

CISC 3320

C21a Race Condition and Critical Section

Hui Chen

Department of Computer & Information Science
CUNY Brooklyn College

Acknowledgement

- These slides are a revision of the slides provided by the authors of the textbook via the publisher of the textbook

Outline

- Race condition
- The Critical-Section Problem
- Peterson's Solution

- Hardware Support for Synchronization
- Mutex Locks
- Semaphores
- Monitors
- Liveness
- Evaluation

Background

- Processes can execute concurrently
 - May be interrupted at any time, partially completing execution
- Concurrent access to shared data may result in data inconsistency
- Maintaining data consistency requires mechanisms to ensure the orderly execution of cooperating processes
- Otherwise, a race condition occurs
 - Examining two examples

Race Condition: Example 1

- Illustration of the problem using the consumer-producer problem
 - Suppose that we wanted to provide a solution to the consumer-producer problem that fills all the buffers.
 - We can do so by having an integer counter that keeps track of the number of full buffers.
 - Initially, counter is set to 0.
 - It is incremented by the producer after it produces a new buffer and is decremented by the consumer after it consumes a buffer.

Producer

```
while (true) {  
    /* produce an item in next produced */  
  
    while (counter == BUFFER_SIZE)  
        ; /* do nothing */  
  
    buffer[in] = next_produced;  
    in = (in + 1) % BUFFER_SIZE;  
    counter++;  
}
```

Consumer

```
while (true) {  
    while (counter == 0)  
        ; /* do nothing */  
    next_consumed = buffer[out];  
    out = (out + 1) % BUFFER_SIZE;  
    counter--;  
    /* consume the item in next consumed */  
}
```

Increment and Decrement Counter

- counter++ could be implemented as

```
register1 = counter  
register1 = register1 + 1  
counter = register1
```

- counter-- could be implemented as

```
register2 = counter  
register2 = register2 - 1  
counter = register2
```

