Announcements

- Reading for this class
  - finish Chapter 10
  - http://www.sgi.com/Products/hardware/servers/technology/overview.html

KSR-1

- COMA - Cache Only Memory Architecture
- second level cache replaces main memory
  – called the “all cache” design
  – cache line size is 128 bytes
- Interconnect is a “ring of rings”
  – first level ring
    • 32 processors
    • ARD - Allcache Routing and Directory cell
  – second level ring
    • 1,2,4 GB/sec
    • moves cache lines between rings
  – ARD
    • full directory of all lines on its ring
    • forwards requests from nodes in its ring
    • forwards requests for lines stored on its ring

KSR-1 (cont.)

- Processor
  - custom designed 64 bit microprocessor
    - a big mistake - never could keep pace with technology
    - needed 64 bits since memory is large
  - 20 MIPS per processor, superscalar design
  - FPU supports chained multiply/add instructions
- OS
  - OSF-1
  - uses shared memory so footprint is not as large as Paragon
  - extensive used of thread based programming model

Stanford Dash

- Structure
  - collection of bus based multi-processors
  - interconnect network and cache controller connect nodes
- Cache System
  - snoopy protocol within in a single SMP node
  - directory based cache controller between nodes
    - misses on local cluster go to home cluster of memory "owner"
    - owner may have current copy or could be cached on another cluster
- Processors
  - 4 MIPS R3000 (33 Mhz) per node
- Interconnect
  - 2 dimensional mesh

Stanford Dash (cont.)

- Performance
  - level 0 cache (1 clock)
  - remote clutser load (132 clocks)
- New Directions
  - FLASH
  - use a full micro-processor for the cache controller
    - permits customization of cache protocols
    - makes the hardware simpler

SGI Origin Servers

- Commercialization of Stanford DASH
  - SMP nodes
  - directory based cache controller
- Changes
  - processors are R10000
  - only 2 nodes per bus
    - slightly cheaper bus than DASH
    - faster processors require more bus bandwidth
  - interconnection network
    - hypercube (to 32 nodes)
    - re-configurale routers beyond

SGI Origin Structure

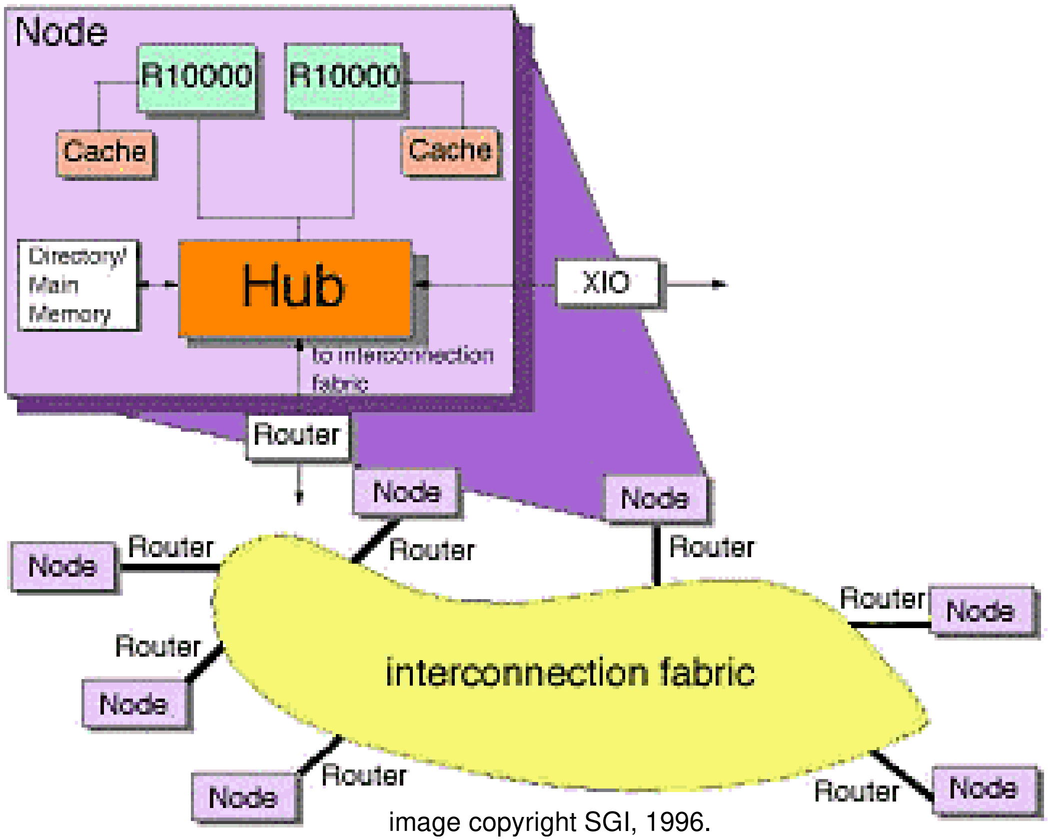

image copyright SGI, 1996.

Cray T3D

- Globally accesable memory
  - each node has is own private memory
  - can read/write the memory of other nodes
    - no cache coherency is provided
- Processors
  - Alpha: 64 bit processors
  - added extra prefetch hardware
    - just off processor circuit
    - permits requests for memory in advance of need
- Interconnect
  - 3 dimension torus
  - support scatter/gather operations with “address centrifuge”
- OS
  - slave micro-kernel on nodes

Tera

- Uses threading to hide latency
  - decendant of HEP
  - rapid hardware context switching
    - if a cache miss occurs, switch threads
  - similar do dataflow
    - ready values drive execution
    - however, but programming model is qvthreaded von Neuman
- Interconnect is 3-D mesh
  - processors, memory, and I/O devices are all on the mesh
- Processors
  - fully custom (using GaS technology)
  - superscalar
  - instructions include a count of the number insns that follow that do not depend on the current instruction