Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Extrae & Paraver Hands-On

Judit Giménez, Germán Llort, Lau Mercadal, Sandra Méndez

tools@bsc.es

19/04/2021

POP Performance Analysis

Extrae features

- Platforms
  - Intel, Cray, BlueGene, MIC, ARM, Android, Fujitsu Sparc ...
- Parallel programming models
  - MPI, OpenMP, pthreads, OmpSs, CUDA, OpenCL, Java, Python ...
- Performance Counters
  - Using PAPI interface
- Link to source code
  - Callstack at MPI routines
  - OpenMP outlined routines
  - Selected user functions (Dyninst)
- Periodic sampling
- User events (Extrae API)

No need to recompile or relink!

How does Extrae work?

- Symbol substitution through LD_PRELOAD
  - Specific libraries for each combination of runtimes
    - MPI
    - OpenMP
    - OpenMP+MPI
    - …

Recommended

- Dynamic instrumentation
  - Based on Dyninst (developed by U.Wisconsin / U.Maryland)
    - Instrumentation in memory
    - Binary rewriting

- Alternatives
  - Static link (i.e., PMPI, Extrae API)

Using Extrae in 3 steps

1. **Adapt** your job submission scripts
2. **Configure** what to trace
   - XML configuration file
   - Example configurations at `$EXTRAE_HOME/share/example`
3. **Run** it!

- For further reference check the **Extrae User Guide:**
  - https://tools.bsc.es/doc/html/extrae
  - Also distributed with Extrae at `$EXTRAE_HOME/share/doc`

Log in to MareNostrum4

```
laptop$ ssh -Y <USER>@mn1.bsc.es
```

- The following directory contains all the examples:

```
mn$ cp -R /home/nct00/nct00003/tools-material
$HOME
mn$ ls -l $HOME/tools-material
bin/
clustering/
extrae/
slides/
traces/
```

Here you have a copy of this slides.
Copy them to your laptop
or open remotely with:

```
> evince slides/Extrae-Paraver-Hands-On.pdf
```

Step 1: Adapt the job script to load Extrae

```
mn$ vi $HOME/tools-material/extrae/job_27p.sh
```

```
#!/usr/bin/env bash

#SBATCH --job-name=lulesh2.0_i_27p
#SBATCH --output=lulesh2.0_i_27p.out
#SBATCH --error=lulesh2.0_i_27p.err
#SBATCH --ntasks=27
#SBATCH --cpus-per-task=1
#SBATCH --time=00:10:00
#SBATCH --exclusive
#SBATCH --reservation=Training21

export OMP_NUM_THREADS=1

srun --cpu-bind=cores ../bin/lulesh2.0_i -i 10 -s
65 -p
```

Request resources

Run the program

Step 1: Adapt the job script to load Extrae

```
mn$ vi $HOME/tools-material/extrae/job_27p.sh
```

```
#!/usr/bin/env bash

#SBATCH --job-name=lulesh2.0_i_27p
#SBATCH --output=lulesh2.0_i_27p.out
#SBATCH --error=lulesh2.0_i_27p.err
#SBATCH --ntasks=27
#SBATCH --cpus-per-task=1
#SBATCH --time=00:10:00
#SBATCH --exclusive
#SBATCH --reservation=Training21

export OMP_NUM_THREADS=1

srun --cpu-bind=cores ./trace.sh ../bin/lulesh2.0_i
…
```

Run with Extrae

Step 1: Adapt the job script to load Extrae

```
mn$ vi $HOME/tools-material/extrae/trace.sh
```

```
#!/usr/bin/env bash

#SBATCH --job-name=lulesh2.0_i_27p
#SBATCH --output=lulesh2.0_i_27p.out
#SBATCH --error=lulesh2.0_i_27p.err
#SBATCH --ntasks=27
#SBATCH --cpus-per-task=1
#SBATCH --time=00:10:00
#SBATCH --exclusive
#SBATCH --reservation=Training21

export OMP_NUM_THREADS=1

srun --cpu-bind=cores ./trace.sh ../bin/lulesh2.
…
```

```
#!/usr/bin/env bash

module load gcc/7.2.0 EXTRAE/3.8.3
export TRACE_NAME=lulesh2.0_27p.prv

# Configure Extrae
export EXTRAE_CONFIG_FILE=extrae.xml

# Load the tracing library (choose C/Fortran)
export
LD_PRELOAD=$EXTRAE_HOME/lib/libmpitrace.so
#export
LD_PRELOAD=$EXTRAE_HOME/li

# Run the program
$*
```

Load Extrae

What to trace?

Type of application

Step 1: Which tracing library?

- Choose depending on the application type

Library	Serial	MPI	OpenMP	pthread	CUDA
libseqtrace	x				
libmpitrace[f][1]		x			
libomptrace			x		
libpttrace				x	
libcudatrace					x
libompitrace[f] [1]		x	x		
libptmpitrace[f] [1]		x		x	
libcudampitrace[f] [1]		x			x

[1] include suffix "f" in Fortran codes

Step 3: Run it!

- Submit your job

```
mn$ cd $HOME/tools-material/extrae
mn$ sbatch job_27p.sh
```

- Easy!

Step 2: Extrae XML configuration

```
mn$ vi $HOME/tools-material/extrae/extrae.xml
```

```
<mpi enabled="yes">
  <counters enabled="yes" />
</mpi>

<openmp enabled="no">
  <locks enabled="no" />
  <counters enabled="yes" />
</openmp>

<pthread enabled="no">
  <locks enabled="no" />
  <counters enabled="yes" />
</pthread>

<callers enabled="yes">
  <mpi enabled="yes">1-3</mpi>
  <sampling enabled="no">1-5</sampling>
</callers>
```

Trace the MPI calls
(What's the program doing?)

Trace the call-stack
(Where in my code?)

Step 2: Extrae XML configuration (II)

```
mn$ vi $HOME/tools-material/extrae/extrae.xml
```

```
<counters enabled="yes">
    <cpu enabled="yes" starting-set-distribution="1">
      <set enabled="yes" domain="all" changeat-time="0">
        PAPI_TOT_INS,PAPI_TOT_CYC
      </set>
    </cpu>
    <network enabled="no" />
    <resource-usage enabled="no" />
    <memory-usage enabled="no" />
  </counters>
```

Select which HW counters are measured
(How's the machine doing?)

Step 2: Extrae XML configuration (III)

```
mn$ vi $HOME/tools-material/extrae/extrae.xml
```

```
<buffer enabled="yes">
  <size enabled="yes">500000</size>
  <circular enabled="no" />
</buffer>

<sampling enabled="no" type="default" period="50m" variability="10m" />

<merge enabled="yes"
  synchronization="default"
  tree-fan-out="16"
  max-memory="512"
  joint-states="yes"
  keep-mpits="yes"
  sort-addresses="yes"
  overwrite="yes">
  $TRACE_NAME$
</merge>
```

Trace buffer size
(Flush/memory trade-off)

Enable sampling
(Want more details?)

Automatic post-processing to generate the Paraver trace

All done! Check your resulting trace

- Once finished (check with “squeue”) you will have the trace (3 files):

```
mn$ ls -l $HOME/tools-material/extrae
lulesh2.0_i_27p.pcf
lulesh2.0_i_27p.prv
lulesh2.0_i_27p.row
```

- Any trouble? Traces already generated here:

```
mn$ ls $HOME/tools-material/traces
```

- Now let’s look into it !

Install Paraver

- Download from https://tools.bsc.es/downloads

Pick your version

wxparaver-4.9.2-win.zip

wxparaver-4.9.2-mac.zip

wxparaver-4.9.2-Linux_i686.tar.gz (32-bits)
wxparaver-4.9.2-Linux_x86_64.tar.gz (64-bits)

```
laptop$ scp <USER>@mn1.bsc.es:tools-packages/<PKG>
```

- Also available @MN4
  - $HOME/tools-packages

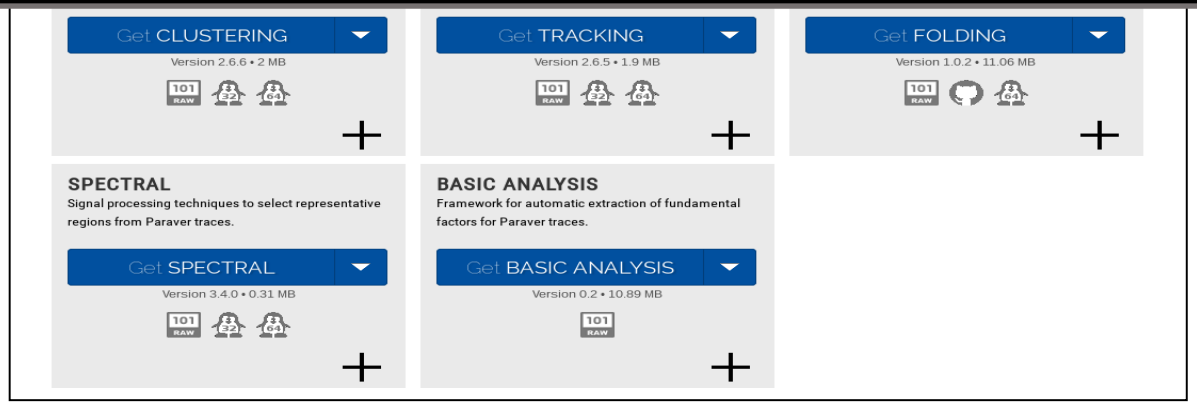

Install Paraver (II)

```
laptop$ tar xf wxparaver-4.9.2-linux-x86_64.tar.gz
laptop$ mv wxparaver-4.9.2-linux-x86_64 paraver
```

