

OpenACC Tools

Sandra Wienke, M.Sc.

wienke@rz.rwth-aachen.de

Center for Computing and Communication
RWTH Aachen University

PPCES, March 2013

Agenda

■ IDE

→ Eclipse

■ Debugging

→ Tips & Tricks

■ Profiling

→ NVIDIA Visual Profiler

IDE - Eclipse

■ Eclipse (+ Plugins): IDE for GPU programming

- Parallel Nsight: OpenACC profiling, CUDA syntax highlighting/debugging/profiling
- Parallel Tools Platform (PTP): OpenACC programming

■ Using Eclipse

- Load your needed modules first, e.g. `module switch intel pgi` (for OpenACC)
- Default Eclipse version in RWTH Cluster environment (not able to do GPU profiling or syntax highlighting): `module load MISC eclipse; eclipse &`

recommended

- Parallel Nsight in RWTH Cluster environment: `module load cuda; nsight`
(or download from <http://www.nvidia.com/object/nsight.html>)

- PTP: Download from <http://www.eclipse.org/ptp/> ,
extract and start `eclipse &`

For the OpenACC labs, we put Eclipse+PTP into your Home directory.

IDE - Eclipse

■ Create Makefile projects with Eclipse

1. Start `eclipse` and chose workspace
2. File → New → Makefile Project with Existing code
3. Chose file directory of your project
4. [If CUDA is used: Toolchain: CUDA Toolkit 5.0]
5. Create Makefile targets (Makefile Tab in right pane), e.g. release, run
6. Double click on Makefile target for execution

■ PTP & OpenACC

→ Doc: <http://help.eclipse.org/juno/topic/org.eclipse.ptp.doc.user/html/toc.html>

→ Recognizing PGI compiler errors:

http://help.eclipse.org/juno/topic/org.eclipse.ptp.doc.user/html/errorParsers.html?cp=51_8

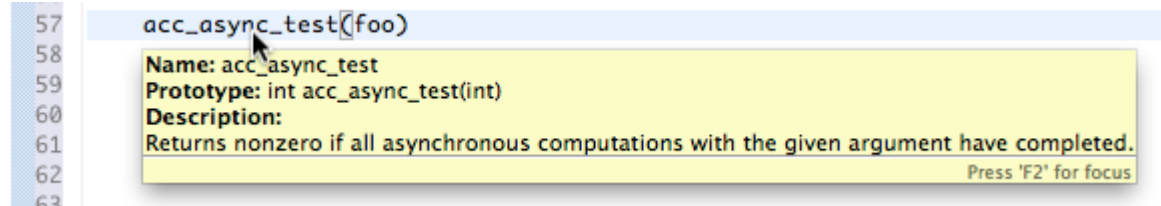
→ Running PTP OpenACC tools:

http://help.eclipse.org/juno/topic/org.eclipse.ptp.pldt.doc.user/html/runOpenACC.html?cp=51_14_8

■ PTP & OpenACC (taken from http://help.eclipse.org/juno/topic/org.eclipse.ptp.pldt.doc.user/html/runOpenACC.html?cp=51_14_8)

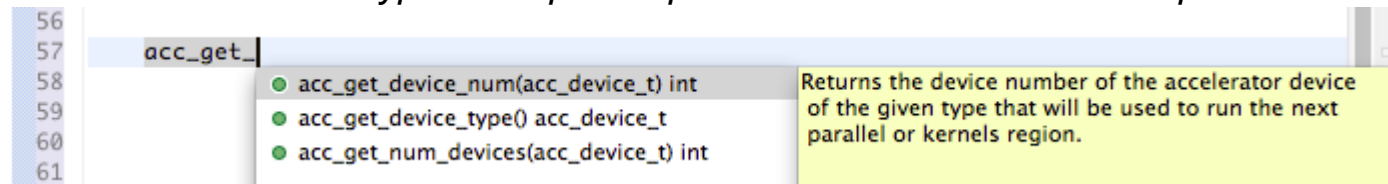
→ Hover help: *In the editor, hover an OpenACC API.*

```
57 acc_async_test(foo)
58
59
60
61
62
63
```



