

Advanced Programming in  
the UNIX Environment  
Third Edition

# UNIX 环境高级编程

第3版·英文版 | 下

[美] W. 理查德·史蒂文斯 (W. Richard Stevens) 著  
史蒂芬·A. 拉戈 (Stephen A. Rago)

Advanced Programming in  
the UNIX Environment  
Third Edition

# UNIX

## 环境高级编程

第3版·英文版 | 下

[美] W. 理查德·史蒂文斯 (W. Richard Stevens) 著  
史蒂芬·A. 拉戈 (Stephen A. Rago)



人民邮电出版社  
北京


20多年来，严谨的C程序员都是依靠一本书来深入了解驱动UNIX和Linux内核的编程接口的实用知识的，这本书就是W. Richard Stevens所著的《UNIX环境高级编程》。现在，Stevens的同事Stephen Rago彻底更新了这本经典著作。新的第3版支持当今领先的系统平台，反映了最新技术进展和最佳实践，并且符合最新的Single UNIX Specification第4版（SUSv4）。

Rago保留了使本书前版成为经典之作的精髓和方法。他在Stevens原著的基础上，从基础的文件、目录和进程讲起，并给诸如信号处理和终端I/O之类的先进技术保留较大的篇幅。他还深入讨论了线程和多线程编程、使用套接字接口驱动进程间通信（IPC）等方面的内容。

这一版涵盖了70多个新接口，包括POSIX异步I/O、旋转锁、屏障（barrier）和POSIX信号量。此外，这一版删除了许多过时的接口，保留了一些广泛使用的接口。书中几乎所有实例都已经在主流的4个平台上测试过，包括Solaris10、Mac OS X 10.6.8（Darwin 10.8.0）、FreeBSD 8.0和Ubuntu 12.04（基于Linux 3.2）。

与前两版一样，读者仍可以通过实例学习，这些实例包括了1万多行可下载的ISO C源代码，书中通过简明但完整的程序阐述了400多个系统调用和函数，清楚地说明它们的用法、参数和返回值。为了使读者能融会贯通，书中还提供了几个贯穿整章的案例，每个案例都根据现在的技术环境进行了全面更新。

本书帮助几代程序员写出了可靠、强大、高性能的代码。第3版根据当今主流系统进行更新，更具实用价值。



# 目录

<b>Chapter 14. Advanced I/O / 高级 I/O</b>	<b>481</b>
14.1 Introduction / 引言	481
14.2 Nonblocking I/O / 非阻塞 I/O	481
14.3 Record Locking / 记录锁	485
14.4 I/O Multiplexing / I/O 多路转接	500
14.4.1 select and pselect Functions / 函数 select 和 pselect	502
14.4.2 poll Function / 函数 poll	506
14.5 Asynchronous I/O / 异步 I/O	509
14.5.1 System V Asynchronous I/O / System V 异步 I/O	510
14.5.2 BSD Asynchronous I/O / BSD 异步 I/O	510
14.5.3 POSIX Asynchronous I/O / POSIX 异步 I/O	511
14.6 readv and writev Functions / 函数 readv 和 writev	521
14.7 readn and writen Functions / 函数 readn 和 writen	523
14.8 Memory-Mapped I/O / 存储映射 I/O	525
14.9 Summary / 小结	531
Exercises / 习题	532
<b>Chapter 15. Interprocess Communication / 进程间通信</b>	<b>533</b>
15.1 Introduction / 引言	533
15.2 Pipes / 管道	534

15.3	popen and pclose Functions / 函数 popen 和 pclose	541
15.4	Coprocesses / 协同进程	548
15.5	FIFOs	552
15.6	XSI IPC	556
15.6.1	Identifiers and Keys / 标识符和键	556
15.6.2	Permission Structure / 权限结构	558
15.6.3	Configuration Limits / 结构限制	559
15.6.4	Advantages and Disadvantages / 优点和缺点	559
15.7	Message Queues / 消息队列	561
15.8	Semaphores / 信号量	565
15.9	Shared Memory / 共享存储	571
15.10	POSIX Semaphores / POSIX 信号量	579
15.11	Client-Server Properties / 客户进程-服务器进程属性	585
15.12	Summary / 小结	587
	Exercises / 习题	587
<b>Chapter 16.</b>	<b>Network IPC: Sockets / 网络 IPC: 套接字</b>	<b>589</b>
16.1	Introduction / 引言	589
16.2	Socket Descriptors / 套接字描述符	590
16.3	Addressing / 寻址	593
16.3.1	Byte Ordering / 字节序	593
16.3.2	Address Formats / 地址格式	595
16.3.3	Address Lookup / 地址查询	597
16.3.4	Associating Addresses with Sockets / 将套接字与地址关联	604
16.4	Connection Establishment / 建立连接	605
16.5	Data Transfer / 数据传输	610
16.6	Socket Options / 套接字选项	623
16.7	Out-of-Band Data / 带外数据	626
16.8	Nonblocking and Asynchronous I/O / 非阻塞和异步 I/O	627
16.9	Summary / 小结	628
	Exercises / 习题	628
<b>Chapter 17.</b>	<b>Advanced IPC / 高级进程间通信</b>	<b>629</b>

