



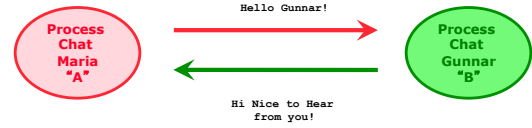
CSCI 6730/ 4730 Operating Systems

Dup & Pipe



Maria Hyönnette, UGA

Two Communicating Processes



- Concept that we want to implement

Maria Hyönnette, UGA

On the path to communication...

- Want: Communicating processes
 - » We start with 2
- Have so far: Forking – to create processes
- Problem:
 - » After fork() is called we end up with two **independent** processes.
 - » Separate Address Spaces
- Solution? How do we communicate?


Maria Hyönnette, UGA

Review 1730 - File: The Unix Way

- One easy way to communicate is to use **files**.
 - » Process A writes to a file and process B reads from it
- File descriptors
 - » Mechanism to work with files
 - » Used by low level I/O
 - Open(), close(), read(), write()
 - » **file descriptors** (the UNIX way) are generalized to other communication devices such as **pipes** and **sockets**

Maria Hyönnette, UGA

Big Picture (more on this later)



PCB

Maria Hyönnette, UGA

Pipe: Producer & Consumer


- Simple example: **who** | **sort**
 - » Both the writing process (**who**) and the reading process (**sort**) of a pipeline that executes concurrently.
- A pipe is usually implemented as an internal OS **buffer** with **2 file descriptors**.
 - » It is a resource that is concurrently accessed
 - by the reader and the writer, so it must be managed carefully (by the Kernel)

Maria Hyönnette, UGA

Buffering: Programming with Pipes

```
#include <unistd.h>
int pipe( int fd[2] );
```

- pipe() binds fd[] to two file descriptors:
 - » fds[0] used to read from pipe
 - » fds[1] used to write (stuff) to pipe
- Half-Duplex (one way) Communication
- Returns 0 if OK and -1 on error.



Example: pipe-yourself.c

```
#include <stdio.h>
#include <unistd.h>
#define MSGSIZE 16 /* null */

char *msg1="hello, world #1";
char *msg2="hello, world #2";
char *msg3="hello, world #3";


int main()
{
    char inbuf[MSGSIZE];
    int p[2], i;

    if( pipe( p ) < 0 )
        { /* open pipe */
            perror( "pipe" );
            exit( 1 );
        }

    write( p[1], msg1, MSGSIZE );
    write( p[1], msg2, MSGSIZE );
    write( p[1], msg3, MSGSIZE );

    for( i=0; i < 3; i++ )
        { /* read pipe */
            read( p[0], inbuf, MSGSIZE );
            printf( "%s\n", inbuf );
        }

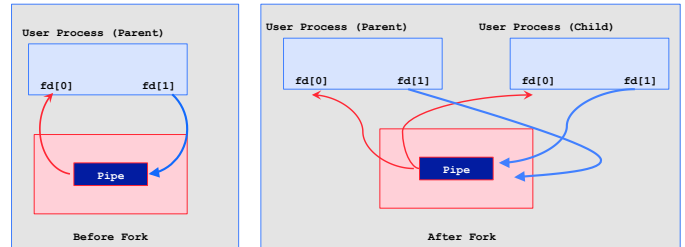
    return 0;
}
```



Things to Note

- Pipes uses FIFO ordering: *first-in first-out*.
- Read / write amounts **do not** need to be the same, but then text will be split differently.
- Pipes are most useful with fork() which creates an IPC connection between the parent and the child (or between the parents children)

What Happens After Fork?



- Design Question:
 - » Decide on : Direction of data flow – then close appropriate ends of pipe (at both parent and child)

Example: Parent Writes/Child Reads pipe-fork-close.c

```
#include <stdio.h>
#include <sys/wait.h>
#include <unistd.h>
#define MSGSIZE 16

char *msg1="hello, world #1";
char *msg2="hello, world #2";
char *msg3="hello, world #3";

int main()
{
    char inbuf[MSGSIZE];
    int p[2], i, pid;


    if( pipe( p ) < 0 )
        { /* open pipe */
            perror( "pipe" );
            exit( 1 );
        }

    if( (pid = fork()) < 0 )
        {
            perror( "fork" );
            exit( 2 );
        }

    if( pid > 0 ) /* parent */
        {
            close( p[0] ); /* read link */
            write( p[1], msg1, MSGSIZE );
            write( p[1], msg2, MSGSIZE );
            write( p[1], msg3, MSGSIZE );
            wait( (int *) 0 );
        }

    if( pid == 0 ) /* child */
        {
            close( p[1] ); /* write link */
            for( i=0; i < 3; i++ )
                {
                    read( p[0], inbuf, MSGSIZE );
                    printf( "%s\n", inbuf );
                }

            return 0;
        }
}
```



Some Rules of Pipes

- Every pipe has a size limit
 - » POSIX minimum is 512 bytes -- most systems makes this figure larger
- `read()` blocks if pipe is empty **and** there is a `write` link open to that pipe [it hangs]
- `read()` from a pipe whose `write()` end is closed **and** is empty returns 0 (indicates EOF) [but it doesn't hang]
 - » Lesson Learned:
 - Close write links or `read()` will never return *****
- `write()` to a pipe with no `read()` ends returns -1 and generates `SIGPIPE` and `errno` is set to `EPIPE`
- `write()` blocks if the pipe is full or there is not enough room to support the `write()`.
 - » May block in the middle of a `write()`

Maria Hyömette, UGA

1:

Pipes and `exec()`

How can we code `who | sort` ?

