

Scaling MATLAB on Compute Canada

compute | **calcul**
canada | canada

<https://docs.computecanada.ca/wiki/MATLAB>



Ways to Run MATLAB

- Single node (on the cluster)
 - Interactively
 - Batch mode
- Multi-node
 - On the cluster (interactively or batch mode)
 - Remotely (batch mode from your desktop)

Calculate π

```

% Query for available cores
sz = str2num(getenv('SLURM_CPUS_PER_TASK'));
if isempty(sz), sz = maxNumCompThreads; end

p = parpool('local',sz);

spmd
    a = (labindex - 1)/numlabs;
    b = labindex/numlabs;
    fprintf('Subinterval: [%-4g, %-4g]\n', a, b)

    myIntegral = integral(@quadpi, a, b);
    fprintf('Subinterval: [%-4g, %-4g]   Integral: %4g\n', a, b, myIntegral)


    piApprox = gplus(myIntegral);
end

approx1 = piApprox{1}; % 1st element holds value on worker 1
fprintf('pi           : %.18f\n', pi)
fprintf('Approximation: %.18f\n', approx1)
fprintf('Error          : %g\n',    abs(pi - approx1))

function y = quadpi(x)
y = 4./(1 + x.^2);

```

--cpus-per-task=X



Connecting to Compute Canada Resources

```
[14:17 1] ssh -i ~/.ssh/id_rsa username@cedar.computecanada.ca
```

```
Last login: Sun Jul 5 11:14:59 2020 from 1.2.3.4
```

```
=====
```

```
Welcome to Cedar!
```

```
For information see: https://docs.computecanada.ca/wiki/Cedar Email  
support@computecanada.ca for assistance and/or to report problems.
```

```
=====
```

Cedar Graham Beluga

Copy Source Files to Your Scratch Folder

```
cp -frp /scratch/rsnorris/scratch/matlab-demos-files  
~/scratch
```

```
cd ~/scratch/matlab-demos-files
```

Single Node – Running MATLAB Interactively

```
salloc -A def-training-wa_cpu --reservation=wgss3-wr_cpu -n 1 --cpus-per-task=8
      --mem-per-cpu=2gb -t 00:30:00
```

```
module load matlab/2020a
matlab -nodisplay
```

```
>> calcp_i
Starting parallel pool (parpool) using the 'local' profile ...
Connected to the parallel pool (number of workers: 8).
Lab 1:
  Subinterval: [0 , 0.125]
...
pi           : 3.141592653589793116
Approximation: 3.141592653589793116
Error        : 0
```

Single Node – Running MATLAB Batch Mode (1)

```
#!/bin/sh

#SBATCH -A def-training-wa_cpu
#SBATCH -n 1                # 1 instance of MATLAB
#SBATCH --cpus-per-task=32  # 32 cores (1 node)
#SBATCH --mem-per-cpu=2gb
#SBATCH --time=00:10:00

#SBATCH --reservation=wgss3-wr_cpu

# Add MATLAB to system path
module load matlab/2020a

# Run code
matlab -batch calcpi
```

Single Node – Running MATLAB Batch Mode (2)

```
[cedar5:~] cd ~/scratch/matlab-demos-files
[cedar5:~/scratch/matlab-demo-files] sbatch matlab.slurm
Submitted batch job 45452890
[cedar5:~/scratch/matlab-demo-files] head slurm-45452890.out
Opening log file: /tmp/java.log.30847
Starting parallel pool (parpool) using the 'local' profile ...
Connected to the parallel pool (number of workers: 32).
Lab 1:
  Subinterval: [0 , 0.03125]
Lab 2:
  Subinterval: [0.03125, 0.0625]
Lab 3:
  Subinterval: [0.0625, 0.09375]
Lab 4:
[cedar5:~/scratch/matlab-demo-files]
[cedar5:~/scratch/matlab-demo-files] tail slurm-45452890.out
  Subinterval: [0.90625, 0.9375]  Integral: 0.0675754
Lab 31:
  Subinterval: [0.9375, 0.96875]  Integral: 0.0655008
Lab 32:
  Subinterval: [0.96875, 1 ]  Integral: 0.0634867
Lab 25:
  Subinterval: [0.75, 0.78125]  Integral: 0.0788075
pi      : 3.141592653589793116
Approximation: 3.141592653589793116
Error   : 0
```


What About GPUs?

```
len = 2^14; % 16384

% A: 2 GB
% b: 128 KB

A = rand(len);
b = rand(len,1);
cpu = @() (A\b);
timeit(cpu)

gpuDevice
gA = gpuArray.rand(len);
gb = gpuArray.rand(len,1);
gpu = @() (gA\gb);
gpstimeit(gpu)
```

