Systemprogrammering Föreläsning 10 Concurrent Servers

Topics

- Limitations of iterative servers
- Process-based concurrent servers
- Event-based concurrent servers
- Threads-based concurrent servers

Process-Based Concurrent Server

* echoserverp.c - A concurrent echo server based on processes * Usage: echoserverp <port> */ #include <ics.h> #define BUFSIZE 1024 void echo(int connfd); void sigchld_handler(int sig); int main(int argc, char **argv) { int listenfd, connfd, port, clientlen; struct sockaddr_in clientaddr; if (argc != 2) { fprintf(stderr, "usage: %s <port>\n", argv[0]); exit(0); 3 port = atoi(argv[1]); listenfd = Open_listenfd(port);

Three Basic Mechanisms for Creating Concurrent Flows

1. Processes

- Kernel automatically interleaves multiple logical flows.
- Each flow has its own private address space.

2. I/O multiplexing with select() or poll()

- User manually interleaves multiple logical flows.
- Each flow shares the same address space.
- Popular for high-performance server designs.

3. Threads

- Kernel automatically interleaves multiple logical flows.
- Each flow shares the same address space.
- Hybrid of processes and I/O multiplexing!

F10 - 2 -

Systemprogrammering 2007

Process-Based Concurrent Server (cont)

	Signal(SIGCHLD, sigchld_handler); /* parent must reap children! */
	/* main server loop */
	while (1) {
	<pre>clientlen = sizeof(struct sockaddr_in);</pre>
	<pre>connfd = Accept(listenfd, (struct sockaddr *) &clientaddr,</pre>
	<pre>&clientlen));</pre>
	if $(Fork() == 0)$ {
	Close(listenfd); /* child closes its listening socket */
	<pre>echo(connfd); /* child reads and echoes input line */</pre>
	Close(connfd); /* child is done with this client */
	exit(0); /* child exits */
	}
	Close(connfd); /* parent must close connected socket! */
	}
ı	<i>,</i>

Systemprogrammering 2007

Process-Based Concurrent Server (cont)

/* sigchld_handler - reaps children as they terminate */
void sigchld_handler(int sig) {
 while ((waitpid(-1, NULL, WNOHANG)) > 0)
 ;
 return;

F10 - 5 -

Systemprogrammering 2007

Pros and Cons of Process-Based Designs

- + Handles multiple connections concurrently
- + Clean sharing model
 - descriptors (no)
 - file tables (yes)
 - global variables (no)
- + Simple and straightforward.
- Additional overhead for process control.
- Nontrivial to share data between processes.
 - Requires IPC (interprocess communication) mechanisms FIFO's (named pipes), System V shared memory and semaphores

I/O multiplexing provides more control with less overhead...

Implementation Issues With Process-Based Designs

Server should restart accept call if it is interrupted by a transfer of control to the SIGCHLD handler

- Not necessary for systems with POSIX signal handling.
 - Our Signal wrapper tells kernel to automatically restart accept
- Required for portability on some older Unix systems.

Server must reap zombie children

- to avoid fatal memory leak.
- Can be avoided by the "Double fork" trick.

Server must close its copy of connfd.

- Kernel keeps reference for each socket.
- After fork, refcnt(connfd) = 2.
- Connection will not be closed until refcnt(connfd)=0.

```
F10 - 6 -
```

Systemprogrammering 2007

Event-Based Concurrent Servers Using I/O Multiplexing

Maintain a pool of connected descriptors.

Repeat the following forever:

- Use the Unix select function to block until:
 - (a) New connection request arrives on the listening descriptor.
 - (b) New data arrives on an existing connected descriptor.
- If (a), add the new connection to the pool of connections.
- If (b), read any available data from the connection
 - Close connection on EOF and remove it from the pool.

The select Function

select() sleeps until one or more file descriptors in the set readset are ready for reading.

#include <sys/select.h>

int select(int maxfdp1, fd_set *readset, NULL, NULL, NULL);

readset

• Opaque bit vector (max FD_SETSIZE bits) that indicates membership in a *descriptor set.*

• If bit k is 1, then descriptor k is a member of the descriptor set.

maxfdp1

- Maximum descriptor in descriptor set plus 1.
- Tests descriptors 0, 1, 2, ..., maxfdp1 1 for set membership.

select() returns the number of ready descriptors and sets each bit of readset to indicate the ready status of its corresponding descriptor.

