



SOFTWARE

- +  19.56 updatex
- +  399.70 updateien
- +  0.00 gene
- 0.00 <<iteration loop>>
- +  447.52 genbc



FAST SOLUTIONS

- PAPI\_L1\_ICM
- PAPI\_L2\_DCM
- PAPI\_L2\_ICM
- PAPI\_L1\_TCM

# Periscope

## Score-P Online Access Tutorial Exercise

### NPB-MZ-MPI/BT

Isaías A. Comprés

Technische Universität München

[periscope@in.tum.de](mailto:periscope@in.tum.de)

September 2012

## **Installation Checklist**

Benchmark Instrumentation

Periscope Online Analysis

Performance Properties Exploration

- Confirm that the configuration file `.periscope` is in the home directory.
- It should look like:

```
MACHINE = localhost
SITE = LiveDVD
REGSERVICE_HOST = latest64.vi-hps.org
REGSERVICE_PORT = 50027
REGSERVICE_HOST_INIT = localhost
REGSERVICE_PORT_INIT = 50001
APPL_BASEPORT = 51000
AGENT_BASEPORT = 50002
```

- Verify that Periscope and Score-P commands are available:

```
livetau$ which scorep
/usr/local/packages/scorep-1.1/bin/scorep
livetau$ which psc_frontend
/usr/local/packages/periscope/1.5/bin/psc_frontend
```

Installation Checklist

**Benchmark Instrumentation**

Periscope Online Analysis

Performance Properties Exploration

- Change directory to the prepared NPB3.3 folder:

```
livetau$ cd $HOME/workshop-vihps/NPB3.3-MZ-MPI
```

- Open the make configuration file with an editor:

```
livetau$ vi config/make.def
```

- Uncomment the MPIF77 definition for Score-P:

```
# Alternative variants to perform instrumentation
#MPIF77 = psc_instrument mpif77
#MPIF77 = scalasca -instrument mpif77
#MPIF77 = tau_f90.sh
#MPIF77 = vtf77 -vt:hyb -vt:f77 mpif77
MPIF77 = scorep --user mpif77
```

- Open the BT-MZ Makefile with an editor:

```
livetau$ vi BT-MZ/Makefile
```

- Comment out the MAIN definitions that are not for Score-P:

```
#MAIN = bt_epik
MAIN = bt_scorep
#MAIN = bt
```

- Take note of the Score-P Online Access phase definition in the prepared BT benchmark. Open the file `BT-MZ/bt_scorep.F` and move to line 219:

```
    if (mod(step, 20) .eq. 0 .or. step .eq. 1) then
      if (myid .eq. root) write(6, 200) step
      if (myid .eq. root) call flush(6)
200    format(' Time step ', i4)
    endif

    SCOREP_USER_OA_PHASE_BEGIN(OA_Phase, "OA_Phase", SCOREP_USER_REGION_TYPE_COMMON)
```

- Make sure that the suite file defines the correct benchmark to build:

```
livetau$ cat config/suite.def
<comments>
bt-mz      W          4
```

- Build and instrument the benchmark:

```
livetau$ make suite
```

- Alternatively, the build can also be made as follows:

```
livetau$ make bt-mz CLASS=W NPROCS=4
```

Installation Checklist

Benchmark Instrumentation

**Periscope Online Analysis**

Performance Properties Exploration

- Periscope is started via its frontend. It automatically starts application and hierarchy of analysis agents.
- Run `psc_frontend --help` for brief usage information

```
% psc_frontend --help
Usage: psc_frontend <options>
  [--help]                (displays this help message)
  [--quiet]               (do not display debug messages)
  [--registry=host:port] (address of the registry service, optional)
  [--port=n]              (local port number, optional)
  [--maxfan=n]            (max. number of child agents, default=4)
  [--timeout=secs]       (timeout for startup of agent hierarchy)
  [--delay=n]             (search delay in phase executions)
  [--appname=name]
  [--apprun=commandline]
  [--mpinumprocs=number of MPI processes]
  [--ompnumthreads=number of OpenMP threads]
...
  [--strategy=name]
  [--sir=name]
  [--phase=(FileID,RFL)]
  [--debug=level]
```



