




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

9th VI-HPS Tuning Workshop at UVSQ
April, 2012

Bert Wesarg and Jens Doleschal

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

- Copy the updated jobscript from the beckett gateway

```
beckett: % scp /share/vt.msub <npb on curie>/jobscripts
```

- Move into tutorial directory on curie again

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

- Select the VampirTrace compiler wrappers

```
% vim config/make.def
-> comment out line 32, resulting in:
    ...
    32: #MPIF77 = mpif77
    ...
-> remove the comment from line 40, resulting in:
    ...
    39: MPIF77 = vtf77 -vt:hyb -vt:f77 mpif77
    ...
```

- Set up modules

```
% module purge  
% module load vampirtrace/5.12.2
```

- Build benchmark

```
% make clean; make suite
```

- Go to bin directory

```
% cd bin.vampir
```

- Submit your application to the batch system

```
% ccc_msub ../jobscript/vt.msub
```

We will skip this step to don't stress you with to much measurements run!

- Investigate the output file

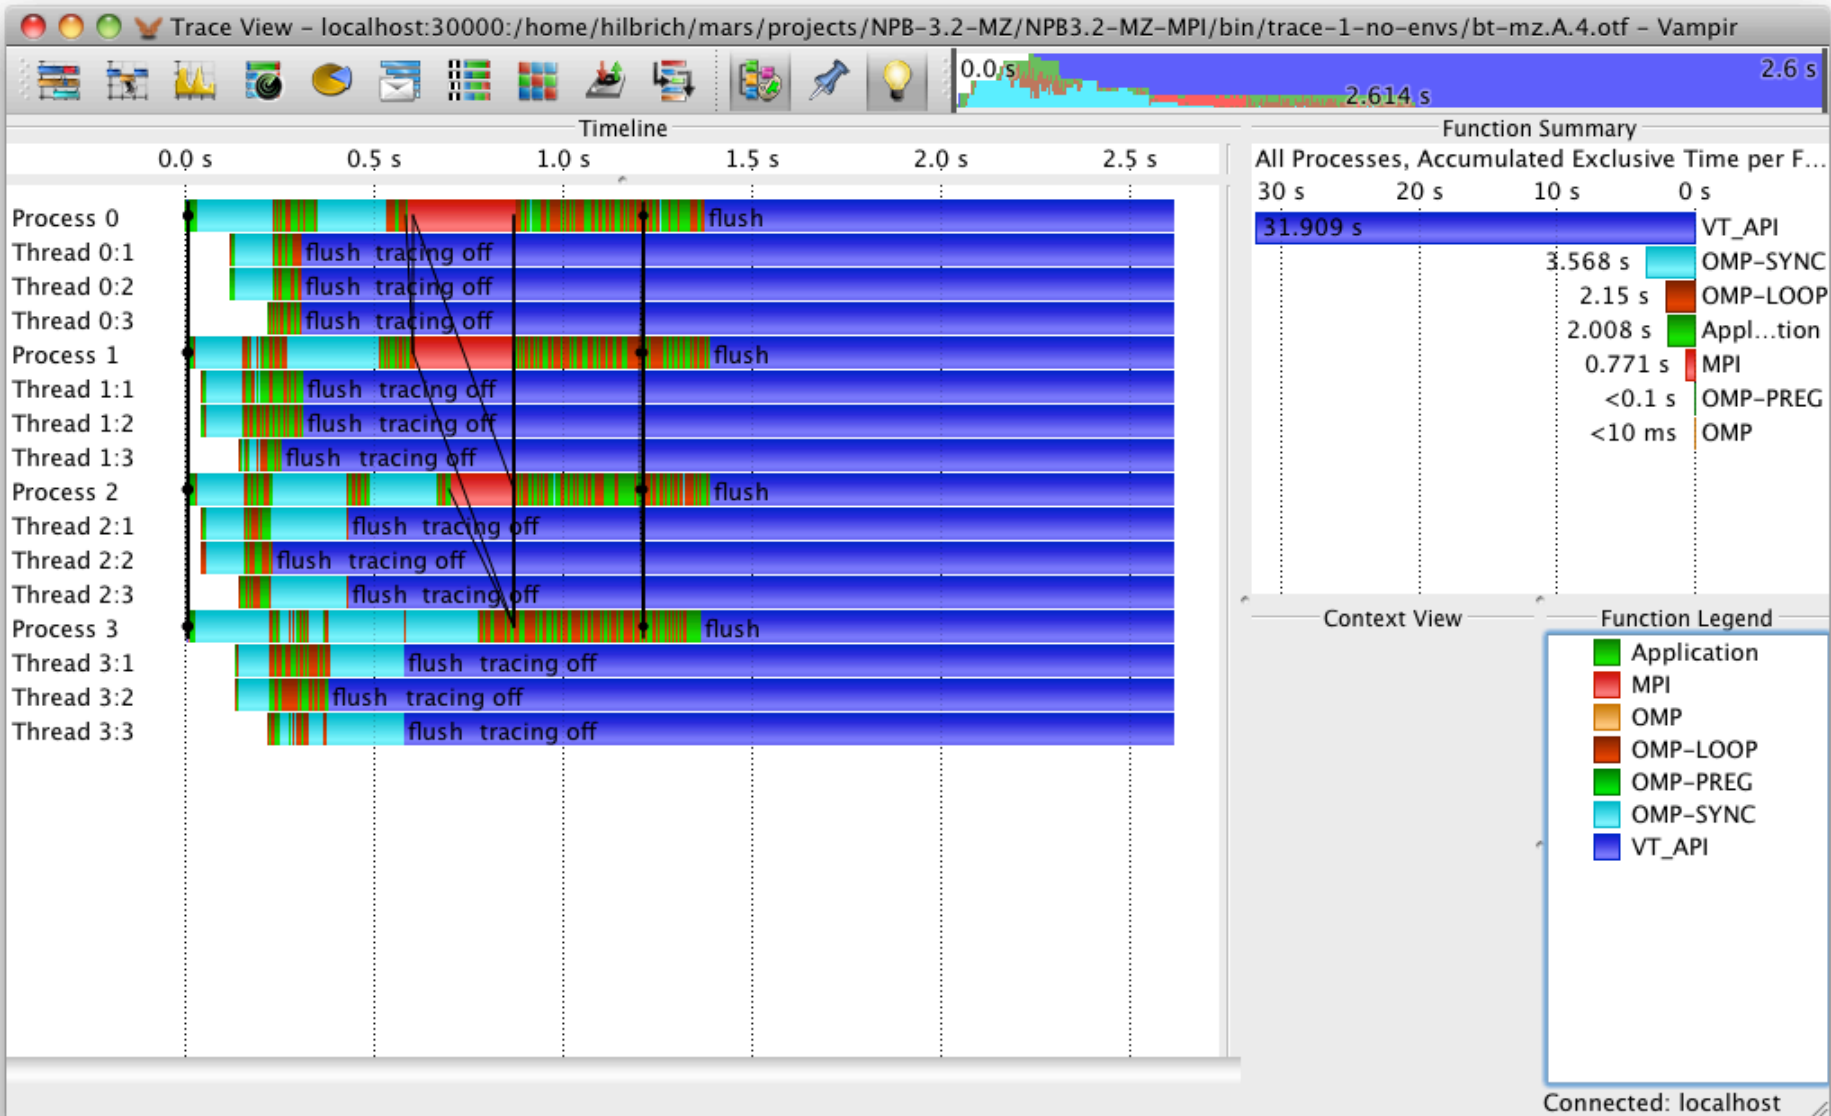
```
