Chapter 5   Interactive Programs and Signals

Concepts Covered

Software tools, daemons,
interactive programs
non-blocking reads, signals,

signal, sigaction, kill, fflush, raise
sigemptyset, sigfillset, sigaddset,
sigdelset,sigprocmask, sleep

5.1   The Different Types of UNIX Programs

Programs that run in a UNIX environment can be classified by their relationship to terminal devices and by their input/output streams. They generally fall into one of three categories: *software tools*, *daemons*, and *interactive user programs*.

5.1.1   Software Tools (Filters)

A *software tool* is a program that

- receives its input from either

    - one or more files given as command-line arguments, or
    - from standard input if no filename arguments are present,

- expects its input to be an unstructured stream of bytes, almost always treated as plain text, and

- puts its output, which is also a stream of bytes, usually plain text, on standard output.

Because software tools can read from standard input and write to standard output, they can be connected via shell pipes to form pipelines, like factory assembly lines. UNIX has many software tools, including `awk`, `cat`, `cut`, `du`, `fold`, `grep`, `od`, `sed`, `sort`, `tr`, and `uniq`.

5.1.2   Daemons

Another class of programs are device drivers, which are not even attached to a terminal. A program that is not attached to a terminal is called a *daemon* in UNIX. A commonly used, but inaccurate, definition of a daemon is that it is a "background" process. To be precise, it is a process that executes without an associated terminal or login shell, waiting for an event to occur. The event might be a user request for a service such as printing or connecting to the Internet, or a clock tick indicating that it is time to run. The word "daemon" is from Greek mythology, and refers to a lesser god who did helpful tasks for the people he or she protected. Daemons are like these lesser gods; they are created at boot time, and exist, in hiding, ready to provide services when called upon.

Because daemons must not be connected to a terminal, one of their first tasks is to close all open file descriptors (in particular, standard input, standard output and standard error). They usually make their working directory the root of the file system. They then take additional steps to break their association with any shell or terminal, among which are leaving their process group and registering their intent to ignore all incoming signals. The concept of process groups will be discussed in Chapter 8. Signals are covered later in this chapter.

Daemon process names typically end in 'd'. This is one way to identify a daemon process in the output of the `ps -ef` command; find the names ending in "d" as in `ftpd`, `httpd`, `lpd`, `sshd`, `syslogd`, and `telnetd`. Daemons will be covered thoroughly in Chapter 8.

5.1.3 Interactive User Programs

Another category of programs are those that are tied to the user terminal in an inextricable way, because they customize the terminal for their own use. Programs that interact with the user through the terminal, such as text editors (`vi`, `nano`, or `emacs`), pagers (`more` and `less`), terminal-based administrative tools such as `top`, games (`snake`, `worm`, and `chess`), and terminal-based mail clients (`pine` and `mailx`), are tightly coupled to the terminal and must control its settings and attributes. They cannot use the standard input and output streams for communicating with the user because these lack the types of controls that a terminal has. These types of programs usually need to control

- whether or not characters are echoed,

- the number of characters that are buffered, if any,

- the movement of the cursor on the screen,

- whether certain key presses should have their default meaning or have application-defined meaning,

- whether timeouts should occur on input,

- whether signals such as `Ctrl-C` should be ignored, queued, or handled immediately.

We already saw how to control the state of the terminal using `stty` at the command level and the `tcgetattr()` and `tcsetattr()` functions at the programming level. Here we will explore the various modes into which we can put the terminal for the benefit of creating interactive programs.

5.2 Designing Interactive User Programs

Most interactive user programs are event-driven or menu-driven, which means that they perform some short task and then wait for user input to do the next task. All window-based applications are event-driven; they are idle, often blocked on input, while they wait for mouse, keyboard, or other events to be delivered to them by the window manager.

Here we go through the steps that are necessary to design and develop an interactive, terminal-based application. We will begin by understanding the problem, and then we will go through successive stages of making a program more and more responsive to user inputs.

5.2.1   Two Different Paradigms

Consider the kind of terminal-based program in which the program repeatedly prompts the user for input and takes an action accordingly. One of the user responses such programs expect is some type of quit signal, which typically breaks the loop and ends the program. This can be modeled by a control structure such as the following:

```
while (true) {
    prompt the user to do something
    wait for the user to respond
    if user's response is to quit
        break the loop and quit
    handle the response
}
```

The input part of this loop usually results in the process's being blocked on input, but it does not have to be designed this way. It might look like

```
while (true) {
    prompt the user to do something
    if the user responded
        handle the response
    otherwise
        do other things
}
```

In this paradigm, the program checks whether there is input and if there is, it responds to it, and if not, it does something else. This approach requires the ability to check if there is input without blocking while waiting for it. In short, it requires a type of input operation known as non-blocking input, which will be discussed below.

Regardless of which input method is used, programs such as video games, ATM machines, and text editors respond immediately to user key presses, rather than waiting for the **Enter** key to be pressed. They run in non-canonical mode, so they do not buffer input characters. Usually, they do not echo the input characters when these characters behave like function keys[1]. Also, they usually ignore illegal key presses. Thus, one task in designing interactive programs is to determine how to control the state of the terminal in this way.

But this is not enough. There is a big difference between a video game and a text editor, having to do with their relationship to time. We can distinguish between two kinds of interactive programs: those whose state is independent of time, and those whose state depends upon time. Any program that animates, in any way, is time-dependent; its state changes as a function of time. Programs that terminate or advance to a different state if the user does not respond within a certain amount of time are also time-dependent, because their state changes as a function of time. Video games are time-dependent. In contrast, a text editor is usually time-independent; it has no time-outs and no animation of any kind.

Programming with time is more complex than programming in a time-independent way because it requires knowledge of several different parts of the kernel API. Before we tackle this problem, we will explore a simpler one, namely how to write a text editor.

---

[1] Think about `vi` for example, and how it behaves in Command mode; you type a 'j' and it moves the cursor without displaying the letter, or more, when you type a space character and it advances a screen's worth of lines.

5.2.2   A Simple Text Editor: `simplevi`

The `vi` text editor is a very complicated piece of software, but we can create an extremely stripped-down version of it and still learn quite a bit in the process. We will call our editor `simplevi`. This simple editor will allow the user to insert characters anywhere in a text document, but it will not provide a means of deleting characters. This is just a minor extension to the program. Also, it will not open an existing text file, but will instead creating a new file each time it is invoked. Adding a feature to open a file does not provide much more insight into the interactive design of the program, but it does make the program larger. As much as possible, the features that `simplevi` does implement will have the same semantics as those in `vi`.

5.2.2.1   Features of simplevi

The `simplevi` program allows a user to create a file in a manner similar to `vi`. It has only an insert command and does not presently support deletion. Like `vi`, it is always in exactly one of three possible modes:

- input mode

- command mode, and

- last_line mode

Each mode is now described briefly.

Command Mode   The initial mode is command mode. In command_mode one can enter the following keystrokes/commands:

Key	Semantics
i	changes to input mode.
:	changes to lastline mode.
h, backspace, or left-arrow key	moves left until the first character in the text line. If a line wraps, it follows the wrapped text backwards.
l, spacebar, or right-arrow key	moves right until the last character in the text line. If a line wraps, it follows the wrapped text backwards.
k or up-arrow key	moves to the text line above, to the same position relative to the start of the text line, unless that text line is shorter than the current one, in which case it is placed at the end of the text line. Scrolling takes place as necessary.
j or down-arrow key	moves to the next text line below, using the same rule as for the up arrow when deciding the horizontal position. Scrolling takes place as necessary.
Ctrl-D	Does nothing except display the message that it was pressed.
Ctrl-C	Does nothing except display the message that it was pressed.
Ctrl-H	Can be used to display a help screen, but at present just shows a one-line text message.

Notes

- Navigation keys use rules similar, but not identical, to `vi`'s rules. The only difference is that vi remembers the last horizontal movement, and when a vertical movement is indicated, the cursor moves to the position in the new text line that is the same as the last remembered horizontal movement.

- Because there is only an insert operation, and not an append operation, the cursor is allowed to be one position to the right of the rightmost character in a line.

Lastline Mode   The user can quit the editor in lastline mode, and/or save the current buffer contents. The allowed commands are a subset of the commands possible in `vi`, but the syntax mimics `vi`. Specifically, in lastline mode one can enter:

`w` *filename* to save the buffer to a file named *filename*

`wq` *filename* to save the buffer to a file named *filename* and quit after

`q`   to quit without saving. It does not warn that the buffer is not saved.

Any amount of white space is allowed before the command and between the command and the filename. Any characters other than these generate an error message and terminate lastline mode.

Filenames can be any combination of alphanumeric characters and underscores. Punctuation and whitespace are not allowed. Typing a command such as

    : w q

creates a file named 'q', as would happen in `vi`.

Letting S denote a space character, F denote a valid filename character, parentheses for grouping, and | for alternation, the language of acceptable newline-terminated strings in lastline mode is

$$S*(\text{w}|\text{w}S*\text{q})SS*FF*S* \ | \ S*\text{q}S*$$

Input Mode   There is only an insert operation, which inserts at the cursor position. The backspace is not implemented. Typing a backspace has no effect. Typing characters other than graphical characters (those found on the keyboard) has unspecified behavior. Graphical characters are inserted to the left of the cursor. Lines wrap as necessary, and the screen scrolls as needed as well.

5.2.2.2   Program Design

It would be much easier to write this program if we used the NCurses library, but as we have not yet covered the NCurses API, we will do it the hard way, the way it was done before NCurses existed. This will give you an appreciation of NCurses when we cover it in the next chapter. Here, we will use the ANSI escape sequences that we covered in Chapter 1.

The design challenge in writing this simple text editor is synchronizing the three major objects that the program must manage:

- a behind-the-scenes text buffer,

- the visible screen, and

- the cursor.

The Buffer Object   The text that the user types must be stored in a buffer of some kind. There are many possible ways to organize this buffer, with time-space trade-offs associated to each. A reasonable solution would be to create an array of pointers to the individual lines of the buffer, each of which would be a dynamically allocated fixed-size array. The array of pointers could be replaced by a doubly-linked list of pointers with added programming complexity. The arrays holding the lines of text could be started smaller and reallocated to larger sizes as the lines grow in size. These details are not important to us now, as performance is less of a concern than understanding the principles. Therefore, we take an even simpler approach: the text buffer is just a large linear array of characters.

Because `vi` is a line-oriented editor and cursor movement follows text lines, it is convenient for the text buffer to keep track of the starts of each line as well as its length, as well as the number of lines and the line in which the cursor is currently located and the cursor's offset from the beginning of that line. It will also be convenient for the buffer to store the cursor's position as the offset from the beginning of the linear array itself. Thus there is some redundancy in the representation, and the buffer object must have functions to synchronize all members. The buffer object is defined as follows:

```
typedef struct _buffer
{
    char text[BUFSIZ];
    int  line_len[MAXLINES]; /* lengths of text lines, including newline
                                    characters */
    int  line_start[MAXLINES]; /* starts of each line */
    int  num_lines;          /* number of text lines in buffer. This
                                    includes lines that have not yet been
                                    terminated with a newline character. It is
                                    the number of newline characters + 1 if the
                                    last character in the buffer is not a
                                    newline. */
    int  index;              /* index of cursor in text buffer */
    int  size;               /* total chars in buffer */
    int  cur_line;           /* current text line, not screen line */
    int  index_in_cur_line;  /* index in current line of cursor */
} Buffer;
```

The `BUFSIZ` constant is a system constant, and `MAXLINES` is defined in the application header file.

The Window Object   The window is a display object. The program needs to find its dimensions on start up and store them in the window object. Because the buffer contents may be larger than can fit in the window, at any given time, there is a set of text lines that is visible in the window,

and this is a subsequence of the lines in the buffer. The smallest index of this sequence is stored as a member of the window to facilitate redrawing the buffer during scrolling operations and insertions. The last line index is recalculated rather than storing it. It could also be stored to save running time. Finally, because the erase character is a terminal setting, its value is stored within the window object (a window is really a terminal), and will be determined at program startup by querying the terminal driver. The window object is defined by

```
typedef struct _window
{
    unsigned short rows;
    unsigned short cols;
    int            line_at_top;
    char           erase_char;
} Window;
```

The window is divided into two regions: the line at the bottom of the screen and everything above it. The bottom line is reserved by the application for writing status messages and for *lastline mode*, which is explained below. The text is never allowed to be written onto the window's last line.

The Cursor   The cursor is simple: it is a structure with a row and column index, representing a position in a two-dimensional array whose upper-left corner is position (0,0). The cursor is defined by

```
typedef struct _cursor
{
    int r;
    int c;
} Cursor;
```

Although the cursor position is in zero-based two-dimensional coordinates, the ANSI escape sequences use a 1-based set of coordinates: the upper-left corner of the screen is (1,1). Therefore, the cursor functions must account for this difference. For example, the function to move the cursor to a new row and column position is

```
void moveto(int line, int column )
{
    char seq_str[20];

    sprintf(seq_str, "\033[%d;%dH", line+1, column+1);
    write(1, seq_str, strlen(seq_str));
}
```

Some of the cursor functions are relatively simple. When **simplevi** is in input mode and a character is typed on the keyboard, three things must happen:

- the character must be inserted into the text buffer;

- the cursor must be advanced to the next insertion point (because echo is controlled by the application, not by the terminal driver), and

- the window must be redrawn to account for the inserted character.

The function to advance the cursor is in the listing below.