- Change to the `bin.scorep` directory:

```
livetau$ cd bin.scorep
```

- Run Periscope by executing `psc_frontend` with the following command and options:

```
livetau$ psc_frontend --apprun=./bt-mz_B.4 --strategy=MPI  
--mpinumprocs=4 -ompnumthreads=1 -phase="OA_phase"
```

```
[psc_frontend][DBG0:fe] Agent network UP and RUNNING. Starting search.
```

```
NAS Parallel Benchmarks 3.3 -- BT Benchmark
```

```
[...]
```

```
Time step 200
```

```
BT Benchmark Completed.
```

```
-----  
End Periscope run! Search took 60.5 seconds (33.3 seconds for startup)
```

- Experiment with strategies and OpenMP thread counts.
  - Strategies: MPI, OMP, scalability\_OMP

Installation Checklist

Program Instrumentation

Periscope Online Analysis

**Performance Properties Exploration**

- Start `Eclipse` with Periscope GUI from console

```
% eclipse
```

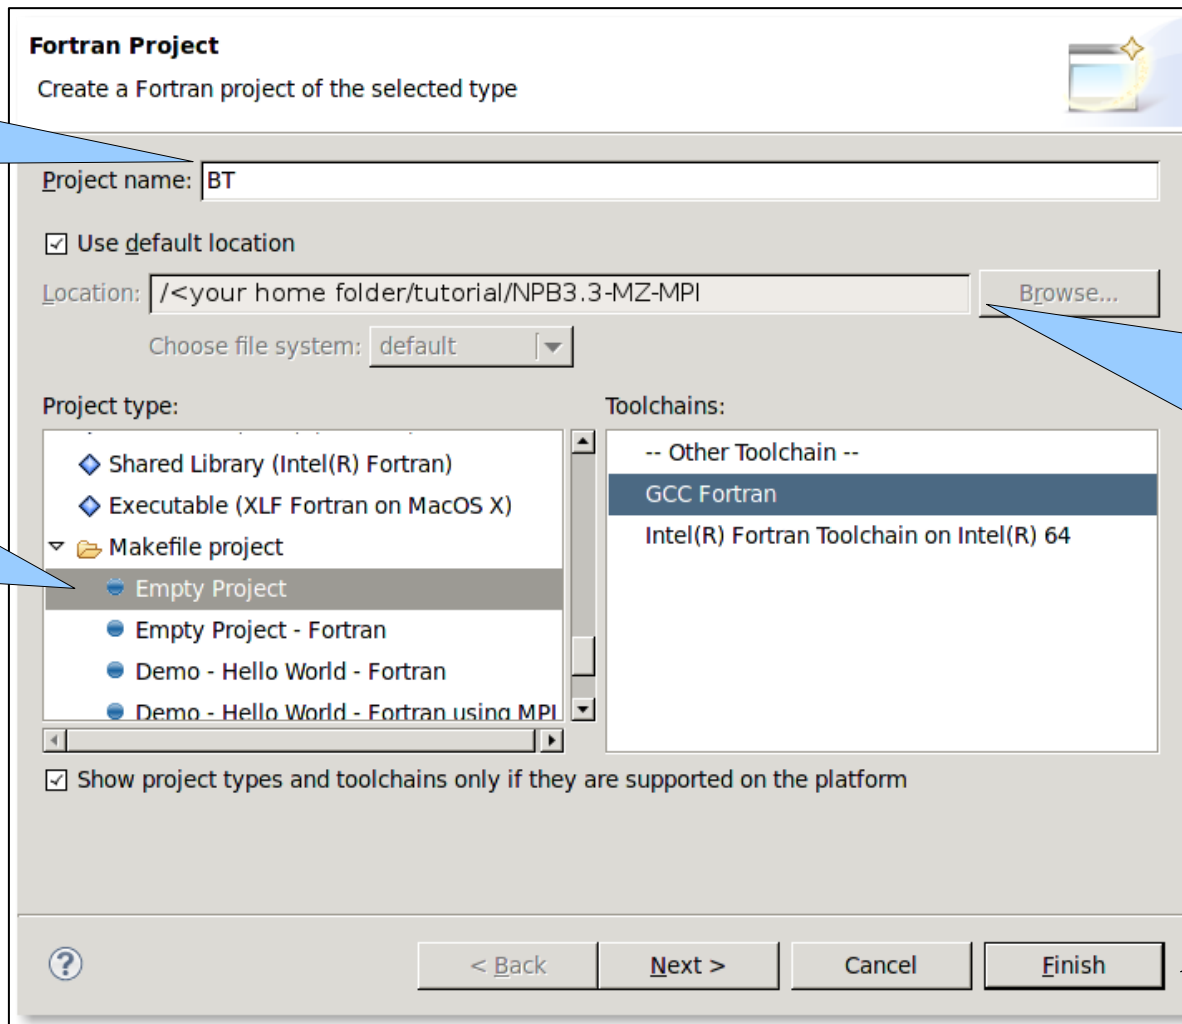
- Or by double-click on Eclipse pictogram on the Desktop



