

S9347: Performance Analysis for Large Scale GPU Applications and DL Frameworks



Dr. Guido Juckeland

/ Robert Henschel

Head Computational Science Dept. / Director of Science CommunityTools at Indiana University

Agenda

What to expect from the next 80 minutes

- Motivation
- Generating profiles and trace files with Score-P
- Visualizing trace files with Vampir
- Looking into Deep Learning Frameworks

Disclaimer

It's extremely easy to waste performance

- Poor/no GPU usage (80-90%)
- Bad MPI (50-90%)
- **Total: 1% of peak (or worse)**
- **Performance tools will not “automagically” make your code faster – they just point to “areas of interest”**

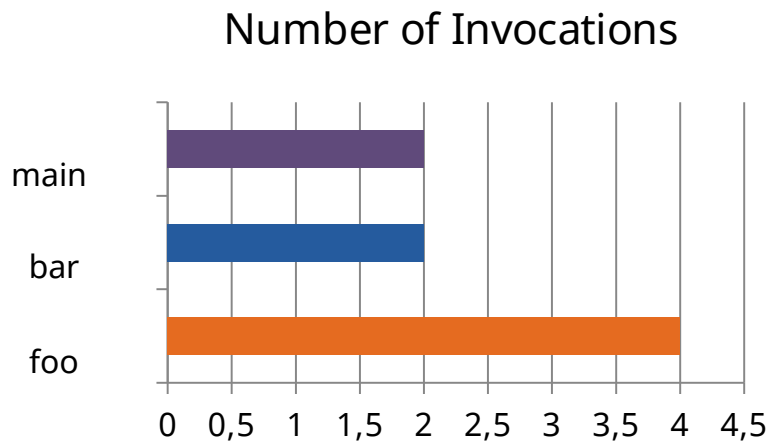
Motivation

Performance Tuning 101


Profiling vs. Tracing

Preserving the details


Statistics



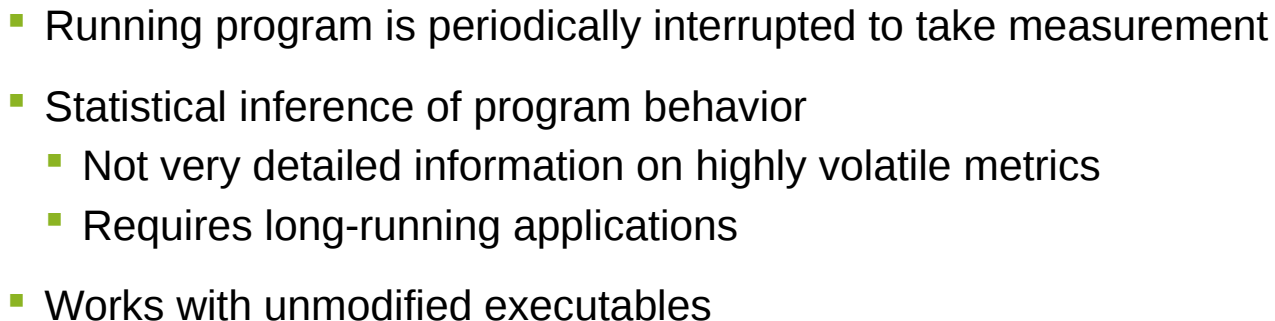
Execution Time



Timelines




Periodic observations of your application (Pull)



Instrumentation

Modify application to deliver information (Push)




- Measurement code is inserted such that every event of interest is captured directly
- Advantage:
 - Much more detailed information
- Disadvantage:
 - Processing of source-code / executable necessary
 - Large relative overheads for small functions

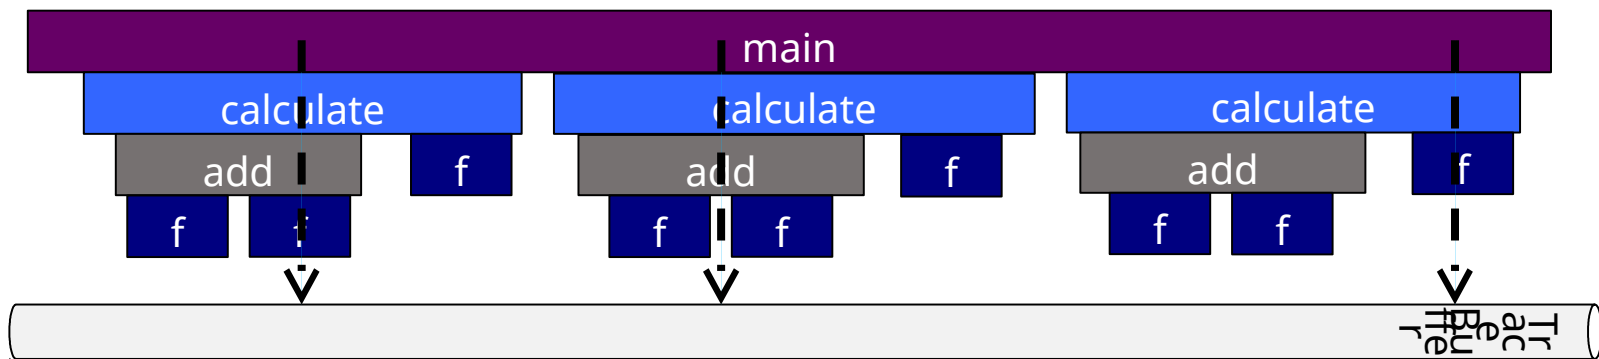
Sampling vs. Tracing

Comparing both approaches visually

Function
Instrumen-
tation:

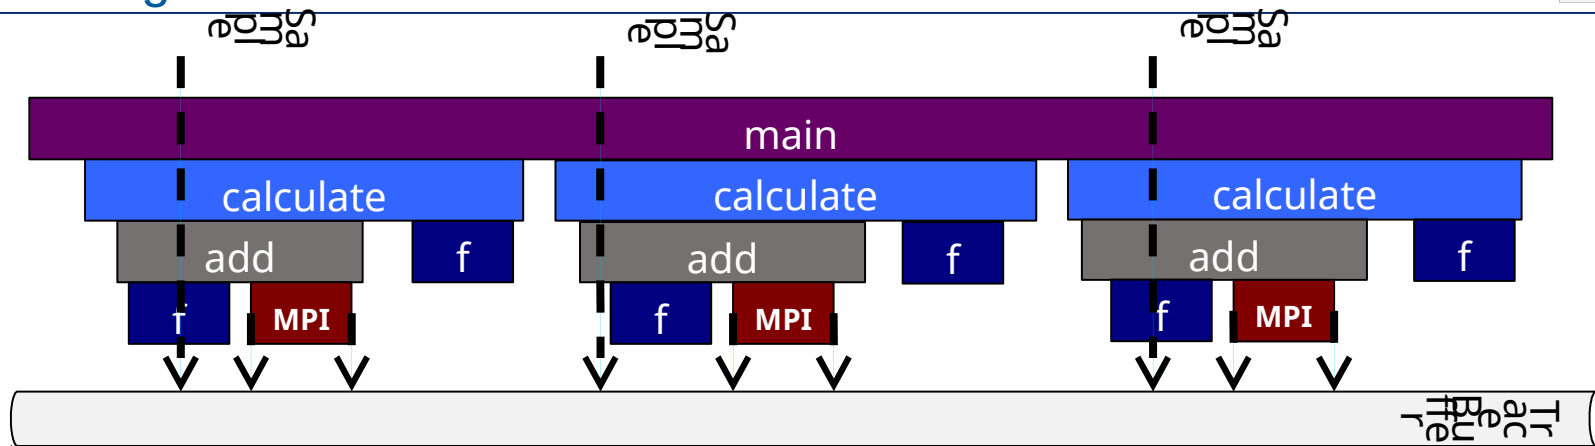


Sampling:



Sampling + Instrumentation

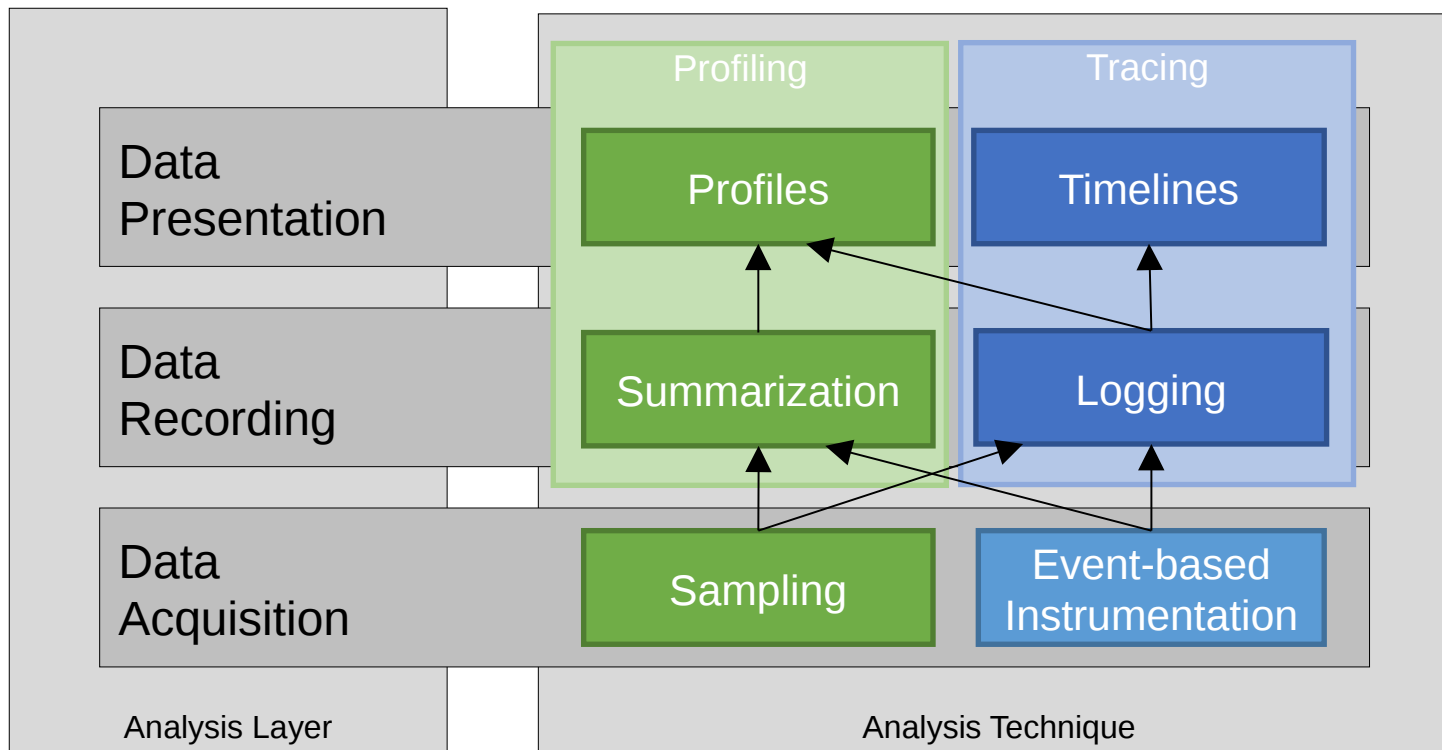
Combining the best of both worlds



- Long running applications:
 - Requires large buffers or heavy filtering
 - Creating a filter requires runs in advance
- Codes with many small functions (e.g.: C++):
 - Function instrumentation a challenge
- **Score-P: Sampling+Tracing**

