The Stack Trace Analysis Tool

Enabling Million-way Debugging

Extreme-Scale Debugging Challenges

- Many control channels
- Large data volumes
- Excessive data analysis overhead
- Scalable results presentation

STAT Incrementally Debugs To Narrow Down the Search Space

- Spatial stack traces across tasks
- Temporal stack traces over time
- Traces gathered with varying level of detail
  - Function name, source line, PC

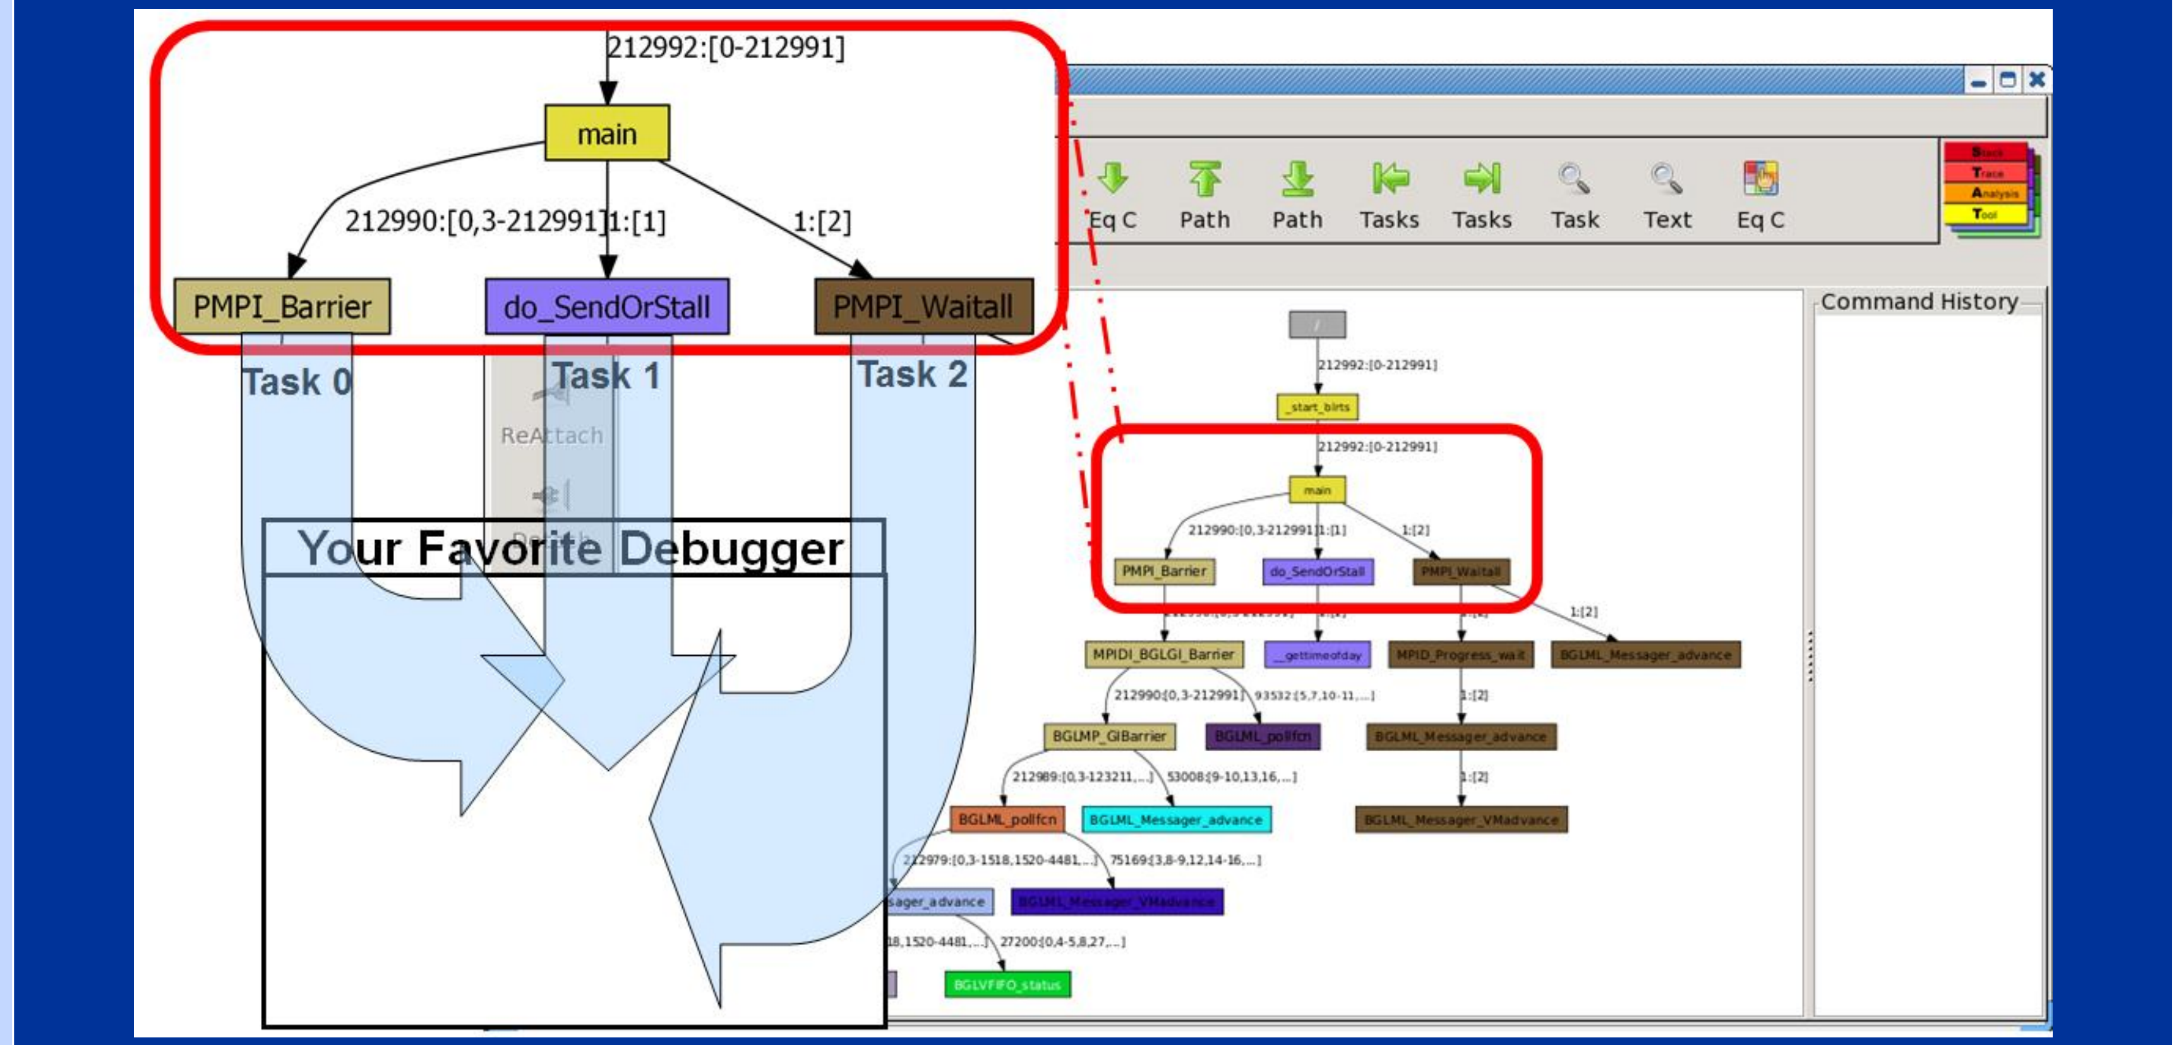

- Identifies processes with similar traces
  - Equivalence classes color encoded
  - Representative subset fed into traditional debugger for root cause analysis