- Start Paraver

```
laptop$ paraver/bin/wxparaver
```

- Trouble installing locally? Remote open from MareNostrum4

```
mn$ module load paraver
mn$ wxparaver
```

Install Paraver tutorials (I)

- Download tutorials:

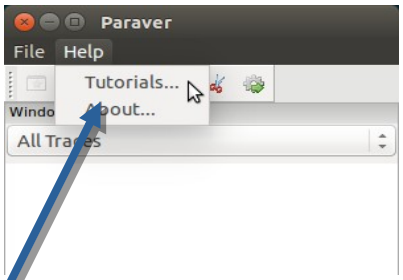

Click on Help → Tutorials

Tutorials window will pop-up

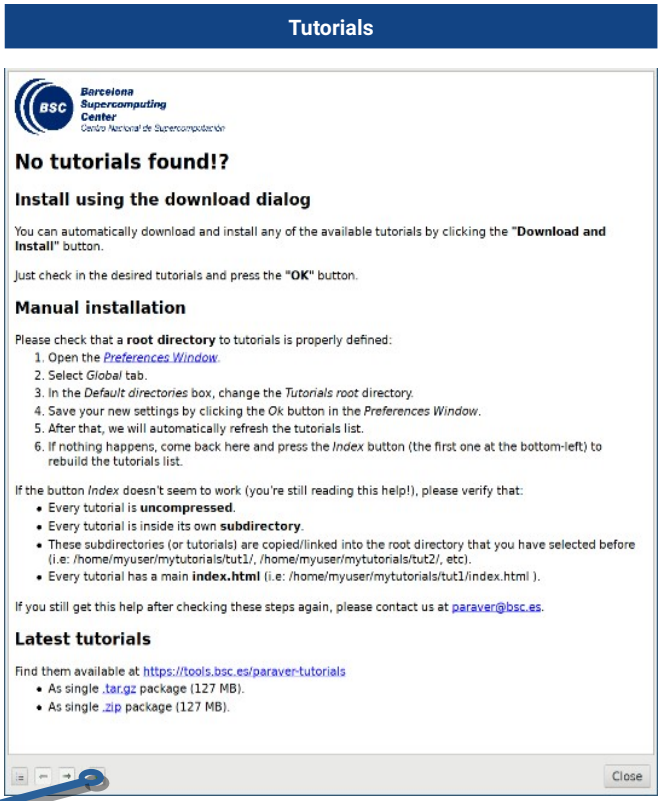

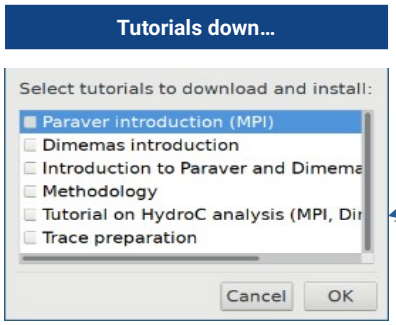

Follow these tutorials by clicking on the hyperlinks and reading the explanations. When you click on a link, multiple views will open.

Install Paraver tutorials – alternative methods(II)

- Download tutorials archive
  - https://tools.bsc.es/paraver-tutorials

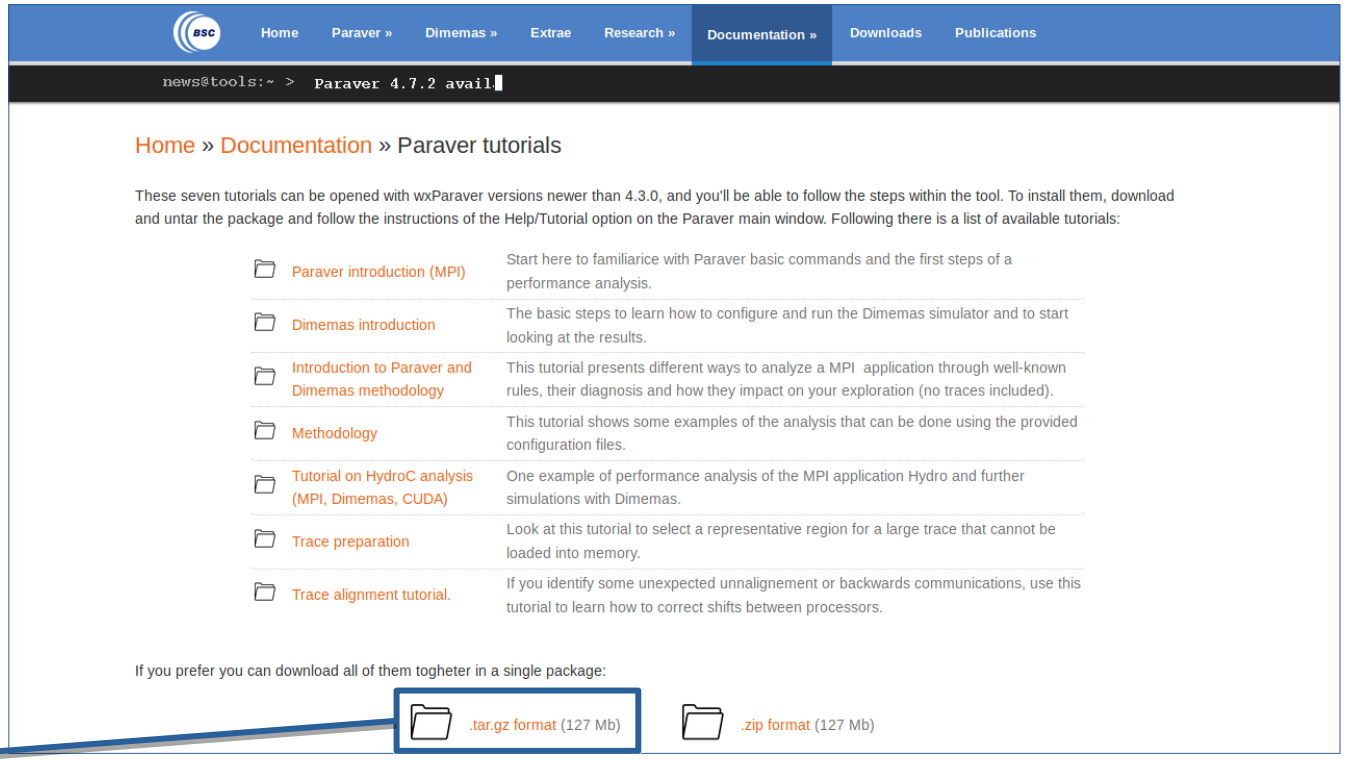

All tutorials

```
paraver-tutorials-20150526.tar.gz
```

Install Paraver tutorials – alternative methods(III)

- Start Paraver:
  - Linux: Run the command:

```
laptop$ paraver/bin/wxparaver
```

  - Windows: Double-click on paraver/wxparaver.exe
  - MAC: Double click on paraver/wxparaver.app

- Open File → Preferences

Setup the “Tutorials root” pointing to your folder “tutorials”

Click Browse and select your folder “tutorials”

