Operating Systems "Inter-Process Communication (IPC) and synchronization"

Mathieu Delalandre University of Tours, Tours city, France <u>mathieu.delalandre@univ-tours.fr</u>

Lecture available at http://mathieu.delalandre.free.fr/teachings/operating1.html

Operating Systems "IPC and synchronization"

- 1. Introduction
- 2. Synchronization for mutual exclusion
 - 2.1. Principles of concurrency
 - 2.2. Synchronization methods for mutual exclusion
- 3. Synchronization for coordination
 - 3.1. Some problems of coordination
 - 3.2. Solving the Producer/Consumer problem
 - 3.3. Solving the multiple Producer/Consumer problem

Introduction (1)

Cooperating / independent process: a process is cooperating if it can affect (or be affected) by the other processes. Clearly, any process than shares data is a cooperating process. Any process that does not share data with any other process is independent.

Inter-process communication (IPC) refers to the set of techniques for the exchange of data among different processes. There are several reasons for providing an environment allowing IPC.

 \checkmark Information sharing: several processes could be interested in the same piece of information, we must provide a framework to allow a concurrent access to this information.

 \checkmark Modularity: we may to construct the system in a modular fashion, dividing a function of the system into separate blocks.

 \checkmark Convenience: even an individual user may work on many related tasks at the same time e.g. editing, printing and compiling a program.

 \checkmark Speedup: with parallelism, if we are interested to run faster a particular task, we must break it into sub-tasks.

Introduction (2)

Process synchronization: refers to the idea that multiple processes join up to reach an agreement or to commit a sequence of action. Clearly, any cooperating process is concerned with synchronization. We can classify the synchronization on the basis of the degree to which the processes are aware of each other.

 \checkmark **Processes unaware of each other:** are independent and not intended to work together. Although the processes are not working together, the OS must deal with the **concurrency** and **mutual exclusion** problems.

 \checkmark **Processes indirectly aware of each other:** are not necessarily aware of each other by their respective ids, but share access to objects such as an I/O buffer. Such processes exhibit **coordination** in sharing objects.

✓ **Processes directly aware of each other:** cooperate and are able to communicate with each other by process ids. These processes are designed to work jointly in some activity. Again, such processes exhibit **coordination**.

Degree of awareness	Synchronization		
Processes unaware of each other	Mutual exclusion		
Processes indirectly aware of each other	Coordination by sharing		
Processes directly aware of each other	Coordination by communication		



Operating Systems "IPC and synchronization"

- 1. Introduction
- 2. Synchronization for mutual exclusion
 - 2.1. Principles of concurrency
 - 2.2. Synchronization methods for mutual exclusion
- 3. Synchronization for coordination
 - 3.1. Some problems of coordination
 - 3.2. Solving the Producer/Consumer problem
 - 3.3. Solving the multiple Producer/Consumer problem

Principles of concurrency (1)

Inter-process communication (IPC) is a set of techniques for the exchange of data among multiple processes or threads.

Race conditions arise when separate processes of execution depend on some shared states. Operations upon shared states could result in harmful collisions between these processes.

Critical section is a piece of code that accesses a shared resource (a data structure or a device) that must not be concurrently accessed by other concurrent/cooperating processes.

Mutual exclusion: two events are mutually exclusive if they cannot occur at the same time. Mutual exclusion algorithms are used to avoid the simultaneous use of a resource by the piece of code of the critical section.

Process synchronization: refers to the idea that multiple processes join up to reach an agreement or to commit a sequence of action.



Principles of concurrency (2)

Race conditions arise when separate processes of execution depend on some shared states. Operations upon shared states could result in harmful collisions between these processes.





