Huygens Professional and Core on HPC Bunya Guide

R. Amor

QBI Advanced Microscopy Facility

Queensland Brain Institute, The University of Queensland

Research Lane, St Lucia, 4072 QLD, AUSTRALIA

<r.amor@uq.edu.au>

December 13, 2024

Abstract

A brief guide on how to use Huygens Professional and Core on HPC Bunya, including running Huygens Core by submitting a job to the Simple Linux Utility for Resource Management (SLURM) scheduler.

Preamble

This is HPC Bunya:
https://rcc.uq.edu.au/systems/high-performance-computing/bunya.

If you haven't done so already, apply for a Bunya account here:
https://services.qriscloud.org.au/services/request/new/ee6def64259741a095c1fed20743e3fb.
Copy questions 2) - 8) into the "Request details" box on the form and provide details for each item. Incomplete applications will be rejected and applicants will be required to fill in a new form.

Bunya training is available here:
https://rcc.uq.edu.au/training-support/training-courses.
Register via email to rcc-support@uq.edu.au to attend UQ-only training sessions.

Read the Bunya user guide here:
https://github.com/UQ-RCC/hpc-docs/blob/main/guides/Bunya-User-Guide.md
and the Bunya OnDemand guide here:
https://github.com/UQ-RCC/hpc-docs/blob/main/guides/OnDemand-Guide.md.

You can then request access to Huygens Professional and Core on HPC Bunya using this link:
https://services.qriscloud.org.au/access/6211d1e3482544b2b73fc85156446e2c/member.

Running Huygens Professional and Core via OnDemand

1. Go to the OnDemand site:
   https://bunya-ondemand.rcc.uq.edu.au/

2. Navigate to Interactive Apps > GPU-Accelerated Desktop. Give the job a name, choose your Bunya account group, and specify the number of GPUs, CPU cores per task, number of tasks, and maximum running time. Click "Launch."

3. OnDemand will prompt: "Please be patient as your job currently sits in queue. The wait time depends on the number of cores as well as time requested."
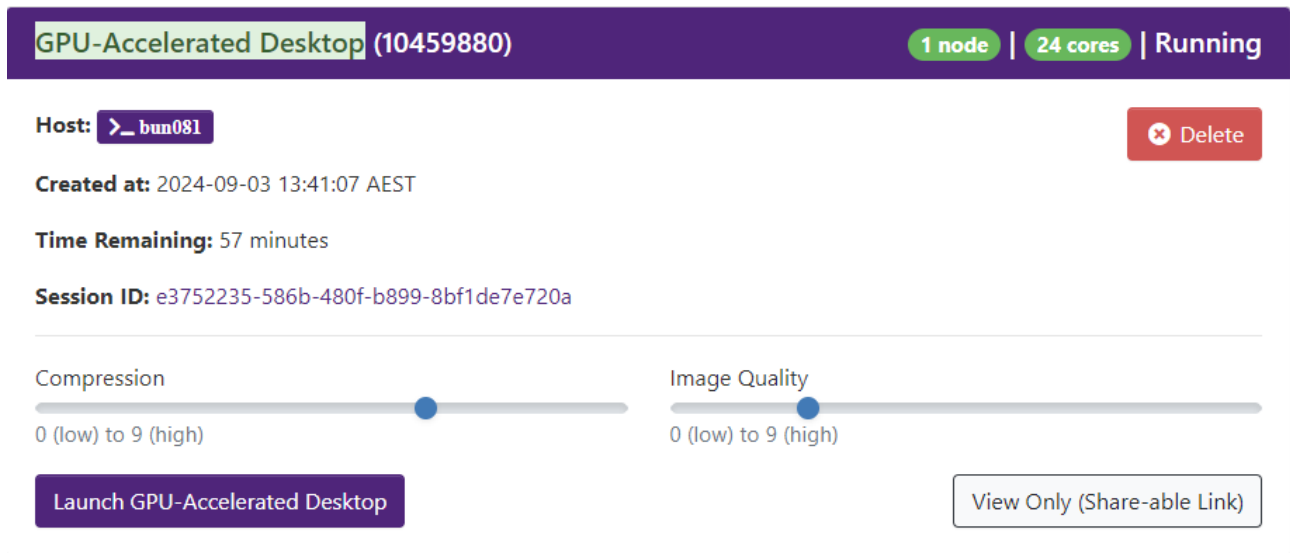   Click "Launch GPU-Accelerated Desktop" when the desktop is running (Fig. 1).

Figure 1: HPC Bunya OnDemand GPU-accelerated desktop running.

