



Trace Management and Analysis with Framesoc

hand on tutorial + demo

Generoso Pagano

generoso.pagano@inria.fr

Inria Bordeaux Seminary, 27/02/2015

Agenda



Getting Started




Trace Analysis



Distribution



Perspectives



Getting Started

Installation in short



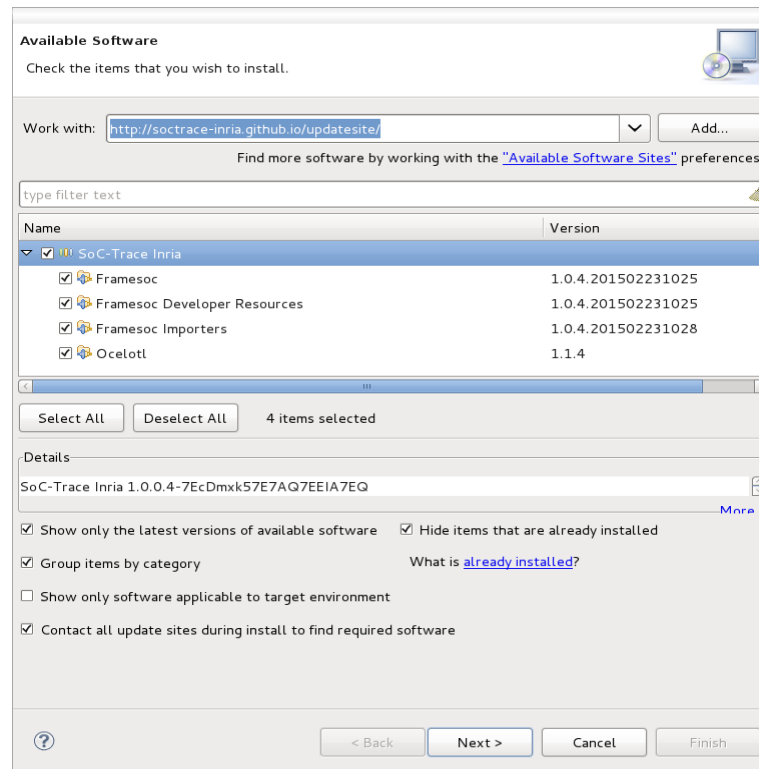
- Installation via Eclipse **update site** mechanism
- *Help > Install New Software... > Work with:*
<http://soctrace-inria.github.io/updatesite/>
- Then follow the automatic wizards and... that's it!

Installation details (1)

- Java
 - Install **JRE 7** or later
- Eclipse
 - Download **Kepler** version
<https://www.eclipse.org/downloads/packages/eclipse-standard-432/keplersr2>
 - Extract the Eclipse archive
 - Launch the **eclipse** executable in the extracted folder

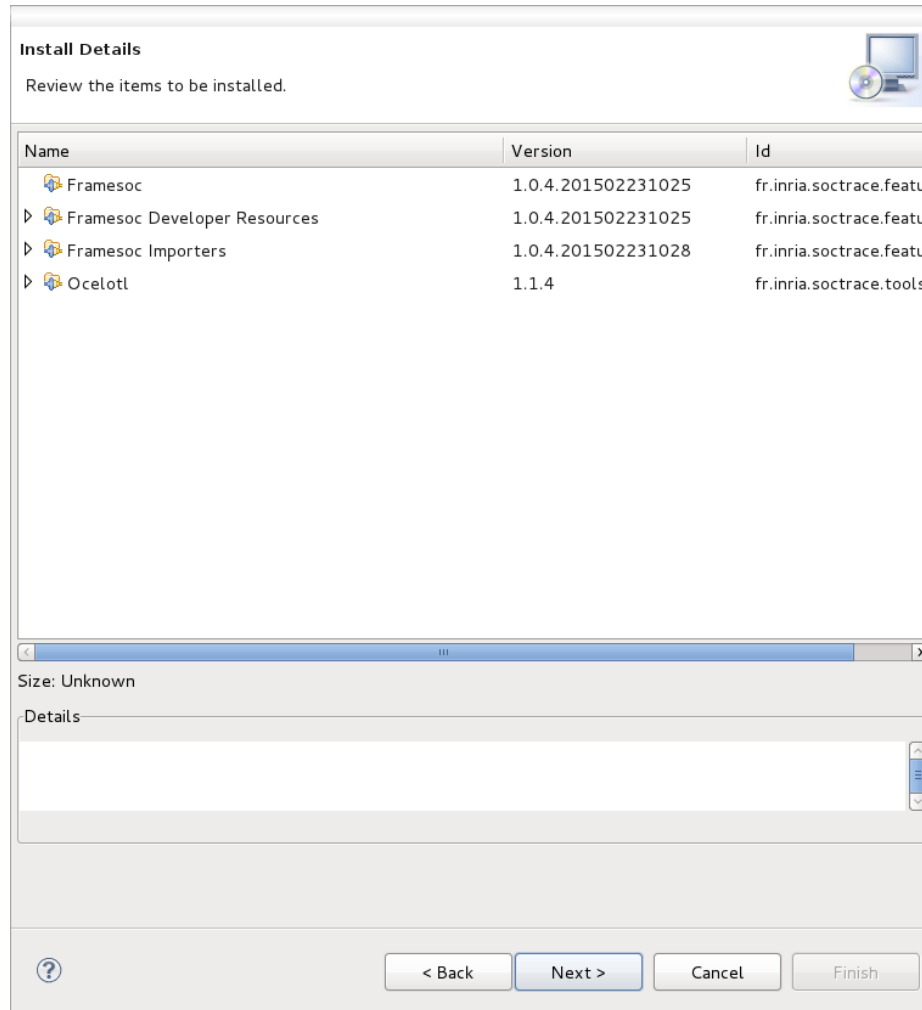
Installation details (2)

