

















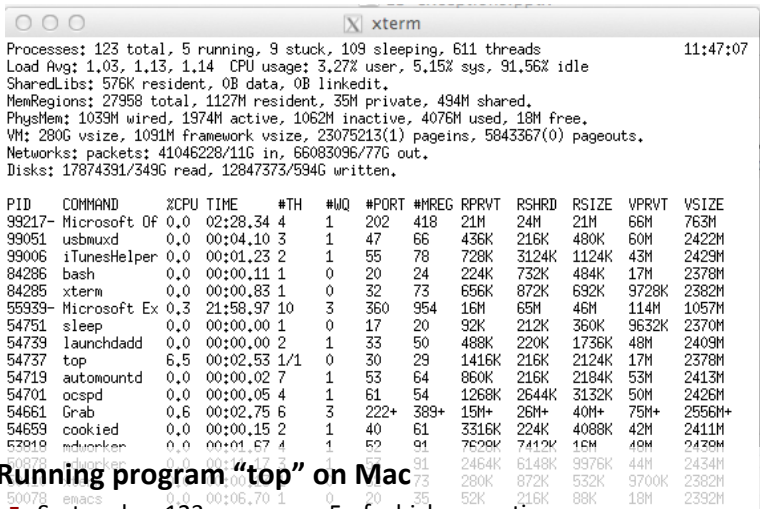








# Multiprocessing example (Mac OS)

A screenshot of a Mac OS terminal window titled 'xterm'. The terminal displays system statistics including total processes (123), active processes (5), CPU usage (3.27%), memory usage (494M shared), and disk activity. Below the statistics is a detailed table of processes. The table has 14 columns: PID, COMMAND, %CPU, TIME, #TH, #WQ, #PORT, #MREG, RPRVT, RSHRD, RSIZE, VPRVT, and VSIZE. The processes listed include various system services like Microsoft Office, usbmuxd, iTunesHelper, bash, xterm, sleep, launchdadd, top, automountd, ocspd, Grab, cookied, mdworker, and emacs. The process with PID 50078 (emacs) is highlighted in red in the original image.

| PID   | COMMAND        | %CPU | TIME     | #TH | #WQ | #PORT | #MREG | RPRVT | RSHRD | RSIZE | VPRVT | VSIZE  |
|-------|----------------|------|----------|-----|-----|-------|-------|-------|-------|-------|-------|--------|
| 99217 | - Microsoft Of | 0.0  | 02:28.34 | 4   | 1   | 202   | 418   | 21M   | 24M   | 21M   | 66M   | 763M   |
| 99051 | usbmuxd        | 0.0  | 00:04.10 | 3   | 1   | 47    | 66    | 436K  | 216K  | 480K  | 60M   | 2422M  |
| 99006 | iTunesHelper   | 0.0  | 00:01.23 | 2   | 1   | 55    | 78    | 728K  | 3124K | 1124K | 43M   | 2429M  |
| 84286 | bash           | 0.0  | 00:00.11 | 1   | 0   | 20    | 24    | 224K  | 732K  | 484K  | 17M   | 2378M  |
| 84285 | xterm          | 0.0  | 00:00.83 | 1   | 0   | 32    | 73    | 656K  | 872K  | 692K  | 9728K | 2382M  |
| 55939 | - Microsoft Ex | 0.3  | 21:58.97 | 10  | 3   | 360   | 954   | 16M   | 65M   | 46M   | 114M  | 1057M  |
| 54751 | sleep          | 0.0  | 00:00.00 | 1   | 0   | 17    | 20    | 92K   | 212K  | 360K  | 9632K | 2370M  |
| 54739 | launchdadd     | 0.0  | 00:00.00 | 2   | 1   | 33    | 50    | 488K  | 220K  | 1736K | 48M   | 2409M  |
| 54737 | top            | 6.5  | 00:02.53 | 1/1 | 0   | 30    | 29    | 1416K | 216K  | 2124K | 17M   | 2378M  |
| 54719 | automountd     | 0.0  | 00:00.02 | 7   | 1   | 53    | 64    | 860K  | 216K  | 2184K | 53M   | 2413M  |
| 54701 | ocspd          | 0.0  | 00:00.05 | 4   | 1   | 61    | 54    | 1268K | 2644K | 3132K | 50M   | 2426M  |
| 54661 | Grab           | 0.6  | 00:02.75 | 6   | 3   | 222+  | 389+  | 15M+  | 26M+  | 40M+  | 75M+  | 2556M+ |
| 54659 | cookied        | 0.0  | 00:00.15 | 2   | 1   | 40    | 61    | 3316K | 224K  | 4088K | 42M   | 2411M  |
| 53919 | mdworker       | 0.0  | 00:01.67 | 4   | 1   | 52    | 91    | 7628K | 7412K | 16M   | 49M   | 2428M  |
| 50978 | mdworker       | 0.0  | 00:14.17 | 3   | 1   | 57    | 91    | 2454K | 6148K | 9976K | 44M   | 2434M  |
| 50078 | emacs          | 0.0  | 00:06.70 | 1   | 0   | 20    | 35    | 52K   | 216K  | 88K   | 18M   | 2392M  |