% cat npb_btmz_<ID>.oe
:
NAS Parallel Benchmarks (NPB3.3-MZ-MPI) - BT-MZ MPI+OpenMP Benchmark

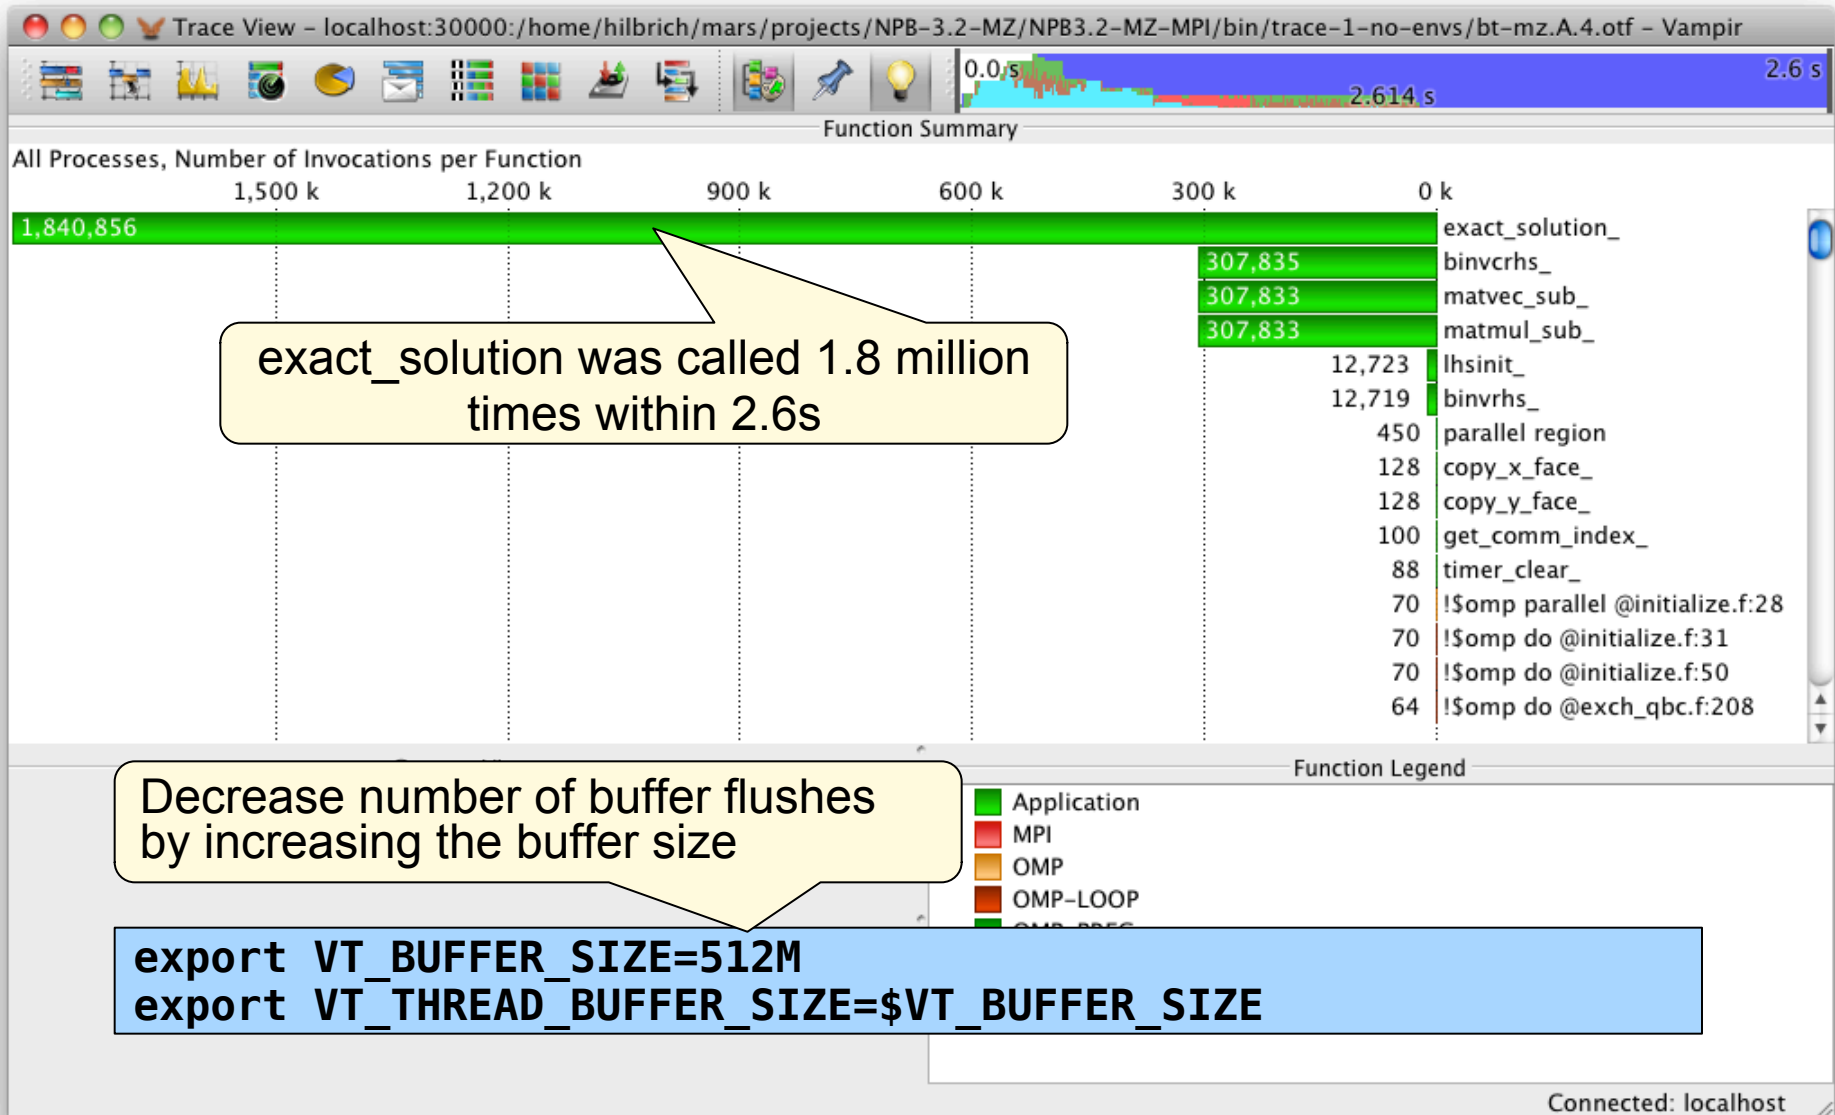
Number of zones:      8 x      8
Iterations: 200      dt:  0.000300
Number of active processes:      4

Use the default load factors with threads
Total number of threads:      16 (  4.0 threads/process)
Calculated speedup =      15.64