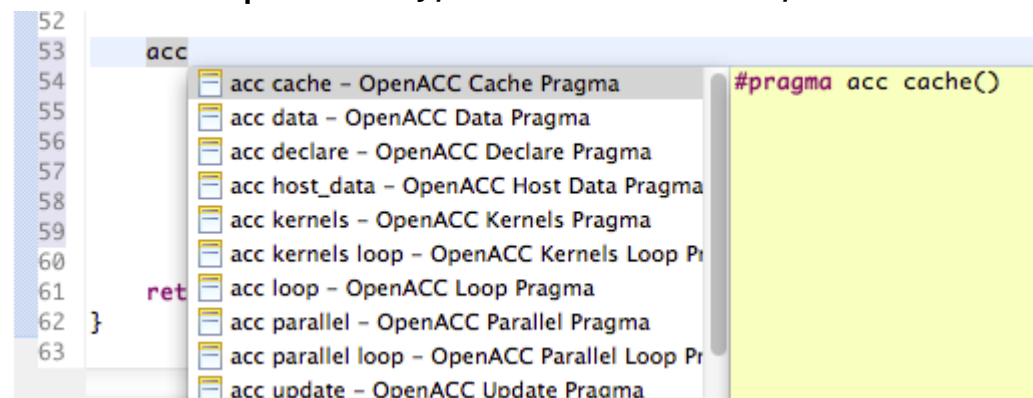
→ Content Assist: *Type incomplete OpenACC function and hit Ctrl-Space.*

```
56
57 acc_get_
58
59
60
61
```



→ Code Templates: *Type acc and hit Ctrl-Space twice.*

```
52
53 acc
54
55
56
57
58
59
60
61 ret
62 }
63
```



Agenda

■ IDE

→ Eclipse

■ Debugging

→ Tips & Tricks

■ Profiling

→ NVIDIA Visual Profiler

- **Debugging host code as usual (e.g. using TotalView)**
- **Debugging GPU kernels requires special tools**
 - CUDA debuggers available
 - OpenACC debuggers (for PGI compiler) not available ☹️

RWTH Cluster
environment:
module load
totalview

■ Tips & Tricks

- Set `ACC_NOTIFY` to nonnegative integer to get a command line message
 - =1: on launch of kernels
 - =2: on moving data between device & host
 - =3: 1 & 2
- Set `PGI_ACC_DEBUG=1` to get detailed debug information

Agenda

■ IDE

→ Eclipse

■ Debugging

→ Tips & Tricks

■ Profiling

→ NVIDIA Visual Profiler

- **Profiling = Analyze behavior of application during runtime**

- e.g. runtime of functions, memory throughput

- **Command line: `export PGI_ACC_TIME=1`**

- **NVIDIA Visual Profiler for CUDA, OpenACC & OpenCL codes**

- Profiles only GPU data movement & computation (not host code)

- Use stand-alone version (`module load cuda; nvvp &`)

- or Eclipse plugin (see IDE Parallel Nsight)

- Start Nsight

- Use the release target of Makefile to create executable

- Press profile button (watch)