Listing 5.1: advance_cursor

```
void advance_cursor( Cursor *cursor, Window win, char ch)
{
    if ( ch == '\n' ) {
        cursor->r++;
        cursor->c = 0;
    }
    else {
        cursor->c++;
        if ( cursor->c == win.cols ) {   /* wrap the line */
            cursor->c = 0;
            cursor->r++;
        }
    }
}
```

The function checks which character was typed. If a newline, it advances the cursor to the next line at column 0, and if not, it checks whether it has to account for wrapping the line: if the cursor is at the right margin, it also advances to the next line, column 0.

Inserting a Character   When the program is in input mode and the use enters a character to be inserted into the buffer, four actions must be taken:

1. physically inserting the character into the buffer, provided no overflow occurs,

2. advancing the cursor to the next logical position in the window,

3. redrawing the buffer contents to the screen, which may cause the cursor's logical position to change, and

4. physically moving the cursor to the final screen position.

The first step is implemented by the `insert()` function in Listing 5.2 below.

Listing 5.2: insert()

```
int insert( Buffer *buf, Window win, char c )
{
    int i;

    if ( ( c == '\n' ) && ( MAXLINES == buf->num_lines ) )
        return OUT_OF_LINES;
    else if ( buf->size == BUFSIZ )
        return OUT_OF_MEM;
```

```
    if ( c == win.erase_char ) {
        /* Not implemented -- just ignore it */
        return UNHANDLEDCHAR;
    }

    for ( i = buf->size; i > buf->index; i-- )
        buf->text[i] = buf->text[i-1];
    buf->text[buf->index] = c;
    buf->size++;
    buf->index++;
    buf->line_len[buf->cur_line]++;

    /* the first character sets line count to 1 */
    if ( buf->size == 1 ) buf->num_lines++;

    if ( c == '\n' ) {
        /* Save the length of the line being split by the newline */
        int temp = buf->line_len[buf->cur_line];

        /* The new length of current line is the current index position + 1. */
        buf->line_len[buf->cur_line] = buf->index_in_cur_line + 1;

        /* Increase number of lines */
        buf->num_lines++;

        /* Shift all line starts and lengths upwards in the array, but
           add 1 to the line starts since they are 1 character further
           than before because of the new newline. Do this from the last line
           down to cur_line+1, which is the line just after the split
           line. */
        for ( i = buf->num_lines-1; i > buf->cur_line+1; i-- ) {
            buf->line_len[i] = buf->line_len[i-1];
            buf->line_start[i] = buf->line_start[i-1]+1;
        }
        /* Set the start of the new line created here. It is the sum of the
           start of cur_line plus the length of cur_line. */
        buf->line_start[buf->cur_line+1] = buf->line_start[buf->cur_line]
                               + buf->line_len[buf->cur_line];

        /* advance to new line */
        buf->cur_line++;
        /* The length of the newly created line is the number
           of characters that were to the right of the current
           index position. */
        buf->line_len[buf->cur_line] = temp - buf->line_len[buf->cur_line-1];
        buf->index_in_cur_line = 0;
    }
    else if ( isprint(c) ) {   /* non-newline character */
        buf->index_in_cur_line++;     // advance index in line
        /* increment all line starts after this line */
        for ( i = buf->cur_line+1; i < buf->num_lines; i++ )
            buf->line_start[i]++;
    }
```

```
        else
            return UNHANDLEDCHAR;
        return 0;
}
```

Skipping some of the details, this function basically shifts all characters upwards in the buffer and adjusts the starts of lines for all lines after the current line. It inserts the new character at the index position, increments the buffer size, the index position, and the length of the current line. It has to determine whether the inserted character is a newline character. If so, it has to split the line and adjust the line starts and line lengths arrays. If not, it checks if it is a printable character, and if so, adjusts the index in the current line and line starts accordingly.

The second step, advancing the cursor, is shown in Listing 5.1. The third step, redrawing the buffer, is handled by the `redraw_buffer()` function, which is found in Listing 5.3 below.

Listing 5.3: redraw_buffer()

```
void redraw_buffer( Buffer buffer, Window *win, Cursor *curs )
{
    int i;
    int lastline;
    int lastchar;
    int firstchar;
    int line_of_cursor;

    /* If the current position in the buffer, buffer.index, is not within the
       visible lines of the window, the window must be shifted.  The shift might
       be up or down, depending on whether index is above or below the window.

    We first need to get the number of the line containing pos. Then we
    check whether that line is between win.line_at_top and lastline.
    We need to calculate the difference and  shift win.line_at_top that
    difference, and recalculate lastline, after which we can draw the
    buffer. */

    /* Compute the last visible complete text line in the buffer */
    get_lastline_in_win( buffer, *win, &lastline );

    /* Get the index of the text line containing the insertion position */
    line_of_cursor = line_in_buffer( buffer, *win, buffer.index ) ;

    /* Check if the window needs to be scrolled */
    if (line_of_cursor < win->line_at_top ) {
        lastline -= (win->line_at_top - line_of_cursor);
        curs->r += (win->line_at_top - line_of_cursor);
        win->line_at_top = line_of_cursor;
    }
    else if (line_of_cursor > lastline ) {
        win->line_at_top += (line_of_cursor - lastline);
        curs->r -= (line_of_cursor - lastline);
        lastline = line_of_cursor;
    }

    /* Get the first and last characters of the visible screen. The lastchar
       is the index of the last character that can appear in the
```

```
            window —— the  last  character  in  the  last  visible  line.  The  first  char
            is  the  start  of  the  line  at  the  top  of  the  screen. */
    lastchar  = buffer.line_start[lastline] + buffer.line_len[lastline];
    firstchar = buffer.line_start[win−>line_at_top];

    /* Prepare to redraw the window. First clear the screen. */
    write(1,CLEAR_SCREEN, lCLEAR_SCREEN);

    /* Do the redraw */
    moveto(0,0);
    for ( i = firstchar; i < lastchar; i++ )
        write(1,&buffer.text[i], 1 );
}
```

The `redraw_buffer()` function could be more efficient. It is written for clarity. The problem is that the line containing the insertion might have grown long enough that it wrapped and part of it is no longer in the visible region of the window. It could also be that the inserted character is outside of the visible region because the cursor was at the right margin and the new character is in the next row of the screen. In these cases, the entire window must be redrawn because the starting line and ending line have changed. In other cases it is not necessary to redraw the entire window, but for simplicity, the entire window is always redrawn.

Redrawing the buffer requires determining whether the current line and the index of the cursor in that line are outside of the visible region. This is accomplished in a few steps:

1. getting the index of the last visible line in the window,

2. getting the index of the text line containing the insertion position,

3. comparing the line of the insertion point to the top line and last line computed above, and if out of bounds, resetting the line at the top of the window so that the insertion point is visible,

4. calculating the first and last characters of the buffer that delimit the visible characters in the window, and

5. moving to the upper left corner and redrawing all characters from the first to the last.

Cursor Movement One of the primary challenges in managing the cursor is that, at any given time, the program must be able to map the cursor position to a position in the text buffer, and vice versa.

The program will always make sure that it knows the current line and the current index in the line as the cursor moves around on the screen. In fact, the cursor movement operations will actually update these variables, and from those recalculate the cursor position. This needs explanation. Text lines may be longer than the screen width. When this happens, they wrap onto two or more lines. In `vi`, when the cursor is moved upward or downward, it "jumps" over the wrapped lines. In other words, it moves from one text line to another, not from one screen row to another. Our `simplevi` program emulates this behavior. Therefore, when the user presses an arrow key up or down, it will move to the preceding or following text line. To make this possible, it will increment or decrement the `cur_line` member of the buffer as needed, and possibly adjust the `index_in_cur_line` variable, as will be explained below. But this implies that the index member is no longer synchronized with

these values. This only matters when an insertion or other operation has to be performed on the text in the buffer at that position. Rather than updating the index each time, the program only updates it when such an action is needed. The following function does this.

Listing 5.4: update_buffer_index()

```
void update_buffer_index( Buffer *buffer )
{
    int totalchars = 0;
    int i = 0;

    while (  i < buffer->cur_line ) {
        totalchars += buffer->line_len[i];
        i++;
    }
    totalchars += buffer->index_in_cur_line;
    buffer->index = totalchars;
}
```

Similarly, it is also necessary to get the cursor's actual window coordinates when all that is know are the current line number and the offset within that line. This is the case, for example, when the down arrow key is pressed and the cursor must be moved to a new text line. It is possible that the cursor would actually be below the last visible line. The only way to know this is to calculate its window coordinates and compare them to the bottom of the screen. The following function is our means of getting the cursor position.

Listing 5.5: get_cursor_at()

```
void get_cursor_at( Buffer buf, Window win, int index, int lineno,
                    Cursor *curs )
{
    int total_lines_before = 0;
    int rows_in_current_textline = 0;
    int i;

    /* The first line is the one at the top of the window, whose index is
       win.line_at_top, initially 0. */
    for ( i = win.line_at_top; i < lineno; i++ ) {
        if ( buf.line_len[i] < win.cols )
            total_lines_before++;
        else
            total_lines_before += (int) ceil((double)buf.line_len[i]/win.cols);
    }
    rows_in_current_textline = index/win.cols;

    curs->r = total_lines_before + rows_in_current_textline;
    curs->c = index - rows_in_current_textline * win.cols;
}
```

There are four movements that the program has to implement: up, down, left, and right. We show just one of these here. The complete program is in the appendix. The up and down movements are a bit more complex than left and right.

When the down arrow key is pressed, or the 'j' key, the cursor must be moved to the position in the text line below at the same offset relative to the beginning of the line, unless that line is shorter

than the current one. In the latter case, it is moved to the end of the line. The complication is that this position may be out of view, in which case the text must be scrolled upwards to put it into view. Therefore, the function to do this must move the cursor logically and then check whether this logical position is out of view. This requires

1. advancing `cur_line` to the line below if not on the last line;

2. adjusting `index_in_cur_line` if the line is too short;

3. getting the index of the last line in the window;

4. getting the actual cursor position;

5. if the cursor is below the last line, calculating by how much scrolling accordingly; and

6. moving the cursor to the new position on the screen and drawing it.

The `moveDown()` function in the listing below encapsulates this logic.

Listing 5.6: moveDown()

```
void moveDown( Buffer *buf, Window *win, Cursor *curs )
{
    int lastline;

    if ( buf->cur_line < buf->num_lines-1 ) {
        buf->cur_line++;
        if ( buf->index_in_cur_line >= buf->line_len[buf->cur_line] ) {
            buf->index_in_cur_line = buf->line_len[buf->cur_line]-1;
        }

        get_lastline_in_win( *buf, *win, &lastline );
        if ( buf->cur_line > lastline ) {
            win->line_at_top += buf->cur_line - lastline;
            get_cursor_at(*buf, *win, buf->index_in_cur_line,
                        buf->cur_line, curs );
            scroll_buffer(*buf, *win );
        }
        else
            get_cursor_at(*buf, *win, buf->index_in_cur_line,
                        buf->cur_line,  curs );
        moveto(curs->r, curs->c);
    }
}
```

5.2.2.3   Terminal Interaction

Terminal interaction includes modifying or querying the terminal state, obtaining the window size and the erase character, and writing various ANSI escape sequences to do things such as clearing parts of the screen or moving the cursor.

Terminal Support Functions Three functions are used for modifying or querying the terminal
state, as shown in the following program listing.

Listing 5.7: Terminal support functions

```
void modify_termios(int fd, int echo, int canon )
{
    struct termios cur_tty;
    tcgetattr(fd, &cur_tty);

    if ( canon )
        cur_tty.c_lflag       |= ICANON;
    else
        cur_tty.c_lflag       &= ~ICANON;
    if ( echo )
        cur_tty.c_lflag       |= ECHO;
    else
        cur_tty.c_lflag       &= ( ~ECHO & ~ECHOE);
    cur_tty.c_lflag       &= ~ISIG;
    cur_tty.c_cc[VMIN]    =   1;
    cur_tty.c_cc[VTIME]   =   0;
    tcsetattr(fd, TCSADRAIN, &cur_tty);
}


void save_restore_tty(int fd, int action)
{
    static struct termios original_state;
    static int               retrieved = FALSE;

    if ( RETRIEVE == action ){
        retrieved = TRUE;
        tcgetattr(fd, &original_state);
    }
    else if (retrieved &&  RESTORE == action ) {
        tcsetattr(fd, TCSADRAIN, &original_state);
    }
    else
        fprintf(stderr, "Illegal action to save_restore_tty().\n");
}

void  init_window( int fd, Window *win )
{
    struct winsize size;

    if (ioctl(fd, TIOCGWINSZ,  &size) < 0) {
        perror("TIOCGWINSZ error");
        return;
    }
    win->rows = size.ws_row;
    win->cols = size.ws_col;
    win->line_at_top = 0;
}


void get_erase_char(int termfd, Window *win )
{
```

```
        struct termios cur_tty;
        tcgetattr(termfd, &cur_tty);

        win->erase_char = cur_tty.c_cc[VERASE];
}
```

`modify_termios()` either enables or disables echo and canonical mode, and it disables keyboard signals. It also sets the `MIN` and `TIME` line discipline values to 1 and 0 respectively so that the program reads a single character at a time and does not time-out. `save_restore_tty()` can be used to save the current terminal state into a local static variable for later restoration. `get_winsize()` uses the `ioctl()` function to get the current window size when the program starts up. If the window is resized while the program is running, all bets are off – it does not handle resizing events.

