

Interconnection Networks for High-Performance Systems

ECE 6115 / CS 8803 – ICN

Instructor: Tushar Krishna (tushar@ece.gatech.edu)

Course Objectives

Interconnection Networks form the backbone of all computer systems today. They occur at various scales across all high-performance systems - systolic-arrays within Google's Deep Learning TPU, high-bandwidth crossbars inside modern GPUs, soft transport macros on FPGAs, mesh networks-on-chip (NoC) in many-core processors, **interposer** fabrics on package, QPI in multi-socket servers, **Infiniband** in supercomputers/clusters, and Fat-Tree **datacenter** networks in the cloud. The growing emphasis on parallelism, scalability, and energy-efficiency across all these systems makes the design of the communication fabric critical to both high-performance and low power consumption.

This course will examine the similarities, differences, and trade-offs in the architecture and implementation of interconnection fabrics across all these systems. Given the breadth of topic areas (computer architecture, VLSI interconnects, computer networks, and distributed systems), students will get a glimpse into designing systems and optimizing for data movement at various scales – from on-chip to cloud-scale.

There will also be a particular focus on the role of interconnection networks in **Deep Learning Accelerators and Systems** during **Inference** (e.g., Google TPU, NVIDIA DLA, MIT Eyeriss) and **Distributed Training** (e.g., Google Cloud TPU, Facebook Zion, NVIDIA DGX-2)

Course Structure

This is an advanced graduate course, structured around a mix of lectures, student presentations, paper critiques, lab assignments and a final project.

In the first half of the course, lectures will cover the fundamentals of interconnection networks: topology, routing, flow-control, microarchitecture, network and system interfaces. In addition, a series of programming-heavy labs will bring everyone up to speed with an interconnection networks simulator Garnet2.0, that is distributed within the gem5 (www.gem5.org) open source full-system multi-core simulator.

In the second half of the course, students will present and critique research papers on interconnection networks from a breadth of computing domains (computer architecture, circuits, HPC, datacenters). Students will also work on a research project focused on addressing a research challenge in interconnection network design. Students will study relevant papers, propose a solution, implement the solution (via simulation) document the project (short paper) and present the paper in a conference format (20 minutes). Projects aligned with students' own graduate research (MS/PhD) will be encouraged if they have an exciting networks component. Projects from past iterations of the course have led to publications in top-tier conferences such as HPCA, ASPLOS, MICRO, ICCAD, ISPASS and NOCS.

Course Text

The material for this course will be derived from the following texts:

1. N. E. Jerger, T. Krishna, and L.-S Peh, “On-Chip Networks, 2nd Edition” Morgan Claypool Publishers, 2017.
2. W. Dally and B. Towles, “Principles and Practices of Interconnection Networks,” Morgan Kauffman Publishers, 2004.
3. Papers from recent conferences: **ISCA, MICRO, HPCA, ASPLOS, SIGCOMM, NSDI, NOCS, DATE, DAC, ISSCC**

Syllabus and Outline

1. Introduction to Interconnection Networks

- Introduction
- Types of Networks
- Evaluation Metrics

2. Topology

- Metrics for comparing topologies
- Direct Topologies
- Indirect Topologies
- Hierarchical Topologies

3. Routing

- Deterministic Routing
- Oblivious Routing
- Adaptive Routing

4. Flow-Control

- Message-based Flow Control
- Packet-based Flow Control
- Flit-based Flow Control
- Virtual Channels

5. Deadlocks

- Channel Dependency Graph
- Turn Model
- Up*/Down* Routing
- Escape Virtual Channels
- Deadlock Recovery

6. Microarchitecture

- Router Organization
- Pipeline
- Optimizations
- Buffer Management
- Crossbar Design
- Allocators and Arbiters

7. System Interface

- Shared Memory Multiprocessors
 - Cache Coherence
 - Deadlocks
- Message Passing

8. Implementation: RTL and Circuits

- Wire Delay
- Router Pipelines
- Power Consumption
- Area Overheads

9. Emerging Technologies

- Silicon Photonics
- Reliability and Faults

10. Interconnection Networks across High-Performance Systems

- Many-core Processors
- GPUs
- FPGA
- Deep Learning Platforms
- On-Package Interconnects
- 2.5D/3D Systems
- Supercomputers
- Datacenters
- IoT Systems

Course Grading

Lab Assignments	33%
Midterm	10%
Paper Critiques	10%
Presentation on Paper/Case Study	10%
Final Project	32% (Report 22%, Presentations 15%)

Absence and Re-Examination Policy

In case students miss a deadline or an exam, the course will abide by the institute policy on student absences (<http://www.catalog.gatech.edu/rules/4/>)

Learning Accommodations

If needed, the course will make classroom accommodations for students with disabilities. These accommodations should be arranged in advance and in accordance with the office of Disability Services (<http://www.adapts.gatech.edu>)

Academic Integrity

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. For information on Georgia Tech's Academic Honor Code, please visit <http://www.catalog.gatech.edu/policies/honor-code/> or <http://www.catalog.gatech.edu/rules/18/>

Any student suspected of cheating or plagiarizing on a quiz, exam, or assignment will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations.