


A Day In The Life Of A Completely Fair Scheduler [CFS]



By - Ribhu Vajpey
Prat Patel

My Parents Were Ideal Fair Schedulers

We give equal time to all processes and all processes run parallelly!

$$\text{Time for Each Process in Time } T = T * \frac{1}{N}$$

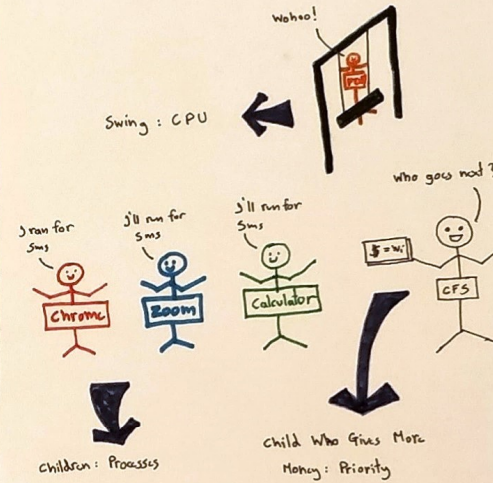
A time quantum Number of processes in Each T

Example: T = 20ms

Process 1 (10ms)	5ms	5ms		
Process 2 (40ms)	5ms	5ms	10ms	20ms
Process 3 (10ms)	5ms	5ms		
Process 4 (20ms)	5ms	5ms	10ms	
	N=4	N=4	N=2	N=1




How Do I Work? Let's Take an Analogy!




How I Choose Who Goes Next?

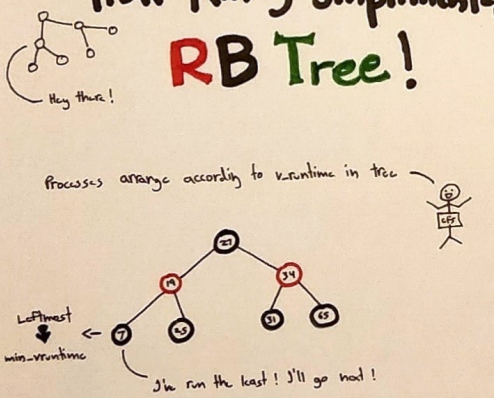
I use Virtual Run Time - A quantity of how long a process has run



When a new kid comes, I allot him vruntime = min_vruntime. New kids get to play first!



How Am I Implimented? RB Tree!



RB Tree is updated at every context switch / timer tick

- New Processes Added
- Finish/Blocked Processes Removal

Why RB Tree?

- Self Balancing Binary Search Tree
- Insert, Delete, Update: $O(\log N)$
- Finding min_vruntime: $O(1)$

Can I Give Priority To Some Processes?


Simply change the update strategies

$$\text{Time Given to Each process } i = T * \frac{W_i}{\sum W_i}$$

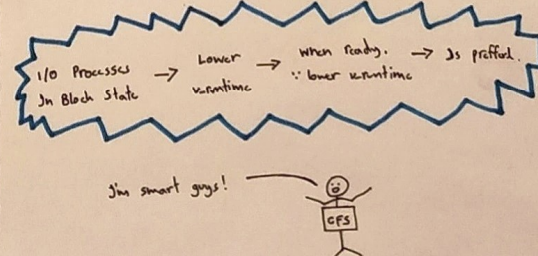
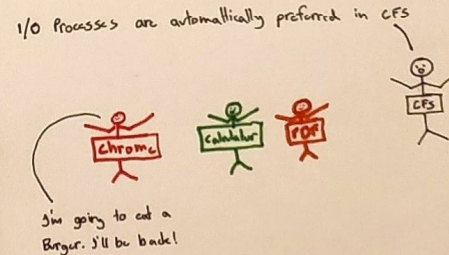
Weight of Process 'i' Summation of all W_i

$$\text{vruntime}_i += t_i * \frac{1}{W_i}$$

$\uparrow W_i \rightarrow \uparrow$ Priority to Process in Each T



What About I/O Processes?



Are You a Nerd?



Want to know more? Check out these:

- <https://opensource.com/article/19/2/fair-scheduling-linux>
- <https://kernel.org/doc/Documentation/scheduler/sched-design-cfs-td>
- <https://developer.ibm.com/technologies/linux/tutorials/>
- <https://cs.columbia.edu/~jimmyj/13fa-w4118/lectures/>

Thanks Folks!