


Researchers mine cache of Intel processors to speed up data packet processing

July 8 2019



Credit: KTH The Royal Institute of Technology

Developed with Ericsson Research, the slice-aware memory-management scheme allows frequently used data to be accessed more quickly via the last-level cache of memory (LLC) of an Intel Xeon CPU. By establishing a key-value store and allocating memory in a way that it maps to the most appropriate LLC slice, they demonstrated both high-speed packet processing and improved performance of a key-value store.

The team used the proposed scheme to implement a tool called CacheDirector, which makes Data Direct I/O (DDIO) slice-aware and published a conference paper, Make the Most out of Last Level Cache in Intel Processors, which was presented at EuroSys 2019 in the spring.

"At the moment, a server receiving 64-byte packets at 100Gbps has just 5.12 nanoseconds to process each packet before the next one arrives," says co-author Alireza Farshin, a doctoral student at KTH's Network Systems Laboratory. But if data is routed to the right cache slice in the CPU, it can be accessed faster—allowing faster processing of more packets, in under 5 nanoseconds.

Data Direct I/O (DDIO) sends packets to random slices, which is far from efficient. Given today's non-uniform cache architecture (NUCA), the cache-management solution is invaluable, says KTH Professor Dejan Kostic, who led the research.

"When combined with the introduction of dynamic headroom in the Data Plane Development Kit (DPDK), the packet's header can be placed in the slice of the LLC that is closest to the relevant processing core. As a result, the core can access packets faster while also reducing queuing time," he says.

"Our work demonstrates that taking advantage of nanosecond improvements in [latency](#) can have a large impact on the performance of applications running on already highly-optimized computer systems," Farshin says. The team found that for a CPU running at 3.2GHz, CacheDirector can save up to around 20 cycles per access to the LLC which amounts to 6.25 nanoseconds. This accelerates packet processing and reduces tail latencies of optimized Network Function Virtualization (NFV) service chains running at 100Gbps by up to 21.5 percent.

More information: Alireza Farshin et al. Make the Most out of Last

Level Cache in Intel Processors, *Proceedings of the Fourteenth EuroSys Conference 2019 CD-ROM on ZZZ - EuroSys '19* (2019). [DOI: 10.1145/3302424.3303977](https://doi.org/10.1145/3302424.3303977)

Provided by KTH Royal Institute of Technology

Citation: Researchers mine cache of Intel processors to speed up data packet processing (2019, July 8) retrieved 18 April 2024 from <https://techxplore.com/news/2019-07-cache-intel-processors-packet.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.