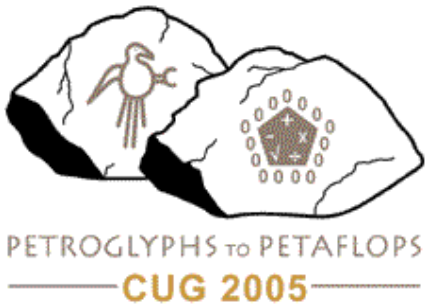

CRAY THE SUPERCOMPUTER COMPANY

ETNUS TOTALVIEW

TotalView on Cray Platforms

John V. DelSignore, Jr.
Chief Architect
Etnus

Luiz DeRose, Ph. D.
Sr. Principal Engineer
PE Tools Manager
Cray, Inc.

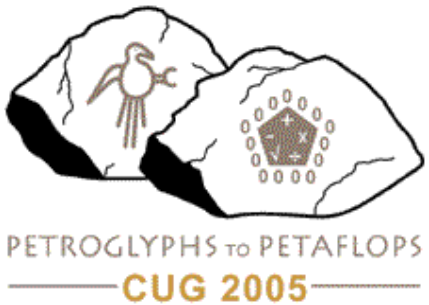

TotalView on Cray Platforms

History

- Cray has utilized TotalView technology since 1991 when Cray ported BBN TotalView to Cray TotalView
- Cray and Etnus began working to port Etnus TotalView to the X1 supercomputer starting in 2001
  - Cray ported, Etnus provided support and updates
- Cray worked with Etnus to provide a debugger on the Sandia Red Storm Machine
  - Collaboratively developed the debugger interface
  - Etnus ported, Cray and Etnus engineering worked together through the validation phase
- Cray and Etnus are discussing Rainier
- Partnership in sales, support, marketing, and engineering

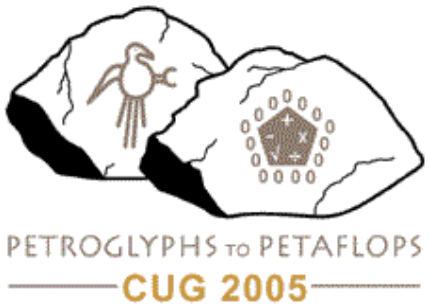

TotalView on Cray Platforms

What is TotalView?

- **Source Code Debugger for serial or parallel development**
  - C, C++, Fortran, Fortran90 supporting complex language features
  - Wide compiler and platform support
  - Multithreaded debugging, including OpenMP
  - Distributed debugging based on a cluster architecture
  - Memory debugging capabilities using Heap Interposition Agent
  - Powerful and Easy GUI and CLI extensible via scripting
  - Also on Linux, all major Unix variants, and Mac OS X

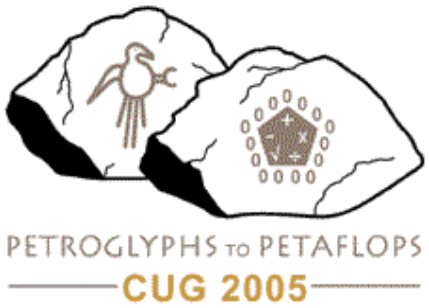

TotalView on Cray Platforms

TotalView as a Parallel Debugger

- TotalView's Architecture
- Automatic Process Acquisition
- Parallel Debugging Features
- MPI Message Queue Debugging
- Scalability

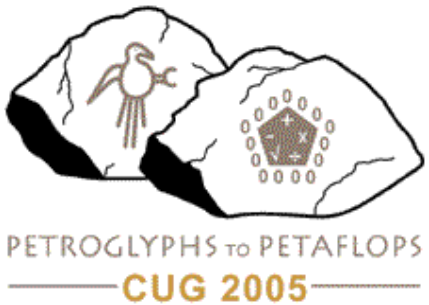

TotalView on Cray Platforms

TotalView's Architecture

- Single Client (TotalView)
  - Implements most of the functionality
  - GUI/CLI and debug engine
- Debugger Servers (tvdsvr)
  - Low overhead
  - One per node
  - Traces multiple rank processes
  - Runs as user
- TotalView communicates directly with tvdsvrs
  - Does not use MPI for communication
  - Uses specialized debugging protocol
- Provides: Robust, Scalable, Minimal Interaction