Terms and How They Relate

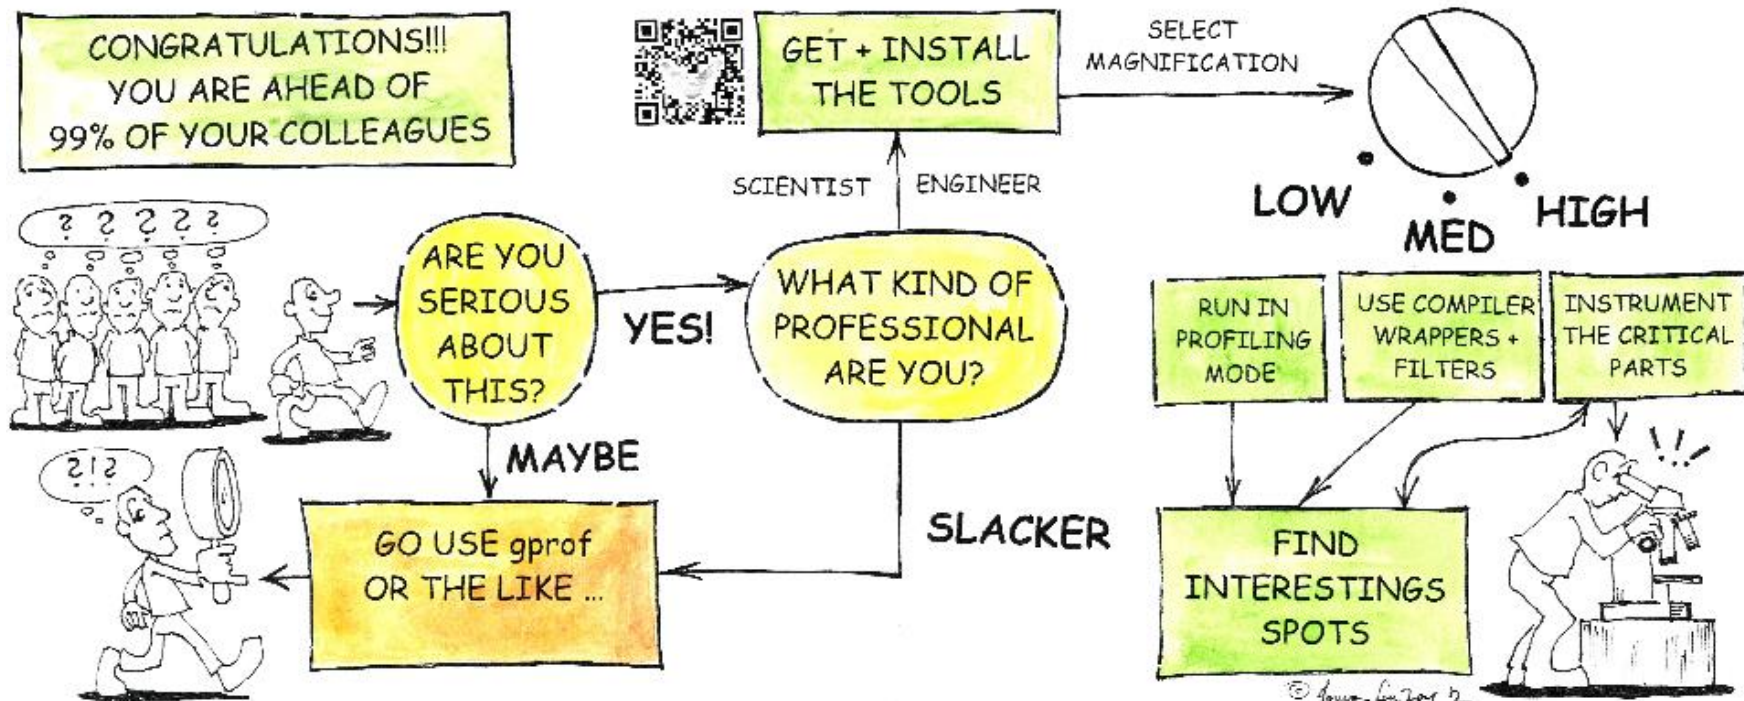
Making sure we use the same words



Summary

Making the “right” choices


SO, YOU HAVE DECIDED TO UNDERSTAND WHAT A PROGRAM EXACTLY DOES?



Generating Traces and Profiles with Score-P

Overall workflow

Recording and studying performance data



- Attach Score-P to application
- Run with attached monitor ==> trace/profile data
- Study trace with Vampir / profile with Cube
- Repeat to:
 - Adapt instrumentation (“what you measure”)
 - Evaluate result of a change

Attaching Score-P

a.k.a. instrumenting your source code

CC = pgcc
CXX = pgCC
F90 = pgf90
MPICC = mpicc
NVCC = nvcc



CC = **scorep** <options> pgcc
CXX = **scorep** <options> pgCC
F90 = **scorep** <options> pgf90
MPICC = **scorep** <options> mpicc
NVCC = **scorep** <options> nvcc

```
$ scorep --help
```

This is the Score-P instrumentation tool. The usage is:
scorep <options> <original command>

Common options are:

```
...
--instrument-filter=<file>
    Specifies the filter file for filtering functions during
    compile-time. It applies the same syntax, as the one
    used by Score-P during run-time.

--user
    Enables user instrumentation.
```

Attaching Score-P

Instrument once – change measurement via runtime variables

```
$ scorep-info config-vars --full
```

```
SCOREP_ENABLE_PROFILING
```

```
[...]
```

```
SCOREP_ENABLE_TRACING
```

```
[...]
```

```
SCOREP_TOTAL_MEMORY
```

```
Description: Total memory in bytes for the measurement system
```

```
[...]
```

```
SCOREP_EXPERIMENT_DIRECTORY
```

```
Description: Name of the experiment directory
```

```
[...]
```

```
$ export SCOREP_ENABLE_PROFILING=true
```

```
$ export SCOREP_ENABLE_TRACING=false
```

```
$ export SCOREP_EXPERIMENT_DIRECTORY=profile
```

```
$ mpirun <instrumented binary>
```

Profiling Example

Combined Sampling+Tracing


Available since Score-P 2.0

```
$ export SCOREP_ENABLE_TRACING=true  
$ export SCOREP_ENABLE_UNWINDING=true  
$ export SCOREP_SAMPLING_EVENTS=perf_cycles@2000000
```

- User code is sampled (pull)
- Runtime libraries with tracing support use events (push):
 - MPI
 - OpenMP / OpenACC / pthreads
 - CUDA / OpenCL
 - I/O

Things to look at

What can Score-P record?



User Functions

- C/C++/Fortran
- Sampling ***NEW***
- Custom regions
- Java
- Python
(*Experimental*)

Parallel Paradigms

- MPI
- Pthreads
- OpenMP
- XeonPhi Native ***NEW***
- CUDA
- OpenACC/OpenCL ***NEW***
- OpenShmem (+Cray)
- I/O (*Experimental*)

Hardware

- Performance counters (PAPI)
- Plugin counters

Operating System

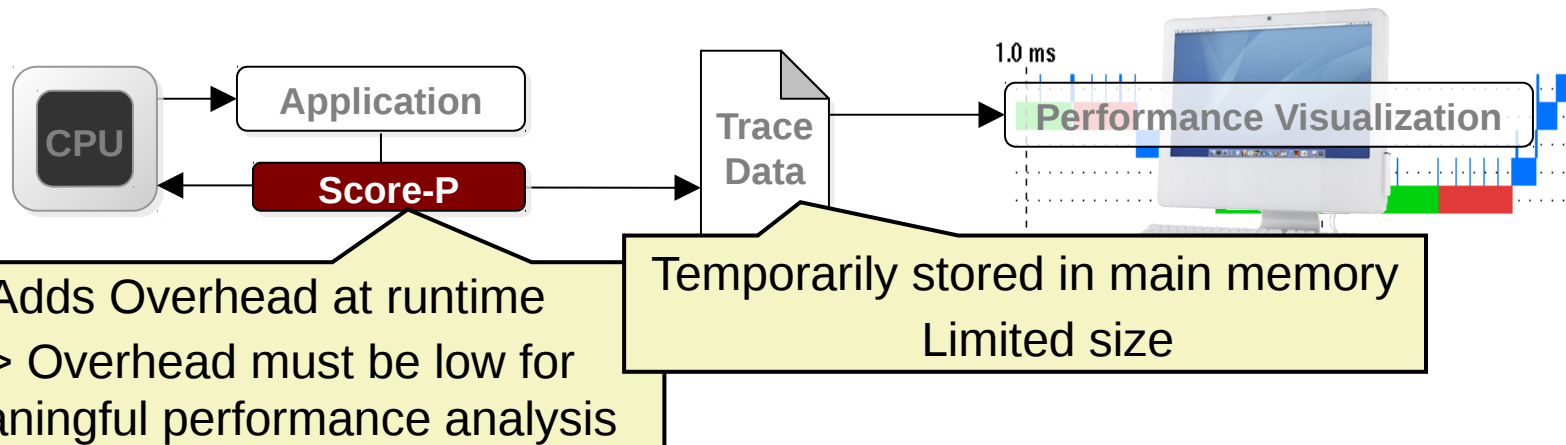
- Resource usage

```
$ export SCOREP_ENABLE_TRACING=yes  
$ export SCOREP_TIMER=clock_gettime  
$ export SCOREP_CUDA_ENABLE=driver, kernel, memcpy, flushatexit  
$ export SCOREP_OPENACC_ENABLE=yes  
$ export ACC_PROFLIB=$SCOREP_LIB/libscorep_adapter_openacc_event.so
```

