Make the Most Out of Last Level Cache in Intel Processors

Alireza Farshin[+], **Amir Roozbeh**[+*], Gerald Q. Maguire Jr.[+], Dejan Kostic[+]

farshin@kth.se
amirrsk@kth.se
maguire@kth.se
dmk@kth.se

[+] KTH Royal Institute of Technology (EECS/COM) [*] Ericsson Research

1 Problem

Faster links exposes processing elements to packets at a higher rate, but the performance of the processors is no longer doubling at the earlier rate, making it harder to keep up with the growth in link speeds. For instance, a server receiving 64 B packets at a rate of 100 Gbps has only **5.12ns** to process a packet before the next packet arrives.

Approach:

It is essential to exploit every oppurtunity to optimize current computer systems. We focus on better management of LLC.

2 Slice-aware Memory Management

In modern (Intel) processors, Last Level Cache (LLC) is divided into multiple slices.

For each core, accessing data stored in a closer slice is faster than accessing data stored in other slices.

We introduce a slice-aware memory management scheme, wherein data can be mapped to specific LLC slice(s).

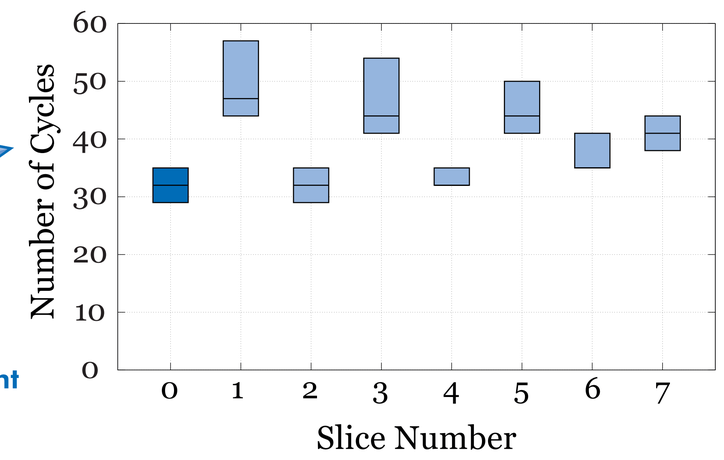

Measuring read access time from Core 0 to different LLC slices.

Accessing the closer LLC Slice can save up to 20 cycles, i.e., 6.25 ns.

