Access to the
Federal High-Performance Computing-Centers

Rolf Rabenseifner

University of Stuttgart
High-Performance Computing-Center Stuttgart (HLRS)
www.hlrs.de

High Performance Computing in Germany

Federal
- HLRS/HWW Stuttgart
- NIC Jülich
- LRZ/HLRB Munich

Regional
- ZIB Berlin
- RRZN Hannover

Thematic (federal)
- RZ Garching
- DKRZ Hamburg
- DWD Offenbach

See also
www.hlrs.de/hpc/centers.html

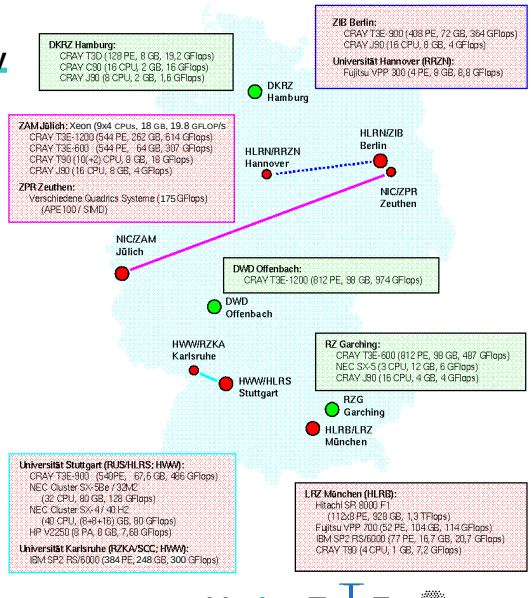

Architectures

Parallel Vector Processor (PVP) & Symmetric Multi Processor (SMP)

- NEC SX-5Be and SX-4
- Cray T90, C90, J90

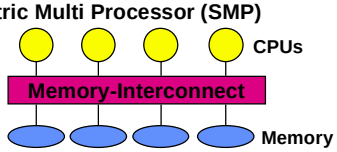

MPP, massively parallel processor achieve

- Cray T3E-1200, 900, 600
- Fujitsu VPP 700, 300

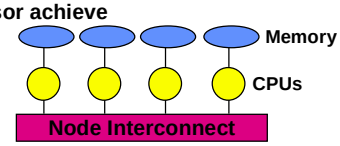

Hierarchical: Cluster of SMP nodes

- Hitachi SR8000
- IBM SP

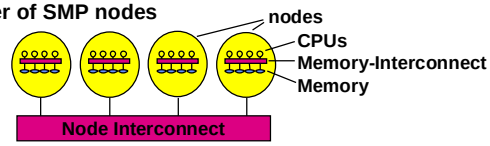

Vector Processing

- most platforms perform vector processing
  - vector processor
  - mainly a question of
    - memory bandwidth
      - how many load/stores per floating point operation?
    - memory-interconnect
  - for which memory access pattern
    - the compiler supports the vector access
    - the hardware supports the vector access

Platforms

Parallel Vector Processor (PVP) & Symmetric Multi Processor (SMP)

- NEC SX-5Be and SX-4: Shared memory parallel vector processor
  - SX-5Be: 16 CPUs per node, up to 4 GFLOP/s per CPU peak, vectorizable codes achieve ~50%
    nodes are clustered via IXS
  - SX-4: up to 32 CPUs per node, 2 GFLOP/s per CPU peak, vectorizable codes achieve ~50%
- Cray T90, C90, J90: Shared memory parallel vector processor
  - older systems

MPP, massively parallel processor achieve

- Cray T3E-1200, 900, 600: Distributed memory, high speed interconnect: 300 MB/s
  - T3E-1200 & -900: stream buffers (vector feature), peak: 1.2 or 0.9 GFLOP/s per CPU,
    good codes achieve ~50 %, sustained ~10%
  - T3E-600: without stream buffers, therefore peak performance only on the cache
- Fujitsu VPP 700, 300: Distributed memory vector processor (PVP, but not SMP)
  - Peak: 2.2 GFLOP/s per CPU

Hierarchical: Cluster of SMP nodes

- Hitachi SR8000: 8 CPUs per node, pseudo vectorization, node-to-node: 950 MB/s
  - Each node can be programmed (a) with automatic parallelization, (b) with OpenMP, (c) as MPP (MPI)
  - SR 8000 F1: peak 12 GFLOP/s per node, 1.5 GFLOP/s per CPU, sustained ~20%
  - SR 8000: peak 8 GFLOP/s per node, 1.0 GFLOP/s per CPU
- IBM SP: typically heterogeneous cluster of different node types: *Wide, Thin, High nodes*
  - High node: peak 12 GFLOP/s per node, 1.5 GFLOP/s per CPU, 8 CPUs per node

Programming Models

Parallel Vector Processor (PVP) & Symmetric Multi Processor (SMP)