17.1	Introduction / 引言	629
17.2	UNIX Domain Sockets / UNIX 域套接字	629
17.3	Unique Connections / 唯一连接	635
17.4	Passing File Descriptors / 传送文件描述符	642
17.5	An Open Server, Version 1 / 打开服务器进程第 1 版	653
17.6	An Open Server, Version 2 / 打开服务器进程第 2 版	659
17.7	Summary / 小结	669
	Exercises / 习题	670
<b>Chapter 18.</b>	<b>Terminal I/O / 终端 I/O</b>	<b>671</b>
18.1	Introduction / 引言	671
18.2	Overview / 概述	671
18.3	Special Input Characters / 特殊输入字符	678
18.4	Getting and Setting Terminal Attributes / 获得和设置终端属性	683
18.5	Terminal Option Flags / 终端选项标志	683
18.6	stty Command / stty 命令	691
18.7	Baud Rate Functions / 波特率函数	692
18.8	Line Control Functions / 行控制函数	693
18.9	Terminal Identification / 终端标识	694
18.10	Canonical Mode / 规范模式	700
18.11	Noncanonical Mode / 非规范模式	703
18.12	Terminal Window Size / 终端窗口大小	710
18.13	termcap, terminfo, and curses / termcap、terminfo 和 curses	712
18.14	Summary / 小结	713
	Exercises / 习题	713
<b>Chapter 19.</b>	<b>Pseudo Terminals / 伪终端</b>	<b>715</b>
19.1	Introduction / 引言	715
19.2	Overview / 概述	715
19.3	Opening Pseudo-Terminal Devices / 打开伪终端设备	722
19.4	pty_fork Function / 函数 pty_fork	726
19.5	pty Program / pty 程序	729
19.6	Using the pty Program / 使用 pty 程序	733
19.7	Advanced Features / 高级特性	740



19.8 Summary / 小结	741
Exercises / 习题	742
<b>Chapter 20. A Database Library / 数据库函数库</b>	<b>743</b>
20.1 Introduction / 引言	743
20.2 History / 历史	743
20.3 The Library / 函数库	744
20.4 Implementation Overview / 实现概述	746
20.5 Centralized or Decentralized? / 集中式还是非集中式?	750
20.6 Concurrency / 并发	752
20.7 Building the Library / 构造函数库	753
20.8 Source Code / 源代码	753
20.9 Performance / 性能	781
20.10 Summary / 小结	786
Exercises / 习题	787
<b>Chapter 21. Communicating with a Network Printer / 与网络打印机通信</b>	<b>789</b>
21.1 Introduction / 引言	789
21.2 The Internet Printing Protocol / 网络打印协议	789
21.3 The Hypertext Transfer Protocol / 超文本传输协议 HTTP	792
21.4 Printer Spooling / 打印假脱机技术	793
21.5 Source Code / 源代码	795
21.6 Summary / 小结	843
Exercises / 习题	843
<b>Appendix A. Function Prototypes / 函数原型</b>	<b>845</b>
<b>Appendix B. Miscellaneous Source Code / 其他源代码</b>	<b>895</b>
B.1 Our Header File / 本书使用的头文件	895
B.2 Standard Error Routines / 标准出错例程	898
<b>Appendix C. Solutions to Selected Exercises / 部分习题答案</b>	<b>905</b>
<b>Bibliography / 参考书目</b>	<b>947</b>



# 14

## Advanced I/O

### 14.1 Introduction

This chapter covers numerous topics and functions that we lump under the term *advanced I/O*: nonblocking I/O, record locking, I/O multiplexing (the `select` and `poll` functions), asynchronous I/O, the `readv` and `writv` functions, and memory-mapped I/O (`mmap`). We need to cover these topics before describing interprocess communication in Chapter 15, Chapter 17, and many of the examples in later chapters.

