



Multicore and Multiprocessor Systems: Part II



Common to all the following commands:

Compiling and linking needs to be performed with
`-pthread`.

The `pthread` functions and related data types are made available in a C program using:

```
#include <pthread.h>
```



```
int pthread_create(pthread_t *thread,  
                  const pthread_attr_t *attr,  
                  void *(*start_routine) (void *),  
                  void *arg);
```

- `thread` unique identifier to distinguish from other threads,
- `attr` attributes for determining thread properties. `NULL` means default properties,
- `start_routine` pointer to the function to be started in the newly created thread,
- `arg` the argument of the above function.



```
int pthread_create(pthread_t *thread,  
                  const pthread_attr_t *attr,  
                  void *(*start_routine) (void *),  
                  void *arg);
```

- `thread` unique identifier to distinguish from other threads,
- `attr` attributes for determining thread properties. `NULL` means default properties,
- `start_routine` pointer to the function to be started in the newly created thread,
- `arg` the argument of the above function.

Note that only a single argument can be passed to the threads start function.



The argument of the start function is a `void` pointer. We can thus define:

```
struct point3d{ double x,y,z; };
struct norm_args{
    struct point3d *P;
    double norm;
};
struct norm_args args;
```

and upon thread creation pass

```
err=pthread_create(tid, NULL, norm, (void *) &args);
```

to a start function

```
void *norm(void *arg) {
    struct norm_args *args=(struct norm_args *)arg;
    struct point3d *P;
    P = args->P;
    args->norm = P->x * P->x + P->y * P->y + P->z * P->z;
    return NULL;
};
```



```
int main(int argc, char* argv[]){
    pthread_t tid1,tid2;

    struct point3d point;
    struct norm_args args;

    args.P = &point;

    point.x=10; point.y=10; point.z=0;
    pthread_create(&tid1, NULL, norm, &args);

    point.x=20; point.y=20; point.z=-50;
    pthread_create(&tid2, NULL, norm, &args);

    pthread_join(tid1, NULL);
    pthread_join(tid2, NULL);
}
```

Depending on the execution of thread `tid1` the argument `point` may get overwritten before it has been fetched, the analogue holds for the `norm` argument inside the function.



CSC

POSIX Threads

Exiting threads and waiting for their termination

Pthreads can exit in different forms:



POSIX Threads

Exiting threads and waiting for their termination

Pthreads can exit in different forms:

- they return from their start function,



POSIX Threads

Exiting threads and waiting for their termination

Pthreads can exit in different forms:

- they return from their start function,
- they call `pthread_exit()` to cleanly exit,



Pthreads can exit in different forms:

- they return from their start function,
- they call `pthread_exit()` to cleanly exit,
- they are aborted by a call to `pthread_cancel()`,



Pthreads can exit in different forms:

- they return from their start function,
- they call `pthread_exit()` to cleanly exit,
- they are aborted by a call to `pthread_cancel()`,
- the process they are associated to is terminated by an `exit()` call.



```
int pthread_exit (void *retval);
```

- `retval` return value of the exiting thread to the calling thread,
- threads exit implicitly when their start function is exited,
- the return value may be evaluated from another thread of the same process via the `pthread_join()` function,
- after the last thread in a process exits the process terminates calling `exit()` with a zero return value. Only then shared resources are released automatically.



```
int pthread_join(pthread_t thread, void **retval);
```

- Waits for a thread to terminate and fetches its return value.
- `thread` the identifier of the thread to wait for,
- `retval` destination to copy the return value (if not `NULL`) to.



The Pthread standard supports four types of synchronization and coordination facilities:

1. `pthread_join()`; we have seen this function above



The Pthread standard supports four types of synchronization and coordination facilities:

1. `pthread_join()`; we have seen this function above
2. **Mutex variable functions** for handling mutexes as defined above



The Pthread standard supports four types of synchronization and coordination facilities:

1. `pthread_join()`; we have seen this function above
2. Mutex variable functions for handling mutexes as defined above
3. **Condition variable functions** treat a condition variable that can be used to indicate a certain event in which the threads are interested. Condition variables may be used to implement semaphore like structures and triggers for special more complex situation that require the threads to act in a certain way.



