

CS4961 Parallel Programming

Lecture 5: Data and Task Parallelism, cont. Data Parallelism in OpenMP

Mary Hall
September 7, 2010

09/07/2010

CS4961

Homework 2, Due Friday, Sept. 10, 11:59 PM

- To submit your homework:
 - Submit a PDF file
 - Use the "handin" program on the CADE machines
 - Use the following command:


```
"handin cs4961 hw2 <prob2file>"
```

Problem 1 (based on #1 in text on p. 59):

Consider the Try2 algorithm for "count3s" from Figure 1.9 of p.19 of the text. Assume you have an input array of 1024 elements, 4 threads, and that the input data is evenly split among the four processors so that accesses to the input array are local and have unit cost. Assume there is an even distribution of appearances of 3 in the elements assigned to each thread which is a constant we call NTPT. What is a bound for the memory cost for a particular thread predicted by the CTA expressed in terms of λ and NTPT.

09/07/2010

CS4961



Homework 2, cont

Problem 2 (based on #2 in text on p. 59), cont.:

Now provide a bound for the memory cost for a particular thread predicted by CTA for the Try4 algorithm of Fig. 114 on p. 23 (or Try3 assuming each element is placed on a separate cache line).

Problem 3:

For these examples, how is algorithm selection impacted by the value of NTPT?

Problem 4 (in general, not specific to this problem):

How is algorithm selection impacted by the value of λ ?

09/07/2010

CS4961



Today's Lecture

- Review Data and Task Parallelism
- Brief Overview of POSIX Threads
- Data Parallelism in OpenMP
 - Expressing Parallel Loops
 - Parallel Regions (SPMD)
 - Scheduling Loops
 - Synchronization
- Sources of material:
 - Jim Demmel and Kathy Yelick, UCB
 - Allan Snaveley, SDSC
 - Larry Snyder, Univ. of Washington

09/07/2010

CS4961



Definitions of Data and Task Parallelism

- Data parallel computation:
 - Perform the same operation to different items of data at the same time; the parallelism grows with the size of the data.
- Task parallel computation:
 - Perform distinct computations -- or tasks -- at the same time; with the number of tasks fixed, the parallelism is not scalable.
- Summary
 - Mostly we will study data parallelism in this class
 - Data parallelism facilitates very high speedups; and scaling to supercomputers.
 - Hybrid (mixing of the two) is increasingly common

09/07/2010

CS4961



Examples of Task and Data Parallelism

- Looking for all the appearances of "University of Utah" on the world-wide web
- A series of signal processing filters applied to an incoming signal
- Same signal processing filters applied to a large known signal
- Climate model from Lecture 1

09/07/2010

CS4961



Review: Predominant Parallel Control Mechanisms

Name	Meaning	Examples
Single Instruction, Multiple Data (SIMD)	A single thread of control, same computation applied across "vector" elts	Array notation as in Fortran 90: <code>A[1:n] = A[1:n] + B[1:n]</code>
Multiple Instruction, Multiple Data (MIMD)	Multiple threads of control, processors periodically synch	Parallel loop: <code>forall (i=0; i<n; i++)</code>
Single Program, Multiple Data (SPMD)	Multiple threads of control, but each processor executes same code	Processor-specific code: <code>if (\$myid == 0) { } }</code>

09/07/2010

CS4961



Programming with Threads

Several Thread Libraries

- PTHREADS is the Posix Standard
 - Solaris threads are very similar
 - Relatively low level
 - Portable but possibly slow
- OpenMP is newer standard
 - Support for scientific programming on shared memory architectures
- P4 (Parnacs) is another portable package
 - Higher level than Pthreads
 - <http://www.netlib.org/p4/index.html>

09/07/2010

CS4961



Overview of POSIX Threads

- **POSIX: Portable Operating System Interface for UNIX**
 - Interface to Operating System utilities
- **PThreads: The POSIX threading interface**
 - System calls to create and synchronize threads
 - Should be relatively uniform across UNIX-like OS platforms
- PThreads contain support for
 - Creating parallelism
 - Synchronizing
 - No explicit support for communication, because shared memory is implicit; a pointer to shared data is passed to a thread

09/07/2010

CS4961



Forking Posix Threads

Signature:

```
int pthread_create(pthread_t *,
                  const pthread_attr_t *,
                  void * (*)(void *),
                  void *);
```