- Can be used in combination
- Also supports CUPTI counters

Limitations

Why tracing is hard



- Event tracing requires trade-offs:
 - Only add the data sources you need
 - Limit granularity (i.e., filtering)
- Score-P is a profiling experiment

DEMO:

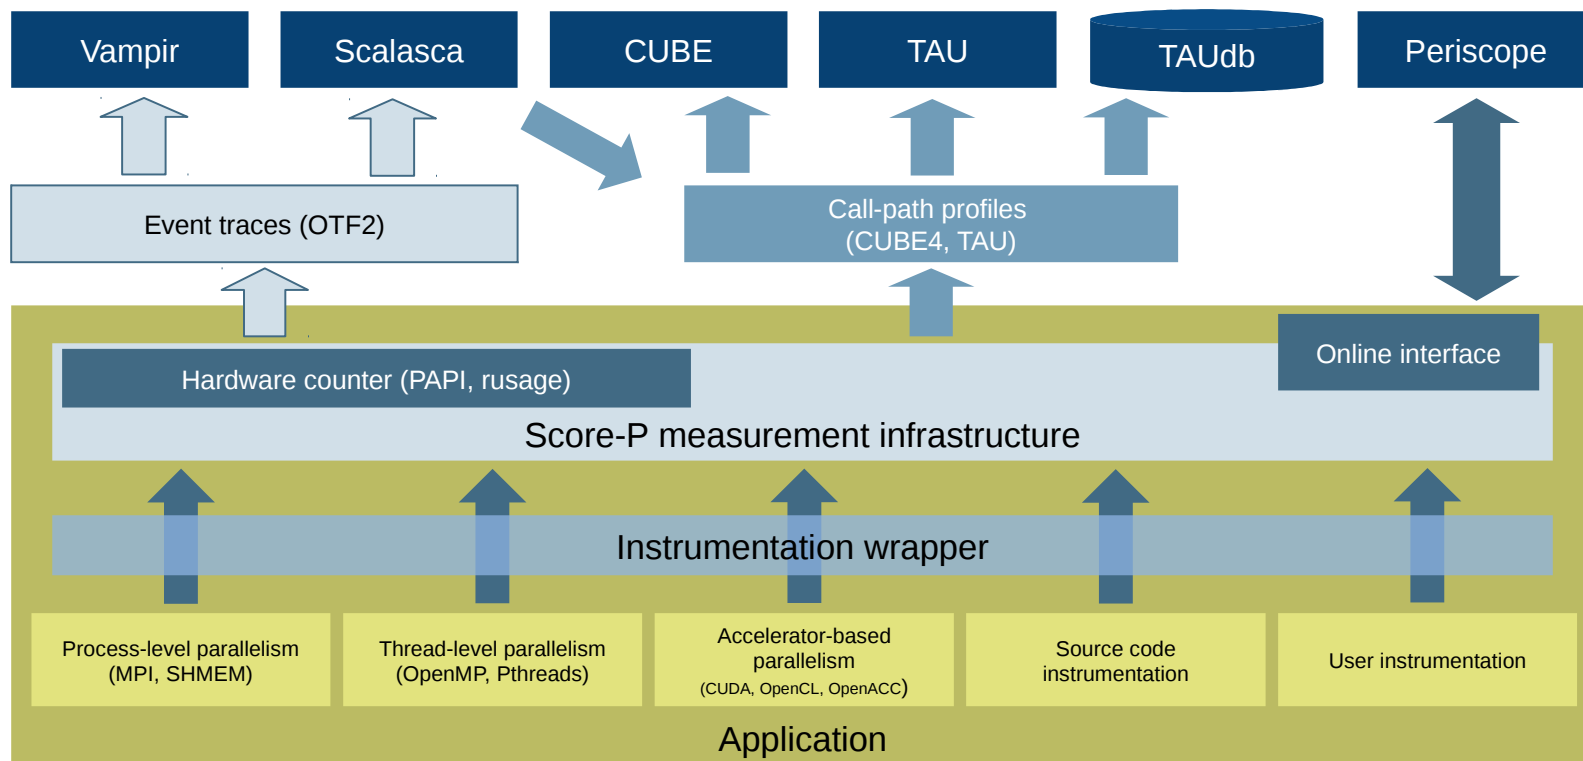
Generating Traces and Profiles with Score-P

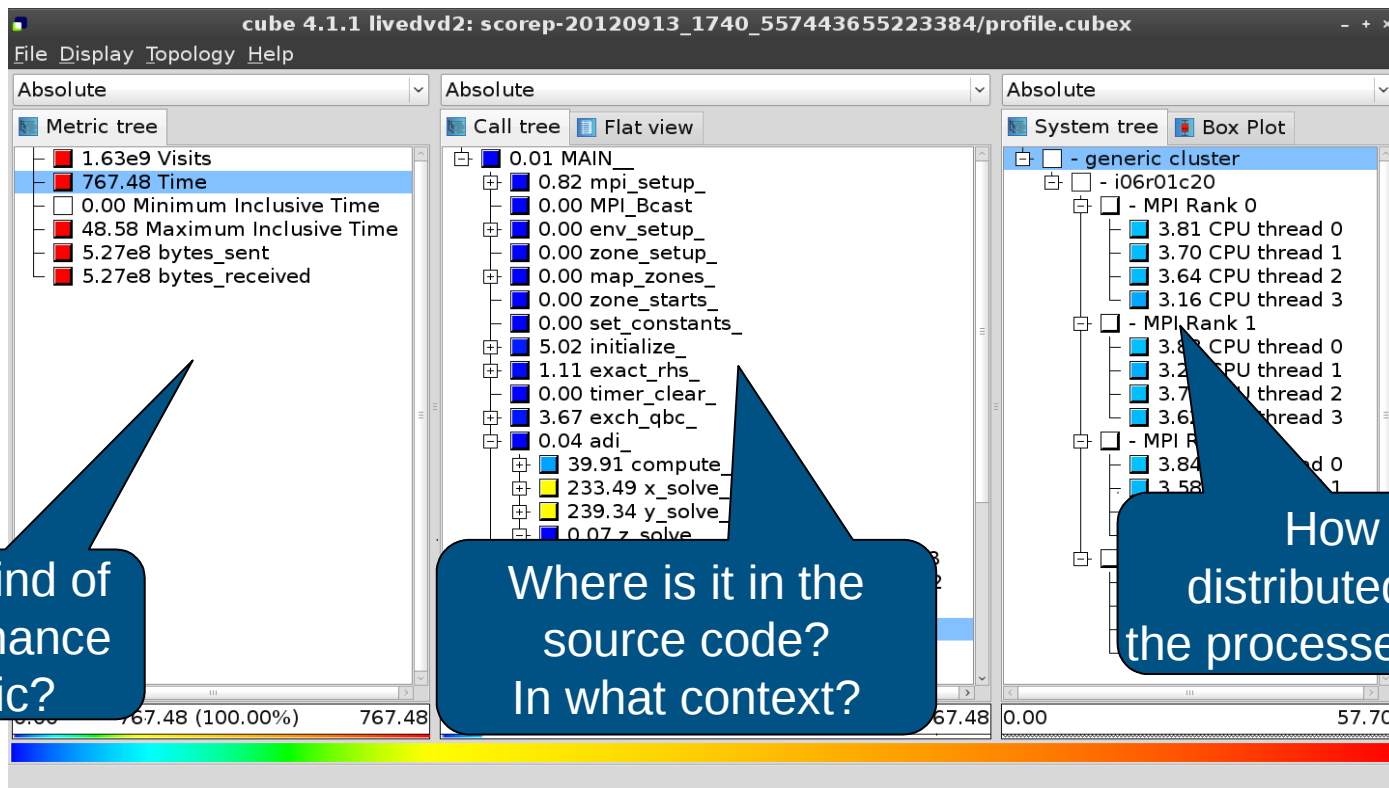
Visualizing

Profiles with CUBE
Traces with Vampir

Bringing it all together

Score-P + Analysis Tools






What kind of performance metric?


Where is it in the source code?
In what context?

How is it distributed across the processes/threads?



