(to be) published in Proceedings, Euro-Par '98, Southampton, GB, Sept. 1-4, 1998.

MPI-GLUE: Interoperable high-performance MPI combining different vendor's MPI worlds

Rolf Rabenseifner

Rechenzentrum der Universität Stuttgart, Allmandring 30, D-70550 Stuttgart,
rabenseifner@rus.uni-stuttgart.de,
WWW home page: http://www.hlrs.de/people/rabenseifner/

Abstract. Several metacomputing projects try to implement MPI for homogeneous and heterogeneous clusters of parallel systems. MPI-GLUE is the first approach which exports nearly full MPI 1.1 to the user's application without losing the efficiency of the vendors' MPI implementations. Inside of each MPP or PVP system the vendor's MPI implementation is used. Between the parallel systems a slightly modified TCP-based MPICH is used, i.e. MPI-GLUE is a layer that combines different vendors' MPIs by using MPICH as a global communication layer. Major design decisions within MPI-GLUE and other metacomputing MPI libraries (PACX-MPI, PVMPI, Globus and PLUS) and their implications for the programming model are compared. The design principles are explained in detail.

Due to the copyright rules of the Springer Verlag, I have to wait for one year before I am allowed to put the full text version on the WWW. Therefore this file will be substituted by the full version of this paper after one year, i.e. after Sept. 1, 1999.

If you need a privat copy earlier then please feel free to send me a mail:

rabenseifner@rus.uni-stuttgart.de