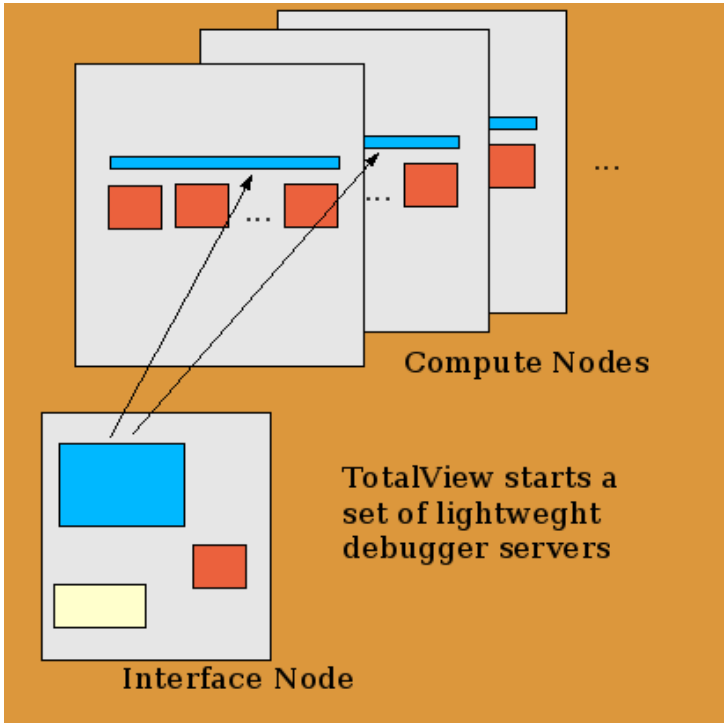

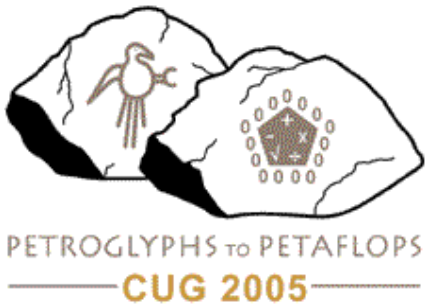

TotalView on Cray Platforms

TotalView Process Acquisition

- Seamlessly attach to all the processes in an MPI job
  - Jobs launched by TotalView, and already running or hung
- Public interface supported by most MPIs
- TotalView supports “subset” attach
  - A scalability feature for large jobs
  - Allows user to debug a subset of the MPI processes
  - Interactively attach and detach from processes
  - Users can “fan out” based on MPI communications

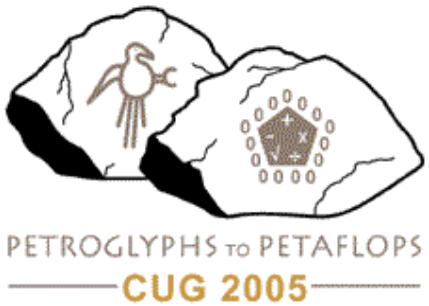

TotalView on Cray Platforms

Parallel Debugging Features

- Designed as a parallel debugger from day 1
- Root and Process windows
  - Status information, source, stack, breakpoints, threads
  - Group controls for processes and threads
- Variable and Expression List windows
  - View data structures, arrays, scalars, memory dumps
  - View SIMD data across all processes from one window
  - Parallel expression evaluation
- Rich set of action points
  - Per thread, process or group granularity
  - Barrier breakpoints
  - Evaluate C, C++, Fortran code fragments

