



Massively Parallel Architectures




MPP Specifics

- No shared memory
- Scales to hundreds or thousands of processors
- Homogeneous sub-components
- Advanced Custom Interconnects




MPP Architectures

- There are numerous approaches to interconnecting CPUs in MPP architectures:
 - Rings
 - Grids
 - Full Interconnect
 - Trees
 - Dancehalls
 - Hypercubes




Rings




Worst case distance
 $n-1$ (one ring)


$\frac{n}{2}$ (bi-directional ring)

Cost n





Chordal Ring 3






Chordal Ring 4







Barrel Shifter



Worst case distance
 $n/2$




Grid/Torus/Illiac Torus




#worst case distance

$$\sqrt{n}$$

Cost n




Fully interconnected




worst case distance
1

Cost $n(n-1)/2$




Trees




worst case $2\log n$

Cost $n \log n$




Fat Trees




worst case $2\log n$

Cost n




Dancehalls/Butterflys




distance

$\log n$

Cost n




Hypercubes



worst case distance


d

Cost d

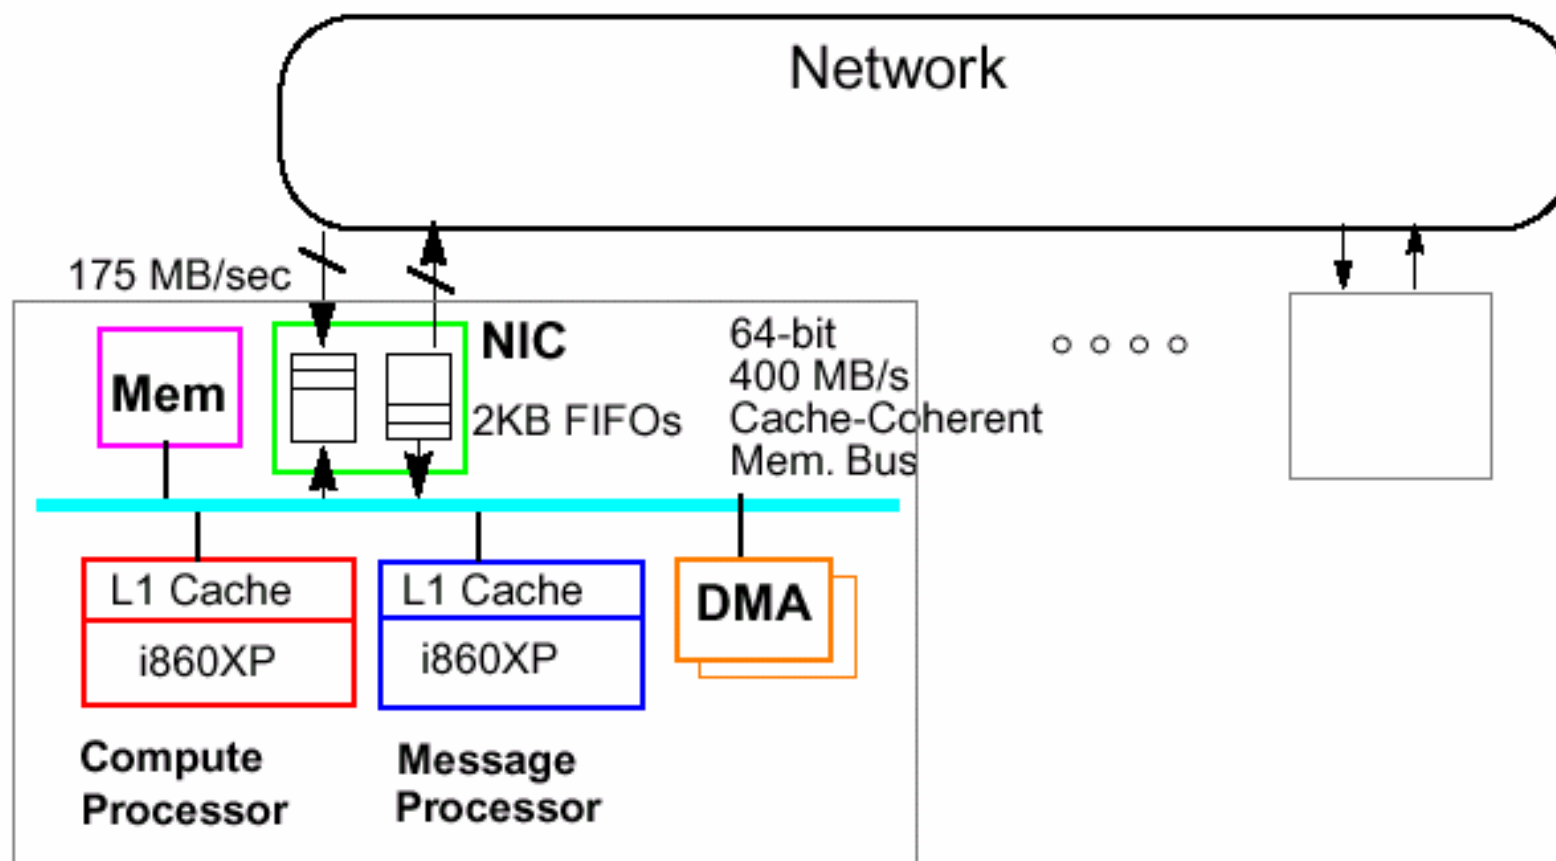



Intel Paragon

- Intel i860 based machine
- “Dual CPU” –
 - 50 MHz CPUs
 - Shares a 400 MB/sec cache coherent bus
- Grid architecture
- Mother of ASCI Red