The Pthread standard supports four types of synchronization and coordination facilities:

1. `pthread_join()`; we have seen this function above
2. Mutex variable functions for handling mutexes as defined above
3. Condition variable functions treat a condition variable that can be used to indicate a certain event in which the threads are interested. Condition variables may be used to implement semaphore like structures and triggers for special more complex situation that require the threads to act in a certain way.
4. `pthread_once()` can be used to make sure that certain initializations are performed by one and only one thread when called by multiple ones.



Dynamic initialization:

```
int pthread_mutex_init(pthread_mutex_t *restrict mutex,  
                       const pthread_mutexattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
```



Dynamic initialization:

```
int pthread_mutex_init(pthread_mutex_t *restrict mutex,  
                      const pthread_mutexattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
```

- `mutex` is the mutex variable to be initialized



Dynamic initialization:

```
int pthread_mutex_init(pthread_mutex_t *restrict mutex,  
                       const pthread_mutexattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
```

- `mutex` is the mutex variable to be initialized
- `attr` can be used to adapt the mutex properties, as for the pthreads `NULL` gives the default attributes,

Dynamic initialization:

```
int pthread_mutex_init(pthread_mutex_t *restrict mutex,  
                      const pthread_mutexattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
```

- `mutex` is the mutex variable to be initialized
- `attr` can be used to adapt the mutex properties, as for the pthreads `NULL` gives the default attributes,
- `restrict`⁴ is a C99-standard keyword limiting the pointer aliasing features and guiding compilers and aiding in the caching optimization.

⁴See also <https://en.wikipedia.org/wiki/Restrict>

Dynamic initialization:

```
int pthread_mutex_init(pthread_mutex_t *restrict mutex,  
                       const pthread_mutexattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
```

- `mutex` is the mutex variable to be initialized
- `attr` can be used to adapt the mutex properties, as for the pthreads `NULL` gives the default attributes,
- `restrict`⁴ is a C99-standard keyword limiting the pointer aliasing features and guiding compilers and aiding in the caching optimization.
- initialization may fail if the system has insufficient memory (error code `ENOMEM`) or other resources (`EAGAIN`)

⁴See also <https://en.wikipedia.org/wiki/Restrict>



```
int pthread_mutex_lock(pthread_mutex_t *mutex);
```

- If `mutex` is unlocked the function returns with the mutex in locked state,



```
int pthread_mutex_lock(pthread_mutex_t *mutex);
```

- If `mutex` is unlocked the function returns with the mutex in locked state,
- If `mutex` is already locked the execution is blocked until the lock is released and it can proceed as above,



```
int pthread_mutex_lock(pthread_mutex_t *mutex);
```

- If `mutex` is unlocked the function returns with the mutex in locked state,
- If `mutex` is already locked the execution is blocked until the lock is released and it can proceed as above,
- Four types of mutexes are defined:
 - `PTHREAD_MUTEX_NORMAL`
 - `PTHREAD_MUTEX_ERRORCHECK`
 - `PTHREAD_MUTEX_RECURSIVE`
 - `PTHREAD_MUTEX_DEFAULT`

All of them show different behavior when locked mutexes should again be locked by the same thread or a thread tries to unlock a previously unlocked mutex and similar unintended situations. This especially regards [error handling](#) and [deadlock detection](#).



```
int pthread_mutex_trylock(pthread_mutex_t *mutex);
```

- The function is equivalent to `pthread_mutex_lock()`, except that it returns immediately in any case.



```
int pthread_mutex_trylock(pthread_mutex_t *mutex);
```

- The function is equivalent to `pthread_mutex_lock()`, except that it returns immediately in any case.
- Success or failure are determined from the return value.



```
int pthread_mutex_trylock(pthread_mutex_t *mutex);
```

- The function is equivalent to `pthread_mutex_lock()`, except that it returns immediately in any case.
- Success or failure are determined from the return value.
- If the mutex type is `PTHREAD_MUTEX_RECURSIVE` the lock count is increased by one and the function returns success.



```
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

- the function releases the lock



```
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

- the function releases the lock
- what exactly “release” means, depends on the properties of the mutex variable

```
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

- the function releases the lock
- what exactly “release” means, depends on the properties of the mutex variable
- e.g., for type `PTHREAD_MUTEX_RECURSIVE` mutexes it means that the counter is decreased by one and they become available once it reaches zero