Since output operations must bypass the C standard I/O library to avoid falling victim to its internal buffering, the program will only use the kernel's `write()` system call for writing to the screen. A write to the terminal device will always place the bytes to be written at the current cursor position. After a write completes, the cursor is advanced to the right of the last character written, unless what was written was an ANSI escape sequence that does not actually write to the screen. Because the `write()` system call needs the length of the string to write, to improve performance, the length of each escape sequence is hard-coded as a constant. The listing below is the complete set of sequences used by the program.

Listing 5.8: ANSI escape sequences

```
/* ANSI Escape Sequences and their lengths (for speed */
const char CLEAR_DOWN[]    = "\033[0J";
const int  lCLEAR_DOWN     =  4;
const char CLEAR_RIGHT[]   = "\033[0K";
const int  lCLEAR_RIGHT    =  4;
const char CURSOR_HOME[]   = "\033[1;1H";
const int  lCURSOR_HOME    =  6;
const char CLEAR_SCREEN[]  = "\033[2J";
const int  lCLEAR_SCREEN   =  4;
const char CLEAR_LINE[]    = "\033[2K";
const int  lCLEAR_LINE     =  4;
const char RIGHT[]         = "\033[1C";
const int  lRIGHT          =  4;
const char LEFT[]          = "\033[1D";
const int  lLEFT           =  4;
const char BACKSPACE[]     = "\033[1D \033[1D";
const int  lBACKSPACE      =  9;
/* The following cannot be declared statically as the dimensions are unknown */
char       PARK[20];        /* string to park cursor at lower left */
int        lPARK;           /* length of PARK string */
```

The string `PARK` for parking the cursor in the last line of the window is defined dynamically, because it depends upon the screen size. It is defined as follows:

```
    get_winsize(STDIN_FILENO, &rows, &cols);
    sprintf(PARK, "\033[%d;1H", rows);
```

5.2.2.4   The Main Program

The main function initializes everything necessary and then enters a loop in which it remains until the user enters a command to quit in lastline mode. Only the main loop is shown in the listing below.

Listing 5.9: main program loop

```
while ( !quit ) {
    if ( in_input_mode ) {
        if (  read(STDIN_FILENO, &c, 1) > 0   ) {
            if ( c == ESCAPE ) {
                in_input_mode = 0;
                write_status_message( "           ", curs);
            }
            else {
                in_input_mode = handle_insertion( &buf, &win, &curs, c);
                if ( in_input_mode == UNHANDLEDCHAR )
                    in_input_mode = 1;
                else
                    write_status_message( INSERT, curs);
            }
        }
    }
    else {
        if ( read(STDIN_FILENO, &c, 1) > 0   ) {
            switch ( c ) {
            case 'i':
                in_input_mode = 1;
                park();
                update_buffer_index(&buf);
                moveto(curs.r, curs.c);
                write_status_message( INSERT, curs);
                break;
            case ':':
                in_lastline_mode = 1;
                park();
                write(1,&prompt, 1);
                status = do_lastline_mode( buf, win, curs);
                if ( status >= 0 ) quit = status;
                moveto(curs.r, curs.c);
                break;
            case '\003':
                write_status_message( CTRLC, curs);
                break;
            case '\004':
                write_status_message( CTRLD, curs);
              break;
            case '\010':
                write_status_message( CTRLH, curs);
                break;
            case 'j':
                moveDown(&buf, &win, &curs );
                break;
            case 'k':
```

```
                    moveUp(&buf, &win, &curs );
                    break;
            case 'l':
            case ' ':
                    moveRight(&buf, &win, &curs );
                    break;
            case 'h':
                    moveLeft(&buf, &win, &curs );
                    break;
            case ESCAPE:
                    handle_escape_char( &buf, &win, &curs);
                    break;
            default:
                    if ( c == win.erase_char )
                        moveLeft(&buf, &win, &curs );
            }
        }
    }
}
```

This program has all of the elements of an interactive application that does not involve time. The remainder of this chapter will introduce features of the UNIX API that are needed for programs whose state changes with time.

5.2.3   Non-Blocking Input

Changing the state of the terminal itself is preferable to changing the attributes of the open file descriptor, because we can exercise more control over it. However, the downside of this is that all programs that use that terminal will be affected by the changes. Usually this is not a problem, since the program you are writing will be the only foreground process, and no other process will be reading from the terminal while it is running. However, it is worthwhile to understand how to exercise some limited control over input through the file descriptor, using the `O_NDELAY` flag (called `O_NONBLOCK` in POSIX).

The `O_NONBLOCK` flag controls whether reads and writes are blocking or non-blocking. When a read is blocking, the process that executes the read waits until input is available, and only then does it continue. This is the semantics that beginning programmers learn. This makes sense; after all, why would you ever want a program to continue past a read instruction if the read did not yield any data?

Exercise. Before reading further, try to answer the preceding question.

Non-blocking I/O is a property of open file connections, not of terminals or devices; when you open a file and get a file descriptor as the return value of the `open()` call, you can specify in the call that the file connection should be non-blocking. In other words, the property of being non-blocking is part of the process's connection to the file or the terminal, not the terminal itself. Two different processes can have the same file open for reading, one using blocking reads and the other, non-blocking.

Remember that when a connection is established between a process and an input source, a buffer is created that holds data from the source on its way to the process. Whether it is a disk file, a terminal, a pointing device, or an audio source, there is some temporary storage area used for

buffering input. When a process opens a non-blocking connection to an input source, whether it is a file or a device, calls to read data from that source retrieve whatever data is in the buffer at the time of the call, up to the amount requested in the read request, and return immediately. If the buffer is empty, they return immediately with no data. To be clear, in the call

```
if ( (bytesRead = read(fd, dataToProcess, bytesToGet) ) > 0 )
    {/* statement block */}
```

if the connection is non-blocking and no data is in the buffer of the file descriptor `fd`, then the call returns immediately with `bytesRead == 0` and the statement block will be skipped.

Similarly, a call to the C library function `getchar()` will return either the character in the front of the buffer or, if the buffer is empty, nothing at all. In either case, `getchar()` returns immediately, i.e., in

```
if ( ( c = getchar())  != EOF )  { /* statement block */ }
```

if there is no character in the buffer, `getchar()` will return `EOF`, which the program can use to decide how to proceed.

Non-blocking input should not be confused with *asynchronous input*. Asynchronous input occurs when the process makes a call to read data and returns immediately, without waiting for the data to be ready. The `read()` call is executed by a separate process, and as soon as the data are available for it, that process performs the I/O and fills the buffer passed to it by the one calling `read()`. This is called asynchronous input because the caller does not synchronize with the process running the `read()` call; they proceed independently once the call is made. Asynchronous input is useful when you may not need the data right away.

Non-blocking input is useful when the lack of input itself is a significant condition to be identified by the program. For example, it might indicate that a user has left the terminal and is no longer responding, or that a connection to a remote host has been broken, or that a pipe is empty. It may also imply that the user is choosing to not supply input because supplying input may mean making something happen that the user does not want to happen, as in a video game. Very often the process requesting the input has other work to do and it can simply check later whether the input is available. For this reason, non-blocking reads are usually placed inside loops where the condition is tested and an appropriate action can be taken. In most programs that use non-blocking input, the state of the program is changing without the user's intervention. This might be because animation is taking place, or a computation is being performed, or something else entirely.

The following listing is of a program that uses non-blocking input and pretends to do a simple animation. It draws "dots" on the screen, nothing more. Because it uses the same `modify_termios()`, `save_restore_tty()` and `get_winsize()` functions from the `simplevi.c` program, their definitions are omitted. It uses a function to put a file descriptor into non-blocking mode, `set_non_block()`.

Listing 5.10: nonblockdemo.c

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <fcntl.h>
```

```c
#include <termios.h>
#include <sys/stat.h>
#ifndef TIOCGWINSZ
#include <sys/ioctl.h>
#endif

#define RETRIEVE        1               /* action for set_tty */
#define RESTORE         2               /* action for set_tty */

// Defined elsewhere:
void modify_termios(int fd, int echo, int canon );
void save_restore_tty(int fd, int action);
void  get_winsize( int fd, int *rows, int *cols );

void set_non_block(int fd )
{
    int flagset;

    flagset    = fcntl(fd, F_GETFL);
    flagset    |= O_NONBLOCK;
    fcntl(fd, F_SETFL, flagset );
}


int main (int argc, char *argv[])
{
    char       ch;            // stores user's char
    char    period = '.';
    size_t bytecount;
    int     count = 0;
    int     done  = 0;    /* to control when to stop loop */
    int     pause = 0;    /* to control pausing of output */
    char    PARK[20];     /* ANSI escape sequence for parking cursor */
    int     numrows;      /* number of rows in window */
    int     numcols;      /* number of columns in window */
    const   char CURSOR_HOME[]  = "\033[1;1H";
    const   char CLEAR_SCREEN[] = "\033[2J";
    const   char SAVE_CURSOR[]  = "\033[s";
    const   char REST_CURSOR[]  = "\033[u";
    const   char MENU[] = "Type q to quit or p to pause or r to resume.";
    char    dots[20];


    /* Check whether input or output has been redirected */
    if ( !isatty(0) || !isatty(1) ) {
        fprintf(stderr,"Output has been redirected!\n");
        exit(EXIT_FAILURE);
    }

    /* Save the original tty state */
    save_restore_tty(STDIN_FILENO, RETRIEVE);

    /* Modify the terminal -
        turn off echo, keybd sigs, and canonical mode */
```

```
    modify_termios( STDIN_FILENO, 0, 0);

    /* Turn off blocking mode */
    set_non_block( STDIN_FILENO );

    /* Get the window's size */
    get_window_size(STDIN_FILENO, &numrows, &numcols);

    /* Create string to park cursor */
    sprintf(PARK, "\033[%d;1H", numrows+1);

    /* Clear the screen and put cursor in upper left corner */
    write(STDOUT_FILENO,CLEAR_SCREEN, strlen(CLEAR_SCREEN));
    write(STDOUT_FILENO,CURSOR_HOME, strlen(CURSOR_HOME));

    /* Start drawing. Stop when the screen is full */
    while ( !done ) {
        if ( ! pause ) {
            count++;
            /* Is screen full except for bottom row? */
            if ( count > (numcols * (numrows-1)) ) {
                pause = 1;
                count--;
            }
            else
                write(STDOUT_FILENO, &period, 1);
        }
        usleep(10000);  /* delay a bit */
        sprintf(dots, " dots:  %d ", count);

        /* Save the cursor, park it, write the menu prompt */
        write(STDOUT_FILENO,SAVE_CURSOR, strlen(SAVE_CURSOR));
        write(STDOUT_FILENO, PARK, strlen(PARK));
        write(STDOUT_FILENO, MENU, strlen(MENU) );
        write(STDOUT_FILENO, dots, strlen(dots));
        /* Do the read. If nothing was typed, do nothing */
        if ( (bytecount = read(STDIN_FILENO, &ch, 1) ) > 0 ) {
            if ( ch == 'q' )
                done = 1;
            else if ( ch == 'p' )
                pause = 1;
            else if ( ch == 'r' )
                pause = 0;
        }
        /* Restore the cursor so the next dot follows the previous */
        write(STDOUT_FILENO,REST_CURSOR, strlen(REST_CURSOR));
    }
    /* Cleanup -- flush queue, clear the screen, and restore terminal */
    tcflush(STDIN_FILENO,TCIFLUSH);
    write(1,CLEAR_SCREEN, strlen(CLEAR_SCREEN));
    write(1,CURSOR_HOME, strlen(CURSOR_HOME));
    save_restore_tty(STDIN_FILENO, RESTORE);
    return 0;
}
```

The program has a loop of the form

```
while ( !done ) {
    /* actual work here */
    if ( (bytecount = read(STDIN_FILENO, &ch, 1) ) > 0 ) {
        if ( ch == 'q' )
            done = 1;
        else if ( ch == 'p' )
            pause = 1;
        else if ( ch == 'r' )
            pause = 0;
        }
}
```

This is the second paradigm shown in Section 5.2.1. The user's input has one of three possible effects: (1) entering `'q'` terminates the loop, (2) entering `'p'` allows the loop to continue but stops output, giving the illusion that the program is paused, and (3) entering `'r'` resumes the output if it is paused and has no effect otherwise. This is a very inefficient method of pausing of course, because the program is gobbling up CPU cycles while it is pretending to do nothing. However, we do not know enough as yet to do otherwise.

The part of the loop with the comment labeled "actual work here" is the part in which it

1. increments the period count and checks how many periods have been written so far,

2. delays a bit,

3. if still room for another period, writes a period, and

4. moves the cursor to the bottom row, writing the prompt and a count of the periods.

This section of the code is a form of animation – it is changing the state of the screen as a function of time. In this case, the timing is achieved by pausing a constant amount of time between redraws. This is a simple, and inefficient, form of animation. Later we will see that there are better means of achieving this.

5.2.4   Allowing Time-Outs

Sometimes we would like to write a program that has time-outs: if the user does not respond within a certain amount of time, it will take this condition to be a significant event in itself. Many programs have some kind of time-out or delay feature like this, so that if the user does not respond within a certain amount of time, the program will either terminate or take some other default action.

We can add a time-out feature to a program by setting `MIN` to 0 and `TIME` to the number of deciseconds we would like it to wait before it decides that there is no input to wait for. For the sake of curiosity, though, we will design a program that will allow us to set the `MIN` and `TIME` terminal attributes to any values we choose, so that we can see how it behaves with different values.

5.2.5   A Test Program

We will build a program in which we can test the effects of changing both the terminal driver's attributes and the open connection's attributes. The main program will be a test driver that allows the user to control the state of the terminal and terminal connection by various command line options, and repeatedly runs a simple function, which we will call `get_response()`, that reads user input in the given state of the terminal and connection. The main program will have a few bells and whistles besides.

The program will have separate functions for controlling the state of the control terminal and for changing the attributes of the file connection to the terminal device. The main program will have a loop to allow us to experiment with the `get_response()` function until we are satisfied that we understand how it behaves under the given settings. There are several pieces to the program, which we present in a bottom-up approach.

First, we combine the `save_restore_tty()` and `set_non_block()` functions into a single function that saves and restores both the terminal settings and the file descriptor flags. It uses the same macros as before:

```
void save_restore_tty(int fd, int action, struct termios *copy)
{
    static struct termios original_state;
    static int           original_flags = -1;
    static int           retrieved = FALSE;

    if ( action == RETRIEVE ){
        retrieved = TRUE;
        tcgetattr(fd, &original_state);
        original_flags = fcntl(fd, F_GETFL);
        if ( copy != NULL )
            *copy = original_state;
    }
    else if (retrieved && action == RESTORE ) {
        tcsetattr(fd, TCSADRAIN, &original_state);
        fcntl( fd, F_SETFL, original_flags);
    }
    else
        fprintf(stderr, "Bad action to save_restore_tty().\n");
}
```

We will change our `modify_termios()`, function so that it can be given a structure whose members describe the terminal settings:

```
typedef struct tty_opts_tag {
    int min;     /* value to assign to MIN */
    int time;    /* value to assign to TIME */
    int echo;    /* value to assign to echo [0|1] */
    int canon;   /* value to assign to canon [0|1] */
} tty_opts;
```

`modify_termios()` will options in the `tty_opts` parameter that is passed to it to the given `termios` structure.

