MPI-3 TOOLS: DLL DISCOVERY & MPI HANDLE DEBUGGING

Jeff Squyres, Cisco Systems

Prerequisites

- Read DLL proposal
  - https://svn.mpi-forum.org/trac/mpi-forum-web/wiki/MPI3Tools/dllapi
- Read MPI handle debug proposal
  - https://svn.mpi-forum.org/trac/mpi-forum-web/wiki/MPI3Tools/handleapi
    - https://svn.open-mpi.org/trac/ompi/browser/trunk/ompi/debuggers/MPI_Handles_interface.txt
    - https://svn.open-mpi.org/trac/ompi/browser/trunk/ompi/debuggers/mpihandles_interface.h

DLL discovery: short version

Old way

- MPI provides “char *MPI_dll_name”
- Debugger reads it at some point

New way

- MPI provides argv-style “char **mpimsgq_dll_locations”
- Debugger does not read it until MPIR_debug_gate()
  - Debugger tries to dlopen (etc.) each string
  - Uses first one that is “acceptable”

MPI Handle Debugging

- Debugger can show *value* of MPI handles
  - Some random int or pointer value
- But this tells the user nothing about the underlying MPI object

- Goal: debugger can show meaningful information about the MPI object

New API proposal: 50k feet

- New DLL (independent of msgq DLL)
- Similar in spirit to existing mqs interface
  - Re-uses most of the existing mqs interface
  - Mostly the same startup/shutdown sequence
  - Renamed a few of startup/shutdown the functions (compared to the message queue debugging)
- Main idea: debugger queries MPI for "meaning" of MPI handles

Debugger queries DLL

- DLL tells debugger types of each handle
  - If handles are ints, debugger *may* not be able to automatically know to interpret
  - …unless it can also see “MPI_Comm” …?
- Debugger can query DLL for info about handles

```
int mpidbg_comm_query(mqs_image *image, mqs_image_info *image_info,
                      mqs_process *process, mqs_process_info *process_info,
                      mqs_taddr_t c_comm, struct mpidbg_comm_info_t **info);
```

Debugger queries DLL

- Handle Fortran and C++ handles as well
  - mpidbg_comm_f2c(…)
  - mpidbg_comm_cxx2c(…)
- DLL queries image and returns a Big Struct
  - Specific to each type of MPIhandle
  - E.g., return communicator name, type (and meta data), size, this proc's rank, pending requests, … etc.

Interface Status

- Open MPI implemented four handle types
  - Communicator
  - Errorhandler
  - Request
  - Status
- Intent was to do a proof-of-concept
  - Then talk to community for feedback / feasibility before doing more