PARALLEL R

prof. dr. Janez Povh
EuroHPC Competence Center, May 292024

Outline/next

EURO

- Introduction to HPC
- Introduction to R
- Parallel R within one node
- Parallelization with Rmpi

Introduction to HPC

What is HPC

EURO

Source: https://www.vyzkumne-infrastruktury.cz/en/2022/06/lumi-supercomputer-has-been-inaugurated/

What is HPC

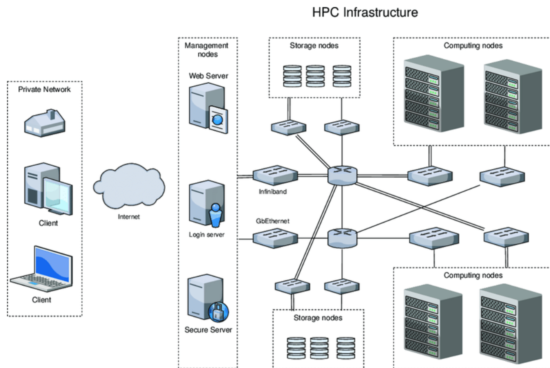

Source: Reghenzani, F. et al, IEEE Access, 8, 208566-208582.

Node structure

EURO

2560x compute nodes

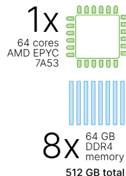

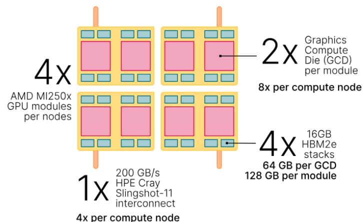

Source: https://docs.lumi-supercomputer.eu/hardware/lumig/

User experience

login.vega.izum.si - PuTTY

```
Using username "jpovh".
Authenticating with public key "rsa-key-20220609"
Further authentication required
Keyboard-interactive authentication prompts from server:
| Verification code:
```

How to access HPC

EURO

HPCs are easily available.

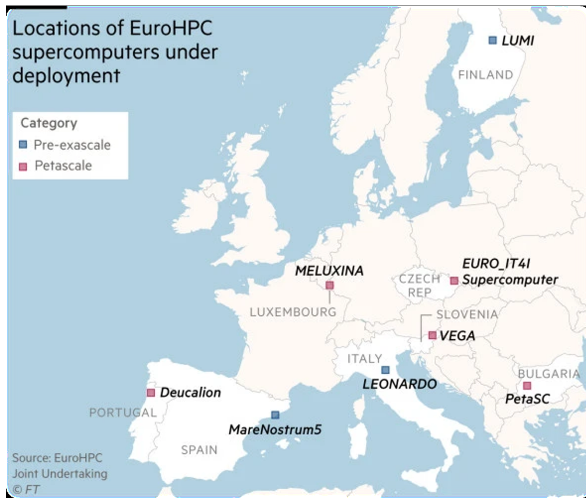

EuroCC2

EURO

- Network of EU Competence centers for HPC.
- All members of EuroHPC JU involved.
- Training, support for industry, talent attraction,...

Outline/next

EURO

- Introduction to HPC
- Introduction to R
- Parallel R within one node
- Parallelization with Rmpi

Introduction to R

What is R

- Software for Statistical Data Analysis
- Based on S
- Programming Environment
- Interpreted Language
- Data Storage, Analysis, Visualization
- Free and Open Source Software

How to obtain R

- R current version 4.4.0 (released April 2024).
- `http://cran.r-project.org`
- Binary/Windows executable code

Pros and Cons

Pros:

- Free and Open Source
- Strong User Community
- Highly extensible, flexible
- Implementation of high-end statistical methods
- Flexible graphics and intelligent defaults

Cons

- Steep learning curve
- Slow for large datasets

Data types

- R Supports virtually any type of data
- Numbers, characters, logicals (TRUE/ FALSE)
- Arrays of virtually unlimited sizes
- Simplest: Vectors and Matrices
- Lists: Can Contain mixed type variables
- Data Frame: Rectangular Data Set

Data structures in R

Linear

- vectors (all same type)
- lists (mixed types)

Rectangular

- data frame
- matrix

Running R

- I recommend RStudio, an IDE for R.
- It is available as RStudio Desktop and **RStudio Server**, which runs on a remote server and allows accessing RStudio using a web browser.