```
void modify_termios( struct termios *cur_tty, tty_opts ts )
{
    if ( ts.canon )
        cur_tty->c_lflag      |= ICANON;
    else
        cur_tty->c_lflag      &= ~ICANON;
    if ( ts.echo )
        cur_tty->c_lflag      |= ECHO;
    else
        cur_tty->c_lflag      &= ~ECHO;
    cur_tty->c_cc[VMIN]       =  ts.min;
    cur_tty->c_cc[VTIME]      =  ts.time;
}
```

The second, `apply_termios_settings()`, applies the values in the `termios` structure to the terminal line associated to the given file descriptor (which should be standard input.)

```
void apply_termios_settings( int fd, struct termios cur_tty )
{
    tcsetattr(fd, TCSANOW, &cur_tty);
}
```

The `set_non_block()` function is the same as the one we used above and is omitted.

The `get_response()` function will prompt the user to type a character and will return a value that indicates what the user typed. For simplicity, it will ask for yes or no answers. It prints a question on the screen and gives the user a chance to give a valid response. If the user types a valid response or `max_tries` attempts were made, it returns.

```
int get_response( FILE* fp, ui_params  uip )
{
    int input, n;
    unsigned char c;
    time_t time0, time_now;

    time(&time0);
    while ( TRUE ){
        printf( "%s (y/n)?", uip.prompt );
        fflush(stdout);
        if ( !uip.isblocking )
            sleep(uip.sleeptime);
        if ( (n = read(fileno(fp), &c, 1)) > 0 ) {
            tcflush(fileno(fp), TCIFLUSH);
            input = tolower(c);
            if ( input == 'y' || input == 'n' ) {
                return input;
```

```
                    }
                    else {
                        printf("\nInvalid input: %c\n", input);
                        continue;
                    }
            }
            time(&time_now);
            printf("\nTimeout waiting for input: %d secs elapsed,"
                    " %d timeouts left\n",
                      (int)(time_now- time0), uip.maxtries);
            if ( uip.maxtries-- == 0 ) {
                printf("\nTime is up.\n");
                return 0;
            }
        }
}
```

Comments

- The `fflush()` call flushes the buffers associated with the file stream passed to it. The C Standard I/O Library provides buffered I/O for file streams. When a program is started, by default, the streams `stdin`, `stdout`, and `stderr` are *line buffered*. This means that the characters are transmitted to the terminal only when a newline is placed onto the stream. Since functions such as `printf()`, `puts()`, and the others that act on `stdout`, act on file streams, they are line buffered. The preceding `printf()` call

  ```
  printf("%s (y/n)?", uip.prompt);
  ```

  sends a string to `stdout` without a terminated newline and therefore this string will not appear immediately. To force the characters to be delivered to the terminal device, we use `fflush(stdout)`, which empties the buffer. If we comment out the `fflush()` call, the prompt will not appear on the screen until after a `read()` runs.

  Note that the buffering provided for streams is independent of the buffering done by the terminal within the line discipline. Even if you put the terminal into non-canonical mode, if you use the higher-level C library functions, C will continue to line buffer. You must use the lower-level file descriptor operations to avoid the buffering.

- The call to flush the terminal's input queue, `tcflush()`, is needed in case the program is run in canonical mode and input is buffered. In this case the user has to enter a newline before the terminal will deliver the characters to the `read()` call, and `get_response()` needs to remove that newline character, otherwise it will be used as the next input character when it is called again.

- `get_response()` calls `sleep()` to block itself for the number of seconds given by `uip.sleeptime`. The `sleep()` function's prototype is

  ```
  unsigned int sleep(unsigned int seconds);
  ```

The process will sleep until either the given time has elapsed or it receives a signal that it does not ignore. (Signals will be covered soon.) This gives the user a chance to type a response in case non-blocking input is in effect. Without this delay, the `read()` call would return faster than the user could blink an eye. The delay is not needed when input is blocking. There are alternatives to `sleep()` with finer granularity, such as `usleep()` and `nanosleep()`.

- Lastly, `get_response()` computes the total time elapsed since the original prompt was displayed by calling the `time()` function initially and each time that the user fails to enter any character before time runs out, and computing the difference in seconds.

The main program, with includes and a few other utilities omitted, follows. This is `term_demo1.c`.

Listing 5.11: term_demo1.c

```c
#define PROMPT               "Do you want to continue?"
#define MAX_TRIES            3          // max tries
#define SLEEPTIME_DEFAULT    2          // delay after prompt
#define BEEP                 putchar('\a')   // alert user
#define RETRIEVE             1          // action for save_restore
#define RESTORE              2          // action for save_restore
#define FALSE                0
#define TRUE                 1


int main(int argc, char* argv[])
{
    int             response;
    tty_opts        ttyopts = {1,0,1,1};
    ui_params       ui_parameters = {
                        SLEEPTIME_DEFAULT, TRUE,
                        MAX_TRIES, PROMPT
                    };
    int             fflags = 0;
    struct termios  ttyinfo;
    int             fd;
    FILE*           fp;

    get_options(&argc, &argv, &fflags, &ttyopts,
                    &(ui_parameters.sleeptime));

    if ( !isatty(0) || !isatty(1) )
        exit(1);
    fp = stdin;

    save_restore_tty(fileno(fp), RETRIEVE, &ttyinfo);
    modify_termios( &ttyinfo, ttyopts );
    apply_termios_settings( fileno(fp), ttyinfo );
```

```
        // If fflags != 0 then the O_NONBLOCK flag is set on fp
        if ( fflags ) {
            ui_parameters.isblocking = FALSE;
            set_non_blocking_mode(fileno(fp), fflags);
        }
        while ( TRUE ) {
            response = get_response(fp, ui_parameters );
            if ( response ) {
                printf("\nMain: Response from user = %c\n",
                        response );
                if ( 'n' == response )
                    break;
            }
            else
                printf("\nMain: No response from user\n" );
        }
        save_restore(fileno(fp), RESTORE, NULL);
        tcflush(0,TCIFLUSH);
        return response;
}
```

Comments

1. One problem with the program as it stands is that it will not respond to any key presses until it wakes up from its sleep. You cannot kick it into waking up. It will sleep for the specified time, come hell or high water. This means that a user with a fast response time will be unhappy with this solution. It is a one-size-fits-all solution to making a program responsive to user response rates. It is, in essence, a polling solution because it just repeatedly checks in on the user. An alternative is to somehow make the program sleep until the user actually does something.

2. Another problem with the program is that if it is terminated abnormally, as when the user types `Ctrl-C`, it will not have a chance to restore the terminal to its original settings. If the program has turned off canonical input and echo, for example, and it is killed before reaching the instruction in the main program to restore the terminal, then when the shell resumes execution upon the program's untimely death, the terminal will still be in non-canonical mode with no echo. Fortunately for present-day UNIX programmers, modern shells such as `bash` and `tcsh` automatically reset the terminal when processes invoked from the shell are killed, so these users will not see this happen.

3. However, even with the shell's ability to reset the terminal, it is still not necessarily immune to the problem that occurs when the program turns on non-blocking I/O[2]. If the program turns on non-blocking reads and is subsequently killed, there is a good possibility that the shell will be killed too. This is because, when a program is invoked from the shell, it shares the file descriptors of its parent shell. In other words, file descriptor 0 in the program points

---

[2]Some shells appear to have fixed this "bug" as well. If a spawned process leaves the `O_NONBLOCK` flag on standard input, they clear it.

to the same file structure in the kernel as it does in the shell. Suppose that the program turns on the `O_NONBLOCK` flag on the standard input connection. It is actually modifying the connection that its parent, i.e., the shell, uses as well. In fact, this standard input connection is shared by all related processes – siblings, cousins, and so on. Any process that is sharing this file connection can potentially make changes that affect all other processes that use this terminal. Once the changes are made, unless they are undone before the process terminates, the shell's connection has those changed properties.

Now think about the shell for a moment. A shell is basically running a loop of the form

```
while ( not end of input ) {
    display the shell prompt
    read the user's command line
    carry out the instruction
}
```

Thus, if the program turns on non-blocking reads and is killed before turning it off, when the shell resumes, it executes a read command. Usually the shell is in blocked input mode, so when it tries to read input but none is there, it enters a blocked state waiting for the user to press the Enter key, which sends input to the shell process. Since non-blocking input is still on, instead of waiting for the user's input, it receives an `EOF` from the function call that it uses to retrieve the input. This `EOF` may cause the shell itself to terminate because most shells are killed by the `EOF` character. You should set `ignoreeof` to prevent this. In `bash`, you add the line

```
set -o ignoreeof
```

in your `.bashrc` to do this. Anyway, the result is that your shell is killed, and if this is a login shell, you will be logged out. On my host machine, `bash` gets caught in a segmentation fault, which should not happen.

4. The problem with non-blocking reads causing the shell to exit can also be solved by opening a new file connection to the terminal instead of using standard input. In the demo program, explore the effect of replacing the line