```
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

- the function releases the lock
- what exactly “release” means, depends on the properties of the mutex variable
- e.g., for type `PTHREAD_MUTEX_RECURSIVE` mutexes it means that the counter is decreased by one and they become available once it reaches zero
- if the mutex becomes available, i.e., unlocked by the function call and there are blocked threads waiting for it, the threading policy decides which thread acquires `mutex` next.



```
int pthread_mutex_destroy(pthread_mutex_t *mutex);
```

- destroys the mutex referenced by `mutex`



```
int pthread_mutex_destroy(pthread_mutex_t *mutex);
```

- destroys the mutex referenced by `mutex`
- the destroyed mutex then becomes uninitialized



```
int pthread_mutex_destroy(pthread_mutex_t *mutex);
```

- destroys the mutex referenced by `mutex`
- the destroyed mutex then becomes uninitialized
- `pthread_mutex_init()` can be used to initialize the same mutex variable again

```
int pthread_mutex_destroy(pthread_mutex_t *mutex);
```

- destroys the mutex referenced by `mutex`
- the destroyed mutex then becomes uninitialized
- `pthread_mutex_init()` can be used to initialize the same mutex variable again
- if mutex is locked or referenced, `pthread_mutex_destroy()` fails with error code `EBUSY`



Example (A deadlock situation when locking multiple mutexes)

Problem:

- Consider two mutex variables m_a and m_b , as well as two threads T1 and T2.



Example (A deadlock situation when locking multiple mutexes)

Problem:

- Consider two mutex variables m_a and m_b , as well as two threads T1 and T2.
- T1 locks m_a first and then m_b ,



Example (A deadlock situation when locking multiple mutexes)

Problem:

- Consider two mutex variables m_a and m_b , as well as two threads T1 and T2.
- T1 locks m_a first and then m_b ,
- T2 locks m_b first and then m_a ,



Example (A deadlock situation when locking multiple mutexes)

Problem:

- Consider two mutex variables m_a and m_b , as well as two threads T1 and T2.
- T1 locks m_a first and then m_b ,
- T2 locks m_b first and then m_a ,
- In case T1 is interrupted by the scheduler after locking m_a , but before locking m_b and in the meantime T2 succeeds in locking it, then the classical deadlock occurs.



Example (A deadlock situation when locking multiple mutexes)

Locking hierarchy solution:

The basic idea here is that all threads need to lock the critical mutexes in the same order. This can easily be guaranteed by hierarchically ordering the mutexes.



Example (A deadlock situation when locking multiple mutexes)

Back off strategy solution:

When we want to keep the differing locking orders, we may use `pthread_mutex_trylock()` with a back off strategy.



Example (A deadlock situation when locking multiple mutexes)

Back off strategy solution:

When we want to keep the differing locking orders, we may use `pthread_mutex_trylock()` with a back off strategy.

- Locking is tried in the desired order,



Example (A deadlock situation when locking multiple mutexes)

Back off strategy solution:

When we want to keep the differing locking orders, we may use `pthread_mutex_trylock()` with a back off strategy.

- Locking is tried in the desired order,
- when a trylock fails, the thread unlocks all previously locked mutexes (it backs off of the protected resources),



Example (A deadlock situation when locking multiple mutexes)

Back off strategy solution:

When we want to keep the differing locking orders, we may use `pthread_mutex_trylock()` with a back off strategy.

- Locking is tried in the desired order,
- when a trylock fails, the thread unlocks all previously locked mutexes (it backs off of the protected resources),
- after the back off it starts over from the first one.

Dynamic initialization:

```
int pthread_cond_init(pthread_cond_t *restrict cond,  
                     const pthread_condattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
```

- cond the condition to be initialized

Dynamic initialization:

```
int pthread_cond_init(pthread_cond_t *restrict cond,  
                      const pthread_condattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
```

- `cond` the condition to be initialized
- `attr` can be used to adapt the condition properties, as for the pthreads `NULL` gives the default attributes,

Dynamic initialization:

```
int pthread_cond_init(pthread_cond_t *restrict cond,  
                     const pthread_condattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
```

- `cond` the condition to be initialized
- `attr` can be used to adapt the condition properties, as for the pthreads `NULL` gives the default attributes,
- `restrict`: see `pthread_mutex_init()`



Dynamic initialization:

```
int pthread_cond_init(pthread_cond_t *restrict cond,  
                     const pthread_condattr_t *restrict attr);
```

