

POSIX Threads: a first step toward parallel programming

George Bosilca
bosilca@icl.utk.edu

Shared/Exclusive Locks


- **ReadWrite Mutual exclusion**
- Extension used by the reader/writer model
- 4 states: write_lock, write_unlock, read_lock and read_unlock.
- multiple threads may hold a shared lock simultaneously, but only one thread may hold an exclusive lock.
- if one thread holds an exclusive lock, no threads may hold a shared lock.

Shared/Exclusive Locks

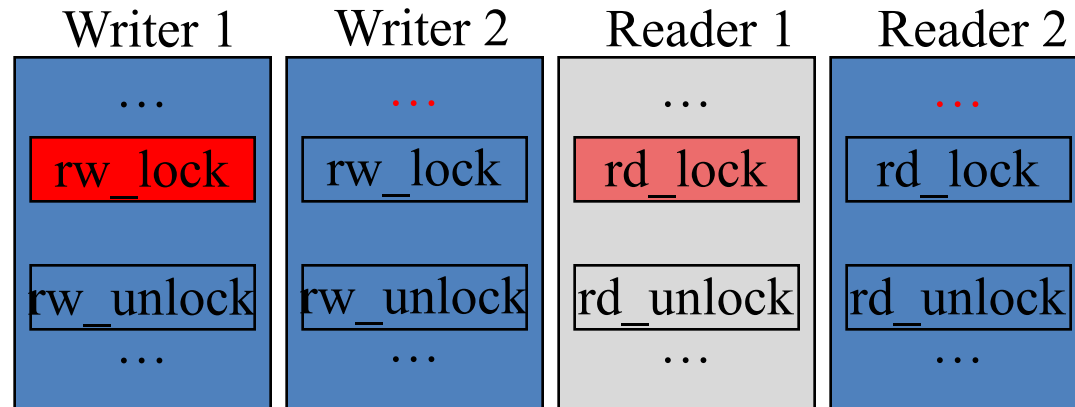
Legend

Active thread

Sleeping thread



Step 1



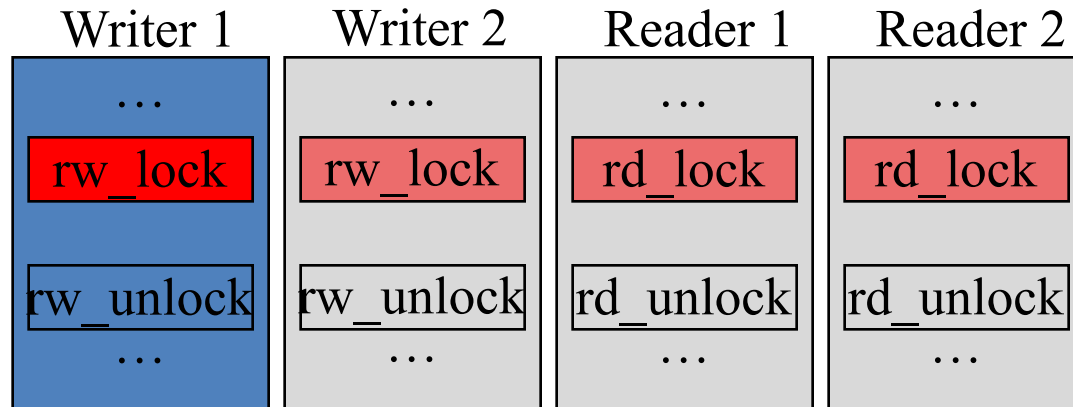
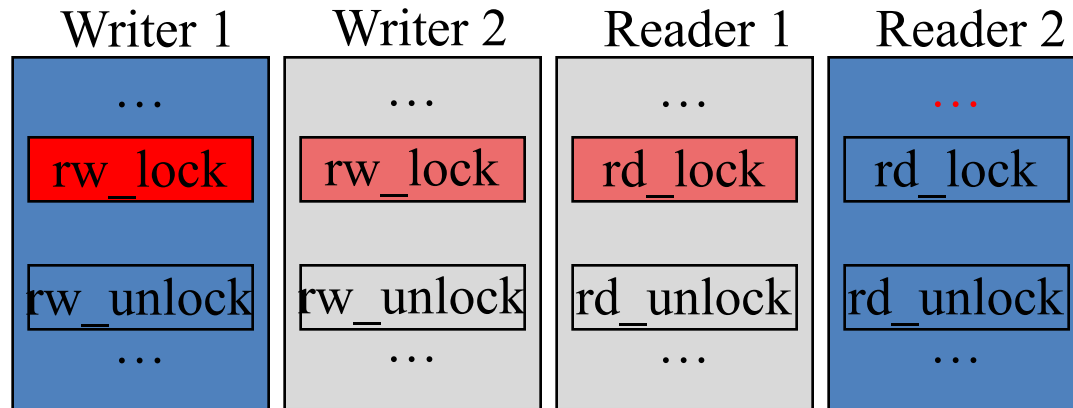
Step 2

