

Parallel Debugging

Matthias Müller, Pavel Neytchev, Rainer Keller, Bettina Krammer

University of Stuttgart
High-Performance Computing-Center Stuttgart (HLRS)
www.hlrs.de



Parallel Debugging Müller, Neytchev, Keller, Krammer
Hochleistungsrechenzentrum Stuttgart

H L R I S 

Outline

- Motivation
- Approaches and Tools
 - Memory Tracing Tools
 - Valgrind
 - MPI-Analysis Tools
 - Marmot
 - Debuggers
 - TotalView
 - DDT



Parallel Debugging Müller, Neytchev, Keller, Krammer
Slide 2 / 60 Hochleistungsrechenzentrum Stuttgart

H L R I S 

Motivation - Problems of Parallel Programming

- All problems of serial programming
- Additional problems:
 - Increased difficulty to verify correctness of program
 - Increased difficulty to debug N parallel processes
 - New parallel problems:
 - **deadlocks**
 - **race conditions**
 - **irreproducibility**



Tools and Techniques to Avoid and Remove Bugs

- Programming techniques
- Static Code analysis
 - Compiler (with `-Wall` flag or similar), lint
- Runtime analysis
 - Memory tracing tools
 - Special OpenMP tools (assure, thread checker)
 - Special MPI tools
- Post mortem analysis
 - Debuggers



skipped

What is a Debugger?

- Common Misconception:
A debugger is a tool to find and remove bugs
- A debugger does:
 - tell you where the program crashed
 - help to gain a better understanding of the program and what is going on
- Consequence:
 - A debugger does not help much if your program does not crash, e.g. just gives wrong results
 - Avoid using a debugger as far as possible.
 - Use it as last resort.



Programming Techniques

- Think about a verbose execution mode of your program
- Use a careful/paranoid programming style
 - check invariants and pre-requisites
(`assert(m>=0)`, `assert(v<c)`)



Static Code Analysis – Compiler Flags

- Use the debugging/assertion techniques of the compiler
 - use debug flags (-g), warnings (-Wall)
 - array bound checks in Fortran
 - use memory debug libraries (-lefence)



skipped

Avoiding Debuggers

- Write portable programs
 - it avoids future problems
 - **architectures/platforms have a short life**
 - **all compilers and libraries have bugs**
 - **all languages and standards include implementation defined behavior**
 - running on different platforms and architectures significantly increases the reliability
- Use verification tools for parallel programming like assure



Valgrind – Debugging Tool

Rainer Keller

University of Stuttgart
High-Performance Computing-Center Stuttgart (HLRS)
<http://www.hlrs.de>



Parallel Debugging Müller, Neytchev, Keller, Krammer
Hochleistungsrechenzentrum Stuttgart

H L R I S 

Valgrind – Overview

- An Open-Source Debugging & Profiling tool.
- Works with any dynamically linked application.
- Emulates CPU, i.e. executes instructions on a synthetic x86.
- Currently it's only available for Linux/IA32.
- Prevents Error-swamping by suppression-files.
- Has been used on many **large** Projects:
KDE, Emacs, Gnome, Mozilla, OpenOffice.
- It's easily configurable to ease debugging & profiling through *skins*:
 - **Memcheck**: Complete Checking (every memory access)
 - Addrcheck: 2xFaster (no uninitialized memory check).
 - Cachegrind: A memory & cache profiler
 - **Callgrind**: A Cache & Call-tree profiler.
 - Helgrind: Find Races in multithreaded programs.



