

The Case Cluster and Cosmomc

Pascal M. Vaudrevange

24.07.2009

1 Modules

2 Queues

3 CosmoMC

Modules

```

pmv9@master:~> module avail
----- /usr/local/share/modulefiles -----
R
abinit-5.3.5cp      intel-compilers-8.0   mkl-7.0.1.6-32
apbs                intel-compilers-8.1   mkl-7.0.1.6-64
apbs-0.5.0         intel-compilers-8.1-64 mkl-9.0-64
apbs-0.5.1         intel-f                mkl-9.0.018-32
apbs-1.0           intel-f-10.0.025-32   mkl-9.0.018-64
apbs-1.0.0         intel-f-10.0.025-64   moab
charm              jdk                    molcas
charmm             jdk-1.6               molcas-28-Nov-07
fftw-2.1.5         jre-1.6               molden4.6
fftw-2.1.5.rab     lammps                molpro
fftw-2.1.5pgi      lammps-20Dec07        mpich
fftw-3.1.2         lammps-20May08        mpich-gcc-32
fftw-3.1.2pgi      mathematica            mpich-icc-32
fluent             mathematica6           mpich-pgi
fluent12.0         matlab                mpich1
fluent6.3          matlab_R2007a         mpich2
gate              matlab_R2007a_serial  mysql-5.0.1b
gaussian           matlab_R2007a_worker  namd
gaussian-E01       matlab_R2007b         namd-2.6-amd64-tcp
grace              matlab_R2007b_worker  neuron
grace-5.1.21      matlab_R2008a         newlammps
gromacs            matlab_R2008a_worker  openmpi
gromacs-4.0.3     matlab_R2008b         pgi
gromacs-4.0.3ser  matlab_R2008b_worker  pgi-7.0-7
gromacs-4.0.5     matlab_R2009a         pgi-8.0-5
gsl-1.10          matlab_R2009a_worker  rosetta
idl               matlab_worker         scalapack
idl-7.0           maui                  vasp
intel-c           mkl                   vmd
intel-c-10.0.025  mkl-10.0.3.020       vmd-1.8.6

----- /usr/local/modules-env-3.2.5/Modules/versions -----
3.2.5

----- /usr/local/modules-env-3.2.5/Modules/3.2.5/modulefiles -----
dot      module-cvs  module-info  modules      null      use.own
pmv9@master:~>

```

Modules

```
pmv9@master:~> module list
Currently Loaded Modulefiles:
  1) modules  2) moab      3) mpich
pmv9@master:~> module show moab
-----
/usr/local/share/modulefiles/moab:

module-whatis  loads the Moab and Silver related tools
setenv         MAUIHOMEDIR /usr/local/moab
prepend-path   PATH /usr/local/moab/bin
prepend-path   MANPATH /usr/local/moab/man
-----

pmv9@master:~> ifort
-bash: ifort: command not found
pmv9@master:~> module load intel-f
pmv9@master:~> ifort
ifort: command line error: no files specified; for help type "ifort -help"
pmv9@master:~> module list
Currently Loaded Modulefiles:
  1) modules  2) moab      3) mpich      4) intel-f
pmv9@master:~> module unload mpich
pmv9@master:~> module list
Currently Loaded Modulefiles:
  1) modules  2) moab      3) intel-f
pmv9@master:~>
```

Who is running what?

```
pmv9@master:~> showq
```

```
active jobs-----
JOBID          USERNAME      STATE PROCS   REMAINING          STARTTIME
221933         jar36   Canceling    4 -3:14:04:26 Sun Jul 19 17:20:46
225017         pmv9     Running      6   3:04:13 Thu Jul 23 12:29:25
225333         jxm343   Running      4   7:27:19 Thu Jul 23 14:52:31
225098         pmv9     Running      6   7:56:40 Thu Jul 23 17:21:52
225100         pmv9     Running      6   7:58:18 Thu Jul 23 17:23:30
225101         pmv9     Running      6   7:58:46 Thu Jul 23 17:23:58
.
.
.
224962         jxz82    Running      1 12:06:18:59 Thu Jul 23 11:44:11
```

```
313 active jobs          630 of 716 processors in use by local jobs (87.99%)
                        125 of 136 nodes active          (91.91%)
```

```
eligible jobs-----
JOBID          USERNAME      STATE PROCS   WCLIMIT          QUEUEUETIME
```

```
0 eligible jobs
```

```
blocked jobs-----
JOBID          USERNAME      STATE PROCS   WCLIMIT          QUEUEUETIME
225031         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:29
225032         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:30
225033         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:30
225034         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:30
225035         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:30
225036         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:30
225037         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:30
225038         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:30
225039         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:31
225040         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:31
225041         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:31
225042         pmv9     Idle        6   8:00:00 Thu Jul 23 12:29:31
.
```

