



Trace management and analysis with FrameSoC

Generoso Pagano

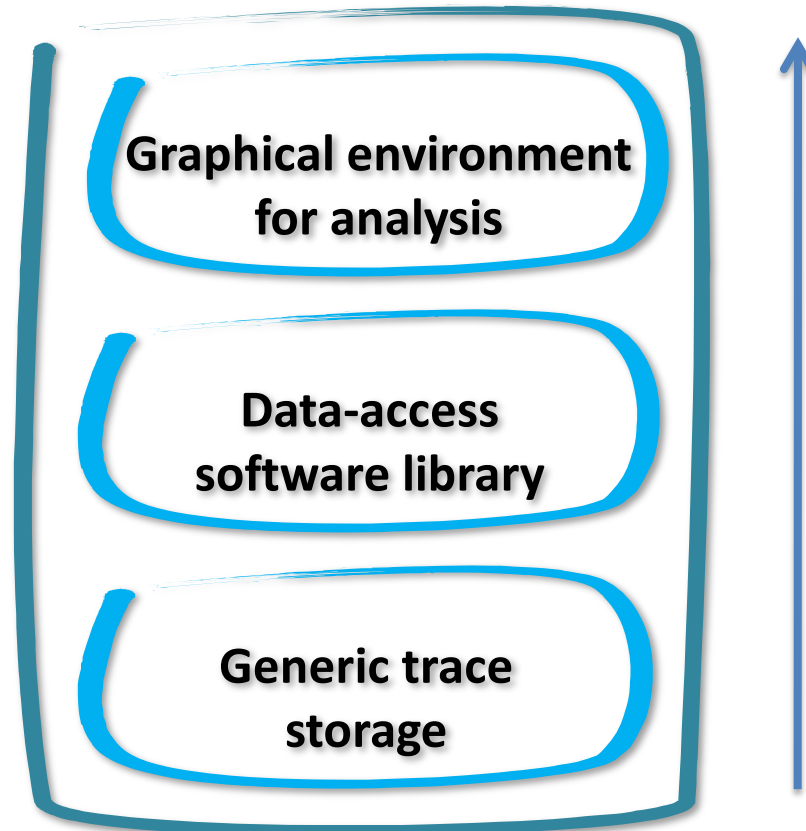
generoso.pagano@inria.fr

Journées Mescal

Autrans, 18/03/2014

FrameSoC: trace analysis infrastructure

- To tackle the above issues we propose **FrameSoC**




- FrameSoC is developed within the **SoC-Trace** project
 - **INRIA**, **UJF**, **STMicroelectronics**, ProbaYes, Magillem

Graphical environment
for analysis


Data-access
software library

**Generic trace
storage**


Many formats, many tools



Many formats, many tools




FrameSoC



- Represent **different trace formats**
 - Self defining approach
 - Event categories
- Store **various trace metadata**
- Store trace **analysis results**

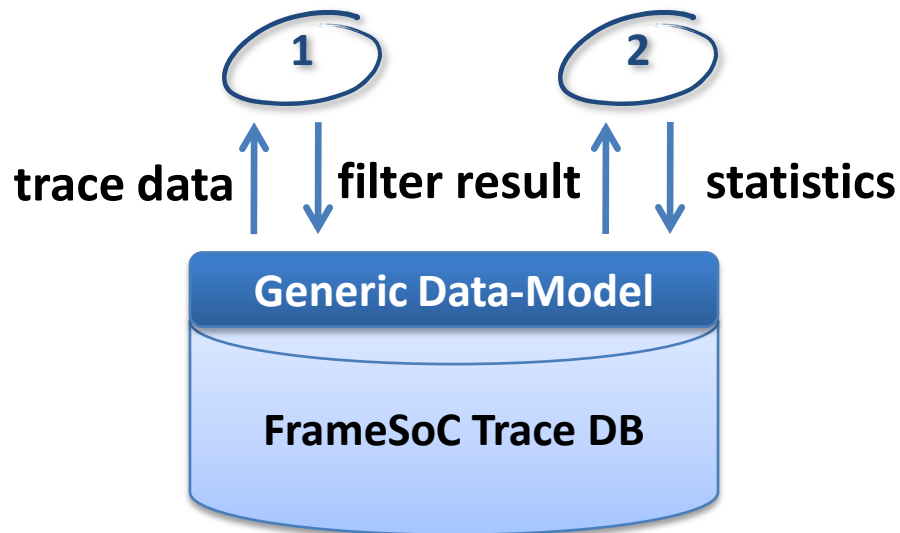
Analysis workflow



Analysis workflow



FrameSoC



Store analysis results

- Tool cooperation
- Avoid long recomputations

Graphical environment
for analysis

**Data-access
software library**

Generic trace
storage

SoC-Trace Library



- It is a set of **Eclipse** plugins
- Read/Write access to data abstracts storage details
 - We read and write Java objects
- **Factorization** of trace-access functionalities
 - Basic bricks to build analysis tools (e.g. *OcelotI*)

