



SOFTWARE

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

PRODUCTIVITY



FAST SOLUTIONS

- PAPI_L1_ICM
- PAPI_L2_DCM
- PAPI_L2_ICM
- PAPI_L1_TCM

VAMPIR & VAMPIRTRACE Hands On

DKRZ Tutorial 2012 in Hamburg
June, 2012

Ronny Tschüter

Slides by: Andreas Knüpfer, Jens Doleschal,
ZIH, Technische Universität Dresden

- Copy NPB sources to your home directory

```
% cp -rp /work/kg0166/tutorial/NPB3.3-MZ-MPI $WRKDIR
```

- Move into source directory

```
% cd NPB3.3-MZ-MPI
```

- Select the VampirTrace compiler wrappers

```
% vim config/make.def  
-> comment out line 34, resulting in:  
...  
34: #MPIF77 = mpxf  
...  
-> remove the comment from line 44, resulting in:  
...  
44: MPIF77 = $(PREP) mpxf  
...
```

- Set up VampirTrace environment

```
% module load vampirtrace
```

- Build benchmark

```
% make clean; make suite PREP="vtf77 -vt:f77"
```

- Go to bin.vampir directory

```
% cd bin.vampir
```

- Submit your application to the batch system

```
% cp ../jobscript/blizzard/vt.ll .  
% llsubmit vt.ll
```

- Investigate the output file

```
% cat nas_bt_mz.job.e <ID>
:
[0]VampirTrace: Maximum number of buffer flushes reached \
(VT_MAX_FLUSHES=1)
[0]VampirTrace: Tracing switched off permanently
...
```

Usually occurs while tracing your application for the first time – don't worry !

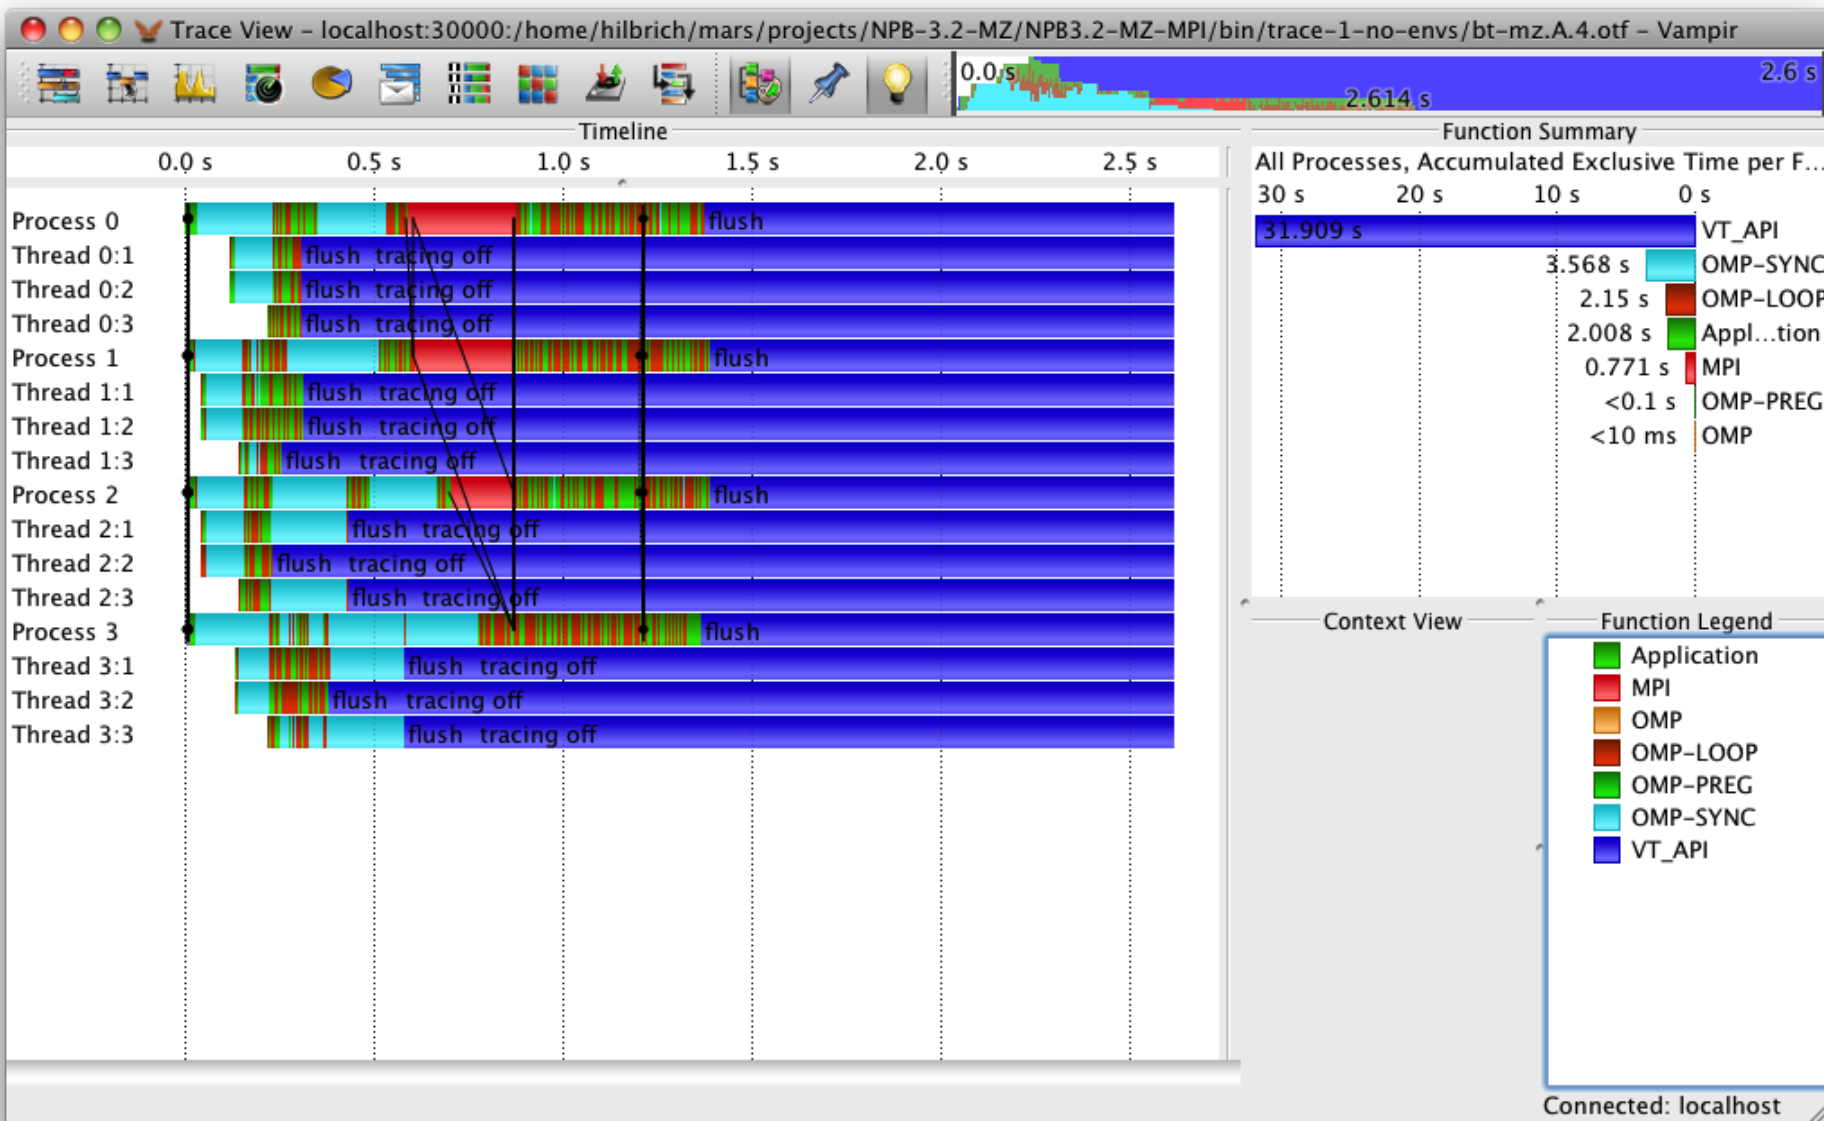
- Resulting trace files

```
% ls  
bt-mz_W.4  
nas_bt_mz.0.def.z  
nas_bt_mz.1.events.z  
nas_bt_mz.10001.events.z  
nas_bt_mz.10002.events.z  
...  
nas_bt_mz.40004.events.z  
nas_bt_mz.50001.events.z  
nas_bt_mz.50002.events.z  
nas_bt_mz.50003.events.z  
nas_bt_mz.50004.events.z  
nas_bt_mz.otf
```

