

Outline

23	Alternative I/O Models	23-1
23.1	Overview	23-3
23.2	Nonblocking I/O	23-5
23.3	Signal-driven I/O	23-11
23.4	I/O multiplexing: <code>poll()</code>	23-14
23.5	Problems with <code>poll()</code> and <code>select()</code>	23-31
23.6	The <code>epoll</code> API	23-34
23.7	<code>epoll</code> events	23-45
23.8	<code>epoll</code> : edge-triggered notification	23-59
23.9	<code>epoll</code> : API quirks	23-70
23.10	Event-loop programming	23-76

Overview

- Like `select()` and `poll()`, `epoll` can monitor multiple FDs
- `epoll` returns readiness information in similar manner to `poll()`
- Two main **advantages**:
 - `epoll` provides **much better performance** when monitoring large numbers of FDs (see TLPI §63.4.5)
 - `epoll` provides two **notification modes**: **level-triggered** and **edge-triggered**
 - Default is level-triggered notification
 - `select()` and `poll()` provide only level-triggered notification
 - (Signal-driven I/O provides only edge-triggered notification)
- Linux-specific, since kernel 2.6.0

[TLPI §63.4]


Central data structure of *epoll* API is an *epoll* instance

- **Persistent** data structure **maintained in kernel space**
 - Referred to in user space via file descriptor
- Can (abstractly) be considered as container for two lists:
 - **Interest list**: list of FDs to be monitored
 - **Ready list**: list of FDs that are ready for I/O
 - Ready list is (dynamic) subset of interest list

The key *epoll* APIs are:

- *epoll_create()*: create a new *epoll* instance and return FD referring to instance
 - FD is used in the calls below
- *epoll_ctl()*: modify interest list of *epoll* instance
 - Add FDs to/remove FDs from interest list
 - Modify events mask for FDs currently in interest list
- *epoll_wait()*: return items from ready list of *epoll* instance

epoll kernel data structures and APIs



Creating an *epoll* instance: *epoll_create()*

```
#include <sys/epoll.h>
int epoll_create(int size);
```

- Creates an *epoll* instance
- *size*:
 - Since Linux 2.6.8: serves no purpose, but must be > 0
 - Before Linux 2.6.8: an *estimate* of number of FDs to be monitored via this *epoll* instance
- Returns file descriptor on success, or -1 on error
 - When FD is no longer required, it should be closed via *close()*
- Since Linux 2.6.27, *epoll_create1()* provides improved API
 - See the man page

Modifying the *epoll* interest list: *epoll_ctl()*

```
#include <sys/epoll.h>
int epoll_ctl(int epfd, int op, int fd,
              struct epoll_event *ev);
```

- Modifies the interest list associated with *epoll* FD, *epfd*
- *fd*: identifies which FD in interest list is to have its settings modified
 - E.g., FD for pipe, FIFO, terminal, socket, POSIX MQ, or even another *epoll* FD
 - (Can't be FD for a regular file or directory)
- *op*: operation to perform on interest list
- *ev*: (Later)

[TLPI §63.4.2]

epoll_ctl() *op* argument

The *epoll_ctl()* *op* argument is one of:

- **EPOLL_CTL_ADD**: add *fd* to interest list of *epfd*
 - *ev* specifies events to be monitored for *fd*
 - If *fd* is already in interest list ⇒ **EEXIST**
- **EPOLL_CTL_MOD**: modify settings of *fd* in interest list of *epfd*
 - *ev* specifies new settings to be associated with *fd*
 - If *fd* is not in interest list ⇒ **ENOENT**
- **EPOLL_CTL_DEL**: remove *fd* from interest list of *epfd*
 - Also removes corresponding entry in ready list, if present
 - *ev* is ignored
 - If *fd* is not in interest list ⇒ **ENOENT**
 - **Closing an FD automatically removes it from all *epoll* interest lists**
 - ⚠ But see later! Manual deletion is sometimes required