Intel Paragon



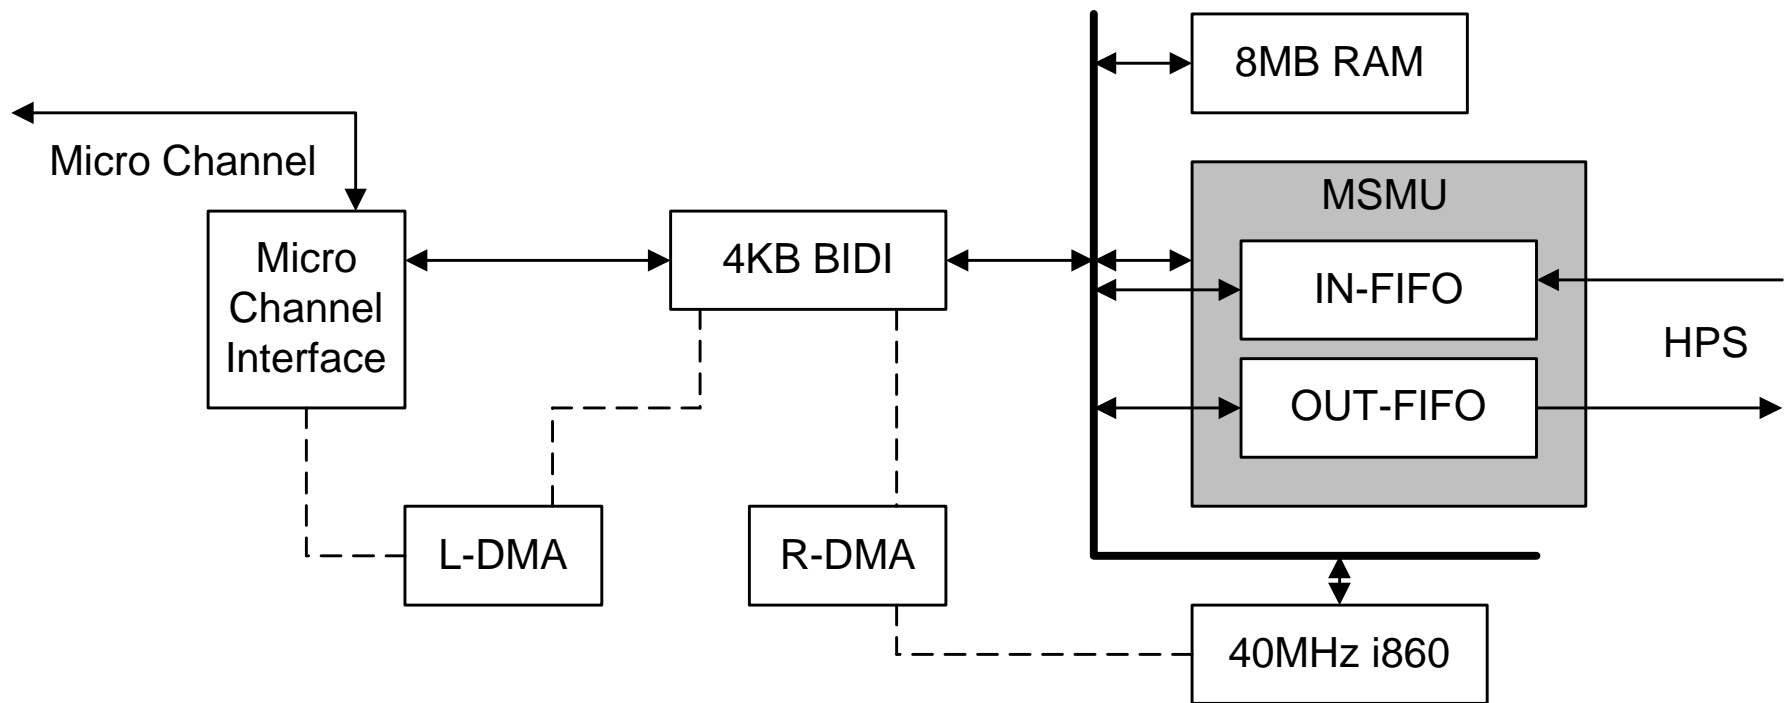



SP2

- Based in RS/6000 nodes
 - POWER2 processors
- Special NIC: MSMU on the micro-channel bus
- Standard Ethernet on the micro-channel bus
- MSMUs interconnected via a HPS backplane




SP2 MSMU







SP2 HPS

- Links are 8 bit parallel
- Contention free latency is 5 ns per stage
 - 875 ns latency for 512 nodes