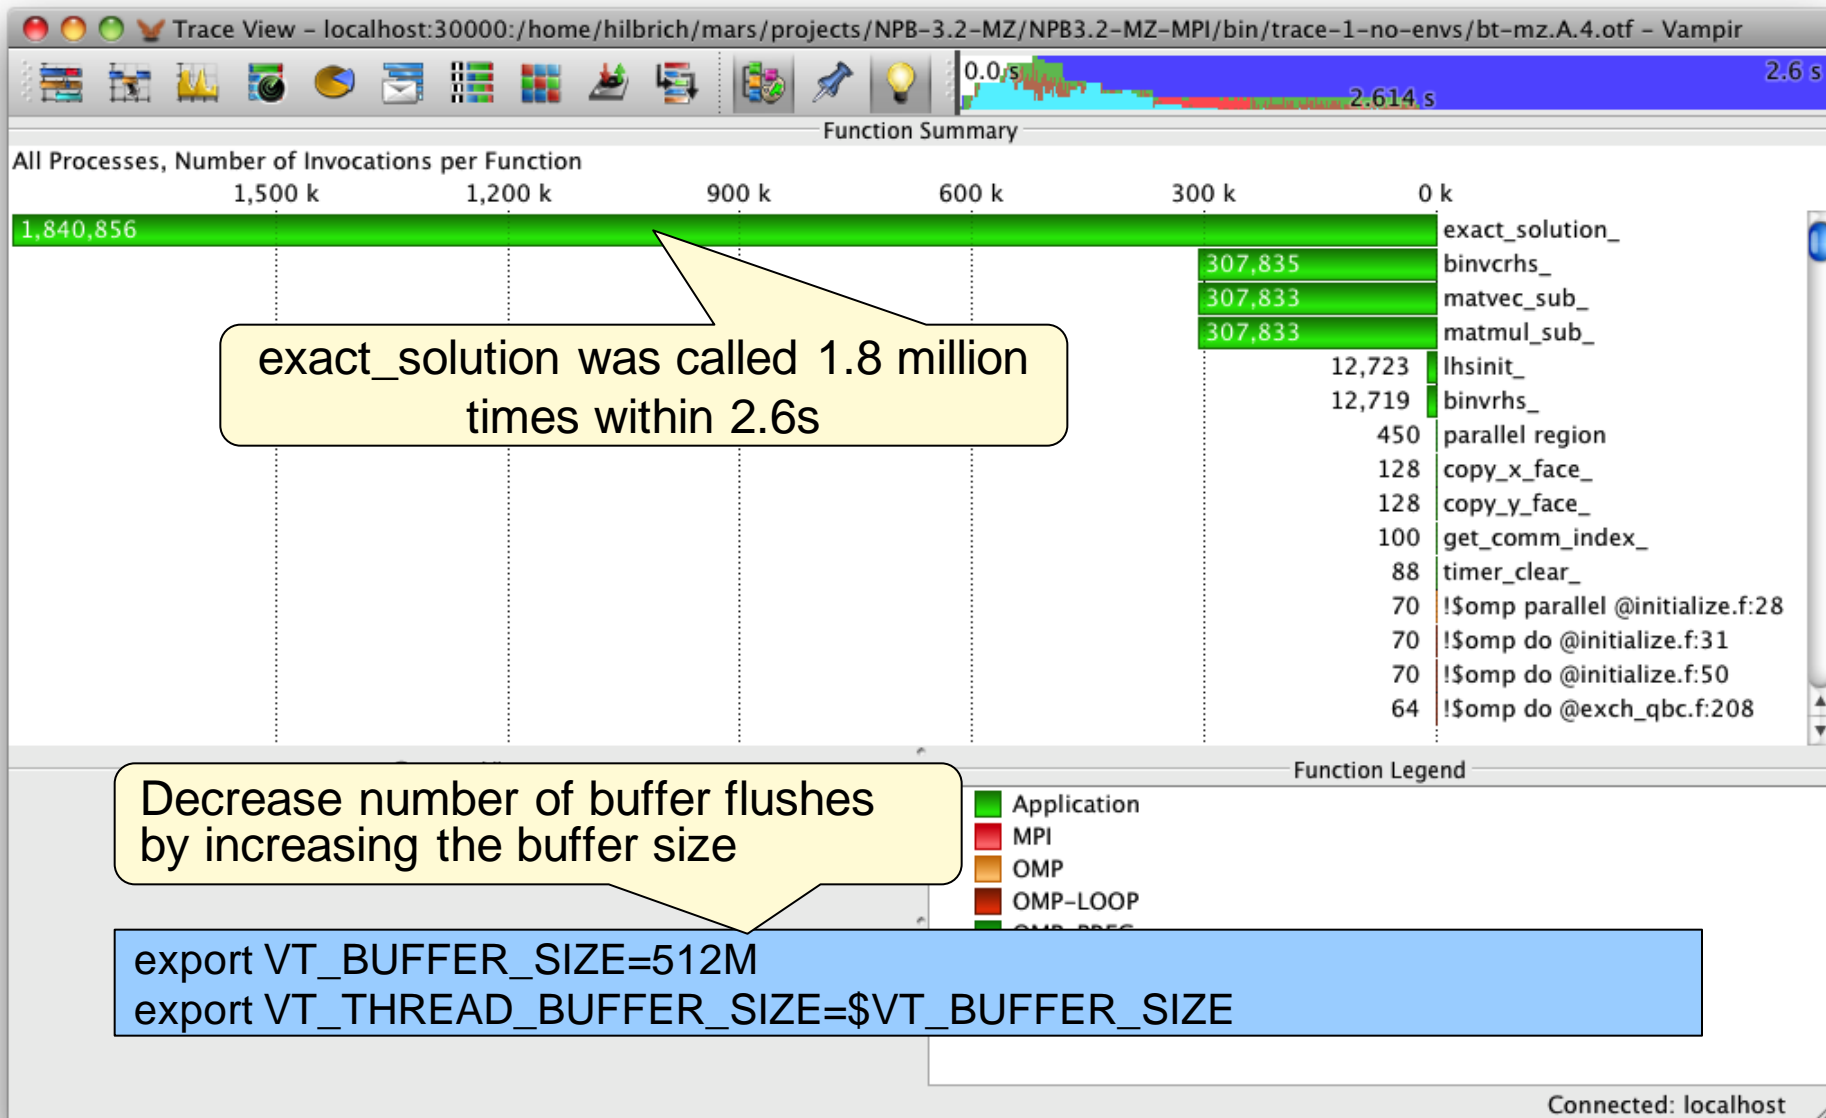
- Visualization with Vampir7 on lizard:

```
% ssh -X lizard.dkrz.de  
% module load vampir  
% vampir
```