`solve_sys_linear_eqns.m`

Submitting GPU Jobs (1)

```
#!/bin/sh

#SBATCH -A def-razoumov-ac_cpu
#SBATCH -n 1                # 1 instance of MATLAB
#SBATCH --cpus-per-task=24  # 24 cores (1 node)
#SBATCH --mem-per-cpu=4gb
#SBATCH --time=00:10:00

#SBATCH --gres=gpu:p100:1  # Request a single GPU

# Add MATLAB to system path
module load matlab/2020a

# Run code
matlab -batch solve_sys_linear_eqns

[cedar5:~/scratch/matlab-demo-files] sbatch matlab-gpu.slurm
```

matlab-gpu.slurm

Submitting GPU Jobs (2)

```
Opening log file: /tmp/java.log.24269
```

```
ans =  
    9.8840
```

```
ans =  
    CUDADevice with properties:
```

```
        Name: 'Tesla P100-PCIE-12GB'  
        Index: 1  
    ComputeCapability: '6.0'  
        SupportsDouble: 1  
    MaxThreadsPerBlock: 1024  
        MaxShmemPerBlock: 49152  
    MaxThreadBlockSize: [1024 1024 64]  
        MaxGridSize: [2.1475e+09 65535 65535]  
        SIMDWidth: 32  
        TotalMemory: 1.2791e+10  
        AvailableMemory: 1.2399e+10  
    MultiprocessorCount: 56  
        ClockRateKHz: 1328500  
        ComputeMode: 'Default'
```

```
ans =  
    1.0806
```

Configure MATLAB for Multi-node Jobs

```
[cdr681:/scratch/rsnorris] module load matlab/2020a  
[cdr681:/scratch/rsnorris] matlab -nodisplay  
Opening log file: /tmp/java.log.18160
```

```
< M A T L A B (R) >  
Copyright 1984-2020 The MathWorks, Inc.  
R2020a (9.8.0.1323502) 64-bit (glnxa64)  
February 25, 2020
```

To get started, type doc.
For product information, visit www.mathworks.com.

```
>> configCluster
```

Must set AccountName, WallTime, and MemUsage before submitting jobs to CEDAR. E.g.

```
>> c = parcluster;  
>> c.AdditionalProperties.AccountName = 'account-name';  
>> % 5 hour walltime  
>> c.AdditionalProperties.WallTime = '05:00:00';  
>> % 2 GB per core  
>> c.AdditionalProperties.MemUsage = '2GB';  
>> c.saveProfile
```

Multi-node

```
function calcp_i

% Query for available cores
sz = str2num(getenv('SLURM_CPUS_PER_TASK'));
if isempty(sz), sz = maxNumCompThreads; end

p = parpool('local',sz);

spmd
    a = (labindex - 1)/numlabs;
    b = labindex/numlabs;
    fprintf('Subinterval: [%-4g, %-4g]\n', a, b)

    myIntegral = integral(@quadpi, a, b);
    fprintf('Subinterval: [%-4g, %-4g]   Integral: %4g\n', a, b, myIntegral)

    piApprox = gplus(myIntegral);
end

approx1 = piApprox{1}; % 1st element holds value on worker 1
fprintf('pi           : %.18f\n', pi)
fprintf('Approximation: %.18f\n', approx1)
fprintf('Error          : %g\n',   abs(pi - approx1))

function y = quadpi(x)
y = 4./(1 + x.^2);
```

Multi-node

```
function calcpi_multi_node
c = parcluster;
c.AdditionalProperties.AccountName = 'def-razoumov-ac';
c.AdditionalProperties.MemUsage = '2gb';
c.AdditionalProperties.WallTime = '00:10:00';

p = c.parpool(64);

spmd
    a = (labindex - 1)/numlabs;
    b = labindex/numlabs;
    fprintf('Subinterval: [%-4g, %-4g]\n', a, b)

    myIntegral = integral(@quadpi, a, b);
    fprintf('Subinterval: [%-4g, %-4g]   Integral: %4g\n', a, b, myIntegral)

    piApprox = gplus(myIntegral);
end

approx1 = piApprox{1}; % 1st element holds value on worker 1
fprintf('pi           : %.18f\n', pi)
fprintf('Approximation: %.18f\n', approx1)
fprintf('Error          : %g\n',   abs(pi - approx1))

function y = quadpi(x)
y = 4./(1 + x.^2);
```

Submitting Multi-node Jobs (1)

```
#!/bin/sh

#SBATCH -A def-training-wa_cpu
#SBATCH -n 1 # 1 instance of MATLAB
#SBATCH --cpus-per-task=32 # 32 cores (1 node)
#SBATCH --mem-per-cpu=2gb
#SBATCH --time=00:10:00

#SBATCH --reservation=wgss3-wr_cpu

# Add MATLAB to system path
module load matlab/2020a

# Run code
matlab -batch calcpu
```