- From the Eclipse main menu:
 - **Help > Install New Software...**



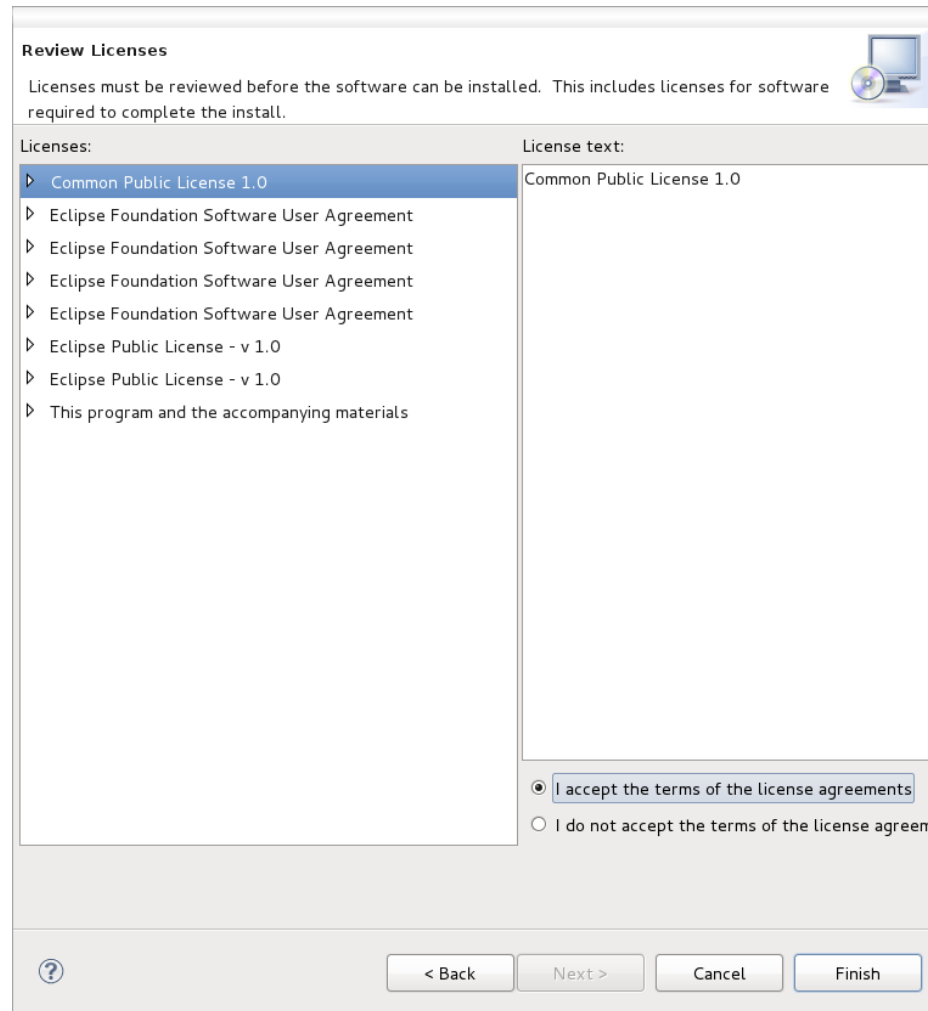
- **Work with:** <http://soctrace-inria.github.io/updatesite/>
- **Select SoC-Trace Inria modules, then press Next.**

Installation details (3)



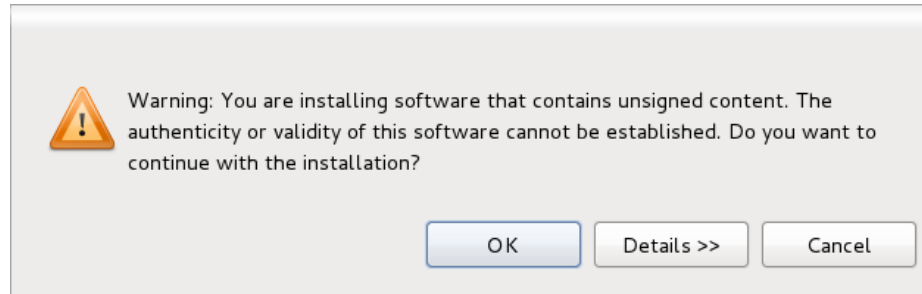
– Press **Next** again.

