


Synchronization (Part 2)



Learning Objectives

Synchronization (Part 2)

- ▶ Understand and apply the monitor concept to solve synchronization problems in concurrent programs
- ▶ Understand and apply the concept of condition variables in monitors
- ▶ Understand and use the readers-writers concurrency design pattern
- ▶ Use synchronization constructs from Pthreads library and Java packages
- ▶ Understand synchronization issues with respect to signals
- ▶ Learn about the client-server model using pipes and named pipes

Monitors

A **monitor** is designed to allow concurrency while retaining the advantage of a structured construct.

- ▶ Each monitor is entrusted with a specific task and in turn it will have its own privileged data and instructions.
- ▶ Entry to the monitor by one process excludes entry by any other process.
- ▶ Since all processing is encapsulated inside the monitor, the programmer cannot accidentally change the monitor variables.
- ▶ A system would be more robust and efficient if we have a monitor for each specific task that involves multiple processes.

Condition Variables

Synchronization (Part 2)

Monitors include condition variables to allow processes to signal each other. A **condition variable** is a global structure inside the monitor that supports three operations.

- ▶ **wait()**: Suspends the invoking process until it receives another signal from another process.
- ▶ **signal()**: Resumes exactly one process if any processes are waiting on the condition variable. If no process is waiting, then the signal has no effect.
- ▶ **queue()**: Returns **TRUE** if there is at least one process suspended on the condition variable, and **FALSE** otherwise.

Queue Management in the Monitor

Synchronization (Part 2)

Signal takes effect immediately (Tony Hoare Semantics).

- ▶ The signal statement should always be the last statement before a process leaves the monitor. This allows the signaled process to be woken up immediately. So a process that gets past the `wait()` can be guaranteed that the condition is true.
- ▶ In case we have processes waiting to enter the monitor as well as processes waiting on a condition variable, then the next process to be allowed in the monitor is taken from the queue of processes waiting on the condition variable.
- ▶ The processes waiting in the queue on a condition variable can be woken by the scheduler in either **arbitrary**, **first-come-first-served**, or **priority** order. In the priority scheduling case we will modify the `wait()` to have an argument `wait(priority)`, where **priority** is an integer with smaller integers implying higher priority.

Signal does not have effect immediately (Per Brinch Hansen Semantics).

The signal is recorded and delivered after the signaling process leaves the monitor. In this case a process on waking up has to test the condition again before proceeding. This would lead to fewer context switches.

Bank Account Monitor (pseudocode)

Synchronization
(Part 2)

```
monitor sharedBalance
{
    int balance;
    public:
        credit(int amount) { balance = balance + amount;}
        debit(int amount) { balance = balance - amount;}
}
```

Simulation of a Binary Semaphore by a Monitor

Synchronization
(Part 2)

```
monitor SemaphoreSimulation {
    boolean busy = FALSE;
    condition notbusy;
    public:
        down() {
            if (busy) notbusy.wait();
            busy = TRUE;
        }

        up() {
            busy = FALSE;
            notbusy.signal();
        }
}
```

Readers and Writers

Synchronization (Part 2)

Suppose a resource is to be shared among a community of processes of two types: *readers* and *writers*.

- ▶ A *reader* process can share the resource with any other reader process, but not with any writer process.
- ▶ A *writer* process requires exclusive access to the resource whenever it acquires any access to the resource.
- ▶ Several different policies can be implemented for managing the shared resource.
 - ▶ *Policy 1: Readers preference.* As long as a reader holds the resource and there are new readers arriving, any writer must wait for the resource to become available. Writer can starve.
 - ▶ *Policy 2: Writers preference.* When a writer process requests access to the shared resource, any subsequent reader process must wait for the writer to gain access to the shared resource and then release it. Reader can starve.
 - ▶ *Policy 3: Give neither preference.* Neither readers nor writers can starve.

