

Laboratory Exercise 3

Character Device Drivers

Part I

This exercise is about character device drivers, which are kernel modules that provide a file-based interface to user-level programs. When a character device driver is inserted into the Linux kernel, a special type of file associated with the driver is created, usually in the filesystem folder */dev*. For example, if the driver is named *chardev* then the associated file would be */dev/chardev*. A user-level program can read or write to this file to communicate with the driver.

An example of code for a character device driver is given in Figure 1. This is a trivial example that is meant to illustrate how the code for a driver has to be structured. If a user reads from this driver, it simply replies with the text *Hello from chardev*. One way to read from the driver is to execute the Linux command `cat /dev/chardev`. This driver supports both reading and writing. The user can change the character string provided by the driver by writing to it. One way to write to the driver is to use the Linux command `echo "New message" > /dev/chardev`.

The driver is created by declaring a number of variables and data structures, and calling several functions that interact with the Linux kernel. Important lines of code in Figure 1 are described below. Lines 1 to 7 include some header files that are required for character device drivers. Lines 10 to 13 provide prototype declarations for the functions `device_open`, `device_release`, `device_read`, and `device_write`. These are the functions that will be called by the Linux kernel when a user program performs an open, close, read, or write operation, respectively, on the file */dev/chardev*.

Lines 18 to 20 declare some special types of variables. All character device drivers must declare these variables, which are of type `dev_t`, `cdev`, and `class`. They are initialized in lines 32 to 56, in the function `start_chardev`, which calls a number of kernel functions. The first step to initialize these variables, shown in line 36, is to obtain a *device number* by calling the kernel function `alloc_chrdev_region`. Then, the class structure is set up in line 40 by calling `class_create` and the `cdev` structure is initialized in line 43 by calling the `cdev_alloc` function. Line 44 initializes the `cdev` structure with a pointer to the `fops` structure, which specifies the names of the functions that open, close, read, and write to the character device driver. Finally, lines 48 and 52 add the `cdev` device and the class to the kernel. Line 53 initializes the `chardev_msg` array with the default text message. The `start_chardev` function is executed when the character device driver is inserted into the Linux kernel. When this driver is removed from the kernel, the function `stop_chardev` is called, shown in lines 57 to 63. Note that it is possible to print error or debug information on the Linux terminal window, by calling the `printk` kernel function. This function works similarly to the C library function `printf`, and has the same types of formatting options. Examples of `printk` are shown in Lines 37 and 49.

Lines 65 to 75 give the code for `device_open` and `device_release`. For most character device drivers nothing needs to be done in these functions, and hence the code simply *returns*.

Lines 77 to 89 show the `device_read` function. The first argument to this function, `filp`, is not used in this example. The second argument, `buffer`, is used to pass character data from the driver back to the user-level process that read from */dev/chardev*. The `length` argument specifies the maximum number of bytes that can be stored in the `buffer`. The final argument, `offset`, will be discussed later. Lines 82 to 83 determine how much data can be sent back to the kernel. If `length` is greater than `bytes` then the whole message can be sent at once. Otherwise only `length` characters can be sent. The value of `bytes` is calculated in Line 82. When a user opens the file */dev/chardev*, and then reads from the file, `*offset` will be 0. Thus, `bytes` will be set to the length of the character string that is returned when the driver is read.

```

1  #include <linux/module.h>
2  #include <linux/kernel.h>
3  #include <linux/fs.h>
4  #include <linux/cdev.h>
5  #include <linux/device.h>
6  #include <asm/io.h>
7  #include <asm/uaccess.h>
8
9  /* Kernel character device driver /dev/chardev. */
10 static int device_open (struct inode *, struct file *);
11 static int device_release (struct inode *, struct file *);
12 static ssize_t device_read (struct file *, char *, size_t, loff_t *);
13 static ssize_t device_write(struct file *, const char *, size_t, loff_t *);
14
15 #define SUCCESS 0
16 #define DEVICE_NAME "chardev"
17
18 static dev_t dev_no = 0;
19 static struct cdev *chardev_cdev = NULL;
20 static struct class *chardev_class = NULL;
21 #define MAX_SIZE 256 // assume that no message longer than this will be used
22 static char chardev_msg[MAX_SIZE]; // the string that can be read or written
23
24 static struct file_operations fops = {
25     .owner = THIS_MODULE,
26     .open = device_open,
27     .release = device_release,
28     .read = device_read,
29     .write = device_write
30 };
31
32 static int __init start_chardev(void)
33 {
34     int err = 0;
35     /* Get a device number. Get one minor number (0) */
36     if ((err = alloc_chrdev_region (&dev_no, 0, 1, DEVICE_NAME)) < 0) {
37         printk (KERN_ERR "chardev: alloc_chrdev_region() error %d\n", err);
38         return err;
39     }
40     chardev_class = class_create (THIS_MODULE, DEVICE_NAME);
41
42     // Allocate and initialize the char device
43     chardev_cdev = cdev_alloc ();
44     chardev_cdev->ops = &fops;
45     chardev_cdev->owner = THIS_MODULE;
46
47     // Add the character device to the kernel
48     if ((err = cdev_add (chardev_cdev, dev_no, 1)) < 0) {
49         printk (KERN_ERR "chardev: cdev_add() error %d\n", err);
50         return err;
51     }
52     device_create (chardev_class, NULL, dev_no, NULL, DEVICE_NAME );
53     strcpy (chardev_msg, "Hello from chardev\n");
54
55     return 0;
56 }