Installation details (4)

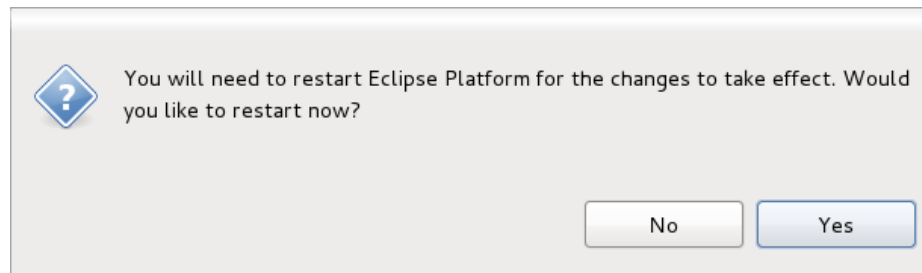


– **Accept** the license agreements and press **Next** again.

Installation details (5)



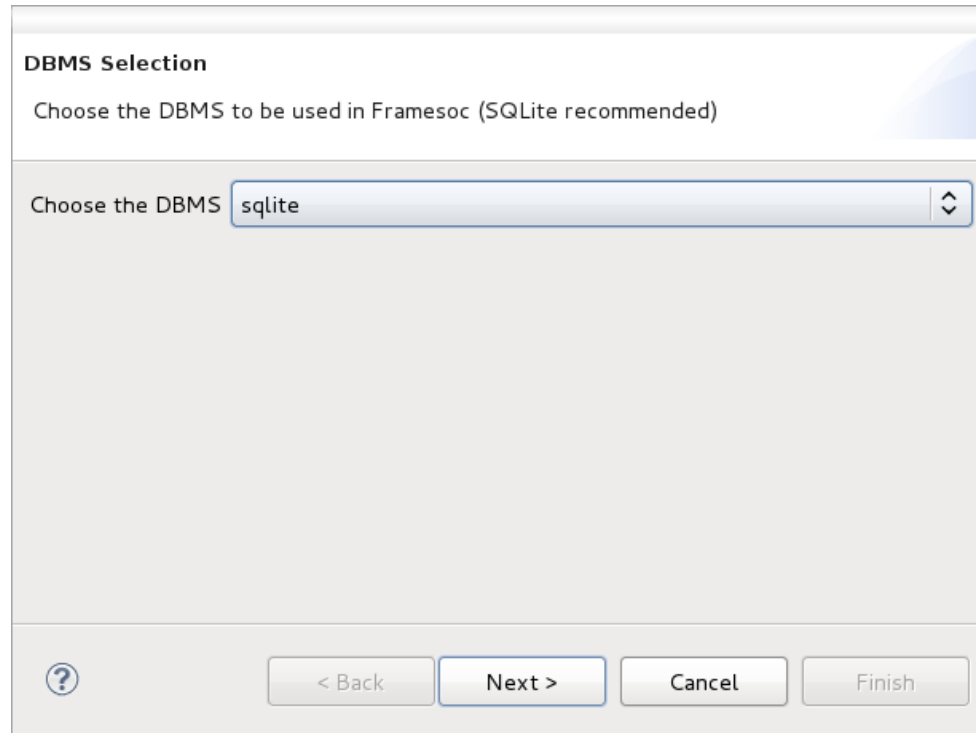
– Press **OK** in the warning dialog.



– Press **Yes** to restart Eclipse.

Initialization (1)

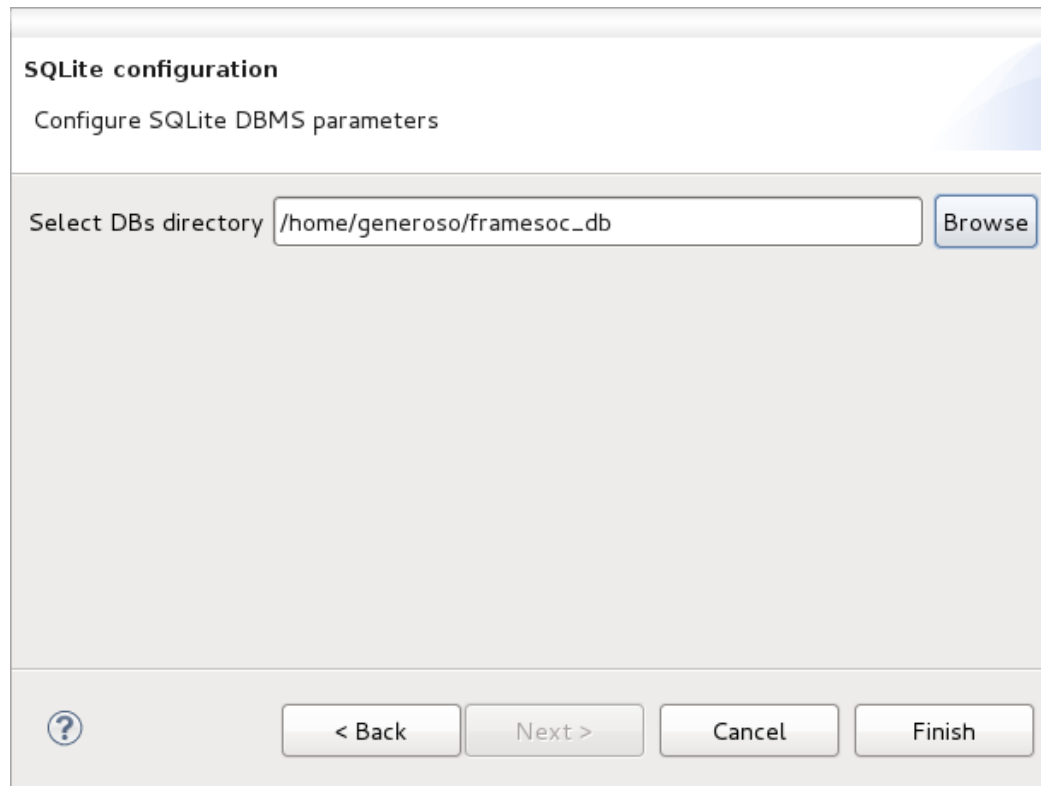
- After restarting, a configuration wizard is shown
 - If not shown: *Framesoc > Management > Initialize System*