Readers and Writers Monitor: Policy 1

Synchronization (Part 2)

```
monitor ReadersWriters {
    int numberOfReaders = 0;
    boolean busy = FALSE;
    condition okToRead, okToWrite;
public:
    void startRead() {
        if (busy) okToRead.wait();
        numberOfReaders = numberOfReaders + 1;
        okToRead.signal();
    }
    void finishRead() {
        numberOfReaders = numberOfReaders - 1;
        if (numberOfReaders == 0) okToWrite.signal();
    }
    void startWrite() {
        if (numberOfReaders > 0 || busy) okToWrite.wait();
        busy = TRUE;
    }
    void finishWrite() {
        busy = FALSE;
        if (okToRead.queue())
            okToRead.signal();
        else
            okToWrite.signal();
    }
}
```

Readers and Writers Monitor: Policy 2

Synchronization (Part 2)

```
monitor ReadersWriters {
    int numberOfReaders = 0;
    boolean busy = FALSE;
    condition okToRead, okToWrite;
public:
    void startRead() {
        if (busy || okToWrite.queue()) okToRead.wait();
        numberOfReaders = numberOfReaders + 1;
        okToRead.signal();
    }
    void finishRead() {
        numberOfReaders = numberOfReaders - 1;
        if (numberOfReaders == 0) okToWrite.signal();
    }
    void startWrite() {
        if (numberOfReaders > 0 || busy) okToWrite.wait();
        busy = TRUE;
    }
    void finishWrite() {
        busy = FALSE;
        if (okToWrite.queue())
            okToWrite.signal();
        else
            okToRead.signal();
    }
}
```

A Resource Allocation Monitor

Synchronization
(Part 2)

```
monitor ResourceAllocator {
    boolean busy = FALSE;
    condition resource;

    public:
        acquire(time: integer) {
            // time = How long is the resource needed for
            if (busy) resource.wait(time);
            busy = TRUE;
        }

        release() {
            busy = FALSE;
            resource.signal();
        }
}
```

-We assume priority scheduling for the wait queue.

Examples of using Monitors

Synchronization
(Part 2)

Other monitor examples:

- ▶ solve the producer consumer problem using monitors.
- ▶ an alarm clock monitor.
- ▶ a monitor that allows access to a file by no more than n processes.
- ▶ a monitor that allows access to a group of processes as long as the sum total of their process ids doesn't exceed a certain integer, say n .

The idea is to use the weakest scheduling that we can get away with.

(arbitrary < FIFO < Priority)

An Alarm Monitor

Synchronization
(Part 2)

Write a monitor that implements an *alarm clock* that enables a calling program to delay itself for a specified number of time units (*ticks*). You may assume the existence of a real hardware clock, which invokes a procedure *tick* in your monitor at regular intervals.

```
monitor AlarmClock {
    int now = 0;
    condition TimeToWakeup;
public:
    void SetAlarm (int period) {
        (* fill in the details here ... *)
    }
    void tick() {
        now++;
        TimeToWakeup.signal();
    }
}
```

An Alarm Monitor (solution)

Synchronization
(Part 2)

```
monitor AlarmClock{
    int now = 0;
    condition TimeToWakeup;
public:
    void SetAlarm(int period){
        period = period + now; //at what time to wake up
        while (now < period) TimeToWakeup.wait(period);
        TimeToWakeup.signal(); //Wake up other processes due now
                                //A cascaded wake up like the one
                                //we used in the Readers and Writers problem
    }
    void tick(void) {
        now = now + 1;
        TimeToWakeup.signal();
    }
}
```

Note: The variable `now` and `period` may overflow (if your operating system is not rebooted for many decades!).

See [synchronization-part2/monitors/java/alarm.monitor](#) for a Java solution.

