

Zestawienie konferencji USENIX: <https://www.usenix.org/conferences/past>

Nowe:

- <https://ossna2024.sched.com/event/1aBOT/optimizing-scheduler-for-linux-gaming-changwoo-min-igalia>
<https://www.slideshare.net/slideshow/optimizing-scheduler-for-linux-gamingpdf/267643346>

Optimizing Scheduler for Linux Gaming.pdf

Open Source Summit North America 2024, April 16-18, 2024, Seattle, Washington (US)

This talk dives into how the scheduler impacts your gameplay on Linux and unveils our journey to smoother gameplay. How does task scheduling impact Linux gaming? Suboptimal task scheduling can cause stuttering while playing games on the Steam Deck game console. First, we nail down the enemy. What exactly is "stuttering," and how can we measure its impact on your gameplay? Next, we extensively analyzed the characteristics of game tasks from the scheduler's point of view. Characterizing task behavior in Linux gaming helps to understand why some schedulers create much stuttering and others create less and to unveil the secrets behind smooth vs. choppy performance. Lastly, we will share our progress on the optimized scheduler for reducing the stuttering problems in Linux gaming, especially Steam Deck. We implemented the scheduling policy based on sched_ext, a BPF-based extensible scheduling framework.

- The 2024 Linux Storage, Filesystem, Memory-Management, and BPF Summit, May 13, <https://lwn.net/Articles/973683/>
Playlist: <https://www.youtube.com/playlist?list=PLbzoR-pLrL6oj1rVTXlnV7cOuetvjKn9q>
- LPC 2024, Linux Plumbers Conference, <https://lpc.events/event/18/timetable/?view=lpc>
 - "Hey, psst, try this." The underground culture around custom CPU schedulers. <https://lpc.events/event/18/contributions/1772/>
-

2024

1. Systor 2024, The ACM International Systems and Storage Conference

Accepted papers:

<https://www.systor.org/2024/accepted-papers/#1630482835121-31ef82a8-bd1f>

Na przykład:

- ~~TwinPilots: A New Computing Paradigm for GPU-CPU Parallel LLM Inference (best paper)~~
<https://dl.acm.org/doi/10.1145/3688351.3689164>
- Performance Characterization of SmartNIC NVMe-over-Fabrics Target Offloading
<https://dl.acm.org/doi/10.1145/3688351.3689154>
- Offloading Datacenter Jobs to RISC-V Hardware for Improved Performance and Power Efficiency
<https://dl.acm.org/doi/10.1145/3688351.3689152>

- Can OS Specialization give new life to old carbon in the cloud?
<https://dl.acm.org/doi/10.1145/3688351.3689158> aaa
2. OSDI 2024, 18th USENIX Symposium on Operating Systems Design and Implementation
List of papers: <https://www.usenix.org/conference/osdi24/technical-sessions>
Full proceedings: https://www.usenix.org/system/files/osdi24_full_proceedings.pdf
Na przykład:
 - Nomad: Non-Exclusive Memory Tiering via Transactional Page Migration
<https://www.usenix.org/conference/osdi24/presentation/xiang>
 - ~~Harvesting Memory-bound CPU Stall Cycles in Software with MSH~~
<https://www.usenix.org/conference/osdi24/presentation/luo>
 - ~~DRust: Language-Guided Distributed Shared Memory with Fine Granularity, Full Transparency, and Ultra Efficiency~~
<https://www.usenix.org/conference/osdi24/presentation/ma-haoran>
 - Llumnix: Dynamic Scheduling for Large Language Model Serving
<https://www.usenix.org/conference/osdi24/presentation/sun-biao> aa
 3. USENIX ATC 2024, 2024 USENIX Annual Technical Conference
List of papers: <https://www.usenix.org/conference/atc24/technical-sessions>
Na przykład:
 - ~~Limitations and Opportunities of Modern Hardware Isolation Mechanisms~~
<https://www.usenix.org/conference/atc24/presentation/chen-xiangdong>
 - ExtMem: Enabling Application-Aware Virtual Memory Management for Data-Intensive Applications
<https://www.usenix.org/conference/atc24/presentation/jalalian>
 - ~~PeRF: Preemption-enabled RDMA Framework~~
<https://www.usenix.org/conference/atc24/presentation/lee>
 - ~~An Empirical Study of Rust for Linux: The Success, Dissatisfaction, and Compromise~~
<https://www.usenix.org/conference/atc24/presentation/li-hongyu>
 - Scalable and Effective Page-table and TLB management on NUMA Systems
<https://www.usenix.org/conference/atc24/presentation/gao-bin-scalable>
 - SuperBench: Improving Cloud AI Infrastructure Reliability with Proactive Validation
<https://www.usenix.org/conference/atc24/presentation/xiong>
 - StreamCache: Revisiting Page Cache for File Scanning on Fast Storage Devices
<https://www.usenix.org/conference/atc24/presentation/li-zhiyue>
 - ScalaCache: Scalable User-Space Page Cache Management with Software-Hardware Coordination
<https://www.usenix.org/conference/atc24/presentation/peng>
 - Every Mapping Counts in Large Amounts: Folio Accounting
<https://www.usenix.org/conference/atc24/presentation/hildenbrand>
 4. HotStorage 2024, The 16th ACM Workshop on Hot Topics in Storage and File Systems
<https://www.hotstorage.org/2024/>
Accepted papers: <https://www.hotstorage.org/2024/accepted.html>
Na przykład:
 - Dictionary Based Cache Line Compression
<https://dl.acm.org/doi/10.1145/3655038.3665941>
 - Improving Virtualized I/O Performance by Expanding the Polled I/O Path of Linux

