




Minho
Advanced
Computing
Center



OpenMP

Advanced Concepts


André Pereira

Minho Advanced Computing Center

ampereira@macc.fccn.pt

OpenMP – Race Conditions

- Inadequate management of concurrent accesses may create wrong results
 - Data races
 - Read-after-write and write-after-read dependencies must be handled by the programmer
- OpenMP provides pragmas to control access to shared data
 - `#pragma omp critical`
- Algorithms and/or data structures may need to be redesigned



OpenMP – Reduction

Threads often compute a subset of a workload

- Partial results stored privately to each thread
 - Improves performance (less contention to access shared data)
 - Reduces the risk of data races
- Partial results must be merged after the parallel section
 - can be implemented manually
 - `#pragma omp reduction` does this automatically (restrictions apply)

```
1  int partial_result[NUMBER_OF_THREADS];
2  int final_result = 0;
3
4  #pragma omp parallel
5  {
6      #pragma omp for
7      for (int i = 0; i < n; i++) {
8          // computes partial result x
9          // ...
10         partial_result[thread_id] = x;
11     }
12 }
13 for (int i = 0; i < NUMBER_OF_THREADS; i++)
14     final_result += partial_result[i];
```





```
1  int result;
2
3  #pragma omp parallel
4  {
5      #pragma omp for reduction(+:result)
6      for (int i = 0; i < n; i++) {
7          // computes partial result x
8          // ...
9          result += x;
10     }
11 }
```

OpenMP – Reduction

```
1 int partial_result[NUMBER_OF_THREADS];
2 int final_result = 0;
3
4 #pragma omp parallel
5 {
6     #pragma omp for
7     for (int i = 0; i < n; i++) {
8         // computes partial result x
9         // ...
10        partial_result[thread_id] = x;
11    }
12 }
13 for (int i = 0; i < NUMBER_OF_THREADS; i++)
14     final_result += partial_result[i];
```



```
1 int result;
2
3 #pragma omp parallel
4 {
5     #pragma omp for reduction(+:result)
6     for (int i = 0; i < n; i++) {
7         // computes partial result x
8         // ...
9         result += x;
10    }
11 }
```



OpenMP – Thread Management

Controlling the sections of code that each thread executes can be achieved through:

- Each thread individual id
- OpenMP pragma directives
 - **Master**: only the master thread executes the code section
 - **Single**: the first thread to arrive executes the code section

```
1  int sum;  
2  
3  #pragma omp parallel  
4  {  
5      if (thread_id == 0)  
6          sum = 0;  
7  
8      // ...  
9  }
```

=

```
1  int sum;  
2  
3  #pragma omp parallel  
4  {  
5      #pragma omp master  
6          sum = 0;  
7  
8      // ...  
9  }
```

≠

```
1  int sum;  
2  
3  #pragma omp parallel  
4  {  
5      #pragma omp single  
6          sum = 0;  
7  
8      // ...  
9  }
```

OpenMP – Loop Scheduling

Workloads can be classified as

- **Regular**: each loop iteration takes the same amount of time
- **Irregular**: loop iterations take different amount of time

The execution time of a parallel region is the time of the slowest thread

```
1  #pragma omp parallel
2  {
3      #pragma omp for
4      for (int i = 0; i < n_particles; i++) {
5          e = calculateEnergy (i);
6      }
7  }
```

REGULAR WORKLOAD

```
1  #pragma omp parallel
2  {
3      #pragma omp for
4      for (int i = 0; i < n_particles; i++) {
5          float d = distanceToParticle (i);
6
7          if (d > 1000)
8              e = calculateEnergy (i);
9          else
10             e = 0;
11     }
12 }
```



IRREGULAR WORKLOAD

OpenMP – Loop Scheduling

Workloads can be classified as

- **Regular**: each loop iteration takes the same amount of time
- **Irregular**: loop iterations take different amount of time