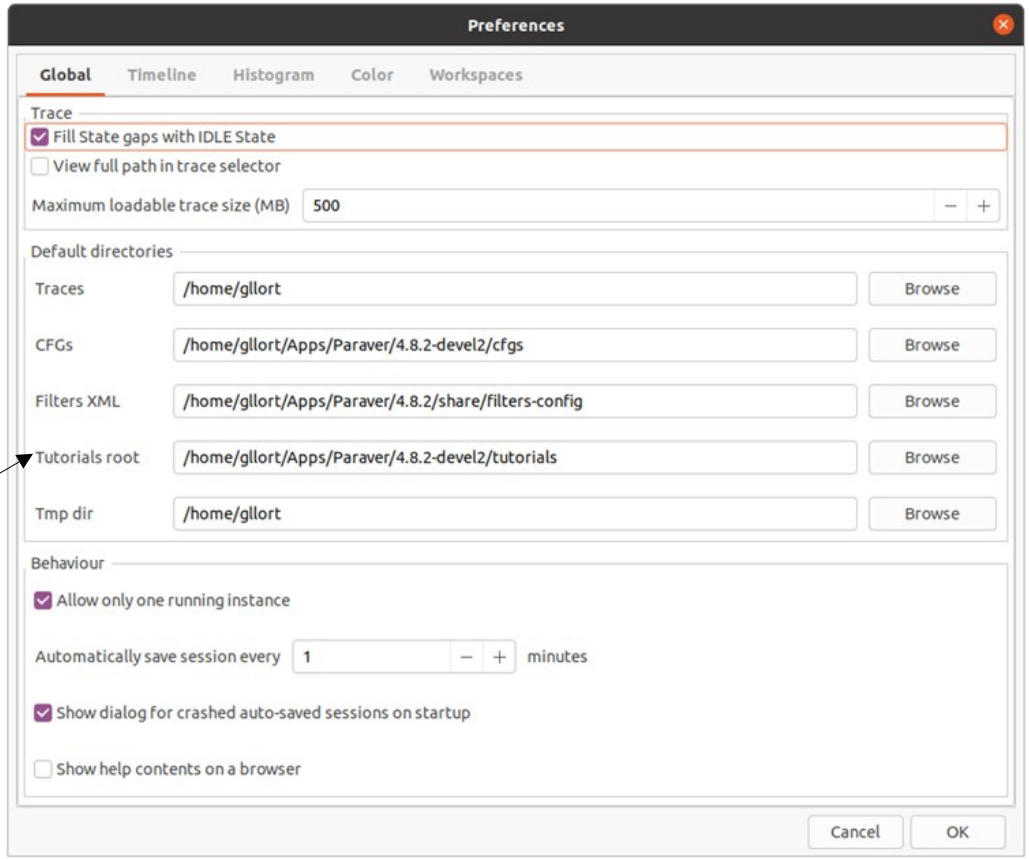

First steps of analysis

- Copy the trace to your laptop

```
laptop$ scp <USER>@mn1.bsc.es: \
tools-material/extrae/lulesh2.0_i_27p.* $HOME
```

- Load the trace with Paraver

Click on File → Load Trace → Browse to "lulesh2.0_i_27p.prv"

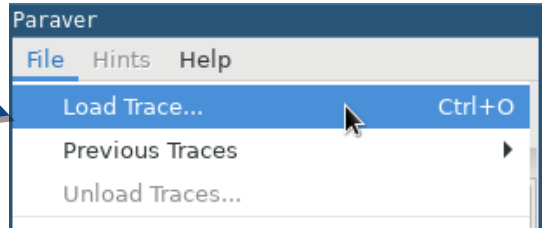

First steps of analysis

- Follow Tutorial #3
  - Introduction to Paraver and Dimemas methodology

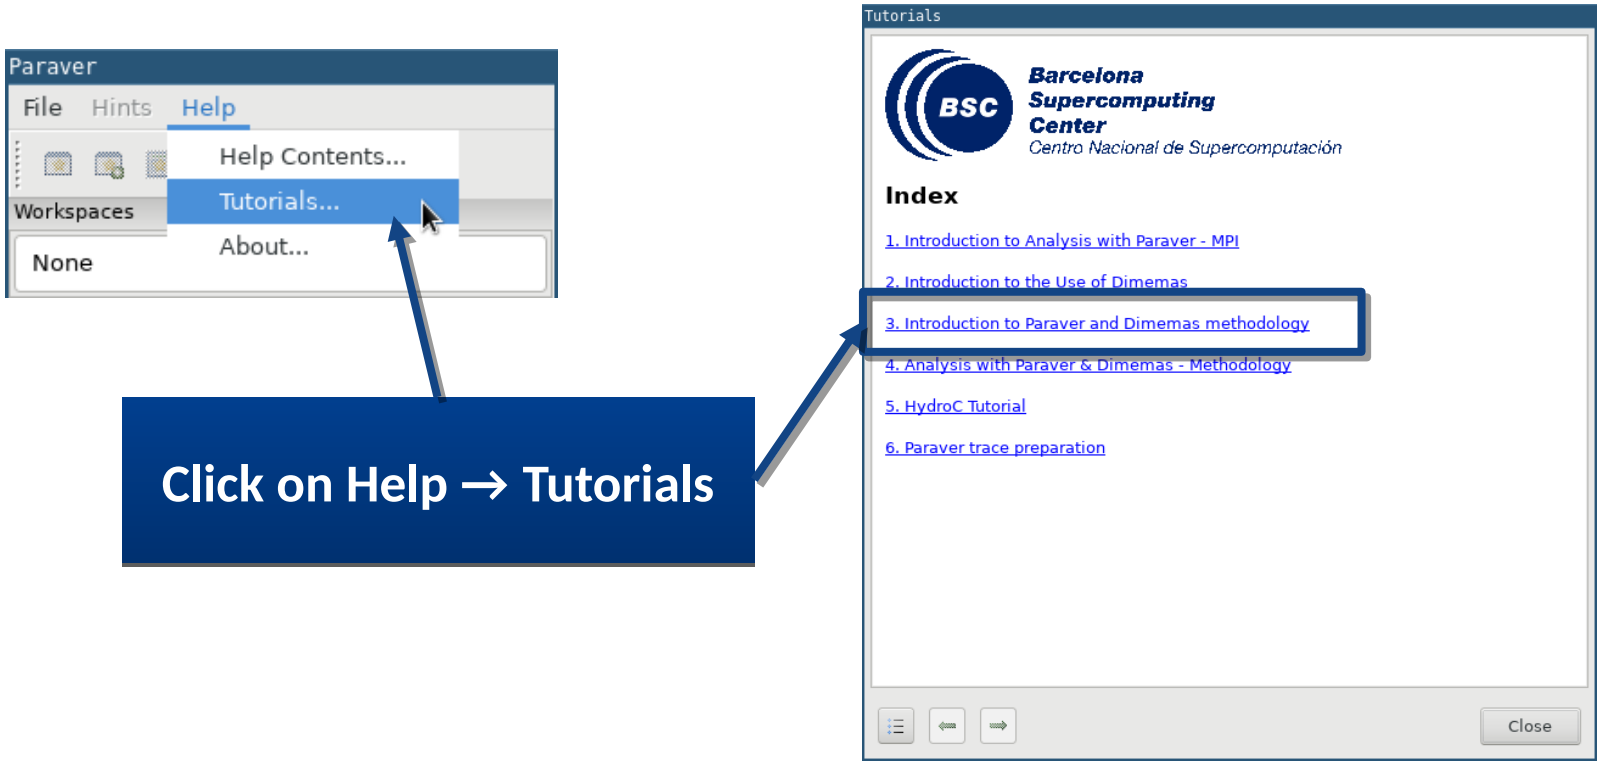

Measure the parallel efficiency

- Click on “mpi_stats.cfg”
  - Check the **Average** for the column labeled “**Outside MPI**”

