
LITMUS^{RT}: A Status Report*

Created Fall 2019

Before - After - During



Sections

LITMUS^{RT}: A Status Report written in **2007**

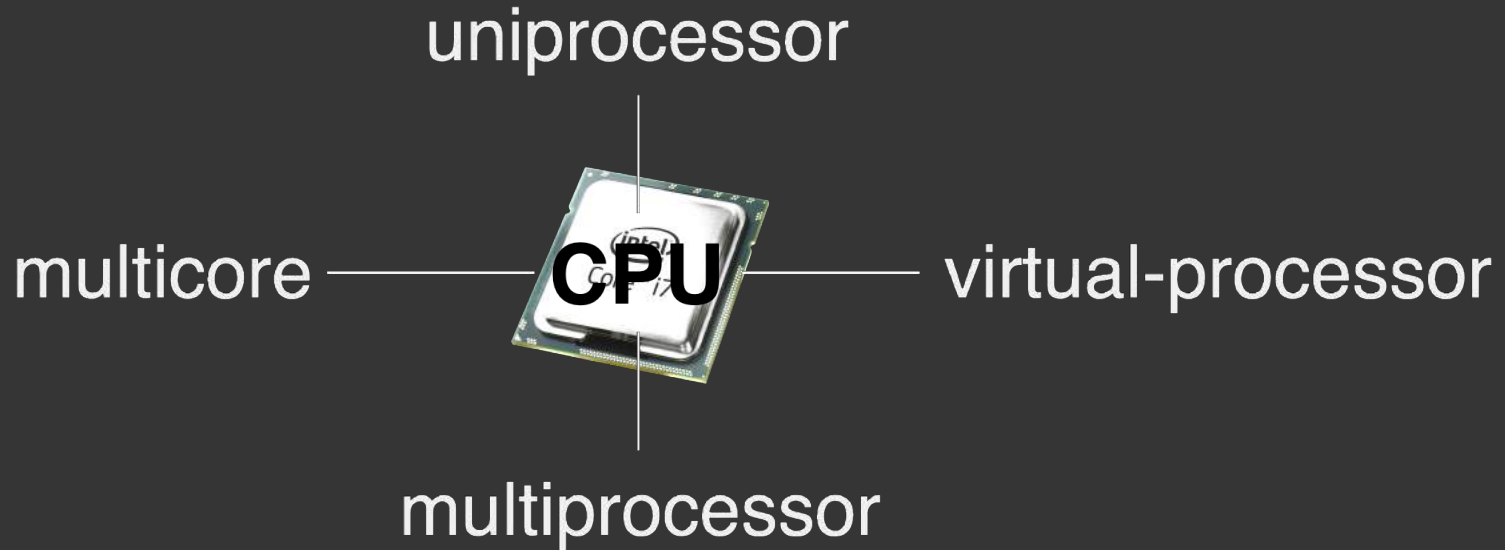
- **Before 2007**
Fundamentals
- **After 2007**
Status of LITMUS^{RT} 2007 - 2019
- **During 2007**
LITMUS^{RT}: A Status Report

SECTION 1 - BEFORE 2007

Background Before LITMUS^{RT}

“Many ... real-time Linux variants under development will be deployable on multicore & multiprocessor platforms” ~ LITMUS^{RT}: A Status Report