What about my jobs?

```
pmv9@master:~> qstat -u pmv9
```

```
master.priv.cwru.edu:
```

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	S	Elap Time
225017.master.priv.c	pmv9	batch	scramble_a	20865	1	--	--	08:00	R	04:56
225031.master.priv.c	pmv9	batch	scramble_a	--	1	--	--	08:00	Q	--
225032.master.priv.c	pmv9	batch	scramble_a	--	1	--	--	08:00	Q	--
225033.master.priv.c	pmv9	batch	scramble_a	--	1	--	--	08:00	Q	--
225034.master.priv.c	pmv9	batch	scramble_a	--	1	--	--	08:00	Q	--
.										
.										
.										

More Detailed Info

```
pmv9@master:~> qstat -f 225017
Job Id: 225017.master.priv.cwru.edu
Job_Name = scramble_alms_open_60_67
Job_Owner = pmv9@master.priv.cwru.edu
resources_used.cput = 36:27:43
resources_used.mem = 452432kb
resources_used.vmem = 613128kb
resources_used.walltime = 04:57:29
job_state = R
queue = batch
server = master.priv.cwru.edu
Checkpoint = u
ctime = Thu Jul 23 12:29:25 2009
Error_Path = master:/home/pmv9/projects/circles/scramble_alms_open_60/scr
mble_alms_open_60_67.e$PBS_JOBID
exec_host = quad32/6+quad32/5+quad32/4+quad32/3+quad32/2+quad32/1
Hold_Types = n
Join_Path = n
Keep_Files = n
Mail_Points = a
mtime = Thu Jul 23 12:29:26 2009
Output_Path = master:/home/pmv9/projects/circles/scramble_alms_open_60/scr
amble_alms_open_60_67.o$PBS_JOBID
Priority = 0
qtime = Thu Jul 23 12:29:25 2009
Rerunable = True
Resource_List.nodect = 1
Resource_List.nodes = 1:ppn=6
Resource_List.walltime = 08:00:00
session_id = 20865
Shell_Path_List = /bin/bash
Variable_List = PBS_O_HOME=/home/pmv9,PBS_O_LANG=en_SG.UTF-8,
PBS_O_LOGNAME=pmv9,
PBS_O_PATH=/home/pmv9/bin:/usr/local/mpich-1.2.7/bin:/usr/local/mpi/
bin:/usr/local/modules-env-3.2.5/Modules/3.2.5/bin:/usr/kerberos/bin:/
usr/local/cuda/bin:/usr/local/bin:/bin:/usr/bin:/usr/X11R6/bin:/home/p
mv9/bin,PBS_O_MAIL=/var/spool/mail/pmv9,PBS_O_SHELL=/bin/bash,
PBS_O_HOST=master.priv.cwru.edu,
PBS_O_WORKDIR=/home/pmv9/projects/circles,PBS_O_QUEUE=batch
etime = Thu Jul 23 12:29:25 2009
```

Most Detailed Info

```
pmv9@master:~> ssh quad32
pmv9@master:~> top
top - 17:31:56 up 62 days, 22:58, 1 user, load average: 9.99, 9.92, 9.84
Tasks: 113 total, 4 running, 109 sleeping, 0 stopped, 0 zombie
Cpu(s): 100.0% us, 0.0% sy, 0.0% ni, 0.0% id, 0.0% wa, 0.0% hi, 0.0% si
Mem: 16396860k total, 10632392k used, 5764468k free, 156488k buffers
Swap: 8385920k total, 0k used, 8385920k free, 9706716k cached

  PID USER      PR  NI  %CPU   TIME+  %MEM  VIRT  RES  SHR  S  COMMAND
20931 pmv9       25   0   756  2225:20  2.7  533m 437m 3288 R  circles
20786 zx145     25   0   44   1183:08  0.1  21124 12m 1436 R  lmp_serial
20988 pmv9       15   0   2    0:00.11  0.0  5364 1112 808 R  top
  1 root       15   0   0    0:02.57  0.0  4824 648 540 S  init
  2 root       RT   0   0    0:00.22  0.0  0 0 0 S  migration/0
  3 root       34  19  0    0:00.05  0.0  0 0 0 S  ksoftirqd/0
  4 root       RT   0   0    0:00.00  0.0  0 0 0 S  watchdog/0
  5 root       RT   0   0    0:00.23  0.0  0 0 0 S  migration/1
  6 root       34  19  0    0:00.04  0.0  0 0 0 S  ksoftirqd/1
  7 root       RT   0   0    0:00.00  0.0  0 0 0 S  watchdog/1
  8 root       RT   0   0    0:00.27  0.0  0 0 0 S  migration/2
  9 root       34  19  0    0:00.08  0.0  0 0 0 S  ksoftirqd/2
 10 root       RT   0   0    0:00.00  0.0  0 0 0 S  watchdog/2
 11 root       RT   0   0    0:00.22  0.0  0 0 0 S  migration/3
pmv9@quad32:~> logout
Connection to quad32 closed.
pmv9@master:~> qdel 225017
```


