

Chapter 6

Startup Parameters

In many of the previous examples, we had to hard-wire something into the kernel module, such as the file name for `/proc` files or the major device number for the device so we can have `ioctl`'s to it. This goes against the grain of the Unix, and Linux, philosophy which is to write flexible program the user can customize.

The way to tell a program, or a kernel module, something it needs before it can start working is by command line parameters. In the case of kernel modules, we don't get `argc` and `argv` — instead, we get something better. We can define global variables in the kernel module and `insmod` will fill them for us.

In this kernel module, we define two of them: `str1` and `str2`. All you need to do is compile the kernel module and then run `insmod str1=xxx str2=yyy`. When `init_module` is called, `str1` will point to the string 'xxx' and `str2` to the string 'yyy'.

In version 2.0 there is no type checking on these arguments¹. If the first character of `str1` or `str2` is a digit the kernel will fill the variable with the value of the integer, rather than a pointer to the string. If a real life situation you have to check for this.

On the other hand, in version 2.2 you use the macro `MACRO_PARM` to tell `insmod` that you expect a parameters, its name *and its type*. This solves the type problem and allows kernel modules to receive strings which begin with a digit, for example.

param.c

¹There can't be, since under C the object file only has the location of global variables, not their type. That is why header files are necessary

```
/* param.c
 *
 * Receive command line parameters at module installation
 */

/* Copyright (C) 1998-99 by Ori Pomerantz */

/* The necessary header files */

/* Standard in kernel modules */
#include <linux/kernel.h> /* We're doing kernel work */
#include <linux/module.h> /* Specifically, a module */

/* Deal with CONFIG_MODVERSIONS */
#if CONFIG_MODVERSIONS==1
#define MODVERSIONS
#include <linux/modversions.h>
#endif

#include <stdio.h> /* I need NULL */

/* In 2.2.3 /usr/include/linux/version.h includes a
 * macro for this, but 2.0.35 doesn't - so I add it
 * here if necessary. */
#ifndef KERNEL_VERSION
#define KERNEL_VERSION(a,b,c) ((a)*65536+(b)*256+(c))
#endif
```

```

/* Emmanuel Papirakis:
 *
 * Parameter names are now (2.2) handled in a macro.
 * The kernel doesn't resolve the symbol names
 * like it seems to have once did.
 *
 * To pass parameters to a module, you have to use a macro
 * defined in include/linux/modules.h (line 176).
 * The macro takes two parameters. The parameter's name and
 * it's type. The type is a letter in double quotes.
 * For example, "i" should be an integer and "s" should
 * be a string.
 */

```

```

char *str1, *str2;

```

```

#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
MODULE_PARM(str1, "s");
MODULE_PARM(str2, "s");
#endif

```

```

/* Initialize the module - show the parameters */
int init_module()
{
    if (str1 == NULL || str2 == NULL) {
        printk("Next time, do insmod param str1=<something>");
        printk("str2=<something>\n");
    } else
        printk("Strings:%s and %s\n", str1, str2);
}

#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
    printk("If you try to insmod this module twice,");
    printk("(without rmmod'ing\n");

```

```
    printk("it first), you might get the wrong");
    printk("error message:\n");
    printk("'symbol for parameters str1 not found'.\n");
#endif
```

```
    return 0;
}
```

```
/* Cleanup */
void cleanup_module()
{
}
```

Chapter 7

System Calls

So far, the only thing we've done was to use well defined kernel mechanisms to register `/proc` files and device handlers. This is fine if you want to do something the kernel programmers thought you'd want, such as write a device driver. But what if you want to do something unusual, to change the behavior of the system in some way? Then, you're mostly on your own.

This is where kernel programming gets dangerous. While writing the example below, I killed the `open` system call. This meant I couldn't open any files, I couldn't run any programs, and I couldn't shutdown the computer. I had to pull the power switch. Luckily, no files died. To ensure you won't lose any files either, please run `sync` right before you do the `insmod` and the `rmmod`.

Forget about `/proc` files, forget about device files. They're just minor details. The *real* process to kernel communication mechanism, the one used by all processes, is system calls. When a process requests a service from the kernel (such as opening a file, forking to a new process, or requesting more memory), this is the mechanism used. If you want to change the behaviour of the kernel in interesting ways, this is the place to do it. By the way, if you want to see which system calls a program uses, run `strace <command> <arguments>`.

In general, a process is not supposed to be able to access the kernel. It can't access kernel memory and it can't call kernel functions. The hardware of the CPU enforces this (that's the reason why it's called 'protected mode'). System calls are an exception to this general rule. What happens is that the process fills the registers with the appropriate values and then calls a special instruction which jumps to a previously defined location in the

kernel (of course, that location is readable by user processes, it is not writable by them). Under Intel CPUs, this is done by means of interrupt 0x80. The hardware knows that once you jump to this location, you are no longer running in restricted user mode, but as the operating system kernel — and therefore you're allowed to do whatever you want.

The location in the kernel a process can jump to is called `system_call`. The procedure at that location checks the system call number, which tells the kernel what service the process requested. Then, it looks at the table of system calls (`sys_call_table`) to see the address of the kernel function to call. Then it calls the function, and after it returns, does a few system checks and then return back to the process (or to a different process, if the process time ran out). If you want to read this code, it's at the source file `arch/<architecture>/kernel/entry.S`, after the line `ENTRY(system_call)`.

So, if we want to change the way a certain system call works, what we need to do is to write our own function to implement it (usually by adding a bit of our own code, and then calling the original function) and then change the pointer at `sys_call_table` to point to our function. Because we might be removed later and we don't want to leave the system in an unstable state, it's important for `cleanup_module` to restore the table to its original state.

The source code here is an example of such a kernel module. We want to 'spy' on a certain user, and to `printk` a message whenever that user opens a file. Towards this end, we replace the system call to open a file with our own function, called `our_sys_open`. This function checks the `uid` (user's id) of the current process, and if it's equal to the `uid` we spy on, it calls `printk` to display the name of the file to be opened. Then, either way, it calls the original `open` function with the same parameters, to actually open the file.

The `init_module` function replaces the appropriate location in `sys_call_table` and keeps the original pointer in a variable. The `cleanup_module` function uses that variable to restore everything back to normal. This approach is dangerous, because of the possibility of two kernel modules changing the same system call. Imagine we have two kernel modules, A and B. A's open system call will be `A_open` and B's will be `B_open`. Now, when A is inserted into the kernel, the system call is replaced with `A_open`, which will call the original `sys_open` when it's done. Next, B is inserted into the kernel, which replaces the system call with `B_open`, which will call what it thinks is the original system call, `A_open`, when it's done.

Now, if B is removed first, everything will be well — it will simply restore the system call to `A_open`, which calls the original. However, if A is removed and then B is removed, the system will crash. A's removal will restore the system call to the original, `sys_open`,

cutting B out of the loop. Then, when B is removed, it will restore the system call to what **it** thinks is the original, `A_open`, which is no longer in memory. At first glance, it appears we could solve this particular problem by checking if the system call is equal to our open function and if so not changing it at all (so that B won't change the system call when it's removed), but that will cause an even worse problem. When A is removed, it sees that the system call was changed to `B_open` so that it is no longer pointing to `A_open`, so it won't restore it to `sys_open` before it is removed from memory. Unfortunately, `B_open` will still try to call `A_open` which is no longer there, so that even without removing B the system would crash.

