

COP 5611: Operating Systems Design Principles

Presentation by:

Fahd Rafi

Saad Ali

Software Interrupts

Chapter 13

Software Interrupts

- Method to interrupt user mode operation by other processes or due to error
- Software Interrupt – Signal
- 20 software interrupts in UNIX
- 0 is no interrupt
- `u.u_signal[n]` specifies action on interrupt n

u.u_signal[n]

- Structure u lies in user.h
- It is the per process data area which is swapped out along with process
- Always contains data for the currently running process

u.u_signal[n]

- Operation to be performed on signal

u_signal[n]	when interrupt #n occurs
zero	the process will terminate itself;
odd non-zero	the software interrupt is ignored;
even non-zero	the value is taken as the address in user space of a procedure which should be executed forthwith.

SIGKIL

- $n=9$
- Distinguished from other interrupts and process always terminates on SIGKIL
- Supposed to remain '0' until the end of process

User Setup

- User can set up the action to be taken for any signal using the 'signal()' system call.
 `signal(2,1) // sets u_signal[2]=1;`
 (meaning it will be ignored due to odd number)
- `u_signal[SIGKIL]` cannot be modified

Causing Interrupt

- Set “p_sig” in process “proc” entry to interrupt number;
 - For example: `p->p_sig=SIGINT;`
- Since only one p_sig is provided, only one and most recent signal can be maintained.

Handling Interrupt

- The interrupt is always handled when the target process becomes active
 - Interrupts must wait till process becomes active
- If user-mode action is to be performed, the user mode stack is used

Tracing

- Tracing is implemented using software interrupts.
 - SIGTRC
- Parent can monitor the progress of a child process

Implementation

- Specify signal action:
 - `ssig()` – Specify action for signal
- Send signal:
 - `kill()` – Send signal to some process
- Other functions:
 - `psignal()` – Send signal to a process
 - `signal()` – Send signal to all processes from a terminal
 - `issig()` – To check if there is an outstanding interrupt
 - `psig()` – To implement action when `issig` returns true
 - `core()` – When core dump is indicated for a terminating process
 - `grow()` – To grow stack size when needed
 - `exit()` – Terminates the currently active process
 - `ptrace()` – Implements `ptrace` system call
 - `stop()` – To stop a process for debugging
 - `procxmt()` – Child carries out certain operations for parent when stopped

Code

ssig()

```
3614 ssig()
3615 {
3616     register a;
3617
3618     a = u.u_arg[0];
3619     if(a<=0 || a>=NSIG || a ==SIGKIL) {
3620         u.u_error = EINVAL;
3621         return;
3622     }
3623     u.u_ar0[R0] = u.u_signal[a];
3624     u.u_signal[a] = u.u_arg[1];
3625     if(u.u_procp->p_sig == a)
3626         u.u_procp->p_sig = 0;
3627 }
```

kill()

```
3630 kill()
3631 {
3632     register struct proc *p, *q;
3633     register a;
3634     int f;
3635
3636     f = 0;
3637     a = u.u_ar0[R0];
3638     q = u.u_procp;
3639     for(p = &proc[0]; p < &proc[NPROC]; p++) {
3640         if(p == q)
3641             continue;
3642         if(a != 0 && p->p_pid != a)
3643             continue;
3644         if(a==0&&(p->p_ttyp!=q->p_ttyp||p<=&proc[1]))
3645             continue;
3646         if(u.u_uid != 0 && u.u_uid != p->p_uid)
3647             continue;
3648         f++;
3649         psignal(p, u.u_arg[0]);
3650     }
3651     if(f == 0)
3652         u.u_error = ESRCH;
3653 }
3654 /* ----- */
```

psignal()

```
3963 psignal(p, sig)
3964 int *p;
3965 {
3966     register *rp;
3967
3968     if(sig >= NSIG)
3969         return;
3970     rp = p;
3971     if(rp->p_sig != SIGKIL)
3972         rp->p_sig = sig;
3973     if(rp->p_stat > PUSER)
3974         rp->p_stat = PUSER;
3975     if(rp->p_stat == SWAIT)
3976         setrun(rp);
3977 }
```

