



Using Tracing to Analyze Hard Disk Performance

Housseem Daoud

December 10, 2015

École Polytechnique de Montréal

Laboratoire **DORSAL**

Agenda

Introduction and objectives


Implementation

- Block Layer Structure
- Relevant Tracepoints
- State History Tree

Developed Views

Use Cases – Demo

- Writeback Algorithm
- Lttng I/O Behavior
- Flush requests



Introduction


The performance of block devices has a big impact on overall system performance



Many optimization techniques have been developed

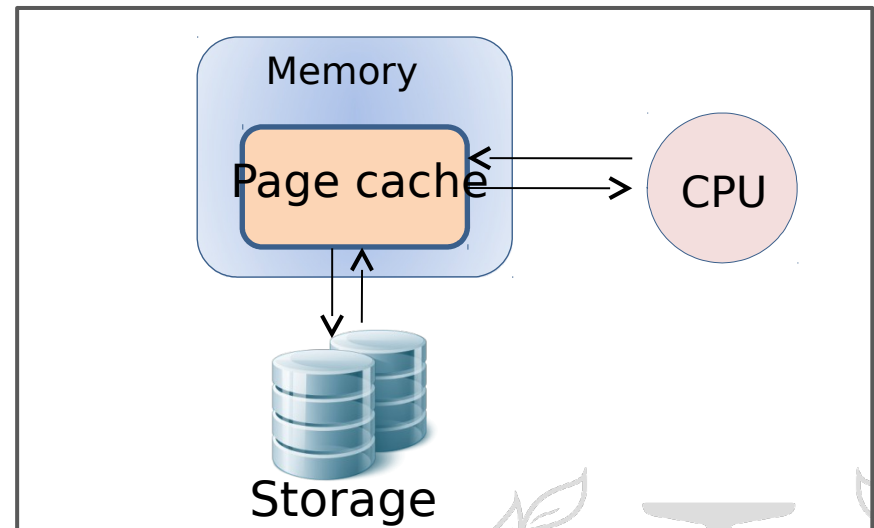
Scheduling

- Request sorting
- Request merging




Caching

- Page cache
- Writeback




Introduction


The Block Layer is very complex



It is difficult to detect disk I/O problems





The need to develop an efficient tool to analyze block layer behavior



Introduction and Objectives

Use tracing to Analyse Hard Disk Performance

- Insert Tracepoints in the block device layer
- Create views to visualize important metrics



