Exposing MPI Objects for Debugging

Laust Brock-Nannestad [1] **John DelSignore** [2] **Jeffrey M. Squyres** [3] **Sven Karlsson** [1] **Kathryn Mohror** [4]

[1]Technical University of Denmark [2]Rogue Wave Software, Inc. [3]Cisco Systems, Inc. [4]Lawrence Livermore National Laboratory

Motivation

- Debugging MPI applications is difficult as developers cannot inspect the state of the MPI runtime
- The MPI Tools Working Group has proposed a standard interface between the debugger and MPI runtime [1]
- With this interface, debuggers can easily present MPI state to the developer

Contributions

- Support for MPI handle introspection in the *TotalView* debugger
- A reference introspection implementation in *Open MPI*
- A demonstration of simplified MPI debugging

Debugger access to MPI state

- MPI implementation provides functions for *introspection* as a library
- Introspection functions extract data from MPI handles
- Debugger and library interact through mutual callback functions [2]

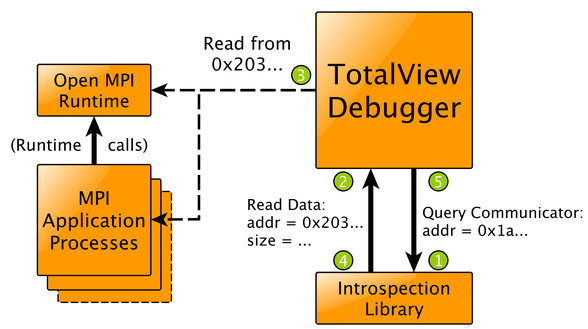

Figure 1 : Example use case to fetch communicator information.

1. Debugger queries a communicator
2. Library requests raw data from process
3. Debugger extracts raw data from process
4. Library parses raw data and returns standardized communicator information
5. Debugger presents data to developer

Complex MPI Debugging Session with a Debugger

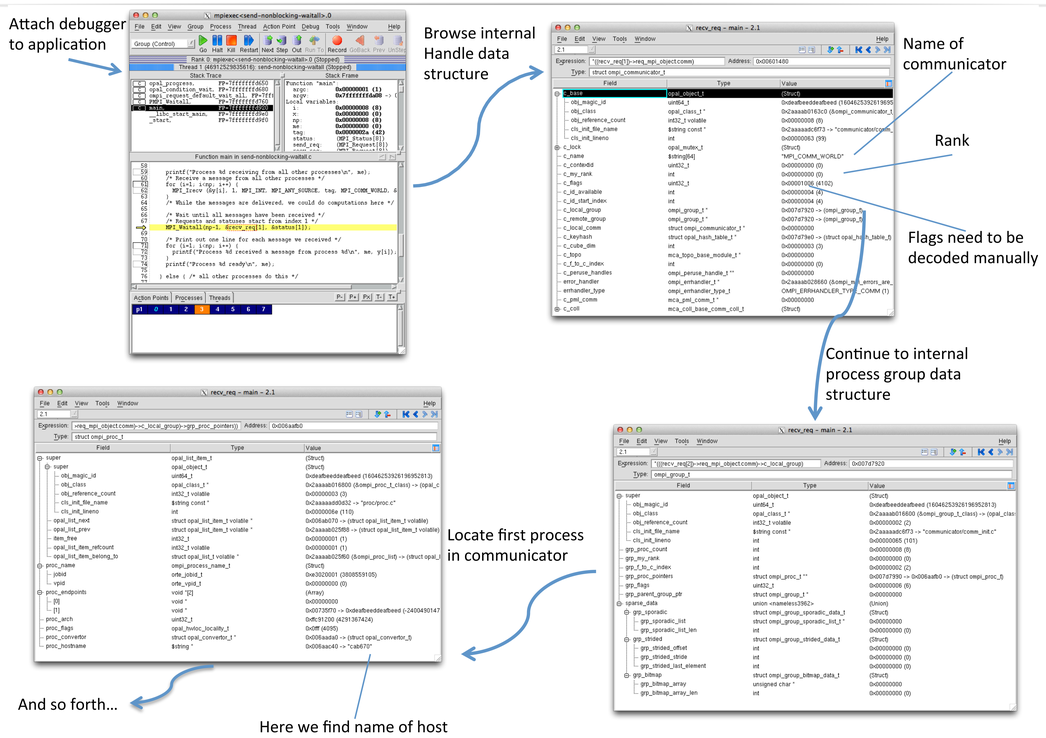

Figure 2 : *Best case* scenario for existing MPI debugging. Developer must understand internal data structures. For closed source MPI implementations, even this is not possible.

Simplified Debugging with MPI Handle Introspection in TotalView

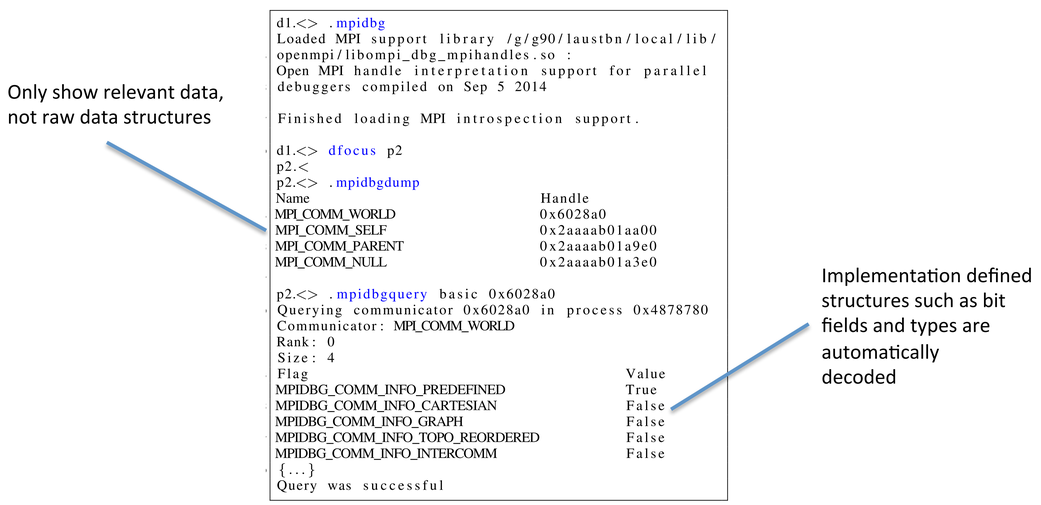

Figure 3 : Debugging with introspection support thorugh the TotalView command line. TotalView displays the status of a communicator.

- MPI implementation traverses data structures – assisted by the Debugger
- Developer sees MPI *API level* information

Future Work

- Integration into TotalView's graphical user interface
- Validation against other MPI implementations
- Visualization of MPI communication and processes
- Querying MPI handle state from the debugger; filtering and showing handles matching certain criteria

Acknowledgements

The authors would like to thank Adam Moody at Lawrence Livermore National Laboratory for useful insights on MPI debugging.

This work (LLNL-POST-658417) was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, and co-funded by the European Artemis PaPP Project nr. 295440 and COPCAMS project nr. 332913.

References

[1] MPI Tools Working Group.
`https://svn.mpi-forum.org/trac/mpi-forum-web/wiki/MPI3Tools`, October 2014.

[2] James Cownie and William Gropp.
A Standard Interface for Debugger Access to Message Queue Information in MPI.
In *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 51–58. Springer, 1999.

Contact Information

Laust Brock-Nannestad
Technical University of Denmark
Email: laub@dtu.dk
Phone: +45 4525 9223

`http://www.compute.dtu.dk/~laub/`