



## Named POSIX Semaphores (3)

```
/* named-sem-init.c */
#include <semaphore.h>
#include <asm/fcntl.h>
#define POSIX_LOCKED 0
#define POSIX_UNLOCKED 1

sem_t *posix_sem;

main () {
    posix_sem = sem_open("/MySemaphore", O_CREAT, 0644, POSIX_UNLOCKED);
    sem_init(posix_sem, 0, 5); /* init: 5 */
}

$ ./named-sem-init
$ ls -l /dev/shm/
-rw-r----- 1 esser users 16 2006-12-05 15:46 sem.MySemaphore
$ hexdump /dev/shm/sem.MySemaphore
00000000 0005 0000 0000 0000 0000 0000 0000 0000
00000010
```

## Named POSIX Semaphores (5)

### Signal operation

```
/* named-sem-signal.c */
#include <semaphore.h>
#include <asm/fcntl.h>

#define POSIX_LOCKED 0
#define POSIX_UNLOCKED 1

sem_t *posix_sem;

main () {
    posix_sem = sem_open("/MySemaphore", O_CREAT, 0644, POSIX_UNLOCKED);
    sem_post(posix_sem);
}

$ ./named-sem-query
Semaphore value 5
$ ./named-sem-signal
$ ./named-sem-query
Semaphore value 6
```

## Named POSIX Semaphores (4)

```
/* named-sem-query.c */
#include <semaphore.h>
#include <asm/fcntl.h>
#define POSIX_LOCKED 0
#define POSIX_UNLOCKED 1

sem_t *posix_sem;

main () {
    int ret;
    posix_sem = sem_open("/MySemaphore", O_CREAT, 0644, POSIX_UNLOCKED);
    sem_getvalue(posix_sem, &ret);
    printf ("Semaphore value %d \n", ret);
}

$ ./named-sem-query
Semaphore value 5
```

## Named POSIX Semaphores (6)

### Wait operation

```
/* named-sem-wait.c */
#include <semaphore.h>
#include <asm/fcntl.h>

#define POSIX_LOCKED 0
#define POSIX_UNLOCKED 1

sem_t *posix_sem;

main () {
    posix_sem = sem_open("/MySemaphore", O_CREAT, 0644, POSIX_UNLOCKED);
    sem_wait(posix_sem);
}

$ ./named-sem-query
Semaphore value 2
$ ./named-sem-wait
$ ./named-sem-query
Semaphore value 1
$ ./named-sem-wait
```

here the process blocks until the semaphore is incremented

## Mutex for Processes

- use a named POSIX semaphore  
(reminder: binary semaphore = mutex)

- **SO:**

```
posix_sem = sem_open("/mutex", O_CREAT,  
                    0644, POSIX_UNLOCKED);  
sem_init(&posix_sem, 0, 1);      /* 1: mutex */
```

## System V IPC Semaphores (2)

- create semaphore: `semget()`
- initialize semaphore: `semctl()`
- use semaphore: `semop()`

– includes **signal** and **wait** operations:

```
struct sembuf          /* in <sys/sem.h> */  
{  
    unsigned short int sem_num; /* semaphore number */  
    short int sem_op;          /* semaphore operation */  
    short int sem_flg;        /* operation flag */  
};
```

`sem_op`: -1 = wait, 1 = signal

`sem_flg`: IPC\_NOWAIT → error instead of waiting

## System V IPC Semaphores (1)

- Alternative to named Posix semaphores:  
System V IPC Semaphores
- System V IPC: Methods for Inter Process  
Communication (IPC)  
(more about this: Ch. 6, IPC)
- A little more complex:
  - semaphore sets (can contain several semaphores)
  - private semaphores (only for process and children)
  - public semaphores (with identifiers)