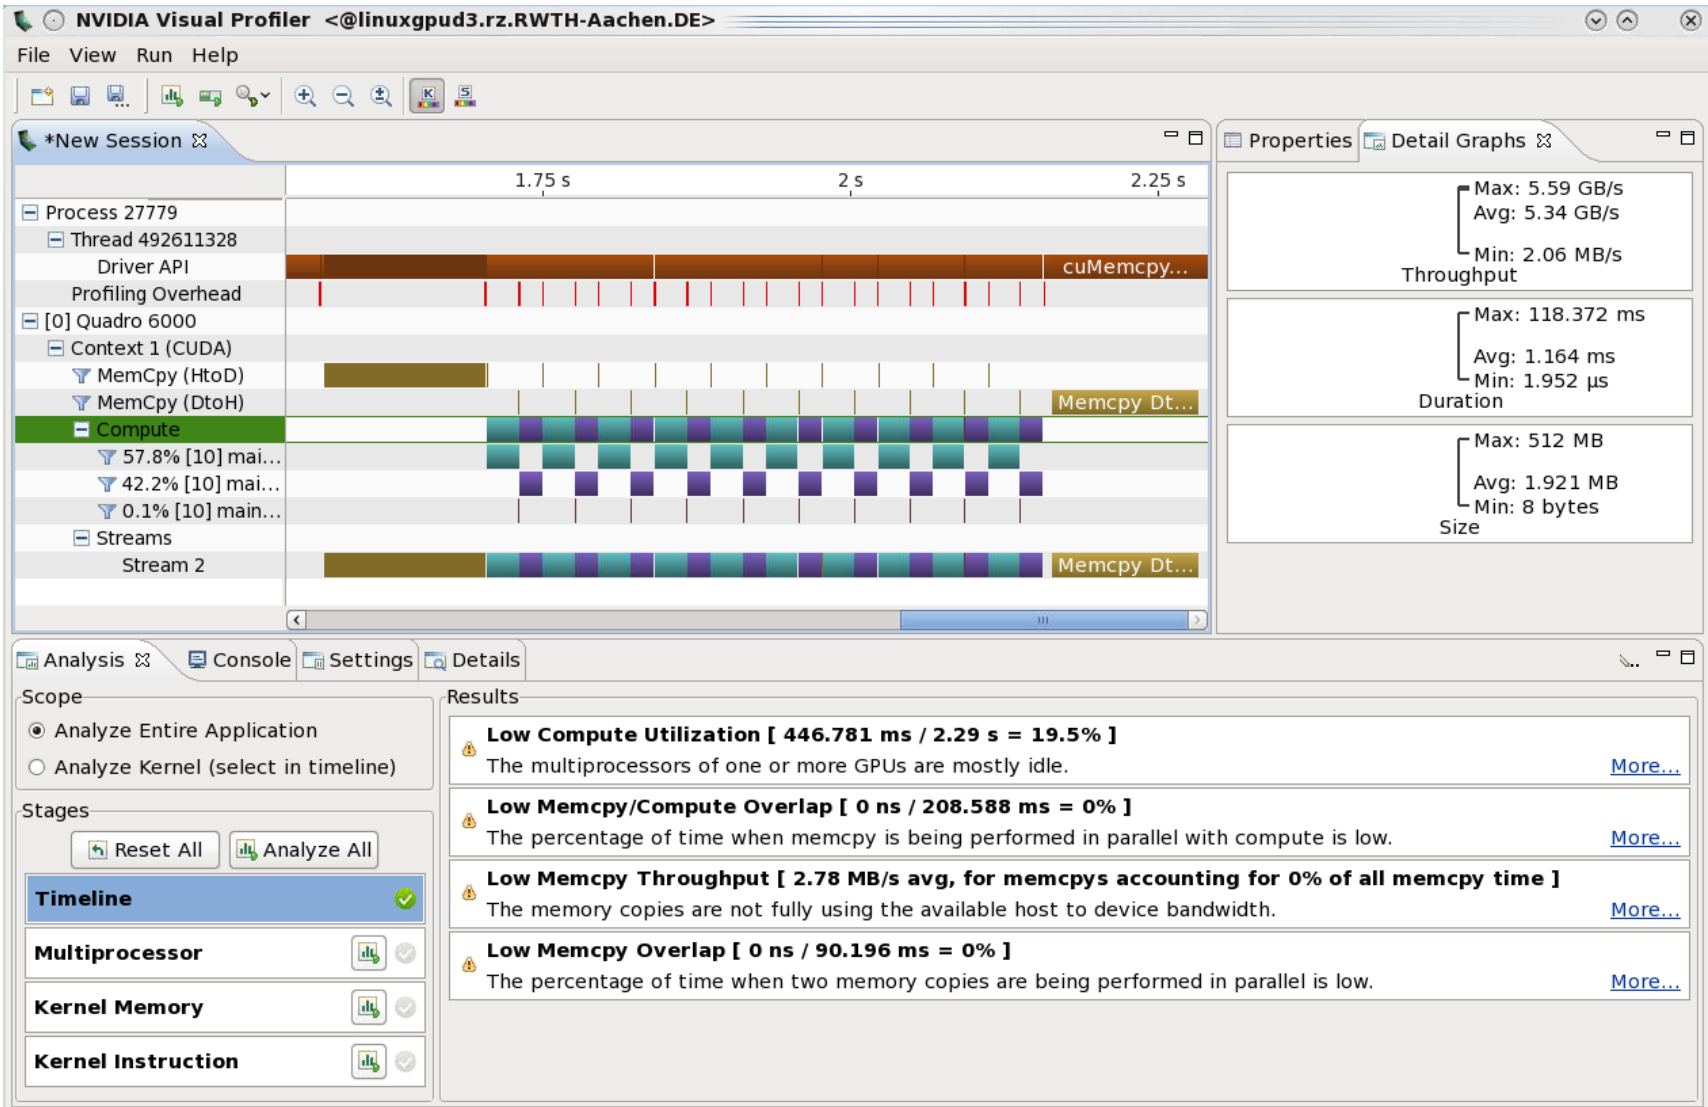
- Proceed (even if errors)

