A Highly Efficient I/O-based Out-of-Core Stencil Algorithm with Globally Optimized Temporal Blocking

Hiroko Midorikawa, Hideyuki Tan
Department of Computer and Information Science
Seikei University, JST CREST
Tokyo, Japan
midori@st.seikei.ac.jp

ABSTRACT

Stencil computation is one of the most important computation kernels in various scientific and engineering simulations. Stencil computations often require significant amounts of memory for addressing large-scale problems and/or for higher resolution data analysis. However, there is a limit to how much DRAM can be increased in main memory because of the number of memory slots on server boards, power consumption constraints, and other resource limitations. One of the common solutions used to satisfy this requirement is using cost-effective and large-capacity non-volatile memory (NVM).

We have been investigating out-of-core stencil computation algorithms and implementations designed for large-capacity type of NVMs (non-volatile memory), such as NAND flash [1-4]. This paper describes and discusses in the advantages and disadvantages of these out-of-core stencil algorithms. They are different in their schemes of using flash, data structures used in stencil computations, and the way of using blocking technique to increase data access locality for accelerating performance. We propose the most efficient I/O-based out-of-core stencil algorithm in detail as well as the evaluation of its performance. The algorithm uses highly parallel asynchronous I/O (AIO) and a temporal blocking technique without redundant calculations to overcome the latency divide between flash and DRAM. It achieves 80% of the performance of in-core computing using sufficient capacity of the main memory, even if available memory capacity is limited to 6.3% of the data size required in the stencil computation problem. In other words, the proposed algorithm degrades performance within 20% for a stencil problem that requires 2 TiB of data by using only 128 GiB of main memory and flash SSDs.

I/O-based algorithms gains not only higher performance but also more stable behavior, compared to mmap-based algorithms using a file-memory-map of files in the flash SSD. The processes using mmap large capacity of memory are often killed unexpectedly by the operating system (OS), i.e. the out of memory (OOM) killer. The performance of mmap-based algorithms also depends on the size of available unused main memory of DRAM at runtime, because the remaining main memory is used by the OS as page cache area for memory-mapped file accessing. When the size of the remaining unused memory is limited at runtime, mmap-based algorithms exhibit low performance.

Another advantage of I/O-based algorithms is that using explicit I/O provides benefits in calculating globally optimal blocking sizes in the temporal blocking technique. To achieve the maximum performance, the algorithms use appropriate spatial and temporal blocking parameters that minimizes the amount of data transferred between the flash device and the DRAM, which is a dominant factor affecting the performance of out-of-core computing. By an auto-tuning system, such as Blk-Tune [3], the I/O-based algorithm allows users to maximize the performance easily for particular platforms and application settings.

The algorithms proposed here are not limited to flash but are also available to other I/O-based NVMs and read/write-based memories. With the use of the algorithm, large-scale stencil problems can be solved with a limited size of main memory.

REFERENCES

[1] H. Midorikawa, H. Tan, and T. Endo, “An Evaluation of the Potential of Flash SSD as Large and Slow Memory for Stencil Computations,” Proc. 2014 Int. Conf. on High Performance Computing and Simulation IEEE-HPCS2014, Jul. 2014, pp. 268-277.

[2] H. Midorikawa and H. Tan, “Locality-Aware Stencil Computations Using Flash SSDs as Main Memory Extension,” Proc. IEEE/ACM Int. Symp. on Cluster, Cloud and Grid Computing CCGrid2015, May 2015, pp. 1163-1168.

[3] H. Midorikawa: “Blk-Tune: Blocking Parameter Auto-Tuning to Minimize Input-Output Traffic for Flash-Based Out-of-Core Stencil Computations”, The 11th Int. workshop on Automatic Performance Tuning, iWAPT2016, IEEE Int. Parallel and Distributed Process. Symp. Workshops (DOI 10.1109/IPDPSW.2016.48), pp.1516-1526, May 2016

[4] H.Midorikawa, H.Tan: "Evaluation of Flash-based Out-of-core Stencil Computation Algorithms for SSD-Equipped Clusters", The 22nd IEEE International Conference on Parallel and Distributed Systems ICPADS2016 , pp.1031-1040, DOI: 10.1109/ICPADS.2016.0137