## System V IPC Semaphores (3)

### Producer Consumer Problem with SysV semaphores (1)

```
/*  
 * sem-producer-consumer.c  
 */  
  
#include <stdio.h>          /* standard I/O routines.          */  
#include <stdlib.h>         /* rand() and srand() functions    */  
#include <unistd.h>         /* fork(), etc.                   */  
#include <time.h>           /* nanosleep(), etc.              */  
#include <sys/types.h>      /* various type definitions.       */  
#include <sys/ipc.h>        /* general SysV IPC structures    */  
#include <sys/sem.h>        /* semaphore functions and structs.*/  
  
#define NUM_LOOPS 20      /* number of loops to perform.    */  
  
union semun { int val; struct semid_ds *buf; unsigned short *array; };  
  
int main(int argc, char* argv[]) {  
    int sem_set_id;        /* ID of the semaphore set.        */  
    union semun sem_val;  /* semaphore value, for semctl().  */  
    int child_pid;        /* PID of our child process.       */  
    int i;                 /* counter for loop operation.     */  
    struct sembuf sem_op; /* structure for semaphore ops.    */  
    int rc;                /* return value of system calls.   */  
    struct timespec delay; /* used for wasting time.          */
```

# System V IPC Semaphores (4)

## Producer Consumer Problem with SysV semaphores (2)

```
/* create private sem. set with one sem. in it, access only to the owner. */
sem_set_id = semget(IPC_PRIVATE, 1, 0600);
if (sem_set_id == -1) { perror("main: semget"); exit(1); }
printf("semaphore set created, semaphore set id '%d'.\n", sem_set_id);

/* initialize the first (and single) semaphore in our set to '0'. */
sem_val.val = 0;
rc = semctl(sem_set_id, 0, SETVAL, sem_val);

/* fork-off a child process, and start a producer/consumer job. */
child_pid = fork();
switch (child_pid) {
    case -1: perror("fork"); exit(1);
    case 0: /* child process: consumer */
        for (i=0; i<NUM_LOOPS; i++) {
            /* block on the semaphore, unless its value is non-negative. */
            sem_op.sem_num = 0;
            sem_op.sem_op = -1; /* <- -1: count down */
            sem_op.sem_flg = 0;
            semop(sem_set_id, &sem_op, 1); /* wait (semaphore) */
            printf("consumer: '%d'\n", i); fflush(stdout);
        }
        break;
}
```

# System V IPC Semaphores (6)

## Variant with two separate programs:

- common key allows accessing the (same) semaphore
- create key with *ftok()* („filename to key“):
- change *semget()* call:

```
key_t semkey = ftok("/tmp", 'a');

/* create private semaphore set */
sem_set_id = semget(IPC_PRIVATE, 1, 0600);

becomes

/* create public semaphore set */
semkey = ftok("/tmp", 'a');
sem_set_id = semget(semkey, 1, 0);
```

# System V IPC Semaphores (5)

## Producer Consumer Problem with SysV semaphores (3)

```
default: /* parent process: producer */
    for (i=0; i<NUM_LOOPS; i++) {
        printf("producer: '%d'\n", i); fflush(stdout);
        /* increase the value of the semaphore by 1. */
        sem_op.sem_num = 0;
        sem_op.sem_op = 1; /* <- +1: count up */
        sem_op.sem_flg = 0;
        semop(sem_set_id, &sem_op, 1); /* signal (semaphore) */
        /* pause execution for a bit, to allow the child process to run */
        /* and handle some requests. this is done about 25% of the time. */
        if (rand() > 3*(RAND_MAX/4)) {
            delay.tv_sec = 0;
            delay.tv_nsec = 10;
            nanosleep(&delay, NULL);
        }
        break;
    }
    return 0;
}
```

Source: <http://users.actcom.co.il/~choo/lupg/tutorials/multi-process/multi-process.html#semaphores>