- <https://dl.acm.org/doi/10.1145/3655038.3665944>
 - Secure Archival is Hard... Really Hard
<https://dl.acm.org/doi/10.1145/3655038.3666093>
 - Can Storage Devices be Power Adaptive?
<https://dl.acm.org/doi/10.1145/3655038.3665945>
 - Can ZNS SSDs be Better Storage Devices for Persistent Cache?
<https://dl.acm.org/doi/10.1145/3655038.3665946>
 - Quantitative Analysis of Storage Requirement for Autonomous Vehicles
<https://dl.acm.org/doi/10.1145/3655038.3665948>
 - Asymmetric RAID: Rethinking RAID for SSD Heterogeneity
<https://dl.acm.org/doi/10.1145/3655038.3665952>
 - Breaking Barriers: Expanding GPU Memory with Sub-Two Digit Nanosecond Latency CXL Controller
<https://dl.acm.org/doi/10.1145/3655038.3665953> zz
5. FAST 2024, 22nd USENIX Conference on File and Storage Technologies
<https://www.usenix.org/conference/fast24>
Full proceedings: https://www.usenix.org/system/files/fast24-full_proceedings.pdf
Na przykład:
- Lessons Learnt in Trying to Build New Storage Technologies
<https://www.usenix.org/conference/fast24/presentation/rowstron>
 - ~~TeRM: Extending RDMA-Attached Memory with SSD~~
<https://www.usenix.org/conference/fast24/presentation/yang-zhe>
 - Combining Buffered I/O and Direct I/O in Distributed File Systems
<https://www.usenix.org/conference/fast24/presentation/qian>
 - Optimizing File Systems on Heterogeneous Memory by Integrating DRAM Cache with Virtual Memory Management
<https://www.usenix.org/conference/fast24/presentation/liu-yubo>
 - I/O Passthru: Upstreaming a flexible and efficient I/O Path in Linux
<https://www.usenix.org/conference/fast24/presentation/joshi>
 - Storage Systems in the LLM Era
<https://www.usenix.org/conference/fast24/presentation/panel-storage-systemsaaa>
6. NSDI 2024, 21st USENIX Symposium on Networked Systems Design and Implementation
Full proceedings: https://www.usenix.org/sites/default/files/nsdi24-full_proceedings.pdf
List of papers: <https://www.usenix.org/conference/nsdi24/technical-sessions>
Na przykład:
- ~~Horus: Granular In-Network Task Scheduler for Cloud Datacenters~~
<https://www.usenix.org/conference/nsdi24/presentation/yassini>
 - DINT: Fast In-Kernel Distributed Transactions with eBPF
<https://www.usenix.org/conference/nsdi24/presentation/zhou-yang>
 - Netcastle: Network Infrastructure Testing At Scale
<https://www.usenix.org/conference/nsdi24/presentation/sherwood>
 - SIEVE is Simpler than LRU: an Efficient Turn-Key Eviction Algorithm for Web Caches
<https://www.usenix.org/conference/nsdi24/presentation/zhang-yazhuo>
 - ~~Load is not what you should balance: Introducing Prequal~~
<https://www.usenix.org/conference/nsdi24/presentation/wydrowski>
 - ~~MobileConfig: Remote Configuration Management for Mobile Apps at Hyperscale~~
<https://www.usenix.org/conference/nsdi24/presentation/guo>

