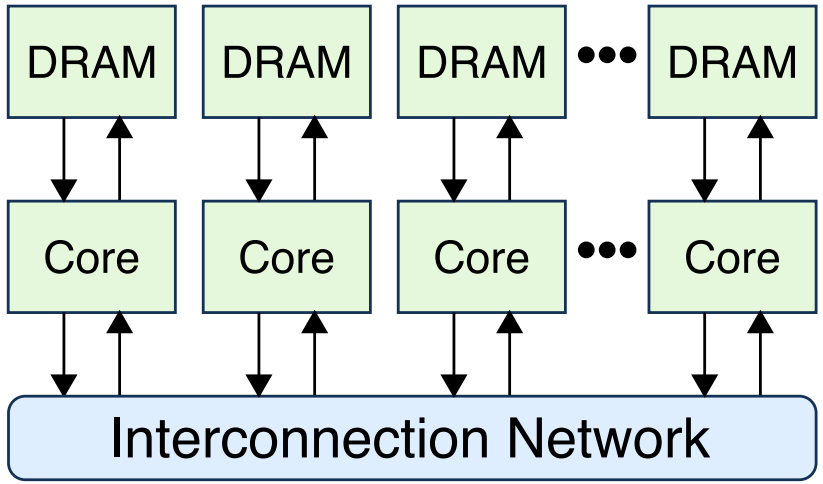

Application Requirements

Message Size	4MB
Average Bandwidth	120 MB/s
Peak Bandwidth	120 MB/s
Latency	Tolerant
Traffic Pattern	Streaming

What network design meets these requirements within the technology constraints and with the least area, power, and maximum bandwidth?

Goal: Flexible Networks Capable of Running Many Different Kinds of Applications

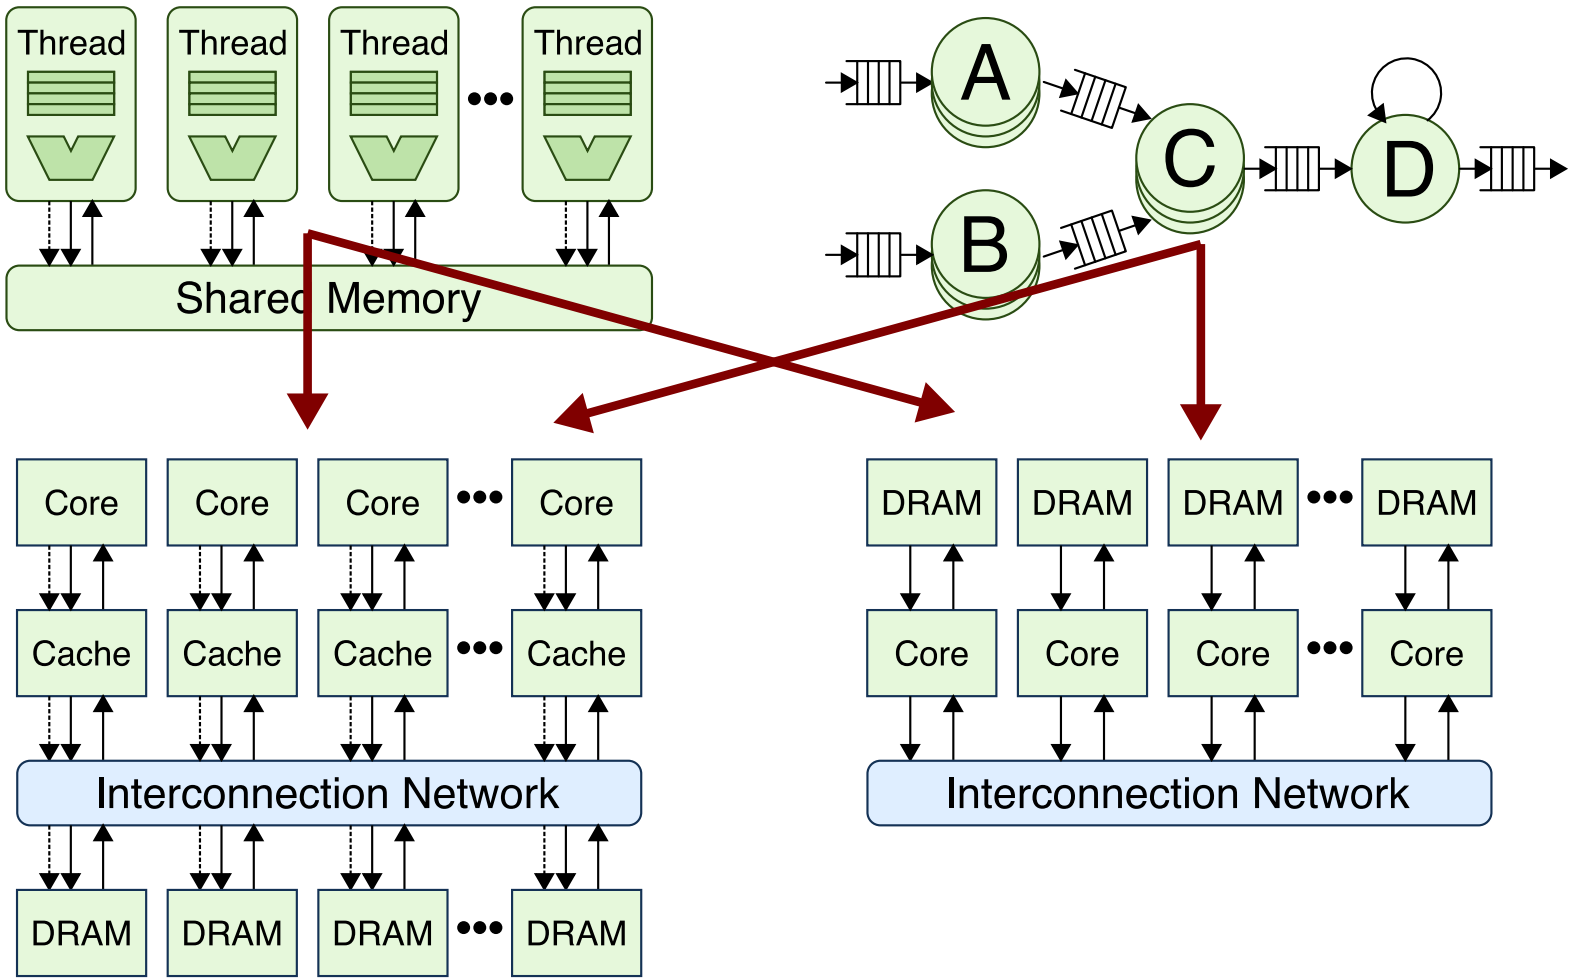

Why Study Chip-Level Networks Now?

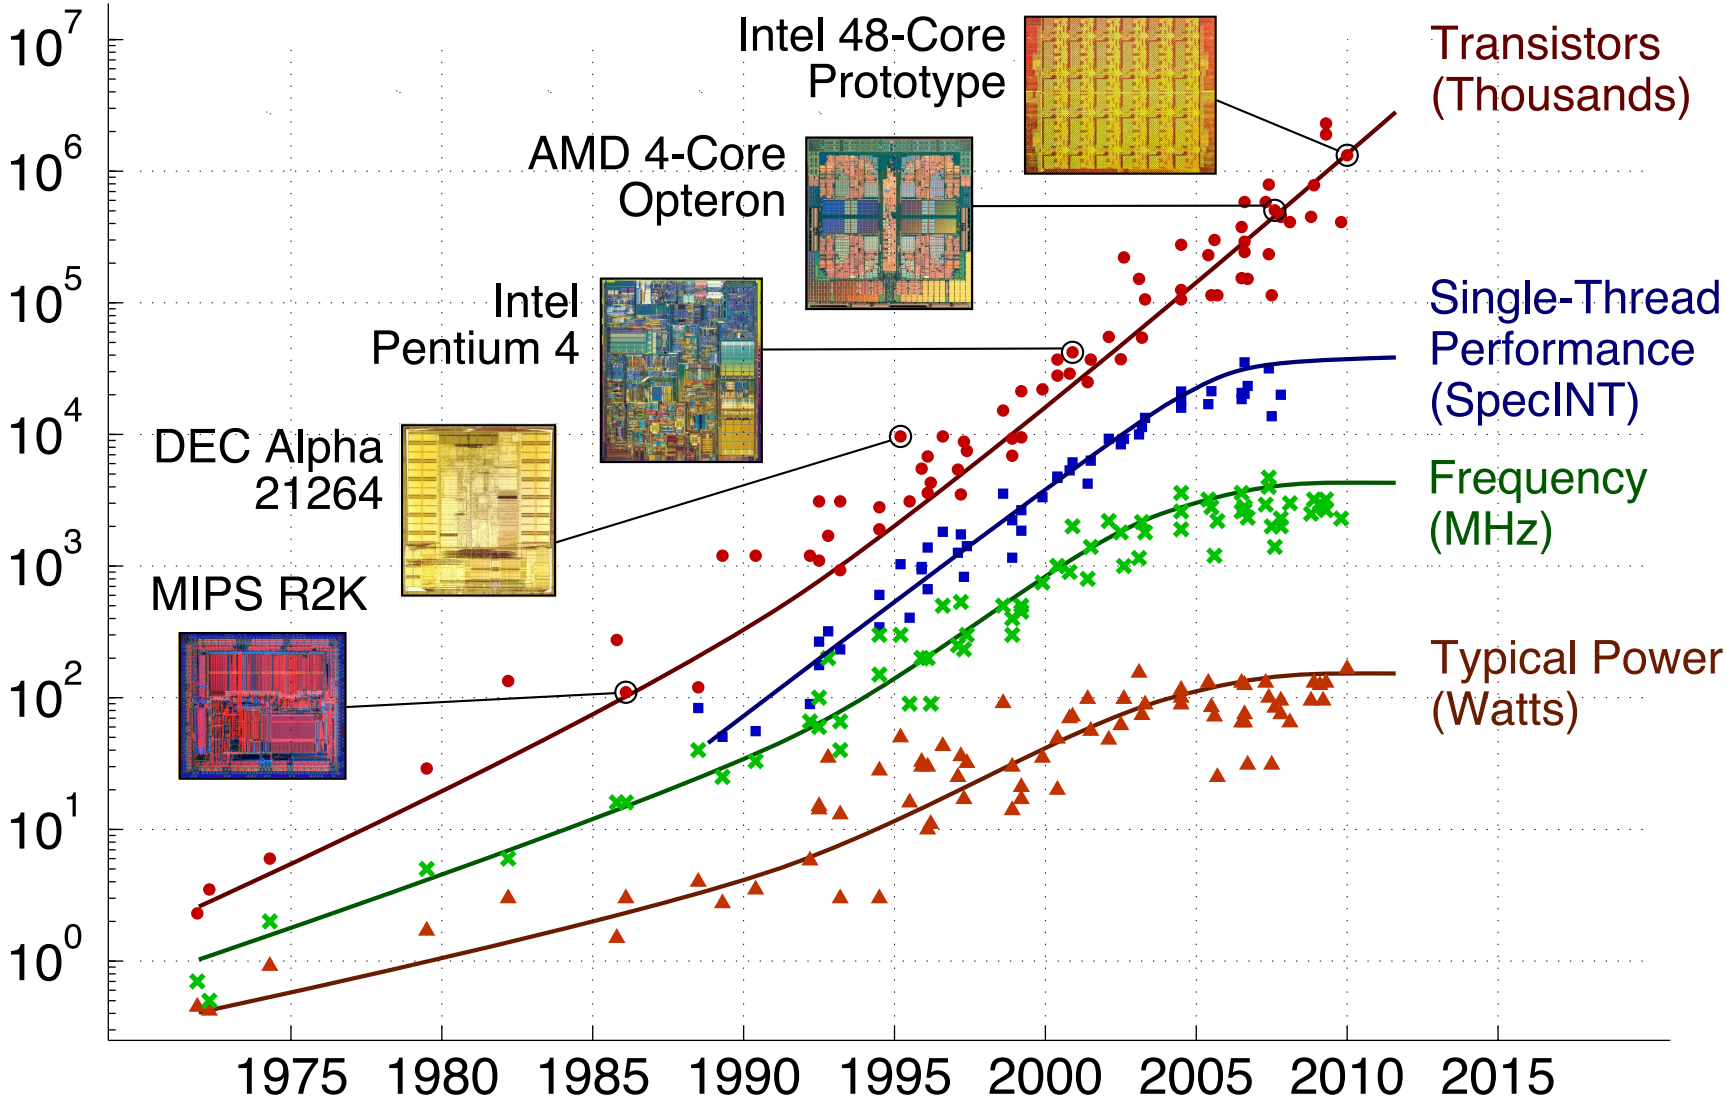

Data partially collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond

Why Study Chip-Level Networks Now?

