


HPC-Europa2 TAM 2012 tutorial

**It doesn't work..
What do I do?**

Jeroen Engelberts & John Donners



GEEKS

Any questions?



About us

■ Jeroen Engelberts

- 1991-1997 Msc in chemometrics
- 1997-2000 Intermezzo: commercial ICT
- 2000-2004 PhD research in theoretical chemistry
- 2005-2010 SARA – eScience Support
- 2010-now SARA – Supercomputing

■ John Donners

- 1996-2001 MSc in physics.
- 2001-2005 PhD in oceanography. Analysis of high-resolution ocean simulations.
- 2005-2008 Climate simulations on the Earth Simulator in Japan.
- 2008-now SARA: Support team for high-performance computing. Debugging, parallellizing, optimizing and scaling of user applications.

A bit about you.. Who uses

- ▀ Fortran?
- ▀ C?
- ▀ MPI?
- ▀ OpenMP?
- ▀ print-statements for debugging?
- ▀ gdb?
- ▀ Totalview?
- ▀ Valgrind?



Acknowledgements & references

- Many examples were taken from the Cyberinfrastructure Tutor (or CITutor), the web-based training site for HPC hosted by NCSA:

<http://www.citutor.org>

It provides on-line training material for debugging, but also for MPI, OpenMP, optimisation, visualisation and more

- HLRS provides the "Online Parallel Programming Workshop", with complete presentations (including audio):

<http://www.hlr.de/organization/sos/par/services/training/par-prog-ws/>

HLRS also has excellent off-line courses in Stuttgart:

<http://www.hlr.de/organization/sos/par/services/training/course-list/>

<http://www.hlr.de/events/>

- TotalView kindly provided their presentations and training material about the TotalView Parallel Debugger.

The age of debugging

- ▶ **Mark II**
- ▶ **Built at Harvard University, financed by the navy.**
- ▶ **about 5 flops**



Assumptions..

- ▶ Question: “The same program works for my colleague.”
- ▶ Possible problem: do you really use the same executable?
- ▶ Check: use 'which [program]' to see the location.

- ▶ Possible problem: dynamic executables load libraries at runtime. These are loaded from directories in your `$LD_LIBRARY_PATH`. Maybe you use different versions of libraries without knowing it?
- ▶ Check: use 'ldd' to check which libraries are used.

Assumptions..

- ▶ Question to the helpdesk: the program works fine on my own system (and many others), so something is wrong with your system.

```
program helloworld
  implicit none

  print*, 'Hello World!'
  return
end program
```

- ▶ works fine with GNU and IBM compilers, but won't compile with the Intel compiler. **Why?**


Assumptions..

- ▶ **return statement is not allowed in Fortran main-program**

```
> ifort return.f90
```

```
return.f90(5): error #6353: A RETURN statement is invalid  
in the main program.
```

- ▶ **Solution: check your program to see if it follows the language standard.**



www.phdcomics.com

“Explaining your code to someone else (even a teddy bear) is wonderfully effective.” - Kernighan & Pike

Glossary

- ▶ (physical) processor, core, node
- ▶ logical processor
- ▶ compiling, linking
- ▶ object file, shared object/dynamic library, static library
- ▶ process, thread
- ▶ stack, heap
- ▶ stack trace, core dump
- ▶ runtime environment

Static library

A collection of object files of which the needed symbols are copied into the executable at the linking stage.

■ ***Example: .a files on unix.***

Shared object/Dynamic library

A collection of object files whose filename is linked into the executable at the linking stage. The dynamic library is searched and loaded at runtime.

▶ *Example: .so files on Linux or .dll files on Windows*

file: determine file type

```
donners@p6012:~> file mpicursus.odp
mpicursus.odp: OpenDocument Presentation
donners@p6012:~> file mpicursus.pdf
mpicursus.pdf: PDF document, version 1.4
donners@p6012:~> file /sara/sw/gromacs/4.0.7-sp/bin/mdrun_mpi
/sara/sw/gromacs/4.0.7-sp/bin/mdrun_mpi: ELF 64-bit MSB executable,
64-bit PowerPC or cisco 7500, version 1 (SYSV), for GNU/Linux 2.6.4,
dynamically linked (uses shared libs), not stripped
donners@p6012:~> file .bashrc
.bashrc: ASCII text
```

which: what executable am I running?

