

Optimizing Intel Data Direct I/O Technology for Multi-hundred-gigabit Networks

Alireza Farshin*[†]

KTH Royal Institute of Technology
farshin@kth.se

Gerald Q. Maguire Jr.

KTH Royal Institute of Technology
maguire@kth.se

Amir Roozbeh*[†]

KTH Royal Institute of Technology
Ericsson Research
amirrsk@kth.se

Dejan Kostić

KTH Royal Institute of Technology
dmk@kth.se

Introduction

Digitalization across society is expected to produce a massive amount of data, leading to the introduction of faster network interconnects. In addition, many Internet services require high throughput and low latency. However, having *only* faster links does not guarantee throughput or low latency. Therefore, it is essential to perform holistic system optimization to fully take advantage of the faster links to provide high-performance services [1]. Intel® Data Direct I/O (DDIO) [2] is a recent technology that was introduced to facilitate the deployment of high-performance services based on fast interconnects. We evaluated the effectiveness of DDIO for multi-hundred-gigabit networks. This paper briefly discusses our findings on DDIO, which show the necessity of optimizing/adapting it to address the challenges of multi-hundred-gigabit-per-second links.

Data Direct I/O Technology (DDIO)

DDIO was introduced to improve the performance of I/O applications by mitigating expensive DRAM accesses. It transfers packets directly to the Last Level Cache (LLC), as opposed to traditional techniques that target main memory. DDIO updates a cache line, a 64-B chunk of a packet, if its address already available in LLC. Otherwise, it allocates the cache line in a limited portion of LLC (i.e., 2 ways in an n-way set-associative cache).

DDIO can become a Bottleneck

Our study shows that DDIO does *not* provide the expected benefits in some scenarios, specifically when incoming packets repeatedly evict previously received packets (i.e., not-yet-processed and already-processed packets) from the LLC. The probability of eviction is high when:

1. Applications are I/O bound,
2. Receiving rate gets close to ~ 100 Gbps,
3. Packet size is larger than ~ 512 B,
4. Number of receive (RX) descriptors is large, and
5. Load is not distributed equally among cores.

*Poster Presenter(s)

[†]Doctoral Student(s)

Fine-tuning DDIO

A little-discussed register called “IIO LLC WAYS” can be used to tune the capacity of DDIO. We show that fine-tuning DDIO enables us to process packets more efficiently at multi-hundred-gigabit rates, thereby eliminating packet evictions from LLC. Fig. 1 demonstrates the effectiveness of tuning DDIO while forwarding packets at 100 Gbps.




Figure 1. Using more DDIO ways (“W”) enables 2 cores to forward 1500-B packets at 100 Gbps with a larger number of descriptors while achieving better or similar tail latency.

Conclusion and Future Work

This paper shows that it is important to understand the details of DDIO and to optimize it appropriately for a given Internet service to achieve high-performance, especially with the introduction of multi-hundred-gigabit networks. Completing our study on DDIO and providing optimization guidelines remain as our future work.

References

- [1] Alireza Farshin, Amir Roozbeh, Gerald Q. Maguire Jr., and Dejan Kostić. 2019. Make the Most out of Last Level Cache in Intel Processors. In *Proceedings of the Fourteenth EuroSys Conference 2019* (Dresden, Germany) (*EuroSys '19*). ACM, New York, NY, USA, Article 8, 17 pages. <https://doi.org/10.1145/3302424.3303977>
- [2] Intel. 2012. Intel Data Direct I/O Technology Overview. <https://www.intel.com/content/www/us/en/io/data-direct-i-o-technology-brief.html>, accessed 2019-07-26.