(1) to (7) are atomic instructions

	loop	
Р	(1)	P.in=in
	(2)	S[P.in] = P.name
	(3)	in = P.in+1
	loop	
	(4)	D.out=out
D	(5)	D.name=S[D.out]
	(6)	out = D.out+1
	(7)	print

- S the spooling directory
- in current writing index of S
- out current reading index of S
- a process
- D the printer daemon process
- X.a A data a part of a process X

Principles of concurrency (3)

Race conditions arise when separate processes of execution depend on some shared states. Operations upon shared states could result in harmful collisions between these processes. e.g. spooling with 2 processes A, B and a Daemon D

	in	A.in	B.in	S[7]	out	D.out	D.name
	7	Ø	Ø	Ø	7	6	X.name
A→1	7	7	Ø	Ø	7	6	X.name
B→1,2,3	8	7	7	B.name	7	6	X.name
A→2,3	8	7	7	A.name	7	6	X.name
D→4,5,6,7	8	7	7	A.name	8	7	A.name

initial states

A reads "in"

B reads "in", writes in "S" and increments "in"

A writes in "S", and increments "in", the harmful collision is here

D prints the file of A, the B one will be never processed

 $P \rightarrow x, y$ process P executes the instructions x, y



- Notation
 - the spooling directory
- in current writing index of S
- at current reading index of S
- a process
- the printer daemon process
- a A data a part of a process X







Principles of concurrency (4)

Critical section is a piece of code that accesses a shared resource (a data structure or a device) that must not be concurrently accessed by other concurrent/cooperating processes. A critical section will usually terminate within a fixed time, a process will have to wait a fixed time to enter it.



Principles of concurrency (5)



Mutual exclusion: two events are mutually exclusive if they cannot occur at the same time. Mutual exclusion algorithms are used to avoid the simultaneous use of a resource by the piece of code of the critical section.

Process synchronization: refers to the idea that multiple processes join up to reach an agreement or to commit a sequence of action.

Operating Systems "IPC and synchronization"

- 1. Introduction
- 2. Synchronization for mutual exclusion
 - 2.1. Principles of concurrency
 - 2.2. Synchronization methods for mutual exclusion
- 3. Synchronization for coordination
 - 3.1. Some problems of coordination
 - 3.2. Solving the Producer/Consumer problem
 - 3.3. Solving the multiple Producer/Consumer problem

Synchronization methods for mutual exclusion "Introduction" (1) 2. Busy-waiting

1. Scheduling while disabling the interrupts







Correspond to the areas of critical sections

3. Sleep and wakeup



Synchronization methods for mutual exclusion "Introduction" (2)

Methods	Approach	Туре	Process	Ordering	Starvation
disabling the interrupts	disabling the interrupts	1	>2	yes	no
Swap, TSL, CAS	1	nardware	22		
Perterson's algorithm	busy-waiting	software	2	no	possible
binary semaphore / mutex	sleep and wakeup		≥2	yes	no

Synchronization methods for mutual exclusion "Disabling the interrupts"

Disabling the interrupts: within an uniprocessor system, processes cannot have an overlapped execution. To guarantee a mutual exclusion, it is sufficient to prevent a process from being interrupted. This capability can be provided in the form of primitives defined in the OS kernel, for disabling and enabling the interrupts when entering in a critical section. e.g.



Synchronization methods for mutual exclusion

Methods	Approach	Туре	Process	Ordering	Starvation
disabling the interrupts	disabling the interrupts	1	>2	yes	no
Swap, TSL, CAS	1	nardware	22		
Perterson's algorithm	busy-waiting	software	2	no	possible
binary semaphore / mutex	sleep and wakeup		≥2	yes	no

Synchronization methods for mutual exclusion "Swap, TSL and CAS" (1)

Swap (or exchar one-shot the cont	ige) is an ent of tw	hardware to location	e instruction, exchanging i ns, atomically.	$\begin{pmatrix} n \\ \vdots \\$	1) Requ 2) set K 3) do S ^y 4) while	est the EY at wap KI e KEY	critical 1 E Y , LC equals	l section DCK 1	n with P	
SWAP KEY,LOCK	KEY	(1) copy LO (1) copy	CK atomic instruction	e.g. three proce	Run in do s 5) Relea 6) se esses A	the crit omethin ase the et LOC , B and	ical sec ng critical K at 0 C cons	etion wi sectior	th P with P	eduling
					KEYA	KEYB	KEYC	LOCK	Section	
	KEY	LOCK			Ø	Ø	Ø	0	Ø	
	\$1	\$0	LOCK at 0, KEY at 1	B→1,2,3	Ø	0	Ø	1	В	B accesses the section
SWAP KEY,LOCK	\$0	\$1	both shift their values	A→1,2,3,4,3,4,3	1	0	Ø	1	В	A is blocked
				B→4,5,6	1	0	Ø	0	Ø	B releases the section
	KEY	LOCK	Busy weiting	A→4,3	0	0	Ø	1	Α	A can access
	\$1	\$1	LOCK and KEY at 1	C→1,2,3,4,3,4	0	0	1	1	А	C is blocked
SWAP KEY,LOCK	\$1	\$1	both keep their values	A→4,5,6	0	0	1	0	Ø	A releases the section
				C→3,4	0	0	0	1	С	C can access
										1

