High Performance Computing Systems (CMSC714)

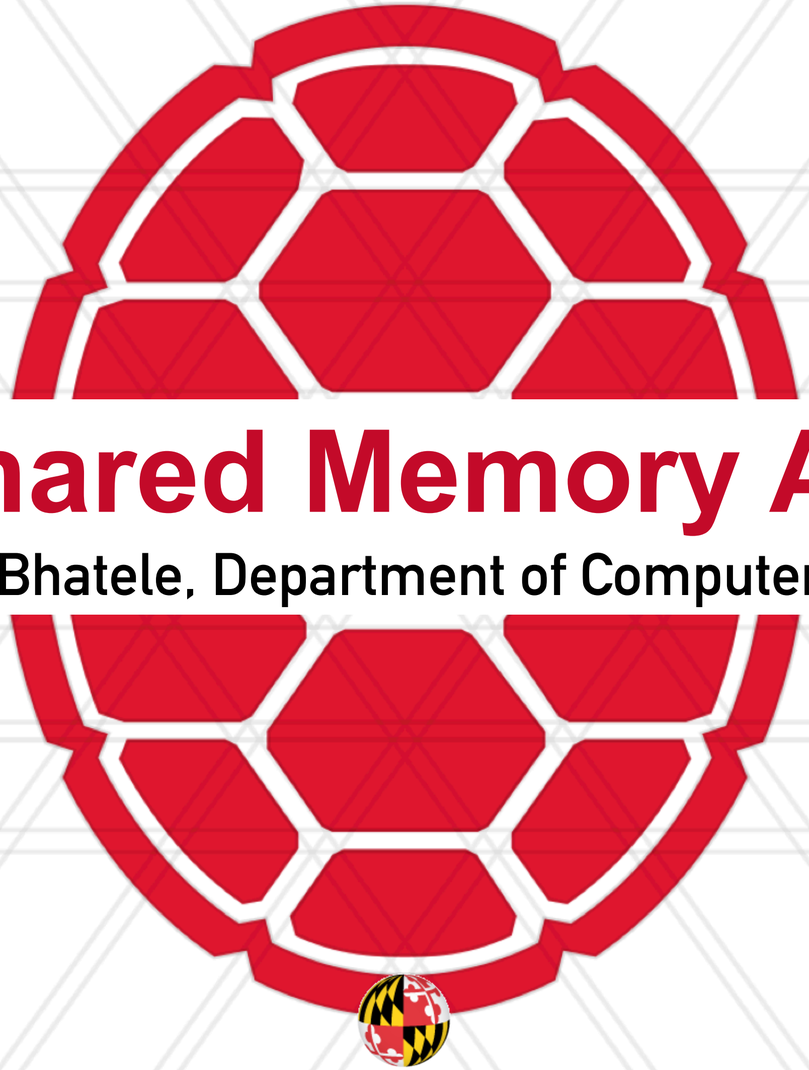

Lecture 8: Shared Memory Architectures

Abhinav Bhatele, Department of Computer Science

UNIVERSITY OF MARYLAND

Announcements

- Assignment 1 is due on Monday (Sept 23) @ 5:00 pm
- Assignment 2 on OpenMP will be published online on Sept 23
  - Due on October 7

Summary of last lecture

- Single node architecture is fairly complex
  - Two product lines: fast processors, low frequency low power processors
- IBM Blue Gene/Q Compute Chip
- Accelerators: IBM Cell BE, AMD APUs, NVIDIA GPGPUs, Intel XE

Shared memory in hardware

- Cache coherent globally addressable memory
- Older machines had bus-based symmetric multi-processing
- Origin was a different architecture: distributed shared memory with cache coherence