Parallel Debugging Müller, Neytchev, Keller, Krammer
Slide 10 / 60 Hochleistungsrechenzentrum Stuttgart

H L R I S 

Valgrind – Usage

- Programs should be compiled with
 - Debugging support (to get position of bug in code)
 - Possibly without Optimization (for accuracy of position & less false positives):

```
gcc -O0 -g -o test test.c
```
- Run the application as normal as parameter to valgrind:

```
valgrind ./test
```
- Then start the MPI-Application as with TV as debugger:

```
mpirun -dbg=valgrind ./mpi_test
```



Valgrind – Memcheck

- Checks for:
 - Use of uninitialized memory
 - Malloc Errors:
 - **Usage of free'd memory**
 - **Double free**
 - **Reading/writing past malloced memory**
 - **Lost memory pointers**
 - **Mismatched malloc/new & free/delete**
 - Stack write errors
 - Overlapping arguments to system functions like `memcpy`.



Valgrind – Example 1/2

```
array = malloc (SIZE * sizeof (int));

if (rank == 0)
    MPI_Recv (array, SIZE+1, MPI_INT, 1, 4711, MPI_CO
else
    MPI_Send (array, SIZE+1, MPI_INT, 0, 4711, MPI_CO
printf ("(Rank:%d) array[0]:%d\n", rank, array[0]);
```



Valgrind – Example 2/2

- PID • With Valgrind `mpirun -dbg=valgrind -np 2 ./mpi_murks:`

```
==11278== Invalid read of size 1
==11278== at 0x4002321E: memcpy (.../memcheck/mac_replace_strmem.c:256)
==11278== by 0x80690F6: MPID_SHMEM_Eagerb_send_short (mpich/./shmshort.c:70)
.. 2 lines of calls to MPICH-functions deleted ...
==11278== by 0x80492BA: MPI_Send (/usr/src/mpich/src/pt2pt/send.c:91)
==11278== by 0x8048F28: main (mpi_murks.c:44)
==11278== Address 0x4158B0EF is 3 bytes after a block of size 40 alloc'd
==11278== at 0x4002BBCE: malloc (.../coregrind/vg_replace_malloc.c:160)
==11278== by 0x8048EB0: main (mpi_murks.c:39)
```

← Buffer-Overrun by 4 Bytes in MPI_Send

```
....
==11278== Conditional jump or move depends on uninitialised value(s)
==11278== at 0x402985C4: _IO_vfprintf_internal (in /lib/libc-2.3.2.so)
==11278== by 0x402A15BD: _IO_printf (in /lib/libc-2.3.2.so)
==11278== by 0x8048F44: main (mpi_murks.c:46)
```

← Printing of uninitialized variable

- It can not find:
 - May be run with 1 process: One pending Recv (Marmot).
 - May be run with >2 processes: Unmatched Sends (Marmot).



skipped

Valgrind – Calltree 1/2

- The Calltree skin (like the cachegrind skin):
 - Tracks memory accesses to check Cache-hit/misses.
 - Additionally records call-tree information.
- ```

==11745== Calltree-0.9.6, a cache profiler for x86-linux.
==11745== Copyright (C) 2002, and GNU GPL'd, by N.Nethercote and J.Weider
==11745== Using valgrind-2.1.0, a program supervision framework for x86-]
==11745== Copyright (C) 2000-2003, and GNU GPL'd, by Julian Seward.
--11745-- warning: Pentium with 12 K micro-op instruction trace cache
--11745-- Simulating a 16 KB cache with 32 B lines
==11745== Estimated CPU clock rate is 1410 MHz
==11745== For more details, rerun with: -v
--11745--

