UF Research Computing

Environment Modules for trouble-free biocomputing and the many faces of UF HPC batch systems

Oleksandr Moskalenko <om@hpc.ufl.edu>
Matt Gitzendanner <magitz@ufl.edu>

February 2012 UF HPC Biocomputing Training Session

Outline

- ✦ Environment Modules
  - Need for speed i.e. reason for Env. Modules
  - Introduction and basic use
  - Module Tree
  - Writing modulefiles
- ✦ HPC PBS with modules
  - How to use modules in PBS pipelines

Environment Modules

UF Research Computing/HPC Center

The biggest ball of twine in…

App versions

Library versions

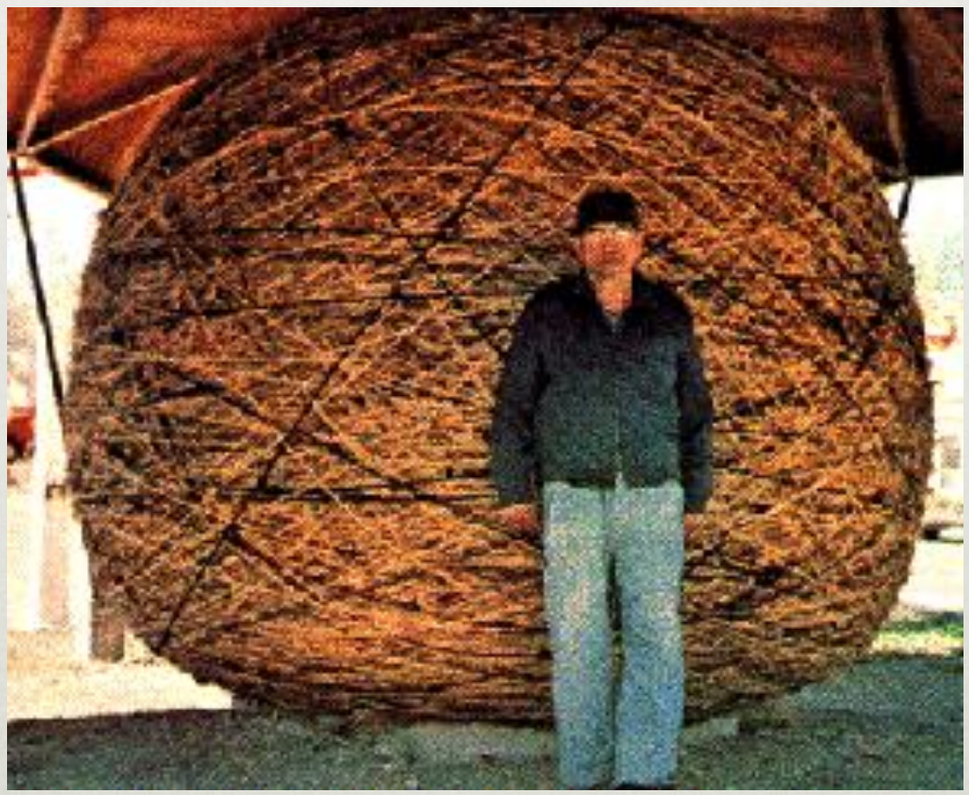

Compilers

ENV variables

Dependencies

PATH

Python versions

LD_LIBRARY_PATH

MPI implementations

The Ghost of the CLI

- Impossible to maintain many versions
- Environment setup is very complex and must be done by every user
- Things break on every upgrade
- No stable pipelines
- No way to switch Comp/MPI easily (the bad dream of the mpi-selector)
- Hard to find the right application
- Result: the nightmare of CLI

Untangle Compilers/MPIs

- mpich2_intel-1.0.8
- mvapich2_intel-1.4.1
- mvapich2_intel11-1.5.1
- mvapich_intel-0.9.9
- mvapich_intel10-0.9.9
- openmpi_intel-1.2.7
- openmpi_intel-1.3.4
- openmpi_intel11-1.3.4
- openmpi_intel11-1.4.3

2 Compilers

x

7 MPI implementations

=

A swamp nobody wants to walk into – most stick with the single default

The Zen of modules

module load zen

What is a module?