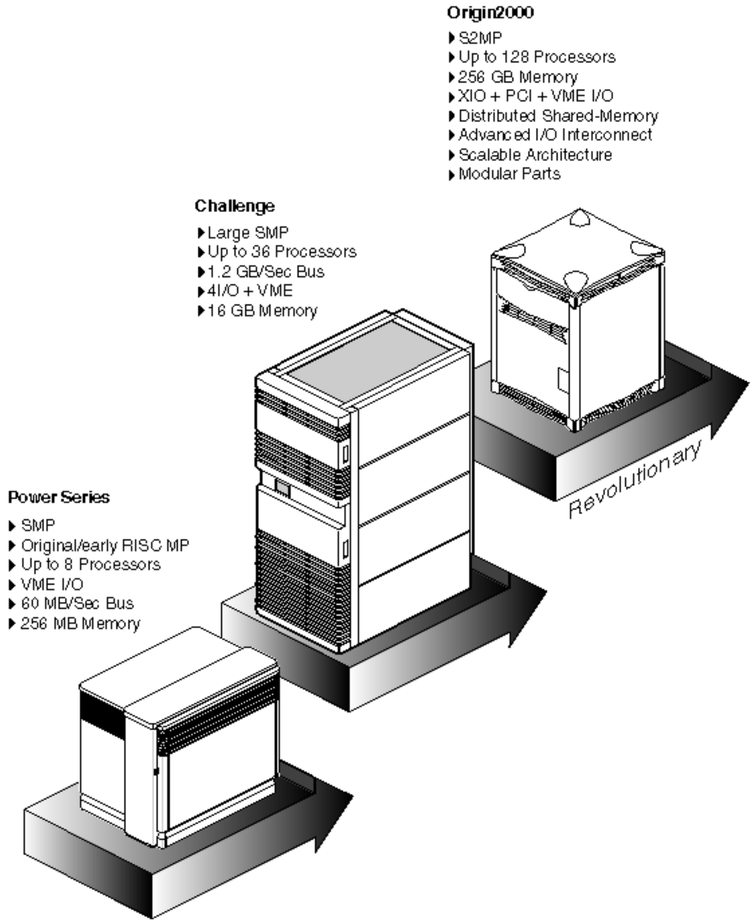

http://csweb.cs.wfu.edu/~torgerse/Kokua/SGI/007-3439-002/sgi_html/ch01.html

SGI Origin 2000

- Up to 512 nodes: 2 processors per node, 4 GB of memory
- Cache coherence maintained via a directory-based protocol
- Distributed directory that keeps track of each data block (page)
  - Implemented in hardware
  - Supports moving entire pages across nodes

