

An introduction to The openMP programming environment

- Introduction : the openMP model
- Basic syntax
- A few examples
- See also the following site for many resources:

<http://openmp.org>

4-1


Threads and the openMP model

- openMP implements the Fork-Join model
- Makes it easy to parallelize loops or parallel sections of codes

Note: A thread is a stream of instructions that can be executed in parallel.

Pros: Arguably the simplest approach to parallel programming.

Cons: Limited to SMPs [Shared memory computers]



4-2

– openMP

4-2

The openMP approach

- Use C (or C++, Fortran, ...) and add directives / pragmas to:
 - * Indicate parallel loops,
 - * Parallel regions of code, ..
 - * .. and more
- Plus a few library routines [e.g., OMP_GET_THREAD_NUM()]
- Intrinsically designed for Shared Memory SMP machines.
- Portable – supported by all High Performance computer vendors:
<http://openmp.org>
- ... and implemented in GNU compilers (gcc,...).

4-3

– openMP

4-3

- Directives/ pragmas:

In C: `#pragma omp ... directives`

In Fortran: `!$OMP ... directives`

```
#pragma omp parallel
{
  ...
  // structured block
  ...
}
```

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

- These notes will illustrate only a few directives

See <http://openmp.org/> for additional details:

- Reference guide for a quick overview
- Specifications [a pdf file] for details

4-4

– openMP

4-4

Basic functions

```
omp_get_thread_num() - get thread number
omp_set_num_threads(nthreads) - set # of threads
omp_get_num_threads() - get number of threads used
```

Example:

```
#include <omp.h>
int omp_get_thread_num();
int main(){
    # pragma omp parallel
    {
        printf("Thread number : %d\n",
            omp_get_thread_num());
    }
}
```

4-5

– openMP

4-5

1. Compile with `gcc -o test.ex -fopenmp test.c`
2. Set number of threads with environment variable:
`setenv OMP_NUM_THREADS 4`
3. Run

```
./test.ex
Thread number: 0
Thread number: 3
Thread number: 2
Thread number: 1
```

4-6

– openMP

4-6

Hello World in openMP: pragma parallel

 1 Compile and run this other version of the previous example

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 int main () {
4     int i;
5     int omp_get_thread_num();
6     printf("Entering parallel threads: \n");
7     #pragma omp parallel
8     {
9         i = omp_get_thread_num();
10        printf(" -->> Hello from thread : %d \n",i);
11    }
12    printf(" <<-- Out of threads \n");
13 }
```

4-7

– openMP

4-7

Hello World in openMP: pragma parallel for

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 int main () {
4     int i, p;
5     int omp_get_thread_num();
6     printf(" Entering parallel threads: \n");
7     #pragma omp parallel for
8     for (i=0;i<12; i++) {
9         p = omp_get_thread_num();
10        printf(" -->> Hello from thread : %d \n",p);
11    }
12    printf(" <<-- Out of threads \n");
13 }
```

4-8

– openMP

4-8

 compile and run:

```
gcc -fopenmp omp_hello.c
```

➤ Can set the number of threads from environment variable...

```
setenv OMP_NUM_THREADS 4
```

➤ ... or in the code with the command
`omp_set_num_threads(nthreads)`

➤ This freezes the number of threads [takes precedence over environment variable OMP_NUM_THREADS]

 What is the difference between the two examples?

4-9

– openMP


4-9

Scoping of variables

➤ Variables can be shared among threads as in

```
#pragma omp parallel shared(var1, var2, ...)
```

➤ Beware of racing between variables.. [no guaranteed order of modifications]

 What can happen if several threads write to the same shared variable? See situation in following example.

```
#define N_MAX 10000
int main() {
    int i;
    double fx, fsum=0.0;
    #pragma omp parallel for
    for (i = 1; i <= N_MAX; i++) {
        fx = (double)i ;
        fsum += fx;
    }
    printf("-- sum %f \n", fsum);
}
```

Program race.c:

4-10

– openMP

4-10

Private variables

Variables can be private – local copies of variables made for each thread – Note: when copies are made they are **not** initialized

```
#pragma omp private(var1, var2, ...)
```

➤ Can set default for scoping of variable by

```
#pragma omp default(DEF)
```

where DEF == one of private, shared, or none.

