Numerical Quantum Physics – Foundations

S. Paeckel[1,2]

[1] Arnold Sommerfeld Center of Theoretical Physics, Department of Physics, University of Munich

[2] Munich Center for Quantum Science and Technology (MCQST)

April 16, 2023

Organization

- No lectures on: 2023-05-01, 2023-05-18, 2023-05-29, 2023-06-08
- Final examination: One week coding project, most probably 2023-07-24 – 2023-07-30 (more information later)
- Exercises:
  - Group assignment this week
  - Every second week new sheet
  - First sheet: 2023-04-21
  - Total of 50% points on exercise sheets yields +0.3 bonus for final examination (only if passed)
  - Exercise `git`: git@gitlab.physik.uni-muenchen.de:nqp/nqp-exercises.git
  - `HOME` directory for lecture: `/project/cip/2023-SS-NQP`
- LMU Jupyterhub
  - https://jupyter.physik.uni-muenchen.de
  - Requires two-factor authentification (for instance Google Authenticator)
- Lecture homepage:

Article

Quantum supremacy using a programmable superconducting processor

https://doi.org/10.1038/s41586-019-1666-5

Received: 22 July 2019

Accepted: 20 September 2019

Published online: 23 October 2019

Frank Arute[1], Kunal Arya[1], Ryan Babbush[1], Dave Bacon[1], Joseph C. Bardin[1,2], Rami Barends[1], Rupak Biswas[3], Sergio Boixo[1], Fernando G. S. L. Brandao[1,4], David A. Buell[1], Brian Burkett[1], Yu Chen[1], Zijun Chen[1], Ben Chiaro[5], Roberto Collins[1], William Courtney[1], Andrew Dunsworth[1], Edward Farhi[1], Brooks Foxen[1,5], Austin Fowler[1], Craig Gidney[1], Marissa Giustina[1], Rob Graff[1], Keith Guerin[1], Steve Habegger[1], Matthew P. Harrigan[1], Michael J. Hartmann[1,6], Alan Ho[1], Markus Hoffmann[1], Trent Huang[1], Travis S. Humble[7], Sergei V. Isakov[1], Evan Jeffrey[1], Zhang Jiang[1], Dvir Kafri[1], Kostyantyn Kechedzhi[1], Julian Kelly[1], Paul V. Klimov[1], Sergey Knysh[1], Alexander Korotkov[1,8], Fedor Kostritsa[1], David Landhuis[1], Mike Lindmark[1], Erik Lucero[1], Dmitry Lyakh[9], Salvatore Mandrà[3,10], Jarrod R. McClean[1], Matthew McEwen[5], Anthony Megrant[1], Xiao Mi[1], Kristel Michielsen[11,12], Masoud Mohseni[1], Josh Mutus[1], Ofer Naaman[1], Matthew Neeley[1], Charles Neill[1], Murphy Yuezhen Niu[1], Eric Ostby[1], Andre Petukhov[1], John C. Platt[1], Chris Quintana[1], Eleanor G. Rieffel[3], Pedram Roushan[1], Nicholas C. Rubin[1], Daniel Sank[1], Kevin J. Satzinger[1], Vadim Smelyanskiy[1], Kevin J. Sung[1,13], Matthew D. Trevithick[1], Amit Vainsencher[1], Benjamin Villalonga[1,14], Theodore White[1], Z. Jamie Yao[1], Ping Yeh[1], Adam Zalcman[1], Hartmut Neven[1] & John M. Martinis[1,5*]

