

Shared Memory: Virtual Shared Memory, Threads & OpenMP

Eugen Betke

University of Hamburg
Department Informatik
Scientific Computing

09.01.2012

Agenda

- 1 Introduction
 - Architectures of Memory Systems
- 2 Virtual Shared Memory
 - Memory Mapping
 - Separated Address Space
 - Fundamental Problems
- 3 Pthreads
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 4 OpenMP
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 5 Summary

Table Of Content

1 Introduction

Architectures of Memory Systems

2 Virtual Shared Memory

Memory Mapping

Separated Address Space

Fundamental Problems

3 Pthreads

Example

Race Conditions and Critical Sections

Locks

4 OpenMP

Example


Race Conditions and Critical Sections

Locks


5 Summary

Shared Memory vs. Distributed Memory [1]


▶ Shared Memory



▶ Distributed Memory




Shared Memory: Uniform Memory Access System



- ▶ Same access times for all the cores
- ▶ Direct connection to a block of memory
- ▶ Relative easy to program

Shared Memory: Nonuniform Memory Access System




- ▶ Different access times
- ▶ Different memory locations

Table Of Content

- 1 Introduction
 - Architectures of Memory Systems
- 2 Virtual Shared Memory
 - Memory Mapping
 - Separated Address Space
 - Fundamental Problems
- 3 Pthreads
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 4 OpenMP
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 5 Summary

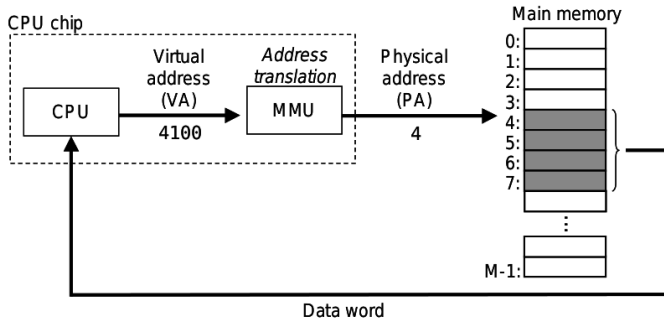
Memory: Physical Address Space [2]

- ▶ A system that uses physical addressing

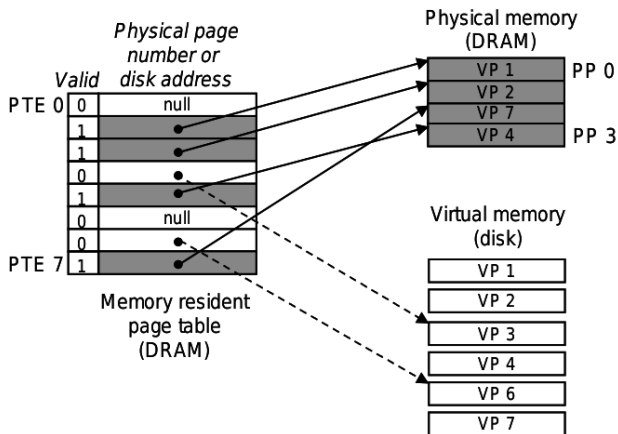


Virtual Memory: Mapping I

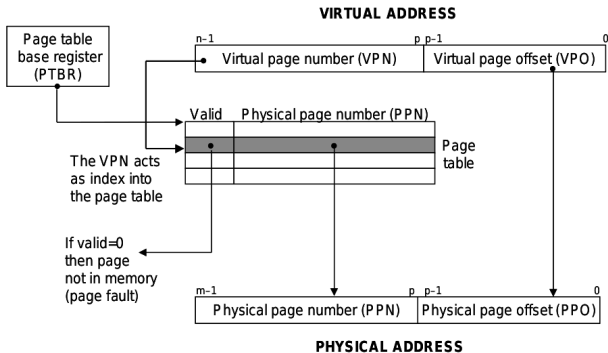
- ▶ A system that uses virtual addressing
- ▶ MMU - Memory Management Unit