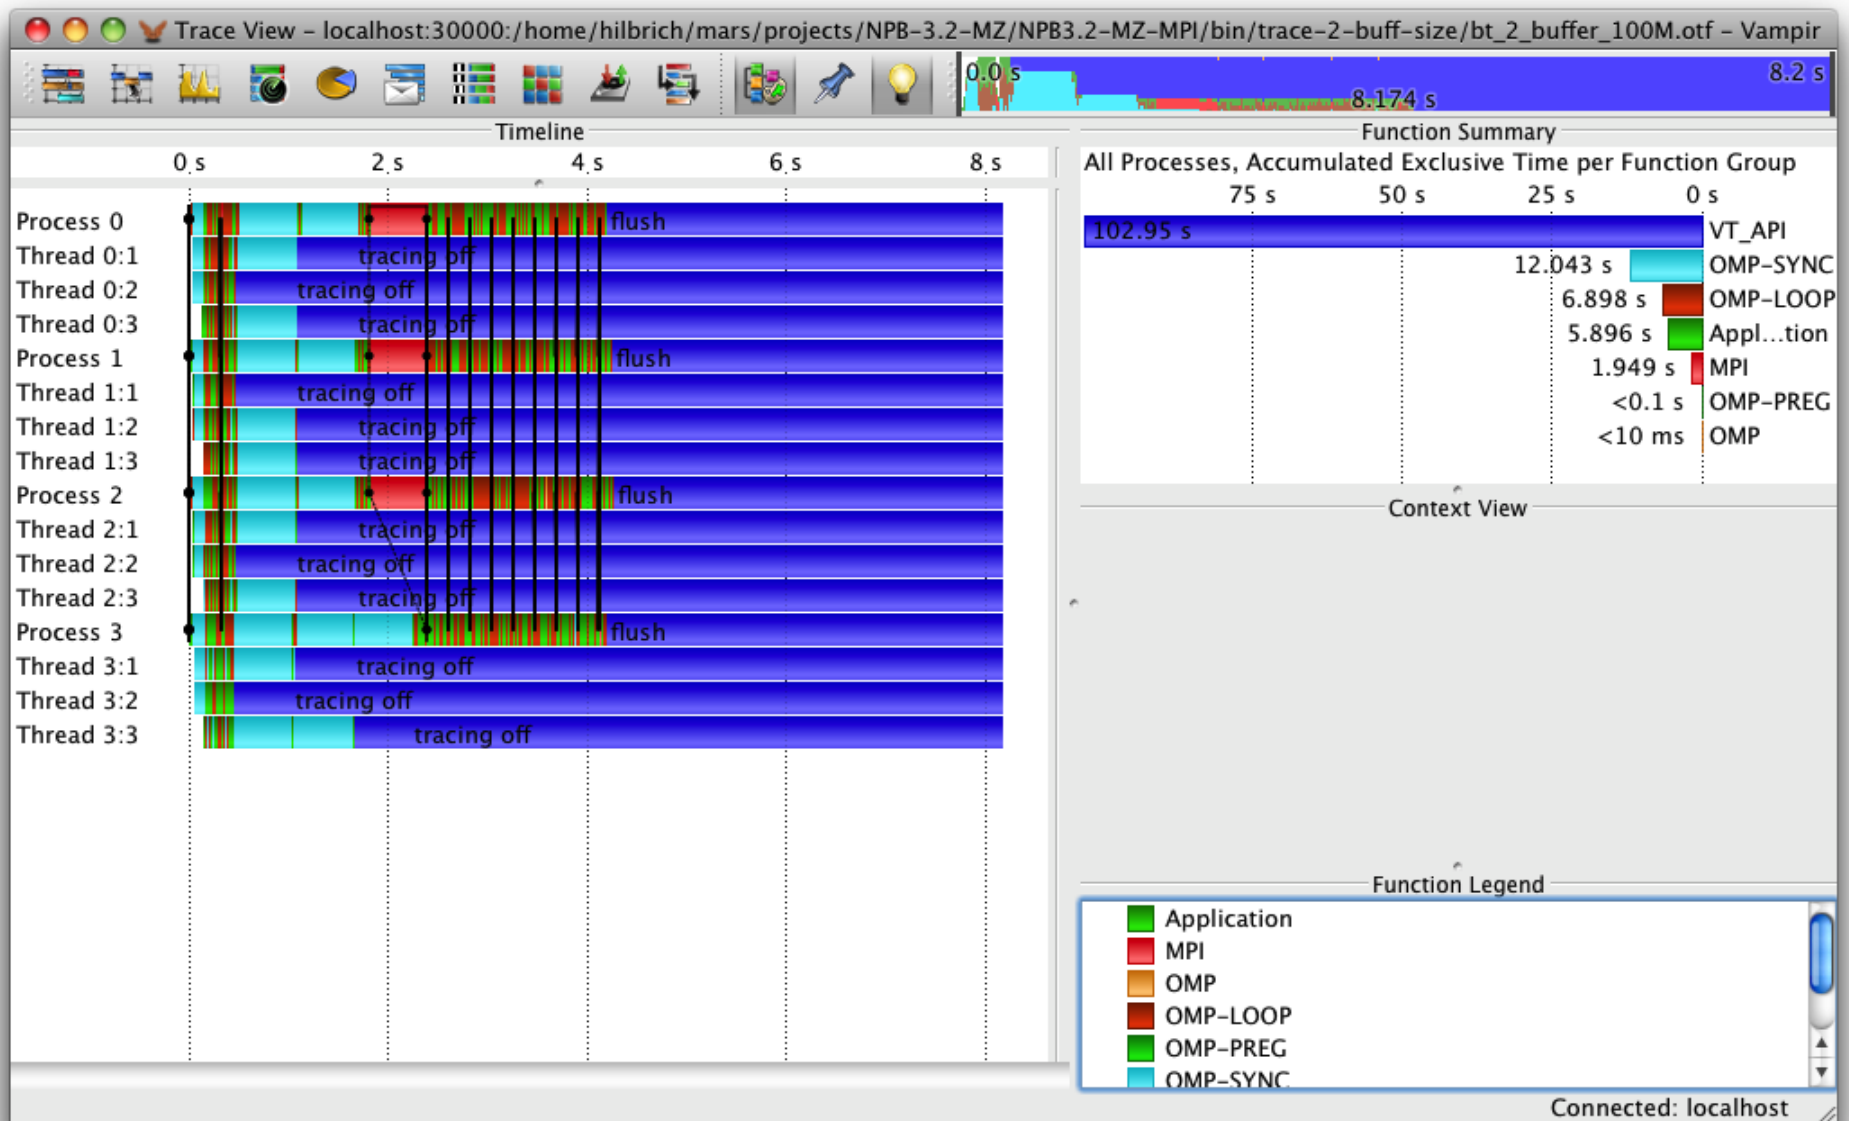
Hands-on: NPB – Initial Trace



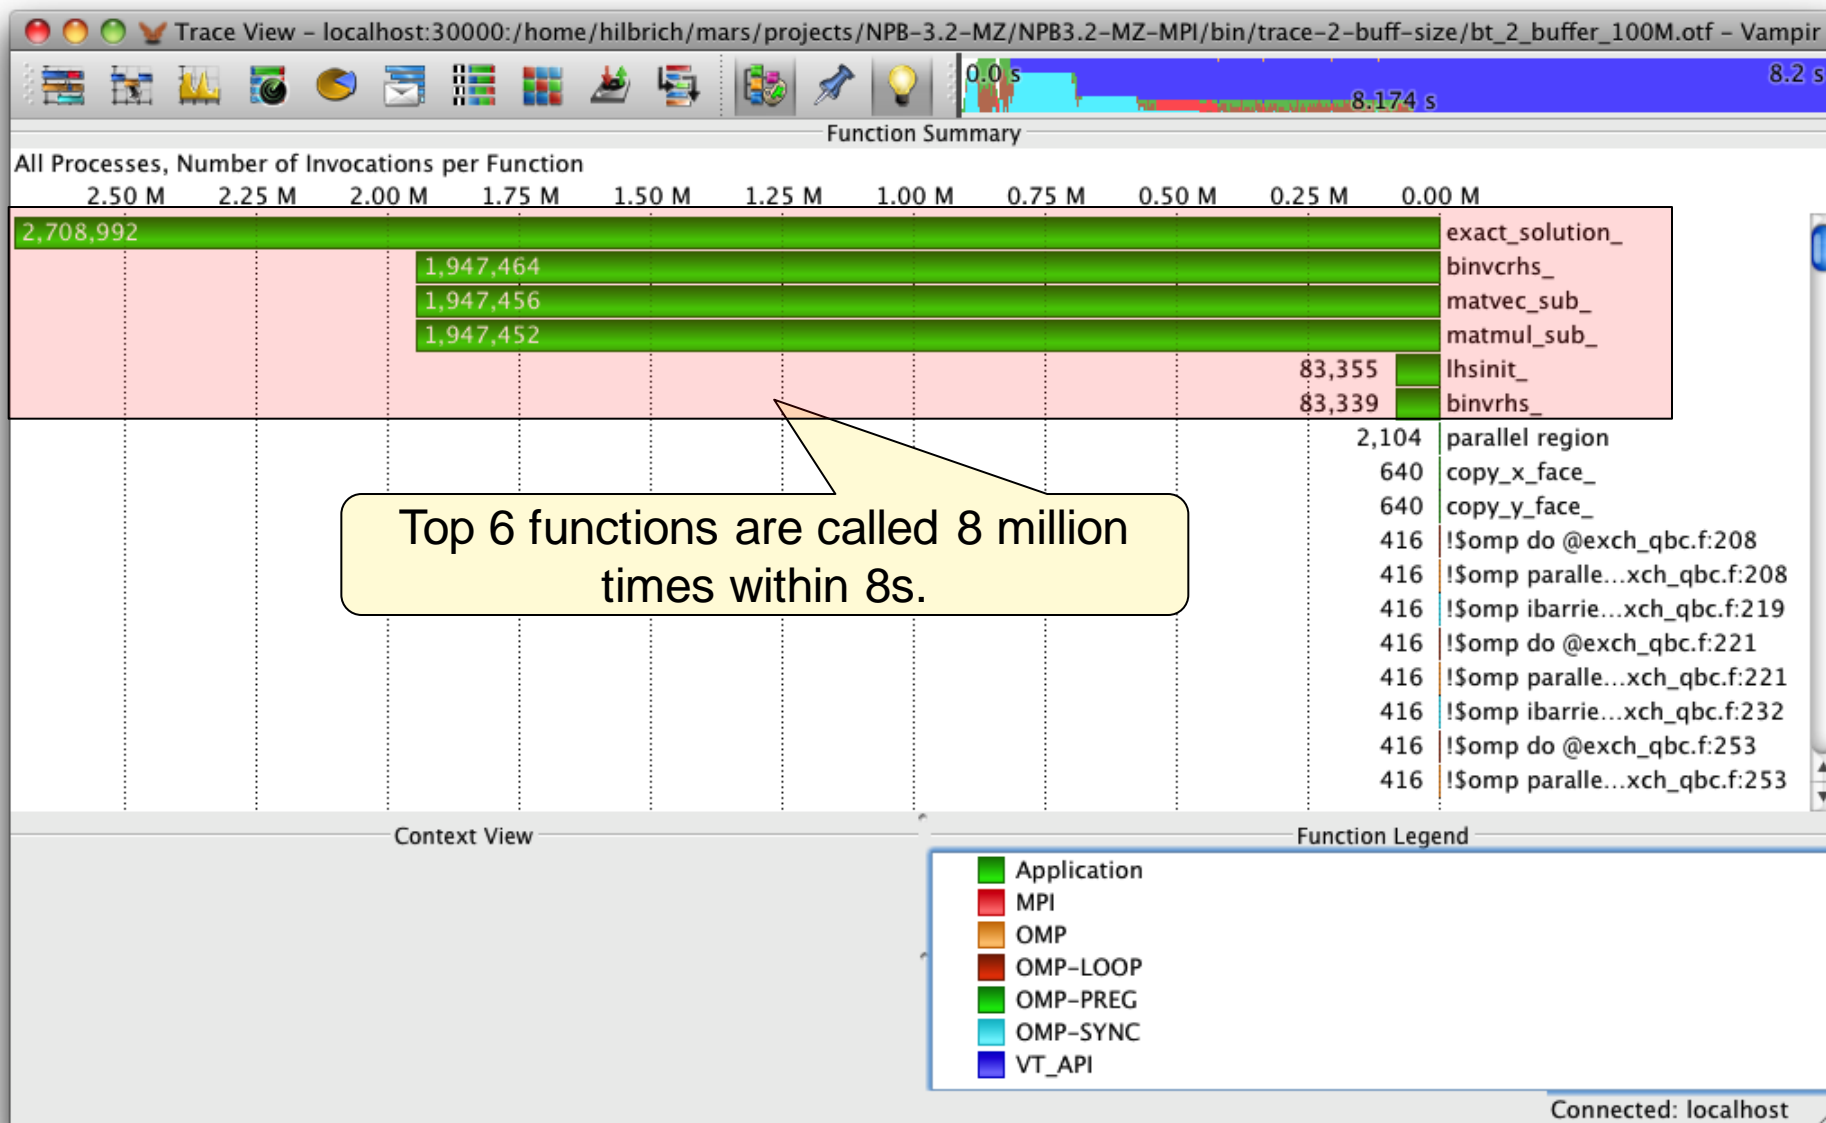
Hands-on: NPB – Initial Trace



Hands-on: NPB – Second Trace



Hands-on: NPB – Second Trace



- Limit trace size with filtering
- Environment variable **VT_FILTER_SPEC**

Example, don't type it into the shell

```
% export VT_FILTER_SPEC=/home/user/filter.spec
```

- Filter definition file contains a list of filters

```
my_*;test_* -- 1000  
debug_* -- 0  
calculate -- -1  
* -- 1000000
```

- See also the **vtfilter** tool