# System V IPC Semaphores (7)

```
1 /*
2  * producer.c
3  */
4
5 #include <stdio.h> /* standard I/O routines. */
6 #include <stdlib.h> /* rand() and srand() functions. */
7 #include <unistd.h> /* fork(), etc. */
8 #include <time.h> /* nanosleep(), etc. */
9 #include <sys/types.h> /* various type definitions. */
10 #include <sys/ipc.h> /* general sysV IPC structures. */
11 #include <sys/sem.h> /* semaphore functions and structs. */
12
13 #define NUM_LOOPS 20 /* number of loops to perform. */
14
15 union semun { int val; struct semid_ds *buf; unsigned short *array; };
16
17 int main(int argc, char* argv[])
18 {
19     int sem_set_id; /* ID of the semaphore set. */
20     key_t semkey; /* key for named semaphore set. */
21     union semun sem_val; /* semaphore value, for semctl(). */
22     int child_pid; /* PID of our child process. */
23     int i; /* counter for loop operation. */
24     struct sembuf sem_op; /* structure for semaphore ops. */
25     int rc; /* return value of system calls. */
26     struct timespec delay; /* used for wasting time. */
27
28     /* create a public semaphore set with one semaphore in it,
29     /* with access only to the owner. */
30     semkey = ftok("/tmp", 'a');
31     sem_set_id = semget(semkey, 0, 0);
32     if (sem_set_id == -1) {
33         perror("main: semget");
34         exit(1);
35     }
36     printf("semaphore set created, semaphore set id '%d'.\n", sem_set_id);
37
38     /* initialize the first (and single) semaphore in our set to '0'. */
39     sem_val.val = 0;
40     rc = semctl(sem_set_id, 0, SETVAL, sem_val);
41
42     for (i=0; i<NUM_LOOPS; i++) {
43         printf("producer: '%d'\n", i);
44         fflush(stdout);
45         /* increase the value of the semaphore by 1. */
46         sem_op.sem_num = 0;
47         sem_op.sem_op = 1;
48         sem_op.sem_flg = 0;
49         semop(sem_set_id, &sem_op, 1);
50         /* pause execution for a little bit, to allow the
51         /* child process to run and handle some requests. */
52         /* this is done about 25% of the time. */
53         if (rand() > 3*(RAND_MAX/4)) {
54             delay.tv_sec = 0;
55             delay.tv_nsec = 10;
56             nanosleep(&delay, NULL);
57         }
58     }
59     return 0;
60 }
```

# System V IPC Semaphores (8)

Terminal 1

```
$ gcc -o consumer consumer.c
$ gcc -o producer producer.c
$ ./consumer
semaphore set created,
semaphore set id '374964228'.
consumer: '0'
consumer: '1'
consumer: '2'
consumer: '3'
consumer: '4'
consumer: '5'
consumer: '6'
consumer: '7'
consumer: '8'
consumer: '9'
consumer: '10'
consumer: '11'
consumer: '12'
consumer: '13'
consumer: '14'
consumer: '15'
consumer: '16'
consumer: '17'
consumer: '18'
consumer: '19'
$ _
```

Terminal 2

```
$ ./producer
semaphore set created,
semaphore set id '374964228'.
producer: '0'
producer: '1'
producer: '2'
producer: '3'
producer: '4'
producer: '5'
producer: '6'
producer: '7'
producer: '8'
producer: '9'
producer: '10'
producer: '11'
producer: '12'
producer: '13'
producer: '14'
producer: '15'
producer: '16'
producer: '17'
producer: '18'
producer: '19'
$ _
```