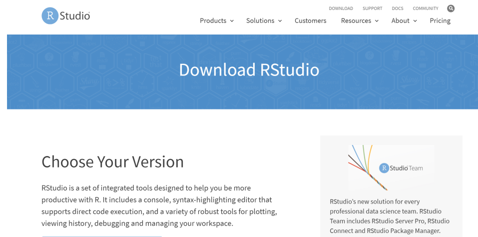

Figure 1: https://rstudio.com/products/rstudio/download/

RStudio on IT4I

- Run RStudio on VM.
- Connect to `shiny.vsb.cz/auth`

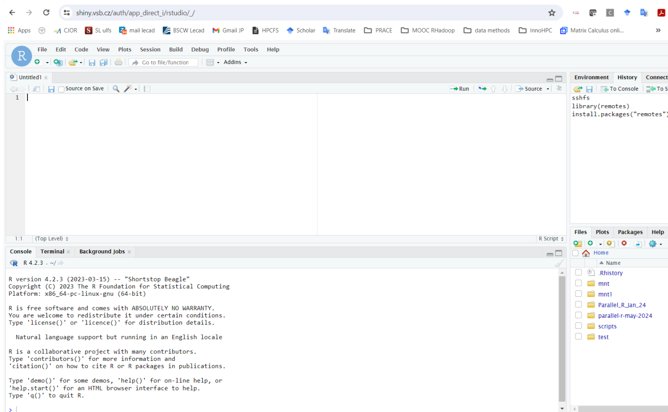

Clone the data

EURO

Clone project from GITHUB

https://github.com/It4innovations/parallel-r-may-2024

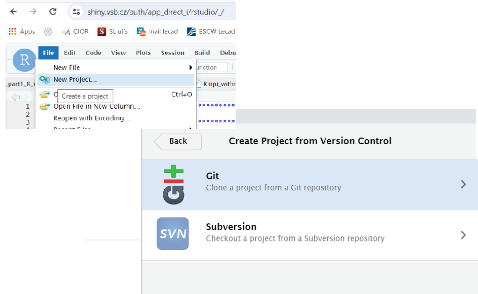

Creating the first script file

Create and save simple data file

```
N=1000;
Data=data.frame(group=character(N),ints=numeric(N),reals=numeric(N))
Data$group=sample(c("a","b","c"), 1000, replace=TRUE);
Data$ints=rbinom(N,10,0.5);
Data$reals=rnorm(N);

head(Data)
Data

write.table(Data, file='data/Data_Ex_1.txt', append = FALSE, dec = ".",col.names =
 TRUE)

ls()
rm(list = ls())
```

Load and analyse the data

EURO

Load data

```
Data_read<-read.table(file='data/Data_Ex_1.txt',header = TRUE)
# first few rows
head(Data_read)
#10 th row
Data_read[10,]
# column group
Data_read$group
Data_read[,1]
```

Analyse the data

EURO

Simple analysis

```
# compute means and counts by groups
group count_ints mean_ints
a | 337 | 5.014837
b | 338 | 5.032544
c | 325 | 4.990769

# primitive solution
Group_lev=levels(Data_read$group)

Tab_summary=data.frame(group=character(3),count_ints=integer(3),mean_ints=numeric(3))
Tab_summary$group<-Group_lev
for (i in c(1:3)){
  sub_data = subset(Data_read,group==Group_lev[i])
  Tab_summary$count_ints[i]<-nrow(sub_data)
  Tab_summary$mean_ints[i]<-mean(sub_data$ints)
}
```

Analyze the data by split, aggregate, sapply

split, aggregate, sapply

```
s <- split(Data_read, Data_read$group)
Tab_summary1<-t(sapply(s, function(x) return(c(mean(x$ints),length(x$group)) )))

Tab_summary2<-cbind(aggregate(ints~group,data = Data_read,FUN=length),aggregate(ints~
  group,data = Data_read,FUN=mean))
Tab_summary2<-Tab_summary2[,-3]
```

Outline/next

EURO

- Introduction to HPC
- Introduction to R
- Parallel R within one node
- Parallelization with Rmpi