Large imbalance
instantly visible

>50% time wasted

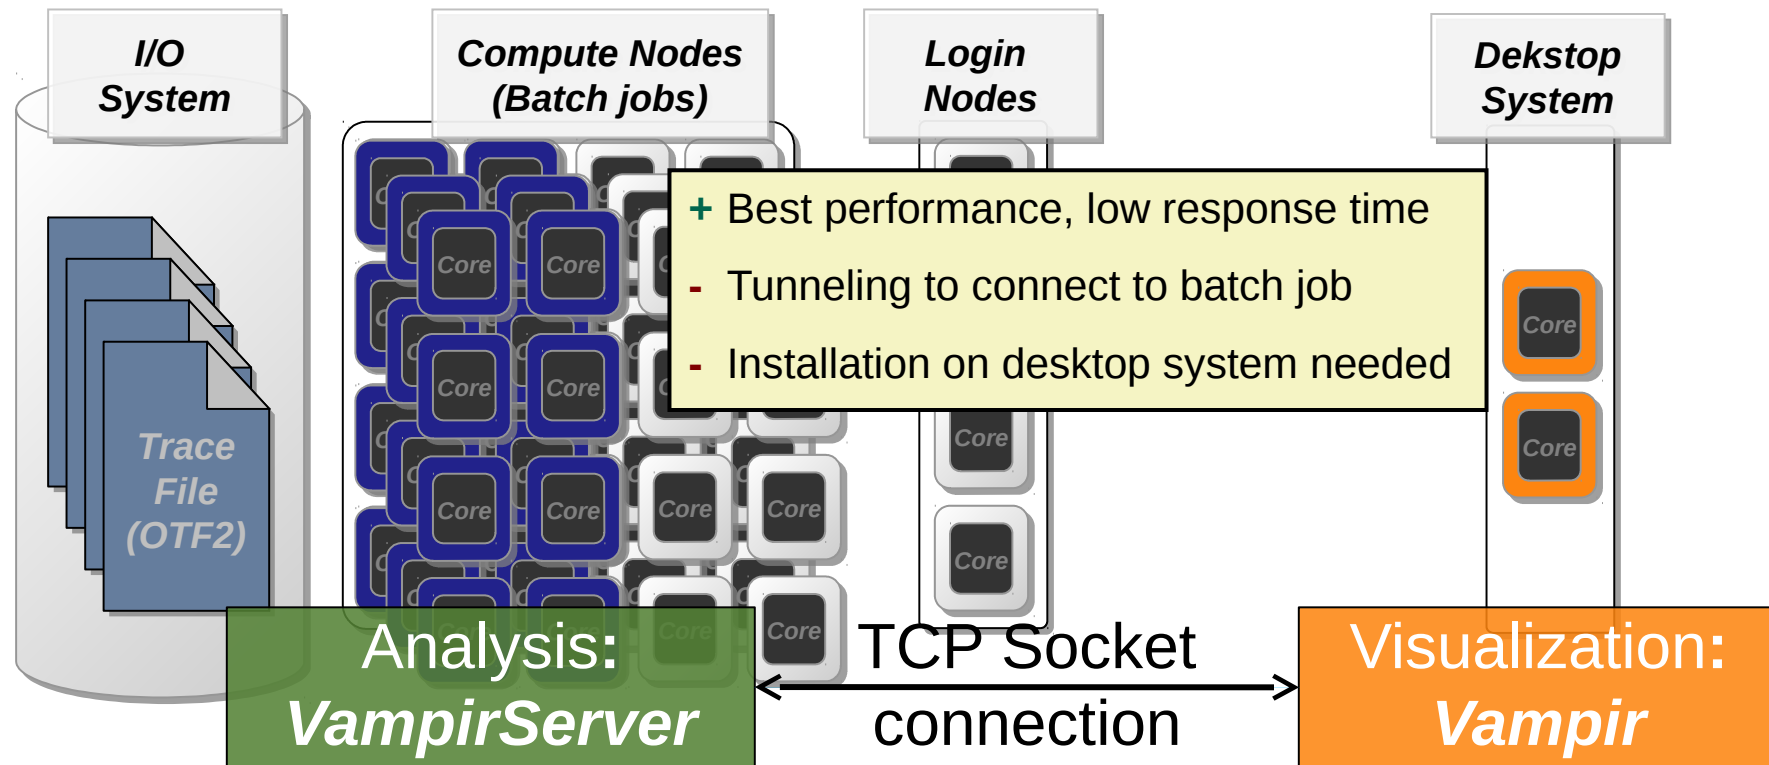


Use your destop system



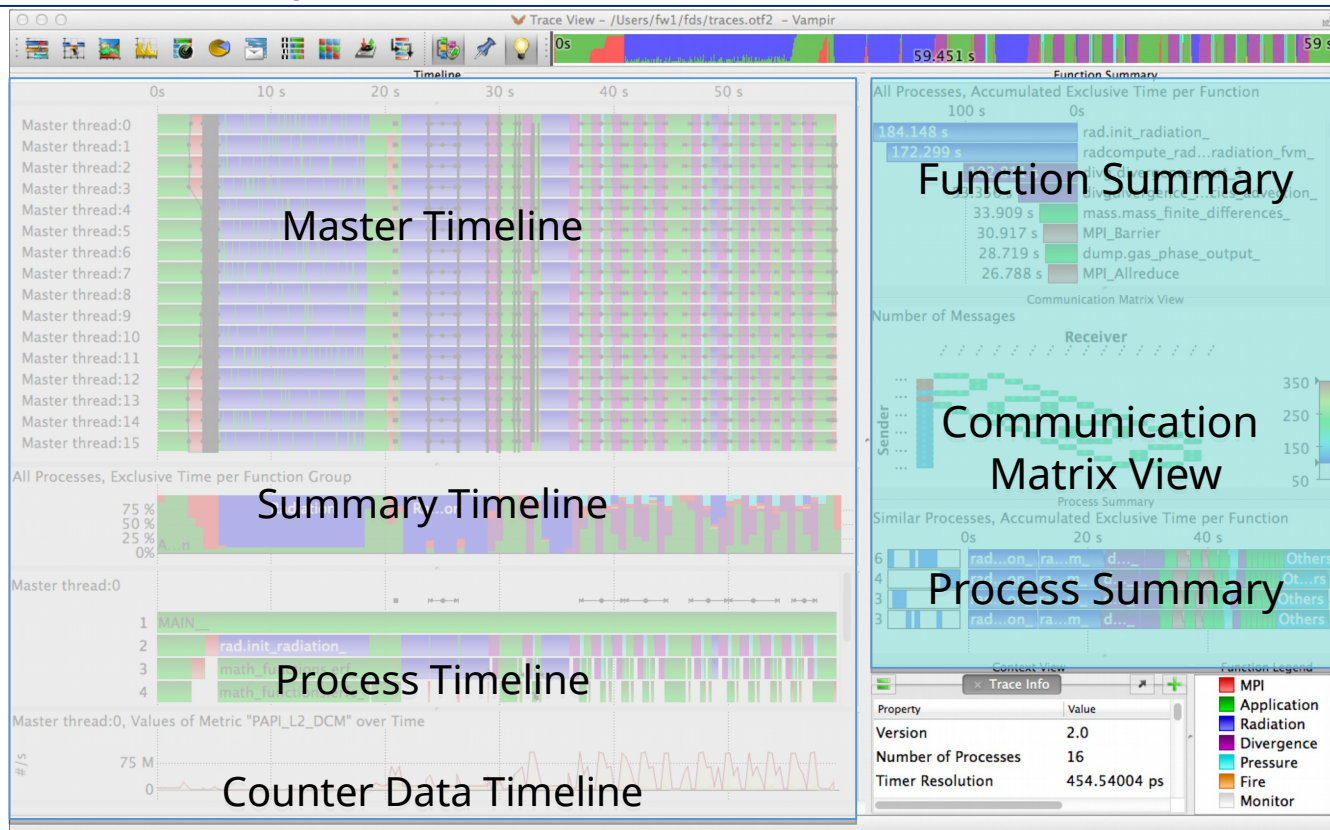
(Re)Using the HPC Resources

Run analysis engine on compute nodes, GUI on desktop



Vampir GUI

What do the fancy colors mean?





