


Processes

Creating and using multiple processes

Process State

Only one process can be running on a uniprocessor

The scheduler decides which of the ready processes to run.


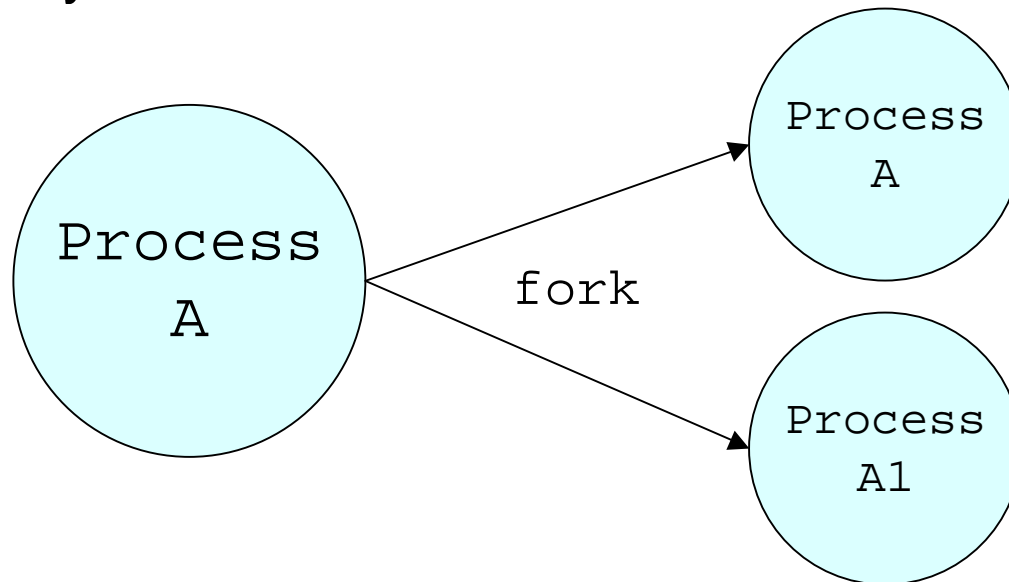


A process is ready if it could use the CPU immediately.

A process is blocked if it is waiting for an event (I/O, signal)

Fork

- The fork system call creates a duplicate of the currently running program.
- The duplicate (child process) and the original (parent process) both proceed from the point of the fork with exactly the same data.
- The only difference is the return value from the fork call.



Fork example

```
int main ()
{
    pid_t pid;
    pid = fork();
    if (pid < 0) {
        perror("fork()");
    } else if (pid > 0) {
        printf("parent\n");
    } else { /* pid == 0 */
        printf("child\n");
    }
    return 0;
}
```

Fork: PIDs and PPIDs

- System call: `int fork(void)`
 - If `fork()` succeeds it returns the child PID to the parent and returns 0 to the child;
 - If `fork()` fails, it returns -1 to the parent (no child is created) and sets `errno`
- Related system calls:
 - `int getpid()` – returns the PID of current process
 - `int getppid()` – returns the PID of parent process (ppid of 1 is 1)

When `fork()` fails

- There is a limit to the maximum number of processes a user can create.
- Once this limit is reached, subsequent calls to `fork()` return -1.

fork () properties

- Properties of parent inherited by child:
 - UID, GID
 - controlling terminal
 - CWD, root directory
 - signal mask, environment, resource limits
 - shared memory segments
- Differences between parent and child
 - PID, PPID, return value from fork()
 - pending alarms cleared for child
 - pending signals are cleared for child

Fork example

```
int i;
pid_t pid;

i = 5;
printf("%d\n", i);
pid = fork();

if (pid > 0)
    i = 6; /* only parent gets here */
else if (pid == 0)
    i = 4; /* only child gets here */
printf("%d\n", i);
```


Fork Example

Original process (parent)

```
int i; pid_t pid;
i = 5;
printf("%d\n", i);
/* prints 5 */
pid = fork();
/* pid == 677 */
if (pid > 0)
    i = 6;
else (pid == 0)
    i = 4;
printf("%d\n", i);
/* prints 6 */
```

Child process

```
int i; pid_t pid;
i = 5;
printf("%d\n", i);

pid = fork();
/* pid == 0 */
if (pid > 0)
    i = 6;
else if (pid == 0)
    i = 4;
printf("%d\n", i);
/* prints 4 */
```

PID/PPID Example

```
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>
int main() {
    pid_t pid;
    printf("ORIG: PID=%d PPID=%d\n",
        getpid(), getppid());
    pid = fork();
    if (pid > 0)
        printf("PARENT: PID=%d PPID=%d\n",
            getpid(), getppid());
    else if (pid == 0)
        printf("CHILD: PID=%d PPID=%d\n",
            getpid(), getppid());
    return(0);
}
```

Process Termination

Ch 5.6

- **Orphan** process:
 - a process whose parent is the init process (PID 1) because its original parent died before it did.
- Terminating a process: `exit()`
- Every normal process is a child of some parent, a terminating process sends its parent a `SIGCHLD` signal and waits for its termination status to be accepted.
- The Bourne shell stores the termination code of the last command in `$?`.

wait()



Ch 5.7

- System call to wait for a child
 - `pid_t wait(int *status)`
- A process that calls `wait()` can:
 - block (if all of its children are still running)
 - return immediately with the termination status of a child (if a child has terminated and is waiting for its termination status to be fetched)
 - return immediately with an error (if it doesn't have any child processes.)

Zombies

- A **zombie** process:
 - a process that is “waiting” for its parent to accept its return code
 - a parent accepts a child’s return code by executing `wait()`
 - shows up as `Z` in `ps -a`
 - A terminating process may be a (multiple) parent; the kernel ensures all of its children are orphaned and adopted by `init`.

wait and waitpid

- `wait()` can
 - block
 - return with termination status
 - return with error
- If there is more than one child `wait()` returns on termination of any children
- `waitpid` can be used to wait for a specific child pid.
- `waitpid` also has an option to block or not to block

wait and waitpid

- `waitpid` has an option to block or not to block
- `pid_t waitpid(pid, &status, option);`
 - if `pid == -1` → wait for any child
 - if `option == WNOHANG` → non-blocking
 - if `option == 0` → blocking
- `waitpid(-1, &status, 0);`
is equivalent to `wait(&status);`

Example of wait

```
#include <sys/types.h>
#include <sys/wait.h>
int main() {
    int status;
    if(fork() == 0) exit(7); /*normal*/
    wait(&status); prExit(status);
    if(fork() == 0) abort(); /*SIGABRT*/
    wait(&status); prExit(status);
    if(fork() == 0) status /= 0; /*FPE*/
    wait(&status); prExit(status);
    return 0;
}
```



prExit.c

```
#include <sys/types.h>
void prExit(int status) {
    if(WIFEXITED( status ) )
        printf("normal termination\n");
    else if(WIFSTOPPED( status ))
        printf("child stopped, signal no.= %d\n",
            WSTOPSIG(status));
    else if(WIFSIGNALED( status ) )
        printf("abnormal termination, "
            "signal no.= %d\n", WTERMSIG(status));
}
```

Exec

Ch 5.3

- The exec system call replaces the program being run by a process by a different one.
- The new program starts executing from the beginning.
- On success, exec never returns, on failure, exec returns -1.



Exec example

Program X

```
int i = 5;  
printf( "%d\n", i );  
exec( "Y" );  
printf( "%d\n", i );
```

Program Y

```
printf( "hello\n" );
```

exec properties

- New process inherits from calling process:
 - PID and PPID, real UID, GID
 - controlling terminal
 - CWD, root directory, resource limits
 - pending signals
 - pending alarms

exec ()

- Six versions exec():

```
execl(char *path, char *arg0, ..., (char *)NULL);
```

```
execv(char *path, char *argv[]);
```

```
execle(char *path, char *arg0, ..., (char *)NULL,  
        char *envp[]);
```

```
execve(char *pathname, char *argv[],  
        char *envp[]);
```

