Demande de stage de recherche 2022

Master thesis subject title

LinuxTransformers : Understanding Linux scheduling bottlenecks

Laboratory name : INRIA Rennes
Team : WIDE
Localization : Rennes - University of Rennes 1 - INRIA - FRANCE

Supervisor
Name : **David Bromberg**
Position : **Professeur des universités**
Email : **david.bromberg@irisa.fr**
Phone number : **+ 33 2 99 84 22 05**

Co-supervisor
Name : **Djob Mvondo**
Position : **Maître de conferences**
Email : **barbe-thystere.mvondodjob@univ-rennes1.fr**

External collaborators
Name : **Julia Lawall**
Position : **Directrice de recherche**

Keywords. System, Large scale distributed system,Linux Kernel, Linux scheduling,

Research axis. Understanding and optimizing the Linux scheduler.

1 Master thesis proposal

(1233 words)

Context. The process scheduler is the part of an operating system that decides what thread runs on what core at what time. As such, it has a critical impact on application performance, particularly for multithreaded applications. While some applications simply spread out across the available cores, with each thread running continuously without disturbance on its preferred core, others involve frequent synchronization, I/O, etc. In practice, whenever a thread gives up access to a core, another thread can replace it, leading to migrations, loss of locality, and degraded performance. At the same

time, evolutions in the scheduler can introduce errors, that can also degrade performance on specific workloads. Tools exist to trace scheduling behavior, that can help identify the presence of such problems, but the sheer volume of information available makes it difficult to map a scheduling trace to a root cause.

Objectives. The main aim of the project is to understand the behavior of the multicore Linux process scheduler under heavy load. We will particularly focus on the diagnosis of scheduling anomalies, as can occur in specific runs with a given scheduler implementation and as can be introduced over time by bugs in the scheduler implementation. Our starting point is the observation that we can produce an unlimited number of execution traces, including over multiple versions of the Linux kernel. Analogous to the diagnosis of illnesses using image processing, we would like to develop a model of the expected behavior of a scheduler on a given application through the use of Transformers models [1][6], and then to detect as anomalous execution traces that deviate from these models. The goal is to detect scheduling problems quickly, when bugs are introduced into the source code, to detect scheduling problems from short-running examples where the problem may not easily be visible to a person looking at a trace, and to connect scheduling problems reflected in a trace to specific elements of the scheduler or the application source code. Based on the results, we will consider how to improve the Linux scheduler to provide better performance.

Expected background : This PhD requires familiarity with the C language and operating systems concepts, as well as a background in machine learning.

References

[1] Ashish Vaswani et al. "Attention is All you Need". In : *Advances in Neural Information Processing Systems 30 : Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*. Sous la dir. d'Isabelle Guyon et al. 2017, p. 5998-6008. URL : https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

[2] Justinien Bouron et al. "The Battle of the Schedulers : FreeBSD ULE vs. Linux CFS". In : *2018 USENIX Annual Technical Conference, USENIX ATC 2018, Boston, MA, USA, July 11-13, 2018*. Sous la dir. d'Haryadi S. Gunawi et Benjamin Reed. USENIX Association, 2018, p. 85-96. URL : https://www.usenix.org/conference/atc18/presentation/bouron.

[3] Cédric Courtaud et al. "Improving Prediction Accuracy of Memory Interferences for Multicore Platforms". In : *IEEE Real-Time Systems Symposium, RTSS 2019, Hong Kong, SAR, China, December 3-6, 2019*. IEEE, 2019, p. 246-259. DOI : 10.1109/RTSS46320.2019.00031. URL : https://doi.org/10.1109/RTSS46320.2019.00031.

[4] Redha Gouicem et al. "Fewer Cores, More Hertz : Leveraging High-Frequency Cores in the OS Scheduler for Improved Application Performance". In : *2020 USENIX Annual Technical Conference, USENIX ATC 2020, July 15-17, 2020*. Sous la dir. d'Ada Gavrilovska et Erez Zadok. USENIX Association, 2020, p. 435-448. URL : https://www.usenix.org/conference/atc20/presentation/gouicern.

[5] Baptiste Lepers et al. "Provable multicore schedulers with Ipanema : application to work conservation". In : *EuroSys '20 : Fifteenth EuroSys Conference 2020, Heraklion, Greece, April 27-30, 2020*. Sous la dir. d'Angelos Bilas et al. ACM, 2020, 3:1-3:16. DOI : 10.1145/3342195.3387544. URL : https://doi.org/10.1145/3342195.3387544.

[6] Rishi Bommasani et al. “On the Opportunities and Risks of Foundation Models”. In : *CoRR* abs/2108.07258 (2021). arXiv : 2108.07258. URL : https://arxiv.org/abs/2108.07258.