```
- After the run, it reports overall program statistics:
 

```

==11810== D refs: 497.790.574 (386.176.612 rd + 111.613.962 wr)
==11810== D1 misses: 863.493 (369.495 rd + 493.998 wr)
==11810== L2d misses: 282.232 (98.857 rd + 183.375 wr)
==11810== D1 miss rate: 0.1% (0.0% + 0.4%)
==11810== L2d miss rate: 0.0% (0.0% + 0.1%)

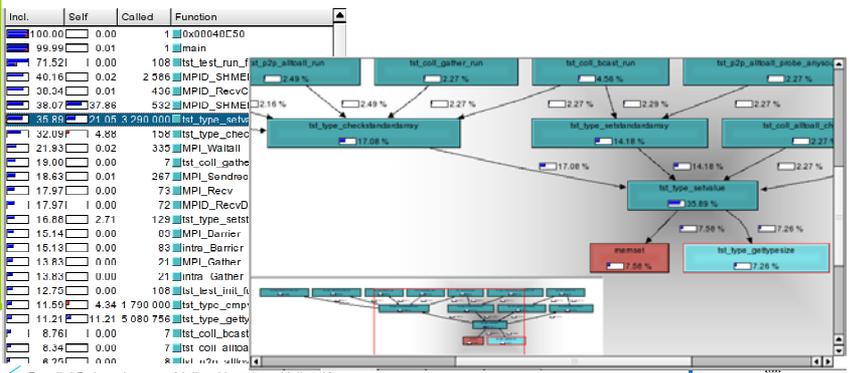
```



skipped

## Valgrind – Calltree 2/2

- Even more interesting: the output trace-file.
- With the help of kcachegrind, one may:
  - Investigate, where Instr/L1/L2-cache misses happened2-2.
  - Which functions were called where & how often.



## Valgrind – Deficiencies

- Valgrind cannot find of these Error-Classes:
  - Semantic Errors
  - Timing-critical errors
  - Uninitialised stack-memory not detected.
  - Problems with new instruction sets  
(e.g. SSE/SSE2 is supported, certain Opcodes are **not**).  
When using the Intel-Compiler: `-tpp5` for Pentium optimisation.



## MARMOT

Bettina Krammer



University of Stuttgart  
High-Performance Computing-Center Stuttgart (HLRS)  
[www.hlrs.de](http://www.hlrs.de)



## What is MARMOT?

- MPI analysis and checking tool to verify at runtime if an application conforms to the MPI standard.
- Library written in C++ that will be linked to the application.
- No source code modification of the application is required.
- Additional process working as debug server, i.e. the application will have to be run with `mpirun` for  $n+1$  instead of  $n$  processes.
- Implementation of C and Fortran language binding of MPI-1.2 standard.
- Environment variables for tool behaviour and output (report of errors, warnings and/or remarks, trace-back, etc.).
- After the execution of the program the user can read a logfile to check for potential problems.



## Availability of MARMOT

- Tests on different platforms, using different compilers and MPI implementations, e.g.
  - IA32/IA64 clusters (Intel, g++ compiler) mpich
  - IBM Regatta
  - NEC SX5
  - Hitachi SR8000
- Download and further information  
<http://www.hlrs.de/organization/tsc/projects/marmot/>



## Example 1: request-reuse (source code)

```
/*
** Here we re-use a request we didn't free before
*/

#include <stdio.h>
#include <assert.h>
#include "mpi.h"

int main(int argc, char **argv) {
 int size = -1;
 int rank = -1;
 int value = -1;
 int value2 = -1;
 MPI_Status send_status, recv_status;
 MPI_Request send_request, recv_request;

 printf("We call Irecv and Isend with non-freed requests.\n");
 MPI_Init(&argc, &argv);
 MPI_Comm_size(MPI_COMM_WORLD, &size);
 MPI_Comm_rank(MPI_COMM_WORLD, &rank);
 printf(" I am rank %d of %d PEs\n", rank, size);
```



## Example 1: request-reuse (source code cont'd)

```
if(rank == 0){
 /** this is just to get the request used **/
 MPI_Irecv(&value, 1, MPI_INT, 1, 18, MPI_COMM_WORLD, &recv_request);
 /** going to receive the message and reuse a non-freed request **/
 MPI_Irecv(&value, 1, MPI_INT, 1, 17, MPI_COMM_WORLD, &recv_request);
 MPI_Wait(&recv_request, &recv_status);
 assert(value = 19);
}
if(rank == 1){
 value2 = 19;
 /** this is just to use the request **/
 MPI_Isend(&value, 1, MPI_INT, 0, 18, MPI_COMM_WORLD, &send_request);
 /** going to send the message **/
 MPI_Isend(&value2, 1, MPI_INT, 0, 17, MPI_COMM_WORLD, &send_request);
 MPI_Wait(&send_request, &send_status);
}
MPI_Finalize();
return 0;
}
```