Parallel R within one node

What system do I have

How many cores

```
library(parallel)
detectCores()
> detectCores()
[1] 20
```

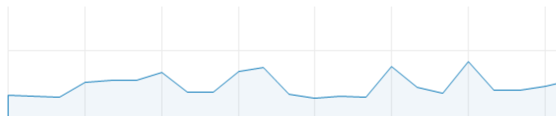

60 seconds

Utilization	Speed	
5%	1,88 GHz	
Processes	Threads	Handles
346	5406	1761616
Up time		
5:17:11:23		

Base speed:	2,40 GHz
Sockets:	1
Cores:	14
Logical processors:	20
Virtualization:	Enabled
L1 cache:	1,2 MB
L2 cache:	11,5 MB
L3 cache:	24,0 MB

apply, lapply, sapply

apply, lapply, sapply

```
apply(X, MARGIN, FUN)
Here:
-x: an array or matrix
-MARGIN=1: the manipulation is performed on rows
-MARGIN=2: the manipulation is performed on columns
-MARGIN=c(1,2): the manipulation is performed on rows and columns
-FUN: tells which function to apply. Built functions like mean, median, sum, min, max and
      even
user-defined functions can be applied
```

apply

For data constructed above (Data_read) compute row and columns means using apply

apply

```
Data_read<-read.table(file='data/Data_Ex_1.txt',header = TRUE)

Data_col_means_1 <- colMeans(Data_read[,-1])
Data_col_means_2 <- apply(Data_read[,-1],2,FUN =mean)

Data_row_means_1 <- rowMeans(Data_read[,-1])
Data_row_means_2 <- apply(Data_read[,-1],1,FUN =mean)

Data_both_squares <- apply(Data_read[,-1],c(1,2),FUN = function(x) return(x^2))
```

lapply

- `lapply` function takes list, vector or data frame as input and returns only list as output
- `sapply` function takes list, vector or data frame as input. It is similar to lapply function but returns only vector as output.

For data constructed above (`Data_read`) compute row and columns sums using `lapply`

lapply

```
Data_col_sums_1 <- apply(Data_read[,-1],2,FUN =sum)
Data_col_sums_2 <- lapply(Data_read[,-1],FUN =sum)

typeof(Data_col_sums_1)
typeof(Data_col_sums_2)


Data_abs <- lapply(Data_read[,-1],FUN =abs)
Data_sq <- lapply(Data_read[,-1],FUN = function(x){x^2})

typeof(Data_abs)
length(Data_abs)
```

sapply

For data constructed above (Data_read) compute row and columns sums using sapply

sapply

```
Data_col_sums_1 <- apply(Data_read[,-1],2,FUN =sum)
Data_col_sums_2 <- lapply(Data_read[,-1],FUN =sum)
Data_col_sums_3 <- sapply(Data_read[,-1],FUN =sum)

typeof(Data_col_sums_1)
typeof(Data_col_sums_2)
typeof(Data_col_sums_3)


Data_col_sums_4 <- lapply(list(Data_read$ints,Data_read$reals),FUN =sum)
Data_col_sums_5 <- sapply(list(Data_read$ints,Data_read$reals),FUN =sum)
Data_col_len_1 <- lapply(list(Data_read$ints,Data_read$reals),FUN =length)
Data_col_len_2 <- sapply(list(Data_read$ints,Data_read$reals),FUN =length)
```

for loop

Let us compute sums of all elements of K random matrices of order $N \times N$

for

```
N=1000
K=60
set.seed(2021)
sum_rand=rep(0,K-1);
tic()
time_for_sys=system.time({
  for (i in c(1:K)){
    A=rand(N,N)
    sum_rand[i]=sum(A)
  }
})
time_for=toc()
```

foreach do loop

Let us compute sums of all elements of K random matrices of order $N \times N$

for

```
set.seed(2021)
sum_rand=rep(0,K-1);
tic()
time_foreach_sys=system.time({
  foreach (i = c(1:K)) %do% {
    A=rand(N,N)
    sum_rand[i]=sum(A)
  }
})
time_foreach=toc()
```

Libraries parallel, doParallel

- `parallel` package comes in the base R installation
- `parallel` works great for any task that you pass to the apply family (e.g., lapply becomes parLapply).
- `doParallel` package works great when you want to use parallel variant of for-loops (foreach -do), and might be a little easier to use.