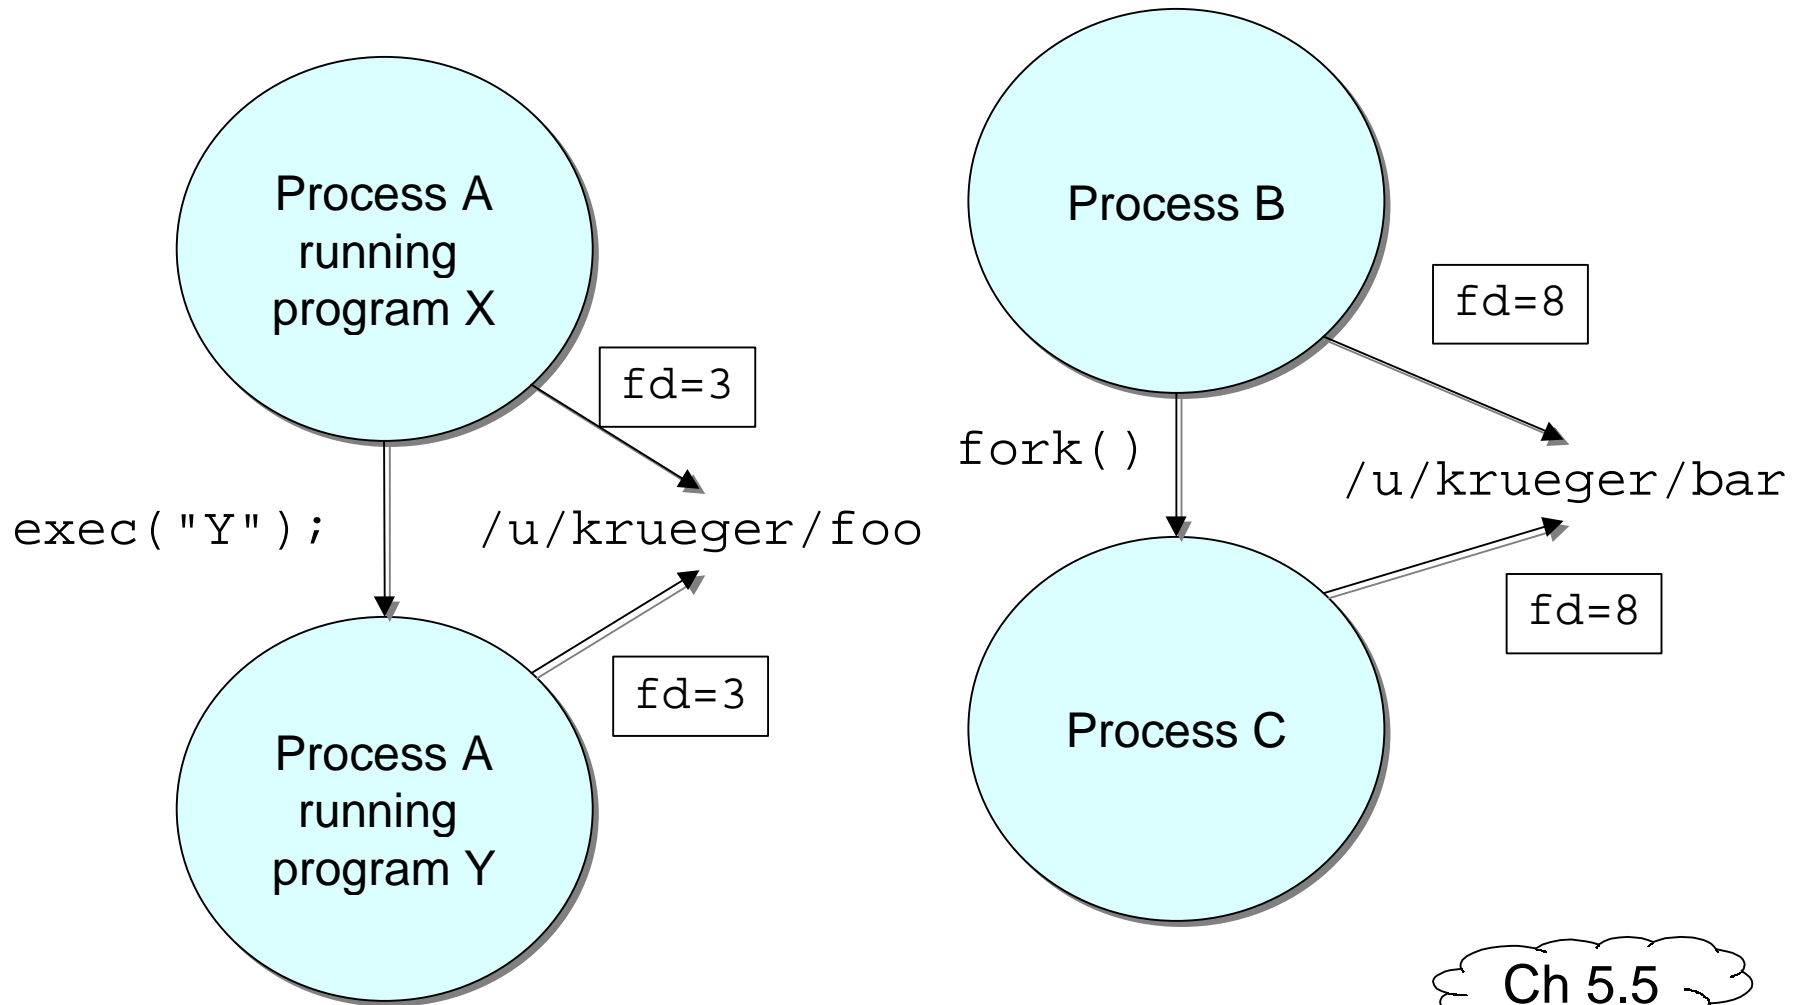
```
execlp(char *file, char *arg0, ..., (char *)NULL);
```

```
execvp(char *file, char *argv[]);
```