Submitting your own job

```
pmv9@master:~/projects/circles> cat scramble_alms_open_45_1.sub
#!/bin/bash
#PBS -l walltime=08:00:00
#PBS -l nodes=1:ppn=6
#PBS -N scramble_alms_open_45_1
#PBS -S /bin/bash
#PBS -e "/home/pmv9/projects/circles/scramble_alms_open_45/scramble_alms_open_45_1.e$PBS_JOBID" -o "/home/pmv9/p

cp -r /home/pmv9/projects/circles/maps $TMPDIR
cp -r /home/pmv9/projects/circles/./circles /home/pmv9/projects/circles/scramble_alms_open_45_1.ini $TMPDIR

cd $TMPDIR
module load intel-f
module load fftw-3.1.2
module list
ulimit -s unlimited
ulimit -c unlimited
export decfort_dump_flag=y
time /usr/local/bin/mpiexec -pernode ./circles scramble_alms_open_45_1.ini > scramble_alms_open_45_1.log
cp -ru scramble_alms_open_45_1* /home/pmv9/projects/circles/scramble_alms_open_45

pmv9@master:~/projects/circles>
pmv9@master:~/projects/circles> qsub scramble_alms_open_45_1.sub
```

CosmoMC/Camb (<http://cosmocooffee.info>)

- Camb (<http://camb.info>)
 - Boltzman solver
 - evolves primordial perturbation power spectra defined in
`powertilt.f90:ScalarPower()`
`powertilt.f90:TensorPower()`
to today's C_ℓ
 - stand-alone program and library, OpenMP
- CosmoMC (<http://cosmologist.info/cosmomc>)
 - Parameter estimation package for cosmological parameters
 - uses Camb + MPI
 - needs WMAP5yr likelihood software + cfitsio

Installing Camb

```
pmv9@master:~> wget http://cosmologist.info/cosmomc/cosmomc.tar.gz
--09:41:55-- http://cosmologist.info/cosmomc/cosmomc.tar.gz
=> `cosmomc.tar.gz'
Resolving cosmologist.info... 66.39.96.252
Connecting to cosmologist.info[66.39.96.252]:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 4,794,135 (4.6M) [application/x-gzip]

100%[=====] 4,794,135 1.35M/s ETA 00:00

09:41:59 (1.34 MB/s) - `cosmomc.tar.gz' saved [4794135/4794135]

pmv9@master:~> tar -xzf cosmomc.tar.gz
pmv9@master:~> cd cosmomc/camb
pmv9@master:~/cosmomc/camb> make
ifort -openmp -O2 -ip -W0 -WB -fpp2 -vec_report0 -c utils.F90
ifort -openmp -O2 -ip -W0 -WB -fpp2 -vec_report0 -c subroutines.f90
ifort -openmp -O2 -ip -W0 -WB -fpp2 -vec_report0 -c inifile.f90
.
.
cmbmain.f90 (2060): (col. 7) remark: OpenMP DEFINED LOOP WAS PARALLELIZED.
cmbmain.f90 (1011): (col. 7) remark: OpenMP DEFINED LOOP WAS PARALLELIZED.
ifort -openmp -O2 -ip -W0 -WB -fpp2 -vec_report0 -c camb.f90
ifort -openmp -O2 -ip -W0 -WB -fpp2 -vec_report0 utils.o subroutines.o
inifile.o power_tilt.o recfast.o reionization.o modules.o bessels.o
equations.o halofit.o lensing.o cmbmain.o camb.o inidriver.F90 -o camb
pmv9@master:~/cosmomc/camb> make all
ar -r libcamb.a subroutines.o inifile.o power_tilt.o recfast.o
reionization.o modules.o bessels.o equations.o halofit.o lensing.o
cmbmain.o camb.o
ar: creating libcamb.a
```