Multicore vs Multiprocessor



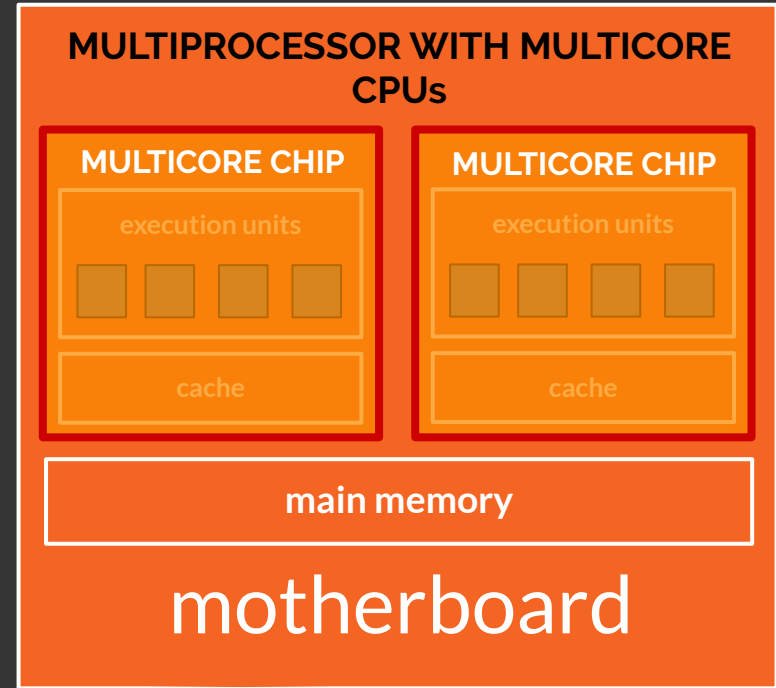
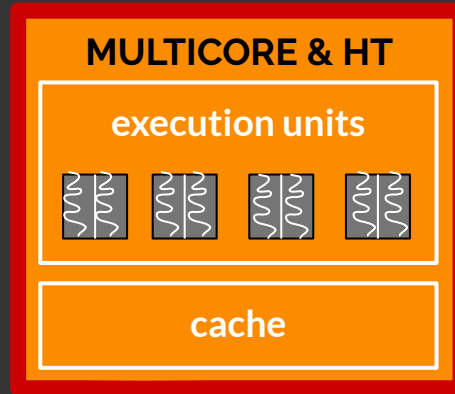
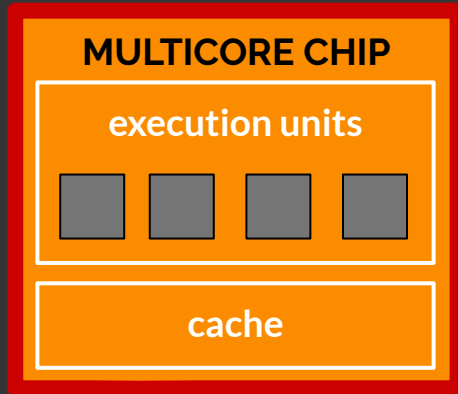
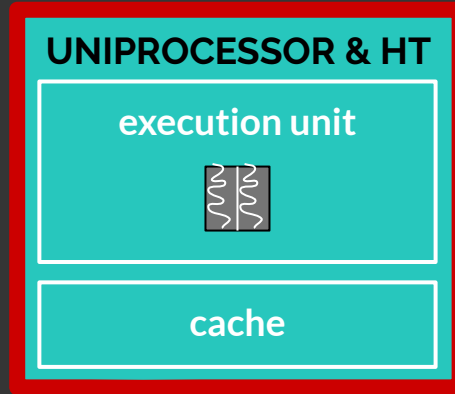
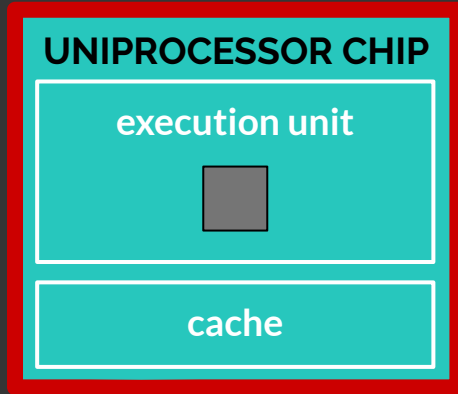
Multicore vs Multiprocessor

- **Uniprocessor**
 - 1 CPU with 1 execution-unit
- **Multicore**
 - 1 CPU with >1 execution-units sharing cache
- **Virtual/Logical Processor** (Hyper-Threading)
 - turns each physical execution-unit into 2 virtual units
- **Multiprocessor**
 - multiple CPUs sharing main memory