4. On the visualisation desktop, open a MATE terminal by navigating to Applications > System Tools > MATE Terminal and type:
```
module load huygens
huygenspro
```
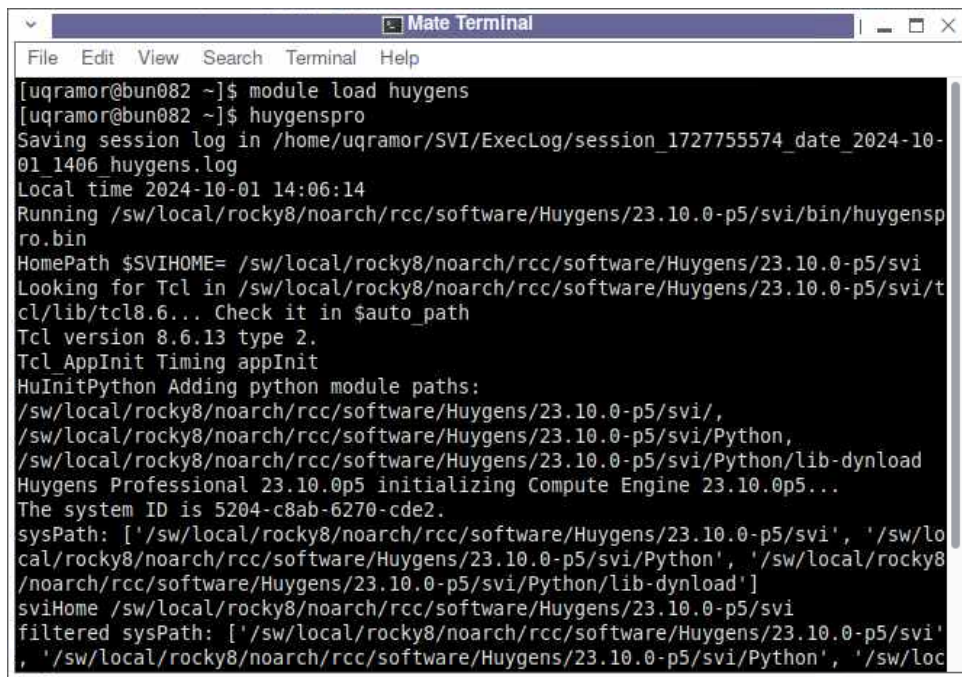(Fig. 2).

Figure 2: Running Huygens Professional via the MATE terminal.

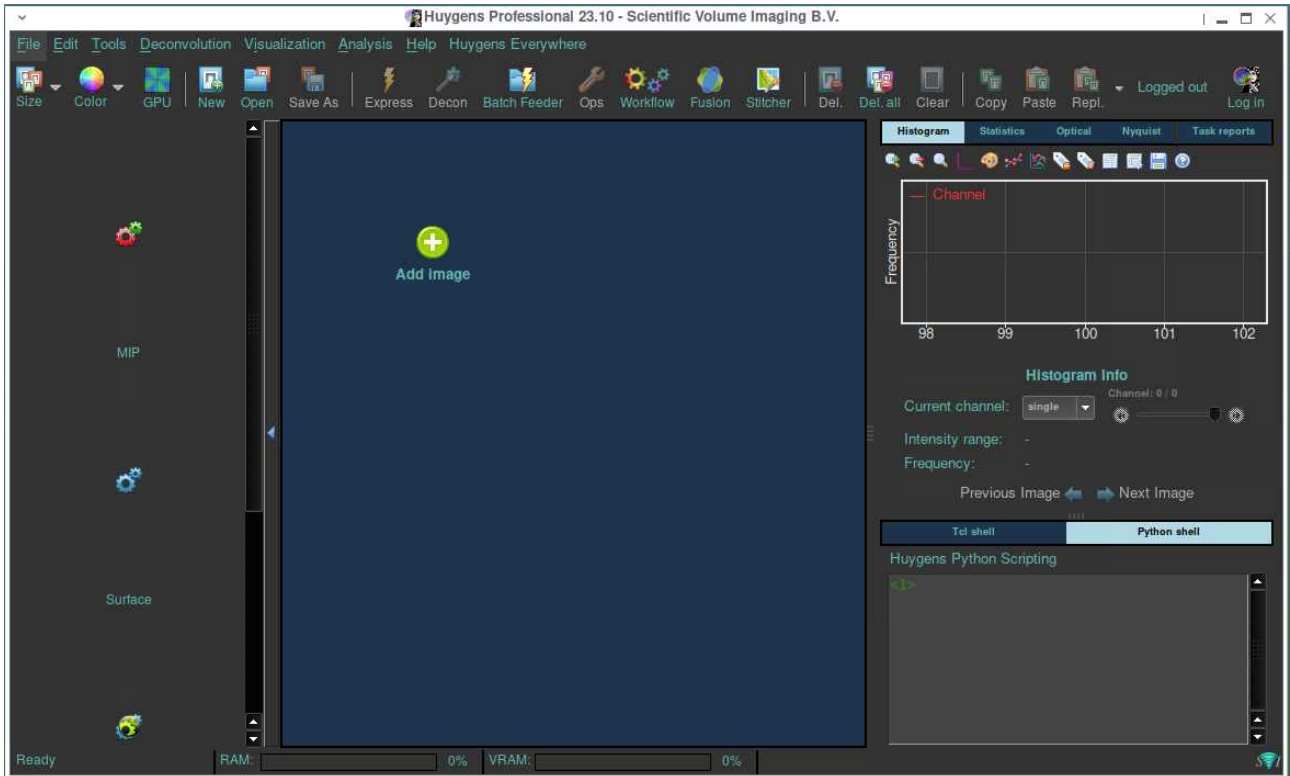5. This launches the Huygens Professional graphical user interface (GUI, Fig. 3).