Time step      1
[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 !

Hands-on: NPB – Initial Trace

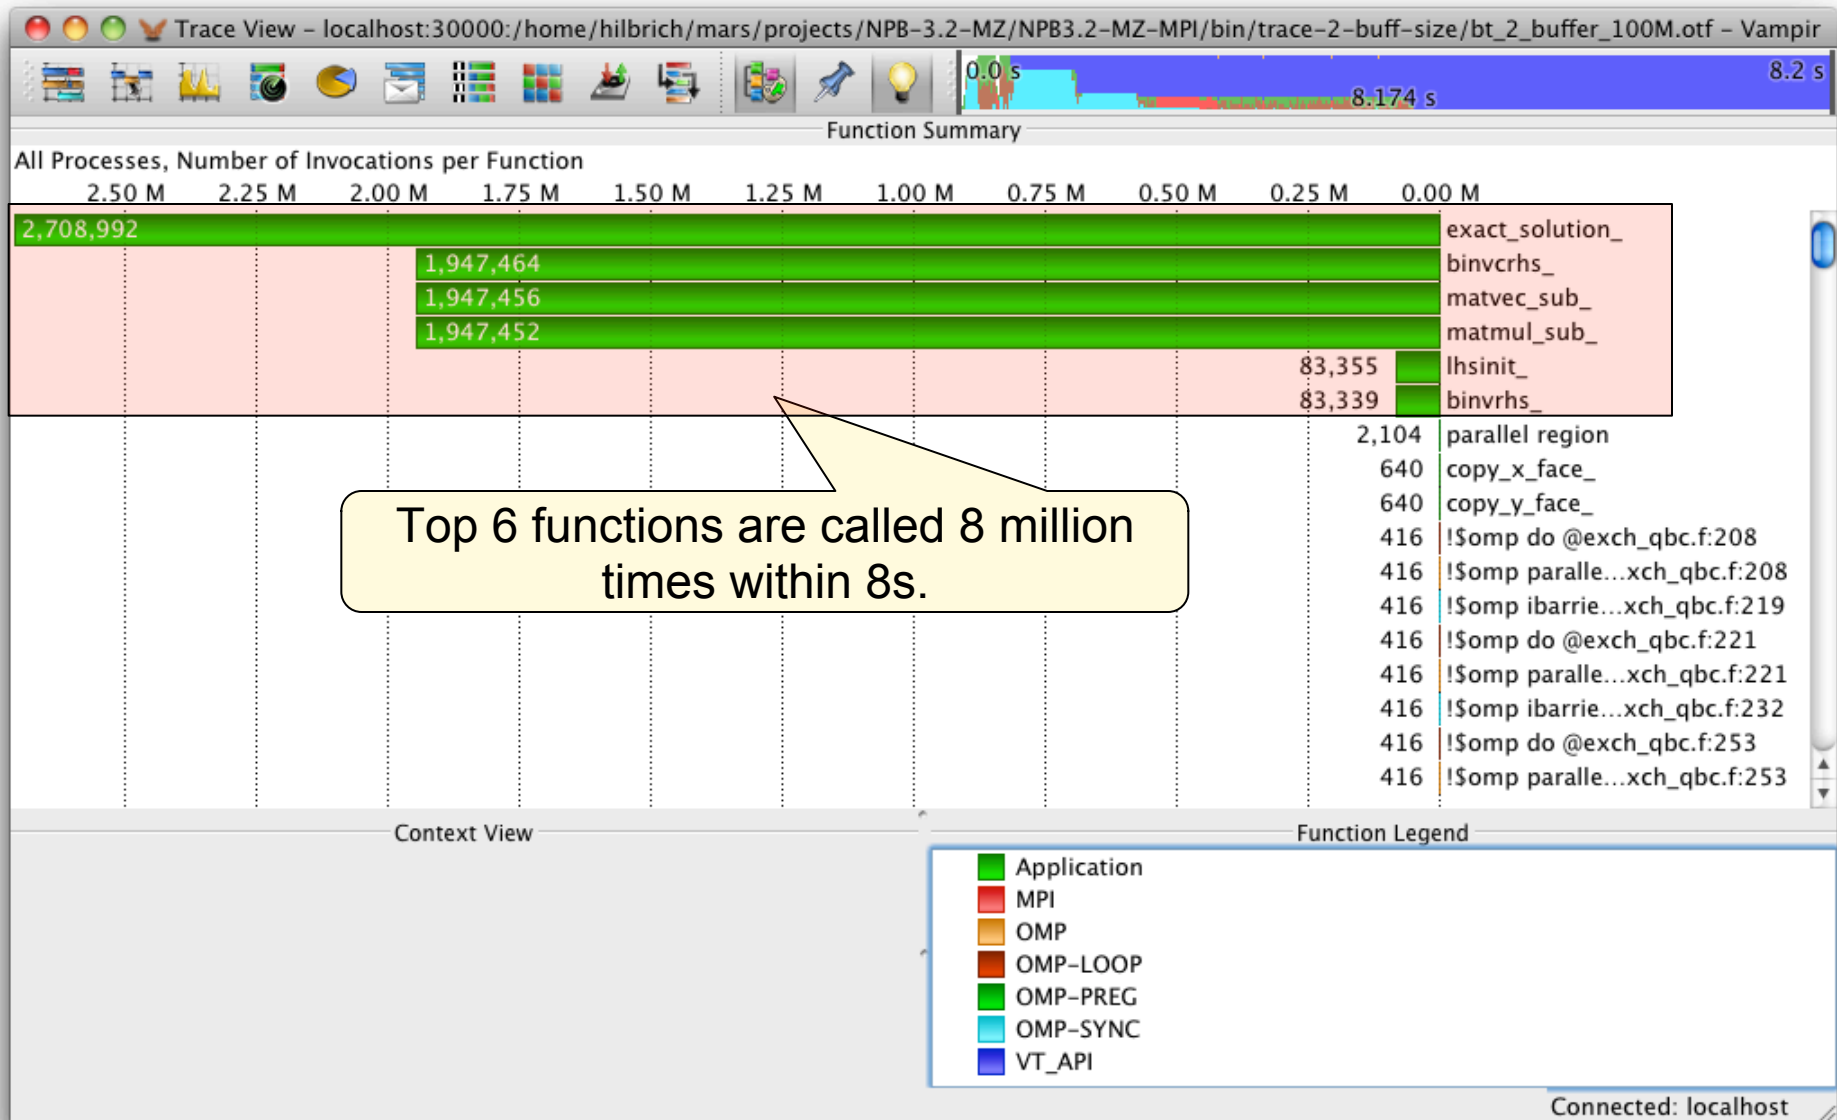




Hands-on: NPB – Second Trace



Hands-on: NPB – Second Trace



Example, don't type it into the shell

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


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