— SECTION - BEFORE 2007



HT = with Hyper-Threading

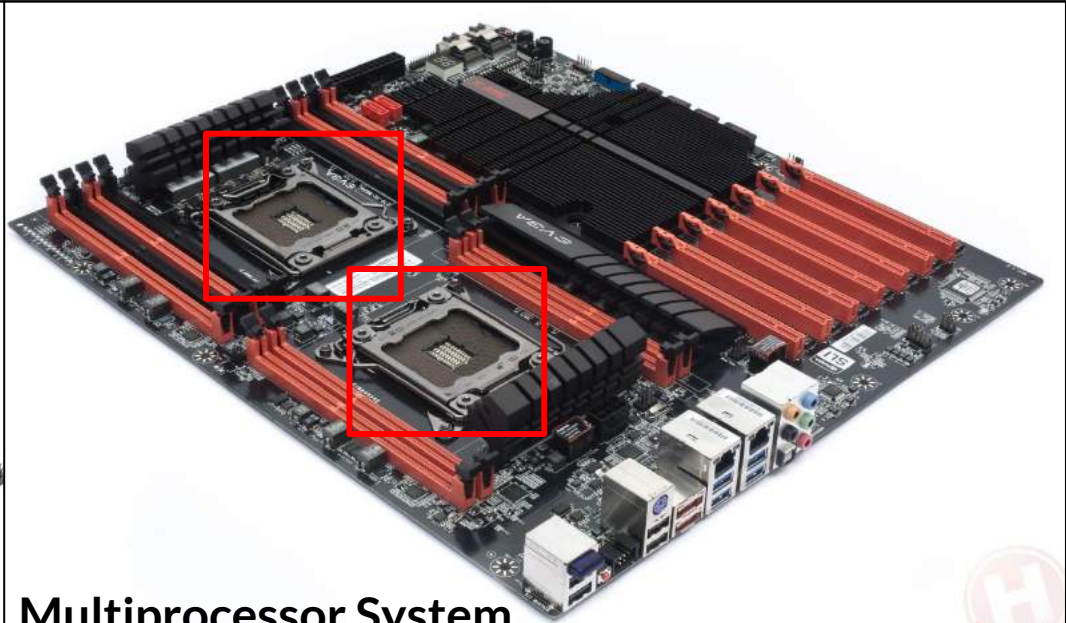


— SECTION - BEFORE 2007

Single vs Multi Socket



Multicore System



Multiprocessor System

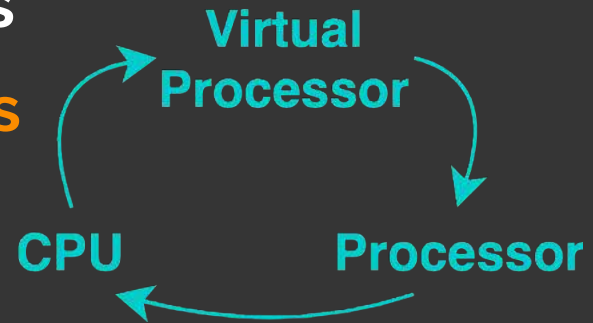


Terms Are Interchangeable

Sometimes:

- **Multicores** are called **Multiprocessors**
- **Virtual Processors** are just called **Processors**
- **Processors** are called **CPUs**
- **CPUs** are called **Virtual/Logical CPUs**

In Real-Time Systems, we use **processors** to denote **virtual processors**



SECTION - BEFORE 2007

PUTTING IT ALL TOGETHER

linux command:

`lscpu`

Socket(s) \approx # of CPU

Core(s) \approx # of cores per chip

Thread(s) \approx Hyper-Threading

CPU(s) \approx # of Virtual Units

LITMUS^{RT} command:

`rtspin -v 1 1000 1 -p 0`

```
terminal — fish /root — ssh 192.168.86.32 — 69x38
root@ubuntu ~# lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:             Little Endian
CPU(s):                 4
On-line CPU(s) list:   0-3
Thread(s) per core:    2
Core(s) per socket:    2
Socket(s):              1
NUMA node(s):          1
Vendor ID:              AuthenticAMD
CPU family:             21
Model:                 48
Model name:             AMD A10-7800 Radeon R7, 12 Compute Cores 4C+8G
Stepping:              1
CPU MHz:               3500.003
BogoMIPS:              6999.58
Virtualization:        AMD-V
L1d cache:             16K
L1i cache:             96K
L2 cache:              2048K
NUMA node0 CPU(s):     0-3
Flags:                 fpu vme de pse tsc msr pae mce cx8 apic sep mt
rr pge mca cmov pat pse36 clflush mmx fxsr sse sse2 ht syscall nx mmx
ext fxsr_opt pdpe1gb rdtscp lm constant_tsc rep_good nopl nonstop_tsc
extd_apicid aperfmperf eagerfpu pni pclmulqdq monitor ssse3 fma cx16
sse4_1 sse4_2 popcnt aes xsave avx f16c lahf_lm cmp_legacy svm extap
ic cr8_legacy abm sse4a misalignsse 3dnowprefetch osvw ibs xop skinit
wdt lwp fma4 tce nodeid_msr tbm topoext perfctr_core perfctr_nb bpex
t ptsc cpb hw_pstate vmmcall fsgsbase bmi1 xsaveopt arat npt lbrv svm
_lock nrip_save tsc_scale vmcb_clean flushbyasid decodeassists pausef
ilter pfthreshold overflow_recov
root@ubuntu ~# █
```