2023

7. FAST'23, 21st USENIX Conference on File and Storage Technologies

<https://www.usenix.org/conference/fast23>

<https://www.usenix.org/conference/fast23/technical-sessions>

Na przykład:

- Building and Operating a Pretty Big Storage System (My Adventures in Amazon S3)
<https://www.usenix.org/conference/fast23/presentation/warfield>
- GL-Cache: Group-level learning for efficient and high-performance caching
<https://www.usenix.org/conference/fast23/presentation/yang-juncheng>
- ~~HadaFS: A File System Bridging the Local and Shared Burst Buffer for Exascale Superecomputers~~
<https://www.usenix.org/conference/fast23/presentation/he>
- ~~Citron: Distributed Range Lock Management with One-sided RDMA~~
<https://www.usenix.org/conference/fast23/presentation/gao>
- NVMeVirt: A Versatile Software-defined Virtual NVMe Device
<https://www.usenix.org/conference/fast23/presentation/kim-sang-hoon>
- Multi-view Feature-based SSD Failure Prediction: What, When, and Why
<https://www.usenix.org/conference/fast23/presentation/zhang>
- Integrated Host-SSD Mapping Table Management for Improving User Experience of Smartphones
<https://www.usenix.org/conference/fast23/presentation/kim-yoona>

8. OSDI'23, 17th USENIX Symposium on Operating Systems Design and Implementation

<https://www.usenix.org/conference/osdi23>

<https://www.usenix.org/conference/osdi23/technical-sessions>

Na przykład:

- ~~RON: One-Way Circular Shortest Routing to Achieve Efficient and Bounded-waiting Spinlocks~~
<https://www.usenix.org/conference/osdi23/presentation/lo>
- ~~Userspace Bypass: Accelerating Syscall-intensive Applications~~
<https://www.usenix.org/conference/osdi23/presentation/zhou-zhe>
- ~~Honeycomb: Secure and Efficient GPU Executions via Static Validation~~
<https://www.usenix.org/conference/osdi23/presentation/mai>
- ~~Core slicing: closing the gap between leaky confidential VMs and bare-metal cloud~~
<https://www.usenix.org/conference/osdi23/presentation/zhou-ziqiao>
- Johnny Cache: the End of DRAM Cache Conflicts (in Tiered Main Memory Systems)
<https://www.usenix.org/conference/osdi23/presentation/lepers>
- SMART: A High-Performance Adaptive Radix Tree for Disaggregated Memory
<https://www.usenix.org/conference/osdi23/presentation/luo>

9. SYSTOR, 23 The 16th ACM International System and Storage Conference

<https://www.systor.org/2023/>

Na przykład

- ~~F3: Serving Files Efficiently in Serverless Computing (Best paper)~~
<https://dl.acm.org/doi/10.1145/3579370.3594771>
- Elastic RAID: Implementing RAID over SSDs with Built-in Transparent Compression
<https://dl.acm.org/doi/10.1145/3579370.3594773>
- BOOSTER: Rethinking the erase operation of low-latency SSDs to achieve high throughput and less long latency
<https://dl.acm.org/doi/10.1145/3579370.3594774>
- DPFS: DPU-Powered File System Virtualization
<https://dl.acm.org/doi/10.1145/3579370.3594769>

10. NSDI'23 20th USENIX Symposium on Networked Systems Design and Implementation

<https://www.usenix.org/conference/nsdi23>

<https://www.usenix.org/conference/nsdi23/technical-sessions>

Na przykład:

- CausalSim: A Causal Framework for Unbiased Trace-Driven Simulation (Best paper)
<https://www.usenix.org/conference/nsdi23/presentation/alomar>
- ~~Empowering Azure Storage with RDMA~~
<https://www.usenix.org/conference/nsdi23/presentation/bai>

