



Gesellschaft für Parallele Anwendungen und Systeme mbH

OpenMP Tools

Hans-Joachim Plum, Pallas GmbH
edited by Matthias Müller, HLRS

Pallas GmbH
Hermülheimer Straße 10
D-50321 Brühl, Germany

info@pallas.de
<http://www.pallas.com>

Introduction: OpenMP is clear and easy ..



- 😊 Easy directive based handling
- 😊 Portable code
- 😊 Single source for sequential/parallel
- 😊 Incremental parallelization
- 😊 Well-defined semantics

So, why tools ?

Introduction: ... but



- ☹ Performance tuning can be tricky
Losses are transparent
- ☹ 'Easy' to introduce data conflict bugs
These are tedious to fix
- ☹ Data mapping
Will be important for clusters. Heavy inter-thread data dependencies often difficult to detect.
- OpenMP tools must
 - bring out the strengths of OpenMP
 - help to overcome the problems

Introduction: KAP/Pro Toolset



- Supported Platforms:
 - Compaq alpha (True64 Unix, WinNT4.0)
 - SGI MIPS (IRIX 6.5)
 - Sun SPARC (Solaris 2.5, 2.6, 2.7)
 - Intel IA32 (WinNT4.0, Linux86, Solaris86)
 - HP PA-RISC 2.0 (HP-UX 11)
 - IBM RS/6000 (AIX 4.1.* - 4.3)

Overview: KAP/Pro Toolset



- Guide: fully OpenMP compliant F77/F90/C/C++ compiler
- GuideView: graphical performance analysis tool
- Assure: fully OpenMP compliant F77/F90/C/C++ compiler for checking semantical code correctness
- AssureView: graphical backend of Assure
- PerView: online performance monitor

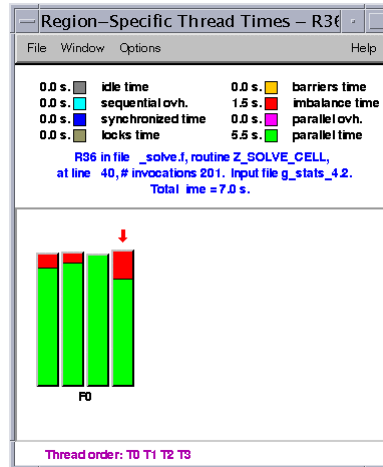
Overview: Guide and GuideView



- Guide: fully OpenMP compliant C/C++/F77/F90 compiler
- Based on native compiler => native performance, native options are passed to backend compiler

PLUS:

- An additional statistics library on user request, traces OpenMP events at runtime
- After run, visualize tracefile with GuideView



What can be detected?



Guide (and its alarm colors):

- Loadimbalance ●
- Too many synchronization points ●
- Too much 'communication' ●
- Heavy critical sections ●
- Acquiring locks ●
- Heavy sequential sections (overhead) ● (●)
- Idle times ●
- And the good color is ●
(independent work of thread)



- Compile code with Guide compiler:

```
guidef90 [options] -WGstats -o myprog myprof.f90
```

- Execute

```
setenv OMP_NUM_THREADS ..  
./myprog  
=> Guide statistics file guide.gvs has been produced
```

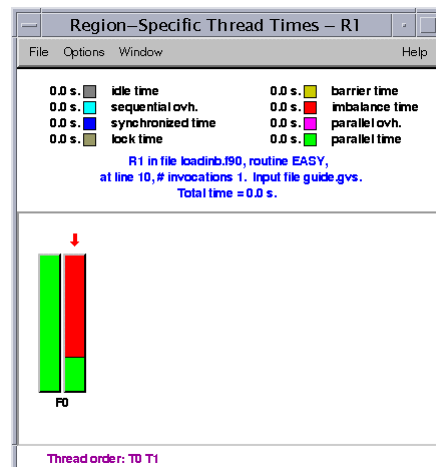
- Visualize

```
guideview guide.gvs&
```

