



EEE4120F

High Performance Embedded Systems

REFERENCE GUIDE

MATLAB Parallel Computing Toolkit(PCT)

Course Team

Course Convener: Associate Professor Simon Winberg

Teaching Assistant: Travimadox Webb

MATLAB Parallel Computing Toolkit

A Practical Guide for Practical 2

1 Introduction

The MATLAB Parallel Computing Toolbox (PCT) enables you to leverage multicore processors, GPUs, and computer clusters to accelerate computations and handle large datasets efficiently. Parallel computing involves breaking down problems into smaller parts that can be solved simultaneously. When this decomposition is straightforward, we call it “**embarrassingly parallel**”.

Tip

Installation: Ensure you have MATLAB installed with the Parallel Computing Toolbox add-on. This usually requires a separate installation via the “Add-Ons” menu.

1.1 Parallel Workers

Parallel Workers are computational engines (or threads) that execute code in parallel. They are the fundamental processing units in PCT and are essentially headless MATLAB instances running in the background.

2 Converting Serial Code to Parallel Code

The simplest way to parallelise MATLAB code is replacing for loops with parfor (parallel for) loops. MATLAB intelligently analyses the loop body, checking data dependencies to ensure correct parallel execution without breaking dependencies or overwriting data mid-process.

```
1 % Serial code
2 for i = 1:1000
3     A(i) = computeIntensiveTask(i);
4 end
5
6 % Parallel code - automatically parallelised
7 parfor i = 1:1000
8     A(i) = computeIntensiveTask(i);
9 end
```

3 Parallel Performance

Parallel computing doesn't guarantee linear speedup. Several factors affect performance:

- **Overhead:** Starting workers and transferring data takes time
- **Amdahl's Law:** Speedup limited by serial portions of code
- **Load Balancing:** Unequal iteration times reduce efficiency

Tip

Performance Tip: Use parallel computing when the benefits outweigh the overhead costs. Do not include pool startup time in benchmarking!

4 Parallel Pools and MATLAB Workers

A parallel pool is a collection of MATLAB workers on your local machine or cluster. See example below:

```

1 % Start a parallel pool with N workers on your local machine
2 parpool('local', N)
3
4 % Check if a pool exists (without creating one)
5 poolobj = gcp('nocreate');
6
7 % Create a pool only if one doesn't exist
8 if isempty(poolobj)
9     parpool('local', N);
10 end
11
12 % Alternative: Using parcluster for more control
13 c = parcluster;
14 p = c.parpool(4); % Creates pool with 4 workers
15
16 % Delete the current pool
17 delete(gcp('nocreate'))
18 % or
19 p.delete % If you created pool with parcluster

```

Tip

Best Practice: MATLAB automatically creates a pool when you first use `parfor`. Use `gcp('nocreate')` to work with existing pools and avoid unnecessary overhead.

5 Understanding the Rules of `parfor`

5.1 Parallel Execution Order

Unlike regular for-loops, **`parfor` iterations execute in random order**. Each iteration must be independent. Note that you cannot query which worker thread is executing a particular iteration within `parfor`.

5.2 SPMD Blocks: Single Program Multiple Data

SPMD stands for “Single Program Multiple Data” i.e the same code runs on different workers but operates on different data. This is similar to how GPU cores operate in lockstep.

```

1 % SPMD example - all workers run the same code
2 spmd
3     fprintf('Worker %d says Hello World!\n', labindex);
4     % labindex tells you which worker number (1 = master thread)
5     % Note: spmdIndex is preferred in MATLAB 2022+
6 end

```

Within `spmd` blocks, use `labindex` (or `spmdIndex` in MATLAB 2022+) to identify which worker is executing. If `labindex` returns 1, you're in the master thread.

5.3 Non-Parallelisable Code

Code that cannot be parallelised with `parfor` includes:

- Iterations that depend on previous iterations
- Code that modifies the same variable in unpredictable ways
- Nested `parfor` loops (outer loop can be `parfor`, inner must be `for`)

Warning

Invalid Example: Each iteration depends on the previous one

```
1 parfor i = 2:n
2     A(i) = A(i-1) + B(i); % Error!
3 end
```

6 Types of Variables in `parfor`-Loops

6.1 Loop Index Variables

The loop index (e.g., `i` in `parfor i = 1:n`) is automatically managed and each worker gets its assigned values.