 - can generate a customized filter file
 - can reduce the size of existing trace files

- Go to bin.vampir directory
- Submit your application to the batch system

```
% cp ../jobscript/blizzard/vt_filter.ll .  
% llsubmit vt_filter.ll
```

Uses a list of filters:

```
binvrhs*  
matmul_sub*  
matvec_sub*  
exact_solution*  
binvrhs*  
lhs*init*  
timer_*
```

- Use profiling mode to get an overview

```
export VT_MODE=STAT
```

Add this to your jobscript

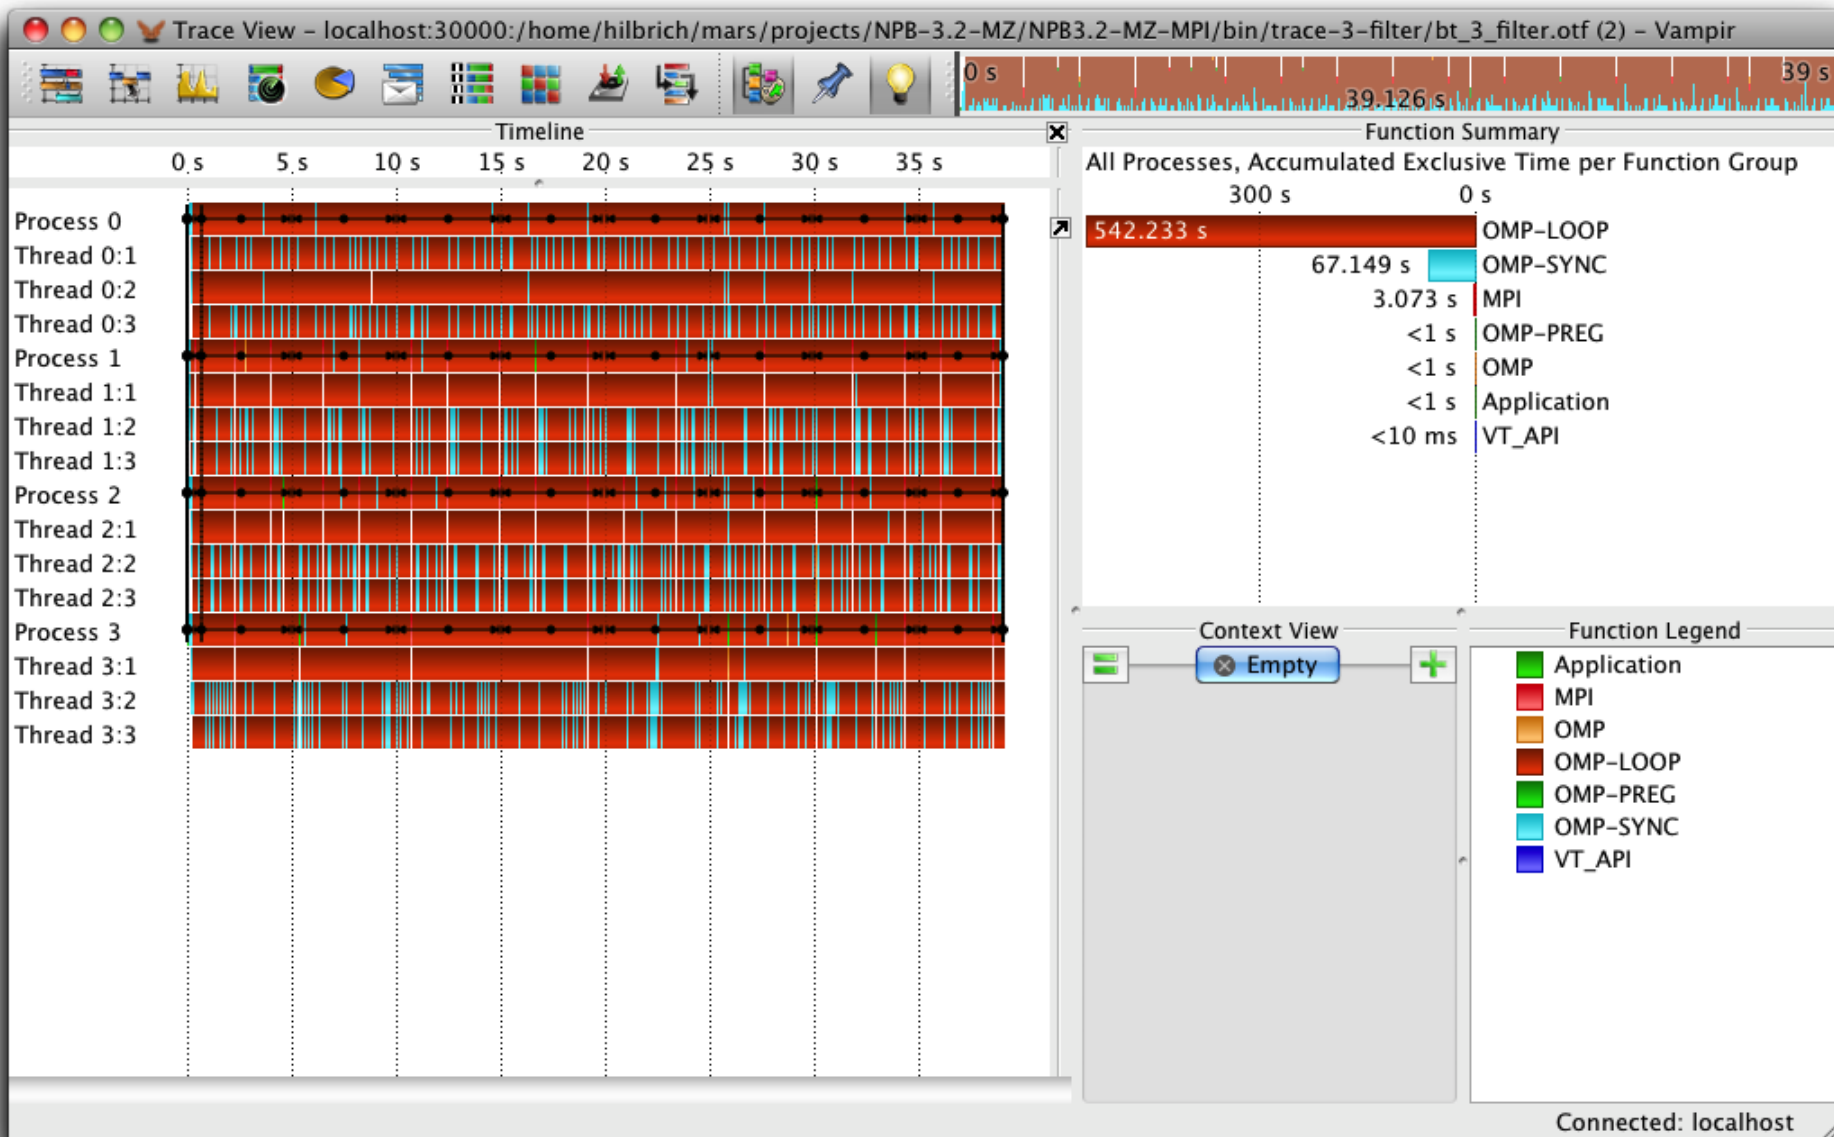
- Investigate result file

```
% cat nas_bt_mz.prof.txt
```