C→5,6

 $P \rightarrow x, y$ process P executes the instructions x, y

0

0

0

Ø

0

C releases the section

Synchronization methods for mutual exclusion "Swap, TSL and CAS" (2)



process P executes the instructions x,y $P \rightarrow x.v$

0

0

0

Ø

0

C releases the section

Synchronization methods for mutual exclusion "Swap, TSL and CAS" (3)

CAS is a tradeoff to the TSL instruction checking a memory location LOCK against a test value TEST. If they are same, a swap occurs between the LOCK and a KEY value. The old **LOCK** value (before the swapping) is still returned.



\$1

\$1

nothing happens

R ← CAS LOCK,0,1

Request (1) Request the critical section with **P** (2) do R equals CAS LOCK, 0, 1 (3) while key \mathbf{R} equals 1 Run in the critical section with **P** do something Release (4) Release the critical section with **P**

set LOCK at 0 (5)

e.g. three processes A, B and C considering the scheduling

	RA	RB	RC	LOCK	Section	
	Ø	Ø	Ø	0	Ø	
B→1,2	Ø	0	Ø	1	В	B accesses the section
A→1,2,3,2,3,2	1	0	Ø	1	В	A is blocked
B→3,4,5	1	0	Ø	0	Ø	B releases the section
A→3,2	0	0	Ø	1	А	A can access
C→1,2,3,2,3	0	0	1	1	А	C is blocked
A→3,4,5	0	0	1	0	Ø	A releases the section
C→2,3	0	0	0	1	С	C can access
C→4,5	0	0	0	0	Ø	C releases the section

 $P \rightarrow x, y$ process P executes the instructions x, y

Synchronization methods for mutual exclusion

Methods	Approach	Туре	Process	Ordering	Starvation
disabling the interrupts	disabling the interrupts	1	>2	yes	no
Swap, TSL, CAS	1	nardware	22		
Perterson's algorithm	busy-waiting	software	2	no	possible
binary semaphore / mutex	sleep and wakeup		≥2	yes	no

Synchronization methods for mutual exclusion "Peterson's algorithm" (1)

The Peterson's algorithm solves the mutual exclusion problem between two processes. Entrance in the critical section is granted for a process P if the other process doesn't want to enter, or if it has given previously the priority to P.



while

wait

access

access

access

Synchronization methods for mutual exclusion "Peterson's algorithm" (2)



(6) Release the critical section with P_i (7) flag[i] = false

$\mathbf{P}_{\mathbf{i}}$ accesses the critical section if						
or	 (1) P_j sets its flags at false (2) P_j sets the turn for P_i 					

e.g. two	processes A,	В	considering	the schee	duling
0	1 ,		0		\mathcal{O}

	tarmo	fla	Section	
	turn	А	В	Section
	Ø	false	false	Ø
B→1,2	Ø	false	true	Ø
A→1,2,3,4,5,4,5	В	true	true	Ø
B→3	Α	true	true	Ø
A→4,6,7	А	false	true	A-Ø
В→4,6,7	А	false	false	B-Ø

B sets its flag at true

A is blocked because the flag of B is true and turn is set with the B value

B sets the turn variable to A

A accesses the section as the turn variable is set to A

B accesses the section as the flag of A is false

 $P \rightarrow x, y$ process P executes the instructions x,y

Synchronization methods for mutual exclusion

Methods	Approach	Туре	Process	Ordering	Starvation
disabling the interrupts	disabling the interrupts	1	>2	yes	no
Swap, TSL, CAS	1	nardware	22		
Perterson's algorithm	busy-waiting	software	2	no	possible
binary semaphore / mutex	sleep and wakeup		≥2	yes	no

Synchronization methods for mutual exclusion "binary semaphores / mutex" (1)

Semaphore is a synchronization primitive composed of a blocking queue and a variable controlled with two operations down / up.