doParallel

- This library is meant for use with `foreach`, which lets you use a particular type of for-loop, that looks like:

```
foreach(i=list_of_elements) %do% {thing with i}.
```

- Foreach allows this to be parallelized, using `dopar`:

```
foreach(i=listOfThings) %dopar% {thing with i}.
```

- Note that: parallelization with `dopar` depends on which backend you use.
  - `doParallel` is one such backend - it tells `foreach` to use `parallel`.
  - There are others: `doFuture`, `doMPI` (another parallel backend, using message passing interface), `doSnow` (another backend, using the `snow` package for creating parallel processes),...
  - By default, `doParallel` uses multicore functionality on Unix-like systems and snow functionality on Windows.

Parallel foreach-dopar loop

EURO

Let us compute sums of all elements of K random matrices of order $N \times N$ using `foreach ...dopar` from `foreach` and `doParallel`

for

```
N=3000
set.seed(2021)
sum_rand=rep(0,11);
tic()
foreach (i = c(1:12)) %dopar% {
  A=randn(N,N)
  sum_rand[i]=sum(A)
}
time_foreach_dopar=toc()
```

Do you observe any difference?

Creating cluster with doParallel

- Option 1: (use doParallel)

Create cluster

```
clust <- makeCluster(n_cores-1)
registerDoParallel(clust)
getDoParName()
.
.
stopCluster(clust)
#registerDoSEQ()  # alternative - register sequential mode
```

Creating cluster with doParallel

- Option 2: use parallel

Create cluster

```
registerDoParallel(cores=n_cores-1)
getDoParName()
.
.
registerDoSEQ()  # alternative - register sequential mode
```

- In linux: the first option uses snow library, second multicore library. In windows: both use snow.

Parallel foreach dopar loop - option 1

Let us compute sums of all elements of K random matrices of order $N \times N$ using `foreach` ...`dopar` from `foreach`, `doParallel`. Create cluster!

Option 1

```
set.seed(2021)
clust <- makeCluster(n_cores-1)
registerDoParallel(clust)  # use multicore, set to the number of our cores - needed for
     foerach dopar
getDoParName()
sum_rand=rep(0,K-1);
tic()
time_foreachdopar_1_sys=system.time({
  print("for each-dopar (cluster allocated)")
  foreach (i = c(1:K)) %dopar% {
    library(pracma)
    A=rand(N)
    sum_rand[i]=sum(A)
}}
)
time_foreach_dopar_1=toc()
stopCluster(clust)
```

Parallel foreach dopar loop - option 2

Let us compute sums of all elements of K random matrices of order $N \times N$ using `foreach ...dopar` from `foreach`, `doParallel`. Create cluster!

Option 2

```
set.seed(2021)
registerDoParallel(n_cores-1)  # use multicore, set to the number of our cores - needed
     for foerach dopar
getDoParName()
sum_rand=rep(0,K-1);
tic()
time_foreachdopar_2_sys=system.time({
  print("for each-dopar (cluster allocated)")
  foreach (i = c(1:K)) %dopar% {
    library(pracma)
    A=rand(N)
    sum_rand[i]=sum(A)
}}
)
time_foreach_dopar_1=toc()
registerDoSEQ()   #this registers sequential mode - equivalent
```

Timings

Timings

	user.self	sys.self	elapsed
time_for_sys	4.16	0.50	4.93
time_for_each_sys	4.21	0.37	4.92
time_for_each_dopar_sys	4.11	0.55	4.98
time_for_each_dopar_1_sys	0.13	0.01	1.90
time_for_each_dopar_2_sys	0.11	0.00	1.80

Library parallel

- encapsulates existing libraries `multicore`, `snow`
- `multicore` functionality supports multiple workers only on those operating systems that support the **fork** system call - this excludes Windows.
- two ways of parallelization:
  - The socket approach: launches a new version of R on each core via networking (e.g. the same as if you connected to a remote server), but the connection is happening all on your own computer.
    - pros: (i) Works on any system (including Windows); (ii) Each process on each node is unique so it can't cross-contaminate.
    - cons: (i) Each process is unique so it will be slower (ii) Things such as package loading need to be done in each process separately. Variables defined on your main version of R don't exist on each core unless explicitly placed there. (iii) More complicated to implement.
  - use `parLapply`, `parSapply`