			excl. time	incl. time	
*excl. time	incl. time	calls	/ call	/ call	name
13.915s	13.915s	32677776	0.425us	0.425us	binvcrhs_
9.358s	20.359s	3216	2.910ms	6.331ms	!\$omp do @z_solve.f:52
9.141s	19.822s	3216	2.842ms	6.164ms	!\$omp do @y_solve.f:52
8.916s	8.916s	32677776	0.272us	0.272us	matmul_sub_
...					

- Generate filter specification

Hands-on: NPB – Filtered Trace



Hands-on: NPB – Filtered Trace



Trace View - localhost:30000:/home/hilbrich/mars/projects/NPB-3.2-MZ/NPB3.2-MZ-MPI/bin/trace-3-filter/bt_3_filter.otf (2) - Vampir

Timeline: 0 s, 5 s, 10 s, 15 s, 20 s, 25 s, 30 s, 35 s

Process 0
Thread 0:1
Thread 0:2
Thread 0:3
Process 1
Thread 1:1
Thread 1:2
Thread 1:3
Process 2
Thread 2:1
Thread 2:2
Thread 2:3
Process 3
Thread 3:1
Thread 3:2
Thread 3:3

Function Summary
All Processes, Accumulated Exclusive Time per Function Group

Function Group	Time
OMP-LOOP	542.233 s
OMP-SYNC	67.149 s
MPI	3.073 s
OMP-PREG	<1 s
OMP	<1 s
Application	<1 s
VT_API	<10 ms

Context View

Property	Value
Display	Master Timelin...
Type	Function
Function	!\$omp do @y_
Function Group	OMP-LOOP
Interval Begin	15.418903 s
Interval End	15.44888 s
Duration	29.9778 ms
Source Location	y_solve.f:51

Function Legend

- Application
- MPI
- OMP
- OMP-LOOP
- OMP-PREG
- OMP-SYNC
- VT_API

16.7 s - 18.3 s (1.6 s)

Connected: localhost

Hands-on: NPB – Filtered Trace



Trace View - localhost:30000:/home/hilbrich/mars/projects/NPB-3.2-MZ/NPB3.2-MZ-MPI/bin/trace-3-filter/bt_3_filter.otf (2) - Vampir

Timeline: 17.0 s, 17.5 s, 18.0 s

Function Summary: All Processes, Accumulated Exclusive Time per Function Group

Function Group	Time
OMP-LOOP	22.804 s
OMP-SYNC	2.776 s
MPI	0.113 s
OMP-PREG	<0.1 s
OMP	<0.1 s
Application	<10 ms

Context View: Empty

Function Legend:

- Application
- MPI
- OMP
- OMP-LOOP
- OMP-PREG
- OMP-SYNC

Connected: localhost

Hands-on: NPB – Filtered Trace



Trace View - localhost:30000:/home/hilbrich/mars/projects/NPB-3.2-MZ/NPB3.2-MZ-MPI/bin/trace-3-filter/bt_3_filter.otf (2) - Vampir

Timeline

Function Summary

All Processes, Accumulated Exclusive Time per Function Group

Function Group	Time
OMP-LOOP	22.804 s
OMP-SYNC	2.776 s
MPI	0.113 s
OMP-PREG	<0.1 s
OMP	<0.1 s
Application	<10 ms

Context View

Property	Value
Display	Master Timelin
Type	Message
Origin	Process 3
Destination	Process 2
Start Time	17.404491 s
Arrival Time	17.404546 s
Duration	55.1 µs
Tag	10003
Size	66.171875 KiB
Data Rate	1.145307 GiB/
Communicator	MPI Communic