A binary semaphore takes only the values 0 and 1. A mutex is a binary semaphore for which a process that locks the semaphore must be the process that unlocks it.

The **down** operation decreases the value of the semaphore or sleeps the current process and pushes it into the queue.



regular down

blocking down

	before	after		before	after
value	false	true	value	true	true
queue	Ø	Ø	queue	Ø	Р





Synchronization methods for mutual exclusion "binary semaphores / mutex" (2)

Semaphore is a synchronization primitive composed of a blocking queue and a variable controlled with two operations **down / up**.

A **binary semaphore** takes only the values 0 and 1. A **mutex** is a binary semaphore for which a process that locks the semaphore must be the process that unlocks it.

The **up** operation increases the value of the semaphore or wakeups a process in the queue.



	before	after		before	after
value	true	false	value	true	true
queue	Ø	Ø	queue	Р	Ø





Synchronization methods for mutual exclusion "binary semaphores / mutex" (3)

The algorithm for mutual exclusion using a binary semaphore is

sem is a semaphore, **P** is the process, (1) to (5) the instructions

- (1) before the request do something
- (2) down sem
- (3) run in the critical section with **P** do something
- (4) before the release do something
- (5) up **sem**

e.g. three processes A, B and C considering the scheduling, the solution is presented with a table

	sen	1	Section	A state	D state	Catata
	value	Q	Section	A state	D state	C state
	false	Ø	Ø	ready	ready	ready
A→1,2,3	true	Ø	А	ready	ready	ready
B→1,2	true	В	А	ready	blocked	ready
C→1,2	true	C,B	А	ready	blocked	blocked
A→4,5	true	С	A-B	ready	ready	blocked
B→3,4,5	true	Ø	B-C	ready	ready	ready
C→3,4,5	false	Ø	C-Ø	ready	ready	ready

A accesses the section, sem becomes true while accessing the semaphore, **B** blocks while accessing the semaphore, **C** blocks **A** exits and pops up **B**, **B** holds the section **B** exits and pops up **C**, **C** holds the section **C** exits and puts the semaphore to false

 $P \rightarrow x, y$ process P executes the instructions x, y

regula	r down		block	blocking down					
	before	after		before	after				
value	false	true	value	true	true				
queue	Ø	Ø	queue	Ø	Р				

|--|

unblocking up

	before	after		before	after
value	true	false	value	true	true
queue	Ø	Ø	queue	Р	Ø

Synchronization methods for mutual exclusion "binary semaphores / mutex" (4)

The algorithm for mutual exclusion using a binary semaphore is

e.g. three processes A, B and C considering the scheduling, the solution is diagram



Operating Systems "IPC and synchronization"

- 1. Introduction
- 2. Synchronization for mutual exclusion
 - 2.1. Principles of concurrency
 - 2.2. Synchronization methods for mutual exclusion
- 3. Synchronization for coordination
 - 3.1. Some problems of coordination
 - 3.2. Solving the Producer/Consumer problem
 - 3.3. Solving the multiple Producer/Consumer problem

Some problems of coordination (1)

The dinning-philosophers problem is summarized as:

- (1) five philosophers sitting at a table are doing one of the two things: eating or thinking.
- (2) a fork is placed in between each pair of adjacent philosophers.
- (3) while eating, they are not thinking, and while thinking, they are not eating.
- (4) a philosopher must eat with two forks (i.e. if thinking, none fork are used).
- (5) each philosopher can only use the forks on his immediate left and immediate right.



The **readers-writer problem** concerns synchronization of processes when accessing the same database in a R/W mode. It is summarized as:

- (1) several processes can read the database at the same time.
- (2) when at least a process reads, no one can write.
- (3) only a single process can write at the same time.
- (4) when a process writes, no one can read.



Some problems of coordination (2)

The producer-consumer (i.e. **bounded buffer**) describes how processes share a common buffer. The problem is to make sure that a process will not try to add data into the buffer if it's full, or to remove data if empty. The problem is summarized as:

(1) the producer and the consumer share a common, fixed-size, buffer,
(2) the producer puts information into the buffer, and the consumer takes it out,
(3) processes are blocked when the size limit is reached, empty for the consumer or full for the producer,
(4) processes are unblocked when the buffer recovers a regular size,
(5) we can generalize the problem to m producers and n consumers, but this extends the synchronization with a mutual exclusion while accessing the buffer.