Writing Monitors with Pthreads library

Synchronization (Part 2)

- ▶ The POSIX threads (Pthreads) library provides locks, semaphores and condition variables. Combining these concepts makes it possible to write monitors in C/C++.
- ▶ The idea is to have a monitor class (or class like file in C). For example, we may have `monitor.c` and `monitor.h`. Associate a mutex lock with it. In each method of the class, we acquire the lock upon entering the method. We unlock the mutex before returning from each method.
- ▶ The condition variables in Pthreads library are the same as the condition variables used in the monitor concept.
- ▶ How to protect a class where public functions can call each other? Use a wrapper function for each public function. The wrapper function locks the mutex, calls the real function and then unlocks the mutex.

Account: A Monitor Example

Synchronization
(Part 2)

- ▶ Account.h
- ▶ Account.c
- ▶ TestAccount.c

Monitors Support in POSIX threads

Synchronization
(Part 2)

Mutexes and Semaphores have already been discussed previously. The following are the functions provided in Pthreads library for supporting condition variables.

```
pthread_cond_t condvar;  
int pthread_cond_init(pthread_cond_t *restrict cond,  
                      const pthread_condattr_t *restrict attr);  
int pthread_cond_destroy(pthread_cond_t *cond);  
int pthread_cond_wait(pthread_cond_t *restrict cond,  
                      pthread_mutex_t *restrict mutex);  
int pthread_cond_timedwait(...);  
int pthread_cond_broadcast(pthread_cond_t *cond);  
int pthread_cond_signal(pthread_cond_t *cond);
```

See man pages for more information.

Notes on pthread_cond_wait and pthread_cond_timedwait:

These functions always have an associated mutex (the one that is being used to enforce mutual exclusion in the monitor). These functions atomically release the mutex and cause the calling thread to block on the condition variable. Upon successful return, the mutex shall have been locked and shall be owned by the calling thread.

Alarm Monitor in Pthreads

Synchronization
(Part 2)

- ▶ AlarmClock.h
- ▶ AlarmClock.c
- ▶ TestAlarm.c

Producers Consumers in Pthreads

Synchronization (Part 2)

- ▶ A simple producer consumer example that uses an array of buffers and conditional variables: [producers-consumers.c](#)
- ▶ How would you modify the producer-consumer example so that it stops cleanly on receiving SIGTERM, SIGINT or SIGHUP signals?
- ▶ SIGINT is generated by typing Ctrl-c on the keyboard. SIGTERM or SIGHUP can be generated via the [kill](#) command.

```
kill -s TERM pid  
kill -s HUP pid
```

where the process can be found with the [ps aux](#) command.

Java Synchronization (1)

Synchronization (Part 2)

- ▶ Java threads are **preemptible** and **native**. The programmer should not make any timing assumptions.
- ▶ Threads have **priorities** that can be changed (increased or decreased).
- ▶ This implies that multiple threads will have race conditions (read/write conflicts based on time of access) when they run. The programmer has to resolve these conflicts.
- ▶ Example of a race condition. See **Account.java** and **TestAccount.java** files.

Java Synchronization (2)

Synchronization (Part 2)

- ▶ Java has `synchronized` keyword for guaranteeing mutually exclusive access to a method or a block of code. Only one thread can be active among all synchronized methods and synchronized blocks of code in a class.

```
// Only one thread can execute the update method at  
// a time in the class.  
synchronized void update() { //... }
```

```
// Access to individual datum can also be synchronized.  
// The object buffer can be used in several classes,  
// implying synchronization across classes.  
  
synchronized(buffer) {  
    this.value = buffer.getValue();  
    this.count = buffer.length();  
}
```

- ▶ Every Java object has an implicit monitor associated with it to implement the `synchronized` keyword. Inner class has a separate monitor than the containing outer class.
- ▶ Java allows **Reentrant Synchronization**, that is, a thread can reacquire a lock it already owns. For example, a `synchronized` method can call another `synchronized` method.