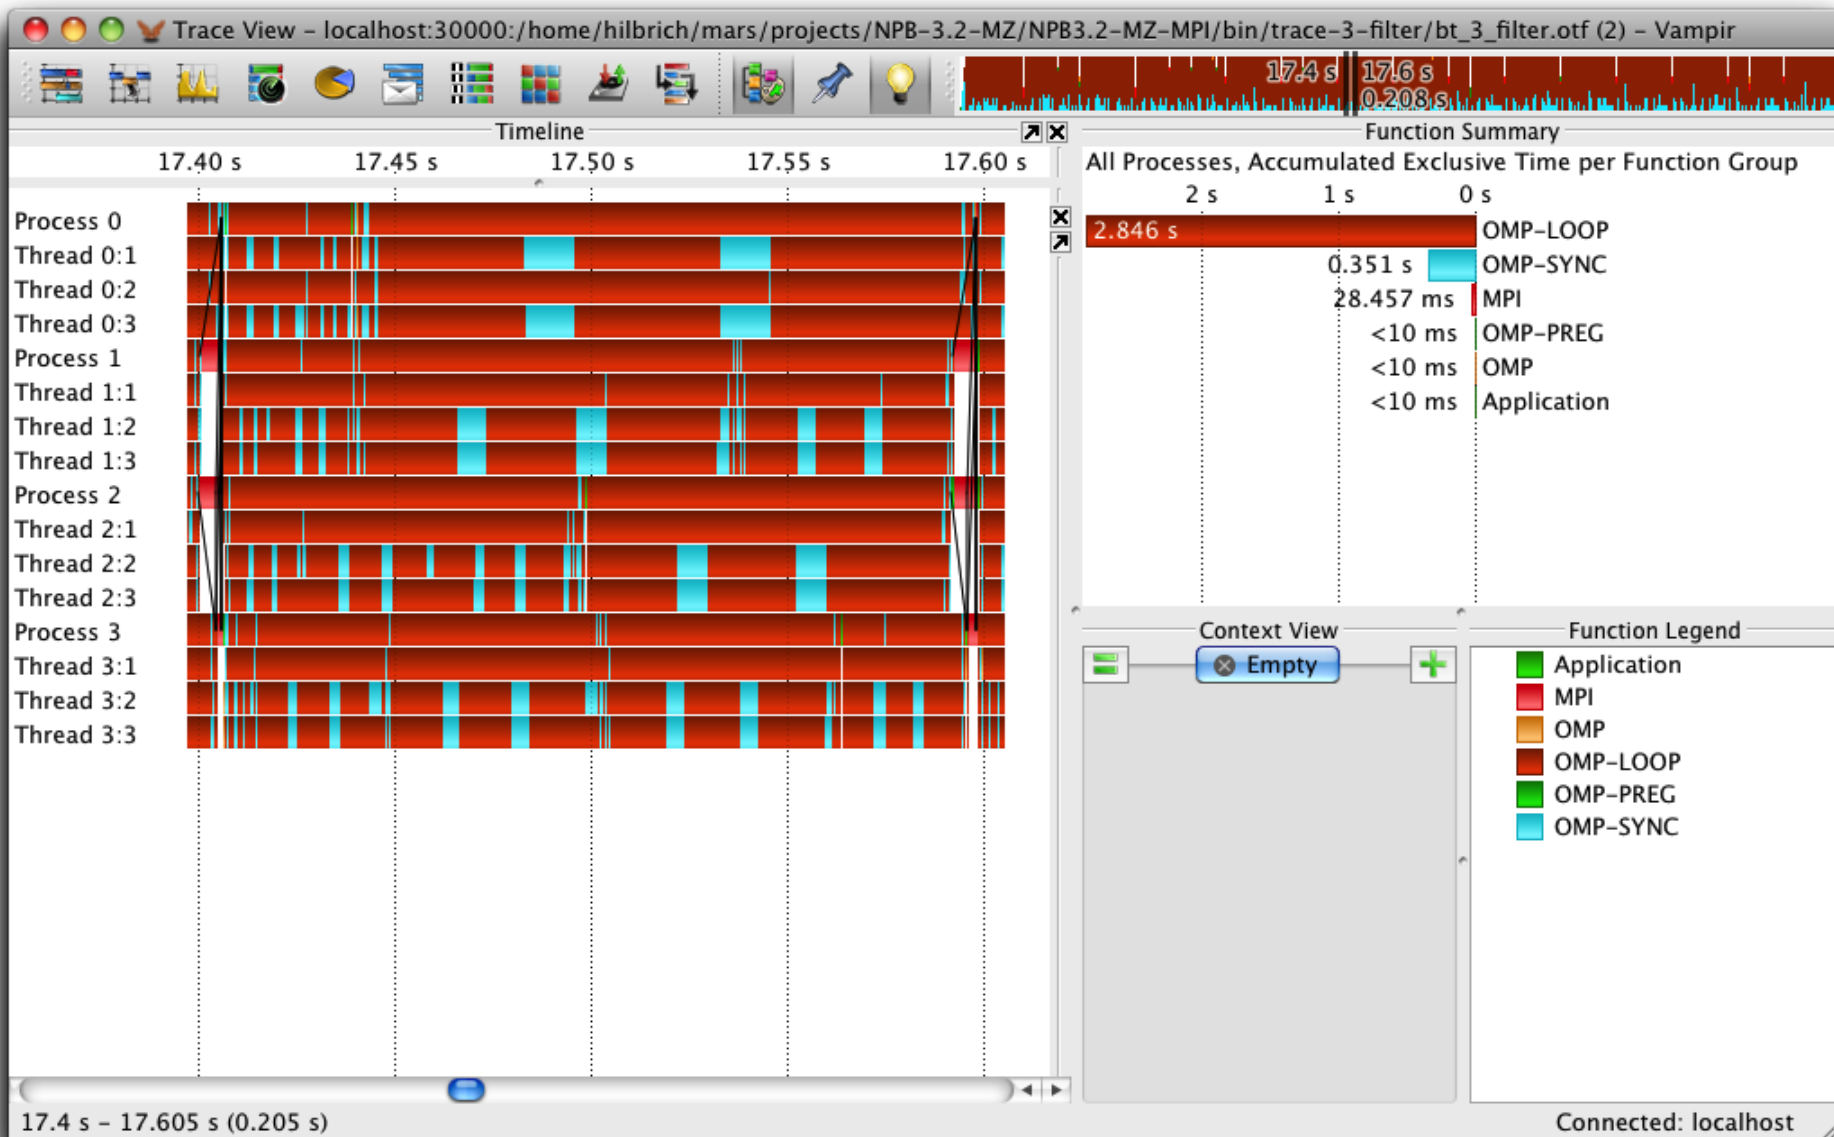
Function Legend

- Application
- MPI
- OMP
- OMP-LOOP
- OMP-PREG
- OMP-SYNC

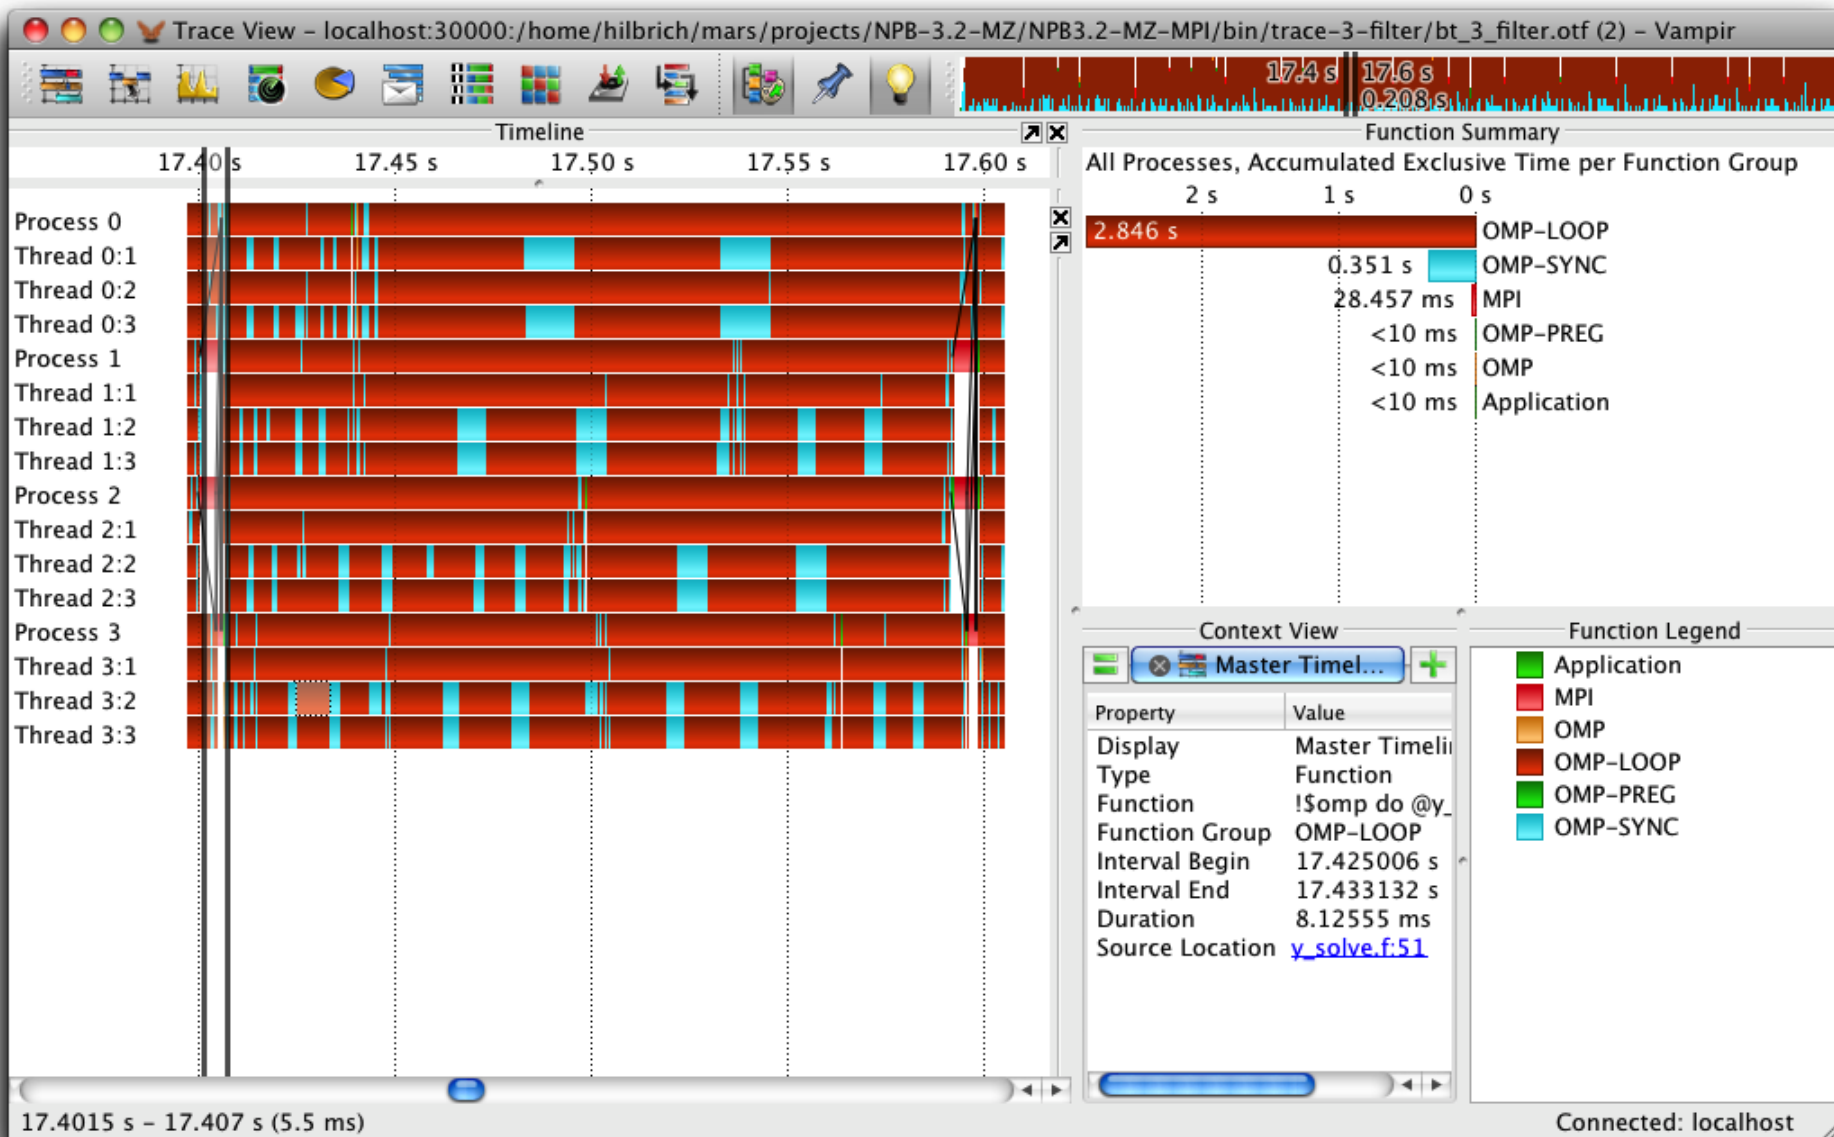
17.4 s - 17.605 s (0.205 s)