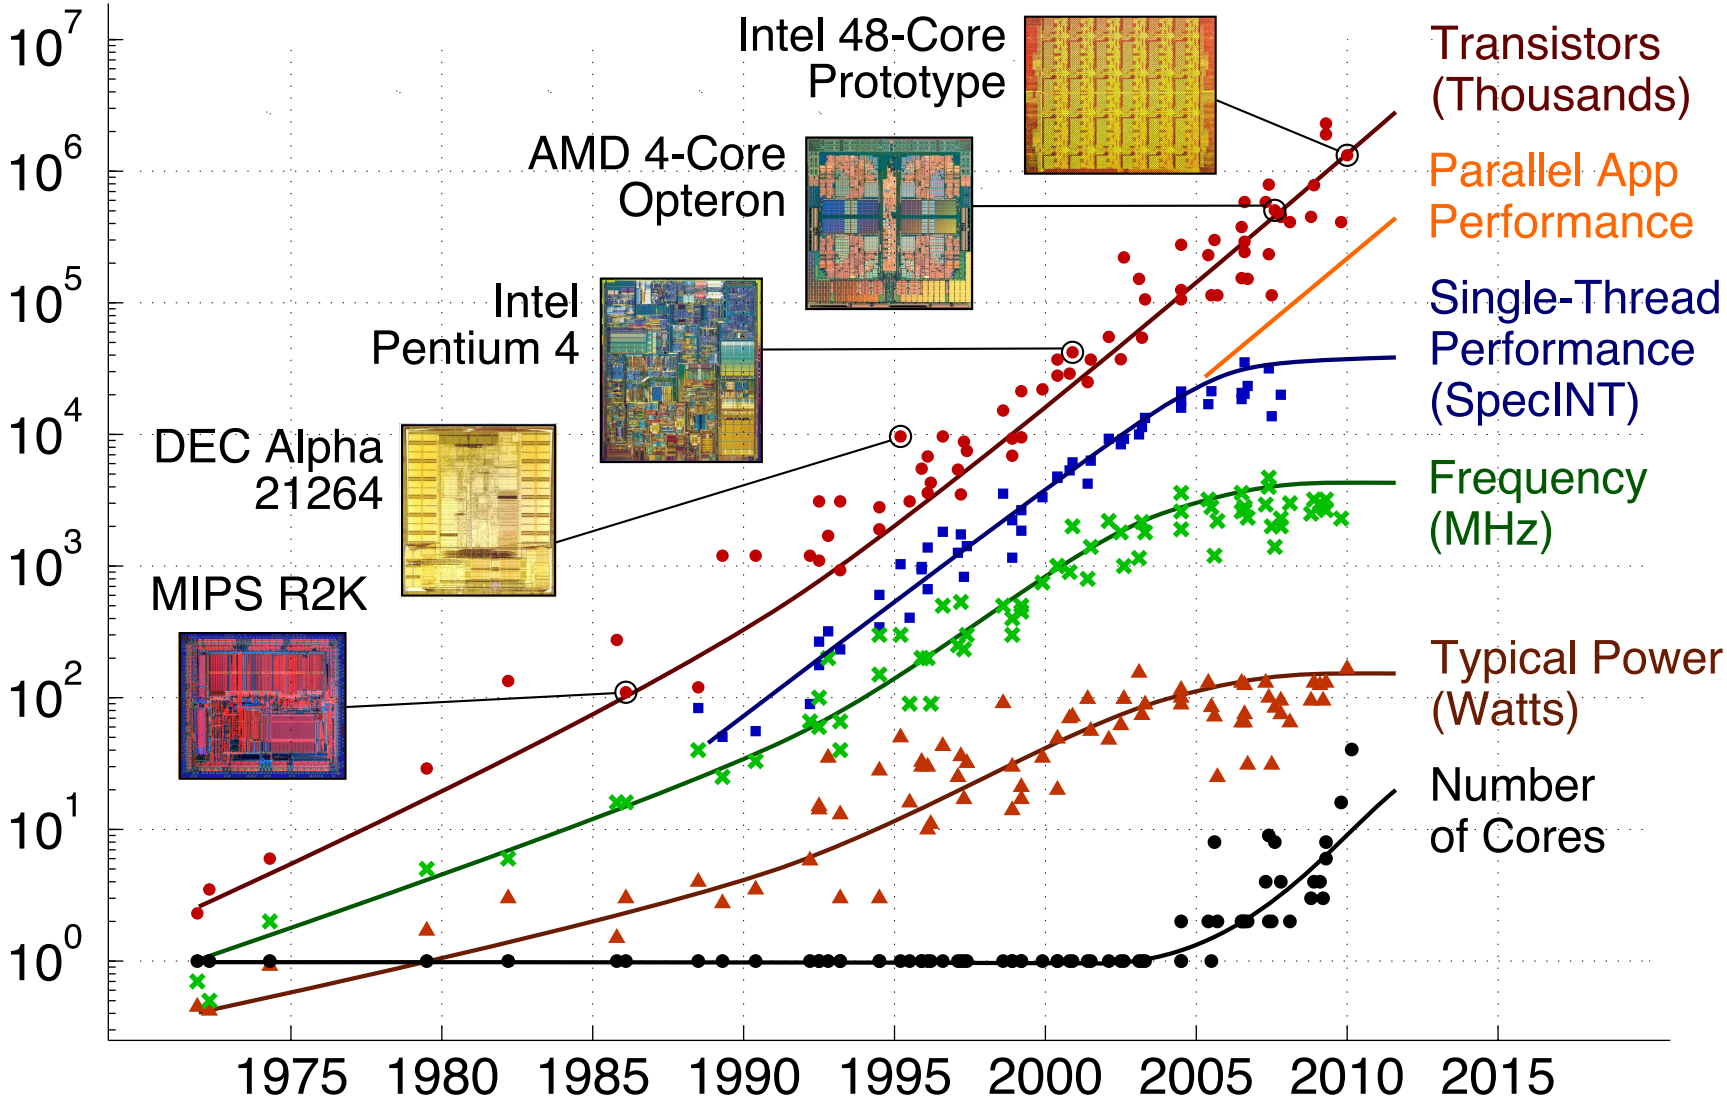

Data partially collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond

Examples of Multicore and Manycore Processors

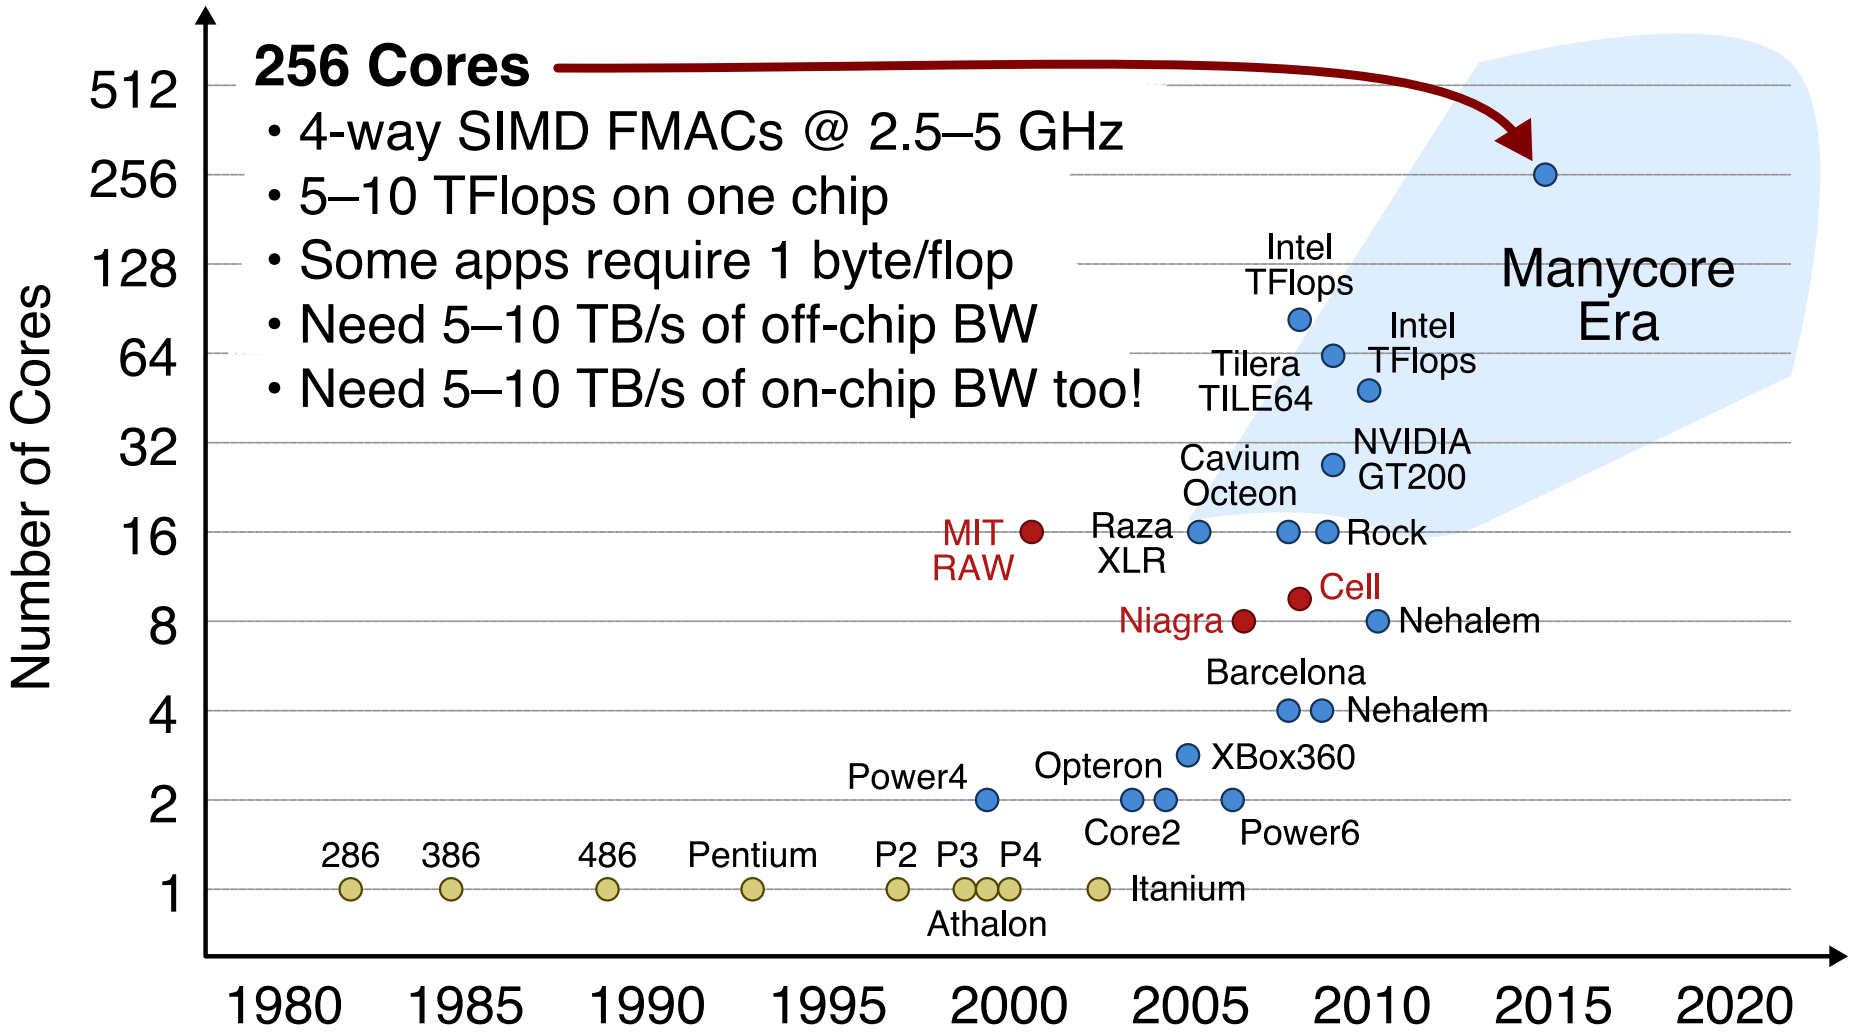