Submitting Multi-node Jobs (2)

```
#!/bin/sh

#SBATCH -A def-training-wa_cpu
#SBATCH -n 1 # 1 instance of MATLAB
#SBATCH --cpus-per-task=1 # 1 cores (1 node)
#SBATCH --mem-per-cpu=2gb
#SBATCH --time=00:20:00

#SBATCH --reservation=wgss3-wr_cpu

# Add MATLAB to system path
module load matlab/2020a

# Run code
matlab -batch calcpu_multi_node
```

matlab-multi-node.slurm

Submitting Multi-node Jobs (3)

```
[cdr681:/scratch/rsnorris/matlab-demo-files] sbatch matlab-multi-node.slurm
Submitted batch job 45457455
[cdr681:/scratch/rsnorris/matlab-demo-files]
[cdr681:/scratch/rsnorris/matlab-demo-files] head slurm-45457455.out
Opening log file: /tmp/java.log.19109
Starting parallel pool (parpool) using the 'cedar R2020a' profile ...
```

```
additionalSubmitArgs =
```

```
'--ntasks=64 --cpus-per-task=1 --ntasks-per-core=1 -A def-razoumov-ac --mem-per-cpu=2gb -t 00:10:00'
```

```
Connected to the parallel pool (number of workers: 64).
```

```
Lab 2:
```

```
Subinterval: [0.015625, 0.03125]
```

```
[cdr681:/scratch/rsnorris/matlab-demo-files]
```

```
[cdr681:/scratch/rsnorris/matlab-demo-files] scontrol show job 45457712 | grep NumCPUs
```

```
NumNodes=19 NumCPUs=64 NumTasks=64 CPUs/Task=1 ReqB:S:C:T=0:0:*:*
```

```
[cdr681:/scratch/rsnorris/matlab-demo-files]
```

```
[cdr681:/scratch/rsnorris/matlab-demo-files] scontrol show job 45457712 | grep NodeList
```

```
ReqNodeList=(null) ExcNodeList=(null)
```

```
NodeList=cdr[501,504-505,513,519,540,567,616,622,624,651,659,680-681,688,737,743,845,868]
```

Configure MATLAB When Running on Your Local Desktop



```
>> configCluster
```

```
Username on CEDAR (e.g. joe): rsnorris
```

Must set AccountName, WallTime, and MemUsage before submitting jobs to CEDAR. E.g.

```
>> c = parcluster;  
>> c.AdditionalProperties.AccountName = 'account-name';  
>> % 5 hour walltime  
>> c.AdditionalProperties.WallTime = '05:00:00';  
>> % 2 GB per core  
>> c.AdditionalProperties.MemUsage = '2GB';  
>> c.saveProfile
```



```
>> c = parcluster;  
>> c.AdditionalProperties.AccountName = 'def-razoumov-ac';  
>> c.AdditionalProperties.WallTime = '00:10:00';  
>> c.AdditionalProperties.MemUsage = '2gb';  
>> c.saveProfile  
>>
```

Remote Submission

```
>> j = c.batch(@calcp_i, 0, {}, 'CurrentFolder','.', 'AutoAddClientPath',false, 'pool',63);
```

```
additionalSubmitArgs =
```

```
'--ntasks=64 --cpus-per-task=1 --ntasks-per-core=1 -A def-razoumov-ac --mem-per-cpu=2gb -t 00:10:00'
```



```
>> j.wait  
>> j.diary
```

Getting Prior Results – Job Monitor

Job Monitor
↑ - □ ×

Select Profile: cedar R2020a (default) Show jobs from all users

ID	Username	Submit Time	Finish Time	Tasks	State	Description
3	rayn	Wed Jul 01 19:42:31 EDT 2020	Wed Jul 01 19:42:5...	1	finished	Batch job running function
5	rayn	Wed Jul 01 19:44:12 EDT 2020	Thu Jul 02 02:44:5...	3	finished	Batch job running function
6	rayn	Wed Jul 01 19:45:18 EDT 2020		10	failed	Independent job
7	rayn	Wed Jul 01 19:56:57 EDT 2020	Wed Jul 01 19:58:2...	1	finished	Batch job running function
8	rayn	Wed Jul 01 19:57:35 EDT 2020	Thu Jul 02 02:59:1...	5	finished	Batch job running function
9	rayn	Wed Jul 01 19:58:09 EDT 2020		10	failed	Independent job
10	rayn	Wed Jul 01 20:02:51 EDT 2020		10	pending	Independent job
11	rayn	Wed Jul 01 20:05:10 EDT 2020		10	pending	Independent job
12	rayn	Mon Jul 06 00:06:05 EDT 2020	Mon Jul 06 07:07:0...	64	finished	Batch job running function

Last updated at Mon Jul 06 00:23:57 EDT 2020
Auto update: Every 5 minutes