Master Timeline



all threads' activities over time per thread



Summary Timeline



all threads' activities over time per activity



Performance Radar



all threads' perf-metric over time



Process Timeline



single thread's activities over time



Counter Data Timeline



single threads perf-metric over time



Function Summary

➔ *runtime/invocation summaries*



Message Summary

➔ *data transfer statistics*



I/O Summary

➔ *I/O statistics*



Process Summary

➔ *Clustering of similar event streams*

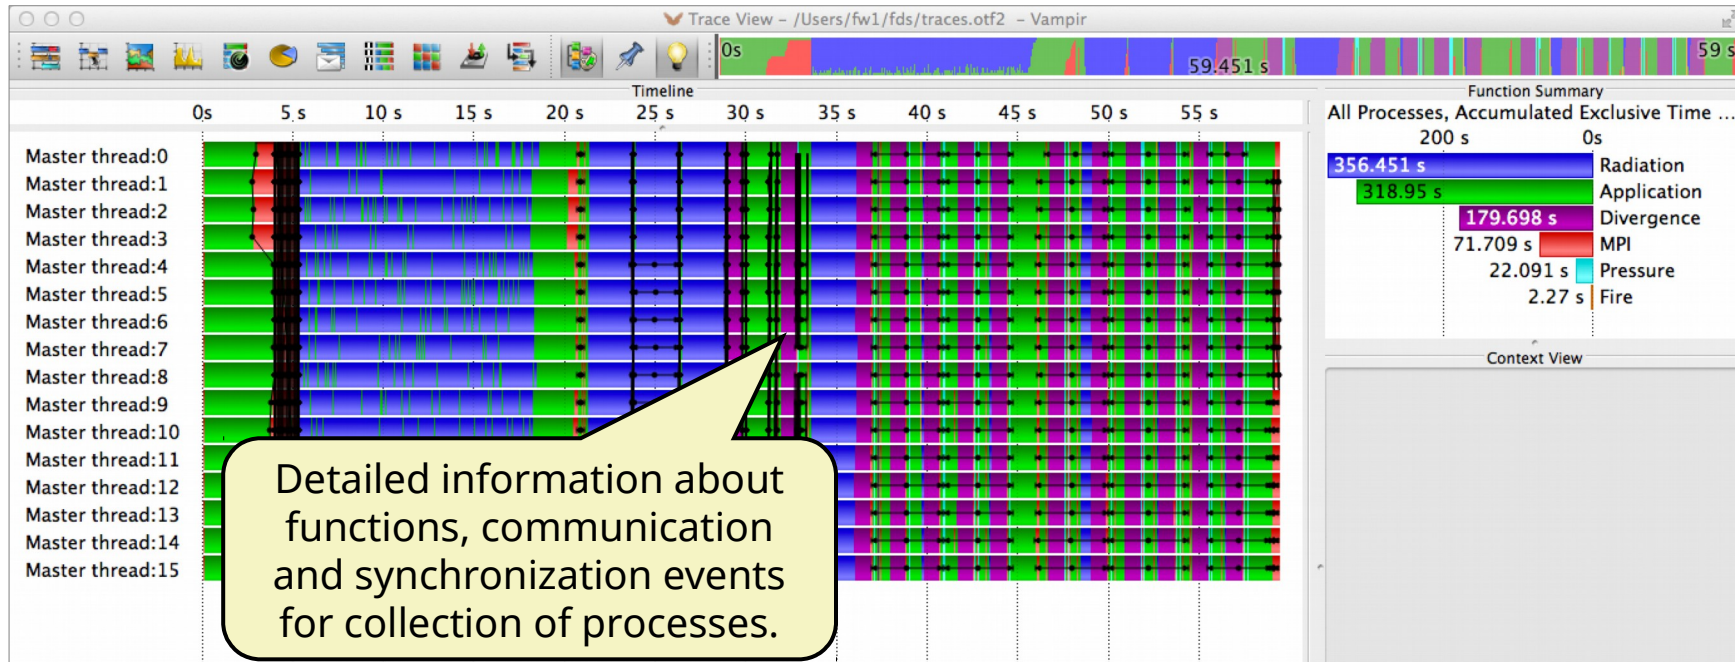