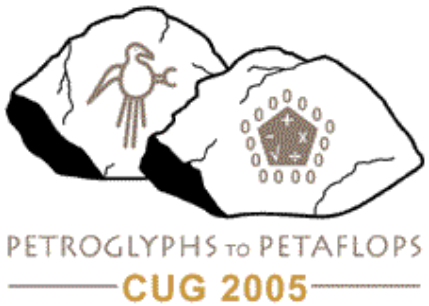

TotalView on Cray Platforms

MPI Message Queue Information

- Shows MPI communications state
- MPI programs can suffer deadlocks
  - State information in MPI library
- TotalView can expose that information
  - Quickly debug deadlocks
  - Public interface that many MPI vendors support
- Message Queue graph
  - Patterns easy to spot
- Message Queue detail windows
  - Provides access to message data

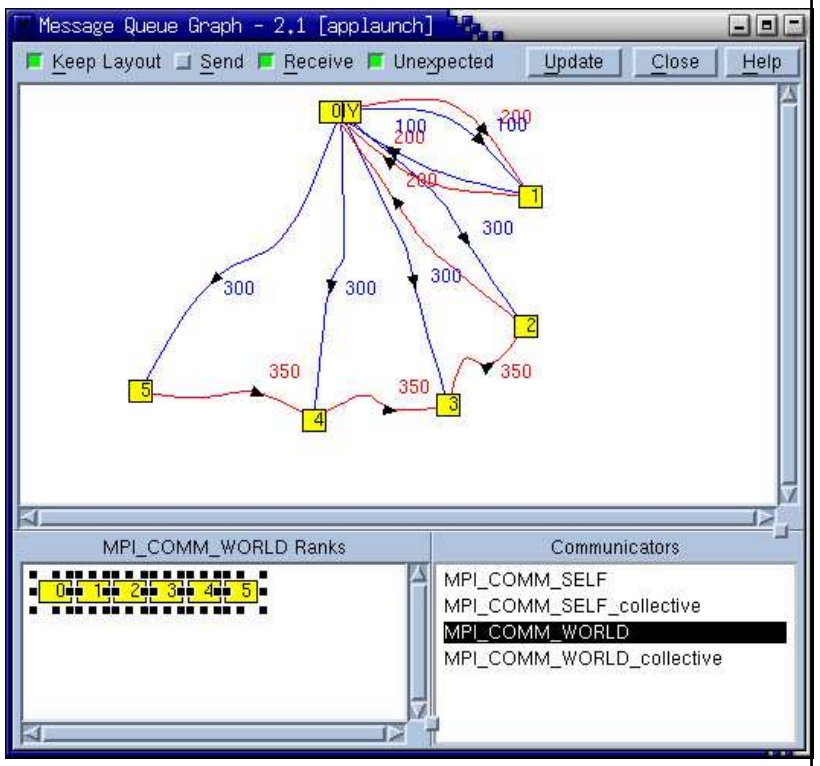

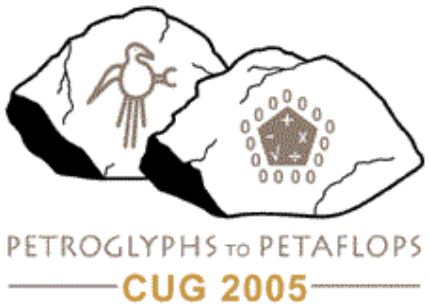

TotalView on Cray Platforms

Scalability

- Scalability means many things
  - Startup and runtime performance / responsiveness
  - Memory usage
  - Status and data representation
  - Control Issues
  - Program size/complexity also grows
- Practical scalability
  - 10s to 100s of processes trivially
  - 1,000s of processes can be debugged currently with TotalView
- Scalability is an ongoing concern