consumer.c und producer.c auf der Vorlesungs-Web-Seite verfügbar

# Synchronization in the Linux Kernel

- atomic operations
  - on Integer variables (atomic\_set, atomic\_add, atomic\_inc, ...)
  - bit operations on bit vectors (set\_bit, clear\_bit, test\_and\_set, ...)
- Spin Locks / Reader Writer Spin Locks
- Semaphores / Reader Writer Semaphores
- „Big Kernel Lock“

```
Sep 19 14:20:18 amd64 sbhd[20494]: Accepted rsa for esser from ::ffff:87.234.201.207 port 61557
Sep 19 14:27:41 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 20 01:00:01 amd64 /usr/sbin/cron[29278]: (root) CMD (/sbin/evlogmgr -c 'severity=DEBUG')
Sep 20 01:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 20 02:00:01 amd64 /usr/sbin/cron[16033]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 20 02:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 20 12:46:44 amd64 sbhd[6516]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62004
Sep 20 12:46:44 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 20 12:54:41 amd64 sbhd[6691]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62105
Sep 20 12:54:44 amd64 sbhd[6694]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62514
Sep 20 15:27:35 amd64 sbhd[9077]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64242
Sep 20 15:27:35 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 20 16:37:11 amd64 sbhd[10102]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63375
Sep 20 16:37:11 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 20 16:38:10 amd64 sbhd[10141]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63546
Sep 21 01:00:01 amd64 /usr/sbin/cron[17055]: (root) CMD (/sbin/evlogmgr -c 'severity=DEBUG')
Sep 21 01:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 21 02:00:01 amd64 /usr/sbin/cron[17878]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 21 02:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 21 17:43:26 amd64 sbhd[11088]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63397
Sep 21 17:43:26 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 21 17:53:39 amd64 sbhd[11069]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64391
Sep 21 18:43:26 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 21 19:43:26 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 22 01:00:01 amd64 /usr/sbin/cron[4674]: (root) CMD (/sbin/evlogmgr -c 'severity=DEBUG')
Sep 22 01:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 22 02:00:01 amd64 /usr/sbin/cron[5191]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 22 02:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 22 20:23:21 amd64 /usr/sbin/cron[13191]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 23 01:00:01 amd64 /usr/sbin/cron[21111]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 23 01:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 23 02:00:01 amd64 /usr/sbin/cron[25555]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 23 02:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 23 18:04:05 amd64 sbhd[5541]: Accepted rsa for esser from ::ffff:87.234.201.207 port 69771
Sep 23 18:04:05 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 23 18:04:34 amd64 sbhd[6606]: Accepted rsa for esser from ::ffff:87.234.201.207 port 69771
Sep 24 01:00:01 amd64 /usr/sbin/cron[31313]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 24 01:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 02:00:01 amd64 /usr/sbin/cron[13253]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 24 02:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 11:15:48 amd64 sbhd[2098]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64456
Sep 24 11:15:48 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 13:49:08 amd64 sbhd[23197]: Accepted rsa for esser from ::ffff:87.234.201.207 port 61330
Sep 24 13:49:08 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 15:42:07 amd64 kernel: snd_seq_midi_event: unsupported module, tainting kernel.
Sep 24 15:42:07 amd64 kernel: snd_seq_midi_event: unsupported module, tainting kernel.
Sep 24 15:42:07 amd64 kernel: snd_seq_midi_event: unsupported module, tainting kernel.
Sep 24 15:42:07 amd64 kernel: snd_seq_midi_event: unsupported module, tainting kernel.
Sep 24 20:25:31 amd64 sbhd[28399]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62566
Sep 24 20:25:31 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 01:00:02 amd64 /usr/sbin/cron[1621]: (root) CMD (/sbin/evlogmgr -c 'severity=DEBUG')
Sep 25 01:00:02 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 02:00:01 amd64 /usr/sbin/cron[1484]: (root) CMD (/sbin/evlogmgr -c 'age > *30d*')
Sep 25 02:00:02 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 10:59:25 amd64 sbhd[8889]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64183
Sep 25 10:59:25 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 10:59:47 amd64 sbhd[8921]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64253
Sep 25 11:30:02 amd64 sbhd[9172]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62029
Sep 25 11:59:25 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 14:05:37 amd64 sbhd[11554]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62822
Sep 25 14:05:37 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 14:06:10 amd64 sbhd[11586]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62951
Sep 25 14:07:17 amd64 sbhd[11608]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63392
Sep 25 14:08:33 amd64 sbhd[11630]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63709
Sep 25 15:25:33 amd64 sbhd[12930]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62778
```

# Linux: Synchronization in the Kernel

## Atomic Integer Operations (1)

- new type *atomic\_t* (24 bit integer)
- initialization: *atomic\_t var = ATOMIC\_INIT(0);*
- set value: *atomic\_set(&var, wert);*
- add: *atomic\_add(wert, &var);*
- increment: *atomic\_inc(&var);*
- subtract: *atomic\_sub(wert, &var);*
- decrement: *atomic\_dec(&var);*
- read: *int i = atomic\_read(&var);*