I can think of two ways to prevent this problem. The first is to restore the call to the original value, `sys_open`. Unfortunately, `sys_open` is not part of the kernel system table in `/proc/ksyms`, so we can't access it. The other solution is to use the reference count to prevent root from `rmmod`'ing the module once it is loaded. This is good for production modules, but bad for an educational sample — which is why I didn't do it here.

syscall.c

```
/* syscall.c
 *
 * System call "stealing" sample
 */

/* Copyright (C) 1998-99 by Ori Pomerantz */

/* The necessary header files */

/* Standard in kernel modules */
#include <linux/kernel.h> /* We're doing kernel work */
#include <linux/module.h> /* Specifically, a module */

/* Deal with CONFIG_MODVERSIONS */
#if CONFIG_MODVERSIONS==1
#define MODVERSIONS
#include <linux/modversions.h>
#endif
```

```
#include <sys/syscall.h> /* The list of system calls */

/* For the current (process) structure, we need
 * this to know who the current user is. */
#include <linux/sched.h>

/* In 2.2.3 /usr/include/linux/version.h includes a
 * macro for this, but 2.0.35 doesn't - so I add it
 * here if necessary. */
#ifndef KERNEL_VERSION
#define KERNEL_VERSION(a,b,c) ((a)*65536+(b)*256+(c))
#endif

#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
#include <asm/uaccess.h>
#endif

/* The system call table (a table of functions). We
 * just define this as external, and the kernel will
 * fill it up for us when we are insmod'ed
 */
extern void *sys_call_table[];

/* UID we want to spy on - will be filled from the
 * command line */
int uid;
```

```

#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
MODULE_PARM(uid, "i");
#endif

/* A pointer to the original system call. The reason
 * we keep this, rather than call the original function
 * (sys_open), is because somebody else might have
 * replaced the system call before us. Note that this
 * is not 100% safe, because if another module
 * replaced sys_open before us, then when we're inserted
 * we'll call the function in that module - and it
 * might be removed before we are.
 *
 * Another reason for this is that we can't get sys_open.
 * It's a static variable, so it is not exported. */
asmlinkage int (*original_call)(const char *, int, int);

/* For some reason, in 2.2.3 current->uid gave me
 * zero, not the real user ID. I tried to find what went
 * wrong, but I couldn't do it in a short time, and
 * I'm lazy - so I'll just use the system call to get the
 * uid, the way a process would.
 *
 * For some reason, after I recompiled the kernel this
 * problem went away.
 */
asmlinkage int (*getuid_call)();

/* The function we'll replace sys_open (the function
 * called when you call the open system call) with. To
 * find the exact prototype, with the number and type
 * of arguments, we find the original function first

```

```

* (it's at fs/open.c).
*
* In theory, this means that we're tied to the
* current version of the kernel. In practice, the
* system calls almost never change (it would wreck havoc
* and require programs to be recompiled, since the system
* calls are the interface between the kernel and the
* processes).
*/
asmlinkage int our_sys_open(const char *filename,
                           int flags,
                           int mode)
{
    int i = 0;
    char ch;

    /* Check if this is the user we're spying on */
    if (uid == getuid_call()) {
        /* getuid_call is the getuid system call,
        * which gives the uid of the user who
        * ran the process which called the system
        * call we got */

        /* Report the file, if relevant */
        printk("Opened file by %d: ", uid);
        do {
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
            get_user(ch, filename+i);
#else
            ch = get_user(filename+i);
#endif
            i++;
            printk("%c", ch);
        } while (ch != 0);
        printk("\n");
    }
}

```

```

/* Call the original sys_open - otherwise, we lose
 * the ability to open files */
return original_call(filename, flags, mode);
}

/* Initialize the module - replace the system call */
int init_module()
{
/* Warning - too late for it now, but maybe for
 * next time... */
printk("I'm dangerous. I hope you did a ");
printk("sync before you insmod'ed me.\n");
printk("My counterpart, cleanup_module(), is even");
printk("more dangerous. If\n");
printk("you value your file system, it will ");
printk("be \"sync; rmmmod\" \n");
printk("when you remove this module.\n");

/* Keep a pointer to the original function in
 * original_call, and then replace the system call
 * in the system call table with our_sys_open */
original_call = sys_call_table[__NR_open];
sys_call_table[__NR_open] = our_sys_open;

/* To get the address of the function for system
 * call foo, go to sys_call_table[__NR_foo]. */

printk("Spying on UID:%d\n", uid);

/* Get the system call for getuid */
getuid_call = sys_call_table[__NR_getuid];

return 0;

```

```
}
```

```
/* Cleanup - unregister the appropriate file from /proc */
```

```
void cleanup_module()
```

```
{
```

```
    /* Return the system call back to normal */
```

```
    if (sys_call_table[__NR_open] != our_sys_open) {
```

```
        printk("Somebody else also played with the ");
```

```
        printk("open system call\n");
```

```
        printk("The system may be left in ");
```

```
        printk("an unstable state.\n");
```

```
    }
```

```
    sys_call_table[__NR_open] = original_call;
```

```
}
```

Chapter 8

Blocking Processes

What do you do when somebody asks you for something you can't do right away? If you're a human being and you're bothered by a human being, the only thing you can say is: 'Not right now, I'm busy. *Go away!*'. But if you're a kernel module and you're bothered by a process, you have another possibility. You can put the process to sleep until you can service it. After all, processes are being put to sleep by the kernel and woken up all the time (that's the way multiple processes appear to run on the same time on a single CPU).

This kernel module is an example of this. The file (called `/proc/sleep`) can only be opened by a single process at a time. If the file is already open, the kernel module calls `module_interruptible_sleep_on`¹. This function changes the status of the task (a task is the kernel data structure which holds information about a process and the system call it's in, if any) to `TASK_INTERRUPTIBLE`, which means that the task will not run until it is woken up somehow, and adds it to `WaitQ`, the queue of tasks waiting to access the file. Then, the function calls the scheduler to context switch to a different process, one which has some use for the CPU.

When a process is done with the file, it closes it, and `module_close` is called. That function wakes up all the processes in the queue (there's no mechanism to only wake up one of them). It then returns and the process which just closed the file can continue to run. In time, the scheduler decides that that process has had enough and gives control of the CPU to another process. Eventually, one of the processes which was in the queue will be given control of the CPU by the scheduler. It starts at the point right after the call to `module_interruptible_sleep_on`². It can then proceed to set a global variable to

¹The easiest way to keep a file open is to open it with `tail -f`.

²This means that the process is still in kernel mode — as far as the process is concerned, it issued the open

tell all the other processes that the file is still open and go on with its life. When the other processes get a piece of the CPU, they'll see that global variable and go back to sleep.

To make our life more interesting, `module_close` doesn't have a monopoly on waking up the processes which wait to access the file. A signal, such as Ctrl-C (SIGINT) can also wake up a process³. In that case, we want to return with `-EINTR` immediately. This is important so users can, for example, kill the process before it receives the file.

