FLIN: Enabling Fairness and Enhancing Performance in Modern NVMe Solid State Drives

Arash Tavakkol, Mohammad Sadrosadati, Saugata Ghose,
Jeremie S. Kim, Yixin Luo, Yaohua Wang, Nika Mansouri Ghiasi,
Lois Orosa, Juan Gómez-Luna, Onur Mutlu

Motivation

- SSDs are **widely used** as a storage medium
- SSDs initially adopted **conventional** host interface protocols (e.g., SATA)
  - Designed for magnetic hard disk drives: only **thousands of IOPS** per device

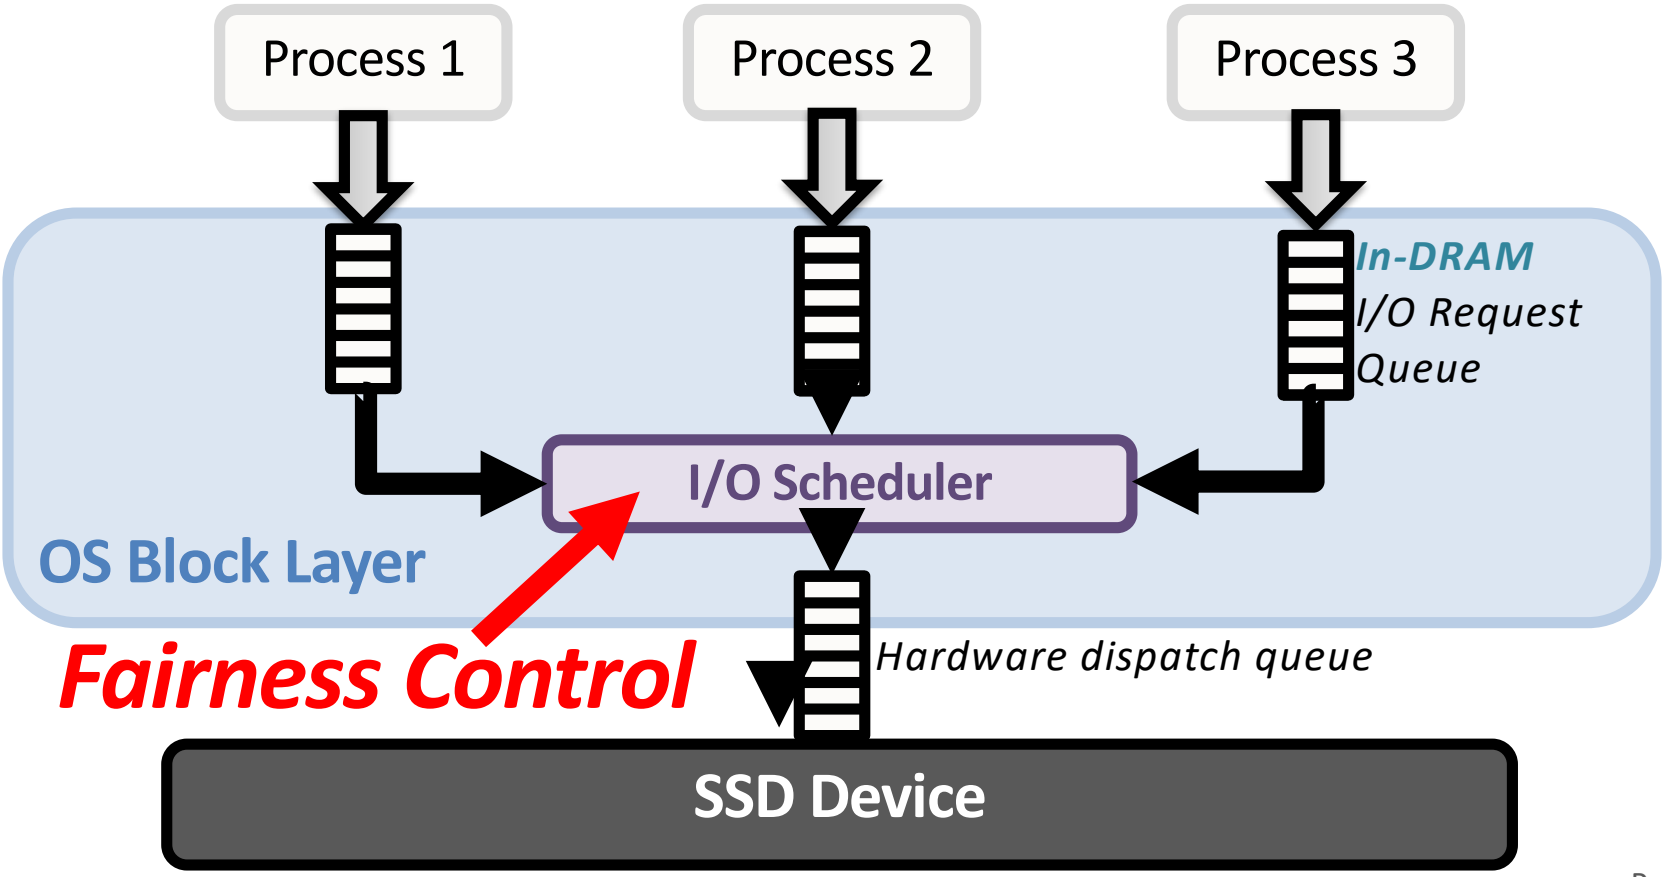

Motivation

- **Modern SSDs use high-performance host interface protocols (e.g., NVMe)**
  - Takes advantage of SSD throughput: enables **millions of IOPS** per device
  - Bypasses OS intervention: **SSD must perform scheduling, ensure fairness**

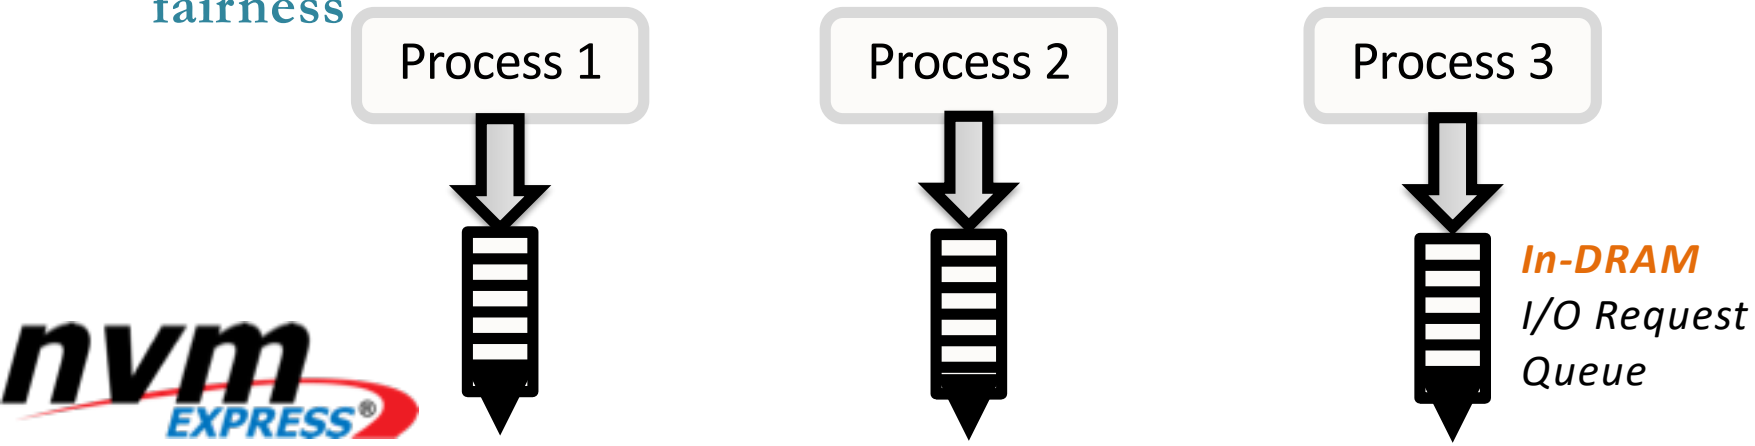

Fairness should be provided by the SSD itself. Do modern SSDs provide fairness?