MRNet and Efficient Data Structures Enable Scalable Analysis

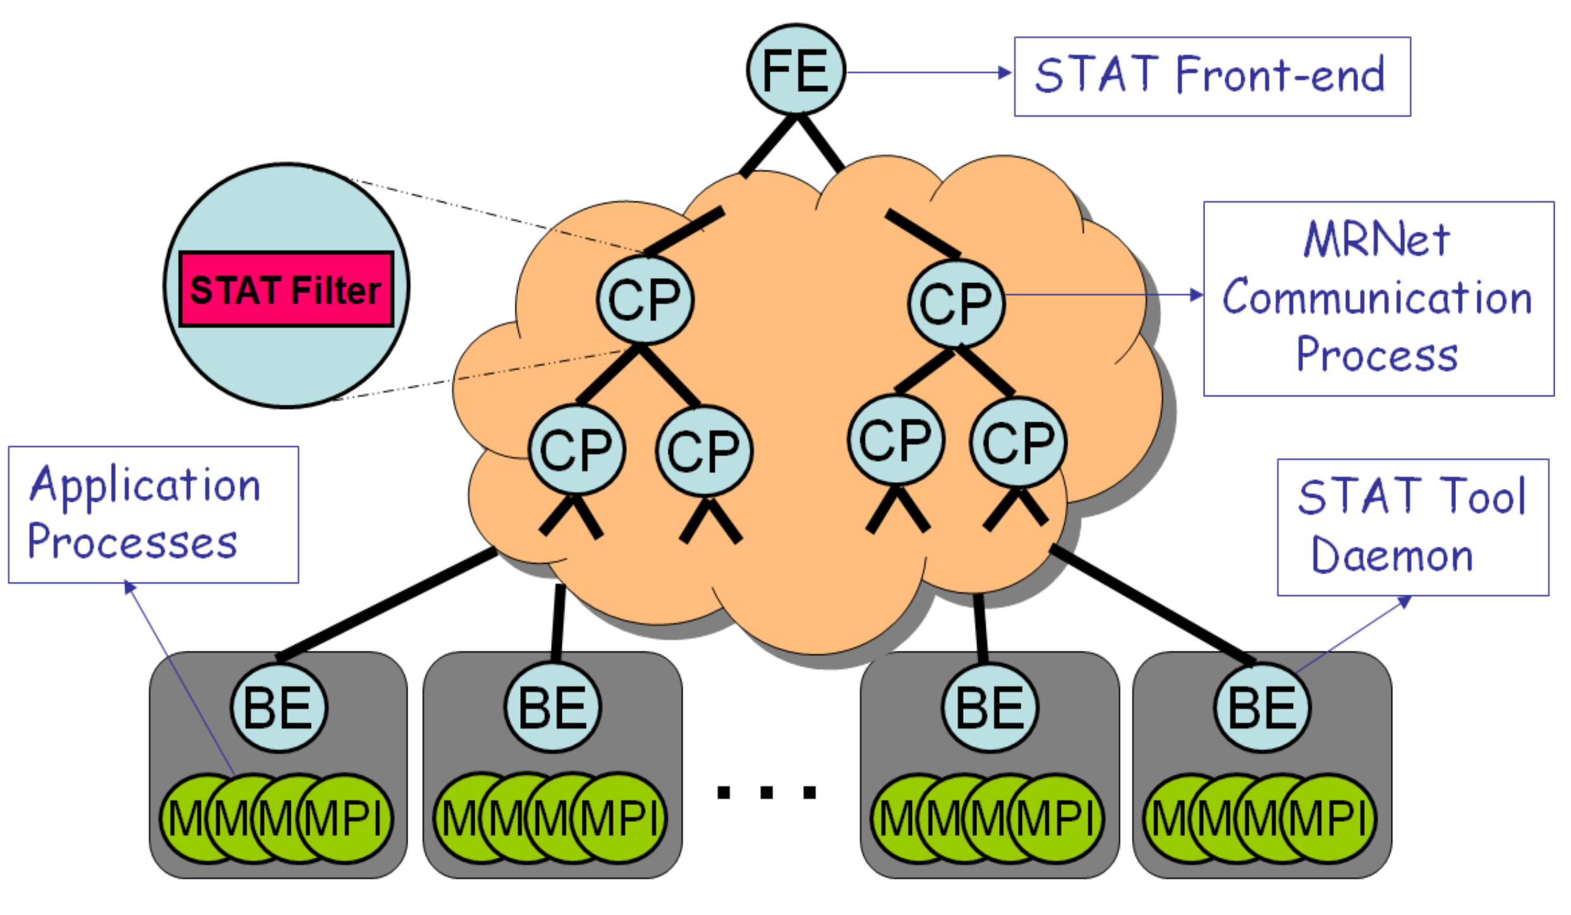

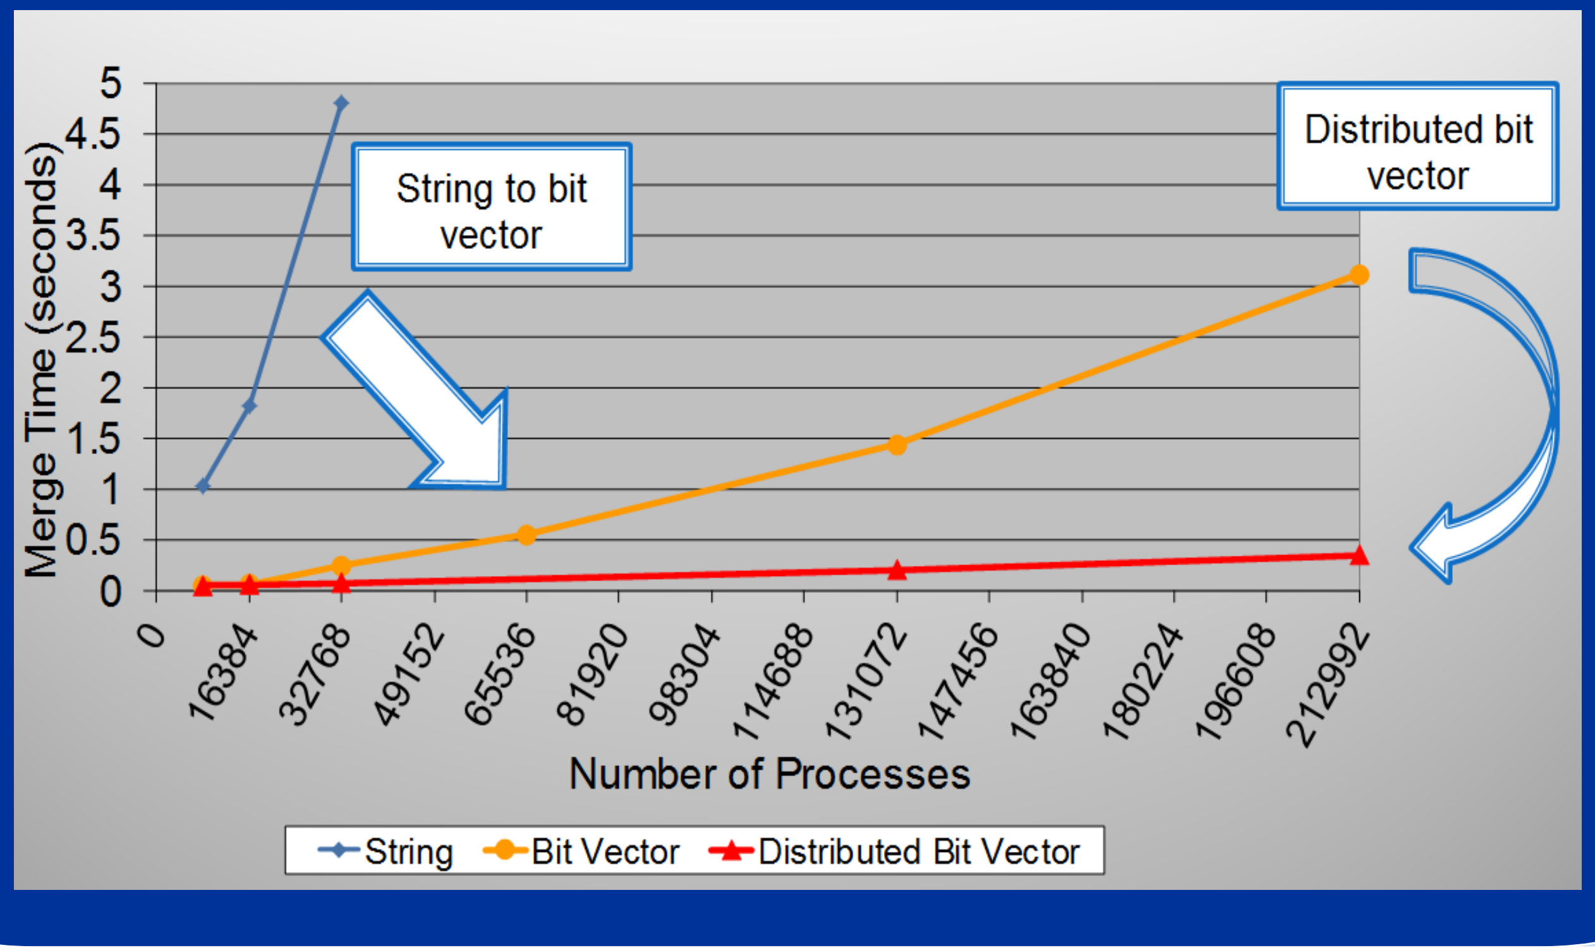

STAT Successfully Identified Bug at 1 Million MPI Tasks on Sequoia

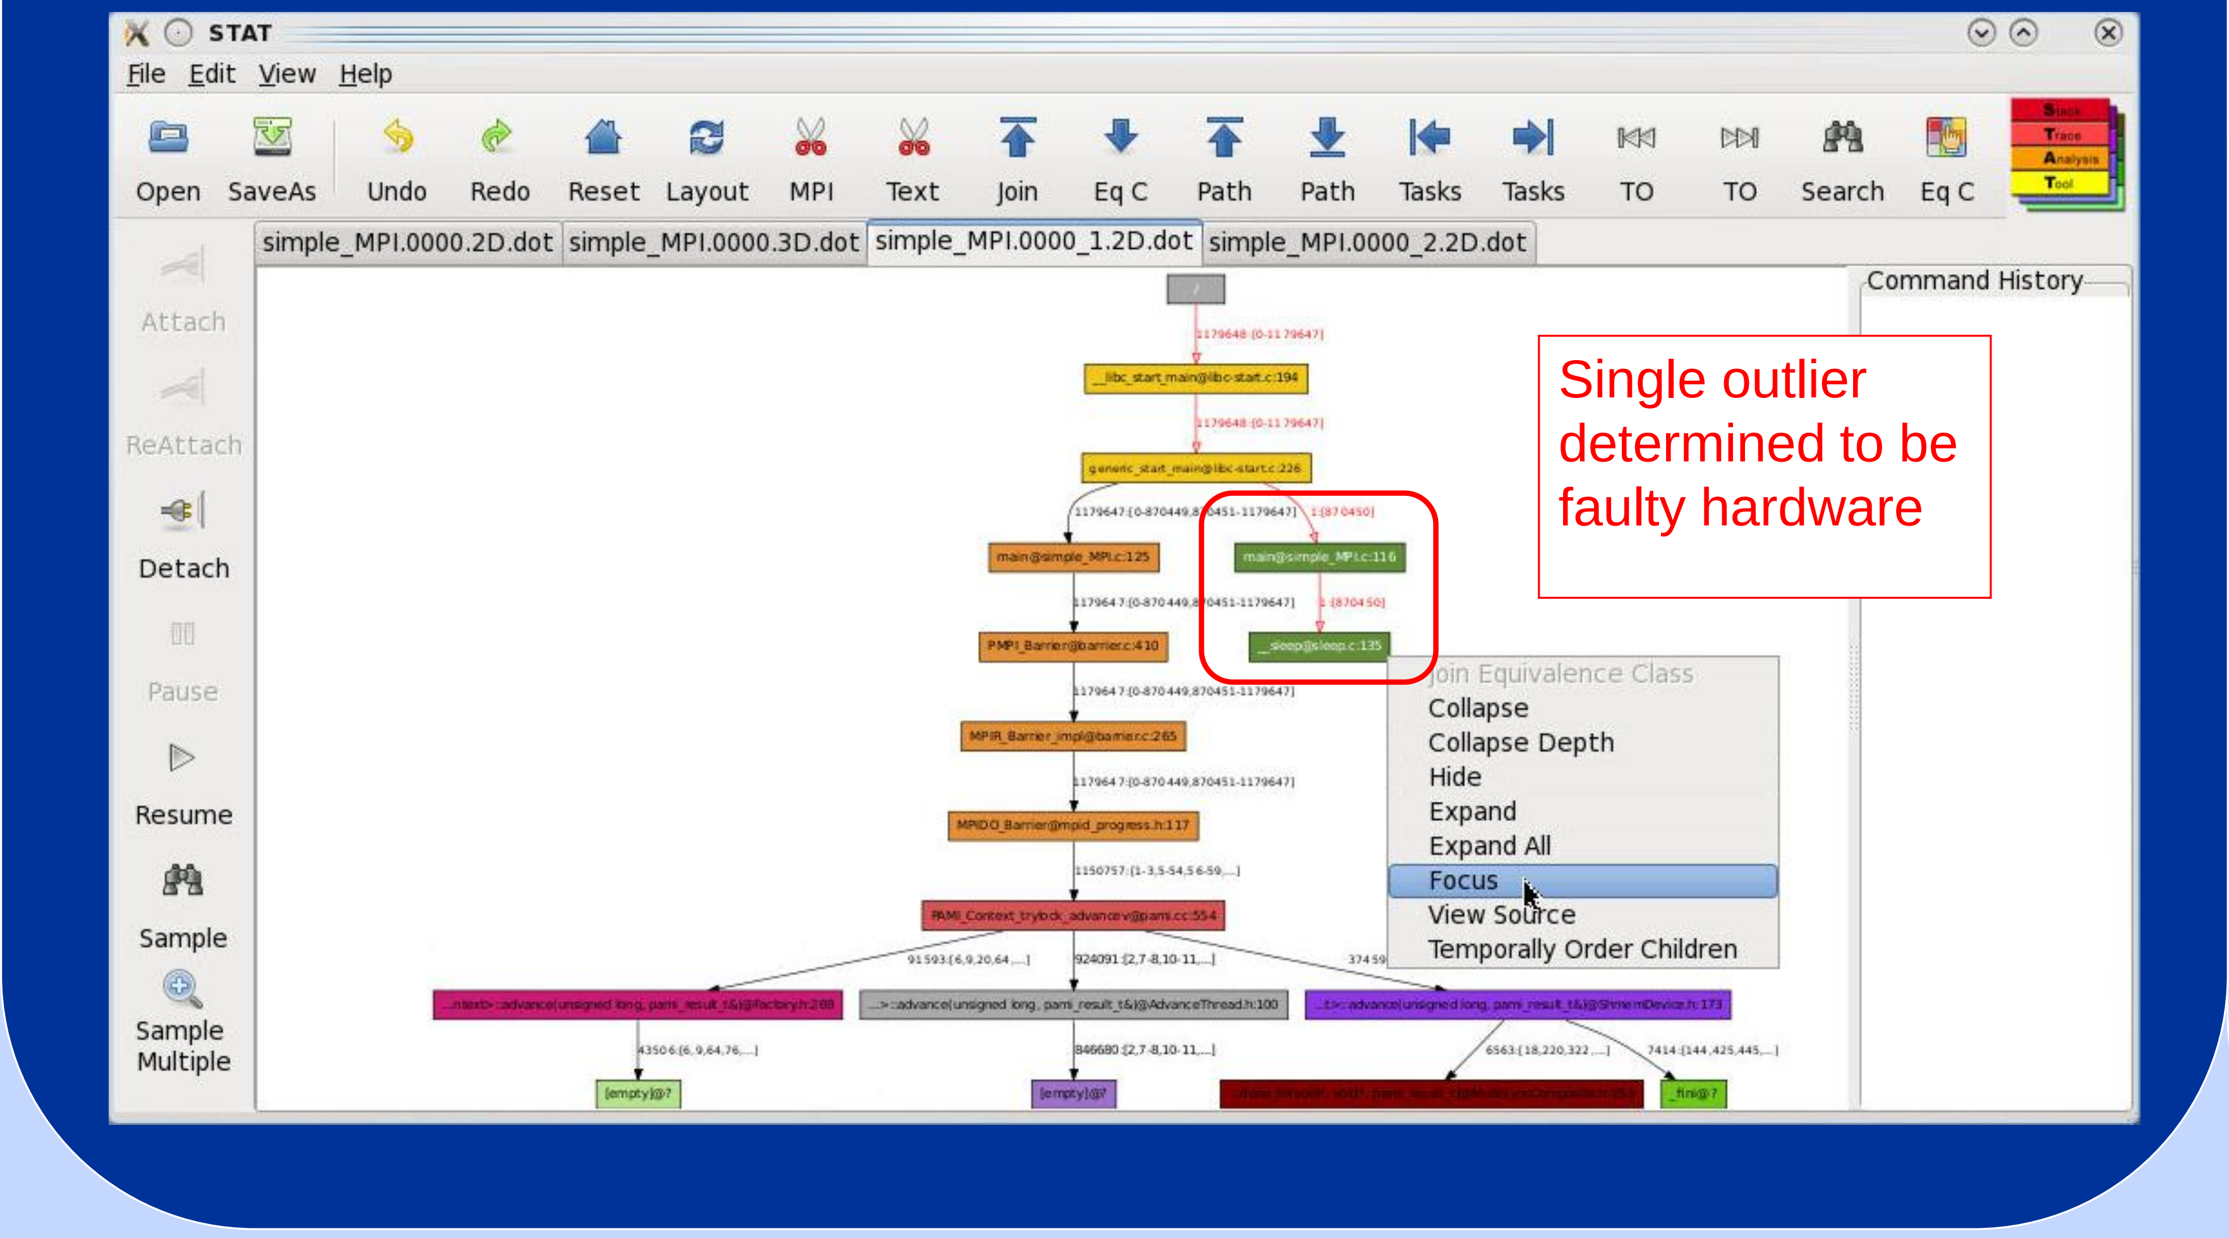

Single outlier determined to be faulty hardware