### 14.2 Nonblocking I/O

In Section 10.5, we said that system calls are divided into two categories: the “slow” ones and all the others. The slow system calls are those that can block forever. They include

- Reads that can block the caller forever if data isn’t present with certain file types (pipes, terminal devices, and network devices)
- Writes that can block the caller forever if the data can’t be accepted immediately by these same file types (e.g., no room in the pipe, network flow control)
- Opens that block until some condition occurs on certain file types (such as an open of a terminal device that waits until an attached modem answers the phone, or an open of a FIFO for writing only, when no other process has the FIFO open for reading)
- Reads and writes of files that have mandatory record locking enabled

- Certain `ioctl` operations
- Some of the interprocess communication functions (Chapter 15)

We also said that system calls related to disk I/O are not considered slow, even though the read or write of a disk file can block the caller temporarily.

Nonblocking I/O lets us issue an I/O operation, such as an `open`, `read`, or `write`, and not have it block forever. If the operation cannot be completed, the call returns immediately with an error noting that the operation would have blocked.

There are two ways to specify nonblocking I/O for a given descriptor.

1. If we call `open` to get the descriptor, we can specify the `O_NONBLOCK` flag (Section 3.3).
2. For a descriptor that is already open, we call `fcntl` to turn on the `O_NONBLOCK` file status flag (Section 3.14). Figure 3.12 shows a function that we can call to turn on any of the file status flags for a descriptor.

Earlier versions of System V used the flag `O_NDELAY` to specify nonblocking mode. These versions of System V returned a value of 0 from the `read` function if there wasn't any data to be read. Since this use of a return value of 0 overlapped with the normal UNIX System convention of 0 meaning the end of file, POSIX.1 chose to provide a nonblocking flag with a different name and different semantics. Indeed, with these older versions of System V, when we get a return of 0 from `read`, we don't know whether the call would have blocked or whether the end of file was encountered. We'll see that POSIX.1 requires that `read` return -1 with `errno` set to `EAGAIN` if there is no data to read from a nonblocking descriptor. Some platforms derived from System V support both the older `O_NDELAY` and the POSIX.1 `O_NONBLOCK`, but in this text we'll use only the POSIX.1 feature. The older `O_NDELAY` is intended for backward compatibility and should not be used in new applications.

4.3BSD provided the `FNDELAY` flag for `fcntl`, and its semantics were slightly different. Instead of affecting only the file status flags for the descriptor, the flags for either the terminal device or the socket were also changed to be nonblocking, thereby affecting all users of the terminal or socket, not just the users sharing the same file table entry (4.3BSD nonblocking I/O worked only on terminals and sockets). Also, 4.3BSD returned `EWOULDBLOCK` if an operation on a nonblocking descriptor could not complete without blocking. Today, BSD-based systems provide the POSIX.1 `O_NONBLOCK` flag and define `EWOULDBLOCK` to be the same as `EAGAIN`. These systems provide nonblocking semantics consistent with other POSIX-compatible systems: changes in file status flags affect all users of the same file table entry, but are independent of accesses to the same device through other file table entries. (Refer to Figures 3.7 and 3.9.)

## Example

Let's look at an example of nonblocking I/O. The program in Figure 14.1 reads up to 500,000 bytes from the standard input and attempts to write it to the standard output. The standard output is first set to be nonblocking. The output is in a loop, with the results of each `write` being printed on the standard error. The function `clr_fl` is similar to the function `set_fl` that we showed in Figure 3.12. This new function simply clears one or more of the flag bits.

---

```

#include "apue.h"
#include <errno.h>
#include <fcntl.h>

char    buf[500000];

int
main(void)
{
    int    ntowrite, nwrite;
    char    *ptr;

    ntowrite = read(STDIN_FILENO, buf, sizeof(buf));
    fprintf(stderr, "read %d bytes\n", ntowrite);

    set_fl(STDOUT_FILENO, O_NONBLOCK); /* set nonblocking */

    ptr = buf;
    while (ntowrite > 0) {
        errno = 0;
        nwrite = write(STDOUT_FILENO, ptr, ntowrite);
        fprintf(stderr, "nwrite = %d, errno = %d\n", nwrite, errno);

        if (nwrite > 0) {
            ptr += nwrite;
            ntowrite -= nwrite;
        }
    }

    clr_fl(STDOUT_FILENO, O_NONBLOCK); /* clear nonblocking */

    exit(0);
}