Installing cfitsio

```
pmv9@master:~> wget ftp://heasarc.gsfc.nasa.gov/software/fitsio/c/cfitsio3140.tar.gz
--10:15:08-- ftp://heasarc.gsfc.nasa.gov/software/fitsio/c/cfitsio3140.tar.gz
=> 'cfitsio3140.tar.gz'
Resolving heasarc.gsfc.nasa.gov... 128.183.240.223
Connecting to heasarc.gsfc.nasa.gov[128.183.240.223]:21... connected.
Logging in as anonymous ... Logged in!
==> SYST ... done.      ==> PWD ... done.
==> TYPE I ... done.   ==> CWD /software/fitsio/c ... done.
==> PASV ... done.    ==> RETR cfitsio3140.tar.gz ... done.
Length: 2,903,378 (2.8M) (unauthoritative)

100%[=====] 2,903,378      4.47M/s

10:15:09 (4.46 MB/s) - 'cfitsio3140.tar.gz' saved [2903378]

pmv9@master:~> tar -xzf cfitsio3140.tar.gz
pmv9@master:~> cd cfitsio
pmv9@master:~/cfitsio> ./configure
checking for uname... 1
.
.
.
config.status: creating cfitsio.pc

    Congratulations, Makefile update was successful.
    You may want to run "make" now.

pmv9@master:~/cfitsio> make
pmv9@master:~/cfitsio> make
make[1]: Entering directory '/home/pmv9/cfitsio'
cc -c -g -O2 -Dg77Fortran -fPIC -DPACKAGE_NAME="\ " -DPACKAGE_TARNAME="\ " -DPACKAGE_VER
.
.
.
a - f77_wrap3.o
a - f77_wrap4.o
make[1]: Leaving directory '/home/pmv9/cfitsio'
pmv9@master:~/cfitsio> make shared
cc -shared -o libcfitsio.so buffers.o cfileio.o checksum.o compress.o drvrfiler.o drvrmem.o drvrmnet.o drvrmem.o
pmv9@master:~/cfitsio> cp libcfitsio.* ~/opt/lib/
```

Installig WMAP5yr likelihood code

```
pmv9@master:~> wget http://lambda.gsfc.nasa.gov/data/map/dr3/dcp/wmap_likelihood_sub_v3p2.tar.gz
--09:55:36-- http://lambda.gsfc.nasa.gov/data/map/dr3/dcp/wmap_likelihood_sub_v3p2.tar.gz
=> 'wmap_likelihood_sub_v3p2.tar.gz'
Resolving lambda.gsfc.nasa.gov... 128.183.175.60
Connecting to lambda.gsfc.nasa.gov|128.183.175.60|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 400,825,770 (382M) [application/x-tar]

100%[=====] 400,825,770 5.20M/s ETA 00:00

09:56:53 (4.97 MB/s) - 'wmap_likelihood_sub_v3p2.tar.gz' saved [400825770/400825770]

pmv9@master:~> tar -xzf wmap_likelihood_sub_v3p2.tar.gz
pmv9@master:~> cd likelihood_v3
pmv9@master:~/likelihood_v3> vi Makefile
#Makefile --- WMAP Likelihood code...

#CFITSIO=/usr/local
CFITSIO=/home/pmv9/software
.
.
F90 = ifort
FFLAGS = -O $(WMAFFLAGS)
INCS = -I. -I$(CFITSIO)/include
LIBS = -L. -L$(CFITSIO)/lib -lcfitsio -L$(MKL_HOME)/lib/em64t -lmkl_intel_lp64 -lmkl_intel_thread
-lmkl_core -lguidt -lpthread

pmv9@master:~/likelihood_v3> module load mkl-10.0.3.020
pmv9@master:~/likelihood_v3> make
ifort -O -DOPTIMIZE -I. -I/home/pmv9/opt/include -c -o read_archive_map.o
read_archive_map.f90
read_archive_map.f90(56): (col. 42) remark: LOOP WAS VECTORIZED.
read_archive_map.f90(57): (col. 42) remark: LOOP WAS VECTORIZED.
.
.
/usr/local/intel/fce/10.0.025/lib/libimf.so: warning: warning: feupdateenv is not implemented
and will always fail
pmv9@master:~/likelihood_v3>
```