Figure 3: The Huygens Professional GUI.
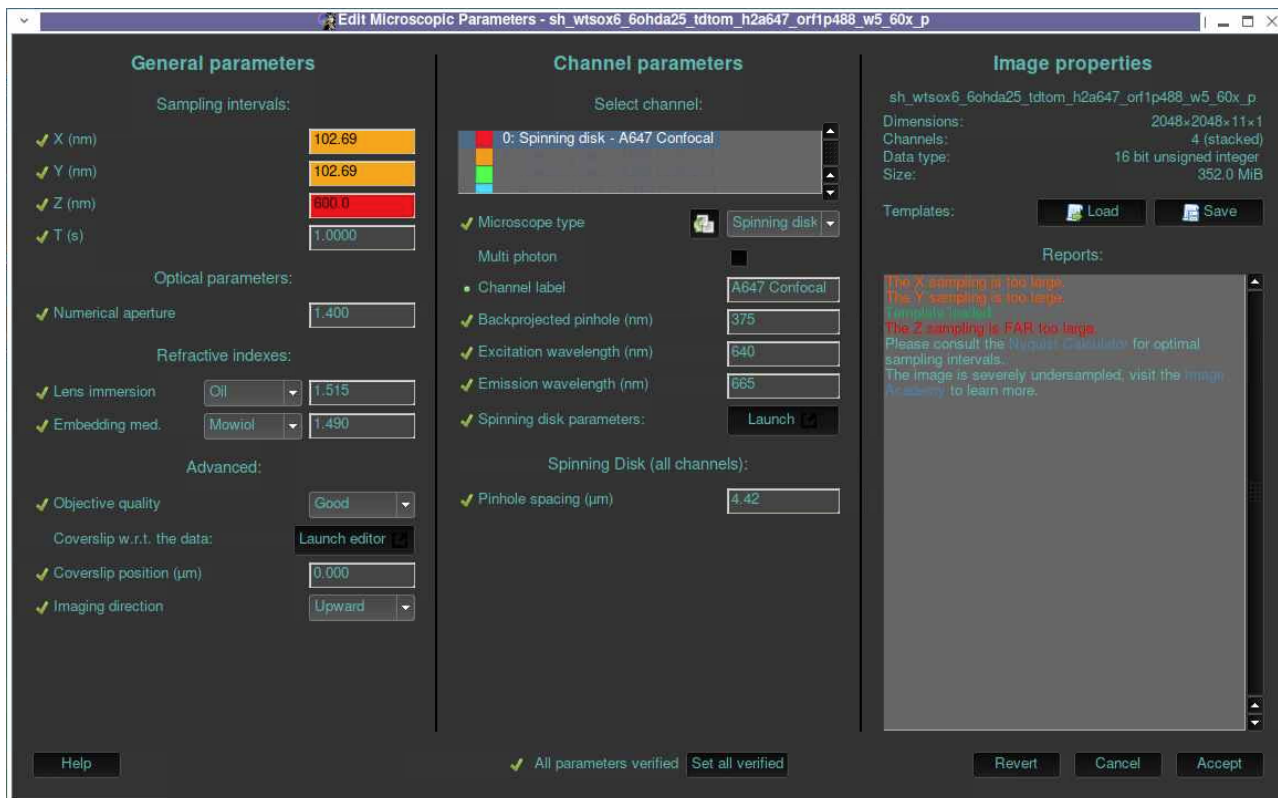
6. Open an image and create a microscope template as normal (Fig. 4).

Figure 4: Creating a microscope template.