Process vs Threads

Both are sequences-of-instructions that can be executed within a processor

Difference

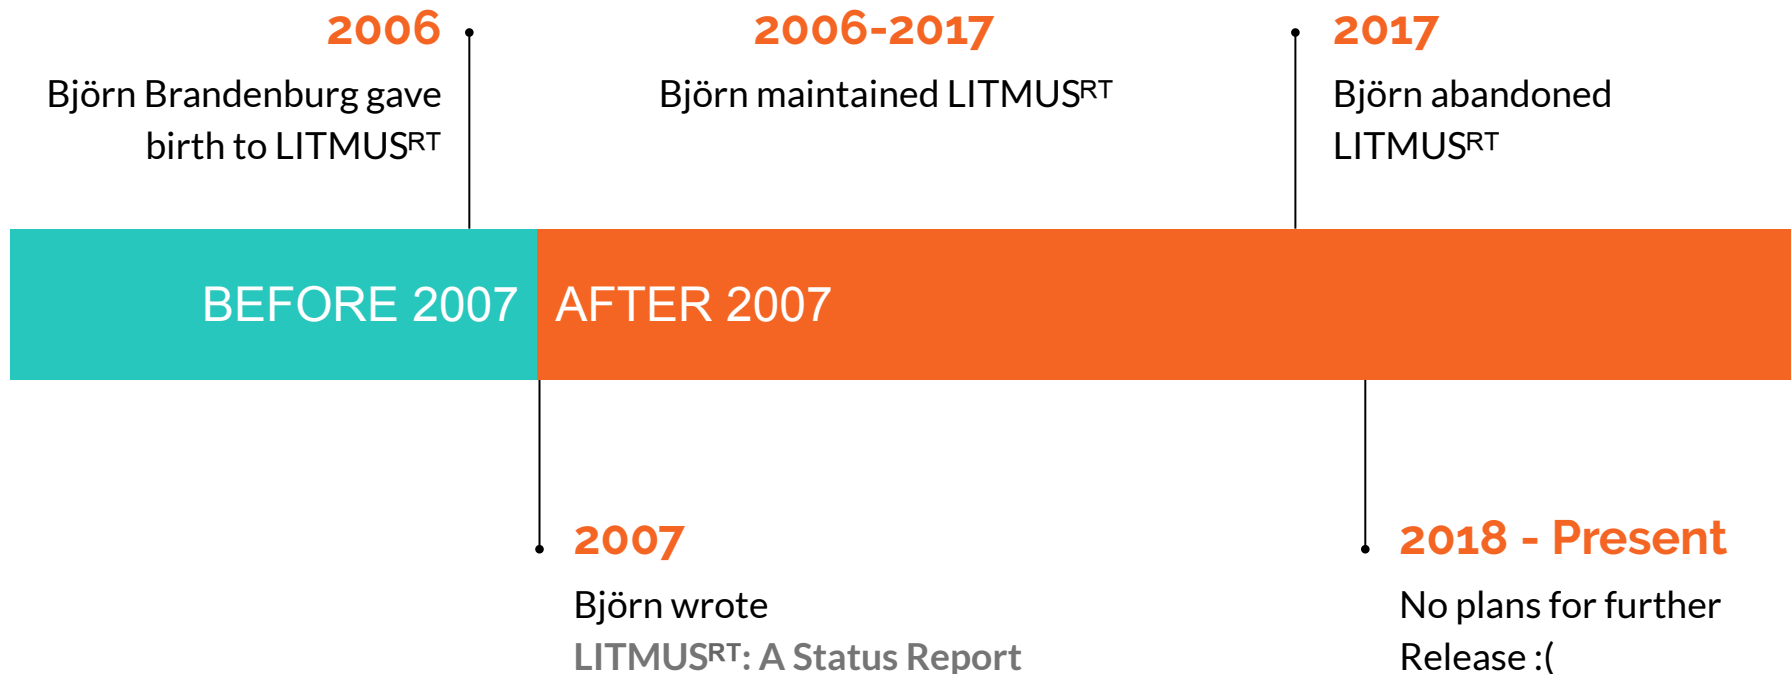
- **Process** - processes do not share memory-stack
- **Threads** - threads share memory