Shared/Exclusive Locks

Legend

Active thread

Sleeping thread

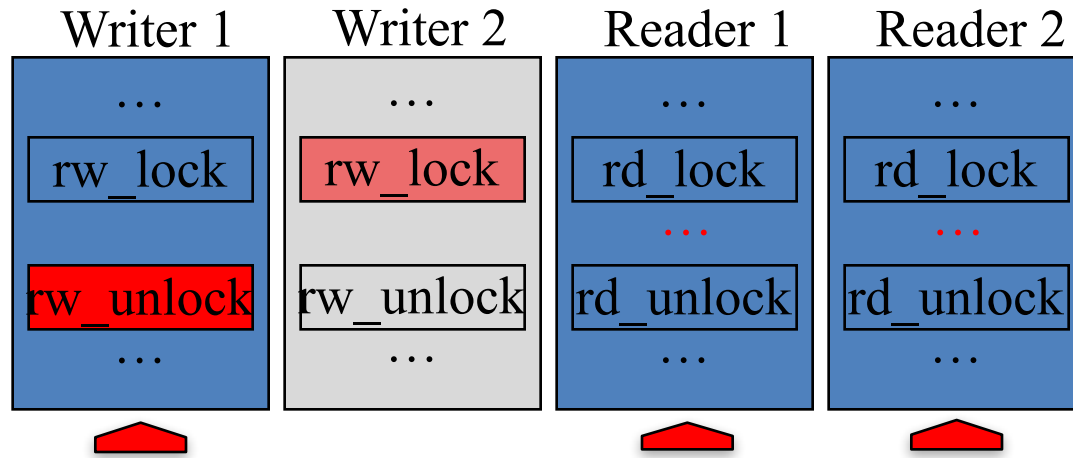


Shared/Exclusive Locks

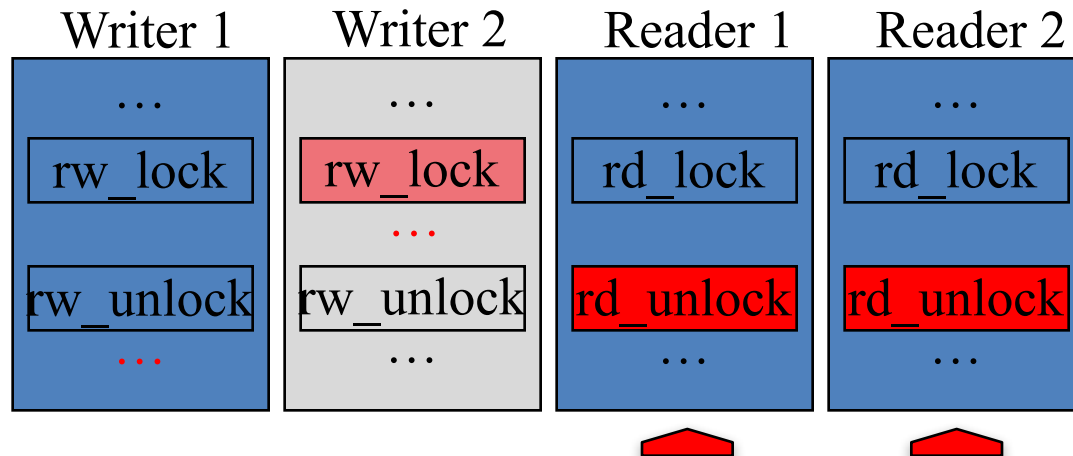
Legend

Active thread

Sleeping thread



Step 5



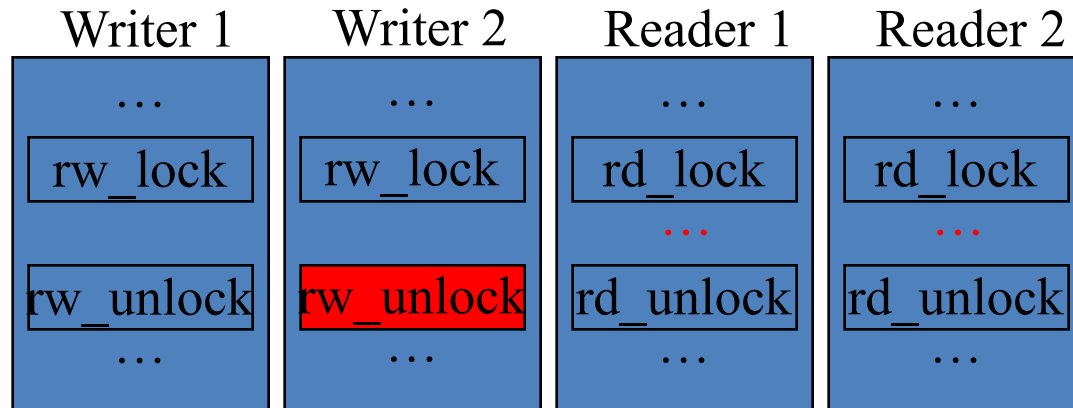
Step 6

Shared/Exclusive Locks

Legend

Active thread


Sleeping thread



Step 7


Condition Variable

- Block a thread while waiting for a condition
- Condition_wait / condition_signal
- Several thread can wait for the same condition, they all get the signal




Condition Variable

- Block a thread while waiting for a condition
- Condition_wait / condition_signal
- Several thread can wait for the same condition, they all get the signal




Condition Variable

- Block a thread while waiting for a condition
- Condition_wait / condition_signal
- Several thread can wait for the same condition, they all get the signal




Condition Variable

- Block a thread while waiting for a condition
- Condition_wait / condition_signal
- Several thread can wait for the same condition, they all get the signal



