

Intra-page Cache Update in SLC Mode with Partial Programming in High Density SSDs

Jun Li¹, Minjun Li¹, Zhigang Cai¹, Francois Trahay², Mohamed Wahib^{3,4}, Balazs Gerofi⁴, Zhiming Liu¹, Min Huang¹, Jianwei Liao¹

1. Southwest University of China

2. Telecom SudParis

3. National Institute of Advanced Industrial Science and Technology

4. RIKEN Center for Computational Science

- Background and motivation
- Design
- Evaluation and analysis
- Conclusion

Background(1/2)

➢ NAND Flash memory based Solid State Drives (SSDs) have been

widely applied in multiple storage devices.



High access performance, low power consumption, random access, et al.

Background(2/2)

High density SSDs reduce the price cost, becoming the mainstream storage device.

- Performance and lifetime reduce lacksquare
- SLC(Single-level cell)-mode applied in high density SSDs.
 - better performance and ۲ endurance, as the cache



Data Migration

Wear-Leveling

Address Mapping

Garbage Collection

Motivation(1/2)

- A large number of requests with small size cause space fragmentation in SSD In-page
 - Write granularity: page
- Partial programming
 - Multiple programming with different offset in the same page
 - Space utilization increases
 - Program disturbance

Eliminate the error caused by in-page program disturbance



Motivation(1/2)

- A large number of requests with small size cause space fragmentation in SSD No in-particular
 - Write granularity: page
- Partial programming
 - Multiple programming with different offset in the same page
 - Space utilization increases
 - Program disturbance

Towards update requests utilizing partial programming in the same page



Motivation(2/2)

- > The requests with 4K account for the majority
 - It is possible to use partial programming for update requests on the same page!

Trace	Update Ratio	SZ∈(0, 4K]	SZ∈(4K, 8K]	SZ>8K
ts0	78.3%	69.8%	17.9%	12.3%
wdev0	76.5%	73.2%	6.8%	20.1%
lun1	49.8%	85.2%	7.3%	7.5%
usr0	54.4%	66.3%	12.1%	21.6%
lun2	7.4%	92.6%	2.5%	4.9%
ads	7.6%	74.5%	14.1%	11.4%

> Update ratio has different degrees

• Hold hot write (frequent accessed) data in SLC cache

• Background and motivation

• Design

- Evaluation and analysis
- Conclusion



> Updated requests utilizes partial programming in the same page

- Three-level SLC-mode blocks (Work Block<Monitor Block<Hot Block)
- Hot write (frequent accessed) data "upgrade"
- Holding hot write data in SLC cache

Design(2/3)

- Traditional GC policy: Greedy
 - Traverse target Plane, and select the block having the least number of invalid page
- Proposed GC policy: subpage granularity
 - Traverse target Plane, and select the block having the least value of Invalid Subpage Ratio (ISR)
 - ISR: measure the cold degree of block.



Tradition: 0无效页

Proposal: 6 invalid subpages 4 valid subpages

GC candidate B

GC candidate A



Tradition: 0无效页

Proposal: 6 invalid subpages 4 valid subpages (including 1 cold valid subpage)

Design(3/3)

- ➢ Valid data migration
 - **DEGRADE** cold write data
 - Hold the level of other data

- Free up the SLC region occupied by cold data
 - Hold cold data in high density SSD blocks



invalid subpage

Block level (Ascending):

High Density Block < Work Block < Monitor Block < Hot Block

- Background and motivation
- Design
- Evaluation and analysis
- Conclusion

Experiment setup(1/2)

EXPERIMENTAL SETTINGS OF SSDsim

SSDSim settings	Parameters	Values	Parameters	Values (ms)
	Page size	16KB	SLC read time	0.025
	Page per block	1024	MLC read time	0.05
	SLC mode ratio	5%	ECC min decode time	0.0005
	GC threshold	0.05	ECC max decode time	0.0968
	GC policy	greedy	SLC write time	0.3
	Wear-leveling	static	MLC write time	0.9
	FTL scheme	Page	Erase time	10

Specifications on Selected Disk Traces

➢ Benchmark

Trace	# of Request	Write ratio	Write size	Hot write
ts0	1,801,734	82.4%	8.0KB	50.5%
wdev0	1,143,261	79.9%	8.2KB	58.2%
lun1	1,073,405	73.1%	7.6KB	10.0%
usr0	2,237,889	59.6%	10.3KB	36.5%
lun2	1,758,887	19.3%	9.7KB	8.5%
ads	1,532,120	9.5%	7.0KB	18.3%

Experiment setup(2/2)

- ≻Comparison methods:
 - *Baseline*: hybrid SLC/MLC SSD architecture.
 - *MGA*(<u>Mapping Granularity A</u>daptive)^[1]: utilizes subpage granularity management with partial programming, by employing two-level mapping table to record subpage information in same pages
 - *IPU*(Intra-page Update): proposed method, which performs update operations with partial programming in the same pages which previously holds the old version of data, and ejects cold write data during GC operations.

[1] Yazhi Feng, Dan Feng, Chenye Yu,Wei Tong, and Jingning Liu. Mapping granularity adaptive ftl based on flash page re-programming. In *DATE*, 2017.

Result(1/3): I/O performance



Our proposal of IPU can decrease I/O time by 11.9% and

7.3%, compared with Baseline and MGA.

Result(2/3): Read error rate



Our proposal of IPU can decrease read error rate, and thus present less read response time (ECC decode time).

Result(3/3): Three-level blocks



32.9% write are the most frequently updated data on average, and hold in SLC-mode blocks to utilize the I/O benefits.

- Background and motivation
- Design
- Evaluation and analysis
- Conclusion

Conclusion

- Intra-page update (IPU) utilizes partial programming towards update requests in the same page.
 - Mitigate program disturbance caused by partial programming.
- We further propose three-level SLC-mode blocks and degraded GC policy for SLC mode blocks.
 - Hold hot write data in SLC mode cache (block).
- Preliminary evaluation prove that, proposed method can reduce the I/O latency and read error rate, compared with the related work.



Thank you very much! Q&A

Contact: junli95@email.swu.edu.cn