- File->New->Project... → Fortran->Fortran Project

Input project name

Project type



The screenshot shows the 'Fortran Project' dialog box with the following settings:

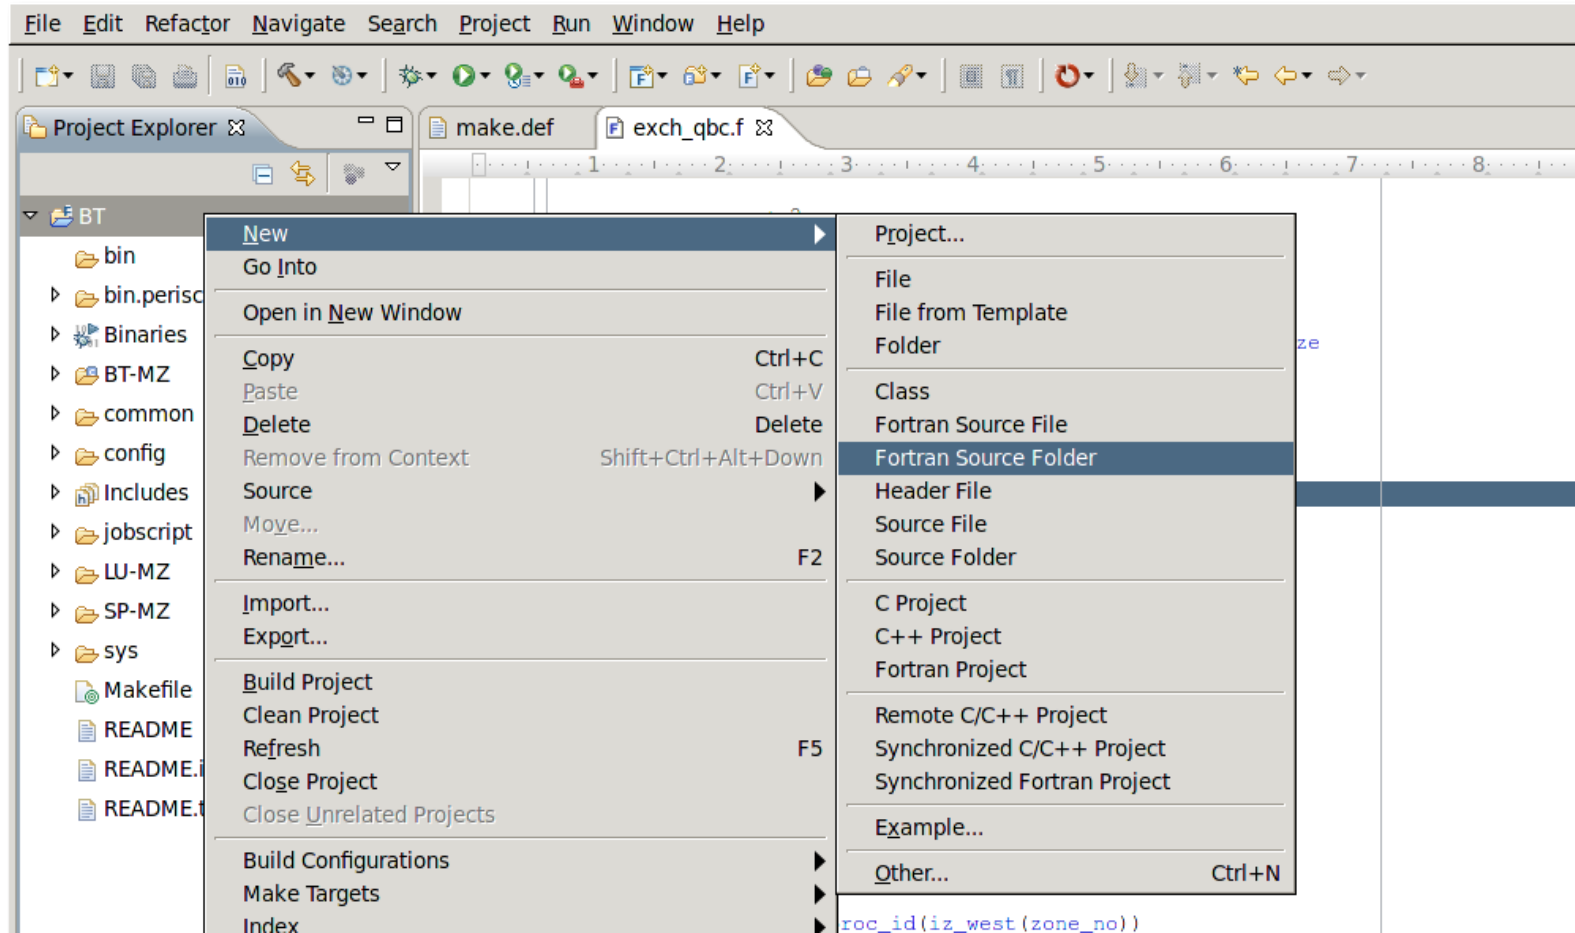
- Project name:** BT
- Use default location
- Location:** /<your home folder/tutorial/NPB3.3-MZ-MPI
- Choose file system:** default
- Project type:** Makefile project (expanded), Empty Project (selected)
- Toolchains:** GCC Fortran (selected)
- Show project types and toolchains only if they are supported on the platform