SP2 HPS



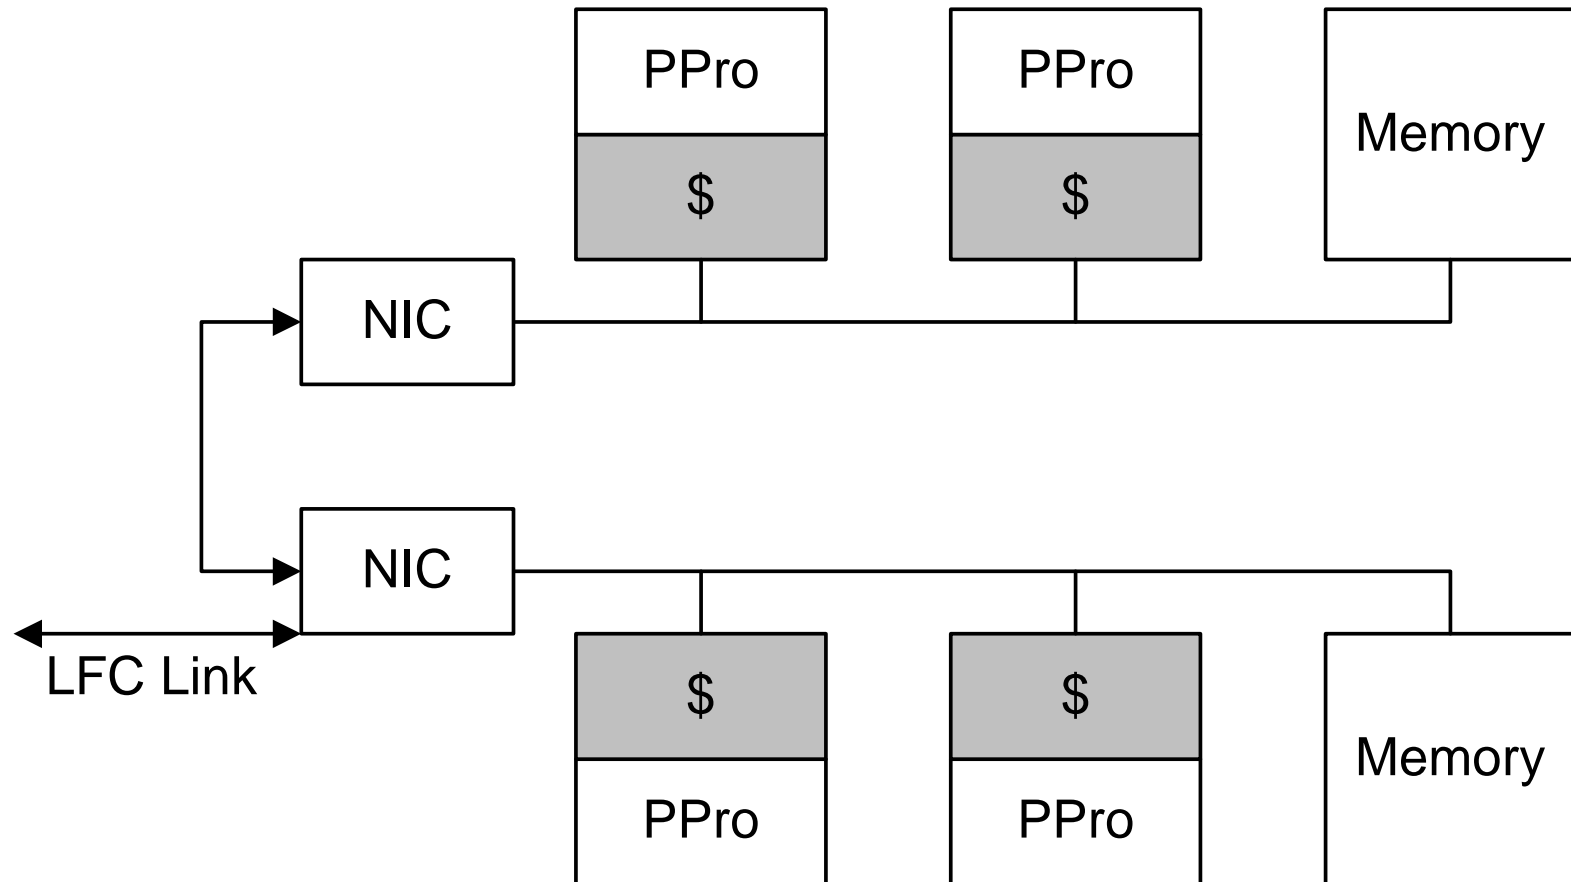



ASCI Red

- Build by Intel for the Department of “Energy”
- Consist of almost 5000 dual PPro boards with a special adaptation for user-level message-passing
- Special support for internal ‘firewalls’

