

## Overview

**System Programs:** associated with the operating system but are not necessarily part of the kernel.  
**Middleware:** software frameworks that provide additional services to application developers.

## Interrupts

Signal to CPU there is a task requiring attention, usually I/O. Trap is s/w generated interrupt caused by error or user request.

**Synchronous:** user process waits for I/O to finish, must know time it takes. **Asynchronous:** request returns before I/O completion, interrupts when done.

**Polling system:** system calls all interrupt handlers to determine which made the interrupt. (slow) **vectorized system:** handler is found via interrupt vector.

Interrupt is **maskable** if it can be ignored by CPU otherwise **nonmaskable**.

## I/O Techniques

**Programmed:** Requested I/O action performed, module sets I/O status register. Processor periodically checks status until action completed.

**Interrupt-driven:** Process issues I/O command to module which interrupts processor when done. Processor then transfers data.

**DMA:** Data is transferred directly to memory, processor only needed for initialization. Bus access slowed for processor during the transfer.

## Processes

Possible states of a process: new, running, waiting, ready, terminated.

## Process Creation Code

```
int main()
{
    pid_t pid;
    /* Fork a child process */
    pid = fork();
    if (pid < 0) { /* error occurred */
        printf(stderr, "Fork Failed!");
        return;
    }
    else if (pid == 0) { /* child process */
        /* parent will wait for the child
           to complete */
        wait(NULL);
        printf("Child Complete");
    }
    return 0;
}
```

`Fork()` returns 0 for the child process, and returns the pid of the child process for the parent. The `exit(1)` system call can be used to terminate a process with status 1. The parent can obtain the status of a terminated child with `pid = wait(&status)` where `status` is the integer status of the child process. A process that has terminated, but whose parent has not yet called `wait()`, is a **zombie** process. If the parent of a process terminated without invoking `wait()`, the process is an **orphan**.

## Posix Shared Memory

```
/* create the shared memory object */
fd = shm_open(name, O_CREAT | O_RDWR, 0666);
/* configure the size of the shared
   memory object */
ftruncate(fd, SIZE);
/* memory map the shared memory object */
ptr = (char *)
mmap(0, SIZE, PROT_READ | PROT_WRITE,
     MAP_SHARED, fd, 0);
/* write to the shared memory object */
sprintf(ptr, "$s", message_0);
ptr += strlen(message_0);
sprintf(ptr, "$s", message_1);
ptr += strlen(message_1);
/* read from the shared memory object */
printf("$s", (char *)ptr);
/* remove the shared memory object */
shm_unlink(name);
/* wait for the thread to exit */
```

## Pipes (message passing)

```
pthread_join(tid,NULL);
pthread_exit(0);

Synchronization
Critical section problem
A solution to the critical section problem must satisfy mutual exclusion, progress (processes are eventually allowed to access their critical sections after requesting it), and bounded waiting (there is a limit on the number of times other processes can access their critical sections after one process has requested access) PETERSON's Solution: processes use a flag variable to indicate if a process is ready to enter its critical section and a shared turn variable to indicate which process' turn it is to enter their critical section.

Semaphores
wait() = P(), signal() = V(), wait() decrements the semaphore, signal() increments the semaphore.

Atomic Instructions
boolean test_and_set(boolean *target) {
    boolean rv = *target;
    *target = true;
    return rv;
}
do {
    while (test_and_set(&lock))
        /* do nothing */
    /* critical section */
    lock = false;
    /* remainder section */
} while (true);
int compare_and_swap(int *value, int
    expected, int new_value) {
    int temp = *value;
    if (value == expected)
        *value = new_value;
    return temp;
}
while (true) {
    while (compare_and_swap(&lock, 0, 1)
        != 0)
        /* do nothing */
    /* critical section */
    lock = 0;
    /* remainder section */
}
```

## Threads

### Amdahl's Law

$$speedup \leq \frac{1}{S + (1-S)/N}$$
 where  $S$  is the serial portion of the application and  $N$  is the number of processing cores

**Threads of a process share** the code section, data section, and files of the process. The registers, stack, and program counter of the threads differ.

**Data Parallelism:** distribute subsets of the same data across multiple computing cores.

**Task Parallelism:** distribute tasks (threads) across multiple computing cores.

**Five areas of challenge in multicore programming:** identifying tasks, balancing the amount of work tasks do, data splitting, data dependency, and testing and debugging.

**Synchronous Signals:** Delivered to the same process that caused the signal.

**Asynchronous Signals:** Delivered to a process that did not cause the signal.

**Thread Local Storage (TLS):** Data local to a specific thread. Similar to the concept of static data and local variables.

**Lightweight Processes (LWP):** Intermediate data structure between user and kernel threads, often used in systems implementing many-to-many or two-level thread models. Looks like a virtual process to the user-thread library. The application can schedule a user thread to run on an LWP. Each LWP is attached to a kernel thread, and it is kernel threads that the operating system schedules to run on physical processors. If a kernel thread blocks, the LWP blocks and also blocks the user thread down the chain.

**Benefits of Threads:** Responsiveness, resource sharing (threads share data by default), economy (threads are take less time and memory to create threads than processes), scalability.