The `epoll_event` structure

`epoll_ctl()` *ev* argument is pointer to an `epoll_event` structure:

```
struct epoll_event {
    uint32_t      events; /* epoll events (bit mask) */
    epoll_data_t data;    /* User data */
};

typedef union epoll_data {
    void      *ptr; /* Pointer to user-defined data */
    int       fd;   /* File descriptor */
    uint32_t  u32;  /* 32-bit integer */
    uint64_t  u64;  /* 64-bit integer */
} epoll_data_t;
```

- ***ev.events***: bit mask of events to monitor for *fd*
 - (Similar to *events* mask given to `poll()`)
- ***data***: info to be passed back to caller of `epoll_wait()` when *fd* later becomes ready
 - **Union field**: value is specified in *one* of the members

Example: using `epoll_create()` and `epoll_ctl()`

```
int epfd;
struct epoll_event ev;

epfd = epoll_create(5);

ev.data.fd = fd;
ev.events = EPOLLIN; /* Monitor for input available */
epoll_ctl(epfd, EPOLL_CTL_ADD, fd, &ev);
```

Outline

23	Alternative I/O Models	23-1
23.1	Overview	23-3
23.2	Nonblocking I/O	23-5
23.3	Signal-driven I/O	23-11
23.4	I/O multiplexing: poll()	23-14
23.5	Problems with poll() and select()	23-31
23.6	The epoll API	23-34
23.7	epoll events	23-45
23.8	epoll: edge-triggered notification	23-59
23.9	epoll: API quirks	23-70
23.10	Event-loop programming	23-76

Waiting for events: *epoll_wait()*

```
#include <sys/epoll.h>
int epoll_wait(int epfd, struct epoll_event *evlist,
               int maxevents, int timeout);
```

- Returns info about ready FDs in interest list of *epoll* instance of *epfd*
- Blocks until at least one FD is ready
- Info about ready FDs is returned in array *evlist*
 - I.e., can get information about multiple ready FDs with one *epoll_wait()* call
 - (Caller allocates the *evlist* array)
- *maxevents*: size of the *evlist* array

[TLPI §63.4.3]

Waiting for events: *epoll_wait()*

```
#include <sys/epoll.h>
int epoll_wait(int epfd, struct epoll_event *evlist,
               int maxevents, int timeout);
```

- *timeout* specifies a timeout for call:
 - -1: block until an FD in interest list becomes ready
 - 0: perform a nonblocking “poll” to see if any FDs in interest list are ready
 - > 0: block for up to *timeout* milliseconds or until an FD in interest list becomes ready
- Return value:
 - > 0: number of items placed in *evlist*
 - 0: no FDs became ready within interval specified by *timeout*
 - -1: an error occurred

Waiting for events: *epoll_wait()*

```
#include <sys/epoll.h>
int epoll_wait(int epfd, struct epoll_event *evlist,
               int maxevents, int timeout);
```

- Info about **multiple** FDs can be returned in the array *evlist*
- Each element of *evlist* returns info about one file descriptor:
 - *events* is a bit mask of events that have occurred for FD
 - *data* is *ev.data* value *currently* associated with FD in the interest list
- **NB:** the FD itself is **not** returned!
 - Instead, we put FD into *ev.data.fd* when calling *epoll_ctl()*, so that it is returned via *epoll_wait()*
 - (Or, put FD into a structure pointed to by *ev.data.ptr*)

Waiting for events: *epoll_wait()*

```
#include <sys/epoll.h>
int epoll_wait(int epfd, struct epoll_event *evlist,
               int maxevents, int timeout);
```

- 👍 If $> \text{maxevents}$ FDs are ready, successive *epoll_wait()* calls round-robin through FDs
 - Helps prevent file descriptor starvation
- 👍 In multithreaded programs:
 - One thread can modify interest list (*epoll_ctl()*) while another thread is blocked in *epoll_wait()*
 - *epoll_wait()* call will return if a newly added FD becomes ready

epoll events

Following table shows:

- Bits given in *ev.events* to *epoll_ctl()*
- Bits returned in *evlist[i].events* by *epoll_wait()*