11. HotStorage'23, The 15th ACM Workshop on Hot Topics in Storage and File Systems

<https://www.hotstorage.org/2023/>

Na przykład

- Do we still need IO schedulers for low-latency disks?
<https://dl.acm.org/doi/10.1145/3599691.3603400>
- When F2FS Meets Compression-Based SSD!
<https://dl.acm.org/doi/10.1145/3599691.3603402>
- Excessive SSD-Internal Parallelism Considered Harmful
<https://dl.acm.org/doi/10.1145/3599691.3603412>
- When Caching Systems Meet Emerging Storage Devices: A Case Study
<https://dl.acm.org/doi/10.1145/3599691.3603413>

12. LinuxCon@OSS EU 2023

https://www.youtube.com/playlist?list=PLbzoR-pLrL6p445pIqTLUIzsceFp_puSy

13. The 2023 [Linux Storage, Filesystem, Memory-Management, and BPF Summit](#)

Na YouTube

https://www.youtube.com/playlist?list=PLbzoR-pLrL6rlmdpJ3-oMgU_zxc1wAhjS

i lwn.net

- <https://lwn.net/Articles/932748/>, Zoned storage and filesystems
- <https://lwn.net/Articles/931668/>, Peer-to-peer DMA
- itp.

14. Storage Developer Conference (SDC 2023), September 2023

<https://storagedeveloper.org/>

Prezentacje na YouTube (SDC 2023):

<https://www.youtube.com/watch?v=SHR2uIavnbk>

15. ACM Symposium on Operating Systems Principles (SOSP 2023), October 2023
<https://sosp2023.mpi-sws.org/>

Różne

- Memory deduplication (od Marcina Copika z ETH Zurich)
Artykuł: https://mcpik.github.io/assets/pdf/2023_bigdata_upm.pdf
Slajdy: https://mcpik.github.io/assets/pdf/2023_bigdata_upm_slides.pdf

 - ‘Mind-blowing’ IBM chip speeds up AI, https://www.nature.com/articles/d41586-023-03267-0?utm_source=Live+Audience&utm_campaign=2437e496f0-briefing-dy-20231020&utm_medium=email&utm_term=0_b27a691814-2437e496f0-51328812
(trzeba poszukać w sieci jakiegoś sensownego artykułu i najlepiej zrobić przegląd najnowszych architektur wspomaających AI)

 - Flame graphs:
<https://www.brendangregg.com/flamegraphs.html>,
<https://queue.acm.org/detail.cfm?id=2927301>
<https://queue.acm.org/downloads/2016/Gregg4.svg>
<https://www.youtube.com/watch?v=VMpTU15rIZY>
<https://www.youtube.com/watch?v=HKQR7wVapgk>

 - Storage Trends 2024, <https://www.youtube.com/watch?v=l8qPrFtY6qE>
 - Storage trends in 2023 and beyond, <https://www.youtube.com/watch?v=QmmVgz5zhgg>
-

Od Andrzeja Jackowskiego

- The Fungible DPU A New Category of Microprocessor

What is New in the Linux kernel

<https://lwn.net/Kernel/>

<https://lwn.net/Kernel/Index/> i wyszukać po ostatnim roku (2024)

- [Support for the TSO memory model on Arm CPUs](#)
- [Atomic writes without tears](#)
- [Virtual machine scheduling with BPF](#)
- [Modernizing BPF for the next 10 years](#)
- [Another push for sched_ext](#)
- [What's scheduled for sched_ext](#)
- [Sched_ext at LPC 2024](#)
- [A proposal for shared memory in BPF programs](#)
- [Measuring and improving buffered I/O](#)
- [Supporting larger block sizes in filesystems](#)
- [Memory-management: tiered memory, huge pages, and EROFS](#)
- itp

1. Prace z Research at Google, Distributed Systems and Parallel Computing

<https://research.google/research-areas/distributed-systems-and-parallel-computing/>

Zdecydowana większość prac bardzo dobrze pasuje do tematyki seminarium, oczywiście lepsze są te najnowsze, z ostatniego roku.

Pełna lista publikacji: <https://research.google/pubs/?area=distributed-systems-and-parallel-computing>