In real-time systems, processes and threads are called jobs

SECTION 2 - AFTER 2007

Status of LITMUS^{RT} 2007 - 2019

LITMUS^{RT}: A Status Report - UPDATED



SECTION 3 - DURING 2007

LITMUS^{RT}: A Status Report - 2007



LITMUS^{RT}: A Status Report

SUB-SECTIONS

- **Basics**
Fundamentals of Real-Time Systems
- **Implementations**
TODO
- **TODO**
TODO
- **TODO**
TODO

Tasks vs Jobs - Scheduler

Tasks ≠ Jobs

Difference

- **Tasks** - generates jobs
- **Jobs** - are scheduled, then executed

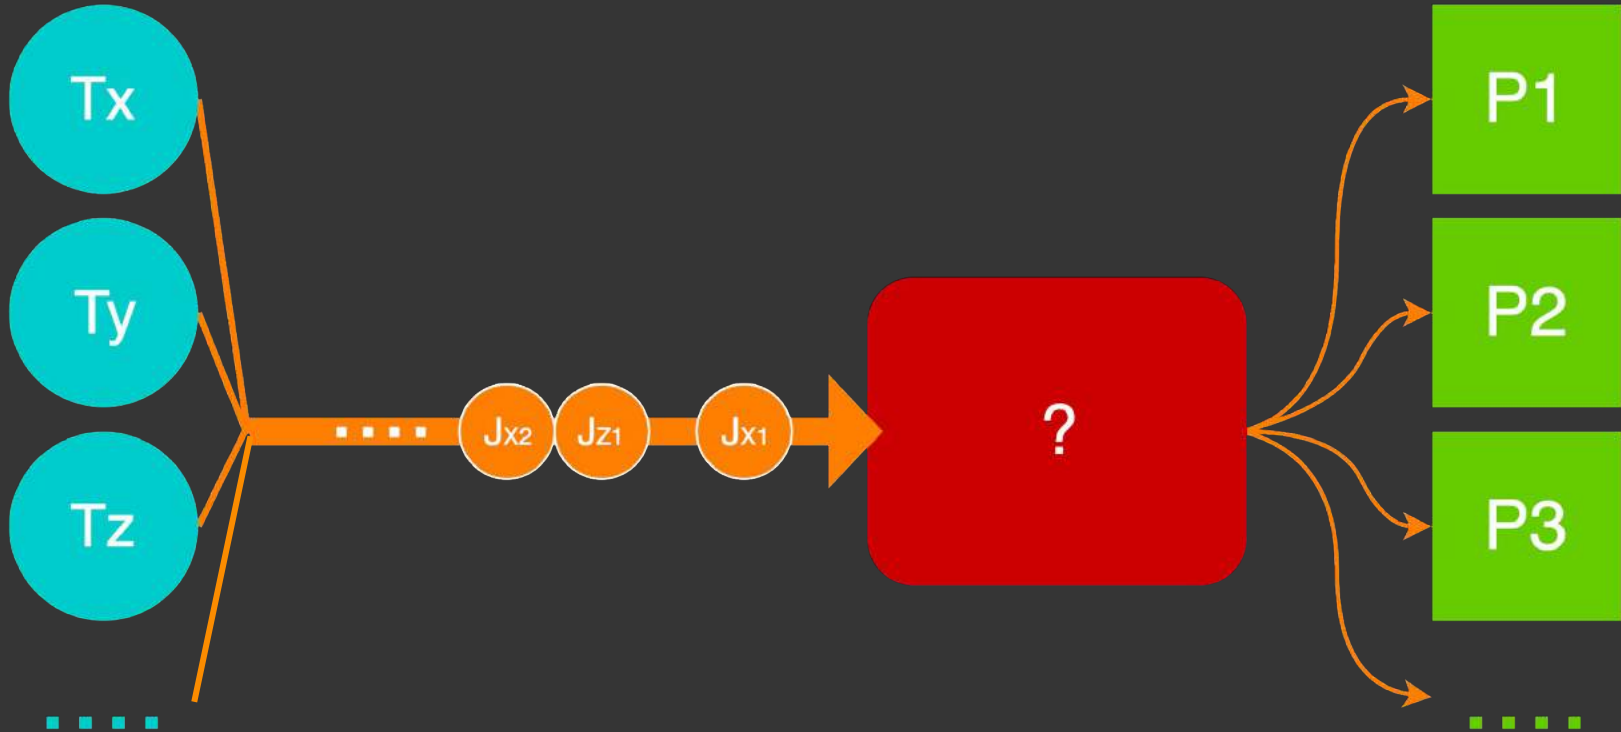
Jobs are processes. In other cases jobs can be threads