There is one more point to remember. Some times processes don't want to sleep, they want either to get what they want immediately, or to be told it cannot be done. Such processes use the `O_NONBLOCK` flag when opening the file. The kernel is supposed to respond by returning with the error code `-EAGAIN` from operations which would otherwise block, such as opening the file in this example. The program `cat_noblock`, available in the source directory for this chapter, can be used to open a file with `O_NONBLOCK`.

sleep.c

```
/* sleep.c - create a /proc file, and if several
 * processes try to open it at the same time, put all
 * but one to sleep */

/* Copyright (C) 1998-99 by Ori Pomerantz */

/* The necessary header files */

/* Standard in kernel modules */
#include <linux/kernel.h> /* We're doing kernel work */
#include <linux/module.h> /* Specifically, a module */

/* Deal with CONFIG_MODVERSIONS */
#if CONFIG_MODVERSIONS==1
#define MODVERSIONS
#include <linux/modversions.h>
#endif
```

system call and the system call hasn't returned yet. The process doesn't know somebody else used the CPU for most of the time between the moment it issued the call and the moment it returned.

³This is because we used `module_interruptible_sleep_on`. We could have used `module_sleep_on` instead, but that would have resulted in extremely angry users whose control C's are ignored.

```

/* Necessary because we use proc fs */
#include <linux/proc_fs.h>

/* For putting processes to sleep and waking them up */
#include <linux/sched.h>
#include <linux/wrapper.h>

/* In 2.2.3 /usr/include/linux/version.h includes a
 * macro for this, but 2.0.35 doesn't - so I add it
 * here if necessary. */
#ifndef KERNEL_VERSION
#define KERNEL_VERSION(a,b,c) ((a)*65536+(b)*256+(c))
#endif

#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
#include <asm/uaccess.h> /* for get_user and put_user */
#endif

/* The module's file functions ***** */

/* Here we keep the last message received, to prove
 * that we can process our input */
#define MESSAGE_LENGTH 80
static char Message[MESSAGE_LENGTH];

/* Since we use the file operations struct, we can't use
 * the special proc output provisions - we have to use
 * a standard read function, which is this function */
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)

```

```

static ssize_t module_output(
    struct file *file, /* The file read */
    char *buf, /* The buffer to put data to (in the
                * user segment) */
    size_t len, /* The length of the buffer */
    loff_t *offset) /* Offset in the file - ignore */
#else
static int module_output(
    struct inode *inode, /* The inode read */
    struct file *file, /* The file read */
    char *buf, /* The buffer to put data to (in the
                * user segment) */
    int len) /* The length of the buffer */
#endif
{
    static int finished = 0;
    int i;
    char message[MESSAGE_LENGTH+30];

    /* Return 0 to signify end of file - that we have
     * nothing more to say at this point. */
    if (finished) {
        finished = 0;
        return 0;
    }

    /* If you don't understand this by now, you're
     * hopeless as a kernel programmer. */
    sprintf(message, "Last input:%s\n", Message);
    for(i=0; i<len && message[i]; i++)
        put_user(message[i], buf+i);

    finished = 1;
    return i; /* Return the number of bytes "read" */
}

```

```

/* This function receives input from the user when
 * the user writes to the /proc file. */
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
static ssize_t module_input(
    struct file *file, /* The file itself */
    const char *buf, /* The buffer with input */
    size_t length, /* The buffer's length */
    loff_t *offset) /* offset to file - ignore */
#else
static int module_input(
    struct inode *inode, /* The file's inode */
    struct file *file, /* The file itself */
    const char *buf, /* The buffer with the input */
    int length) /* The buffer's length */
#endif
{
    int i;

    /* Put the input into Message, where module_output
     * will later be able to use it */
    for(i=0; i<MESSAGE_LENGTH-1 && i<length; i++)
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
        get_user(Message[i], buf+i);
#else
        Message[i] = get_user(buf+i);
#endif
    /* we want a standard, zero terminated string */
    Message[i] = '\0';

    /* We need to return the number of input
     * characters used */
    return i;
}

/* 1 if the file is currently open by somebody */

```

```

int Already_Open = 0;

/* Queue of processes who want our file */
static struct wait_queue *WaitQ = NULL;

/* Called when the /proc file is opened */
static int module_open(struct inode *inode,
                      struct file *file)
{
    /* If the file's flags include O_NONBLOCK, it means
     * the process doesn't want to wait for the file.
     * In this case, if the file is already open, we
     * should fail with -EAGAIN, meaning "you'll have to
     * try again", instead of blocking a process which
     * would rather stay awake. */
    if ((file->f_flags & O_NONBLOCK) && Already_Open)
        return -EAGAIN;

    /* This is the correct place for MOD_INC_USE_COUNT
     * because if a process is in the loop, which is
     * within the kernel module, the kernel module must
     * not be removed. */
    MOD_INC_USE_COUNT;

    /* If the file is already open, wait until it isn't */
    while (Already_Open)
    {
#ifdef LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
        int i, is_sig=0;
#endif
    }

    /* This function puts the current process,
     * including any system calls, such as us, to sleep.
     * Execution will be resumed right after the function
     * call, either because somebody called

```

```

    * wake_up(&WaitQ) (only module_close does that,
    * when the file is closed) or when a signal, such
    * as Ctrl-C, is sent to the process */
module_interruptible_sleep_on(&WaitQ);

/* If we woke up because we got a signal we're not
 * blocking, return -EINTR (fail the system call).
 * This allows processes to be killed or stopped. */

/*
 * Emmanuel Papirakis:
 *
 * This is a little update to work with 2.2.*. Signals
 * now are contained in two words (64 bits) and are
 * stored in a structure that contains an array of two
 * unsigned longs. We now have to make 2 checks in our if.
 *
 * Ori Pomerantz:
 *
 * Nobody promised me they'll never use more than 64
 * bits, or that this book won't be used for a version
 * of Linux with a word size of 16 bits. This code
 * would work in any case.
 */
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)

    for(i=0; i<_NSIG_WORDS && !is_sig; i++)
        is_sig = current->signal.sig[i] &
            ~current->blocked.sig[i];
    if (is_sig) {
#else
    if (current->signal & ~current->blocked) {
#endif
        /* It's important to put MOD_DEC_USE_COUNT here,

```

```

    * because for processes where the open is
    * interrupted there will never be a corresponding
    * close. If we don't decrement the usage count
    * here, we will be left with a positive usage
    * count which we'll have no way to bring down to
    * zero, giving us an immortal module, which can
    * only be killed by rebooting the machine. */
    MOD_DEC_USE_COUNT;
    return -EINTR;
}
}

/* If we got here, Already_Open must be zero */

/* Open the file */
Already_Open = 1;
return 0; /* Allow the access */
}

/* Called when the /proc file is closed */
#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
int module_close(struct inode *inode, struct file *file)
#else
void module_close(struct inode *inode, struct file *file)
#endif
{
    /* Set Already_Open to zero, so one of the processes
    * in the WaitQ will be able to set Already_Open back
    * to one and to open the file. All the other processes
    * will be called when Already_Open is back to one, so
    * they'll go back to sleep. */
    Already_Open = 0;

    /* Wake up all the processes in WaitQ, so if anybody

```

```

    * is waiting for the file, they can have it. */
    module_wake_up(&WaitQ);

    MOD_DEC_USE_COUNT;

#if LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
    return 0; /* success */
#endif
}

/* This function decides whether to allow an operation
 * (return zero) or not allow it (return a non-zero
 * which indicates why it is not allowed).
 *
 * The operation can be one of the following values:
 * 0 - Execute (run the "file" - meaningless in our case)
 * 2 - Write (input to the kernel module)
 * 4 - Read (output from the kernel module)
 *
 * This is the real function that checks file
 * permissions. The permissions returned by ls -l are
 * for referece only, and can be overridden here.
 */
static int module_permission(struct inode *inode, int op)
{
    /* We allow everybody to read from our module, but
     * only root (uid 0) may write to it */
    if (op == 4 || (op == 2 && current->euid == 0))
        return 0;

    /* If it's anything else, access is denied */
    return -EACCES;
}

