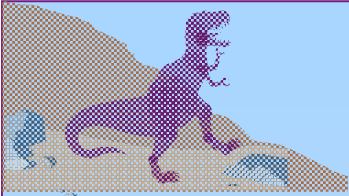


Chapter 7: Process Synchronization


- The Critical-Section Problem
- Synchronization Hardware
- Semaphores
- Synchronization in Solaris 2 & Windows 2000






The Critical-Section Problem

- n processes all competing to use some shared data
- Each process has a code segment, called *critical section*, in which the shared data is accessed.
- Problem – ensure that when one process is executing in its critical section, no other process is allowed to execute in its critical section.

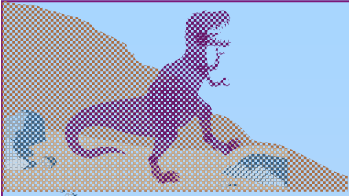




Solution to Critical-Section Problem

1. **Mutual Exclusion.** If process P_i is executing in its critical section, then no other processes can be executing in their critical sections.
2. **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.
3. **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.
 - ☐ Assume that each process executes at a nonzero speed
 - ☐ No assumption concerning relative speed of the n processes.





Initial Attempts to Solve Problem


- Only 2 processes, P_0 and P_1
- General structure of process P_i (other process P_j)
 - do** {
 - entry section*
 - critical section
 - exit section*
 - reminder section
 - } while (1);**
- Processes may share some common variables to synchronize their actions.





Algorithm 1

- Shared variables:
 - ◆ **int turn;**
initially **turn = 0**
 - ◆ **turn = i** $\Rightarrow P_i$ can enter its critical section
- Process P_i
 - do {**
 - while (turn != i) ;**
critical section
 - turn = i;**
remainder section
 - } while (1);**
- Satisfies mutual exclusion, but not progress





Algorithm 2

- Shared variables

- ◆ **boolean flag[2];**

- initially **flag [0] = flag [1] = false.**

- ◆ **flag [i] = true** $\Rightarrow P_i$ ready to enter its critical section

- Process P_i

- do {**

- flag[i] = true;**

- while (flag[j]) ;**


- critical section

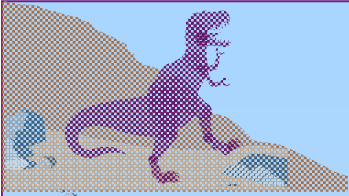
- flag [i] = false;**

- remainder section

- } while (1);**

- Satisfies mutual exclusion, but not progress requirement.






Algorithm 3

- Combined shared variables of algorithms 1 and 2.
- Process P_i
 - do** {
 - flag** [i] = true;
 - turn** = j;
 - while** (flag [j] && turn == j) ;
 - critical section
 - flag** [i] = false;
 - remainder section
 - } while** (1);
- Meets all three requirements; solves the critical-section problem for two processes.




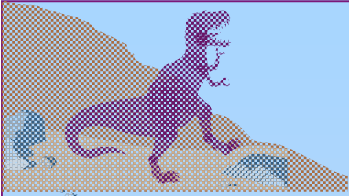


Bakery Algorithm

Critical section for n processes

- Before entering its critical section, process receives a number. Holder of the smallest number enters the critical section.
- If processes P_i and P_j receive the same number, if $i < j$, then P_i is served first; else P_j is served first.
- The numbering scheme always generates numbers in increasing order of enumeration; i.e., 1,2,3,3,3,3,4,5...





Bakery Algorithm

- Notation \leq lexicographical order (ticket #, process id #)
 - ◆ $(a,b) < c,d$ if $a < c$ or if $a = c$ and $b < d$
 - ◆ $\max(a_0, \dots, a_{n-1})$ is a number, k , such that $k \geq a_i$ for $i = 0, \dots, n-1$

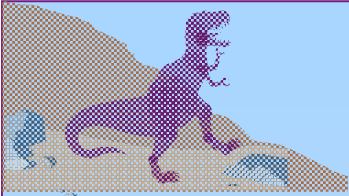
- Shared data

boolean choosing[n];

int number[n];

Data structures are initialized to **false** and **0** respectively






Bakery Algorithm

```
do {  
    choosing[i] = true;  
    number[i] = max(number[0], number[1], ..., number [n - 1])+1;  
    choosing[i] = false;  
    for (j = 0; j < n; j++) {  
        while (choosing[j]) ;  
        while ((number[j] != 0) && (number[j] < number[i])) ;  
    }  
    critical section  
    number[i] = 0;  
    remainder section  
} while (1);
```





Synchronization Hardware

- Test and modify the content of a word atomically

```
boolean TestAndSet(boolean &target) {  
    boolean rv = target;  
    target = true;  
  
    return rv;  
}
```





Mutual Exclusion with Test-and-Set

- Shared data:

boolean lock = false;

- Process P_i

do {

while (TestAndSet(lock)) ;