Condition Variable

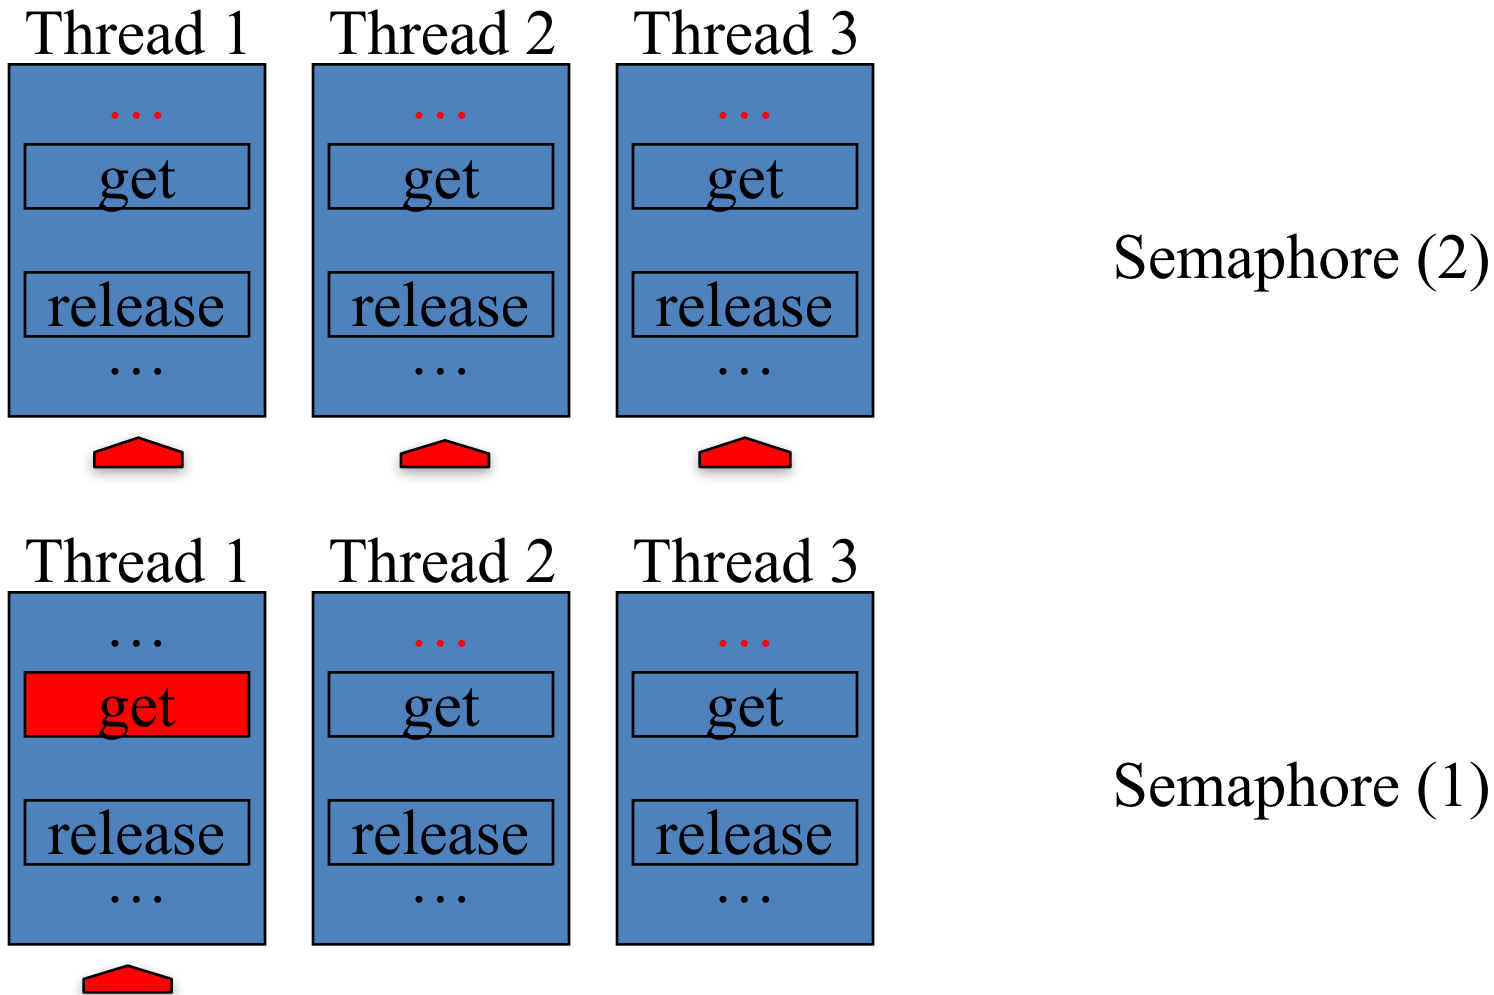
- Block a thread while waiting for a condition
- Condition_wait / condition_signal
- Several thread can wait for the same condition, they all get the signal