Article

Quantum supremacy using a programmable superconducting processor

https://doi.org/10.1038/s41586-019-1666-5

Received: 22 July 2019

Accepted: 20 September 2019

Published online: 23 October 2019

Frank Arute[1], Kunal Arya[1], Ryan Babbush[1], Dave Bacon[1], Joseph C. Bardin[1,2], Rami Barends[1], Rupak Biswas[3], Sergio Boixo[1], Fernando G. S. L. Brandao[1,4], David A. Buell[1], Brian Burkett[1], Yu Chen[1], Zijun Chen[1], Ben Chiaro[5], Roberto Collins[1], William Courtney[1], Andrew Dunsworth[1], Edward Farhi[1], Brooks Foxen[1,5], Austin Fowler[1], Craig Gidney[1], Marissa Giustina[1], Rob Graff[1], Keith Guerin[1], Steve Habegger[1], Matthew P. Harrigan[1], Michael J. Hartmann[1,6], Alan Ho[1], Markus Hoffmann[1], Trent Huang[1], Travis S. Humble[7], Sergei V. Isakov[1], Evan Jeffrey[1], Zhang Jiang[1], Dvir Kafri[1], Kostyantyn Kechedzhi[1], Julian Kelly[1], Paul V. Klimov[1], Sergey Knysh[1], Alexander Korotkov[1,8], Fedor Kostritsa[1], David Landhuis[1], Mike Lindmark[1], Erik Lucero[1], Dmitry Lyakh[9], Salvatore Mandrà[3,10], Jarrod R. McClean[1], Matthew McEwen[5], Anthony Megrant[1], Xiao Mi[1], Kristel Michielsen[11,12], Masoud Mohseni[1], Josh Mutus[1], Ofer Naaman[1], Matthew Neeley[1], Charles Neill[1], Murphy Yuezhen Niu[1], Eric Ostby[1], Andre Petukhov[1], John C. Platt[1], Chris Quintana[1], Eleanor G. Rieffel[3], Pedram Roushan[1], Nicholas C. Rubin[1], Daniel Sank[1], Kevin J. Satzinger[1], Vadim Smelyanskiy[1], Kevin J. Sung[1,13], Matthew D. Trevithick[1], Amit Vainsencher[1], Benjamin Villalonga[1,14], Theodore White[1], Z. Jamie Yao[1], Ping Yeh[1], Adam Zalcman[1], Hartmut Neven[1] & John M. Martinis[1,5*]

- 69 authors
- 3 Supercomputing clusters
- Data analysis with external collaborators

Large Scale Numerics

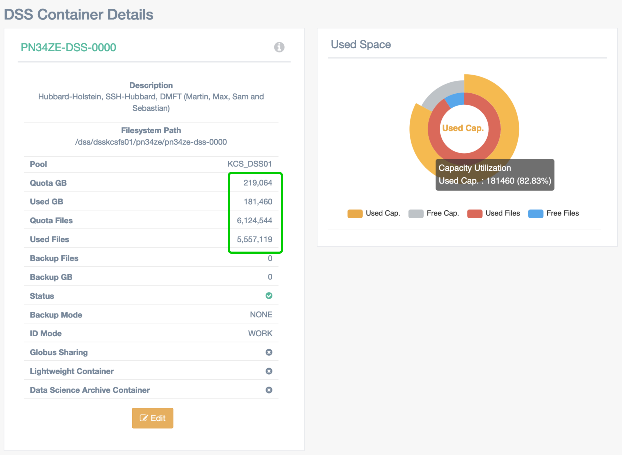

Exploration of the ground-state phase diagram of electron-phonon system:

- 180TB raw data
- 5M files of raw data
- 2 Supercomputing clusters

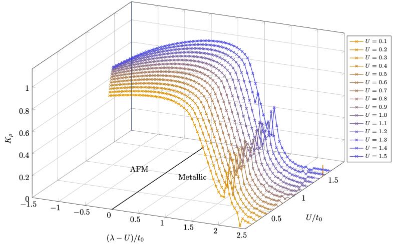

Analysis by Vincent Mourik (UNSW), Sergey Frolov (Pittsburgh)

Data manipulation and omission in **'Quantized Majorana conductance', Zhang et al, Nature 2018**

On November 24th 2019 we received from one of the authors of the now retracted 'Quantized Majorana Conductance' paper in Nature a pdf file with experimental notebook-quality data. Within this pdf file, we found data that appear to contradict the central claim of the paper. We have found that the original source experimental data may have been manipulated, namely cut, as well as cut out and pasted together. Furthermore, entire datasets that contradict the central claim of the Nature paper were suppressed. Our analysis was carried out between December 2019 and March 2020. We are able to publish our analysis now, in March 2021, after the authors of the Nature paper have finalized their retraction and deposited data from their experiment in full on Zenodo.

