UL HPC School 2016

Closing Remarks

The UL HPC Team

Nov. 25th, 2016, MSA auditorium 3.330

University of Luxembourg (UL), Luxembourg

http://hpc.uni.lu

Main sources of information

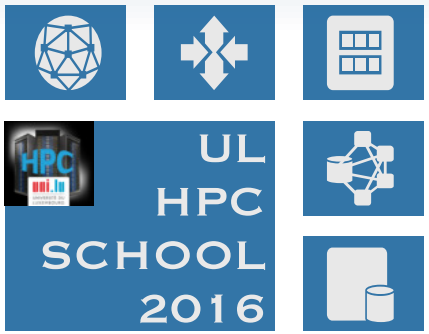

- **Thanks for attending!**
  - ↪ we hope you found it useful
  - ↪ Feedback / comments welcome
  - ↪ hpc-sysadmins@uni.lu
  - ↪ Improvement suggestion
- Something confusing? badly presented? topics not covered?
  - ↪ notify us and we will do our best to have it covered
  - ↪ format suggestion are also welcome

Looking forward meeting you next year, **March, 2017**

Fill the Survey!

`https://goo.gl/OflPyf`

A Few Take Away Messages

PC $\neq$ HPC $\neq$ Cloud

- does not mean they cannot work in synergy

Context	Local PC	HPC
Sequential	$T_1(\text{local}) = 100$	$T_1(\text{hpc}) = \mathbf{120s}$
Parallel/Distributed	$T_2(\text{local}) = 70\text{s}$	$T_2(\text{hpc}) = 80\text{s}$
		$T_8(\text{hpc}) = \mathbf{60s}$

- **Parallel/Distributed** runs **DO NOT COME FOR FREE**
  - ↪ runs **will be sequential** even if you reserve ≥ 2 cores/nodes
  - ↪ you have to **explicitly** adapt your jobs to benefit from the multi-cores/nodes

Other Take Away Tools

- SSH is your new friend (as the UL HPC platform)
  - ↪ as `vim`, `git`, `rsync`, `wget`, `make`...
- Always **check what you are doing on the platform!**
  - ↪ common pitfalls : **Out of memory**
  - ↪ tools for you: `htop`, Ganglia, `valgrind`, `ddt`, `map`...
  - ↪ good practice: **benchmark** your code
    - ✓ helps to anticipate on **appropriate walltime**
- **write launchers files**! `oarsub -S <launcher.sh>`
  - ↪ `-S` to interpret `#OAR` comments as default job options
  - ↪ take advantage of hierarchy of resources `-l nodes=N/core=C,walltime=H`

Pattern	Description
`enclosure=N`	number of enclosure
`nodes=N`	number of nodes
`core=N`	number of cores
`walltime=hh:mm:ss`	job's max duration

Main UL HPC commands

oarsub	submit/reserve a job	(by default: **1 core for 2 hours**)
oarstat	shows information about running or planned jobs	
module avail	list available modules	
module load	load a given module	
htop	a nicer top	

Main UL HPC resources

Reference `http://hpc.uni.lu`

Tutorials `http://github.com/ULHPC/tutorials`

Getting Help `hpc-users@uni.lu`

Typical Workflow on UL HPC resources

1. Connect to the frontend of a site/cluster `ssh`
2. *(eventually)* synchronize you code `scp/rsync/svn/git`
3. *(eventually)* Reserve a few interactive resources `oarsub -I`
   - ↪ *(eventually)* Configure the resources `kadeploy`
   - ↪ *(eventually)* Prepare your experiments `gcc/icc/mpicc/javac/...`
   - ↪ Test your experiment on small size problem `mpirun/java/bash....`
   - ↪ Free the resources
4. Reserve some resources `oarsub`
5. Run your experiment via a launcher script `bash/python/perl/ruby...`
6. Grab the results `scp/rsync`
7. Free the resources

