

Improving Application Responsiveness with the BFQ Disk I/O Scheduler

Paolo Valente

Mauro Andreolini

Dipartimento di Ingegneria dell'Informazione
University of Modena and Reggio Emilia
Italy

- Production-quality disk scheduler for Linux
 - Each application can be reserved the desired fraction of the disk throughput
 - Guaranteed even if the overall throughput fluctuates
 - High disk throughput
 - Low latency for applications performing little and sporadic I/O, such as multimedia ones
- Used on PCs and smartphones
http://algogroup.unimore.it/people/paolo/disk_sched/

Contributions

- A set of heuristics added to BFQ
 - The resulting new version of BFQ is labeled *v1*
- A benchmark suite
http://algogroup.unimore.it/people/paolo/disk_sched/
- Experimental results on three different systems with a single rotational disk

The new heuristics

- The heuristics improve:
 - Responsiveness
 - Application start-up time
 - Time to accomplish the jobs of interactive applications
 - Throughput
 - Robustness across heterogeneous devices

Contents of this presentation

- Focus on
 - Responsiveness
 - Throughput
- Live demo of the performance of the new version of BFQ (recorded version [here](#))
 - Compared against CFQ
 - One of the most effective production-quality schedulers in terms of throughput boosting and latency
- The demo also briefly shows the performance of BFQ in terms of latency for soft real-time applications and of disk throughput

Downside

- Newly-created and interactive applications enjoy a low latency because
 - BFQ provides them with more than their fair share of the disk throughput
- This extra share is necessarily *stolen* to non-interactive, long-lived applications
 - Applications most sensitive to this problem:
 - Soft real-time applications, such as multimedia

Latency for soft real-time

- **Key point:**
 - The heuristics fit the original accurate service provided by BFQ well enough to still guarantee to soft real-time applications a latency comparable to CFQ
 - Even better, from BFQ-v2 on, a new simple heuristic guarantees to soft real-time applications a latency three times as low as CFQ
 - Without breaking high-responsiveness guarantees

Throughput

- In addition to low latency for both interactive and soft real-time applications
- **High disk throughput**
 - Up to 30% higher throughput than CFQ under most workloads in our experiments
 - The same throughput as CFQ under the other workloads
- Ongoing work
 - High throughput with virtual machines (KVM)
 - High throughput and service guarantees with SSD and RAID