2021 IEEE 28th International Conference on High Performance Computing, Data, and Analytics (HiPC) | 978-1-6654-1016-8/21/$31.00 ©2021 IEEE | DOI: 10.1109/HIPC53243.2021.00011

Improving Efficiency and Performance Through Faster Scheduling Mechanisms

Adam Belay

Massachusetts Institute of Technology

Abstract:

Today's high-performance applications demand microsecond-scale tail latencies and high request rates from operating systems, and most applications handle loads that have high variance over multiple timescales. Achieving these goals in an efficient way, where multiple applications must share the resources within each machine, is a difficult problem for today's operating systems. Their primary shortcoming is that they adjust resource allocations too slowly—over 10s or 100s of milliseconds. This leaves them vulnerable to extreme spikes in tail latency under shifts in resource demand while, at the same time, leaving resources underutilized.

In this talk, I will present Caladan, a new approach to CPU scheduling that can achieve significantly better tail latency, throughput, and resource utilization through a collection of control signals and policies that rely on faster core allocation. Caladan consists of a centralized scheduler core that actively manages resource contention in the memory hierarchy and between hyperthreads, and a kernel module that bypasses the standard Linux Kernel scheduler to support microsecond-scale monitoring and placement of tasks. Caladan can maintain nearly perfect performance isolation between colocated applications with phased behaviors, and it lays the groundwork toward applying faster scheduling techniques to a broader range of systems problems.

Bio:

Adam Belay is an Assistant Professor of Computer Science at the Massachusetts Institute of Technology, where he works on operating systems, networking, and systems security. During his Ph.D. at Stanford, he developed Dune, a system that safely exposes privileged CPU instructions to userspace; and IX, a dataplane operating system that significantly accelerates I/O performance. Dr. Belay's current research interests lie in developing systems that cut across hardware and software layers to increase datacenter efficiency and performance. He is a member of the Parallel and Distributed Operating Systems Group, and a recipient of a Google Faculty Award, a Facebook Research Award, and the Stanford Graduate Fellowship.