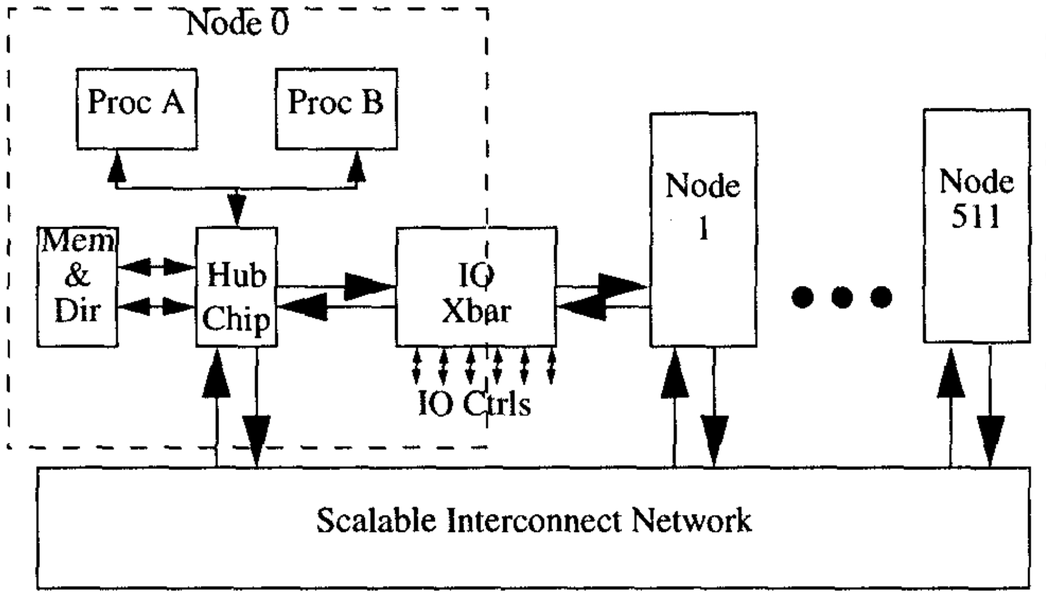

Hypercube network

- General topology: *k*-ary *n*-cube networks
- Hypercube: k=2

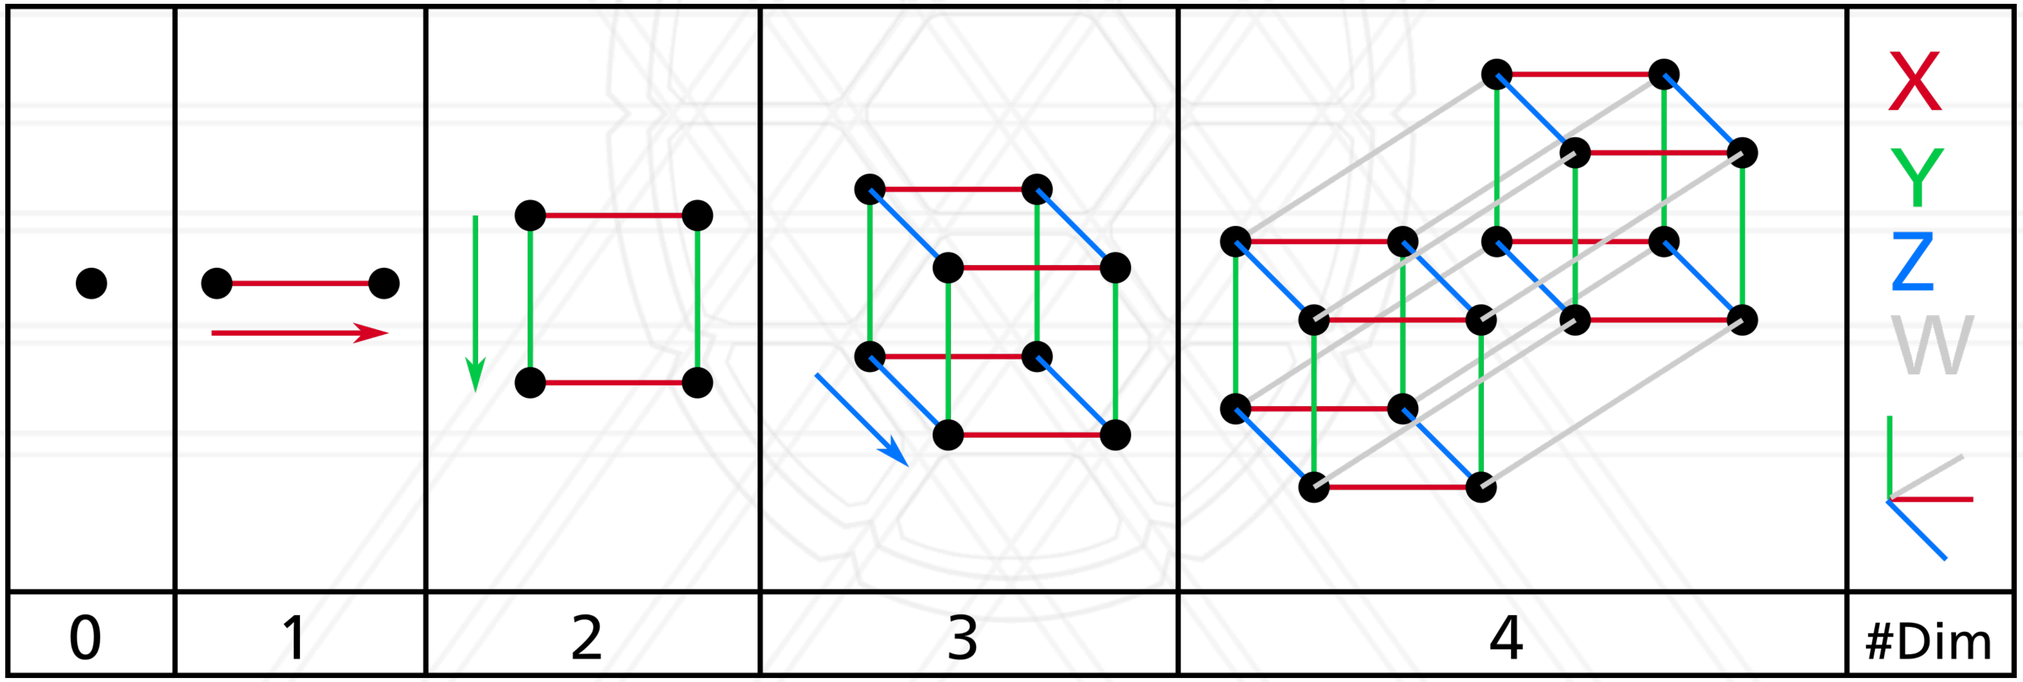

SGI Altix 3000

- Based on Intel Itanium 2 processors and Linux
- 4 processors and up to 32 GB of memory