- Select **sqlite**, then press **Next**

Initialization (2)

- Specify a directory for the trace databases



- Press **Finish**



Trace Analysis

Framesoc perspective

The screenshot displays the Framesoc perspective in Eclipse Platform, showing a detailed analysis of MPI traces. The interface is divided into several panes:

- Traces:** A tree view on the left showing the loaded traces, including CTF traces and GStreamer traces.
- State duration:** A pie chart and a table showing the distribution of MPI states. The table is as follows:

| Name | Percentage | Value |
|-------------------|------------|--------|
| MPI_Wait | 76.89 % | 5.88E |
| MPI_Send | 21.80 % | 1.67E |
| MPI_Init | 1.05 % | 80.5 s |
| Aggregated slices | 0.26 % | 20.2 s |
| MPI_Barrier | 0.15 % | 11.3 s |
| MPI_Irecv | 0.11 % | 8.7 s |
| MPI_Finalize | 0.00 % | 180 m |
| MPI_Bcast | 0.00 % | 43.5 n |
| MPI_Reduce | 0.00 % | 20 ms |
- Event Types:** A bar chart showing the frequency of various MPI events over time, with a legend on the right for event types like MPI_Barrier, MPI_Bcast, etc.
- Trace Details:** A table on the bottom left providing metadata for the selected trace, such as alias, tracing date, application, board, and number of CPUs.
- Events:** A detailed timeline on the bottom right showing the percentage of displayed links (100.0%) and the sequence of events for each MPI rank (rank 0 to rank 9) across different processes (graphene-96 to graphene-98).

- *Window > Open Perspective > Other... > Framesoc*
- Management of multiple traces
- Different analysis views