The execution time of a parallel region is the time of the slowest thread

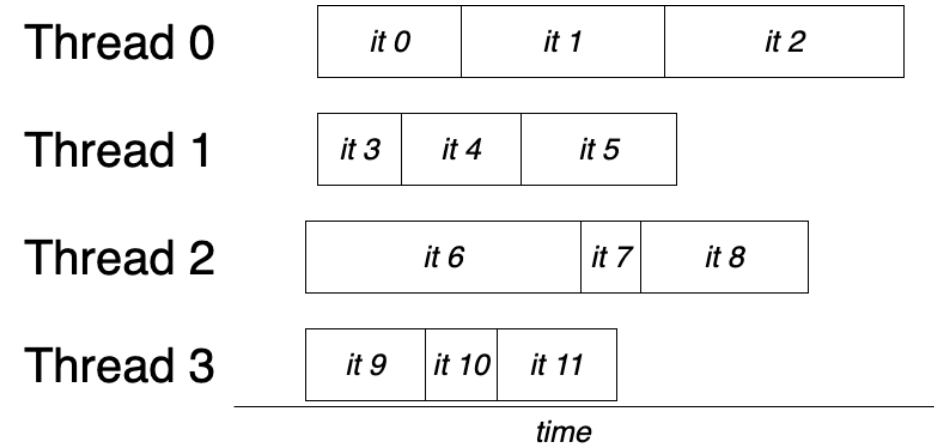


OpenMP – Loop Scheduling

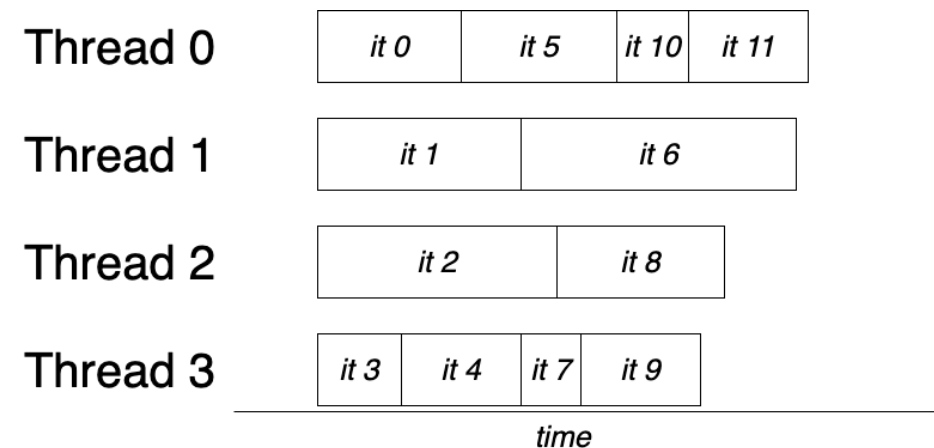
OpenMP loop scheduling options

- **static**
 - Iterations are assigned equally among threads prior to the parallel section
 - Small overhead
 - Static user-defined chunk size
- **dynamic**
 - Each thread gets iterations according to their throughput
 - Higher overhead
 - Static user-defined chunk size
- **guided**
 - Similar to dynamic but adjusts the chunk size during runtime
 - Less overhead than dynamic

STATIC SCHEDULING



DYNAMIC SCHEDULING



OpenMP – Loop Scheduling

OpenMP loop scheduling options

- static
- dynamic
- guided

Chunk size can also be user defined

- Set as a single iteration by default
- Larger chunks require less scheduling overhead
- Smaller chunks are better to schedule irregular workloads

```
1  #pragma omp parallel
2  {
3      #pragma omp for schedule(dynamic,10)
4      for (i=1; i<n; i++)
5          b[i] = (a[i] + a[i-1]) / 2.0;
6  }
```



Lab Session

Vector Sum

1. A simple parallelization

- Parallelize the code using `#pragma omp parallel`
- Distribute the iterations using a `#pragma omp for` directive
- Use the `critical` directive on shared data if necessary
- Execute and measure the performance of the code

2. Privatize shared data

- Privatize the accesses to shared data
 - each thread should contain it's copy of sum
- Implement a manual merge of the partial results
- Execute and measure the performance of the code

3. Removing implicit synchronizations



- Replace the manual merge of the results with OpenMP's reduction directive
- Execute and measure the performance of the code

```
4  int vector_sum (int array[], int size) {
5      int sum = 0;
6
7      for (int i = 0; i < size; i++)
8          sum += array[i];
9
10     return sum;
11 }
```

Pi Calculation - Integral Approximation

Pi can be calculated through the area of a circle

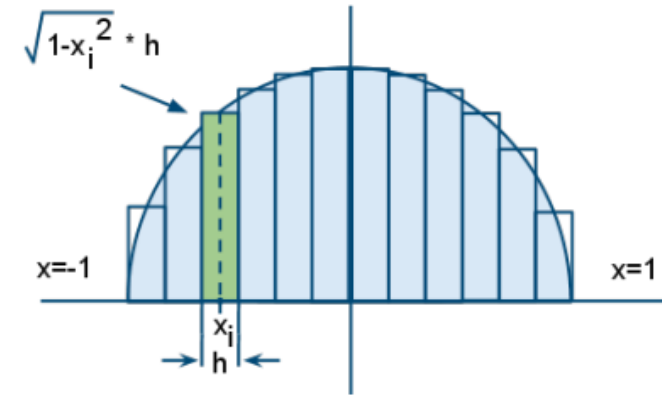
- The integral of the function of a circle is equivalent
- Integrals can be approximated through iterative methods



Pi Calculation - Integral Approximation

For loop parallelization

- Parallelize the code using the `omp parallel` and `omp for` directives
- Privatize variables if needed
- Each thread should compute a partial sum (avoid race conditions!)
- Merge the partial results using OpenMP directives
- Execute the code and measure its performance



```
4 double pi_integration (long num_steps) {
5     int i;
6     double x, pi, sum = 0.0;
7     double step = 1.0 / (double) num_steps;
8
9     for (i = 0; i < num_steps; i++) {
10         x = (i + 0.5) * step;
11         sum = sum + 4.0 / (1.0 + x * x);
12     }
13
14     pi = step * sum;
15
16     return pi;
17 }
```

Thank you for attending!



hello@macc.fccn.pt

macc.fccn.pt



@minhoacc