



Lecture 12: OpenMP

Abhinav Bhatele, Department of Computer Science



UNIVERSITY OF
MARYLAND

Announcements

- Use office hours
- If you foresee not being able to complete assignments for a valid reason, email me asap instead of after the deadline

saxpy (single precision $a*x+y$) example

```
for (int i = 0; i < n; i++) {  
    z[i] = a * x[i] + y[i];  
}
```

saxpy (single precision $a*x+y$) example

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

Overriding defaults using clauses

- Specify how data is shared between threads executing a parallel region
- `private(list)`
- `shared(list)`
- `default(shared | none)`
- `reduction(operator: list)`
- `firstprivate(list)`
- `lastprivate(list)`

<https://www.openmp.org/spec-html/5.0/openmpsul06.html#x139-5540002.19.4>

private clause

- Each thread has its own copy of the variables in the list
- Private variables are uninitialized when a thread starts
- The value of a private variable is unavailable to the master thread after the parallel region has been executed

default clause

- Determines the data sharing attributes for variables for which this would be implicitly determined otherwise

Anything wrong with this example?

```
val = 5;

#pragma omp parallel for private(val)
for (int i = 0; i < n; i++) {
    ... = val + 1;
}
```


Anything wrong with this example?

```
val = 5;
```

```
#pragma omp parallel for private(val)
for (int i = 0; i < n; i++) {
    ... = val + 1;
}
```

The value of val will not be available to threads inside the loop

Anything wrong with this example?

```
#pragma omp parallel for private(val)
for (int i = 0; i < n; i++) {
    val = i + 1;
}

printf("%d\n", val);
```

Anything wrong with this example?

```
#pragma omp parallel for private(val)
for (int i = 0; i < n; i++) {
    val = i + 1;
}

printf("%d\n", val);
```

The value of val will not be available to the master thread outside the loop

firstprivate clause

- Initializes each thread's private copy to the value of the master thread's copy

```
val = 5;

#pragma omp parallel for firstprivate(val)
for (int i = 0; i < n; i++) {
    ... = val + 1;
}
```

lastprivate clause

- Writes the value belonging to the thread that executed the last iteration of the loop to the master's copy
- Last iteration determined by sequential order

lastprivate clause

- Writes the value belonging to the thread that executed the last iteration of the loop to the master's copy
- Last iteration determined by sequential order

```
#pragma omp parallel for lastprivate(val)
for (int i = 0; i < n; i++) {
    val = i + 1;
}

printf("%d\n", val);
```

reduction(operator: list) clause

- Reduce values across private copies of a variable
- Operators: +, -, *, &, |, ^, &&, ||, max, min

```
#pragma omp parallel for
for (int i = 0; i < n; i++) {
    val += i;
}

printf("%d\n", val);
```

<https://www.openmp.org/spec-html/5.0/openmpsul07.html#x140-5800002.19.5>

reduction(operator: list) clause

- Reduce values across private copies of a variable
- Operators: +, -, *, &, |, ^, &&, ||, max, min

```
#pragma omp parallel for reduction(+: val)
for (int i = 0; i < n; i++) {
    val += i;
}

printf("%d\n", val);
```

<https://www.openmp.org/spec-html/5.0/openmpsul07.html#x140-5800002.19.5>

User-specified loop scheduling

- Schedule clause

`schedule (type[, chunk])`

- `type`: static, dynamic, guided, runtime
- `static`: iterations divided as evenly as possible ($\#iterations/\#threads$)
 - $chunk < \#iterations/\#threads$ can be used to interleave threads
- `dynamic`: assign a chunk size block to each thread
 - When a thread is finished, it retrieves the next block from an internal work queue
 - Default chunk size = 1

Other schedules

- guided: similar to dynamic but start with a large chunk size and gradually decrease it for handling load imbalance between iterations
- auto: scheduling delegated to the compiler
- runtime: use the `OMP_SCHEDULE` environment variable

<https://software.intel.com/content/www/us/en/develop/articles/openmp-loop-scheduling.html>

Calculate the value of $\pi = \int_0^1 \frac{4}{1+x^2}$

```
int main(int argc, char *argv[])
{
    ...

    n = 10000;

    h = 1.0 / (double) n;
    sum = 0.0;

    for (i = 1; i <= n; i += 1) {
        x = h * ((double)i - 0.5);
        sum += (4.0 / (1.0 + x * x));
    }
    pi = h * sum;

    ...
}
```

Calculate the value of $\pi = \int_0^1 \frac{4}{1+x^2}$

```
int main(int argc, char *argv[])
{
    ...

    n = 10000;
    h = 1.0 / (double) n;
    sum = 0.0;

    #pragma omp parallel for firstprivate(h) private(x) reduction(+: sum)
    for (i = 1; i <= n; i += 1) {
        x = h * ((double)i - 0.5);
        sum += (4.0 / (1.0 + x * x));
    }
    pi = h * sum;

    ...
}
```

Parallel region

- All threads execute the structured block

```
#pragma omp parallel [clause [clause] ... ]  
    structured block
```

- Number of threads can be specified just like the parallel for directive

Synchronization

- Concurrent access to shared data may result in inconsistencies
- Use mutual exclusion to avoid that
 - critical directive
 - atomic directive
- Library lock routines

<https://software.intel.com/content/www/us/en/develop/documentation/advisor-user-guide/top/appendix/adding-parallelism-to-your-program/replacing-annotations-with-openmp-code/adding-openmp-code-to-synchronize-the-shared-resources.html>



UNIVERSITY OF
MARYLAND

Abhinav Bhatele

5218 Brendan Iribe Center (IRB) / College Park, MD 20742

phone: 301.405.4507 / e-mail: bhatele@cs.umd.edu