Communication Matrix View ➔ *Pairwise communication statistics*

Vampir Performance Charts in Detail



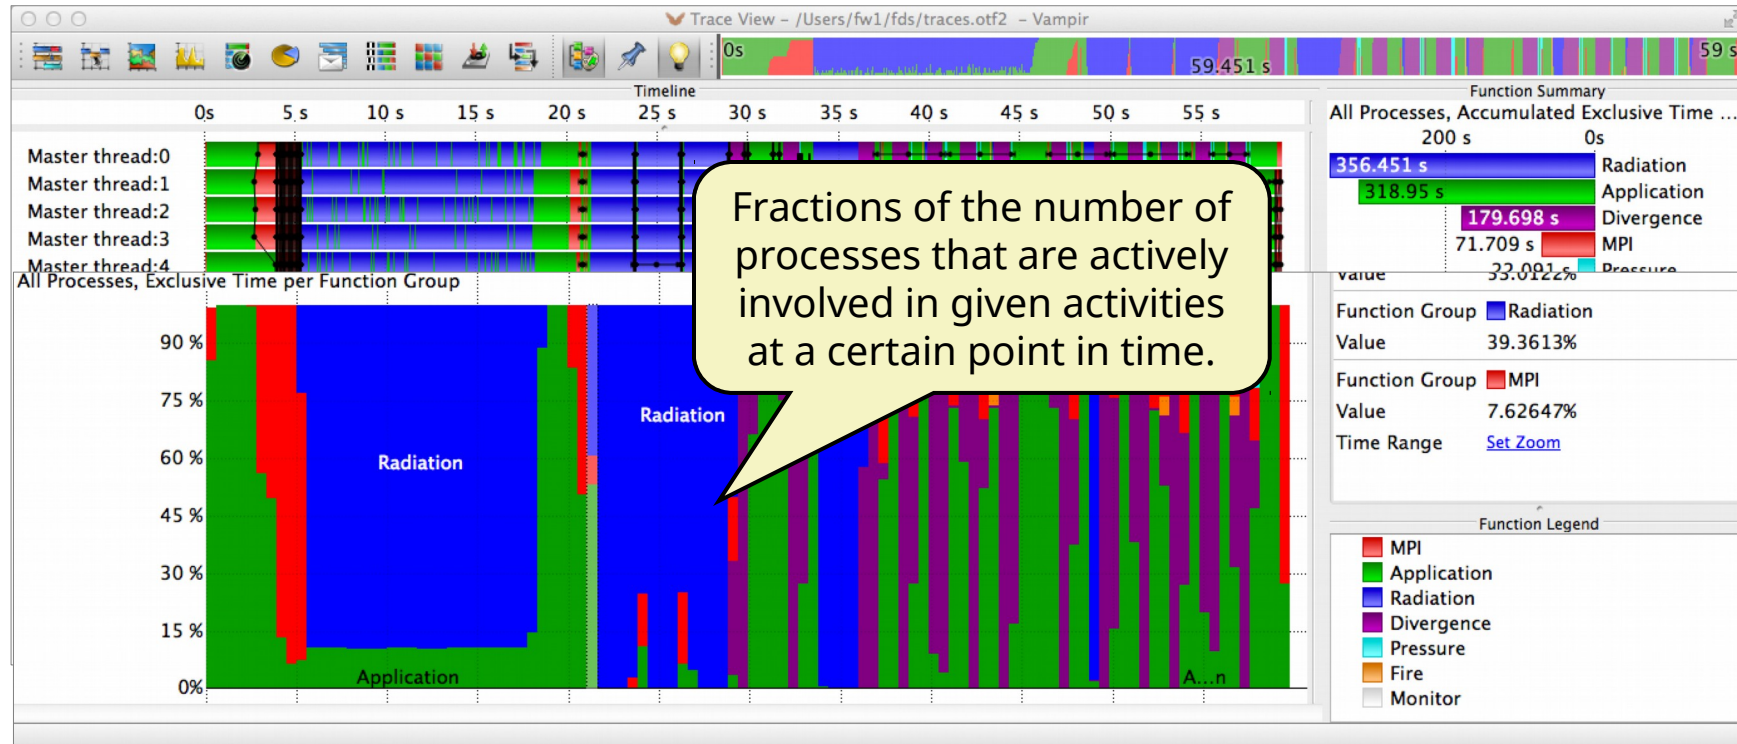
Master Timeline



Vampir Performance Charts in Detail



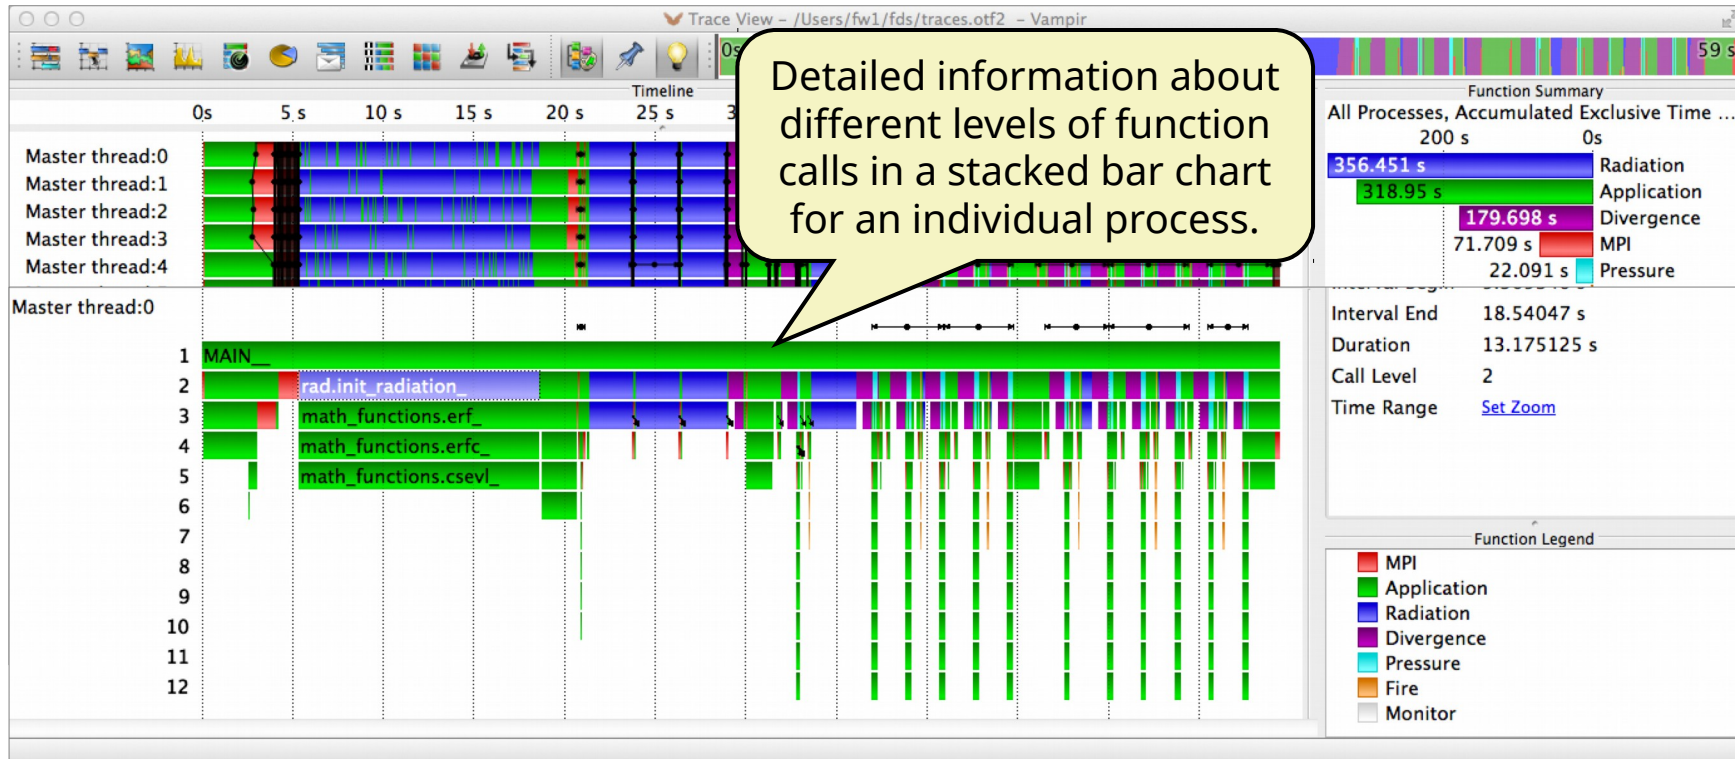
Summary Timeline



Vampir Performance Charts in Detail



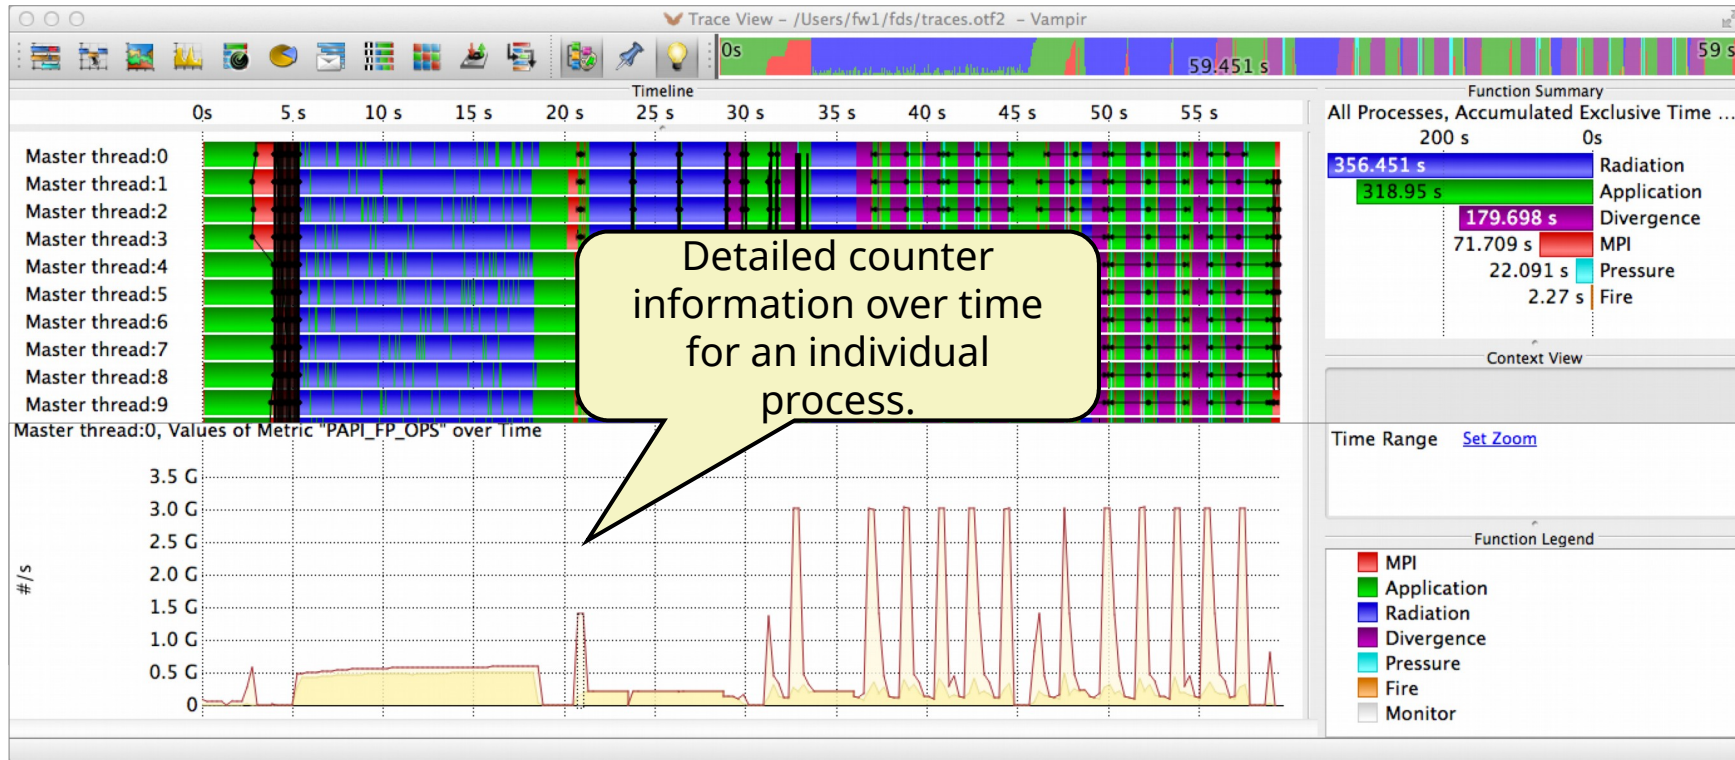
Process Timeline



Vampir Performance Charts in Detail



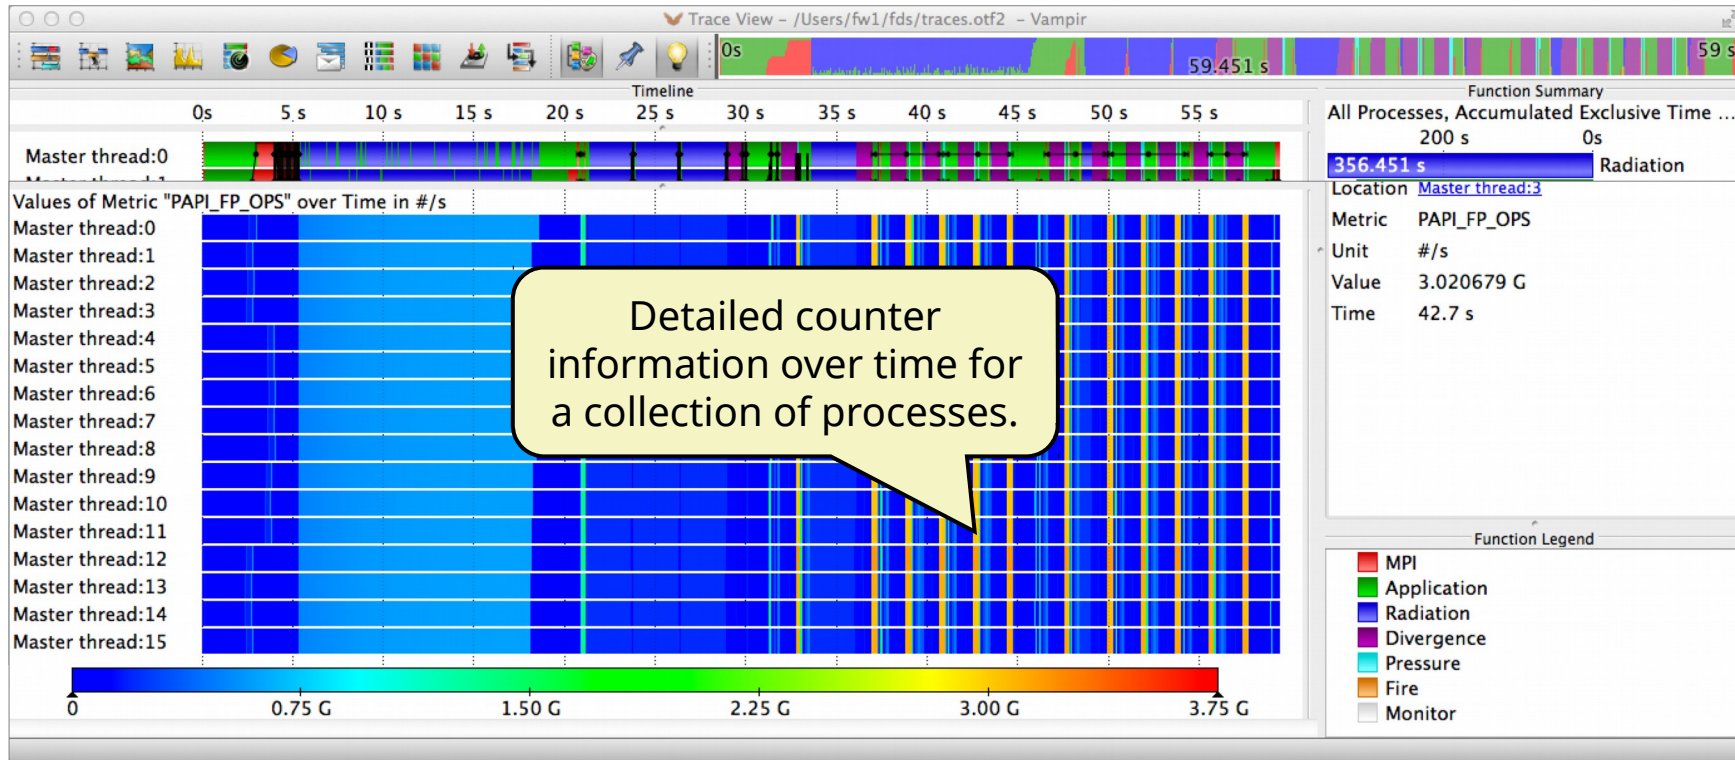
Counter Timeline



Vampir Performance Charts in Detail



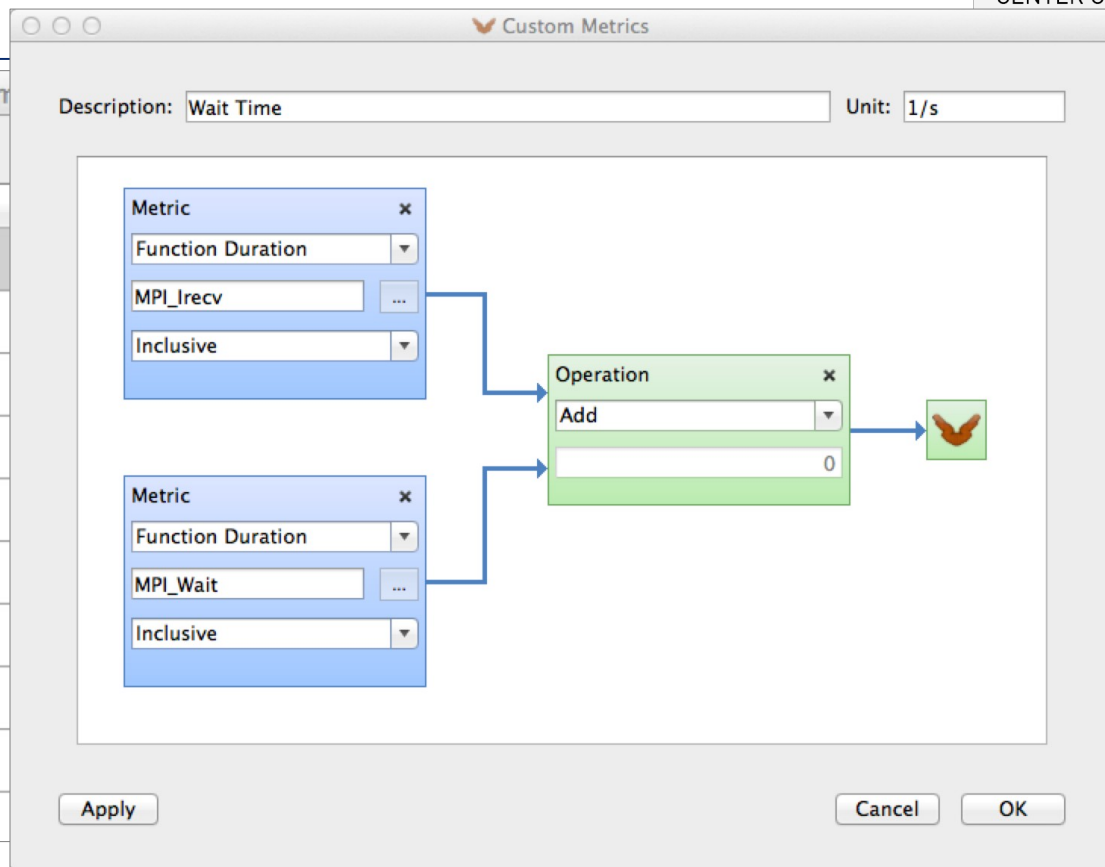
Performance Radar



Vampir Performance Metrics

Where do they come from?