Import a trace

- From Eclipse main menu:
 - Framesoc > Trace Analysis > Import Trace

Import a new trace

Tool

Double timestamps Long timestamps

Timestamp Shift

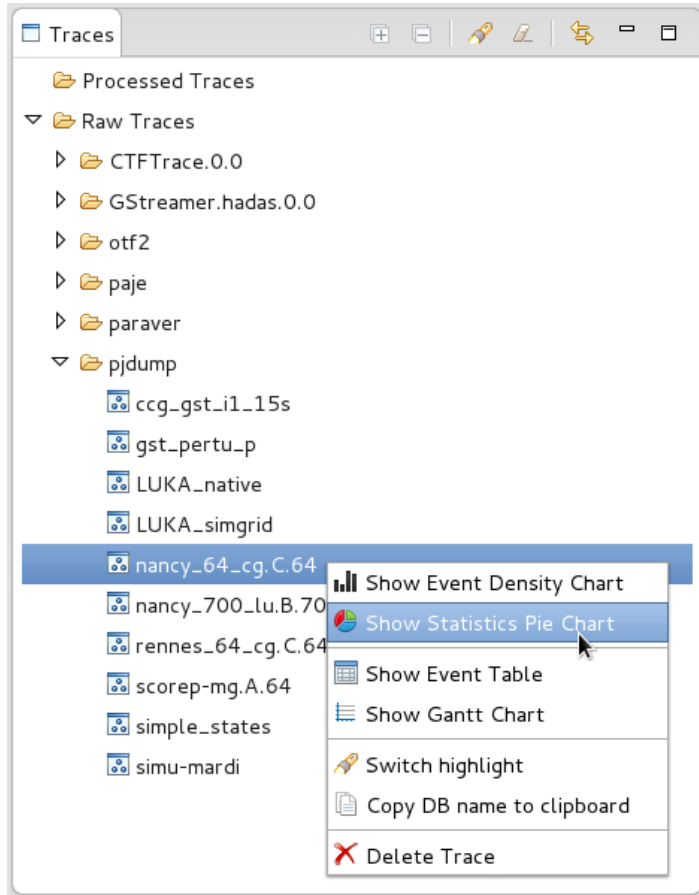
Time Unit

Trace files

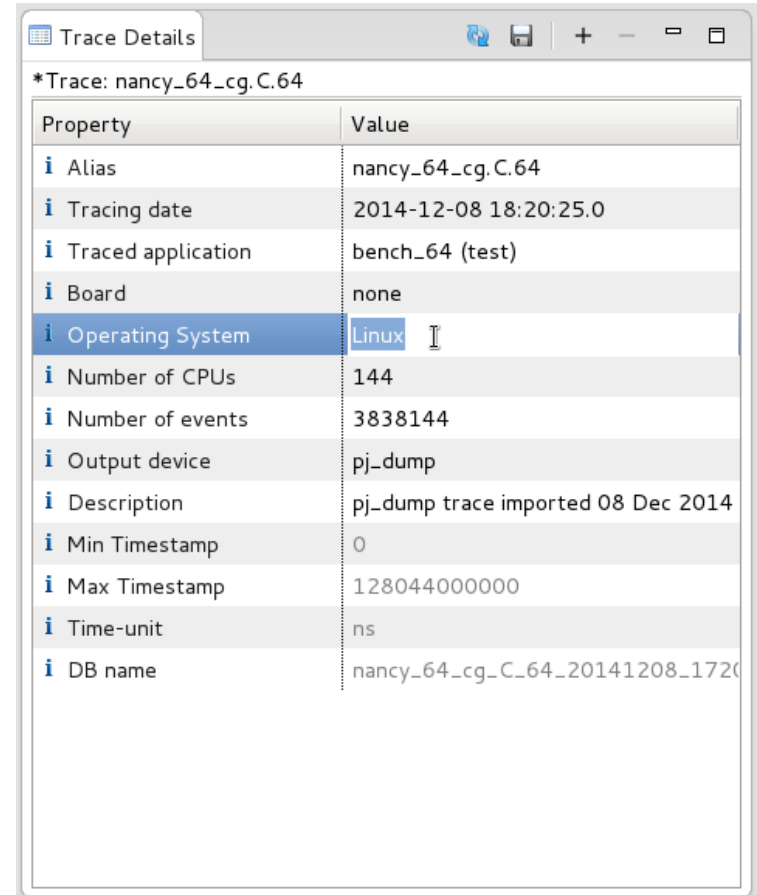
Tool message

Specify at least one trace file.

Browse traces and their metadata

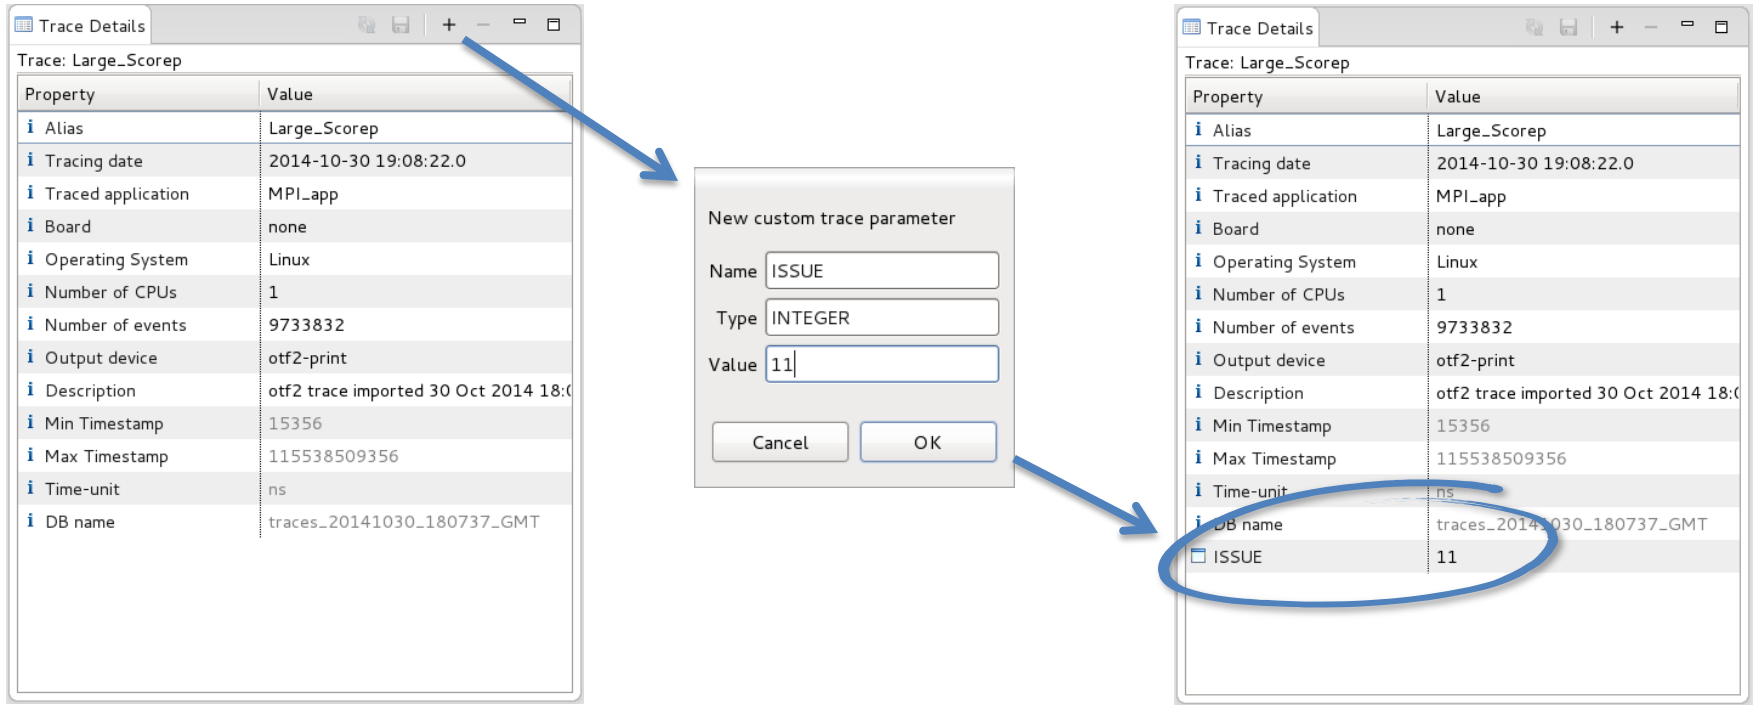


- Trace browser



- Metadata editor

Add custom metadata



- Adding custom metadata to traces
- Tag traces

Trace filtering and highlighting

Check the traces to highlight them in the Traces view.