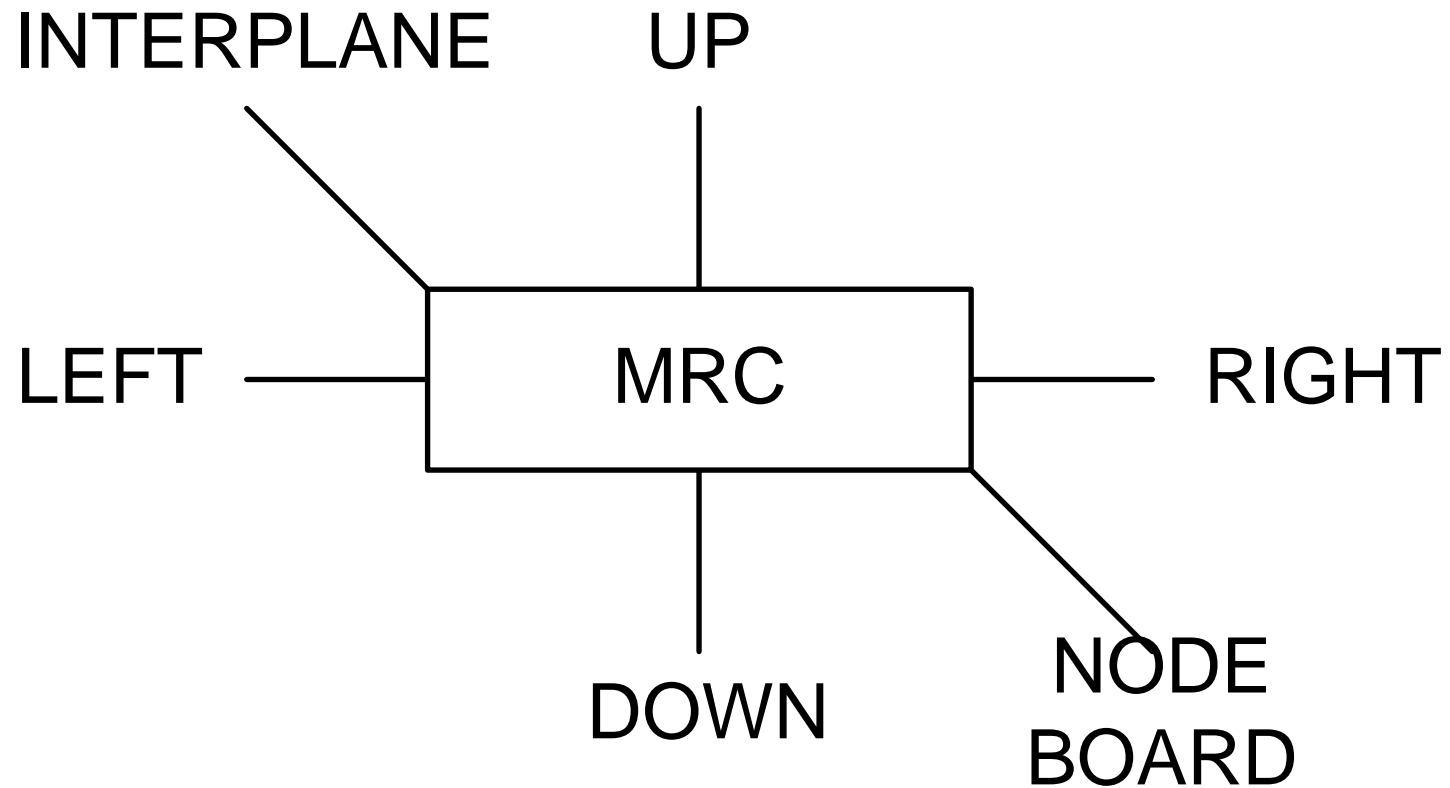



ASCI Red Node







ASCI Red MRC






ASCI Red Grid







Scali

- Based on Intel or Sparc based nodes
- Nodes are connected by a Dolphin SCI interface, using a grid of rings
- Very high performance MPI and support for commodity operating systems




Performance???