- Filter definition file contains a list of filters

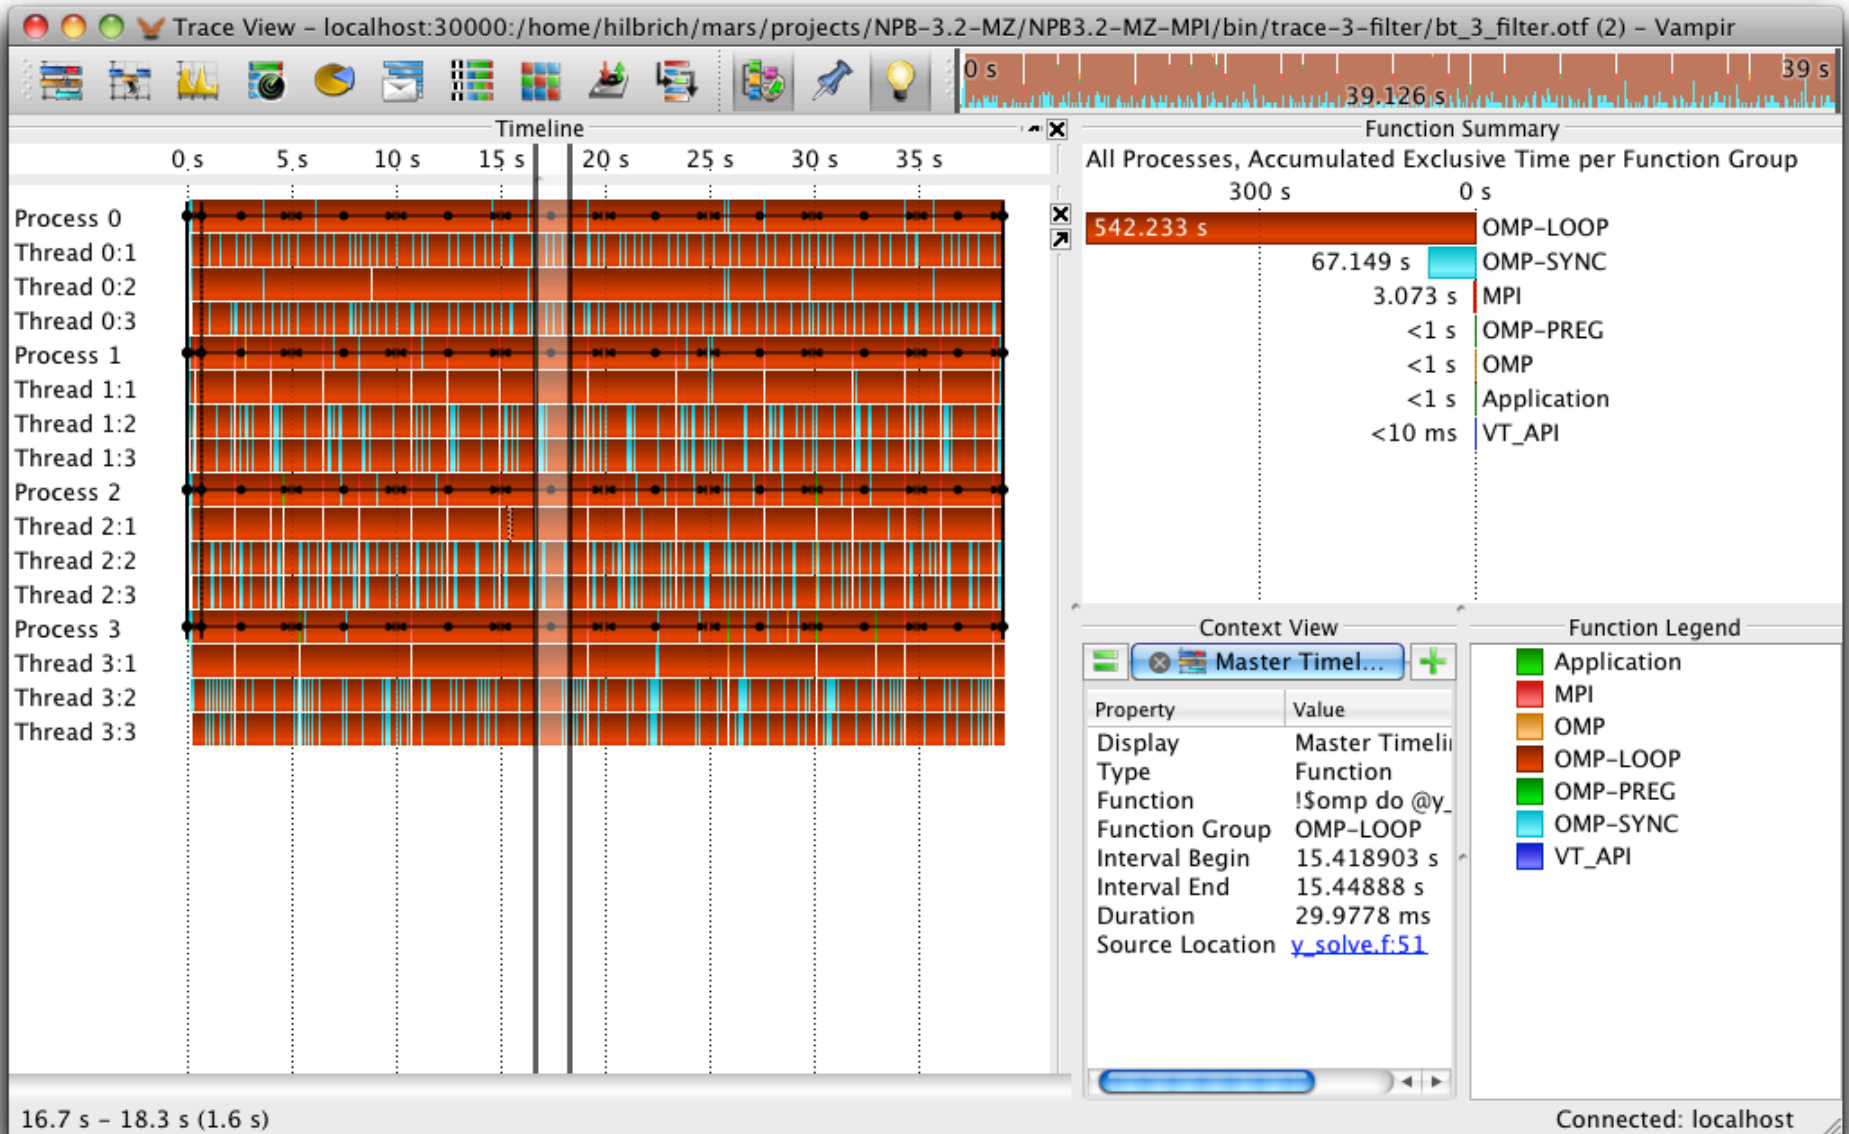
```
my_*;test_* -- 1000  
debug_* -- 0  
calculate -- -1  
* -- 1000000
```

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

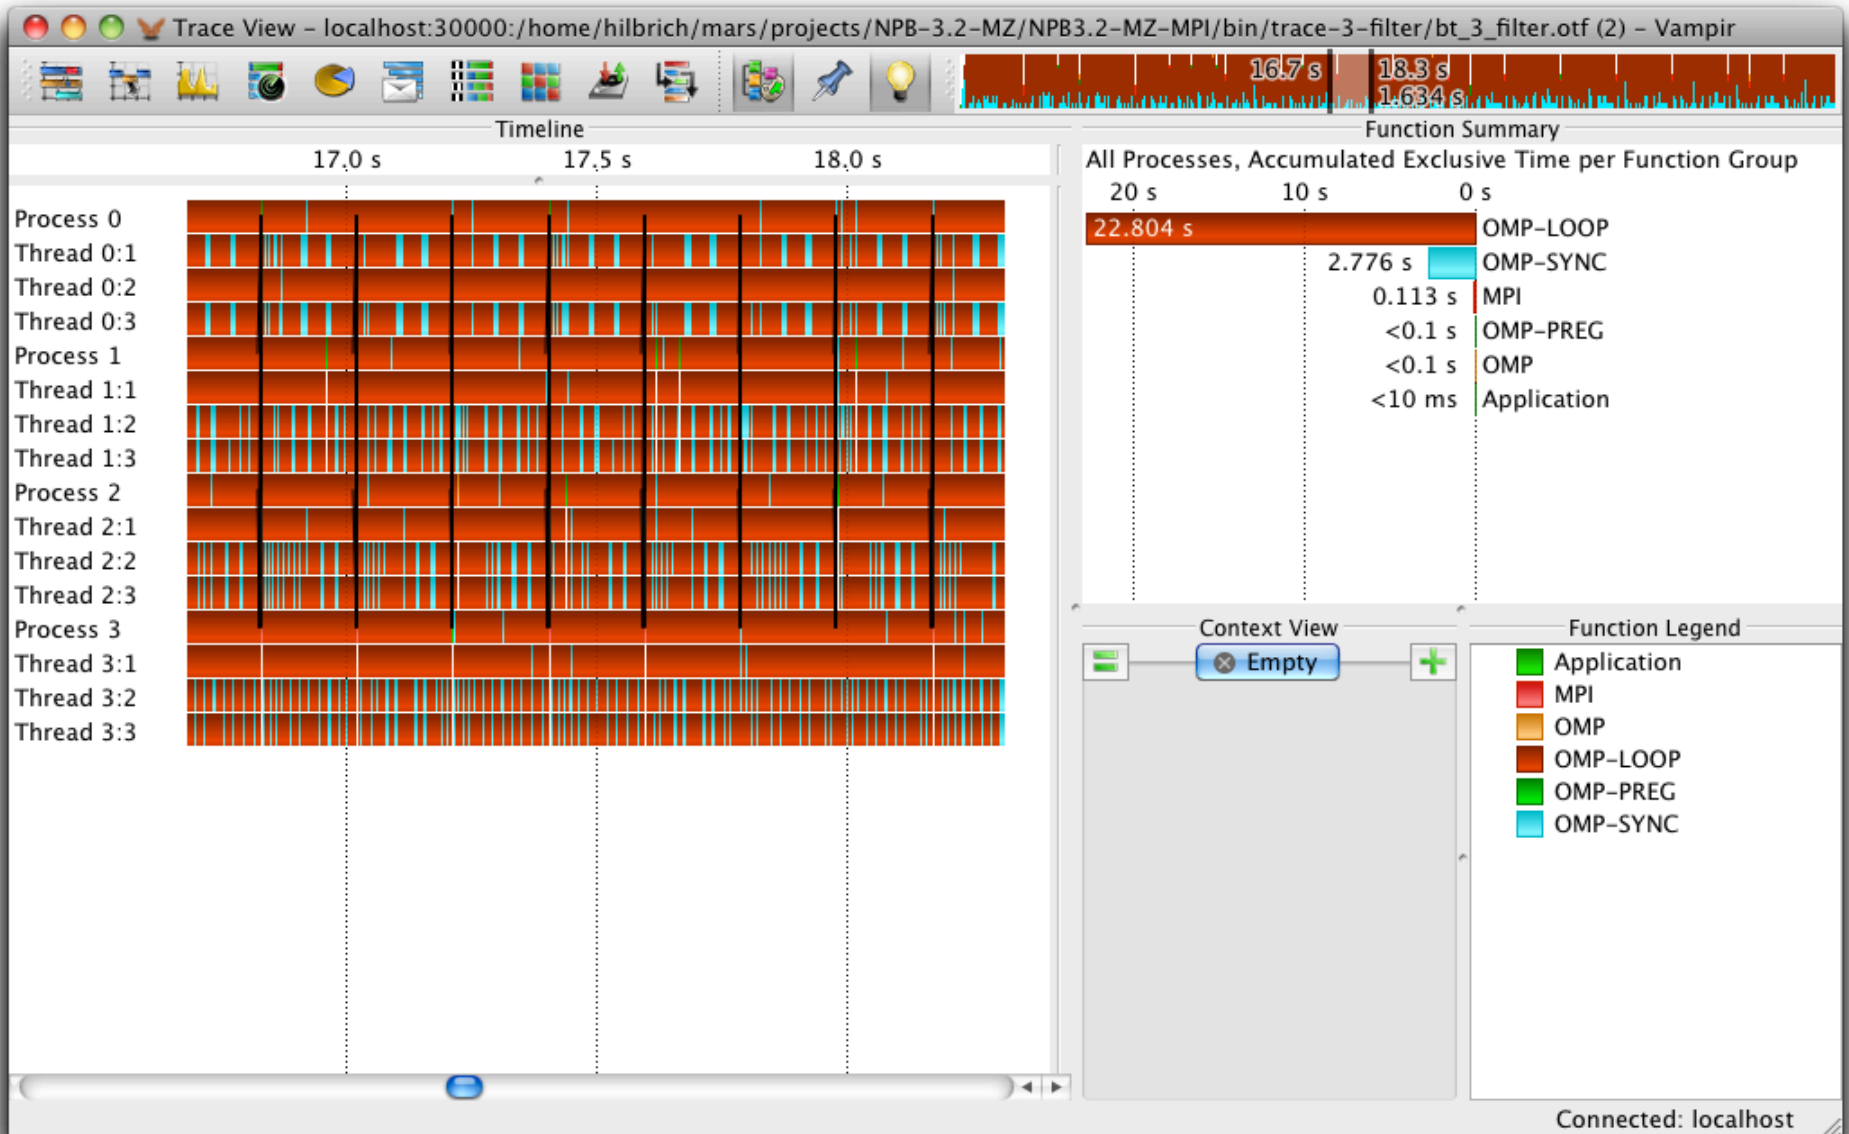
Hands-on: NPB – Filtered Trace



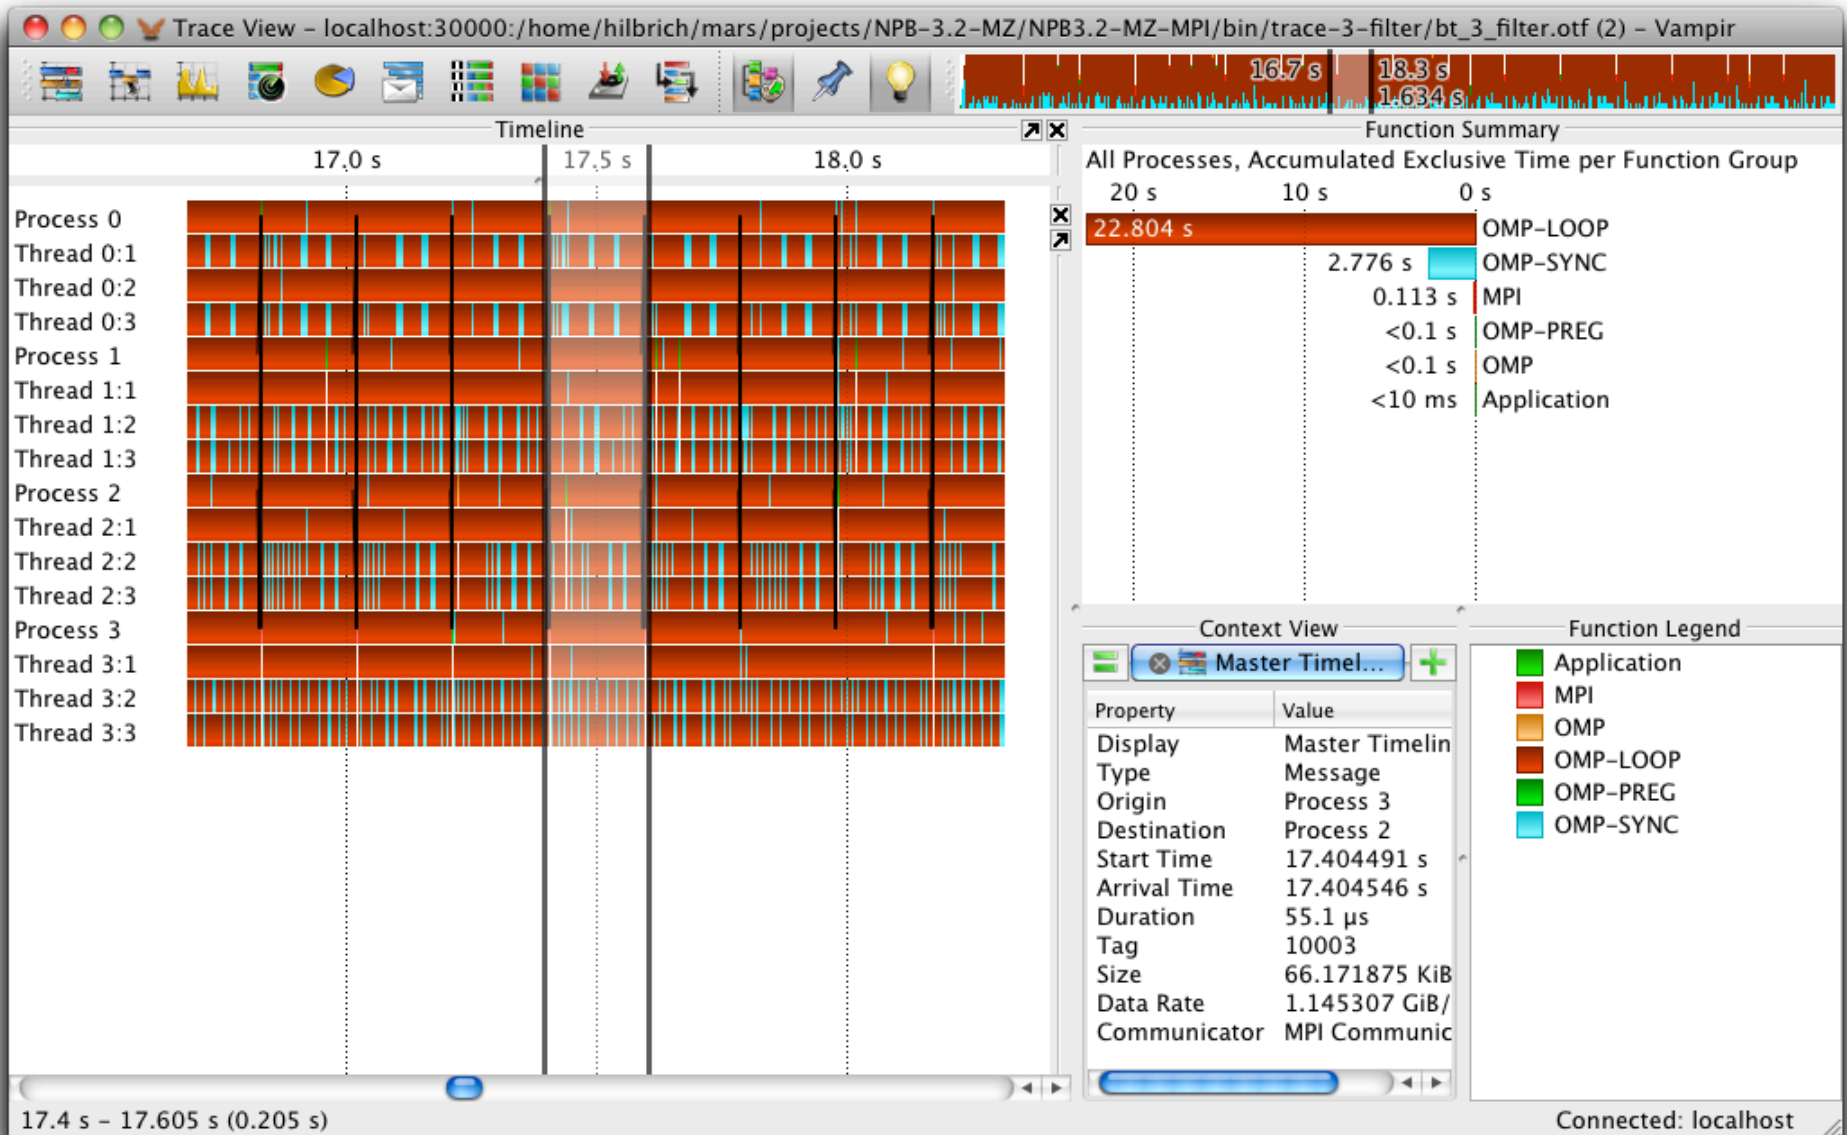
Hands-on: NPB – Filtered Trace