Serial tasks: BAD and NAIVE approach

```
# Example 1: run in sequence $TASK 1...$TASK $NB_TASKS
for i in `seq 1 $NB_TASKS`; do
    $TASK $i
done
# Example 2: For each line of $ARG_TASK_FILE, run in sequence
#   $TASK <line1>... $TASK <lastline>
while read line; do
    $TASK $line
done < $ARG_TASK_FILE
```

Serial tasks: A better approach (fork & wait)

```
# Example 1: run in sequence $TASK 1...$TASK $NB_TASKS
for i in `seq 1 $NB_TASKS`; do
    $TASK $i &
done
wait

# Example 2: For each line of $ARG_TASK_FILE, run in sequence
#   $TASK <line1>... $TASK <lastline>
while read line; do
    $TASK $line &
done < $ARG_TASK_FILE
fi
wait
```

Serial tasks: A better approach (fork & wait)

Different runs may not take the same time: **load imbalance**.

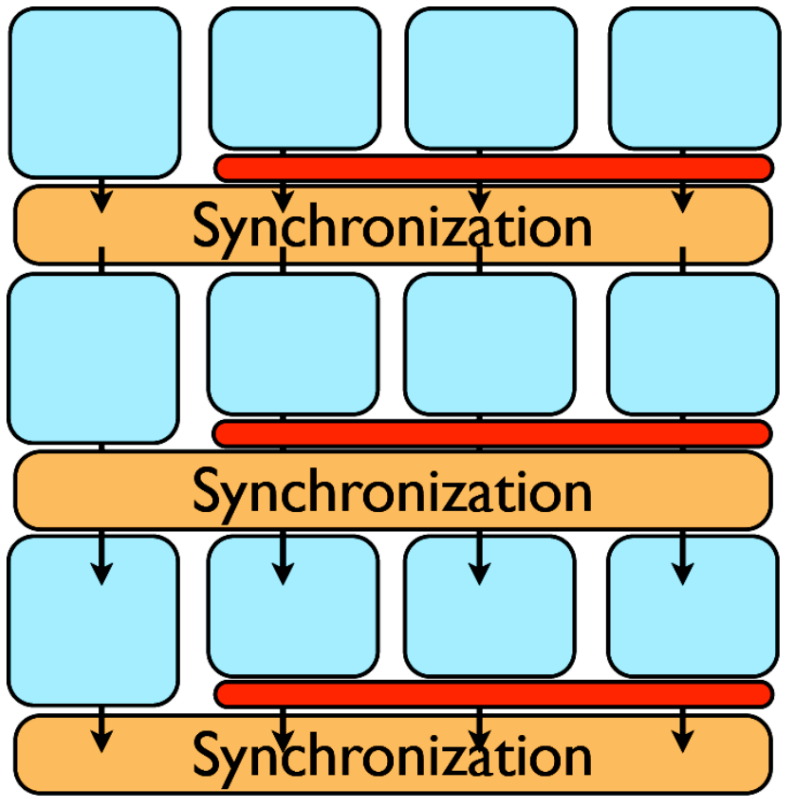

Serial tasks with GNU Parallel

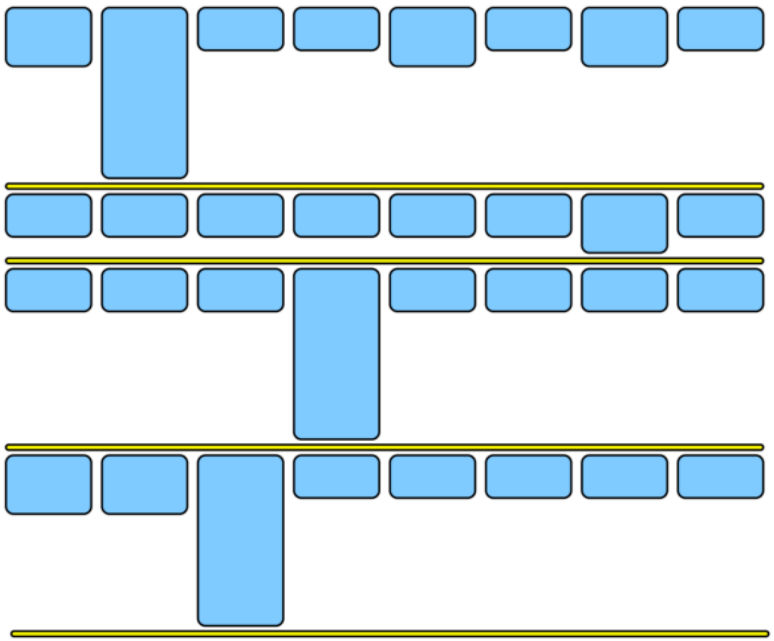

17 hours
42% utilization

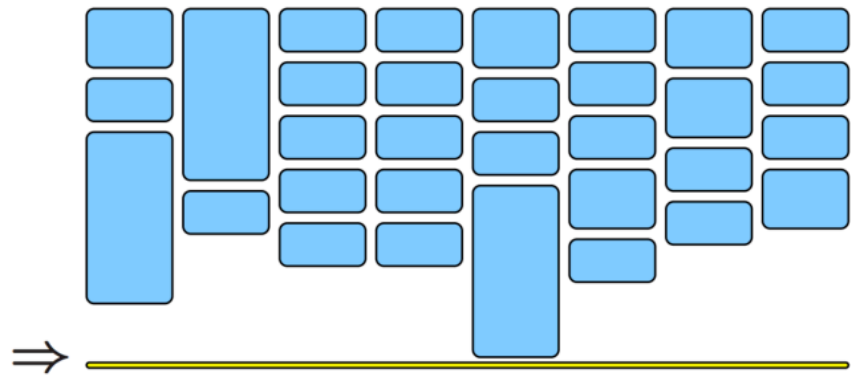

10 hours
72% utilization

Reporting your usage of the platform

https://hpc.uni.lu/users/AUP.html
https://hpc.uni.lu/about/publis.html

- In your scientific publications:
  - ↪ **acknowledge** your usage of the UL HPC platform
  - ↪ cf Acceptable Use Policy
  - ↪ **More importantly**: tag your publication with ULHPC!!

```
Acknowledgment: Experiments presented in this paper were carried out using the
HPC facilities of University of Luxembourg~\cite{VBCG_HPCS14}