## Atomic Integer Operations (2)

- ***res = atomic\_sub\_and\_test (i, &var);***  
atomically subtracts *i* from *var*
  - return value true, if result is 0;
  - return value false, if result is not 0
- ***res = atomic\_dec\_and\_test (&var);***  
***res = atomic\_inc\_and\_test (&var);***  
atomically decrements or increments *var*
  - return value true, if result is 0;
  - return value false, if result is not 0

## Atomic Bit Operations (1)

- set and clear single bits in bit vector
- datatype: arbitrary, e.g. *unsigned long bitvector = 0;*
  - access through pointer
  - number of settable / clearable bits depends on size of the used datatype
- ***set\_bit (i, &bitvector);*** set *i*-th bit
- ***clear\_bit (i, &bitvector);*** clear *i*-th bit
- ***change\_bit (i, &bitvector);*** change *i*-th bit

## Atomic Integer Operations (3)

- ***res = atomic\_add\_negative (i, &var);***  
atomically adds *i* to *var*.
  - return value true, if result is negative;
  - return value false, if result is  $\geq 0$

## Atomic Bit Operations (2)

- Test-and-Set operations return the previous value of the bit
  - ***b = test\_and\_set\_bit (i, &bitvector);***
  - ***b = test\_and\_clear\_bit (i, &bitvector);***
  - ***b = test\_and\_change\_bit (i, &bitvector);***
- read single bit
  - ***b = test\_bit (i, &bitvector);***
- search functions
  - ***pos = find\_first\_bit (&bitvector, length);***
  - ***pos = find\_first\_zero\_bit (&bitvector, length);***

## Spin Locks (1)

- Lock with Mutex functionality: mutual exclusion
- Code which requests a Spin Lock but is not successful, will spin in a loop until the lock becomes available
- Type: `spinlock_t`

```
spinlock_t xy_lock = SPIN_LOCK_UNLOCKED

spin_lock (&xy_lock);
/* critical region */
spin_unlock (&xy_lock);
```

## Spin Locks (3)

- when all interrupts are on before acquiring the spin lock, there's a simpler method:

```
spinlock_t xy_lock = SPIN_LOCK_UNLOCKED

spin_lock_irq (&xy_lock);
/* critical section */
spin_unlock_irq (&xy_lock);
```

disables / reenables all interrupts

- spin locks are not „recursive“, i.e.: it's impossible to acquire the same spin lock twice, e.g. when calling a function recursively

## Spin Locks (2)

- since Spin Locks don't sleep they can be used inside interrupt handlers
- In that case: also disable interrupts:

```
spinlock_t xy_lock = SPIN_LOCK_UNLOCKED
unsigned long flags;

spin_lock_irqsave (&xy_lock, flags);
/* critical region */
spin_unlock_irqrestore (&xy_lock, flags);
```

(save current interrupt settings in *flags*, then disable; restore original state)

## Spin Locks (4)

- to avoid blocking, it is possible to query the spin lock state with `spin_is_locked (&xy_lock)`;

- locking attempt with `spin_try_lock`:

```
if ( spin_try_lock (&xy_lock) ) {
    /* critical region */
    spin_unlock (&xy_lock);
} else {
    /* was not allowed to enter the critical region */
}
```

- both functions should not be used: either you need the lock (and in that case will possibly have to wait), or you don't need it ...

## Reader Writer Locks (1)

- alternative to normal locks, allowing several concurrent read accesses – but exclusive for write access (like a standard lock):

```
rwlock_t xy_rwlock = RW_LOCK_UNLOCKED;
```

reading code

```
read_lock (&xy_rwlock) ) {
    /* critical region,
       read-only */
read_unlock (&xy_rwlock);
```

writing code

```
write_lock (&xy_rwlock) ) {
    /* critical region,
       read & write */
write_unlock (&xy_rwlock);
```

- only use in case of strict separation of reading and writing code parts

## Semaphores (1)

- Kernel semaphores are „sleeping“ locks
- if a semaphore is already locked, new requesters are put into a queue.
- when freeing a semaphore, the first waiting thread of the queue will be awoken
- semaphores are useful for locks that shall be kept over a longer period of time
  - no waste of CPU time