Buttons at the bottom: < Back, Next >, Cancel, Finish

Unmark "Use default location" and provide path to *BT* folder



Press Finish

- Right-click -> File-> New -> Fortran Source Folder



- Choose BT-MZ as a source folder


**Source folder**

  Exclusion patterns of 1 source folder(s) updated to solve nesting.

Project name:

Folder name:

Update exclusion filters in other source folders to solve nesting.



# Loading properties



File Edit Refactor Navigate Search Project Run Window Help

Project Explorer

- BT
  - bin
    - bin.periscope
      - bt-mz\_B.4
      - bt-mz\_B.4.sir
      - DebugOut
      - npb\_btmz\_358118.oe
      - npb\_btmz\_358124.oe
      - properties\_MPI\_43987.psc

```
nr = nr + 2
qoffset = qoffset + m_size
tag = tag + num_procs
end do
else if (c_size .lt. 0) then
write(*,*) 'error: integer overflow', myid, ip, c_size
call mpi_abort(MPI_COMM_WORLD, 1, ierror)
stop
endif

if (nr .gt. 0) then
call mpi_waitall(nr, requests, statuses, ierror)
endif
```

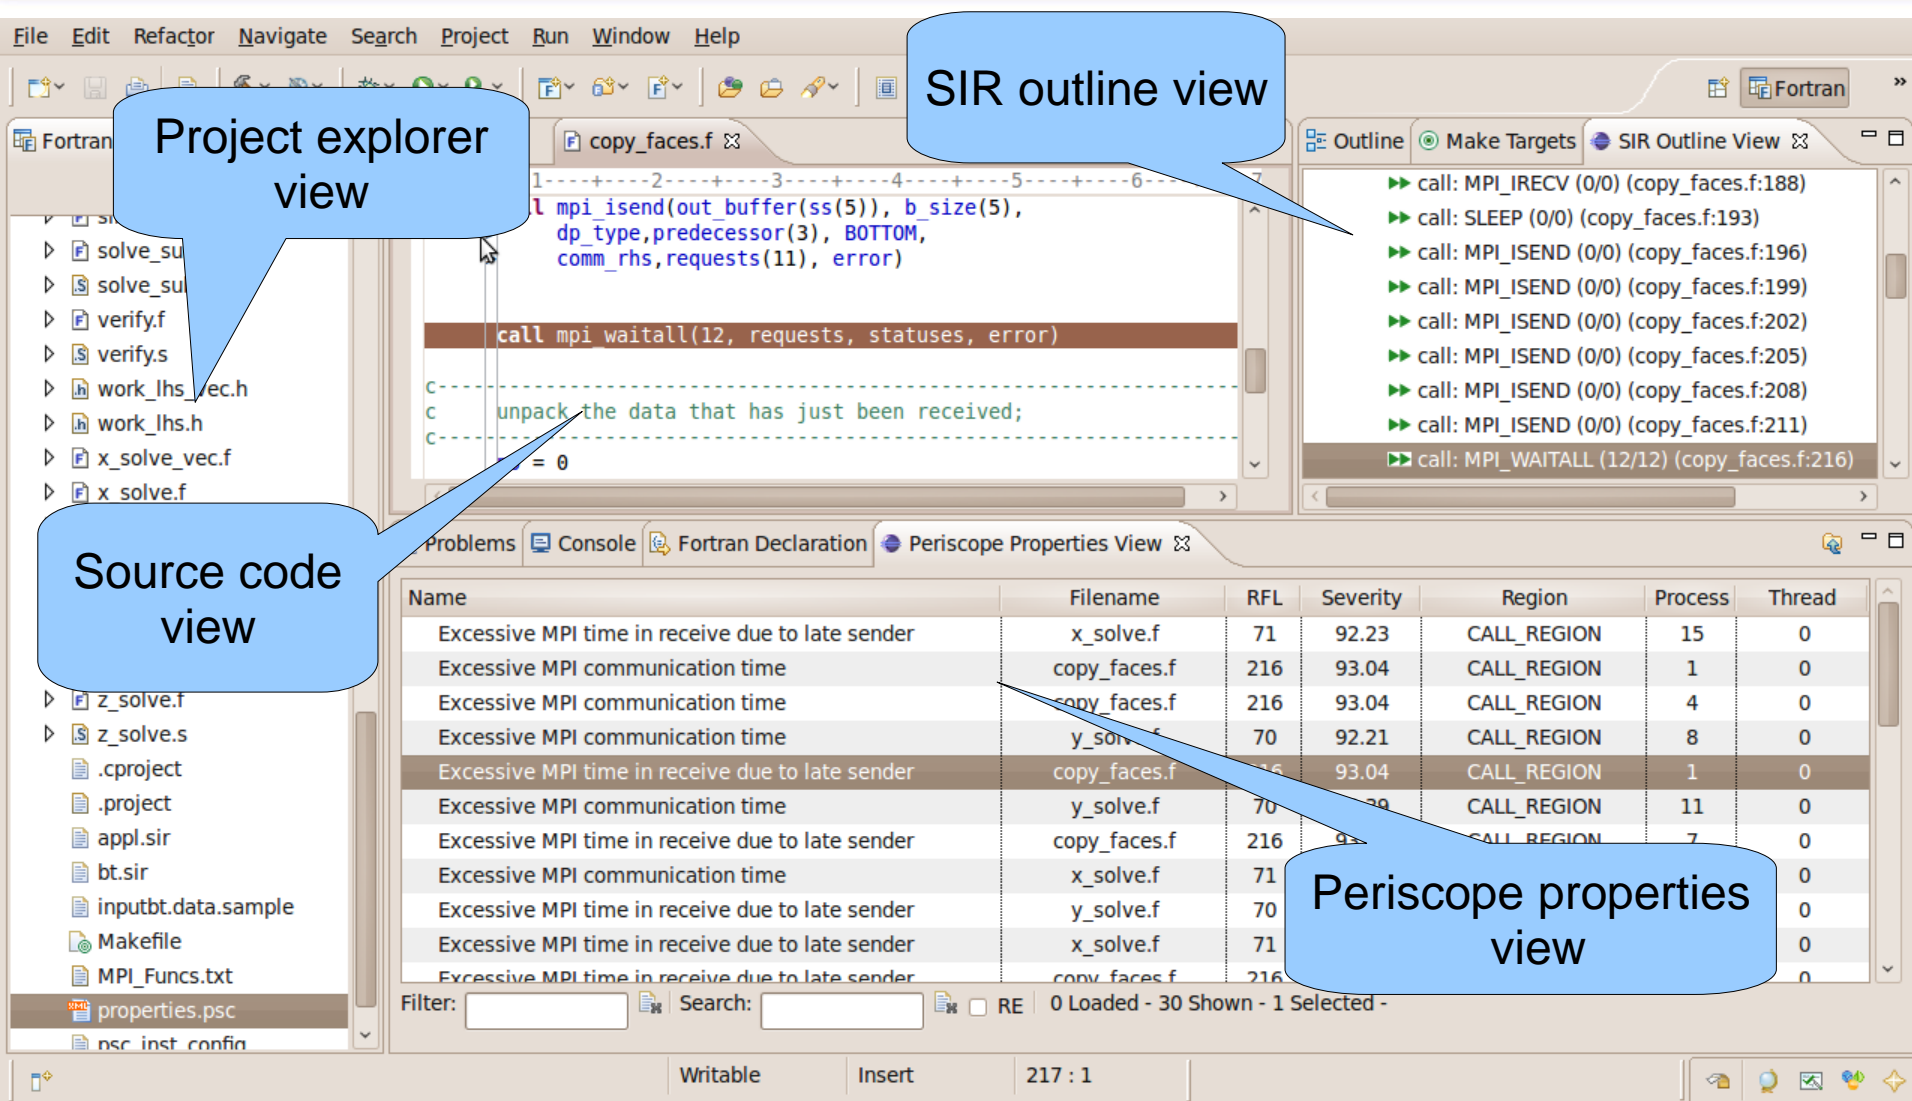
Context Menu:

- New
- Open F3
- Open With
- Copy Ctrl+C
- Paste Ctrl+V
- Delete Delete
- Remove from Context Shift+Ctrl+Alt+Down
- Mark as Landmark Shift+Ctrl+Alt+Up
- Move...
- Rename... F2
- Import...
- Export...
- Refresh F5
- Make Targets
- Validate
- Show in Remote Systems view
- Clean Selected File(s)
- Build Selected File(s)
- Run As
- Debug As
- Profile As
- Team
- Compare With
- Replace With
- Run C/C++ Code Analysis
- Periscope
  - Load all properties
  - Load properties above a severity
  - Load and cluster properties
- Properties Alt+Enter

Periscope Table View

name	RFL	Severity	Percentage	Region	Process
h_qbc.f	121	11.16	1.00	Call region	3
h_qbc.f	121	11.71	1.00	Call region	2
h_qbc.f	121	11.45	1.00	Call region	1
h_qbc.f	121	11.45	1.00	Call region	3

Expand BT project,  
search for \*.psc  
and  
Right click->Periscope->  
Load all properties



The screenshot displays the Periscope GUI interface with three callout boxes highlighting key features:

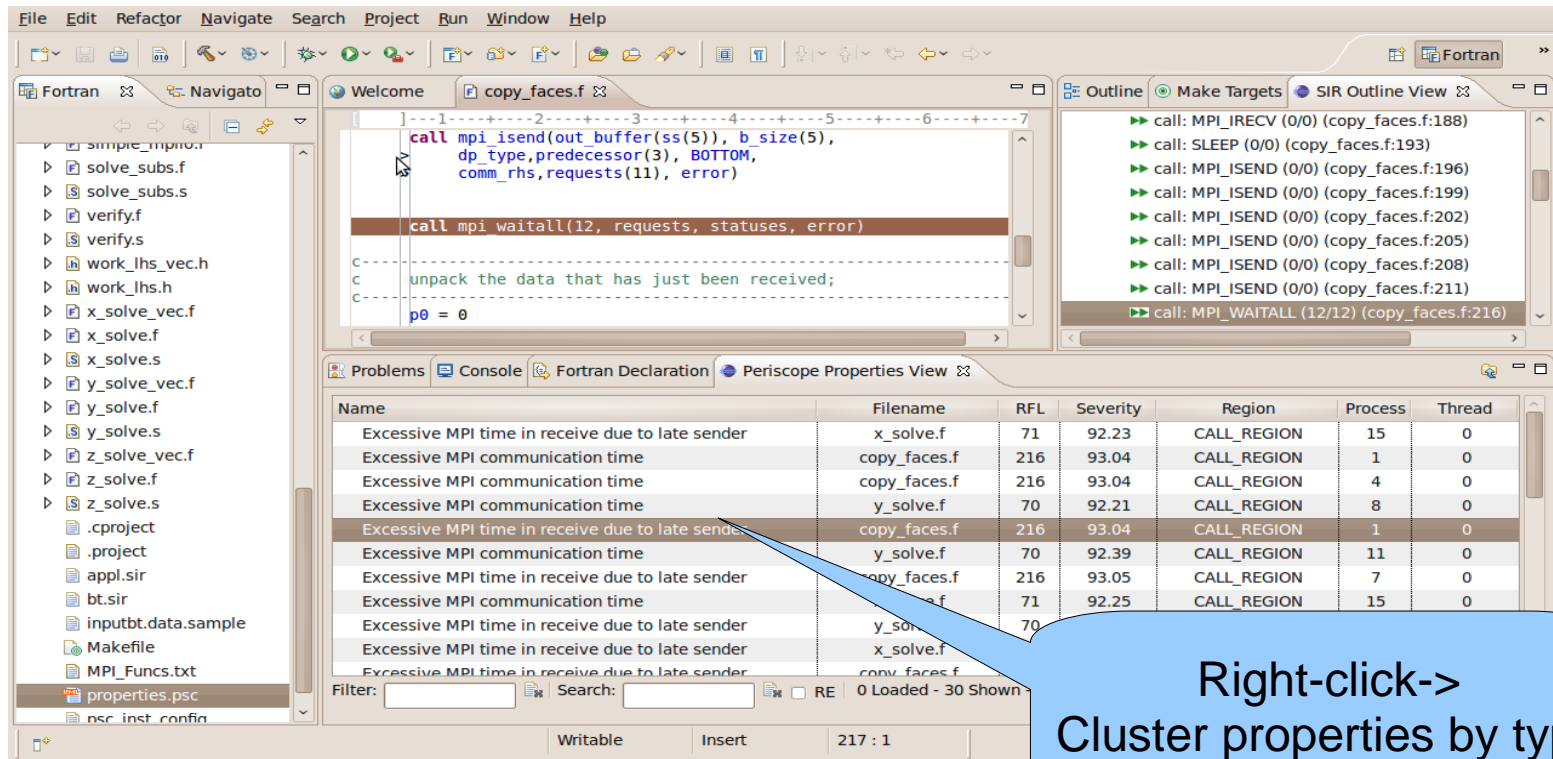
- Project explorer view:** Located on the left side, showing a tree view of project files including `solve_su`, `verify.f`, `work_lhs_vec.h`, `x_solve_vec.f`, `z_solve.f`, and `properties.psc`.