Virtual Memory: Mapping II



Virtual Memory: Mapping III



Virtual Shared Memory: Separated Address Space I

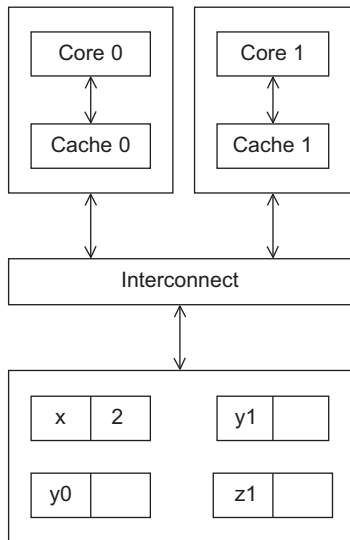


- ▶ Each process has its own Virtual Address Space

Virtual Shared Memory: Separated Address Space II

- ▶ Simplifying linking
 - ▶ Allows a linker to produce fully linked executables that are independent of the ultimate location of the code and data in the physical memory
- ▶ Simplifying loading
 - ▶ Easy to load executable and shared objects into memory
- ▶ Simplifying sharing
 - ▶ Mapping to the same physical page
- ▶ Simplifying memory allocation
 - ▶ Allocation of contiguous space

Cache Coherence



Example

x is initialized to 2

y0 is private and owned by core 0

y1, z1 are private and owned by core 1

Time	Thread 0	Thread 1
0	y0 = x;	y1 = 3 * x;
1	x = 7;	
2		z1 = 4 * x;

What is the value of z1?

z1 = 4 * 7 = 28

or

z1 = 4 * 2 = 8

Problem: Cache Coherence

Snooping cache coherence

Cores share a bus. When core 0 updates the copy of x stored in its cache, it broadcasts this information across the bus. Core 1 is “snooping” the bus and it will see that x has been updated. He marks his copy of x as invalid.

Directory-based cache coherence

A directory stores the status of each cache line. A directory is typically distributed and each CPU/Memory pair is responsible to update the status of its own local memory. If core 0 reads a cache line in its own local cache, it writes in the directory, that he has a copy of this cache line in his local cache. When an other core modifies a variable, that lies in that cache line, the cache controller invalidates the copies in corresponding local caches.

Problem: False Sharing I

Example (Serial program)

As an example, suppose we want to repeatedly call a function $f(i, j)$ and add the computed values into a vector.

```
1 int i, j, m, n;  
2 double y[m];  
3 /* Assign y = 0 */  
4 ...  
5 for (i = 0; i < m; i++)  
6     for (j = 0; j < n; j++)  
7         y[i] += f(i, j);
```


Problem: False Sharing II

Example (Parallel program)

```
1  /* Private variable */
2  int i, j, iter_count;
3  /* Shared variables initialized by one core */
4  int m, n, core_count;
5  double y[m];
6
7  iter_count = m/core_count;
8
9  /* Core 0 does this */
10 for (i = 0; i < iter_count; i++)
11     for (j = 0; j < n; j++)
12         y[i] += f(i,j);
13
14 /* Core 1 does this */
15 for (i = iter_count + 1; i < 2 * iter_count; i++)
16     for (j = 0; j < n; j++)
17         y[i] += f(i,j);
```

- ▶ $m = 8$
- ▶ doubles are 8 bytes
- ▶ cache line can store eight doubles (64 bytes)
- ▶ y takes one full cache line

What happens?

Table Of Content

- 1 Introduction
 - Architectures of Memory Systems
- 2 Virtual Shared Memory
 - Memory Mapping
 - Separated Address Space
 - Fundamental Problems
- 3 Pthreads**
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 4 OpenMP
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 5 Summary

- ▶ Pthreads - POSIX threads
- ▶ Specifies a library for Unix-like systems
- ▶ It exists other specifications like: Java threads, Windows threads, Solaris threads. All of the specifications support the same basic ideas.