Processes and File Descriptors

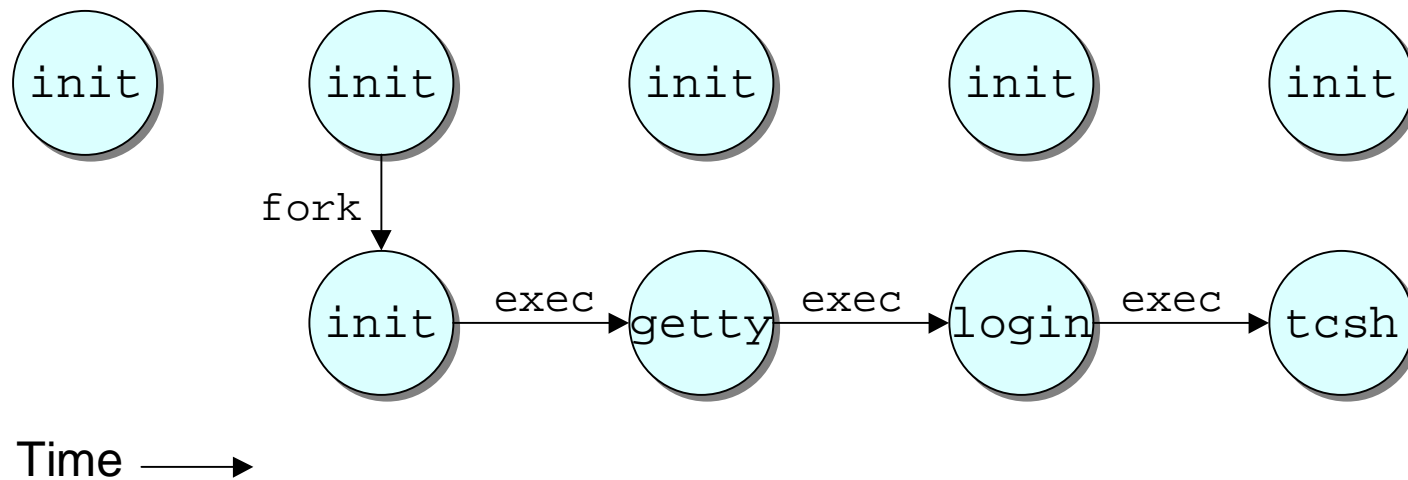
- File descriptors are handles to open files.
- They belong to processes not programs.
- They are a process's link to the outside world.

FDs preserved across fork and exec




Ch 5.5

Initializing Unix



- See “`top`”, “`ps -aux`” to see what’s running
- The **only** way to create a new process is to duplicate an existing process. Therefore the ancestor of **all** processes is `init` with `pid = 1`
- The **only** way to run a program is with `exec`

How csh runs commands



- When a command is typed, csh forks and then execs the typed command.
- After the fork, file descriptors 0, 1, and 2 still refer to stdin, stdout, and stderr in the new process.
- By convention, the executed program will use these descriptors appropriately.

How csh runs

