

# Lecture 22: Processes in xv6

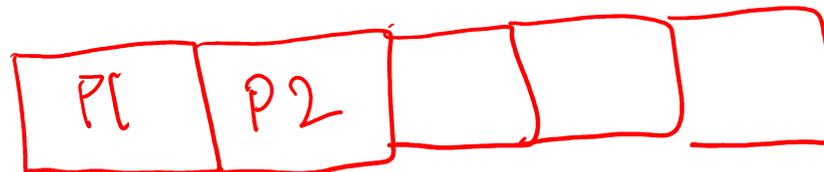
Mythili Vutukuru

IIT Bombay

<https://www.cse.iitb.ac.in/~mythili/os/>

# The process abstraction

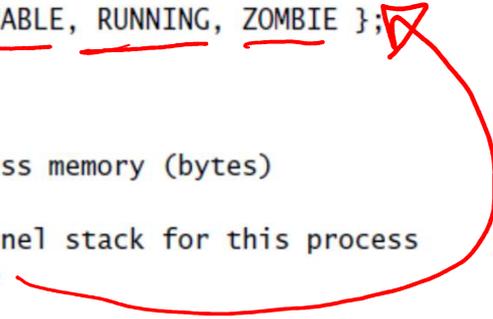
- The OS is responsible for concurrently running multiple processes (on one or more CPU cores/processors)
  - Create, run, terminate a process
  - Context switch from one process to another
  - Handle any events (e.g., system calls from process)
- OS maintains all information about an active process in a process control block (PCB)
  - Set of PCBs of all active processes is a critical kernel data structure
  - Maintained as part of kernel memory (part of RAM that stores kernel code and data, more on this later)
- PCB is known by different names in different OS
  - struct proc in xv6
  - task\_struct in Linux



# PCB in xv6: struct proc

- Page 23, process structure and process states

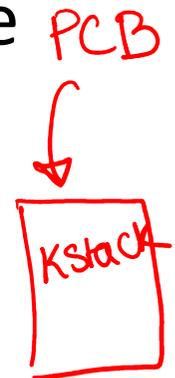
```
2334 enum procstate { UNUSED, EMBRYO, SLEEPING, RUNNABLE, RUNNING, ZOMBIE };
2335
2336 // Per-process state
2337 struct proc {
2338     uint sz;                // Size of process memory (bytes)
2339     pde_t* pgdir;          // Page table
2340     char *kstack;          // Bottom of kernel stack for this process
2341     enum procstate state; // Process state
2342     int pid;                // Process ID
2343     struct proc *parent;    // Parent process
2344     struct trapframe *tf;   // Trap frame for current syscall
2345     struct context *context; // swtch() here to run process
2346     void *chan;             // If non-zero, sleeping on chan
2347     int killed;             // If non-zero, have been killed
2348     struct file *ofile[NOFILE]; // Open files
2349     struct inode *cwd;      // Current directory
2350     char name[16];          // Process name (debugging)
2351 };
2352
```



# struct proc: kernel stack

```
2340 char *kstack; // Bottom of kernel stack for this process
```

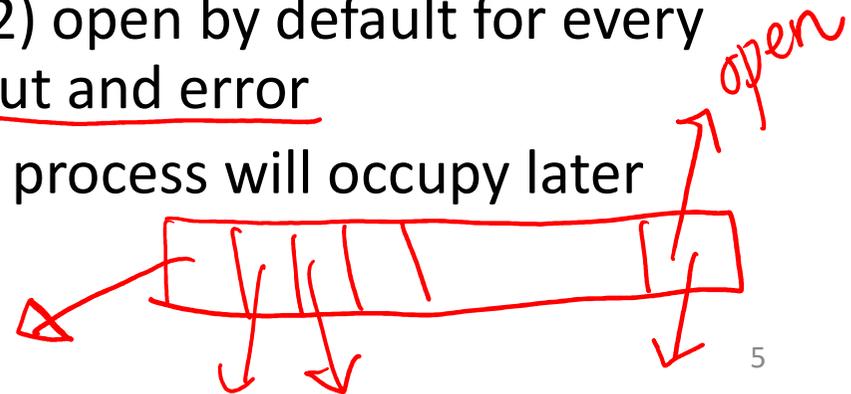
- Recall: register state (CPU context) saved on user stack during function calls, to restore/resume later
- Likewise, CPU context stored on kernel stack when process jumps into OS to run kernel code
  - Why separate stack? OS does not trust user stack
  - Separate area of memory per process within the kernel, not accessible by regular user code
  - Linked from struct proc of a process



# struct proc: list of open files

```
2348  struct file *ofile[NOFILE]; // Open files
```

- Array of pointers to open files (struct file has information about the open file, more on this later)
  - When user opens a file, a new entry is created in this array, and the index of that entry is passed as a file descriptor to user
  - Subsequent read/write calls on a file use this file descriptor to refer to the file
  - First 3 files (array indices 0,1,2) open by default for every process: standard input, output and error
  - Subsequent files opened by a process will occupy later entries in the array



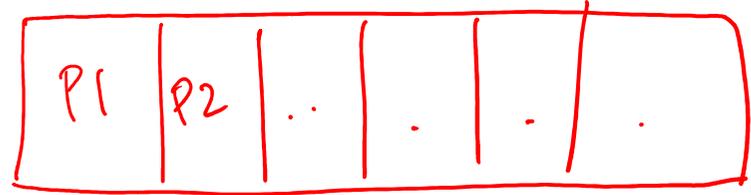
# struct proc: page table

```
2339  pde_t* pgdir;           // Page table
```

- Every instruction or data item in the memory image of process (code/data, stack, heap, etc.) has an address
  - Virtual addresses, starting from 0
  - Actual physical addresses in memory can be different (all processes cannot store their first instruction at address 0)
- Page table of a process maintains a mapping between the virtual addresses and physical addresses (more on this later)

# Process table (ptable) in xv6 *lock*

```
2409 struct {
2410     struct spinlock lock;
2411     struct proc proc[NPROC];
2412 } ptable;
```



- ptable: Fixed-size array of all processes
  - Real kernels have dynamic-sized data structures
- CPU scheduler in the OS loops over all runnable processes, picks one, and sets it running on the CPU

```
2768     // Loop over process table looking for process to run.
2769     acquire(&ptable.lock);
2770     for(p = ptable.proc; p < &ptable.proc[NPROC]; p++){
2771         if(p->state != RUNNABLE)
2772             continue;
2773
2774         // Switch to chosen process. It is the process's job
2775         // to release ptable.lock and then reacquire it
2776         // before jumping back to us.
2777         c->proc = p;
2778         switchvm(p);
2779         p->state = RUNNING;
```

# Process state transition examples

- A process that needs to sleep (e.g., for disk I/O) will set its state to SLEEPING and invoke scheduler
- A process that has run for its fair share will set itself to RUNNABLE (from RUNNING) and invoke scheduler
- Scheduler will once again find another RUNNABLE process and set it to RUNNING

```
2826 // Give up the CPU for one scheduling round.
2827 void
2828 yield(void)
2829 {
2830     acquire(&ptable.lock);
2831     myproc()->state = RUNNABLE;
2832     sched();
2833     release(&ptable.lock);
```

```
2873 void
2874 sleep(void *chan, struct spinlock *lk)
2875 {
2876     struct proc *p = myproc();
2877
2878     if(p == 0)
2879         panic("sleep");
2880
2881     if(lk == 0)
2882         panic("sleep without lk");
2883
2884     // Must acquire ptable.lock in order to
2885     // change p->state and then call sched.
2886     // Once we hold ptable.lock, we can be
2887     // guaranteed that we won't miss any wakeup
2888     // (wakeup runs with ptable.lock locked),
2889     // so it's okay to release lk.
2890     if(lk != &ptable.lock){
2891         acquire(&ptable.lock);
2892         release(lk);
2893     }
2894     // Go to sleep.
2895     p->chan = chan;
2896     p->state = SLEEPING;
2897     sched();
2898 }
2899
```

# Summary of xv6 processes

- We have seen basics of PCB structure (struct proc), list of processes (ptable), scheduler code, state transitions
- We will keep revisiting this xv6 code multiple times to understand it better
  - Each concept will deepen understanding further