Bit	<i>epoll_ctl()</i> ?	<i>epoll_wait()</i> ?	Description
EPOLLIN	•	•	Normal-priority data can be read
EPOLLPRI	•	•	High-priority data can be read
EPOLLRDHUP	•	•	Shutdown on peer socket
EPOLLOUT	•	•	Data can be written
EPOLLONESHOT	•		Disable monitoring after event notification
EPOLLET	•		Employ edge-triggered notification
EPOLLERR		•	An error has occurred
EPOLLHUP		•	A hangup occurred

- Other than **EPOLLOUT** and **EPOLLET**, bits have same meaning as similarly named *poll()* bit flags

Example: altio/epoll_input.c

```
./epoll_input file...
```

- Monitors one or more files using *epoll* API to see if input is possible
- Suitable files to give as arguments are:
 - FIFOs
 - Terminal device names
 - (May need to run *sleep* command in FG on the other terminal, to prevent shell stealing input)
 - Standard input
 - `/dev/stdin`

Example: altio/epoll_input.c (1)

```
#define MAX_BUF      1000      /* Max. bytes for read() */
#define MAX_EVENTS    5
    /* Max. number of events to be returned from
       a single epoll_wait() call */

int epfd, ready, fd, s, j, numOpenFds;
struct epoll_event ev;
struct epoll_event evlist[MAX_EVENTS];
char buf[MAX_BUF];

epfd = epoll_create(argc - 1);
```

- Declarations for various variables
- Create an *epoll* instance, obtaining *epoll* FD

Example: altio/epoll_input.c (2)

```
for (j = 1; j < argc; j++) {
    fd = open(argv[j], O_RDONLY);
    printf("Opened \"%s\" on fd %d\n", argv[j], fd);

    ev.events = EPOLLIN;
    ev.data.fd = fd;
    epoll_ctl(epfd, EPOLL_CTL_ADD, fd, &ev);
}

numOpenFds = argc - 1;
```

- Open each of the files named on command line
- Each file is monitored for input (**EPOLLIN**)
- *fd* placed in *ev.data*, so it is returned by *epoll_wait()*
- Add the FD to *epoll* interest list (*epoll_ctl()*)
- Track the number of open FDs

Example: altio/epoll_input.c (3)

```
while (numOpenFds > 0) {
    printf("About to epoll_wait()\n");
    ready = epoll_wait(epfd, evlist, MAX_EVENTS, -1);
    if (ready == -1) {
        if (errno == EINTR)
            continue;          /* Restart if interrupted
                                by signal */
        else
            errExit("epoll_wait");
    }
    printf("Ready: %d\n", ready);
}
```

- Loop, fetching *epoll* events and analyzing results
- Loop terminates when all FDs has been closed
- *epoll_wait()* call places up to **MAX_EVENTS** events in *evlist*
 - *timeout == -1* ⇒ infinite timeout
- Return value of *epoll_wait()* is number of ready FDs

Example: altio/epoll_input.c (4)

```
for (j = 0; j < ready; j++) {
    printf("  fd=%d; events: %s%s%s\n", evlist[j].data.fd,
        (evlist[j].events & EPOLLIN) ? "EPOLLIN " : "",
        (evlist[j].events & EPOLLHUP) ? "EPOLLHUP " : "",
        (evlist[j].events & EPOLLERR) ? "EPOLLERR " : "");
    if (evlist[j].events & EPOLLIN) {
        s = read(evlist[j].data.fd, buf, MAX_BUF);
        printf("    read %d bytes: %.*s\n", s, s, buf);
    } else if (evlist[j].events & (EPOLLHUP | EPOLLERR)) {
        printf("    closing fd %d\n", evlist[j].data.fd);
        close(evlist[j].data.fd);
        numOpenFds--;
    }
}
```