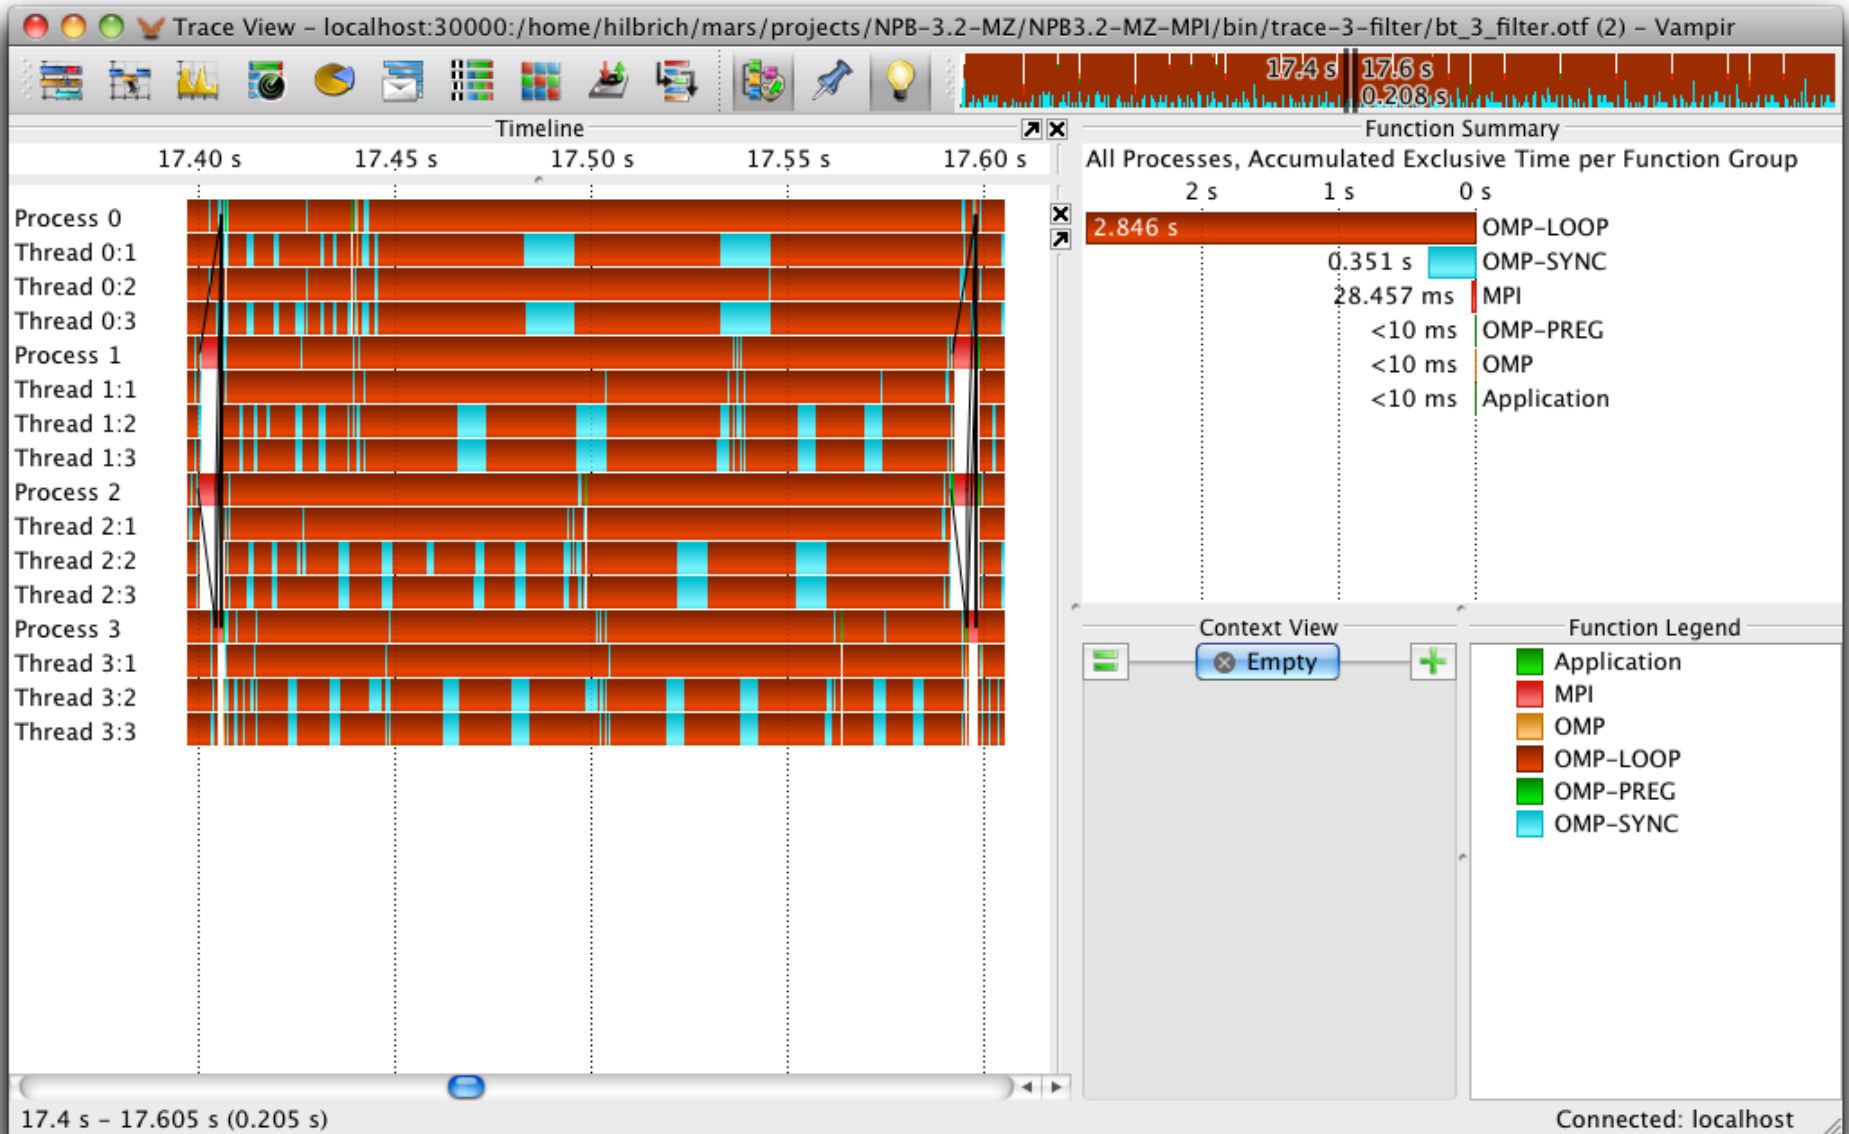
Hands-on: NPB – Filtered Trace



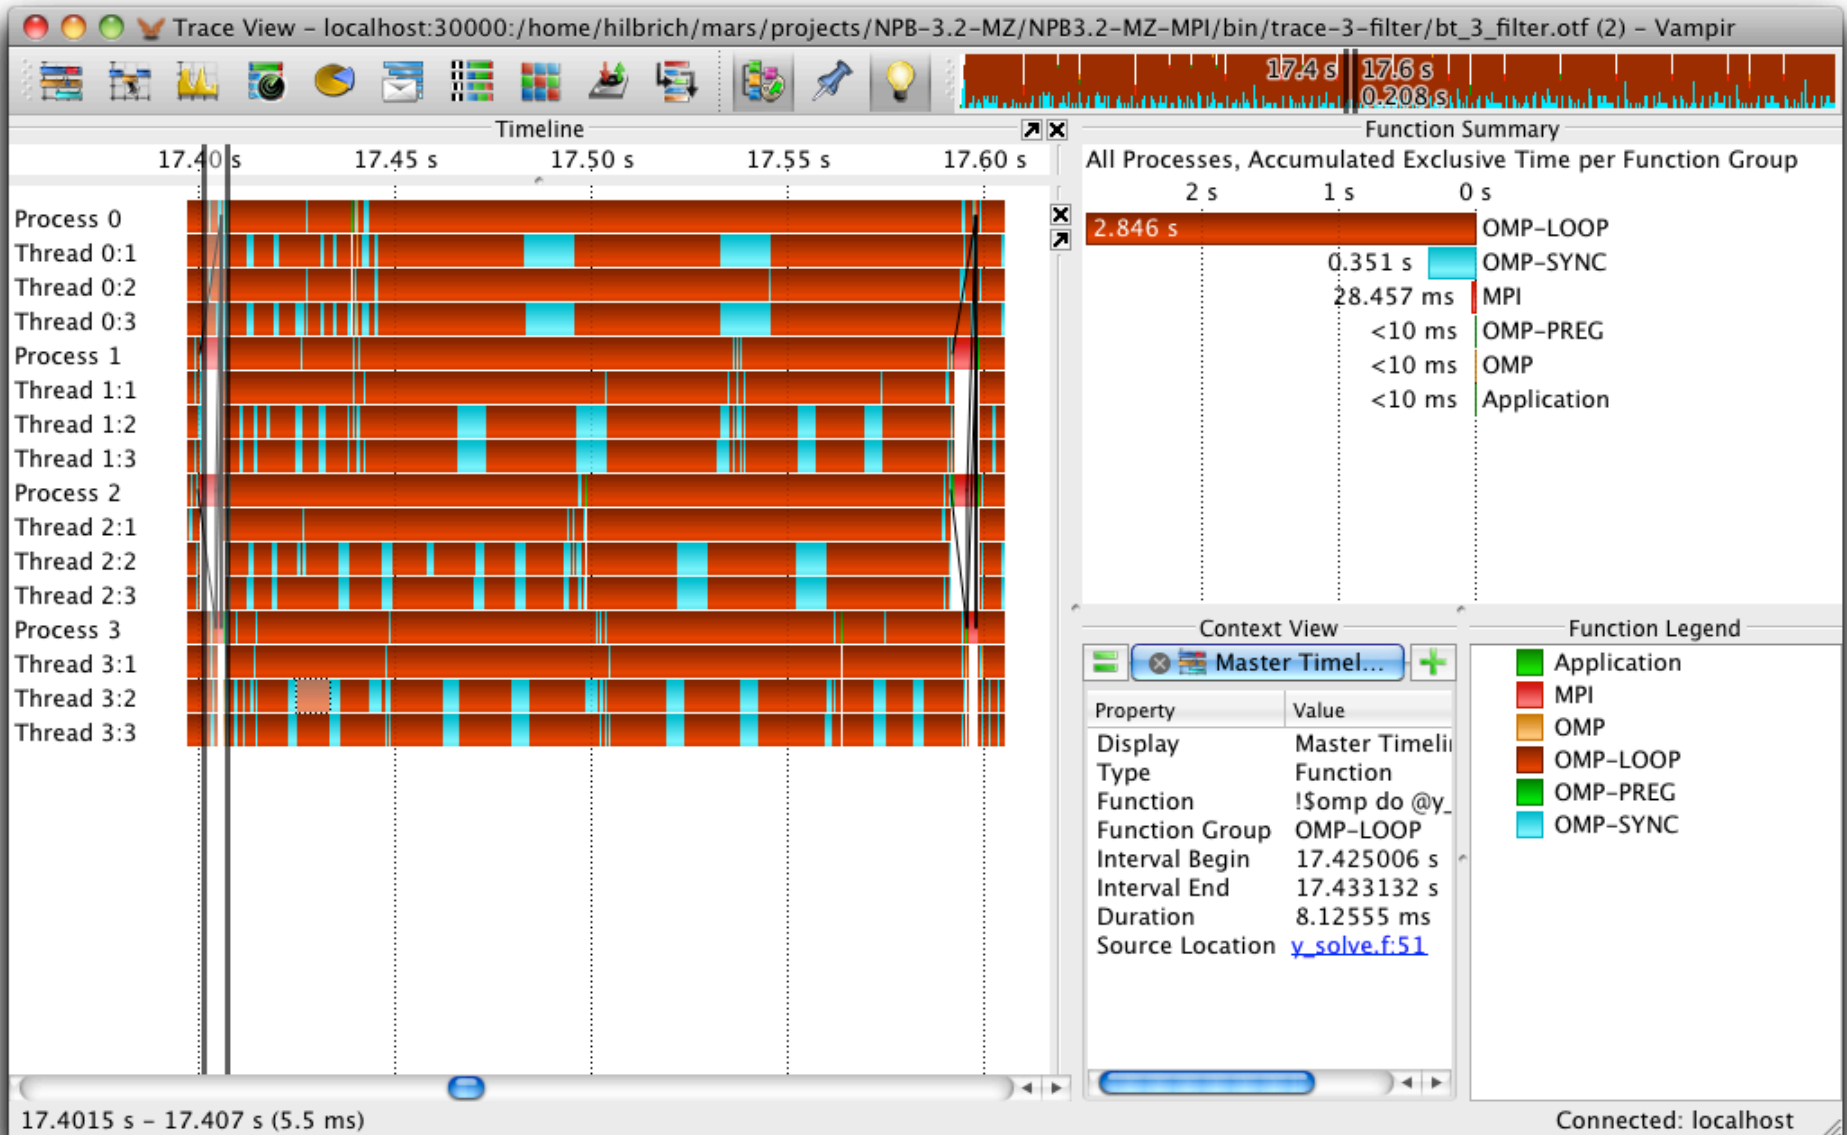
Hands-on: NPB – Filtered Trace




Hands-on: NPB – Filtered Trace



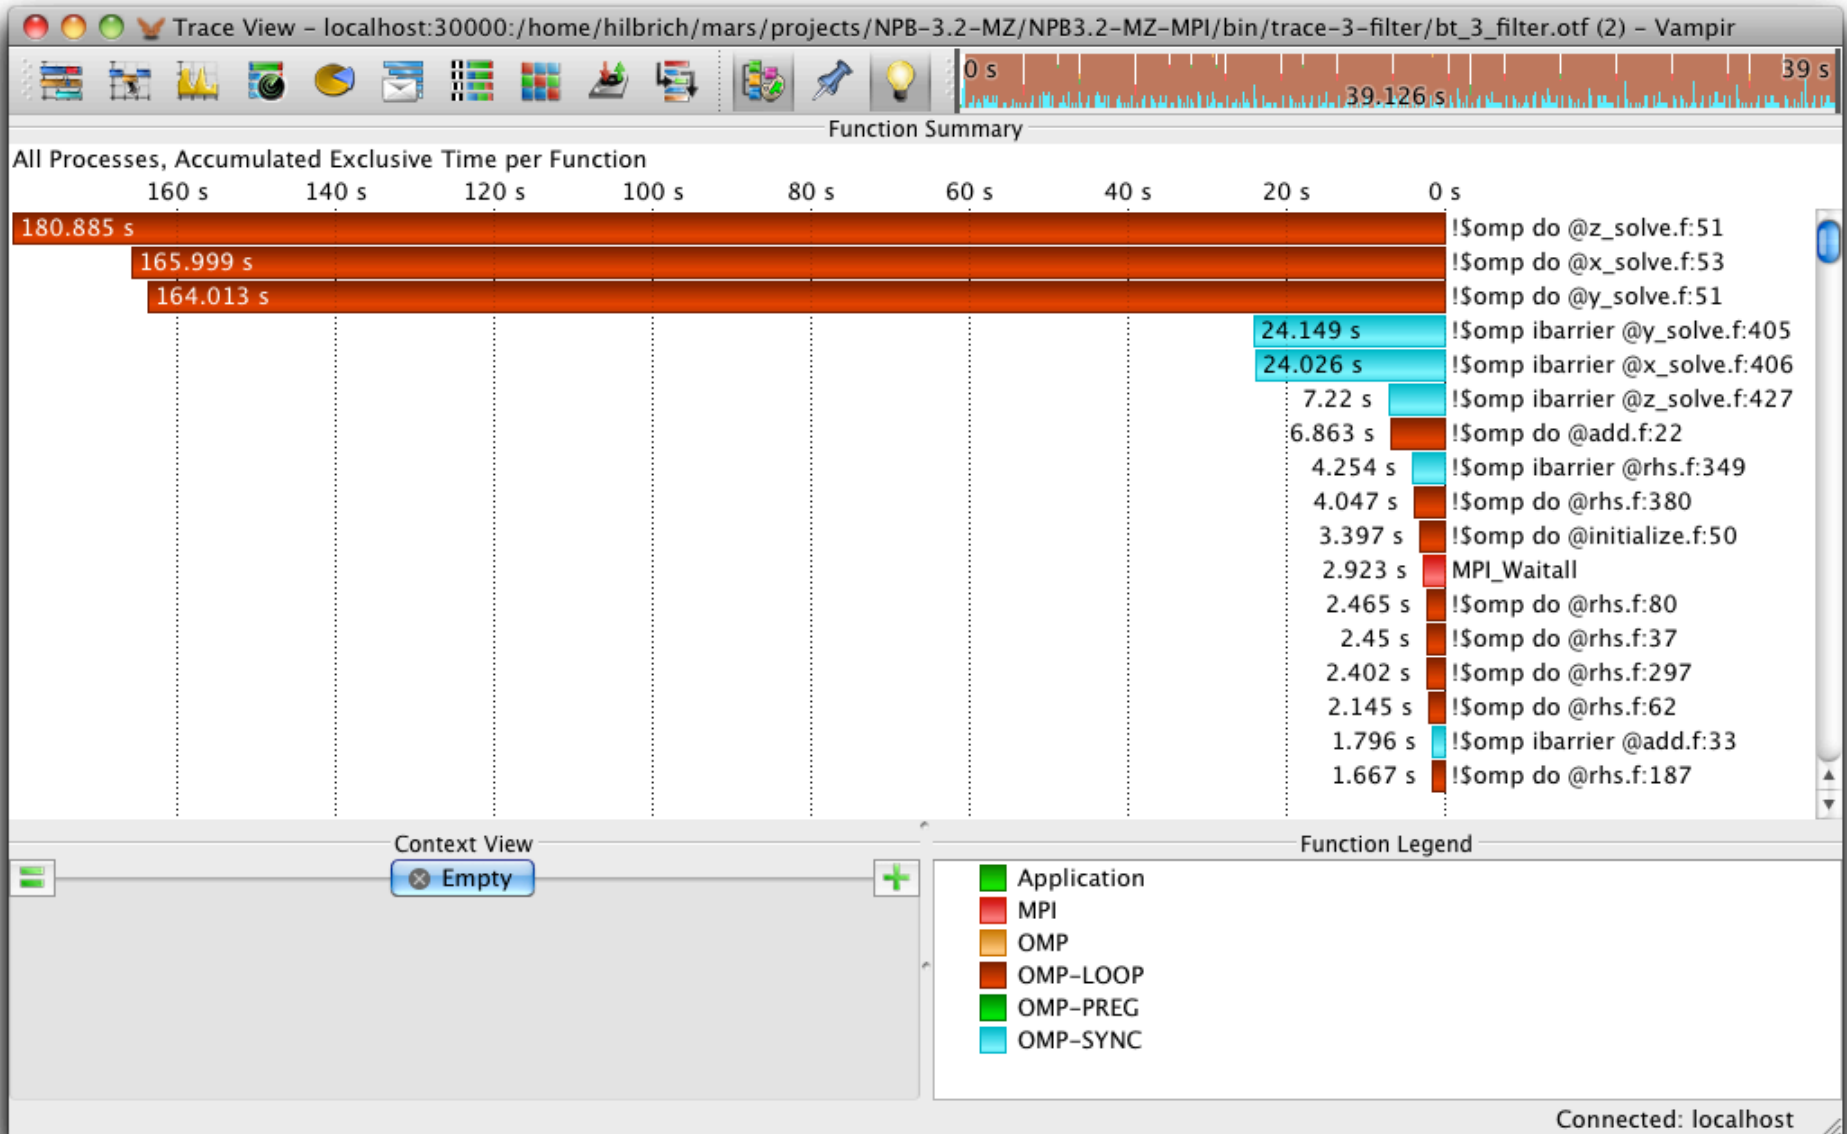
Hands-on: NPB – Filtered Trace



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_TOT_INS:PAPI_FP_OPS:PAPI_L2_TCM:PAPI_L3_TCM
```

