Performance Characterization of Modern Storage Stacks

Zebin Ren and Animesh Trivedi

VU Amsterdam

z.ren@vu.nl
a.trivedi@vu.nl

GitHub

@Large Research
Massivizing Computer Systems

VU VRIJE UNIVERSITEIT AMSTERDAM

Paper: https://atlarge-research.com/pdfs/2023-cheops-iostack.pdf

Source code: https://github.com/atlarge- research/Performance-Characterization-Storage-Stacks

The Development of Storage Devices

New Devices

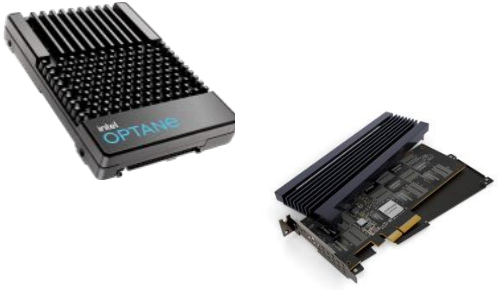

Less than **1k** I/O per Second
Latency: **~5ms**

550-1000K I/O per Second
Latency: **~7us**

More than 1000x speed up

New Interfaces

CPU is the Bottleneck

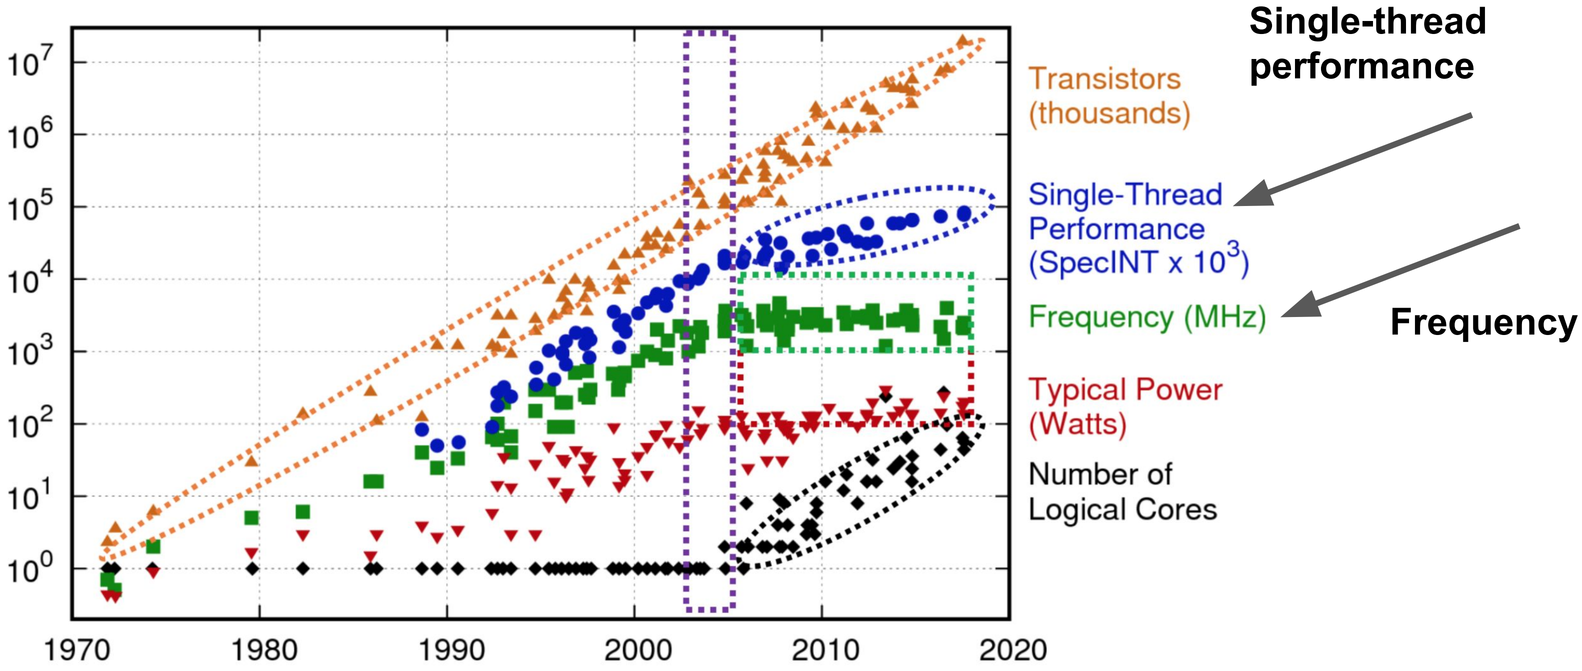

CPU has become the bottleneck !

https://www.researchgate.net/figure/42-Years-of-Microprocessor-Trend-Data-6-Orange-Moores-Law-trend-Purpule-Dennard_fig1_336577121