Java Synchronization (3)

Synchronization (Part 2)

- ▶ The `wait()` and `notify()` methods (of the `Object` class) allow a thread to give up its hold on a lock at an arbitrary point, and then wait for another thread to give it back before continuing.
- ▶ Another thread must call `notify()` for the waiting thread to wakeup. If there are other threads around, then there is no guarantee that the waiting thread gets the lock next. **Starvation** is a possibility. We can use an overloaded version of `wait()` that has a timeout.
- ▶ The method `notifyAll()` wakes up all waiting threads instead of just one waiting thread.

Example with wait()/notify()

Synchronization
(Part 2)

```
class MyThing {
    synchronized void waiterMethod() {
        // do something and then wait
        // This gives up the lock and puts the calling
        // thread to sleep
        wait();
        // continue where we left off
    }

    synchronized void notifierMethod() {
        // do something
        // notifier the waiter that we've done it
        notify();
        //do more things
    }

    synchronized void relatedMethod() {
        // do some related stuff
    }
}
```

Writing Monitors with Java

Synchronization (Part 2)

- ▶ Java directly implements the notion of monitors with the `synchronized` keyword. Any Java class can be made into a monitor by following two rules.
 - ▶ Add the `synchronized` keyword in front of each method declaration.
 - ▶ Make sure that there are no directly accessible class variables. For example, make all class variables be private (which is the recommended Object-Oriented practice anyways).
- ▶ Every java object has a `wait()`, `notify()`, `notifyAll()` methods that correspond to monitor concepts.

Example 1: Bank Account

Synchronization
(Part 2)

See examples in the folder: `monitors/java/account`

- ▶ Race conditions: `Account.java`, `TestAccount.java`
- ▶ Thread safe version using `synchronized` keyword: `ThreadsafeAccount.java`

Example 2: Producer/Consumer Problem

Synchronization (Part 2)

- ▶ A *producer* thread creates messages and places them into a queue, while a *consumer* thread reads them out and displays them.
- ▶ The queue has a maximum depth.
- ▶ The producer and consumer don't operate at the same speed. In the example, the producer creates messages every second but the consumer reads and displays only every two seconds.

How long will it take for the queue to fill up? What will happen when it does?

Example 2 (contd.)

See examples in the folder:

`monitors/java/producers-consumers`

- ▶ Producer and Consumer sharing a synchronized queue.
- ▶ See examples: `SharedQueue.java`, `Producer.java`, `Consumer.java`, `PC.java`

Example 3: Alarm Monitor

Synchronization
(Part 2)

- ▶ See example in the folder `monitors/java/alarm.monitor`

Condition variables support in MS Windows API

Synchronization
(Part 2)

Condition variables were introduced in Vista and Windows Server 2008. Condition variables aren't supported in Windows Server 2003 and Windows XP/2000.

`CONDITION_VARIABLE` `condvar`;

`InitializeConditionVariable`: Initializes a condition variable.

`SleepConditionVariableCS`: Sleeps on the specified condition variable and releases the specified critical section as an atomic operation.

`SleepConditionVariableSRW`: Sleeps on the specified condition variable and releases the specified SRW lock as an atomic operation.

`WakeAllConditionVariable`: Wakes all threads waiting on the specified condition variable.

`WakeConditionVariable`: Wakes a single thread waiting on the specified condition variable.

CS stands for critical section. SRW stands for Slim Read Write locks.

From the MSDN Library documentation.

Condition variables example in MS Windows API

Synchronization (Part 2)

```
CRITICAL_SECTION CritSection;
CONDITION_VARIABLE ConditionVar;
void PerformOperationOnSharedData()
{
    EnterCriticalSection(&CritSection);
    // Wait until the predicate is TRUE
    while( TestPredicate() == FALSE ) {
        SleepConditionVariableCS(&ConditionVar, &CritSection, INFINITE)
    }
    // The data can be changed safely because we own the
    // critical section
    ChangeSharedData();
    LeaveCriticalSection(&CritSection);
    // If necessary, signal the condition variable by calling
    // WakeConditionVariable or WakeAllConditionVariable
    // so other threads can wake
}
```

From the MSDN Library documentation.