issig()

```
3991 issig()
3992 {
3993     register n;
3994     register struct proc *p;
3995
3996     p = u.u_procp;
3997     if(n = p->p_sig) {
3998         if (p->p_flag&STRC) {
3999             stop();
4000             if ((n = p->p_sig) == 0)
4001                 return(0);
4002         }
4003         if((u.u_signal[n]&1) == 0)
4004             return(n);
4005     }
4006     return(0);
4007 }
4008 /* ----- */
```


psig ()

```
4043 psig()
4044 {
4045     register n, p;
4046     register *rp;
4047
4048     rp = u.u_procp;
4049     n = rp->p_sig;
4050     rp->p_sig = 0;
4051     if((p=u.u_signal[n]) != 0) {
4052         u.u_error = 0;
4053         if(n != SIGINS && n != SIGTRC)
4054             u.u_signal[n] = 0;
4055         n = u.u_ar0[R6] - 4;
4056         grow(n);
4057         suword(n+2, u.u_ar0[RPS]);
4058         suword(n, u.u_ar0[R7]);
4059         u.u_ar0[R6] = n;
4060         u.u_ar0[RPS] = & ~TBIT;
4061         u.u_ar0[R7] = p;
4062         return;
4063     }
4064     switch(n) {
4065
4066     case SIGQUIT:
4067     case SIGINS:
4068     case SIGTRC:
4069     case SIGIOT:
4070     case SIGEMT:
4071     case SIGFPT:
4072     case SIGBUS:
4073     case SIGSEG:
4074     case SIGSYS:
4075         u.u_arg[0] = n;
4076         if(core())
4077             n += 0200;
4078     }
4079     u.u_arg[0] = (u.u_ar0[R0]<<8) | n;
4080     exit();
4081 }
```

Pipes – Chapter 21

“Pipe.c”

Pipes

- Used for creating Pipes
 - Pipe is a FIFO character list
 - One group of processes write other read
 - Intercommunication

Pipe.c

- Global Variable
 - PIPESZ (4096)
- Functions
 - pipe()
 - readp()
 - writep()
 - plock()
 - prele()

Structures


INODE – Focus of all file activities – Unique inode for each file

```
5659 struct      inode
5660 {
5661     char      i_flag;
5662     char      i_count;    /* reference count */
5663     int       i_dev;      /* device where inode resides */
5664     int       i_number;   /* i number, 1-to-1 with device
5665                          address */
5666     int       i_mode;
5667     char      i_nlink;    /* directory entries */
5668     char      i_uid;      /* owner */
5669     char      i_gid;      /* group of owner */
5670     char      i_size0;    /* most significant of size */
5671     char      *i_size1;   /* least sig */
5672     int       i_addr[8]; /* device addresses constituting file */
5673     int       i_lastr;    /* last logical block read (for
5674                          read-ahead) */
5675 } inode [NINODE];
```

```
5678 /* flags */
5679 #define ILOCK 01 /* inode is locked */
5680 #define IUPD 02 /* inode has been modified */
5681 #define IACC 04 /* inode access time to be updated */
5682 #define IMOUNT 010 /* inode is mounted on */
5683 #define IWANT 020 /* some process waiting on lock */
5684 #define ITEXT 040 /* inode is pure text prototype */
```

Structure .. File

- One file structure is allocated for each pipe call. It holds read write pointers associated with each open file/pipe



```
5507 struct    file
5508 {
5509     char    f_flag;
5510     char    f_count;        /* reference count */
5511     int     f_inode;        /* pointer to inode structure */
5512     char    *f_offset[2];   /* read/write character pointer */
5513 } file[NFILE];
```