Example call:

```
errcode = pthread_create(&thread_id,
                        &thread_attribute,
                        &thread_fun; &fun_arg);
```

- **thread_id** is the thread id or handle (used to halt, etc.)
- **thread_attribute** various attributes
 - standard default values obtained by passing a NULL pointer
- **thread_fun** the function to be run (takes and returns void*)
- **fun_arg** an argument can be passed to thread_fun when it starts
- **errorcode** will be set nonzero if the create operation fails

09/07/2010

CS4961



Simple Threading Example

```
void* SayHello(void *foo) {
    printf( "Hello, world!\n" );
    return NULL;
}

int main() {
    pthread_t threads[16];
    int tn;
    for(tn=0; tn<16; tn++) {
        pthread_create(&threads[tn], NULL, SayHello, NULL);
    }
    for(tn=0; tn<16 ; tn++) {
        pthread_join(threads[tn], NULL);
    }
    return 0;
}
```

Compile using gcc -pthread

But overhead of thread creation is nontrivial
SayHello should have a significant amount of work

09/07/2010

CS4961



Shared Data and Threads

- Variables declared outside of main are shared
- Object allocated on the heap may be shared (if pointer is passed)
- Variables on the stack are private; passing pointer to these around to other threads can cause problems
- Often done by creating a large "thread data" struct
 - Passed into all threads as argument
 - Simple example:

```
char *message = "Hello World!\n";
pthread_create( &thread1,
              NULL,
              (void*) &print_fun,
              (void*) message);
```

09/07/2010

CS4961



Posix Thread Example

```
#include <pthread.h>
void print_fun( void *message ) {
    printf("%s \n", message);
}

main() {
    pthread_t thread1, thread2;
    char *message1 = "Hello";
    char *message2 = "World";

    pthread_create( &thread1,
        NULL,
        (void*)&print_fun,
        (void*) message1);
    pthread_create(&thread2,
        NULL,
        (void*)&print_fun,
        (void*) message2);
    return (0);
}
```

Compile using `gcc -lpthread`

Note: There is a race condition in the print statements

09/07/2010

CS4961



Explicit Synchronization: Creating and Initializing a Barrier

- To (dynamically) initialize a barrier, use code similar to this (which sets the number of threads to 3):


```
pthread_barrier_t b;
pthread_barrier_init(&b, NULL, 3);
```
- The second argument specifies an object attribute; using NULL yields the default attributes.
- To wait at a barrier, a process executes:


```
pthread_barrier_wait(&b);
```
- This barrier could have been statically initialized by assigning an initial value created using the macro `PTHREAD_BARRIER_INITIALIZER(3)`.

09/07/2010

CS4961



Mutexes (aka Locks) in POSIX Threads

- To create a mutex:


```
#include <pthread.h>
pthread_mutex_t amutex = PTHREAD_MUTEX_INITIALIZER;
pthread_mutex_init(&amutex, NULL);
```
- To use it:


```
int pthread_mutex_lock(amutex);
int pthread_mutex_unlock(amutex);
```
- To deallocate a mutex


```
int pthread_mutex_destroy(pthread_mutex_t *mutex);
```
- Multiple mutexes may be held, but can lead to deadlock:

thread1	thread2
lock(a)	lock(b)
lock(b)	lock(a)

09/07/2010

CS4961



Summary of Programming with Threads

- POSIX Threads are based on OS features
 - Can be used from multiple languages (need appropriate header)
 - Familiar language for most programmers
 - Ability to shared data is convenient
- Pitfalls
 - Data race bugs are very nasty to find because they can be intermittent
 - Deadlocks are usually easier, but can also be intermittent
- OpenMP** is commonly used today as a simpler alternative, but it is more restrictive

09/07/2010

CS4961



OpenMP Motivation

- Thread libraries are hard to use
 - P-Threads/Solaris threads have many library calls for initialization, synchronization, thread creation, condition variables, etc.
 - Programmer must code with multiple threads in mind
- Synchronization between threads introduces a new dimension of program correctness
- Wouldn't it be nice to write serial programs and somehow parallelize them "automatically"?
 - OpenMP can parallelize many serial programs with relatively few annotations that specify parallelism and independence
 - It is not automatic: you can still make errors in your annotations