```

```
/* Structures to register as the /proc file, with
 * pointers to all the relevant functions. ***** */
```

```
/* File operations for our proc file. This is where
 * we place pointers to all the functions called when
 * somebody tries to do something to our file. NULL
 * means we don't want to deal with something. */
```

```
static struct file_operations File_Ops_4_Our_Proc_File =
{
    NULL, /* lseek */
    module_output, /* "read" from the file */
    module_input, /* "write" to the file */
    NULL, /* readdir */
    NULL, /* select */
    NULL, /* ioctl */
    NULL, /* mmap */
    module_open, /* called when the /proc file is opened */
#ifdef LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
    NULL, /* flush */
#endif
    module_close /* called when it's closed */
};
```

```
/* Inode operations for our proc file. We need it so
 * we'll have somewhere to specify the file operations
 * structure we want to use, and the function we use for
 * permissions. It's also possible to specify functions
 * to be called for anything else which could be done to an
 * inode (although we don't bother, we just put NULL). */
```

```
static struct inode_operations Inode_Ops_4_Our_Proc_File =
{
```

```

    &File_Ops_4_Our_Proc_File,
    NULL, /* create */
    NULL, /* lookup */
    NULL, /* link */
    NULL, /* unlink */
    NULL, /* symlink */
    NULL, /* mkdir */
    NULL, /* rmdir */
    NULL, /* mknod */
    NULL, /* rename */
    NULL, /* readlink */
    NULL, /* follow_link */
    NULL, /* readpage */
    NULL, /* writepage */
    NULL, /* bmap */
    NULL, /* truncate */
    module_permission /* check for permissions */
};

/* Directory entry */
static struct proc_dir_entry Our_Proc_File =
{
    0, /* Inode number - ignore, it will be filled by
        * proc_register[_dynamic] */
    5, /* Length of the file name */
    "sleep", /* The file name */
    S_IFREG | S_IRUGO | S_IWUSR,
    /* File mode - this is a regular file which
     * can be read by its owner, its group, and everybody
     * else. Also, its owner can write to it.
     *
     * Actually, this field is just for reference, it's
     * module_permission that does the actual check. It
     * could use this field, but in our implementation it
     * doesn't, for simplicity. */
    1, /* Number of links (directories where the

```

```

        * file is referenced) */
0, 0, /* The uid and gid for the file - we give
      * it to root */
80, /* The size of the file reported by ls. */
&Inode_Ops_4_Our_Proc_File,
/* A pointer to the inode structure for
 * the file, if we need it. In our case we
 * do, because we need a write function. */
NULL /* The read function for the file.
      * Irrelevant, because we put it
      * in the inode structure above */
};

/* Module initialization and cleanup ***** */

/* Initialize the module - register the proc file */
int init_module()
{
    /* Success if proc_register_dynamic is a success,
     * failure otherwise */
#ifdef LINUX_VERSION_CODE >= KERNEL_VERSION(2,2,0)
    return proc_register(&proc_root, &Our_Proc_File);
#else
    return proc_register_dynamic(&proc_root, &Our_Proc_File);
#endif

    /* proc_root is the root directory for the proc
     * fs (/proc). This is where we want our file to be
     * located.
     */
}

```

```
/* Cleanup - unregister our file from /proc. This could
 * get dangerous if there are still processes waiting in
 * WaitQ, because they are inside our open function,
 * which will get unloaded. I'll explain how to avoid
 * removal of a kernel module in such a case in
 * chapter 10. */
void cleanup_module()
{
    proc_unregister(&proc_root, Our_Proc_File.low_ino);
}
```

Chapter 9

Replacing `printk`'s

In the beginning (chapter 1), I said that X and kernel module programming don't mix. That's true while developing the kernel module, but in actual use you want to be able to send messages to whichever tty¹ the command to the module came from. This is important for identifying errors after the kernel module is released, because it will be used through all of them.

The way this is done is by using `current`, a pointer to the currently running task, to get the current task's tty structure. Then, we look inside that tty structure to find a pointer to a string write function, which we use to write a string to the tty.

`printk.c`

```
/* printk.c - send textual output to the tty you're
 * running on, regardless of whether it's passed
 * through X11, telnet, etc. */
```

```
/* Copyright (C) 1998 by Ori Pomerantz */
```

```
/* The necessary header files */
```

¹Teletype, originally a combination keyboard–printer used to communicate with a Unix system, and today an abstraction for the text stream used for a Unix program, whether it's a physical terminal, an xterm on an X display, a network connection used with telnet, etc.

```

/* Standard in kernel modules */
#include <linux/kernel.h> /* We're doing kernel work */
#include <linux/module.h> /* Specifically, a module */

/* Deal with CONFIG_MODVERSIONS */
#if CONFIG_MODVERSIONS==1
#define MODVERSIONS
#include <linux/modversions.h>
#endif

/* Necessary here */
#include <linux/sched.h> /* For current */
#include <linux/tty.h> /* For the tty declarations */

/* Print the string to the appropriate tty, the one
 * the current task uses */
void print_string(char *str)
{
    struct tty_struct *my_tty;

    /* The tty for the current task */
    my_tty = current->tty;

    /* If my_tty is NULL, it means that the current task
     * has no tty you can print to (this is possible, for
     * example, if it's a daemon). In this case, there's
     * nothing we can do. */
    if (my_tty != NULL) {

        /* my_tty->driver is a struct which holds the tty's
         * functions, one of which (write) is used to
         * write strings to the tty. It can be used to take
         * a string either from the user's memory segment
         * or the kernel's memory segment.
         */
    }
}

```

```
* The function's first parameter is the tty to
* write to, because the same function would
* normally be used for all tty's of a certain type.
* The second parameter controls whether the
* function receives a string from kernel memory
* (false, 0) or from user memory (true, non zero).
* The third parameter is a pointer to a string,
* and the fourth parameter is the length of
* the string.
*/
(*(my_tty->driver).write)(
    my_tty, /* The tty itself */
    0, /* We don't take the string from user space */
str, /* String */
strlen(str)); /* Length */
```

```
/* ttys were originally hardware devices, which
* (usually) adhered strictly to the ASCII standard.
* According to ASCII, to move to a new line you
* need two characters, a carriage return and a
* line feed. In Unix, on the other hand, the
* ASCII line feed is used for both purposes - so
* we can't just use \n, because it wouldn't have
* a carriage return and the next line will
* start at the column right
*
*         after the line feed.
*
* BTW, this is the reason why the text file
* is different between Unix and Windows.
* In CP/M and its derivatives, such as MS-DOS and
* Windows, the ASCII standard was strictly
* adhered to, and therefore a new line requires
* both a line feed and a carriage return.
*/
```

```
*(my_tty->driver).write)(
    my_tty,
```

```
    0,  
    "\015\012",  
    2);  
}  
}
```

```
/* Module initialization and cleanup ***** */
```

```
/* Initialize the module - register the proc file */
```

```
int init_module()  
{  
    print_string("Module Inserted");  
  
    return 0;  
}
```

```
/* Cleanup - unregister our file from /proc */
```

```
void cleanup_module()  
{  
    print_string("Module Removed");  
}
```

Chapter 10

Scheduling Tasks

Very often, we have ‘housekeeping’ tasks which have to be done at a certain time, or every so often. If the task is to be done by a process, we do it by putting it in the `crontab` file. If the task is to be done by a kernel module, we have two possibilities. The first is to put a process in the `crontab` file which will wake up the module by a system call when necessary, for example by opening a file. This is terribly inefficient, however — we run a new process off of `crontab`, read a new executable to memory, and all this just to wake up a kernel module which is in memory anyway.