- Source code view:** The central editor window displays Fortran code from `copy_faces.f`, including `mpi_isend` and `call mpi_waitall` statements.
- SIR outline view:** A panel on the right showing a list of MPI-related function calls such as `call: MPI_IRecv`, `call: SLEEP`, and `call: MPI_WAITALL`.
- Periscope properties view:** A table at the bottom right showing performance metrics for MPI operations.

Name	Filename	RFL	Severity	Region	Process	Thread
Excessive MPI time in receive due to late sender	x_solve.f	71	92.23	CALL_REGION	15	0
Excessive MPI communication time	copy_faces.f	216	93.04	CALL_REGION	1	0
Excessive MPI communication time	copy_faces.f	216	93.04	CALL_REGION	4	0
Excessive MPI communication time	y_solve.f	70	92.21	CALL_REGION	8	0
Excessive MPI time in receive due to late sender	copy_faces.f	216	93.04	CALL_REGION	1	0
Excessive MPI communication time	y_solve.f	70	92.21	CALL_REGION	11	0
Excessive MPI time in receive due to late sender	copy_faces.f	216	93.04	CALL_REGION	7	0
Excessive MPI communication time	x_solve.f	71	92.23	CALL_REGION	0	0
Excessive MPI time in receive due to late sender	y_solve.f	70	92.21	CALL_REGION	0	0
Excessive MPI time in receive due to late sender	x_solve.f	71	92.23	CALL_REGION	0	0
Excessive MPI time in receive due to late sender	copy_faces.f	216	93.04	CALL_REGION	0	0