Central claim from the bold paragraph of the Nature paper:

recent observation[7] of a peak height close to $2e^2/h$. Here we report a quantized conductance plateau at $2e^2/h$ in the zero-bias conductance measured in indium antimonide semiconductor nanowires covered with an aluminium superconducting shell. The height of our zero-bias peak remains constant despite changing parameters such as the magnetic field and tunnel coupling, indicating that it is a quantized conductance plateau. We distinguish this quantized Majorana peak from possible non-Majorana origins by investigating its robustness

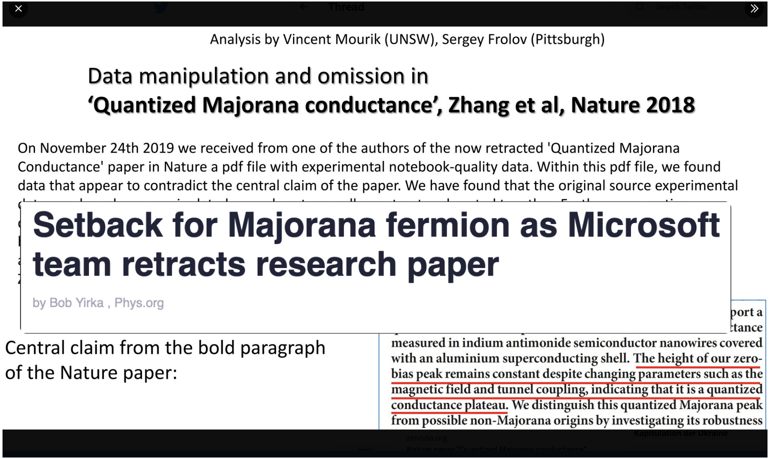

But things can go horribly wrong...
For peer review we need:

- Documentation
- Reproducibility
- Data consistency

And only after that comes the good story!

Large Scale Numerics

Reproducibility

- Distributed version control systems: Git, Tortoise, ...
- Software documentation: Doxygen, Github, Read the Docs, ...
- Collaborative editing: Gitlab, Overleaf, Nextcloud, Dropbox, ...
- Containerization: Singularity, Docker, ...

Documentation

Data consistency

Distributed version control systems

So how do you organize your project development?

- Control system
  - Content tracking
  - Backup functionality

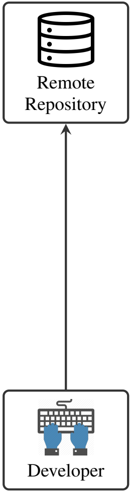

Distributed version control systems

So how do you organize your project development?

- Control system
  - Content tracking
  - Backup functionality
- Version control system
  - Structured project evolution
  - Roll-back/Push-forward functionality

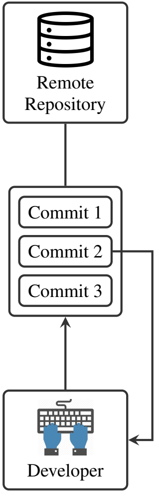

Distributed version control systems

So how do you organize your project development?

- Control system
  - Content tracking
  - Backup functionality
- Version control system
  - Structured project evolution
  - Roll-back/Push-forward functionality
- Distributed version control system
  - Parallel development branches
  - Branching/Merging from common project base
  - Issue fixing and task assignments

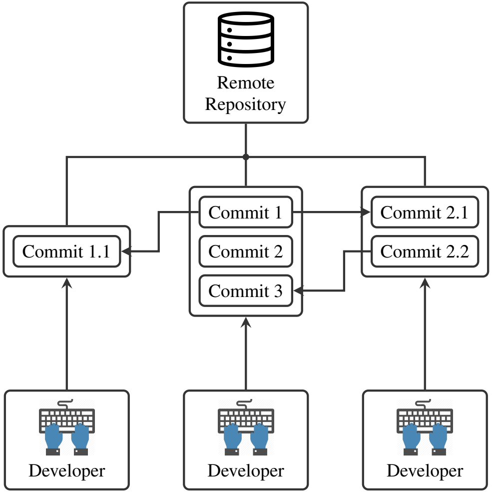

Distributed version control systems at the example of **Git**

So how do you organize your project development?

- Control system
  - `git add`, `git stash`
  - `git clone`, `git push`, `git pull`

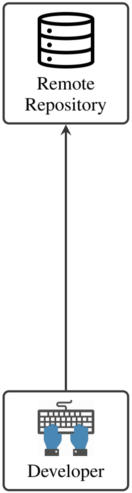

Distributed version control systems at the example of **Git**

So how do you organize your project development?