Instead of doing that, we can create a function that will be called once for every timer interrupt. The way we do this is we create a task, held in a `struct tq_struct`, which will hold a pointer to the function. Then, we use `queue_task` to put that task on a task list called `tq_timer`, which is the list of tasks to be executed on the next timer interrupt. Because we want the function to keep on being executed, we need to put it back on `tq_timer` whenever it is called, for the next timer interrupt.

There’s one more point we need to remember here. When a module is removed by `rmmmod`, first its reference count is checked. If it is zero, `module_cleanup` is called. Then, the module is removed from memory with all its functions. Nobody checks to see if the timer’s task list happens to contain a pointer to one of those functions, which will no longer be available. Ages later (from the computer’s perspective, from a human perspective it’s nothing, less than a hundredth of a second), the kernel has a timer interrupt and tries to call the function on the task list. Unfortunately, the function is no longer there. In most cases, the memory page where it sat is unused, and you get an ugly error message. But if some other code is now sitting at the same memory location, things could get **very** ugly.

Unfortunately, we don't have an easy way to unregister a task from a task list.

Since `cleanup_module` can't return with an error code (it's a void function), the solution is to not let it return at all. Instead, it calls `sleep_on` or `module_sleep_on`¹ to put the `rmmmod` process to sleep. Before that, it informs the function called on the timer interrupt to stop attaching itself by setting a global variable. Then, on the next timer interrupt, the `rmmmod` process will be woken up, when our function is no longer in the queue and it's safe to remove the module.

sched.c

```
/* sched.c - schedule a function to be called on
 * every timer interrupt. */

/* Copyright (C) 1998 by Ori Pomerantz */

/* The necessary header files */

/* Standard in kernel modules */
#include <linux/kernel.h> /* We're doing kernel work */
#include <linux/module.h> /* Specifically, a module */

/* Deal with CONFIG_MODVERSIONS */
#if CONFIG_MODVERSIONS==1
#define MODVERSIONS
#include <linux/modversions.h>
#endif

/* Necessary because we use the proc fs */
#include <linux/proc_fs.h>

/* We schedule tasks here */
#include <linux/tqueue.h>
```

¹They're really the same.

```

/* We also need the ability to put ourselves to sleep
 * and wake up later */
#include <linux/sched.h>

/* In 2.2.3 /usr/include/linux/version.h includes a
 * macro for this, but 2.0.35 doesn't - so I add it
 * here if necessary. */
#ifndef KERNEL_VERSION
#define KERNEL_VERSION(a,b,c) ((a)*65536+(b)*256+(c))
#endif

/* The number of times the timer interrupt has been
 * called so far */
static int TimerIntrpt = 0;

/* This is used by cleanup, to prevent the module from
 * being unloaded while intrpt_routine is still in
 * the task queue */
static struct wait_queue *WaitQ = NULL;

static void intrpt_routine(void *);

/* The task queue structure for this task, from tqueue.h */
static struct tq_struct Task = {
    NULL,    /* Next item in list - queue_task will do
              * this for us */
    0,      /* A flag meaning we haven't been inserted
              * into a task queue yet */
    intrpt_routine, /* The function to run */
    NULL    /* The void* parameter for that function */
};

```

```

/* This function will be called on every timer
 * interrupt. Notice the void* pointer - task functions
 * can be used for more than one purpose, each time
 * getting a different parameter. */
static void intrpt_routine(void *irrelevant)
{
    /* Increment the counter */
    TimerIntrpt++;

    /* If cleanup wants us to die */
    if (WaitQ != NULL)
        wake_up(&WaitQ); /* Now cleanup_module can return */
    else
        /* Put ourselves back in the task queue */
        queue_task(&Task, &tq_timer);
}

```

```

/* Put data into the proc fs file. */
int procfile_read(char *buffer,
                  char **buffer_location, off_t offset,
                  int buffer_length, int zero)
{
    int len; /* The number of bytes actually used */

    /* This is static so it will still be in memory
     * when we leave this function */
    static char my_buffer[80];

    static int count = 1;

    /* We give all of our information in one go, so if

```

```

    * the anybody asks us if we have more information
    * the answer should always be no.
    */
if (offset > 0)
    return 0;

/* Fill the buffer and get its length */
len = sprintf(my_buffer,
              "Timer was called %d times so far\n",
              TimerIntrpt);
count++;

/* Tell the function which called us where the
 * buffer is */
*buffer_location = my_buffer;

/* Return the length */
return len;
}

```

```

struct proc_dir_entry Our_Proc_File =
{
    0, /* Inode number - ignore, it will be filled by
        * proc_register_dynamic */
    5, /* Length of the file name */
    "sched", /* The file name */
    S_IFREG | S_IRUGO,
    /* File mode - this is a regular file which can
     * be read by its owner, its group, and everybody
     * else */
    1, /* Number of links (directories where
        * the file is referenced) */
    0, 0, /* The uid and gid for the file - we give
        * it to root */
    80, /* The size of the file reported by ls. */

```

```

NULL, /* functions which can be done on the
      * inode (linking, removing, etc.) - we don't
      * support any. */
procfile_read,
/* The read function for this file, the function called
 * when somebody tries to read something from it. */
NULL
/* We could have here a function to fill the
 * file's inode, to enable us to play with
 * permissions, ownership, etc. */
};

/* Initialize the module - register the proc file */
int init_module()
{
    /* Put the task in the tq_timer task queue, so it
     * will be executed at next timer interrupt */
    queue_task(&Task, &tq_timer);

    /* Success if proc_register_dynamic is a success,
     * failure otherwise */
#ifdef LINUX_VERSION_CODE > KERNEL_VERSION(2,2,0)
    return proc_register(&proc_root, &Our_Proc_File);
#else
    return proc_register_dynamic(&proc_root, &Our_Proc_File);
#endif
}

/* Cleanup */
void cleanup_module()
{
    /* Unregister our /proc file */
    proc_unregister(&proc_root, Our_Proc_File.low_ino);
}

```

```
/* Sleep until intrpt_routine is called one last
 * time. This is necessary, because otherwise we'll
 * deallocate the memory holding intrpt_routine and
 * Task while tq_timer still references them.
 * Notice that here we don't allow signals to
 * interrupt us.
 *
 * Since WaitQ is now not NULL, this automatically
 * tells the interrupt routine it's time to die. */
sleep_on(&WaitQ);
}
```

Chapter 11

Interrupt Handlers

Except for the last chapter, everything we did in the kernel so far we've done as a response to a process asking for it, either by dealing with a special file, sending an `ioctl`, or issuing a system call. But the job of the kernel isn't just to respond to process requests. Another job, which is every bit as important, is to speak to the hardware connected to the machine.

