

Center for Computational Sciences, University of Tsukuba www.ccs.tsukuba.ac.jp

PACS-CS System Hardware and Software



number of nodes	2560 (16 x 16 x 10)
Peak performance	14.3 TFLOPS
Node configuration	Single CPU / node
CPU	Intel LV Xeon EM64T 2.8GHz 1MB L2 cache
Memory	DDR2/400 2GB/node (5.12 TB/system)
Network	3-dimensional Hyper-Crossbar Network
Link bandwidth	750MB/s (3-D simultaneous trans.)
Local HDD	160 GB/node (RAID-1, usable space)
Total system size	59 racks
Power consumption	550 kW
Operating system	Linux (Fedore Core 3) + SCore middleware
MPI	YAMPI and MPICH

- Single CPU per node for high memory bandwidth (1.1 Byte/FLOP)
- Using commodity GbEthernet NICs and switches with software trunking for cost-effective wide band-
- width interconnection (0.13 Byte/FLOP)
- 3-dimensional Hyper-Crossbar network for wide
- aggregated bandwidth per node to support various configuration of nearest neighboring mesh models
- Wide bisection bandwidth on any of three dimensions (640 GB/s on each dimension)
- Fault tolerant local hard disk drives in RAID-1 configuration for system and user space

Separated dual nodes on 1-U chassis in the same

Specially designed high-throughput and

system

density with 2-socket configured dual-Xeon

latency network layer (PM/Ethernet-HXB) oper-



low-

Hardware implementation: Hitachi Co. Ltd. ated under SCore cluster middleware Software implementation: Fujitsu Co. Ltd. (PM/Ethernet - HXB)

3-D Hyper-Crossbar Network



A computation node is equipped with three (X-, Y- and Z-dimension) of paired on-board GbE NICs (6 ports in total) for data communication. Nodes on a single line of a dimension are connected by an L2 GbE switch. For 3-D nearest neighboring communication, the node can communicate with surrounding nodes simultaneously with aggregated 750 MB/s of theoretical peak bandwidth.

A dedicated network layer PM/Ethernet-HXB provides the feature of network trunk with a paired links and high-speed routing on 3-D. This network is suitable for direct physical mapping of problems with spatial domain decomposition.

Dedicated Mother Board and Chassis