- ✦ Lmod software (Robert McLay, Ph.D. @ TACC)
- ✦ Shell environment
- ✦ Modulefiles – systemwide and personal

- ✦ Discover, evaluate, use, swap

Modules with benefits

- Forget the gory details
- Many application versions can coexist
- Easy switching of compilers/MPI versions
- Easy addition of dependencies
- Automated environment setup (paths, variables, licenses, etc….)
- Powerful search and help system
- Stable pipelines

Focus on Science, not busywork

Basic usage - Discover

- module spider
- module spider clu
- module spider netcdf/3.6.3
- module avail
- module avail clu (what does it find?)
- module key clu

Output of “module spider”

…

weblogo: weblogo/3.0

generation of sequence logos

weka: weka/3.6.2

WEKA - algorithms for data mining

wgs: wgs/6.1

Celera Whole Genome Shotgun Assembler

wublast: wublast/2.0

Basic Local Alignment Search Tool - WUSTL version

Output of “module spider clu”

clustalw: clustalw/2.1

Description: multiple sequence alignment software

Help: This module enables the use of the ClustalW software

Clustal is popular multiple sequence alignment software.

ClustalW produces biologically meaningful multiple sequence alignments of divergent sequences. It calculates the best match for selected sequences and lines them up, so identities, similarities, and differences can be identified.

Version 2.1

Discover – module avail

```
-------------- /apps/lmod/modulefiles/compiler/intel/10_1 ---------------
  amber/10        metis/4.0       mrbayes/3.1.2     openmpi/1.2.7
  fftw/3.2.2      mkl/10.0        mvapich/0.9.9     openmpi/1.3.4 (default)
  gromacs/4.0.7   mpich2/1.0.8    mvapich2/1.4.1
-------------- /apps/lmod/modulefiles/core ---------
  HPC                      gnuplot/4.4.3          python/2.6.4
  R/2.13.1                 hdf5/1.8.7             python/2.6.5
  abaqus/6.9               hyphy/2.0020110824beta  python/2.7.2 (default)
  add_scores/r1            hypre/2.0.0            raxml/1.0.5-light
  admb/10.0                intel/10.1             raxml/7.3.0 (default)
  admb/9.1 (default)       intel/11.1 (default)   repeatmasker/3.3.0
  amos/3.0.0               lastz/1.02.00          rosetta/2.1.1
  amos/runAmos.log (default)  lps/2.2             rosetta/2.1.2 (default)
```

Branches of the moduletree

Core:

---------- /apps/lmod/modulefiles/core ----------

clustalw/2.1

Compiler

---------- /apps/lmod/modulefiles/compiler/intel/10_1 ----------

mrbayes/3.1.2

MPI

---------- /apps/lmod/modulefiles/mpi/intel10_1/openmpi_1_2 ----------

amber/10 dock/64 gromacs/4.0.7 mpiblast/1.6.0 mrbayes/3.1.2 raxml/1.0.5-light

Basic usage - Evaluate

- module spider clustalw
- module whatis clustalw
- module show clustalw

Discover – module spider clustalw

clustalw: clustalw/2.1

Description: multiple sequence alignment software

Help: This module enables the use of the ClustalW software

Clustal is popular multiple sequence alignment software.

ClustalW produces biologically meaningful multiple sequence alignments of divergent sequences. It calculates the best match for selected sequences and lines them up, so identities, similarities, and differences can be identified.

Version 2.1

Evaluate – module whatis clustalw

```
clustalw/2.1       : Name: ClustalW
clustalw/2.1       : Version: 2.1
clustalw/2.1       : Category: biology,bioinformatics,alignment
clustalw/2.1       : URL: http://www.clustal.org/
clustalw/2.1       : Description: multiple sequence alignment software
```

Evaluate – module show clustalw

Top part – like module spider clustalw

Bottom part:

```
whatis("Name: ClustalW")
whatis("Version: 2.1")
whatis("Category: biology,bioinformatics,alignment")
whatis("URL: http://www.clustal.org/")
whatis("Description: multiple sequence alignment software")
prepend_path("PATH", "/apps/clustalw/2.1/bin")
setenv("HPC_CLUSTALW_DIR", "/apps/clustalw/2.1")
setenv("HPC_CLUSTALW_BIN", "/apps/clustalw/2.1/bin")
set_alias("clustalw", "clustalw2")
```