| Alias | Date | Events | Out device | Description | Ts unit | Parameters | DB |
|---|-----------------------|----------|------------|-------------------|----------|----------------|-----------|
| <input checked="" type="checkbox"/> <filter> | <filter> | <filter> | <filter> | <filter> | <filter> | .*ISSUE='11'.* | <filter> |
| <input checked="" type="checkbox"/> Large_Scorep | 2014-10-30 19:08:22.0 | 9733832 | otf2-print | otf2 trace imp ns | | ISSUE='11' | traces_20 |
| <input checked="" type="checkbox"/> scorep-dt.C.BH.85 | 2014-11-14 11:18:24.0 | 678 | paraver | paje trace imp ns | | ISSUE='11' | scorep_dt |
| <input checked="" type="checkbox"/> ccg_gst_i1_15s | 2014-11-14 11:20:51.0 | 1500315 | paraver | paje trace imp ns | | ISSUE='11' | ccg_gst_i |
| <input checked="" type="checkbox"/> scorep-dt.C.BH.85 | 2014-11-14 11:21:18.0 | 1500315 | paraver | paje trace imp ns | | ISSUE='11' | scorep_dt |
| <input checked="" type="checkbox"/> scorep-mg.A.64 | 2014-11-14 11:21:18.0 | 0 | paraver | paje trace imp ns | | ISSUE='11' | scorep_m |
| <input checked="" type="checkbox"/> simu-mardi | 2014-11-14 11:21:19.0 | 0 | paraver | paje trace imp ns | | ISSUE='11' | simu_mari |
| <input checked="" type="checkbox"/> scorep-dt.C.BH.85 | 2014-11-14 11:24:40.0 | 678 | paraver | paje trace imp ns | | ISSUE='11' | scorep_dt |
| <input checked="" type="checkbox"/> scorep-mg.A.64 | 2014-11-14 11:24:40.0 | 678 | paraver | paje trace imp ns | | ISSUE='11' | scorep_m |
| <input checked="" type="checkbox"/> scorep-mg.A.64 | 2014-11-14 11:26:39.0 | 136608 | paraver | paje trace imp ns | | ISSUE='11' | scorep_m |


Cancel OK

Traces

- video_decoding_1
- video_decoding_2
- Large_Scorep
- Large_scorep_traces_index
- traces
- traces [novar]
- traces [novar]
- paje
 - ccg_gst_i1_15s
 - scorep-dt.C.BH.85
 - scorep-dt.C.BH.85
 - scorep-dt.C.BH.85
 - scorep-mg.A.64
 - scorep-mg.A.64
 - scorep-mg.A.64
 - simu-mardi
- paraver
 - EXTRAE_Paraver_trace_mpich
 - EXTRAE_Paraver_trace_mpich
 - EXTRAE_Paraver_trace_mpich
- pjdump
 - ccg_gst_i1_15s
 - gst_pertu_p
 - links