F10 - 9 -

Systemprogrammering 2007

select Example

*
* main loop: wait for connection request or stdin command.
* If connection request, then echo input line
* and close connection. If stdin command, then process.
*/
<pre>printf("server> ");</pre>
fflush(stdout);
fdset readfs;
while (notdone) {
/*
* select: check if the user typed something to stdin or
* if a connection request arrived.
*/
<pre>FD_ZERO(&readfds); /* initialize the fd set */</pre>
<pre>FD_SET(listenfd, &readfds); /* add socket fd */</pre>
<pre>FD_SET(0, &readfds); /* add stdin fd (0) */</pre>
<pre>Select(listenfd+1, &readfds, NULL, NULL, NULL);</pre>

F10 - 11 -

Systemprogrammering 2007

Macros for Manipulating Set Descriptors

void FD_SET(int fd, fd_set *fdset);
 Turn on bit fd in fdset.

F10 - 10 -

Systemprogrammering 2007

select Example (cont)

First we check for a pending event on stdin.

<pre>/* if the user has typed a command, process it */ if (FD_ISSET(0, &readfds)) { forts(buf, BUFSIZE, stdin);</pre>
switch (buf[0]) {
case $'c'$: /* print the connection count */
printf("Received %d conn. requests so far.\n", connectcnt);
<pre>printf("server> ");</pre>
<pre>fflush(stdout);</pre>
break;
<pre>case 'q': /* terminate the server */</pre>
<pre>notdone = 0;</pre>
break;
default: /* bad input */
<pre>printf("ERROR: unknown command\n");</pre>
<pre>printf("server> ");</pre>
<pre>fflush(stdout);</pre>
}
}

Systemprogrammering 2007

select Example (cont)

Next we check for a pending connection request.

```
/* if a connection request has arrived, process it */
if (FD_ISSET(listenfd, &readfds)) {
    connfd = Accept(listenfd,
                           (struct sockaddr *) &clientaddr, &clientlen);
    connectcnt++;
    bzero(buf, BUFSIZE);
    Rio_readn(connfd, buf, BUFSIZE);
    Rio_writen(connfd, buf, strlen(buf));
    Close(connfd);
    }
} /* while */
```

F10 - 13 -

Systemprogrammering 2007

Event-based Concurrent Server (cont)

```
int main(int argc, char **argv)
     int listenfd, connfd, clientlen;
     struct sockaddr_in clientaddr;
     static pool pool;
     if (argc != 2) {
          fprintf(stderr, "usage: %s <port>\n", argv[0]);
          exit(0);
     ł
     port = atoi(argv[1]);
     listenfd = Open_listenfd(port);
     init_pool(listenfd, &pool);
      while (1) {
          pool.ready set = pool.read set;
          pool.nready = Select(pool.maxfd+1, &pool.ready_set,
                               NULL, NULL, NULL);
          if (FD_ISSET(listenfd, &pool.ready_set)) {
              clientlen = sizeof(struct sockaddr_in);
              connfd = Accept(listenfd, (SA *)&clientaddr, &clientlen);
              add_client(connfd, &pool);
          check_clients(&pool);
      ł
                                                                    Systemprogrammering 2007
F10 - 15 -
```

Event-based Concurrent Echo Server

```
* echoservers.c - A concurrent echo server based on select
#include "csapp.h"
typedef struct { /* represents a pool of connected descriptors */
   int maxfd:
                  /* largest descriptor in read set */
   fd_set read_set; /* set of all active descriptors */
   fd set ready set; /* subset of descriptors ready for reading */
                   /* number of ready descriptors from select */
   int nready;
   int maxi;
                   /* highwater index into client array */
   rio t clientrio[FD SETSIZE]; /* set of active read buffers */
} pool;
void init_pool(int listenfd, pool *p);
void add_clent(int connfd, pool *p);
void check_clients(pool *p);
```

int byte_cnt = 0; /* counts total bytes received by server */

F10 - 14 -

Systemprogrammering 2007

Event-based Concurrent Server (cont)

```
/* initialize the descriptor pool */
void init_pool(int listenfd, pool *p)
{
    /* Initially, there are no connected descriptors */
    int i;
    p->maxi = -1;
    for (i=0; i< FD_SETSIZE; i++)
        p->clientfd[i] = -1;
    /* Initially, listenfd is only member of select read set */
    p->maxfd = listenfd;
    FD_ZERO(&p->read_set);
    FD_SET(listenfd, &p->read_set);
}
```



Pro and Cons of Event-Based Designs

- + One logical control flow.
- + Can single-step with a debugger.
- + No process or thread control overhead.
 - Design of choice for high-performance Web servers and search engines.
- Significantly more complex to code than process- or thread-based designs.
- Can be vulnerable to denial of service attack = How?