Static/Macro initialization:

```
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
```

- `cond` the condition to be initialized
- `attr` can be used to adapt the condition properties, as for the pthreads `NULL` gives the default attributes,
- `restrict`: see `pthread_mutex_init()`
- every condition variable is associated to a mutex.



```
int pthread_cond_destroy(pthread_cond_t *cond);
```

- destroys the condition variable referenced by `cond`



```
int pthread_cond_destroy(pthread_cond_t *cond);
```

- destroys the condition variable referenced by `cond`
- the destroyed condition then becomes uninitialized



```
int pthread_cond_destroy(pthread_cond_t *cond);
```

- destroys the condition variable referenced by `cond`
- the destroyed condition then becomes uninitialized
- `pthread_cond_init()` can reinitialize the same condition variable



```
int pthread_cond_destroy(pthread_cond_t *cond);
```

- destroys the condition variable referenced by `cond`
- the destroyed condition then becomes uninitialized
- `pthread_cond_init()` can reinitialize the same condition variable
- if `cond` is blocking threads when destroyed the standard does not specify the behavior of `pthread_cond_destroy()`.



```
int pthread_cond_wait(pthread_cond_t *restrict cond,  
                      pthread_mutex_t *restrict mutex);
```

- assumes that `mutex` was locked before by the calling thread,



```
int pthread_cond_wait(pthread_cond_t *restrict cond,  
                     pthread_mutex_t *restrict mutex);
```

- assumes that `mutex` was locked before by the calling thread,
- results in the thread getting blocked and at the same time (atomically) releasing `mutex`



```
int pthread_cond_wait(pthread_cond_t *restrict cond,  
                      pthread_mutex_t *restrict mutex);
```

- assumes that `mutex` was locked before by the calling thread,
- results in the thread getting blocked and at the same time (atomically) releasing `mutex`
- another thread may evaluate this to wake up the now blocked thread (see `pthread_cond_signal()`)



```
int pthread_cond_wait(pthread_cond_t *restrict cond,  
                     pthread_mutex_t *restrict mutex);
```

- assumes that `mutex` was locked before by the calling thread,
- results in the thread getting blocked and at the same time (atomically) releasing `mutex`
- another thread may evaluate this to wake up the now blocked thread (see `pthread_cond_signal()`)
- upon waking up the thread automatically tries to gain access to `mutex` again,



```
int pthread_cond_wait(pthread_cond_t *restrict cond,  
                     pthread_mutex_t *restrict mutex);
```

- assumes that `mutex` was locked before by the calling thread,
- results in the thread getting blocked and at the same time (atomically) releasing `mutex`
- another thread may evaluate this to wake up the now blocked thread (see `pthread_cond_signal()`)
- upon waking up the thread automatically tries to gain access to `mutex` again,
- if it succeeds it should test the condition again to check whether another thread changed it in the meantime.



```
int pthread_cond_signal(pthread_cond_t *cond);
```

- if no thread is blocked on the condition variable `cond` there is no effect,
- otherwise, **one** of the waiting threads is woken up and proceeds as described above.

```
int pthread_cond_broadcast(pthread_cond_t *cond);
```

- wakes up **all** threads blocking on `cond`,
- all of them try to acquire the associated mutex,
- only **one** of them can succeed,
- the others get blocked on the mutex now.



```
int pthread_cond_timedwait(pthread_cond_t *restrict cond,  
                           pthread_mutex_t *restrict mutex,  
                           const struct timespec *restrict abstime);
```

- equivalent to `pthread_cond_wait()` except that it only blocks for the period specified by `abstime`,



```
int pthread_cond_timedwait(pthread_cond_t *restrict cond,  
                           pthread_mutex_t *restrict mutex,  
                           const struct timespec *restrict abstime);
```

- equivalent to `pthread_cond_wait()` except that it only blocks for the period specified by `abstime`,
- if the thread did not get signaled or broadcast before `abstime` expires it returns with error code `ETIMEDOUT`.



Pthread coordination mechanisms

A counting semaphore for Pthreads

Semaphores are not available in the POSIX Threads standard.



Semaphores are not available in the POSIX Threads standard.

However, they can be created using the existing mechanisms of mutexes and conditions.



Semaphores are not available in the POSIX Threads standard.

However, they can be created using the existing mechanisms of mutexes and conditions.