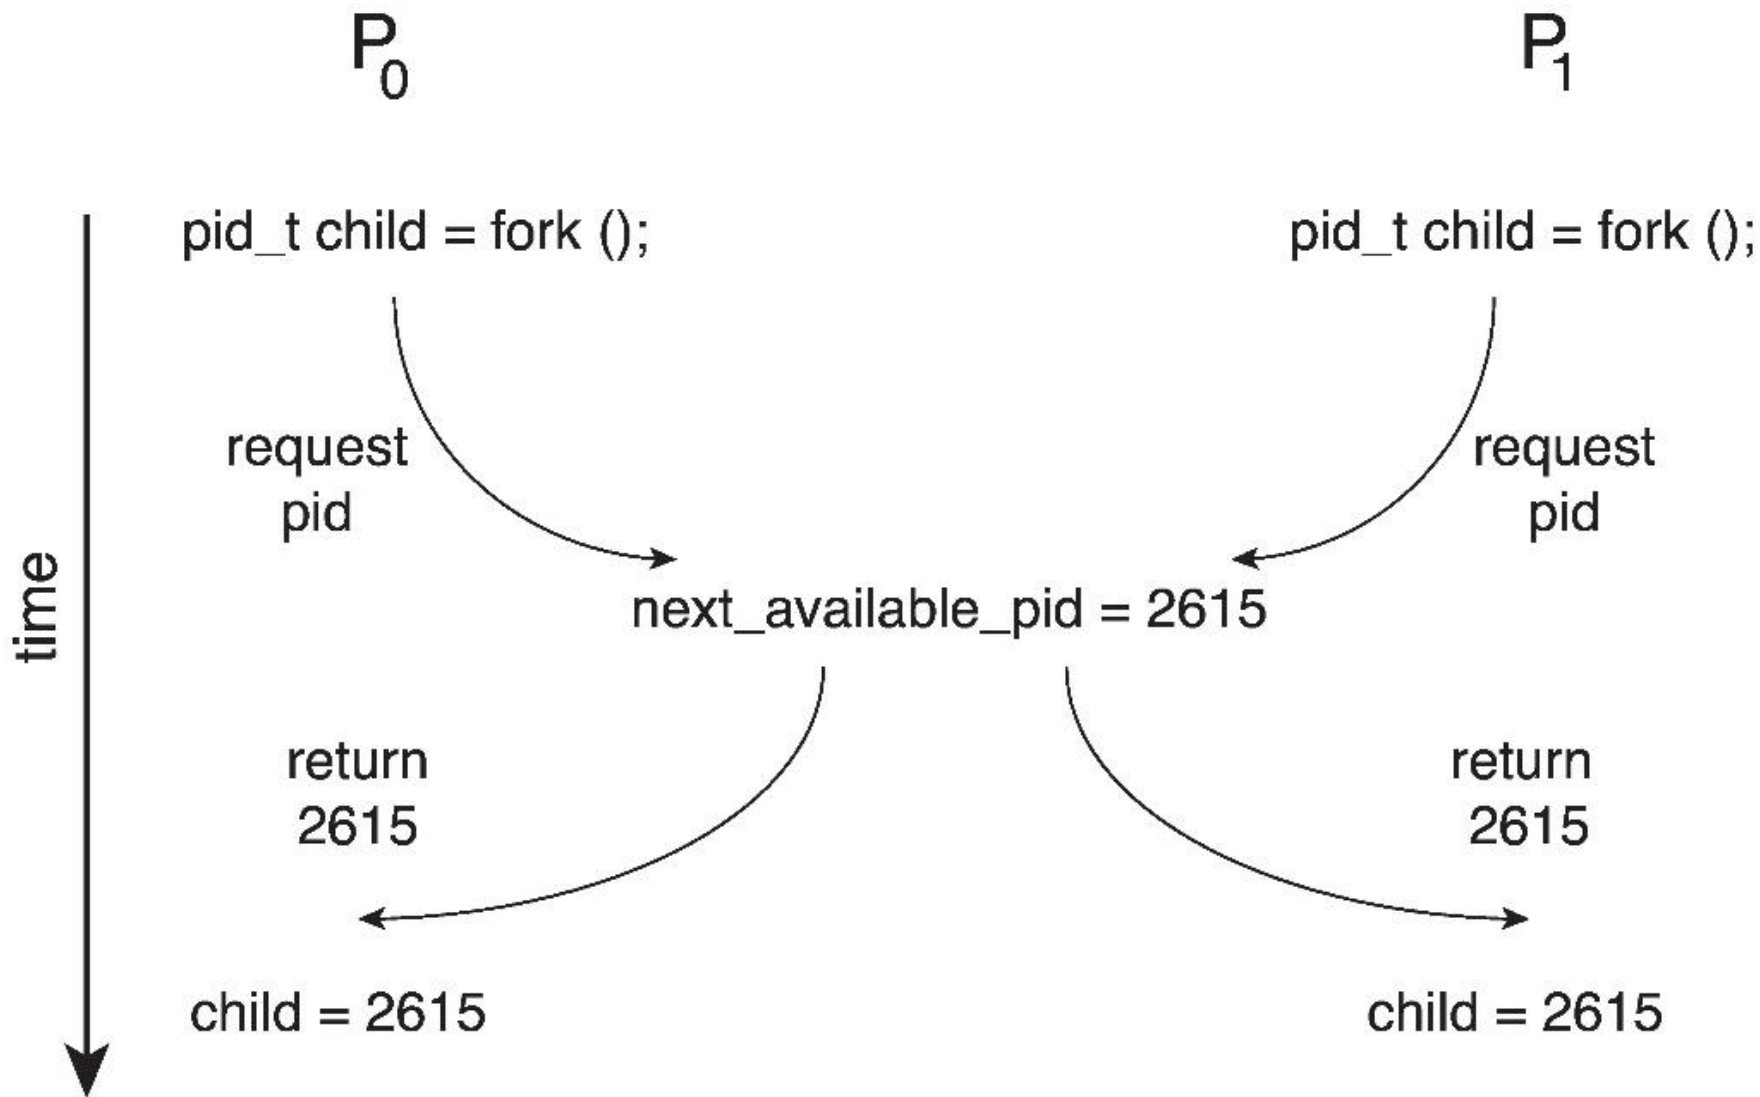

Race Condition

- Consider this execution interleaving with "count = 5" initially:

