


DDT, a Global Debugger Solution at KIT

Frank.Schmitz@kit.edu



Karlsruhe Institute of Technology (KIT)

University

11 fakulties
 120 institutes
 4,000 employees
 18,500 students
 250 mio € budget

Forschungszentrum

10 programs (research areas)
 education
 21 large institutes
 structures for the administration
 4,000 employees
 institutes
 310 mio € budget
 services

infrastructure

n Karlsruhe Institute of Technology - two strong partners united

- α Universität Karlsruhe (TH)
 - n One of the most research oriented universities in Germany
 - n Nominated to one of three top universities in October 2006
 - n Capitalized with over 100 million € over 5 years
- α Forschungszentrum Karlsruhe
 - n Member of the Helmholtz Association
 - n Internationally acknowledged non-university institution

n Under the roof of KIT – Steinbuch Centre for Computing

- α Computing Center of the Universität Karlsruhe (TH)
- α Institute of Scientific Computing at the Forschungszentrum Karlsruhe

Virtual Computing Centre Karlsruhe since 1996



Cooperation in
High Performance Computing

Virtual Computing Centre since 1996

- n IT-Services for UKA
- n High Performance Computing
- n Numerical Methods on HPC
- n Integrated Information Mgmt.

- n IT-Services for FZK
- n Scientific Computing
- n GridKa, CampusGrid, D-Grid
- n Nat. & Intl. Grid Projects

HP XC4000 – Baden-Württemberg-State-Supercomputer, a part of SCC




Phase 2 (Q3 2006)

- »750 4-way nodes (two sockets)
 - dual core AMD Opteron 2.6 Ghz
 - 16 GB main memory
- »10 server nodes
- »10 file server nodes
 - Approx. 56 TB storage system
- »Infiniband DDR interconnect
- »20 water cooled racks (HP MCS)

- » Total of **3000 processor cores**
- » Total of **15.6 TFlop/s peak performance**
- » Total of **12 TB** of main memory
- » Total **110 TB** of local disk space (scratch)

Parts of SCC



a SX-9 node will be added in
autumn 2008

DDT versions at KIT

- n 1.9.1c in the Linux based OPUS^{IB} environment
- n SX8beta6.RELEASE on SX8R and front-ends (Version 2.3)
- n 1.10 in the AIX environment
- n 2.2 on the HP XC4000 system

Concept: Like the solution with TotalView 5-10 years ago, SCC will give the users the possibility to use only one interface for solving problems in a parallel and heterogeneous environment.
à in the future only one version should be in production

But: user's don't want to use a debugger. Help from the computer centre is needed!

DDT at HP XC4000

- n Interactive mode is limited to 8 tasks
- n No limitation for batch, but wait time is required if the machine is fully loaded

DDT at OPUS^{IB}

- n Interactive mode with only limitations in memory and CPU-time
- n No limitation for batch, but wait time is required if the machine is fully loaded

DDT under AIX

- n Interactive mode is limited to 8 tasks
- n No limitation for batch, but wait time is required if the machine is fully loaded

DDT at SX-8R

- n Interactive mode with no limitations in memory and CPU-time (interactive mode preferred), limitation is one node
- n No limitation for batch, but wait time is required because the machine is normally fully loaded


DDT at SX-8R, some features

- n Breakpoints
- n attaching running processes
- n OpenMP
- n visualize data
- n Watchpoints

DDT problems at SX-8R

- n timeout because debugging is not a high priority process, you have to set `DDT_NO_TIMEOUT=1`, it's not the solution but it helps!
- n selection of message queues doesn't work
- n the scalar unit of the SX-8R is very slow compared with x86 hardware
- n add a close icon to close windows (some have it!)
- n evaluating for visualization in 3D takes a lot of time
- n adding static fields to a list of watchpoints is not possible
- n 3-D view isn't very helpful, but the table helps a lot

3-D view of data, helpful?



Data table

Array Expression:

Range of \$i

From:

To:

Display:

Range of \$j

From:

To:

Display:

Auto-update

Aggregate Function:

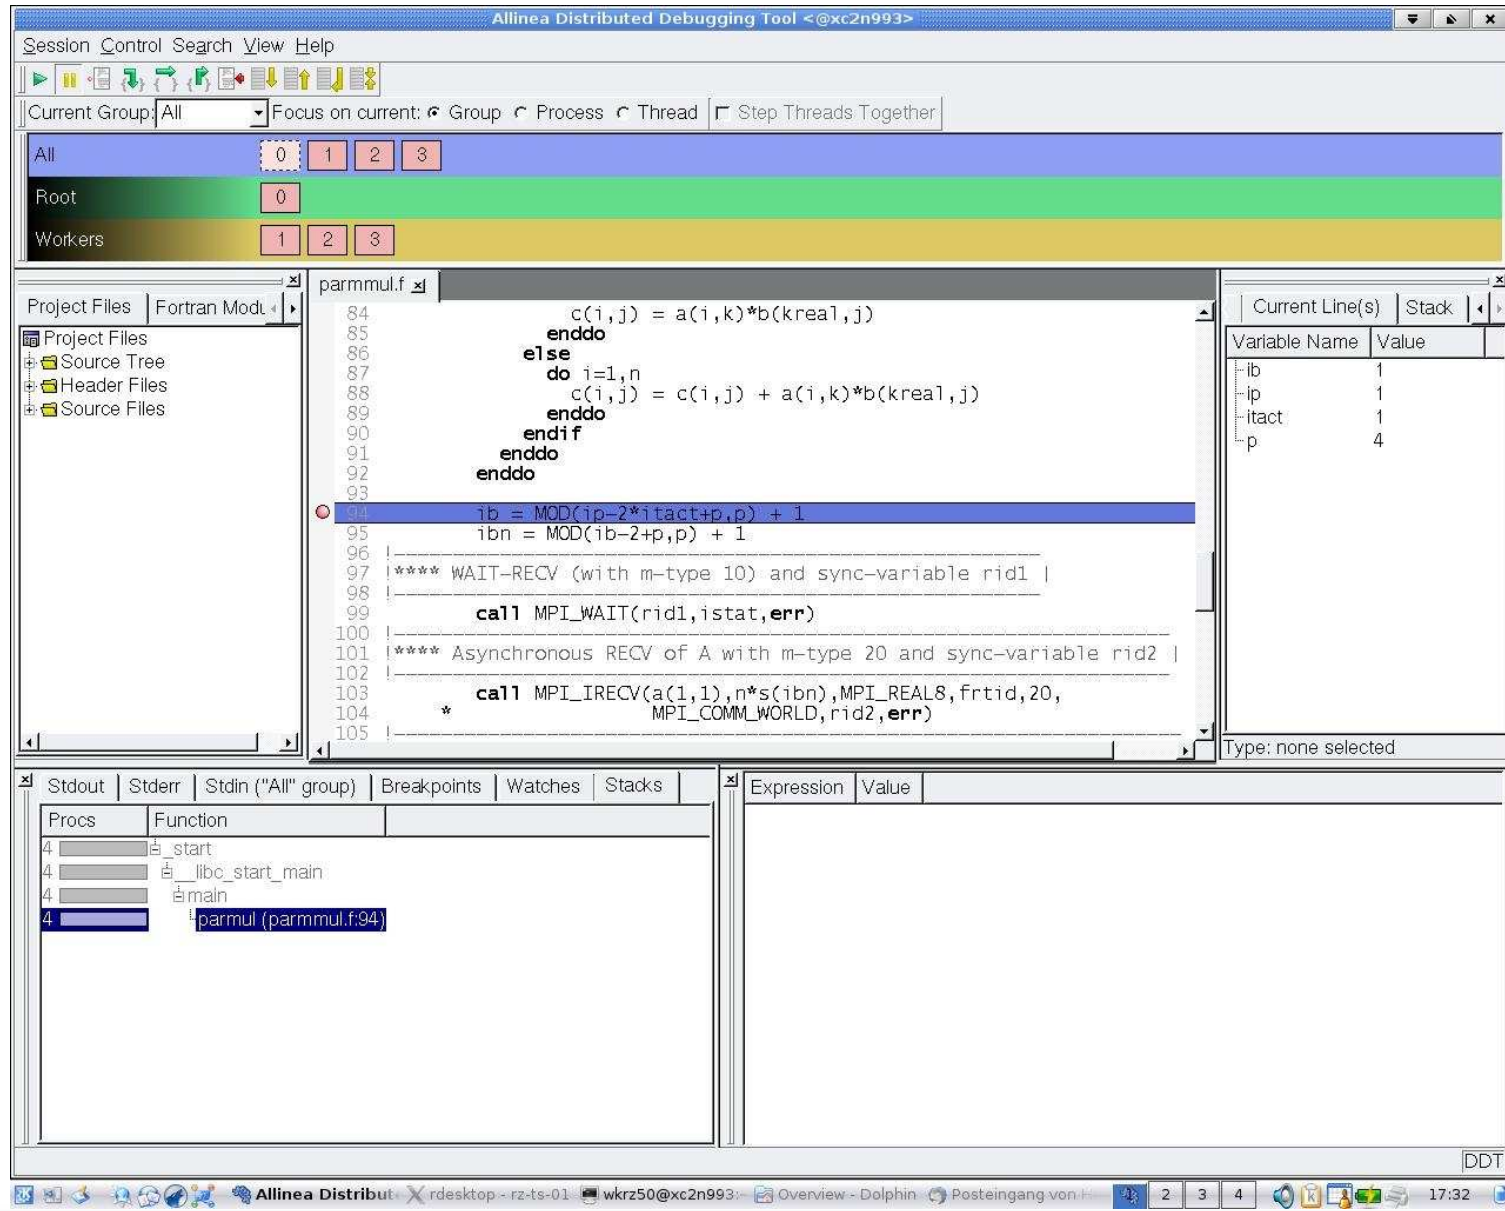
Filter: =

Data Table

| | 1 | 2 | 3 | 4 | 5 |
|----|-------------------|-------------------|-------------------|-------------------|----------|
| 1 | 10925.54870298208 | 1166.444316954099 | 571.1687855806970 | 512.7761596216023 | 785.6856 |
| 2 | 1166.444316954099 | 571.7380015093366 | 513.7971280909828 | 788.0304735593255 | 2938.907 |
| 3 | 571.1687855806970 | 513.7971280909828 | 788.8120804799773 | 2944.738150080539 | 87049.37 |
| 4 | 512.7761596216023 | 788.0304735593255 | 2944.738150080539 | 87142.89830155349 | 1733.850 |
| 5 | 785.6856527973699 | 2938.907474176382 | 87049.37528013957 | 1733.850981586058 | 664.0792 |
| 6 | 2927.246122368068 | 86768.80621589771 | 1730.448935629626 | 663.4315815664968 | 124890.6 |
| 7 | 86301.19110882796 | 1723.644843716763 | 661.4886028997657 | 124890.6107246275 | 673.8041 |
| 8 | 1713.438705847467 | 658.2503051218807 | 124890.6009732201 | 671.8517675277851 | 1792.370 |
| 9 | 653.7166882328418 | 124890.5863461092 | 668.5977728192163 | 1785.471708258071 | 100729.3 |
| 10 | 124890.5668432945 | 664.0421802272199 | 1775.123900701252 | 100201.8088397234 | 2997.151 |