## Pthread

```
pthread_t tid; /* thread identifier */
pthread_attr_t attr; /* set of thread
   attributes */
/* set default attributes of thread */
pthread_attr_init(&attr);
/* create the thread */
pthread_create(&tid, &attr, runner, arg
    + [1]);
/* cancel the thread */
pthread_cancel(tid);
/* check if there is a cancellation
   request */
pthread_testcancel();
/* wait for the thread to exit */
```

## 3SH3 Final – Linus Torvalds Edition

### Reader-writer locks

If only reader-writer locks are used in a multithreaded application, then deadlock is still possible. Locks cannot be shared if there is a writer (mutual exclusion).

A thread can hold one reader-writer lock while waiting to acquire another (hold and wait). You cannot take a lock away, so no pre-emption is upheld. A circular wait among all threads is possible.

**Banker's Algorithm**  
 Need = Max - Allocation. Safe state = all threads can finish executing at the boundary of an instruction cycle. Fine-grained systems include logic for thread switching, so the cost of switching between threads is small.

**Load Balancing**

**Pull migration:** when an idle processor

switch to another thread to begin executing the latter is free or allocated and, if it is free, to which page of which process (or processes). Used by the OS to keep track of physical memory.

**Reentrant code:** it never changes during execution and can be shared among processes. Can be implemented as shared pages.

### Number of Bits and/or Entries

**Important:** if any size is not in bytes convert to bytes first.

If the address takes  $x$  bits to store, page size is  $2^y$  bytes, and number of frames or pages is  $2^z$ , then  $x = y + z$ .

**Conventional single level** page table stores an entry for each virtual page or page number. **Inverted page table** stores an entry for each physical frame.

### Processor Affinity

A process has an "affinity" for the processor it is currently running on because the processor will usually have values it uses often in the cache of the current processor.

**Soft Affinity:** OS will attempt to keep a process on a single processor, but it is possible for a process to migrate between processors during load balancing.

**Hard Affinity:** OS allows a process to specify a subset of processors on which it can run using sys calls.

**Resource Allocation Graphs**

A directed edge from thread  $T_i$  to resource type  $R_j$  is denoted by  $T_i \rightarrow R_j$  and signifies that thread  $T_i$  has requested an instance of resource type  $R_j$ .

A directed edge from resource type  $R_j$  to thread  $T_i$  is denoted by  $R_j \rightarrow T_i$  signifies that an instance of resource type  $R_j$  has been allocated to thread  $T_i$ .

**Methods for Handling Deadlock**

### Real time Scheduling

The scheduler for a real-time operating system must support a priority based algorithm with preemption

**Deadlock Prevention:** prevent circular wait from happening (only practical solution)

**Deadlock Avoidance:** don't grant resources if it leads the system to an unsafe state.

**Deadlock Detection and Recovery:** Elbow terminate or preempt resources.

**Scheduling**

**Predicting next CPU burst Formula**

$$\tau_{n+1} = \alpha \cdot t_n + (1-\alpha) \cdot \tau_n$$
 where  $t_n$  is the value of the nth CPU burst,  $0 \leq \alpha \leq 1$

$\tau_0$  affects the starting value of the predictions, and  $\alpha$  affects how much the last CPU burst vs the last predicted CPU burst is weighted.

If  $\alpha = 0$ , then  $\tau_{n+1} = \tau_n$  and recent history has no effect on the future CPU burst. If  $\alpha = 1$ , then  $\tau_{n+1} = t_n$  and only the most recent CPU bursts matter (history is assumed to be old and irrelevant).  $\alpha = 1/2$  weights recent and past history evenly.

**Pthread Synchronization**

**Mutex Locks**

`pthread_mutex_t mutex;`

`/* create and initialize mutex lock */`

`pthread_mutex_init(&mutex, NULL);`

`/* acquire the mutex lock */`

`pthread_mutex_lock(&mutex);`

`/* critical section */`

`/* release the mutex lock */`

`pthread_mutex_unlock(&mutex);`

**Condition Variables**

`pthread_mutex_t mutex;`

`pthread_cond_t cond_var;`

`pthread_mutex_init(&mutex, NULL);`

`pthread_cond_init(&cond_var, NULL);`

`/* thread waiting on condition variable`

`pthread_cond_wait(&cond_var, &mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

`pthread_mutex_unlock(&mutex);`

`/* thread that modifies shared data,`

`/* signalling waiting thread`

`pthread_cond_signal(&cond_var);`

## C-SCAN Scheduling

C-SCAN moves the head from one end of the disk to the other, servicing requests along the way. When the head reaches the other end, however, it immediately returns to the beginning of the disk without servicing any requests on the return trip. Is a variant of SCAN scheduling.

**Write amplification:** The creation of I/O requests not by applications but by the NVM device doing garbage collection and space management. Can greatly impact the write performance of the device.

## latency & I/O time

Latency based on spindle speed:  $1/(RPM/60) = 60/RPM$

Average latency = 1/2 latency

Access Latency = average seek time + average latency

Average I/O time = access latency + (amount to transfer / transfer rate) + controller overhead

## Conversions

1 s = 1000 ms, 1 ms = 1000  $\mu$ s (microsecond), 1  $\mu$ s = 1000 ns

1 kB =  $2^{10}$  B = 1024 B,

1 MB =  $2^{20}$  B = 1,048,576 B,

1 GB =  $2^{30}$  B = 1,073,741,824 B,

1 TB =  $2^{40}$  B = 1,099,511,627,776 B