Scheduler - schedules jobs to be executed on processor(s)

Tasks

Jobs

Job Scheduler Processor(s)



Scheduler Properties

A **Scheduler** can be described with 3 properties:

1. Number of queues

Global, Partitioned, or Clustered

2. Scheduling Order of each queue

FIFO, EDF, Priority-Based, Round-Robin, etc

3. Whether jobs are PREEMPTIBLE or NON-PREEMPTIBLE

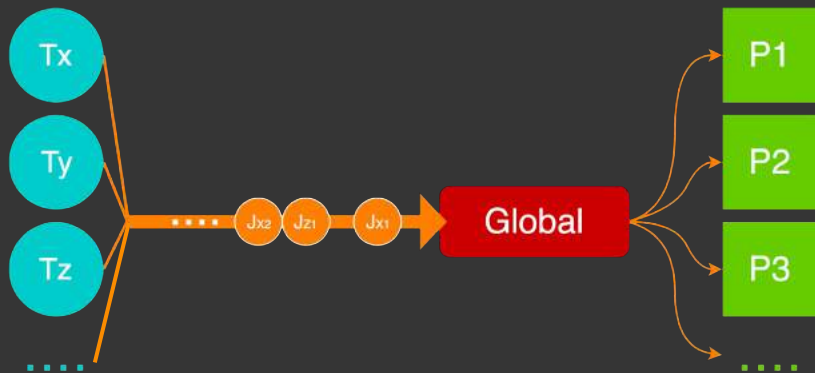
Scheduler - # of Queues

Global - one queue for all processors

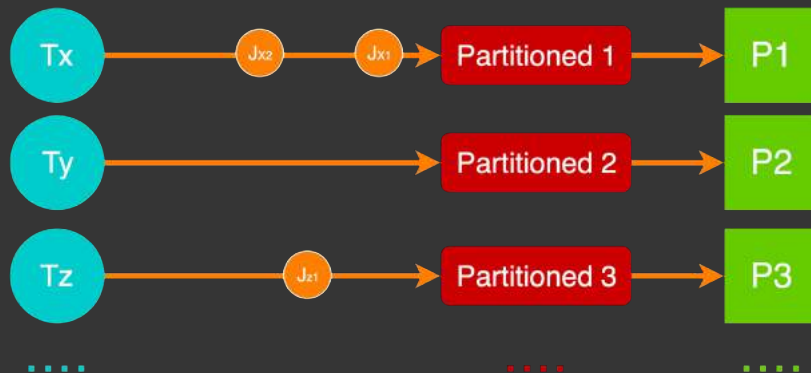
Partitioned - a queue for each processor