```
fp = stdin;
```

by

```
fp = fopen(ctermid(NULL), "r");
```

5.3 Signals

Signals are, as Richard Stevens once put it, software interrupts. They are a mechanism for handling asynchronous events, such as when a user types `Ctrl-C` at a terminal. Most non-trivial applications need to handle signals. In this section we provide an overview of signals, including what they are, how they are generated, how they are named, and how processes can deal with them.

From a strictly technical point of view, a signal is a message that has a type but does not have content. Messages are usually defined to be containers for data. Signals are not containers. Signals

outside of the computer are things like traffic lights, hand gestures that mean "please stop, taxi driver" or "give me the check", alarm clock rings, or warning lights like "you're about to run out of gas". They do not have contents. It is enough that they have identity, so a particular signal type has well-defined meaning. When a combatant raises a white flag, the enemy knows that this signal means "I give up." In UNIX, a signal is simply an integer with a mnemonic name. For example, `SIGINT` is the interrupt signal.

5.3.1   Typing Ctrl-C at a Terminal

When you type a `Ctrl-C`, the effect is to terminate the currently running process. Why? When you type the `Ctrl-C`, the character code for it is sent by hardware and then, within the kernel, to the terminal's device driver. The device driver checks the character code and sees that it matches the `INTR` code[3]. Since it knows this is a control character that is supposed to cause delivery of an interrupt signal, it checks whether the `isig` attribute is set for the terminal. If `isig` is set, then the corresponding signal must be delivered, which is `SIGINT`, so it calls the signal subsystem of the kernel to notify it to send the `SIGINT` signal to all processes whose control terminal[4] is the one that received the `Ctrl-C`. If any of these processes has not explicitly notified the kernel of how it wants to handle this signal, it will terminate upon receipt of the signal, because by default, processes are killed if they do not catch `SIGINT`. Soon we shall see that a process can declare what the disposition of a delivered signal should be while it is running.

5.3.2   Sources of Signals

All signals are sent by the kernel to processes. There is no other way for a process to receive a signal; the kernel is like the central signal processing station inside the machine. The kernel will send a signal to one or more processes if it receives a request to do so. Requests can come from a few different types of sources.

The terminal A user can type a key combination that causes the terminal driver to ask the kernel to send a signal. This is an asynchronous signal, since it can arrive at a process at any time, independent of what the process might be doing. Examples include `Ctrl-C`, `Ctrl-Z`, `Ctrl-S`.

Hardware Hardware exceptions can generate signals. The kernel detects when the exception occurs and sends a signal to the offending process. These may be synchronous or asynchronous. Synchronous events are things such as floating-point exceptions, illegal instructions, addressing exceptions (such as attempts to access addresses outside of the process's address space), and other events generally caused by the process itself. They are synchronous because if the process is run again, they will occur again at the same point in the process's execution. Asynchronous events are things like power loss and terminal hang-ups.

Software Software conditions can generate signals when something noteworthy happens. This can happen when out-of-band data arrives over a network connection, or when a process writes to a pipe after the reader of the pipe has terminated, or when an alarm clock set by the process expires.

---

[3]In SunOS it is `INTR`. On other systems, it might be `VINTR`.

[4]The operating system keeps track of which processes are attached to which terminals.

Processes Processes themselves can request the kernel to send signals to processes, even them-
selves. This is not so strange; an alarm clock is a way for you to send a signal to yourself
in the future; you set the alarm and wait for it to signal you. Similarly, a process can
ask the kernel to send a wake-up call to itself at some future time. A process can also
ask the kernel to send a signal to other processes to which it has permission to send
signals. For ordinary user processes, these include any processes with the same real or
effective user-id.

5.3.3   Signal Types

UNIX systems define signals in the header file `<signal.h>`. More accurately, the header file
`<signal.h>` includes the header file that contains the signal definitions. Different UNIX systems
store this file in different places. In Linux, the file `<bits/signum.h>` is where the signals are defined.
The kernel includes this header file, but user level program are supposed to include `<signal.h>`.
The idea is to keep separate headers for the kernel and the user-level programs.

The exact set of signals varies from one system to another, but some of them are standard across
all systems. Signal names are just names for small integers such as `SIGINT`, `SIGKILL`, `SIGHUP`, and
`SIGCHLD`. All names begin with the prefix `SIG`. `SIGHUP` is the *hang-up* signal. It is sent to a process
when its control terminal has been disconnected. `SIGCHLD` is the signal sent to a parent by the
kernel when it detects that one of its child processes has terminated. There are typically about 30
to 35 different signals defined in any UNIX system. The list of signals has changed over the years.
The first 30 signals listed below are found in Linux and Solaris 9 ; the last 4 only in Solaris 9. The
constant `NSIG` is the the total number of signals defined. Since the signal numbers are allocated
consecutively, `NSIG` is also one greater than the largest defined signal number.

Name	Value	Default	Event	Note	Category
SIGHUP	1	Exit	Hangup		Termination
SIGINT	2	Exit	Interrupt		Termination
SIGQUIT	3	Core	Quit		Termination
SIGILL	4	Core	Illegal Instruction		Program Error
SIGTRAP	5	Core	Trace or Breakpoint Trap		Program Error
SIGABRT	6	Core	Abort		Program Error
SIGEMT	7	Core	Emulation Trap		Program Error
SIGFPE	8	Core	Arithmetic Exception		Program Error
SIGKILL	9	Exit	Killed		Termination
SIGBUS	10	Core	Bus Error	1	Program Error
SIGSEGV	11	Core	Segmentation Fault		Program Error
SIGSYS	12	Core	Bad System Call	1	Program Error
SIGPIPE	13	Exit	Broken Pipe		Operation Error
SIGALRM	14	Exit	Alarm Clock		Alarm
SIGTERM	15	Exit	Terminated		Termination
SIGUSR1	16	Exit	User Signal 1	1	Miscellaneous
SIGUSR2	17	Exit	User Signal 2	1	Miscellaneous
SIGCHLD	18	Ignore	Child Status Changed	1	Job Control
SIGPWR	19	Ignore	Power Fail or Restart		Hardware
SIGURG	21	Ignore	Urgent Socket Condition		Asynchronous I/O
SIGPOLL	22	Exit	Pollable Event		Asynchronous I/O
SIGSTOP	23	Stop	Stopped (signal)	1	Job Control
SIGTSTP	24	Stop	Stopped (user)	1	Job Control
SIGCONT	25	Ignore	Continued	1	Job Control
SIGTTIN	26	Stop	Stopped (tty input)	1	Job Control
SIGTTOU	27	Stop	Stopped (tty output	1	Job Control
SIGVTALRM	28	Exit	Virtual Timer Expired	1	Alarm
SIGPROF	29	Exit	Profiling Timer Expired	1	Alarm
SIGXCPU	30	Core	CPU time limit exceeded	1	Operation Error
SIGXFSZ	31	Core	File size limit exceeded	1	Operation Error
SIGWINCH	20	Ignore	Window Size Change	2	Miscellaneous
SIGWAITING	32	Ignore	Concurrency signal	3	Miscellaneous
SIGLWP	33	Ignore	Inter-LWP signal	3	Miscellaneous
SIGFREEZE	34	Ignore	Check point Freeze	3	Miscellaneous
SIGTHAW	35	Ignore	Check point Thaw	3	Miscellaneous

Notes

1. In Linux the numerical value of the signal is architecture-dependent.

2. This is only found in Sun OS and BSD.

3. These are in Solaris 9.

The above list has four columns. The first is the mnemonic name for the signal, i.e., the name that can be used in a program. The second is the integer value, which you do not need to know. The third is the default action that happens to a process. For example, SIGCHLD is ignored by default,

`SIGSTOP` causes the program to stop by default, and `SIGINT` causes the process to terminate. The last column indicates the cause or condition that leads to this signal.

5.3.4 Sending Signals

In UNIX, the kernel can send a signal to a process when some hardware error condition arises. For example, if a program attempts to execute an illegal instruction, the kernel will receive the hardware notification and will send the `SIGILL` (illegal instruction signal) to the offending process. A process can also send a signal to one or more processes (or even itself) by using the `kill()` system call. The form of the call is

```
int kill(int processid, int signal);
```

The first parameter stores a means to specify the process id of the process to receive the signal. The second parameter is the kind of signal to send. In the simplest case,

```
kill(942, SIGTERM);
```

sends the `SIGTERM` signal to the process whose process-id is 942. A process cannot send a signal to another process if they do not share the same real or effective user-id1. If a process does not have permission to issue the kill call, `kill()` returns –1.

`processid` can be 0, -1, or a negative number, and it means something different in each case. If `processid` is 0, the signal will be sent to all processes in the same process group, whereas if it is –1, and the sender is not the superuser, it is sent to all processes for which it has permission to send signals, which are those processes with the same real or effective user-id. If `processid` < -1, it is sent to all processes in the process group with id `-processid`.

A process can also send a signal to itself using

```
int raise( int signal);
```

which is equivalent to

```
kill(getpid(), signal);
```

The call to `raise()` will return only when the process has handled the signal.

5.3.5 Signal Generation and Delivery

UNIX systems generally distinguish between the generation of a signal and its delivery. According to the *Open Group Base Specification Issue 6* (IEEE Std 1003.1),

"A signal is said to be ***'generated'*** for (or sent to) a process or thread when the event that causes the signal first occurs. Examples of such events include detection of hardware faults, timer expiration, signals generated via the `sigevent` structure and terminal activity, as well as invocations of the `kill()` and `sigqueue()` functions. In some circumstances, the same event generates signals for multiple processes.

"A signal is ***delivered*** to a process when the appropriate action for that process and signal is taken."

What this means is that delivery takes place when the process receives the signal and responds by either

- *ignoring* it,

- taking the *default* action, or

- executing a *signal handler* for it.

Signal handling is described below. The point to remember now is that from the moment that a signal is generated for a process until the moment that the signal is delivered, the signal is *pending*. Pending signals are managed by the operating system[5].

5.3.6   Signal Handling

A process does not have to accept the default action caused by a signal. It can choose to respond differently to all signals except for `SIGKILL` and `SIGSTOP`. These signals always terminate the process. To handle a signal, the programmer defines a function called a *signal handler*. The signal handler is executed when the signal is received, provided that it has been installed.

The program notifies the operating system that it has a handler for a specific signal by executing a system call to install the handler. The original system call for installing a signal handler was the `signal()` system call. The `signal()` system call was unreliable because it was possible to miss signals when using it. It was replaced by a reliable signal installing call named `sigaction()`. We will explore the `sigaction()` call later. For now, we start with the `signal()` system call, partly because it is an easier one to use, and partly so that you understand its weaknesses. Once you do, you should avoid using it.

The `signal()` call has the form

```
signal( signal_number, handler_action)
```

The first parameter is the number of a signal, but you should always use its mnemonic name such as `SIGINT` or `SIGQUIT`. The second parameter is one of the following:

`SIG_DFL`   Take the default action, which is usually to terminate the process.

`SIG_IGN`   Ignore the signal completely and continue.

user-defined function Address of a user-defined function

---

[5]This discussion of signals overlooks the complexity entailed because of threads and multi-threaded processes. Until we discuss threads in general, we have to overlook this topic. But you should bear in mind that the operating system has to make decisions when signals are generated as to whether they are to be sent to every thread in a process or just to a single thread in particular, and that certain signals must always be sent to one choice or the other.

Examples

The following program, `signal_demo1.c`, shows how signal handlers for `Ctrl-C` (SIGINT) and `Ctrl-\` (SIGQUIT) are installed.

Listing 5.12: signal_demo1.c

```c
#include          <stdio.h>
#include          <signal.h>

void catch1(int signum)
{
    printf("You can do better than that!\n");
}

void catch2(int signum)
{
    printf("I'm no quitter!\n");
}

int main()
{
    int i;

    signal( SIGINT, catch1 );
    signal( SIGQUIT, catch2 );
    for(i = 20; i > 0; i— ) {
        printf("Try to kill me with ^C or ^\\.\n");
        sleep(1);
    }
    return 0;
}
```

The call `signal(signum, f)` installs `f()` as the signal handler for the signal `signum`. When `signal()` is executed, `f()` is installed. Until that point, `f()` is not installed. When you run `signal_demo1` and enter a `Ctrl-C`, the SIGINT signal is sent to the process executing `signal_demo1`; as a result, the handler `f()` runs, and when it terminates, the program resumes execution. In `signal_demo1.c`, the only action taken by either handler is to print a message on the screen, simply to show that the function was executed.

The next program, `signal_demo2.c`, is almost the same as `signal_demo1.c` with one exception: SIGINT and SIGQUIT are ignored by calling `signal()` with `SIG_IGN` as the second argument.

```c
#include          <stdio.h>
#include          <signal.h>

int main()
{
```

```
        signal( SIGINT, SIG_IGN );     // ignore Ctrl-C
        signal( SIGQUIT, SIG_IGN );   // ignore Ctrl-\

        for(i = 20; i > 0; i-- ) {
            printf("Try to kill me with ^C or ^\\.\n");
            sleep(1);
        }
        return 0;
}
```

5.3.7   Putting It Together

We revise `term_demo1.c` so that it handles `Ctrl-C` and `Ctrl-\` interrupts (whether from the keyboard or sent via a `kill()` from another process.) This program is `term_demo2.c`. The listing below shows only the changed portion of the code.

```
....
#include           <signal.h>
....   (snip)

    if ( fflags ) {
        ui_parameters.isblocking = FALSE;
        set_nodelay(fileno(fp), fflags);
    }

    signal(SIGINT, interrupt_handler);
    signal(SIGQUIT, interrupt_handler);

    while ( TRUE ) {
        response = get_response(fp, ui_parameters );
....  (snip)

char * signame( int signo)
{
    static char  name[16];
    switch ( signo ) {
    case   SIGINT:
        strcpy(name, "SIGINT");
        break;
    case   SIGQUIT:
        strcpy(name, "SIGQUIT");
    }
    return name;
}

void interrupt_handler(int signum)
```

```
{
    printf("Exiting with signal %s\n", signame(signum));
    save_restore(fileno(stdin), RESTORE, NULL);
    exit(2);
}
```

The major changes are that the `signal()` function is used to install handlers for `SIGINT` and
`SIGQUIT`. In this program they share the same handler, named `interrupt_handler()`. This handler
prints a message on the standard output and then restores the `termios` structure's flags and the
file status flags to what they were before the program was run.

5.3.8   Weaknesses of the Signal Mechanism

Signals in this form do not carry any information other than their particular values. Therefore, they
are of limited use. They were never intended to be a robust form of communication, and they are
still not completely reliable. The early form of signal handling using the signal() system call was
very unreliable. While a process was in the midst of catching a signal, it was unable to detect the
arrival of another signal of the same type; any new signals of that type were lost. This means that
if two signals of the same type were sent in rapid succession to a process, the second might be lost.
Later versions of the signal() function in BSD and in SVR corrected this problem in different ways,
so that it now has at least two different behaviors. The modern version in Linux 2.6 combines the
semantics of each. It is best to avoid the signal() call for that reason.

5.3.9   Signal Handling The "Right" Way

Signals are a primitive form of inter-process communication by today's standards, but at the time
they were conceived, they provided a simple, efficient method of solving the most important inter-
process communication problems. The `signal()` system call was early UNIX's method of defining
and installing signal handlers. One problem with the `signal()` call is that it needs to be reset each
time, like a mouse trap – once it catches a signal, arriving signals are missed. Another problem
with `signal()` is that its behavior was left unspecified in the case when multiple signals arrived,
and different implementations of UNIX provided different semantics to handle multiple signals.

5.3.10   Multiple Signals

Suppose that a signal handler is in the midst of handling a signal that has been delivered when a
second signal is generated and is pending. There are a few possible ways to dispose of this new
signal:

- Ignore it completely, effectively losing the new signal;

- Put it in a queue and handle it when the current signal has been handled completely, effectively
  blocking pending signals while handling the current one;

- Interrupt the processing of the current signal, handle the new signal, and return to the old
  signal when the new one has been handled, effectively treating the handler like an involuntary
  recursive function;

In any one UNIX system, you might have found one of these solutions employed rather than the others, without any consistency. The POSIX standard introduced a uniform solution to the problem in the `sigaction()` interface.

5.3.11   The sigaction() call

The `sigaction()` system call allows a process to install a signal handler and to specify how it will respond to multiple arriving signals. Its prototype is