- Control system
  - `git add, git stash`
  - `git clone, git push, git pull`
- Version control system
  - `git commit`
  - `git checkout`

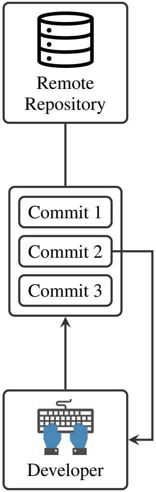

Distributed version control systems at the example of **Git**

So how do you organize your project development?

- Control system
  - `git add`, `git stash`
  - `git clone`, `git push`, `git pull`
- Version control system
  - `git commit`
  - `git checkout`
- Distributed version control system
  - `git branch`, `git merge`, `git cherry-pick`, `git reset`
  - `git log`

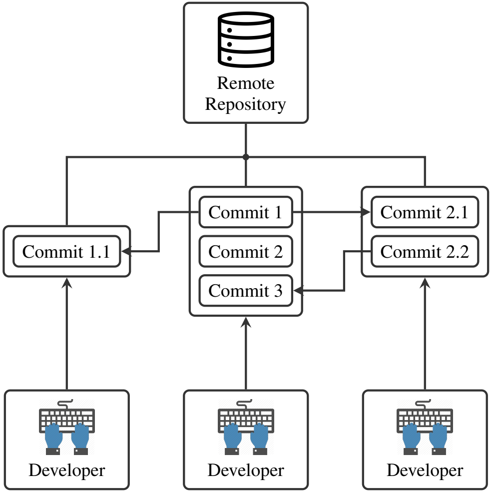

High Performance Computing (HPC): When do you really need a cluster?

Parallelization! But it's not all the same...

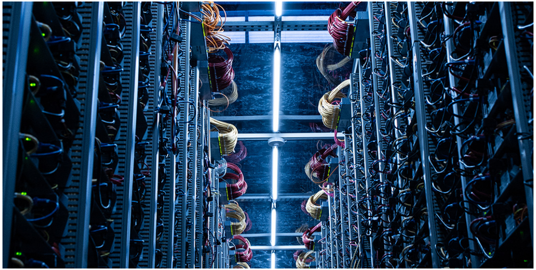

High Performance Computing (HPC): When do you really need a cluster?

Parallelization! But it's not all the same...

- OpenMP parallelization (single node computation)
  - Shared memory, dependent processes
  - Large RAM per CPU ratio

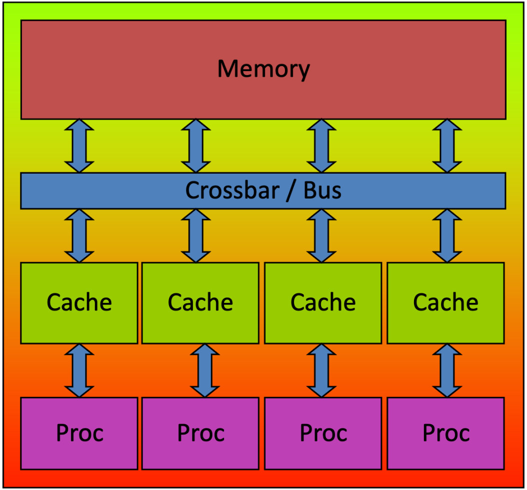

High Performance Computing (HPC): When do you really need a cluster?

Parallelization! But it's not all the same...

- OpenMP parallelization (single node computation)
  - Shared memory, dependent processes
  - Large RAM per CPU ratio
- MPI parallelization (multi node computation)
  - Independent processes with independent memory
  - Fast interconnect between nodes

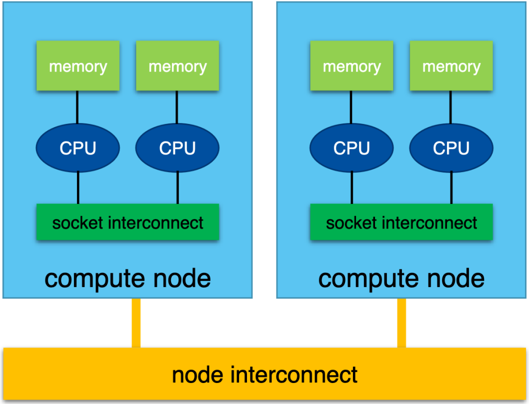

Large Scale Numerics

High Performance Computing (HPC): When do you really need a cluster?

