

RMINC on HPF

RMINC is now ready for use on the HPF cluster. This provides a means to rapidly run computationally intensive code like mincLmer on 100's of cores.

Step-by-step guide

1. Prior to logging in to the cluster you should organize your files such that HPF has access to them. So move your files over to /hpf if they are not already there.
2. Log in to the cluster

```
ssh -AX hpf.ccm.sickkids.ca
```

3. Start a larger than usual qlogin session

```
qlogin -X -l vmem=20G,walltime=40:00:00
```

4. Then you are ready to load up the required modules

```
module use /hpf/largeprojects/MICe/tools/modulefiles && \  
module load mice-env
```

5. Then you can start an R session as usual and load up RMINC

```
R  
library(RMINC)
```

6. To run an example mincLmer from Darren and Lily's work you can run the following

```
df <- read.csv("/hpf/largeprojects/MICe/dfernandes/lqiu_long/all_relative_jacobians/files.csv")  
mask <- "/hpf/largeprojects/MICe/dfernandes/lqiu_long/test_p65dc2/mask/mask_dimorder_eroded.mnc"  
  
vslmer <- mincLmer(files~sex*age+(1|ID), parallel = c("pbs", 100), data=df, mask=mask)
```

This example should run in 10-20 minutes. Note the parallel argument, this is where you set the number of jobs to run. In this case it is 100. The element ("pbs") is now checked only to see if it is "local", local tells RMINC to run the job on multiple cores of a single machine. Anything else defers to our underlying job scheduling package BatchJobs. See https://rawgit.com/Mouse-Imaging-Centre/RMINC/master/inst/documentation/RMINC_Parallelism.html for more details, as well as a complete list of functions that can be parallelized in this way. Darren has also written a page on how to [parallelize voxelwise statistics](#).