Consumer

Operating Systems "IPC and synchronization"

- 1. Introduction
- 2. Synchronization for mutual exclusion
 - 2.1. Principles of concurrency
 - 2.2. Synchronization methods for mutual exclusion
- 3. Synchronization for coordination
 - 3.1. Some problems of coordination
 - 3.2. Solving the Producer/Consumer problem
 - 3.3. Solving the multiple Producer/Consumer problem

Solving the Producer/Consumer problem "Introduction"

Methods	Approach Type		Application problem	Coordination type	
sleep and wakeup	wakeup		Producer / Consumer	coordination by communication	
semaphore	sleep and	software		coordination by sharing	
semaphore / mutex	wakeup		Multiple		
monitor			Producers / Consumers		

Solving the Producer/Consumer problem "sleep wakeup" (1)

Sleep and wakeup are atomic actions to change the states of processes for synchronization.

The producer/consumer algorithm is



e.g. two processes P, C with a successful synchronization

consumer, producer are processes

consumer

loop

- (1) if **buffer** is empty
- (2) sleep
- (3) pop item from buffer
- (4) if **buffer** was full (i.e. actual size = n-1)
- (5) wakeup **producer**

			Ŧ	
	buffer	P state	C state	
	0	awake	sleepy	
P→1,3,4,5	1	awake	awake	P wakeups C
C→3,4	0	awake	awake	C restarts at Program Counter
P→1,3,4,5	1	awake	awake	here is a lost wakeup
C→1,3,4,1,2	0	awake	sleepy	when empty, C will sleep

 $P \rightarrow x, y$ process P executes the instructions x, y

producer

loop

- (1) if **buffer** is full
- (2) sleep
- (3) push a new **item** in **buffer**
- (4) if **buffer** was empty (i.e. actual size =1)
- (5) wakeup **consumer**

Solving the Producer/Consumer problem "sleep wakeup" (2)

Sleep and wakeup are atomic actions to change the states of processes for synchronization.

The producer/consumer algorithm is



consumer, producer are processes

e.g. two processes P, C with a synchronization and failure

consu	mer		buffer	P state	C state	
lo	op		0	awake	awake	
(1)	if buffer is empty	C→1	0	awake	awake	
(2) (3)	sleep	P→1,3,4,5	1	awake	awake	here is a lost wakeup C blocked on sleep
(3) (4)	4) if buffer was full (i.e. actual size = $n-1$)	C→2	1	awake	sleepy	
(5)	wakeup producer	P→1,3,4	2	awake	sleepy]]
		P→1,3,4	3	awake	sleepy	← fill in buffer
produ	cer					
loop (1) if buffer is full		P→1,2	n	sleepy	sleepy	P, C will sleep
		$P \rightarrow x, y$ process	P executes	the instruct	ions x,y	for always

(1) (2) sleep

- (3) push a new item in buffer
- (4) if **buffer** was empty (i.e. actual size =1)
- (5) wakeup consumer

Solving the Producer/Consumer problem "sleep wakeup" (3)

Sleep and wakeup with a wakeup waiting bit is an extension of the method to support the lost wakeups.

The producer/consumer algorithm is

consumer, producer are processes



e.g. two processes P, C with a successful synchronization

consumer loop			1 66	ww	bits	D ()		
			buffer	Р	C	P state	C state	
(1)	if buffer is empty		0	0	0	awake	awake	
(2) (3)	sleep pop item from buffer	C→1	0	0	0	awake	awake	
(4)	if buffer was full (i.e. actual size = $n-$	1) $P \rightarrow 1, 3, 4, 5$	1	0	1	awake	awake	C gets a bit
(5)	wakeup producer	C→2	1	0	0	awake	awake	C uses the bit
		P→1,3,4	2	0	0	awake	awake	
produ	cer	C→3,4	1	0	0	awake	awake	41
lc	oop							will go on

(1) 11 **buffer** 1s full

(2) sleep

- (3) push a new item in buffer
- (4) if **buffer** was empty (i.e. actual size =1)
- (5) wakeup consumer

 $P \rightarrow x, y$ process P executes the instructions x, y

Solving the Producer/Consumer problem

Methods	Approach	Туре	Application problem	Coordination type	
sleep and wakeup			Producer / Consumer	coordination by communication	
semaphore	sleep and	software		coordination by sharing	
semaphore / mutex	wakeup		Multiple		
monitor			Producers / Consumers		

Solving the Producer/Consumer problem "semaphores" (1)

Semaphore is a synchronization primitive composed of a blocking queue and a variable controlled with two operations **down / up**.

