

Systems Software &  
Architecture Lab.  
Seoul National University


2025.10.22

# Project #3: The Road to Balance



# CPU Affinity

- **What?**
  - **CPU affinity** ensures that a process (or kernel thread) executes only on a specified subset of CPUs.
- **Why?**
  - Cache locality and performance
  - NUMA locality
  - Task isolation
  - Balancing multi-threaded application
- **How?**
  - Let's see...



# CPU Affinity in Linux

- “task\_struct” structure
  - nr\_cpus\_allowed
  - user\_cpus\_ptr
  - cpus\_mask
- System calls
  - sched\_setaffinity()
  - sched\_getaffinity()

include/linux/sched.h - struct task\_struct

```
836
837     unsigned int         policy;
838     int                 nr_cpus_allowed;
839     const cpumask_t      *cpus_ptr;
840     cpumask_t            *user_cpus_ptr;
841     cpumask_t            cpus_mask;
842     void                 *migration_pending;
843 #ifdef CONFIG_SMP
844     unsigned short       migration_disabled;
845 #endif
846     unsigned short       migration_flags;
847
```

\*linux v6.9

Field	Description	Type
nr_cpus_allowed	Numbers of allowed cpus in cpus_mask	int
cpus_ptr	Effective mask of cpus allowed	cpumask_t
users_cpus_ptr	Mask set by user, set by sched_setaffinity()	cpumaks_t

# CPU Affinity in Linux

## ■ Check allowed cpus before

- Scheduling
- Migrating
- Load balancing
- ...

kernel/sched/core.c

```
2475 static inline bool is_cpu_allowed(struct task_struct *p, int cpu)
2476 {
2477     /* When not in the task's cpumask, no point in looking further. */
2478     if (!cpumask_test_cpu(cpu, p->cpus_ptr))
2479         return false;
2480
2481     /* migrate_disabled() must be allowed to finish. */
2482     if (is_migration_disabled(p))
2483         return cpu_online(cpu);
2484
2485     /* Non kernel threads are not allowed during either online or offline. */
2486     if (!(p->flags & PF_KTHREAD))
2487         return cpu_active(cpu) && task_cpu_possible(cpu, p);
2488
2489     /* KTHREAD_IS_PER_CPU is always allowed. */
2490     if (kthread_is_per_cpu(p))
2491         return cpu_online(cpu);
2492
2493     /* Regular kernel threads don't get to stay during offline. */
2494     if (cpu_dying(cpu))
```

# Taskset in Linux

```
Usage: taskset [options] [mask | cpu-list] [pid|cmd [args...]]
```

Show or change the CPU affinity of a process.

Options:

-a, --all-tasks	operate on all the tasks (threads) for a given pid
-p, --pid	operate on existing given pid
-c, --cpu-list	display and specify cpus in list format
-h, --help	display this help
-V, --version	display version

The default behavior is to run a new command:

```
taskset 03 sshd -b 1024
```

You can retrieve the mask of an existing task: → **sched\_getaffinity**

```
taskset -p 700
```

Or set it:

```
taskset -p 03 700 → sched_setaffinity
```

List format uses a comma-separated list instead of a mask:

```
taskset -pc 0,3,7-11 700
```

# Xv6 Scheduler

- Round-Robin
  - Scan for RUNNABLE proc
  - If found, set RUNNING, switch context
- Single global queue
  - Lock before checking each proc
  - Not scalable!

kernel/proc.c

```
10
11 struct proc proc_free_list[NPROC];
12
```

```
453 for(p = proc; p < &proc[NPROC]; p++) {
454     acquire(&p->lock);
455     if(p->state == RUNNABLE) {
456         // Switch to chosen process. It is the process's job
457         // to release its lock and then reacquire it
458         // before jumping back to us.
459         p->state = RUNNING;
460         c->proc = p;
461 #ifdef LOG
462     PRINTLOG_START;
463 #endif
464     swtch(&c->context, &p->context);
465
466     // Process is done running for now.
467     // It should have changed its p->state before coming back.
468     c->proc = 0;
469     found = 1;
470     }
471     release(&p->lock);
472 }
473 if(found == 0) {
474     // nothing to run; stop running on this core until an interrupt.
475     asm volatile("wfi");
476 }
```



# Part 1. CPU affinity with the pin() system call (10 points)

- `int pin(int mask)` system call
  - Sets or queries CPU affinity of the calling process
  - System call number assigned to 30 (no modifications required)

Mode	Mask Value	Description
SET	<code>mask &gt; 0</code>	The process may run only on RISC-V harts whose bits are set in mask
UNPIN	<code>mask == 0</code>	Allow the process to run on all online harts
QUERY	<code>mask == -1</code>	Returns the current CPU affinity mask of the calling process



## Part 2. Per-core run queues & process migration (30 points)

- Implement Per hart run-queue
  - Expand the existing single global run queue of xv6
  - Single queues maximum capacity is set the QPROC (kernel/param.h)
  - Round-robin policy inside its local queue
- Implement process migration
  - Implement a method to move tasks from one queue to another
  - Child process inherits parent's affinity mask
  - After pin(), if the current hart is no longer permitted, the process **must** be migrated





## Part 3. Load Balancing Policy

- Design and implement your own load balancing policy!
- Minimize the completion time of a given batch of CPU-intensive processes
- 10ms penalty whenever it runs on a different hart, so avoid needless migrations
  - Do not modify the skeleton code parts that refers to this
- Things you might consider...
  - Periodic rebalancing
  - Idle-time stealing
  - ...



# BONUS

- Using multiple test cases, completion time will be ranked
- 20% bonus to the top 10
- 10% bonus to the top 11 ~ 20



# Additional Notes

- You must pull the latest version of skeleton code. (Date: **10/23**)
  - An additional “mask” field to track the allowed harts to run the process
- Every details about your code should be in the document, and all the descriptions done in the document should be implemented in your code.



# Due date

- Due

- 11:59 PM, November 6 (Thursday)

- Submission

- Run the make submit command to generate a tarball named xv6-pa3-{STUDENTID}.tar.gz in the xv6-riscv-snu directory
- Upload the compressed file to the submission server
- The total number of submissions for this project will be limited to 30
- Only the version marked FINAL will be considered for the project score



Thank you!