Implementation

Block Layer Structure

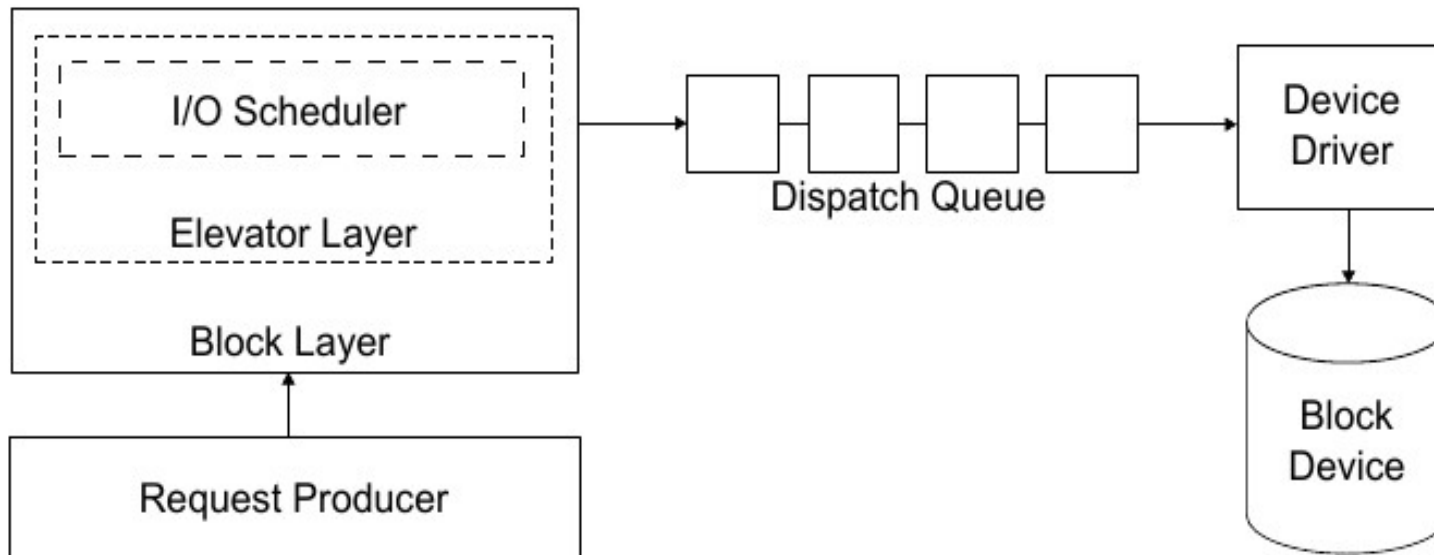




Figure 1: Linux block layer structure



Implementation


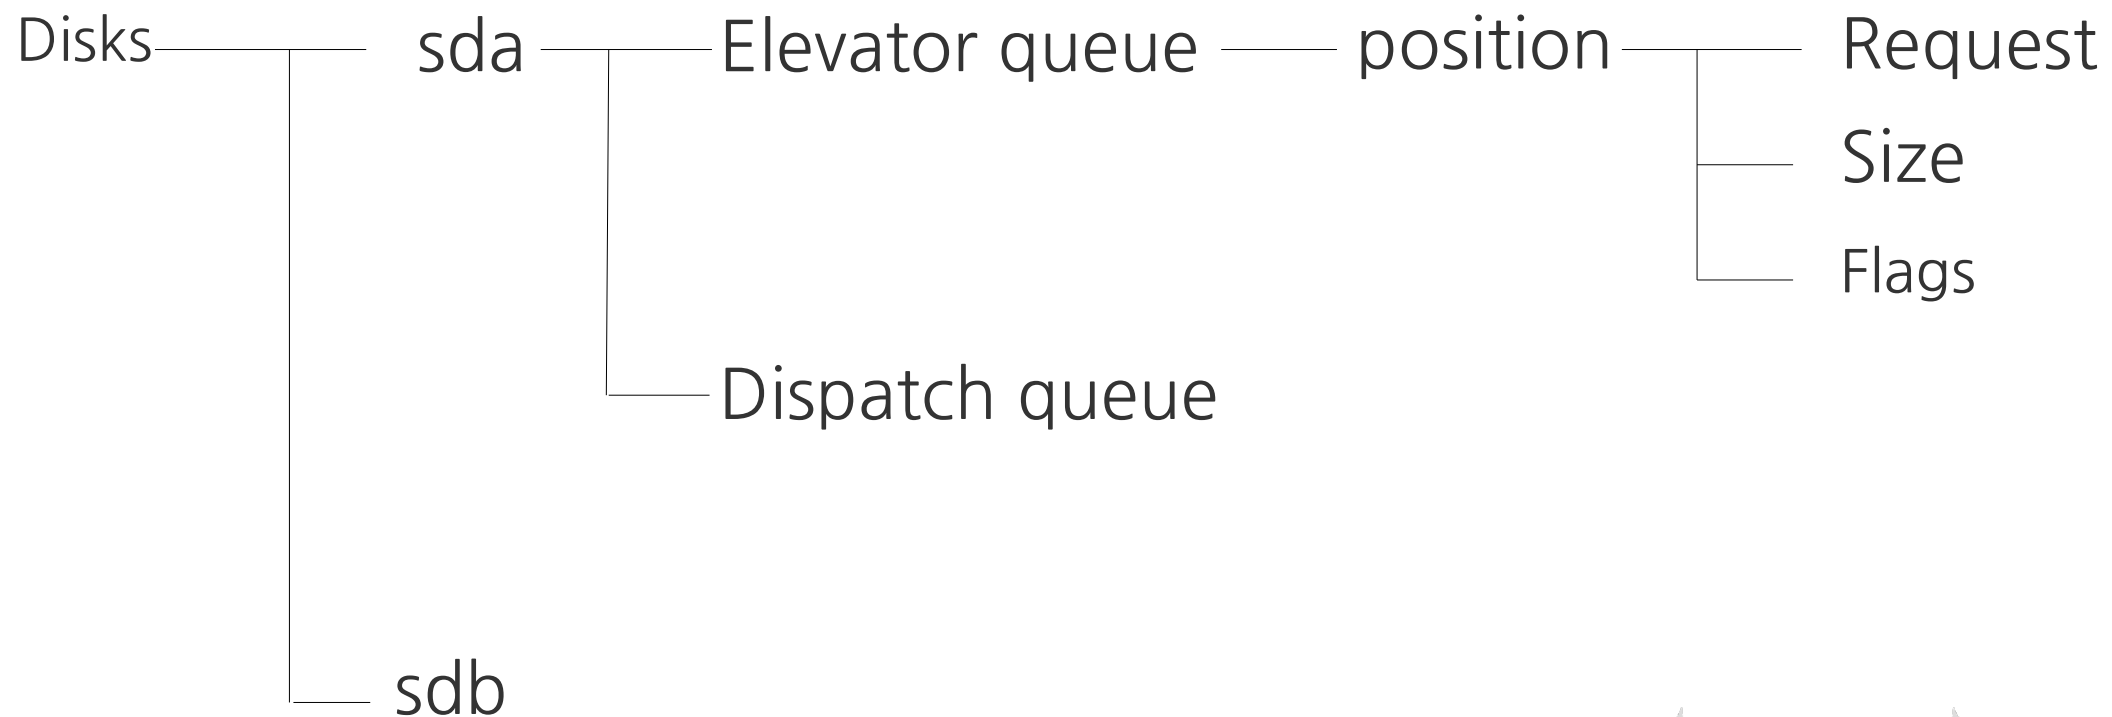
Relevant tracepoints

- `ltnng_statedump_block_device`
- `block_rq_insert`
- **`block_rq_merge`**
- `block_rq_issue`
- `block_rq_complete`