**Graphical environment
for analysis**

Data-access
software library

Generic trace
storage

FrameSoC Workbench

FrameSoC - Eclipse Platform

File Edit Navigate Search Project FrameSoC Run Window Help

Quick Access Resource FrameSoC

Traces

- Processed Traces
- Raw Traces
 - com.st.framesoc.kptrace.0.2
 - spear
 - ts_record
 - GStreamer.hadas.0.0
 - gst_good_decoding
 - gst_video_decoding
 - pjdump.inria.0.0
 - luka_native
 - luka_simgrid
 - paje_otf
 - scorep-mg.A.64**
 - simu-mardi
 - testsmall

<Event Density Chart>

Trace: scorep-mg.A.64

<Statistics Pie Chart>

Trace: scorep-mg.A.64

Event Types

Name	Percentage
MPI_Irecv	31.69 %
MPI_Wait	31.69 %
MPI_Send	31.69 %
MPI_Allreduce	4.12 %
Aggregated slices	0.80 %
MPI_Barrier	0.28 %
MPI_Bcast	0.28 %
MPI_Init	0.05 %
MPI_Comm_size	0.05 %
MPI_Finalize	0.05 %

Filter matched 11 of 11 items

Trace Details

Trace: scorep-mg.A.64

Property	Value
Tracing date	2014-02-06 15:04:41.0
Traced application	mpi_bench
Board	none
Operating System	linux
Number of CPUs	1
Number of events	136608
Output device	pj_dump
Description	pj_dump trace imported 06 Feb
Alias	scorep-mg.A.64
DB name	PJDUMP_20140206_140439
Time-unit	ns

<Events>

Trace: scorep-mg.A.64

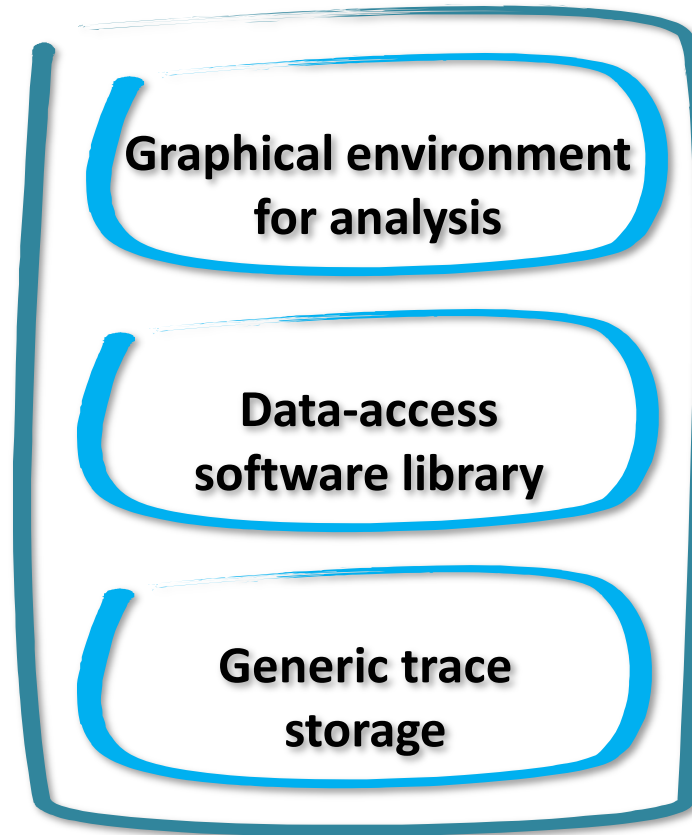
Timestamp	CPU	Event Producer	Category	Event Type	Parameters
6815920000	0	rank10	State	MPI_Comm_rank	END_TIMESTAMP='6815920000'
6815930000	0	rank10	State	MPI_Comm_size	END_TIMESTAMP='6815930000'
6815930000	0	rank1	State	MPI_Comm_rank	END_TIMESTAMP='6815930000'
6815940000	0	rank1	State	MPI_Comm_size	END_TIMESTAMP='6815940000'
6815950000	0	rank57	State	MPI_Comm_rank	END_TIMESTAMP='6815950000'
6815950000	0	rank57	State	MPI_Comm_size	END_TIMESTAMP='6815950000'
6815960000	0	rank41	State	MPI_Comm_rank	END_TIMESTAMP='6815960000'
6815970000	0	rank41	State	MPI_Comm_size	END_TIMESTAMP='6815970000'
6815970000	0	rank10	State	MPI_Barrier	END_TIMESTAMP='6815970000'
6815970000	0	rank1	State	MPI_Barrier	END_TIMESTAMP='6815970000'
6815980000	0	rank57	State	MPI_Barrier	END_TIMESTAMP='6815980000'
6815990000	0	rank41	State	MPI_Barrier	END_TIMESTAMP='6815990000'