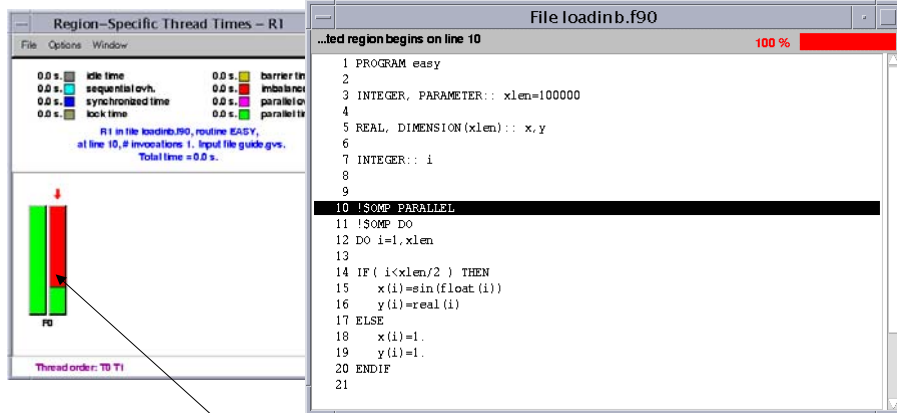
What can be detected: Loadimbalance



```
!$OMP PARALLEL  
!$OMP DO  
DO i=1,xlen  
IF( i<xlen/2 ) THEN  
    x(i)=sin(float(i))  
    y(i)=real(i)  
ELSE  
    x(i)=1.  
    y(i)=1.  
ENDIF  
ENDDO  
!$OMP END PARALLEL
```



Source code view is available!



Right mouse click
 => Show code

What can be detected: Barriers

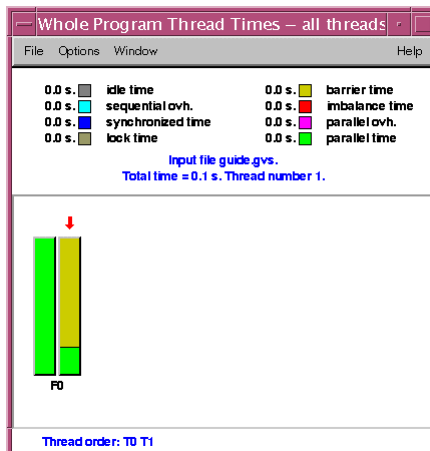


```

!$OMP PARALLEL
!$OMP DO
DO i=1,xlen
<imbalanced as above>
ENDDO

!$OMP DO
DO i=2,xlen
  z(i) = x(i-1)+y(i)
ENDDO

!$OMP END PARALLEL
    
```



What can be detected: Overhead



Parallelization of
innermost loop:

```
DO i=1,1000
```

```
  !$OMP PARALLEL DO
```

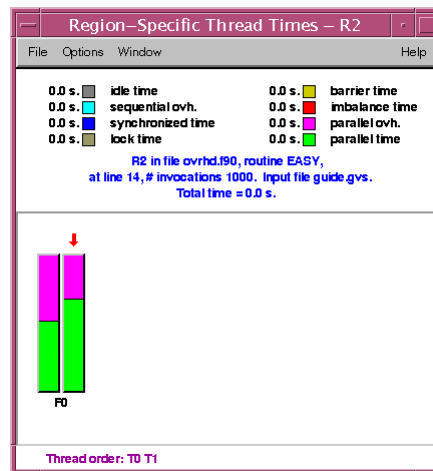
```
    DO j=1,100
```

```
      temp(j) = &
```

```
      temp(j)+1./i
```

```
    ENDDO
```

```
ENDDO
```



What can be detected: Critical + Locks



```
!$OMP PARALLEL PRIVATE(temp)
```

```
!$OMP DO
```

```
DO i=1,100000
```

```
  ! PARALLEL PRVIATE
```

```
  ! CALCULATIONS ..
```

```
  temp = 1.
```

```
  ! CRITICAL UPDATES
```


```
  !$OMP CRITICAL
```

```
    sum = sum+sin(temp)
```

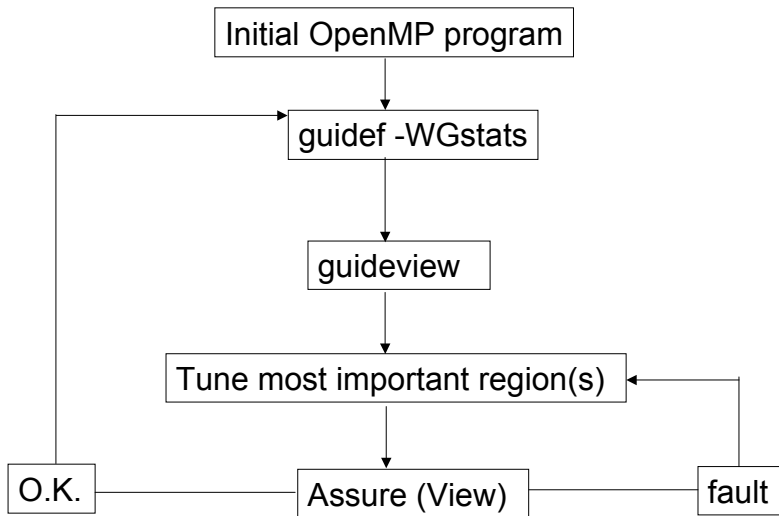
```
  !$OMP END CRITICAL
```

```
ENDDO
```

```
!$OMP END PARALLEL
```