There are two types of interaction between the CPU and the rest of the computer's hardware. The first type is when the CPU gives orders to the hardware, the other is when the hardware needs to tell the CPU something. The second, called interrupts, is much harder to implement because it has to be dealt with when convenient for the hardware, not the CPU. Hardware devices typically have a very small amount of ram, and if you don't read their information when available, it is lost.

Under Linux, hardware interrupts are called IRQs (short for **I**nterrupt **R**equests)¹. There are two types of IRQs, short and long. A short IRQ is one which is expected to take a **very** short period of time, during which the rest of the machine will be blocked and no other interrupts will be handled. A long IRQ is one which can take longer, and during which other interrupts may occur (but not interrupts from the same device). If at all possible, it's better to declare an interrupt handler to be long.

When the CPU receives an interrupt, it stops whatever it's doing (unless it's processing a more important interrupt, in which case it will deal with this one only when the more important one is done), saves certain parameters on the stack and calls the interrupt handler. This means that certain things are not allowed in the interrupt handler itself, because the

¹This is standard nomenclature on the Intel architecture where Linux originated.

system is in an unknown state. The solution to this problem is for the interrupt handler to do what needs to be done immediately, usually read something from the hardware or send something to the hardware, and then schedule the handling of the new information at a later time (this is called the ‘bottom half’) and return. The kernel is then guaranteed to call the bottom half as soon as possible — and when it does, everything allowed in kernel modules will be allowed.

The way to implement this is to call `request_irq` to get your interrupt handler called when the relevant IRQ is received (there are 16 of them on Intel platforms). This function receives the IRQ number, the name of the function, flags, a name for `/proc/interrupts` and a parameter to pass to the interrupt handler. The flags can include `SA_SHIRQ` to indicate you’re willing to share the IRQ with other interrupt handlers (usually because a number of hardware devices sit on the same IRQ) and `SA_INTERRUPT` to indicate this is a fast interrupt. This function will only succeed if there isn’t already a handler on this IRQ, or if you’re both willing to share.

Then, from within the interrupt handler, we communicate with the hardware and then use `queue_task_irq` with `tq_immediate` and `mark_bh(BH_IMMEDIATE)` to schedule the bottom half. The reason we can’t use the standard `queue_task` in version 2.0 is that the interrupt might happen right in the middle of somebody else’s `queue_task`². We need `mark_bh` because earlier versions of Linux only had an array of 32 bottom halves, and now one of them (`BH_IMMEDIATE`) is used for the linked list of bottom halves for drivers which didn’t get a bottom half entry assigned to them.

11.1 Keyboards on the Intel Architecture

Warning: The rest of this chapter is completely Intel specific. If you’re not running on an Intel platform, it will not work. Don’t even try to compile the code here.

I had a problem with writing the sample code for this chapter. On one hand, for an example to be useful it has to run on everybody’s computer with meaningful results. On the other hand, the kernel already includes device drivers for all of the common devices, and those device drivers won’t coexist with what I’m going to write. The solution I’ve found was to write something for the keyboard interrupt, and disable the regular keyboard interrupt handler first. Since it is defined as a static symbol in the kernel source files (specifically, `drivers/char/keyboard.c`), there is no way to restore it. Before `insmod`’ing this code, do on another terminal `sleep 120 ; reboot` if you value your file system.

²`queue_task_irq` is protected from this by a global lock — in 2.2 there is no `queue_task_irq` and `queue_task` is protected by a lock.

This code binds itself to IRQ 1, which is the IRQ of the keyboard controlled under Intel architectures. Then, when it receives a keyboard interrupt, it reads the keyboard's status (that's the purpose of the `inb(0x64)`) and the scan code, which is the value returned by the keyboard. Then, as soon as the kernel think it's feasible, it runs `got_char` which gives the code of the key used (the first seven bits of the scan code) and whether it has been pressed (if the 8th bit is zero) or released (if it's one).

intrpt.c

```
/* intrpt.c - An interrupt handler. */

/* Copyright (C) 1998 by Ori Pomerantz */

/* The necessary header files */

/* Standard in kernel modules */
#include <linux/kernel.h> /* We're doing kernel work */
#include <linux/module.h> /* Specifically, a module */

/* Deal with CONFIG_MODVERSIONS */
#if CONFIG_MODVERSIONS==1
#define MODVERSIONS
#include <linux/modversions.h>
#endif

#include <linux/sched.h>
#include <linux/tqueue.h>

/* We want an interrupt */
#include <linux/interrupt.h>

#include <asm/io.h>
```

```

/* In 2.2.3 /usr/include/linux/version.h includes a
 * macro for this, but 2.0.35 doesn't - so I add it
 * here if necessary. */
#ifndef KERNEL_VERSION
#define KERNEL_VERSION(a,b,c) ((a)*65536+(b)*256+(c))
#endif

/* Bottom Half - this will get called by the kernel
 * as soon as it's safe to do everything normally
 * allowed by kernel modules. */
static void got_char(void *scancode)
{
    printk("Scan Code %x %s.\n",
        (int) *((char *) scancode) & 0x7F,
        *((char *) scancode) & 0x80 ? "Released" : "Pressed");
}

/* This function services keyboard interrupts. It reads
 * the relevant information from the keyboard and then
 * schedules the bottom half to run when the kernel
 * considers it safe. */
void irq_handler(int irq,
                void *dev_id,
                struct pt_regs *regs)
{
    /* This variables are static because they need to be
     * accessible (through pointers) to the bottom
     * half routine. */
    static unsigned char scancode;
    static struct tq_struct task =
        {NULL, 0, got_char, &scancode};
    unsigned char status;

```

```

    /* Read keyboard status */
    status = inb(0x64);
    scancode = inb(0x60);

    /* Schedule bottom half to run */
#if LINUX_VERSION_CODE > KERNEL_VERSION(2,2,0)
    queue_task(&task, &tq_immediate);
#else
    queue_task_irq(&task, &tq_immediate);
#endif
    mark_bh(IMMEDIATE_BH);
}

/* Initialize the module - register the IRQ handler */
int init_module()
{
    /* Since the keyboard handler won't co-exist with
     * another handler, such as us, we have to disable
     * it (free its IRQ) before we do anything. Since we
     * don't know where it is, there's no way to
     * reinstate it later - so the computer will have to
     * be rebooted when we're done.
     */
    free_irq(1, NULL);

    /* Request IRQ 1, the keyboard IRQ, to go to our
     * irq_handler. */
    return request_irq(
        1, /* The number of the keyboard IRQ on PCs */
        irq_handler, /* our handler */
        SA_SHIRQ,
        /* SA_SHIRQ means we're willing to have othe
         * handlers on this IRQ.
         */

```

```
    * SA_INTERRUPT can be used to make the
    * handler into a fast interrupt.
    */
    "test_keyboard_irq_handler", NULL);
}

/* Cleanup */
void cleanup_module()
{
    /* This is only here for completeness. It's totally
    * irrelevant, since we don't have a way to restore
    * the normal keyboard interrupt so the computer
    * is completely useless and has to be rebooted. */
    free_irq(1, NULL);
}
```


Chapter 12

Symmetrical Multi-Processing

One of the easiest (read, cheapest) ways to improve hardware performance is to put more than one CPU on the board. This can be done either making the different CPUs take on different jobs (asymmetrical multi-processing) or by making them all run in parallel, doing the same job (symmetrical multi-processing, a.k.a. SMP). Doing asymmetrical multi-processing effectively requires specialized knowledge about the tasks the computer should do, which is unavailable in a general purpose operating system such as Linux. On the other hand, symmetrical multi-processing is relatively easy to implement.