Motivation

SAFARI

- **We study fairness control in real state-of-the-art SSDs**
  - An example of two datacenter workloads running concurrently

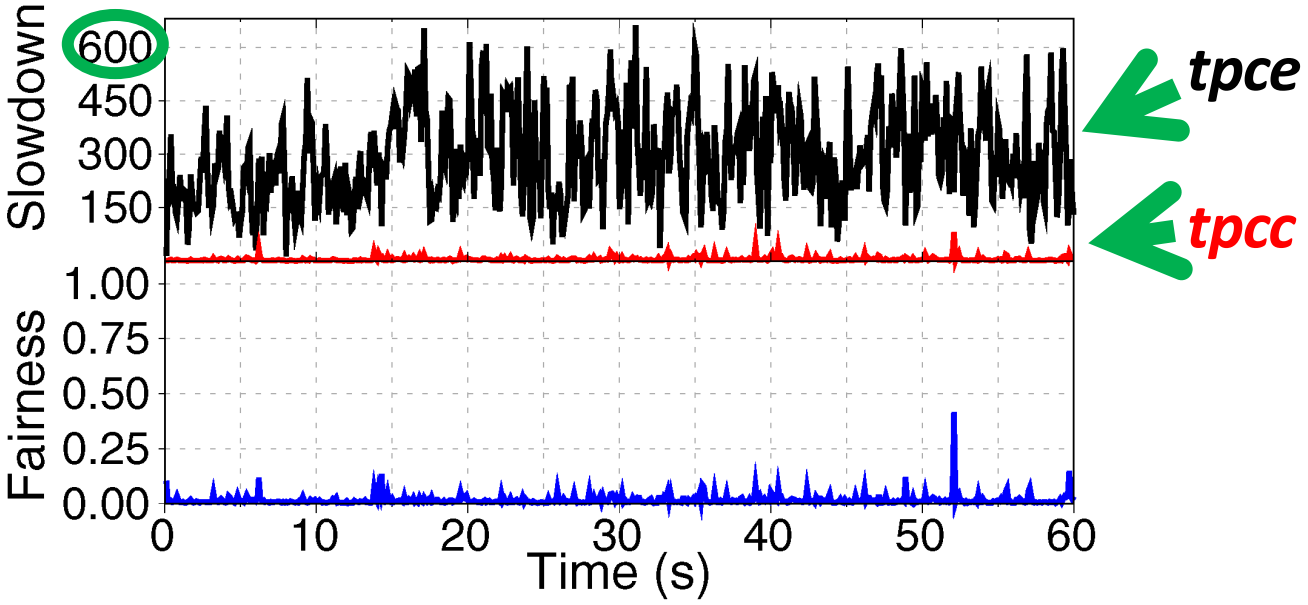

Modern NVMe SSDs focus on providing high performance at the expense of large amounts of unfairness

Our First Contribution

- **We perform a comprehensive analysis of inter-application interference in state-of-the-art SSDs**
  1. The intensity of requests sent by each application
  2. Differences in request access patterns
  3. The ratio of reads to writes
  4. Garbage collection

Our Second Contribution

- We propose the Flash-Level INterference-aware scheduler (FLIN)
- FLIN is a lightweight device-level I/O request scheduling mechanism that provides fairness among requests from different applications
- FLIN carefully reorders transactions within the SSD controller to balance the slowdowns incurred by concurrent applications

Our Third Contribution

- We comprehensively evaluate FLIN using a wide variety of enterprise and datacenter storage workloads
- On average, 70% fairness and 47% performance improvement over a state-of-the-art device-level I/O request scheduler
- FLIN is implemented fully within the SSD firmware with a very modest DRAM overhead (< 0.06%)

FLIN: Enabling Fairness and Enhancing Performance in Modern NVMe Solid State Drives

Arash Tavakkol, Mohammad Sadrosadati, Saugata Ghose,
Jeremie S. Kim, Yixin Luo, Yaohua Wang, Nika Mansouri Ghiasi,
Lois Orosa, Juan Gómez-Luna, Onur Mutlu