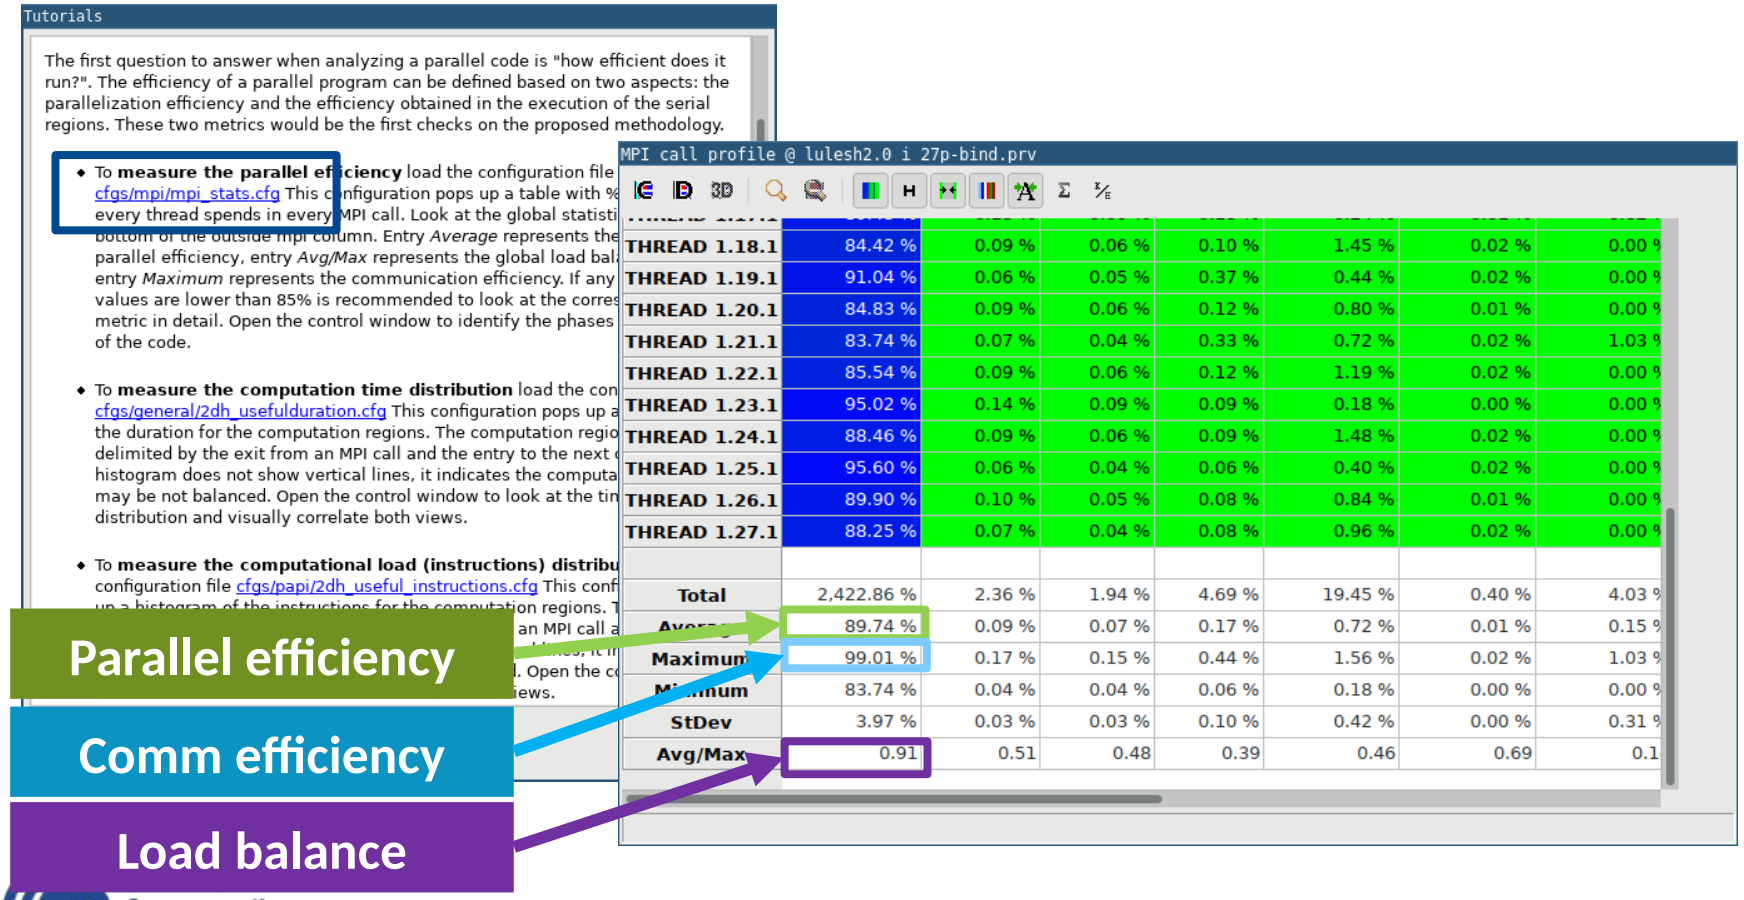

Computation time distribution

- Click on “2dh_usefulduration.cfg” (2nd link) Shows **time computing**

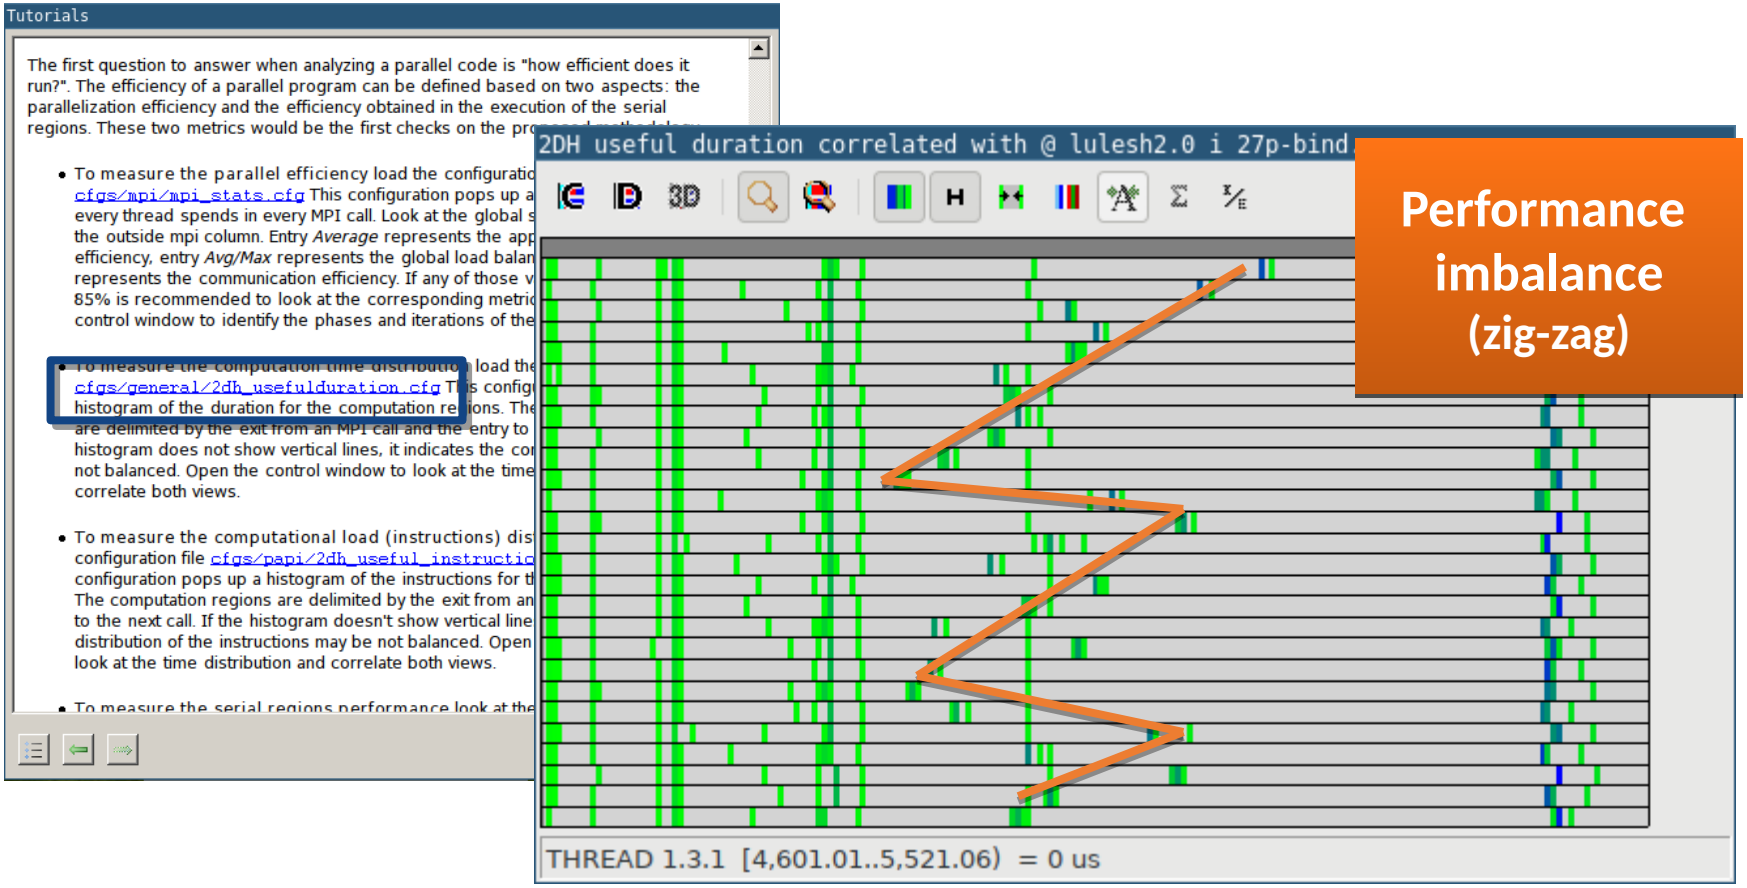