```

Figure 1: The character device driver code. (Part a)

```

57 static void __exit stop_chardev(void)
58 {
59     device_destroy (chardev_class, dev_no);
60     cdev_del (chardev_cdev);
61     class_destroy (chardev_class);
62     unregister_chrdev_region (dev_no, 1);
63 }
64
65 /* Called when a process opens chardev */
66 static int device_open(struct inode *inode, struct file *file)
67 {
68     return SUCCESS;
69 }
70
71 /* Called when a process closes chardev */
72 static int device_release(struct inode *inode, struct file *file)
73 {
74     return 0;
75 }
76
77 /* Called when a process reads from chardev. Provides character data from
78 * chardev_msg. Returns, and sets *offset to, the number of bytes read. */
79 static ssize_t device_read(struct file *filp, char *buffer, size_t length,
80     loff_t *offset)
81 {
82     size_t bytes;
83     bytes = strlen (chardev_msg) - (*offset);    // how many bytes not yet sent?
84     bytes = bytes > length ? length : bytes;    // too much to send at once?
85
86     if (bytes)
87         (void) copy_to_user (buffer, &chardev_msg[*offset], bytes);
88     *offset = bytes;    // keep track of number of bytes sent to the user
89     return bytes;
90 }
91
92 /* Called when a process writes to chardev. Stores the data received into
93 * chardev_msg, and returns the number of bytes stored. */
94 static ssize_t device_write(struct file *filp, const char *buffer, size_t
95     length, loff_t *offset)
96 {
97     size_t bytes;
98     bytes = length;
99
100     if (bytes > MAX_SIZE - 1)    // can copy all at once, or not?
101         bytes = MAX_SIZE - 1;
102     (void) copy_from_user (chardev_msg, buffer, bytes);
103     chardev_msg[bytes] = '\0';    // NULL terminate
104     // Note: we do NOT update *offset; we keep the last MAX_SIZE or fewer bytes
105     return bytes;
106 }
107
108 MODULE_LICENSE("GPL");
109 module_init (start_chardev);
110 module_exit (stop_chardev);

```

Figure 1. The character device driver code. (Part b)

In Line 86 the kernel function `copy_to_user` is called to return character data back to the user via the `buffer`. The kernel variable `offset` is set to the value of the variable `bytes` in line 87. This value of `offset`, associated with `/dev/chardev`, is stored in the kernel as long as the file remains open. In line 88 the `device_read` function provides as a return value the number of bytes stored into the `buffer`. A typical user-level program that reads from the device driver (e.g., `cat /dev/chardev`) will read from the file until `device_read` returns 0, which indicates *end-of-file*. The way this works in our example is as follows. The first time `device_read` is called it copies the character string back to the `buffer` and returns the *string length*. But a second call to `device_read` will copy nothing into the `buffer` and will return 0. This mechanism is facilitated by the way in which `offset` is used in the code.

Lines 91 to 104 show the `device_write` function. The first argument to this function, `filp`, is not used in this example. The second argument, `buffer`, is used to get character data from the user program into the device driver. The `length` argument specifies the amount of data that is to be transferred. The `offset` argument is not used in this example. Data transferred from the user is stored into `chardev_msg`, overwriting the previous message. Line 98 checks for overflow, so that the amount of data can be reduced if needed. In line 100 the kernel function `copy_from_user` is called to get the user-data and copy it into `chardev_msg`.

The `device_write` function provides as a return value the number of bytes copied into `chardev_msg`. A typical user-level program (e.g., `echo "New message" > /dev/chardev`) will continue to call the `device_write` function until a total of `length` bytes have been received by the driver. In our example each call to `device_write` overwrites the data stored in `chardev_msg`.