Energy and Performance of Multicore and Manycore

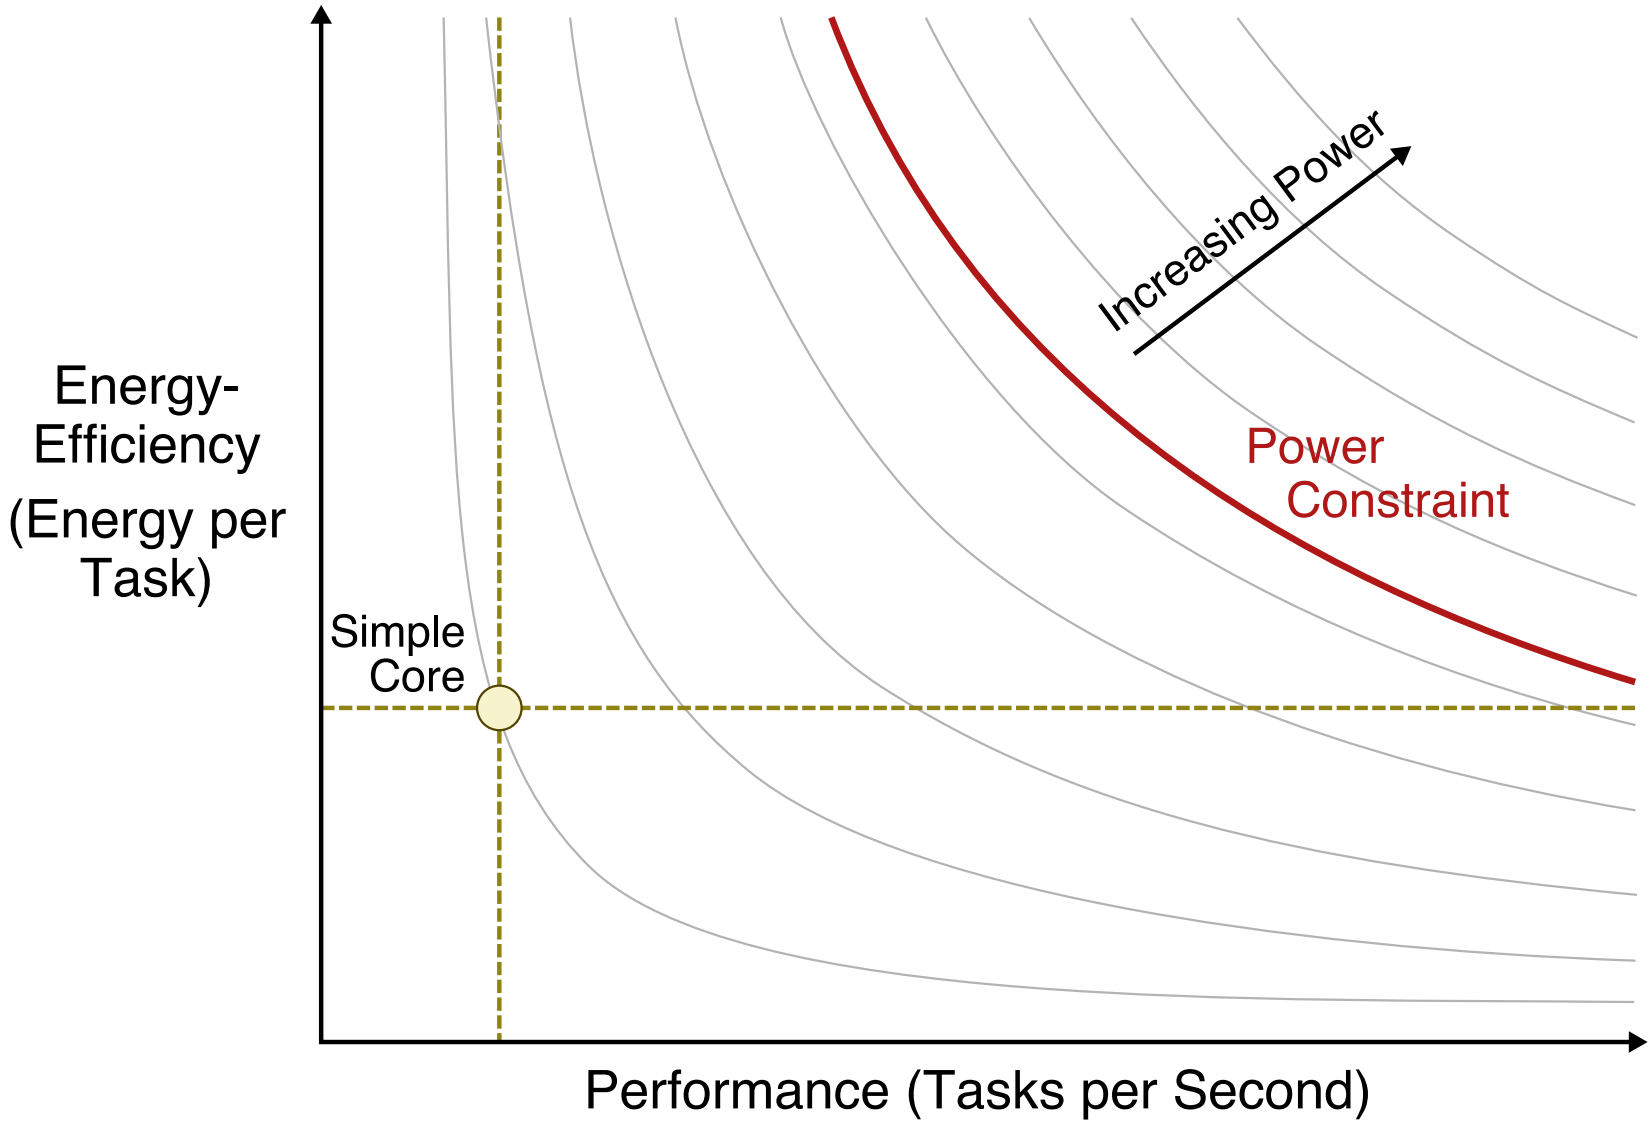

Energy and Performance of Multicore and Manycore

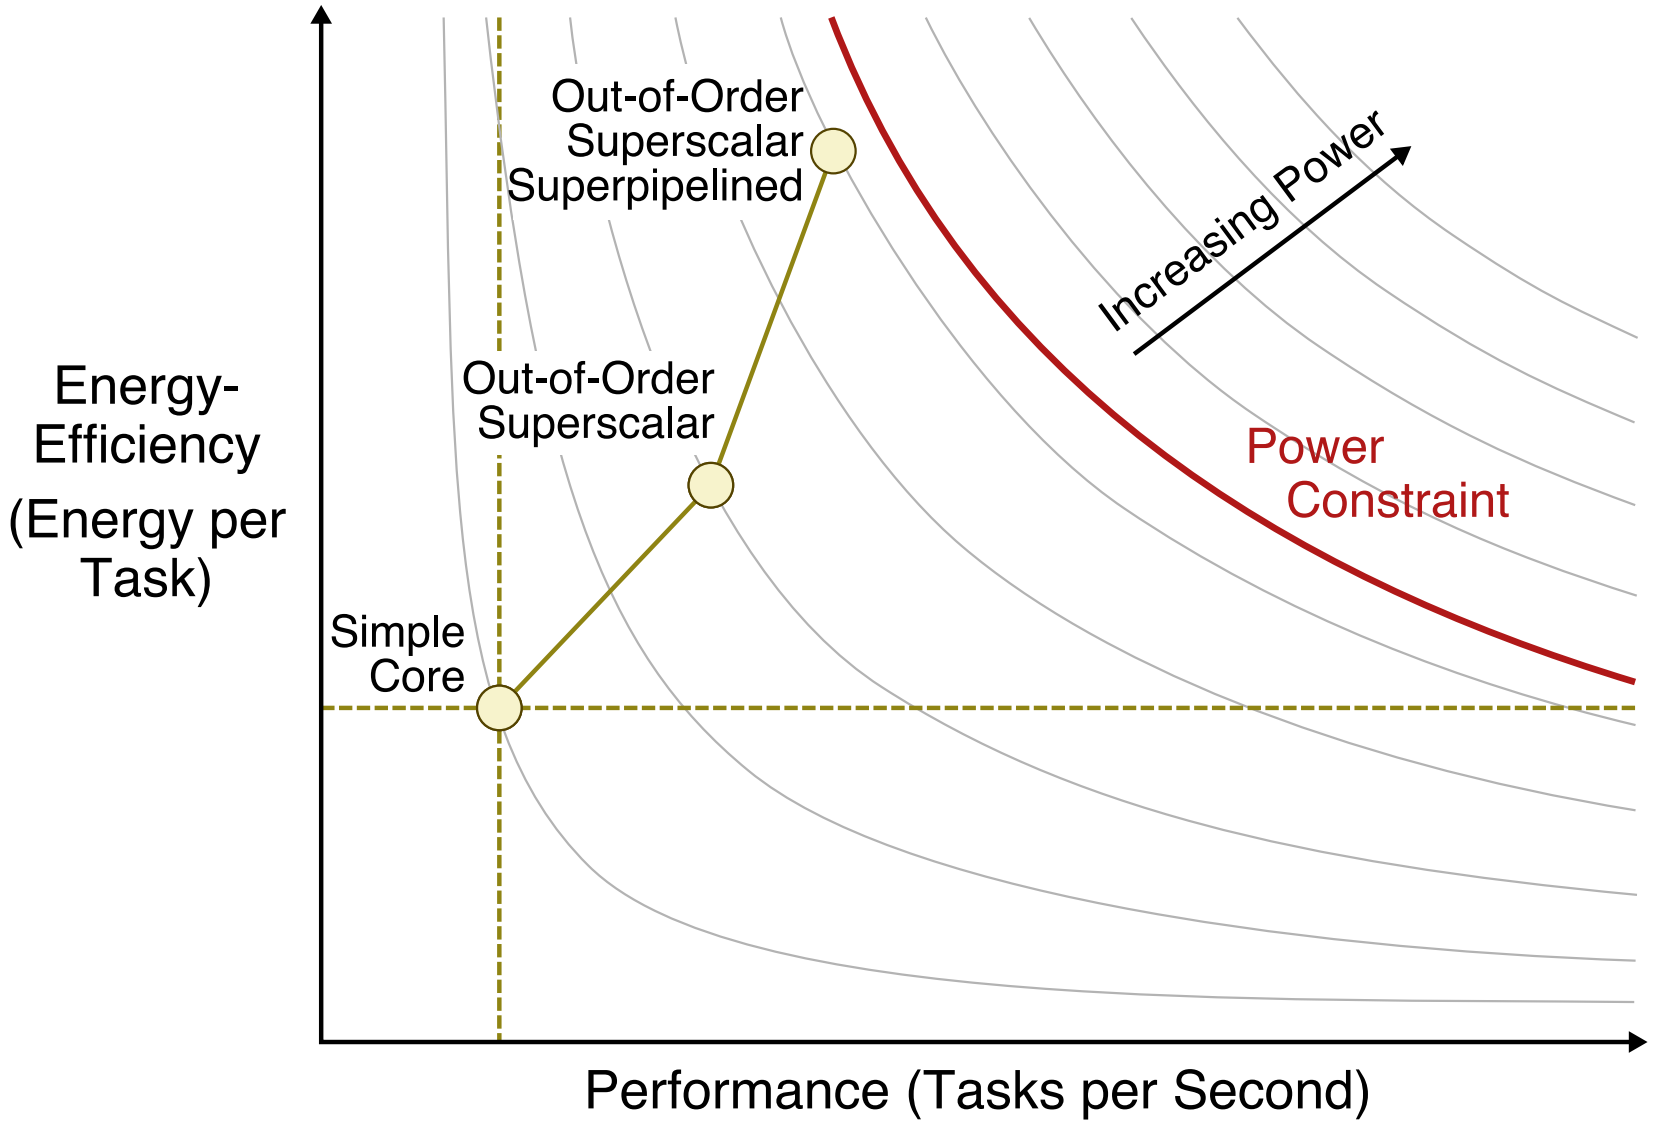

Energy and Performance of Multicore and Manycore