Computation load distribution

- Click on “2dh_useful_instructions.cfg” (3rd link) → Shows **amount of work**

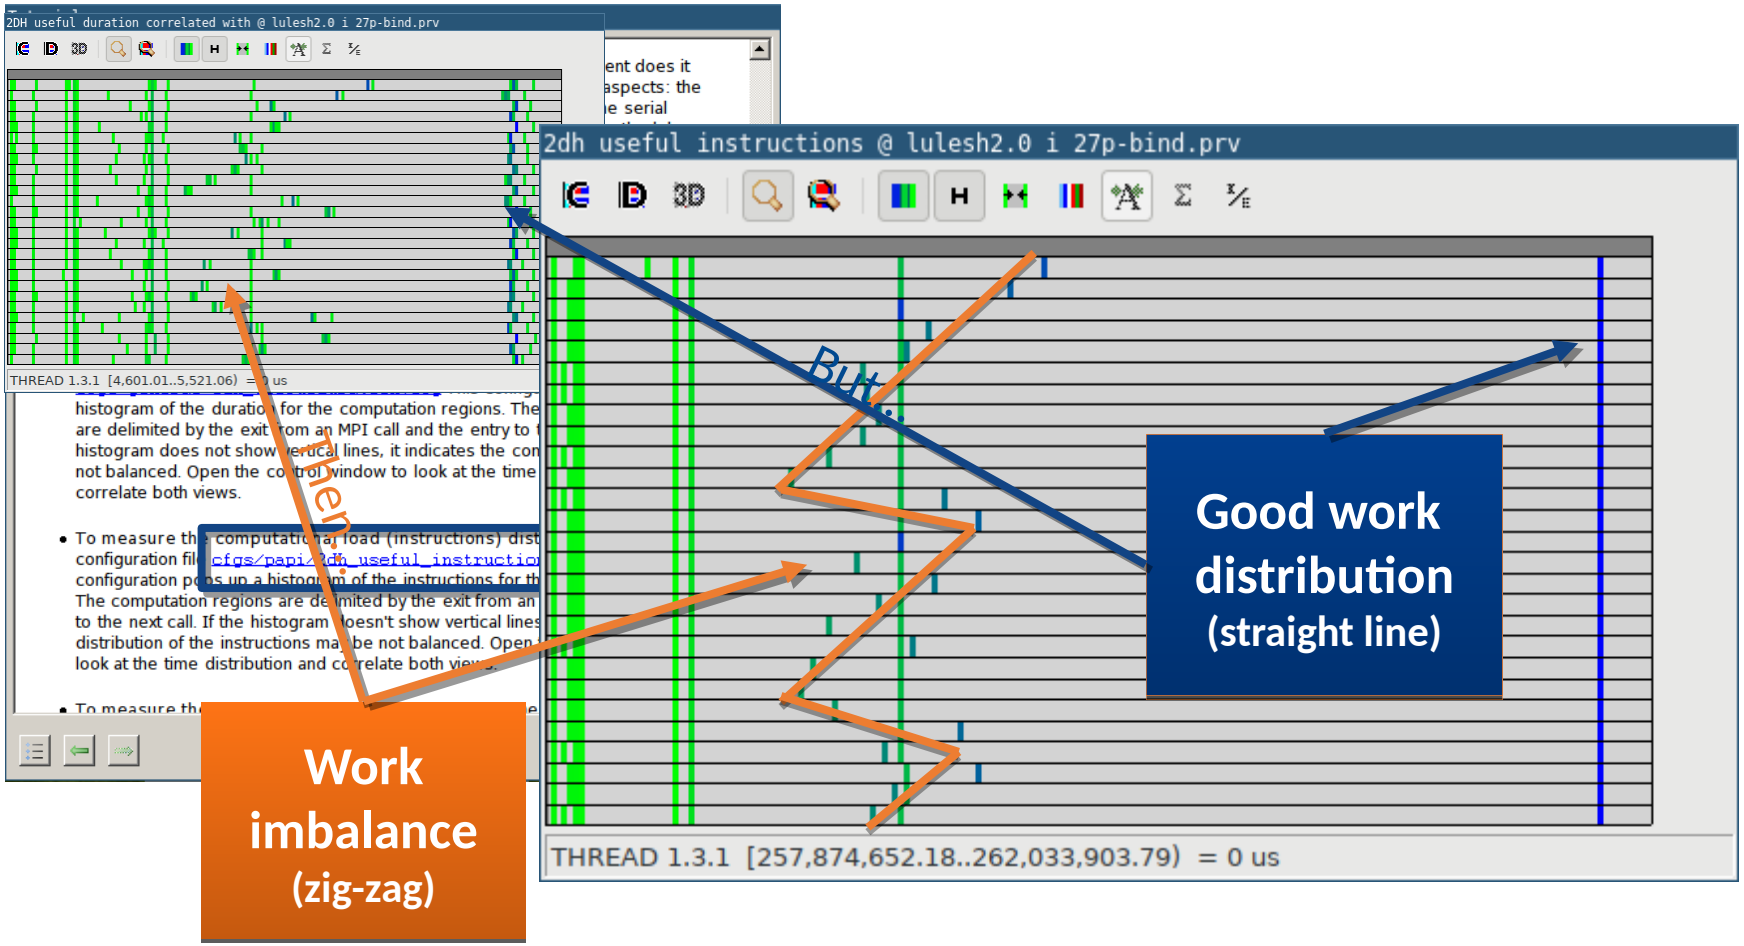

Where does this happen?