09/07/2010

CS4961



OpenMP: Prevailing Shared Memory Programming Approach

- Model for parallel programming
- Shared-memory parallelism
- Portable across shared-memory architectures
- Scalable
- Incremental parallelization
- Compiler based
- Extensions to existing programming languages (Fortran, C and C++)
 - mainly by directives
 - a few library routines

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

09/07/2010

CS4961



A Programmer's View of OpenMP

- OpenMP is a portable, threaded, shared-memory programming *specification* with "light" syntax
 - Exact behavior depends on OpenMP implementation!
 - Requires compiler support (C/C++ or Fortran)
- OpenMP will:
 - Allow a programmer to separate a program into *serial regions* and *parallel regions*, rather than concurrently-executing threads.
 - Hide stack management
 - Provide synchronization constructs
- OpenMP will not:
 - Parallelize automatically
 - Guarantee speedup
 - Provide freedom from data races

09/07/2010

CS4961



OpenMP Data Parallel Construct: Parallel Loop

- All pragmas begin: `#pragma`
- Compiler calculates loop bounds for each thread directly from *serial* source (computation decomposition)
- Compiler also manages data partitioning of Res
- Synchronization also automatic (barrier)

Serial Program:	Parallel Program:
<pre>void main() { double Res[1000]; for(int i=0;i<1000;i++) { do_huge_comp(Res[i]); } }</pre>	<pre>void main() { double Res[1000]; #pragma omp parallel for for(int i=0;i<1000;i++) { do_huge_comp(Res[i]); } }</pre>

09/07/2010

CS4961



OpenMP Execution Model

- Fork-join model of parallel execution
- Begin execution as a single process (master thread)
- Start of a parallel construct:
 - Master thread creates team of threads
- Completion of a parallel construct:
 - Threads in the team synchronize -- **implicit barrier**
- Only master thread continues execution
- Implementation optimization:
 - Worker threads spin waiting on next fork

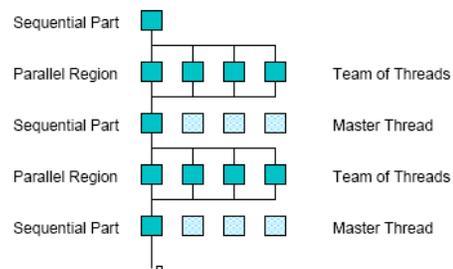


09/07/2010

CS4961



OpenMP Execution Model



09/07/2010

CS4961



Count 3s Example? (see textbook)

- What do we need to worry about?

09/07/2010

CS4961



OpenMP directive format C (also Fortran and C++ bindings)

- Pragmas, format


```
#pragma omp directive_name [ clause [ clause ] ... ] new-line
```
- Conditional compilation


```
#ifdef _OPENMP
  block,
  e.g., printf("%d avail.processors\n", omp_get_num_procs());
#endif
```
- Case sensitive
- Include file for library routines


```
#ifdef _OPENMP
#include <omp.h>
#endif
```

09/07/2010

CS4961



Limitations and Semantics

- Not all "element-wise" loops can be ||ized

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

- Loop index: signed integer
 - Termination Test: <, <=, >, >= with loop invariant int
 - Incr/Decr by loop invariant int; change each iteration
 - Count up for <, <=; count down for >, >=
 - Basic block body: no control in/out except at top
- Threads are created and iterations divvied up; requirements ensure iteration count is predictable

09/07/2010

CS4961



OpenMP Synchronization

- Implicit barrier
 - At beginning and end of parallel constructs
 - At end of all other control constructs
 - Implicit synchronization can be removed with `nowait` clause
- Explicit synchronization
 - `critical`
 - `atomic`

09/07/2010

CS4961



OpenMp Reductions

- OpenMP has reduce

```
sum = 0;
#pragma omp parallel for reduction(+:sum)
for (i=0; i < 100; i++) {
  sum += array[i];
}
```

- Reduce ops and init() values:

+	0	bitwise &	~0	logical &	1
-	0	bitwise	0	logical	0
*	1	bitwise ^	0		

09/07/2010

CS4961



OpenMP parallel region construct

- Block of code to be executed by multiple threads in parallel
- Each thread executes the **same code redundantly (SPMD)**
 - Work within work-sharing constructs is distributed among the threads in a team
- Example with C/C++ syntax


```
#pragma omp parallel [ clause [ clause ] ... ] new-line
structured-block
```
- clause can include the following:
 - `private (list)`
 - `shared (list)`