{\small --- see \url{http://hpc.uni.lu}}.
```

```
@InProceedings{VBCG_HPCS14,
  author =       {S. Varrette and P. Bouvry and H. Cartiaux and F. Georgatos},
  title =        {Management of an Academic HPC Cluster: The UL Experience},
  booktitle =    {Proc. of the 2014 Intl. Conf. on High Performance Computing \& Simulation (HPCS 2014)},
  year =         {2014},
  pages =        {959--967},
  month =        {July},
  address =      {Bologna, Italy},
  publisher =    {IEEE},
}
```

Spread the world!

- Report the UL HPC to your colleagues / boss / students
- Help us to raise attention / funds!
  - ↪ if you're involved in projects preparation:
    - ✓ save budget for UL HPC usage
    - ✓ $\simeq$ **0.03 €** per CPUhour
- if you're involved in discussion with hierarchy
  - ↪ raise awareness / interest to the support of our platform

Reproducible Research Tutorial...

IEEE CloudCom 2016 - Dec 12th-15th, 2016

http://2016.cloudcom.org

- One of the top IEEE conference on Cloud Computing
  - world-class keynotes and technical papers
  - practical tutorials and business panels
- **Come and Register Now !**

Contacts:

Sebastien.Varrette@uni.lu
Pascal.Bouvry@uni.lu

FACULTY OF SCIENCES, TECHNOLOGY AND COMMUNICATION

1

Questions?

http://hpc.uni.lu

Sebastien Varrette
mail: sebastien.varrette@uni.lu
Office E-007
Campus Kirchberg
6, rue Coudenhove-Kalergi
L-1359 Luxembourg

UL HPC Management Team
mail: hpc-sysadmins@uni.lu