- Go from the table to the timeline

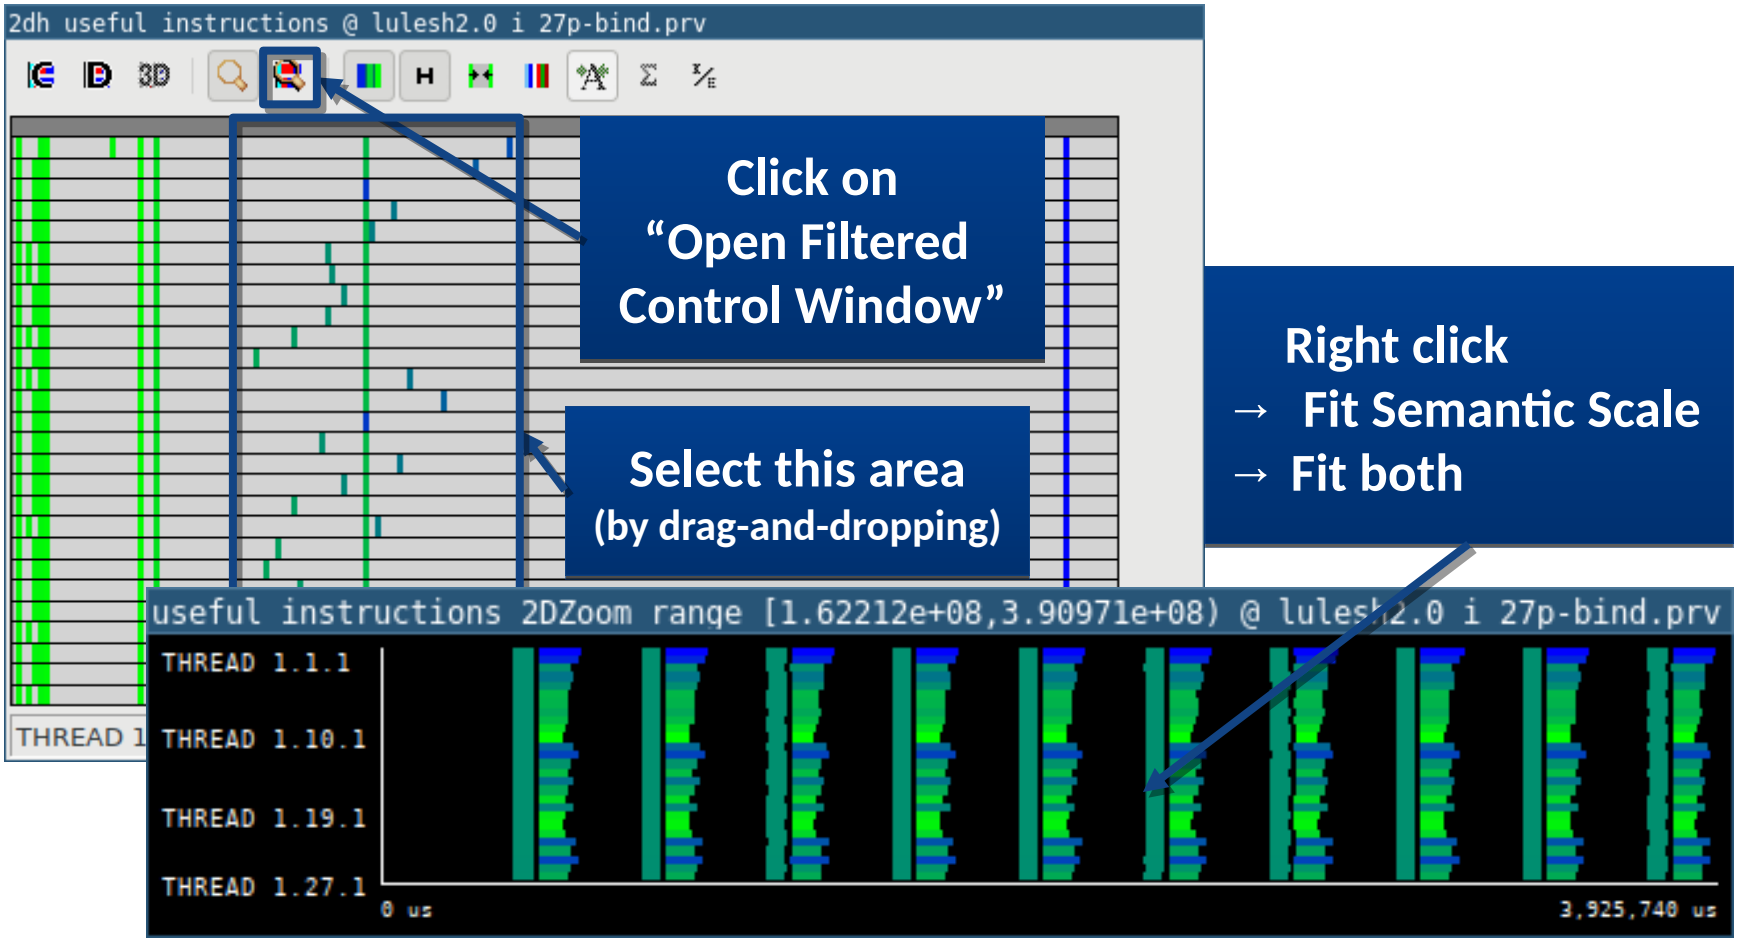

Where does this happen?

- **Slow** & **Fast** at the same time? ⇒ Imbalance

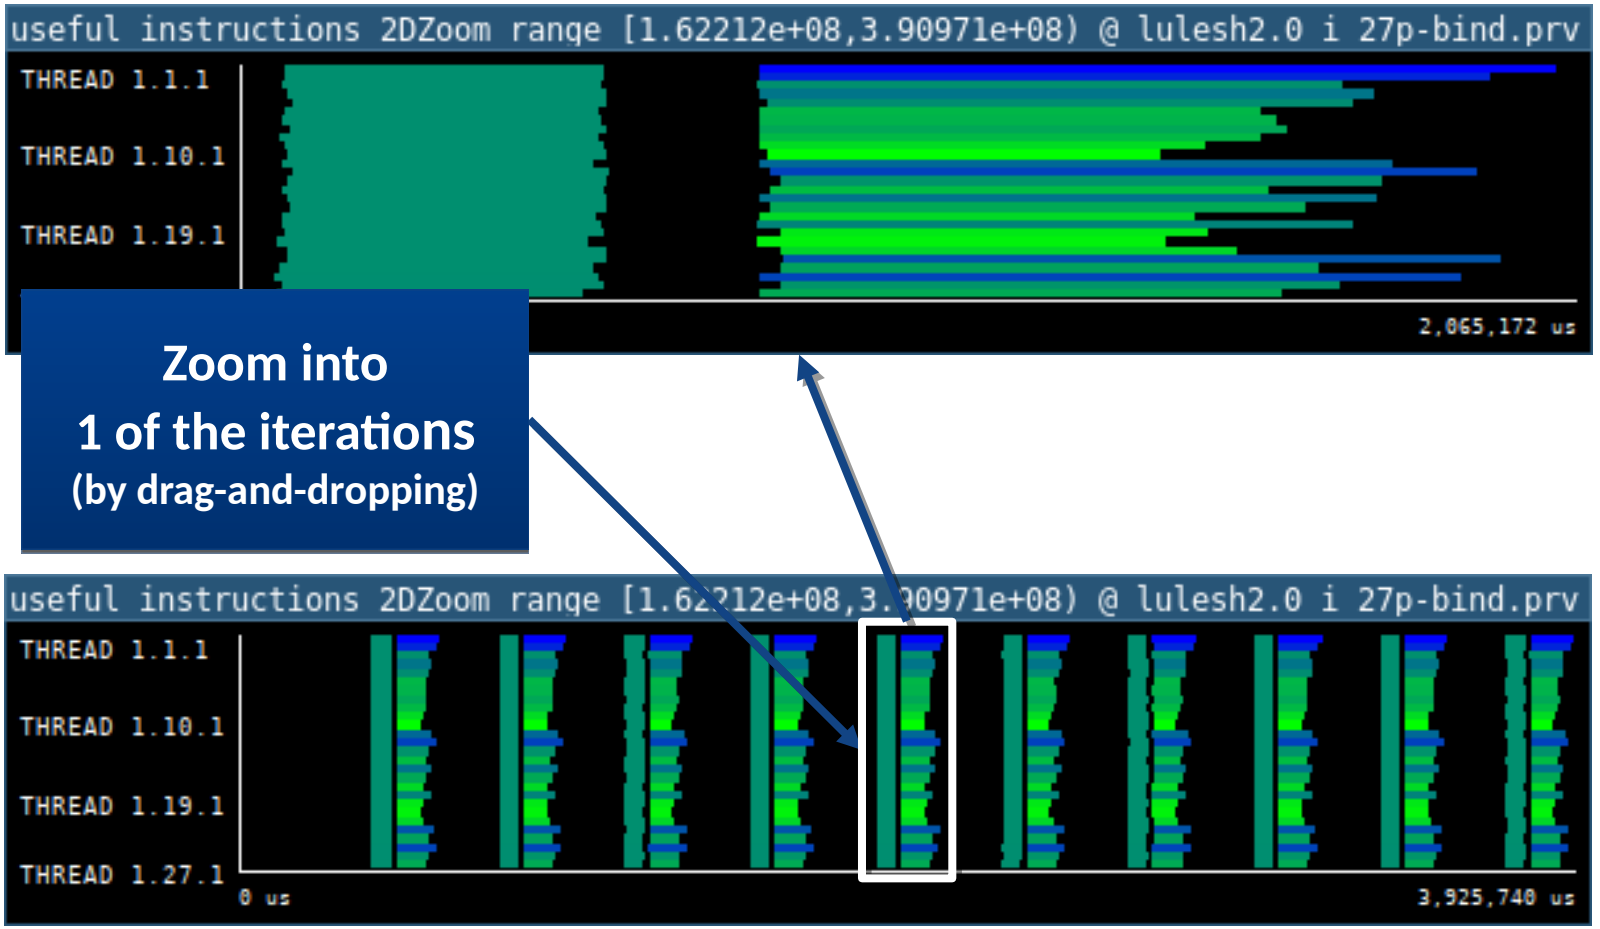

Where does this happen?