- Switch perspective

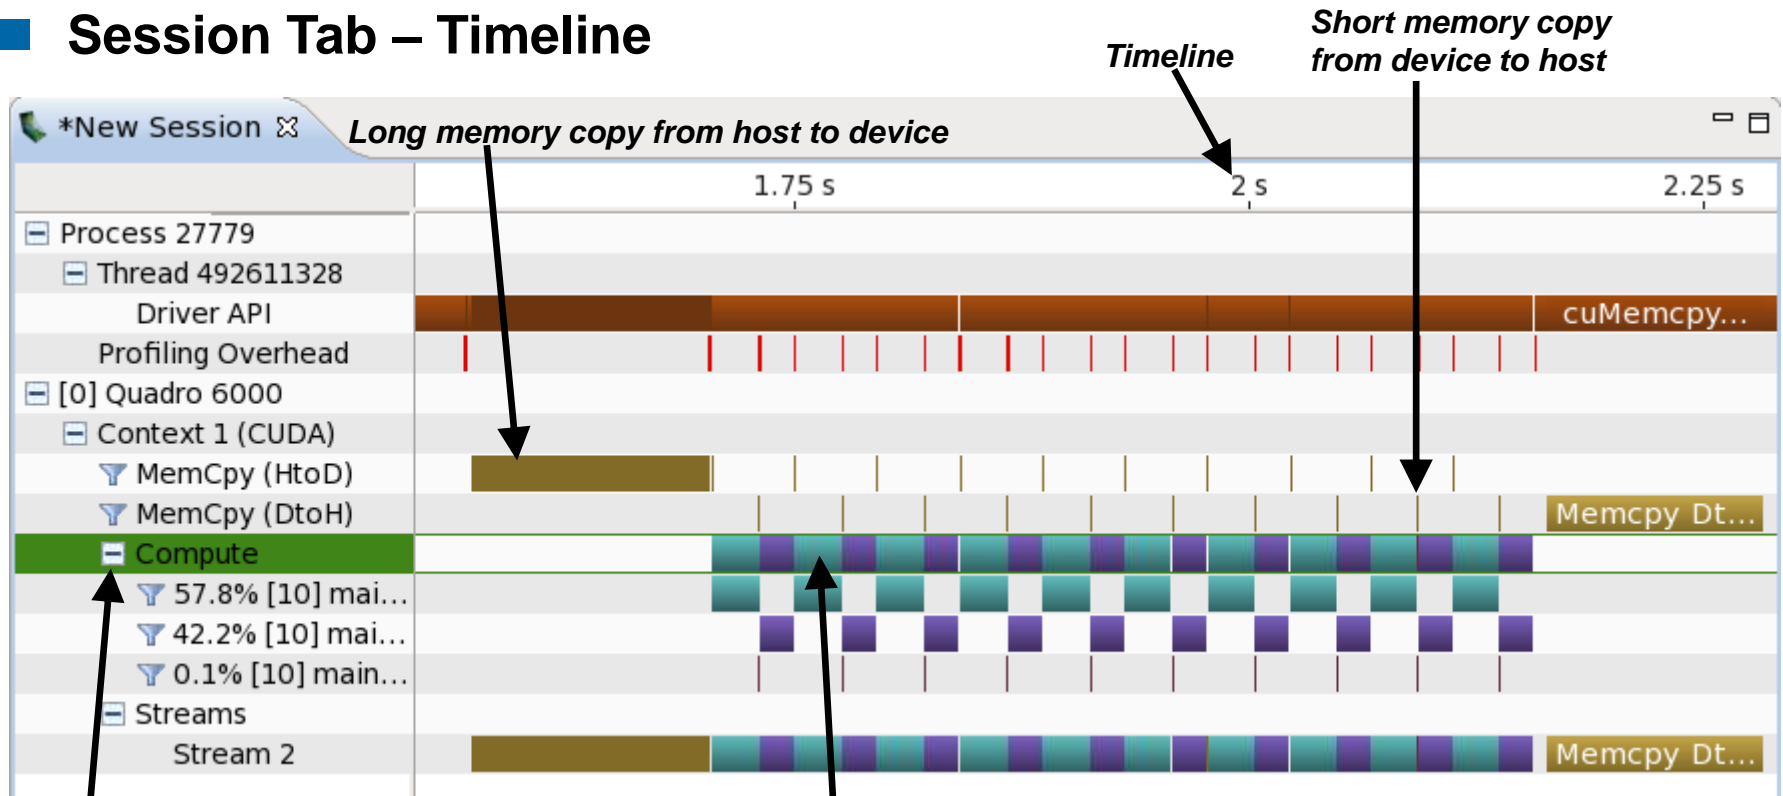
- Output/ interpretation see chapter NVIDIA Visual Profiler (next slide)