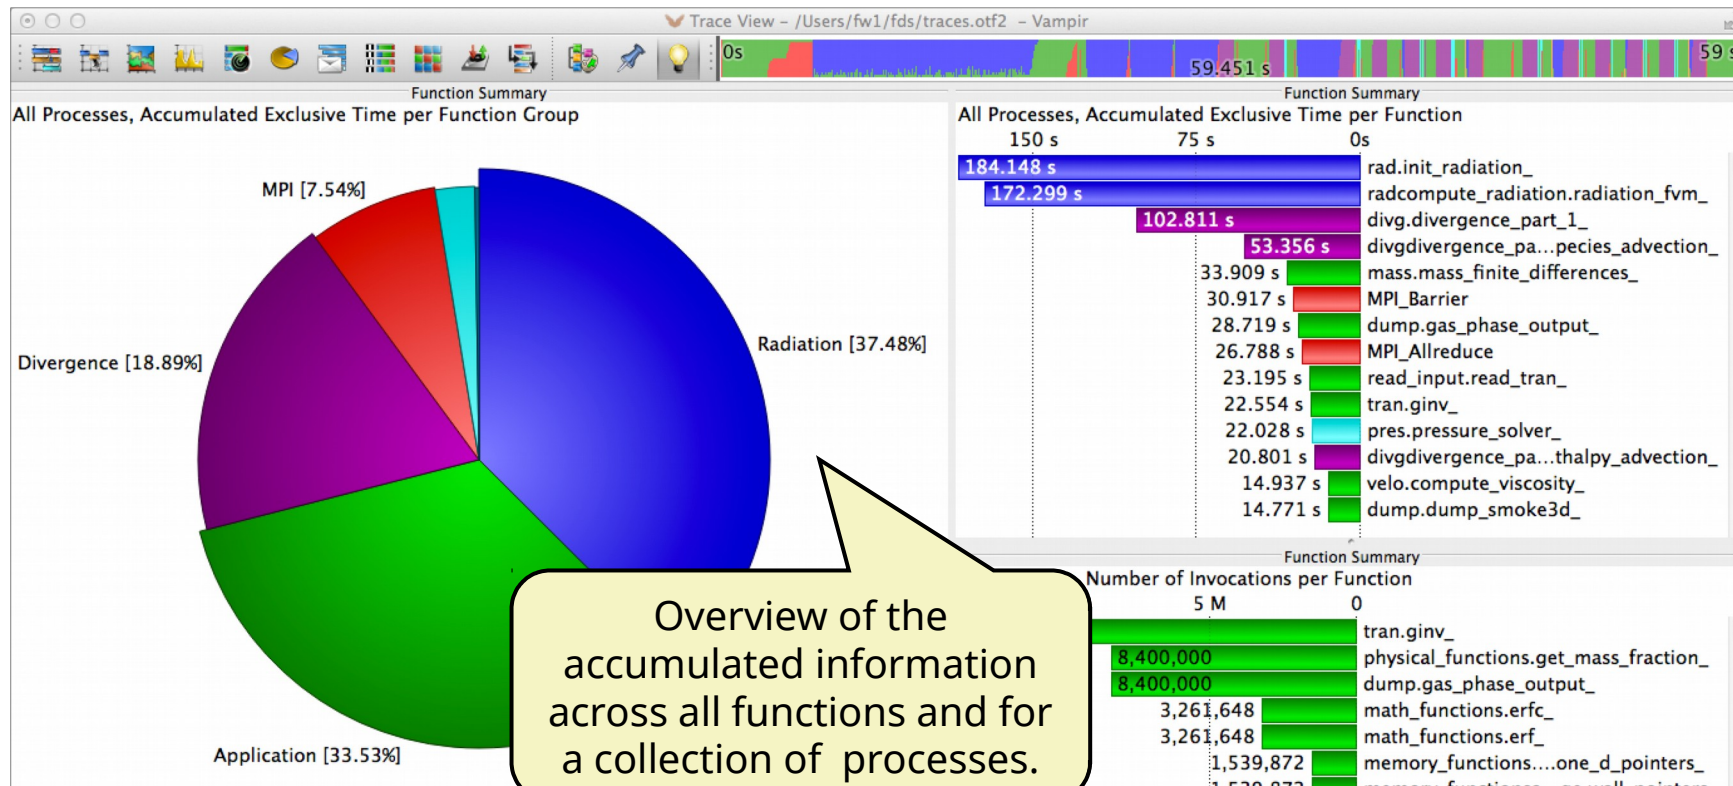
Active	Description
<input checked="" type="checkbox"/>	FLOPS in User Defined Function
<input checked="" type="checkbox"/>	MPI Latencies
<input checked="" type="checkbox"/>	Message Data Rate
<input checked="" type="checkbox"/>	Message Transfer Times
<input checked="" type="checkbox"/>	Message Volume in Transit
<input checked="" type="checkbox"/>	Number of Hits
<input checked="" type="checkbox"/>	Number of Invocations
<input checked="" type="checkbox"/>	Simultaneous Messages
<input checked="" type="checkbox"/>	Time Spent in MPI_Wait