## Example 1: request-reuse (output log)

```
We call Irecv and Isend with non-freed requests.
1 rank 0 performs MPI_Init
2 rank 1 performs MPI_Init
3 rank 0 performs MPI_Comm_size
4 rank 1 performs MPI_Comm_size
5 rank 0 performs MPI_Comm_rank
6 rank 1 performs MPI_Comm_rank
 I am rank 0 of 2 PEs
7 rank 0 performs MPI_Irecv
 I am rank 1 of 2 PEs
8 rank 1 performs MPI_Isend
9 rank 0 performs MPI_Irecv
10 rank 1 performs MPI_Isend
 ERROR: MPI_Irecv Request is still in use !!
11 rank 0 performs MPI_Wait
 ERROR: MPI_Isend Request is still in use !!
12 rank 1 performs MPI_Wait
13 rank 0 performs MPI_Finalize
14 rank 1 performs MPI_Finalize
```



## Parallel Debuggers



## Parallel Debuggers

- Most vendor debuggers have some support
- gdb has basic support for threads
- Debugging MPI programs with a “scalar” debugger is hard but possible
  - MPiCh supports debugging with gdb attached to one process
  - manual attaching to the processes is possible
- TotalView is a good but expensive tool
- DDT is an alternative



## TOTALVIEW



## What is TotalView?

- Parallel debugger
- Source level debugging for C, C++, F77, F90, HPF
- **MPI, OpenMP**, Pthreads, PVM, shmem
- SMPs, MPPs, PVPs, Clusters
- Available on all major Unix Platforms and most Supercomputers
- GUI (independent of platform, exception Cray T3E)
  - TotalView 4.x based on tcl/tk
  - TotalView 5.x based on Motif



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 27 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

## Availability of TotalView

- [Compaq Digital Alpha](#)
- [HP-UX](#)
- [IBM RS6000 and SP Power](#)
- [SGI MIPS](#)
- [Sun SPARC SunOS 5](#)
- [Linux Intel IA32 \(RedHat\)](#)
- [Linux Alpha \(RedHat\)](#)
- Linux IA64
- Linux Opteron
- Cray T3E by Cray
- Hitachi SR2201 by SofTek, SR8000
- NEC SX Series



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 28 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

## Availability of TotalView at HWW

| Platform       | Availability | Remarks    |
|----------------|--------------|------------|
| Volvox         | Yes          | V6         |
| Hitachi SR8000 | Yes          | V 4.0      |
| Cray T3E       | Yes          | Cray 3.0.0 |
| NEC SX         | Yes          |            |
| SGI Onyx       | No           | Use cvd    |

Development  
Platforms

More information:

<http://www.hlrs.de/organization/tsc/services/tools/debugger/totalview>



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 29 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

## Availability of TotalView at University Stuttgart

- Two user 8 CPU Floating License for University Stuttgart:
  1. Download Software from <http://www.etnus.com>
  2. Set environment variable for license:  
`LM_LICENSE_FILE=7244@servint1.rus.uni-stuttgart.de`

More information about campus licenses available at  
<http://www.hlrs.de/organization/tsc/services/tools/campus>



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 30 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

## TotalView usage at HLRS

- Set USE\_TOTALVIEW in your login scripts
- CRAY T3E: set USE\_PROG\_ENV
- Compile with -g compiler switch  
CRAY T3E: compiler switch -G
- command name: **totalview**



## Starting TotalView

On a new process:

```
% totalview myprog -a arguments to myprog
```

To debug MPI programs:

```
% totalview mpirun -a -nprocs 3 myprog
```

```
% mpirun -tv -np 3 myprog
```