➤ If no default is set, and there is no explicit clause for scoping, variables are assumed to be shared

4-11

– openMP

4-11

Example: Dot-Product

```
omp_set_num_threads(nt); // nt = # threads
m = n/nt; // assumes n divisible by nt (!)
#pragma omp parallel for private(t, i1, i2, i)
for (it = 0; it < nt; it++) {
    i1 = it*m;
    i2 = i1+m;
    if (i2 > n) i2 = n;
    t = 0.0;
    for (i = i1; i < i2; i++)
        t += x[i]*y[i];
    tt[it] = t;
}
t = 0.0;
for (it = 0; it < nt; it++)
    t += tt[it];
```

4-12

– openMP

4-12

Critical sections

- Solutions to race conditions: `critical` sections which permit a code fragment to be executed by one thread only

```
#pragma omp critical [name]
{
    ... structured block ...
}
```

- Go back to program `race.c` seen earlier..

Here is how it can be corrected..

- First declare `fx` as private..
- Then summation should be critical. [loss of parallelism]

4-13

– openMP

4-13

Program `race_cor.c`:

```
1  #define N_MAX 10000
2  int main(){
3      int i;
4      double fx, fsum;
5      #pragma omp parallel for private(fx)
6      for (i = 1; i <= N_MAX; i++) {
7          fx = (double)i ;
8          #pragma omp critical
9              fsum += fx;
10     }
11     printf("-- sum %f \n", fsum);
12 }
```

- Better solution: Reduction operation


4-14

– openMP

4-14

Reduction operations

- A reduction does a global operation (e.g. a sum) on an array down to one single result. For example $a = \sum_{i=0}^{n-1} x_i$ or $a = \max_{i=0}^{n-1} x_i$, ..



Reduction

Clause syntax: `reduction(<op>: variable)`

4-15

– openMP

4-15

Example: Dot product computation seen earlier

```
omp_set_num_threads(nt);
...
t = 0.0;
#pragma omp parallel for reduction(+:t)
for (i = 0; i < n; i++)
    t += x[i]*y[i];
```

- Private copy of `t` (in clause) is created for each thread.
- At the end of reduction, reduction operation `+` (in clause) is applied to private variable `t` (in clause) –
- Result of this reduction written to 'master' thread (shared variable) `t`

4-16

– openMP

4-16

Sections

- Each section executed by one thread
- Cannot branch into and out of block of sections

```
#pragma omp sections
{
  // section 1
  #pragma omp
  {
    block
  }
  // section 2
  #pragma omp
  {
    block
  }
  ...
}
```

4-17

– openMP

4-17

Schedule clauses

- Consider the example

```
#pragma omp parallel for
for (i=0; i<n; i++) {
  /*--- cost of function varies
        a lot with i */
  x[i] = some_function(i);
}
```

Result: poor load balancing. Solution schedule work dynamically.

schedule (type [,chunk])

type is one of static, dynamic, guided or runtime

4-18

– openMP

4-18

runtime = set by an environment variable

setenv OMP_SCHEDULE ‘‘type,chunk’’ command.

```
#pragma omp parallel for schedule(dynamic)
for (i=0; i<n; i++) {
  /* cost of function varies a lot with i */
  x[i] = some_function(i);
}
```

4-19

– openMP

4-19

A few runtime functions

- omp_get_num_threads () returns current number of threads

Note: this is always one in a sequential section.

- omp_set_num_threads (int) sets number of threads in code

2 other methods for this: environment variable (seen before), and clause [#pragma omp parallel num.threads(3)]

- omp_get_num_procs () returns current number of processors available

4-20

– openMP

4-20