Clustered - a queue for a subset of processors

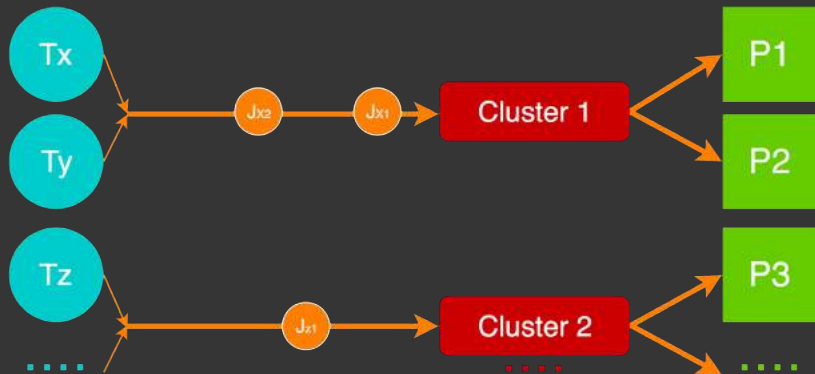
— SECTION - DURING 2007 - BASICS
GLOBAL



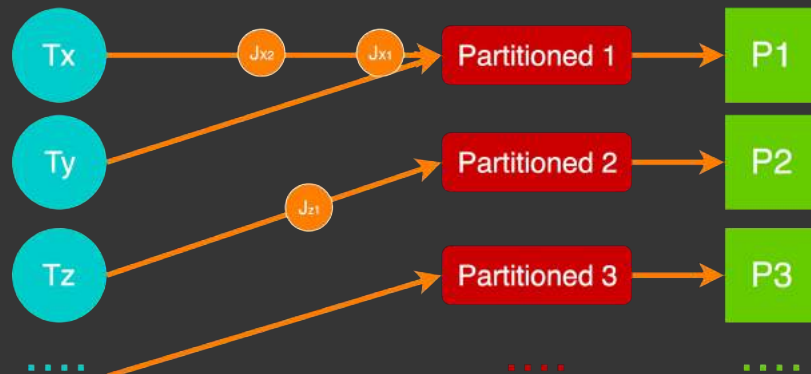
PARTITIONED Version 1



CLUSTERED



PARTITIONED Version 2



Scheduler - Scheduling Orders

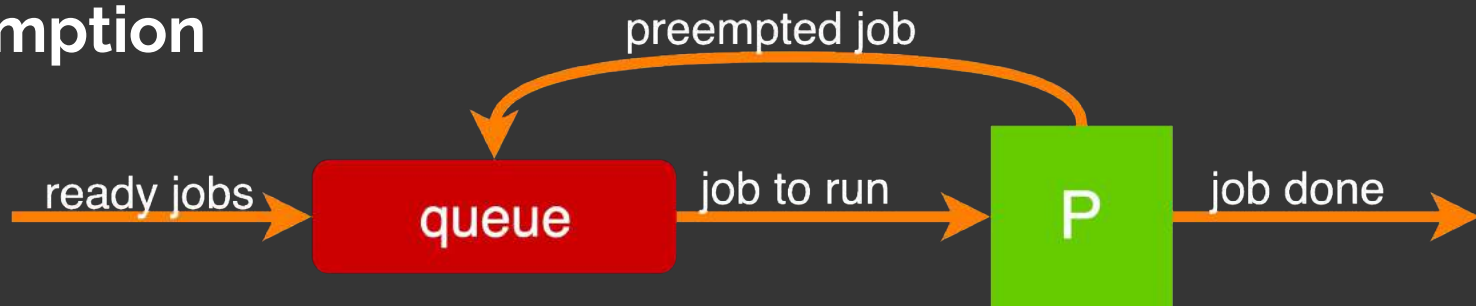
Each **queue** has a scheduling order

Common Scheduling Orders

- **FIFO** - First In First Out
- **EDF** - Earliest Deadline First
- **Priority Based** -
- **Round-Robin** - each job is allocated a time-slice (inherently preemptive)

Scheduler - Preemption & Non-Preemption

Preemption



Non-Preemption



Scheduling Model

Choose value for each scheduling property and you will get a **Scheduling Model**

Some Common Scheduling Models

- Partitioned EDF
- Preemptive Global EDF
- Non-Preemptive Global EDF
- etc

Pfair Scheduling & PD₂

Motivation: We have learned that a Preemptive EDF scheduler for a SINGLE processor is always optimal. But is Preemptive EDF also optimal for a multiprocessor system?

NO

Pfair Scheduling - solves this by dividing each task into equal length subtasks. Subtasks are scheduled via EDF

An implementation of Pfair Algorithm is called **PD₂**

Which Scheduling Model is the Best?

It really depends on the types of tasks at hand

"... for hard-deadline tasks, P-EDF and PD2 are usually preferable, while for soft-deadline tasks, G-EDF and G-NP-EDF are better" ~ LITMUS^{RT}: A Status Report



Thanks!

For more information check out
<http://confluence.marcuschiu.com>