

# Getting Started with Serial and Parallel MATLAB on Tintin

---

## CONFIGURATION

Start MATLAB. Configure MATLAB to run parallel jobs on your cluster by calling `configCluster`. For each cluster, `configCluster` only needs to be called once per version of MATLAB.

```
>> configCluster
```

Jobs will now default to the cluster rather than submit to the local machine.

Before submitting a job, you must specify the name of the charge account and the walltime, via `ClusterInfo`, which will be explained in more detail below.

## CONFIGURING JOBS

Prior to submitting the job, we can specify various parameters to pass to our jobs, such as queue, e-mail, etc.

Specification is done with `ClusterInfo`. The `ClusterInfo` class supports tab completion to ease recollection of method names.

NOTE: Any parameters set with `ClusterInfo` will be persistent between MATLAB sessions.

```
>> % Specify a charge account [REQUIRED]
>> ClusterInfo.setProjectName('name-of-account')
```

```
>> % Set walltime to 1 hour [REQUIRED]
>> ClusterInfo.setWallTime('01:00:00')
```

```
>> % Specify a partition to run on [REQUIRED]
>> ClusterInfo.setQueueName('partition-name')
```

```
>> % Specify e-mail address to receive notifications about your job
>> ClusterInfo.setEmailAddress('user@company.com')
```

```
>> % Request a 64 GB node
>> ClusterInfo.setMemUsage('mem64GB')
```

Additional parameters that can be supplied are:

- `ProcsPerNode`
- `RequireExclusiveNodes`
- `Reservation`
- `UseGpu`

To see the values of the current configuration options, call the `state` method. To clear a value, assign the property an empty value (`''`, `[]`, or `false`), or call the `clear` method to clear all values.

```
>> % To view current configurations
>> ClusterInfo.state

>> % To clear a configuration that takes a string as an input argument
>> ClusterInfo.setEmailAddress(' ')

>> % To clear all configurations
>> ClusterInfo.clear
```

## SERIAL JOBS

Use the `batch` command to submit asynchronous jobs to the cluster. The `batch` command will return a job object which is used to access the output of the submitted job. See the MATLAB documentation for more help on `batch`.

```
>> % Get a handle to the cluster
>> c = parcluster;

>> % Submit a job to query where MATLAB is running on the cluster
>> j = c.batch(@pwd, 1, {});

>> % Query the job for the state
>> j.State

>> % If the state of the job is finished, fetch the results
>> j.fetchOutputs{:}

>> % Delete the job after the results are no longer needed
>> j.delete
```

To retrieve a list of running or completed jobs, call `parcluster` to retrieve the cluster object. The cluster object stores an array of jobs that were run, are running, or are queued to run. This allows us to fetch the results of completed jobs. Retrieve and view the list of jobs as shown below.

```
>> c = parcluster;

>> % Retrieve the results of past jobs from the cluster
>> jobs = c.Jobs
```

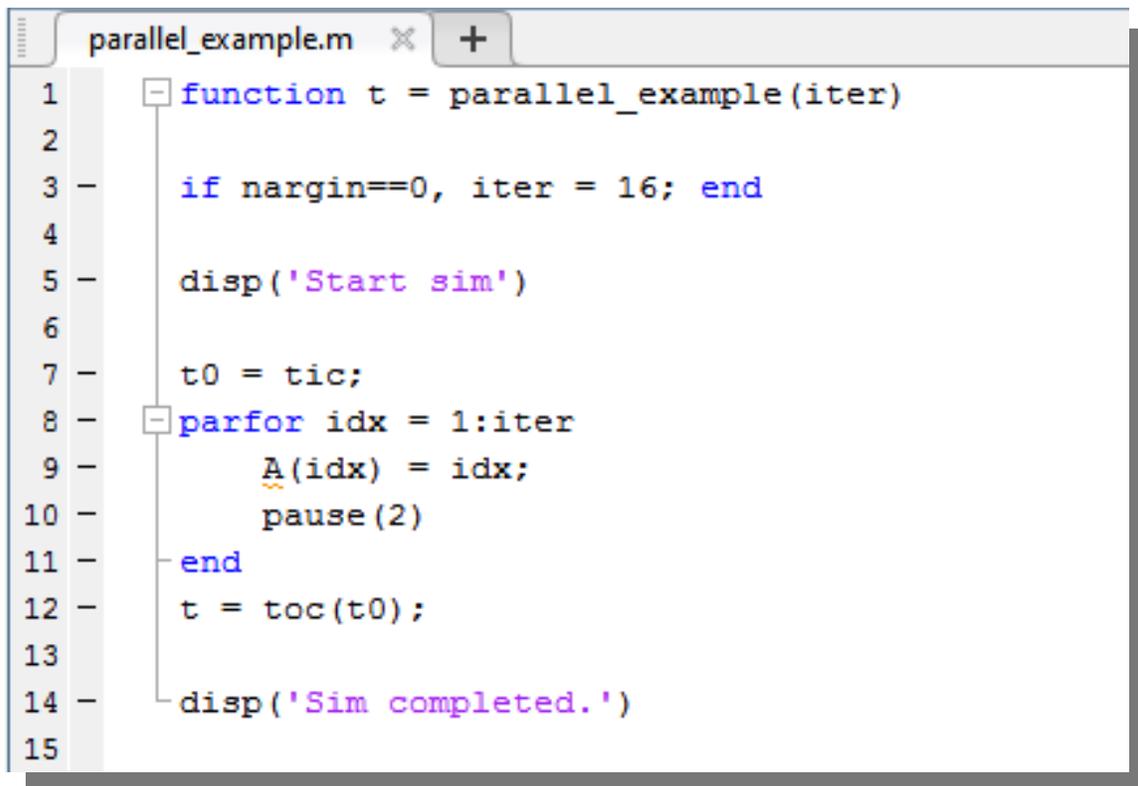
Once we've identified the job we want, we can retrieve the results as we've done previously.

```
>> % Retrieve the results of the 3rd job
>> j3 = jobs(3);
>> j3.fetchOutputs{:}
```

`fetchOutputs` is used to retrieve function output arguments; if calling `batch` with a script, use `load` instead. Data that has been written to files on the cluster needs be retrieved directly from the file system.