7. Do a deconvolution run using the wizard and save the deconvolution template as normal (Fig. 5).
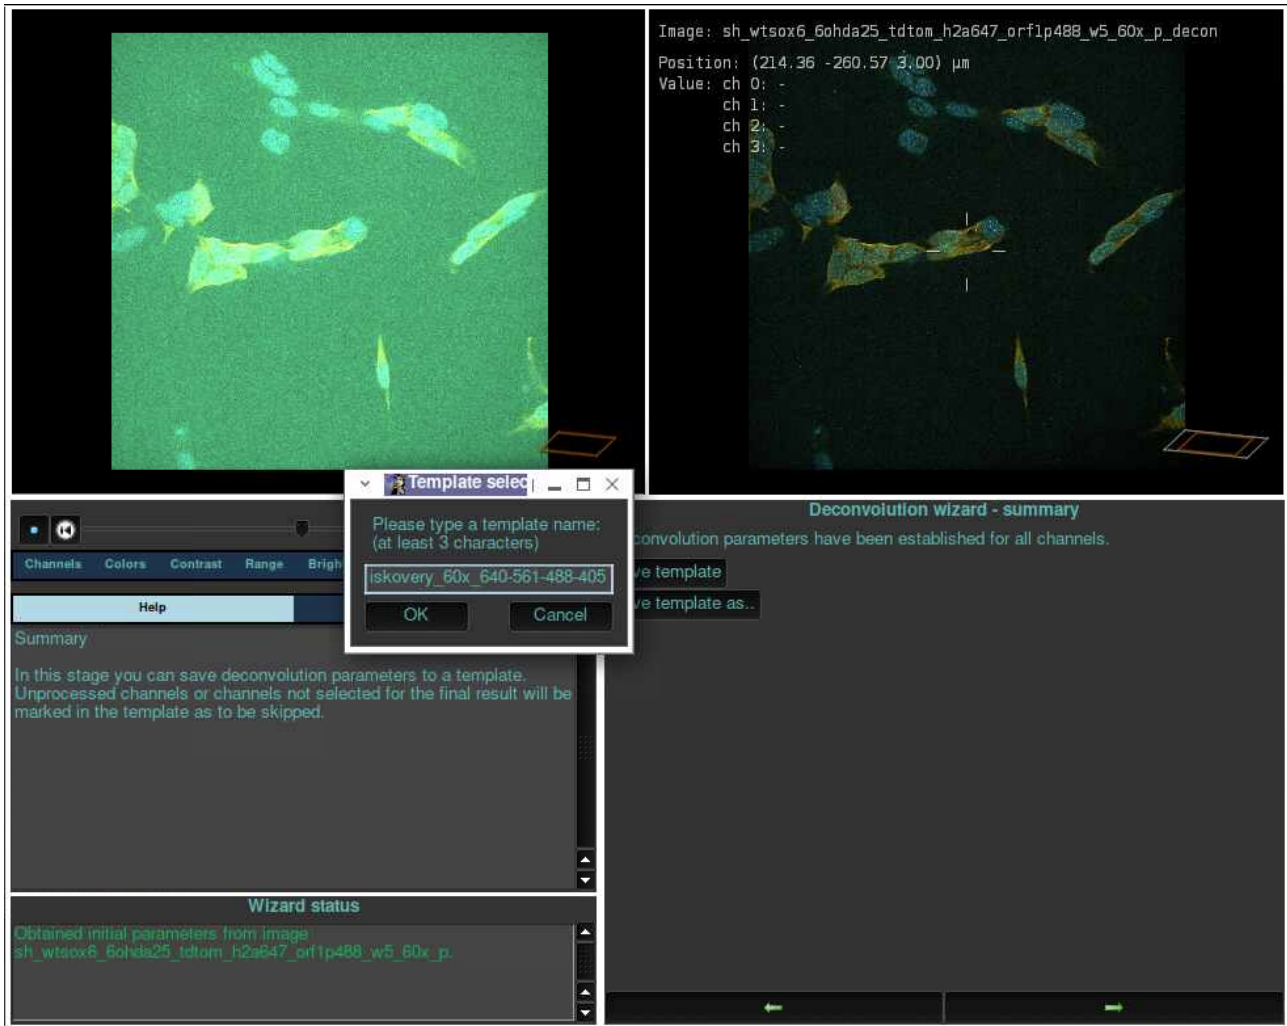
Figure 5: Saving a deconvolution template.