- Scan up to *ready* items in *evlist*
- Display *events* bits
- If *EPOLLIN* event occurred, read some input and display it on *stdout*
 - *%.s* ⇒ print string with field width taken from argument list (*s*)
- Otherwise, if error or hangup, close FD and decrements FD count
- Code correctly handles case where both *EPOLLIN* and *EPOLLHUP* are set in *evlist[j].events*

Exercises

- 1 Write a client ([*template: altio/ex.is_chat_cl.c*]) that communicates with the TCP chat server program, *is_chat_sv.c*. The program should be run with the following command line:

```
./is_chat_cl <host> <port> [<nickname>]
```

The program should create a connection to the server, and then use the *epoll* API to monitor both the terminal and the TCP socket for input. All input that becomes available on the socket should be written to the terminal and vice versa.

- Each time the program sends input from the terminal to the socket, that input should be prepended by the nickname supplied on the command line. If no nickname is supplied, then use the string returned by *getlogin(3)*. (*snprintf(3)* provides an easy way to concatenate the strings.)
- The program should terminate if it detects end-of-file or an error condition on either file descriptor.
- Calling *epoll_wait()* with *maxevents==1* will simplify the code!
- Bonus points if you find a way to crash the server (reproducibly)!

Exercises

- ② Write the chat server ([template: altio/ex.is_chat_sv.c]).

Note the following points:

- The program should take one command-line argument: the port number to which it should bind its listening socket.
- The program should accept and handle multiple simultaneous client connections. Input read from any client should be broadcast to all other clients.
- Use the *epoll* API to manage the file descriptors.
- You should use nonblocking file descriptors to ensure that the server never blocks when accepting connections or when reading or writing to clients.
- When the server detects end-of file or an error (other than *EAGAIN*) while reading or writing on a client connection, it should close that connection. (Remember that closing a file descriptor automatically removes it from any *epoll* interest lists of which it is a member.)

Exercises

- ③ Write a program ([template: altio/ex.epoll_pipes.c]) which performs the same task as the *altio/poll_pipes.c* program, but uses the *epoll* API instead of *poll()*.

Hints:

- After writing to the pipes, you will need to call *epoll_wait()* in a loop. The loop should be terminated when *epoll_wait()* indicates that there are no more ready file descriptors.
- After each call to *epoll_wait()*, you should display each ready pipe read file descriptor and then drain all input from that file descriptor so that it does not indicate as ready in future calls to *epoll_wait()*.
- In order to drain a pipe without blocking, you will need to make the file descriptor for the read end of the pipe nonblocking.

23	Alternative I/O Models	23-1
23.1	Overview	23-3
23.2	Nonblocking I/O	23-5
23.3	Signal-driven I/O	23-11
23.4	I/O multiplexing: <code>poll()</code>	23-14
23.5	Problems with <code>poll()</code> and <code>select()</code>	23-31
23.6	The <code>epoll</code> API	23-34
23.7	<code>epoll</code> events	23-45
23.8	<code>epoll</code> : edge-triggered notification	23-59
23.9	<code>epoll</code> : API quirks	23-70
23.10	Event-loop programming	23-76

Edge-triggered notification

- By default, `epoll` provides **level-triggered** (LT) notification
 - Tells us whether an **I/O operation can be performed on FD without blocking**
 - Like `poll()` and `select()`
- `EPOLLET` provides **edge-triggered** (ET) notification
 - Has I/O **activity occurred since `epoll_wait()` last notified about this FD?**
 - Or, if no `epoll_wait()` since FD was added/modified by `epoll_ctl()`, then: is FD ready?
- Example:

```
struct epoll_event ev;
ev.data.fd = fd
ev.events = EPOLLIN | EPOLLET;
epoll_ctl(epfd, EPOLL_CTL_ADD, fd, &ev);
```

Edge-triggered notification

- Illustration of difference between LT and ET notification:
 - 1 Monitoring a socket for input possible (**EPOLLIN**)
 - 2 Input arrives on socket
 - 3 We call `epoll_wait()`, which informs us that FD is ready
 - We *perhaps* consume some (**but not all**) available input
 - **No further input arrives on socket**
 - 4 We call `epoll_wait()` again
- LT notification: second `epoll_wait()` would (again) report FD as ready
 - Because outstanding data is still available for reading
- ET notification: second `epoll_wait()` would **not** report FD as ready
 - Because no I/O activity occurred since previous `epoll_wait()`

