Conference in Numerical Analysis 2014 (NumAn 2014)
September 2-5, 2014
Chania,Greece

Exploring the Performance of Out-of-Core Linear Algebra Algorithms in Flash based Storage[1]

Athanasios Fevgas, Panagiota Tsompanopoulou and Panayiotis Bozanis
Department of Electrical and Computer Engineering, University of Thesssaly,
Volos, Thessaly, Greece
fevgas@inf.uth.gr,yota@inf.uth.gr,pbozanis@inf.uth.gr

Abstract

In the recent years, flash memory has been widely utilized as storage medium to mobile and embedded systems, laptops and servers. The outstanding efficiency of flash based storages motivated us to study the performance of out-of-core linear algebra algorithms in flash SSDs.
Flash memory is a non-volatile electronic storage that can be electrically erased and reprogrammed. There are two types of flash, NOR and NAND with the later utilized as mass storage medium. In the rest of this document the term flash denotes the NAND flash. Storages based on flash lack of mechanical and moving parts, providing low power consumption, shock resistance and high read/write performance. Flash consists of cells which store one or more bits. Cells are organized to pages and pages to blocks. Reads and writes are performed at page level, while erases at block level. Write operations are slower than reads and erases are even slower. Moreover, pages have to be erased before are re-written and flash endurance is limited by a finite number of write/erase cycles (wear out). Solid state drives (SSDs) are block devices compatible with traditional hard disk drives (HDDs) relying in flash memories. The main components of an SSD are the flash memory chips and a controller which emulates the block interface using FTL (Flash Translation Layer). FTL remaps logical addresses, used by the upper layers, to physical addresses in flash chips. It incorporates out-of-place-updates, wear leveling and garbage collection mechanisms aiming to improve write performance and prevent wear out. All the mentioned flash characteristics make data structures and algorithms designed for hard disks not performing well in it. Many recent studies, mostly in databases, aim to design new approaches suitable for flash. Some of them are using deltas instead of performing expensive page rewrites while others deferring operations in the future in order to reduce random writes.
The development of efficient external memory (out-of-core) algorithms for solving linear equations systems or calculating eigenvalues of large matrices has been a popular research topic. Several algorithms have been proposed aiming to accelerate calculations by efficiently partitioning and managing large disk-resident datasets (matrices) into main memory blocks (sub-matrices). Alternative approaches require clusters with distributed memory, large enough for the entire dataset, and high bandwidth interconnections. Nowadays, the emergence of multi-processor, multi-core and GPU accelerated computers provides high processing power at low cost. On the other hand, flash storages are capable to accelerate the storage layer. Considering the specifics of the flash memory, we present a study of the performance of few out-of-core algorithms for numerical linear algebra problems in flash based storages.

Key words: Out-of-Core algorithms, linear algebra, scientific data, flash memory, SSD

[1]The present research work has been co-financed by the European Union (European Social Fund ESF) and Greek national funds through the Operational Program *Education and Lifelong Learning* of the National Strategic Reference Framework (NSRF) - Research Funding Program: THALIS. Investing in knowledge society through the European Social Fund.