Producers Consumers in MS Windows API

Synchronization
(Part 2)

- ▶ Code example: [producers-consumers.c](#)

Inter-Process Communication (IPC)

Synchronization (Part 2)

- ▶ **Semaphores** and **Signals** are useful for synchronization but not for sending information among processes.
- ▶ **Monitors** allow communication using *shared memory*.
- ▶ **Message passing** allows communication without using shared memory. The message passing mechanism copies the information from the address space of one process to the address space of another process.
- ▶ What's a message? How do we send a message? Why is the operating system needed as an intermediary to send messages? (for independent processes, since for threads we can always use shared memory).

Message Passing Using Mailboxes

Synchronization (Part 2)

- ▶ Mailboxes in user space versus in system space.
- ▶ **Message protocols**. A message may contain an instance of a C structure, may be a string of ASCII characters etc. The format is between the sender and receiver to decide a priori or negotiate.
- ▶ **Message headers** may contain information such as the sending process id, the receiver process id, number of bytes in the message, message type etc. Messages may not even have a header depending on how they are implemented.
- ▶ Performance issues. **Copy-on-write** mechanism can be used to improve the performance.

Message Passing Primitives

Synchronization (Part 2)

- ▶ `send()`.
 - ▶ *synchronous send*. Blocks the sending process until the message is received by the destination process. This synchronization is an example of a producer-consumer problem.
 - ▶ *asynchronous send*. The message is delivered to the receiver's mailbox and then the sender is allowed to continue without waiting for the receiver to read the message. The receiver doesn't even have to ever retrieve the message.
- ▶ `receive()`.
 - ▶ *blocking receive*. If there is no message in the mailbox, the process suspends until a message is placed in the mailbox. A receive operation is analogous to a resource request.
 - ▶ *non-blocking receive*. The process is allowed to query the mailbox and then control is returned immediately either with a message if there is one in the mailbox or with a status that no message is available.

Message passing and Semaphores

Synchronization (Part 2)


- ▶ How to simulate a semaphore using message passing?
- ▶ How to simulate a monitor using message passing?
- ▶ How to implement message passing using mailboxes (with all processes having access to shared memory) and semaphores?
- ▶ How to implement message passing using mailboxes (with all processes having access to shared memory) and monitors?



Message Passing using Semaphores

Synchronization (Part 2)

- ▶ A mailbox area in shared memory will be used to hold the mailboxes. Each mailbox contains an array of message slots, named *mail*. Each mailbox has variables *numSlots* and *nextSlot*, to keep track of the number of slots and the next available slot. It also has two queues for waiting threads.
- ▶ Each process has an associated semaphore s_i (initially zero) on which it will block when a *send/receive* must block.
- ▶ A global *mutex* is used to ensure mutual exclusion in accessing the mailbox area in shared memory to prevent race conditions.



Based on the above, now sketch out the pseudo-code for *send/receive*.


How can we improve the performance of the above implementation?



Message Passing using Monitors

Synchronization (Part 2)

- ▶ Each *mailbox* is managed by a monitor. Each mailbox contains an array of message slots, named *mail*. Each mailbox has variables *numSlots* and *nextSlot*, to keep track of the number of slots and the next available slot. It also has two queues for waiting threads.
- ▶ Each process has a condition variable (say *mboxCheck*) on which it will block when a **send/receive** must block.



Based on the above, sketch out the pseudo-code for **send/receive** methods.

More on Signals

- ▶ Two main system calls that deal with signals: `signal()` and `sigaction()`. The `sigaction()` call is consistent across various systems whereas `signal()` is not necessarily consistent across systems but is simpler to use. See man page for details on both.
- ▶ Examples:
 - ▶ `signal-ex1.c`
 - ▶ `signal-ex2.c`

Signal Handling after Exec

Synchronization
(Part 2)

When a program is exec'd the status of all signals is either default or ignore. The exec functions changes the disposition of any signals that are being caught to their default action (why?) and leaves the status of all other signals alone.