8. Open the Workflow Processor and create a batch job, specifying the input images or folder, the microscope template, deconvolution template, and output file format and directory (Fig. 6).
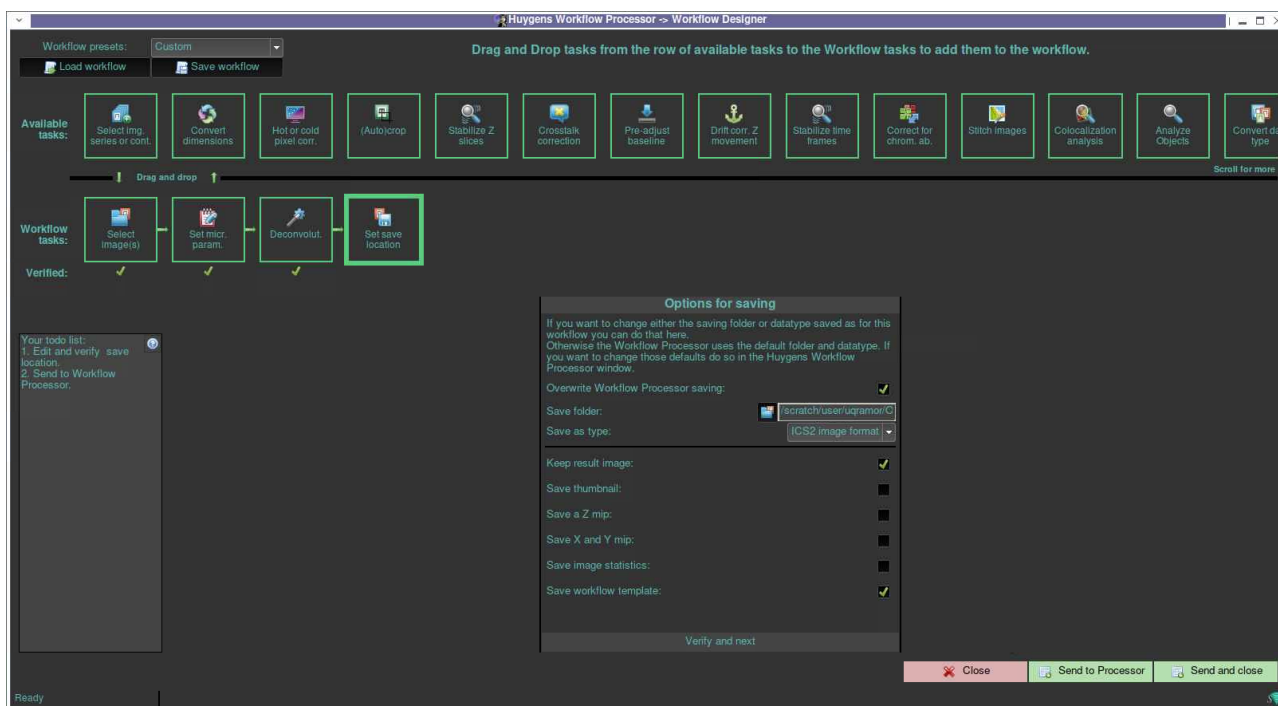
Figure 6: Creating a batch job on the Workflow Processor.

9. Save the batch template (`.hgsb` file, Fig. 7). You will use this template file to run batch jobs with Huygens Core, on the visualisation desktop and on SLURM.
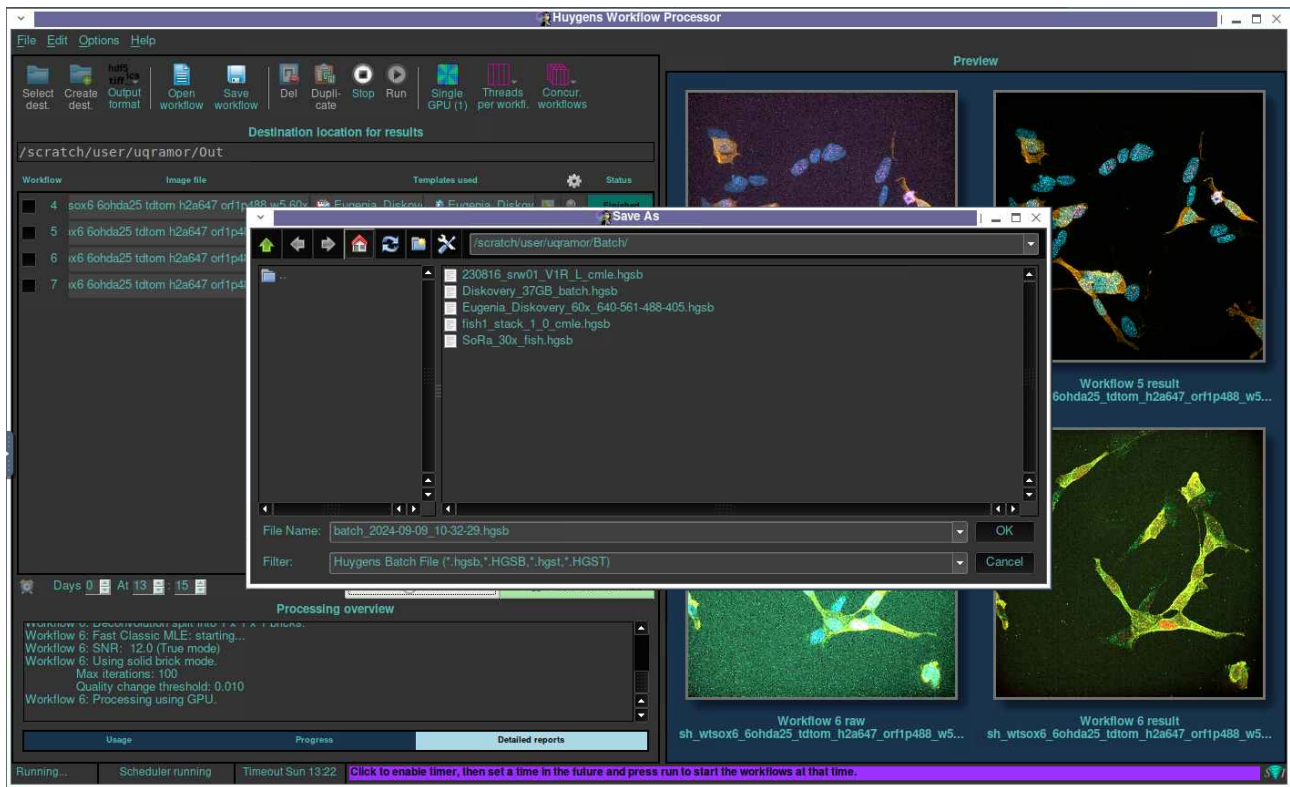
Figure 7: Saving a batch template.