Storage Stack

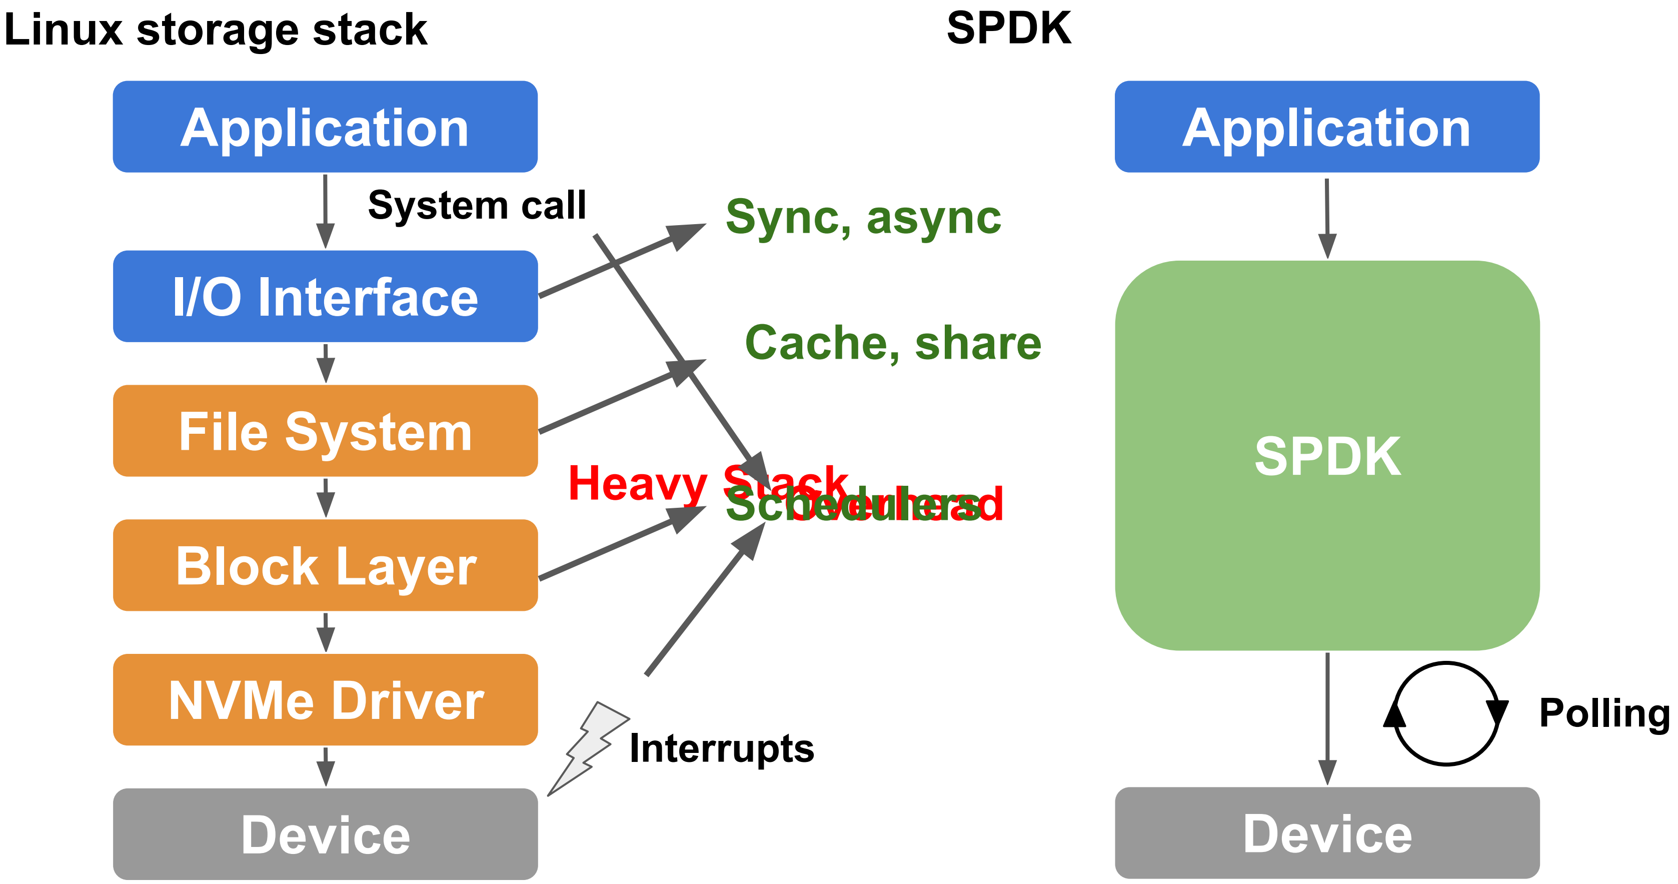

I/O Interfaces

POSIX IO (**psync**)

- Synchronous interface
- Widely used

Asynchronous I/O (**libaio**)

- Asynchronous I/O interface for Linux

io_uring (**iou**)

