LX-SSD: Enhancing the Lifespan of NAND Flash-based Memory via Recycling Invalid Pages

Ke Zhou, Shaofu Hu, Ping Huang, Yuhong Zhao

Wuhan National Laboratory for Optoelectronics, Huazhong University of Science and Technology

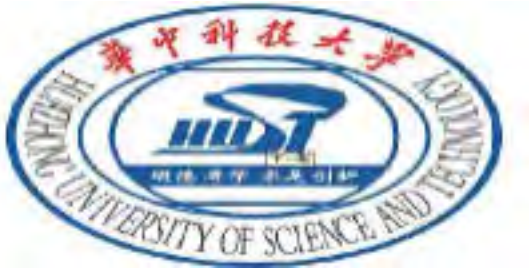

Presented by: Shaofu Hu

Email: hushaofu@hust.edu.cn

Outline

- **Backgrounds and Motivations**
- **Design of Recyclable SSD**
- **Performance Evaluation**
- **Conclusion**

Backgrounds

- **NAND flash**

  NAND Flash-based SSDs have unique merits outpacing traditional HDDs

  high-density scaling technology

	SLC	MLC	TLC
Bits per cell	1	2	3
P/E Cycles	10,000	3,000	1,000

Backgrounds

- **PE cycle**

The PE cycle of a block is composed of a series of events:

N programmings

One erase

Backgrounds

- **PE cycle**

The PE cycle of a block is composed of a series of events:

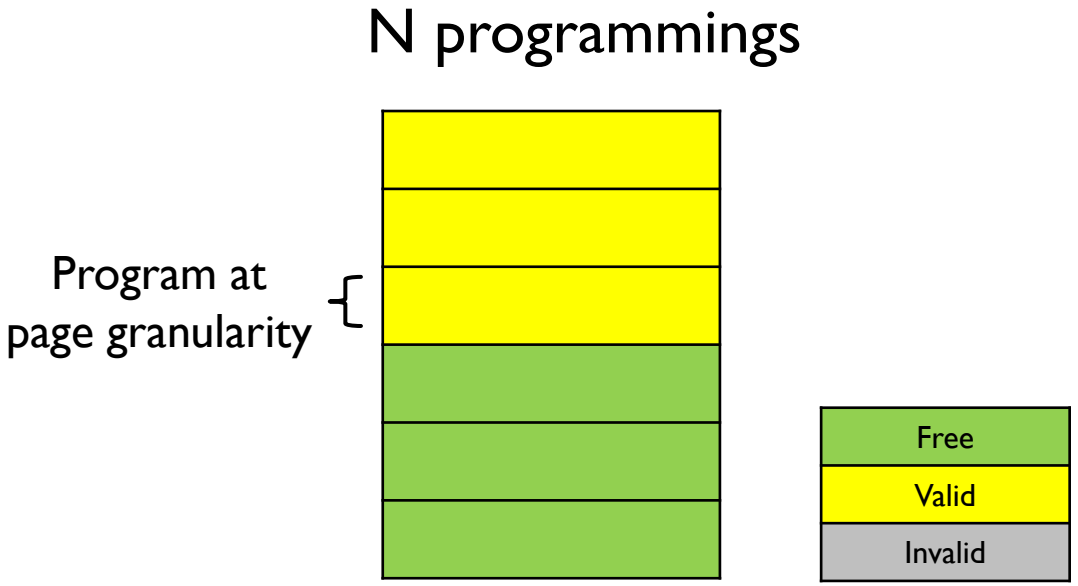

Backgrounds

- **PE cycle**

The PE cycle of a block is composed of a series of events:

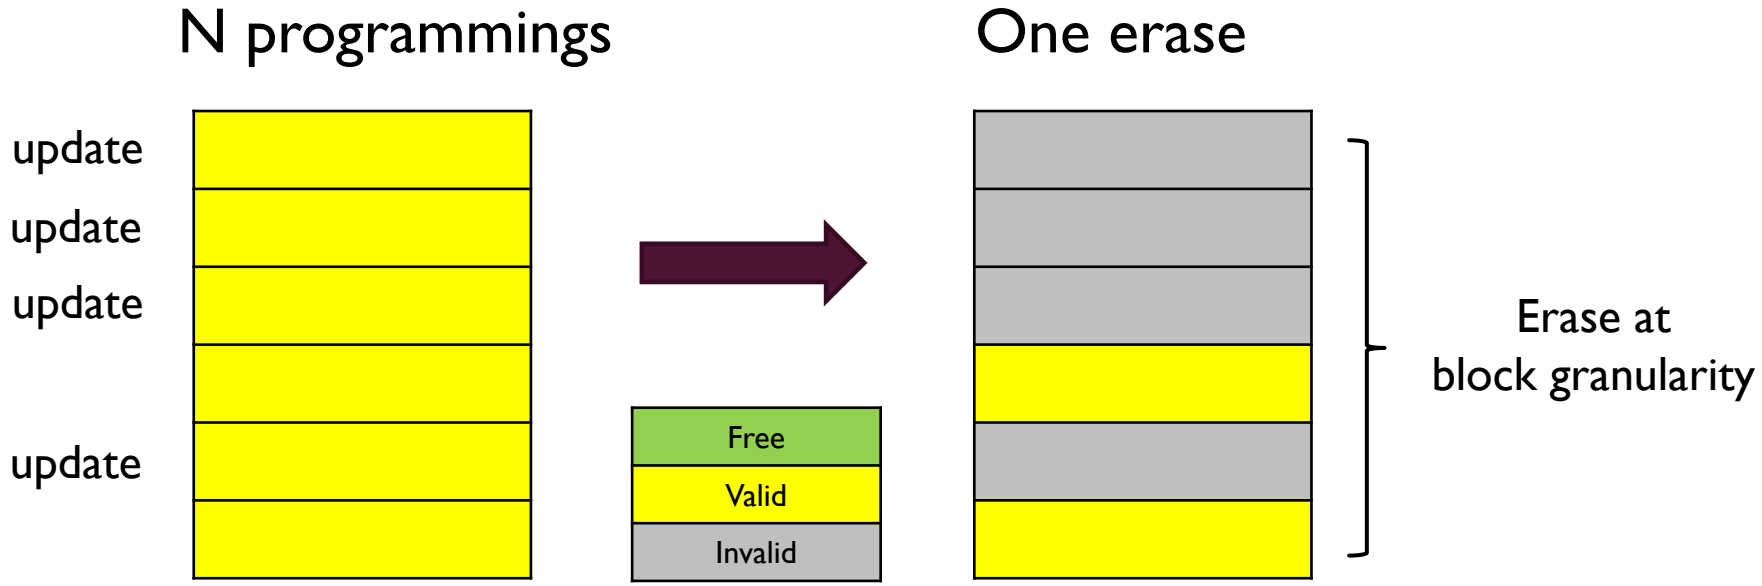

- **Block recycle**

making the entire block reusable again, but at the cost of :

(1) migrating valid pages (2) need erase operation

Backgrounds

- **Block recycle**

  all invalid pages in the block are considered to be useless

CAN NOT BE ACCESSED

BE CONSIDERED TO BE AS TRASH

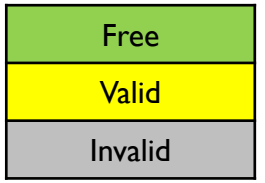

Are them truly useless?

