

OpenMP

Open Multi Processing

Philipp Quach


Seminar "Effiziente Programmierung"
University of Hamburg

December 15th 2016


Content

- 1 Introduction
 - Parallel Programming
 - Introduction to OpenMP
- 2 Features
 - Parallel construct
 - Loops
 - Sections
 - Shared, unshared variables
 - Offloading
 - Thread-Safety
 - Synchronization
- 3 Compilation
- 4 Performance
- 5 Conclusion
 - How good is OpenMP
 - Summary
- 6 Sources

- Simpler code is serial
 - One instruction at a time
 - executed one after the other
 - run on a single machine



- Performant code should be parallelized
 - concurrent execution



[Bar]

Introduction to OpenMP

- Supports C, C++ and Fortran
- Comes with the compiler
- Programmer directed
- High-level

Low vs high-level approach

PThreads (low-level)

```
1  #include<stdio.h>
2  #include<pthread.h>
3
4  void* say_hello(void* data)
5  {
6      char *str;
7      str = (char*)data;
8      printf("%s\n", str);
9  }
10
11 void main()
12 {
13     pthread_t t1,t2;
14     pthread_create(&t1,NULL,say_hello,"Hello Seminar");
15     pthread_create(&t2,NULL,say_hello,"Hello Seminar");
16     pthread_join(t1,NULL);
17     pthread_join(t2,NULL);
18 }
```

Low vs high-level approach

OpenMP (high-level)

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 void main()
5 {
6     #pragma omp parallel num_threads(2)
7     printf("Hello Seminar\n");
8 }
```

```
PhilippQuach@DESKTOP-VEMDLHR /cygdrive/c
$ gcc -fopenmp -o omp omp.c; ./omp
Hello Seminar
Hello Seminar

PhilippQuach@DESKTOP-VEMDLHR /cygdrive/c
$ gcc -o pthreads pthreads.c; ./pthreads
Hello Seminar
Hello Seminar
```

Syntax

- Preprocessor directive begins with `#pragma omp`
- Followed by a specification as to what feature is being applied
- The parallelism is applied to the block of code following the preprocessor directive

```
1  foo() {  
2      #pragma omp <command specifier>  
3      {  
4          //some block of code that runs parallel  
5      }  
6  }
```

- `!$OMP <COMMAND SPECIFIER>` in Fortran
- unknown pragmas are ignored by the compiler

omp.h

- `#include <omp.h>`
- provides many helpful functions
 - e.g. `omp_get_thread_num()`
- not required to run OpenMP code

Parallel construct

■ #pragma omp parallel

```
1  int main(void) {  
2      #pragma omp parallel  
3      printf("hello Seminar\n");  
4  
5      return EXIT_SUCCESS;  
6  }
```

- Creates a team of n threads
- n usually depends on the number of cpu cores unless specified otherwise
- Parallelized block is executed once by every thread

```
PhilippQuach@DESKTOP-VEMDLHR /cygdrive/c/Users/PhilippQuach/Documents/OpenMP  
$ gcc -fopenmp -o omp omp.c; ./omp  
Hello Seminar  
Hello Seminar  
Hello Seminar  
Hello Seminar
```

num_threads

- `#pragma omp parallel num_threads(int)`
- alternative: `omp_set_num_threads(int)` from `omp.h`

```
1 int main(void) {  
2     #pragma omp parallel num_threads(3)  
3     printf("hello Seminar\n");  
4  
5     return EXIT_SUCCESS;  
6 }
```

- Let's you specify the number of threads to be created

```
PhilippQuach@DESKTOP-VEMDLHR /cygdrive/c/Users/PhilippQuach/Documents/OpenMP  
$ gcc -fopenmp -o omp omp.c; ./omp  
hello Seminar  
hello Seminar  
hello Seminar
```

Parallel if

- #pragma omp parallel if(bool)

```
1  int main(void) {
2      #pragma omp parallel if(0)
3      printf("hello Seminar\n");
4
5      return EXIT_SUCCESS;
6  }
```

- parallelizes only if the boolean within the if clause is true

```
PhilippQuach@DESKTOP-VEMDLHR /cygdrive/c/Users/PhilippQuach/Documents/OpenMP
$ gcc -fopenmp -o omp omp.c; ./omp
hello Seminar
```

For construct

■ #pragma omp for