- A new asynchronous I/O interface
- Designed for performance

io_uring

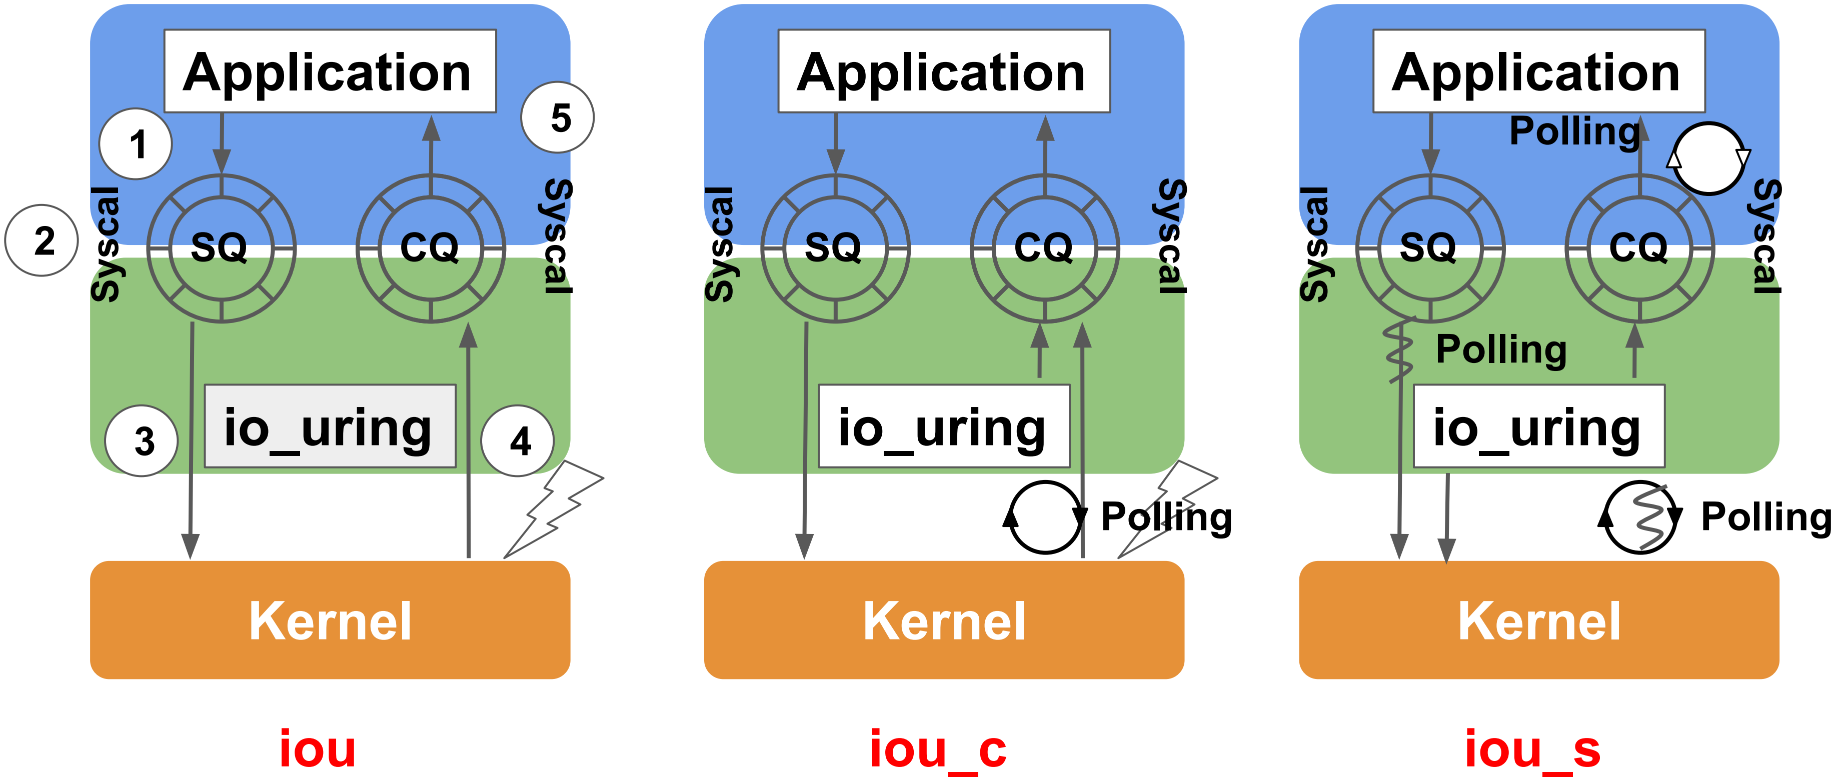

Completion Polling

Submission Polling

Research Problems

Q1: What is the **performance gap** between different **I/O API** and **storage stacks**?

Q2: What is the **cause** of the **performance gap**?

Q3: How does the performance gap **scale** with the number of processes?

Setup

Devices

Intel Optane * 7 → **3.8 Million** IOPS

Workload generator

fio → Widely used + flexible

Workload

4KB random read → to maximize software overhead

Low workload → **1** outstanding request

High workload → **128** outstanding request

What is the performance gap between different I/O APIs and storage stacks?

Performance: Low Workload (Queue Depth = 1)

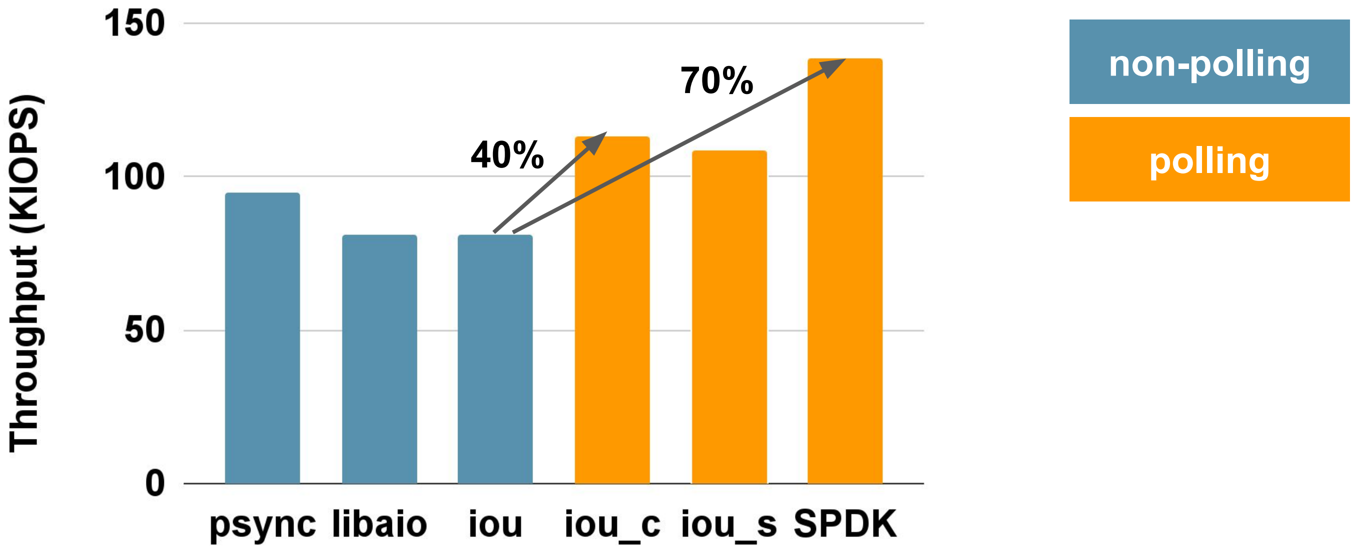

psync has better throughput than libaio and iou

Polling improves the throughput

SPDK has better throughput than the Linux storage stack

Performance: High Workload (Queue Depth = 128)