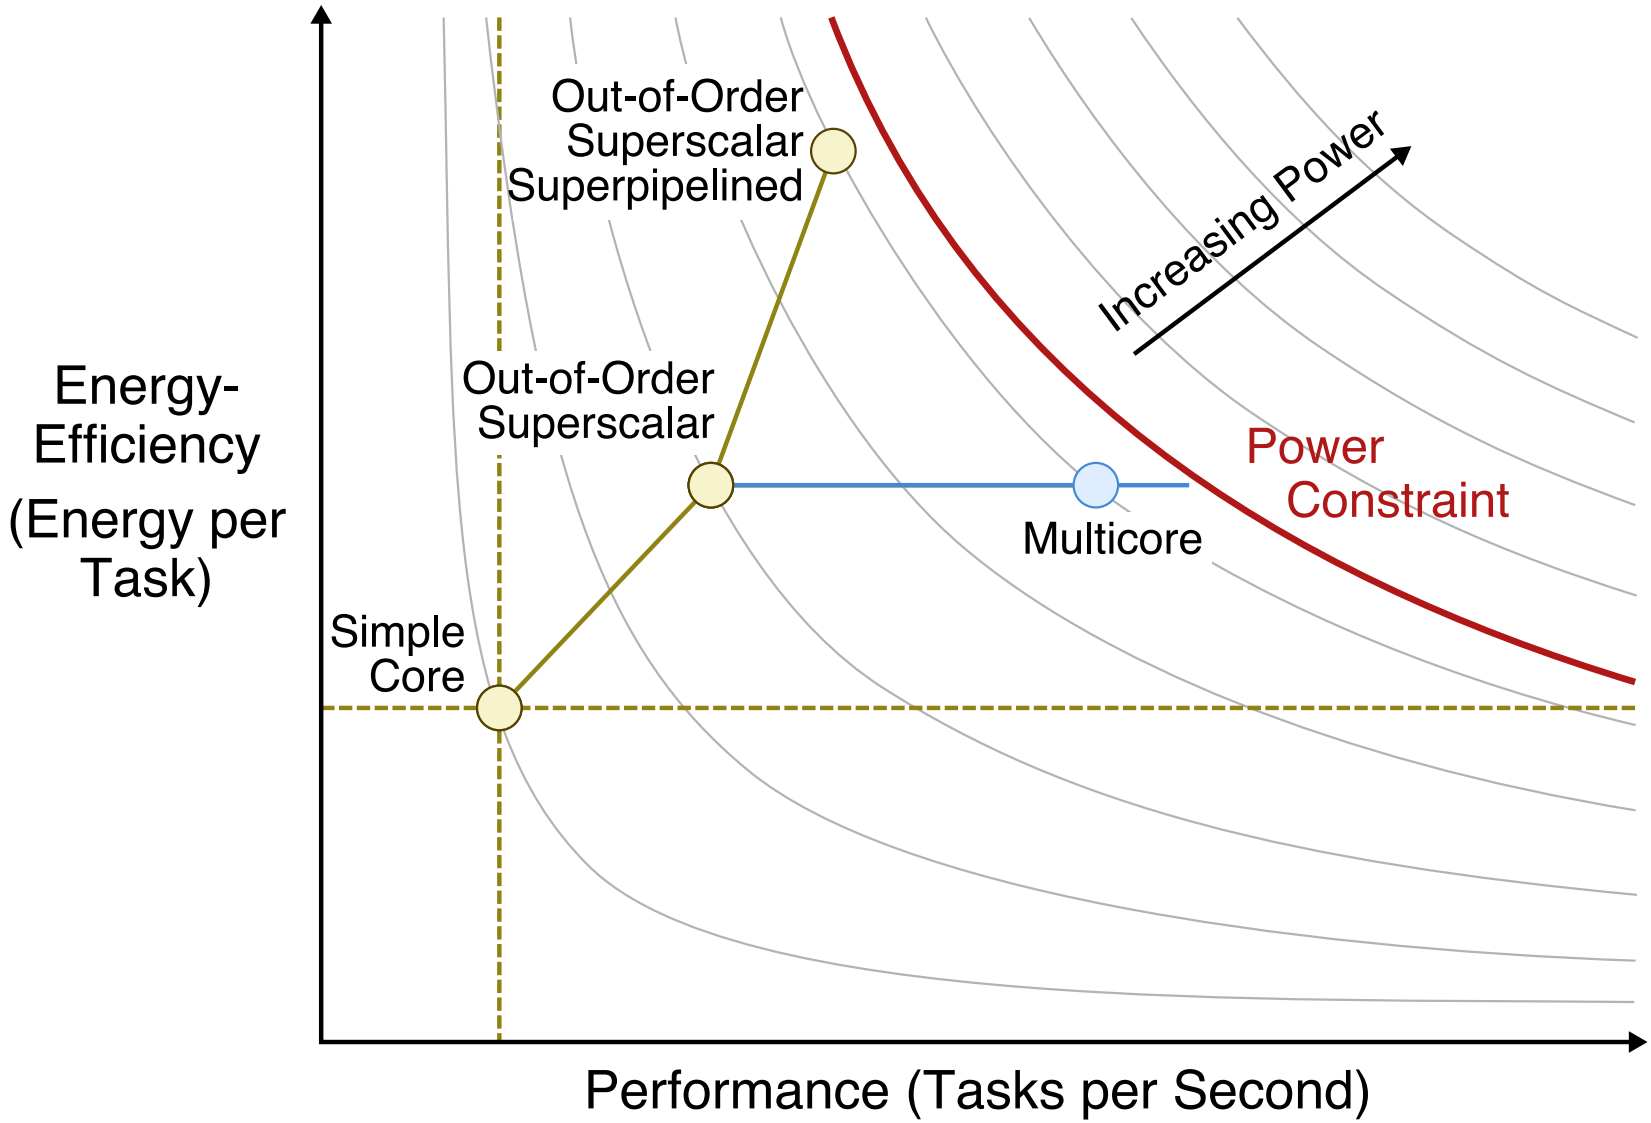

Energy and Performance of Multicore and Manycore

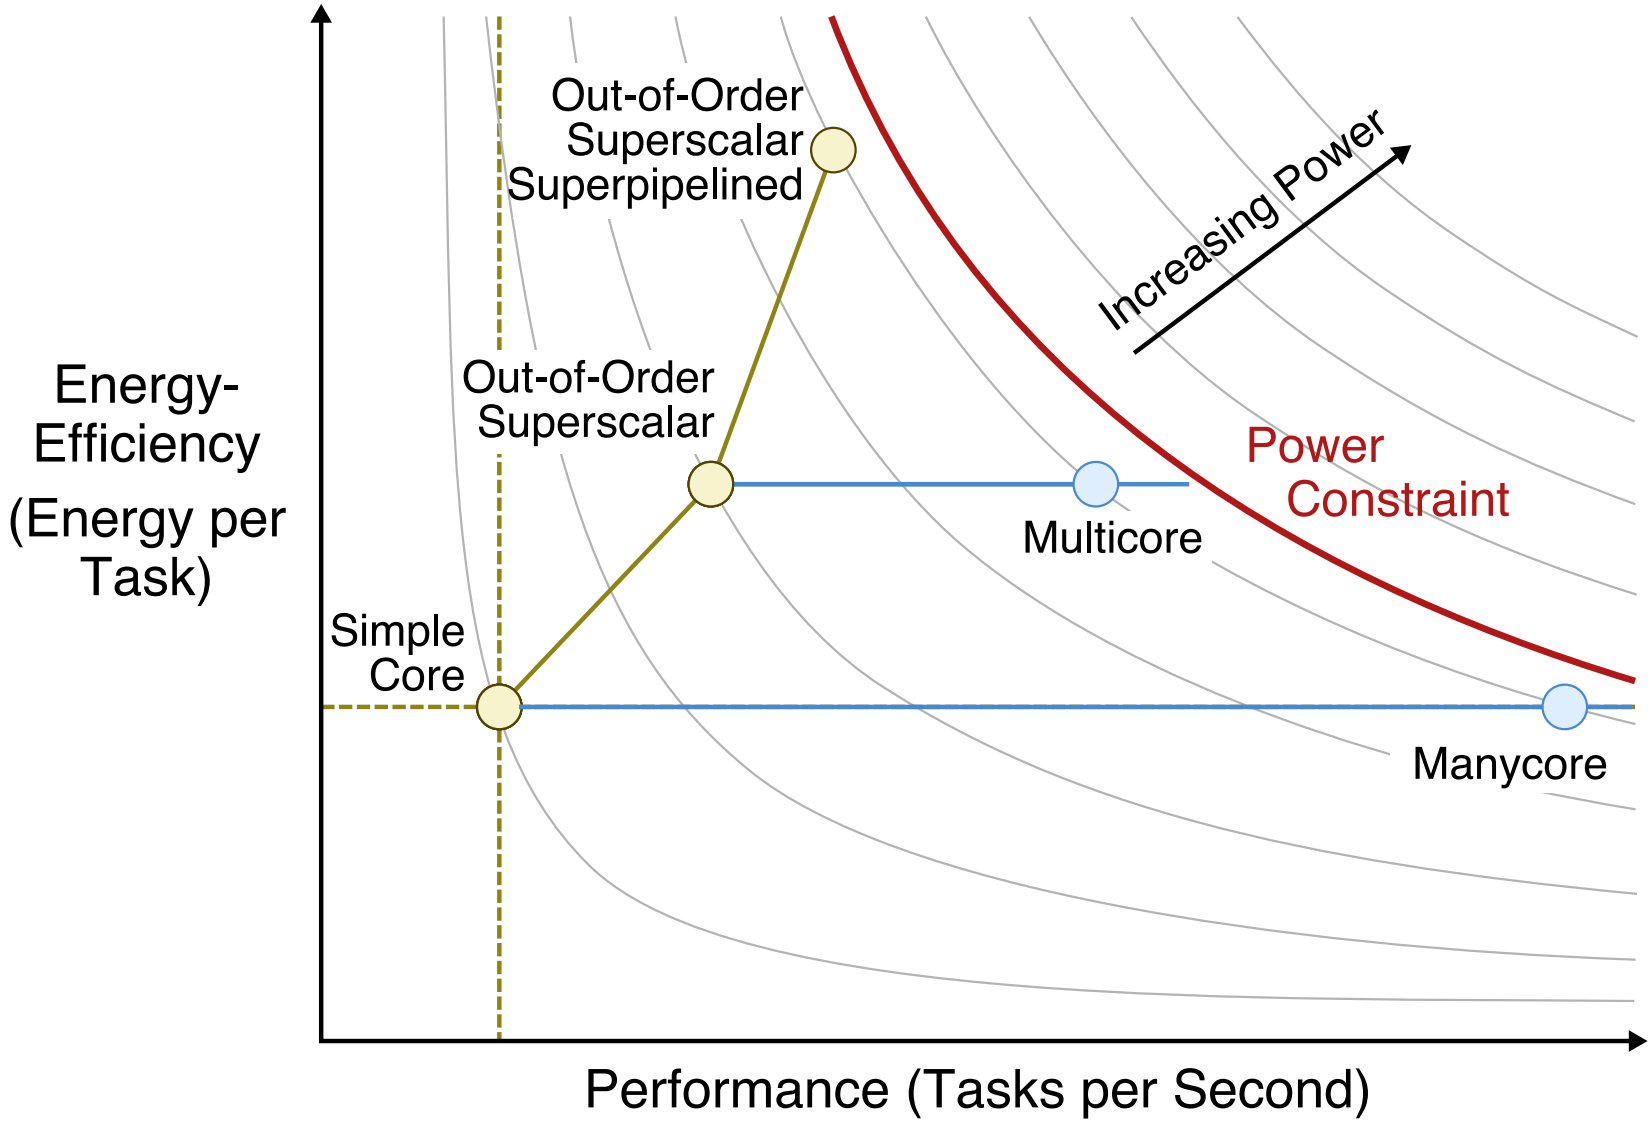

Energy and Performance of Multicore and Manycore

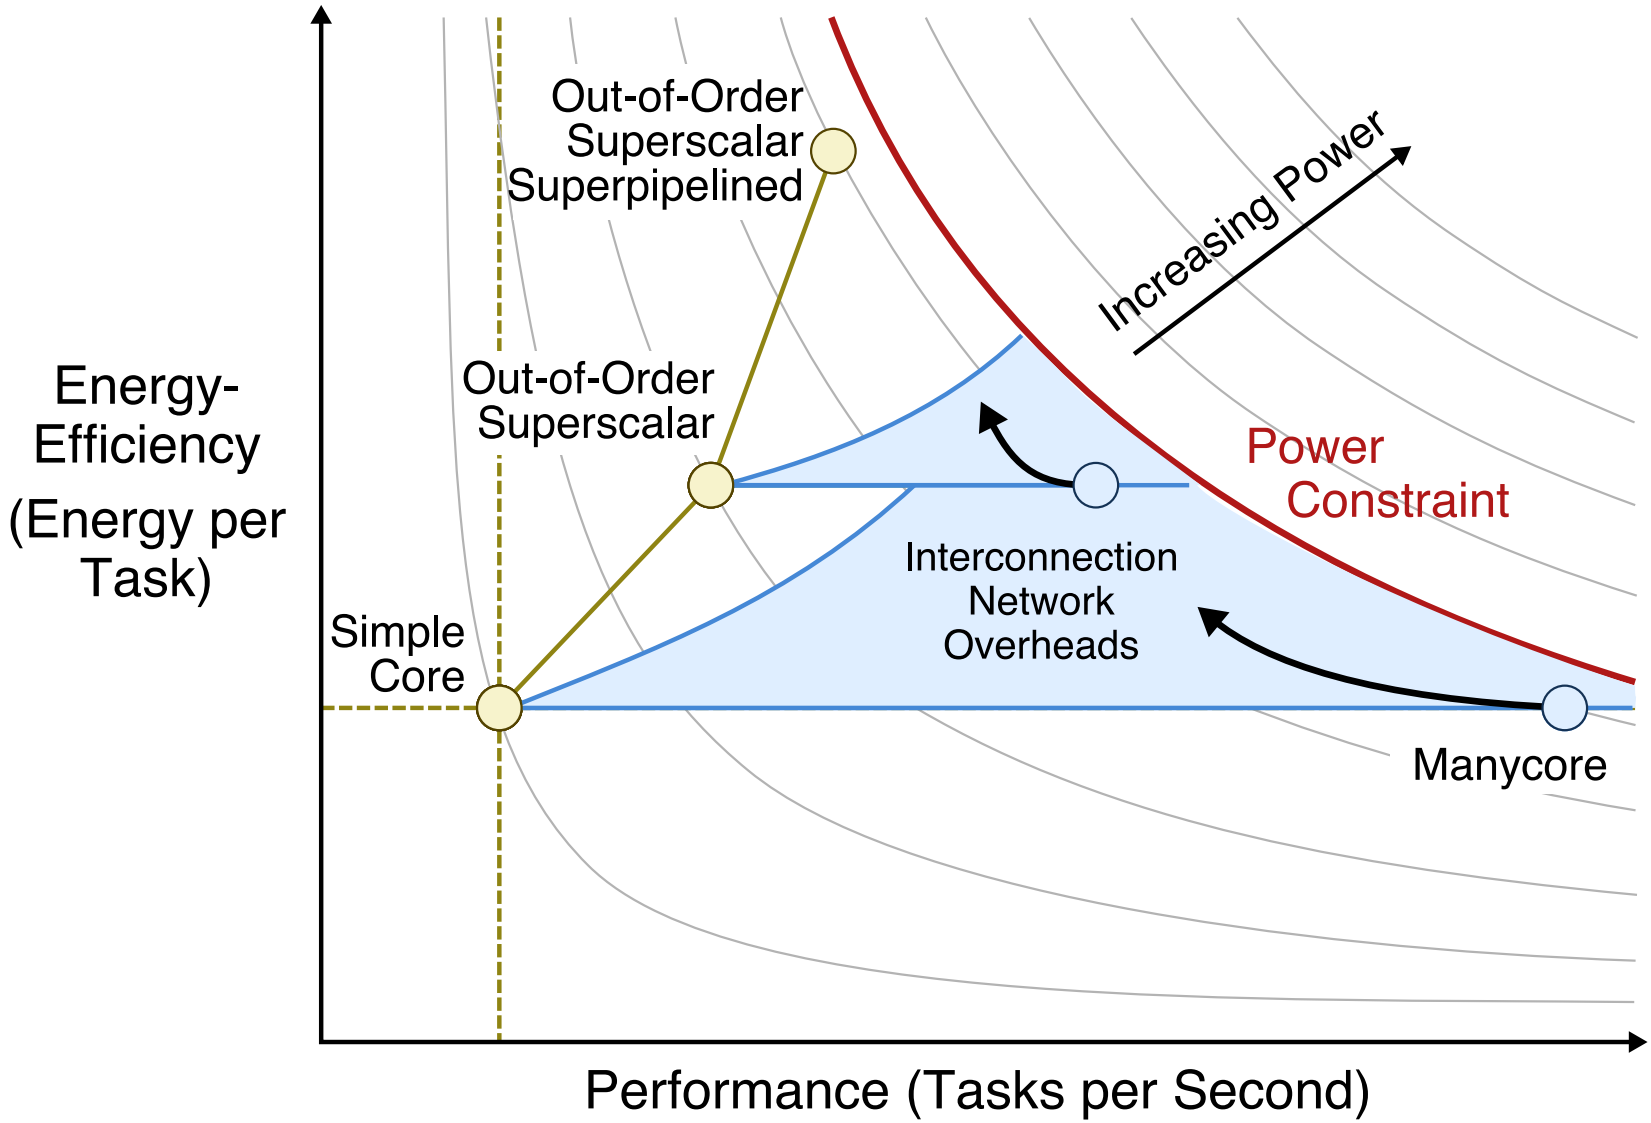