```
#include <signal.h>
int (sigaction (int signum, const struct sigaction *act,
struct sigaction *oldaction);
```

where

signum   is the value of the signal to be handled

act   is a pointer to a `sigaction` structure that specifies the handler, masks, and flags for the
            signal

oldact   is a pointer to a structure to hold the currently active `sigaction` data.

We will examine the `sigaction` structure first to see how flexible this interface is. Notice that the function name is the same as the name of the structure whose address is passed to it, like the `stat()` function and the `stat` structure.

5.3.12   The sigaction struct

The `sigaction` structure is defined in `<signal.h>`. The definition is unusual because it has two members (`sa_handler` and `sa_sigaction`) that are allowed to overlap in memory and must be used in mutual exclusion. The simplest way to present it is as if it were two different overloaded definitions of the same structure:

```
struct sigaction                  // backward-compatible, old-style handler
{
    void (*sa_handler) (int); // the action to take
    sigset_t sa_mask;             // additional signals to block
                                  // during handling of the signal
    int  sa_flags;                // flags that affect behavior
};
```

or

```
struct sigaction              // POSIX compliant, new-style handler
{
void (*sa_sigaction) (int, siginfo_t *, void *);
```

```
                                        // pointer to signal handler
    sigset_t sa_mask;                   // additional signals to block
                                        // during handling of the signal
    int  sa_flags;                      // flags that affect behavior
    };
```

In other words, there are two different forms of the `sigaction` structure. The first one uses the old-style of handler, and the second uses the newer POSIX compliant method. The structs are otherwise identical.

Notes

In the old-style, the `sa_handler` member does not have to be a pointer to a function. It can also be one of the two flags `SIG_IGN` or `SIG_DFL`. If it is `SIG_IGN`, the signal is ignored; if `SIG_DFL`, then the default action is taken. If a pointer to a handler is supplied, that handler will be run. The handler must have a single integer argument.

In the new style, the `sa_handler` is replaced by a pointer to a function that has three parameters as follows:

- An `sa_mask`, which defines which signals should be blocked while the handler is processing the signal. By default, the signal that caused the handler to run will always be blocked. Adding signals to `SA_MASK` is a way to block other signals as well.

- The `sa_flags` is a flag set that can be used to control how subsequent signals of the same type as the one that caused the handler to run are handled. For example, if a handler is handling a `SIGINT` signal and another `SIGINT` arrives while the process is in the handler, the `sa_flags` will determine how to dispose of the second `SIGINT`. It has no effect on other arriving signals. The flags determine how arriving signals are handled after the handler has been invoked. All flags are independent and `sa_flags` is their bitwise-or. The most important flags are:

    SA_RESETHAND If this bit is set, the signal action is reset to `SIG_DFL`. This means that as soon as the signal is delivered, the default action will take place. This flag implies the `SA_NODEFER` flag; signals are not blocked, instead causing the process to take whatever is the default action for the type of signal. The intention is to make the handler behave like the old-style `signal()` handler, since any signal arriving after the first will cause the default behavior.

    SA_NODEFER If this bit is set, the kernel will not automatically block the signal while it is being caught. This means that an arriving signal will cause the handler itself to be interrupted and re-entered with the second signal. This is involuntary recursion.

    SA_RESTART If this bit is set, certain system calls that would otherwise be terminated if a signal were delivered during their execution, will be automatically restarted. This bit allows the BSD style handling.

    SA_SIGINFO If this bit is set, two additional arguments are passed to the signal-catching function. If the second argument is not `NULL`, it points to a `siginfo_t` structure containing the reason why the signal was generated; the third argument points to a `ucontext_t` structure containing the receiving process's context when the signal was delivered.

Note. If multiple instances of an individual signal are delivered while that signal is currently blocked, then only one instance is queued. For example, if you issue multiple `SIGINT` signals to a process, without setting the `SA_NODEFER` flag in the handler, then the first one will cause the handler to run, the second will be queued, but all those after that will be lost.

5.3.12.1   Example

The first example, sigaction_demo1.c, shows how the `SA_SIGINFO` flag can be used. A signal handler for `SIGINT` is installed with the `SA_SIGINFO` flag set. The program delays itself using `sleep(60)` so that the user has time to enter a `Ctrl-C`. When the `Ctrl-C` is entered and the program is delivered a `SIGINT` signal, the handler runs and accesses the `siginfo_t` structure to print the values of its members as a result of the signal.

Listing 5.13: sigaction_demo1.c

```c
#include <unistd.h>
#include <sys/types.h>
#include <signal.h>
#include <bits/siginfo.h>
#include <stdio.h>
#include <stdlib.h>

void sig_handler (int signo, siginfo_t *info,  void *context)
{
    printf ("Signal number: %d\n", info->si_signo);
    printf ("Error number:  %d\n", info->si_errno);
    printf ("PID of sender: %d\n", info->si_pid);
    printf ("UID of sender: %d\n", info->si_uid);
    exit (1);
}

int main (int argc, char* argv[])
{
    struct sigaction the_action;

    the_action.sa_flags = SA_SIGINFO;
    the_action.sa_sigaction = sig_handler;

    sigaction (SIGINT, &the_action, NULL);

    printf ("Type Ctrl-C within the next minute"
            "or send signal 2. \n");
    sleep (60);
    return 0;
}
```

5.3.12.2   Example

The following program will demonstrate the use of the `sigaction` structure for old-style handlers
as well as the new-style handlers. The program begins by defining the `sigaction` structure that
will be passed to the `sigaction` function. The handler function, `interrupt_handler`, is assigned
to the `sa_handler` field. The `sa_flags` field is initialized with the bitwise-or of the flags referenced
in the command-line. The `sa_mask` is built step by step. There are a few ways to do this. You can
start with an empty set and add to it or start with a full set and remove from it. In this case, I am
allowing the user to interactively add signal numbers to the blocked set, so I start with an empty
set and add to it. The method of creating these sets is discussed below.

Listing 5.14: sigaction_demo2.c

```
/*
 * Usage:
 *          sigaction_demo  [ reset | noblock | restart ]*
 *
 * i.e., 0 or more of the words reset, noblock, and restart may appear
 * on the command line, and multiple instances of the same word as the same
 * effect as a single instance.
 *
 * reset    turns on SA_RESETHAND
 * noblock turns on SA_NODEFER
 *  * restart turns on SA_RESTART
 *
 * NOTES
 * (1) If you supply the word "reset" on the command line it will set the
 *     handling to SIG_DFL for signals that arrive when the process is
 *     in the handler. If noblock is also set, the signal will have the
 *     default behavior immediately. If it is not set, the default will
 *     delay until after the handler exits. If noblock is set but reset is
 *     not, it will recursively enter the handler.
 * (2) The interrupt_handler purposely delays for a few seconds in order to
 *     give the user time to enter a few interrupts on the keyboard.
 * (3) interrupt_handler is the handler for both SIGINT and SIGQUIT, so if it
 *     is not reset, neither Ctrl-C nor Ctrl-\ will kill it .
 * (4) It will ask you to enter the numeric values of signals to block. If
 *     you don't give any, no signals are blocked.
 *
 */


#include    <stdio.h>
#include    <string.h>
#include    <signal.h>
#include    <time.h>
#include    <fcntl.h>

#define INPUTLEN         100

int main(int ac, char* av[])
{
    struct sigaction newhandler;        /* new settings         */
    sigset_t       blocked;             /* set of blocked sigs */
    void           inthandler();        /* the handler          */
```

```
char            x[INPUTLEN];
int             flags = 0;
int             signo, n;
char            s[] = "Entered text: ";
int             s_len = strlen(s);


while ( 1 < ac ) {
    if (0 == strncmp("reset", av[ac-1], strlen(av[ac-1])) )
        flags |= SA_RESETHAND;
    else if (0 == strncmp("noblock", av[ac-1], strlen(av[ac-1])) )
        flags |= SA_NODEFER;
    else if (0 == strncmp("restart", av[ac-1], strlen(av[ac-1])) )
        flags |= SA_RESTART;
    ac--;
}
/* load these two members first */
newhandler.sa_sigaction = interrupt_handler;/* handler function    */
newhandler.sa_flags = SA_SIGINFO | flags;    /* new style handler   */

/* then build the list of blocked signals */
sigemptyset(&blocked);                       /* clear all bits      */

printf("Type the numeric value of a signal to block (0 to stop):");
while ( 1 ) {
    scanf( "%d", &signo );
    if ( 0 == signo )
        break;
    sigaddset(&blocked, signo);          /* add signo to list */
    printf("next signal number (0 to stop): ");
}
newhandler.sa_mask = blocked;            /* store blockmask      */

// install inthandler as the SIGINT handler
if ( sigaction(SIGINT, &newhandler, NULL) == -1 )
    perror("sigaction");

// if successful, install inthandler as the SIGQUIT handler also
else if ( sigaction(SIGQUIT, &newhandler, NULL) == -1 )
    perror("sigaction");
else
    while( 1 ) {
        x[0] = '\0';
        tcflush(0,TCIOFLUSH);
        printf("Enter text then <RET>: (quit to quit)\n");
        n = read(0, &x, INPUTLEN);
        if ( n == -1 && errno == EINTR ) {
            printf("read call was interrupted\n");
            x[n] = '\0';
            write(1, &x, n+1);
        }
        else if ( strncmp("quit", x, 4) == 0 )
            break;
        else {
            x[n] = '\0';
```

```
                        write (1 , &s , s_len );
                        write (1 , &x , n+1);
                        printf ("\n");
                }
        } // while
    return 0;
}


void interrupt_handler (int signo , siginfo_t *info , void *context)
{
    int       localid ;     /* stores a number to uniquely identify signal */
    time_t    timenow ;     /* current time -- used to generate id */
    static    int ticker = 0;  /* used for id also */
    struct tm *tp ;

    time(&timenow );
    tp = localtime (&timenow );
    localid = 36000*tp->tm_hour + 600*tp->tm_min + 10*tp->tm_sec +
                ticker++ % 10;
    printf ("Entered handler : sig = %d \tid = %d\n",
            info->si_signo , localid );
    sleep (3);
    printf ("Leaving handler : sig = %d \tid = %d\n",
            info->si_signo , localid );
}
```

The while loop in the main program exists just to demonstrate that the system calls to perform
I/O are restarted when the interrupts occur. You can run this program with any combination
of SA_RESTART, SA_RESETHAND, and SA_NODEFER to see the combined effect of the flags. You can
add any signal to the blocked set, but you will only be able to send SIGINT and SIGQUIT unless
you open a separate window and use the kill command to send arbitrary signals to the process.
You can try this – put signal 4 in the blocked set and issue kill -4 with the process id from a
second window while the process is handling a Ctrl-C. You will see that signal 4 is blocked until
interrupt_handler() finishes.

The while loop is designed so that you do not have to kill this program to terminate it. If you type
"quit" it will terminate.

5.3.13   Creating Signal Mask Sets

There are four functions that modify signal mask sets:

```
    sigemptyset(sigset_t *setp);
    sigfillset(sigset_t *setp);
    sigaddset(sigset_t *setp, int signum);
    sigdelset(sigset_t *setp, int signum);
```

The first two create empty and full mask sets respectively. The next two add or delete specific
signals from the specified sets. You can either create an empty mask and add to it, or create a full
mask and delete from it. If you plan on having more than half of the signals in it, then do the
latter, otherwise do the former.

5.3.14   Blocking Signals Temporarily around Critical Sections

The `sigprocmask()` system call can be used to block or unblock signals sent to a process. This is useful if you need to temporarily turn off all signals in a small section of code. This is one way to create something like a critical section, in the sense that the process will not be interrupted in the middle of the code fragment. It does not prevent the kernel from preempting the process and letting another process run on the CPU, so it does not provide a means of protecting shared variables that might be accessed while the process is removed from the CPU, but it does allow the process to complete some critical sequence of statements without any signal handlers running in the interim, and without being terminated in the midst of it. The prototype is

```
int  sigprocmask( int how, const sigset_t *sigs, sigset_t *prev);
```

where `how` is one of `SIG_BLOCK`, `SIG_UNBLOCK`, or `SIG_SETMASK`. `SIG_BLOCK` will block the specified signal set, i.e., prevent those signals from reaching until the procmask is changed. `SIG_UNBLOCK` allows the signals in the set to be unblocked, and `SIG_SETMASK` is used to change the mask completely, i.e., assign a new mask to the procmask. The following program can be run to demonstrate how the blocking works. If you type `Ctrl-C` during the first loop, the process will continue to loop and it will exit before the second loop is executed. If you change the `SIG_BLOCK` to `SIG_UNBLOCK` then the `Ctrl-C` will kill the process when you type it.

Listing 5.15: procmask_demo.c

```c
#include <signal.h>
#include <stdio.h>

int main()
{
    int i;
    sigset_t sigs, prevsigs;

    sigemptyset(&sigs);
    sigaddset(&sigs, SIGINT);
    sigprocmask(SIG_BLOCK, &sigs, &prevsigs);

    for ( i = 0; i < 5; i++) {
        printf("Waiting %d\n", i);
        sleep(1);
    }
    sigprocmask(SIG_SETMASK, &prevsigs, NULL);
    for ( i = 0; i < 5; i++) {
        printf("After %d\n", i);
        sleep(1);
    }
}
```

5.4   Summary

An understanding of terminals and signal handling is essential to being able to write UNIX System applications. You now have almost all of the tools at your disposal to write terminal-based interactive programs that can use the full terminal window and allow the user to interact with it at will. As you will discover in the next chapter, you are still missing a few more pieces. For one, we still need to manage timing more accurately and allow events to happen at specific times as controlled by the program. For another, we do not have the means of placing characters anywhere on the screen. This is what comes next.

The entire simplevi.c program is listed below. Many comments have been deleted to reduce the space.

Listing 16: simplevi.c