Uses for edge-triggered notification

- Can be more efficient: application is not repeatedly reminded that FD is ready
- Example: application that (periodically) generates data to be written to a socket
 - Application does not always have data to write
 - Application monitors socket for writability (**EPOLLOUT**)
 - Application is also monitoring other FDs for I/O possible
 - At some point, socket is full (output not possible)
 - Peer drains some data, socket becomes writable
 - LT notification: every `epoll_wait()` would (immediately) wake and say FD is writable
 - ET notification: only first `epoll_wait()` would say FD is writable
 - Application could cache that info for later action (e.g., when data is generated)

Edge-triggered notification provides an optimization

- Scenario: multiple threads/processes are `epoll_wait()`-ing on same `epoll` FD
 - E.g., `epoll` FD is monitoring listening socket
 - LT notification: **all** waiters are woken up when connection request arrives
 - ET notification: only **one** waiter is woken up
 - Avoids thundering herd problem
 - Code examples: `altio/multithread_epoll_wait.c`, `altio/epoll_flags_fork.c`
 - The `EPOLLEXCLUSIVE` flag provides a similar behavior in some scenarios when using level-triggered notification
 - Since Linux 4.5
 - See `epoll_ctl(2)` and `altio/epoll_flags_fork.c`

Edge-triggered notification and EPOLLONESHOT

- Scenario: monitoring socket for input available with `EPOLLET`
 - Assumption: we want to know when input is available, but don't want to read it **yet**
 - (So, we use `EPOLLET` to avoid repeated notifications)
- New input keeps appearing \Rightarrow ET notification keeps notifying
 - Really, we needed only **first** notification
- Solution: `EPOLLONESHOT`

One-shot monitoring: EPOLLONESHOT

- Specifying `EPOLLONESHOT` in *events* causes FD to be reported just once by `epoll_wait()`
- FD is then marked inactive in interest list
- FD remains in interest list, and can be reactivated using `epoll_ctl(EPOLL_CTL_MOD)`
 - Continuing previous example: reenable notification after draining input from socket

[TLPI §63.4.3]

Using edge-triggered notification

- Normally **employed with nonblocking I/O**
 - Can't monitor "I/O level", so must do nonblocking I/O calls until no more I/O is possible
 - Otherwise: risk blocking when doing I/O
- **Beware of FD starvation**
 - Scenarios where responding to a busy FD leaves other ready FDs starved of attention
 - (Starvation scenarios can also occur with level-triggered notification)
 - See TLPI §63.4.6

Exercises

The `altio/i_epoll.c` program can be used to perform epoll monitoring and file I/O operations on the objects named in its command-line arguments. The program is interactive, and supports the following commands:

```
p [<timeout>]
    Do epoll_wait() with millisecond timeout (default: 0)
e <fd> [<flags>]
    Modify epoll settings of <fd>; <flags> can include:
    'r' - EPOLLIN
    'w' - EPOLLOUT
    'e' - EPOLLET
    'o' - EPOLLONESHOT
    If no flags are given, disable <fd> in the interest list
r <fd> <size>
    Blocking read of <size> bytes from <fd>
R <fd> <size>
    Nonblocking read of <size> bytes from <fd>
w <fd> <size> [<char>]
    Blocking write of <size> bytes to <fd>; <char> is character
    to write (default: 'x')
W <fd> <size> [<char>]
    Nonblocking write of <size> bytes to <fd>
```

Each command-line argument has the form `<path>[:<flags>]` (to open a file) or `s%<host>%<port>[:<flags>]` (to connect a socket to a specified host/port). `<flags>` is as described above, and defaults to "r". (If testing with sockets, you will find the command `ncat -l <port>` useful, in order to create a server that you can connect to.)