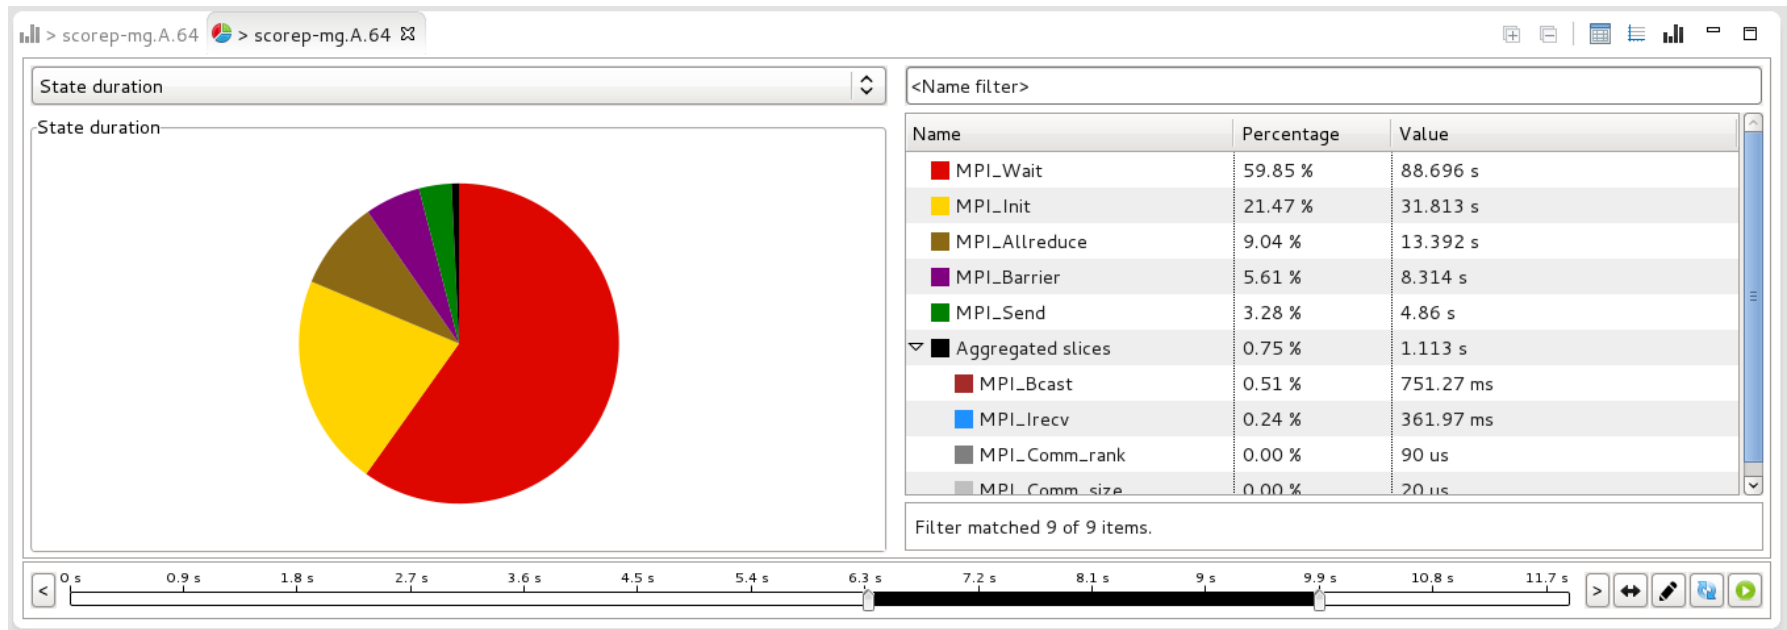
- Filter and highlight interesting traces

Gantt Chart




- State aggregation / link filtering
- Event producer / event type filtering

Statistics Pie Chart



- Several statistics operators
- Statistics computed on generic time interval
- Exclude items from statistics computation
- Group different items

Event Density Chart



- Event type filtering
- Event producer filtering

Event Table

scorep-mg.A.64 > scorep-mg.A.64

| Timestamp | CPU | Event Producer | Category | Event Type | Parameters |
|------------|----------|----------------|----------|---------------|---|
| <filter> | <filter> | <filter> | <filter> | <filter> | <filter> |
| 6815920000 | 0 | rank10 | State | MPI_Comm_rank | END_TIMESTAMP='6815920000', IMBRICATION='0' |
| 6815930000 | 0 | rank10 | State | MPI_Comm_size | END_TIMESTAMP='6815930000', IMBRICATION='0' |
| 6815930000 | 0 | rank1 | State | MPI_Comm_rank | END_TIMESTAMP='6815930000', IMBRICATION='0' |
| 6815940000 | 0 | rank1 | State | MPI_Comm_size | END_TIMESTAMP='6815940000', IMBRICATION='0' |
| 6815950000 | 0 | rank57 | State | MPI_Comm_rank | END_TIMESTAMP='6815950000', IMBRICATION='0' |
| 6815950000 | 0 | rank57 | State | MPI_Comm_size | END_TIMESTAMP='6815950000', IMBRICATION='0' |
| 6815960000 | 0 | rank41 | State | MPI_Comm_rank | END_TIMESTAMP='6815960000', IMBRICATION='0' |
| 6815970000 | 0 | rank41 | State | MPI_Comm_size | END_TIMESTAMP='6815970000', IMBRICATION='0' |
| 6815970000 | 0 | rank10 | State | MPI_Barrier | END_TIMESTAMP='6817760000', IMBRICATION='0' |
| 6815970000 | 0 | rank1 | State | MPI_Barrier | END_TIMESTAMP='6817820000', IMBRICATION='0' |
| 6815980000 | 0 | rank57 | State | MPI_Barrier | END_TIMESTAMP='6817960000', IMBRICATION='0' |
| 6815990000 | 0 | rank41 | State | MPI_Barrier | END_TIMESTAMP='6817790000', IMBRICATION='0' |
| 6816000000 | 0 | rank36 | State | MPI_Comm_rank | END_TIMESTAMP='6816000000', IMBRICATION='0' |
| 6816000000 | 0 | rank36 | State | MPI_Comm_size | END_TIMESTAMP='6816000000', IMBRICATION='0' |