```c
#include <stdio.h>
#include <ctype.h>
#include <stdlib.h>
#include <termios.h>
#include <fcntl.h>
#include <time.h>
#include <string.h>
#include <unistd.h>
#include <math.h>
#include <getopt.h>
#ifndef TIOCGWINSZ
#include <sys/ioctl.h>
#endif

#define RETRIEVE        1
#define RESTORE         2
#define FALSE           0
#define TRUE            1
#define ESCAPE          '\033'
#define CONTROL_C       '\003'
#define CONTROL_D       '\004'
#define CONTROL_H       '\010'
#define KEY_UP          65
#define KEY_DOWN        66
#define KEY_RIGHT       67
#define KEY_LEFT        68
#define MAXLINES        1000
#define MAXCHARS        255
#define OUT_OF_LINES   -1
#define OUT_OF_MEM     -2
#define UNHANDLEDCHAR  -3

const  char CLEAR_DOWN[]    = "\033[0J";
const  int  lCLEAR_DOWN     =   4;
const  char CLEAR_RIGHT[]   = "\033[0K";
const  int  lCLEAR_RIGHT    =   4;
const  char CURSOR_HOME[]   = "\033[1;1H";
const  int  lCURSOR_HOME    =   6;
const  char CLEAR_SCREEN[]  = "\033[2J";
```

```c
const  int   lCLEAR_SCREEN   =   4;
const  char  CLEAR_LINE[]    = "\033[2K";
const  int   lCLEAR_LINE     =   4;
const  char  RIGHT[]         = "\033[1C";
const  int   lRIGHT          =   4;
const  char  LEFT[]          = "\033[1D";
const  int   lLEFT           =   4;
const  char  BACKSPACE[]     = "\033[1D \033[1D";
const  int   lBACKSPACE      =   9;
char         PARK[20];           /* string to park cursor at lower left */
int          lPARK;              /* length of PARK string */


/* Miscellaneous strings for output */
const  char  CTRLC[]         = "You typed Control-C.";
const  char  CTRLD[]         = "You typed Control-D.";
const  char  CTRLH[]         = "This is the Help Command. Not much help, sorry!";
const  char  BLANK           = ' ';
const  char  INSERT[]        = "--INSERT--";
const  int   lINSERT         = 10;
const  char  MAXLINES_MSSGE[]
                             = "You reached the maximum number of lines."
                               " Exiting input mode.";
const  char  OUT_OF_MEM_MSSGE[]
                             = "You reached the maximum buffer size."
                               " Exiting input mode.";
const  char  UNHANDLEDCHAR_MSSGE[]
                             = "This input not yet implemented."
                               " Exiting input mode.";

inline void clearandhome()
{
    write(1,CLEAR_SCREEN, lCLEAR_SCREEN);
    write(1,CURSOR_HOME, lCURSOR_HOME);
}

inline void park()
{
    write(1,PARK, lPARK);
    write(1, CLEAR_LINE, lCLEAR_LINE);
}


/**************************************************************************/
/*                          Data Types                                  */
/**************************************************************************/

typedef struct _cursor
{
    int r;
    int c;
} Cursor;


typedef struct _window
{
```

```
    unsigned short rows;
    unsigned short cols;
    int            line_at_top;
    char           erase_char;
} Window;


typedef struct _buffer
{
    char text[BUFSIZ];
    int  line_len[MAXLINES];  /* lengths of text lines, including newline
                                 characters */
    int  line_start[MAXLINES]; /* starts of each line */
    int  num_lines;           /* number of text lines in buffer. This
                                 includes lines that have not yet been
                                 terminated with a newline character. It is
                                 the number of newline characters + 1 if the
                                 last character in the buffer is not a
                                 newline. */
    int  index;               /* index of cursor in text buffer */
    int  size;                /* total chars in buffer */
    int  cur_line;            /* current text line, not screen line */
    int  index_in_cur_line;   /* index in current line of cursor */
} Buffer;

/***************************************************************************/
/*                          Function Prototypes                            */
/***************************************************************************/


/* Window/Terminal Functions */
void init_window          ( int fd,  Window *win );
void moveto               ( int line, int column );
void write_status_message( const char *message, Cursor curs );
void save_restore_tty(int fd, int action);
void modify_termios( int fd, int echo, int canon);
void set_erase_char(int termfd, Window *win );

/* Buffer Functions */
int  insert( Buffer *buf, Window win, char ch );
void init_buffer( Buffer *buffer);
void update_buffer_index( Buffer *buffer );
int  buffer_index( int index_in_line, int cur_line, int linelength[] );
void redraw_buffer( Buffer buffer,  Window *win, Cursor *curs );
void scroll_buffer( Buffer buf, Window win );
int  line_in_buffer( Buffer buf, Window win, int pos );
void save_buffer ( const char path[], Buffer buf, char *statusstr );
int  handle_insertion( Buffer *buf, Window *win, Cursor *curs, char c);
void get_lastline_in_win( Buffer buffer, Window win,  int *lastline );

/* Lastline Mode */
int parselastline( char *str, int len, Buffer buf, char *statusstr );
int do_lastline_mode(Buffer buf, Window win, Cursor curs );
```

```
/* Cursor Functions */
void init_cursor    ( Cursor *cursor );
void show_cursor    ( Buffer buf, Window win, Cursor cursor,
                        int index_in_line, int line_number );
void advance_cursor( Cursor *cursor, Window win, char ch );
void get_cursor_at (Buffer buf, Window win, int index,
                    int lineno, Cursor *curs);
void handle_escape_char( Buffer *bin, Window *win, Cursor *curs );
void moveUp    ( Buffer *buf, Window *win, Cursor *curs );
void moveDown  ( Buffer *buf, Window *win, Cursor *curs );
void moveRight( Buffer *buf, Window *win, Cursor *curs );
void moveLeft  ( Buffer *buf, Window *win, Cursor *curs );


/****************************************************************************/
/*                                Main                                      */
/****************************************************************************/

int main(int argc, char* argv[])
{
    int      quit             = 0;
    int      in_input_mode    = 0;
    int      in_lastline_mode = 0;
    Buffer   buf;
    Window   win;
    Cursor   curs;                      /* cursor position (0,0) is upper left */
    char     prompt = ':';              /* prompt character */
    char     c;
    int status;

    if ( !isatty(STDIN_FILENO) || !isatty(STDOUT_FILENO) ) {
        fprintf(stderr, "Not a terminal\n");
        exit(1);
    }

    save_restore_tty(STDIN_FILENO, RETRIEVE);
    modify_termios( STDIN_FILENO, 0, 0);
    init_buffer(&buf);
    init_cursor(&curs);
    init_window(STDIN_FILENO, &win);

    sprintf(PARK, "\033[%d;1H", win.rows);
    lPARK = strlen(PARK);

    /* Clear the screen and put cursor in upper left corner */
    clearandhome();

    while ( !quit ) {
        if ( in_input_mode ) {
            if ( read(STDIN_FILENO, &c, 1) > 0  ) {
                if ( c == ESCAPE ) {
                    in_input_mode = 0;
                    write_status_message( "            ", curs );
                }
```

```
            else {
                /* insert typed char and echo it */
                in_input_mode = handle_insertion( &buf, &win, &curs, c);
                if ( in_input_mode == UNHANDLEDCHAR )
                    in_input_mode = 1;
                else
                    write_status_message( INSERT, curs);
            }
        }
    }
    else {
        if ( read(STDIN_FILENO, &c, 1) > 0  ) {
        switch ( c ) {
        case 'i':
            in_input_mode = 1;
            park();
            update_buffer_index(&buf);
            moveto(curs.r, curs.c);
            write_status_message( INSERT, curs);
            break;
        case ':':
            in_lastline_mode = 1;
            park();
            write(1,&prompt, 1);
            status = do_lastline_mode( buf, win, curs);
            if ( status >= 0 ) quit = status;
            moveto(curs.r, curs.c);
            break;
        case '\003':
            write_status_message( CTRLC, curs);
            break;
        case '\004':
            write_status_message( CTRLD, curs);
          break;
        case '\010':
            write_status_message( CTRLH, curs);
            break;
        case 'j':
            moveDown(&buf, &win, &curs );
            break;
        case 'k':
            moveUp(&buf, &win, &curs );
            break;
        case 'l':
        case ' ':
            moveRight(&buf, &win, &curs );
            break;
        case 'h':
            moveLeft(&buf, &win, &curs );
            break;
        case ESCAPE:
            handle_escape_char( &buf, &win, &curs);
            break;
        default:
```

```
                        if ( c == win.erase_char )
                            moveLeft(&buf, &win, &curs );
                    }
                }
            }
        }
        printf("\n");
        tcflush(STDIN_FILENO,TCIFLUSH);
        clearandhome();
        save_restore_tty(STDIN_FILENO, RESTORE);
        return 0;
}
/**************************************************************************/

int parselastline( char *str, int len, Buffer buf, char *statusstr  )
{
        int i           = 0;
        int foundquit   = 0;
        int foundwrite = 0;
        int badchar     = 0;
        int done        = 0;
        char *filename = NULL;
        int state;

        state = 1;
        while ( ( i < len) && !done ) {
            switch ( state ) {
            case 1:
                if ( str[i] == ' ' )
                    state = 1;
                else if ( str[i] == 'w' ) {
                    foundwrite = 1;
                    state = 2;
                }
                else if ( str[i] == 'q' ) {
                    foundquit = 1;
                    state = 7;
                }
                else
                    state = 5;
                i++;
                break;
            case 2:
                if ( str[i] == 'q' ) {
                    foundquit = 1;
                    state = 3;
                }
                else if ( str[i] == ' ' )
                    state = 4;
                else
                    state = 5;
                i++;
                break;
            case 3:
```

```
                    if  (  str [ i ]  ==  ' '  )
                            state  =  4;
                    else
                            state  =  5;
                    i++;
                    break ;
              case  4:
                    if  (  str [ i ]  ==  ' '  )
                            state  =  4;
                    else  if  (  isalnum ( str [ i ])  ||  str [ i ]  ==  '_'  )  {
                            filename  =  &( str [ i ]);
                            state  =  6;
                    }
                    else
                            state  =  5;
                    i++;
                    break ;
              case  5:
                    badchar  =  1;
                    sprintf ( statusstr ,  "\033[7m: %s  Not  an  editor  command\033[27m",  str );
                    return  −1;
              case  6:
                    filename  =  &( str [ i −1]);
                    while  (  (  i  <  len )  &&  ( isalnum ( str [ i ])  ||  str [ i ]  ==  '_'  )  )
                            i++;
                    str [ i ]  =  '\0 ';
                    done  =  1;
                    break ;
              case  7:
                    if  (  str [ i ]  ==  ' '  )
                            state  =  7;
                    else  {
                            badchar  =  1;
                            sprintf ( statusstr ,
                                    "\033[7m: %s  Not  an  editor  command\033[27m",  str );
                            return  −1;
                    }
                    i++;
              }
        }
        if  (  foundwrite  )  {
              if  (  filename  !=  NULL  )
                    save_buffer ( filename ,  buf ,  statusstr  );
              else  {
                    sprintf ( statusstr ,
                            "\033[7m: %s  Not  an  editor  command\033[27m",  str );
                    return  −1;
              }
        }
        if  (  foundquit  )
              return  1;
        else
              return  0;
}
```

```c
/***************************************************************************/

int do_lastline_mode(Buffer buf, Window win, Cursor curs )
{
    char tempstr[MAXCHARS];
    char statusstr[MAXCHARS];
    char c;
    int i = 0;
    int status;
    int in_lastline_mode = 1;

    while ( in_lastline_mode ) {
        read(STDIN_FILENO, &c, 1);
        if ( c == '\n' ) {
            tempstr[i] = '\0';
            status = parselastline(tempstr, strlen(tempstr),
                                    buf, statusstr);
            in_lastline_mode = 0;
            write_status_message( statusstr, curs);
            statusstr[0] = '\0';
        }
        else if ( c == win.erase_char ) {
            write(1,BACKSPACE, lBACKSPACE );
            if ( i > 0 )
                i--;
            else
                in_lastline_mode = 0;
        }
        else {
            tempstr[i++] = c;
            write(1,&c,1);
        }
    }
    return status;
}
/***************************************************************************/

void handle_escape_char( Buffer *buf, Window *win, Cursor *curs )
{
    char c;

    read(STDIN_FILENO, &c, 1);
    if ( c == 91 ) {
        read(STDIN_FILENO, &c, 1);
        switch ( c ) {
        case KEY_UP:
            moveUp(buf, win, curs );
            break;
        case KEY_DOWN:
            moveDown(buf, win, curs );
            break;
        case KEY_RIGHT:
            moveRight(buf, win, curs );
            break;
```

```
            case KEY_LEFT:
                moveLeft(buf, win, curs );
                break;
        }
    }
}
/*****************************************************************************/

void set_erase_char(int termfd, Window *win )
{
    struct termios cur_tty;
    tcgetattr(termfd, &cur_tty);

    win->erase_char = cur_tty.c_cc[VERASE];
}
/*****************************************************************************/

void modify_termios(int fd, int echo, int canon )
{
    struct termios cur_tty;
    tcgetattr(fd, &cur_tty);


    if ( canon )
        cur_tty.c_lflag       |= ICANON;
    else
        cur_tty.c_lflag       &= ~ICANON;
    if ( echo )
        cur_tty.c_lflag       |= ECHO;
    else
        cur_tty.c_lflag       &= ( ~ECHO & ~ECHOE);

    cur_tty.c_lflag       &= ~ISIG;
    cur_tty.c_cc[VMIN]    =   1;
    cur_tty.c_cc[VTIME]   =   0;

    tcsetattr(fd, TCSADRAIN, &cur_tty);
}
/*****************************************************************************/

void save_restore_tty(int fd, int action)
{
    static struct termios original_state;
    static int            retrieved = FALSE;

    if ( RETRIEVE == action ){
        retrieved = TRUE;
        tcgetattr(fd, &original_state);
    }
    else if (retrieved && RESTORE == action ) {
        tcsetattr(fd, TCSADRAIN, &original_state);
    }
    else
        fprintf(stderr, "Illegal action to save_restore_tty().\n");
```