10. Run Huygens Core by opening a new terminal and typing:
    `module load huygens`
    `hucore -template ''batch_template_location''`
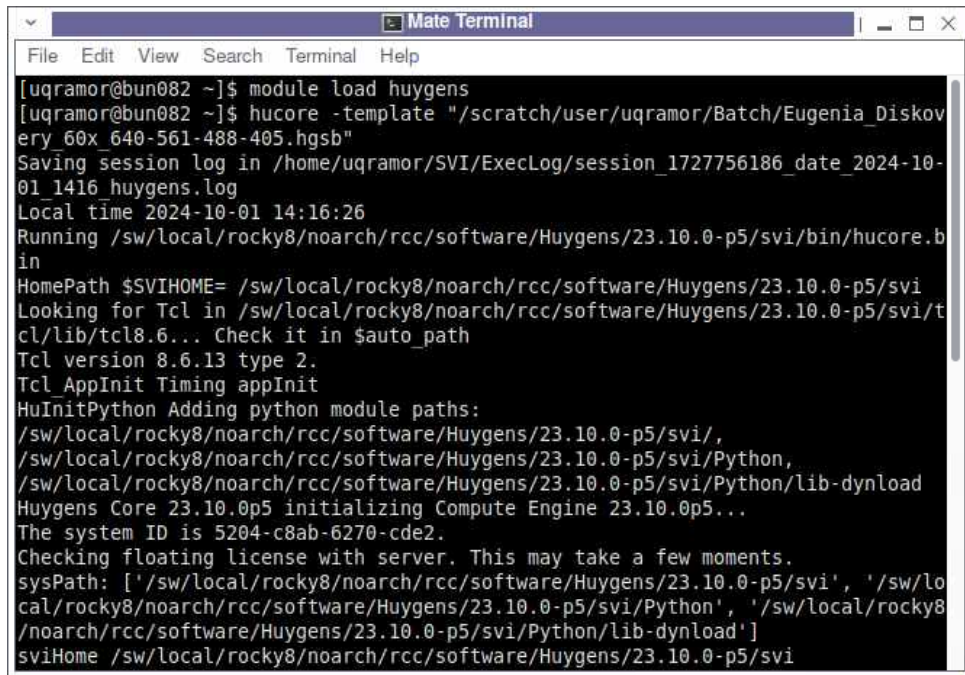
    for example,

    `hucore -template ''/scratch/user/uqramor/Batch/Eugenia_Diskovery_60x_640-561-488-405.hgsb''`
    (Fig. 8).

    Huygens Core will then run the batch job and save the outputs in the directory specified on the batch template.

Figure 8: Running Huygens Core via OnDemand.

Huygens SLURM on Bunya

1. On your machine, open a terminal window. Log on to Bunya by typing:
   `ssh USERID@bunya.rcc.uq.edu.au`

   Enter your UQ password. Go through the Duo authentication prompts. You will then be logged on to Bunya (Fig. 9).