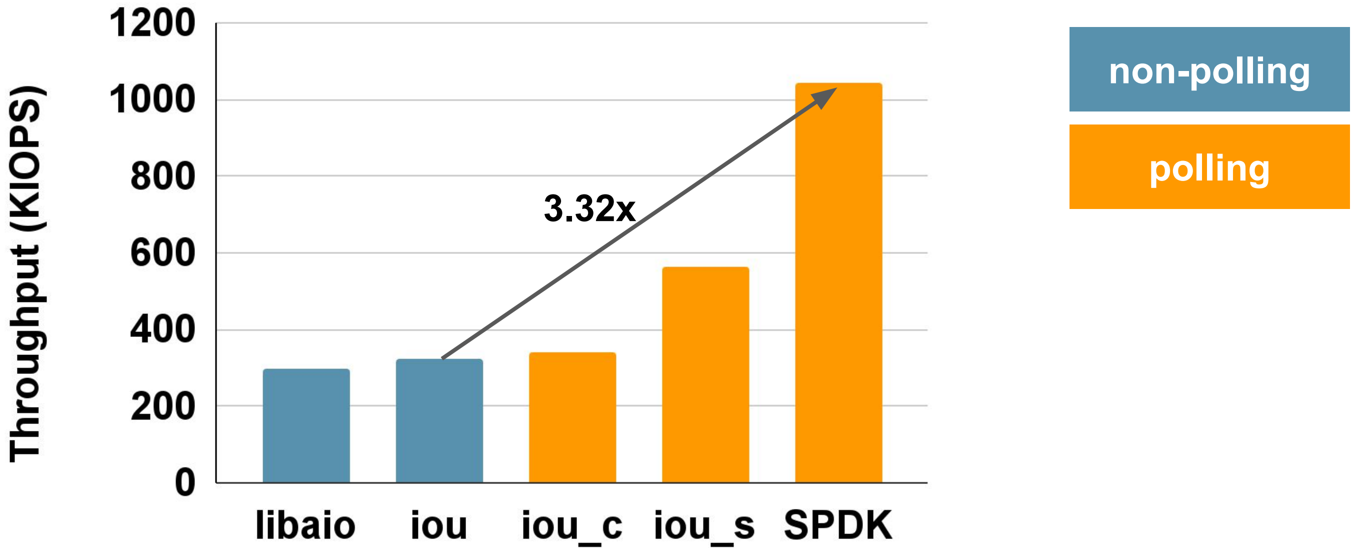

iou is better than libaio

Polling improves throughput, slightly

SPDK has much better throughput than the linux storage stack

Why there is a performance gap?

Number of instructions per I/O

Instructions per cycle (IPC)