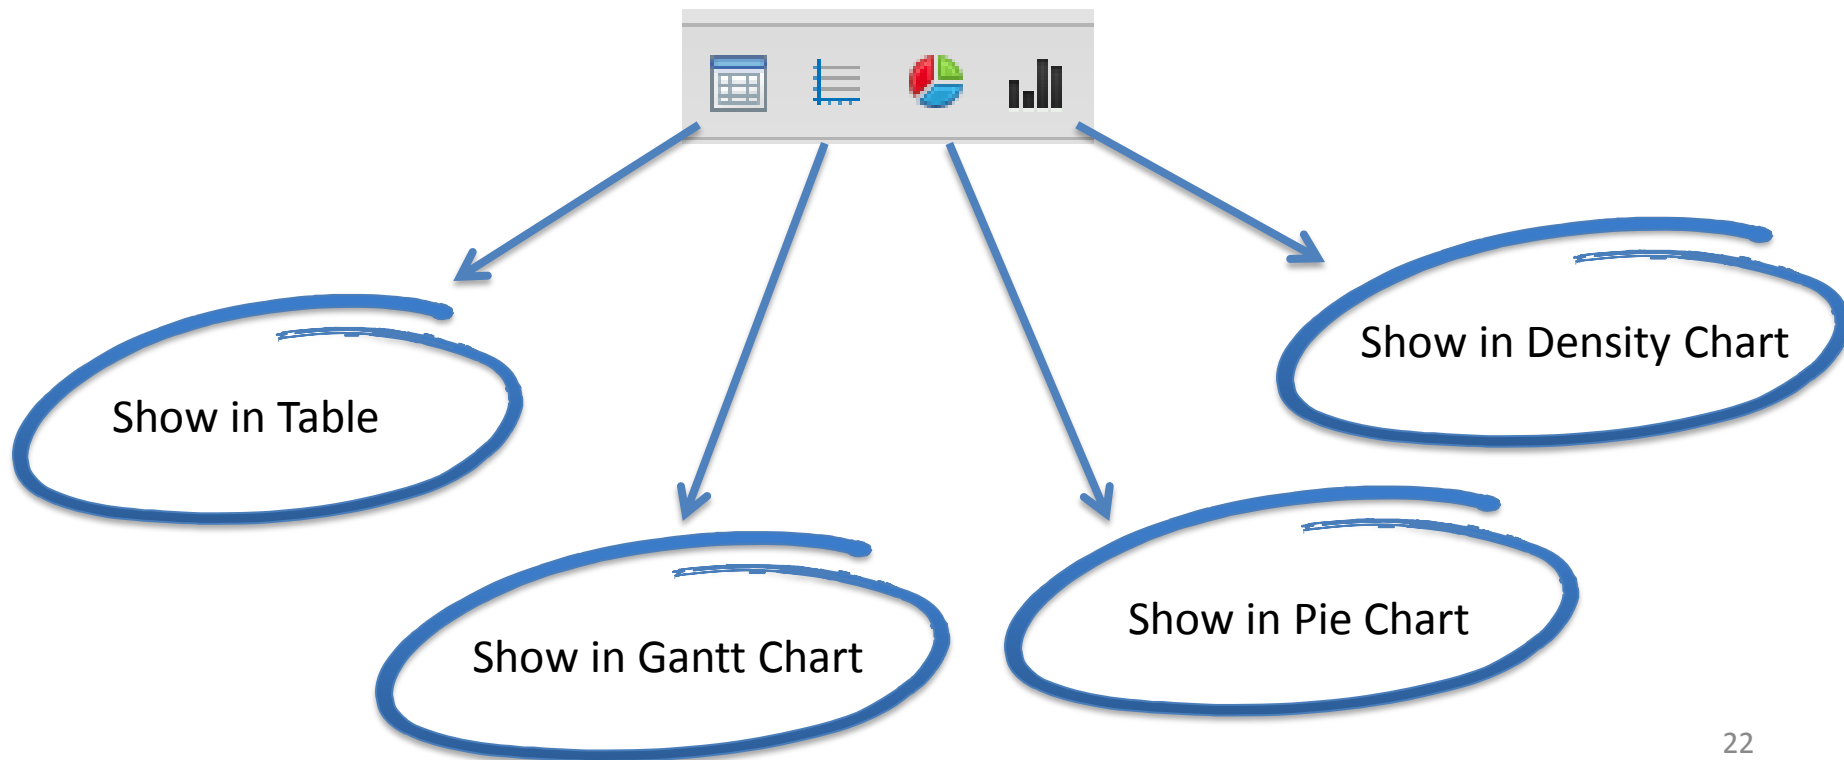
Filter matched 82720 of 82720 loaded events


Timeline: 0 s, 0.9 s, 1.8 s, 2.7 s, 3.6 s, 4.5 s, 5.4 s, 6.3 s, 7.2 s, 8.1 s, 9 s, 9.9 s, 10.8 s, 11.7 s

- Filtering column values using regular expression

View switching and synchronization

- Each analysis view toolbar contains **3** of these **4** buttons
 - Each button allows to **switch** to the corresponding view
 - It is possible to **synchronize** views on the time interval





Distribution

Source repository

GitHub

- Code hosted on GitHub **soctrace-inria** organization
 - <https://github.com/soctrace-inria/>
- Contributions via pull requests

Useful links

- **Framesoc wiki**

<https://github.com/soctrace-inria/framesoc/wiki>

- **Framesoc website**


<http://soctrace-inria.github.io/framesoc/>

- **Ocelotl website**

<http://soctrace-inria.github.io/ocelotl/>

- **Test traces**

<http://moais.imag.fr/membres/damien.dosimont/files/traces/>



Perspectives

Some perspectives

- Storage using NoSQL distributed solutions
 - Cassandra
- Multidimensional filtering in all view
 - Time / Event Producers / Event Types
- ...

Just have a look at our **github:issues**

<https://github.com/soctrace-inria/framesoc/issues>

Questions?

