



## Exposing MPI Objects for Debugging

**Brock-Nannestad, Laust; DeSignore, John; Squyres, Jeffrey M.; Karlsson, Sven ; Mohror, Kathryn**

*Publication date:*  
2014

*Document Version*  
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

*Citation (APA):*

Brock-Nannestad, L., DeSignore, J., Squyres, J. M., Karlsson, S., & Mohror, K. (2014). *Exposing MPI Objects for Debugging*. Poster session presented at International Conference for High Performance Computing, Networking, Storage and Analysis, SC14, New Orleans, United States.  
[http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/tech\\_poster/tech\\_poster\\_pages/post284.html](http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/tech_poster/tech_poster_pages/post284.html)

---

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



# Exposing MPI Objects for Debugging

## 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]

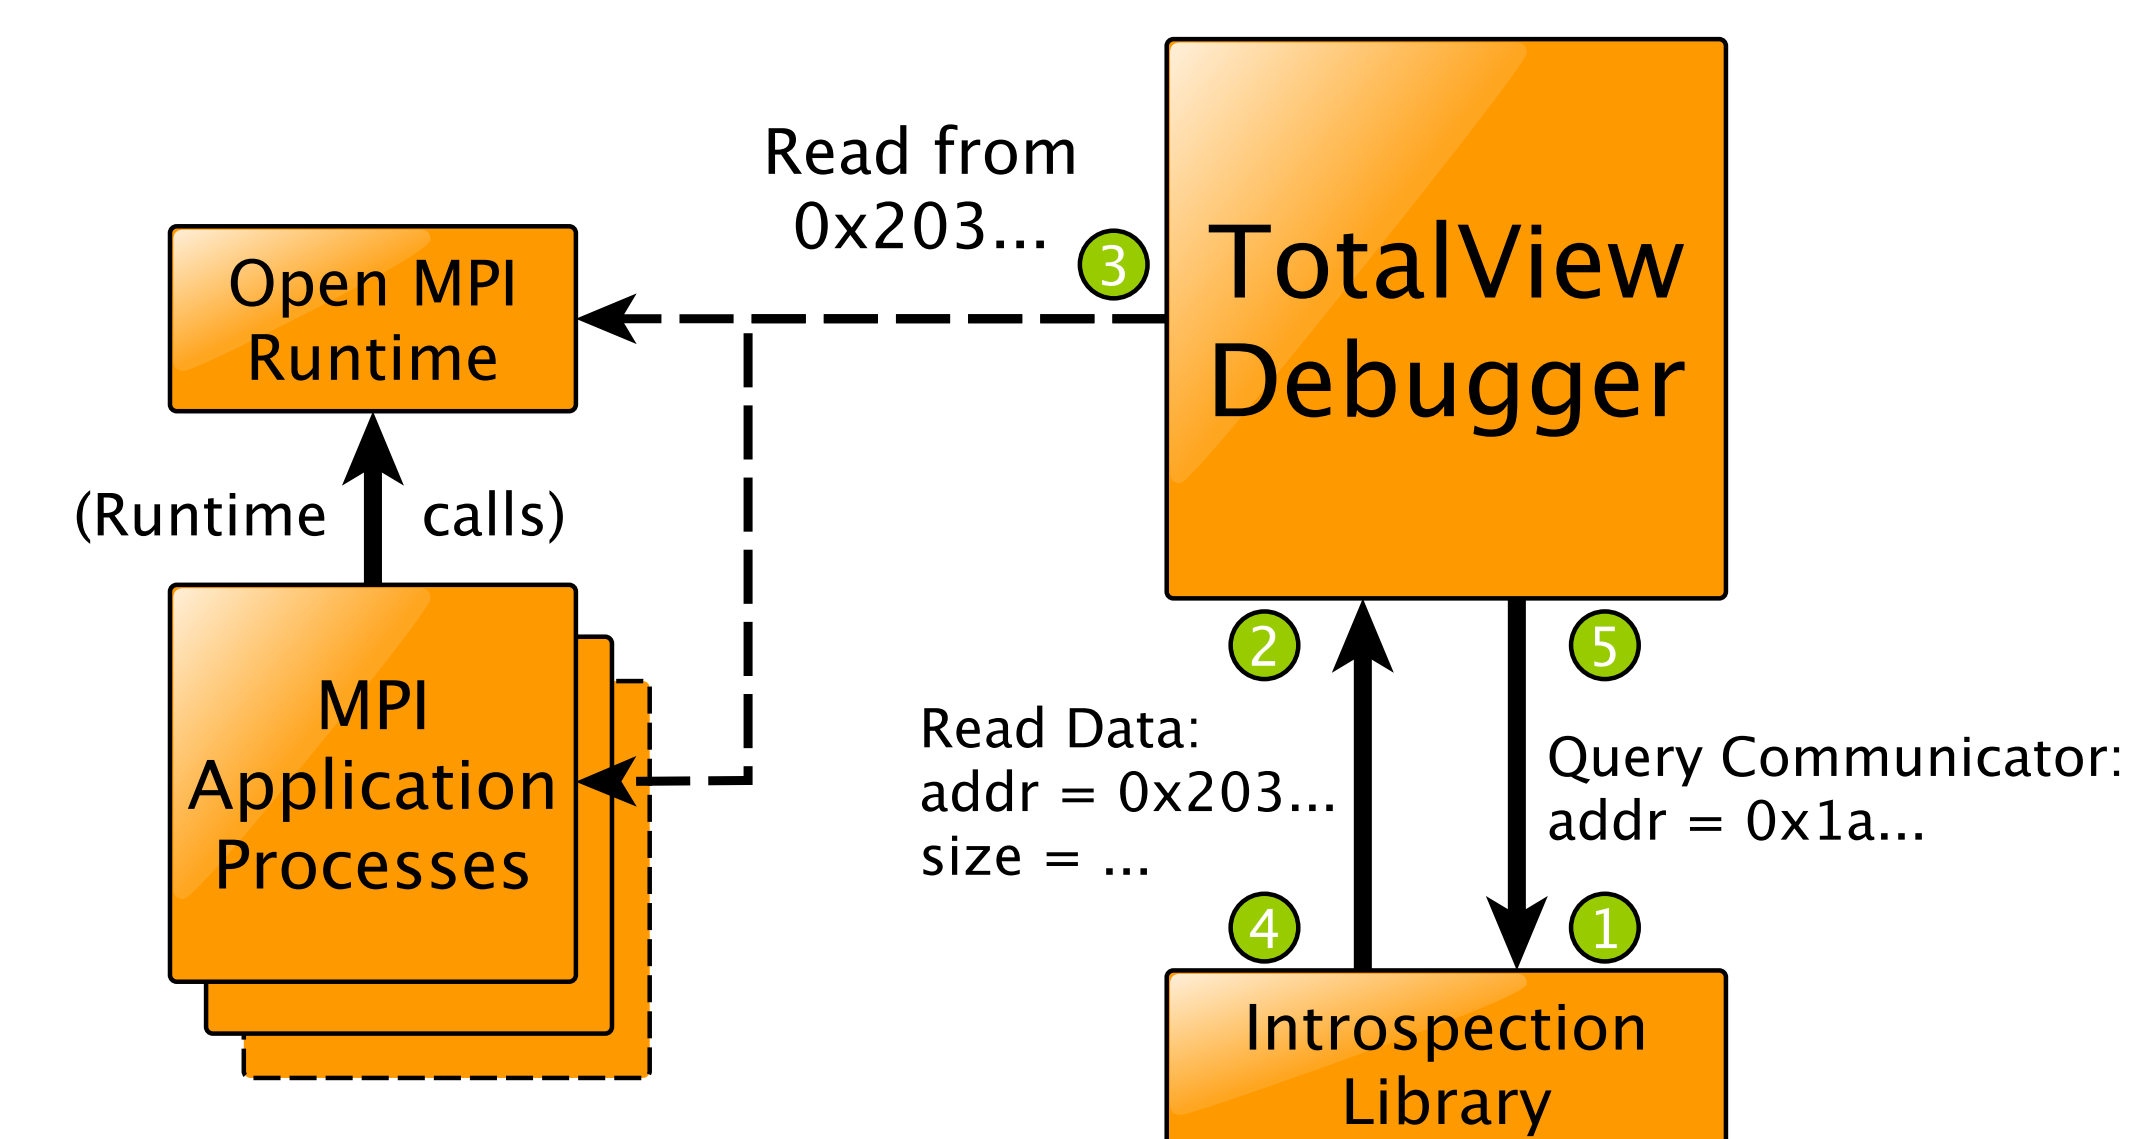


Figure 1 : Example use case to fetch communicator information.

- Debugger queries a communicator
- Library requests raw data from process
- Debugger extracts raw data from process
- Library parses raw data and returns standardized communicator information
- 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 through 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

- MPI Tools Working Group. <https://svn.mpi-forum.org/trac/mpi-forum-wiki/wiki/MPI3Tools>, October 2014.
- 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/>