

狼

Ookami User Group Meeting 03/23/2023

Parallel Programming Models on Ookami

Tony <anthony.curtis@stonybrook.edu>


Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- Ookami is a high performance computing (HPC) cluster of Fujitsu A64FX processors
 - It only has CPUs, not GPUs!
 - Nodes connected with Infiniband
 - So how do we program the system?
 - Let's take a look at the nodes

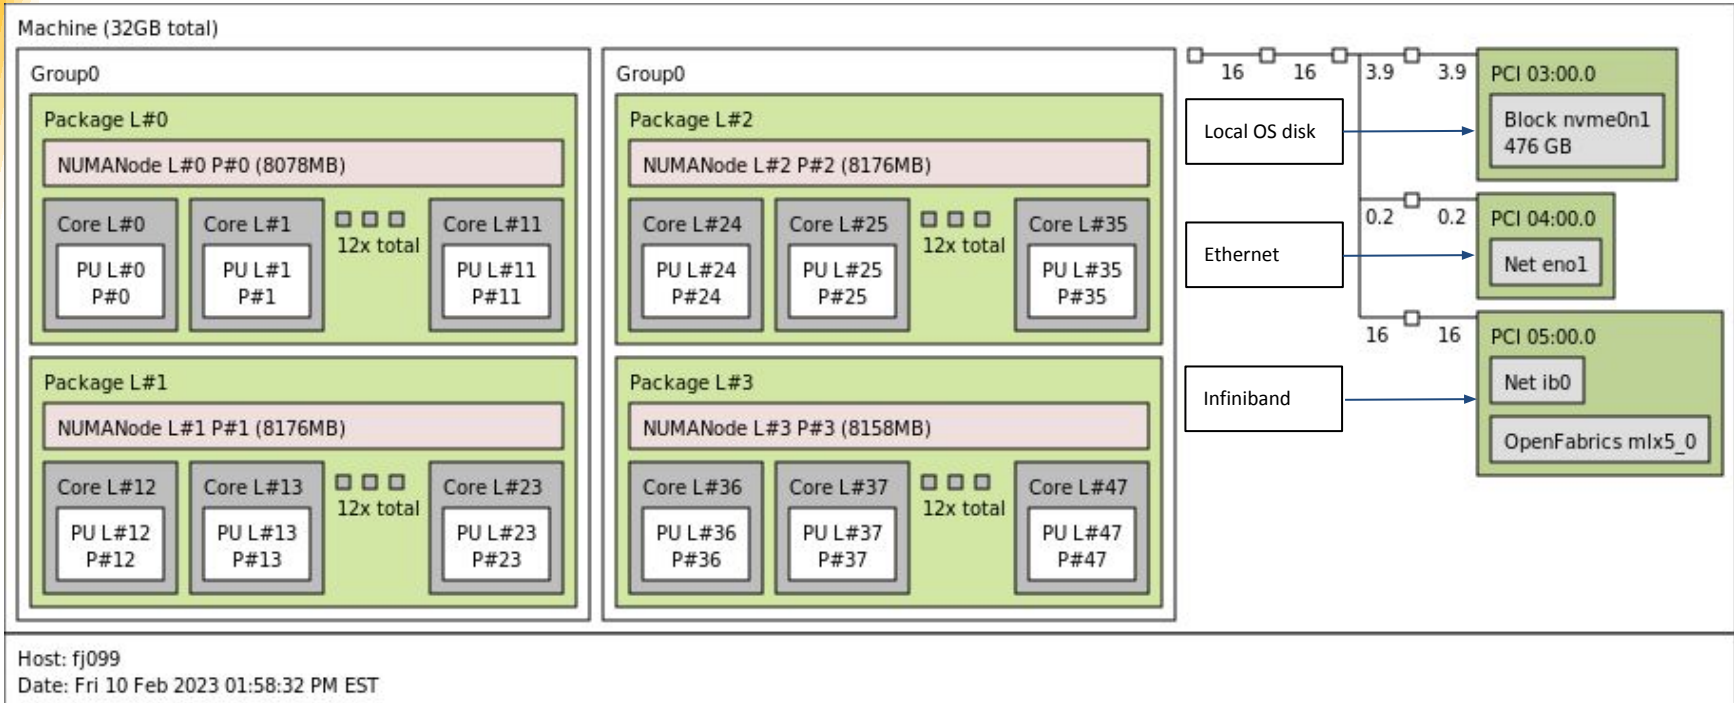
Ookami User Group Meeting 2023

Parallel Programming Models on Ookami



Ookami User Group Meeting 2023

Parallel Programming Models on Ookami



Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- Ookami's compute nodes have an interesting architecture:
 - 4 NUMA groups (Core Management Groups) of 12 processors
 - Each CMG has 8GB of high bandwidth memory (HBM)
 - On-chip communication ring that connects all the groups & memories
 - Groups 0 and 1 are paired, 2 and 3 are paired
 - Communication metrics are
 - 10 to self
 - 20 to pair/twin
 - 30 to other

```
fj-debug1$ numactl -H
node  0  1  2  3
0:    10  20  30  30
1:    20  10  30  30
2:    30  30  10  20
3:    30  30  20  10
```