Perform the following:

1. Create a C source-code file named `chardev.c` for the device driver code in Figure 1. The source code can be obtained alongside this lab exercise on the Intel FPGA University Program website.
2. Create a Makefile for your character device driver, following the format given in the tutorial *Using Linux on the DE1-SoC*. Compile the code to create the kernel module `chardev.ko`, and insert this module into the Linux kernel.
3. Check the filesystem folder `/dev` to see that the file `/dev/chardev` was created as a result of inserting your kernel module. Type the command `cat /dev/chardev` and observe that your character device driver responds with the message `Hello from chardev`. Overwrite the default message by typing a command such as `echo "New Message" > /dev/chardev`. Then, issue the `cat` command again to see that the driver responds with the new message.
4. In addition to using commands like `cat` and `echo`, you can write your own user-level programs that read and write to your character device driver. An example program is given in Figure 2. It uses the kernel functions `open`, `read`, `write`, and `close` to communicate with the character device driver through the file `/dev/chardev`. Create a C source-code file called `part1.c` for this program, and compile it using a command such as `gcc -Wall -o part1 part1.c`. You can obtain the source code file from the Intel University Program website. Run the resulting executable program and observe the output that it produces.

Instead of using the kernel functions `open`, `read`, `write`, and `close`, an alternative is to use the C library functions `fopen`, `fread`, `fwrite`, and `fclose`. A version of the code in `part1.c` based on using `fopen` is also provided on the Intel FPGA University Program website. Whether a programmer uses the `open` or `fopen` families of functions is mostly a matter of personal preference.

5. An alternative version of the `read_device` function is given in Figure 3. Instead of calling the kernel function `copy_to_user`, this code copies one byte at a time by calling the kernel function `put_user`. Similarly, an alternative version of the `device_write` function that transfers one byte at a time instead of calling `copy_from_user` is shown in Figure 3. These alternative versions of the code are meant to provide additional examples of how device driver code can be written. The two versions of the code perform exactly the same functions.

```

#include <stdio.h>
#include <signal.h>
#include <string.h>
#include <errno.h>
#include <fcntl.h>
#include <unistd.h>

#define BYTES 256 // max # of bytes to read from /dev/chardev

volatile sig_atomic_t stop;
void catchSIGINT(int signum){
    stop = 1;
}

/* This code uses the character device driver /dev/chardev. The code reads the
 * default message from the driver and then prints it. After this the code
 * changes the message in a loop by writing to the driver, and prints each new
 * message. The program exits if it receives a kill signal (for example, ^C
 * typed on stdin). */
int main(int argc, char *argv[]){
    int chardev_FD; // file descriptor
    char chardev_buffer[BYTES]; // buffer for chardev character data
    int ret_val, chars_read; // number of characters read
    char new_msg[128]; // space for the new message
    int i_msg;

    // catch SIGINT from ^C, instead of having it abruptly close this program
    signal(SIGINT, catchSIGINT);

    // Open the character device driver for read/write
    if ((chardev_FD = open("/dev/chardev", O_RDWR)) == -1) {
        printf("Error opening /dev/chardev: %s\n", strerror(errno));
        return -1;
    }

    i_msg = 0;
    while (!stop) {
        chars_read = 0;
        while ((ret_val = read (chardev_FD, chardev_buffer, BYTES)) != 0)
            chars_read += ret_val; // read the driver until EOF
        chardev_buffer[chars_read] = '\0'; // NULL terminate
        printf ("%s", chardev_buffer);

        sprintf (new_msg, "New message %d\n", i_msg);
        i_msg++;
        write (chardev_FD, new_msg, strlen(new_msg));

        sleep (1);
    }

    close (chardev_FD);
    return 0;
}

```

Figure 2: A program that communicates with */dev/chardev*.

```

/* Called when a process reads from chardev. */
static ssize_t device_read(struct file *filp, char *buffer, size_t length,
    loff_t *offset)
{
    size_t bytes_read = 0;
    char *msg_Ptr = &(chardev_msg[*offset]);

    // Write to user buffer
    while (length && *msg_Ptr) {
        put_user>(*msg_Ptr++, buffer++);
        length--;
        bytes_read++;
    }
    (*offset) = bytes_read;
    return bytes_read;
}

/* Called when a process writes to chardev */
static ssize_t device_write(struct file *filp, const char *buffer, size_t
    length, loff_t *offset)
{
    int i;
    for (i = 0; i < length; ++i)
        chardev_msg[i] = buffer[i]; // assume that data won't overflow
    chardev_msg[i] = '\0'; // NULL terminate
    return length;
}

```

Figure 3: Alternative versions of device_read and device_write

Part II

In some Linux systems user-level programs are not permitted to access the physical memory addresses of I/O devices. In such systems I/O devices can only be accessed via device drivers. In this part you are to implement two character device drivers. One driver provides the state of the KEY pushbutton switches in the DE1-SoC Computer, via the file `/dev/KEY`. The other driver provides the state of the SW slider switches, via the file `/dev/SW`.