Parallelization! But it's not all the same...

- OpenMP parallelization (single node computation)
  - Shared memory, dependent processes
  - Large RAM per CPU ratio
- MPI parallelization (multi node computation)
  - Independent processes with independent memory
  - Fast interconnect between nodes
- Architecture specific solutions
  - Graphics Processing Unit – cluster
  - Tensor Processing Unit – cluster

PRX QUANTUM **3,** 020331 (2022)

Editors' Suggestion

Simulation of Quantum Many-Body Dynamics with Tensor Processing Units: Floquet Prethermalization

Alan Morningstar,[1,2,*] Markus Hauru,[2] Jackson Beall,[2] Martin Ganahl,[2] Adam G.M. Lewis,[2] Vedika Khemani,[3] and Guifre Vidal[2]

[1] *Department of Physics, Princeton University, Princeton, New Jersey 08544, USA*
[2] *Sandbox@Alphabet, Mountain View, California 94043, USA*
[3] *Department of Physics, Stanford University, Stanford, California 94305, USA*

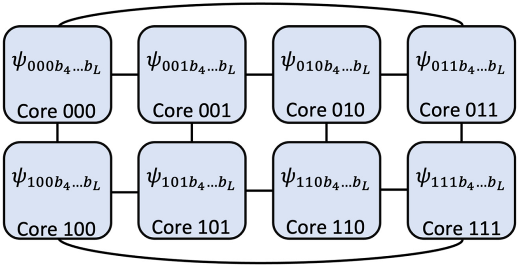

High Performance Computing (HPC): How to use cluster?

You have a code but how to use it?

- Write jobscript which performs desired computation
  - Wrap computation into a few commands
  - Specify hardware/software requirements, runtime, used resources, ...
  - Setup/Change job directories

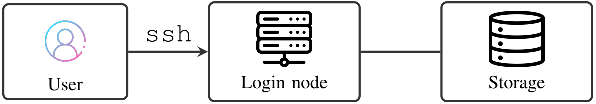

High Performance Computing (HPC): How to use cluster?

You have a code but how to use it?

- Write jobscript which performs desired computation
  - Wrap computation into a few commands
  - Specify hardware/software requirements, runtime, used resources, ...
  - Setup/Change job directories
- Submit jobscript to queueing system
  - Simple Linux Utility for Resource Management (SLURM)
  - Specify job dependencies
  - Wait ☕ ...

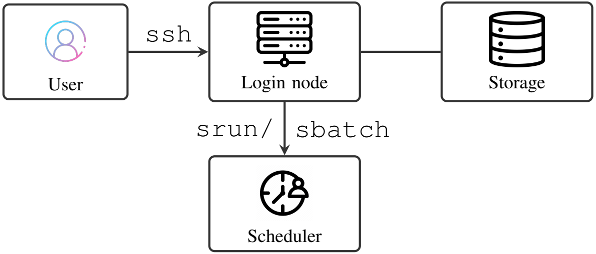

High Performance Computing (HPC): How to use cluster?

You have a code but how to use it?

- Write jobscript which performs desired computation
  - Wrap computation into a few commands
  - Specify hardware/software requirements, runtime, used resources, ...
  - Setup/Change job directories
- Submit jobscript to queueing system
  - Simple Linux Utility for Resource Management (SLURM)
  - Specify job dependencies
  - Wait ...
  - ...until ...finished
  - `sinfo, squeue, scontrol, ...`

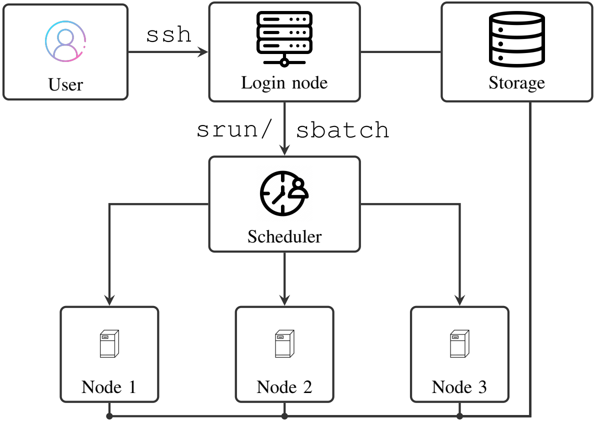

High Performance Computing (HPC): How to use cluster?