Observation: Writes to `stdout` and reads from `stdin`.

1. Use `exec()` to 'run' code in two processes (one runs `who` [child] and the other `sort` [parent]) which share a pipe (exec in child starts a new program within a copy of the 'parent' process).
2. Connect the pipe to `stdin` and `stdout` using `dup2()` .

Maria Hyömette, UGA

1:

Dup2

- Duplicate a pipe file descriptor to `stdin` or `stdout` (whichever is appropriate), e.g.,
 - » `dup2(pipefd, stdin)`, or
 - » `dup2(pipefd, stdout)`
- Now processes connected to pipe can read and write like it is from `stdin` and `stdout`
 - » **Caveat:** Beware of hanging on the 'pipe'
 - **Solution:** Close all file descriptors that comprise its pipes so that the pipes don't hang.

Maria Hyömette, UGA

1:

Duplicate File Descriptors

```
#include <unistd.h>
int dup2( int old-fd, int new-fd );
```

- Set one FD to the value of another.
- `new-fd` and `old-fd` now refer to the same file
- if `new-fd` is open [before copied over], it is first automatically closed
- Note that `dup2()` refer to fds not streams
- Example:

```
» dup2( fd[1], fileno(stdout));
```



Pipeline.c 11

Maria Hyömette, UGA


Example: `sort < file1.txt | uniq`

- What does this look like? How would a shell be programmed to process this?
 - » Well we know we need a parent & child to communicate through the pipe!
 - » Parent
 - » Child
 - » We need to open a file and read from it – and then read it as we read it from standard input.

Maria Hyömette, UGA

1:

Want: `sort < file1.txt | uniq`




- **Want:** How do we get there?

Maria Hyömette, UGA

1:

Want: "sort < file1 | uniq"




```
fileDES = open( "file1.txt", O_RDONLY );
```

Maria Hyönnelie, UGA

1!

Want: "sort < file1 | uniq"




```
fileDES = open( "myfile.txt", O_RDONLY );
dup2( fileDES, fileno( stdin ) );
```

Maria Hyönnelie, UGA

2!

Want: "sort < file1 | uniq"




```
fileDES = open( "myfile.txt", O_RDONLY );
dup2( fileDES, fileno( stdin ) );
close( fileDES );
```

Maria Hyönnelie, UGA

2

Want: "sort < file1 | uniq"




```
pipe( fd );
... fork() ...
```



Maria Hyönnelie, UGA

2!

Want: "sort < file1 | uniq"




```
fork();
/* now do the plumbing */
```

Maria Hyönnelie, UGA

2:

Want: "sort < file1 | uniq"




```
fork();
/* decide who does what */
```

Maria Hyönnelie, UGA

2:

Want: "sort < file1 | uniq"

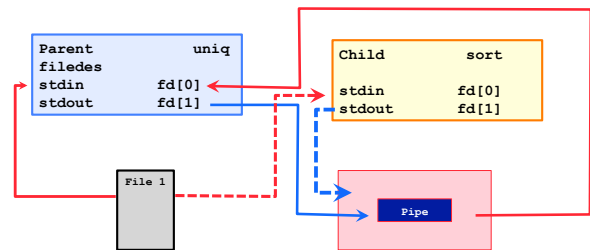


```
/* make writing to the pipe the same
/* as writing to stdout */
dup2( fd[1], fileno(stdout)); /* in green */
```

Maria Hyönnelie, UGA

2!

Want: "sort < file1 | uniq"




```
close(fd[0]); close(fd[1]); /* child */
/* leaving the ---- connections for child */
```

Maria Hyönnelie, UGA

2!

Want: "sort < file1 | uniq"




```
dup2(fd[0], fileno(stdin)); /* parent */
/* parent reads from pipe */
```

Maria Hyönnelie, UGA

2!

Want: "sort < file1 | uniq"



```
close(fd[1]); close(fd[0]); /* parent */
```

Maria Hyönnelie, UGA

2!

Example: "sort < file1 | uniq"

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

/* child | parent */
/* sort < file1.txt | uniq */
int main()
{
    int status;
    int fileDES;
    int pipeDES[2];
    pid_t pid;

    fileDES = open( "myfile.txt", O_RDONLY );
    dup2( fileDES, fileno( stdin ) );
    /* don't need to read via this one anymore */
    close( fileDES );

    /* create a child that communicate via a pipe */
    /* parent reads from pipe, child writes to pipe */
    pipe( pipeDES );

    pid = fork();
    if( pid < 0 )
    {
        perror( "fork" );
        exit( 1 );
    }
    else if( pid == 0 ) // child
    {
        close( pipeDES[ 0 ] );
        dup2( pipeDES[ 1 ], fileno( stdout ) );
        close( pipeDES[ 1 ] );
        execl( "/usr/bin/sort", "sort", (char *) 0 );
    }
    else if( pid > 0 ) // parent
    {
        close( pipeDES[ 1 ] );
        dup2( pipeDES[ 0 ], fileno( stdin ) );
        close( pipeDES[ 0 ] );
        execl( "/usr/bin/uniq", "uniq", (char *) 0 );
    }
}
```

Maria Hyönnelie, UGA

2!

Thought questions

- Other ways of designing this task?

Maria Hyönnelie, UGA

3!