- Multi-functional table is used in the GUI for Eclipse for the visualization of bottlenecks
  - Multiple criteria sorting algorithm
  - Complex categorization utility
  - Searching engine using Regular Expressions
  - Filtering operations
  - Direct navigation from the bottlenecks to their precise source location using the default IDE editor for that source file type (e.g. CDT/Photran editor).
- SIR outline view shows a combination of the standard intermediate representation (SIR) of the analysed application and the distribution of its bottlenecks. The main goals of this view are to assist the navigation in the source code and attract developer's attention to the most problematic code areas.

- Clustering can effectively summarize displayed properties and identify a similar performance behaviour possibly hidden in the large amount of data



The screenshot shows an IDE with a Fortran project. The main editor displays code from `copy_faces.f`, including MPI send and waitall calls. The `Periscope Properties View` window is open, showing a table of performance properties. A blue callout bubble points to the table with the text "Right-click-> Cluster properties by type".

Name	Filename	RFL	Severity	Region	Process	Thread
Excessive MPI time in receive due to late sender	x_solve.f	71	92.23	CALL_REGION	15	0
Excessive MPI communication time	copy_faces.f	216	93.04	CALL_REGION	1	0
Excessive MPI communication time	copy_faces.f	216	93.04	CALL_REGION	4	0
Excessive MPI communication time	y_solve.f	70	92.21	CALL_REGION	8	0
Excessive MPI time in receive due to late sender	copy_faces.f	216	93.04	CALL_REGION	1	0
Excessive MPI communication time	y_solve.f	70	92.39	CALL_REGION	11	0
Excessive MPI time in receive due to late sender	copy_faces.f	216	93.05	CALL_REGION	7	0
Excessive MPI communication time	copy_faces.f	71	92.25	CALL_REGION	15	0
Excessive MPI time in receive due to late sender	y_solve.f	70				
Excessive MPI time in receive due to late sender	x_solve.f					
Excessive MPI time in receive due to late sender	copy_faces.f					

Right-click->  
Cluster properties by type

# Properties clustering



File Edit Refactor Navigate Search Project Run Window Help

Problems Console Fortran Declaration Periscope Properties View Clustering Results View

Name	Filename	RFL	Severity	Confidence	Processes	Threads	Clustering Error
call: MPI_WAIT (8) (y_solve.f:70)	y_solve.f	70	92.35	1.00	Regions Group		
Excessive MPI time in receive due to late send					Types Group		Clustering squared error: 0.13/0.50
Cluster 1					8 9		
Cluster 2					10 11		
Excessive MPI communication time (4)					Types Group		Clustering squared error: 0.17/0.50
Cluster 1	y_solve.f		92.45		10 11		
Cluster 2	y_solve.f	70	92.28		8 9		
call: MPI_WAITALL (12) (copy_faces.f:216)	copy_faces.f	216	93.01	1.00	Regions Group		
Excessive MPI time in receive due to late send	copy_faces.f	216			Types Group		Clustering squared error: 0.11/0.50
Cluster 1	copy_faces.f	216	92.98		3 12 13		
Cluster 2	copy_faces.f	216	93.04		1 7		
Excessive MPI communication time (6)	copy_faces.f	216			Types Group		Clustering squared error: 0.11/0.50
Cluster 1	copy_faces.f	216	92.98		3 1		
Cluster 2	copy_faces.f	216	93.04		1 4		
call: MPI_WAIT (x_solve.f:71)	x_solve.f	71	92.40	1.00	Regions		
Excessive MPI time in receive due to late send	x_solve.f	71			Types		Clustering squared error: 0.12/0.50
Cluster 1	x_solve.f	71	92.60		14		
Cluster 2	x_solve.f	71	92.34		2 5 6		
Excessive MPI communication time (6)	x_solve.f	71			Types		Clustering squared error: 0.13/0.50
Cluster 1	x_solve.f	71	92.36				
Cluster 2	x_solve.f	71	92.62				

Filter: Search: RE 0 Loaded - 21 Shown - 1 Selected -

Severity value of the Cluster 1

Region and property where clustering performed

Processes belonging To the Cluster1