Types of Interconnection Networks

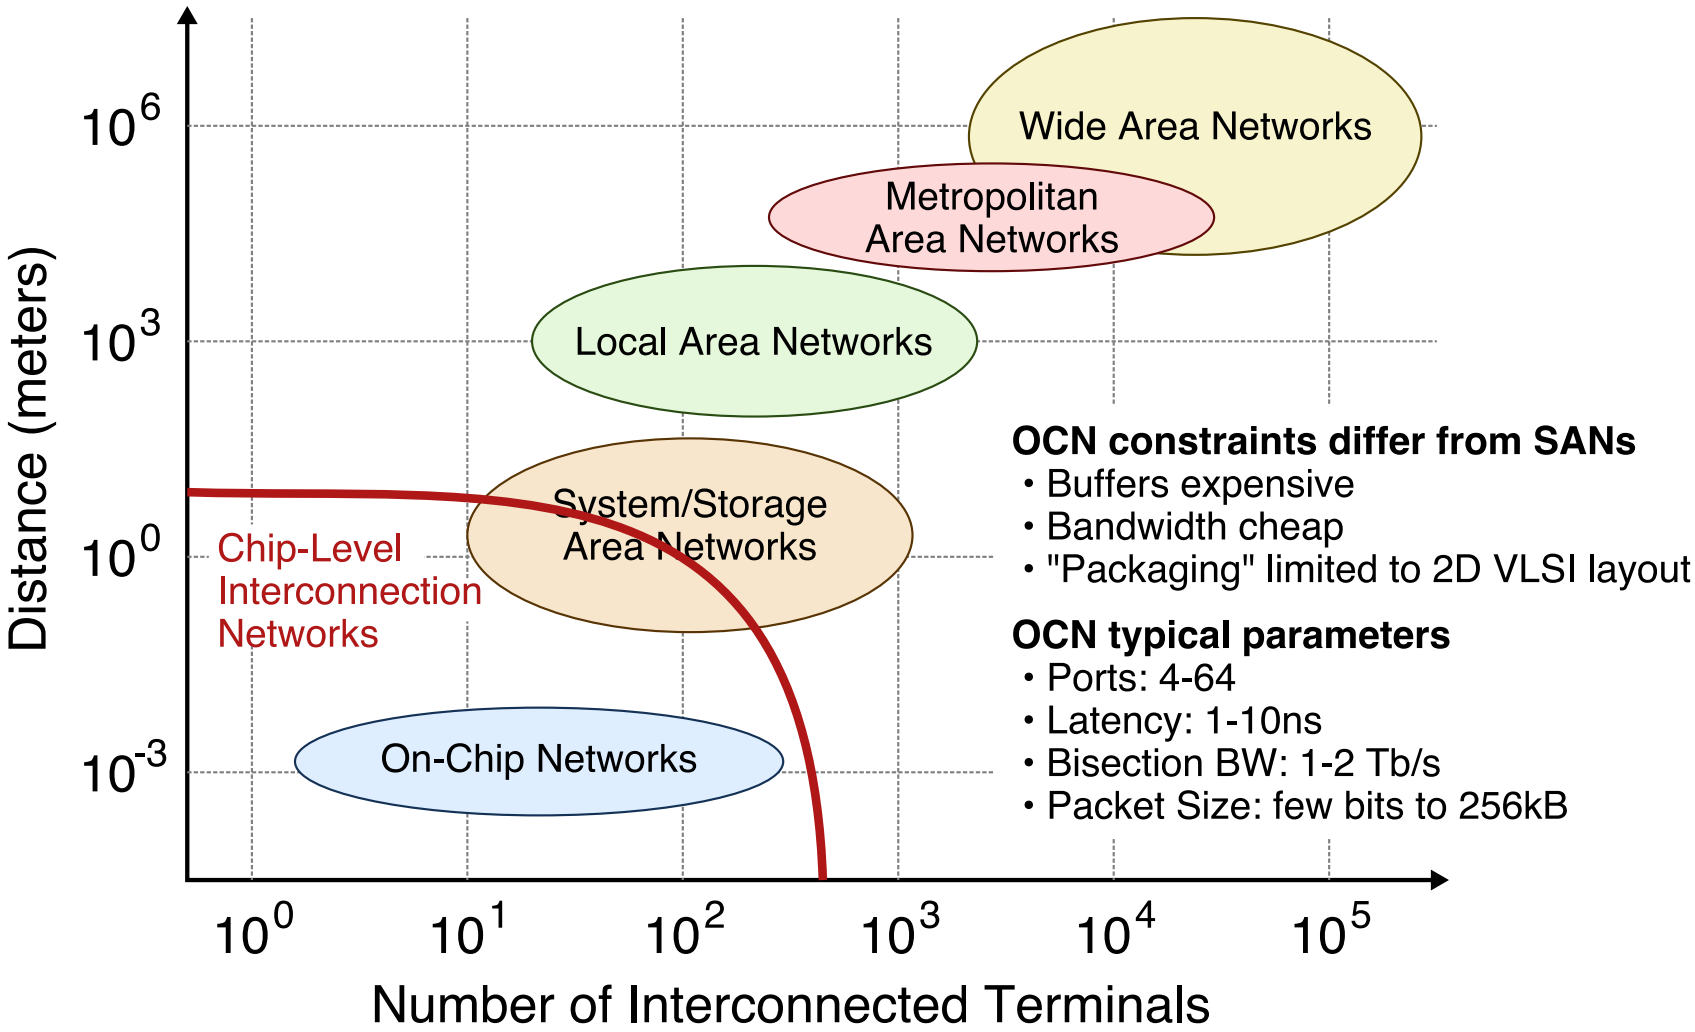

Agenda

Course Motivation

Interconnection Network Basics
- Topology
- Routing
- Flow Control
- Router Microarchitecture
- Examples