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- Which parallel programming models are available?
- Let's break it down into 1 node and many nodes...

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On 1 node
 - Threads (low-level)
 - Pthreads
 - Linux/POSIX standard
 - Basis of most OpenMP implementations
 - Qthreads
 - Alternative from Sandia
 - E.g. implementation choice in Chapel

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On 1 node
 - OpenMP & compilers for C, C++, Fortran
 - GCC
 - `-fopenmp`
 - ARM
 - `-fopenmp`
 - Fujitsu
 - `-Kopenmp` (traditional mode) and/or `-fopenmp` (Clang mode)
 - Cray
 - `-homp` and/or `-fopenmp` (newest version)
 - NVIDIA
 - `-fopenmp`

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On many nodes
 - MPI
 - Open-MPI
 - Has multiple modules that understand both intra- and inter-node communication, uses UCX
 - » e.g. XPMEM, KNEM, Infiniband, TCP, Collectives
 - We have builds with different compilers, esp. for different Fortran capabilities (e.g. for LLVM/ARM flang)
 - De-facto wrapper names: mpicc, mpicxx, mpif90

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On many nodes
 - MPI
 - MVAPICH2
 - The well-known MPI implementation from Ohio State
 - Integrated with SLURM, launch via `srun` not `mpirun/mpiexec`
 - We have versions installed by Cray for use with their compilers, *and* local “vanilla” versions

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On many nodes
 - MPI
 - MPICH
 - One of the “original” MPI implementations from Argonne (ANL)
 - Also supports various intra- and inter-connects, via UCX or libfabric
 - N.B. Cray’s MPI on their “big” machines is based on MPICH, but this isn’t it!

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On many nodes
 - MPI
 - Fujitsu
 - Vendor-customization of Open-MPI
 - » Meant mostly for the TOFU interconnect
 - E.g. as on Fugaku
 - » But also works on Infiniband
 - Custom wrappers: can confuse build systems
 - » `mpifcc`, `mpiFCC`, `mpifrt`

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On many nodes
 - OpenSHMEM
 - Available as part of Open-MPI
 - Accent on sparse/irregular 1-sided communication
 - » Put/get data with no corresponding receive call
 - » Relaxed synchronization

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- Both? YES!
 - Hybrid: usually known as MPI + X
 - MPI to distribute work across nodes
 - OpenMP (often the “X”) to parallelize computation on a node
 - Ookami is a great target for this model
 - Understand thread and process placement
 - Ask on Slack any time, and visit office hours for hands-on

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- On many nodes
 - Chapel
 - Programming language from HPE/Cray
 - High-level local- and global-views of parallelism
 - https://www.stonybrook.edu/commcms/ookami/support/index_links_and_docs.php

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- Newer languages/environments to consider/test
 - Chapel (HPE/Cray language) *ookami modules*
 - Kokkos (a64fx-aware!) *ookami modules*
 - Legion/Regent
 - Go
 - Rust
 - Julia (a64fx-aware!) *ookami modules*
 - Native C++ parallelism

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

- Stony Brook's community involvement?
 - Barbara Chapman's group is involved in:
 - OpenMP ARB
 - LLVM (clang and GPU-offload friends)
 - SPEC HPG benchmarks
 - PMIx, PR RTE, XPMEM, UCX, Open-MPI
 - OpenACC (not overly relevant for Ookami)

Ookami User Group Meeting 2023

Parallel Programming Models on Ookami

. Selected links

- <https://www.open-mpi.org/>
- <https://mvapich.cse.ohio-state.edu/downloads/>
- <https://www.mpich.org/downloads/>
- <https://www.openmp.org/>
- <https://chapel-lang.org/>
- <https://kokkos.github.io/>
- <https://legion.stanford.edu/>
- <https://www.stonybrook.edu/commcms/ookami/support/faq/ookami-fujitsu-compilers>
- <https://www.stonybrook.edu/commcms/ookami/support/faq/index.php>
- <http://www.openshmem.org/>