Write locality and its optimization for persistent memory[*]

Dong Chen, Chencheng Ye, Chen Ding
Department of Computer Science, University of Rochester
dchen39@cs.rochester.edu, cye@cs.rochester.edu, cding@cs.rochester.edu

Motivation

DRAM vs NVM

DRAM: faster but smaller and higher energy
NVM: larger, lower energy but limited write endurance

Hybrid memory architecture

(1) hierarchical model
(2) side-by-side model

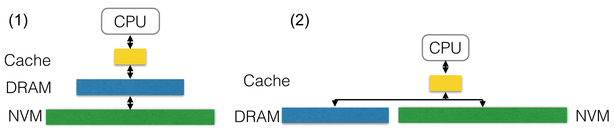

Leverage cache to minimize writes to NVM

Traditional locality theory

Reuse distance(rd)

Def: The number of distinct memory accesses between use and reuse.
if rd > cache size, then reuse is a cache miss

reference time	1	2	3	4	5	6	7	8	9	10
address	a	b	c	a	b	c	c	c	a	b

Footprint(fp)[1]

Def: The averaged Working Set Size (WSS) for some specific window size.
if fp(rt) > cache size, reuse is a cache miss

reference time	1	2	3	4	5	6	7	8	9	10	
address	a	b	c	a	b	c	c	c	a	b	
WWS(1)	1	1	1	1	1	1	1	1	1	1	fp(1) = 1
WWS(2)		2	2	2	2	2	1	1	2	2	fp(2) = 1.78
WWS(3)			3	3	3	3	2	1	2	3	fp(3) = 2.5

a ? ? a
3 > fp(3) > 2

Write locality

Write interval

Def: The time window between two *consecutive* writes to the same location. No intervening writes to that location. Reads to that location or reads/writes to other locations are allowed.

reference time	1	2	3	4	5	6	7	8	9	10
address	a	b	c	a	b	c	c	c	a	b
access type	w	r	w	r	w	r	w	w	w	w
write interval a	w	-	-	r	-	-	-	-	w	

Write reuse distance(wrd)

Def: The maximum reuse distance for w-r, r-r, r-w pairs inside each write interval.
if wrd > cache size, then the first write causes a writeback

Write reuse time(wrt)

Def: The maximum reuse time for w-r, r-r, r-w pairs inside each write interval.
if fp(wrt) > cache size, then the first write causes a writeback

Write back ratio curves for 29 SPEC CPU2006 benchmarks

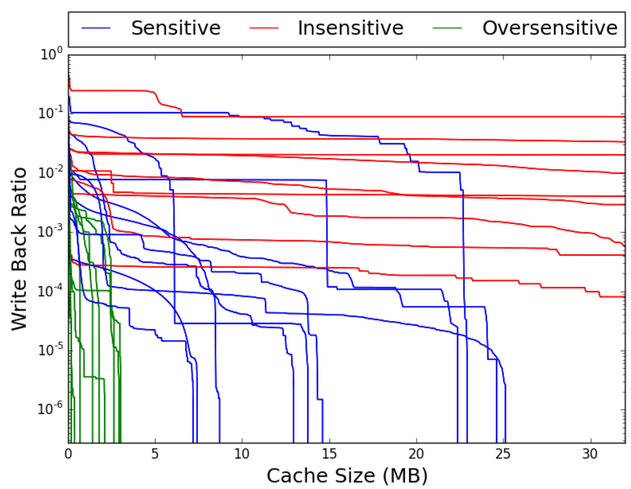

Write locality optimization

Optimal cache partition to minimize co-run writebacks

How to partition the cache with p cache ways for m co-running programs to minimize the writebacks? (For generalized partition sharing, see [2])

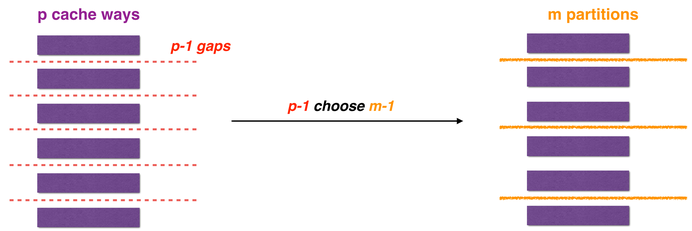

Write back reduction by optimal partitioned cache VS shared cache

elasticity: the maximum miss ratio increment for each individual program in co-run groups.

elasticity	0%	10%	100%	10x	100x	∞
2-prog.	4.09%	8.93%	9.95%	10.45%	10.8%	12.22%
3-prog.	23.97%	26.31%	26.61%	26.98%	27.21%	27.49%
4-prog.	33.93%	35.01%	35.24%	35.42%	35.45%	35.57%

Summary

Write locality is measured by write reuse distance and write reuse time and used to minimize cache write backs.

Reference

[1] "HOTL: a higher order theory of locality", Xiaoya Xiang, Chen Ding, Bin Bao, Hao Luo, The Symposium on Architectural Support for Programming Languages and Operating Systems, 2013.
[2] "Optimal cache partition-sharing", Brock, Jacob, Chencheng Ye, Chen Ding, Yechen Li, Xiaolin Wang, Yingwei Luo. International Conference on Parallel Processing, 2015.
[*] "Write locality and its optimization for persistent memory". Dong Chen, Chencheng Ye, Chen Ding. International symposium on memory systems, 2016.