Caveats

```
pmv9@master:~> module show mkl-10.0.3.020
-----
/usr/local/share/modulefiles/mkl-10.0.3.020:

module-whatis    loads the Intel MKL mkl-10.0.3.020 for 64 bits
setenv           OMP_NUM_THREADS 1
prepend-path     LD_LIBRARY_PATH /usr/local/intel/mkl/10.0.3.020/lib/em64t
setenv           MKL_HOME /usr/local/intel/mkl/10.0.3.020
setenv           MKL -L/usr/local/intel/mkl/10.0.3.020/lib/em64t -lmkl_intel_lp64 -lmkl_intel_thread
                 -lmkl_core -Wl,--end-group -Wl,-Bdynamic -lguide -lpthread
setenv           MKL_STATIC -L/usr/local/intel/mkl/10.0.3.020/lib/em64t -lmkl_em64t -lmkl_lapack
                 -lmkl_scalapack -lguide
-----

pmv9@master:~>
```

Installing CosmoMC

```
pmv9@master:~> cd cosmomc/source/
pmv9@master:~/cosmomc/source> vi Makefile
#You may need to edit the library paths for MKL for Intel
#Beware of using optimizations that lose accuracy - may give errors when running

#Edit for CFITSIO directories
cfitsio = /home/pmv9/opt
#/usr/local/Cluster-Users/cpac/cmb/2.1.0/cfitsio

WMAP = /home/pmv9/likelihood_v3

IFLAG = -I
INCLUDE=

#Intel MPI (assuming mpif77 set to point to ifort)
#change -lmkl_ia32 to -lmkl_p3 for MKL versions earlier than 6 (6 needed for ifc 8+)
F90C = mpif90
FFLAGS = -O2 -ip -W0 -WB -openmp -fpp -DMPI -vec_report0
#LAPACKL = -L/usr/local/intel/mkl/9.1/lib/em64t -lmkl_lapack -lmkl_em64t -lguide -lpthread
LAPACKL = -L$(MKL_HOME)/lib/em64t -lmkl_intel_lp64 -lmkl_intel_thread -lmkl_core -lguide -lpthread
.
.
.

pmv9@master:~/cosmomc/source> make
mpif90 -DMATRIX_SINGLE -O2 -ip -W0 -WB -openmp -fpp -DMPI -vec_report0
-I/home/pmv9/opt/include -I/home/pmv9/likelihood_v3 -I../camb -c utils.F90
.
.
mpif90 -o ../cosmomc /home/pmv9/likelihood_v3/read_archive_map.o /home/pmv9/likelihood_v3/read_fits.o /home/pmv9/likelihood_v3/read_archive.o
pmv9@master:~/cosmomc/source> make getdist
mpif90 -DMATRIX_SINGLE -O2 -ip -W0 -WB -openmp -fpp -DMPI -vec_report0 -I/home/pmv9/opt/include -I/home/pmv9/likelihood_v3 -I../camb -c getdist.F90
mpif90 -o ../getdist utils.o Matrix_utils.o settings.o GetDist.o -L/home/pmv9/opt/lib -L/home/pmv9/likelihood_v3/lib -L../camb -c getdist.F90
pmv9@master:~/cosmomc/source> cd ..
pmv9@master:~/cosmomc> ./cosmomc params.ini
Number of MPI processes: 1
Random seeds: 28388, 9014 rand_inst: 1
Computing tensors: F
Doing CMB lensing: T
lmax = 2100
```

Debugging CosmoMC

```

pmv9@master:~/cosmomc> mpirun ./cosmomc params.ini
forrtl: severe (59): list-directed I/O syntax error, unit -5, file Internal List-Directed Read
Image      PC          Routine      Line        Source
cosmomc    00000000005F2B2A  Unknown     Unknown     Unknown
cosmomc    00000000005F1D2A  Unknown     Unknown     Unknown
cosmomc    00000000005A34D6  Unknown     Unknown     Unknown
cosmomc    00000000005572A2  Unknown     Unknown     Unknown
cosmomc    00000000005568AE  Unknown     Unknown     Unknown
cosmomc    000000000057203B  Unknown     Unknown     Unknown
cosmomc    0000000000571069  Unknown     Unknown     Unknown
cosmomc    00000000004B1AE9  Unknown     Unknown     Unknown
cosmomc    0000000000409A22  Unknown     Unknown     Unknown
libc.so.6  0000003DD9E1C3FB  Unknown     Unknown     Unknown
cosmomc    000000000040996A  Unknown     Unknown     Unknown
pmv9@master:~/cosmomc>

```


And now..

Let's do it!