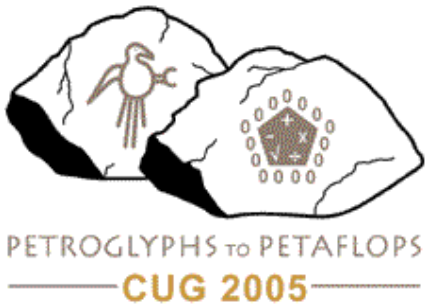

TotalView on Cray Platforms

TotalView 2005 Releases

- Two major feature releases in 2005 (June and Q4 '05)
- Four minor point releases throughout 2005
- Feature Releases
  - TotalView 7.0
    - C++ and F90 Expression system
    - MPI usability update
    - Memory Debugging Improvements
  - TotalView 7.x
    - More C++ features: exception handling, dynamic class hierarchy casting
    - MPI, UI, and memory debugging improvements

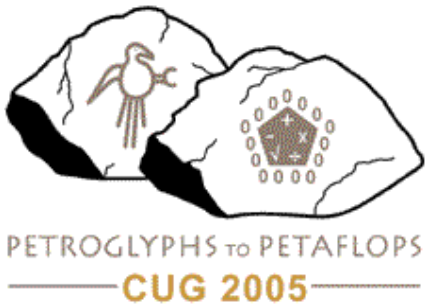

TotalView on Cray Platforms

TotalView runs on the full range of Cray Supercomputer offerings

- Entry Level and Midrange Supercomputers
  - Cray XD1 (available through Etnus or Cray)
- High End HPC Clusters Supercomputers
  - Cray XT3 (available through Etnus or Cray)
- High End HPC Scalable Vector Supercomputers
  - Cray X1 Series (available through Cray)
- First line support
  - Via Etnus: XT3 & XD1
  - Via Cray: X1 Series

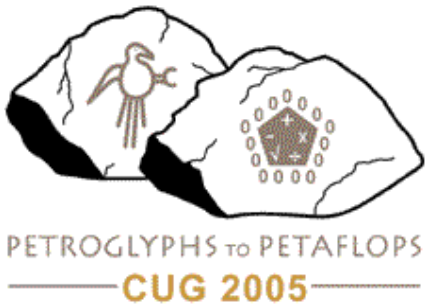

TotalView on Cray Platforms

Cray XD1 Supercomputer

- Same product offering as Etnus standard AMD/X86_64 product
- Supports Cray HPC Enhanced Linux
- Compilers
  - gcc, intel, pgi, pathscale, absoft
- Parallel paradigms
  - MPI, OpenMP

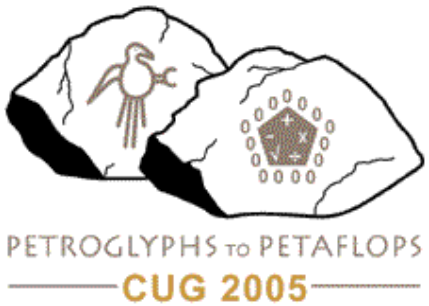

TotalView on Cray Platforms

Cray XT3 Supercomputer

- An Etnus and Cray collaborative effort
- TotalView Architecture on RedStorm
  - Utilizes the standard commodity TotalView for Linux X86_64
  - Utilizes Red Storm X64_64 UNICOS/lc environment
    - Catamount microkernel for compute nodes
    - Full Linux for Server nodes
  - PGI Compilers
  - MPICH 2 implementation
    - At the MPICH 1 level of functionality

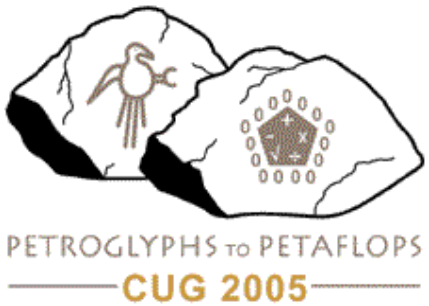

TotalView on Cray Platforms

X1 Series High End HPC Scalable Vector Systems

- Operating System:
  - UNICOS/mp
- Compilers
  - Cray Fortran
  - Cray C and C++
  - Cray Assembler
- Parallel paradigms
  - MPI
  - OpenMP
  - SHMEM
  - UPC / CAF (partial support)

