



# Advanced OpenMP

Lecture 5: Tasks

---

- The task construct defines a section of code
- Inside a parallel region, a thread encountering a task construct will package up the task for execution
- Some thread in the parallel region will execute the task at some point in the future
- Tasks can be nested: i.e. a task may itself generate tasks.

Syntax:

Fortran:

```
! $OMP TASK [clauses]
```

```
structured block
```

```
! $OMP END TASK
```

C/C++:

```
#pragma omp task [clauses]
```

```
structured-block
```

- At thread barriers (explicit or implicit)
  - applies to all tasks generated in the current parallel region up to the barrier
  
- At taskwait directive
  - i.e. Wait until all tasks defined in the current task have completed.
  - Fortran: **!\$OMP TASKWAIT**
  - C/C++: **#pragma omp taskwait**
  
  - Note: applies only to tasks generated in the current task, not to “descendants” .

```
p = listhead ;  
while (p) {  
    process (p) ;  
    p=next (p) ;  
}
```

- Classic linked list traversal
- Do some work on each item in the list
- Assume that items can be processed independently
- Cannot use an OpenMP loop directive

# Parallel pointer chasing

```
#pragma omp parallel
{
  #pragma omp single private(p)
  {
    p = listhead ;
    while (p) {
      #pragma omp task firstprivate(p)
      {
        process (p) ;
      }
      p=next (p) ;
    }
  }
}
```

Only one thread packages tasks

makes a copy of **p** when the task is packaged

```
#pragma omp parallel
{
    #pragma omp for private(p)
    for ( int i =0; i <numlists ; i++) {
        p = listheads [ i ] ;
        while (p ) {
            #pragma omp task firstprivate(p)
            {
                process (p) ;
            }
            p=next (p ) ;
        }
    }
}
```

All threads package tasks

- The default for tasks is usually firstprivate, because the task may not be executed until later (and variables may have gone out of scope).
- Variables that are shared in all constructs starting from the innermost enclosing parallel construct are shared.

```
#pragma omp parallel shared(A) private(B)
{
    ...
    #pragma omp task
    {
        int C;
        compute(A, B, C);
    }
}
```

A is shared  
B is firstprivate  
C is private



- Things can get rather complicated with nested tasks....

```
#pragma omp task private(B)
{
    B = ...
    #pragma omp task shared (B)
    {
        compute(B);
    }
    ...
#pragma omp taskwait
}
```

- Every outer task has its own copy of B
- All inner tasks use their parent task's copy of B
- Taskwait ensures these don't go out of scope....

- Binary tree of tasks
- Traversed using a recursive function
- A task cannot complete until all tasks below it in the tree are complete

```
void postorder(node *p) {  
    if (p->left)  
        #pragma omp task  
        { postorder(p->left); }  
    if (p->right)  
        #pragma omp task  
        { postorder(p->right); }  
    #pragma omp taskwait  
    process(p->data);  
}
```

← Parent task suspended until children tasks complete

- Certain constructs have task scheduling points at defined locations within them
- When a thread encounters a task scheduling point, it is allowed to suspend the current task and execute another (called *task switching*)
- It can then return to the original task and resume

```
#pragma omp single
{
    for (i=0; i<ONEZILLION; i++)
        #pragma omp task
            process(item[i]);
}
```

- Risk of generating too many tasks
- Generating task will have to suspend for a while
- With task switching, the executing thread can:
  - execute an already generated task (draining the “*task pool*”)
  - execute the encountered task

```
#pragma omp single
{
    #pragma omp task untied
    for (i=0; i<ONEZILLION; i++)
        #pragma omp task
        process(item[i]);
}
```

- Eventually, too many tasks are generated
- Generating task is suspended and executing thread switches to a long and boring task
- Other threads get rid of all already generated tasks, and start starving...
- With thread switching, the generating task can be resumed by a different thread, and starvation is over
- Too strange to be the default: the programmer is responsible!
  - beware of using thread IDs and threadprivate data!

- Getting the data attribute scoping right can be quite tricky
  - default scoping rules different from other constructs
  - as ever, using **default(none)** is a good idea
- Don't use tasks for things already well supported by OpenMP
  - e.g. standard do/for loops
  - the overhead of using tasks is greater
- Don't expect miracles from the runtime
  - best results usually obtained where the user controls the number and granularity of tasks

# Parallel pointer chasing again

```
#pragma omp parallel
{
    #pragma omp single private(p)
    {
        p = listhead ;
        while (p) {
            #pragma omp task firstprivate(p)
            {
                process (p,nitems) ;
            }
            for (i=0; i<nitems &&p; i++){
                p=next (p) ;
            }
        }
    }
}
```

process  
nitems at  
a time

skip nitems ahead  
in the list

- Task can be *included* (executed now by the thread that encounters them) or *deferred* (executed later by some thread in the team).
- The task construct can take an **if (expr)** clause, which if the expression evaluates to false, means the task will be included.
- The task construct can take a **final (expr)** clause, which if the expression evaluates to true, means any tasks generated inside this task will be included.
- The task construct can take a **mergeable** clause, which indicate it can be safely executed by reusing its parent data environment. Most useful if used in conjunction with **final**



- Mandelbrot example using tasks.