6.2 Reduction Operations

Reduction variables accumulate values across iterations using associative operations (`+`, `*`, `min`, `max`, etc.).

```
1 total = 0;
2 parfor i = 1:1000
3     total = total + computeValue(i); % Valid reduction
4 end
```

6.3 Sliced and Broadcast Variables

Sliced variables: Arrays where each iteration accesses different elements in a predictable pattern.

```
1 A = zeros(100,1);
2 parfor i = 1:100
3     A(i) = i^2; % A is sliced (each worker gets different indices)
4 end
```

Broadcast variables: Read-only variables sent to all workers. Large broadcast variables reduce performance.

6.4 Distributed Arrays

Arrays can be automatically distributed across multiple workers for processing and then gathered back to the master thread.

```

1 % Distribute array A across workers
2 D = distributed(A); % D is distributed version accessible by workers
3
4 % Example with SPMD and distributed data
5 p = parpool('Processes', 4);
6 spmd
7     C = rand(3, labindex-1); % Each worker creates different size array
8 end
9
10 % View how data is distributed (each lab has different portion)
11 C % Shows: Lab 1: [3x0], Lab 2: [3x1], Lab 3: [3x2], Lab 4: [3x3]
12
13 % Gather distributed data back to master thread
14 finalResult = gather(C); % Consolidates all data into single array
15
16 delete(p);

```

Tip

Key Functions: Use `distributed()` to split data across workers, and `gather()` (or `gplus/spmdPlus`) to collect results back to the master thread.

6.5 Limit the Size of Broadcast Variables

Tip

Optimisation: Extract only needed data before the `parfor` loop to minimize broadcast size.

```

1 % Instead of broadcasting entire structure:
2 params = largeStruct.neededParams;
3 parfor i = 1:n
4     result(i) = compute(params, i);
5 end

```

6.6 Improve Efficiency with Sliced and Temporary Variables

Use temporary variables inside the loop to reduce classification complexity and improve performance.

```

1 parfor i = 1:n
2     temp = expensiveComputation(data(i));
3     result(i) = processTemp(temp);
4 end

```

7 Nested Loops with Shared Memory

Only the outermost loop can be a `parfor`. Inner loops must remain serial.

```

1 % Parallelize outer loop
2 parfor i = 1:m

```

```
3     for j = 1:n
4         C(i,j) = A(i,j) * B(i,j);
5     end
6 end
```

Tip

Performance Strategy: Parallelise the loop with the most iterations or longest execution time. Test both configurations to find the optimal approach.

8 GPU Computing

MATLAB can utilise GPUs for certain computations and may do so automatically for various functions, particularly when data size and processing type merit GPU use.

8.1 Checking GPU Availability

To verify if your GPU is accessible to MATLAB, call `gpuArray` without arguments:

```
>> gpuArray
ans = [] % Empty array indicates GPU is working
```

If you see an error indicating no GPU available or drivers not found, and you have a compatible GPU (e.g., not too ancient NVIDIA card), you may need to update your graphics drivers, reboot, and restart MATLAB.

You can also check the number of available GPUs:

```
>> gpuDeviceCount
ans = 1
```

8.2 Basic GPU Workflow

The typical approach involves creating/loading data on the CPU, transferring to GPU, performing computations, then gathering results back.

```
1 % Create an array on CPU
2 X = [1, 2, 3];
3
4 % Transfer data to GPU
5 G = gpuArray(X);
6
7 % Verify data is on GPU
8 isgpuarray(G) % Returns logical 1 (true)
9
10 % Perform element-wise computation (SPMD operation on GPU)
11 GSq = G.^2; % GPU computes what needs to be done
12
13 % Transfer result back to CPU
14 XSq = gather(GSq) % Gathers distributed data blocks
15 % Result: XSq = [1, 4, 9]
16
17 % Verify data is back on CPU
18 isgpuarray(XSq) % Returns logical 0 (false)
```

Tip

Key Point: You do NOT need to put GPU code inside an `spmd` block. The `gpuArray` function handles data transfer, and operations on GPU arrays automatically execute on the GPU. The GPU inherently operates in SPMD fashion across its many cores.

Tip

When to use GPU: Matrix operations, element-wise operations, and algorithms with high arithmetic intensity. GPUs excel at operations on large arrays but have overhead for small data.

Warning

Note: GPU computing requires a compatible NVIDIA GPU and the Parallel Computing Toolbox.

9 Practice & Next Steps

For more information, see: [MATLAB Quick Start Tutorial](#).