Library `parallel`

- The forking approach copies the entire current version of R and moves it to a new core.
  - (i) Faster than sockets. (ii) Because it copies the existing version of R, your entire workspace exists in each process. (iii) Easy to implement.
  - Cons (i) Only works on POSIX systems (Mac, Linux, Unix, BSD) and not Windows. (ii) it can cause issues specifically with random number generation or when running in a GUI (such as RStudio). This doesn't come up often.
- use `mclapply`

Parallel versions of lapply

By using library parallel and parSapply, mclapply compute sums of all elements of K random matrices of order $N \times N$. Create cluster!

parallel versions of apply

```
mat_sum<-function(x){
  library(pracma)
  A=rand(x)
  return(sum(A))
}

time_lapply<-system.time({
  set.seed(2021)
  sum_rand_lapply=lapply(rep(N,K),FUN=mat_sum)
})

time_sapply<-system.time({
  set.seed(2021)
  sum_rand_sapply=sapply(rep(N,K),FUN=mat_sum)
})
```

Parallel versions of lapply

parallel versions of apply

```
#forking
time_mcLapply<-system.time({
  set.seed(2021)
  sum_rand_mcLapply=mclapply(X=rep(N,K),FUN=mat_sum,mc.cores = n_cores)
})

# socketing
clust <- makeCluster(n_cores, type="PSOCK")
time_parLapply<-system.time({
  set.seed(2021)
  sum_rand_parLapply=parLapply(clust,rep(N,K),fun=mat_sum)
})
stopCluster(clust)

clust <- makeCluster(n_cores, type="PSOCK")
time_parSapply<-system.time({
  set.seed(2021)
  sum_rand_parSapply=parSapply(clust,rep(N,K),FUN=mat_sum)
})
stopCluster(clust)
```

Parallel versions of lapply

parallel versions of apply

```
    times_apply<-rbind(time_lapply,time_sapply,time_parLapply,time_parSapply,time_
     mcLapply)

    > times_apply[,1:3]
                   user.self  sys.self elapsed
time_lapply        1.741      0.011    1.751
time_sapply        1.726      0.007    1.731
time_parLapply     0.007      0.004    1.940
time_parSapply     0.005      0.005    1.842
time_mcLapply      0.004      0.238    1.679
```

Libraries for shared memory parallelization in R

- Parallel for-loop (`foreach...dopar`). Cluster created by `registerDoParallel(N)` and `registerDoSEQ()`. Library `foreach`, `doParalel` needed.
- Parallel apply: `parLapply`, `parSapply`, `mcLapply` need library parallel.

Very paralleizable task

Perfectly paralelizable computing task

```
# simple very parallel
library(parallel)
library(tictoc)

f <- function(...) {
  Sys.sleep(1)
  "DONE"
}

tic()
res <- lapply(1:25, f)
t1=toc()
#> 5.025 sec elapsed

tic()
res <- mclapply(1:25, f, mc.cores = 25)
t2=toc()
#> 1.019 sec elapsed
```

Outline/next

EURO

- Introduction to HPC
- Introduction to R
- Parallel R within one node
- Parallelization with Rmpi

Parallelization with Rmpi

What is Rmpi

- Rmpi library: Interface for MPI (Message Passing Interface) in R.
- Enables parallel and distributed computing in the R programming language.
- Facilitates communication and coordination between R processes across multiple nodes.
- Particularly useful for parallelizing computationally intensive tasks like simulations or data processing.
- Users can harness the power of parallel computing for improved performance in certain applications.
- Latest version from Dec 2023, see https://cran.r-project.org/web/packages/Rmpi/Rmpi.pdf

Few basic command

- `Rmpi::mpi.comm.size(0)`: returns the number of active processes in current computing task/job
- `Rmpi::mpi.comm.rank(0)`: returns the ID of current process (number from $\{0, 1, 2, \ldots, \text{size} - 1\}$
- `Rmpi::mpi.get.processor.name()` - returns the name of compute node where the process