Pthreads: Forking and Joining



- ▶ Threads are often called light-weight processes
- ▶ Master thread forks slave threads
- ▶ Slave threads joins to the master thread
- ▶ Typical approaches to thread startup:
 - ▶ Static threads: all threads are created before computation
 - ▶ Example: Computation of scalar product
 - ▶ Dynamic threads: threads are created at demand
 - ▶ Example: Web server applications, that responds to client requests

Pthreads: Example I

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <pthread.h>
4
5 /* Global variable: accessible to all threads */
6 int thread_count;
7
8 void *Hello(void* rank); /* Thread function */
```

Pthreads: Example II

```
1 int main(int argc, char* argv[]) {
2     long      thread; /* Use long in case of a 64-bit system */
3     pthread_t* thread_handles;
4
5     /* Get number of threads from command line */
6     thread_count = strtol(argv[1], NULL, 10);
7
8     thread_handles = malloc (thread_count*sizeof(pthread_t));
9
10    for (thread = 0; thread < thread_count; thread++)
11        pthread_create(&thread_handles[thread], NULL, Hello, (void*) thread);
12
13    printf("Hello from the main thread\n");
14
15    for (thread = 0; thread < thread_count; thread++)
16        pthread_join(thread_handles[thread], NULL);
17
18    free(thread_handles);
19    return 0;
20 }
```

Pthreads: Example III

```
1 void *Hello(void* rank) {
2     long my_rank = (long) rank; /* Use long in case of 64-bit system */
3
4     printf("Hello from thread %ld of %d\n", my_rank, thread_count);
5
6     return NULL;
7 } /* Hello */
```

Pthreads: Forking and Joining functions

▶ Starting/forking threads

```
1 int pthread_create(  
2     pthread_t* thread_p,  
3     const pthread_attr_t* attr_p,  
4     void* (*start_routine)(void*),  
5     void* arg_p);
```

▶ Thread function

```
1 void* thread_function(void* args_p);
```

▶ Stopping/joining threads

```
1 int pthread_join(  
2     pthread_t thread,  
3     void** ret_val_p);
```


Pthreads: Example execution

► Compilation

```
1 gcc -g -Wall -o pthread_hello pthread_hello.c -lpthreads
```

► Execution with 1 thread

```
1 ./pthread_hello 1
```

```
1 Hello from the main thread
2 Hello from thread 0 of 1
```

► Execution with 4 threads results in a non-deterministic output

```
1 ./pthread_hello 4
```

```
1 Hello from the main thread
2 Hello from thread 0 of 4
3 Hello from thread 1 of 4
4 Hello from thread 2 of 4
5 Hello from thread 3 of 4
```

```
1 Hello from thread 0 of 4
2 Hello from thread 2 of 4
3 Hello from thread 1 of 4
4 Hello from the main thread
5 Hello from thread 3 of 4
```

Pthreads: Race conditions and critical sections

Race condition

When several threads attempt to access a shared resource such as a shared variable or a shared file, at least one of the accesses is an update, and the access can result in an error, we have a race condition.

Pthreads: Example of Race Condition

Example

Suppose that we have two threads. Each thread computes a value and stores it in a private variable y . We want add these both values to the shared variable x .

```
1 y = Compute(my_rank);  
2 x = x + y;
```

Time	Thread 0	Thread 1
1	Started by main thread	
2	Call Compute()	Started by main thread
3	Assign $y = 1$	Call Compute()
4	Put $x = 0$ and $y = 1$ into registers	Assign $y = 2$
5	Add 0 and 1	Put $x = 0$ and $y = 2$ into registers
6	Store 1 in memory location x	Add 0 and 2
7		Store 2 in memory location x

The “winner” result will be overwritten by the “loser”.

- ▶ Locks
 - ▶ busy-waiting
 - ▶ mutexes (mutual exclusions)
 - ▶ semaphores
 - ▶ read-write locks

Pthreads: Example of Critical Section

Formula for computing π

$$\pi = 4 \left(1 - \frac{1}{3} + \frac{1}{5} - \frac{1}{7} + \dots + (-1)^n \frac{1}{2n+1} + \dots \right) \quad (1)$$

► Serial computation of π