By relatively easy, I mean exactly that — not that it's *really* easy. In a symmetrical multi-processing environment, the CPUs share the same memory, and as a result code running in one CPU can affect the memory used by another. You can no longer be certain that a variable you've set to a certain value in the previous line still has that value — the other CPU might have played with it while you weren't looking. Obviously, it's impossible to program like this.

In the case of process programming this normally isn't an issue, because a process will normally only run on one CPU at a time¹. The kernel, on the other hand, could be called by different processes running on different CPUs.

In version 2.0.x, this isn't a problem because the entire kernel is in one big spinlock. This means that if one CPU is in the kernel and another CPU wants to get in, for example because of a system call, it has to wait until the first CPU is done. This makes Linux SMP safe², but terribly inefficient.

¹The exception is threaded processes, which can run on several CPUs at once.

²Meaning it is safe to use it with SMP

In version 2.2.x, several CPUs can be in the kernel at the same time. This is something module writers need to be aware of. I got somebody to give me access to an SMP box, so hopefully the next version of this book will include more information.

Chapter 13

Common Pitfalls

Before I send you on your way to go out into the world and write kernel modules, there are a few things I need to warn you about. If I fail to warn you and something bad happens, please report the problem to me for a full refund of the amount I got paid for your copy of the book.

1. **Using standard libraries** You can't do that. In a kernel module you can only use kernel functions, which are the functions you can see in `/proc/ksyms`.
2. **Disabling interrupts** You might need to do this for a short time and that is OK, but if you don't enable them afterwards, your system will be stuck and you'll have to power it off.
3. **Sticking your head inside a large carnivore** I probably don't have to warn you about this, but I figured I will anyway, just in case.

Appendix A

Changes between 2.0 and 2.2

I don't know the entire kernel well enough to document all of the changes. In the course of converting the examples (or actually, adapting Emmanuel Papirakis's changes) I came across the following differences. I listed all of them here together to help module programmers, especially those who learned from previous versions of this book and are most familiar with the techniques I use, convert to the new version.

An additional resource for people who wish to convert to 2.2 is in <http://www.atnf.csiro.au/~rgooch/linux/docs/porting-to-2.2.html>.

1. **asm/uaccess.h** If you need `put_user` or `get_user` you have to `#include` it.
2. **get_user** In version 2.2, `get_user` receives both the pointer into user memory and the variable in kernel memory to fill with the information. The reason for this is that `get_user` can now read two or four bytes at a time if the variable we read is two or four bytes long.
3. **file_operations** This structure now has a `flush` function between the `open` and `close` functions.
4. **close in file_operations** In version 2.2, the `close` function returns an integer, so it's allowed to fail.
5. **read and write in file_operations** The headers for these functions changed. They now return `ssize_t` instead of an integer, and their parameter list is different. The `inode` is no longer a parameter, and on the other hand the offset into the file is.

6. **proc_register_dynamic** This function no longer exists. Instead, you call the regular `proc_register` and put zero in the `inode` field of the structure.
7. **Signals** The signals in the task structure are no longer a 32 bit integer, but an array of `_NSIG_WORDS` integers.
8. **queue_task_irq** Even if you want to schedule a task to happen from inside an interrupt handler, you use `queue_task`, not `queue_task_irq`.
9. **Module Parameters** You no longer just declare module parameters as global variables. In 2.2 you have to also use `MODULE_PARM` to declare their type. This is a big improvement, because it allows the module to receive string parameters which start with a digit, for example, without getting confused.
10. **Symmetrical Multi-Processing** The kernel is no longer inside one huge spinlock, which means that kernel modules have to be aware of SMP.

Appendix B

Where From Here?

I could easily have squeezed a few more chapters into this book. I could have added a chapter about creating new file systems, or about adding new protocols stacks (as if there's a need for that — you'd have to dig under ground to find a protocol stack not supported by Linux). I could have added explanations of the kernel mechanisms we haven't touched upon, such as bootstrapping or the disk interface.

However, I chose not to. My purpose in writing this book was to provide initiation into the mysteries of kernel module programming and to teach the common techniques for that purpose. For people seriously interested in kernel programming, I recommend the list of kernel resources in <http://jungla.dit.upm.es/~jmseyas/linux/kernel/hackers-docs.html>. Also, as Linus said, the best way to learn the kernel is to read the source code yourself.

If you're interested in more examples of short kernel modules, I recommend Phrack magazine. Even if you're not interested in security, and as a programmer you should be, the kernel modules there are good examples of what you can do inside the kernel, and they're short enough not to require too much effort to understand.

I hope I have helped you in your quest to become a better programmer, or at least to have fun through technology. And, if you do write useful kernel modules, I hope you publish them under the GPL, so I can use them too.

Appendix C

Goods and Services

I hope nobody minds the shameless promotions here. They are all things which are likely to be of use to beginning Linux Kernel Module programmers.

C.1 Getting this Book in Print

The Coriolis group is going to print this book sometimes in the summer of '99. If this is already summer, and you want this book in print, you can go easy on your printer and buy it in a nice, bound form.

Appendix D

Showing Your Appreciation

This is a free document. You have no obligations beyond those given in the GNU Public License (Appendix E). However, if you want to do something in return for getting this book, there are a few things you could do.

- Send me a postcard to

Ori Pomerantz
Apt. #1032
2355 N Hwy 360
Grand Prairie
TX 75050
USA

If you want to receive a thank-you from me, include your e-mail address.

- Contribute money, or better yet, time, to the free software community. Write a program or a document and publish it under the GPL. Teach other people how to use free software, such as Linux or Perl.
- Explain to people how being selfish is *not* incompatible with living in a society or with helping other people. I enjoyed writing this document, and I believe publishing it will contribute to me in the future. At the same time, I wrote a book which, if you got this far, helps you. Remember that happy people are usually more useful to oneself than unhappy people, and able people are *way* better than people of low ability.

- **Be happy.** If I get to meet you, it will make the encounter better for me, it will make you more useful for me ;-).

Appendix E

The GNU General Public License

Printed below is the GNU General Public License (the *GPL* or *copyleft*), under which this book is licensed.