Hello word example

Compute smallest eigenvalue of $n \times n$ random symmetric matrices

```
library(Rmpi)
n=30
size <- Rmpi::mpi.comm.size(0)
rank <- Rmpi::mpi.comm.rank(0)
host <- Rmpi::mpi.get.processor.name()
if (rank == 0){
  cat("size ","rank ","host ","max_eigen_value\n")
  cat(size,rank,host,"NaN\n")
} else {
  where=getwd()
  A=matrix(rnorm(n^2),nrow=n)
  A=A+t(A)
  a = max(eigen(A)$values)
  cat(size,rank,host,a,"\n")
}
```

How to distribute this task across cluster

- Save the scripts from previous slide into separate file, called e.g. `Rmpi_master_slave.R`
- Create separate `.batch` file, where the parallelization is defined, e.g., `Job_Rmpi_master_slave.sbatch`

How to distribute this task across cluster

Compute smallest eigenvalue of n symmetric matrices of size $N \times N$

```
#!/bin/bash
#SBATCH --export=ALL,LD_PRELOAD=
#SBATCH --job-name Rmpi
#SBATCH --partition=rome --mem=24GB --time=02:00
#SBATCH --nodes=8
#SBATCH --ntasks-per-node 48 ## maximum is 48
#SBATCH --output=logs/%x_%j.out

module load OpenMPI/4.1.4-GCC-11.3.0
module load R/4.2.1-foss-2022a
srun  Rscript Rmpi_master_slave.R
```

Go to Barbora

- Create directory
  `mkdir /home/rstudio/mnt/`
- copy to it files
  `Job_Rmpi_master_slave.sbatch, Rmpi_master_slave.R`
- mount this directory
  `sshfs -o IdentityFile=/home/rstudio/.ssh/id_ed25519 it4i-jpovh@barbora.it4i.cz:. /home/rstudio/mnt/`
- connect to barbora with ssh
  `ssh -i /home/rstudio/.ssh/id_ed25519 it4i-jpovh@barbora.it4i.cz`

Connect to Barbora

```
rstudio@9cfcdfe2562f:~$ sshfs -o IdentityFile=/home/rstudio/.ssh/id_ed25519 it4i-jpovh@barbora.it4i.cz:. /home/rstudio/mnt/
rstudio@9cfcdfe2562f:~$ ssh -i /home/rstudio/.ssh/id_ed25519 it4i-jpovh@barbora.it4i.cz
client_global_hostkeys_private_confirm: server gave bad signature for ED25519 key 1: incorrect signature
Last login: Tue May 28 00:13:59 2024 from 195.113.175.60

                  Barbora

                  ...running on Red Hat Enterprise Linux 8.4


        Public Service Announcement: Apptainer on the Karolina cluster
Posted:                                                     (2024-05-10 10:23:47)

Apptainer is now a part of the operating system, you do not need to load the
module.

$ apptainer --version
apptainer version 1.3.1-1.el8

[it4i-jpovh@login2.barbora ~]$
```

Run

`sbatch Job_Rmpi_master_slave.sbatch`

Results in log file

```
[1] "size rank host max_eigen_value"
[1] "384 0 cn48 NaN"
[1] "384 110 cn50 8.25199803297607"
[1] "384 173 cn52 8.01187455128492"
[1] "384 68 cn49 8.05800653948316"
[1] "384 200 cn53 8.81600769867893"
[1] "384 258 cn54 8.12244071842822"
[1] "384 332 cn55 7.61927646789373"
[1] "384 338 cn56 4.9472190383247"
```

How parallelise without slurm?

Compute smallest eigenvalue of n symmetric matrices of size $N \times N$

```
rm(list=ls())  # R code: parallel version
library(snow)
library(Rmpi)
nclus=6
cl <-snow::makeMPIcluster(nclus)  #alter either n or mc to affect run time
n=30
N_per_proc=100
#x=matrix(runif(n),n,1)
#x=cbind(1,x)
min_eig_values=function(n,N){
  a=c()
  for (ind in 1:N){
    A=matrix(rnorm(n^2),nrow=n)
    A=A+t(A)
    a[ind] = max(eigen(A)$values)
  }
  return(a)
}
ptim=proc.time()[3]
b=clusterCall(cl,min_eig_values,n=n,N=N_per_proc)
b=unlist(b)
hist(b)
tim=proc.time()[3]-ptim
#Rmpi::mpi.quit()
snow::stopCluster(cl)
```