```
1 double Serial_pi(long long n) {
2     double sum = 0.0;
3     long long i;
4     double factor = 1.0;
5
6     for (i = 0; i < n; i++, factor = -factor) {
7         sum += factor/(2*i+1);
8     }
9     return 4.0*sum;
10 } /* Serial_pi */
```

Pthreads: Example of Critical Section I

```
1 void* Thread_sum(void* rank) {
2     long my_rank = (long) rank;
3     double factor;
4     long long i;
5     long long my_n = n/thread_count;
6     long long my_first_i = my_n*my_rank;
7     long long my_last_i = my_first_i + my_n;
8
9     if (my_first_i % 2 == 0)
10        factor = 1.0;
11    else
12        factor = -1.0;
13
14    for (i = my_first_i; i < my_last_i; i++, factor = -factor) {
15        sum += factor/(2*i+1); // Critical Section
16    }
17
18    return NULL;
19 } /* Thread_sum */
```

"Concept of Busy-Waiting"

```
1 y = Compute(my_rank);  
2 while = (flag != my_rank);  
3 x = x + y;  
4 flag++;
```

- ▶ Thread keep re-executing the test until the test is false
- ▶ Simple implementation with a busy-wait loop
- ▶ Programmer can control the order of execution of threads
- ▶ Consumes CPU cycles
- ▶ Can seriously degrade performance

Pthreads: Example of Busy-Waiting

```
1 void* Thread_sum(void* rank) {
2     long my_rank = (long) rank;
3     double factor, my_sum = 0.0;
4     long long i;
5     long long my_n = n/thread_count;
6     long long my_first_i = my_n*my_rank;
7     long long my_last_i = my_first_i + my_n;
8
9     if (my_first_i % 2 == 0)
10        factor = 1.0;
11    else
12        factor = -1.0;
13
14    for (i = my_first_i; i < my_last_i; i++, factor = -factor)
15        my_sum += factor(2*i+1);
16
17    while (flag != my_rank);
18    sum += my_sum;
19    flag = (flag+1) % thread_count;
20
21    return NULL;
22 } /* Thread_sum */
```


- ▶ Abbreviation of mutual exclusions
- ▶ A mutex is a special type of a variable

Pthreads: Mutexes

A variable of type `pthread_mutex_t` needs to be initialized before it (a mutex) can be used.

► Initialization of a mutex

```
1 int pthread_mutex_init(  
2     pthread_mutex_t* mutex_p,  
3     const pthread_mutexattr_t* attr_p);
```

► Destruction of a mutex

```
1 int pthread_mutex_destroy(pthread_mutex_t* mutex_p);
```

► Gain access to a critical section

```
1 int pthread_mutex_lock(pthread_mutex_t* mutex_p);
```

► Unlock critical section

```
1 int pthread_mutex_unlock(pthread_mutex_t* mutex_p);
```

Pthreads: Semaphores

- ▶ Semaphores can be thought as a special type of unsigned `int`
 - ▶ They can take values 0, 1, 2, ...
 - ▶ Binary semaphore takes 0 and 1 as values
 - ▶ Value 0 means “locked” and 1 means “unlocked”
- ▶ Semaphores are not a part of Pthreads and it's necessary to add the following preprocessor directive:

```
1 #include <semaphore.h>
```

- ▶ It's possible to control the order in which the threads execute the critical section

Pthreads: Semaphores in C-Language

A variable of type `sem_t` needs to be initialized before it (a semaphore) can be used.

► Initialization of a semaphore

```
1 int sem_init(  
2     sem_t* semaphore_p,  
3     int shared,  
4     unsigned initial_val);
```

► Destruction of a mutex