Use

- module list
- module load clustalw
- module load python/2.6.5
- module add
- module del, rm, unload clustalw – pick one!
- module load intel/10.1 openmpi/1.2.7 mrbayes

Loading and Unloading

/Put the module load command into the job script/

- module load clustalw
- module list

```
Currently Loaded Modules:

  1) clustalw/2.1
```

- module unload clustalw OR module purge
- module list

```
Warning: No modules installed
```

Dependencies

- module spider mrbayes/3.1.2

```
This module can only be loaded through the
following modules:

   intel/10.1

   intel/10.1, mvapich/0.9.9

   intel/10.1, openmpi/1.2.7
```

- module load intel/10.1 openmpi/1.2.7 mrbayes

Swap – module swap

- module load intel openmpi mrbayes
- module list

Currently Loaded Modules:

1) intel/10.1 2) openmpi/1.2.7 3) mrbayes/3.1.2

- module swap openmpi mvapich

Due to MODULEPATH changes the follow modules have been reloaded:

1) mrbayes

- module list

Currently Loaded Modules:

1) intel/10.1 2) mvapich/0.9.9 3) mrbayes/3.1.2

Your own environment

- module list

Currently Loaded Modules:

1) intel/10.1 2) mvapich/0.9.9 3) mrbayes/3.1.2

- module setdefault mymodz1
- module listdefault

Possible defaults:

1) mymodz1 2) test1 3) pipeline2

- module getdefault mymodz1
- module list

Currently Loaded Modules:

1) intel/10.1 2) mvapich/0.9.9 3) mrbayes/3.1.2

Your own modulefiles

- Create a directory, for example ~/.modules
- Create subdirectories named after modules, for example ~/.modules/happiness
- Put modulefiles into those subdirectories, for example ~/.modules/happiness/1.0.lua
- Use your own modules
  - $ module use ~/.modules
  - module {spider, load, unload} happiness
- Examples are in "/apps/lmod/modulefiles"

Job Scheduling and Usage

PBS

UF Research Computing/HPC Center

The tale of two systems

PBS vs. Grid Engine

UF Research Computing/HPC Center

PBS (main cluster)

- Access the cluster via ssh
  - $ ssh submit.hpc.ufl.edu
- Access test nodes via ssh from submit
  - test01, test04, test05

GE (test biocluster)

- Access the cluster via ssh
  - $ ssh bio.hpc.ufl.edu
- Access test/interactive node from bio
  - $ ssh biotest

Basic Submission Scripts

PBS

```
#!/bin/bash
#
#PBS -N My_Job_Name
#PBS -r n
#PBS -M Jane_Doe@ufl.edu
#PBS -m abe
#PBS -o My_Job_Name.out
#PBS -j oe
#PBS -l nodes=1:ppn=1
#PBS -l pmem=900mb
#PBS -l walltime=00:05:00
#PBS -q testq

cd $PBS_O_WORKDIR
date
module load perl
perl -version
```

GE

```
#!/bin/bash
#
#$ -N My_Job_Name
#$ -r n
#$ -M Jane_Doe@ufl.edu
#$ -m abe
#$ -o My_Job_Name.out
#$ -j y
#$ -l threaded 1
#$ -l h_vmem=900m
#$ -l h_rt=00:05:00
#$ -q test


date
module load perl
perl –version
```

Interactive sessions

PBS

- qsub -I

GE

- qlogin

PBS/GE Default Resources

PBS
(Main)

- Cores: 1
- Memory: 600M
- Walltime: 12:00:00

GE
(Biocluster)

- Cores: 1
- Memory: 8G
- Walltime: 24:00:00

PBS/GE Queues

- testq: queue for small and short test jobs

Automatic:

- investor: dedicated queue for investors
- other: job queue for non-investors

Array jobs

PBS

```
#!/bin/bash
#
#PBS -N array_example
#PBS -M om@hpc.ufl.edu
#PBS -m a
#PBS -j oe
#PBS -o
array_example_pbs.out
#PBS -l nodes=1:ppn=1
#PBS -l pmem=1gb
#PBS -l walltime=00:05:00
#PBS -t 1-10

cd $PBS_O_WORKDIR
date
echo "This is a task #
$PBS_ARRAYID." >
$PBS_JOBNAME.out
```