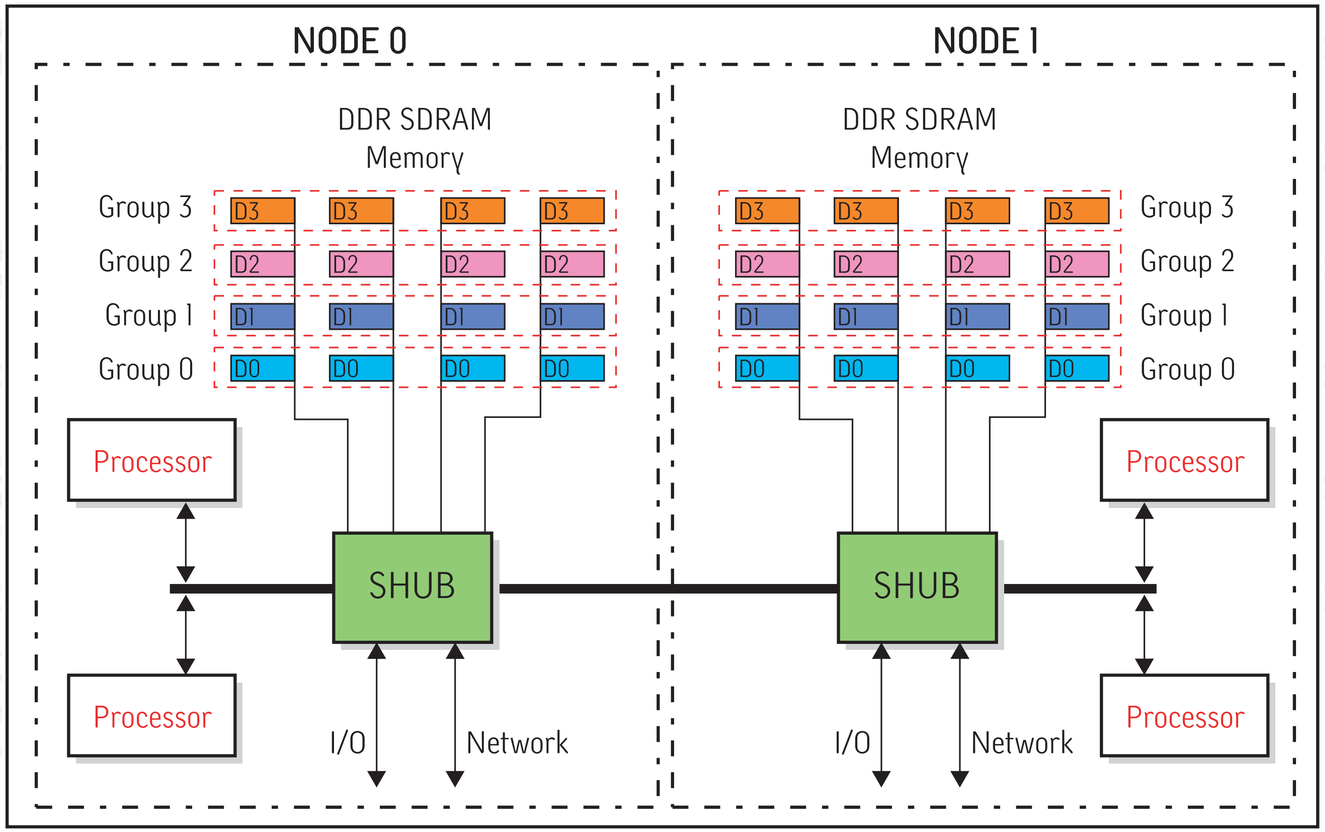

Fat-tree network

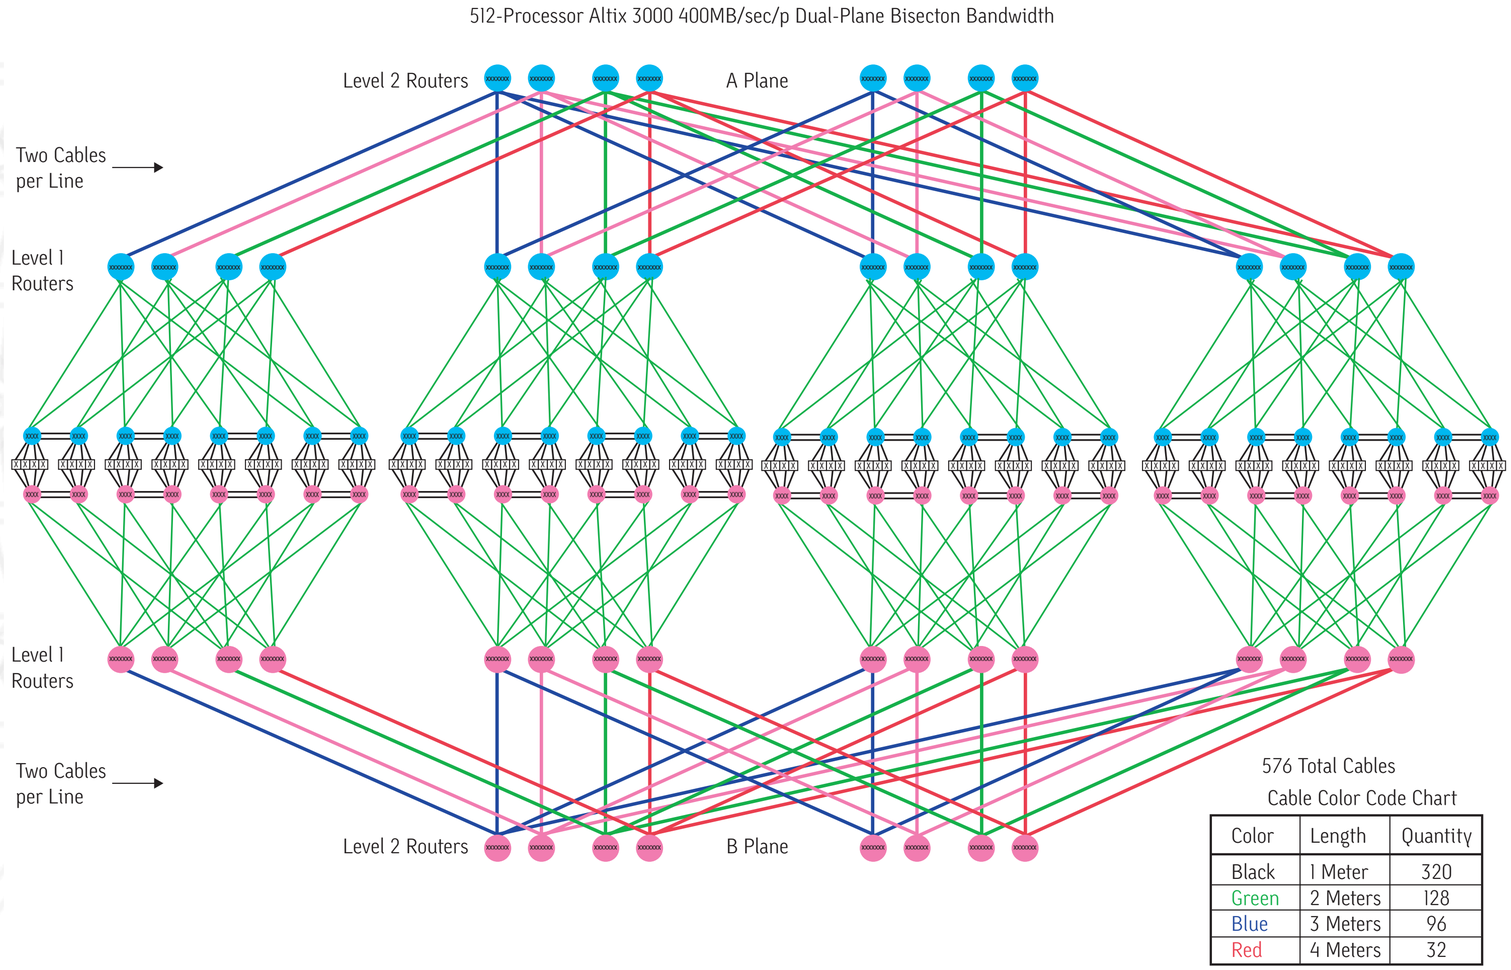

Color	Length	Quantity
Black	1 Meter	320
Green	2 Meters	128
Blue	3 Meters	96
Red	4 Meters	32

Questions

The SGI Origin: A ccNUMA Highly Scalable Server

- Can we use OpenMP on NUMA machine like SGI Origin? Is it hard to implement? Any disadvantage or advantage?
- Does the SGI Origin have any disadvantages?
- The Origin's cache coherence protocol is non-blocking. Would that not imply that it could be possible a request for memory would be made and the processor requesting may not receive access to the latest version of that memory?
- How does a directory-based cache compare in performance to a non-directory-based cache?
- Is directory poisoning commonly used in NUMA systems? Where is the "state" of a directory stored?

Questions

The SGI Altix 3000 Global Shared-Memory Architecture

- What are the limitations of this system?
- Altix 3000 works well with weather and climate codes, why is that? How are its libraries?
- What's so super about a SHUB?
- What is DIMM and when is it used?
- What is "chip-kill" error detection?

Questions?

UNIVERSITY OF
MARYLAND

Abhinav Bhatele

5218 Brendan Iribe Center (IRB) / College Park, MD 20742

phone: 301.405.4507 / e-mail: bhatele@cs.umd.edu