The incremental optimization cycle



© Pallas GmbH

Guide



Output file names can be controlled by environment variables:

KMP_STATSFILE (guide.gvs file)

KDD_OUTPUT (.kdd file of assure, see below)

Three meta characters are available to use in the names:

%P #threads of the run

%l unix process id of the run

%H hostname to run the program

© Pallas GmbH



e.g.

```
setenv KMP_STATSFILE gs%P_%I_%H
```

gives Guide statsfile similar to

```
gs2_11074_vision.gvs
```

GuideView: Displays



Usage:

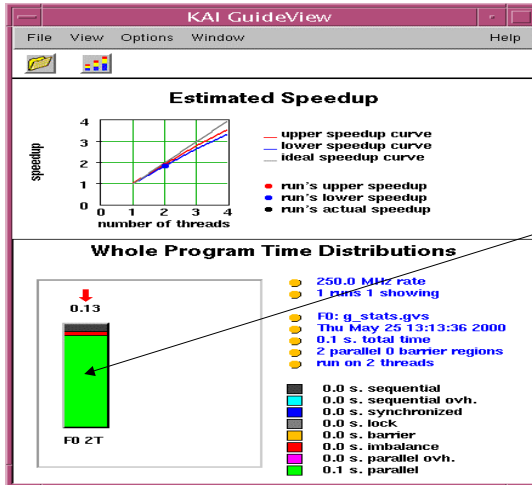
```
guideview <stats_file>
```

Main display

- Estimated speedup
- Whole program timing breakdown



Whole timing breakdown (averages of threads)

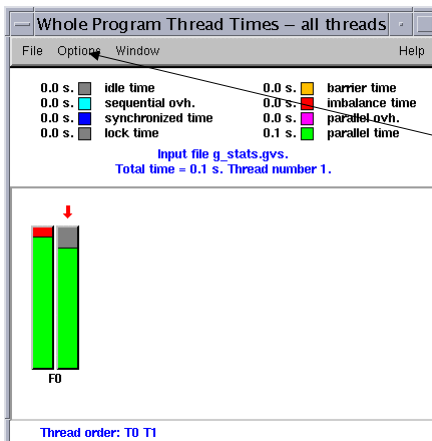


Left button
doubleclick

=> ...



... => Thread timing breakdown

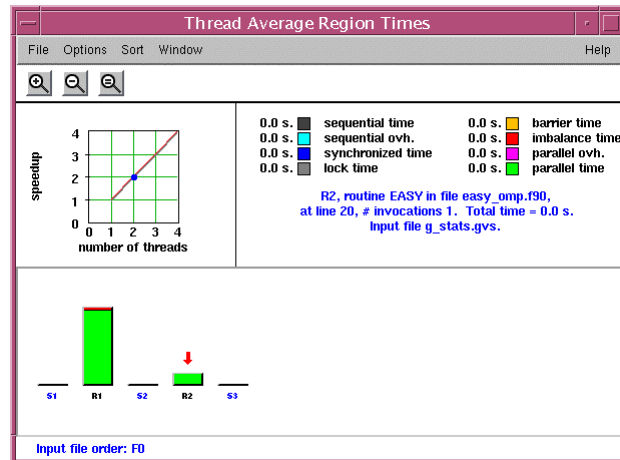


Region View

=> ...



...=> Region timing breakdown (again average or per thread)



Region timing breakdown

- Each parallel region (denoted 'R1, R2, ...')
- Each sequential part (denoted 'S1, S2, ...')
- Inside each region (optional):

Breakdown into sections between barrier points (denoted 'RxB1, RxB2 ...').