```

---

Figure 14.1 Large nonblocking write

If the standard output is a regular file, we expect the write to be executed once:

<code>\$ ls -l /etc/services</code>	<i>print file size</i>
<code>-rw-r--r-- 1 root 677959 Jun 23 2009 /etc/services</code>	
<code>\$ ./a.out &lt; /etc/services &gt; temp.file</code>	<i>try a regular file first</i>
<code>read 500000 bytes</code>	
<code>nwrite = 500000, errno = 0</code>	<i>a single write</i>
<code>\$ ls -l temp.file</code>	<i>verify size of output file</i>
<code>-rw-rw-r-- 1 sar 500000 Apr 1 13:03 temp.file</code>	

But if the standard output is a terminal, we expect the write to return a partial count sometimes and an error at other times. This is what we see:

```

$ ./a.out < /etc/services 2>stderr.out
$ cat stderr.out
read 500000 bytes
nwrite = 999, errno = 0
nwrite = -1, errno = 35
nwrite = -1, errno = 35
nwrite = -1, errno = 35
nwrite = -1, errno = 35
nwrite = 1001, errno = 0
nwrite = -1, errno = 35
nwrite = 1002, errno = 0
nwrite = 1004, errno = 0
nwrite = 1003, errno = 0
nwrite = 1003, errno = 0
nwrite = 1005, errno = 0
nwrite = -1, errno = 35
:
nwrite = 1006, errno = 0
nwrite = 1004, errno = 0
nwrite = 1005, errno = 0
nwrite = 1006, errno = 0
nwrite = -1, errno = 35
:
nwrite = 1006, errno = 0
nwrite = 1005, errno = 0
nwrite = 1005, errno = 0
nwrite = -1, errno = 35
:
nwrite = 347, errno = 0

```

*output to terminal*  
*lots of output to terminal ...*

*61 of these errors*

*108 of these errors*

*681 of these errors*

*and so on ...*

On this system, the `errno` of 35 is `EAGAIN`. The amount of data accepted by the terminal driver varies from system to system. The results will also vary depending on how you are logged in to the system: on the system console, on a hard-wired terminal, on a network connection using a pseudo terminal. If you are running a windowing system on your terminal, you are also going through a pseudo terminal device. □

In this example, the program issues more than 9,000 `write` calls, even though only 500 are needed to output the data. The rest just return an error. This type of loop, called *polling*, is a waste of CPU time on a multiuser system. In Section 14.4, we'll see that I/O multiplexing with a nonblocking descriptor is a more efficient way to do this.

Sometimes, we can avoid using nonblocking I/O by designing our applications to use multiple threads (see Chapter 11). We can allow individual threads to block in I/O calls if we can continue to make progress in other threads. This can sometimes simplify the design, as we shall see in Chapter 21; at other times, however, the overhead of synchronization can add more complexity than is saved from using threads.

## 14.3 Record Locking

What happens when two people edit the same file at the same time? In most UNIX systems, the final state of the file corresponds to the last process that wrote the file. In some applications, however, such as a database system, a process needs to be certain that it alone is writing to a file. To provide this capability for processes that need it, commercial UNIX systems provide record locking. (In Chapter 20, we develop a database library that uses record locking.)

*Record locking* is the term normally used to describe the ability of a process to prevent other processes from modifying a region of a file while the first process is reading or modifying that portion of the file. Under the UNIX System, “record” is a misnomer; the UNIX kernel does not have a notion of records in a file. A better term is *byte-range locking*, given that it is a range of a file (possibly the entire file) that is locked.

### History

One of the criticisms of early UNIX systems was that they couldn’t be used to run database systems, because they did not support locking portions of files. As UNIX systems found their way into business computing environments, various groups added support for record locking (differently, of course).

Early Berkeley releases supported only the `flock` function. This function locks only entire files, not regions of a file.

Record locking was added to System V Release 3 through the `fcntl` function. The `lockf` function was built on top of this, providing a simplified interface. These functions allowed callers to lock arbitrary byte ranges in a file, ranging from the entire file down to a single byte within the file.

POSIX.1 chose to standardize on the `fcntl` approach. Figure 14.2 shows the forms of record locking provided by various systems. Note that the Single UNIX Specification includes `lockf` in the XSI option.