A counting semaphore should be a data type that acts like a counter with non-negative values and for which two operations are defined:

1. A **signal operation** increments the counter and wakes up a task blocked on the semaphore if one exists.
2. A **wait operation** simply decrements the counter if it is positive. If it was zero already the thread is blocking on the semaphore.



- data structure for the semaphore:

```
typedef struct _sema_t{
    int count;
    pthread_mutex_t m;
    pthread_cond_t c;
} sema_t;
```

- the initialization

```
void InitSema(sema_t *ps) {
    pthread_mutex_init(&ps->m, NULL);
    pthread_cond_init(&ps->c, NULL);
}
```

- and the cleanup

```
void CleanupSema(void *arg) {
    pthread_mutex_unlock((pthread_mutex_t *) arg);
}
```

source: [RAUBER/RÜNGER'10]



```
void ReleaseSema(sema_t *ps){ // signal operation
    pthread_mutex_lock(&ps->m) ;
    pthread_cleanup_push(CleanupSema, &ps->m) ;
    {
        ps->count++;
        pthread_cond_signal(&ps->c) ;
    }
    pthread_cleanup_pop ( 1 ) ;
}

void AcquireSema(sema_t *ps){ // wait operation
    pthread_mutex_lock (&ps->mutex) ;
    pthread_cleanup_push(CleanupSema, &ps->m) ;
    {
        while (ps->count==0)
            pthread_cond_wait (&ps->c, &ps->m) ;
        ps->count--;
    }
    pthread_cleanup_pop(1) ;
}
```

source: [RAUBER/RÜNGER'10]



Example (Producer/Consumer queue buffer protection)

- A buffer of fixed size n is shared by
- a producer thread generating entries and storing them in the buffer if it is not full,
- a consumer thread removing entries from the same buffer for further processing unless it is empty.

For the realization of the protected access two semaphores are required:

1. Number of entries occupied (initialized by 0),
2. Number of free entries (initialized by n).

The Mechanism works for an arbitrary number of producers and consumers.
(Details will be worked out on exercise sheet 2.)



1. Master/Slave model:

- A master thread is controlling the execution of the program,
- the slave threads are executing the work.



1. Master/Slave model:

- A master thread is controlling the execution of the program,
- the slave threads are executing the work.

2. Client/Server model:

- Client threads produce requests,
- Server threads execute the corresponding work.



1. Master/Slave model:

- A master thread is controlling the execution of the program,
- the slave threads are executing the work.

2. Client/Server model:

- Client threads produce requests,
- Server threads execute the corresponding work.

3. Pipeline model:

- Every thread (except for the first and last in line) produces output that serves as input for another thread,
- after a startup phase (filling the pipeline) the parallel execution is achieved.



1. Master/Slave model:

- A master thread is controlling the execution of the program,
- the slave threads are executing the work.

2. Client/Server model:

- Client threads produce requests,
- Server threads execute the corresponding work.

3. Pipeline model:

- Every thread (except for the first and last in line) produces output that serves as input for another thread,
- after a startup phase (filling the pipeline) the parallel execution is achieved.

4. Worker model:

- equally privileged workers organize their workload,
- an important variant is the task pool treated as detailed example next.



Task Pools

Basic idea of the task pool

Idea:

Creation of a parallel threaded program that can dynamically schedule tasks on the available processors.



Idea:

Creation of a parallel threaded program that can dynamically schedule tasks on the available processors.

Key ingredients in the approach are:

- usage of a fixed number of threads
- organization of the pending tasks in a task pool,
- threads fetch the tasks from the pool and execute them leading to a dynamic assignment of the work load.



Idea:

Creation of a parallel threaded program that can dynamically schedule tasks on the available processors.

Key ingredients in the approach are:

- usage of a fixed number of threads
- organization of the pending tasks in a task pool,
- threads fetch the tasks from the pool and execute them leading to a dynamic assignment of the work load.