- Hints → Call stack references → Caller function

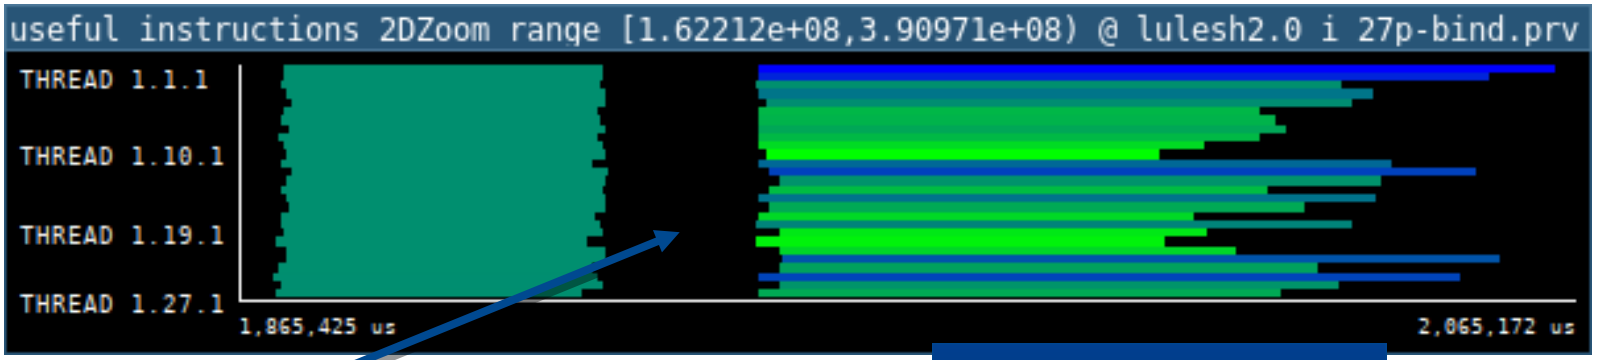

Right click → Copy

Right click → Paste → Time

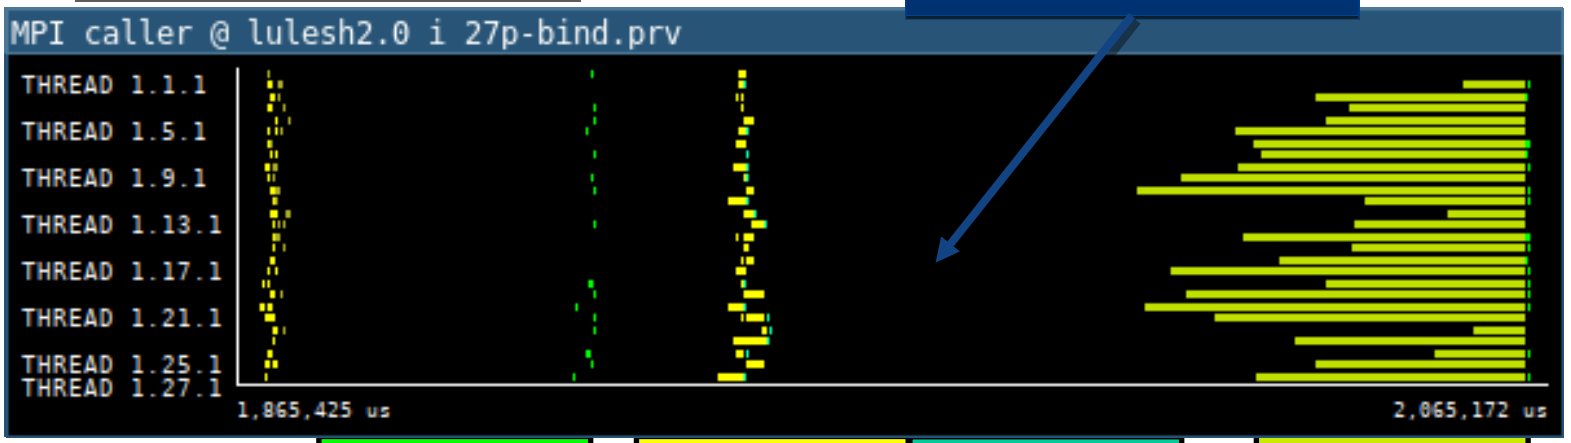

CommRecv | CommSend | CommMonoQ | TimeIncrement

Save CFG's (2 methods)

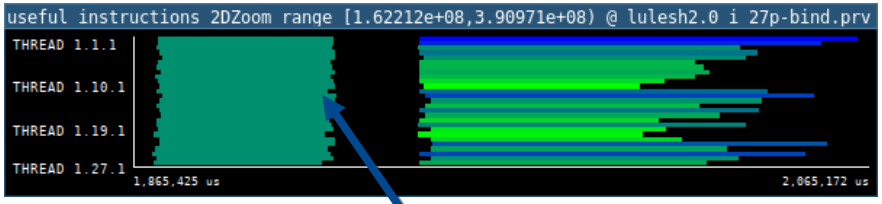

Right click on timeline

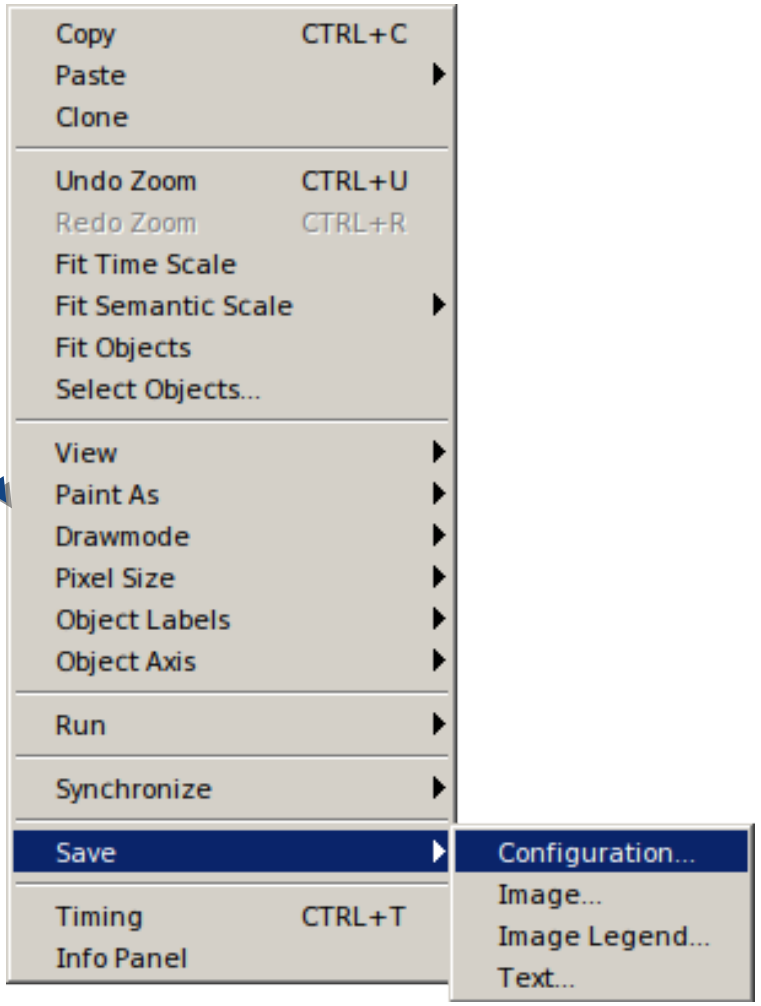

Save CFG's (2 methods)