Micro-architectural Efficiency: # Instructions per I/O

Low Workload

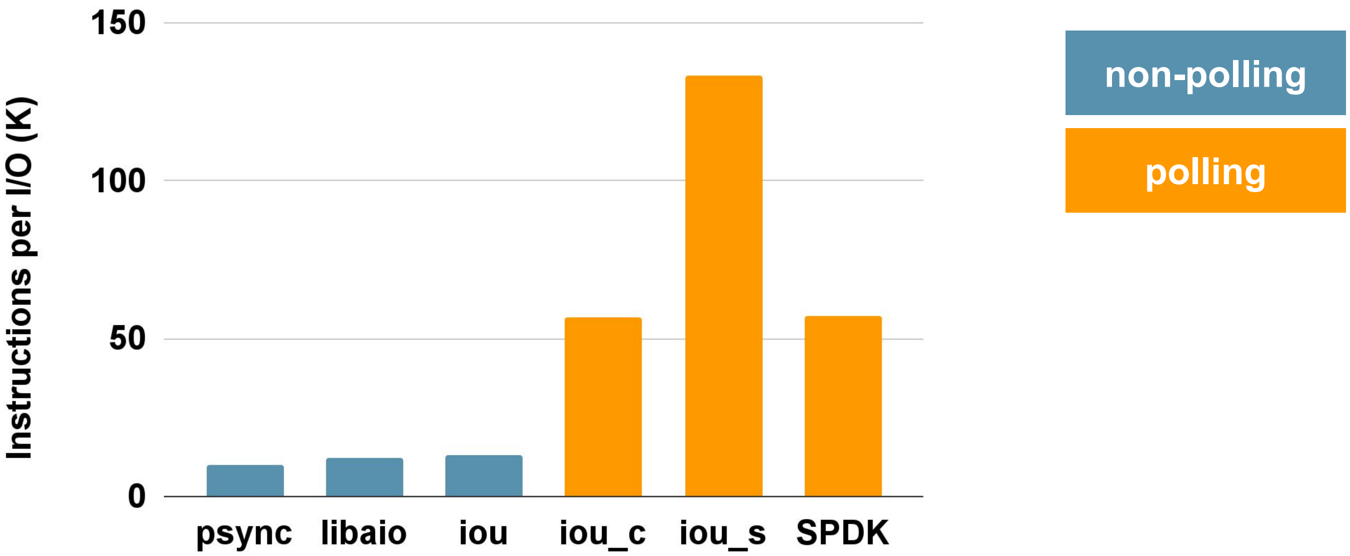

psync is more efficient than libaio and io_uring

Polling wastes instructions at low workload

Micro-architectural Efficiency: # Instructions per I/O
High Workload

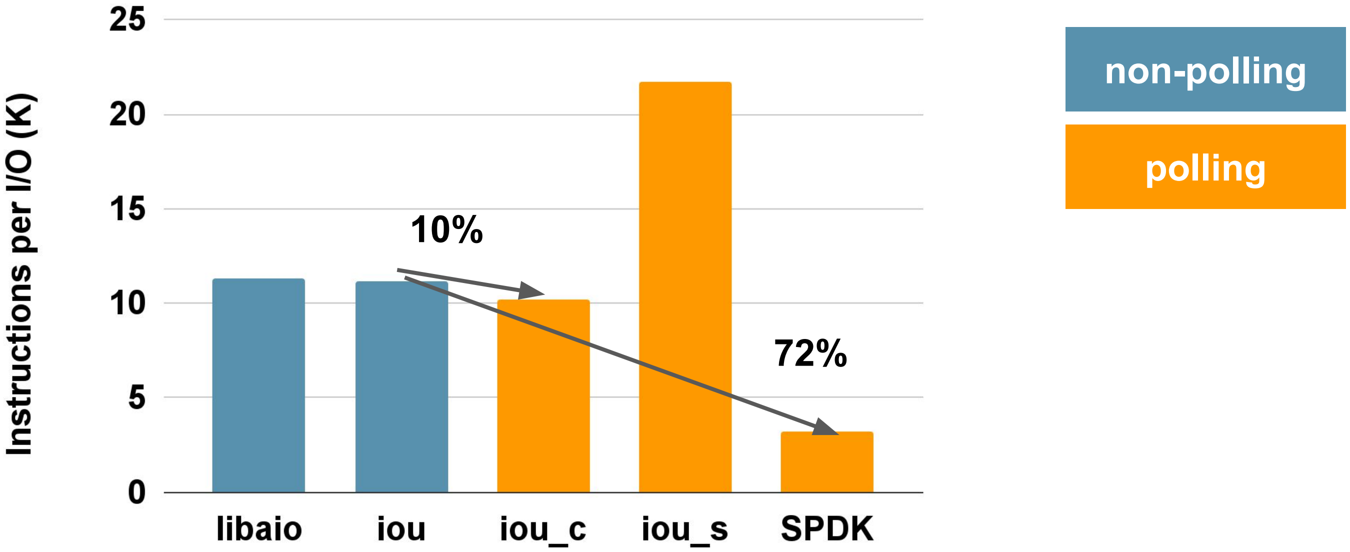

Polling is efficient at high workload

SPDK is much more efficient than the Linux storage stack