S0: producer execute $register1 = counter$ {register1 = 5}
S1: producer execute $register1 = register1 + 1$ {register1 = 6}
S2: consumer execute $register2 = counter$ {register2 = 5}
S3: consumer execute $register2 = register2 - 1$ {register2 = 4}
S4: producer execute $counter = register1$ {counter = 6}
S5: consumer execute $counter = register2$ {counter = 4}

Race Condition: Example 2

- Processes P0 and P1 are creating child processes using the `fork()` system call
- Race condition on kernel variable `next_available_pid` which represents the next available process identifier (pid)
- Unless there is mutual exclusion, the same pid could be assigned to two different processes!



Questions?

- Concept of race condition
- Need of process synchronization

Critical Section Problem

- Consider system of n processes $\{p_0, p_1, \dots, p_{n-1}\}$
- Each process has **critical section** segment of code
 - Process may be changing common variables, updating table, writing file, etc
 - When one process in critical section, no other may be in its critical section
- **Critical section problem** is to design protocol to solve this
- Each process must ask permission to enter critical section in **entry section**, may follow critical section with **exit section**, then **remainder section**

Critical Section

- General structure of process P_i

do {

entry section

critical section

exit section

remainder section

} while (true);

Solution to Critical-Section Problem

- 3 requirements must be met
 - Mutual Exclusion
 - Progress
 - Bounded Waiting
- Assume that each process executes at a nonzero speed
- No assumption concerning relative speed of the n processes

Mutual Exclusion

- If process P_i is executing in its critical section, then no other processes can be executing in their critical sections

Progress

- If no process is executing in its critical section and there exist some processes that wish to enter their critical section, then the selection of the processes that will enter the critical section next cannot be postponed indefinitely

Bounded Waiting

- A bound must exist on the number of times that other processes are allowed to enter their critical sections after a process has made a request to enter its critical section and before that request is granted

Critical-Section Handling in OS

- Two approaches depending on if kernel is preemptive or non-preemptive
 - Preemptive
 - allows preemption of process when running in kernel mode
 - Non-preemptive
 - Runs until exits kernel mode, blocks, or voluntarily yields CPU
 - Essentially free of race conditions in kernel mode

Peterson's Solution

- Not guaranteed to work on modern architectures! (But good algorithmic description of solving the problem)
- Two process solution
 - Assume that the load and store machine-language instructions are atomic; that is, cannot be interrupted
- The two processes share two variables:
 - `int turn; boolean flag[2];`
 - The variable *turn* indicates whose turn it is to enter the critical section
 - The *flag array* is used to indicate if a process is ready to enter the critical section.
 - `flag[i] = true` implies that process P_i is ready!

Algorithm for Process P_i

- Notice "i" in Process P_i

```
while (true){
    flag[i] = true;
    turn = j;
    while (flag[j] && turn == j)
        ;

    /* critical section */

    flag[i] = false;

    /* remainder section */

}
```

Peterson's Solution: 3 Requirements

- Provable that the 3 critical section requirements are met:
 1. Mutual exclusion is preserved
 2. Progress requirement is satisfied
 3. Bounded-waiting requirement is met

Peterson's Solution: Mutual Exclusion

- Note that
 - P_i enters its critical section only if either $\text{flag}[j] == \text{false}$ or $\text{turn} == i$.
 - if both processes can be executing in their critical sections at the same time, then $\text{flag}[0] == \text{flag}[1] == \text{true}$.
- These two observations imply that
 - P_0 and P_1 could not have successfully executed their while statements at about the same time, since the value of turn can be either 0 or 1 but cannot be both.
- Therefore,
 - one of the processes, e.g., P_j must have successfully executed the while statement, whereas P_i had to execute at least one additional statement (" $\text{turn} == j$ ").
 - However, at that time, $\text{flag}[j] == \text{true}$ and $\text{turn} == j$, and this condition will persist as long as P_j is in its critical section; as a result, mutual exclusion is preserved.