Figure 9: Connecting to Bunya via ssh on a terminal window.
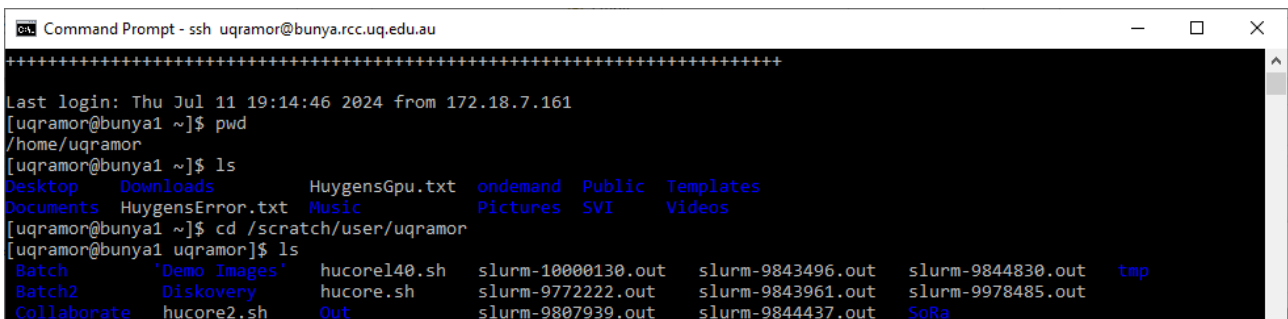
2. By default, you will be on the /home directory. Print the current working directory:
   `$ pwd`

   (Note: "$" indicates a command line and is not part of the command.)

   UQ RCC requires that all jobs are called from the /scratch directory. Change to your /scratch directory:
   `$ cd /scratch/user/USERID`
   (Fig. 10).

Figure 10: Printing the current working directory and changing to the scratch directory.

3. Create a script to submit your job to the SLURM scheduler. Use vi, or a similar editor. Give the script a name, and use .sh as the extension (Fig. 11).

   Watch QCIF's quick overview of vi here:
   https://www.youtube.com/watch?v=TuR9d9Z_Fis

```
[uqramor@bunya1 uqramor]$ cat hucore2.sh
#!/bin/bash -l

#SBATCH --nodes=1
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=24
#SBATCH --mem=16G
#SBATCH --job-name=HUCORE
#SBATCH --time=01:00:00
#SBATCH --qos=gpu
#SBATCH --partition=gpu_cuda
#SBATCH --constraint=cuda48gb
#SBATCH --gres=gpu:1
#SBATCH --account=a_qbi_microscopy

module load huygens
hucore -template "/scratch/user/uqramor/Batch/Eugenia_Diskovery_60x_640-561-488-405.hgsb"
[uqramor@bunya1 uqramor]$
```

Figure 11: An example script to submit a Huygens Core job to SLURM.

A short description of the lines on the script follows. The Bunya User Guide has more information.

`#!/bin/bash -l`
→ We are using BASH (the Bourne Again SHell), a shell program and command language for Linux. The `-l` command makes the shell a login shell.

The `SBATCH` commands are for allocating resources:
`#SBATCH --nodes=1` → number of nodes

`#SBATCH --ntasks-per-node=1` → 1 for single- and multi-thread jobs

`#SBATCH --cpus-per-task=24` → number of threads

`#SBATCH --mem=16G` → RAM per job in MB, GB, TB; see the Bunya guide

`#SBATCH --job-name=HUCORE` → the name you give your job

`#SBATCH --time=01:00:00` → time the job needs to complete; max=168 hrs for gpu_cuda

`#SBATCH --qos=gpu` → the Quality of Service name; see the QoS guide

`#SBATCH --partition=gpu_cuda` → you are using the GPU partition

`#SBATCH --constraint=cuda48gb` → the GPU RAM specified, in aiming to use a particular GPU; for an A100 MIG slice, use "`--constraint=cuda10gb`", for H100 and A100, use "`--constraint=cuda80gb`", or "`--constraint=cuda48gb`"

`#SBATCH --gres=gpu:1` → GPU number

`#SBATCH --account=a_qbi_microscopy` → your accounting group

The next two lines launch Huygens Core for batch runs:
`module load huygens` → launches the Huygens module

`hucore -template ''.hgsb batch file''` → launches Huygens Core and runs the batch job as specified in the template

4. Submit the job to SLURM:
   `$ sbatch SCRIPT_FILENAME`

```
[uqramor@bunya1 uqramor]$ sbatch hucore2.sh
Submitted batch job 10002720
```

Figure 12: Submitting the job to SLURM.

5. Check the status in the job queue:
   ```
   $ squeue --me
   ```

```
[uqramor@bunya1 uqramor]$ squeue --me
         JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
      11014598  gpu_cuda   HUCORE  uqramor PD       0:00      1 (Priority)
      11013736   gpu_viz Viz_Desk  uqramor  R      45:39      1 bun082
      11013730   gpu_viz Viz_Desk  uqramor  R      45:56      1 bun122
[uqramor@bunya1 uqramor]$
```

Figure 13: Checking the status in the queue.