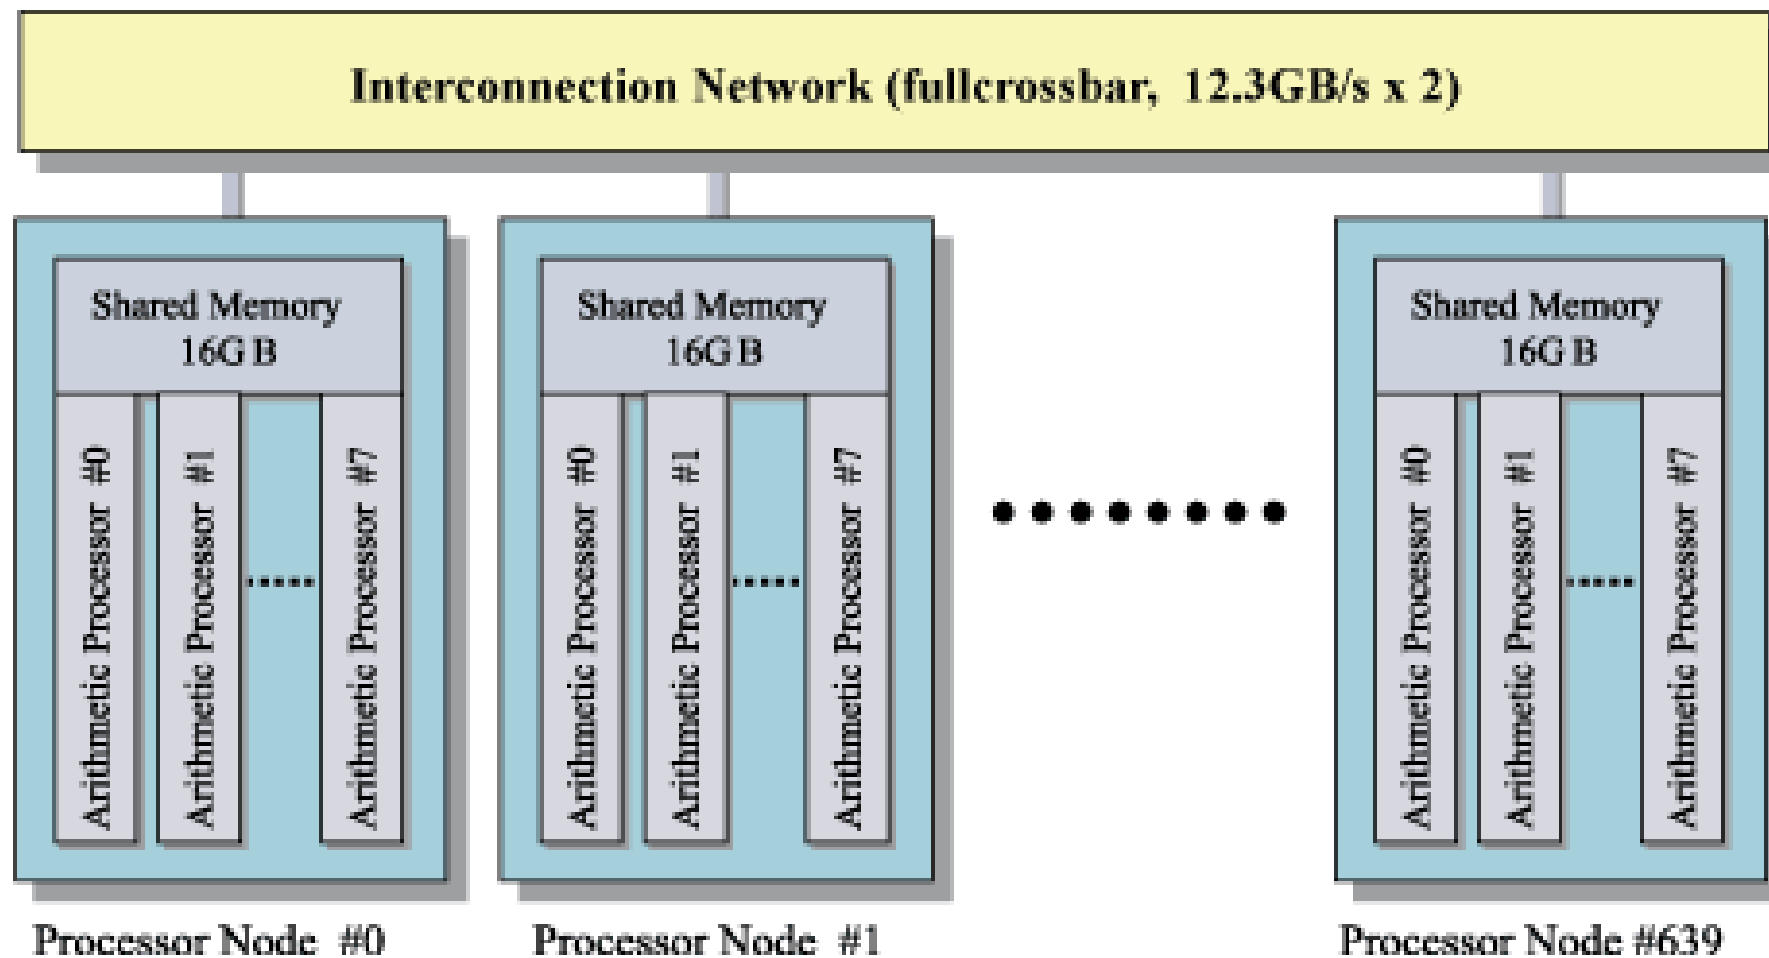



Earth Simulator