Micro-architectural Efficiency: IPC
Low Workload

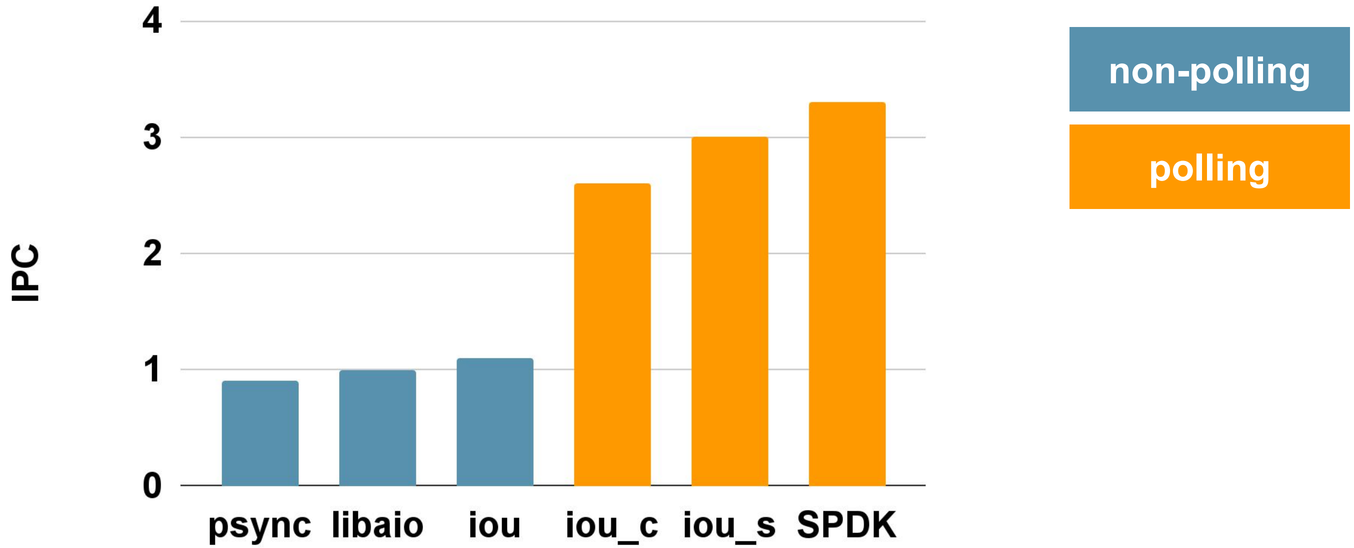

iou has higher IPC

Polling leads to high IPC at low workload

Micro-architectural Efficiency: IPC
High Workload

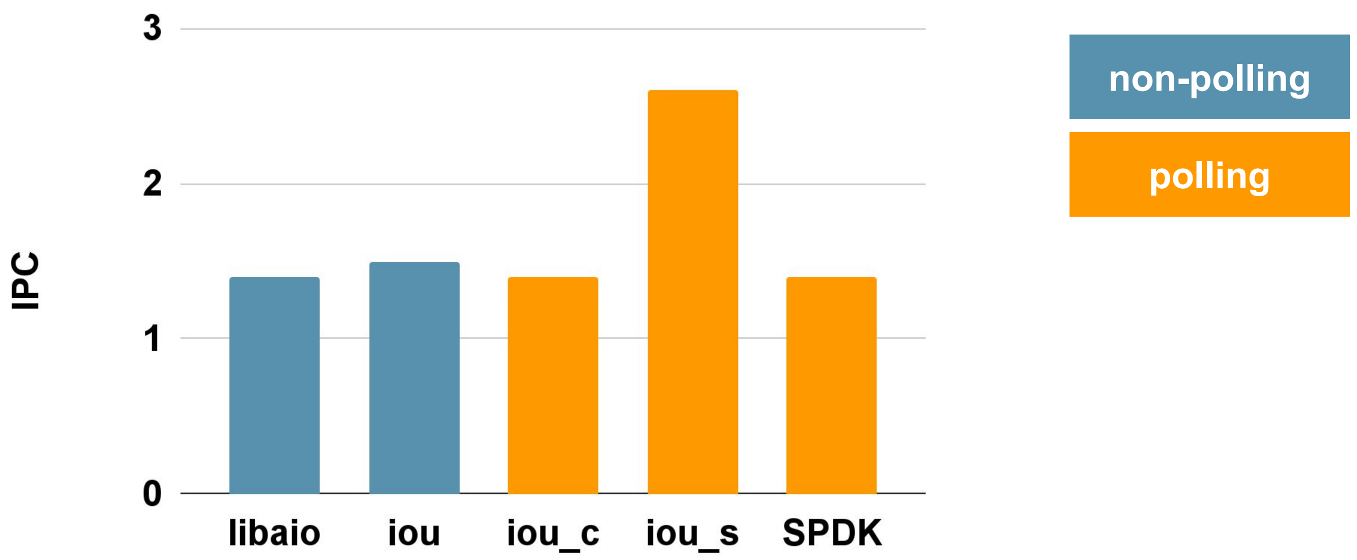

Non-polling delivers to high IPC than low workload

No difference between non-polling and polling APIs

How does the performance gap scale with the number of processes?