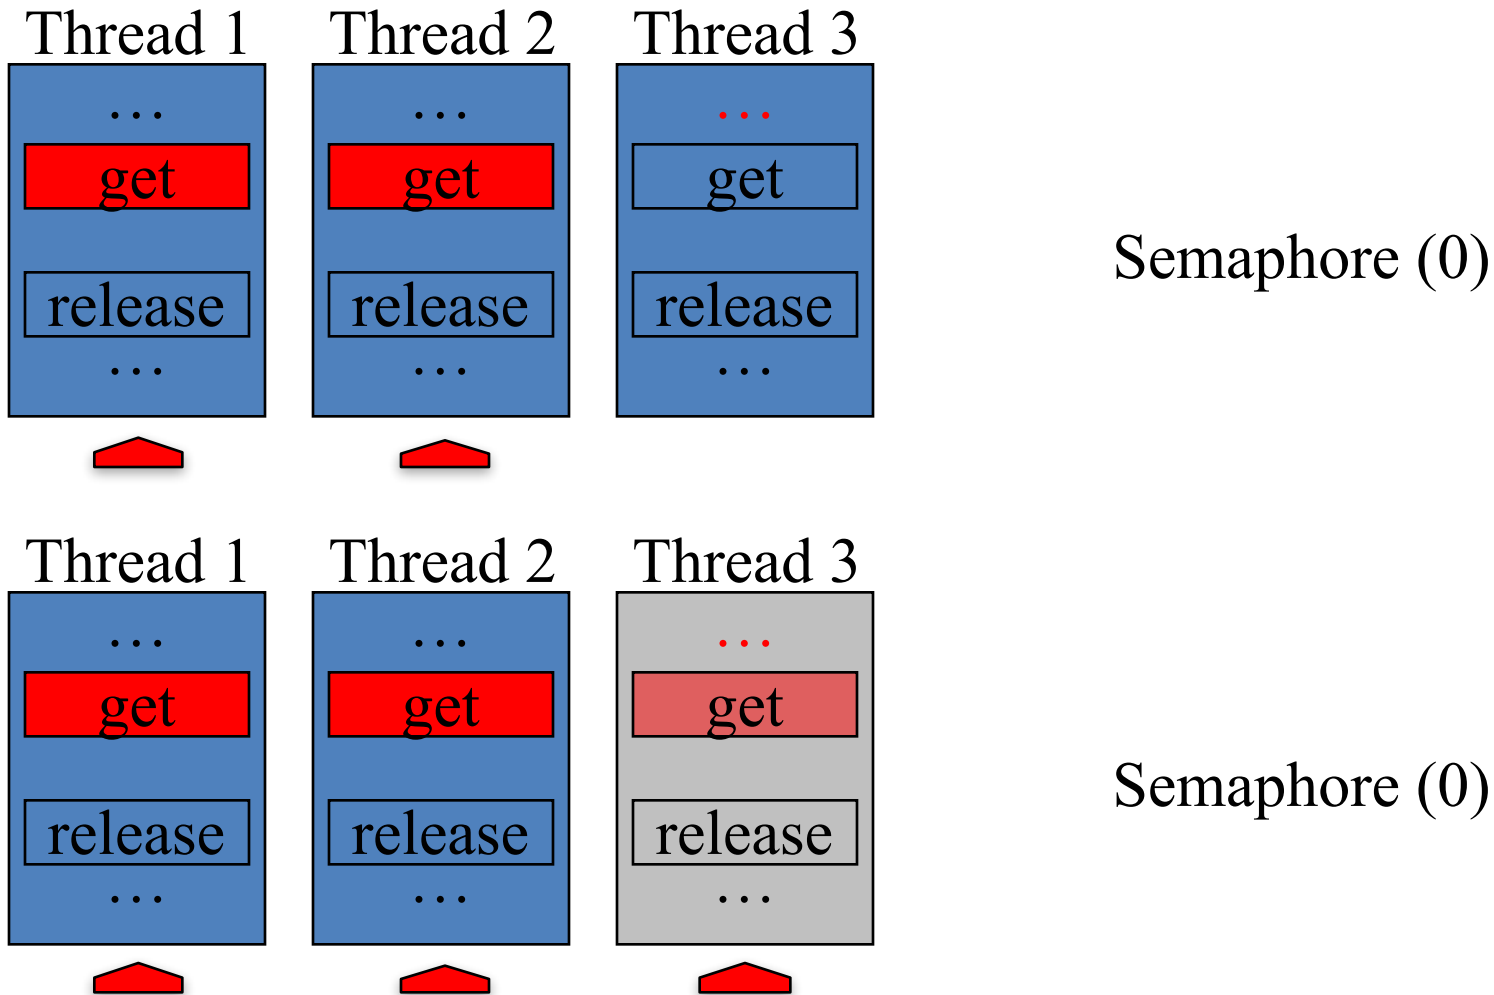
Semaphores

- simple counting mutexes
- The semaphore can be hold by as many threads as the initial value of the semaphore.
- When a thread get the semaphore it decrease the internal value by 1.
- When a thread release the semaphore it increase the internal value by 1.