System	Advisory	Mandatory	<code>fcntl</code>	<code>lockf</code>	<code>flock</code>
SUS	•		•	XSI	
FreeBSD 8.0	•		•	•	•
Linux 3.2.0	•	•	•	•	•
Mac OS X 10.6.8	•		•	•	•
Solaris 10	•	•	•	•	•

**Figure 14.2** Forms of record locking supported by various UNIX systems

We describe the difference between advisory locking and mandatory locking later in this section. In this text, we describe only the POSIX.1 `fcntl` locking.

Record locking was originally added to Version 7 in 1980 by John Bass. The system call entry into the kernel was a function named `locking`. This function provided mandatory record locking and propagated through many versions of System III. Xenix systems picked up this function, and some Intel-based System V derivatives, such as OpenServer 5, continued to support it in a Xenix-compatibility library.

## fcntl Record Locking

Let's repeat the prototype for the `fcntl` function from Section 3.14.

```
#include <fcntl.h>

int fcntl(int fd, int cmd, ... /* struct flock *flockptr */ );

Returns: depends on cmd if OK (see following), -1 on error
```

For record locking, *cmd* is `F_GETLK`, `F_SETLK`, or `F_SETLKW`. The third argument (which we'll call *flockptr*) is a pointer to an `flock` structure.

```
struct flock {
    short l_type; /* F_RDLCK, F_WRLCK, or F_UNLCK */
    short l_whence; /* SEEK_SET, SEEK_CUR, or SEEK_END */
    off_t l_start; /* offset in bytes, relative to l_whence */
    off_t l_len; /* length, in bytes; 0 means lock to EOF */
    pid_t l_pid; /* returned with F_GETLK */
};
```

This structure describes

- The type of lock desired: `F_RDLCK` (a shared read lock), `F_WRLCK` (an exclusive write lock), or `F_UNLCK` (unlocking a region)
- The starting byte offset of the region being locked or unlocked (`l_start` and `l_whence`)
- The size of the region in bytes (`l_len`)
- The ID (`l_pid`) of the process holding the lock that can block the current process (returned by `F_GETLK` only)

Numerous rules apply to the specification of the region to be locked or unlocked.

- The two elements that specify the starting offset of the region are similar to the last two arguments of the `lseek` function (Section 3.6). Indeed, the `l_whence` member is specified as `SEEK_SET`, `SEEK_CUR`, or `SEEK_END`.
- Locks can start and extend beyond the current end of file, but cannot start or extend before the beginning of the file.
- If `l_len` is 0, it means that the lock extends to the largest possible offset of the file. This allows us to lock a region starting anywhere in the file, up through and including any data that is appended to the file. (We don't have to try to guess how many bytes might be appended to the file.)
- To lock the entire file, we set `l_start` and `l_whence` to point to the beginning of the file and specify a length (`l_len`) of 0. (There are several ways to specify the beginning of the file, but most applications specify `l_start` as 0 and `l_whence` as `SEEK_SET`.)



We previously mentioned two types of locks: a shared read lock (`l_type` of `F_RDLCK`) and an exclusive write lock (`F_WRLCK`). The basic rule is that any number of processes can have a shared read lock on a given byte, but only one process can have an exclusive write lock on a given byte. Furthermore, if there are one or more read locks on a byte, there can't be any write locks on that byte; if there is an exclusive write lock on a byte, there can't be any read locks on that byte. We show this compatibility rule in Figure 14.3.

		Request for	
		read lock	write lock
Region currently has	no locks	OK	OK
	one or more read locks	OK	denied
	one write lock	denied	denied

**Figure 14.3** Compatibility between different lock types

The compatibility rule applies to lock requests made from different processes, not to multiple lock requests made by a single process. If a process has an existing lock on a range of a file, a subsequent attempt to place a lock on the same range by the same process will replace the existing lock with the new one. Thus, if a process has a write lock on bytes 16–32 of a file and then tries to place a read lock on bytes 16–32, the request will succeed, and the write lock will be replaced by a read lock.

To obtain a read lock, the descriptor must be open for reading; to obtain a write lock, the descriptor must be open for writing.

We can now describe the three commands for the `fcntl` function.

**F\_GETLK** Determine whether the lock described by *flockptr* is blocked by some other lock. If a lock exists that would prevent ours from being created, the information on that existing lock overwrites the information pointed to by *flockptr*. If no lock exists that would prevent ours from being created, the structure pointed to by *flockptr* is left unchanged except for the `l_type` member, which is set to `F_UNLCK`.