For example:

```
if (fork()==0) { // the child process
    if (signal(SIGINT, SIG_IGN) == SIG_ERR)
        err_ret("failed to set SIGINT behavior");
    if (signal(SIGTSTP, SIG_IGN) == SIG_ERR)
        err_ret("failed to set SIGTSTP behavior");
    execvp(program, argv);
    // the exec'd program will ignore the
    // signals SIGINT and SIGTSTP
```

Signal Handling for an Application

Synchronization
(Part 2)

How an application ought to set its signal handling:

An application process should catch the signal only if the signal is not currently being ignored.

```
int sig_int(), sig_quit();

if (signal(SIGINT, SIG_IGN) != SIG_IGN)
    signal(SIGINT, sig_int);
if (signal(SIGQUIT, SIG_IGN) != SIG_IGN)
    signal(SIGQUIT, sig_quit);
```

See code example: [signal-ex3.c](#)

Other Signal Issues

Synchronization (Part 2)

- ▶ Signals are set to their default after being caught (under Linux and System V UNIX semantics). This can be changed by using `sigaction()` instead of the `signal()` system call.
- ▶ What if the program was updating some complicated data structures... then `exit` or `siglongjmp()` is not a good solution. We can set a flag to indicate the interrupt and continue processing. The program can deal with the interrupt later.
- ▶ Under the POSIX semantics a process can also block signals so that they are delivered later when it is ready to deal with the signals. (See man pages for `sigaction()` and `sigprocmask()` under Linux).
- ▶ What if the parent and the child are both trying to catch an interrupt? They both might then try to read from standard input. That would be confusing, to say the least. The solution is to have the parent program ignore interrupts until the child is done.

Random Rogue and Blocking Signals

Synchronization
(Part 2)

- ▶ Random rogue: shows how to catch and avoid segmentation faults: [signal-ex4.c](#)
- ▶ Shows how to block a signal and receive later at a time of our choosing: [signal-ex5.c](#)

Signal Handling in the Shell

Synchronization
(Part 2)

- ▶ How to handle `Ctrl-c` or `SIGINT`?
 - ▶ The shell ignores it on startup.
 - ▶ The shell catches it on startup and sends appropriate signal to foreground jobs when it catches the signal.
 - ▶ The shell catches it on startup, then ignores it when starting a foreground job and catches it again afterwards.
- ▶ How to handle `Ctrl-z` or `SIGTSTP`?
 - ▶ The shell ignores it on startup.
 - ▶ The shell catches it on startup and send appropriate signal to foreground jobs when it catches the signal.
 - ▶ The shell catches it on startup, then ignores it when starting a foreground job and catches it again afterwards.
- ▶ How do we prevent background jobs from getting affected by `Ctrl-c` and `Ctrl-z`?
- ▶ How to implement `fg` and `bg` built-in commands?

Pipes in Unix

Synchronization (Part 2)

- ▶ A pipe allows communication between two processes that have a common ancestor.
- ▶ A **pipe** is a half-duplex (data flows in only one direction) FIFO buffer with an API similar to file I/O.

```
#include <unistd.h>

int pipe(int filedes[2]);
//returns filedes[0] for reading, filedes[1] for writing
```

- ▶ Reading from a pipe whose write end has been closed causes an End Of File to be returned. Writing to a pipe whose read end has been closed causes the signal SIGPIPE to be generated. The write returns with `errno` set to `EPIPE`.
- ▶ The size of pipe is limited to `PIPE_BUF`. A write of `PIPE_BUF` or less will not interleave with the writes from other processes. The constant `PIPE_BUF` is defined in the file `/usr/include/linux/limits.h`
- ▶ Examples: `hello-pipe1.c`, `hello-pipe2.c`

The Power Of Pipelines

Synchronization
(Part 2)