Course Logistics

Topology: Arrangement of Nodes and Channels

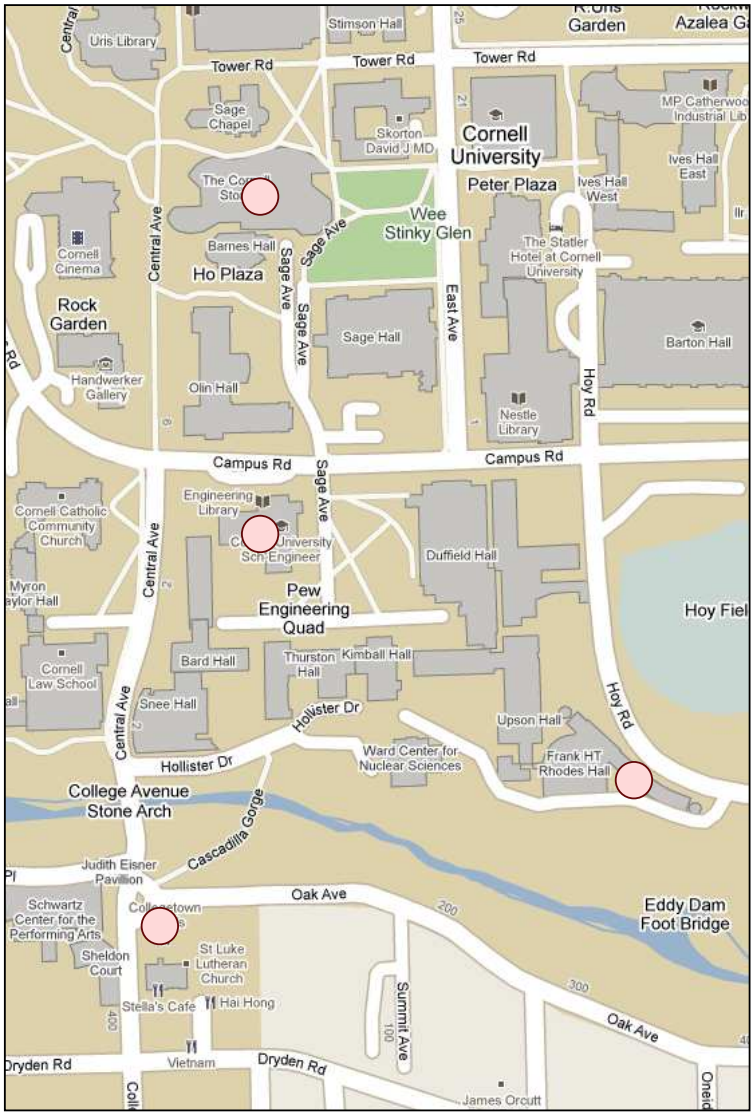

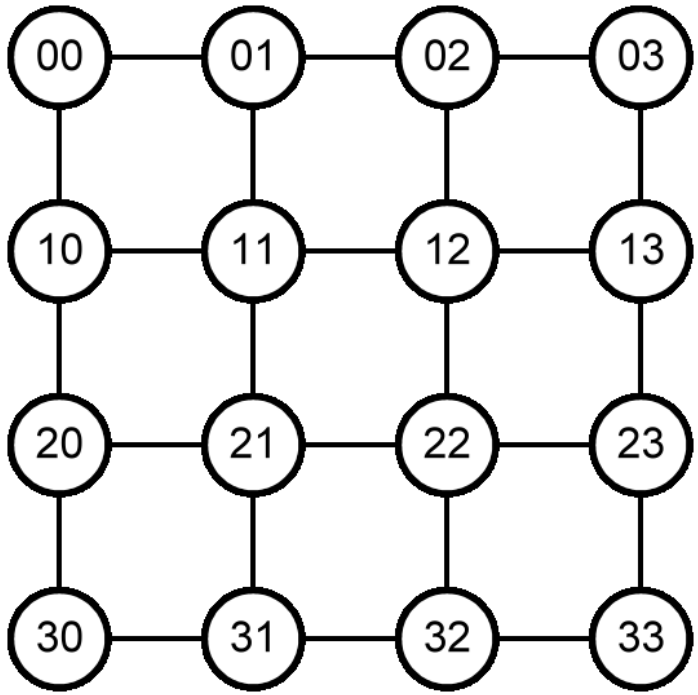

Many Potential Topologies

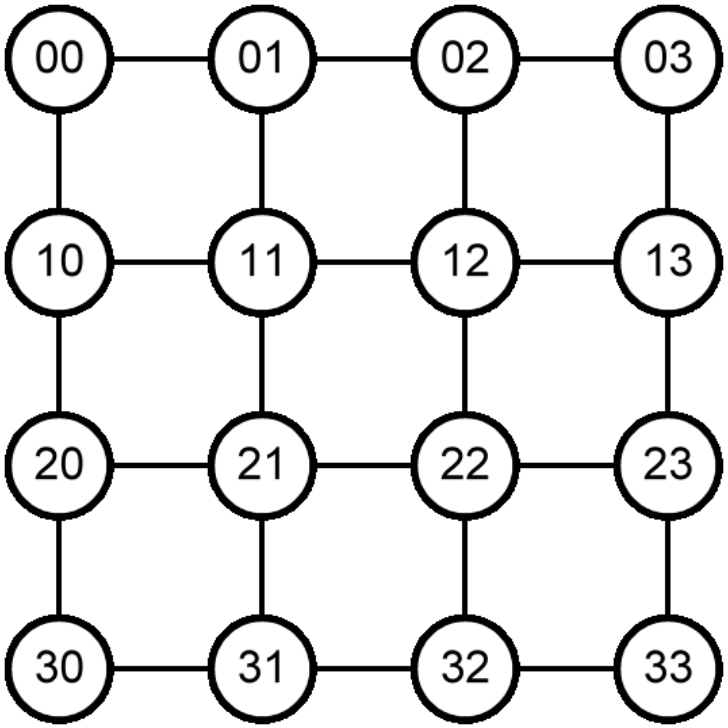

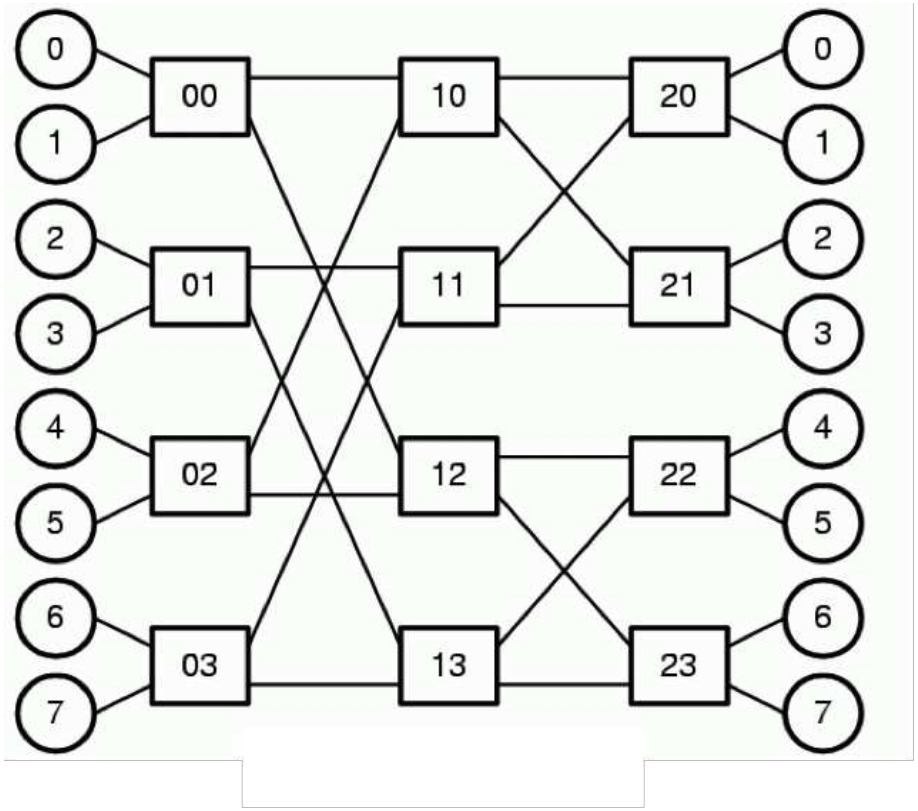

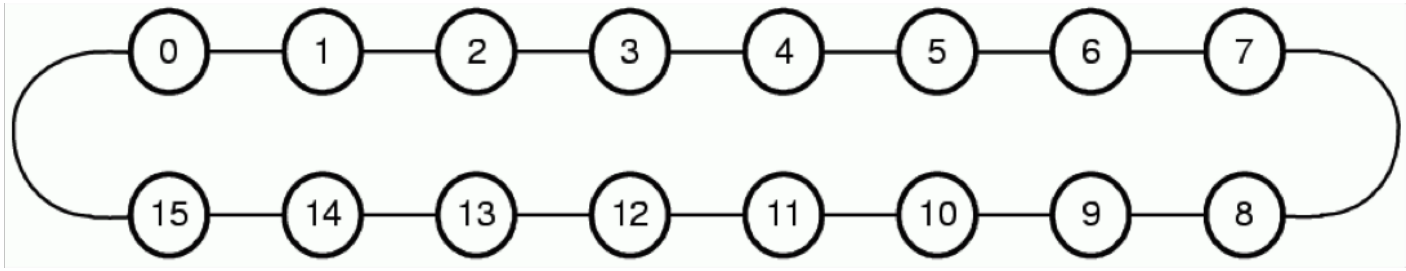

Topology is Constrained By Packaging

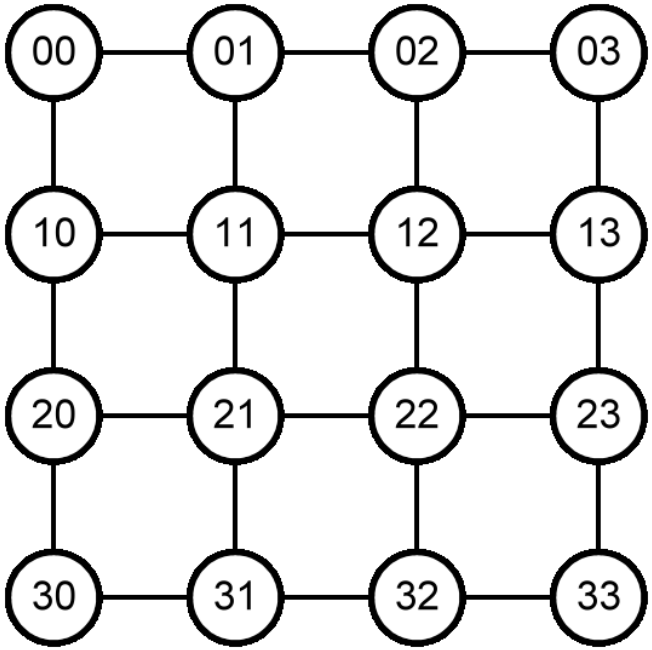

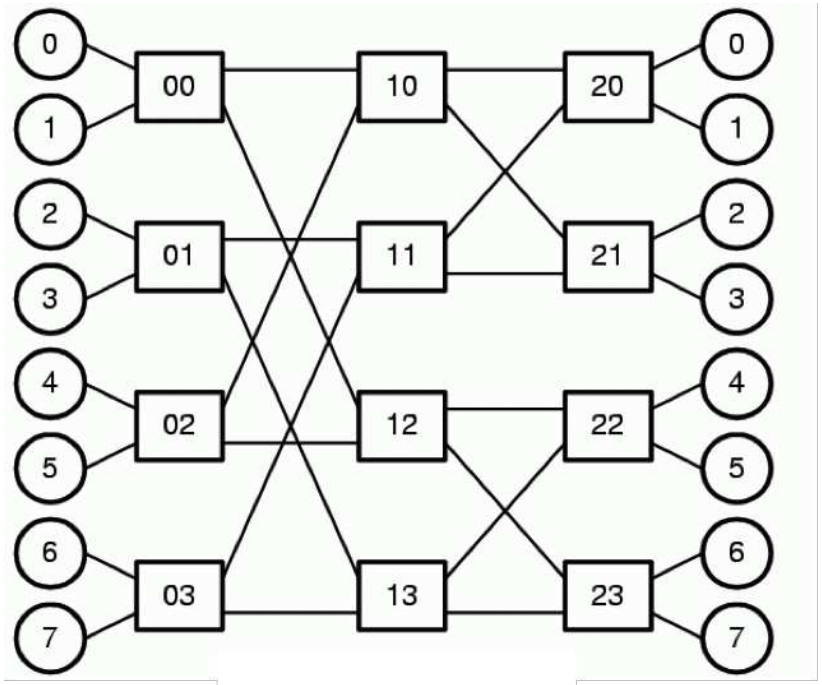

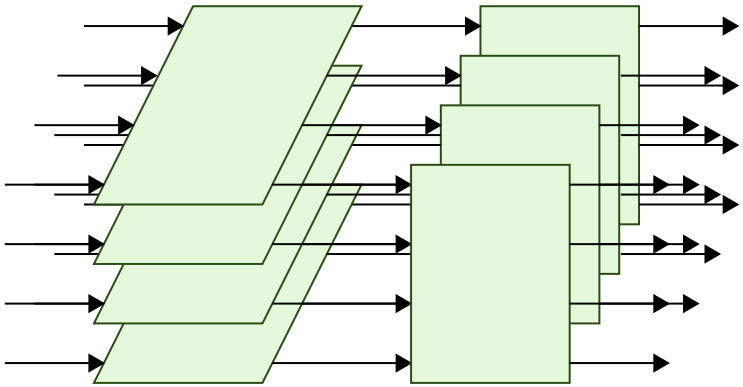

Routing: Determining Path Between Terminals

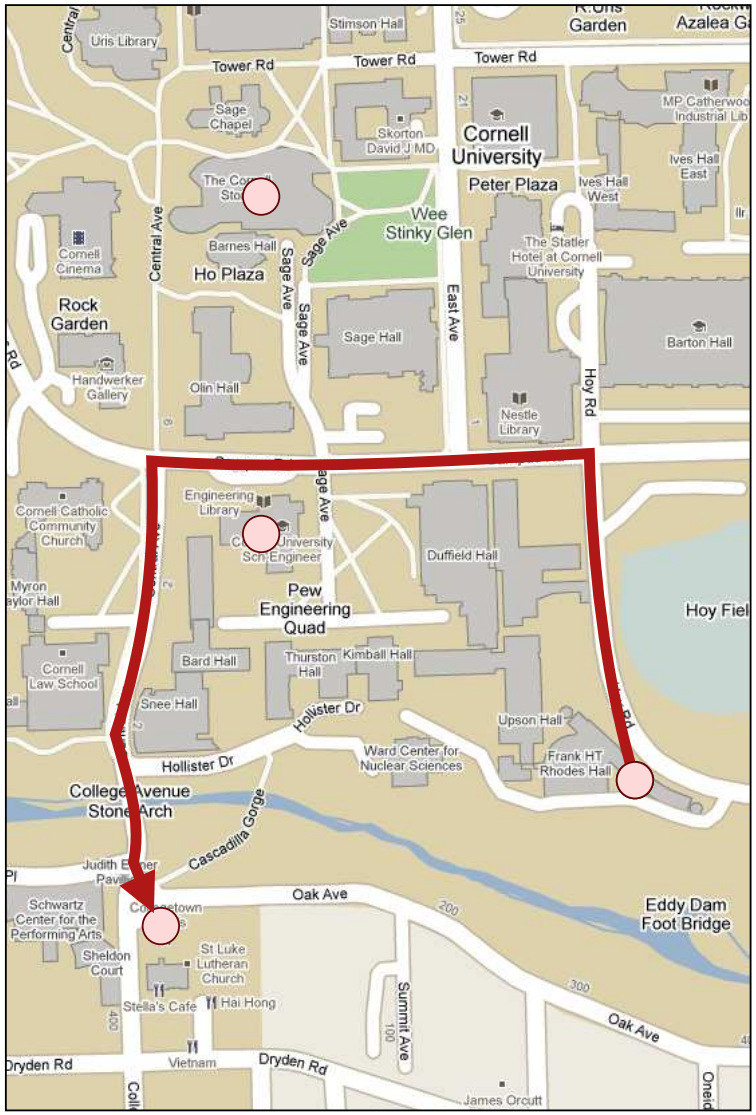

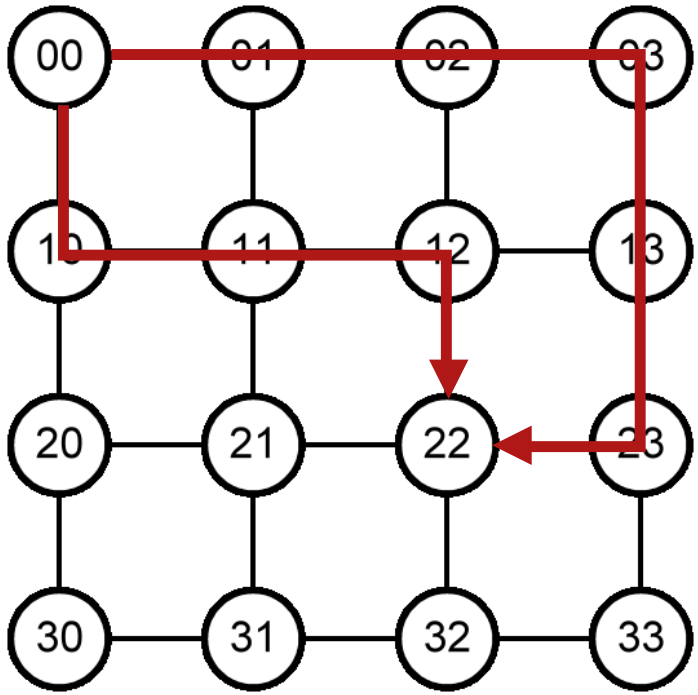

Minmal Routing vs. Non-Minimal Routing
Oblivious vs. Adaptive Routing
Deterministics vs. Randomized Routing

