



Contribution ID: 62

Type: **not specified**

OS Scheduling with Nest: Keeping Tasks Close Together on Warm Cores

Monday, 12 September 2022 12:45 (45 minutes)

To best support highly parallel applications, Linux's CFS scheduler tends to spread tasks across the machine on task creation and wakeup. It has been observed, however, that in a server environment, such a strategy leads to tasks being unnecessarily placed on long-idle cores that are running at lower frequencies, reducing performance, and to tasks being unnecessarily distributed across sockets, consuming more energy. In this talk, we propose to exploit the principle of core reuse, by constructing a nest of cores to be used in priority for task scheduling, thus obtaining higher frequencies and using fewer sockets. We implement the Nest scheduler in the Linux kernel. While performance and energy usage are comparable to CFS for highly parallel applications, for a range of applications using fewer tasks than cores, Nest improves performance 10%-2x and can reduce energy usage.

I agree to abide by the anti-harassment policy

Yes

Primary author: LAWALL, Julia (Inria)

Co-authors: CHHAYA-SHAILESH, Himadri (Inria); LOZI, Jean-Pierre (Oracle Labs); LEPERS, Baptiste (University of Sydney); ZWAENEPOEL, Willy (University of Sydney); MULLER, Gilles (Inria)

Presenter: LAWALL, Julia (Inria)

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track