- NEC SX-5Be and SX-4: Shared memory parallel vector processor
  - Vectorization and
  - Parallelization: automatic, by vendor's directives, OpenMP, (MPI, proc.-to-proc. 8 GB/s)
- Cray T90, C90, J90: Shared memory parallel vector processor
  - Vectorization and (automatic) Parallelization

MPP, massively parallel processor

- Cray T3E-1200, 900, 600: Distributed memory, high speed interconnect: 310 MB/s
  - MPI and shmem, simple vectorization
- Fujitsu VPP 700, 300: Distributed memory vector processor
  - Vectorization and MPI

Hierarchical: Cluster of SMP nodes —> hierarchical programming model

- Hitachi SR8000: 8 CPUs per node, pseudo vectorization, node-to-node: 990 MB/s
  - Pseudo vectorization: automatic by compiler
  - Each node can be programmed (a) with automatic parallelization, (b) with OpenMP, (c) as MPP (MPI)
  - Combined with MPI multi-node parallelization
- IBM SP: typically heterogeneous cluster of different node types: *Wide, Thin, High nodes*
  - Depends on nodes: Cache-based optimization, OpenMP, MPI

HPC Platforms in Germany

SMP, Parallel Vector — State: 6 / 2000, updated 1/2001

- **NEC SX-5Be and SX-4:** Shared memory parallel vector processor, proc.-to-proc. 8 GB/s
  - Stuttgart 208 GFLOP/s, 96 GB, 4 Nodes with 72 CPUs (16+16+32+8)
  - Garching 6 GFLOP/s, 12 GB, 3 CPUs
- **Cray T90, C90, J90:** Shared memory parallel vector processor
  - Jülich 22 GFLOP/s, 16 GB, 28 CPUs
  - Hamburg 18 GFLOP/s, 4 GB, 24 CPUs
  - Munich 7 GFLOP/s, 1 GB, 4 CPUs
  - Garching 4 GFLOP/s, 4 GB, 16 CPUs
  - Berlin 4 GFLOP/s, 8 GB, 16 CPUs

MPP

- **Cray T3E-1200, 900, 600:** Distributed memory, high speed interconnect: 310 MB/s
  - Offenbach 974 GFLOP/s, 98 GB, 812 PEs
  - Jülich 921 GFLOP/s, 326 GB, 1088 PEs
  - Stuttgart 486 GFLOP/s, 67 GB, 540 PEs
  - Garching 487 GFLOP/s, 96 GB, 812 PEs
  - Hamburg 19 GFLOP/s, 8 GB, 128 Pes (T3D)
- **Fujitsu VPP 700, 300:** Distributed memory vector processor
  - Munich 114 GFLOP/s, 17 GB, 77 PEs
  - Hannover 8.8 GFLOP/s, 8 GB, 4 PEs

Hierarchical

- **Hitachi SR8000:** 8 CPUs per node, pseudo vectorization, node-to-node: 990 MB/s
  - Munich 1300 GFLOP/s, 928 GB, 112*8 PEs
  - Stuttgart 128 GFLOP/s, 128 GB, 16*8 PEs
- **IBM SP:** typically heterogeneous cluster of different node types: *Wide, Thin, High nodes*
  - Karlsruhe 300 GFLOP/s, 248 GB, 384 PEs
  - Munich 114 GFLOP/s, 104 GB, 52 CPUs

HLRS Stuttgart: hww configuration

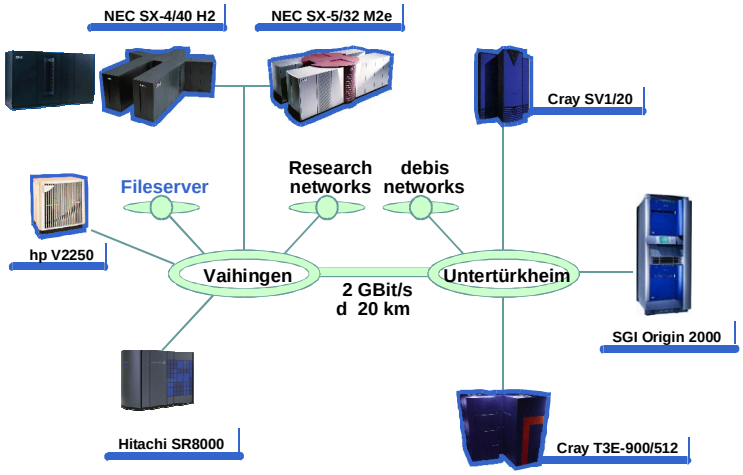

HLRS at Stuttgart

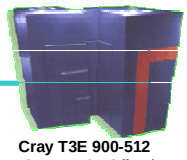

Cray T3E 900-512
512 PEs, 461 Gflop/s
64 GByte

- High-Performance Computing-Center Stuttgart
  Höchstleistungsrechenzentrum Stuttgart (HLRS)
  at Rechenzentrum Universität Stuttgart (RUS)
- Platforms:
  - NEC SX-4/40 and SX-5/32
  - Hitachi SR 8000
  - Cray T3E 900-512
- Online-Proposals:
  - www.hlrs.de
    —> HW&Access —> Access
    —> proposal
    —> test account (***Schnupperaccount***)

