Interconnection Networks: Introduction

Natalie Enright Jerger

Introduction

- How to connect individual devices into a group of communicating devices?
  - A device can be:
    - Component within a chip
    - Component within a computer
    - Computer
    - System of computers
  - Network consists of:
    - End point devices with interface to network
    - Links
    - Interconnect hardware
- Goal: transfer maximum amount of information with the least cost (minimum time, power)

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

Types of Interconnection Networks (2)

- System/Storage Area Network (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 typically on the order of tens of meters, but some with as high as a few hundred meters
    - InfiniBand: 120 Gbps over a distance of 300 m
  - Examples (standards and proprietary)
    - InfiniBand, Myrinet, Quadrics, Advanced Switching Interconnect

Types of Interconnection Networks (3)

- 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 on the order of few kilometers, but some with distance spans of a few tens of kilometers
  - Example (most popular): Ethernet, with 10 Gbps over 40Km

Types of Interconnection Networks (4)

- Wide Area Networks (WANs)
  - Interconnect systems distributed across the globe
  - Internetworking support is required
  - Many millions of devices interconnected
  - Maximum interconnect distance of many thousands of kilometers
  - Example: ATM

Organization

- Next few lectures will focus on On-chip networks
- Concepts applicable to all types of networks
  - Focus on trade-offs and constraints as applicable to NoCs

On-Chip Networks (NoCs)

- Why Network on Chip?
  - 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

- Off-chip: I/O bottlenecks
  - Pin-limited bandwidth
  - Inherent overheads of off-chip I/O transmission
- On-chip
  - Tight area and power budgets
  - Ultra-low on-chip latencies

MulticoreExamples (1)

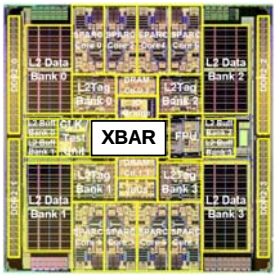

Sun Niagara

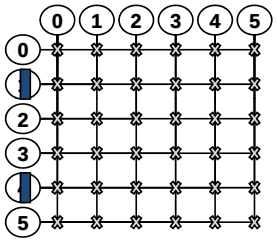

MulticoreExamples (2)

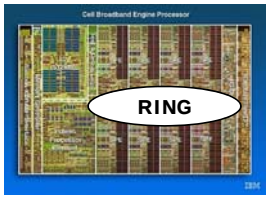

IBM Cell

- Element Interconnect Bus
  - 4 rings
  - Packet size: 16B-128B
  - Credit-based flow control
  - Up to 64 outstanding requests
  - Latency: 1 cycle/hop

Many Core Example

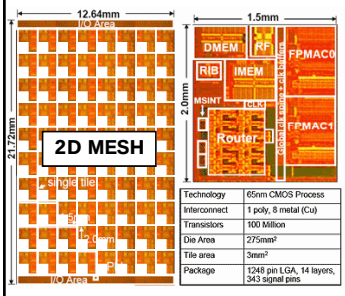

- Intel Polaris
  - 80 core prototype
- Academic Research ex:
  - MIT Raw, TRIPs
    - 2-D Mesh Topology
    - Scalar Operand Networks

Suggested Reading

- William Dally and Brian Towles. Principles and Practices of Interconnection Networks. Morgan Kaufmann Pub., San Francisco, CA, 2003.
- William Dally and Brian Towles, "Route packets not wires: On-chip interconnection networks," in Proceedings of the 38th Annual Design Automation Conference (DAC-38), 2001, pp. 684–689.
- David Wentzlaff, Patrick Griffin, Henry Hoffman, LieweiBao, Bruce Edwards, Carl Ramey, Matthew Mattina, Chi-Chang Miao, John Brown I I I, and AnantAgarwal. On-chip interconnection architecture of the tile processor. IEEE Micro, pages 15–31, 2007.
- Michael Bedford Taylor, Walter Lee, SamanAmarasinghe, and AnantAgarwal. Scalar operand networks: On-chip interconnect for ILP in partitioned architectures. In Proceedings of the International Symposium on High Performance Computer Architecture, February 2003.
- S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, and N. Borkar. An 80-tile 1.28tflops network-on-chip in 65nm cmos. *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International, pages 98–589, Feb. 2007.*