To debug IBM POE programs:

```
% totalview poe -a myprog [args]
```

To debug CRAY T3E programs:

```
% totalview -X #procs myprog [args]
```



skipped

## TotalView on Hitachi SR8000

- Compilation:
  - f90 -g
  - cc -g
  - KCC -g --backend -tv
- OpenMP
  - f90 -g -omp -procnum=8
  - cc -g -omp -parallel=1 -O2
- MPI
  - mpirun -tv



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 33 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

skipped

## TotalView on HPN

- Compilation:
  - f90 -g
  - cc -g
  - KCC -g
- OpenMP
  - guidef90 -g
  - guidec -g
  - guidec++ -g
- MPI
  - mpirun -np #procs -tv ./a.out



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 34 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

skipped

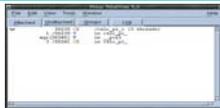
## TotalView Exercise: Basic Look & Feel

- Log into hwwhpn.hww.de
- Use bash as shell
- Change into directory  
~/TOTALVIEW/#NR/TOTALVIEW/SIMPLE
- Compile `calc_pi_{f90,c,cc}.{f90,c,cc}`
- Start totalview with `totalview executable`



## TotalView Windows

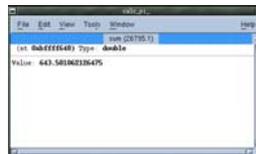
Root Window



Process Window

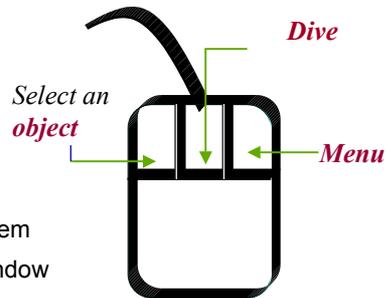


Data Windows



## TotalView Mouse Buttons

- **Left button is Select:**
  - Chooses an item of interest,
  - Starts editing a item
- **Middle button is Dive:**
  - Gets more information about an item
  - **Shift+Dive** forces open a new window
- **Right button is Menu:**
  - Raises a menu of actions
  - All menus have a **Help (^?)** entry



## TotalView Main Window

A screenshot of the TotalView Main Window showing a list of processes and threads. The window title is "TotalView 3.7.7". The list is expanded to show the following data:

| Process ID | Process/thread status | Process name      | Number of threads |
|------------|-----------------------|-------------------|-------------------|
| 3766       | B4                    | ./txsort_t        | (16 threads)      |
| 1/25257    | K                     | at 0x000000       |                   |
| 2/16059    | B4                    | in .sort          |                   |
| 3/15793    | T                     | in .txfork        |                   |
| 4/16823    | K                     | at 0x000000       |                   |
| 5/17069    | K                     | at 0x000000       |                   |
| 6/20403    | K                     | at 0x000000       |                   |
| 7/20927    | T                     | in ._pthread_body |                   |
| 8/21685    | K                     | at 0x000000       |                   |
| 9/22191    | K                     | at 0x000000       |                   |

Annotations in the image include:

- "Expand list" pointing to the expand/collapse icon.
- "Thread list tid/systid" pointing to the first column of the thread list.
- "Process ID" pointing to the first column of the process list.
- "Process/thread status" pointing to the second column of the process list.
- "Process name" pointing to the third column of the process list.
- "Number of threads" pointing to the fourth column of the process list.
- "Function or PC value" pointing to the function names in the thread list.