Scalability of performance

Impact of I/O schedulers

Scalability

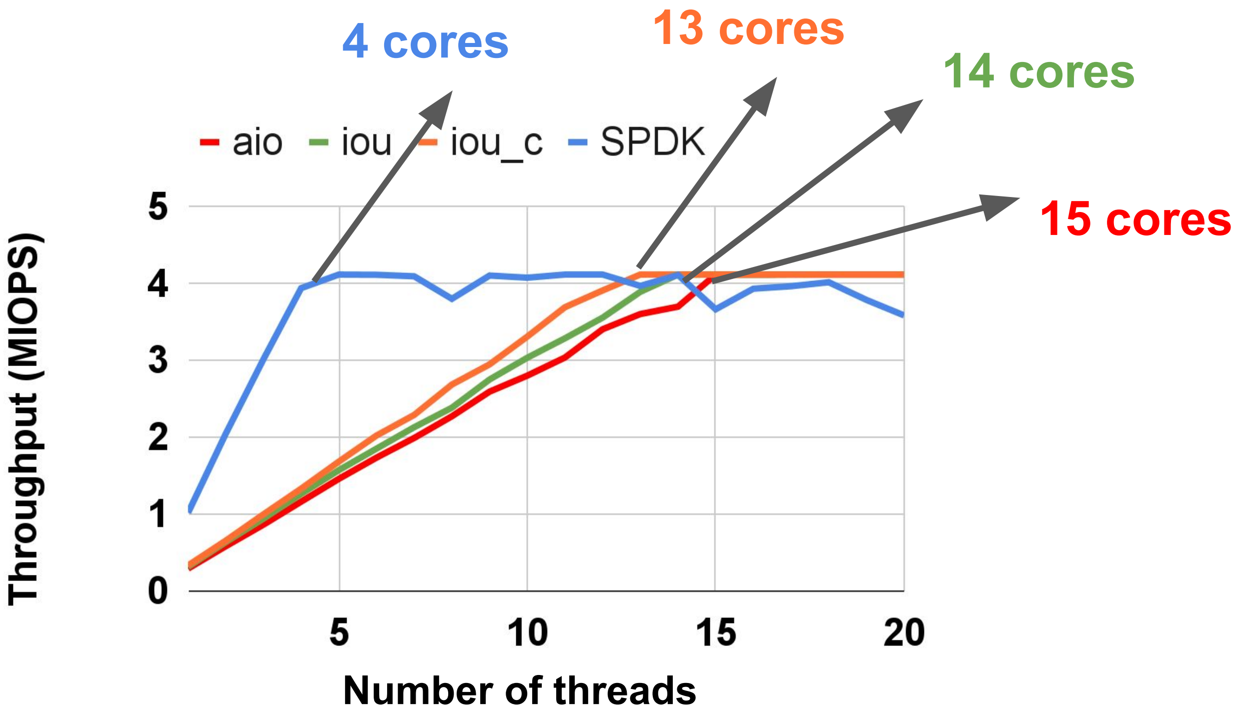

Performance scales linearly for the Linux kernel I/O APIs

io_uring has better performance

SPDK has much higher efficiency than the Linux storage stack

I/O Schedulers

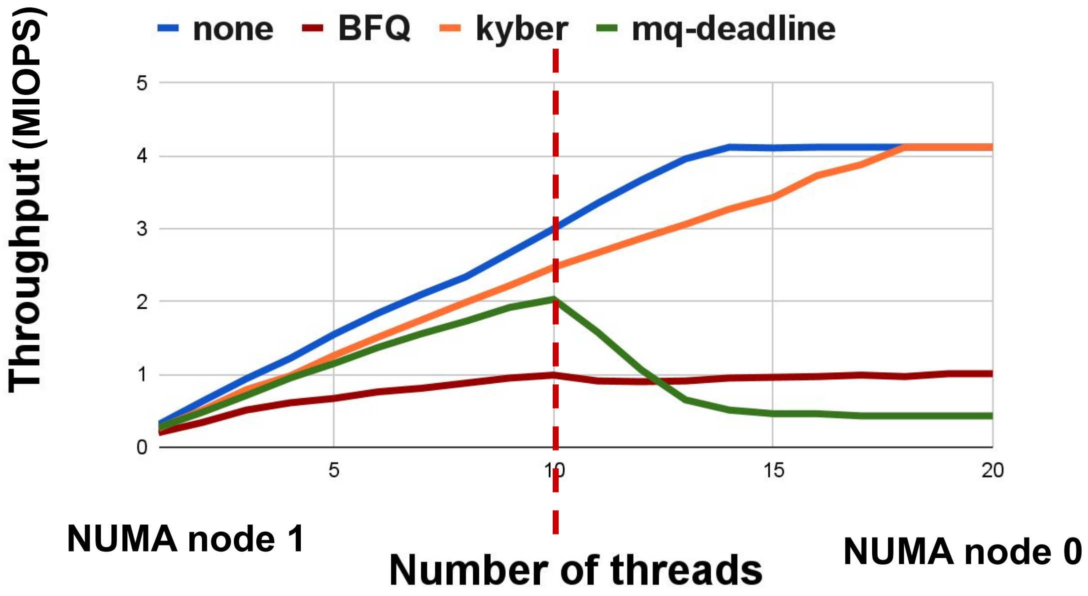

All the I/O schedulers has overhead than the none scheduler

kyber can saturate all the devices with enough CPU resource

mq-deadline and BFQ has bad performance for cross-NUMA access

Take-Home Messages

1. ***Use polling, but carefully***
   Polling wastes CPU time at low I/O workload

2. ***Big gap between Linux storage stack and SPDK***
   SPDK is lightweight and can deliver higher throughput when CPU is the bottleneck

3. The problem of ***Linux I/O stack*** is ***inefficiency***
   Reduce software overhead, scalability of I/O schedulers

GitHub

Source code: https://github.com/atlarge-research/Performance-Characterization-Storage-Stacks

Thank you!
Questions?