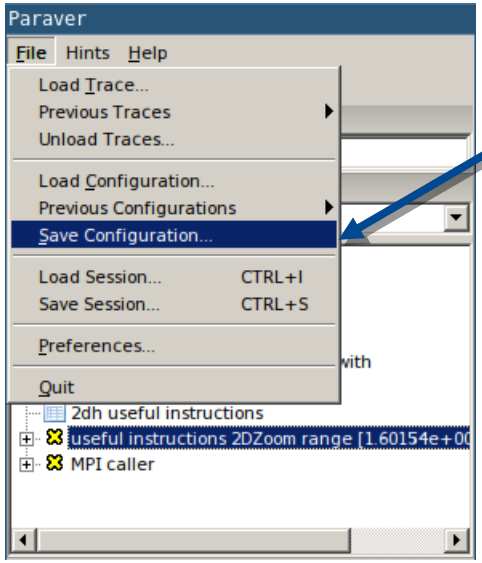

1. Main Paraver window

2. Select

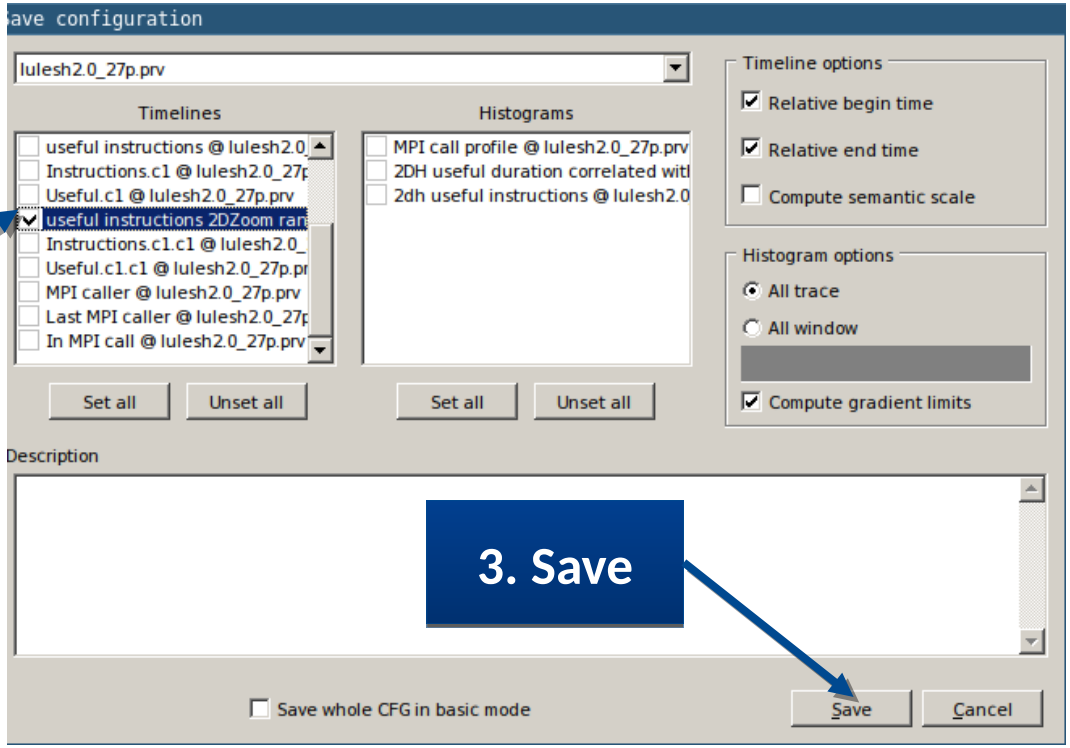

3. Save

CFG's distribution

- Paraver comes with many more included CFG's

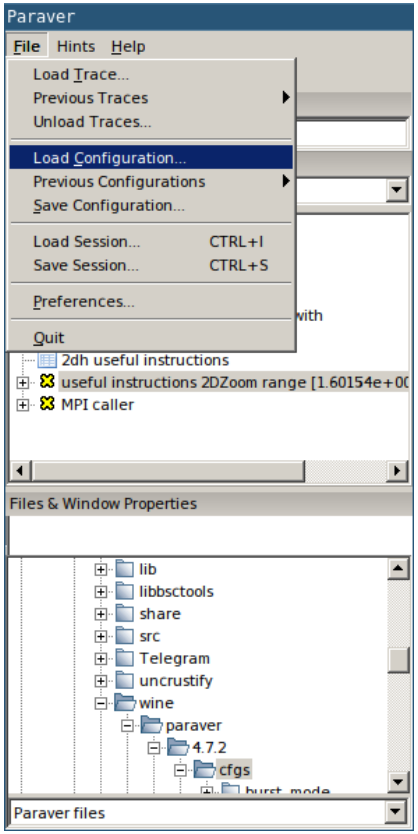

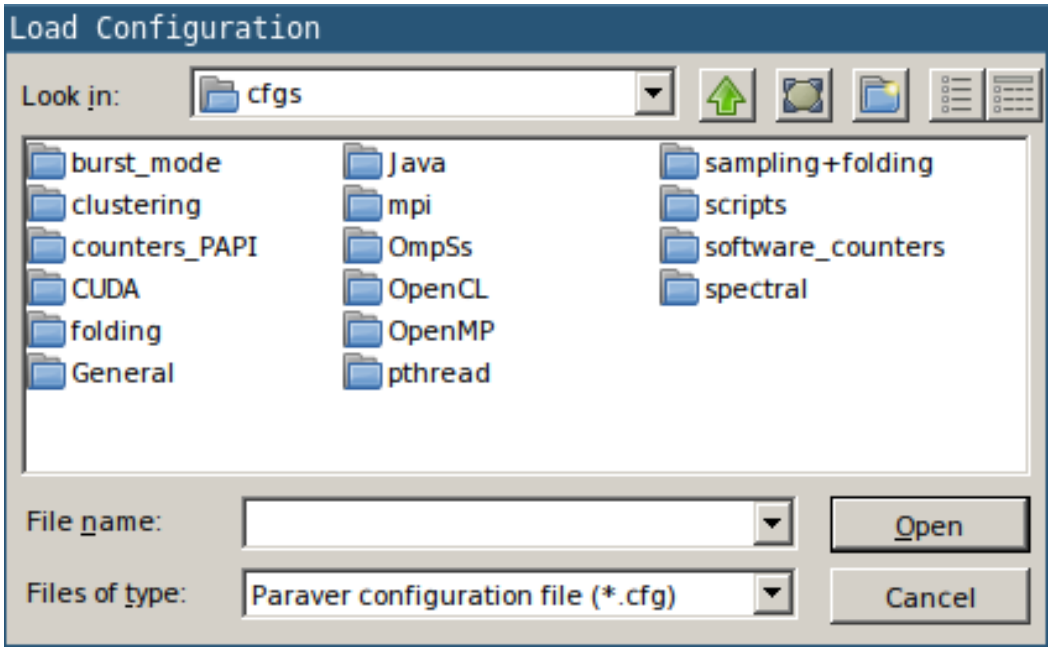

Hints: a good place to start!

- Paraver suggests CFG's based on the contents of the trace

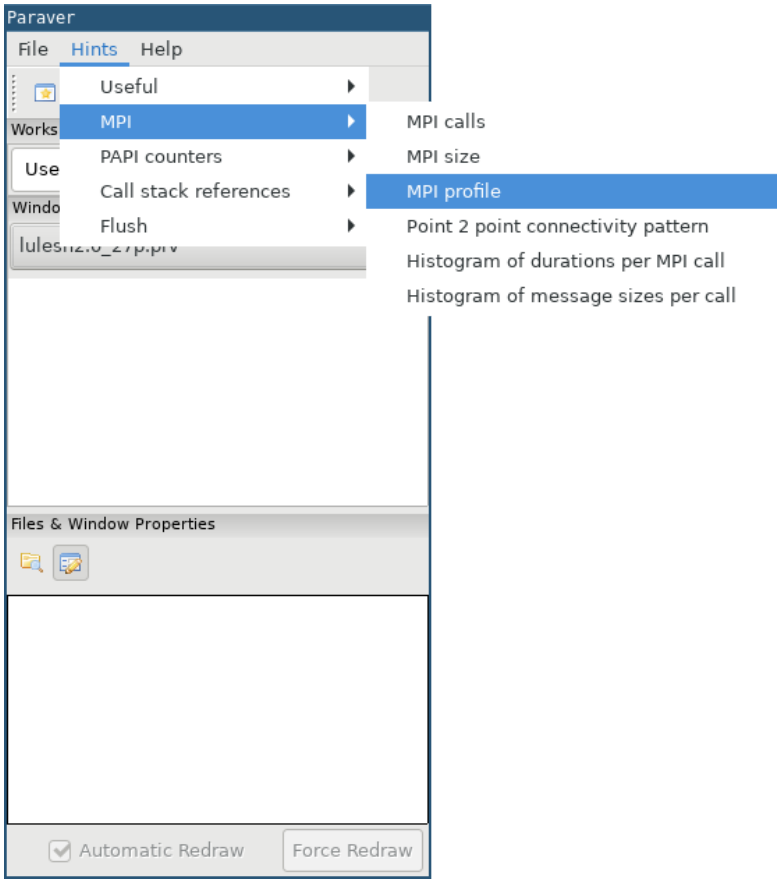

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Takeaway: Analyse efficiencies and unbalances

Judit Giménez, Germán Llort, Lau Mercadal, Sandra Méndez

tools@bsc.es

19/04/2021

POP Performance Analysis

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Clustering Hands-On

Judit Giménez, Germán Llort, Lau Mercadal, Sandra Méndez

tools@bsc.es

19/04/2021

POP Performance Analysis

Cluster-based analysis (I)

- Run the clustering tool on the trace you generated

```
mn$ module load clustering_suite

mn$ cd $HOME/tools-material/clustering

mn$ BurstClustering \
-d cluster.xml \
-i ../extrae/lulesh2.0_i_27p.prv \
-o lulesh2.0_i_27p_clustered.prv
```

- If you didn't get your own trace, use a prepared one from:

```
mn$ ls
$HOME/tools-material/traces/lulesh2.0_i_27p.prv
```

Cluster-based analysis (II)

- Check the clustering scatter plot

```
mn$ gnuplot \
lulesh2.0_i_27p_clustered.IPC.PAPI_TOT_INS.gnuplo
    t
```

- Identify main computing trends
- Work (Y) vs. Performance (X)
- Look at the clusters shape
  - Variability in both axes

  indicate **potential imbalances**

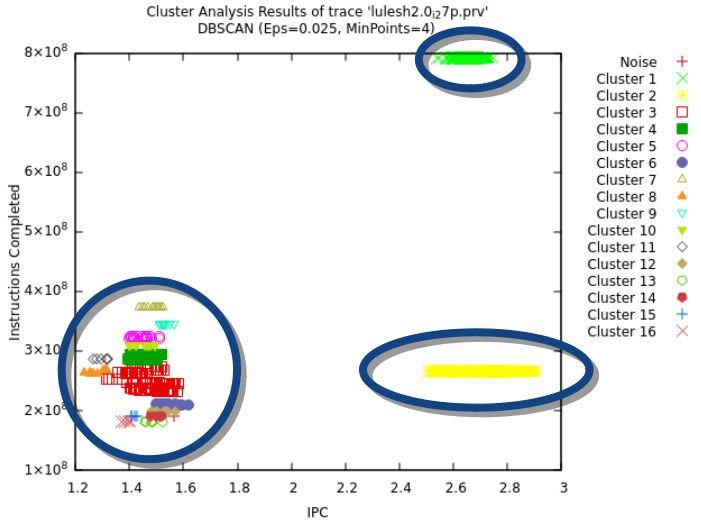

Cluster-based analysis (III)

- Check the clustered trace
  - Copy the trace to your laptop

```
mn$ scp <USER>@mn1.bsc.es: \
tools-material/clustering/*.{pcf,prv,row} $HOME
```

  - Load with Paraver

```
laptop$ paraver/bin/wxparaver \
$HOME/lulesh2.0_i_27p_clustered.prv
```

    - File Load configuration paraver/cfgs/clustering/clusterID_window.cfg

Cluster-based analysis (III)

- Correlate scatter plots & timelines to detect imbalances

Variable work
and/or
Variable speed
+
Simultaneously @ different processes
=
Imbalances