

Abstract

As the controller specification of flash storage becomes luxurious, researches of page-level mapping FTL are again reviving. However, scalability has built a barrier to better performance. In this work, we have presented a design of an efficient garbage collection policy for page-level mapping FTL of large scale solid state disks. It uses block-level information to approximate an update-frequency aware data allocation policy. Its GC policy is enhanced by taking both victim utilization and invalidation recency into account. In addition, with the aid of special data structure, no manual-tuned parameters required. This works has shown competitive performance against previous works and pretty good stability under both pattern switch and large geometry environment.

Acknowledgements

I would like to express my immense gratitude to my advisor Professor, Li-Pin Chang, without whose guidance I will never be able to accomplish this work. Thank you for sparing so much of your precious family time and patience pushing me forward to a place I have never imagined.

I am also deeply appreciated the generous advices and encouragements from Professor Wei-Chung Hsu, who kindly helped me not only with my research and also my career plan again and again.

To all the respected professors that I have learned from, your selfless mentoring has completed me. Professor Shih-Kun Huang, you are a wonderful teacher enlightened both my study and my interpersonal intelligence. Professor Chin, Feng-Der, you have broadened my world that makes me able to imagine an new era of science. Professor Sheau-Ling Hsieh, you have led me through the mysterious door of networking and it is cozy that you always treat us as your family.

Also thanks to all the colleagues of ESS laboratory. For those graduated last year, Jin-Ting Kuo, Wei-Jie Huang, Yi-Hsun Huang and Li-Chun Huang, it is your passion and unselfish share of experience that helped me grow faster and always felt like at home. For Ying-Jie Li, Yi-Cheng Wu and Wei-Han Wang, it is my honor to work and learn with you, thank you a million for your constant timely help. For Wen-Ping Li, Yi-Kang Chang, Tung-Tang Chou, Po-Han Sung and Chen-Yi Wen, fuel up, it is your generation now.


Last but not least, countless thanks to my parents and my lovely friends who never stop trusting me. It is your faith in me and support that makes everything possible. Without your understanding and comfort, I would never pass those darkest nights through. Thank you, I love you.

Contents

Chapter 1	Introduction	1
Chapter 2	Background	3
2.1	page-level mapping and garbage collection	3
2.2	Data-Separation and Garbage Collection	4
2.3	Victim Selection and Garbage Collection	5
Chapter 3	Data Separation Policy	7
3.1	The Basic Data Structure for Garbage Collection	7
3.2	Identifying Hot Data and Non-hot Data	7
3.3	Page Lifetime Computation	10
4	Victim Selection Policy	12
4.1	Block Stability	12
4.2	The Dual Greedy Strategy	13
5	Experimental Results	15
5.1	Setup and Metrics	15
5.2	Host Workloads	16
5.3	Page-Level Mapping vs. Hybrid Mapping: Pattern Switch	18
5.4	Remark: The Needs for Adaptiveness	19
6	Conclusion	21
	Bibliography	23

List of Tables

Table 1.	Prior GC policy for page-level mapping FTL	6
Table 2.	The 3 experimental workloads	16
Table 3.	Comparison between DAC and Dual Greedy	18



List of Figures

2.1 Separating hot data and non-hot data in flash blocks greatly improves garbage collection performance	4
3.1 Multi-LRI list structure	8
3.2 Definition of flash page lifetime	9
3.3 The distribution of page lifetimes after processing 100,000 write requests under the disk workload of Windows XP	9
3.4 Page lifetime approximation	11
4.1 The length distribution of block stable periods after processing 100,000 write requests	13
4.2 GC victim selection under Dual Greedy Policy	14
5.1 Experimental results under (a)-(c) different host workloads and (d) different geometry settings.	17
5.2 Observation under pattern switch	19
5.3 Under the PC workload, DAC delivered sub-optimal performance with the parameter settings which are the best choices for the RND and MM workloads.	20