Backup Slides: Work Breakdown

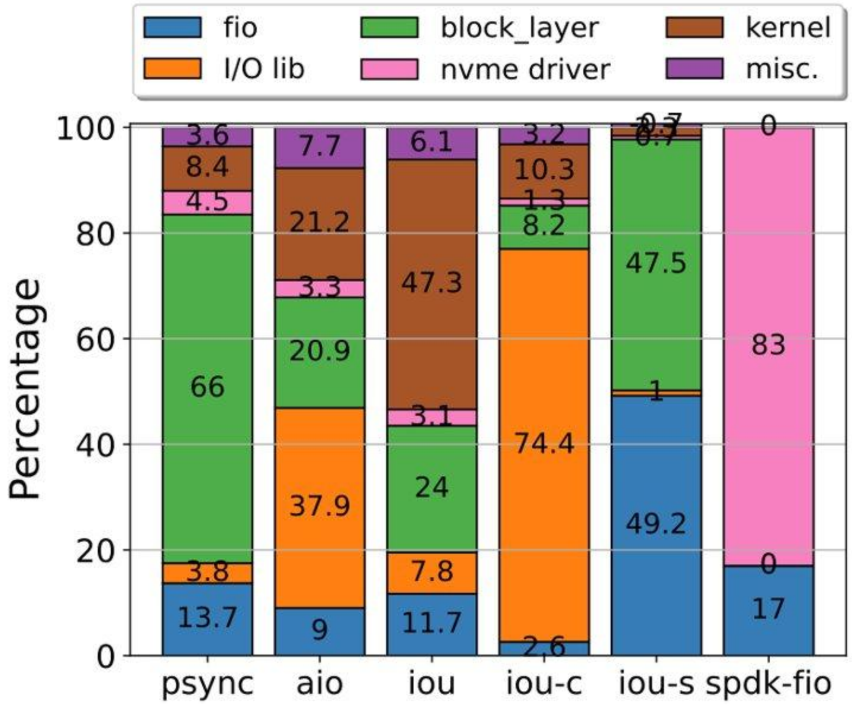

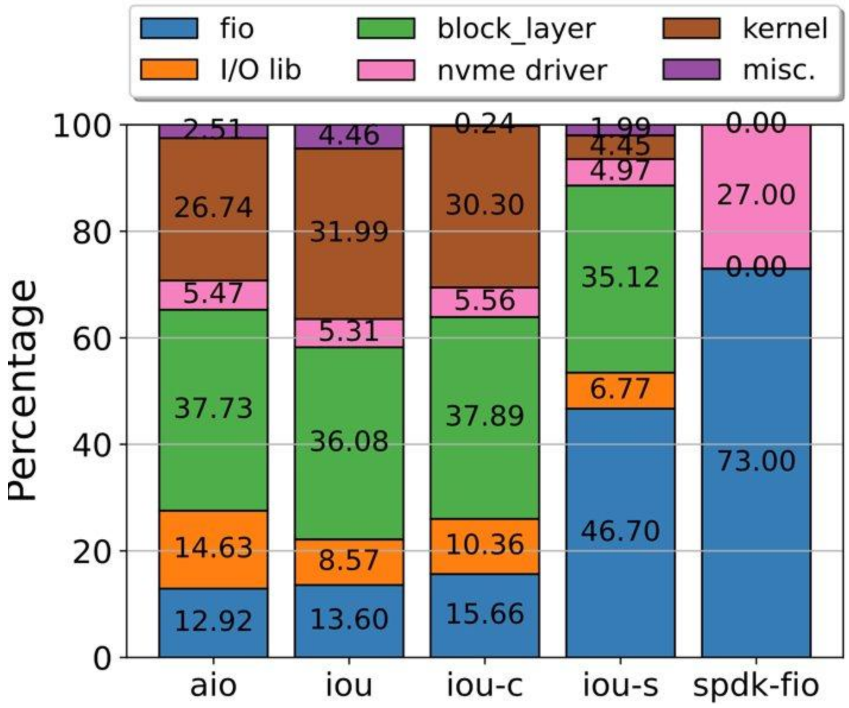

Backup Slides: I/O Scheduler

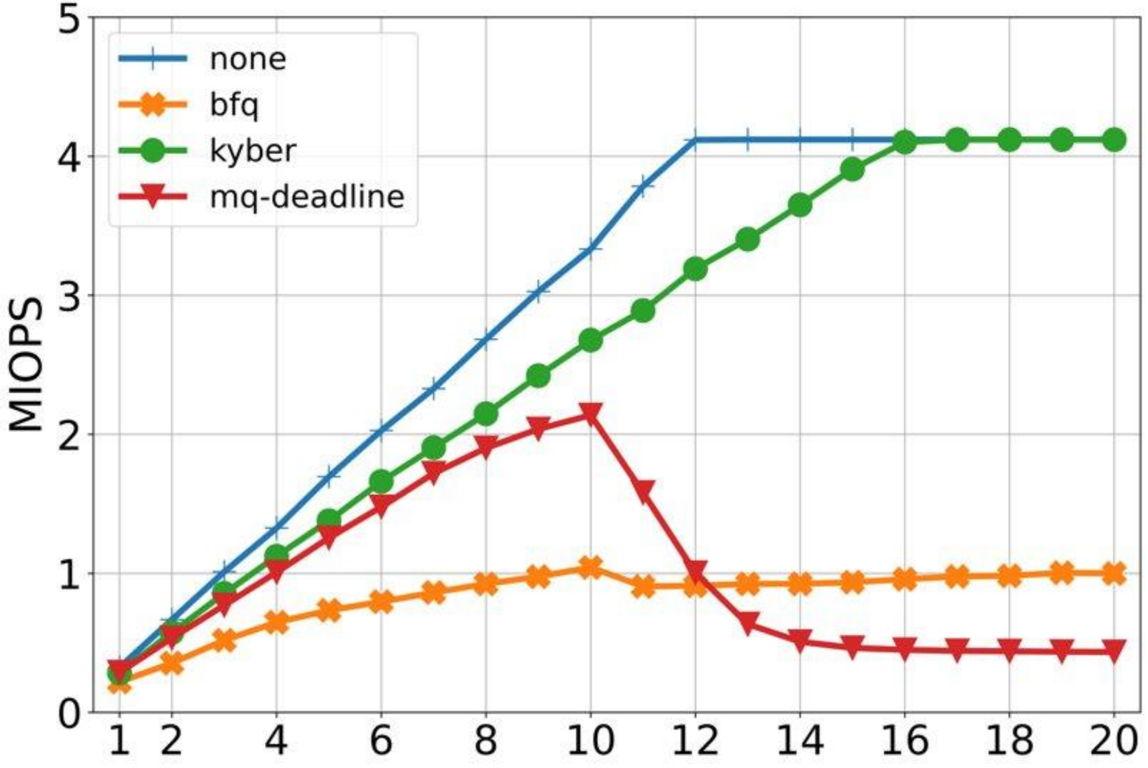

Backup Slides: I/O Scheduler

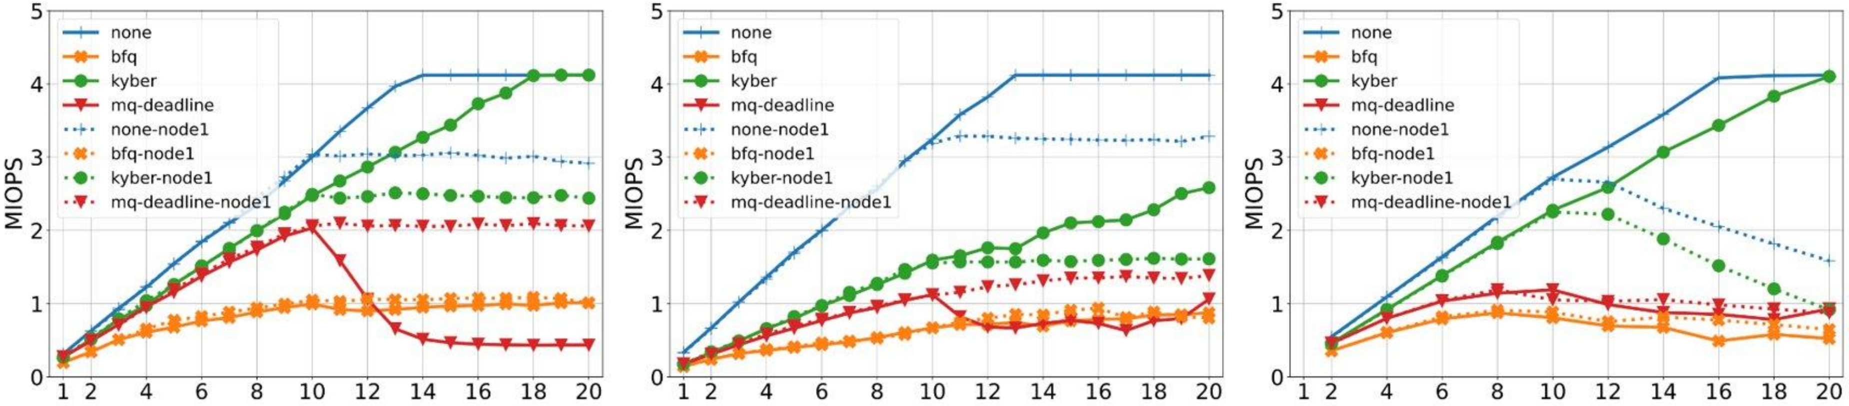

(a) default `io_uring` (iou). **(b)** with completion polling (iou-c). **(c)** with submission polling (iou-s).