Can be unselected in "Options/Show Barrier Regions"



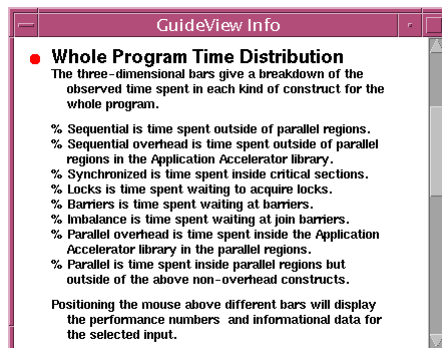
Always remember:

There is a link between each of the displayed sections and the corresponding code

(right mouse button on column => context 'show code')



- Click "**Help/GuideView Info**" for explanations





Sort button

Several sort criteria for the regions, in particular:

sort by time => most important first

sort by overhead => most critical first

Options/Filter button

Filter regions, e.g. only display heavy ones (more than 10% of total)



Compare different #threads for same program (e.g.)

- `guideview <first statsfile>`
- Click "**File/open new file** "
Menu opened => select second statsfile

All views are then comparative between the two files

GuideView: Trace Comparison (1 <-> 2 threads)



© Pallas GmbH

Overview: Assure and AssureView



- Assuref77/f90: OpenMP compliant, restriction in usage of OMP library (in particular: OMP_GET_THREAD_NUM)
- Use as normal compiler, but not for getting performance (small input data set)
- Multithreaded run is simulated sequentially, all memory accesses verified
- Run AssureView to visualize error breakdown. When "No Errors" are reported, multithreaded run is assured free of semantical errors as explained below, but only in the branches touched by the simulation run.

© Pallas GmbH

What can be detected: invoke Assure



- Compile code with yet another compiler:

```
assuref90 [options] -o myprog myprof.f90
```

- Execute (but don't expect performance!!)

```
./myprog
```

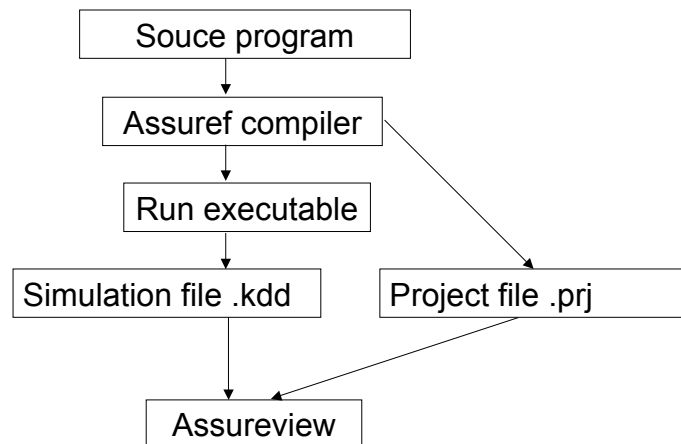
- Visualize

```
assureview
```

Assure



The Assure process



What can be detected: Conflicts



```
real:: a(0:N), b(N)
a(0) = 0.
!$OMP PARALLEL
!$OMP DO
DO i=1,N
    a(i) = 1./i
    b(i) = a(i-1)+a(i)
ENDDO
```

Whats wrong ??

© Pallas GmbH

What can be detected: Conflicts



The screenshot displays two windows from a code analysis tool. The left window, titled 'Source: confl.f90', shows the original Fortran code. The right window, titled 'Sink: confl.f90', shows the same code with red annotations indicating detected conflicts. Red lines connect corresponding lines between the two windows, highlighting areas where conflicts were found. In the 'Sink' window, red dots are placed next to lines 13, 14, and 25, which correspond to the parallelized sections of the original code. The code in both windows is as follows:

```
1 program confl
2
3 integer, parameter:: N=10
4 real:: a(0:N), b(N), aux
5
6 a(0) = 0.
7
8 !$OMP PARALLEL
9
10 !$OMP DO
11 DO i=1,N
12     a(i) = 1./i
13     b(i) = a(i-1)+a(i)
14 ENDDO
15
16 !$OMP END PARALLEL
17
18 print*, b
19
20 !$OMP PARALLEL
21
22 !$OMP DO
23 DO i=1,N
24     b(i) = aux+a(i)
25 ENDDO
26
27 !$OMP ENDDO NOWAIT
28
29 !$OMP DO
30 DO i=1,N
```