**F\_SETLK** Set the lock described by *flockptr*. If we are trying to obtain a read lock (`l_type` of `F_RDLCK`) or a write lock (`l_type` of `F_WRLCK`) and the compatibility rule prevents the system from giving us the lock (Figure 14.3), `fcntl` returns immediately with `errno` set to either `EACCES` or `EAGAIN`.

Although POSIX allows an implementation to return either error code, all four implementations described in this text return `EAGAIN` if the locking request cannot be satisfied.

This command is also used to clear the lock described by *flockptr* (`l_type` of `F_UNLCK`).




**F\_SETLKW** This command is a blocking version of **F\_SETLK**. (The **W** in the command name means *wait*.) If the requested read lock or write lock cannot be granted because another process currently has some part of the requested region locked, the calling process is put to sleep. The process wakes up either when the lock becomes available or when interrupted by a signal.

Be aware that testing for a lock with **F\_GETLK** and then trying to obtain that lock with **F\_SETLK** or **F\_SETLKW** is not an atomic operation. We have no guarantee that, between the two `fcntl` calls, some other process won't come in and obtain the same lock. If we don't want to block while waiting for a lock to become available to us, we must handle the possible error returns from **F\_SETLK**.

Note that POSIX.1 doesn't specify what happens when one process read locks a range of a file, a second process blocks while trying to get a write lock on the same range, and a third process then attempts to get another read lock on the range. If the third process is allowed to place a read lock on the range just because the range is already read locked, then the implementation might starve processes with pending write locks. Thus, as additional requests to read lock the same range arrive, the time that the process with the pending write-lock request has to wait is extended. If the read-lock requests arrive quickly enough without a lull in the arrival rate, then the writer could wait for a long time.

When setting or releasing a lock on a file, the system combines or splits adjacent areas as required. For example, if we lock bytes 100 through 199 and then unlock byte 150, the kernel still maintains the locks on bytes 100 through 149 and bytes 151 through 199. Figure 14.4 illustrates the byte-range locks in this situation.



**Figure 14.4** File byte-range lock diagram

If we were to lock byte 150, the system would coalesce the adjacent locked regions into a single region from byte 100 through 199. The resulting picture would be the first diagram in Figure 14.4, the same as when we started.

## Example — Requesting and Releasing a Lock

To save ourselves from having to allocate an `flock` structure and fill in all the elements each time, the function `lock_reg` in Figure 14.5 handles all these details.

---

```
#include "apue.h"
#include <fcntl.h>

int
lock_reg(int fd, int cmd, int type, off_t offset, int whence, off_t len)
{
    struct flock    lock;

    lock.l_type = type;      /* F_RDLCK, F_WRLCK, F_UNLCK */
    lock.l_start = offset;   /* byte offset, relative to l_whence */
    lock.l_whence = whence; /* SEEK_SET, SEEK_CUR, SEEK_END */
    lock.l_len = len;       /* #bytes (0 means to EOF) */

    return(fcntl(fd, cmd, &lock));
}
```

---

Figure 14.5 Function to lock or unlock a region of a file

Since most locking calls are to lock or unlock a region (the command `F_GETLK` is rarely used), we normally use one of the following five macros, which are defined in `apue.h` (Appendix B).

```
#define read_lock(fd, offset, whence, len) \
    lock_reg((fd), F_SETLK, F_RDLCK, (offset), (whence), (len))
#define readw_lock(fd, offset, whence, len) \
    lock_reg((fd), F_SETLKW, F_RDLCK, (offset), (whence), (len))
#define write_lock(fd, offset, whence, len) \
    lock_reg((fd), F_SETLK, F_WRLCK, (offset), (whence), (len))
#define writew_lock(fd, offset, whence, len) \
    lock_reg((fd), F_SETLKW, F_WRLCK, (offset), (whence), (len))
#define un_lock(fd, offset, whence, len) \
    lock_reg((fd), F_SETLK, F_UNLCK, (offset), (whence), (len))
```

We have purposely defined the first three arguments to these macros in the same order as the `lseek` function. □

## Example — Testing for a Lock

Figure 14.6 defines the function `lock_test` that we'll use to test for a lock.

---

```
#include "apue.h"
#include <fcntl.h>

pid_t
lock_test(int fd, int type, off_t offset, int whence, off_t len)
{
    struct flock    lock;
    int              rc;
```

---