## ■ Running program "top" on Mac

- System has 123 processes, 5 of which are active
- Identified by Process ID (PID)








# Multiprocessing example (Linux)

htop command



| PID   | USER     | PRI | NI  | VIRT  | RES   | SHR   | S | CPU% | MEM% | TIME+    | Command  |
|-------|----------|-----|-----|-------|-------|-------|---|------|------|----------|--|
| 1567  | kjfaneca | 20  | 0   | 2187M | 194M  | 72864 | S | 10.0 | 8.2  | 0:42.56  | /usr/lib/firefox-esr/firefox-esr   |
| 15716 | root     | 20  | 0   | 43164 | 35240 | 16968 | R | 6.0  | 0.0  | 1h39:47  | /opt/SUNWut/lib/Xnewt :49 -auth /var/lib/gdm/:49.Xauth -nolisten tcp -br +bs -terminat |
| 5180  | csadmin  | 20  | 0   | 17168 | 3332  | 1304  | R | 6.0  | 0.0  | 0:04.41  | htop   |
| 7394  | root     | 20  | 0   | 19564 | 11216 | 2692  | S | 5.0  | 0.0  | 2h20:32  | /opt/SUNWut/lib/Xnewt :16 -auth /var/lib/gdm/:16.Xauth -nolisten tcp -br +bs -terminat |
| 6483  | root     | 20  | 0   | 19564 | 11248 | 2688  | S | 4.0  | 0.0  | 2h03:50  | /opt/SUNWut/lib/Xnewt :23 -auth /var/lib/gdm/:23.Xauth -nolisten tcp -br +bs -terminat |
| 7476  | gdm      | 20  | 0   | 143M  | 19824 | 9568  | S | 3.0  | 0.0  | 2h12:43  | /usr/lib/gdm/gdmgreeter  |
| 6565  | gdm      | 20  | 0   | 143M  | 19812 | 9564  | S | 3.0  | 0.0  | 1h49:55  | /usr/lib/gdm/gdmgreeter  |
| 3979  | root     | 20  | 0   | 19564 | 11248 | 2688  | S | 3.0  | 0.0  | 2h03:15  | /opt/SUNWut/lib/Xnewt :37 -auth /var/lib/gdm/:37.Xauth -nolisten tcp -br +bs -terminat |
| 1625  | kjfaneca | 20  | 0   | 1800M | 110M  | 64000 | S | 2.0  | 0.1  | 0:28:87  | /usr/lib/firefox-esr/plugin-container -greomni /usr/lib/firefox-esr/omni.ja -apponmi   |
| 6164  | root     | 20  | 0   | 19564 | 11204 | 2692  | S | 2.0  | 0.0  | 1h37:27  | /opt/SUNWut/lib/Xnewt :46 -auth /var/lib/gdm/:46.Xauth -nolisten tcp -br +bs -terminat |
| 1063  | gdm      | 20  | 0   | 143M  | 19828 | 9576  | S | 2.0  | 0.0  | 1h28:49  | /usr/lib/gdm/gdmgreeter  |
| 6247  | gdm      | 20  | 0   | 143M  | 19816 | 9564  | S | 2.0  | 0.0  | 1h21:32  | /usr/lib/gdm/gdmgreeter  |
| 1607  | kjfaneca | 20  | 0   | 2187M | 194M  | 72864 | S | 0.0  | 0.2  | 0:03.91  | Compositor   |
| 20002 | dtsoniat | 20  | 0   | 2283M | 304M  | 84944 | S | 0.0  | 0.2  | 51:05.98 | /usr/lib/firefox-esr/firefox-esr   |
| 1606  | kjfaneca | 20  | 0   | 2187M | 194M  | 72864 | S | 0.0  | 0.2  | 0:01.48  | SoftwareVsyncTh  |
| 1628  | kjfaneca | 20  | 0   | 1800M | 110M  | 64000 | S | 0.0  | 0.1  | 0:02.29  | Chrome_ChildThr  |
| 1580  | kjfaneca | 20  | 0   | 2187M | 194M  | 72864 | S | 0.0  | 0.2  | 0:02.26  | Gecko_IOThread   |
| 1583  | kjfaneca | 20  | 0   | 2187M | 194M  | 72864 | S | 0.0  | 0.2  | 0:01.08  | Socket Thread  |
| 2024  | root     | 20  | 0   | 771M  | 2684  | 1288  | S | 0.0  | 0.0  | 1h20:55  | /usr/sbin/xe-daemon -p /var/run/xe-daemon.pid  |
| 2029  | root     | 20  | 0   | 771M  | 2684  | 1288  | S | 0.0  | 0.0  | 16:39.72 | xe-daemon  |
| 1     | root     | 20  | 0   | 10656 | 864   | 724   | S | 0.0  | 0.0  | 1:27.98  | init [2]   |
| 32677 | cliretti | 20  | 0   | 38960 | 2036  | 988   | S | 0.0  | 0.0  | 6:14.62  | tmux   |
| 32722 | cliretti | 20  | 0   | 12860 | 3668  | 1720  | S | 0.0  | 0.0  | 0:00.17  | -bash  |
| 32678 | cliretti | 20  | 0   | 12960 | 3776  | 1728  | S | 0.0  | 0.0  | 0:01.35  | -bash  |
| 31618 | root     | 20  | 0   | 3796  | 1020  | 860   | S | 0.0  | 0.0  | 0:00.00  | /opt/SUNWut/bin/utaction -i -t 900 -d /opt/SUNWut/lib/xcleanup                         |
| 31431 | root     | 20  | 0   | 3796  | 1020  | 860   | S | 0.0  | 0.0  | 0:00.00  | /opt/SUNWut/bin/utaction -i -t 900 -d /opt/SUNWut/lib/xcleanup                         |
| 31358 | matoups1 | 9   | -11 | 225M  | 4604  | 3016  | S | 0.0  | 0.0  | 0:00.17  | /usr/bin/pulseaudio --start  |
| 31378 | matoups1 | 9   | -11 | 225M  | 4604  | 3016  | S | 0.0  | 0.0  | 0:00.00  | null-sink  |
| 31377 | matoups1 | 20  | 0   | 116M  | 3228  | 2464  | S | 0.0  | 0.0  | 0:00.01  | /usr/lib/pulseaudio/pulse/gconf-helper   |
| 29687 | jmallery | 20  | 0   | 48496 | 3424  | 1900  | S | 0.0  | 0.0  | 0:01.74  | /usr/lib/x86_64-linux-gnu/gconf/gconfd-2   |
| 29683 | jmallery | 9   | -11 | 225M  | 4588  | 3016  | S | 0.0  | 0.0  | 0:00.08  | /usr/bin/pulseaudio --start  |
| 29697 | jmallery | 9   | -11 | 225M  | 4588  | 3016  | S | 0.0  | 0.0  | 0:00.00  | null-sink  |
| 29696 | jmallery | 20  | 0   | 116M  | 3228  | 2464  | S | 0.0  | 0.0  | 0:00.01  | /usr/lib/pulseaudio/pulse/gconf-helper   |
| 29672 | jmallery | 20  | 0   | 81460 | 3332  | 2724  | S | 0.0  | 0.0  | 0:16.17  | /usr/lib/gvfs/gvfsd-trash -s-spawner :1.9 /org/gtk/gvfs/exec_spaw/0                    |
| 29667 | jmallery | 20  | 0   | 351M  | 12428 | 9676  | S | 0.0  | 0.0  | 0:28.31  | /usr/lib/x86_64-linux-gnu/mate-panel/clock-applet                                      |





# Multiprocessing example (Linux)

glances command

TASKS 259 (936 thr), 1 run, 258 slp, 0 oth sorted automatically by cpu\_percent, tree view

| CPU% | MEM% | VIRT  | RES   | PID   | USER   | NI | S | TIME+    | IOR/s | IOW/s | Command   |
|------|------|-------|-------|-------|--------|----|---|----------|-------|-------|---|
| 0.0  | 0.0  | 181M  | 5.06M | 1     | root   | 0  | S | 3:34.32  | 0     | 0     | /sbin/init splash   |
| 0.0  | 0.0  | 342M  | 3.90M | 1222  | root   | 0  | S | 0:02.93  | 0     | 0     | ├─ /usr/sbin/lightdm  |
| 0.0  | 0.0  | 225M  | 5.69M | 1636  | root   | 0  | S | 0:00.59  | 0     | 0     | │   └─ lightdm --session-child 12 19                                    |
| 0.0  | 0.0  | 45.4M | 3.14M | 1992  | mtoups | 0  | S | 0:10.69  | 0     | 0     | │     └─ /sbin/upstart --user   |
| 0.0  | 5.1  | 3.62G | 811M  | 18749 | mtoups | 0  | S | 58:49.58 | 0     | 0     | │       └─ /usr/lib/ffmpeg/ffmpeg -private                              |
| 10.9 | 4.7  | 3.07G | 743M  | 11153 | mtoups | 0  | S | 11:58.84 | 0     | 0     | │         └─ /usr/lib/ffmpeg/ffmpeg -cont                               |
| 5.3  | 11.8 | 16.0G | 1.84G | 10769 | mtoups | 0  | S | 15:37.80 | 0     | 0     | │           └─ /usr/lib/ffmpeg/ffmpeg -cont                             |
| 0.3  | 3.0  | 2.86G | 471M  | 18813 | mtoups | 0  | S | 21:40.65 | 0     | 0     | │             └─ /usr/lib/ffmpeg/ffmpeg -cont                           |
| 0.0  | 3.7  | 2.79G | 595M  | 11186 | mtoups | 0  | S | 32:42.62 | 0     | 0     | │               └─ /usr/lib/ffmpeg/ffmpeg -cont                         |
| 0.0  | 0.6  | 1.84G | 88.1M | 5698  | mtoups | 0  | S | 1:47.94  | 0     | 0     | │                 └─ /usr/lib/ffmpeg/ffmpeg -cont                       |
| 12.5 | 4.9  | 3.65G | 783M  | 3056  | mtoups | 0  | S | 18:46.47 | 0     | 0     | │         └─ compiz   |
| 2.8  | 0.5  | 732M  | 81.3M | 7687  | mtoups | 0  | S | 1:39.29  | 0     | 0     | │           └─ /usr/lib/gnome-terminal/gnome-ter                        |
| 0.0  | 0.0  | 23.1M | 6.00M | 2303  | mtoups | 0  | S | 0:00.40  | 0     | 0     | │             └─ bash   |
| 6.2  | 0.2  | 93.2M | 29.9M | 4096  | mtoups | 0  | R | 1:04.82  | 0     | 0     | │               └─ /usr/bin/python3 /usr/bin/g                          |
| 0.0  | 0.0  | 23.0M | 6.12M | 947   | mtoups | 0  | S | 0:00.21  | 0     | 0     | │                 └─ bash   |
| 0.3  | 0.0  | 47.9M | 5.09M | 2158  | mtoups | 0  | S | 0:00.12  | 0     | 0     | │                   └─ ssh csadmin@math209.cs.uno.                      |
| 0.0  | 0.0  | 22.9M | 5.97M | 27951 | mtoups | 0  | S | 0:00.60  | 0     | 0     | │                     └─ bash   |
| 0.0  | 0.2  | 192M  | 26.3M | 27974 | mtoups | 0  | S | 0:00.75  | 0     | 0     | │                       └─ vim FAIL                                     |
| 0.0  | 0.0  | 24.0M | 7.07M | 28066 | mtoups | 0  | S | 0:01.98  | 0     | 0     | │                         └─ bash                                       |
| 0.0  | 0.0  | 22.9M | 5.98M | 30713 | mtoups | 0  | S | 0:01.84  | 0     | 0     | │                           └─ bash                                     |
| 0.0  | 0.0  | 48.0M | 5.80M | 18229 | mtoups | 0  | S | 0:00.17  | 0     | 0     | │                             └─ ssh csadmin@autolab.cs.uno.            |
| 0.0  | 0.0  | 22.6M | 2.07M | 9877  | mtoups | 0  | S | 0:00.36  | 0     | 0     | │                               └─ bash                                 |
| 0.0  | 0.0  | 23.4M | 5.57M | 21113 | mtoups | 0  | S | 0:01.22  | 0     | 0     | │                                 └─ bash                               |
| 0.0  | 0.0  | 23.4M | 5.33M | 20982 | mtoups | 0  | S | 0:00.85  | 0     | 0     | │                                   └─ bash                             |
| 0.0  | 0.0  | 22.5M | 3.56M | 26694 | mtoups | 0  | S | 0:00.30  | 0     | 0     | │                                     └─ bash                           |
| 0.0  | 0.0  | 22.9M | 6.02M | 24970 | mtoups | 0  | S | 0:00.80  | 0     | 0     | │                                       └─ bash                         |
| 0.0  | 0.2  | 192M  | 27.4M | 28683 | mtoups | 0  | S | 0:00.78  | 0     | 0     | │   └─ vim README.md              |
| 0.0  | 0.0  | 22.9M | 6.00M | 16403 | mtoups | 0  | S | 0:00.25  | 0     | 0     | │   └─ bash                     |
| 0.0  | 0.0  | 23.0M | 5.86M | 15823 | mtoups | 0  | S | 0:00.40  | 0     | 0     | │   └─ bash                   |
| 0.0  | 0.0  | 22.9M | 6.02M | 28625 | mtoups | 0  | S | 0:00.90  | 0     | 0     | │   └─ bash                 |
| 0.0  | 0.0  | 23.0M | 6.00M | 12036 | mtoups | 0  | S | 0:00.15  | 0     | 0     | │   └─ bash               |
| 0.0  | 0.7  | 1.26G | 115M  | 2080  | mtoups | 0  | S | 0:12.25  | 0     | 0     | │   └─ evince 07ecf.pdf |

Shell lab

Intro to processes




Exceptional Control Flow





# Multiprocessing: The (traditional) reality




- Save current registers in memory






# Concurrent processes

- Each process is a logical control flow.
- Two processes *run concurrently* (are concurrent) if their flows overlap in time
- Otherwise, they are *sequential*
- Examples (running on single core):
  - Concurrent: A & B, A & C
  - Sequential: B & C



# Concurrent processes

- Control flows for concurrent processes are physically disjoint in time
- However, we can think of concurrent processes as running in parallel with each other













**From a programmer's perspective, we can think of a process as being in one of three states**

## ■ Running

- Process is either executing, or waiting to be executed and will eventually be *scheduled* (i.e., chosen to execute) by the kernel

## ■ Stopped

- Process execution is *suspended* and will not be scheduled until further notice (next lecture when we study signals)

## ■ Terminated

- Process is stopped permanently

- **Process becomes terminated for one of three reasons:**
  - Receiving a signal whose default action is to terminate (next lecture)
  - Returning from the `main` routine
  - Calling the `exit` function
- **`void exit(int status)`**
  - Terminates with an *exit status* of `status`
  - Convention: normal return status is 0, nonzero on error
  - Another way to explicitly set the exit status is to return an integer value from the main routine
- **`exit` is called **once** but **never** returns.**

# Example of exit()

You can run this program on systems-lab:

```
$ /home/CSCI2467/labs/misc/ch08/exit
```

```
$ /home/CSCI2467/labs/misc/ch08/exit 4
```

```
int main(int argc, char** argv)
{
    int value;

    if (argc < 2) { exit(7); }
    /* atoi converts char to int */
    value = atoi(argv[1]);

    printf(" value is %d\n", value);
    return value+1;
}
```

Notes:

- argc is count of arguments
- argv is vector of arguments (array of strings)
- atoi() converts characters to int

# Creating processes

The indispensable `fork()`

- ***Parent process* creates a new running *child process* by calling `fork`**

# Creating processes

## The indispensable `fork()`

- ***Parent process* creates a new running *child process* by calling `fork`**
- **`int fork(void)`**
  - Returns 0 to the child process, child's PID to parent process
  - Child is *almost* identical to parent:
    - Child get an identical (but separate) copy of the parent's virtual address space.
    - Child gets identical copies of the parent's open file descriptors
    - Child has a different PID than the parent



# Creating processes

## The indispensable `fork()`

- **Parent process creates a new running *child process* by calling `fork`**
- **`int fork(void)`**
  - Returns 0 to the child process, child's PID to parent process
  - Child is *almost* identical to parent:
    - Child get an identical (but separate) copy of the parent's virtual address space.
    - Child gets identical copies of the parent's open file descriptors
    - Child has a different PID than the parent
- **`fork` is interesting (and often confusing) because it is called *once* but returns *twice***

# Example from lab activity

(one line added)

```
void main()
{
    pid_t mypid;

    fork(); /* this line: crucial difference */
    mypid = getpid();

    printf(" my PID is: %d\n", mypid);
}
```

- What's the point of doing this?

# A more practical `fork()` use

- What's the point of doing this?
  - Program does twice as much stuff

# A more practical `fork()` use

- What's the point of doing this?
  - Program does twice as much stuff
- But we want to be able to create a child process that can do more than just duplicate the parent

# A more practical `fork()` use

- What's the point of doing this?
  - Program does twice as much stuff
- But we want to be able to create a child process that can do more than just duplicate the parent
- Crucial observation: `fork()` has two *distinct* return values

# A more practical `fork()` use

- What's the point of doing this?
  - Program does twice as much stuff
- But we want to be able to create a child process that can do more than just duplicate the parent
- Crucial observation: `fork()` has two *distinct* return values
- Solution: we can use an `if( )` statement in order to have the parent and child do different things

# Example from lab activity

```
void main()
{
    int mypid, forkresult;

    printf("program starts, pid is: %d\n", getpid());

    forkresult = fork();

    mypid = getpid();

    if (forkresult == 0) {
        printf("ZERO fork returned %d, my PID is: %d\n",
            forkresult, mypid);
    } else {
        printf("NONZERO fork returned %d, my PID is: %d\n"
            , forkresult, mypid);
    }
}
```



# Running the fork example

When you ran this in class, you got a unique PID (each time)

```
program starts, pid is: 8805  
NONZERO fork returned 8806, my PID is: 8805  
ZERO fork returned 0, my PID is: 8806
```

# fork() is essential!

- This is not an obscure detail
- This is what makes a multi-process system possible!
- All processes begin with `fork()`  
(called by `init`, `bash`, `sshd`, `firefox`, etc)

... also crucial to Shell Lab

# fork example

See similar code on systems-lab: `/home/CSCI2467/labs/misc/forks.c`

```
int main()
{
    pid_t pid;
    int x = 1;

    pid = Fork();
    if (pid == 0) { /* Child */
        printf("child : x=%d\n", ++x);
        exit(0);
    }

    /* Parent */
    printf("parent: x=%d\n", --x);
    exit(0);
}
```

*fork.c*

```
linux> ./fork
parent: x=0
child : x=2
```

- Call once, return twice
- Concurrent execution
  - Can't predict execution order of parent and child
- Duplicate but separate address space
  - `x` has a value of 1 when fork returns in parent and child
  - Subsequent changes to `x` are independent
- Shared open files
  - `stdout` is the same in both parent and child

- **A *process graph* is a useful tool for capturing the partial ordering of statements in a concurrent program:**
  - Each vertex is the execution of a statement
  - $a \rightarrow b$  means  $a$  happens before  $b$
  - Edges can be labeled with current value of variables
  - `printf` vertices can be labeled with output
  - Each graph begins with a vertex with no inedges
- **Any *topological sort* of the graph corresponds to a feasible total ordering.**
  - Total ordering of vertices where all edges point from left to right


# Process graph example

```
int main()
{
    pid_t pid;
    int x = 1;

    pid = Fork();
    if (pid == 0) { /* Child */
        printf("child : x=%d\n", ++x);
        exit(0);
    }

    /* Parent */
    printf("parent: x=%d\n", --x);
    exit(0);
}
```

*fork.c*




# Process graph example (rotated)

```
int main()
{
    pid_t pid;
    int x = 1;


    pid = Fork();
    if (pid == 0) { /* Child */
        printf("child : x=%d\n", ++x);
        exit(0);
    }

    /* Parent */
    printf("parent: x=%d\n", --x);
    exit(0);
} fork.c
```




# Interpreting process graphs


## ■ Original graph:




## ■ Relabelled graph:



## Feasible total ordering:




## Infeasible total ordering:



# fork() example: two consecutive fork()'s

```
void fork2()
{
    printf("L0\n");
    fork();
    printf("L1\n");
    fork();
    printf("Bye\n");
}
```

*forks.c*



Feasible output:

L0  
L1  
Bye  
Bye  
L1  
Bye  
Bye

Infeasible output:


L0  
Bye  
L1  
Bye  
L1  
Bye  
Bye



# fork() example: nested fork()s in parent

```
void fork4()
{
    printf("L0\n");
    if (fork() != 0) {
        printf("L1\n");
        if (fork() != 0) {
            printf("L2\n");
        }
    }
    printf("Bye\n");
}
```

*forks.c*



**Feasible output:**

L0  
L1  
Bye  
Bye  
L2  
Bye

**Infeasible output:**

L0  
Bye  
L1  
Bye  
Bye  
L2

# fork() example: ./forks 4 has multiple feasible orderings


```
mtoups@x1ubuntu:/tmp$ ./forks 4
L0
L1
L2
Bye
Bye
Bye
Bye
Bye
mtoups@x1ubuntu:/tmp$ ./forks 4
L0
L1
Bye
L2
Bye
Bye
Bye
mtoups@x1ubuntu:/tmp$ █
```

```
mtoups -- -bash -- 80x24
L0
L1
L2
Bye
Bye
Bye
Bye
[Ms-MacBook-Pro:~ mtoups$ ./forks 4
L0
L1
L2
Bye
Bye
Bye
Bye
[Ms-MacBook-Pro:~ mtoups$ ./forks 4
L0
L1
Bye
Bye
L2
Bye
Bye
Bye
Ms-MacBook-Pro:~ mtoups$ █
```

# fork() example: nested fork()s in children

```
void fork5()
{
    printf("L0\n");
    if (fork() == 0) {
        printf("L1\n");
        if (fork() == 0) {
            printf("L2\n");
        }
    }
    printf("Bye\n");
}
```

*forks.c*



Feasible output:

L0  
Bye  
L1  
L2  
Bye  
Bye  
Bye

Infeasible output:

L0  
Bye  
L1  
Bye  
Bye  
L2

## ■ Idea

- When process terminates, it still consumes system resources
  - Examples: Exit status, various OS tables
- Called a “zombie”
  - Living corpse, half alive and half dead

## ■ Reaping

- Performed by parent on terminated child (using `wait` or `waitpid`)
- Parent is given exit status information
- Kernel then deletes zombie child process

## ■ What if parent doesn't reap?

- If any parent terminates without reaping a child, then the orphaned child will be reaped by `init` process (`pid == 1`)
- So, only need explicit reaping in long-running processes
  - e.g., shells and servers

# Zombie Example

```
void fork7() {  
    if (fork() == 0) {  
        /* Child */  
        printf("Terminating Child, PID = %d\n", getpid());  
        exit(0);  
    } else {  
        printf("Running Parent, PID = %d\n", getpid());  
        while (1)  
            ; /* Infinite loop */  
    }  
}
```

*forks.c*

```
linux> ./forks 7 &  
[1] 6639  
Running Parent, PID = 6639  
Terminating Child, PID = 6640  
linux> ps  
  PID TTY          TIME CMD  
 6585 tty9          00:00:00 tcsh  
 6639 tty9          00:00:03 forks  
 6640 tty9          00:00:00 forks <defunct>  
 6641 tty9          00:00:00 ps  
linux> kill 6639  
[1] Terminated  
linux> ps  
  PID TTY          TIME CMD  
 6585 tty9          00:00:00 tcsh  
 6642 tty9          00:00:00 ps
```

■ **ps** shows child process as "defunct" (i.e., a zombie)

■ Killing parent allows child to be reaped by **init**

# Non-terminating Child Example

```
void fork8()
{
    if (fork() == 0) {
        /* Child */
        printf("Running Child, PID = %d\n",
              getpid());
        while (1)
            ; /* Infinite loop */
    } else {
        printf("Terminating Parent, PID = %d\n",
              getpid());
        exit(0);
    }
}
```

*forks.c*

```
linux> ./forks 8
Terminating Parent, PID = 6675
Running Child, PID = 6676
linux> ps
  PID TTY          TIME CMD
 6585 ttyp9        00:00:00 tcsh
 6676 ttyp9        00:00:06 forks
 6677 ttyp9        00:00:00 ps
linux> kill 6676
linux> ps
  PID TTY          TIME CMD
 6585 ttyp9        00:00:00 tcsh
 6678 ttyp9        00:00:00 ps
```

- Child process still active even though parent has terminated
- Must kill child explicitly, or else will keep running indefinitely

# wait(): synchronizing with children

- Parent reaps a child by calling the `wait` function

- `int wait(int *child_status)`

- Suspends current process until one of its children terminates
- Return value is the `pid` of the child process that terminated
- If `child_status != NULL`, then the integer it points to will be set to a value that indicates reason the child terminated and the exit status:
  - Checked using macros defined in `wait.h`
    - `WIFEXITED`, `WEXITSTATUS`, `WIFSIGNALED`,  
`WTERMSIG`, `WIFSTOPPED`, `WSTOPSIG`,  
`WIFCONTINUED`
    - See textbook for details


# wait(): synchronizing with children

```

void fork9() {
  int child_status;

  if (fork() == 0) {
    printf("HC: hello from child\n");
    exit(0);
  } else {
    printf("HP: hello from parent\n");
    wait(&child_status);
    printf("CT: child has terminated\n");
  }
  printf("Bye\n");
}
    
```

*forks.c*



**Feasible output:**

HC  
HP  
CT  
Bye

**Infeasible output:**

HP  
CT  
Bye  
HC





# waitpid(): waiting for a specific process

- `pid_t waitpid(pid_t pid, int &status, int options)`
  - Suspends current process until specific process terminates
  - Various options (see textbook)

```
void fork11() {
    pid_t pid[N];
    int i;
    int child_status;

    for (i = 0; i < N; i++)
        if ((pid[i] = fork()) == 0)
            exit(100+i); /* Child */
    for (i = N-1; i >= 0; i--) {
        pid_t wpid = waitpid(pid[i], &child_status, 0);
        if (WIFEXITED(child_status))
            printf("Child %d terminated with exit status %d\n",
                wpid, WEXITSTATUS(child_status));
        else
            printf("Child %d terminate abnormally\n", wpid);
    }
}
```

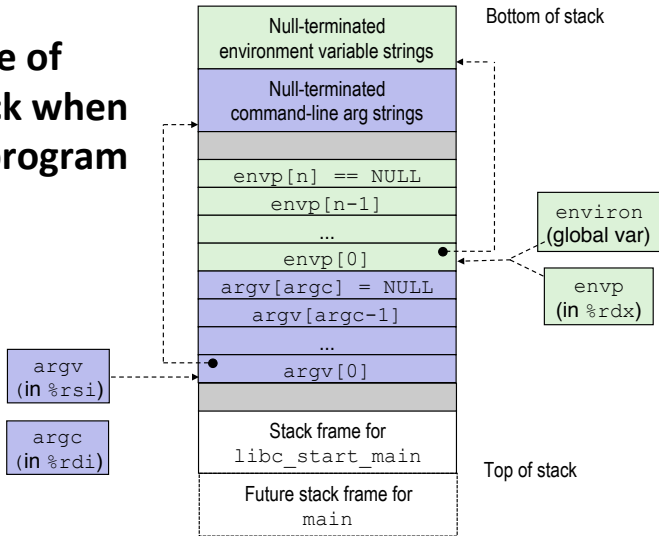
*forks.c*

# execve(): loading and running programs

- `int execve(char *filename, char *argv[], char *envp[])`
- **Loads and runs in the current process:**
  - Executable file **filename**
    - Can be object file or script file beginning with `#!interpreter` (e.g., `#!/bin/bash`)
  - ...with argument list **argv**
    - By convention `argv[0]==filename`
  - ...and environment variable list **envp**
    - “name=value” strings (e.g., `USER=droh`)
    - `getenv`, `putenv`, `printenv`
- **Overwrites code, data, and stack**
  - Retains PID, open files and signal context
- Called **once** and **never** returns
  - ...except if there is an error



# Structure of the stack when a new program starts











- Each x86-64 system call has a unique ID number
- Examples:

| <i>Number</i> | <i>Name</i> | <i>Description</i>     |
|---------------|-------------|------------------------|
| 0             | read        | Read file              |
| 1             | write       | Write file             |
| 2             | open        | Open file              |
| 3             | close       | Close file             |
| 4             | stat        | Get info about file    |
| 57            | fork        | Create process         |
| 59            | execve      | Execute a program      |
| 60            | _exit       | Terminate process      |
| 62            | kill        | Send signal to process |

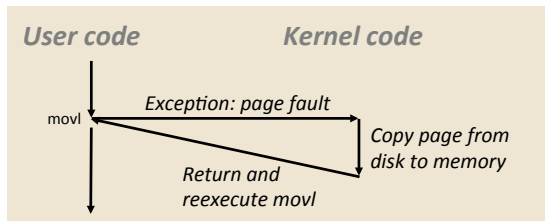


# Fault example (recoverable): page fault

- User writes to memory location
- That portion (page) of user's memory is currently on disk

```
int a[1000];  
main ()  
{  
    a[500] = 13;  
}
```


```
80483b7:    c7 05 10 9d 04 08 0d   movl   $0xd,0x8049d10
```



# Fault example (unrecoverable): invalid memory reference

```
int a[1000];
main ()
{
    a[5000] = 13;
}
```

```
80483b7: c7 05 60 e3 04 08 0d movl $0xd,0x804e360
```



- Sends **SIGSEGV** signal to user process
- User process exits with “segmentation fault”





