

1 Thread-Level Parallelism

As powerful as data level parallelization is, it can be quite inflexible, as not all applications have data that can be vectorized. Multithreading, or running a single piece of software on multiple hardware threads, is much more powerful and versatile.

OpenMP provides an easy interface for using multithreading within C programs. Some examples of OpenMP directives:

- The `parallel` directive indicates that each thread should run a copy of the code within the block. If a for loop is put within the block, **every** thread will run every iteration of the for loop.

```
#pragma omp parallel
{
    ...
}
```

NOTE: The opening curly brace needs to be on a newline or **else** there will be a compile-time error!

- The `parallel for` directive will split up iterations of a for loop over various threads. Every thread will run **different** iterations of the for loop. The following two code snippets are equivalent.

```
#pragma omp parallel for
for (int i = 0; i < n; i++) {
    ...
}
```

```
#pragma omp parallel
{
    #pragma omp for
    for (int i = 0; i < n; i++) { ... }
}
```

There are two functions you can call that may be useful to you:

- `int omp_get_thread_num()` will return the number of the thread executing the code
- `int omp_get_num_threads()` will return the number of total hardware threads executing the code

1.1 For each question below, state and justify whether the program is **sometimes incorrect**, **always incorrect**, **slower than serial**, **faster than serial**, or **none of the above**. Assume the default number of threads is greater than 1. Assume no thread will complete before another thread starts executing. Assume `arr` is an `int[]` of length `n`.

- (a) // Set element `i` of `arr` to `i`
- ```
#pragma omp parallel
```

```

{
 for (int i = 0; i < n; i++)
 arr[i] = i;
}

```

(b) // Set arr to be an array of Fibonacci numbers.

```

arr[0] = 0;
arr[1] = 1;
#pragma omp parallel for
for (int i = 2; i < n; i++)
 arr[i] = arr[i-1] + arr[i - 2];

```

(c) // Set all elements in arr to 0;

```

int i;
#pragma omp parallel for
for (i = 0; i < n; i++)
 arr[i] = 0;

```

**1.2** What potential issue can arise from this code?

```

1 // Decrements element i of arr. n is a multiple of omp_get_num_threads()
2 #pragma omp parallel
3 {
4 int threadCount = omp_get_num_threads();
5 int myThread = omp_get_thread_num();
6 for (int i = 0; i < n; i++) {
7 if (i % threadCount == myThread) arr[i] -= 1;
8 }
9 }

```

**1.3** // Assume n holds the length of arr

```

2 double fast_product(double *arr, int n) {
3 double product = 1;
4 #pragma omp parallel for
5 for (int i = 0; i < n; i++) {
6 product *= arr[i];
7 }
8 return product;
9 }

```

(a) What is wrong with this code?

- (b) Fix the code using `#pragma omp critical`
- (c) Fix the code using `#pragma omp reduction(operation: var)`.

## 2 Amdahl's Law

In the programs we write, there are sections of code that are naturally able to be sped up. However, there are likely sections that just can't be optimized any further to maintain correctness. In the end, the overall program speedup is the number that matters, and we can determine this using Amdahl's Law:

$$\text{True Speedup} = \frac{1}{S + \frac{1-S}{P}}$$

where  $S$  is the non-spiced-up part and  $P$  is the speedup factor (determined by the number of cores, threads, etc.).

**2.1** You are going to run a convolutional network to classify a set of 100,000 images using a computer with 32 threads. You notice that 99% of the execution of your project code can be parallelized on these threads. What is the speedup?

**2.2** You run a profiling program on a different program to find out what percent of this program each function takes. You get the following results:

| Function | % Time |
|----------|--------|
| f        | 30%    |
| g        | 10%    |
| h        | 60%    |

- (a) We don't know if these functions can actually be parallelized. However, assuming all of them can be, which one would benefit the most from parallelism?
- (b) Let's assume that we verified that your chosen function can actually be parallelized. What speedup would you get if you parallelized just this function with 8 threads?