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

- Inspect filter specification

```
% cat ../config/vt_filter.txt
exact_solution_ -- 0
binvrhs_ -- 0
matvec_sub_ -- 0
matmul_sub_ -- 0
lhsinit_ -- 0
binvrhs_ -- 0
```

- Activate filtering, set buffer sizes, and PAPI counter

```
VT_BUFFER_SIZE=512M
VT_THREAD_BUFFER_SIZE=$VT_BUFFER_SIZE
VT_FILTER_SPEC=../config/vt_filter.txt
VT_METRICS=PAPI_TOT_INS:PAPI_FP_OPS:PAPI_L2_TCM:PAPI_L3_TCM
export VT_BUFFER_SIZE VT_THREAD_BUFFER_SIZE VT_FILTER_SPEC \
VT_METRICS
```

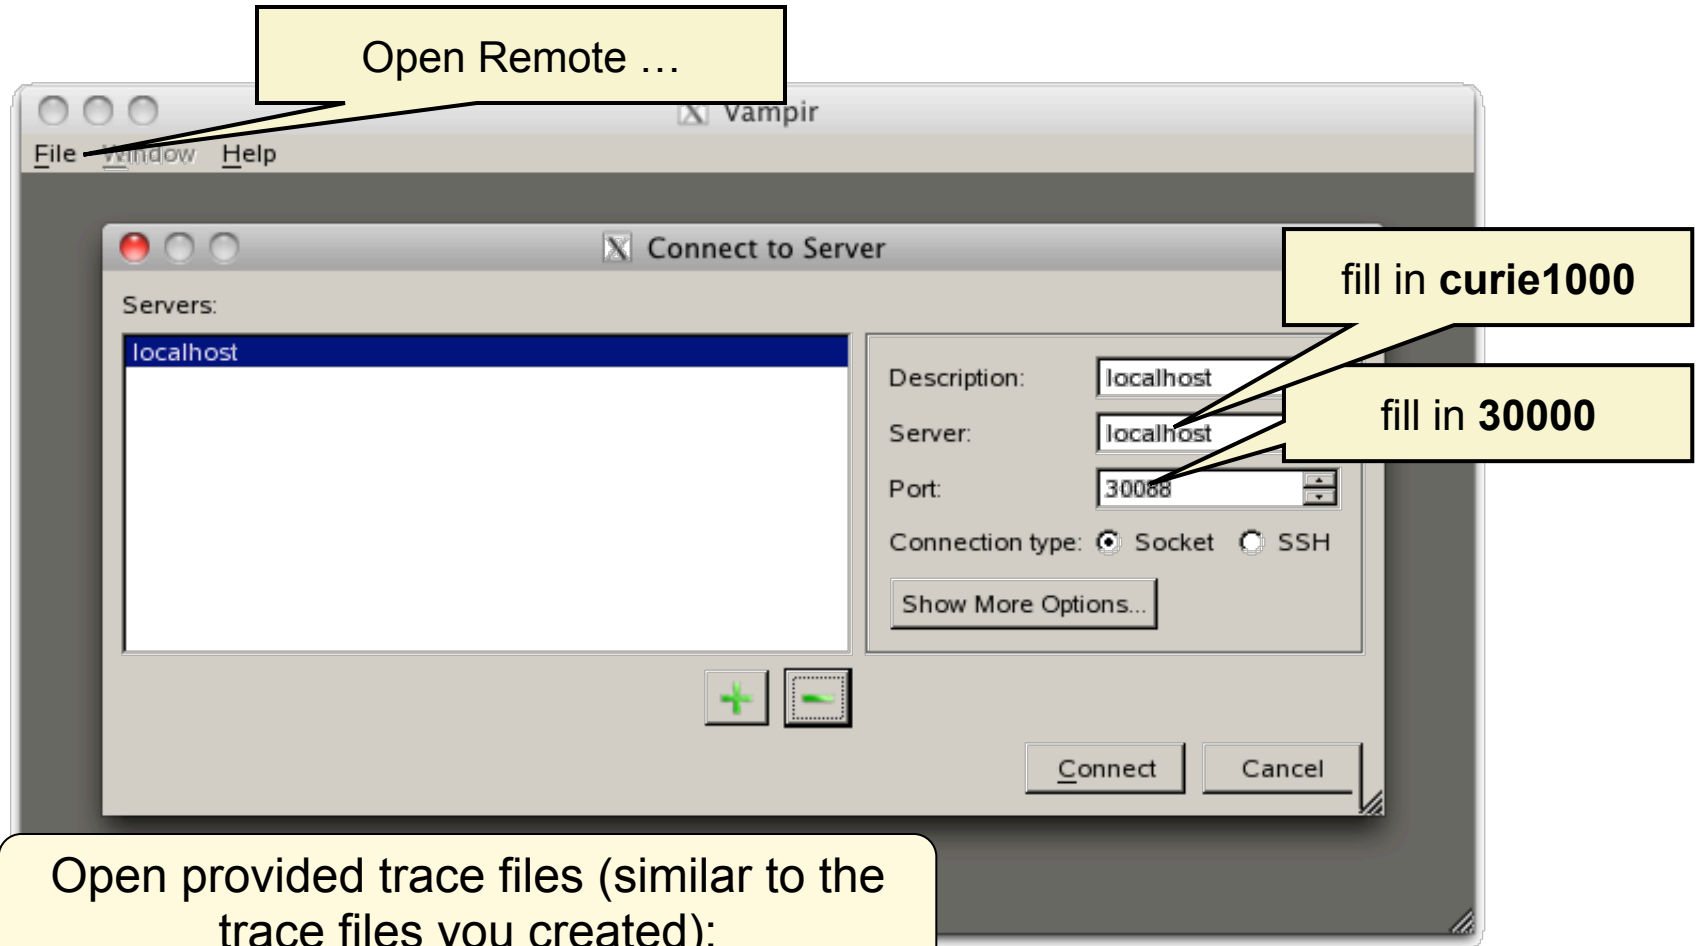
Add this to your jobscript

- Resulting trace files

```
% ls
bt-mz_B.4
bt-mz_B.4.0.def.z
bt-mz_B.4.1.events.z
bt-mz_B.4.10001.events.z
bt-mz_B.4.10002.events.z
...
bt-mz_B.4.40004.events.z
bt-mz_B.4.50001.events.z
bt-mz_B.4.50002.events.z
bt-mz_B.4.50003.events.z
bt-mz_B.4.50004.events.z
bt-mz_B.4.otf
```

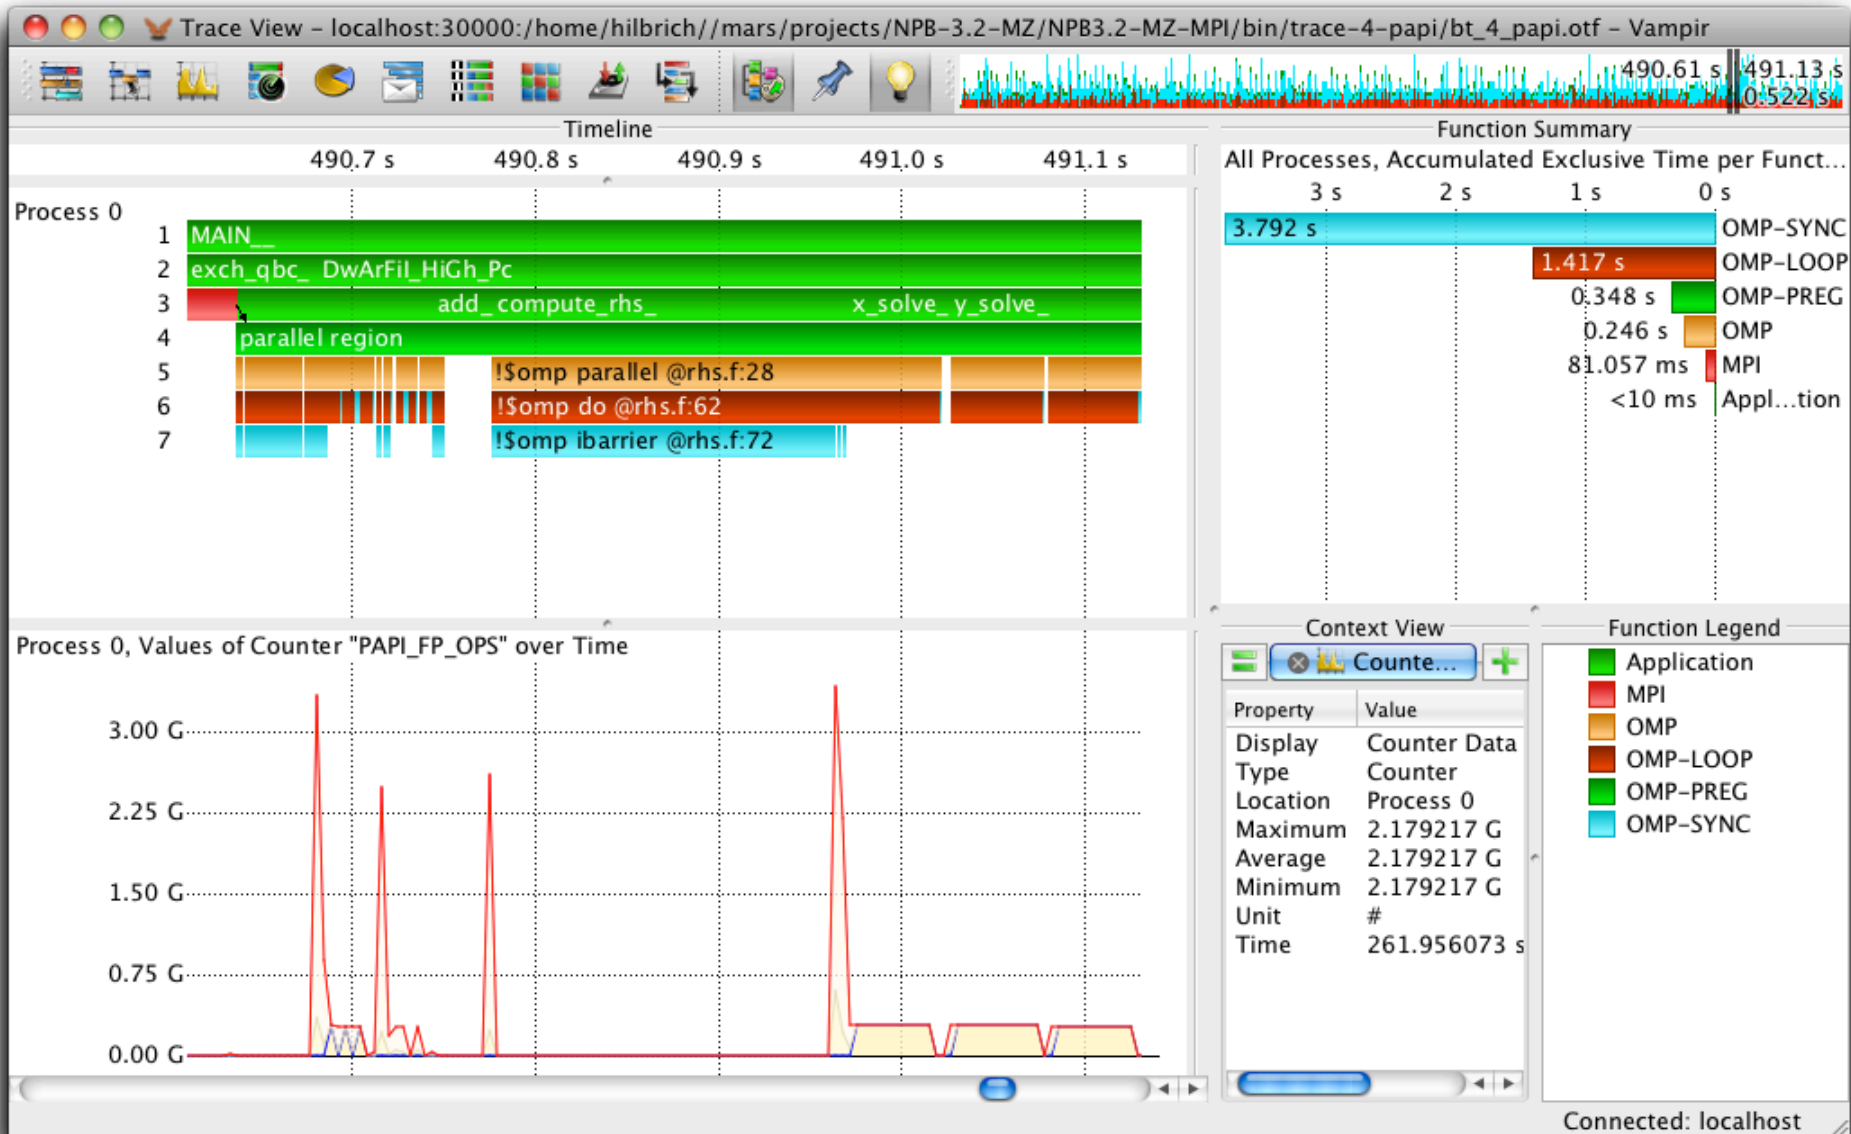
- Visualization with Vampir7

```
% module load vampir
% vampir
```




`/ccc/scratch/cont005/tgcc0007/wesargb/npb_btmz_vt-4.otf`

Hands-on: NPB – Filter & Counter Trace



- 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



Acknowledgement:

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