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright ©1989, 1991 Free Software Foundation, Inc. 675 Mass Ave, Cambridge, MA 02139, USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

PREAMBLE

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation’s software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain

responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The 'Program', below, refers to any such program or work, and a 'work based on the Program' means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term 'modification'.) Each licensee is addressed as 'you'.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact

all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
 - c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under

this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or

she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and 'any later version', you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM 'AS IS' WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR

PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

APPENDIX: HOW TO APPLY THESE TERMS TO YOUR NEW PROGRAMS

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the 'copyright' line and a pointer to where the full notice is found.

one line to give the program's name and a brief idea of what it does. Copyright ©19yy name of author

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) 19yy name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for
details type show w. This is free software, and you are
welcome to redistribute it under certain conditions; type
show c for details.
```

The hypothetical commands `show w` and `show c` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w` and `show c`; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a ‘copyright disclaimer’ for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program Gnomovision (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

Index

- /dev, 12, 13
- /proc file system, 25
- /proc/interrupts, 98
- /proc/ksyms, 106
- /proc/meminfo, 25
- /proc/modules, 8, 14, 25
- /proc
 - using for input, 32
- _IO, 44
- _IOR, 44
- _IOW, 44
- _IOWR, 44
- _NSIG_WORDS, 108
- __KERNEL__, 7
- __NO_VERSION__, 8
- __SMP__, 7
- 2.0.x kernel, 24
- 2.2 changes, 107
- 2.2.x kernel, 24
- access
 - sequential, 13
- argc, 61
- argv, 61
- asm/uaccess.h, 107
- BH_IMMEDIATE, 98
- blocking processes, 73
- blocking, how to avoid, 74
- bottom half, 98
- busy, 73
- calls
 - system, 65
- character device files, 12
- chardev.c, source file, 14, 44
- chardev.h, source file, 55
- cleanup_module, 5, 14
- cleanup_module
 - general purpose, 13
- close, 107
- compilation
 - conditional, 24
- compiling, 6
- conditional compilation, 24
- config.h, 7
- CONFIG_MODVERSIONS, 7
- configuration
 - kernel, 7
- console, 8
- copying Linux, 120
- copyright, 113–120
- CPU
 - multiple, 104
- crontab, 90
- ctrl-c, 74
- current pointer, 33
- current task, 86

- defining ioctl, 57
- development version
 - kernel, 23
- device files
 - block, 13
- device files
 - character, 12, 13
- device files
 - input to, 43
- device number
 - major, 13
- devices
 - physical, 12
- DOS, 2
- EAGAIN, 74
- EINTR, 74
- elf_i386, 9
- ENTRY(system_call), 66
- entry.S, 66
- file system registration, 32
- file system
 - /proc, 25
- file_operations structure, 13, 32
- file_operations
 - structure, 107
- flush, 107
- Free Software Foundation, 113
- General Public License, 113–120
- get_user, 33, 107
- GNU
 - General Public License, 113–120
- handlers
 - interrupt, 97
- hard disk
 - partitions of, 12
- hard wiring, 61
- header file for ioctl, 57
- hello world, 5
- hello.c, source file, 5
- housekeeping, 90
- IDE
 - hard disk, 12
- inb, 99
- init_module, 5
- init_module
 - general purpose, 13
- inode, 25
- inode_operations structure, 32
- input to device files, 43
- Input
 - using /proc for, 32
- insmod, 8, 61, 65
- intel architecture
 - keyboard, 98
- interrupt 0x80, 66
- interrupt handlers, 97
- interruptible_sleep_on, 73
- interrupts, 108
- interrupts
 - disabling, 106
- intrpt.c, source file, 99
- ioctl, 43
- ioctl.c, source file, 57
- ioctl
 - defining, 57
- ioctl
 - header file for, 57
- ioctl
 - official assignment, 44

- ioctl
 - using in a process, 60
- irqs, 108
- kernel configuration, 7
- kernel versions, 23
- KERNEL_VERSION, 24
- kernel_version, 8
- keyboard, 98
- ksyms
 - proc file, 106
- ld, 9
- libraries
 - standard, 106
- LINUX, 7
- Linux
 - copyright, 120
- LINUX_VERSION_CODE, 24
- MACRO_PARM, 61
- major device number, 13
- major number, 12
- makefile, 6
- Makefile, source file, 7, 11
- mark_bh, 98
- memory segments, 33
- minor number, 12
- mknod, 13
- MOD_DEC_USE_COUNT, 14
- MOD_INC_USE_COUNT, 14, 67
- mod_use_count_, 14
- modem, 12, 43
- MODULE, 7
- Module Parameters, 108
- module.h, 8
- module_cleanup, 91
- module_interruptible_sleep_on, 73
- MODULE_PARM, 108
- module_permissions, 33
- module_register_chrdev, 13
- module_sleep_on, 74, 91
- module_wake_up, 74
- modversions.h, 7
- multi tasking, 73
- multi-processing, 104
- multiple source files, 8
- multitasking, 74
- non blocking, 74
- number
 - major (of device driver), 12
- number
 - major (of physical device), 12
- O_NONBLOCK, 74
- official ioctl assignment, 44
- open
 - system call, 66
- param.c, source file, 61
- Parameters
 - Module, 108
- parameters
 - startup, 61
- partition
 - of hard disk, 12
- permissions, 33
- physical devices, 12
- pointer
 - current, 33
- printk, 8
- printk.c, source file, 86

- printk
 - replacing, 86
- proc file system, 25
- proc
 - using for input, 32
- proc_dir_entry structure, 32
- proc_register, 25, 108
- proc_register_dynamic, 25, 108
- processes
 - blocking, 73
- processes
 - killing, 74
- processes
 - putting to sleep, 73
- processes
 - waking up, 74
- processing
 - multi, 104
- procfs.c, source file, 26, 33
- put_user, 33, 107
- putting processes to sleep, 73
- queue_task, 90, 98, 108
- queue_task_irq, 98, 108
- read, 107
- read
 - in the kernel, 33
- reference count, 14, 91
- refund policy, 106
- registration
 - file system, 32
- replacing printk's, 86
- request_irq, 98
- rmmod, 8, 65, 67, 91
- rmmod
 - preventing, 14
- root, 8
- SA_INTERRUPT, 98
- SA_SHIRQ, 98
- salut mundi, 5
- sched.c, source file, 91
- scheduler, 74
- scheduling tasks, 90
- segment
 - memory, 33
- selfishness, 111
- sequential access, 13
- serial port, 43
- shutdown, 65
- SIGINT, 74
- signal, 74
- signals, 108
- sleep.c, source file, 74
- sleep
 - putting processes to, 73
- sleep_on, 74, 91
- SMP, 104, 108
- source files
 - multiple, 8
- source
 - chardev.c, 14, 44
 - chardev.h, 55
 - hello.c, 5
 - intrpt.c, 99
 - ioctl.c, 57
 - Makefile, 7, 11

- source
 - param.c, 61
- source
 - printk.c, 86
- source
 - procf.c, 26, 33
- source
 - sched.c, 91
- source
 - sleep.c, 74
- source
 - start.c, 9
- source
 - stop.c, 10
- source
 - syscall.c, 67
- ssize_t, 107
- stable version
 - kernel, 23
- standard libraries, 106
- start.c, source file, 9
- startup parameters, 61
- stop.c, source file, 10
- strace, 65
- struct file_operations, 13, 32
- struct inode_operations, 32
- struct proc_dir_entry, 32
- struct tq_struct, 90
- struct
 - tty, 86
- structure
 - task, 73
- Symmetrical Multi-Processing, 108
- symmetrical multi-processing, 104
- sync, 65
- sys_call_table, 66
- sys_open, 67
- syscall.c, source file, 67
- system calls, 65
- system_call, 66
- task, 90
- task structure, 73
- task
 - current, 86
- TASK_INTERRUPTIBLE, 73
- tasks
 - scheduling, 90
- terminal, 12
- terminal
 - virtual, 8
- tq_immediate, 98
- tq_struct struct, 90
- tq_timer, 90
- tty_struct, 86
- type checking, 61
- uaccess.h
 - asm, 107
- version.h, 8
- versions supported, 24
- versions
 - kernel, 107
- virtual terminal, 8
- waking up processes, 74
- write, 107
- write
 - in the kernel, 33
- write
 - to device files, 43

X

why you should avoid, 8

xterm -C, 8