SLURM example for ASC cluster

- Open `SSH`-tunnel to login-node

```
Sebastian.Paeckel@th-nb-lswmbp38:~/Documents/Teaching/nqp-lecture/lectures/chapter_1$ ssh Sebastian.Paeckel@cip-sv-login.cip.physik.uni-muenchen.de
Warning: Permanently added the RSA host key for IP address '10.153.21.53' to the list of known hosts.
Sebastian.Paeckel@cip-sv-login.cip.physik.uni-muenchen.de's password:

 ***************************************************************
 *                                                             *
 *  IT - Faculty of Physics, LMU                               *
 *                                                             *
 *    For IT related problems, please consult our Webpage      *
 *      english: <http://www.en.it.physik.uni-muenchen.de/>    *
 *      german:  <http://www.it.physik.uni-muenchen.de/>       *
 *                                                             *
 *    Or contact your local Helpdesk                           *
 *       helpdesk-cip@physik.uni-muenchen.de                   *
 *                                                             *
 ***************************************************************

mkdir: cannot create directory '/scratch-local': Permission denied
mkdir: cannot create directory '/scratch-local': Permission denied
mkdir: cannot create directory '/scratch-local': Permission denied
mkdir: cannot create directory '/scratch-local': Permission denied
sebastian.paeckel@cip-sv-login03:~$
```

High Performance Computing (HPC): How to use cluster?

SLURM example for ASC cluster

- Open SSH-tunnel to login-node
- Setup SLURM jobscript
  - Specify SLURM parameter with preceeding #SBATCH
  - Repeated violation of job constraints reduces job priorities
  - Checkout available node types to avoid impossible constraints

```
#!/bin/bash

#SBATCH --export=ALL
#SBATCH --mail-type=ALL
#SBATCH --mail-user=sebastian.paeckel@physik.uni-muenchen.de
#SBATCH --partition=th-ws,th-cl,cluster,large
#SBATCH --constraint=avx2
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=1

#SBATCH --time=3-00:00:00
#SBATCH --mem=12G
##SBATCH --chdir=/project/th-scratch/s/Sebastian.Paeckel/
##SBATCH --output=/project/th-scratch/s/Sebastian.Paeckel/
##SBATCH --error=/project/th-scratch/s/Sebastian.Paeckel/

echo Hello world from $(hostname)
```

High Performance Computing (HPC): How to use cluster?

SLURM example for ASC cluster

- Open SSH-tunnel to login-node
- Setup SLURM jobscript
  - Specify SLURM parameter with preceeding #SBATCH
  - Repeated violation of job constraints reduces job priorities
  - Checkout available node types to avoid impossible constraints
- Submit job to queue using
  - sbatch for non-interactive jobs
  - srun/salloc for interactive jobs
- Watch job status using squeue

```
Sebastian.Paeckel@th-nb-lswmbp38:~/Documents/Teaching/nqp-lecture/lectures/chapter_1$ ssh Sebastian.Paeckel@cip-sv-login.cip.physik.uni-muenchen.de
Warning: Permanently added the RSA host key for IP address '10.153.21.53' to the list of known hosts.
Sebastian.Paeckel@cip-sv-login.cip.physik.uni-muenchen.de's password:

 **************************************************************
 *                                                            *
 *  IT - Faculty of Physics, LMU                              *
 *                                                            *
 *    For IT related problems, please consult our Webpage     *
 *      english: <http://www.en.it.physik.uni-muenchen.de/>   *
 *      german:  <http://www.it.physik.uni-muenchen.de/>      *
 *                                                            *
 *    Or contact your local Helpdesk                          *
 *       helpdesk-cip@physik.uni-muenchen.de                  *
 *                                                            *
 **************************************************************

mkdir: cannot create directory '/scratch-local': Permission denied
mkdir: cannot create directory '/scratch-local': Permission denied
mkdir: cannot create directory '/scratch-local': Permission denied
mkdir: cannot create directory '/scratch-local': Permission denied
sebastian.paeckel@cip-sv-login03:~$ sbatch -J test-job jobscript.slurm; squeue -u Sebastian.Paeckel
Submitted batch job 3631991
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
           3631983       cip  jupyter sebastia  R      39:47      1 cip-cl-compute7
           3631991 large,th- test-job sebastia PD       0:00      1 (Priority)
sebastian.paeckel@cip-sv-login03:~$
```

High Performance Computing (HPC): How to use cluster?