Structures

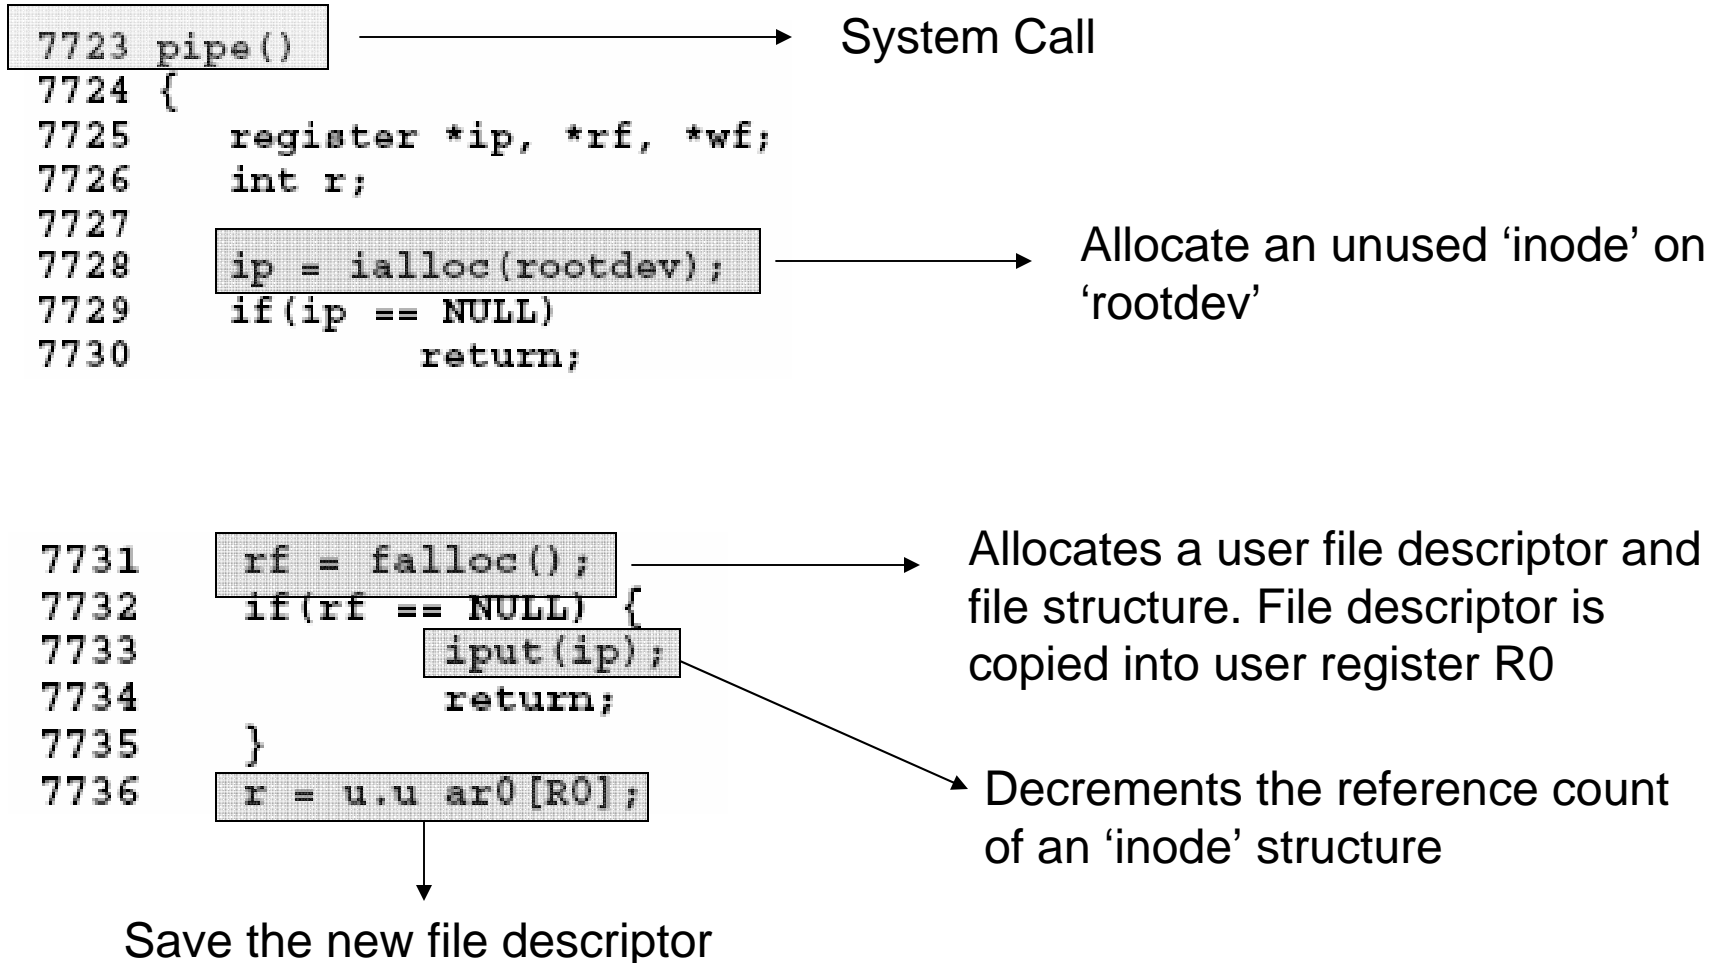
```
0413 struct user
0414 {
0415     int u_rsav[2];      /* save r5,r6 when exchanging stacks */
0416     int u_fsav[25];    /* save fp registers */
0417     /* rsav and fsav must be first in structure */
0418     char u_segflg;     /* flag for IO; user or kernel space */
0419     char u_error;     /* return error code */
0420     char u_uid;       /* effective user id */
0421     char u_gid;       /* effective group id */
0422     char u_ruid;      /* real user id */
0423     char u_rgid;      /* real group id */
0424     int u_procp;      /* pointer to proc structure */
0425     char *u_base;     /* base address for IO */
0426     char *u_count;    /* bytes remaining for IO */
0427     char *u_offset[2]; /* offset in file for IO */
0428     int *u_cdir;     /* pointer to inode for current directory */
0429     char u_dbuf[DIRSIZ]; /* current pathname component */
0430     char *u_dirp;    /* current pointer to inode */
0431     struct {
0432         int u_ino;
0433         char u_name[DIRSIZ];
0434     } u_dent;
0435     int *u_pdir;     /* inode of parent directory of dirp */
0436     int u_uisa[16]; /* prototype segmentation addresses */
0437     int u_uisd[16]; /* prototype segmentation descriptors */
0438     int u_ofile[NOFILE]; /* pointers to file structures of
0439                          open files */
0440     int u_arg[5];    /* arguments to current system call */
0441     int u_tsize;     /* text size (*64) */
0442     int u_dsize;     /* data size (*64) */
0443     int u_ssize;     /* stack size (*64) */
0444     int u_sep;       /* flag for I and D separation */
0445     int u_qsav[2];   /* label variable for quits & interrupts */
0446     int u_ssav[2];   /* label variable for swapping */
0447     int u_signal[NSIG]; /* disposition of signals */
0448     int u_utime;     /* this process user time */
0449     int u_stime;     /* this process system time */