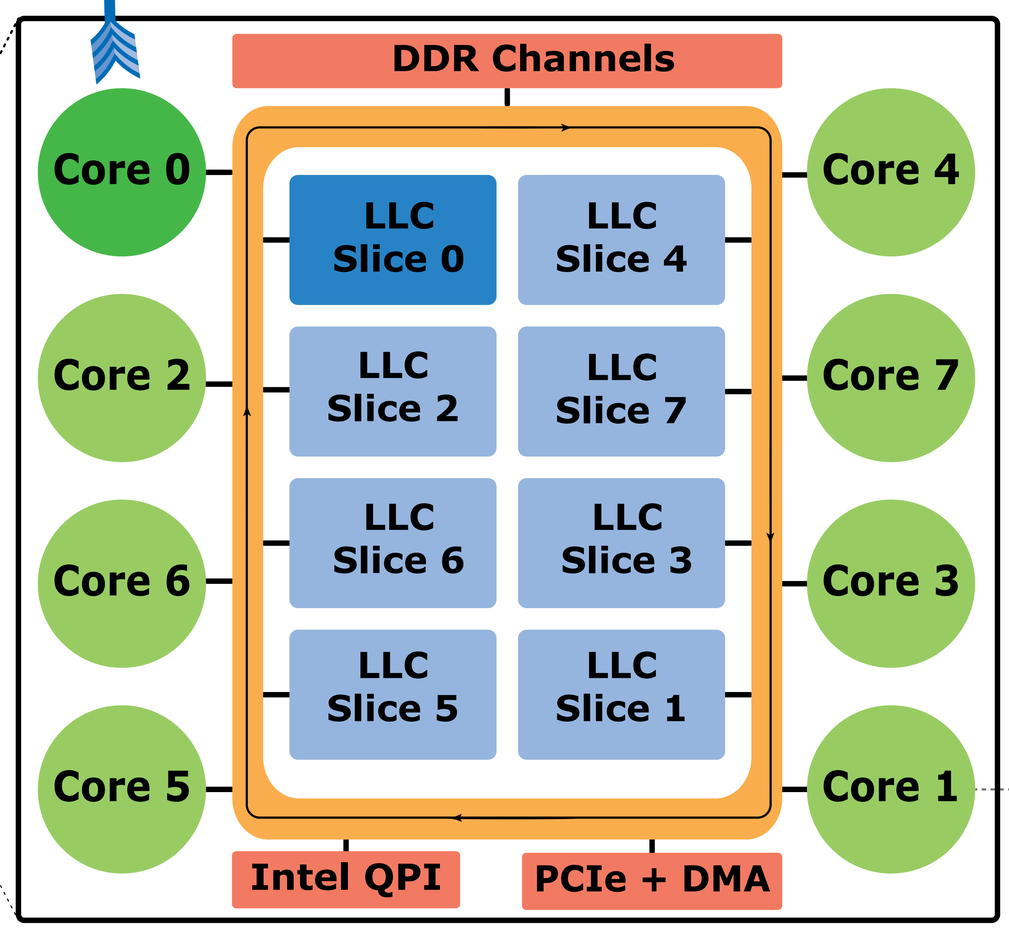

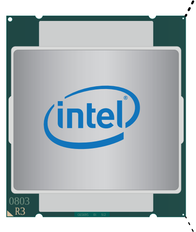

Intel® Xeon® E5-2667 v3

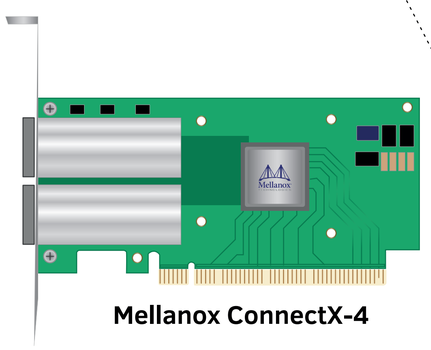

Mellanox ConnectX-4

Sending/Receiving Packets via DDIO

Data Direct I/O (DDIO) technology sends packets to LLC rather than DRAM.

Lower level caches are private to each core.

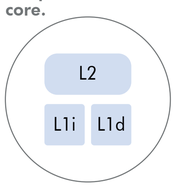

3 CacheDirector

CacheDirector is a network I/O solution that extends DDIO by implementating slice-aware memory management as an extension to Data Plane Development Kit (DPDK).

It sends the first 64 B of a packet (containing the packet's header) directly to the appropriate LLC slice.

The CPU core that is responsible for processing a packet can access the packet header in fewer CPU cycles.

4 Impact on Tail Latency for NFV Service Chains

We evaluate the performance of NFV service chains in the presence of CacheDirector.

CacheDirector reduces tail latencies by up to 21.5% for optimized NFV service chains that are running at 100 Gbps.

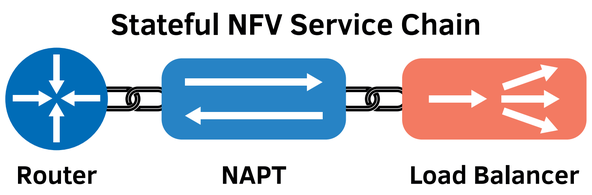

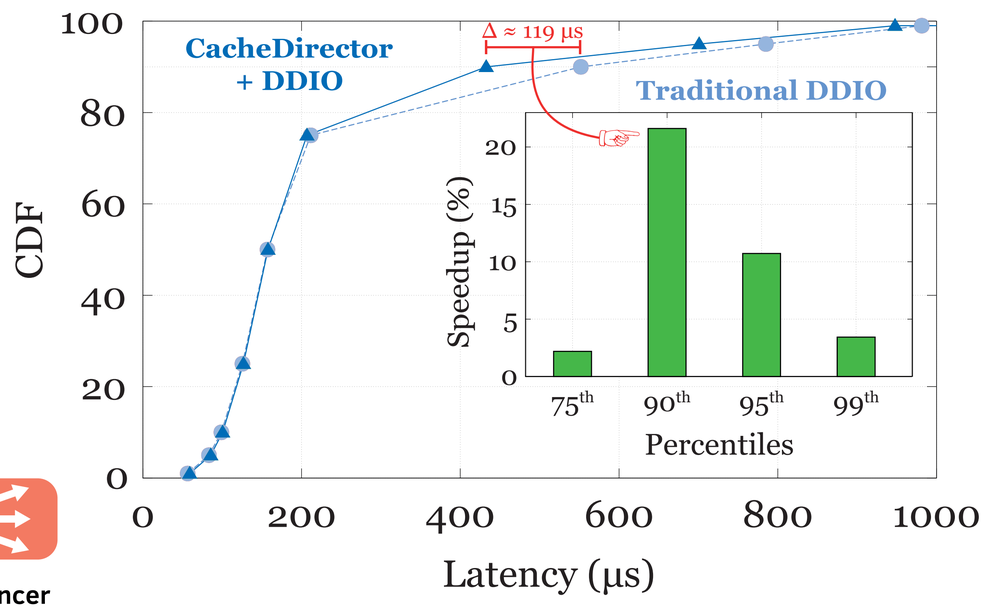

5 Conclusions

With slice-aware memory management scheme, which exploits the latency differences in accessing different LLC slices, it is possible to boost application performance and realize cache isolation.

We proposed CacheDirector, a network I/O solution, which utilizes slice-aware memory management.

CacheDirector increases efficiency and performance, while reducing tail latencies over the state-of-the-art.

Our Paper

Work supported by SSF, WASP, and ERC.