A counting (or a general) semaphore is not a binary semaphore, it embeds a variable covering the range $[0, +\infty]$.





Solving the Producer/Consumer problem "semaphores" (2)

The algorithm for solving the producer/consumer problem with semaphores is



Solving the Producer/Consumer problem "semaphores" (3)

The algorithm for solving the producer/consumer problem with semaphores is



fill = 0, empty = n are semaphores	rocesses P, C with n=2								
buffer is the data structure		1	fill	fill		empty		Catata	
consumer		buffer	value	Q	value	Q	P state	C state	
loop		0	0	Ø	2	Ø	ready	ready	
(1) down fill	$C \rightarrow 1$	0	0	С	2	Ø	ready	blocked	C sleeps at down on fill
(2) pop item from buffer	P→1,2,3	1	0	Ø	1	Ø	ready	ready	P wakeups C at up on fill
(3) up empty	C→2,3	0	0	Ø	2	Ø	ready	ready	next scheduling, C restarts on pop
	P→1,2,3	1	1	Ø	1	Ø	ready	ready	fill in huffor
producer	P→1,2,3	2	2	Ø	0	Ø	ready	ready	
loop	P→1	2	2	Ø	0	Р	blocked	ready	P stopped at down on empty

producer

loop

- (1) down **empty**
- push a new item in buffer (2)
- up fill (3)

regular down	
--------------	--

	blocking	uo

	before	after
value	2	1
queue	Ø	Ø

blocking down	
---------------	--

 $P \rightarrow x, y$ process P executes the instructions x,y

after		before	after
1	value	0	0
Ø	queue	Ø	Р

regular up

	before	after
value	2	3
queue	Ø	Ø

unblocking up

	before	after
value	0	0
queue	Р	Ø

e g two processes P C with n=?

Operating Systems "IPC and synchronization"

- 1. Introduction
- 2. Synchronization for mutual exclusion
 - 2.1. Principles of concurrency
 - 2.2. Synchronization methods for mutual exclusion
- 3. Synchronization for coordination
 - 3.1. Some problems of coordination
 - 3.2. Solving the Producer/Consumer problem
 - 3.3. Solving the multiple Producer/Consumer problem

Solving the multiple Producer/Consumer problem "Introduction"

Methods	Approach	Туре	Application problem	Coordination type											
sleep and wakeup			Producer / Consumer	coordination by communication											
semaphore	sleep and	software	software	software	software	software	software	software	software	software	software	software	software		
semaphore / mutex	wakeup		Multiple	coordination by											
monitor			Producers / Consumers	sharing											

Solving the multiple Producer/Consumer problem "semaphores - mutex" (1)

fill and empty work from a buffer having a size bounded between 0 to n.



The buffer is treated as a circular storage, and pointer values must be expressed modulo the size of the buffer. Therefore, we can have In > Out or In < Out depending the access case.



The pop and push operations are then not atomic.

pop		push	1
(1)	w = b[out]	(1)	b[in] = v
(2)	out = $(out+1)\%(n+1)$	(2)	in = (in+1)%(n+1)

In addition, the buffer slots are data-dependent (e.g. byte, double, data structure, etc.), the read/write operations b[] could be instruction sets.

The one-to-one solution to the Producer/Consumer problem with a bounded-buffer.

fill = 0, empty = n are semaphores
buffer is the data structure

consumer

- loop
- (1) down fill

(2)
$$w=b[out]$$

(3) out =
$$(out+1)\%(n+1)$$

(4) up empty

producer

loop

- (1) down empty
- (2) b[in] = v
- (3) in = (in+1)%(n+1)
- (4) up fill



Solving the multiple Producer/Consumer problem "semaphores - mutex" (3)