■ Use the NVIDIA Visual Profiler

1. Make sure that you have an executable program & start the profiler
2. Select `File` → `New Session`
3. Chose your executable as file
 - Specify arguments (if needed)
 - Specify environment variables (if needed)
4. If you want to shorten the execution time, set a timeout limit
5. `Finish` the session configuration & wait for results



■ Session Tab – Timeline



Collpase to see summarized info only

Compute time for first kernel

- Is data only transfered when needed?
- Which kernel does need the most time?

■ Analysis Tab

→ Gives hints for optimization (not always useful)

■ Details Tab

→ Switch from Analysis Tab to the Details Tab

			<i>Runtime</i>	<i>Grid dimensions</i>						
Name	Start Time	Duration	Grid Size	Block Size	Regs	Static SMem	Dynamic SMem	Size	Throughput	
Memcpy HtoD [async]	2.111 s	3.552 µs	n/a	n/a	n/a	n/a	n/a	n/a	2.15 MB/s	
main_80_gpu	2.092 s	18.829 ms	[128,2048,1]	[64,4,1]	19	0	0	n/a	n/a	
Memcpy DtoH [async]	2.092 s	1.952 µs	n/a	n/a	n/a	n/a	n/a	n/a	3.91 MB/s	
main_73_gpu_red	2.092 s	36.355 µs	[1,1,1]	[256,1,1]	12	0	2048	n/a	n/a	
main_69_gpu	2.067 s	25.349 ms	[16,1024,1]	[32,8,1]	26	0	2048	n/a	n/a	
Memcpy HtoD [async]	2.067 s	3.584 µs	n/a	n/a	n/a	n/a	n/a	n/a	2.13 MB/s	

→ On right hand side, activate summary view



Name	Avg. Duration	Regs	Static SMem	Avg. Dynamic SMem
main_73_gpu_red	36.355 µs	12	0	2048
main_80_gpu	18.833 ms	19	0	0
main_69_gpu	25.808 ms	26	0	2048

Kernel name <func>_<line>_gpu