Perform the following:

1. Create a new kernel module in a file `KEY_SW.c`. Write the code for *both* character device drivers in this module. Declare separate variables of type `dev_t`, `cdev`, and `class` for each driver. Initialize each of these variables by following the steps shown in Figure 1. When setting up the `fops` data structure (see line 24 in Figure 1), each of your character device drivers needs to have a function for opening, releasing, and reading its `/dev/` file. These drivers do not require a function for writing, since a user would not need to write anything to `/dev/KEY` or `/dev/SW`.

Your module needs to have an initialization function. If this function were named `init_drivers`, then it would be declared using the syntax

```
static int __init init_drivers(void)
```

Remember that you have to identify the module initialization function by using the statement

```
module_init (init_drivers);
```

In `KEY_SW.c` declare global variables that will hold virtual addresses for the KEY and SW ports in the DE1-SoC Computer. Initialize these variables in `init_drivers`, using the kernel function `ioremap_nocache`, as illustrated in the tutorial *Using Linux on the DE1-SoC*. Note that you do not need to use interrupts for this part of the exercise.

Write code for the `open`, `release`, and `read` functions for your drivers. For the KEY driver you should read the state of the KEYS from the port's *Edgecapture* register. The programmer registers in the KEY port of the DE1-SoC Computer are illustrated in Figure 4. Return the KEY values to the user as *character* data (ASCII) in the `read` function for the KEY driver. One way to convert binary data into character data is to make use of a library function such as `printf`. For the SW driver, read the slider switch settings from the port's *Data* register and return these values to the user, in the form of *character* data, via the driver's `read` function. The SW port is illustrated in Figure 5.

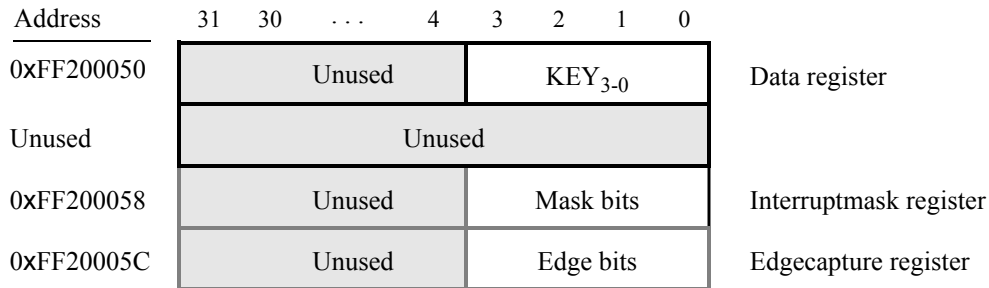


Figure 4: The KEY pushbutton switch port.




Figure 5: The SW slider switch port.

2. Create a *Makefile* that can be used to compile your kernel module to produce the file *KEY_SW.ko*. Insert this module into the kernel and verify that it creates the character device files */dev/KEY* and */dev/SW*.
3. Test your drivers by using the commands `cat /dev/KEY` and `cat /dev/SW`. Verify that the drivers provide correct values for the switches.

Part III

For this part you are to write another two character device drivers. One driver controls the LEDR lights in the DE1-SoC Computer, via the file */dev/LEDR*. The other driver controls the seven-segment displays HEX5–HEX0, via the file */dev/HEX*. The LEDR port in the DE1-SoC Computer is illustrated in Figure 6, and the seven-segment display ports are depicted in Figure 7. Your driver should be able to display decimal digits 0 – 9 on each of the six displays. Perform the following:

1. Create a kernel module source-code file called *LEDR_HEX.c*. Implement both character device drivers in this module. For both drivers you should implement functions for `open`, `release`, and `write` operations. These drivers do not require a `read` function.
2. Create a *Makefile*, compile your module, and insert it into the kernel.
3. Test the LEDR and HEX character device drivers by using the `echo` Linux command.




Figure 6: The LEDR red light port.




Figure 7: The seven-segment display ports.

Part IV

For this part you are to write a user-level program called *part4.c* that makes use of the character device drivers from Parts II and III. Perform the following:

1. In your program open the files */dev/KEY* and */dev/SW* for reading, and open the files */dev/LEDR* and */dev/HEX* for writing. Make a loop in your program that does the following. Whenever a *KEY* is pressed capture the values of the *SW* switches. Display these values on the *LEDR* lights. Also, keep a running accumulation of the *SW* values that have been read, and show this sum on the *HEX* displays, as a decimal value.
2. Compile your program using a command such as `gcc -Wall -o part4 part4.c`.
3. Run your program and make sure that it works properly. Each time a *KEY* is pressed, the values of the *SW* switches should immediately appear on the *LEDR* lights and the sum should appear on the *HEX* displays. As an example, if the *SW* switches are set to 000000101, then the first time a *KEY* is pressed **00005** should be shown on the *HEX* displays, and the second time a *KEY* is pressed **00010** should be displayed.