```

Pipe System Call

- Allocate an inode for the root device
- Allocate a file table entry
- Remember file table entry in 'r' and allocate another file table entry
- Return user file identification in R0 and R1
- Complete the entries in 'file' and 'inode' structure.

Pipe - Code



```
7737 wf = falloc();
7738 if(wf == NULL) {
7739     rf->f_count = 0;
7740     u.u_ofile[r] = NULL;
7741     iput(ip);
7742     return;
7743 }
```

Allocates a user file descriptor and file structure. File descriptor is again copied into user register R0

Set pointer to file structure of read open file to NULL

```
7744 u.u_ar0[R1] = u.u_ar0[R0];
7745 u.u_ar0[R0] = r;
7746 wf->f_flag = FWRITE|FPIPE;
7747 wf->f_inode = ip;
7748 rf->f_flag = FREAD|FPIPE;
7749 rf->f_inode = ip;
```

Register R1 = Write File Descriptor
Register R0 = Read File Descriptor

Make inode pointer of both structures equal to same inode

```
7750 ip->i_count = 2;
7751 ip->i_flag = IACC|IUPD;
7752 ip->i_mode = IALLOC;
7753 }
```

Function – readp

- Two offsets are required:
 - For read
 - For write (write offset = filesize)
- Pass a file pointer to readp → Extract inode pointer from the file structure
- Lock the pipe
- Check if both reader and writer side of pipe is active: If not error
- Read and unlock the pipe

Readp - Code

```
7758 readp(fp)
```

```
7759 int *fp;
```

```
7760 {
```

```
7761     register *rp, *ip;
```

```
7762
```

```
7763     rp = fp;
```

```
7764     ip = rp->f_inode;
```