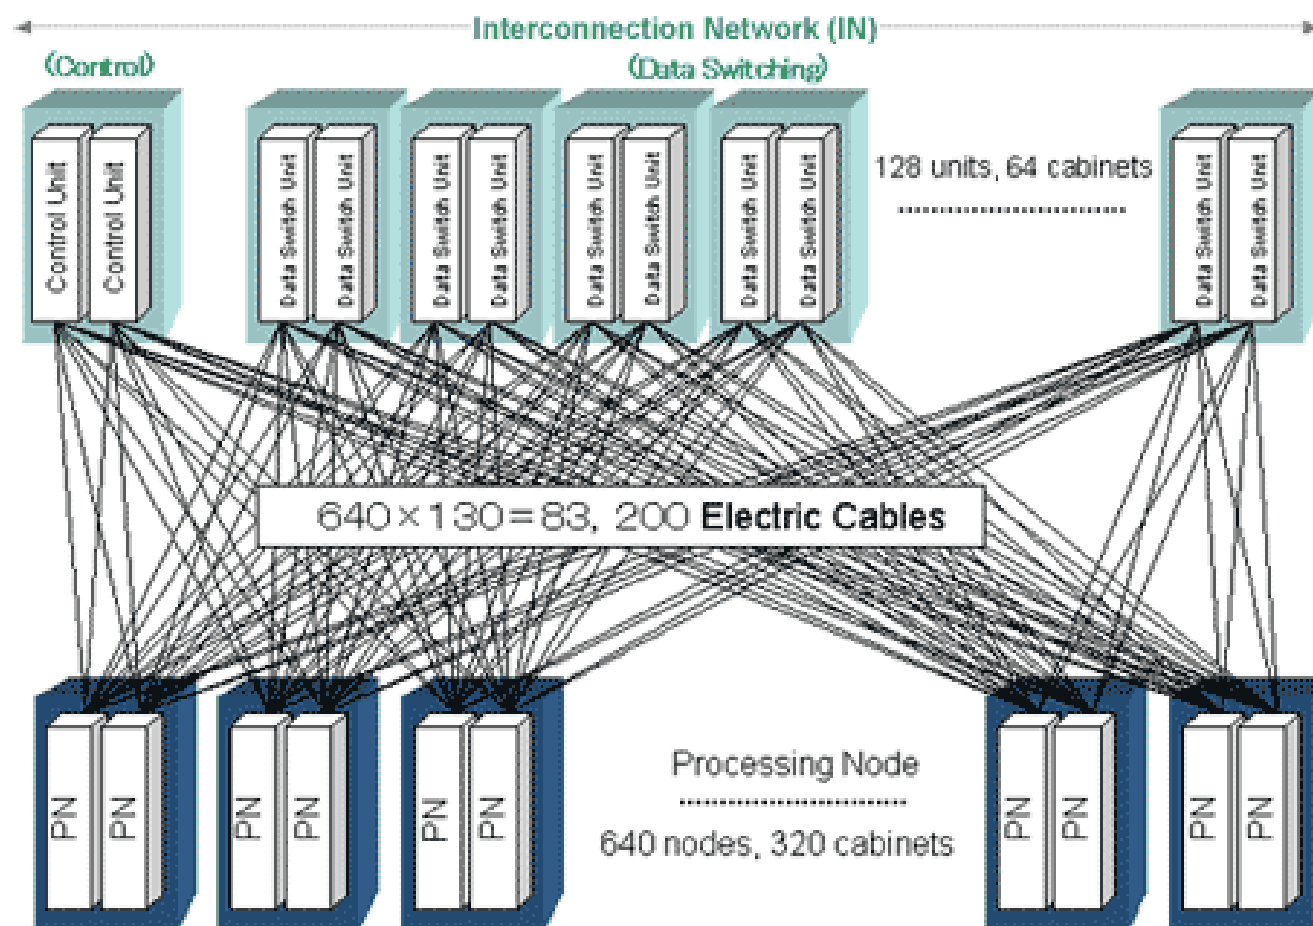
ES







ES


Interconnection Network





ES





ES

Wiring of interconnection network cables



Earth Simulator Research and Development Center