Main advantages

- automatic dynamic load balancing among the threads
- comparably small overhead for the administration of threads



- data structure for one task:

```
typedef struct _work_t{
    void (*routine) (void*); //worker function to call
    void* arg ;
    struct _work_t *next;
} work_t ;
```

- data structure for the task pool:

```
typedef struct _tpool_t{
    int num_threads ; // number of threads
    int max_size, curr_size; // max./cur. number of tasks in pool
    pthread_t *threads; //array of threads
    work_t *head , *tail; // start/end of the task queue
    pthread_mutex_t lock; //access control for the task pool
    pthread_cond_t not_empty ; // tasks are available
    pthread_cond_t not_full ; // tasks may be added
} tpool_t ;
```

source: [RAUBER/RÜNGER'10]



```
tpool_t *tpool_init(int num_threads , int max_size){
    int i;
    tpool_t *tpl;

    tpl=(tpool_t *) malloc (sizeof(tpool_t));
    tpl->num_threads=num_threads ;
    tpl->max_size=max_size ;
    tpl->cur_size=0;
    tpl->head=tpl->tail=NULL;

    pthread_mutex_init (&tpl->lock, NULL);
    pthread_cond_init (&tpl->not_empty, NULL);
    pthread_cond_init (&tpl->not_full, NULL);
    tpl->threads=(pthread_t *) malloc(num_threads *sizeof(pthread_t));
    for(i=0; i<num_threads; i++)
        pthread_create(tpl->threads+i, NULL, tpool_thread, (void *)tpl) ;
    return tpl;
}
```

source: [RAUBER/RÜNGER'10]



```
void *tpool_thread(void *vtpl) {
    tpool_t *tpl=(tpool_t *) vtpl;
    work_t *wl ;

    for ( ; ; ) {
        pthread_mutex_lock(&tpl->lock);
        while(tpl->cur_size==0)
            pthread_cond_wait(&tpl->not_empty , &tpl->lock);
        wl=tpl->head; tpl->cur_size--;
        if(tpl->cur_size==0)
            tpl->head=tpl->tail=NULL;
        else tpl->head = wl->next;
        if (tpl->cur_size==tpl->max_size-1) // pool full
            pthread_cond_signal(&tpl->not_full);
        pthread_mutex_unlock(&tpl->lock);
        (*wl->routine)(wl->arg);
        free(wl);
    }
}
```

source: [RAUBER/RÜNGER'10]



```
void tpool_insert(tpool_t *tpl, void(*f) (void*), void *arg) {
    work_t *wl ;

    pthread_mutex_lock (&tpl->lock);
    while (tpl->cur_size==tpl->max_size)
        pthread_cond_wait (&tpl->not_full, &tpl->lock);
    wl=(work_t *) malloc (sizeof(work_t));
    wl->routine=f; wl->arg=arg; wl->next=NULL ;
    if ( tpl->cur_size==0) {
        tpl->head=tpl->tail=wl;
        pthread_cond_signal (&tpl->not_empty);
    }
    else{
        tpl->tail->next=wl; tpl->tail=wl;
    }
    tpl->cur_size++;
    pthread_mutex_unlock (&tpl->lock);
}
```

source: [RAUBER/RÜNGER'10]



Shared Memory Blocks

General shared memory blocks

In contrast to Threads, different processes do not share their address space. Therefore, different ways to communicate in multiprocessing applications are necessary.

One possible way are **shared memory objects**. Unix-like operating systems provide at least one of:

- **old:** System V Release 4 (SVR4) Shared Memory⁵
- **new:** POSIX Shared Memory⁶.

Both techniques implement shared memory objects, like common memory, semaphores and message queues, which are accessible from different applications with different address spaces.

⁵*System V Interface Definition, AT&T Unix System Laboratories, 1991*

⁶*IEEE Std 1003.1-2001 Portable Operating System Interface System Interfaces*



Common Memory Locations

- They are used to share data between applications.
- They are managed by the kernel and not by the application.
- Each location is represented as a file in `/dev/shm/`.
- They are handled like normal files.
- They are created using `shm_open` and mapped to the memory using `mmap`.
- Exist as long as no application deletes them.
- Even when the creating program exits they stay available,
- See manpage: `man 7 shm_overview`.



POSIX Semaphores

- Counting semaphores are available from different address spaces.
- They correspond to `pthread_mutex_*` in threaded applications.
- They are represented as a file in `/dev/shm/sem.*`.
- See manpage: `man 7 sem_overview`.

Message Queues

- They represent a generalized Signal concept which can transfer a small payload (2 to 4 KiB).
- They correspond to `pthread_cond_*` in threaded applications.
- They can be represented as file in `/dev/mqueue`.
- See manpage: `man 7 mq_overview`.