Filter matched 60594 of 60594 loaded events

<Gantt Chart>

Trace: scorep-mg.A.64


Time Unit: ns | Tick under cursor: 7.2E9 | Duration: -




Let's see FrameSoC in action!

An example: comparison of two traces

- Use case provided by *Luka*




native




simgrid

≠



native



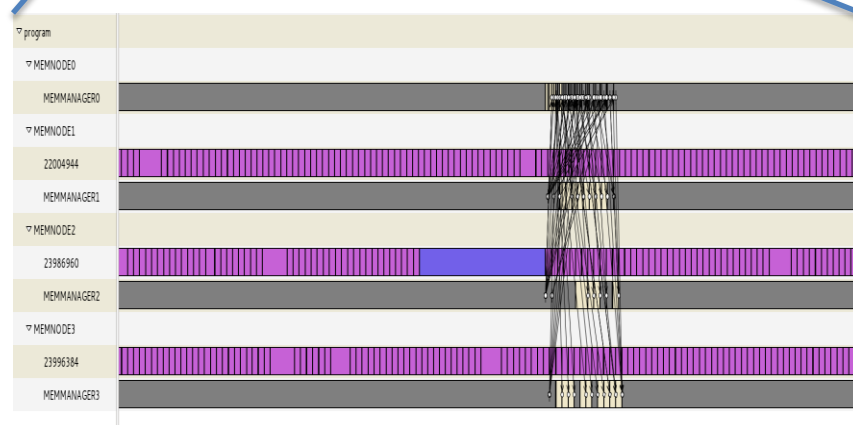

simgrid



native



simgrid



native

Timestamp	CPU	Event Producer	Category	Event Type
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	AllocatingReuse
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	AllocatingReuse
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing


Filter matched 213783 of 218632 loaded events

native

Timestamp	CPU	Event Producer	Category	Event Type
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	AllocatingReuse
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	AllocatingReuse
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Allocating
19865200000000	0	MEMMANAGER2	State	Nothing
19865200000000	0	MEMMANAGER2	State	Reclaiming
19865200000000	0	MEMMANAGER2	State	Nothing

Filter matched 213783 of 218632 loaded events

native




native

Timestamp	CPU	Event Producer	Category	Event Type
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Allocating
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Reclaiming
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	AllocatingReuse
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Allocating
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Reclaiming
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Allocating
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Reclaiming
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Allocating
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Reclaiming
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	AllocatingReuse
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Allocating
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Reclaiming
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Allocating
198652000000000	0	MEMMANAGER2	State	Nothing
198652000000000	0	MEMMANAGER2	State	Reclaiming
198652000000000	0	MEMMANAGER2	State	Nothing

Filter matched 213783 of 218632 loaded events



native



simgrid



- **Actual problem:** *simgrid* assumed infinite GPU memory
 - no swapping to the RAM occurred during simulation

Conclusion...

✓ **FrameSoC can facilitate trace analysis**

- Base analysis tools
- Fast analysis
- Hints to find the problem

...and Perspectives

- ⊕ Type/Producer filtering in Pie and Histogram
- ⊕ Time filtering in Pie
- ⊕ Provide more statistical views
- ⊕ Improve ergonomics
- ⊕ ...



First open-source release of FrameSoC in June 2014

Questions?