Find the 10 most frequent words in a given text file (and their respective counts).

```
cat Shakespeare.txt | tr -cs "[A-Z][a-z][']" "[\012*]" | tr A-Z a-z |  
sort | uniq -c | sort -rn | sed 10q
```

Generate all anagrams from a given dictionary.

```
sign < $2 | sort | squash | awk '{if (NF > 1) print $0}'
```

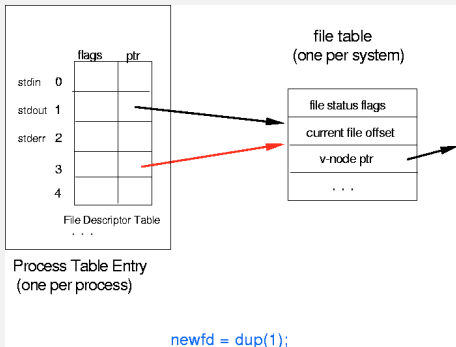
```
---squash---
#!/bin/sh
/usr/bin/awk '
$2 != prev { prev = $2; if (NR > 1) printf "\n"}
             { printf " %s ", $2 }
END         { printf "\n" }'

---sign.c---
#include <stdio.h>
#include <string.h>
#define WORDMAX 101
int compchar(char *x, char *y) { return ((*x) - (*y));}
void main(void)
{
    char thisword[WORDMAX], sign[WORDMAX];

    while (scanf("%s", thisword) != EOF) {
        strcpy(sign, thisword);
        qsort(sign, strlen(sign), 1, compchar);
        printf("%s %s\n", sign, thisword);
    }
}
```

Pipes and I/O redirection

Synchronization
(Part 2)



- ▶ See I/O redirection example 1: [dup1.c](#)
- ▶ See I/O redirection example 2: [dup2.c](#)
- ▶ We can combine I/O redirection with pipes to create pipelines of commands. See code examples
 - ▶ [pipe2.c](#)
 - ▶ [pipe3.c](#)
 - ▶ [pipe4.c](#)

Process Groups

Synchronization (Part 2)

- ▶ A **process group** is a collection of one or more processes. Each process group has a unique process group ID.
- ▶ Each process under Unix belongs to a process group. Each process group *may* have a process group leader. The leader is identified by having its process group ID equal to its process ID.
- ▶ The two related systems calls allow a process to find out its process group's ID and to change its process group.

```
#include <sys/types.h>
#include <unistd.h>

pid_t getpgrp(void);
int set setpgid(pid_t pid, pid_t pgid);
```

- ▶ A process can set the process group ID of itself or its children. Furthermore it cannot change the process group ID of its child after the child calls the `exec()` system call.
- ▶ Only processes in the foreground group can perform I/O from the terminal. A process that is not in the foreground group will receive a **SIGTTIN** or **SIGTTOU** signal when it tries to perform I/O on the terminal, which will stop the process.

See several examples in the folder **process-groups** under **synchronization-part2** lab folder.

Named Pipes (FIFOs)

Synchronization
(Part 2)

- ▶ **Named Pipes (FIFOs)** allow arbitrary processes to communicate.

```
#include <sys/types.h>
#include <sys/stat.h>

int mkfifo ( const char *pathname, mode_t mode );

mkfifo pathname
```

- ▶ If we write to a FIFO that no process has open for reading, the signal **SIGPIPE** is generated. When the last writer for a FIFO closes the FIFO, an end of file is generated for the reader of the FIFO.
- ▶ The reads/writes can be made blocking or non-blocking.
- ▶ If we have multiple writers for a FIFO, atomicity is guaranteed only for writes of size no more than **PIPE_BUF**.