6. Monitor the job by using the `tail` command.
   ```
   $ tail -f slurm-JOBID.out
   ```
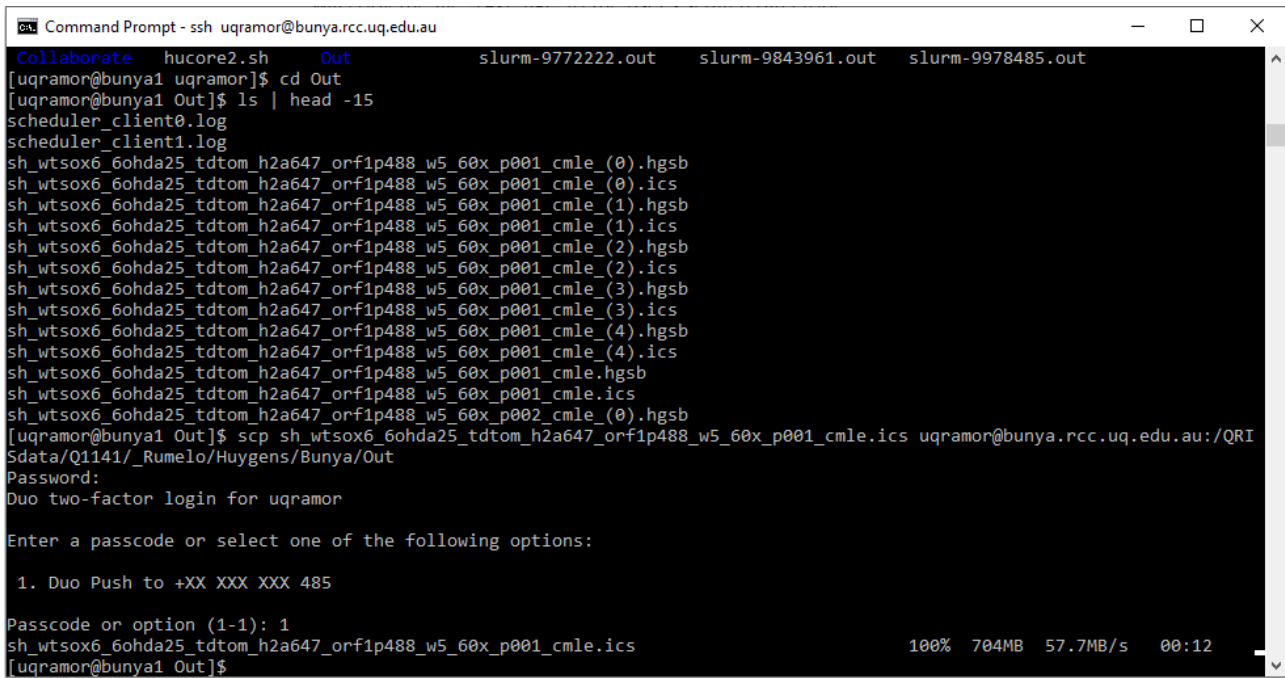
   To exit the `tail` command, press Ctrl + C.

```
[uqramor@bunya1 uqramor]$ tail -f slurm-11014598.out
z dim: 1.000
t dim: 1.000
Doing task imgSave...
Scaling mode while saving: -cmode scale
Scaling mode while saving: -cmode scale
Saving image using command: sh_wtsox6_6ohda25_tdtom_h2a647_orf1p488_w5_60x_p003_cmle__9_ save {/scratch/user/uqramor/Out
/sh_wtsox6_6ohda25_tdtom_h2a647_orf1p488_w5_60x_p003_cmle_(9).ics} -type ics2  -cmode scale
Successfully saved results to: {/scratch/user/uqramor/Out/sh_wtsox6_6ohda25_tdtom_h2a647_orf1p488_w5_60x_p003_cmle_(9).i
cs}
Template /scratch/user/uqramor/Batch/Eugenia_Diskovery_60x_640-561-488-405.hgsb successfully processed
Tcl_AppInit App init time 209.50
Local time 2024-10-01 15:12:27
```

Figure 14: Monitoring the job.

7. When the job is done, use scp (**s**ecure **cop**y) to transfer individual files to your RDM collection:
    ```
    $ scp FILENAME USERID@bunya.rcc.uq.edu.au:/RDM_FILE_PATH
    ```


Figure 15: Transferring a file using scp.

or copy the whole output directory:
```
$ scp -r OUTPUT_DIRECTORY USERID@bunya.rcc.uq.edu.au:/RDM_FILE_PATH
```


Figure 16: Transferring the output directory using scp.

Acknowledgments