



# Advanced OpenMP

Lecture 6: Nested parallelism

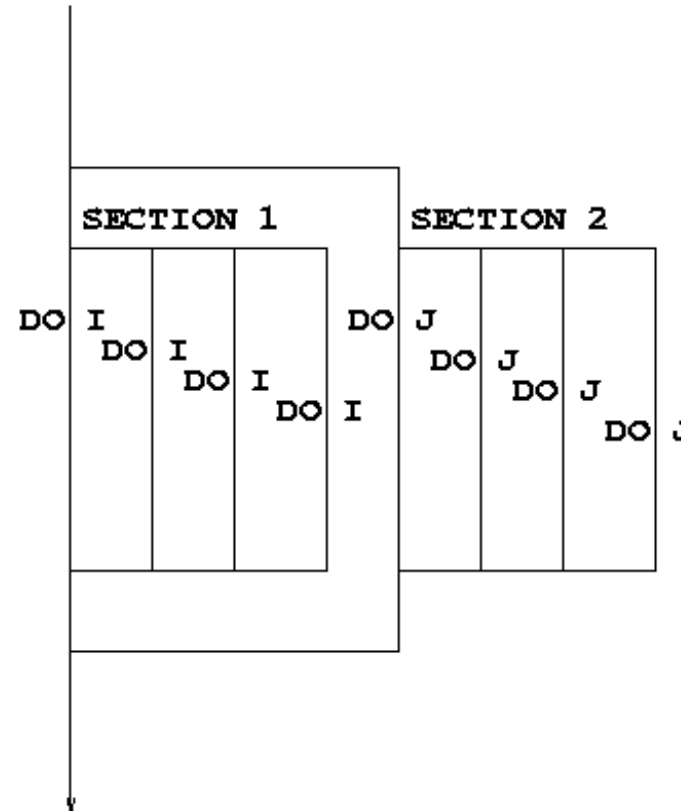
---

- Nested parallelism is supported in OpenMP.
- If a PARALLEL directive is encountered within another PARALLEL directive, a new team of threads will be created.
- This is enabled with the **OMP\_NESTED** environment variable or the **OMP\_SET\_NESTED** routine.
- If nested parallelism is disabled, the code will still be executed, but the inner teams will contain only one thread.

# Nested parallelism (cont)

Example:

```
!$OMP PARALLEL
!$OMP SECTIONS
!$OMP SECTION
!$OMP PARALLEL DO
  do i = 1,n
    x(i) = 1.0
  end do
!$OMP SECTION
!$OMP PARALLEL DO
  do j = 1,n
    y(j) = 2.0
  end do
!$OMP END SECTIONS
!$OMP END PARALLEL
```





- Not often needed, but can be useful to exploit non-scalable parallelism (SECTIONS).
  - Also useful if the outer level does not contain enough parallelism
- Note: nested parallelism isn't supported in some implementations (the code will execute, but as if `OMP_NESTED` is set to `FALSE`).
  - turns out to be hard to do correctly without impacting performance significantly.
  - don't enable nested parallelism unless you are using it!

- Can use the environment variable

```
export OMP_NUM_THREADS=2,4
```

- Will use 2 threads at the outer level and 4 threads for each of the inner teams.
- Can use `omp_set_num_threads()` or the `num_threads` clause on the parallel region.

## omp\_set\_num\_threads()

- Useful if you want inner regions to use different numbers of threads:

```
CALL OMP_SET_NUM_THREADS(2)
!$OMP PARALLEL DO
    DO I = 1,4
CALL OMP_SET_NUM_THREADS(innerthreads(i))
!$OMP PARALLEL DO
    DO J = 1,N
        A(I,J) = B(I,J)
    END DO
    END DO
```

- The value set overrides the value(s) in the environment variable OMP\_NUM\_THREADS

- One way to control the number of threads used at each level is with the NUM\_THREADS clause:

```
!$OMP PARALLEL DO NUM_THREADS(2)
  DO I = 1,4
!$OMP PARALLEL DO NUM_THREADS(innerthreads(i))
  DO J = 1,N
    A(I,J) = B(I,J)
  END DO
END DO
```

- The value set in the clause overrides the value in the environment variable OMP\_NUM\_THREADS and that set by `omp_set_num_threads()`

- Can also control the maximum number of threads running at any one time.

```
export OMP_THREAD_LIMIT=64
```

- ...and the maximum depth of nesting

```
export OMP_MAX_ACTIVE_LEVELS=2
```

or call

```
omp_set_max_active_levels()
```



- `omp_get_level()`
  - returns the level of parallelism of the calling thread
  - returns 0 in the sequential part
- `omp_get_active_level()`
  - returns the level of parallelism of the calling thread, ignoring levels which are inactive (teams only contain one thread)
- `omp_get_ancestor_thread_num(level)`
  - returns the thread ID of this thread's ancestor at a given level
  - ID of my parent:  
`omp_get_ancestor_thread_num(omp_get_level()-1)`
- `omp_get_team_size(level)`
  - returns the number of threads in this thread's ancestor team at a given level

- For perfectly nested rectangular loops we can parallelise multiple loops in the nest with the **collapse** clause:

```
#pragma omp parallel for collapse(2)
for (int i=0; i<N; i++) {
    for (int j=0; j<M; j++) {
        .....
    }
}
```

- Argument is number of loops to collapse starting from the outside
- Will form a single loop of length  $N \times M$  and then parallelise and schedule that.
- Useful if  $N$  is  $O(\text{no. of threads})$  so parallelising the outer loop may not have good load balance
- More efficient than using nested teams

- Note that barriers (explicit or implicit) only affect the innermost enclosing parallel region.
- No way to have a barrier across multiple teams
- In contrast, critical regions, atomics and locks affect all the threads in the program
- If you want mutual exclusion within teams but not between them, need to use locks (or atomics).