Vampir Performance Charts in Detail



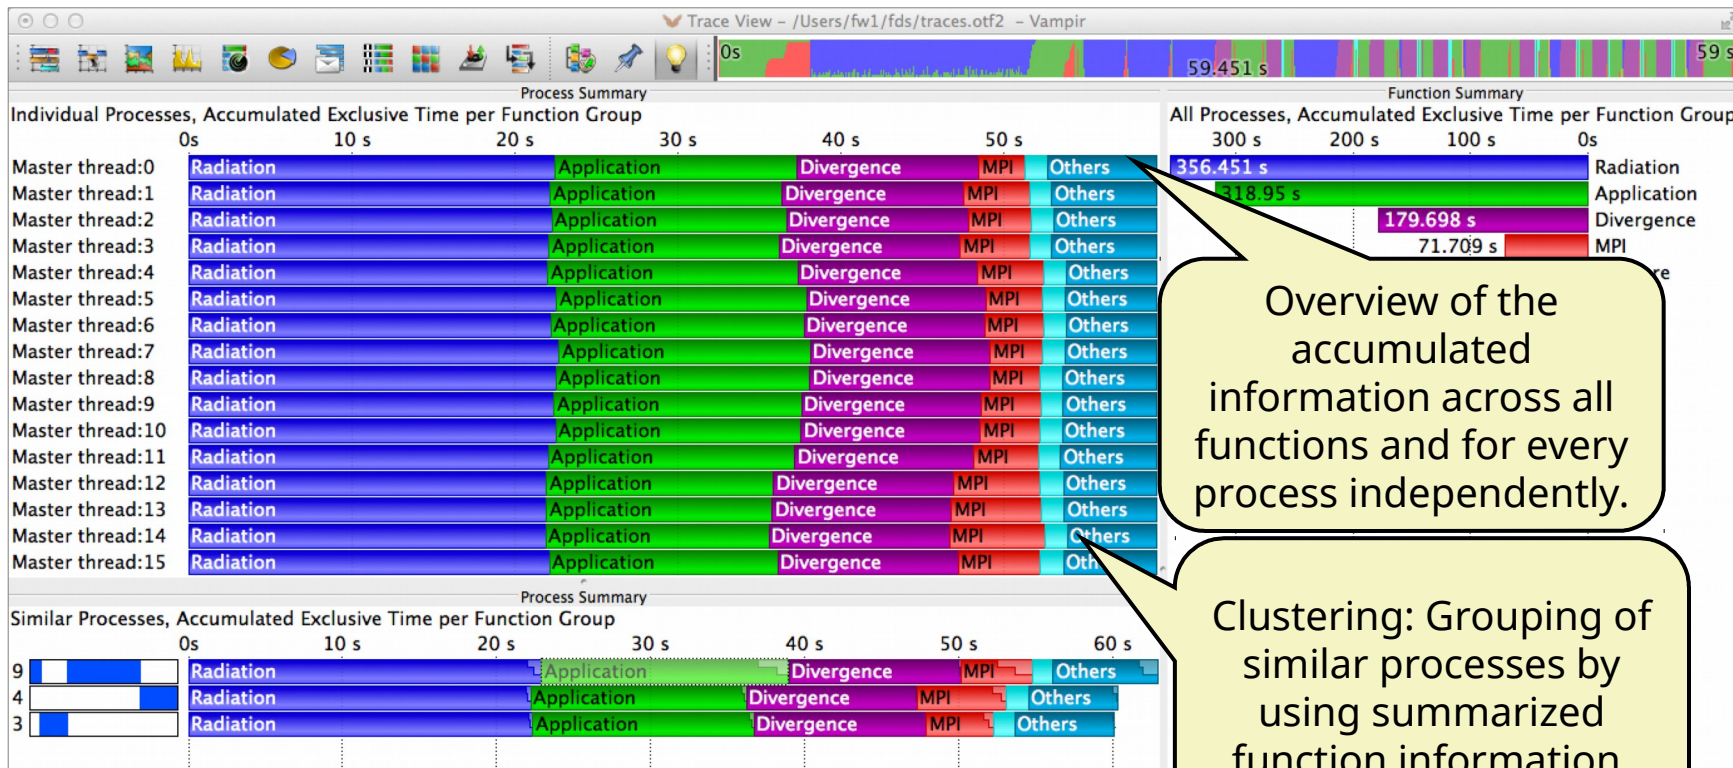
Function Summary



Vampir Performance Charts in Detail



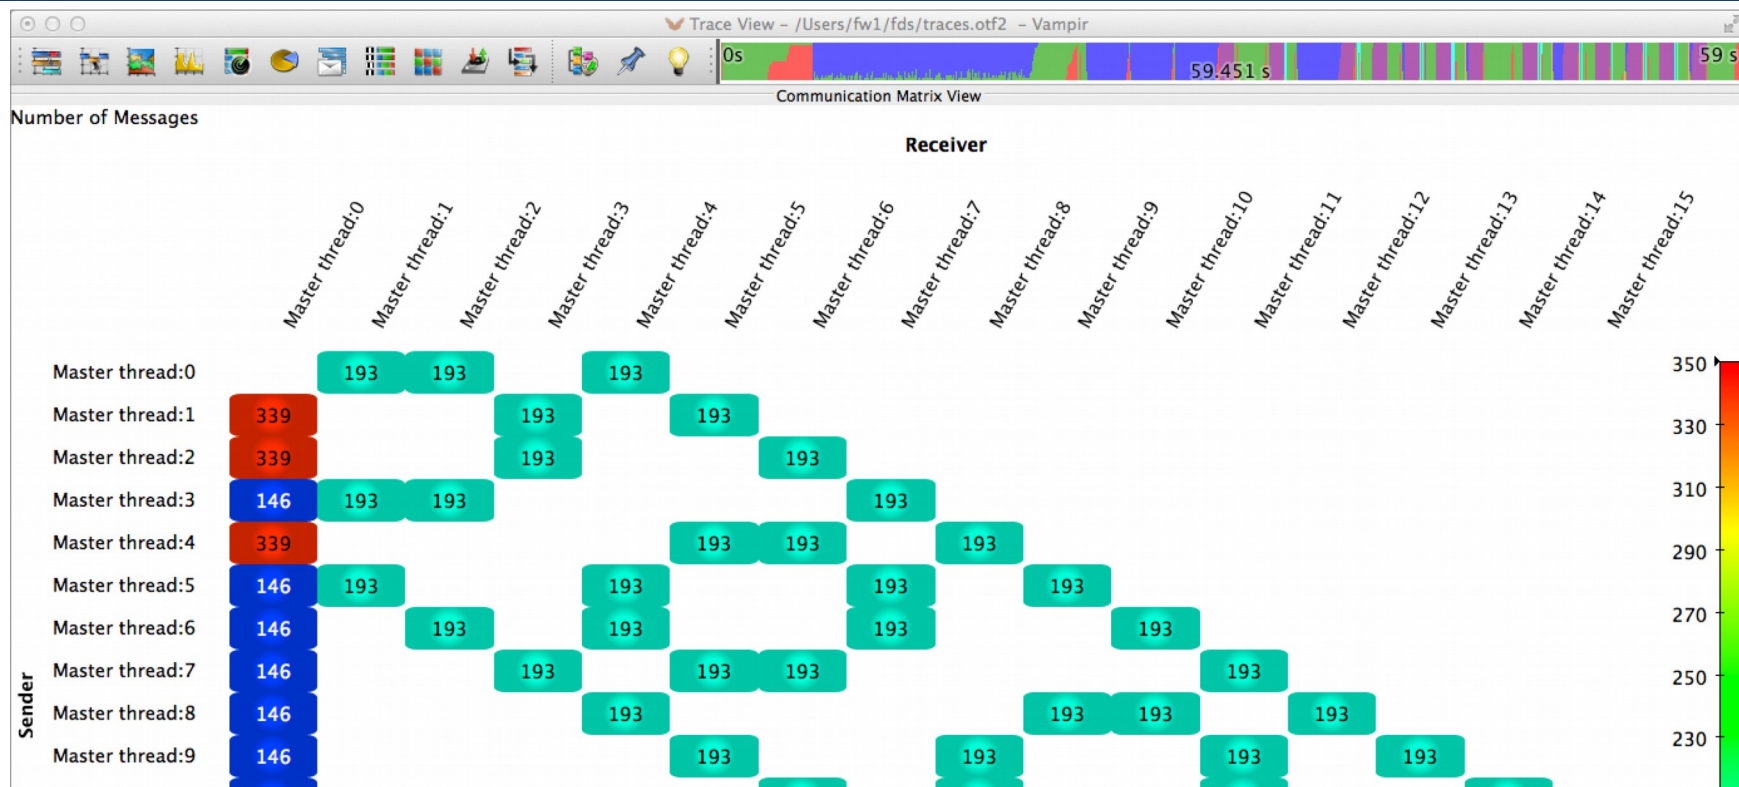
Process Summary



Vampir Performance Charts in Detail




Communication Matrix View



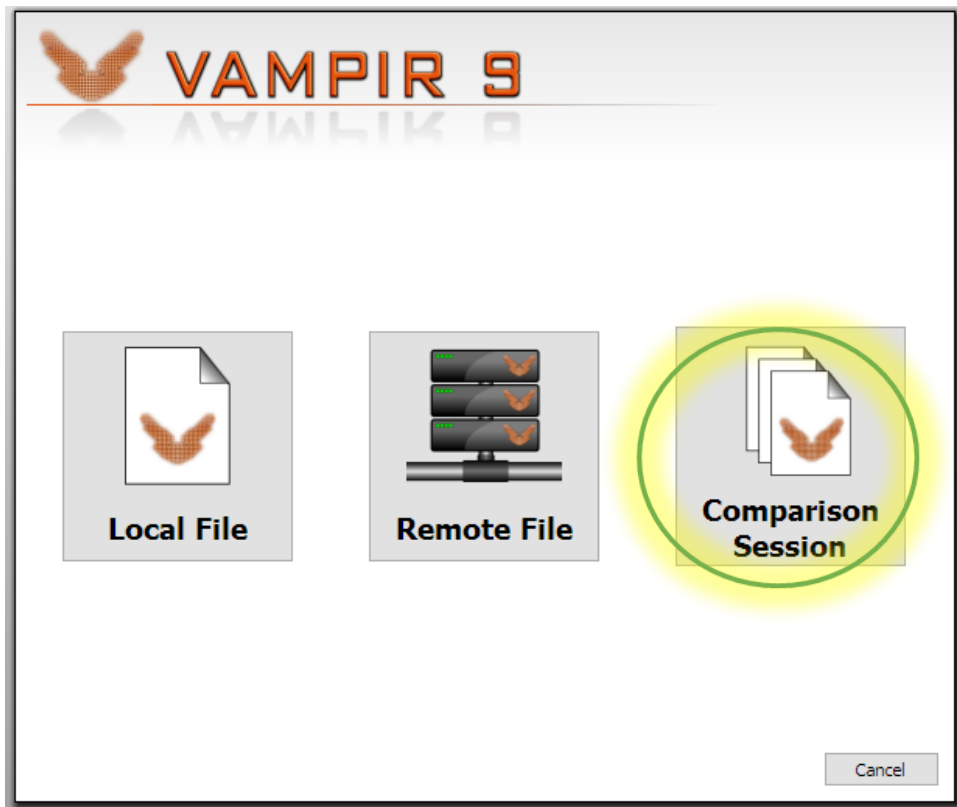
Vampir at Scale

Fit to chart height (feat. 200,000+ event streams)

Process 13730
Process 20594
Process 27459
Process 34324
Process 41188
Process 48053
Process 54918
Process 61782
Process 68647
Process 75512
Process 82376
Process 89241
Process 96106
Process 102970
Process 109835
Process 116700
Process 123564
Process 130429
Process 137294
Process 144158
Process 151023
Process 157888
Process 164752
Process 171617
Process 178482
Process 185346
Process 192211



Comparing Traces with Vampir


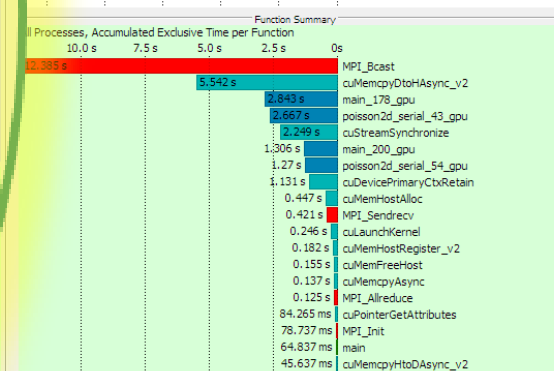








Seeing the differences



GPU


File Edit View Filter Window Help




Zooming in



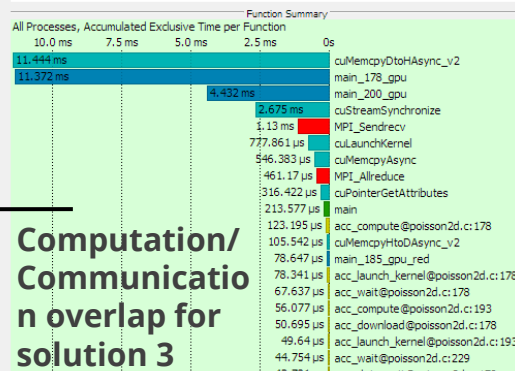
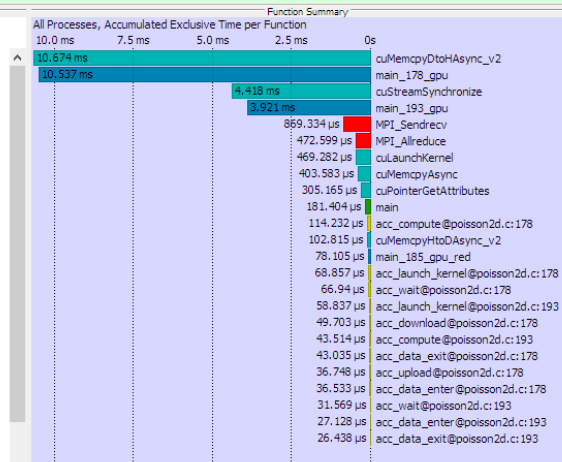


GPU



One iteration of solution1



One iteration of solution2



Computation/
Communication
overlap for
solution 3

DEMO:

Visualizing Trace Files with Vampir

Looking into DL Frameworks

Score-P Python Bindings

Tracing/Profiling for all python programs


- Not yet included in main release
- Available on GitHub:
 - https://github.com/score-p/scorep_binding_python
- NSight/nvvp for single node DL frameworks still better (user instrumentation)
- Score-P only choice for MPI-parallel DL frameworks


```
$ export SCOREP_ENABLE_PROFILING=true  
$ export SCOREP_ENABLE_TRACING=false  
$ export SCOREP_EXPERIMENT_DIRECTORY=profile  
  
$ python -m scorep --mpi <script.py>
```

Profiling Example

Vampir with Python Traces

It looks all the same





Vampir is available at <http://www.vampir.eu>

Vampir at IU: <https://kb.iu.edu/d/awbv>

Get support via vampirsupport@zih.tu-dresden.de

Score-P: <http://www.vi-hps.org/projects/score-p>