Threads provide a middle ground between processes and I/O multiplexing...

Event-based Concurrent Server (cont)



Traditional View of a Process

Process = process context + code, data, and stack



Systemprogrammering 2007

F10 - 20 -

Alternate View of a Process

Process = thread + code, data, and kernel context



Logical View of Threads

Threads associated with a process form a pool of peers.

Unlike processes which form a tree hierarchy



A Process With Multiple Threads

Multiple threads can be associated with a process

- Each thread has its own logical control flow (sequence of PC values)
- Each thread shares the same code, data, and kernel context
- Each thread has its own thread id (TID)



Concurrent Thread Execution

Two threads run concurrently (are concurrent) if their logical flows overlap in time.

Otherwise, they are sequential.



Systemprogrammering 2007

Thread C

- - - - - - -

Threads vs. Processes

How threads and processes are similar

- Each has its own logical control flow.
- Each can run concurrently.
- Each is context switched.

How threads and processes are different

- Threads share code and data, processes (typically) do not.
- Threads are somewhat less expensive than processes.
 - Process control (creating and reaping) is twice as expensive as thread control.
 - Linux/Pentium III numbers:
 - » ~20K cycles to create and reap a process.
 - » ~10K cycles to create and reap a thread.

F10 - 25 -

Systemprogrammering 2007

The Pthreads "hello, world" Program



Posix Threads (Pthreads) Interface

Pthreads: Standard interface for ~60 functions that

manipulate threads from C programs.

- Creating and reaping threads.
 - pthread_create
 - pthread_join
- Determining your thread ID
 - pthread_self
- Terminating threads
 - pthread_cancel
 - pthread_exit

F10 - 26 -

- exit [terminates all threads], ret [terminates current thread]
- Synchronizing access to shared variables
 - pthread_mutex_init
 - pthread_mutex_[un]lock
 - pthread_cond_init
- pthread_cond_[timed]wait

Systemprogrammering 2007

Execution of Threaded"hello, world"



Thread-Based Concurrent Echo Server

```
int main(int argc, char **argv)
ł
   int listenfd, *connfdp, port, clientlen;
   struct sockaddr in clientaddr;
   pthread t tid;
   if (argc != 2) {
        fprintf(stderr, "usage: %s <port>\n", argv[0]);
        exit(0):
   }
   port = atoi(argv[1]);
   listenfd = open listenfd(port);
   while (1) {
        clientlen = sizeof(struct sockaddr_in);
        connfdp = Malloc(sizeof(int));
        *connfdp = Accept(listenfd, (SA *) &clientaddr, &clientlen);
        Pthread_create(&tid, NULL, thread, connfdp);
   ł
```

F10 - 29 -

Systemprogrammering 2007

Systemprogrammering 2007

Issues With Thread-Based Servers

Must run "detached" to avoid memory leak.

- At any point in time, a thread is either joinable or detached.
- Joinable thread can be reaped and killed by other threads.
 - must be reaped (with pthread_join) to free memory resources.
- Detached thread cannot be reaped or killed by other threads.
 - resources are automatically reaped on termination.
- Default state is joinable.
 - use pthread_detach(pthread_self()) to make detached.

Must be careful to avoid unintended sharing.

- For example, what happens if we pass the address of connfd to the thread routine?
 - Pthread_create(&tid, NULL, thread, (void *)&connfd);

All functions called by a thread must be thread-safe

(next lecture)

Thread-Based Concurrent Server (cont)

- * thread routine */
 void *thread(void *vargp)
 {
 int connfd = *((int *)vargp);
 }
 - Pthread_detach(pthread_self());
 Free(vargp);

echo_r(connfd); /* reentrant version of echo() */
Close(connfd);
return NULL;

F10 - 30 -

Systemprogrammering 2007

Pros and Cons of Thread-Based Designs

+ Easy to share data structures between threads

- e.g., logging information, file cache.
- + Threads are more efficient than processes.
- --- Unintentional sharing can introduce subtle and hard-toreproduce errors!
 - The ease with which data can be shared is both the greatest strength and the greatest weakness of threads.
 - (next lecture)