| 11 | 658.1849897517960 | 1761.326823958828 | 99463.19693384752 | 2979.994768807044 | 819.3344 |

Expression "c[\$i,\$j]" evaluated for process 0 at 09:32.

DDT at HP XC4000



Session Control Search View Help

Current Group: All Focus on current: Group Process Thread Step Threads Together

All 0 1 2 3
Root 0
Workers 1 2 3

Project Files Fortran Mod

Project Files
Source Tree
Header Files
Source Files

```
parmmul.f
84      c(i,j) = a(i,k)*b(kreal,j)
85      enddo
86      else
87      do i=1,n
88      c(i,j) = c(i,j) + a(i,k)*b(kreal,j)
89      enddo
90      endif
91      enddo
92      enddo
93
94      ib = MOD(ip-2*itact+p,p) + 1
95      ibn = MOD(ib-2+p,p) + 1
96
97      **** WAIT-RECV (with m-type 10) and sync-variable rid1 |
98      -----
99      call MPI_WAIT(rid1,istat,err)
100     -----
101     **** Asynchronous RECV of A with m-type 20 and sync-variable rid2 |
102     -----
103     * call MPI_IRecv(a(1,1),n*s(ibn),MPI_REAL8,frtid,20,
104       MPI_COMM_WORLD,rid2,err)
105     -----
```


| Current Line(s) | Stack |
|-----------------|-------|
| ib | 1 |
| ip | 1 |
| itact | 1 |
| p | 4 |

| Procs | Function |
|-------|-----------------------|
| 4 | _start |
| 4 | _libc_start_main |
| 4 | main |
| 4 | parmul (parmmul.f:94) |

