

CS 5220: Shared memory programming

David Bindel

2017-09-28

OpenMP: Open spec for MultiProcessing

- Standard API for multi-threaded code
 - Only a spec — multiple implementations
 - Lightweight syntax
 - C or Fortran (with appropriate compiler support)
- High level:
 - Preprocessor/compiler directives (80%)
 - Library calls (19%)
 - Environment variables (1%)
- Basic syntax: `#pragma omp construct [clause ...]`
 - Usually affects structured block (one way in/out)
 - OK to have `exit()` in such a block

- Environmental inquiries with `omp_get_*` functions
- Creating parallel regions with `#pragma omp parallel`
- Annotations for variables (shared, private, reduction)
- Synchronization via critical sections, atomic ops, barriers
- Today: Work sharing, tasks, and some examples

Work sharing constructs split work across a team

- **Parallel for**: split by loop iterations
- **sections**: non-iterative tasks
- **single**: only one thread executes (synchronized)
- **master**: master executes, others skip (no sync)

Parallel iteration

Idea: Map **independent** iterations onto different threads

```
1  #pragma omp parallel for
2  for (int i = 0; i < N; ++i)
3      a[i] += b[i];
4
5  #pragma omp parallel
6  {
7      // Stuff can go here...
8      #pragma omp for
9      for (int i = 0; i < N; ++i)
10         a[i] += b[i];
11 }
```

Implicit barrier at end of loop (unless **nowait** clause)

The iteration can also go across a higher-dim index set

```
1  #pragma omp parallel for collapse(2)
2  for (int i = 0; i < N; ++i)
3      for (int j = 0; j < M; ++j)
4          a[i*M+j] = foo(i,j);
```

Restrictions

- **for** loop must be in “canonical form”
 - Loop var is an integer, pointer, random access iterator (C++)
 - Test compares loop var to loop-invariant expression
 - Increment or decrement by a loop-invariant expression
 - No code between loop starts in **collapse** set
 - Needed to split iteration space (maybe in advance)
- Iterations should be independent
 - Compiler may not stop you if you screw this up!
- Iterations may be assigned out-of-order on one thread!
 - Unless the loop is declared **monotonic**

Reduction loops

How might we parallelize something like this?

```
1     double sum = 0;  
2     for (int i = 0; i < N; ++i)  
3         sum += big_hairy_computation(i);
```


Reduction loops

How might we parallelize something like this?

```
1     double sum = 0;  
2     #pragma omp parallel for reduction(+:sum)  
3     for (int i = 0; i < N; ++i)  
4         sum = big_hairy_computation(i);
```

OK, what about something like this?

```
1     for (int i = 0; i < N; ++i) {  
2         int result = big_hairy_computation(i);  
3         add_to_queue(q, result);  
4     }
```

Work is *mostly* independent, but not wholly.

Solution: **ordered** directive in loop with **ordered** clause

```
1     #pragma omp parallel for ordered
2     for (int i = 0; i < N; ++i) {
3         int result = big_hairy_computation(i);
4         #pragma ordered
5         add_to_queue(q, result);
6     }
```

Ensures the **ordered** code executes in loop order.

SIMD loops

As of OpenMP 4.0:

```
1     #pragma omp parallel simd reduction(+:sum) aligned(a:64)
2     for (int i = 0; i < N; ++i) {
3         a[i] = b[i] * c[i];
4         sum = sum + a[i];
5     }
```

Can also declare vectorized functions:

```
1     #pragma omp declare simd
2     float myfunc(float a, float b, float c)
3     {
4         return a*b + c;
5     }
```

Other parallel work divisions

- `sections`: like `cobegin/coend`
- `single`: do only in one thread (e.g. I/O)
- `master`: do only in master thread; others skip

Sections

```
1     #pragma omp parallel
2     {
3         #pragma omp sections nowait
4         {
5             #pragma omp section
6             do_something();
7
8             #pragma omp section
9             and_something_else();
10
11            #pragma omp section
12            and_this_too();
13            // No implicit barrier here
14        }
15        // Implicit barrier here
16    }
```

sections nowait to kill barrier.

Task-based parallelism

- Work-sharing so far is rather limited
 - Work cannot be produced/consumed dynamically
 - Fine for data parallel array processing...
 - ... but what about tree walks and such?
- Alternate approach (OpenMP 3.0+): Tasks

Task involves:

- Task construct: **task** directive plus structured block
- Task: Task construct + data

Tasks are handled by run time, complete at barriers or **taskwait**.

Example: List traversal

```
1 #pragma omp parallel
2 {
3     #pragma omp single nowait
4     {
5         for (link_t* link = head; link; link = link->next)
6             #pragma omp task firstprivate(link)
7             process(link);
8     }
9     // Implicit barrier
10 }
```

One thread generates tasks, others execute them.

Example: Tree traversal

```
1 int tree_max(node_t* n)
2 {
3     int lmax, rmax;
4     if (n->is_leaf)
5         return n->value;
6
7     #pragma omp task shared(lmax)
8         lmax = tree_max(n->l);
9     #pragma omp task shared(rmax)
10        rmax = tree_max(n->r);
11    #pragma omp taskwait
12
13    return max(lmax, rmax);
14 }
```

The `taskwait` waits for all child tasks.

Task dependencies

What happens if one task produces what another needs?

```
1     #pragma omp task depend(out:x)
2     x = foo();
3     #pragma omp task depend(in:x)
4     y = bar(x);
```

Topics not addressed

- Low-level synchronization (locks, `flush`)
- OpenMP 4.x constructs for accelerator interaction
- A variety of more specialized clauses

See <http://www.openmp.org/>

Some examples (at board)

What are different ways to organize these:

- Dot product?
- Monte Carlo computation with adaptive termination?
- Wave equation time stepper?