- ▶ If you don't use an absolute path to your executable, the shell executes the first one in your \$PATH-variable.
- ▶ but are you sure what is your \$PATH at any moment?
- ▶ Check it with:

`which program`



Idd: print shared library dependencies

```
donners@p6012:~> ldd /sara/sw/gromacs/4.0.7-sp/bin/mdrun_mpi
linux-vdso64.so.1 => (0x0000000000100000)
libgslcblas.so.0 => /sara/sw/gsl/1.11/lib/libgslcblas.so.0
(0x0000040000040000)
libxml2.so.2 => /usr/lib64/libxml2.so.2 (0x00000400000d0000)
libz.so.1 => /lib64/libz.so.1 (0x00000400002d0000)
libgsl.so.0 => /sara/sw/gsl/1.11/lib/libgsl.so.0 (0x0000040000300000)
libnsl.so.1 => /lib64/libnsl.so.1 (0x0000040000560000)
libm.so.6 => /lib64/power6/libm.so.6 (0x00000400005a0000)
libssl.so.1 => /usr/lib64/libssl.so.1 (0x0000040000670000)
libxlf90_r.so.1 => /opt/ibmcmp/lib64/libxlf90_r.so.1 (0x0000040002460000)
...
```

- can be used on dynamic executables and dynamic libraries
- libraries are searched in `$LD_LIBRARY_PATH`
- Long addresses indicate 64-bit code

nm: list symbols

```

...
donners@p6012:~/DEISA/turflame/turflame-openmp3> nm pois.o
...
                U __domaindecomposition_NMOD_usedomain
0000000000000000 D
__multigridpoissonsolver_NMOD_&&multigridpoissonsolver
0000000000000048 D __multigridpoissonsolver_NMOD_maxvalue
0000000000000018 D __multigridpoissonsolver_NMOD_pois
00000000000000a8 D __multigridpoissonsolver_NMOD_pplus
0000000000000030 D __multigridpoissonsolver_NMOD_residu
0000000000000060 D __multigridpoissonsolver_NMOD_restrict
0000000000000090 D __multigridpoissonsolver_NMOD_setdum
0000000000000078 D __multigridpoissonsolver_NMOD_sweep
                U _xlfBeginIO
                U _xlfEndIO
                U mpi_allreduce
...

```

works on object files, executables, libraries

Useful to find undefined references



strace: trace system calls

```
donners@lisa:~$ strace ./les3d.hybrid
```

```
..  
write(1, "Loading Flamelet Generated Manifo"...  
= 38  
etcwd("/home/donners/DEISA/turflame/turflame-openmp3"...  
tat("/home/donners/DEISA/turflame/turflame-openmp3/FGM_DIFF.dat", {st_mode=S_IFREG|0644,  
t_size=102255881, ...}) = 0  
etcwd("/home/donners/DEISA/turflame/turflame-openmp3"...  
pen("/home/donners/DEISA/turflame/turflame-openmp3/FGM_DIFF.dat", O_RDWR|O_CREAT, 0666) = 23  
octl(23, SNDCTL_TMR_TIMEBASE or TCGETS, 0x7fff7056f600) = -1 ENOTTY (Inappropriate ioctl for  
evice)  
stat(23, {st_mode=S_IFREG|0644, st_size=102255881, ...}) = 0  
octl(23, SNDCTL_TMR_TIMEBASE or TCGETS, 0x7fff7056f600) = -1 ENOTTY (Inappropriate ioctl for  
evice)  
read(23, "FLAMELET GENERATED MANIFOLD\n\n[NUM"...  
= 8192  
..
```

- + Often useful for I/O problems.
- + No recompilation needed.
- Output can be overwhelming. Redirect output and search for last output from program.

Questions & Answers