09/07/2010

CS4961



Programming Model - Loop Scheduling

- schedule clause determines how loop iterations are divided among the thread team
 - `static ([chunk])` divides iterations statically between threads
 - Each thread receives [chunk] iterations, rounding as necessary to account for all iterations
 - Default [chunk] is $\text{ceil}(\text{\# iterations} / \text{\# threads})$
 - `dynamic ([chunk])` allocates [chunk] iterations per thread, allocating an additional [chunk] iterations when a thread finishes
 - Forms a logical work queue, consisting of all loop iterations
 - Default [chunk] is 1
 - `guided ([chunk])` allocates dynamically, but [chunk] is exponentially reduced with each allocation

09/07/2010

CS4961



Loop scheduling



09/07/2010

CS4961



OpenMP critical directive

- Enclosed code
 - executed by all threads, but
 - **restricted to only one thread at a time**
- `#pragma omp critical [(name)] new-line`
structured-block
- A thread waits at the beginning of a critical region until no other thread in the team is executing a critical region with the same name.
- All unnamed critical directives map to the same unspecified name.

09/07/2010

CS4961



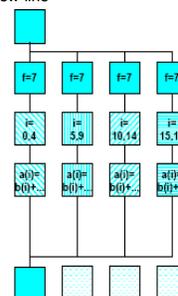
Variation: OpenMP parallel and for directives

Syntax:

```
#pragma omp for [ clause [ clause ] ... ] new-line
for-loop
```

clause can be one of the following:

```
shared( list )
private( list )
reduction( operator: list )
schedule( type [ , chunk ] )
nowait (C/C++: on #pragma omp for)
#pragma omp parallel private(f) {
  f=7;
#pragma omp for
  for (i=0; i<20; i++)
    a[i] = b[i] + f * (i+1);
} /* omp end parallel */
```



09/07/2010

CS4961



Programming Model - Data Sharing

- Parallel programs often employ two types of data

- Shared data, visible to all threads, similarly named
- Private data, visible to a single thread (often stack-allocated)

- PThreads:

- Global-scoped variables are shared
- Stack-allocated variables are private

- OpenMP:

- shared variables are shared
- private variables are private
- Default is shared
- Loop index is private

```
// shared, globals
int bigdata[1024];

void* foo(void* bar) {
    int tid;

    #pragma omp parallel \
        shared ( bigdata ) \
        private ( tid )
    {
        /* Calc. here */
    }
}
```



OpenMP environment variables

OMP_NUM_THREADS

- sets the number of threads to use during execution
- when dynamic adjustment of the number of threads is enabled, the value of this environment variable is the maximum number of threads to use

- For example,

```
setenv OMP_NUM_THREADS 16 [csh, tcsh]
export OMP_NUM_THREADS=16 [sh, ksh, bash]
```

OMP_SCHEDULE

- applies only to `do/for` and `parallel do/for` directives that have the schedule type `RUNTIME`

- sets schedule type and chunk size for all such loops

- For example,

```
setenv OMP_SCHEDULE GUIDED,4 [csh, tcsh]
export OMP_SCHEDULE= GUIDED,4 [sh, ksh, bash]
```

09/07/2010

CS4961



OpenMP runtime library. Query Functions

`omp_get_num_threads`:

Returns the number of threads currently in the team executing the parallel region from which it is called

```
int omp_get_num_threads(void);
```

`omp_get_thread_num`:

Returns the thread number, within the team, that lies between 0 and `omp_get_num_threads()-1`, inclusive. The master thread of the team is thread 0

```
int omp_get_thread_num(void);
```

09/07/2010

CS4961



Summary of Lecture

- OpenMP, data-parallel constructs only
 - Task-parallel constructs later
- What's good?
 - Small changes are required to produce a parallel program from sequential (parallel formulation)
 - Avoid having to express low-level mapping details
 - Portable and scalable, correct on 1 processor
- What is missing?
 - Not completely natural if want to write a parallel code from scratch
 - Not always possible to express certain common parallel constructs
 - Locality management
 - Control of performance

09/07/2010

CS4961