## TotalView Process Window

Stack Trace pane

Source pane

Thread pane

Process/thread motion buttons

Local variables for the selected frame

Action Points pane



## TotalView Source Pane

Gridded box is a possible site for a breakpoint

Select to set one

Current point of execution

Current function and source file

```
Function calc_pi in calc_pi.c
23 void crash(int* ip){
24 ip=0;
25 *ip=7;
26 }
27
28 double calc_pi(int n){
29 double sum=0.0;
30 double w=1.0/n;
31 int i=0;
32 double x=0.0;
33 #pragma omp parallel for reduction(+:sum)
34 for(i=1;i<=n;i++){
35 x = w*(i-0.5);
36 sum=sum+f(x);
37 if(i == 800)
38 crash(&i);
39 }

```

- Dive on a source word to get more information
- Select a line to use **Run to selection** command
- Select or dive on a line number to set an action point



## Parallel Debugging - Philosophy

- By default, TotalView places processes in groups
  - Program Group - Includes parent and all related processes
  - Share Group - Only processes that share the same source code
- Command can act on single process or share group
  - halt process (h) , halt group (H)
  - next step process (n), next step group (N)
  - go process (g), go group (G)



## TotalView Exercise: Debug simple program

- Run calc\_pi inside totalview:
  - Check where the program crashes
- Analyze core file with totalview
  - run calc\_pi
  - execute `totalview calc_pi core`
- For advanced users: choose another programming paradigm:
  - MPI, OpenMP, MPI+OpenMP



## TotalView support for debugging MPI

- Special support for MPI is available depending on your MPI library:
  - display message queue state of a process
- Supported MPI implementations:
  - mpich v1.1.0 or later ( use -debug in configure)
  - HP MPI v1.6
  - Compaq MPI >v1.7
  - IBM, release >2.2 of Parallel Environment, threaded version of MPI
  - SGI MPI v1.3 or later



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 43 / 60 Höchstleistungsrechenzentrum Stuttgart



## TotalView MPI Message Queue Window

Communicator name and info

Non-blocking receive operations

Unmatched incoming messages

Non-blocking send operations

- Dive on source or target to refocus Process window
- Dive on buffer to see message contents

```
testsome.0
Message State for "testsome.0" (1288,1)
MPI_COMM_WORLD
Comm_size 3
Comm_rank 0
Pending receives
[0]
 Status Pending
 Source 2 (testsome.2)
 Tag 0x00000000 (0)
 User Buffer 0x000605c0 -> 0x00000000 (0)
 Buffer Length 0x00000014 (20)
Unexpected messages
[0]
 Status Complete
 Source 2 (testsome.2)
 Target 0x00000002 (2)
 Tag 0x00000000 (0)
 System Buffer 0x00000000
 Buffer Length 0x00000000 (0)
 Received Length 0x00000000 (0)
Non-blocking sends
[0]
 Status Complete
 Target 2 (testsome.2)
 Tag 0x00000000 (0)
 Buffer 0x000605a0 -> 0x00000001 (1)
 Buffer Length 0x00000014 (20)

MPI_COMM_WORLD_collective
Comm_size 3
Comm_rank 0
```



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 44 / 60 Höchstleistungsrechenzentrum Stuttgart



skipped

## TotalView Exercise: Parallel program

- Example in TOTALVIEW/MPI:
  - `deadlock_{c,cc,f90}.{c,cc,f90}`
  - start program with `mpirun -tv -np 2 a.out`
  - interrupt execution after “deadlock”
  - try to find the reason for the deadlock and fix it
- For advanced users:
  - `pending_{c,cc,f90}.{c,cc,f90}`
  - try to find pending message by setting breakpoint at `MPI_Finalize`



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 45 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

## TotalView more information

- <http://www.etnus.com/products/totalview/index.html>
- <http://www.hlr.de/organization/tsc/services/tools/debugger/totalview>
  - User Guide
  - Installation Guide
  - CLI Guide
  - Powerpoint Tutorial
- CRAY T3E: Online Documentation at <http://www.hlr.de/platforms/crayt3e>



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 46 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S 

# Distributed Debugging Tool (DDT)



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Hochleistungsrechenzentrum Stuttgart

H L R I S 

## What is DDT?

- Parallel debugger
- Source level debugging for C, C++, F77, F90
- MPI, OpenMP
- SMPs, Clusters
- Available on Linux distributions and Unix
- GUI (independent of platform, based on QT libraries)