## PARALLEL JOBS

Users can also submit parallel workflows with `batch`. Let's use the following example for our parallel job.



```
parallel_example.m x +
1 function t = parallel_example(iter)
2
3     if nargin==0, iter = 16; end
4
5     disp('Start sim')
6
7     t0 = tic;
8     parfor idx = 1:iter
9         A(idx) = idx;
10        pause(2)
11    end
12    t = toc(t0);
13
14    disp('Sim completed.')
15
```

We'll use the `batch` command again, but since we're running a parallel job, we'll also specify a MATLAB Pool.

```
>> % Get a handle to the cluster
>> c = parcluster;

>> % Submit a batch pool job using 4 workers for 16 simulations
>> j = c.batch(@parallel_example, 1, {}, 'Pool', 4);

>> % View the current job status
>> j.State

>> % Fetch the results after the job is in the finished state
```

```
>> j.fetchOutputs{:}
```

```
ans =
```

```
8.8872
```

The job ran in 8.89 seconds using 4 workers. Note that these jobs will always request N+1 CPU cores, since one worker is required to manage the batch job and pool of workers. For example, a job that needs eight workers will consume nine CPU cores.

We'll run the same simulation, but increase the Pool size. This time, to retrieve the results at a later time, we'll keep track of the job ID.

NOTE: For some applications, there will be a diminishing return when allocating too many workers, as the overhead may exceed computation time.

```
>> % Get a handle to the cluster
```

```
>> c = parcluster;
```

```
>> % Submit a batch pool job using 8 workers for 16 simulations
```

```
>> j = c.batch(@parallel_example, 1, {}, 'Pool', 8);
```

```
>> % Get the job ID
```

```
>> id = j.ID
```

```
id =
```

```
4
```

```
>> % Clear the job variable, as though we quit MATLAB
```

```
>> clear j
```

Once we have a handle to the cluster, we'll call the findJob method to search for the job with the specified job ID.

```
>> % Get a handle to the cluster
```

```
>> c = parcluster;
```

```
>> % Find the old job
```

```
>> j = c.findJob('ID', 4);
```

```
>> % Retrieve the state of the job
```

```
>> j.State
```

```
ans
```

```
finished
```

```
>> % Fetch the results
```

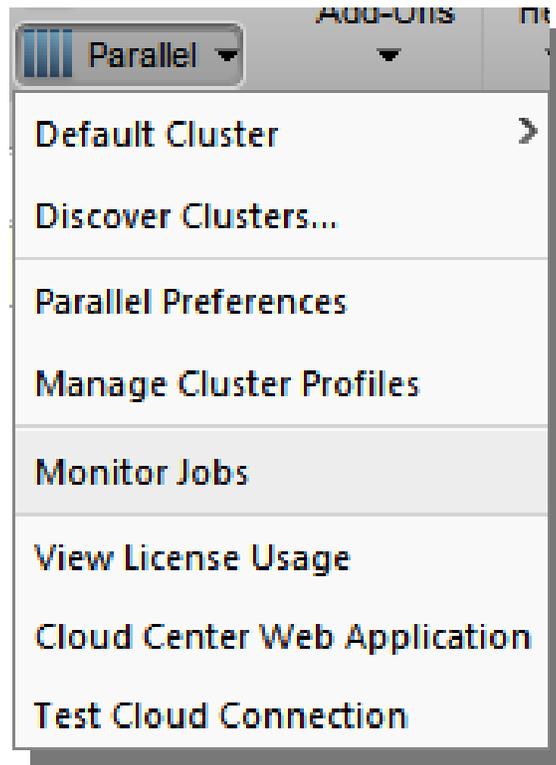
```
>> j.fetchOutputs{:}
```

```
ans =
```

4.7270

The job now runs in 4.73 seconds using 8 workers. Run the code with different number of workers to determine the ideal number to use.

Alternatively, to retrieve job results via a graphical user interface, use the Job Monitor (Parallel > Monitor Jobs).



## DEBUGGING

If a serial job produces an error, we can call the `getDebugLog` method to view the error log file.

```
>> j.Parent.getDebugLog(j.Tasks(1))
```

When submitting independent jobs, with multiple tasks, you will have to specify the task number. For Pool jobs, do not deference into the job object.

```
>> j.Parent.getDebugLog(j)
```

The scheduler ID can be derived by calling schedID

```
>> schedID(j)
```

```
ans
```

```
25539
```

This is how the scheduler identifies the job.

## **TO LEARN MORE**

To learn more about the MATLAB Parallel Computing Toolbox, check out these resources:

- [Parallel Computing Coding Examples](#)
- [Parallel Computing Documentation](#)
- [Parallel Computing Overview](#)
- [Parallel Computing Tutorials](#)
- [Parallel Computing Videos](#)
- [Parallel Computing Webinars](#)