File systems at ASC cluster

- Always check your `quota`!
- If home directory is full → account de-facto dead (with all your data)!

```
sebastian.paeckel@cip-sv-login02:~$ quota
--------------------------------------------------------------------------------------------------------------------
                    path       type |  SPACE: used       soft       hard |  FILES: used       soft       hard
--------------------------------------------------------------------------------------------------------------------
            /project/cip     lustre |       2.4 Gi   300.0 Gi   400.0 Gi |        37043      75000     100000
                   /home        nfs |       5.1 Gi          ~    40.0 Gi |            ~          ~          ~
```

High Performance Computing (HPC): How to use cluster?

File systems at ASC cluster

- Always check your `quota`!
- If home directory is full → account de-facto dead (with all your data)!
- Virtual drives
  - always available
  - smaller bandwidth $\sim 10 - 100$MBit/s
- Local drives
  - only available on current machine
  - large bandwidth > 1GBit/s
- Backups are important but expensive if number of files is large (file limit)

Mount	Speed	Availability	Total Size	def. Quota	Backup	Period of storage
/home	**	**network wide (see details)**	8 TB	40GB	snapshot & tape	unlimited
/project/theorie	**	**network wide**	30 TB	100 GB	snapshot & tape	unlimited
/scratch	***	**network wide**	67 TB	500 GB	snapshot	unlimited
/scratch-local	****	**local**	depending	no	no	Deletion after 30 days
/tmp	*****	**local**	depending	no	no	Deletion after 7 days or the next reboot

Working on problem sheets using the LMU Jupyterhub

LMU Jupyterhub provides a simple user interface for interactive SLURM jobs at the ASC cluster

- Before login set up two-factor authentification

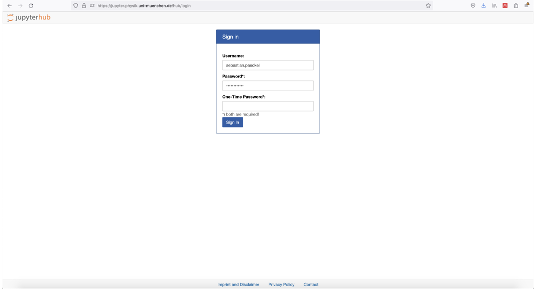

Working on problem sheets using the LMU Jupyterhub

LMU Jupyterhub provides a simple user interface for interactive SLURM jobs at the ASC cluster

- Before login set up two-factor authentification
- Setup job specifics
  - Choose cluster
  - Specify hardware: CPUs, RAM
  - Choose max. runtime

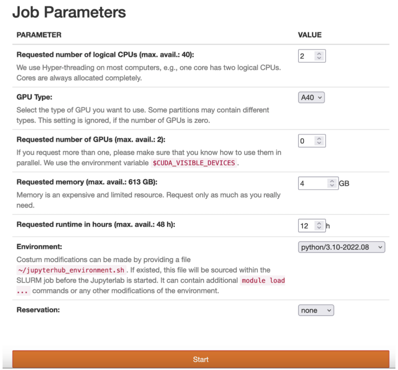

Working on problem sheets using the LMU Jupyterhub

LMU Jupyterhub provides a simple user interface for interactive SLURM jobs at the ASC cluster

- Before login set up two-factor authentification
- Setup job specifics
  - Choose cluster
  - Specify hardware: CPUs, RAM
  - Choose max. runtime

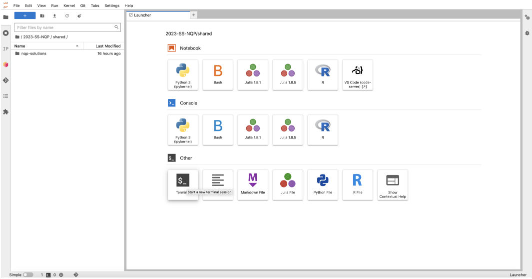

- In terminal: initialize environment
  - First login: Create symbolic link to lecture's `home` directory:
    `ln -s /project/cip/2023-SS-NQP/ ~/2023-SS-NQP`
  - Load default modules: `source ~/2023-SS-NQP/init_modules.sh`