## Reader Writer Locks (2)

- |                  | there is already a reader | there is already a writer | no locks yet |
|------------------|---------------------------|---------------------------|--------------|
| read_lock(&lck)  | successful                | fails                     | successful   |
| write_lock(&lck) | fails                     | fails                     | successful   |

- there are variants for disabling interrupts, too:
  - read\_lock\_irq                      read\_unlock\_irq
  - read\_lock\_irqsave                read\_unlock\_irqrestore
  - write\_lock\_irq                     write\_unlock\_irq
  - write\_lock\_irqsave                write\_unpock\_irqrestore

## Semaphores (2)

- semaphores can only be used in process context, not in interrupt handlers (the scheduler does not deal with interrupt handlers)
- code which wants to use a semaphore, must not hold a normal lock (semaphore access can cause the thread to be put to sleep)
- semaphores can let more than one thread access the resource



## Semaphores (3)

Type: *semaphore*

Static declaration

```
static DECLARE_SEMAPHORE_GENERIC (name, count);
static DECLARE_MUTEX (name);          /* count=1 */
```

dynamic semaphore creation

```
sema_init (&sem, count);
init_MUTEX (&sem);                  /* count=1 */
```

- use with *up()* and *down()*

```
down (&sem);
/* critical section */
up (&sem);
```

## Semaphores (5)

- example for *down\_trylock()*

```
/* taken from /usr/src/linux/kernel/printk.c */
if (!down_trylock(&console_sem)) {
    console_locked = 1;
    /*
     * We own the drivers. We can drop the spinlock and let
     * release_console_sem() print the text
     */
    spin_unlock_irqrestore(&logbuf_lock, flags);
    console_may_schedule = 0;
    release_console_sem();
    /* function release_console_sem() calls up(&console_sem); */
} else {
    /*
     * Someone else owns the drivers. We drop the spinlock, which
     * allows the semaphore holder to proceed and to call the
     * console drivers with the output which we just produced.
     */
    spin_unlock_irqrestore(&logbuf_lock, flags);
}
```

## Semaphores (4)

- Variants of *down()*
  - *down(&sem);*  
non-interruptible sleep, if semaphore is not available
  - *down\_interruptible(&sem);*  
interruptible sleep, if semaphore is not available
  - *down\_trylock(&sem);*  
tries to acquire the semaphore – if that fails, this function will return with a false value

## Reader Writer Semaphores (1)

- similar to Reader Writer Locks:  
Type *rw\_semaphore*, allowing special Up and Down operations for read and write access
- all Reader Writer Semaphores are mutexes  
(counter is always 1 at initialization)

static declaration

```
static DECLARE_RWSEM (name);
```

dynamical semaphore creation

```
init_rwsem (&sem);
```

## Reader Writer Semaphores (2)

```
static DECLARE_RWSEM (xy_rwsem);
```

reading code

```
down_read (&xy_rwsem) ) {
    /* critical region,
       read-only */
up_read (&xy_rwsem);
```

writing code

```
down_write (&xy_rwsem) ) {
    /* critical region,
       read & write */
up_write (&xy_rwsem);
```

just like Reader Writer Locks:

	there is already a reader	there is already a writer	no locks yet
down_read(&sem)	successful	fails	successful
down_write(&sem)	fails	fails	successful

## „Big Kernel Lock“ (BKL) (2)

- BKL can only be used in process context (not in interrupt routines)
- a process holding the BKL may sleep
  - when going to sleep, the BKL is automatically released
  - on wake-up it will be re-acquired
- BKL is recursive: a process which already holds the BKL may call *lock\_kernel()* again
- don't use it!

## „Big Kernel Lock“ (BKL) (1)

- relict from older kernel versions
- global lock for the whole kernel (which affects all kernel regions that protect data access with it)

```
lock_kernel ();
/* critical region */
unlock_kernel ();

if ( kernel_locked() ) {
    ...
}
```