© Pallas GmbH

What can be detected: Conflicts



```
Source: confl.f90
11 !$OMP DO
12 DO i=1,N
13   a(i) = 1./i
14   b(i) = a(i-1)+a(i)
15 ENDDO
16
17 !$OMP END PARALLEL
18
19 print*, b
20
21 !$OMP PARALLEL
22
23 !$OMP DO
24 DO i=1,N
25   b(i) = aux+a(i)
26 ENDDO
27 !$OMP ENDDO NOWAIT
28
29 !$OMP DO
30 DO i=1,N
31   a(i) = b(i-1)
32 ENDDO
33
34
35 !$OMP END PARALLEL
36
37 print*, b
38
39 !$OMP PARALLEL
40
```

```
Sink: confl.f90
17 !$OMP END PARALLEL
18
19 print*, b
20
21 !$OMP PARALLEL
22
23 !$OMP DO
24 DO i=1,N
25   b(i) = aux+a(i)
26 ENDDO
27 !$OMP ENDDO NOWAIT
28
29 !$OMP DO
30 DO i=1,N
31   a(i) = b(i-1)
32 ENDDO
33
34
35 !$OMP END PARALLEL
36
37 print*, b
38
39 !$OMP PARALLEL
40
41 !$OMP DO
42 DO i=1,N
43   aux = a(i-1)+a(i)
44   b(i) = aux*aux
45 ENDDO
46
```

© Pallas GmbH

Assure: Error Types



Write-Read conflicts

```
!$OMP PARALLEL DO
DO i=1,N
    a = b+c(i)
    d(i) = a+e(i)
```

The 2 statements inside the loop have to be executed in that (Write-Read) order, which is not guaranteed in a multithreaded run (a is shared by default).

Repair: private(a)

© Pallas GmbH



Read-Write conflicts

```
!$OMP PARALLEL DO
DO i=1,N
    d(i) = a+e(i)
    a = b+c(i)
```

Repair: private(a)



Write-Write conflicts

```
!$OMP PARALLEL DO
DO i=1,N
    a = b+c(i)
```

Repair: private(a)



Private symbol, used outside loop

```
!$OMP parallel do private(a)
DO i=1,N
    a=c(i)
    d(i) = a*a
ENDDO
PRINT*, a
```

Repair: lastprivate(a)



Uninitialized private

```
firstiter = .TRUE.

!$OMP parallel do private(firstiter)
DO i=1,N
    IF( firstiter ) THEN ...
ENDDO
```

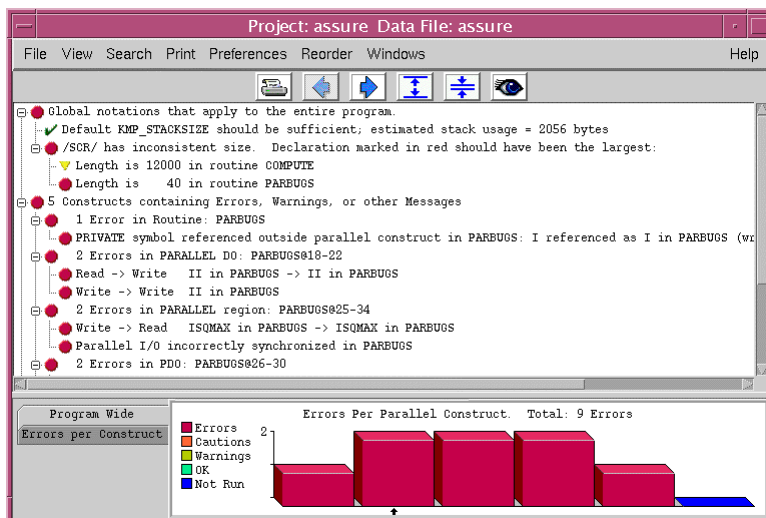
Repair: firstprivate(firstiter)



Main

- Main error list
 - Clickable button for each error
Click to get precise diagnostics
 - Overview chart showing statistics of bugs, different severities
- Call Tree

Assureview: Displays





Reading the diagnostics

Click the "+" buttons to get into the diagnostics

Finally the code sections are shown containing the error locations, (source and sink), both clearly marked.



Inside code windows

- Show Search: normal string search menu
- Show Stack: show the calling sequence for arriving at the location.



Other buttons

View

Select display of the error list

Search

Normal search menu, inside error list

Print

Self explaining



Preferences

Miscellaneous settings. In particular:
source code locations ("finding files")

Reorder

.. error list by different criteria

Thanks for your attention!



Pallas GmbH
Hermülheimer Straße 10
D-50321 Brühl, Germany

info@pallas.de
<http://www.pallas.com>