Temporal Ordering Analysis Identifies the Root Cause of Hangs

- Uses Rose compiler for static analysis of code
  - Identify loops and loop ordering variable
- Uses StackwalkerAPI for runtime information
  - Gather stack traces with source file name and line number
  - Extract program variables for loop ordering

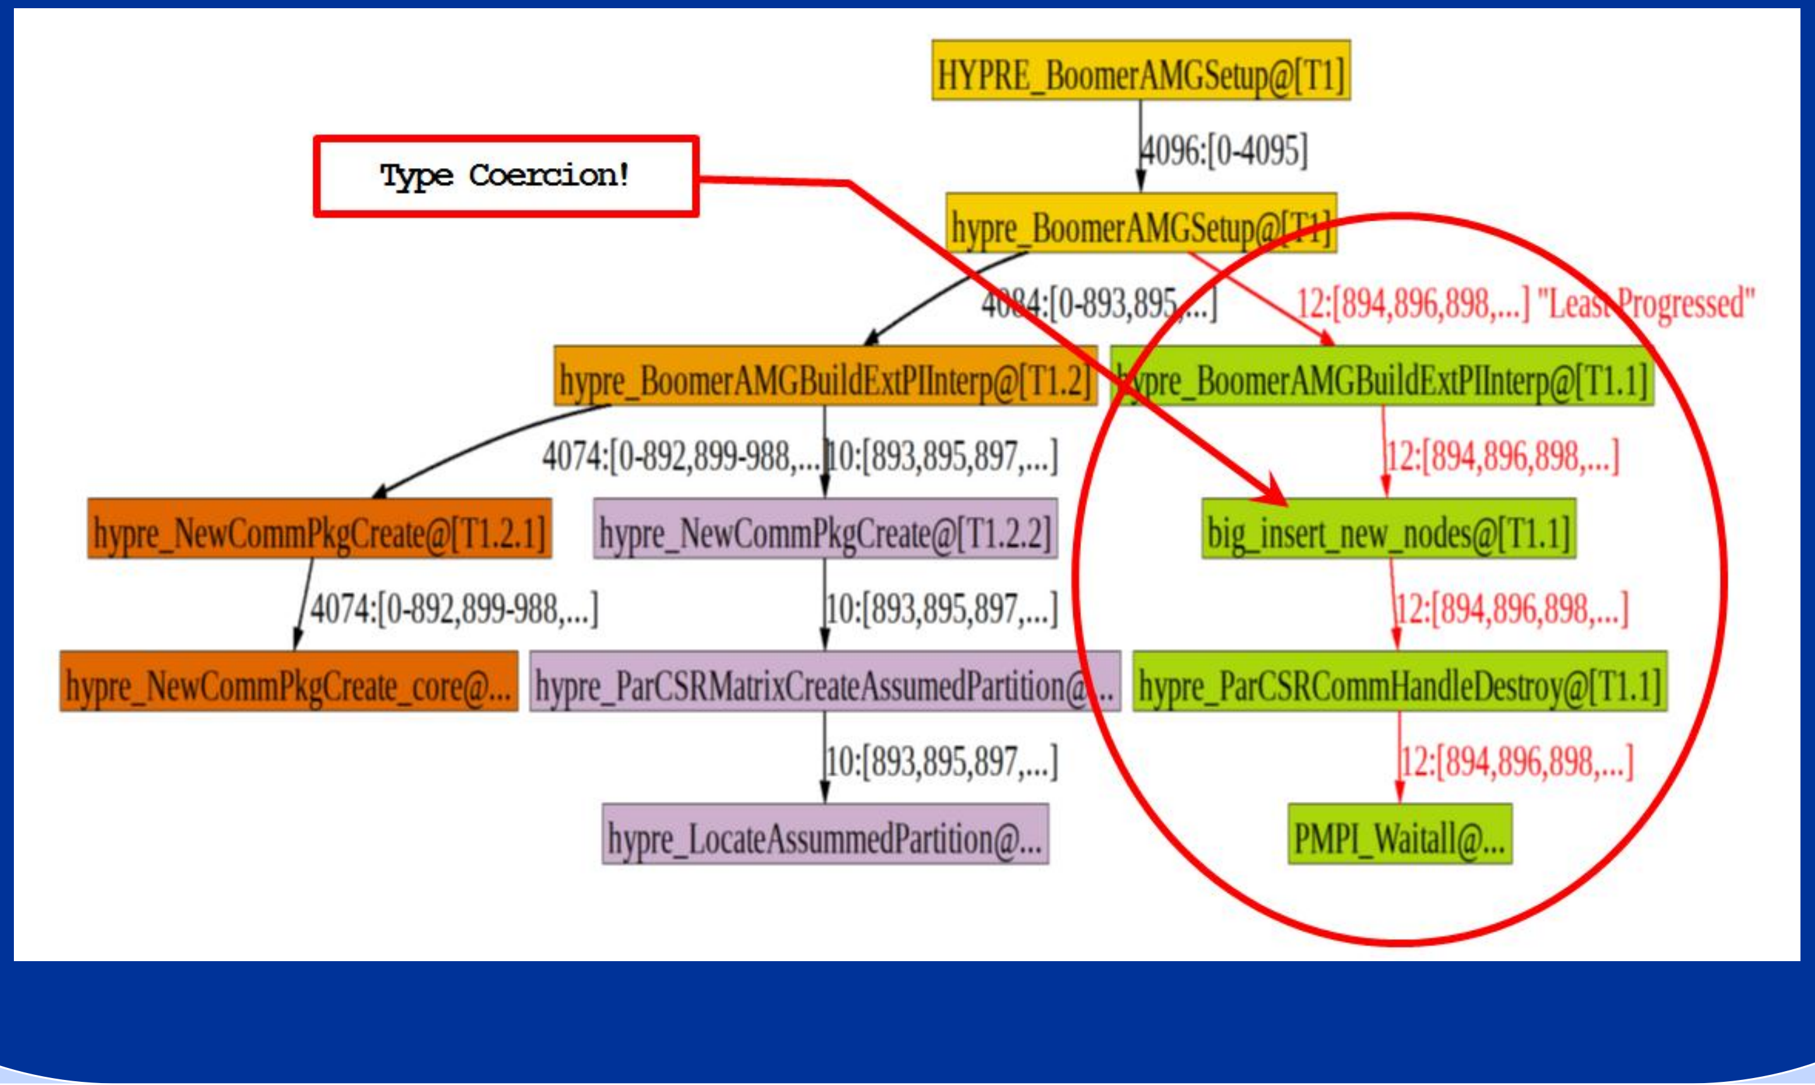