The solution is then to protect access to buffer with a mutex. The general algorithm for solving the multiple producer/consumer problem with semaphore becomes

fill = 0, empty = n are semaphores
mutex is a mutex
buffer is the data structure

consumer

loop

- (1) down fill
- (2) down **mutex**
- (3) pop **item** from **buffer**
- (4) up mutex
- (5) up empty

producer

loop

- (1) down **empty**
- (2) down **mutex**
- (3) push a new item in buffer
- (4) up mutex
- (5) up fill



Solving the multiple Producer/Consumer problem "semaphores - mutex" (4)

The solution is then to protect access to buffer with a mutex. The general algorithm for solving the multiple producer/consumer problem with semaphore becomes

fill = 0, empty = n are semaphores

mutex is a mutex

buffer is the data structure

e.g. two producers P1, P2, one consumer C with n = 4

consumer		huffor	fill	fill		empty		mutex		
		buller	value	Q	value	Q	value	Q	Section	
loop		0	0	Ø	4	Ø	false	Ø	Ø]
(1) down fill	P1→1,2	0	0	Ø	3	Ø	true	Ø	P1	1
(2) down mutex	P2→1,2	0	0	Ø	2	Ø	true	P2	P1]1
 (3) pop item from buffer (4) up mutex 	P1→3,4,5	1	1	Ø	2	Ø	true	Ø	P2] t
(5) up empty	Р2→3	2	1	Ø	2	Ø	true	Ø	P2	1
	C→1,2	2	0	Ø	2	Ø	true	С	P2](
producer	Р2→4,5	2	1	Ø	2	Ø	true	Ø	С] t
loop	C→3,4,5	1	1	Ø	3	Ø	false	Ø	Ø	1
(1) down empty	L									-

 $P \rightarrow x, y$ process P executes the instructions x, y

P2 is blocked on mutex while P1 accesses the buffer, here mutual exclusion applies

C is blocked on mutex while P2 accesses the buffer, here there is no mutual exclusion

(2) down **mutex**

(3) push a new **item** in **buffer**

(4) up mutex

(5) up fill

Solving the multiple Producer/Consumer problem "semaphores - mutex" (5)

Inverting the code for solving the multiple producer/consumer problem will result in a deadlock.

fill = 0, empty = n are semaphores
mutex is a mutex
buffer is the data structure

e.g. one producers P, one consumer C



	empty		mutex			Distate	Catata	
	value	Q	value	Q	Section	P state	C state	
	0	Ø	false	Ø	Ø	ready	ready	
P→1,2	0	Р	true	Ø	Р	blocked	ready	
C→1	0	Р	true	С	Р	waiting	blocked	P,C will sleep for always

 $P \rightarrow x, y$ process P executes the instructions x, y



Solving the multiple Producer/Consumer problem

Methods	Approach Type Application prob		Application problem	Coordination type
sleep and wakeup		software	Producer / Consumer	coordination by communication
semaphore	sleep and			
semaphore / mutex	wakeup		Multiple	coordination by sharing
monitor			Producers / Consumers	

Solving the multiple Producer/Consumer problem "monitor" (1)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.





Solving the multiple Producer/Consumer problem "monitor" (2)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.





Solving the multiple Producer/Consumer problem "monitor" (3)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.

	Mesa	Hoare				
wait	common implementation					
signal	specific to Mesa, also called notify	specific to Hoare				



Solving the multiple Producer/Consumer problem "monitor" (4)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.



The wait operation



After a wait operation, a process moves to the queue of the condition variable.

	in all the cases			
before the wait	P _k in the monitor			
after the wait	P_k in the condition queue			

Solving the multiple Producer/Consumer problem "monitor" (5)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.



The signal operation with a Mesa implementation, also called notify



If at least one process is in the condition queue, it is notified but the signaling process continues. The signaled process will be resumed at some convenient future time, when the monitor will be available.

	if queue empty	otherwise		
before the signal	P_k in the	P_k in the monitor, P_j in the condition queue		
after the signal	normal exit	P_k in the monitor, P_j in the entry queue		

Solving the multiple Producer/Consumer problem "monitor" (6)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.