Implementation

State History Tree



Developed Views


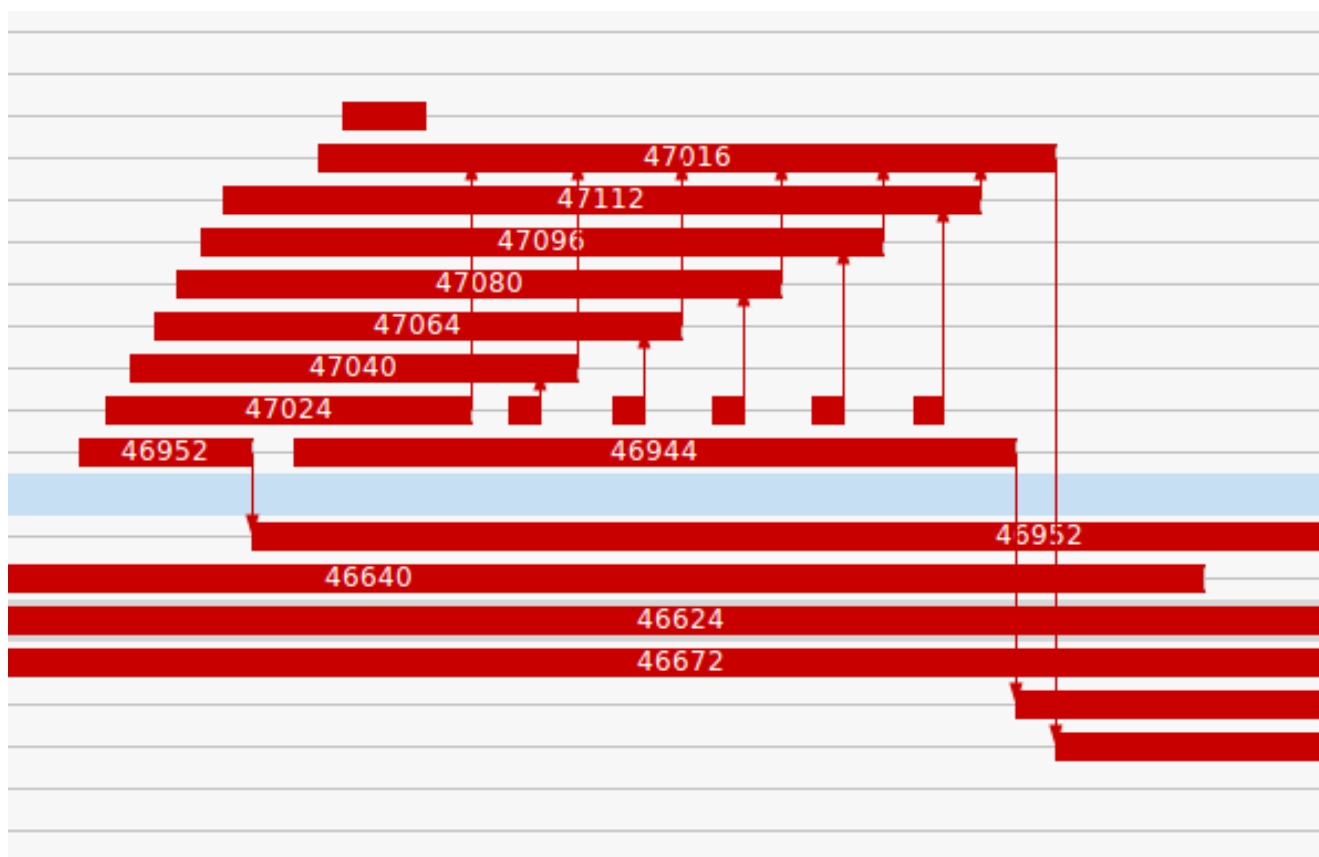
Waiting Queues



Developed Views

Waiting Queues



Merging



Developed Views

Waiting Queues


Queue Length



Developed Views

Request Latencies


Latency Average



Developed Views

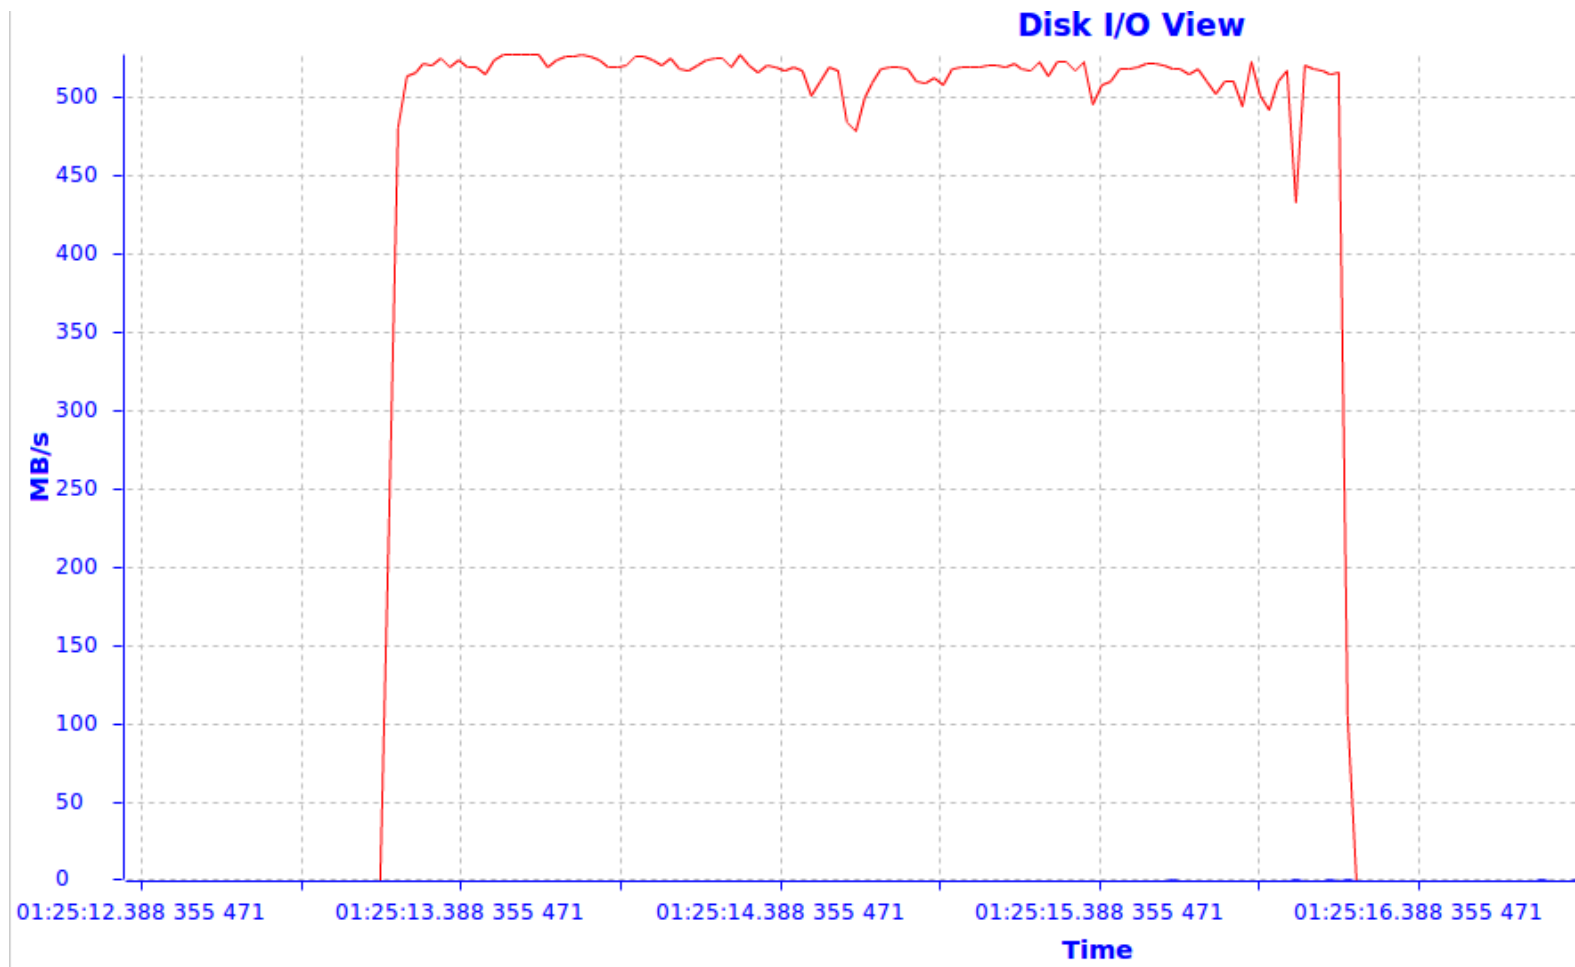
Request Latencies

Latency Distribution



Developed Views



Disk throughput



Developed Views


I/O Activity of Processes

PID	Process	Read (MB)	Write (MB)
1724	Xorg	0.0	0.051
2595	cinnamon	0.0	0.004
8315	lttng	0.0	0.0
6271	/usr/bin/termin	0.0	0.022
2612	nemo	0.0	0.0
1474	irqbalance	0.0	0.0
8212	/usr/bin/termin	0.0	0.0
2029	collectl	0.0	0.0
3194	chrome	0.0	0.0
3014	chrome	0.0	0.007
6761	lttng-consumerd	0.0	0.206
8247	lttng-simple	0.0	0.0
8316	hdparm.sh	0.0	0.0
8317	sudo	0.0	0.0
943	rsyslogd	0.0	0.0
2820	accounts-daemon	0.0	0.0
8318	hdparm	1548.0	0.0
5328	java	0.0	0.0
2446	rtkit-daemon	0.0	0.0
3206	chrome	0.0	0.0
2615	thunderbird	0.0	0.0
3619	chrome	0.0	0.0
3149	chrome	0.0	0.0
4197	okular	0.0	0.0
8153	gnome	0.0	0.001



DEMO

- Flush requests
- Writeback Algorithm
- Lttng I/O Behavior



Thank You !