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 48 / 60 Hochleistungsrechenzentrum Stuttgart

H L R I S 

## Availability of DDT

- Linux:
  - Linux IA32 (Intel and AMD)
  - Linux IA64
  - Linux Opteron
- Unix
  - PowerPC (AIX)
  - SGI Altix
  - SGI Irix
  - SUN Sparc
  - PA-Risc and Itanium Superdome



skipped

## Availability of DDT at HWW

| Platform | Availability | Remarks |
|----------|--------------|---------|
| Volvox   | Yes          | V 1.4   |
| AzusA    | Yes          | V 1.4   |

More information:

<http://www.hlrs.de/organization/tsc/services/tools/debugger/ddt/>



## Availability of DDT at University Stuttgart

- Two user Floating License for University Stuttgart:
  1. Download Software from <http://www.etnus.com>
  2. Set environment variable for license.  
`LM_LICENSE_FILE=7244@servint1.rus.uni-stuttgart.de`

More information about campus licenses available at  
<http://www.hlrs.de/organization/par/services/tools/campus>



skipped

## DDT usage at HLRS

- Set USE\_DDT in your login scripts
- Compile with -g compiler switch
- Command name: **ddt** or **\$DDT**
- To start debugging with DDT simply type:  
`% $DDT myprog arguments to myprog`

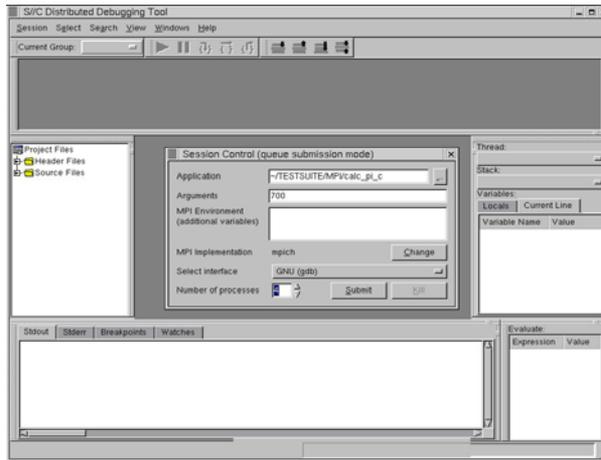


## DDT Look & Feel

DDT Main Window

Configuration Window

and all belonging Panes (Thread, Stack, Output, Source code, etc.)



Parallel Debugging Müller, Neytchev, Keller, Kramer  
Slide 53 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S

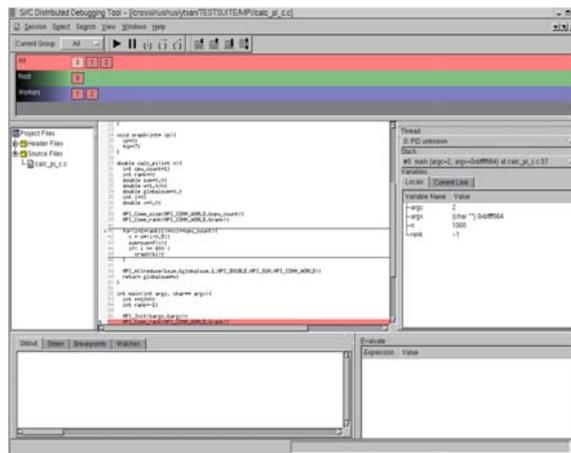
## DDT Main/Process Window

MPI Groups

File browse and Source pane

Output, Breakpoints, Watch

Pane



Thread, Stack, Local and Global Variables Pane Evaluation window



Parallel Debugging Müller, Neytchev, Keller, Kramer  
Slide 54 / 60 Höchstleistungsrechenzentrum Stuttgart