```
}

/***********************************************************************
                            Window  Functions
***********************************************************************/

void   init_window( int fd, Window *win )
{
    struct winsize size;

    if (ioctl(fd, TIOCGWINSZ, &size) < 0) {
        perror("TIOCGWINSZ error");
        return;
    }
    win->rows = size.ws_row;
    win->cols = size.ws_col;
    win->line_at_top = 0;
    set_erase_char( fd, win);
}
/***********************************************************************/

void write_status_message( const char *message, Cursor curs )
{
    write(1,PARK, lPARK);
    write(1, CLEAR_LINE, lCLEAR_LINE);
    write(1, message, strlen(message));
    moveto(curs.r, curs.c);
}
/***********************************************************************/

void moveto(int line, int column )
{
    char seq_str[20];

    sprintf(seq_str, "\033[%d;%dH", line+1, column+1);
    write(1, seq_str, strlen(seq_str));
}


/***********************************************************************
                            Buffer  Functions
***********************************************************************/

void init_buffer( Buffer *buffer)
{
    buffer->num_lines         = 0;
    buffer->cur_line          = 0;
    buffer->line_len[0]       = 0;
    buffer->line_start[0]     = 0;
    buffer->size              = 0;
    buffer->index_in_cur_line = 0;
    buffer->index             = 0;
}
/***********************************************************************/
```

```c
void save_buffer ( const char path[] , Buffer buf , char *statusstr )
{
    char newline = '\n';
    int fd;
    //char statusstr[80];

    fd = creat(path, 0644);
    if ( fd != -1 ) {
        sprintf(statusstr,
          "\"%s\" %dL %dC written",
          path, buf.num_lines+1, buf.size );
        write(fd, buf.text, strlen(buf.text));
        if ( buf.text[buf.size-1] != '\n' )
            write(fd, &newline, 1);
        close(fd);
    }
    else
        exit(1);
}
/*****************************************************************************/

void update_buffer_index( Buffer *buffer )
{
    int totalchars = 0;
    int i = 0;

    while ( i < buffer->cur_line ) {
        totalchars += buffer->line_len[i];
        i++;
    }
    totalchars += buffer->index_in_cur_line;
    buffer->index = totalchars;
}

/*****************************************************************************/

void get_lastline_in_win( Buffer buffer, Window win, int *lastline )
{
    int totallines = 0;
    int i;
    int max_possible = win.rows-1; /* rows less status line */

    i = win.line_at_top;
    while ( i < buffer.num_lines ) {
        if ( buffer.line_len[i] <= win.cols )
            totallines++;
        else
            totallines += (int) ceil((double)buffer.line_len[i]/win.cols);
        if ( totallines > max_possible )
            break;
        else {
            i++;
        }
    }
```

```
        // total > max_possible, so we use previous line, which is i-1.
        *lastline = i-1;
}
/********************************************************************************/

void redraw_buffer( Buffer buffer, Window *win, Cursor *curs )
{
    int i;
    int lastline;
    int lastchar;
    int firstchar;
    int line_of_cursor;

    /* Compute the last visible complete text line in the buffer */
    get_lastline_in_win( buffer, *win, &lastline );

    /* Get the index of the text line containing the insertion position */
    line_of_cursor = line_in_buffer( buffer, *win, buffer.index ) ;

    /* Check if the window needs to be scrolled */
    if (line_of_cursor < win->line_at_top ) {
        lastline -= (win->line_at_top - line_of_cursor);
        curs->r += (win->line_at_top - line_of_cursor);
        win->line_at_top = line_of_cursor;
    }
    else if (line_of_cursor > lastline ) {
        win->line_at_top += (line_of_cursor - lastline);
        curs->r -= (line_of_cursor - lastline);
        lastline = line_of_cursor;
    }

    lastchar  = buffer.line_start[lastline] + buffer.line_len[lastline];
    firstchar = buffer.line_start[win->line_at_top];

    /* Prepare to redraw the window. First clear the screen. */
    write(1,CLEAR_SCREEN, lCLEAR_SCREEN);

    /* Do the redraw */
    moveto(0,0);
    for ( i = firstchar; i < lastchar; i++ )
        write(1,&buffer.text[i], 1 );
}
/********************************************************************************/

void scroll_buffer( Buffer buf, Window win )
{
    /*
        This calculates the position of the first character on the screen
        as the leftmost character in the current line_at_top, and then
        calls get_lastline_in_win() to get the index of the last text line
        that can fit in its entirety within the window. It then computes
        the index of the last character in that line.

        It then clears the screen and writes the contents of the text buffer,
```

```
            starting at the computed firstchar until the lastchar. The cursor
            has to be moved to the upper left-hand corner before starting.
            The caller is responsible for restoring the previous cursor position.
    */

    int i;
    int lastline;
    int lastchar;
    int firstchar = buf.line_start[win.line_at_top];

    get_lastline_in_win( buf, win, &lastline );
    lastchar = buf.line_start[lastline] + buf.line_len[lastline];

    write(1,CLEAR_SCREEN, lCLEAR_SCREEN);
    moveto(0,0);
    for ( i = firstchar; i < lastchar; i++ )
        write(1,&buf.text[i], 1 );
}
/************************************************************************/

int line_in_buffer( Buffer buf, Window win, int pos )
{
    int i = 0;

    while ( i < buf.num_lines )
        if ( buf.line_start[i] <= pos )
            i++;
        else
            break;
    /* If the inserted character is a newline, add
        the extra line */
    if ( ( buf.text[pos] == '\n' )  )
        i++;
    return i-1;
}
/************************************************************************/

int insert( Buffer *buf, Window win, char c )
{
    int i;

    if ( ( c == '\n' ) && ( MAXLINES == buf->num_lines ) )
        return OUT_OF_LINES;
    else if ( buf->size == BUFSIZ )
        return OUT_OF_MEM;

    if ( c == win.erase_char  )
        return UNHANDLEDCHAR;

    for ( i = buf->size; i > buf->index; i-- )
        buf->text[i] = buf->text[i-1];
    buf->text[buf->index] = c;
    buf->size++;
    buf->index++;
```

```
        buf->line_len[buf->cur_line]++;

        /* the first character sets line count to 1 */
        if ( buf->size == 1 ) buf->num_lines++;

        if ( c == '\n' ) {
            int temp = buf->line_len[buf->cur_line];

            buf->line_len[buf->cur_line] = buf->index_in_cur_line + 1;
            buf->num_lines++;

            for ( i = buf->num_lines-1; i > buf->cur_line+1; i-- ) {
                buf->line_len[i] = buf->line_len[i-1];
                buf->line_start[i] = buf->line_start[i-1]+1;
            }
            buf->line_start[buf->cur_line+1] = buf->line_start[buf->cur_line]
                                    + buf->line_len[buf->cur_line];

            buf->cur_line++;
            buf->line_len[buf->cur_line] = temp - buf->line_len[buf->cur_line-1];
            buf->index_in_cur_line = 0;
        }
        else if ( isprint(c) ) {   /* non-newline character */
            buf->index_in_cur_line++;    // advance index in line
            /* increment all line starts after this line */
            for ( i = buf->cur_line+1; i < buf->num_lines; i++ )
                buf->line_start[i]++;
        }
        else
            return UNHANDLEDCHAR;

        return 0;
}

int handle_insertion( Buffer *buf, Window *win, Cursor *curs, char c)
{
        int retvalue;

        /* insert typed char and echo it */
        retvalue = insert(buf, *win,  c);
        if ( retvalue < 0 ) {
            if ( retvalue == OUT_OF_LINES )
                write_status_message( MAXLINES_MSSGE, *curs );
            else if ( retvalue == OUT_OF_MEM )
                write_status_message( OUT_OF_MEM_MSSGE, *curs );
            else if ( retvalue == UNHANDLEDCHAR ) {
                write_status_message( UNHANDLEDCHAR_MSSGE, *curs );
                return retvalue;
            }
            return 0;
        }
        advance_cursor( curs, *win, c );
        redraw_buffer( *buf, win, curs );
        moveto(curs->r, curs->c);
```

```
        return  1;
}

/************************************************************************
                             Cursor  Functions
 ***********************************************************************/

void  get_cursor_at ( Buffer buf , Window win , int index , int lineno ,
                        Cursor *curs )
{
    int  total_lines_before = 0;
    int  rows_in_current_textline = 0;
    int  i ;

    for  (  i = win.line_at_top ;  i < lineno ;  i++ )  {
        if  (  buf.line_len [ i ] < win.cols  )
            total_lines_before++;
        else
            total_lines_before += (int) ceil (( double) buf.line_len [ i ] / win.cols );
    }
    rows_in_current_textline = index / win.cols ;
    curs->r = total_lines_before + rows_in_current_textline ;
    curs->c = index - rows_in_current_textline * win.cols ;
}
/***********************************************************************/

void  advance_cursor ( Cursor *cursor , Window win , char ch)
{
    if  (  ch == '\n'  )  {
        cursor->r++;
        cursor->c = 0;
    }
    else  {
        cursor->c++;
        if  (  cursor->c == win.cols  )  {   /* wrap the line */
            cursor->c = 0;
            cursor->r++;
        }
    }
}
/***********************************************************************/

void  init_cursor ( Cursor *cursor )
{
    cursor->r           = 0;
    cursor->c           = 0;
}
/***********************************************************************/

void  show_cursor ( Buffer buf , Window win , Cursor cursor ,
                    int index_in_line , int line_number )
{
    char  curs_str [ 80 ];
```

```
        sprintf(curs_str, "Cursor: [%d,%d]   line index: %d   win topline: %d "
                          "buf #lines: %d",
                   cursor.r+1, cursor.c+1, line_number, win.line_at_top,
                   buf.num_lines   ) ;
    write(1,PARK, lPARK);
    write(1, CLEAR_LINE, lCLEAR_LINE);
    write(1, curs_str, strlen(curs_str));
    moveto(cursor.r, cursor.c);
}
/**************************************************************************/

void moveUp( Buffer *buf, Window *win, Cursor *curs   )
{
    /* if buf.cur_line == 0, we cannot go up further */
    if ( buf->cur_line > 0 ) {
        buf->cur_line--;

        if ( buf->index_in_cur_line >= buf->line_len[buf->cur_line] ) {
            buf->index_in_cur_line = buf->line_len[buf->cur_line]-1;
        }

        if ( buf->cur_line >= win->line_at_top ) {
            get_cursor_at(*buf, *win, buf->index_in_cur_line, buf->cur_line,
                               curs );
            }
        else {
            win->line_at_top = buf->cur_line;
            get_cursor_at(*buf, *win, buf->index_in_cur_line, buf->cur_line,
                               curs );
            scroll_buffer(*buf, *win );
        }
        moveto(curs->r, curs->c);
    }
}
/**************************************************************************/

void moveDown( Buffer *buf, Window *win, Cursor *curs )
{
    int lastline;

    if ( buf->cur_line < buf->num_lines-1 ) {
        buf->cur_line++;
        /* Check whether the cursor would be past the rightmost character
           of the now current line. If so, position it just past the rightmost
           character. */
        if ( buf->index_in_cur_line >= buf->line_len[buf->cur_line] ) {
            buf->index_in_cur_line = buf->line_len[buf->cur_line]-1;
        }

        get_lastline_in_win( *buf, *win,  &lastline );

        if ( buf->cur_line > lastline ) {
            /* Need to scroll */
            win->line_at_top += buf->cur_line - lastline;
```

```
                get_cursor_at(*buf, *win, buf->index_in_cur_line, buf->cur_line,
                                    curs );
                scroll_buffer(*buf, *win );
        }
        else
                get_cursor_at(*buf, *win, buf->index_in_cur_line,
                                buf->cur_line,   curs );

        moveto(curs->r, curs->c);
    }
}
/*************************************************************************/

void moveRight( Buffer *buf, Window *win, Cursor *curs )
{
    if ( ( buf->index_in_cur_line < buf->line_len[buf->cur_line] -1 ) ||
        ( ( buf->index_in_cur_line < buf->line_len[buf->cur_line] )
            && ( buf->cur_line == buf->num_lines-1 ) ) )

    {
        buf->index_in_cur_line++;
        if ( buf->index_in_cur_line % win->cols == 0 ) {
            curs->r++;
            curs->c = 0;
            if ( curs->r > win->rows-2 ) {
                win->line_at_top += curs->r - (win->rows-2);
                scroll_buffer(*buf, *win );
            }
            moveto(curs->r, curs->c);
        }
        else {
            curs->c++;
            write(1, RIGHT, lRIGHT);
        }
    }
}
/*************************************************************************/

void moveLeft ( Buffer *buf, Window *win, Cursor *curs )
{
    if ( buf->index_in_cur_line > 0 ) {
        if ( buf->index_in_cur_line % win->cols == 0 ) {
            curs->r--;
            curs->c = win->cols -1;
            moveto(curs->r, curs->c);
        }
        else {
            curs->c--;
            write(1, LEFT, lLEFT);
        }
        buf->index_in_cur_line--;
    }
}
```