Connected: localhost

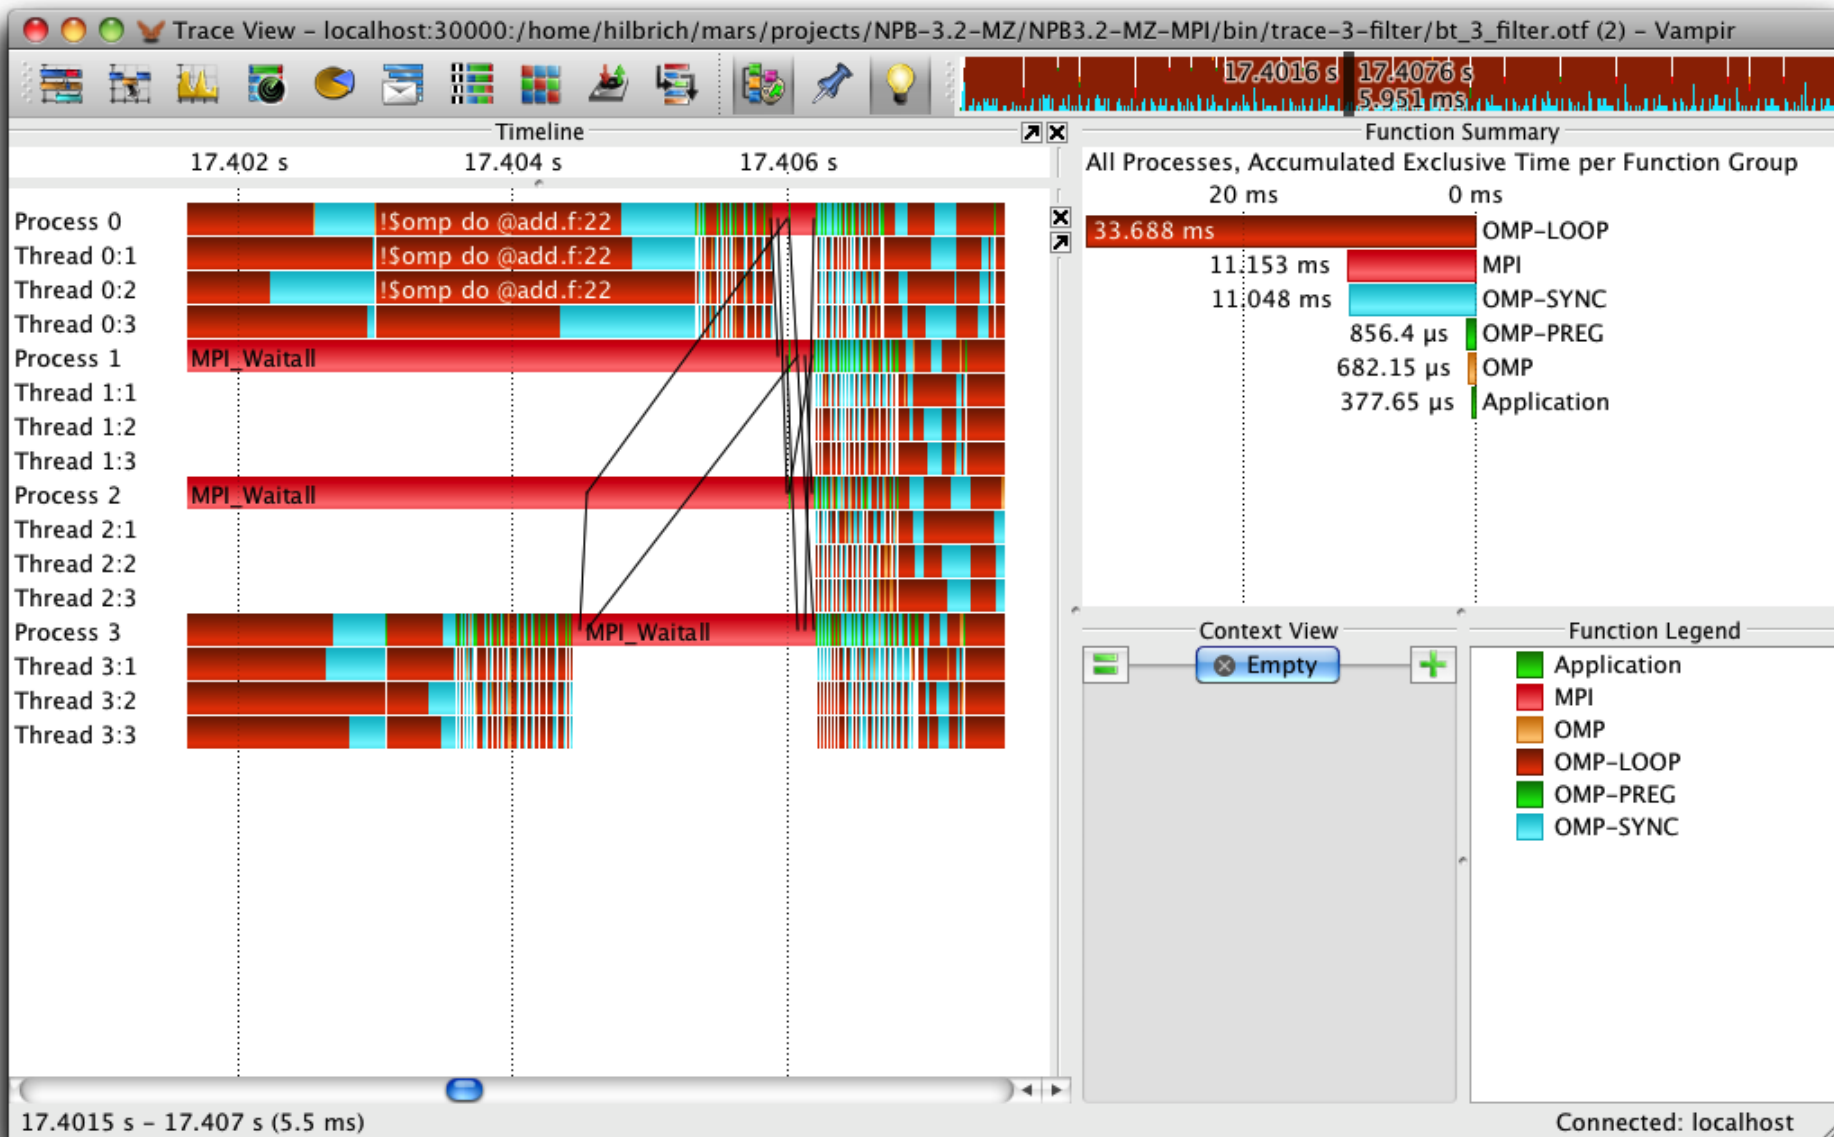
Hands-on: NPB – Filtered Trace



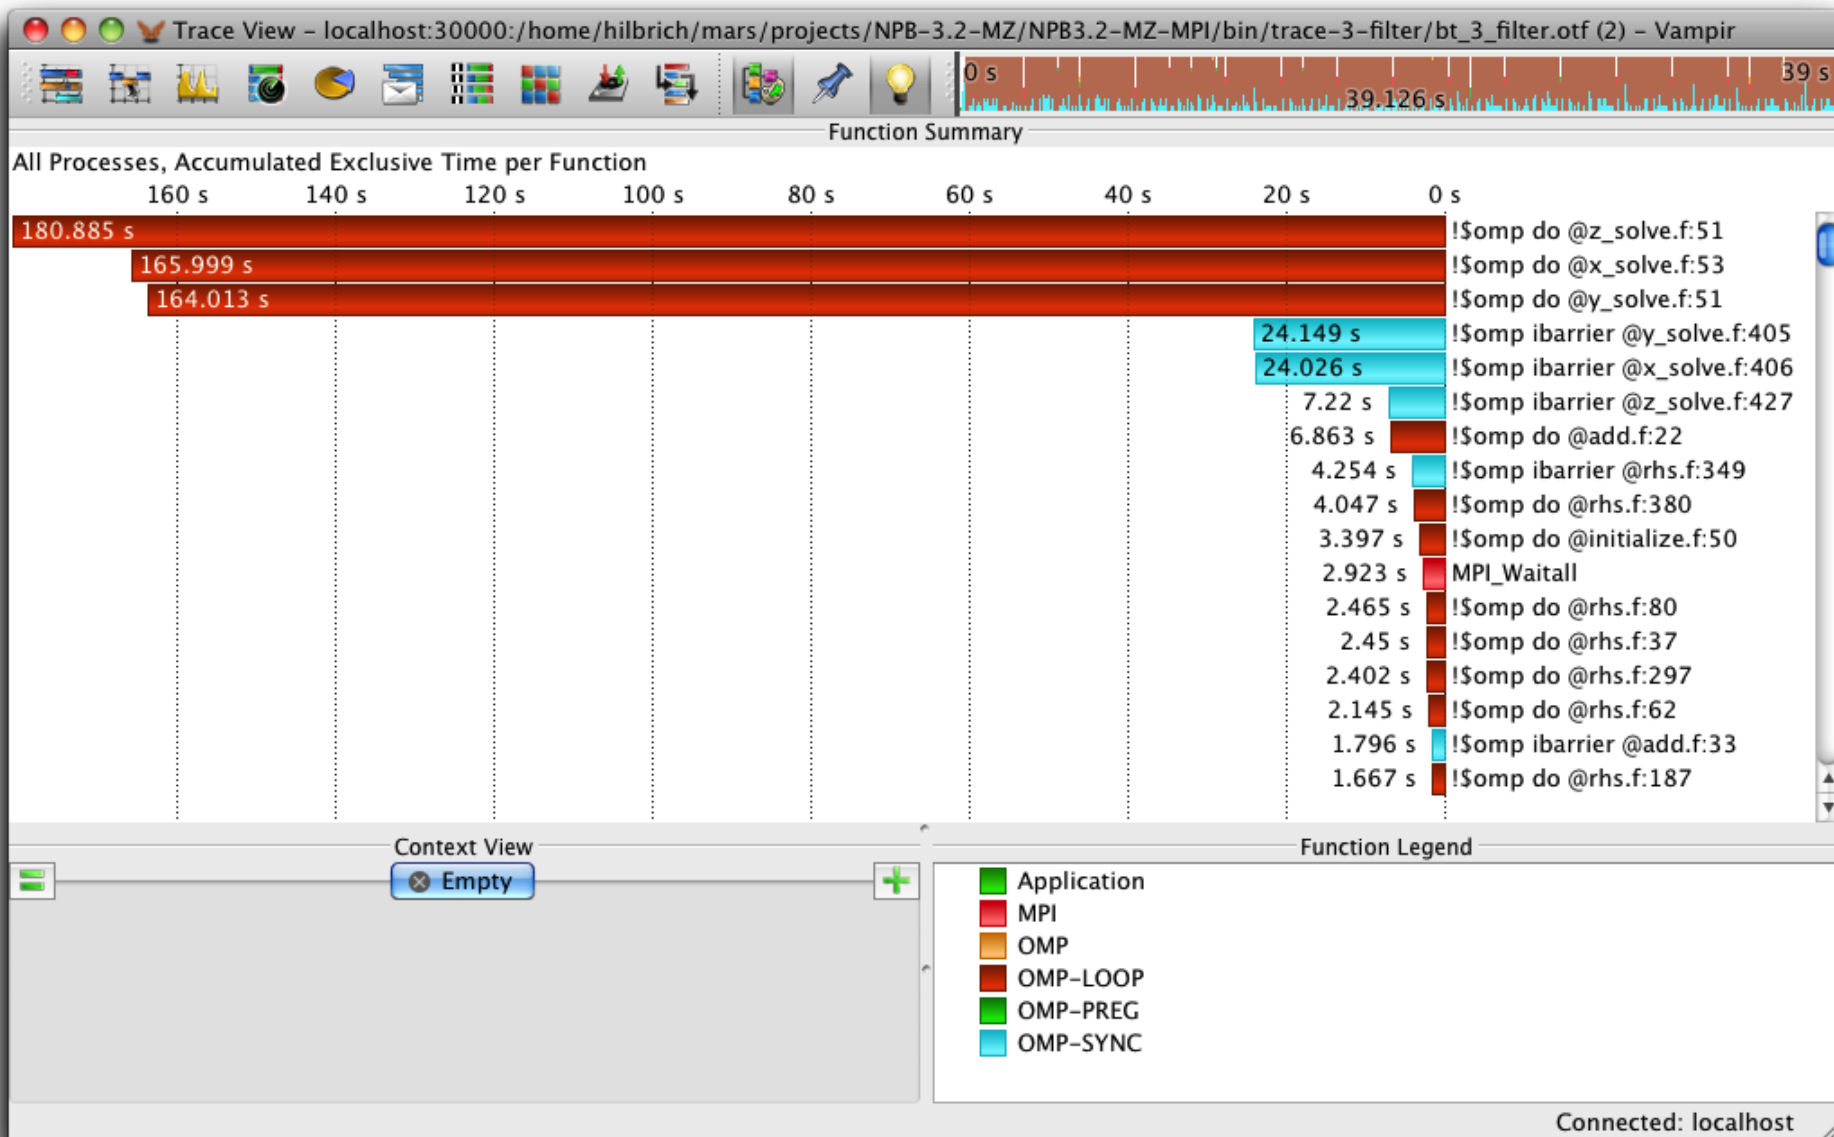
Hands-on: NPB – Filtered Trace



Hands-on: NPB – Filtered Trace



Hands-on: NPB – Filtered Trace



- PAPI counters can be included in traces
 - If VampirTrace was build with PAPI support
 - If PAPI is available on the platform
- **VT_METRICS** specifies a list of PAPI counters

```
VT_METRICS=PAPI_FP_OPS
```

- see also the PAPI commands [papi_avail](#) and [papi_command_line](#)

- Memory allocation counters can be recorded:
 - If VampirTrace build with memory allocation tracing support
 - If GNU glibc is used on the platform
- intercept glibc functions like “malloc” and “free”
- Environment variable **VT_MEMTRACE**

```
% export VT_MEMTRACE = yes
```

- I/O counters can be included in traces
 - If VampirTrace was build with I/O tracing support
- Standard I/O calls like “open” and “read” are recorded
- Environment variable **VT_IOTRACE**

```
% export VT_IOTRACE = yes
```

Hands-on: NPB – Filter & Counter Trace



Trace View - localhost:30000:/home/hilbrich//mars/projects/NPB-3.2-MZ/NPB3.2-MZ-MPI/bin/trace-4-papi/bt_4_papi.otf - Vampir

Timeline: 490.7 s, 490.8 s, 490.9 s, 491.0 s, 491.1 s

Function Summary: All Processes, Accumulated Exclusive Time per Funct...
 3 s, 2 s, 1 s, 0 s
 3.792 s OMP-SYNC
 1.417 s OMP-LOOP
 0.348 s OMP-PREG
 0.246 s OMP
 81.057 ms MPI
 <10 ms Appl...tion

Process 0

- 1 MAIN_
- 2 exch_qbc_DwArFil_HiGh_Pc
- 3 add_compute_rhs_ x_solve_y_solve_
- 4 parallel region
- 5 !\$omp parallel @rhs.f:28
- 6 !\$omp do @rhs.f:62
- 7 !\$omp ibarrier @rhs.f:72

Process 0, Values of Counter "PAPI_FP_OPS" over Time

Context View