```
1 int sem_destroy(sem_t* semaphore_p);
```

► Increment semaphore

```
1 int sem_post(sem_t* semaphore_p);
```

► Decrement semaphore

```
1 int sem_wait(sem_t* semaphore_p);
```

Pthreads: Read-Write Locks

- ▶ Low-level locking
- ▶ Provides two lock-functions
 - ▶ One lock function locks the read-write lock for reading
 - ▶ The other lock function locks the read-write lock for writing

Pthreads: Read-Write Locks in C-Language

A variable of type `pthread_rwlock_t` needs to be initialized before it (a rwlock) can be used.

► Initialization of rwlock

```
1  int pthread_rwlock_init(  
2      pthread_rwlock_t* rwlock_p,  
3      const pthread_rwlockattr_t*);
```

► Destruction of rwlock

```
1  int pthread_rwlock_destroy(pthread_rwlock_t* rwlock_p);
```

► Read-write lock for reading

```
1  int pthread_rwlock_rdlock(pthread_rwlock_t* rwlock_p);
```

► Read-write lock for writing

```
1  int pthread_rwlock_wrlock(pthread_rwlock_t* rwlock_p);
```

► Unlock

```
1  int pthread_rwlock_unlock(pthread_rwlock_t* rwlock_p);
```

Table Of Content

- 1 Introduction
 - Architectures of Memory Systems
- 2 Virtual Shared Memory
 - Memory Mapping
 - Separated Address Space
 - Fundamental Problems
- 3 Pthreads
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 4 OpenMP
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 5 Summary

OpenMP (OpenMultiProcessing)

- ▶ Standard for programming shared memory systems
- ▶ Uses library functions and preprocessor directives (pragmas)
- ▶ Requires compiler support
- ▶ Developers could incrementally parallelize existing serial programs
- ▶ Higher-level than Pthreads

OpenMP: Example 1

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include <omp.h>
4
5  void Hello(void); /* Thread function */
6
7  int main(int argc, char* argv[]) {
8      int thread_count = strtol(argv[1], NULL, 10);
9
10     # pragma omp parallel num_threads(thread_count)
11         Hello();
12
13     return 0;
14 } /* main */
15
16 void Hello(void) {
17     int my_rank = omp_get_thread_num();
18     int thread_count = omp_get_num_threads();
19
20     printf("Hello from thread %d of %d\n", my_rank, thread_count);
21
22 } /* Hello */
```

OpenMP: Example II

- ▶ Compilation

```
1 gcc -g -Wall -fopenmp -o omp_hello omp_hello.c
```

- ▶ Execution with 4 threads

```
1 ./omp_hello 4
```

- ▶ Output is non-deterministic

```
1 Hello from thread 0 of 4
2 Hello from thread 1 of 4
3 Hello from thread 2 of 4
4 Hello from thread 3 of 4
```

```
1 Hello from thread 1 of 4
2 Hello from thread 2 of 4
3 Hello from thread 0 of 4
4 Hello from thread 3 of 4
```

OpenMP: Pragmas

OpenMP pragmas always begin with `# pragma omp`. They are followed by a directive. Strings after a directive are called clauses. Clauses provide additional information for the directive.

Example

```
1 # pragma omp parallel num_threads(thread_count)
2   Hello();
```

- ▶ The `parallel` directive specifies, that the structured block of code that follows should be executed by multiple threads
- ▶ The clause `num_threads(thread_count)` specifies how many threads of the structured block below should be created

- ▶ `#include <omp.h>` provides predefined constants and OpenMP functions
 - ▶ `int omp_get_thread_num(void)`; returns the ID of the current thread
 - ▶ `int omp_get_num_threads(void)`; returns the number of threads in the team
- ▶ `#include omp.h` is only needed, if we use predefined constants or call OpenMP functions

Race Conditions and Critical Sections

Race Condition: multiple threads are attempting to access a shared resource, at least one of the accesses is an update, and the accesses can result in an error.