Peterson's Solution: Remarks

- Although useful for demonstrating an algorithm, Peterson's Solution is not guaranteed to work on modern architectures.
- Understanding why it will not work is also useful for better understanding race conditions.
- To improve performance, processors and/or compilers may reorder operations that have no dependencies.
- For single-threaded this is OK as the result will always be the same.
- For multithreaded the reordering may produce inconsistent or unexpected results!

Peterson's Solution: 2-Thread Example

- Two threads share the data:

```
boolean flag = false;  
int x = 0;
```

- Thread 1 performs

```
while (!flag)  
    ;  
print x
```

- Thread 2 performs

```
x = 100;  
flag = true
```

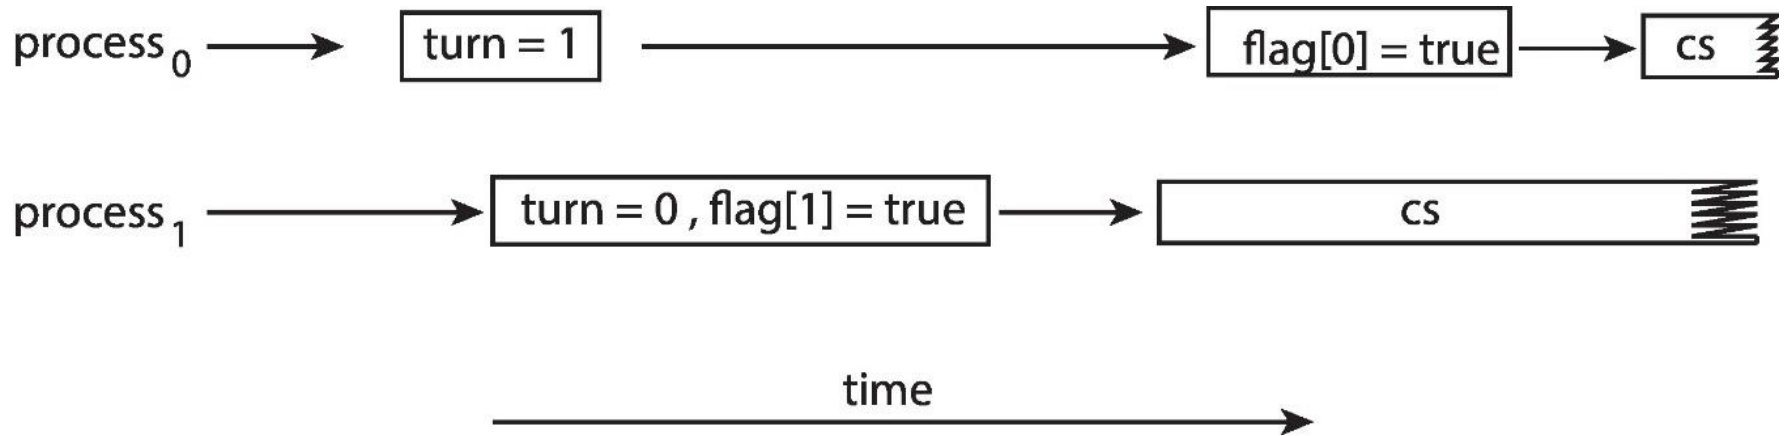
- What is the expected output?

Peterson's Solution

- 100 is the expected output.
- However, the operations for Thread 2 may be reordered:

```
flag = true;  
x = 100;
```

- If this occurs, the output may be 0!
 - The effects of instruction reordering in Peterson's Solution
 - This allows both processes to be in their critical section at the same time!



Questions?

- Peterson's solution
 - A software solution, a good description of an algorithm solving the problem
- Is it guaranteed to work on modern operating systems?