Uses of FIFOs

Synchronization (Part 2)

- ▶ Can be used by shell commands to pass data from one shell pipeline to another, without creating intermediate temporary files.

```
mkfifo fifo1
prog3 < fifo1 &
prog1 < infile | tee fifo1 | prog2
```

A real example of a nonlinear pipeline:

```
wc < fifo1 &
cat /usr/share/dict/words | tee fifo1 | wc -l
```

- ▶ Look at the following simple example using one fifo: [hello-fifo.c](#)
- ▶ Look at the following example for a two-way communication using two fifos: [fifo-talk.c](#)

Client Server Communication Using FIFOs

Synchronization (Part 2)

- ▶ The server creates a FIFO using a pathname known to the clients. Clients write requests into this FIFO.
- ▶ The requests must be *atomic* and of size less than `PIPE_BUF`, which is defined in `limits.h` standard header file.
- ▶ The server replies by writing to a client-specific FIFO. For example, the client specific FIFO could be `/tmp/serv1.xxxxx` where `xxxxx` is the process id of the client.

See [wikipedia: client server model](#) for more information on the client-server model.

Fifo Server Client Example

Synchronization
(Part 2)

- ▶ Shows a simple server that listens on a fifo to any client that connects to it and chats with the client.
- ▶ Server code: `fifo-server.c`
- ▶ Client code: `fifo-client.c`

Daemons and Servers

Synchronization (Part 2)

- ▶ A **Daemon** is a process that lives for a long time. Daemons are often started when the system is bootstrapped and terminate only when the system is shutdown. They run in the background because they don't have a controlling terminal.
- ▶ A **Server** is a process that waits for a client to contact it, requesting some type of service. Typically, the server then sends some reply back to the client.
- ▶ A common use for a daemon process is as a server process.

Writing a Daemon

Synchronization
(Part 2)

Coding rules for a daemon.

1. Call `fork()` and have the parent exit.
2. Call `setsid` to create a new session. Then the process becomes a session leader of a new session, becomes the process group leader of a new process group and has no controlling terminal.
3. Change the current working directory to the root directory. Alternately, some daemons might change the working directory to some specific location.
4. Set the file creation mask to 0, so it does not use anything inherited from its parent process.
5. Unneeded file descriptors should be closed.
6. Use the `syslog` system call to log messages to the `syslogd` daemon, which is a central facility for all daemons to log messages. See man page for `syslog` (`man 3 syslog`) for more details.

Synchronization
(Part 2)

```
#include <stdlib.h>
#include <unistd.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>

int daemon_init(void)
{
    pid_t    pid;

    if ( (pid = fork()) < 0)
        return(-1);
    else if (pid != 0)
        exit(0);    /* parent goes bye-bye */

    /* child continues */
    setsid();      /* become session leader */

    chdir("/");    /* change working directory */

    umask(0);     /* clear our file mode creation mask */

    return(0);
}
```

Synchronization in the Linux kernel

Synchronization (Part 2)

- ▶ **Atomic operations** on bits, 32-bit and 64-bit integers. See [include/asm-generic/atomic.h](#) and [arch/x86/include/asm/atomic.h](#) for architecture specific implementation. See [arch/x86/include/asm/bitops/atomic.h](#) for test-and-set related bit operations.
- ▶ **Spinlocks**. Threads do not sleep if waiting for a lock. Suitable only for short duration locks.
- ▶ **Reader-Writer spinlocks**. Gives preference to readers over writers. Multiple readers can hold the lock but only one writer..
- ▶ **Semaphores**. See [include/linux/semaphore.h](#) and [kernel/locking/semaphore.c](#). Uses wait queues and sleep.
- ▶ **Reader-Writer Semaphores**.
- ▶ **Mutexes**. Similar to a binary semaphore but with a simpler interface, more efficient performance, and additional constraints on its use.
- ▶ **Completion variables**. Similar to condition variables.
- ▶ **Sequential locks**. Reader and writers with preference given to writers.
- ▶ **Read-Write barriers and ordering**.
- ▶ **BKL: Big Kernel Lock**. Removed in kernel version 2.6.39.
- ▶ **Preemption disabling**. Bad thing :-)