

PC² User Meeting 2018

- How to apply for computing time at PC² -

Thomas D. Kühne

Chair of Theoretical Chemistry

PC² - Paderborn Center for Parallel Computing

December, 10th 2018

<https://pc2.uni-paderborn.de>

Noctua: Cray CS 500 Storm

- 256 compute nodes
 - 2x Intel Xeon Gold 6148, 40 cores, 2.4 GHz
 - 192 GiB main memory
- 16 FPGA nodes
 - Intel Xeon 6148+6148F, 192 GiB
 - each with 2 Stratix 10 FPGA cards
- Parallel file system
 - Lustre
 - 720 TB disk capacity
- Interconnect Intel Omni-Path
 - 100 Gbit/s network
 - Blocking factor 1:1.4



OCuLUS: Owl CLUStEr

- 572 compute nodes
 - 2x Intel Xeon E5-2670, 16 cores, 2.6 GHz
 - 64 / 256 GiB main memory
- 40 GPU nodes
 - 2x Intel Xeon E5-2670, 16 cores, 2.6 GHz
 - 64 GiB main memory
 - Nvidia Tesla K20 / GeForce GTX 1080 Ti
- 4 SMP nodes
 - 4x Intel Xeon E5-4670, 32 cores, 2.7 GHz
 - 1024 GiB main memory
- Parallel file system
 - FHG BeeGFS
 - 500 TB disk capacity
- Interconnect QDR InfiniBand
 - 40 Gbit/s network
 - Blocking factor 1:2



Noctua1 vs. OCuLUS (total)

Metric	Noctua1	OCuLUS	Diff.
# nodes	272	616	- 56%
# cores	10.880	9.856	+ 10%
Total memory [TiB]	52.2	41.2	+ 27%
HP-Linpack [TFlop/s]	537	188.7	x 2.8
Accu. STREAM [GiB/s]	50,150	43,700	+ 15%
Accu. SpecFP2006 rate	381,000	295,700	+ 29%
Accu. SpecINT2006 rate	516,800	381.900	+ 35%
MPI network [Gbps]	100 (Omni Path)	40 (InfiniBand QDR)	x 2.5
blocking factor	1:1.4	1:2	
latency [μ s]	1.24	1.9	- 35%
bandwidth [GByte/s]	24,5	7	x 3.5
Storage Capacity [TB]	720	500	+ 44%
max. bandwidth [GiB/s]	20	25	- 20%
Power consumption [kW]	164	230	- 29%

Allocation of Computing Resources

- With the acquisition of Noctua, PC² will be partially opened nationwide
- More than 100mio core-h will be distributed on a competitive basis
 - At least one Flagship project
- All proposals will be reviewed
 - Technical evaluation by PC² staff
 - Scientific evaluation by Peer Review Board (PRB), which is appointed by PC² management board on recommendation of the advisory board
- Big emphasize on hardware-acceleration, in particular usage of FPGAs
- Proposal – “As simple as possible, but not simpler”
 - Noctua / OCuLUS projects vs. Test accounts on either machine
 - Scientific project description
 - Reproducible rationalization of requested core-h
 - Documented scalability efficiency or effective use of hardware-acceleration
 - Website-ready abstract
 - Without report on previous funding period, no new project

Peer Review Board

- Technical evaluation
 - Conducted by PC² staff members
 - Decision will be made based on the proper use of hardware-acceleration, or sufficient massive parallelism
 - Scalability efficiency of massive parallel projects
 - Low-core or embarrassingly parallel projects will be either discarded or forwarded to our high-throughput cluster
 - Only projects that passed the technical evaluation will be passed to the PRB
- Peer Review Board (PRB)
 - Decision of the the PRB are based on referee reports of qualified scientists in the respective field of research, which are either members of the PRB or, if necessary, external referees
 - Apart from scientific relevance, an important criterion is the sensible use of the available HPC hardware
 - I.e. preference will be given to projects that make use of hardware-acceleration, such as FPGAs or GPUs, or massive parallelism

Timeline

- Users
 - Early January: Necessary forms, website and submission system will be ready
 - 17.2.2019: Deadline to submit a proposal for 1st funding period
 - 1.4.2019: Begin of 1st funding period
 - 18.8.2019: Deadline to submit a proposal for 2nd funding period
 - 30.9.2019: Begin of 2nd funding period
- PC2 Staff Members
 - Early January: Necessary forms, website and submission system will be ready
 - Till mid February: Definition of scheduling and backfilling policy
 - End of February: Technical evaluation of received proposals
 - Till end of March: Definition and activation of queuing system
- Peer Review Board
 - Till end of January: Recommendation of PRB member by the PC² advisory board
 - Early February: Appointment of the PRB by the PC² management board
 - March: Evaluation of proposals that passed the technical evaluation

Queue Configuration

Name	Max. Nodes*	Max. Runtime*
short	2	30m
test	50	30m
batch	100	24h
long	50	21d
fpga	16	2h
all	272	12h

* Initial settings

Projects are restricted in maximum number of nodes and to certain queues.
The queue configuration is subject of change.

Applications / Libraries / Tools

CP2K
Gaussian
Turbomole

IntelMPI
OpenMPI
FFTW
MKL
OpenBLAS
ScaLapack

Easybuild
Python3
Valgrind
gdb4hpc

... and further applications on request. Please contact us.