Example

Time	Thread 0	Thread 1
0	global_result = 0 to register	finish my_result
1	my_result = 1 to register	global_result = 0 to register
2	add my_result to global_result	my_result = 2 to register
3	store global_result = 1	add my_result to global_result
4		store global_result = 2

```
1 # pragma omp critical
2   global_result = my_result;
```

- ▶ Only one thread can execute after `# pragma omp critical` following structured block of code.
- ▶ No other thread can start execute this code until the first thread has finished.

- ▶ `# pragma omp critical`
- ▶ `# pragma omp atomic`
- ▶ Lock-functions in `omp.h`

OpenMP: `critical` Directive

```
1 # pragma omp critical(name)
2     <structured block>
```

- ▶ Blocks protected with `critical` directives with different names can be executed simultaneously.

OpenMP: atomic Directive

```
1 # pragma omp atomic
2     <single C assignment statement>
```

- ▶ Can only protect critical sections that consist of a single C assignment statement.
- ▶ Statement must have one of the following form:
 - ▶ `x <op> = <expression>;`
 - ▶ `x++;`
 - ▶ `++x;`
 - ▶ `x--;`
 - ▶ `--x;`
- ▶ `<op>` can be one of the binary operators:
 - ▶ `+, *, -, /, &, ^, |, <<, >>`
- ▶ `<expression>` must not reference `x`

OpenMP: Lock Functions

A variable of type `omp_lock_t` needs to be initialized before it (a lock) can be used.

► Initialization

```
1 void omp_init_lock(omp_lock_t* lock_p);
```

► Set lock

```
1 void omp_set_lock(omp_lock_t* lock_p);
```

► Unset lock

```
1 void omp_unset_lock(omp_lock_t* lock_p);
```

► Destroy lock

```
1 void omp_destroy_lock(omp_lock_t* lock_p);
```

OpenMP: reduction Clause

1 `reduction(<operator>: <variable list>)`

- ▶ `<operator>` can be `+`, `*`, `-`, `&`, `|`, `^`, `&&`, `||`
- ▶ OpenMP creates for each variable in `variable list` a private variable and stores there the result of computation
- ▶ OpenMP creates a critical section, where the results from the private variables are computed with the corresponding variable

OpenMP: Reduction Clause in C-Language

Example

Code without reduction clause:

```
1      global_result = 0.0;
2  #      pragma omp parallel num_threads(thread_count)
3      {
4          double my_result = 0.0 /* private */
5          my_result += Local_trap(double a, double b, int n);
6  #      pragma omp critical
7          global_result += my_result;
8      }
```

Equivalent code with reduction clause:

```
1      global_result = 0.0
2  #      pragma omp parallel num_threads(thread_count) \
3          reduction(+: global_result)
4      global_result += Local_trap(double a, double b, int n);
```

OpenMP provides many more other directives, clauses and library functions.

- ▶ Parallelization of `for`-loops
- ▶ Barriers and condition variables

Table Of Content

- 1 Introduction
 - Architectures of Memory Systems
- 2 Virtual Shared Memory
 - Memory Mapping
 - Separated Address Space
 - Fundamental Problems
- 3 Pthreads
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 4 OpenMP
 - Example
 - Race Conditions and Critical Sections
 - Locks
- 5 Summary

Summary

- ▶ Memory Systems
 - ▶ Shared and distributed systems
 - ▶ Uniform and non-uniform memory systems
- ▶ Virtual shared memory
 - ▶ Virtual addressing
 - ▶ Problems: cache coherence and false sharing
- ▶ Pthreads and OpenMP
 - ▶ Example
 - ▶ Race conditions
 - ▶ Locks



P. Pacheco, *An Introduction to Parallel Programming*.

San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1st ed., 2011.



R. E. Bryant and D. R. O'Hallaron, *Computer Systems: A Programmer's Perspective*.

USA: Addison-Wesley Publishing Company, 2nd ed., 2010.