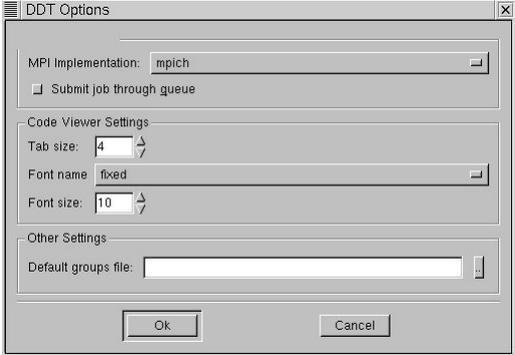
H L R I S

**skipped**

## DDT Options Window

The Options window:  
here is the MPI library implementation selected

The program can also be submitted through some batch system



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 55 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S

**skipped**

## DDT Options Window (Queue)

The Options windows:  
This program uses mpich as MPI implementation and starts the program through PBS batch system

For more information about the Template files and the commands, please read the DDT Users Manual.



Parallel Debugging Müller, Neytchev, Keller, Krammer  
Slide 56 / 60 Höchstleistungsrechenzentrum Stuttgart

H L R I S

skipped

## DDT Thread and Stack Window

### Thread Pane:

Switch between all program threads

### Stack Pane:

Switch between the functions in the selected thread

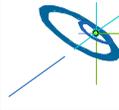
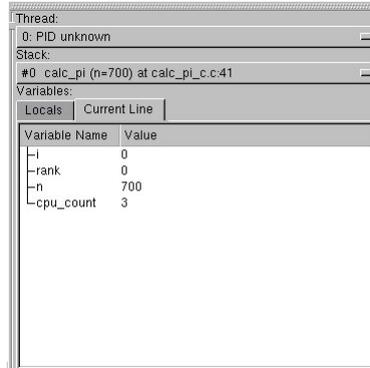
### Variables Pane:

Local variables:

Shows the variables value for the function in which ddt is currently stopped

### Current Line:

Shows every variable value between the two lines selected by the user



## DDT Source Pane

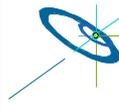
With the right mouse button Set or Remove a breakpoint at the selected line

When the program is running and stopped at a breakpoint, the line is coloured in red (by OpenMP programs) and red, blue or green by programs using MPI

```

1 // Simple OpenMP Program to calculate PI
2 //
3 //
4 // Author: Matthias Mueller-Oueller@irs.de
5 // Thu Dec 28 15:12:09 CET 2006
6 //
7 // Usage: calc_pi [iterations]
8 // The program will crash if iterations >= 800
9 //
10 // This is the intended behavior, because the program
11 // test a debugger.
12 //
13 //
14 #include <stdio.h>
15 #include <stdlib.h>
16 #include <math.h>
17 double f(double a){
18 return 4.0/(1.0 + a*a);
19 }
20
21 void crash(int* ip){
22 ip[0] =
23 *ip++;
24 }
25
26 double calc_pi(int n){
27 double sum=0.0;
28 double w=1.0/n;
29 int i;
30 double w2=w*w;
31
32 #pragma omp parallel for reduction(+:sum)
33 for(i=1; i<=n; i++){
34 w = w*(1-0.5*i);
35 sum+=w*w*f(w);
36 if(i == 800)
37 crash(ip);
38 }
39 return sum*w;
40 }
41
42 int main(int argc, char** argv){
43 int n=800;
44 if(argc>2){
45 fprintf(stderr, " Usage : %s iterations\n", argv[0]);
46 return 1;
47 }
48 if(argc==2){
49 n=atoi(argv[1]);
50 }
51 printf(" Pi = %g\n", calc_pi(n));
52 return 0;
53 }

```



## Parallel Debugging - Philosophy

- By default, DDT places processes in groups
  - All Group - Includes parent and all related processes
  - Root/Workers Group - Only processes that share the same source code
- Command can act on single process or group
  - stop process , stop group
  - next step process , next step group
  - go process, go group



## DDT more information

- [http://www.streamline-computing.com/softwaredivision\\_1.shtml](http://www.streamline-computing.com/softwaredivision_1.shtml)