Pass a pointer of file structure from which has a pointer to inode of the pipe

Extract inode pointer

Readp – Code .. Cont'd

```
7765 loop:
7766     /* Very conservative locking.
7767     */
7768     plock(ip);
7769     /* If the head (read) has caught up with
7770     * the tail (write), reset both to 0.
7771     */
7772     if(rp->f_offset[1] == ip->i_size) {
7773         if(rp->f_offset[1] != 0) {
7774             rp->f_offset[1] = 0;
7775             ip->i_size = 0;
7776             if(ip->i_mode & IWRITE) {
7777                 ip->i_mode = & ~IWRITE;
7778                 wakeup(ip+1);
7779             }
7780         }
7781
7782         /* If there are not both reader and
7783         * writer active, return without
7784         * satisfying read.
7785         */
7786         prele(ip);
7787         if(ip->i_count < 2)
7788             return;
7789         ip->i_mode = | IREAD;
7790         sleep(ip+2, PPIPE);
7791         goto loop;
7792     }
```

→ Lock the inode

→ If offset becomes equal to size of the inode than reset

→ Wake up blocked writer

→ Raise the flag that I want to read and go to sleep

Readp – Code .. Cont'd

If every thing is fine than read and return:

```
7795     u.u_offset[0] = 0;
7796     u.u_offset[1] = rp->f_offset[1];
7797     readi(ip);
7798     rp->f_offset[1] = u.u_offset[1];
7799     prele(ip);
```

Function – writep()

- Lock the pipe
- Check if both reader and writer side of pipe is active: If not error
- If pipe is full wait for reader to consume characters
- Write desired number of bytes

Writep - Code

```
7805 writep(fp)
7806 {
7807     register *rp, *ip, c;
7808
7809     rp = fp;
7810     ip = rp->f_inode;
7811     c = u.u_count;
```


Writep - Code

```
7812 loop:
7813     /* If all done, return.
7814     */
7815     plock(ip);
7816     if(c == 0) {
7817         prele(ip);
7818         u.u_count = 0;
7819         return;
7820     }
7821     /* If there are not both read and
7822     * write sides of the pipe active,
7823     * return error and signal too.
7824     */
7825     if(ip->i_count < 2) {
7826         prele(ip);
7827         u.u_error = EPIPE;
7828         psignal(u.u_procp, SIGPIPE);
7829         return;
7830     }
7831     /* If the pipe is full,
7832     * wait for reads to delete
7833     * and truncate it.
7834     */
7835     if(ip->i_size1 == PIPSI1Z) {
7836         ip->i_mode |= IWRITE;
7837         prele(ip);
7838         sleep(ip+1, PPIPE);
7839         goto loop;
7840     }
```

→ No more bytes to write - return

→ Receive the signal that there are no more readers

→ Size reaches default size – no more writes can be done

Writep - Code

```
7844     u.u_offset[0] = 0;
7845     u.u_offset[1] = ip->i_size1;
7846     u.u_count = min(c, PIPSIZ-u.u_offset[1]);
7847     c -= u.u_count;
7848     writei(ip);
7849     prele(ip);
```

Function – plock()

- Locks a pipe before writing or reading
- If already locked:
 - Set the want bit
 - Sleep
- Otherwise:
 - Set the lock flag

Plock - Code

```
7862 plock(ip)
7863 int *ip;
7864 {
7865     register *rp;
7866
7867     rp = ip;
7868     while(rp->i_flag&ILOCK) {
7869         rp->i_flag = | IWANT;
7870         sleep(rp, PPIPE);
7871     }
7872     rp->i_flag = | ILOCK;
7873 }
```

Pass pointer of inode that we want to lock

Set the IWANT bit

Give up the processor till a wake up occurs on ip, at which the process enters the scheduling queue at priority PIPE.

```
0151 /* priorities: do not alter much */
0152
0153
0154 #define PSWP          -100
0155 #define PINOD         -90
0156 #define PRIBIO        -50
0157 #define PPIPE         1
0158 #define PWAIT         40
0159 #define PSLEP         90
0160 #define PUSER         100
```

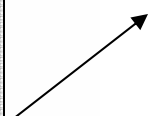
Function – prele()

- Unlocks the pipe after writing or reading
- If WANT bit is on:
 - Wakeup

Prele - Code

```
7882 prele(ip)
7883 int *ip;
7884 {
7885     register *rp;
7886
7887     rp = ip;
7888     rp->i_flag =& ~ILOCK;
7889     if(rp->i_flag&IWANT) {
7890         rp->i_flag =& ~IWANT;
7891         wakeup(rp);
7892     }
7893 }
```

Wake up all
processes waiting
on this inode



End