Flow Control: Managing Allocation of Resources

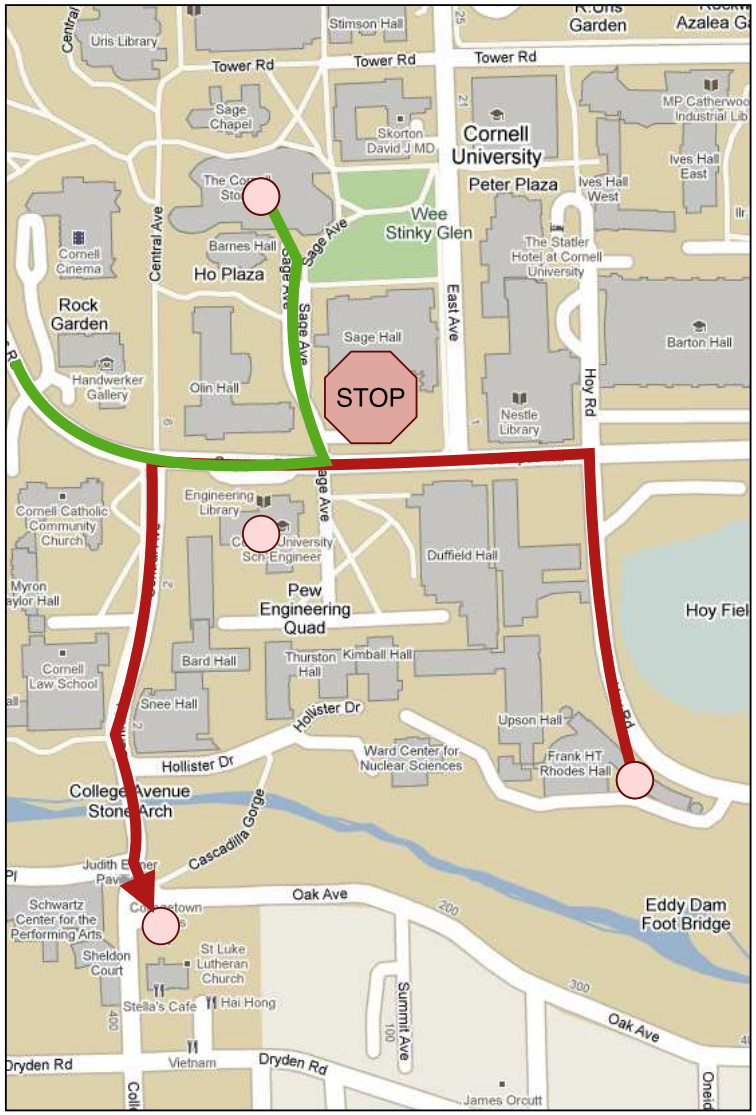

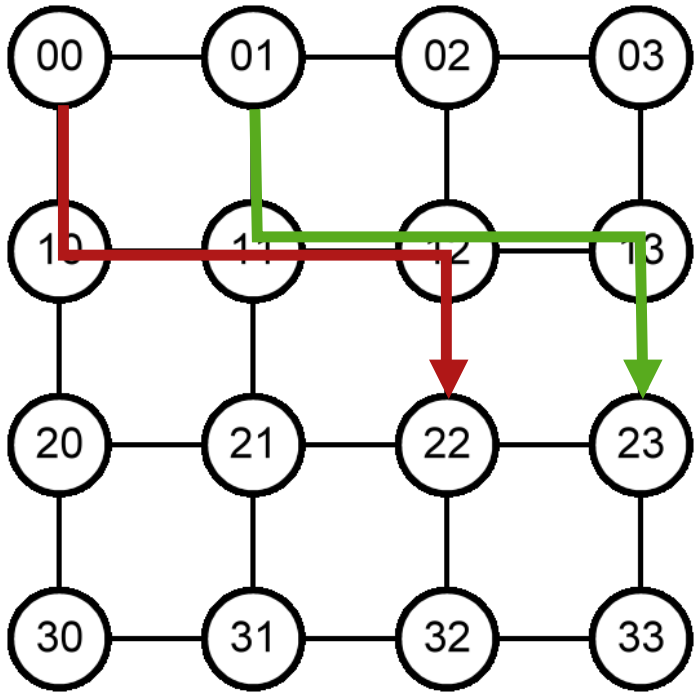

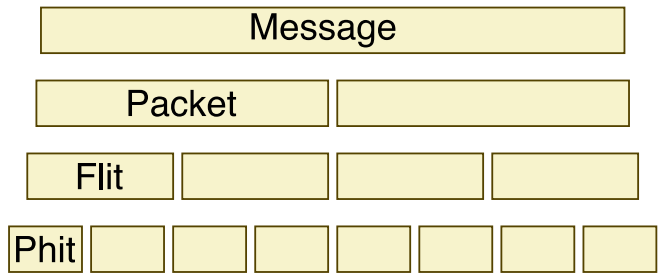

Router Microarchitecture

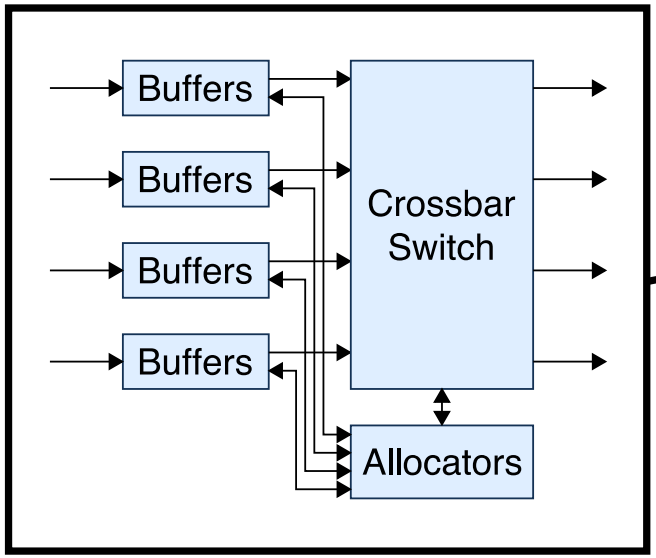

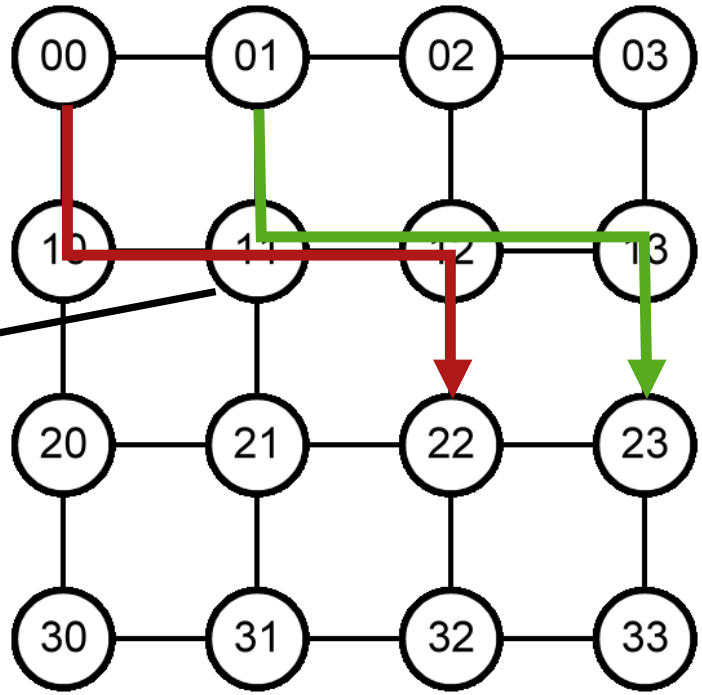

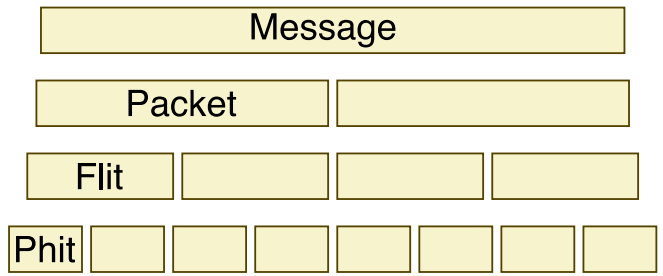

Evaluating A Network Implementation

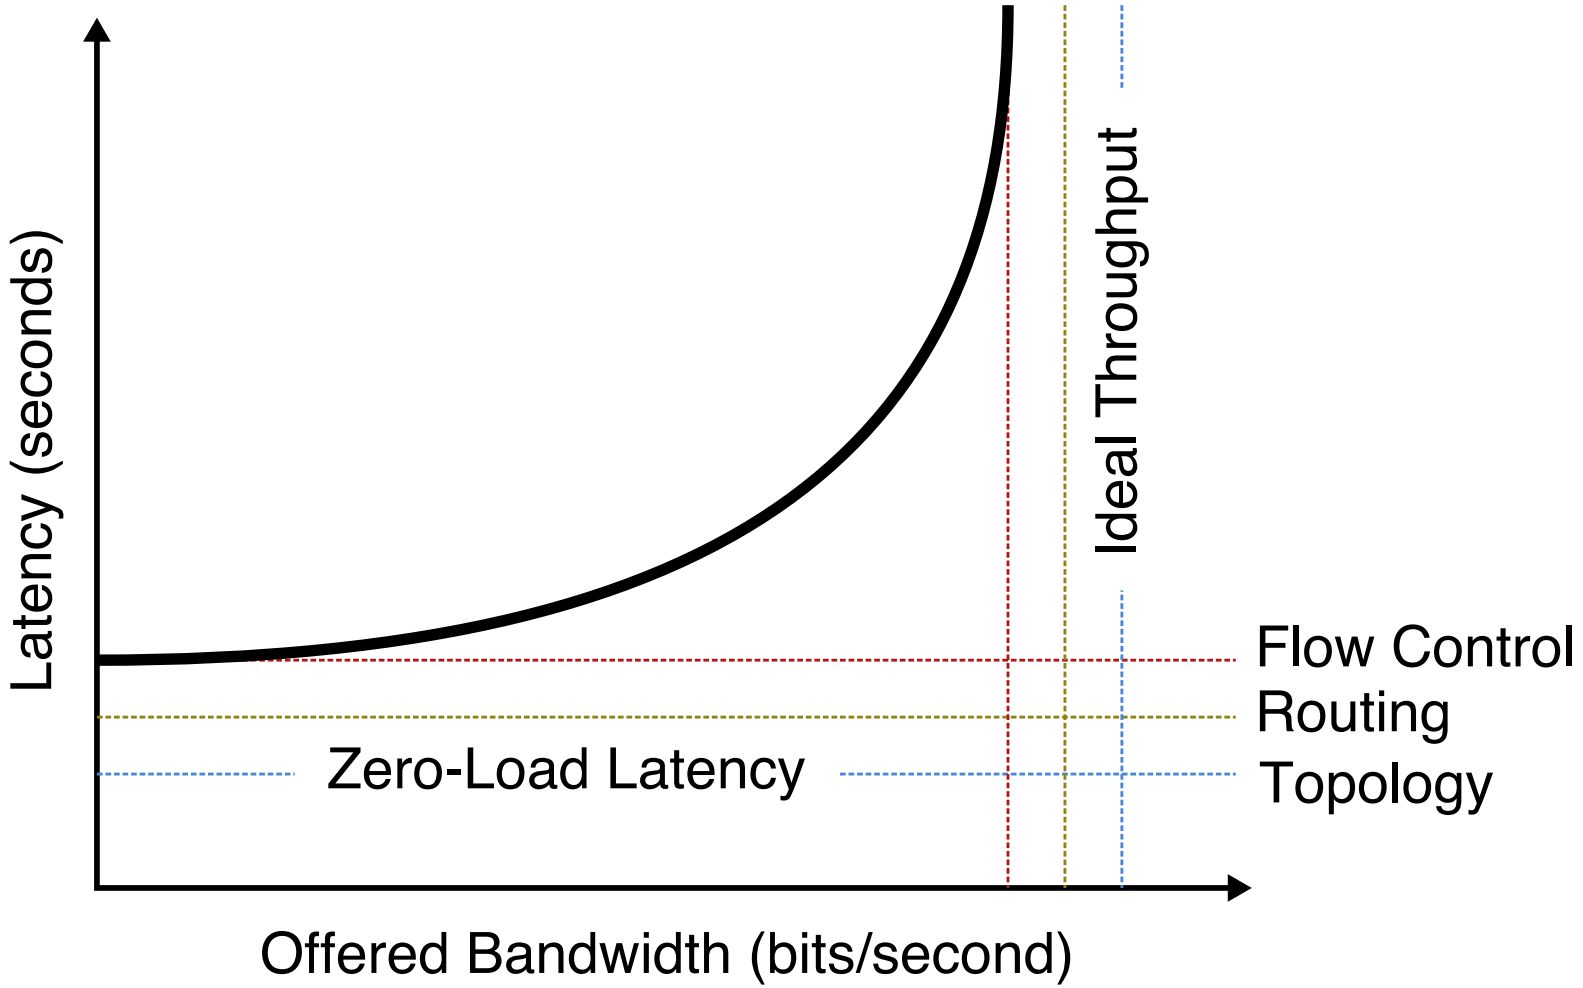

Sun Niagara Processor

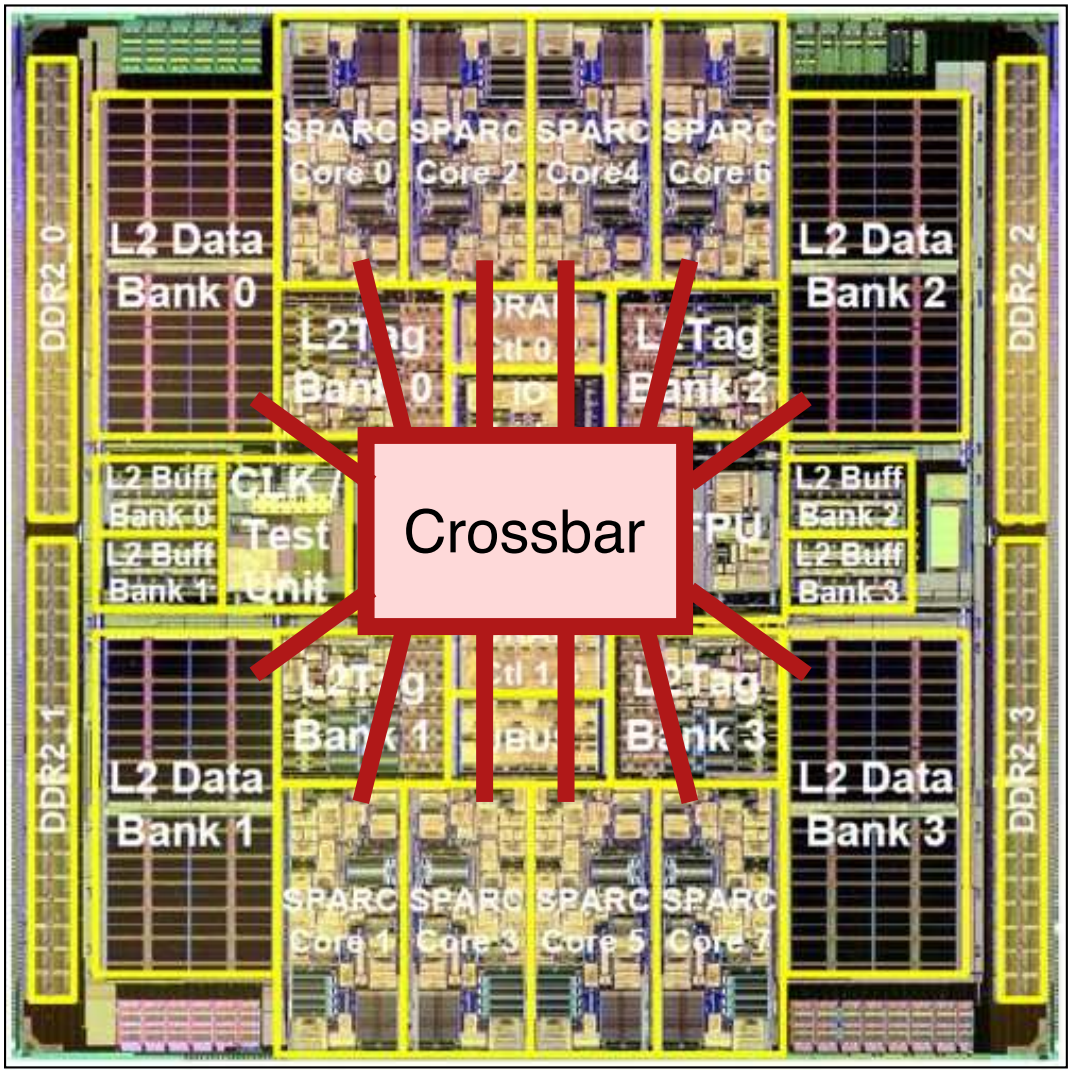

- 8 multithreaded processors
- Single-stage crossbar connecting 8 cores to 4 L2 cache banks
- "200 GB/s" total bisection BW

IBM Cell Processor

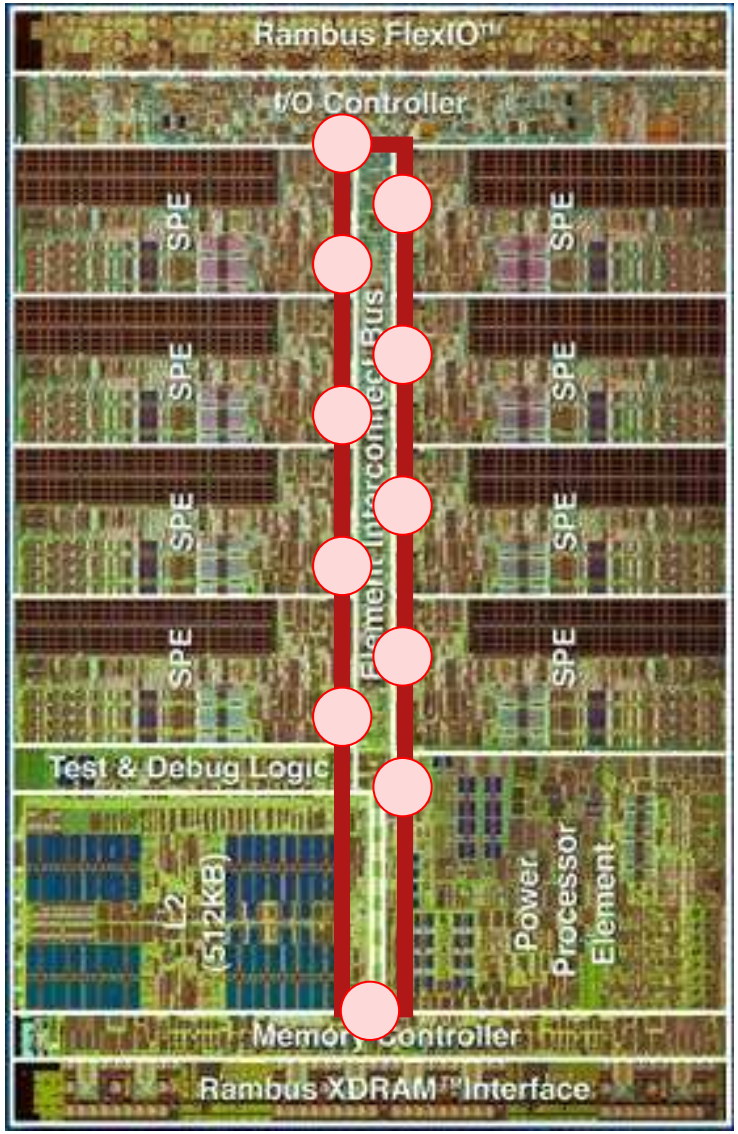

- 1 general-purpose processor
- 8 processors specialized for data-parallelism
- 4 uni-directional rings
- Each ring is 128b wide at 1.6 GHz
- Network Bisection BW = 25.6 GB/s
- Total Bisection = 102.4 GB/s

MIT Raw Processor

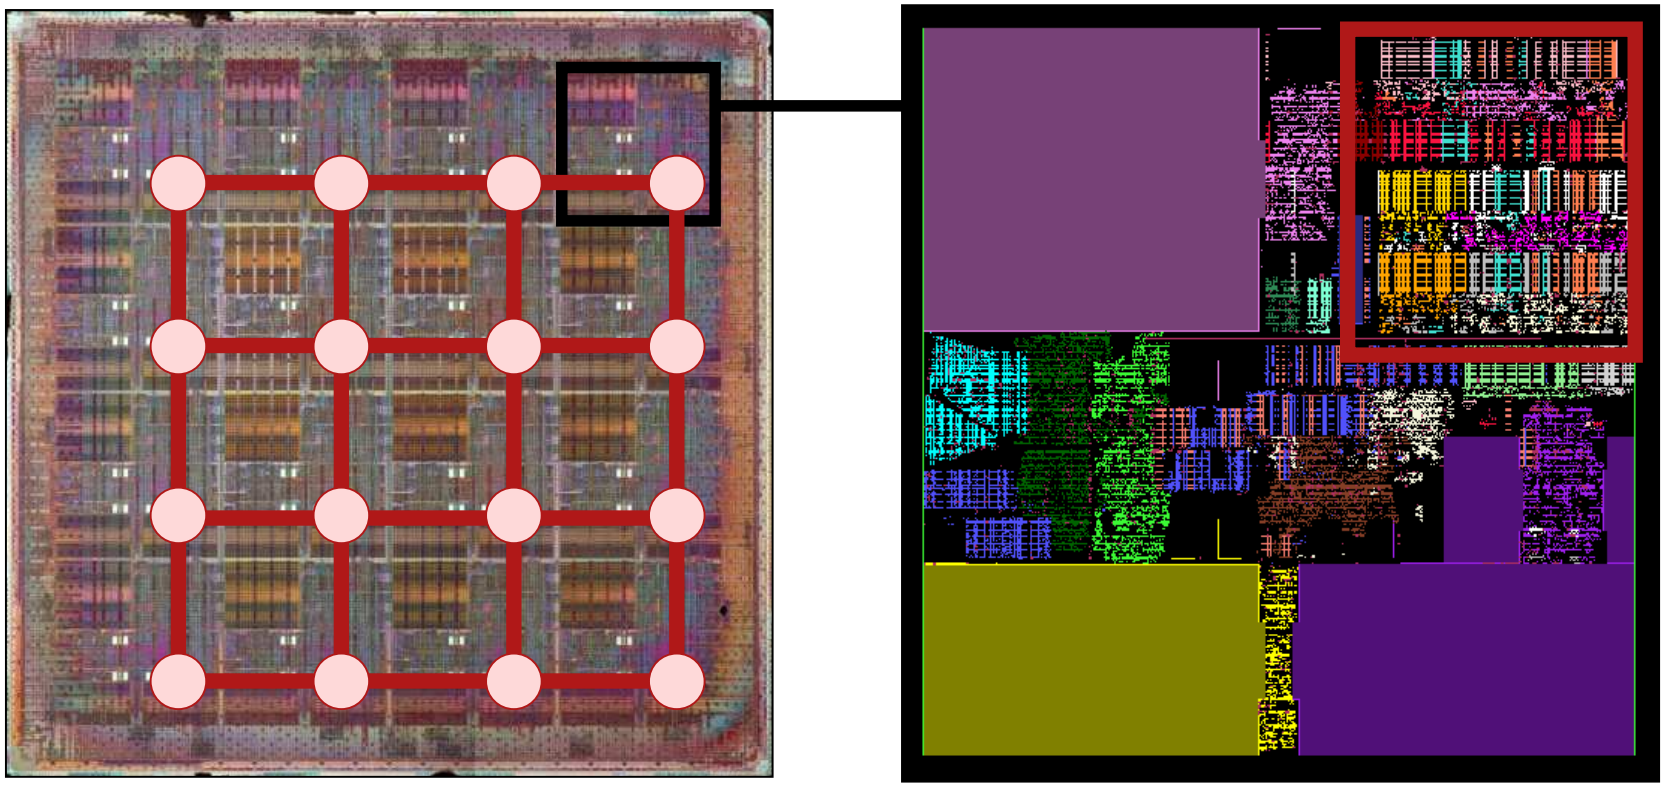

- 16 simple RISC cores
- Two dynamically routed mesh networks (32b/channel)
- Two statically routed mesh networks for message passing (32b/channel)
- Bisection bandwidth per network is 8*32b at 400 MHz 12.8 = 12.8 GB/s
- Total bisection bandwidth is 51.2 GB/s
- Network consumes 20-30% of total chip power

Agenda

Course Motivation

Interconnection Network Basics
- Topology
- Routing
- Flow Control
- Router Microarchitecture
- Examples

Course Logistics

Course Schedule

- Week 1: Overview and Example Network
- Week 2: Topology: Lectures
- Week 3: Topology: Paper Discussions
- Week 4: Routing: Lectures
- Week 5: Routing: Paper Discussions
- Week 6: Flow Control: Lectures
- Week 7: Flow Control: Paper Discussions
- Week 8: Router Microarchitecture: Lectures
- Week 9: Router Microarchitecture: Paper Discussions
- Week 10: Emerging Technologies
- Late March: Project proposal due, Quiz
- April–May: Design project
- May: Final project presentations and report due

We will need to iterate on topics as we go along

Paper Critiques

- Critique should briefly summarize key contribution of paper, provide analysis on the paper's strengths and weaknesses, and propose one new idea or extension
- Everyone is required to submit short one- or two-page critiques for all readings (formal and informal listeners too!)
- Everyone is required to shepherd the discussion for one paper during the semester and prepare a longer four- or five-page critique for that paper synthesizing the other critiques, class discussion, related work, and your own thoughts