Property	Value
Display	Counter Data
Type	Counter
Location	Process 0
Maximum	2.179217 G
Average	2.179217 G
Minimum	2.179217 G
Unit	#
Time	261.956073 s


Function Legend

- Application
- MPI
- OMP
- OMP-LOOP
- OMP-PREG
- OMP-SYNC


Connected: localhost

- control options by environment variables:
 - VT_PFORM_GDIR Directory for final trace files
 - VT_PFORM_LDIR Directory for intermediate files
 - VT_FILE_PREFIX Trace file name
 - VT_BUFFER_SIZE Internal trace buffer size
 - VT_MAX_FLUSHES Max number of buffer flushes
 - VT_MEMTRACE Enable memory allocation tracing
 - VT_MPICHECK Enable MPI checking
 - VT_IOTRACE Enable I/O tracing
 - VT_MPITRACE Enable MPI tracing
 - VT_FILTER_SPEC Name of filter definition file
 - VT_GROUPS_SPEC Name of grouping definition file
 - VT_METRICS PAPI counter selection

- Performance analysis very important in HPC
- Use performance analysis tools for profiling and tracing
- Do not spend effort in DIY solutions, e.g. like printf-debugging
- Use tracing tools with some precautions
 - overhead
 - data volume
- Let us know about problems and about feature wishes
- vampirsupport@zih.tu-dresden.de



Vampir and VampirTraces are available at <http://www.vampir.eu> and <http://www.tu-dresden.de/zih/vampirtrace/> , get support via vampirsupport@zih.tu-dresden.de



Staff at ZIH - TU Dresden:

Ronny Brendel, Holger Brunst, Jens Doleschal,
Ronald Geisler, Daniel Hackenberg, Michael Heyde,
Tobias Hilbrich, Rene Jäkel, Matthias Jurenz,
Michael Kluge, Andreas Knüpfer, Matthias Lieber,
Holger Mickler, Hartmut Mix, Matthias Müller,
Wolfgang E. Nagel, Reinhard Neumann, Michael Peter,
Heide Rohling, Johannes Spazier, Michael Wagner,
Matthias Weber, Bert Wesarg