The Buhr's representation of the Mesa monitor

Inotation	
a.q, b.q	are the queues of the condition variables a, b
e.q	queue of processes that want to enter
m	the monitor with one process at a time
enter	when a process requests the monitor
access	when a process gets the monitor
exit	when a process exits the monitor
wait	when a process moves after a wait operation
notified	when a process leaves the queue of a condition variable following a notify operation

Notation

Solving the multiple Producer/Consumer problem "monitor" (7)

The bounded-buffer algorithm with multiple consumers and producers, using a Mesa monitor.

process P executes the instructions x,y $P \rightarrow x, y$

producer

consumer

add item

remove item

(1)

(2)

(3)

(4)

(5)

(1)

(2)

(3)

(4)

(5)

loop

loop

(0)

Solving the multiple Producer/Consumer problem "monitor" (9)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.

The signal operation with a Hoare implementation

If at least one process is in the condition queue, it runs immediately after the signal operation. The signaling process will be pushed in a specific access queue.

	if queue empty	otherwise		
before the signal	P _k in the	P_k in the monitor, P_j in the condition queue		
after the signal	monitor, normal exit	P_j in the monitor, P_k moves to a specific access queue called signal		

Solving the multiple Producer/Consumer problem "monitor" (10)

A **monitor** is a special piece of code, associated to condition variables, that are providing mutual exclusion within the monitor. Special rules are applied to the scheduler and memory:

- 1. only one process at a time can access the monitor,
- 2. irregular in/out of monitor by processes are controlled with two operations, **wait** and **signal**, to be applied on the condition variables that are close to semaphore mechanisms,
- 3. monitors are given in two implementations, Mesa and Hoare.

The Buhr's representation of a Hoare monitor

are the queues of the condition variables a, b
queue of processes that want to enter
queue of processes that have been pushed out after a signal operation
the monitor with one process at a time
when a process requests the monitor
when a process gets the monitor
when a process exits the monitor
when a process moves after a wait operation
when a process leaves the queue of a condition variable following a signal operation
when a process moves out after a successful signal operation

NT / /*

Solving the multiple Producer/Consumer problem "monitor" (11)

The bounded-buffer algorithm with multiple consumers and producers, using a Hoare monitor.

Solving the multiple Producer/Consumer

producer

loop call add new item (0)

consumer

loop call remove item (0)

monitor ProducerConsumer full = 0, empty = n are conditions **count** is a numerical value

add item

- if count equals N (1)
- (2)wait on full
- push new item in buffer (3)
- increment count (4)
- (5) signal on empty

remove item

- if **count** equals 0 (1)
- (2)wait on empty
- (3) pop item from buffer
- (4) decrement count
- signal on full (5)

problem "monitor" (12)

Extend the previous problem with an Hoare monitor: - Scheduling between the (E)ntry and the (S)ignal queues is a Round Robin with a time slice 3/4 (E) and 1/4 (S). At the turn 1, the time slice starts with (E).

			Conditions		a	entry queue	T		
b	buffer	count	full	empty	signal	Section	\rightarrow	Turn	
	0	0	Ø	C1	Ø	Ø	Ø	Ø	
₽1→0,1,3,4	1	1	Ø	C1	Ø	P1	Ø	1 (E)	P1 enters/accesses
€2→0	1	1	Ø	C1	Ø	P1	C2	1 (E)	C2 enters
P1→5	1	1	Ø	Ø	P1	P1-C1	C2	1 (E)	P1 blocked on (5)
C1→3,4,5,0	0	0	Ø	Ø	P1	C1-Ø	<mark>C1</mark> ,C2	Signalled	C1 signalled on (3)
₽3→0	0	0	Ø	Ø	P1	Ø	P3 ,C1,C2	Ø	P3 enters
P2→0	0	0	Ø	Ø	P1	Ø	P2,P3,C1,C2	Ø	P2 enters
€2→1,2	0	0	Ø	C2	P1	C2-Ø	P2,P3,C1	2 (E)	C2 blocked on (2)
C1→1,2	0	0	Ø	C 1,C2	P1	C1-Ø	P2,P3	3 (E)	C1 blocked on (2)
₽1→1,3,4,5	1	1	Ø	C1	P1-P1	P1-C2	P2,P3	4 (S)	P1 loops on s.q
€2→3,4,5,0	0	0	Ø	C1	P1	C2-Ø	C2,P2,P3	Signalled	C2 signalled on (3)
P3→1,3,4,5	1	1	Ø	Ø	P3 ,P1	P3-C1	C2,P2	1 (E)	P3 blocked on (5)

process P executes the instructions x,y $P \rightarrow x, y$