Energy-efficient Databases using Solid State Disks

John M. Ford – Advisor: Dr. Junping Sun

Introduction

- Large-scale database users spend an increasing amount of money on powering and cooling their data centers.
- Solid state disks promise to reduce energy consumption of large database systems
- They may be inserted in various places in the memory heirarchy
- They require new algorithms that reduce write and erase operations for maximum benefit to be realized
  - Multi-buffer Manager for in-page logging databases
  - Virtual Memory page managers.

Research Direction

Efficient Buffer Management Algorithms

- What Changes Are Needed When Using Solid-state Disks?
- What Changes Will Improve Energy-efficiency

Relative Energy Consumption of NAND Flash operations

Typical Energy Consumption and Access Time of NAND Flash Memory (Park et al. 2006)

	Read	Write	Erase
Energy Consumption (μJ/4kb)	9.4	59.6	16.5
Access Time (μs/4kb)	284.2	1833.0	499.2

It is critical to reduce write and erase operations to reduce power consumption

A New Memory Hierarchy?

- Processor Registers
- Cache Memory
- Main Memory (Local)
- Main Memory (NUMA)
- Solid-state Disk
- Disk
- Tape

Adding the SSD into the memory hierarchy instead of stuffing it into the disk subsystem?

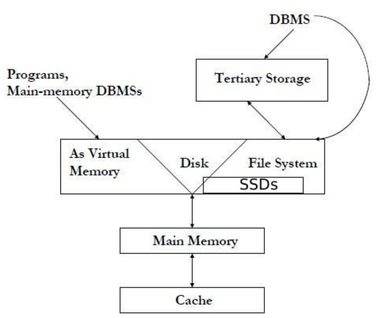

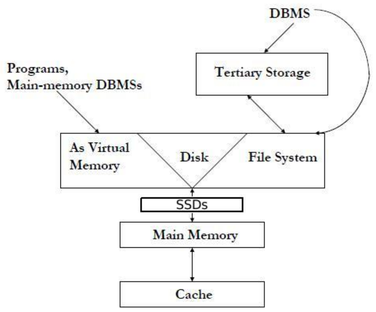

An approach to study buffer management for energy-efficient databases

- Search the literature for prior work on energy-efficient database systems
- Search the literature for research in low-power devices using Flash
  - Phones
  - Music Players
- Examine algorithms developed for buffer management in low-power devices
- Determine applicability to database management systems

Multi-buffer Manager (Cesana and He, 2010)

- Not designed specifically for SSD systems
- Paper describes an implementation for an in-page logging database
- Can significantly outperform traditional disk-based system in Flash
- Experiments found the page replacement policy plays a pivotal role in performance
  - Least Recently Used (LRU) is the basis for comparison
  - Clean First Least Recently Used (CFLRU) provides a decrease in energy required over LRU policy
  - This work (Multi-buffer Manager) provides a 40% decrease in energy consumption over the CFLRU algorithm

Cold-Clean First LRU (Li et al., 2009)

- Builds on the CCLRU of Park et al, (2006)
- Takes into account whether the page is being actively used or not
- "Overwhelms" competing algorithms on various workloads
  - Used a simulation Flash-DBSim
  - Used workloads with varying mix of read/write operations

Table II from Li et al. (2009) The synthesized workload

Type	Total references	Different pages accessed	Read /Write Ratio	Locality
T9182	300,000	10,000	90% / 10%	80% / 20%
T5582	300,000	10,000	50% / 50%	80% / 20%
T1982	300,000	10,000	10% / 90%	80% / 20%
T9155	300,000	10,000	90% / 10%	50% / 50%
T5555	300,000	10,000	50% / 50%	50% / 50%
T1955	300,000	10,000	10% / 90%	50% / 50%

Summary and References

Summary

- A lot of research being done on low-power buffer management for portable electronics as well as for green computing
- Some specialized Flash databases have been developed
- Work to integrate these buffer management algorithms into large-scale databases is still an area of active research

References

[1] Ulpian Cesana and Zhen He.
Multi-buffer manager: Energy-efficient buffer manager for databases on flash memory.
ACM Trans. Embed. Comput. Syst., 9(3):28:1–28:36, March 2010.

[2] Zhi Li, Peiquan Jin, Xuan Su, Kai Cui, and Lihua Yue.
CCF-LRU: a new buffer replacement algorithm for flash memory.
IEEE Trans. on Consum. Electron., 55(3):1351–1359, August 2009.

[3] Xiaofeng Meng, Peiquan Jin, Wei Cao, and Lihua Yue.
Report on the first international workshop on flash-based database systems (FlashDB 2011).
SIGMOD Rec., 40(2):40–44, September 2011.

[4] Zichen Xu.
Building a power-aware database management system.
In *Proceedings of the Fourth SIGMOD PhD Workshop on Innovative Database Research*, IDAR '10, pages 1–6, New York, NY, USA, 2010. ACM.

[5] Zichen Xu, Yi-Cheng Tu, and Xiaorui Wang.
Exploring power-performance tradeoffs in database systems.
In *Data Engineering (ICDE), 2010 IEEE 26th International Conference on*, pages 485–496, March 2010.

CCF-LRU Results, Hit Ratios (Li et al., 2009)

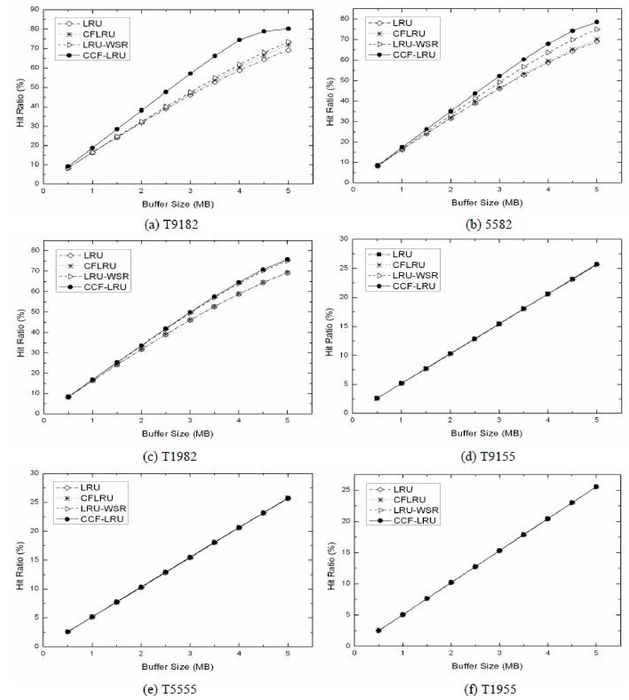

CCF-LRU Results, Total Writes (Li et al., 2009)