Exercises

The following exercises are intended to demonstrate the effect of the `EPOLLET` and `EPOLLONESHOT` flags.

- 1 In separate windows, create two FIFOs and use `cat` to write to each FIFO:

```
mkfifo x
cat > x
```

```
mkfifo y
cat > y
```

- 2 Run the `i_epoll` program, using it to monitor both FIFOs for reading, specifying the `EPOLLET` flag for the FIFO `y`; note the file descriptor numbers used for each FIFO:

```
./i_epoll x:r y:re
```

- 3 Type some input into both `cat` commands, and then use the "p" command to perform an `epoll_wait()`:

```
i_epoll> p
```

You should find that both file descriptors report as ready for reading (`EPOLLIN`).


Exercises

- 4 Enter the “p” command again. You should find that only the FIFO `x` reports `EPOLLIN`. (`y` does not report as ready because no new input has appeared on the FIFO.)
- 5 Type some input into the `cat` command that is writing to the FIFO `y`, and once more use the “p” command to perform an `epoll_wait()`. You should find that both FIFOs report `EPOLLIN`. (`y` reports as ready again because new input has appeared on the FIFO.)
- 6 Switch the monitoring of the FIFO `y` to use `EPOLLET` and `EPOLLONESHOT` with the command “e <fd> reo”.
- 7 Type some input into the FIFO `y`, and then use the “p” command to perform an `epoll_wait()`. You should find that both `x` and `y` report `EPOLLIN`.
- 8 Type some more input into the FIFO `y`, and again use the “p” command to perform an `epoll_wait()`. You should find that `y` does not report as ready (because, after it reported as ready in the previous step, it was disabled in the interest list by `EPOLLONESHOT`).
- 9 Reenable the FIFO `y` in the interest list using the command “e <fd> re” and again use the “p” command to perform an `epoll_wait()`. You should find that `y` reports `EPOLLIN`.
- 10 Try any other experiments you might think of!

Outline

23	Alternative I/O Models	23-1
23.1	Overview	23-3
23.2	Nonblocking I/O	23-5
23.3	Signal-driven I/O	23-11
23.4	I/O multiplexing: poll()	23-14
23.5	Problems with poll() and select()	23-31
23.6	The epoll API	23-34
23.7	epoll events	23-45
23.8	epoll: edge-triggered notification	23-59
23.9	epoll: API quirks	23-70
23.10	Event-loop programming	23-76

epoll and duplication of file descriptors

- Entries in *epoll* interest list are associated with **combination** of file descriptor and open file description
 - Not just FD alone
-  Lifetime of interest list entry == lifetime of OFD
 - Can provide some surprises when FDs are duplicated...

epoll and duplication of file descriptors

- Suppose that *fd* in code below refers to a socket...

```
ev.events = EPOLLIN;
ev.data.fd = fd
epoll_ctl(epfd, EPOLL_CTL_ADD, fd, &ev);
newfd = dup(fd);
close(fd);
epoll_wait(epfd, ...);
```

- What happens if some input now arrives on the socket?
- *epoll_wait()* might still return events for registration of *fd*
 - Because open file description is still alive and present in interest list
 - OFD is kept alive by *newfd*
 - ⚠ Notifications return data given in registration of *fd*!!

epoll and duplication of file descriptors

- Analogous scenarios possible with *fork()*:

```
ev.events = EPOLLIN;
ev.data.fd = fd
epoll_ctl(epfd, EPOLL_CTL_ADD, fd, &ev);
if (fork() == 0) {
    /* Child continues, does not close 'fd' */
} else {
    close(fd);
    epoll_wait(epfd, ...);
}
```

epoll and duplication of file descriptors

- ⚠ Can't `EPOLL_CTL_DEL` *fd* after `close()`
 - \Rightarrow `EBADF`
- Must either:
 - Close duplicate FDs
 - ⚠ But you may not know about duplicate if it was created by a library function that used `dup()` or `fork()`
 - Or manually `EPOLL_CTL_DEL` *fd* **before** closing it