

SORS: Topology-aware placement and load-balancing

Speaker: Emmanel Jeannot, Senior research scientist, LaBRi laboratory at INRIA-Bordeaux

Abstract: Current generation of NUMA nodes clusters feature multicore and many core processors. Programming such architectures efficiently is a challenge because numerous hardware characteristics have to be taken into account, especially the memory hierarchy. One appealing idea to improve the performance of parallel applications is to decrease their communication costs by matching the communication pattern to the underlying hardware architecture. In this task we detail the algorithm and techniques proposed to achieve such a result. First, we gather both the communication pattern information and the hardware details. Then we compute a relevant reordering of the various process ranks of the application. Finally, those new ranks are used to reduce the communication costs of the application. To show the relevance of this approach we have:

- 1) compared the placement to standard MPI ones (round-robin)
- 2) developed two load balancers for Charm++ that take into account topology and communication aspects depending on the fact that the application is compute-bound or communication-bound
- 3) implement it in a batch scheduler (SLURM) to improve the election of the nodes for a given application.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 15 jul 2024 - 08:26): <https://www.bsc.es/ca/research-and-development/research-seminars/sors-topology-aware-placement-and-load-balancing>