

nCUBE Overview

- Distributed-memory MIMD, hypercube interconnect
- History
 - nCUBE 1 — 1985
 - nCUBE 2 — 1989
 - 34 GFLOPS, scalable
 - ?–8192 processors
 - nCUBE 3 — date?
 - 1–6.5 TFLOPS, 65 TB memory, 24 TB/s hypercube interconnect, 1024 3 GB/s I/O channels, scalable
 - 8–65,536 processors
- Operation
 - Can be partitioned into “subcubes”
 - Programming paradigms: SPMD, cooperative processing between subcubes, client/server

1

Fall 1999, Lecture 35

nCUBE 3 Processor

- 0.6 μ m, 3-layer CMOS, 2.7 million transistors, 50 MHz, 16 KB data cache, 16 KB instruction cache, 100 MFLOPS
 - Argument against off-the-shelf processor: shared memory, vector floating-point units, aggressive caches are necessary in workstation market but superfluous here
- ALU, FPU, virtual memory management unit, caches, SDRAM controller, 18-port message router, and 16 DMA channels
 - ALU for integer operations, FPU for floating point operations, both 64 bit
 - Most integer operations execute in one 20ns clock cycle
 - FPU can complete two single- or double-precision operations in one clock cycle
 - Virtual memory pages can be marked as “non-resident”, the system will generate messages automatically for that data

2

Fall 1999, Lecture 35

nCUBE 3 Interconnect

- Hypercube interconnect
 - Added hypercube dimension allows for double the processors, but processors can be added in increments of 8
 - Wormhole routing + adaptive routing around blocked or faulty nodes
- ParaChannel I/O array
 - Separate network of nCUBE processors for load distribution and I/O sharing
 - 8 computational nodes (nCUBE processors plus local memory) connect directly to one ParaChannel node, and can also communicate with those nodes via the regular hypercube network
 - ParaChannel nodes can connect to RAID mass storage, SCSI disks, etc.
 - One I/O array can be connected to more than 400 disks

3

Fall 1999, Lecture 35

nCUBE 3 Software

- Parallel Software Environment
 - nCX microkernel OS — runs on all compute nodes and I/O nodes
 - UNIX functionality
 - Programming languages including FORTRAN 90, C, C++, as well as HPF, Parallel Prolog, and Data Parallel C

MediaCUBE Overview

- Emphasized on their web page; for delivery of interactive video to client devices over a network (from LAN-based training to video-on-demand to homes)
 - MediaCUBE 30 = 270 1.5 Mbps data streams, 750 hours of content
 - MediaCUBE 3000 = 20,000 & 55,000

4

Fall 1999, Lecture 35

Silicon Graphics POWER CHALLENGEarray Overview

- Hierarchical distributed-memory MIMD
- Up to 144 MIPS R8000 processors or up to 288 MISP R1000 processors
 - Up to 109 GFLOPS, 128 GB memory, 4 GB/s disk transfer, and 28 TB of disk
- POWER CHALLENGEarray consists of up to 8 POWER CHALLENGE or POWER Onyx (POWERnode) systems
 - POWERnodes shared-memory multiprocessors of up to 18 MIPS R8000 processors or 36 MIPS R1000 processors, with up to 16 GB of memory and up to 4 320 MB/s I/O channels
 - Programs that fit within a POWERnode can use the shared-memory single-address-space model
 - Larger program can span POWERnodes

Silicon Graphics POWER CHALLENGEarray Programming

- Fine- to medium-grained parallelism
 - Shared-memory techniques within a POWERnode, using parallelizing FORTRAN and C compilers
- Medium- to coarse-grained parallelism
 - Shared-memory within a POWERnode or message-passing between POWERnode
 - Applications based on message-passing will run within a POWERnode, and libraries such as MPI or PVM will use the shared-memory instead
- Large applications
 - Hierarchical programming, using a combination of the two techniques