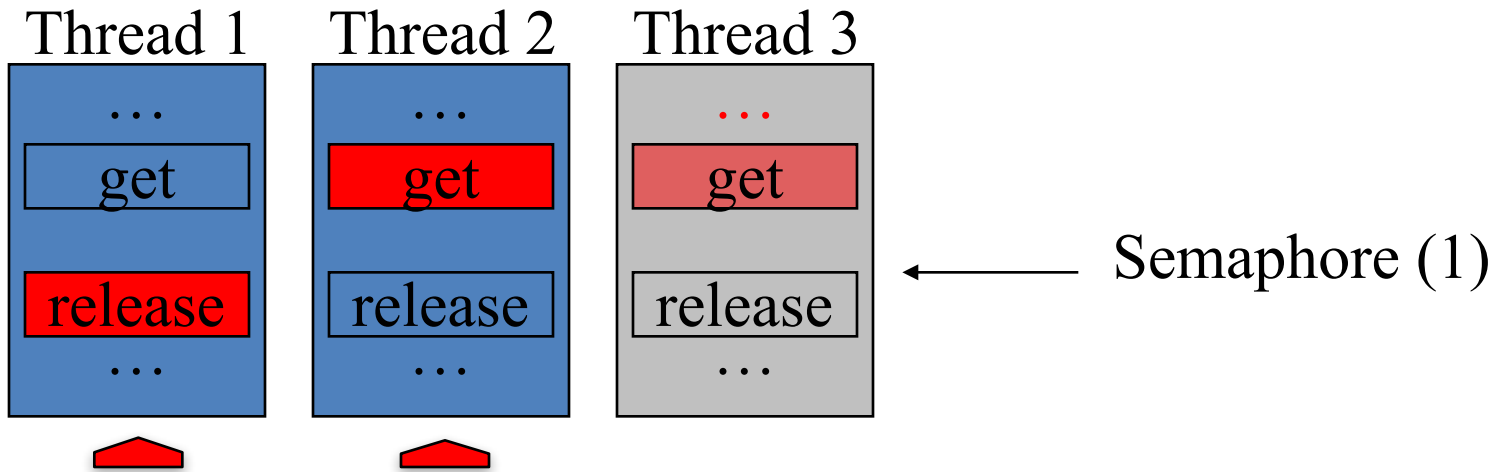
Semaphores



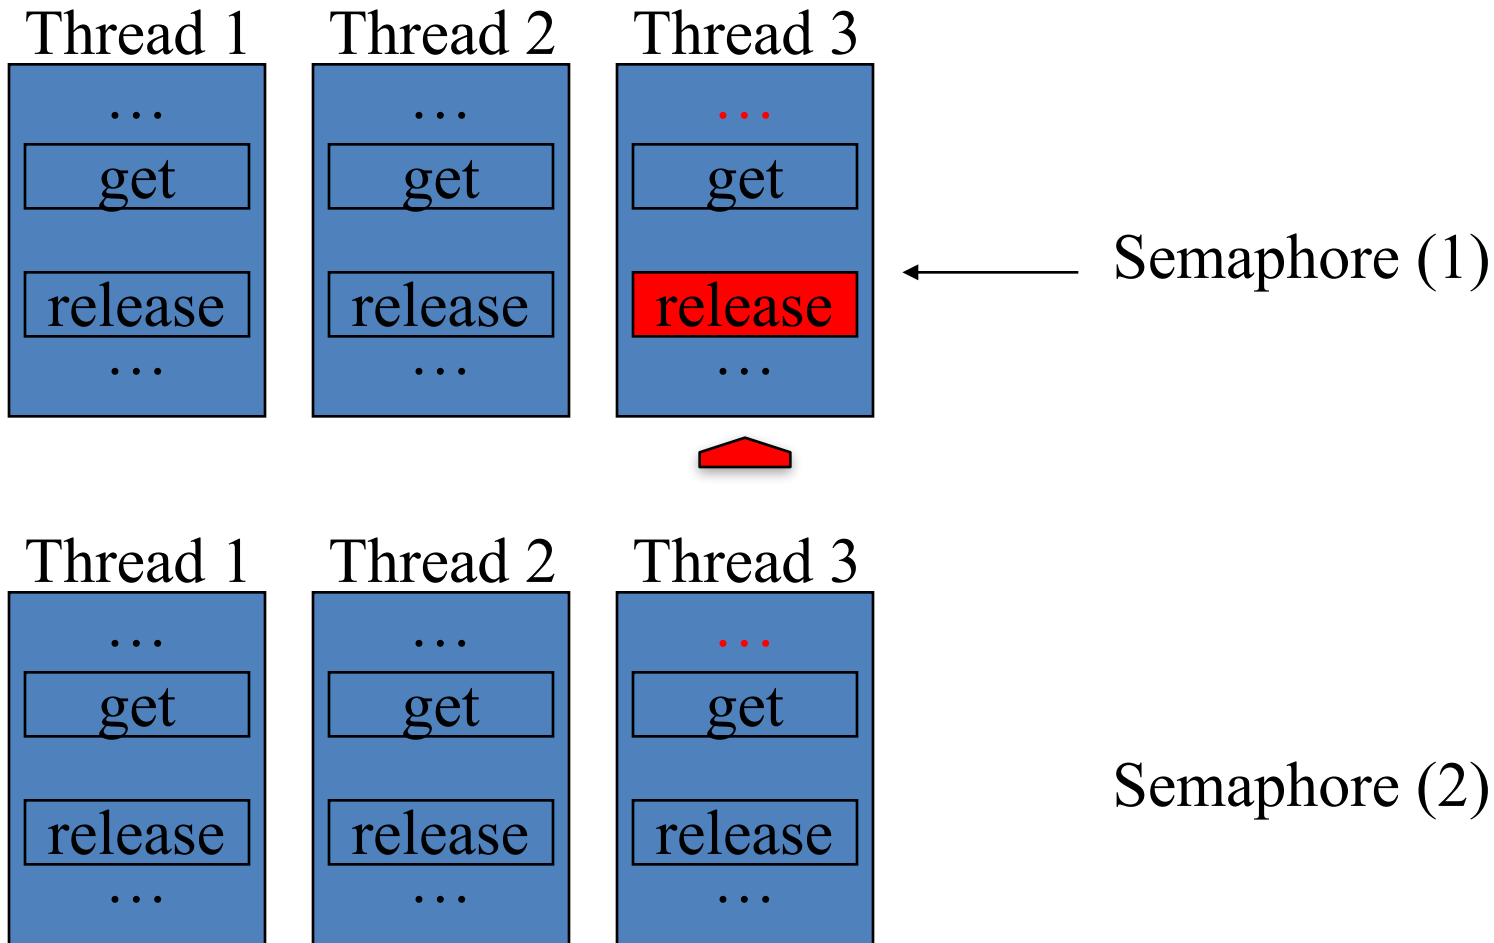
Semaphores



Semaphores



Semaphores



Atomic instruction

- Is any operation that a CPU can perform such that all results will be made visible to each CPU at the same time and whose operation is safe from interference by other CPUs
 - TestAndSet
 - CompareAndSwap
 - DoubleCompareAndSwap
 - Atomic increment
 - Atomic decrement

Pthread API

Prefix	Use
pthread_	Thread management (create/destroy/cancel/join/exit)
pthread_attr_	Thread attributes
pthread_mutex_	Mutexes
pthread_mutexattr_	Mutexes attributes
pthread_cond_	Condition variables
pthread_condattr_	Condition attributes
pthread_key_	Thread-specific data key (TLS)
pthread_rwlock_	Read/write locks
pthread_barrier_	Synchronization barriers

Thread Management (create)

<code>int pthread_create (pthread_t *restrict thread, const pthread_attr_t *restrict attr, void *(*start_routine)(void *), void *restrict arg)</code>	[OUT] thread id [IN] attributes [IN] thread function [IN] argument for thread function
<code>int pthread_exit (void *value_ptr)</code>	[OUT] Return to caller/joiner
<code>int pthread_cancel (pthread_t thread)</code>	[IN] thread to be cancelled
<code>int pthread_attr_init (pthread_attr_t *attr)</code>	[OUT] attributes to be initialized
<code>int pthread_attr_set*(pthread_attr_t *restrict attr, *)</code>	[IN] set attributed (state, stack)
<code>int pthread_attr_destroy (pthread_attr_t *attr)</code>	[IN] attributed to be destroyed

Attributes: Detached or joinable state, Scheduling inheritance, Scheduling policy, Scheduling parameters, Scheduling contention scope, Stack size, Stack address, Stack guard (overflow) size

Questions:

- Once created what will be the status of the thread and how it will be scheduled by the OS ? (use [sched_setscheduler](#))
- Where it will be run ? (use [sched_setaffinity](#) or [HWLOC](#))