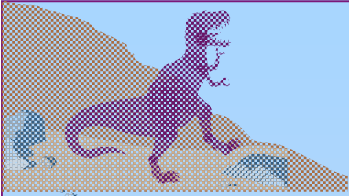
critical section

lock = false;

remainder section

}





Synchronization Hardware

- Atomically swap two variables.


```
void Swap(boolean &a, boolean &b) {  
    boolean temp = a;  
    a = b;  
    b = temp;  
}
```

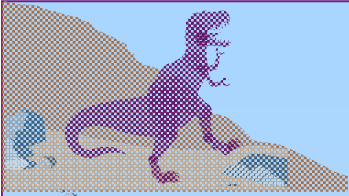




Mutual Exclusion with Swap

- Shared data (initialized to **false**):
 boolean lock;
 boolean waiting[n];
- Process P_i
 do {
 key = true;
 while (key == true)
 Swap(lock, key);
 critical section
 lock = false;
 remainder section
 }





Semaphores

- Synchronization tool that does not require busy waiting.
- Semaphore S – integer variable
- can only be accessed via two indivisible (atomic) operations


wait (S):

while $S \leq 0$ do *no-op*;
 $S--$;

signal (S):

$S++$;





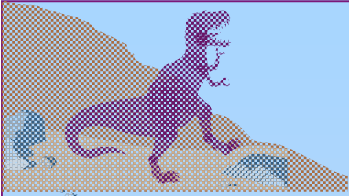
Critical Section of n Processes

- Shared data:
semaphore mutex; //initially *mutex* = 1

- Process P_i :

do {
 wait(mutex);
 critical section
 signal(mutex);
 remainder section
} **while (1);**



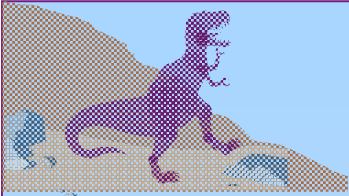


Semaphore Implementation

- Define a semaphore as a record

```
typedef struct {  
    int value;  
    struct process *L;  
} semaphore;
```
- Assume two simple operations:
 - ✦ **block** suspends the process that invokes it.
 - ✦ **wakeup(*P*)** resumes the execution of a blocked process **P**.





Implementation

- Semaphore operations now defined as

wait(S):

```
S.value--;  
if (S.value < 0) {  
    add this process to S.L;  
    block;  
}
```

signal(S):

```
S.value++;  
if (S.value <= 0) {  
    remove a process P from S.L;  
    wakeup(P);  
}
```






Semaphore as a General Synchronization Tool

- Execute B in P_j only after A executed in P_i
- Use semaphore $flag$ initialized to 0
- Code:

P_i	P_j
\vdots	\vdots
A	$wait(flag)$
$signal(flag)$	B





Deadlock and Starvation

- **Deadlock** – two or more processes are waiting indefinitely for an event that can be caused by only one of the waiting processes.

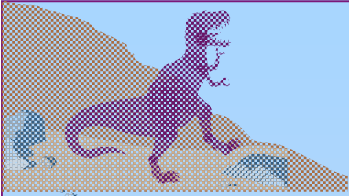
- Let S and Q be two semaphores initialized to 1

P_0
 $wait(S);$
 $wait(Q);$
 \vdots
 $signal(S);$
 $signal(Q)$

P_1
 $wait(Q);$
 $wait(S);$
 \vdots
 $signal(Q);$
 $signal(S);$

- **Starvation** – indefinite blocking. A process may never be removed from the semaphore queue in which it is suspended.






Two Types of Semaphores

- *Counting* semaphore – integer value can range over an unrestricted domain.
- *Binary* semaphore – integer value can range only between 0 and 1; can be simpler to implement.
- Can implement a counting semaphore S as a binary semaphore.





Implementing S as a Binary Semaphore

- Data structures:

```
binary-semaphore S1, S2;  
int C;
```


- Initialization:

```
S1 = 1
```

```
S2 = 0
```

```
C = initial value of semaphore S
```






Implementing S


■ *wait* operation

```
wait(S1);  
C--;  
if (C < 0) {  
    signal(S1);  
    wait(S2);  
}  
signal(S1);
```

■ *signal* operation

```
wait(S1);  
C ++;  
if (C <= 0)  
    signal(S2);  
else  
    signal(S1);
```

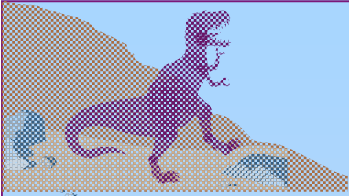




Windows 2000 Synchronization

- Uses interrupt masks to protect access to global resources on uniprocessor systems.
- Uses *spinlocks* on multiprocessor systems.
- Also provides *dispatcher objects* which may act as wither mutexes and semaphores.
- Dispatcher objects may also provide *events*. An event acts much like a condition variable.





Solaris 2 Synchronization

- Implements a variety of locks to support multitasking, multithreading (including real-time threads), and multiprocessing.
- Uses *adaptive mutexes* for efficiency when protecting data from short code segments.
- Uses *condition variables* and *readers-writers* locks when longer sections of code need access to data.
- Uses *turnstile*s to order the list of threads waiting to acquire either an adaptive mutex or reader-writer lock.