```
1  int main(void) {
2      #pragma omp parallel num_threads(2)
3      {
4          #pragma omp for
5          for(int n=0; n<10; ++n)
6              {
7                  printf(" %d", n);
8              }
9      }
10
11     return EXIT_SUCCESS;
12 }
```

■ Each thread of the active team handles a different part of the loop

```
PhilippQuach@DESKTOP-VEMDLHR /cygdrive/c/Users/PhilippQuach/Documents/OpenMP
$ gcc -fopenmp -o omp omp.c; ./omp
0 5 1 6 2 7 3 8 4 9
```

Parallel for

■ #pragma omp parallel for

```
1  int main(void) {
2      #pragma omp parallel for
3      for(int n=0; n<10; ++n)
4      {
5          printf(" %d", n);
6      }
7
8      return EXIT_SUCCESS;
9  }
```

- Combines #pragma omp parallel and #pragma omp for into one line
- Creates a team of threads and assigns each thread a part of the loop

Schedule

- static (default), dynamic, auto, guided, runtime

```

1  ...//2 active threads
2  #pragma omp for schedule(static)
3  for(int n=0; n<10; ++n) printf(" %d", n);

```

```

$ gcc -fopenmp -o omp omp.c; ./omp
0 5 1 6 2 7 3 8 4 9

```

```

1  //2 threads
2  #pragma omp for schedule(dynamic, 3)
3  for(int n=0; n<10; ++n) printf(" %d", n);

```

```

$ gcc -fopenmp -o omp omp.c; ./omp
0 3 1 4 2 5 6 9 7 8

```

Ordered

```
1 //2 threads
2 #pragma omp for ordered schedule(static)
3 for(int n=0; n<10; ++n) {
4     printf(" %d", n);
5 }
```

```
$ gcc -fopenmp -o omp omp.c; ./omp
0 5 1 6 2 7 3 8 4 9
```

```
1 //2 threads
2 #pragma omp for ordered schedule(static)
3 for(int n=0; n<10; ++n) {
4     #pragma omp ordered
5     printf(" %d", n);
6 }
```

```
$ gcc -fopenmp -o omp omp.c; ./omp
0 1 2 3 4 5 6 7 8 9
```

Nested loops and the collapse clause

```

1 //2 threads
2 #pragma omp for
3 for(int n=0; n<3; ++n) {
4     for(int m=0; m<2; ++m) {
5         printf("(%d%d)", n, m);
6     }
7 }

```

```

$ gcc -fopenmp -o omp omp.c; ./omp
(00)(20)(01)(21)(10)(11)

```

```

1 //2 threads
2 #pragma omp for collapse(2)
3 for(int n=0; n<3; ++n) {
4     for(int m=0; m<2; ++m) {
5         printf("(%d%d)", n, m);
6     }
7 }

```

```

$ gcc -fopenmp -o omp omp.c; ./omp
(00)(11)(01)(20)(10)(21)

```


Sections

```
1 //3 threads
2 #pragma omp sections
3 {
4     {
5         printf("a ");
6     }
7     #pragma omp section
8     {
9         printf("b1 ");
10        printf("b2 ");
11    }
12    #pragma omp section
13    {
14        printf("c ");
15    }
16 }
```

```
$ gcc -fopenmp -o omp omp.c; ./omp;./omp ;./omp ;./omp ;./omp ;./omp
a b1 c b2
b1 a c b2
a b1 c b2
b1 a c b2
c a b1 b2
a c b1 b2
```

Shared, unshared variables

- shared: One variable shared by all threads (default)
- private: Each thread has their own variable of this name

```

1  int main(void) {
2      int m, l=0;
3      #pragma omp parallel for num_threads(2) private(l) shared(m)
4      for(int n=0; n<10;n++) {
5          l++;
6          m++;
7          printf("(%d,%d)", l, m);
8      }
9  }

```

```

$ gcc -fopenmp -o omp omp.c; ./omp
(1,1)(103625,2)(2,3)(103626,4)(3,5)(103627,6)(4,7)(103628,8)(5,9)(103629,10)

```

Firstprivate