HYPRE_BoomerAMGSetup@[T1]
4096:[0-4095]
Type Coercion!
hypre_BoomerAMGSetup@[T1]
4084:[0-893,895,...]
12:[894,896,898,...] "Least Progressed"
hypre_BoomerAMGBuildExtPIInterp@[T1.2]
hypre_BoomerAMGBuildExtPIInterp@[T1.1]
4074:[0-892,899-988,...]
10:[893,895,897,...]
12:[894,896,898,...]
hypre_NewCommPkgCreate@[T1.2.1]
hypre_NewCommPkgCreate@[T1.2.2]
big_insert_new_nodes@[T1.1]
4074:[0-892,899-988,...]
10:[893,895,897,...]
12:[894,896,898,...]
hypre_NewCommPkgCreate_core@...
hypre_ParCSRMatrixCreateAssumedPartition@...
hypre_ParCSRCommHandleDestroy@[T1.1]
10:[893,895,897,...]
12:[894,896,898,...]
hypre_LocateAssummedPartition@...
PMPI_Waitall@...

Additional Information

- Runs on Linux Clusters, IBM BlueGene systems, and Cray platforms.
- http://www.paradyn.org/STAT/STAT.html
- Source available at https://outreach.scidac.gov/projects/stat/
- Arnold et al., "Stack Trace Analysis for Large Scale Debugging," *IPDPS 2007.*
- Lee et al., "Lessons Learned at 208K, Towards Debugging Millions of Cores," *SC 2008.*

Barton P. Miller
University of Wisconsin

Dorian C. Arnold
University of New Mexico

Dong H. Ahn, Bronis R. de Supinski, Gregory L. Lee, Matthew LeGendre, Martin Schulz
Lawrence Livermore National Laboratory

LLNL-POST-597172
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC