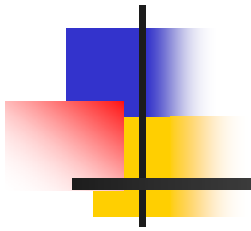

The O(1) Linux Scheduler

Alexandre David

B2-206

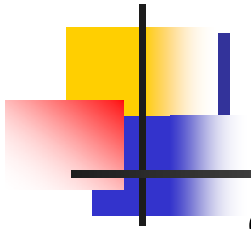

Introduction

- Scheduler (OS): Components that decides which process to run next.
  - How to choose?
  - Processes run for *time slice* units of time = granularity.
  - Scheduler *policy* determines what runs when.
- Different kinds of processes:
  - IO-bound – must be responsive, wait most of the time, little computation.
  - Processor-bound – run most of the time, lots of computations.

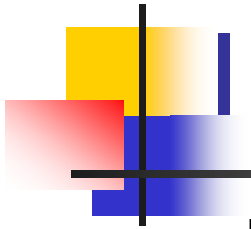

Introduction

- Priority based scheduling: Runnable processes with time slice left and the highest priority always run.
  - Sort them?
  - Which data structure?
  - Which priority? Priority is dynamically modified to fulfill scheduling policies:
    - fast process response or
    - high process throughput.

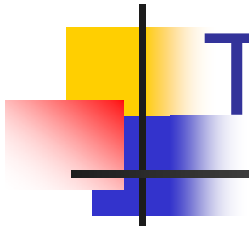

The Scheduler

- List of runnable processes = “runqueue”.
  - 2 priority arrays: one active and one expired, swapped with pointers.
- Priority array:
  - bitmap for priorities,
  - with lists of processes per priority level.
- Finding the highest runnable process = finding the first bit set to 1.
  - Independent on the number of processes.
  - Dependent on the (fixed) number of priority levels.

Priority Bitmap