```

1  int main(void) {
2      int m, l=0;
3      #pragma omp parallel for num_threads(2) firstprivate(l)
      ↪ shared(m)
4      for(int n=0; n<10;n++) {
5          l++;
6          m++;
7          printf("(%d,%d)", l, m);
8      }
9  }

```

```

$ gcc -fopenmp -o omp omp.c; ./omp
(1,1)(1,2)(2,3)(2,4)(3,5)(3,6)(4,7)(4,8)(5,9)(5,10)

```

Offloading

- Execution also on other hardware than the computers CPU

```
1 #pragma omp target device(device_number)
2 {
3     //executed on the device with the number specified
4 }
```

- `omp.h` provides helpful methods e.g. to set a default device or find out device numbers

Atomic

```
1  int count = 0;
2  #pragma omp parallel num_threads(100)
3  {
4      //#pragma omp atomic
5      count++;
6  }
7  printf("Number of threads: %d\n", count);
```

Not atomic:

```
$ gcc -fopenmp -o omp omp.c; ./omp ; ./omp ; ./omp
Number of threads: 98
Number of threads: 100
Number of threads: 99
```

Reduction

```
1  #include <stdio.h>
2  #include <stdlib.h>
3
4  int main()
5  {
6      int count = 0;
7      #pragma omp parallel num_threads(100) reduction(+:count)
8      {
9          count++;
10     }
11     printf("Number of threads: %d\n", count);
12     return 0;
13 }
```

Critical

```
1  #pragma omp parallel num_threads(2)
2  {
3      if(omp_get_thread_num() == 0){
4          #pragma omp critical(loop)
5          for(int n = 0; n < 5; n++) printf("a");
6      } else {
7          #pragma omp critical(loop)
8          for(int n = 0; n < 5; n++) printf("b");
9      }
10 }
```

```
$ gcc -fopenmp -o omp omp.c; ./omp
aaaaabbbbb
```

Barrier

```

1  #pragma omp parallel num_threads(2)
2  {
3      if(omp_get_thread_num() == 1)
4      {
5          for(int n = 0; n < 10; n++) printf("n ");
6      }
7      #pragma omp barrier
8      printf("\npast the barrier");
9  }

```

With barrier:

```

$ gcc -fopenmp -o omp omp.c;./omp
n n n n n n n n n n
past the barrier
past the barrier

```

Without barrier:

```

$ gcc -fopenmp -o omp omp.c;./omp
n
past the barrier n n n n n n n n
past the barrier

```


Nowait

```

1  #pragma omp parallel num_threads(2)
2  {
3      #pragma omp for nowait
4      for(int n = 0; n < 10; n++){
5          printf("%d", omp_get_thread_num());
6          if(omp_get_thread_num() == 1) printf("");
7      }
8      printf("\ndone with the loop");
9  }

```

With nowait:

```

$ gcc -fopenmp -o omp omp.c; ./omp
01001001
done with the loop11
done with the loop

```

Without nowait:

```

$ gcc -fopenmp -o omp omp.c; ./omp
0100100111
done with the loop
done with the loop

```

Requirements

- Compiler supporting OpenMP
- Set compiler flag for OpenMP e.g. `-fopenmp`
 - Produces serial code otherwise
- Link the runtime library `libgomp-1.dll`

Additional compilation

Additionally to the usual compilation:

- Reads omp directives and checks for correctness
- Substitution:
 - Replace sections by Do- and For-constructs
 - Implicit to explicit barrier
- Handles memory
- Applies some optimization
- Creates multithreading code from omp constructs
- Outlines parallel region to function

Transformed code example

Original:

```
1 void main(){
2     #pragma omp parallel
3     {
4         #pragma omp for
5         for( i = 0; i < n; i++ ){...}
6     }
7 }
```

Transformed:

```
1 void outlined(...){
2     tid = ompc_get_thread_num();
3     ompc_static_init(tid, lower, upper, incr, ..);
4     for( i = lower; i < upper; i += incr ){ ... }
5     ompc_barrier();
6 }
7
8 void main(){
9     __ompc_fork(..., &outlined, ...);
10 }
```

Parallel Overhead

- Time spent coordinating threads etc.
 - Initializing threads
 - Terminating threads
 - Coordination such as synchronization
- Aim: Minimize overheads

Speedup

- OpenMP uses a thread-pool
 - Threads are created once
 - Once done with their work, return to dock
 - Then wait for new work
- Speedup over serial code can vary strongly
- $Speedup(P) = \frac{T_{Serial}(P)}{T_{Elapsed}(P)} = \frac{1}{\frac{T}{P} - f + 1 + O_P \cdot P}$ (simplified)
- $Efficiency(P) = \frac{Speedup(P)}{P}$

Bad usage makes it worse

```

1  #pragma omp parallel for private(j)
2  for(i=0;i<=100000;i++)
3  {
4      for(j=0;j<=100000;j++)
5      {
6          #pragma omp atomic
7          a++;
8      }
9  }
10 printf("%lld", a);

```

Serial:

```

$ gcc -o omp omp.c; time ./omp
10000200001
real    0m24.882s

```

Parallel:

```

$ gcc -fopenmp -o omp omp.c; time ./omp
2618243971
real    0m39.045s

```

Thread-safe:

```

$ gcc -fopenmp -o omp omp.c; time ./omp
10000200001
real    3m49.735s

```

Example speedup

Table: Matrix-Vector-Product

Size	Serial time	Parallel Time	Speedup
10000*10000	0.10	0.03	2.95
30000*30000	1.01	0.23	4.33
40000*40000	1.88	0.39	4.73

[App14]

- Execution with 4 cores, 8 logical processors/threads

Optimization

- Minimize Overheads
- Load balance: Threads should have similar runtime
- Thread-Safety causes waiting time
- Don't parallelize in inner loops
- Maximize parallel regions
- The ordered construct is slow
- Optimize barrier and nowait usage
- Avoid memory conflicts

How good is OpenMP

Pro:

- Target audience: general-purpose application programmers
 - portability, maintainability, convenience
- Highly effective for simple loop based code

Contra:

- Too narrow for complex code structures
- Doesn't optimize for the specific hardware the code runs on

Summary

- OpenMP is easy to use
 - Parallelize by adding a few lines
 - Not necessary to rewrite existing code
- Not a substitution of low-level APIs
- Although high level, the many features allow for flexible control
- Possible speedup depends on hardware
- Poor parallelization may even slow down, optimize well!

Sources I



Appentra.

A widely-used algebraic code: Parallel computation of matrix-vector product.

<http://www.appentra.com/>

[parallel-computation-of-matrix-vector-product/](http://www.appentra.com/parallel-computation-of-matrix-vector-product/), 2014.

[Online; accessed 8-December-2016].



Blaise Barney.

Introduction to Parallel Computing.

https://computing.llnl.gov/tutorials/parallel_comp/.

[Online; accessed 8-December-2016].



Lei Huang Barbara Chapman.

How OpenMP is Compiled.

[https:](https://iwomp.zih.tu-dresden.de/downloads/OpenMP-compilation.pdf)

[//iwomp.zih.tu-dresden.de/downloads/OpenMP-compilation.pdf](https://iwomp.zih.tu-dresden.de/downloads/OpenMP-compilation.pdf).

[Online; accessed 8-December-2016].



Keld Helsgaun.

How to Get Good Performance by Using OpenMP.

[http://www.akira.ruc.dk/~keld/teaching/IPDC_f10/Slides/](http://www.akira.ruc.dk/~keld/teaching/IPDC_f10/Slides/pdf4x/4_Performance.4x.pdf)

[pdf4x/4_Performance.4x.pdf](http://www.akira.ruc.dk/~keld/teaching/IPDC_f10/Slides/pdf4x/4_Performance.4x.pdf), 2010.

[Online; accessed 8-December-2016].

Sources II



Timothy G. Mattson.

How Good is OpenMP.

Hindawi Publishing Corporation, 2003.



Microsoft.

atomic.

<https://msdn.microsoft.com/de-de/library/8ztckdts.aspx>, 2016.

[Online; accessed 8-December-2016].



Joel Yliluoma.

Guide into OpenMP: Easy multithreading programming for C++.

<http://bisqwit.iki.fi/story/howto/openmp/>, 2007.

[Online; accessed 8-December-2016].