NEC SX-5/32 M2e
2*16 PEs, 128 Gflop/s
80 GByte

NEC SX-4/40 H2
32+8 PEs, 80 Gflop/s
8+8 GByte

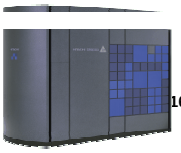

Hitachi SR8000
16*8 PEs, 128 Gflop/s
128 GByte

NIC at Jülich

- John von Neumann-Institut für Computing (NIC)
  at Forschungszentrum Jülich
- Platforms:
  - Cray T3E 1200-512
  - Cray T3E 600-512
- Online-Proposals:
  - www.fz-juelich.de/nic/nic-e.html —> Computing Time
- Deadlines
  - ~ April 30, for accounts valid July – June (next year)
  - ~ Oct. 31, for accounts valid Jan. – Dec. (next year)

LRZ / HLRB at Munich: Hitachi SR8000

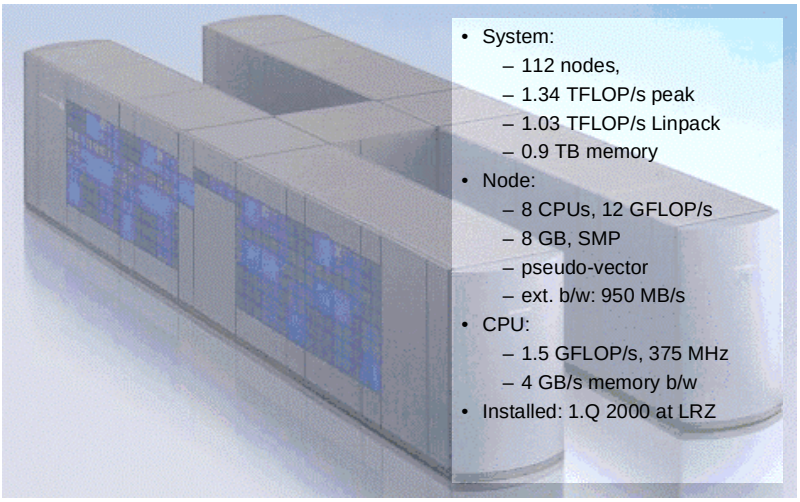

- System:
  - 112 nodes,
  - 1.34 TFLOP/s peak
  - 1.03 TFLOP/s Linpack
  - 0.9 TB memory
- Node:
  - 8 CPUs, 12 GFLOP/s
  - 8 GB, SMP
  - pseudo-vector
  - ext. b/w: 950 MB/s
- CPU:
  - 1.5 GFLOP/s, 375 MHz
  - 4 GB/s memory b/w
- Installed: 1.Q 2000 at LRZ

Hitachi SR 8000-F1/112 (Rank 5 in TOP 500 / June 2000)

- Höchstleistungsrechner in Bayern (HLRB)
  at Leibniz-Rechenzentrum (LRZ)
- Platform:
  - Hitachi SR8000-F1 (112 nodes, each with 8 processors)
- Online-Proposals:
  - www.lrz.de/services/compute/hlrb/
    —> Online Project Proposal Submission

High Performance Computing in Germany — Summary

Federal
- HLRS/HWW Stuttgart
- NIC Jülich
- LRZ/HLRB Munich

Regional
- ZIB Berlin
- RRZN Hannover

Thematic (federal)
- RZ Garching
- DKRZ Hamburg
- DWD Offenbach

See also
www.hlrs.de/hpc/centers.html

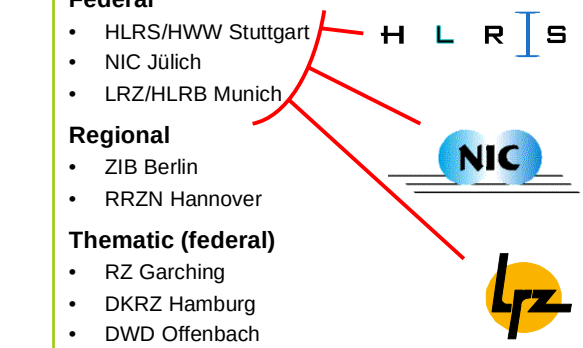

www.hlrs.de
–>HW&Access–>Access

www.fz-juelich.de/nic/nic-e.html

www.lrz.de/services/compute/hlrb/