GE

```
#!/bin/bash
#
#$ -N array_test_ge
#$ -r n
#$ -M om@hpc.ufl.edu
#$ -m a
#$ -o array_ge.out
#$ -j y
#$ -pe threaded 1
#$ -l h_vmem=1g
#$ -l h_rt=00:05:00
#$ -t 1

date
echo "This is an array task #
$SGE_TASK_ID." > array_ge_
$SGE_TASK_ID.out
```

MPI jobs

PBS

```
#!/bin/bash
#
#PBS -N mpi_example
#PBS -M om@hpc.ufl.edu
#PBS -m abe
#PBS -o mpi_pbs.out
#PBS -e mpi_pbs.err
#PBS -l nodes=2:ppn=1
#PBS -l pmem=1gb
#PBS -l walltime=01:00:00
module load intel/11.1 openmpi
raxml
cd $PBS_O_WORKDIR
mkdir mpi_pbs
cd mpi_pbs
mpiexec raxmlHPC-MPI -m GTRGAMMA -
s ../primates.phy -f i -n boot -b
1234 -N 20 -p 12931
```

GE

```
#!/bin/bash
#
#$ -N mpi_example
#$ -M om@hpc.ufl.edu
#$ -m ae
#$ -o mpi_ge.out
#$ -e mpi_ge.err
#$ -j n
#$ -pe mpi 2
#$ -cwd

module load intel/11.1 openmpi
raxml

mkdir mpi_ge
cd mpi_ge
mpiexec raxmlHPC-MPI -m GTRGAMMA -
s ../primates.phy -f i -n boot -b
1234 -N 20 -p 12931
```

Job Management

- qsub: job submission
- qstat: check queue status
- qdel: job deletion
- Documentation
  - http://wiki.hpc.ufl.edu/index.php/Job_Submission_Queues
  - http://wiki.hpc.ufl.edu/index.php/BioCluster

How to get help

Help!!!

UF Research Computing/HPC Center

How to get help

- ✦ Asking for help
  - Support Request Tickets
    - http://support.hpc.ufl.edu
      - Use for everything - not just software bugs but for any questions or help requests
      - Searchable database of solutions
  - When you don't have access to web
    - support@hpc.ufl.edu
    - om@hpc.ufl.edu (Biological Support)
    - magitz@ufl.edu (Bio training and Q/A)

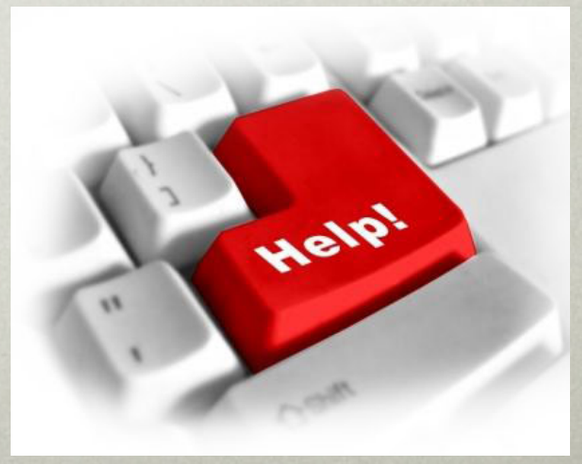

Documentation

✦ UF HPC Encyclopedia

- http://wiki.hpc.ufl.edu
  - Documents on hardware and software resources
  - User guides
  - Sample submission scripts
  - Research-specific sections
- http://hpc.ufl.edu/support
  - Frequently Asked Questions
  - Account set up and maintenance

HPC modules

Thank you!

UF Research Computing/HPC Center

Module Tree

- Root of the tree - /apps/lmod/modulefiles
  - Core – does not depend on a particular Compiler/MPI combination

/apps/lmod/modulefiles/core/clustalw/2.1.lua

  - Compiler – depends on a particular compiler

/apps/lmod/modulefiles/compiler/intel/10_1/mrbayes/3.1.2.lua

  - MPI – depends on a compiler/MPI combination

/apps/lmod/modulefiles/mpi/intel10_1/openmpi_1_2/raxml/1.0.5-light.lua