Course Project

- Proposed your own new research direction, overlap with current research is encouraged
- Re-evaluate recent paper (from class or not) and propose extension: use paper from class, reference from textbook, or take a look at recent conferences such as ISCA, MICRO, HPCA, NOCS
- Responsible for setting up your own simulation infrastructure although I can help, and hopefully students can help each other: Simics+GEMS, Booksim, SESC?, mtlsim
- Final short presentation and final report in the form of a conference-style eight-page paper
- Should work in a group of two students if at all possible

Miscellaneous

Grading Structure

- 10% class discussion
- 10% short paper critiques for each paper
- 20% long paper critique for one paper
- 20% quiz
- 40% final project presentation and report

Textbook and Other Reference Material

- "Principles and Practices of Interconnection Networks" by Dally et al.
- "Interconnection Networks: An Engineering Approach" by Duato et al.
- "On-Chip Networks" by Peh and Jerger available on course website
- Paper readings available on course website

Expectations of Students

- Advanced graduate-level class
- Those students without strong backgrounds in computer architecture and/or interconnection networks are expected to review the listed textbooks and references as necessary
- Students are expected to spend a reasonable amount of time reading and preparing short critiques (even formal and informal listeners)
- Students are expected to spend a significant amount of time preparing for shepherding and writing the long critique
- Students are expected to work hard to identify a new research direction for final project, and spend a significant amount of time studying the related work, evaluating the idea, and preparing the final presentation and report

Course Goals

Specific Goals for Interconnection Networks

- Study main components of interconnection networks including topology, routing, flow control, and microarchitecture
- Understand key classic and modern research papers in the field
- Explore new research direction in interconnection networks through a detailed design project

Broad Goals for Learning How to Conduct Research

- Go back to basics on a research topic
- Critically review and discuss research papers
- Identify and evaluate new research directions
- Clearly present new research ideas