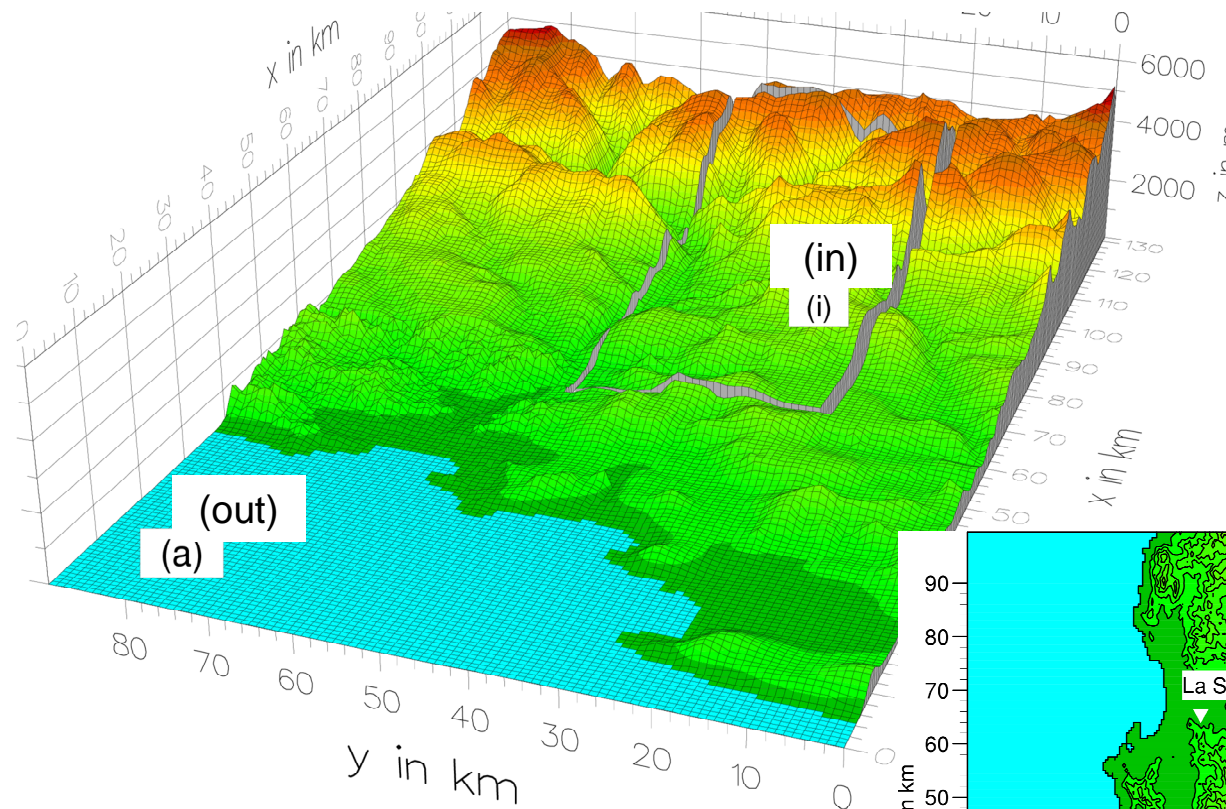
| Expression | Value |
|------------|-------|
|------------|-------|

DDT

Allinea Distribut... rdesktop - rz-ts-01... wkrz50@xc2n993... Overview - Dolphin... Posteingang von... 2 3 4 17:32


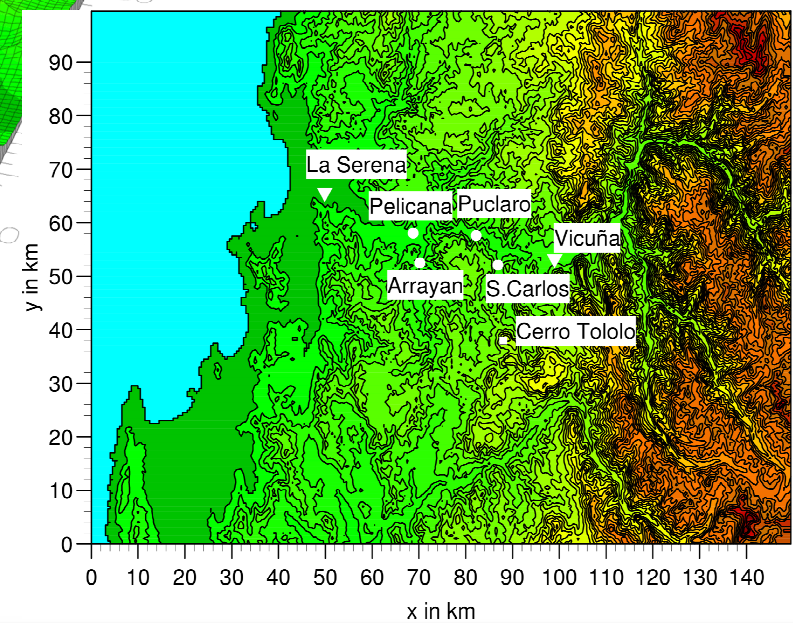


The influence of a Storage Lake in the Arid Elqui Valley in Chile on the Local Climate




225 x 135 horizontal gridpoints
100 vertical levels

out: $\Delta x = 800 \text{ m}$, $\Delta y = 800 \text{ m}$
 in(1): $\Delta x = 400 \text{ m}$, $\Delta y = 400 \text{ m}$
 in(2): $\Delta x = 200 \text{ m}$, $\Delta y = 200 \text{ m}$
 in(3): $\Delta x = 100 \text{ m}$, $\Delta y = 100 \text{ m}$



11:02

→ 10 m s⁻¹



Thank you!

Questions?