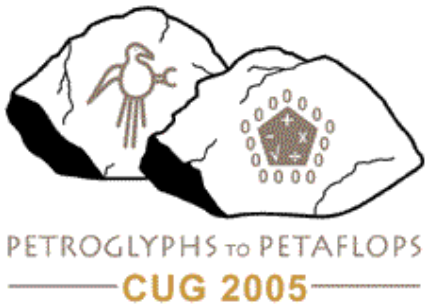

TotalView on Cray Platforms

TotalView 6.5 for the X1 Series

- Released May 10, 2005
- New features:
  - Support for live pthreads and OpenMP
  - All vector registers, vector masks, and bit-matrix multiply registers are now accessible from within a TotalView debugging session
  - Improvement of browse and menu functionality
  - Improvement on the application launching interface

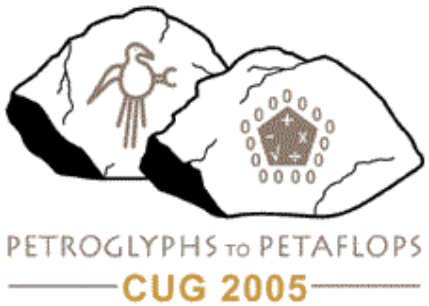

TotalView on Cray Platforms

Cray and Etnus Release Engineering

- Cray XT3 and XD1 are part of the standard TotalView release cycle
- Cray gets a new source drop from Etnus on every release
- Cray X1 Series releases are approximately two releases behind the latest TotalView product releases
- Backline support to Cray via Etnus Support and Engineering

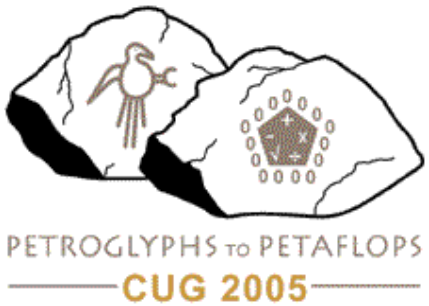

TotalView on Cray Platforms

On-Going Developments

- Etnus is working on a new visualizer technology that can be made available on all Cray Supercomputers
  - Based on VTK/OpenGL
- Etnus and Cray are in business discussions for Rainier
- Finalizing Sandia Red Storm validation
  - TotalView support on the XT3 at scale
- Looking into feasibility of memory debugging on XT3
- Continued inter-company communication and customer support

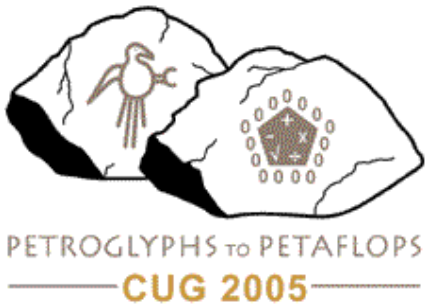

TotalView on Cray Platforms

Future: TotalView on Rainier

- Etnus and Cray have had several technical exchanges
  - Three components:
    - Back end running on the Scalar nodes (Mt. Adams)
    - Back end running on the Vector nodes (Black Widow)
    - Front end running on the Service node:
      - Handles the scheduler interface
      - Triggers TotalView on the different nodes
      - Handles the heterogeneity of the service node and the vector node
      - Supports multi-binary launch
- Etnus and Cray are working on the business arrangement
  - Similar to the model on the X1 Series
  - Single license for the system, independently of the type of nodes

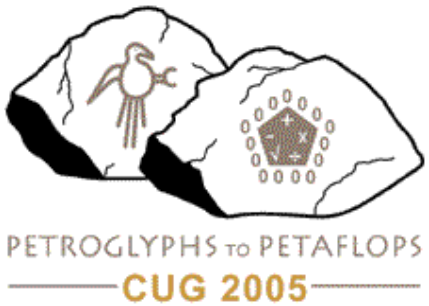

TotalView on Cray Platforms

Questions and Discussion

For more information:

- www.etnus.com
- docs.cray.com (Cray X1 Series)