Backgrounds

- **Page recycle** recycle ⟺ reusable

Block recycle

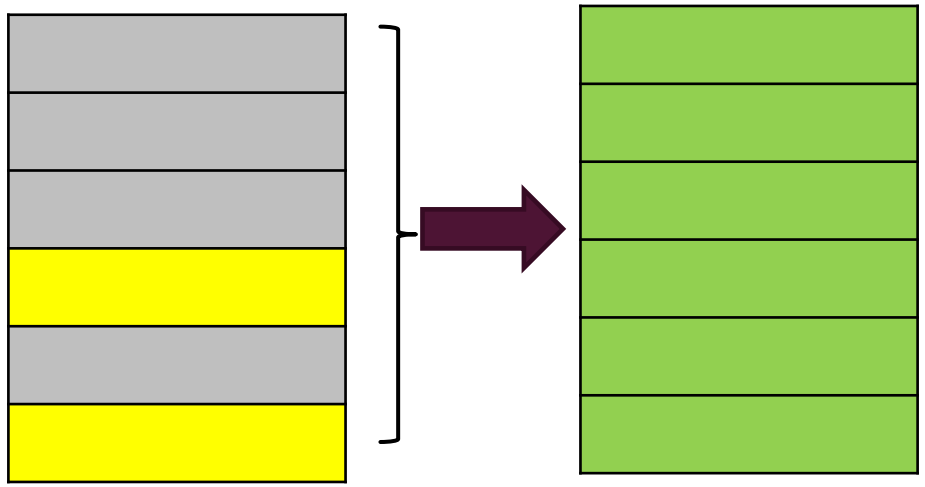

Page recycle

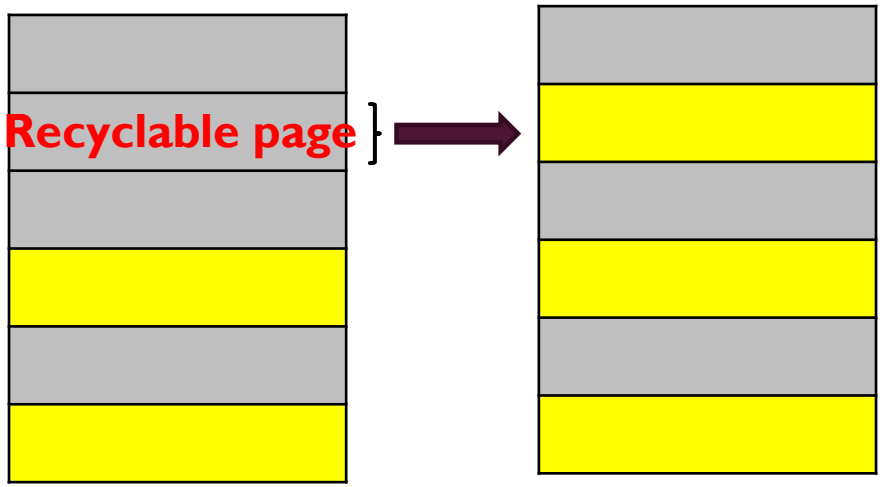

- **Recyclable page:** an invalid page shares the identical content with an incoming write request

Backgrounds

- **Page recycle**

  - recycle the invalid pages
  - act as a supplement to block recycle
  - extend the PE cycle from N-programs-per-erase to N^+-programs-per-erase

Motivations

- **Invalid pages are common**

  A NAND flash page alternates between three states: ***free***, ***valid***, and ***invalid***

  free pages only occupy limited space, typically 5%-15%

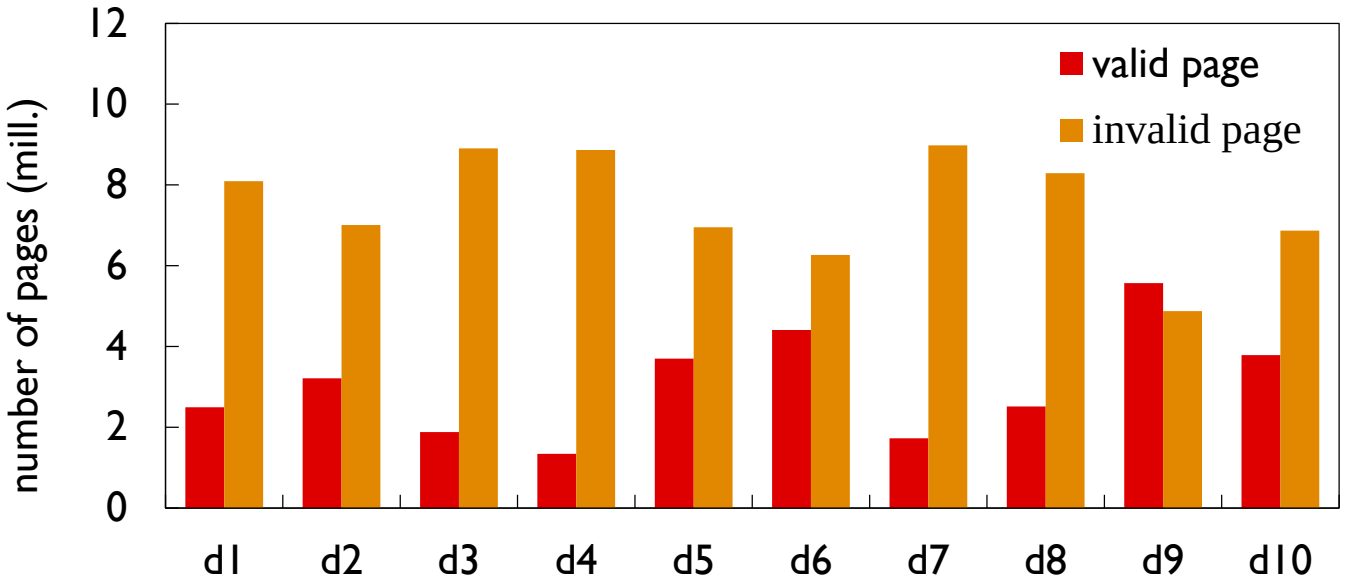

the number of invalid pages is **1-4 times** the number of valid pages

Motivations

- **Invalid pages could be useful**

  three workloads, including ***home***, ***mail,*** and ***web***, from FIU, are introduced to observe the appearance of useful data in invalid pages

Workload	Request (mill.)	Size (GB)	Writes (%)	Unique addr (mill.)
homes	9.3	3.7	93.8	2.4
mail	16.3	4.8	75.3	1.2
web	5.9	2.1	82.5	0.8

- **Recyclable page**

  - 24.1 GB, 42.2 GB, and 15.5 GB invalid pages are observed in *home*, *mail,* and *web* respectively
  - 2.8 GB, 7.4 GB, and 2.5 GB recyclable pages
  - 11.5%-17.6% of invalid pages are recyclable

Motivations

- ***Locality Characterization***

We characterize two types of value locality to effectively assist us in dividing the invalid pages into highly recyclable and less recyclable:

- value popularity

  a characterization on frequency representing how many times an invalid page has been accessed

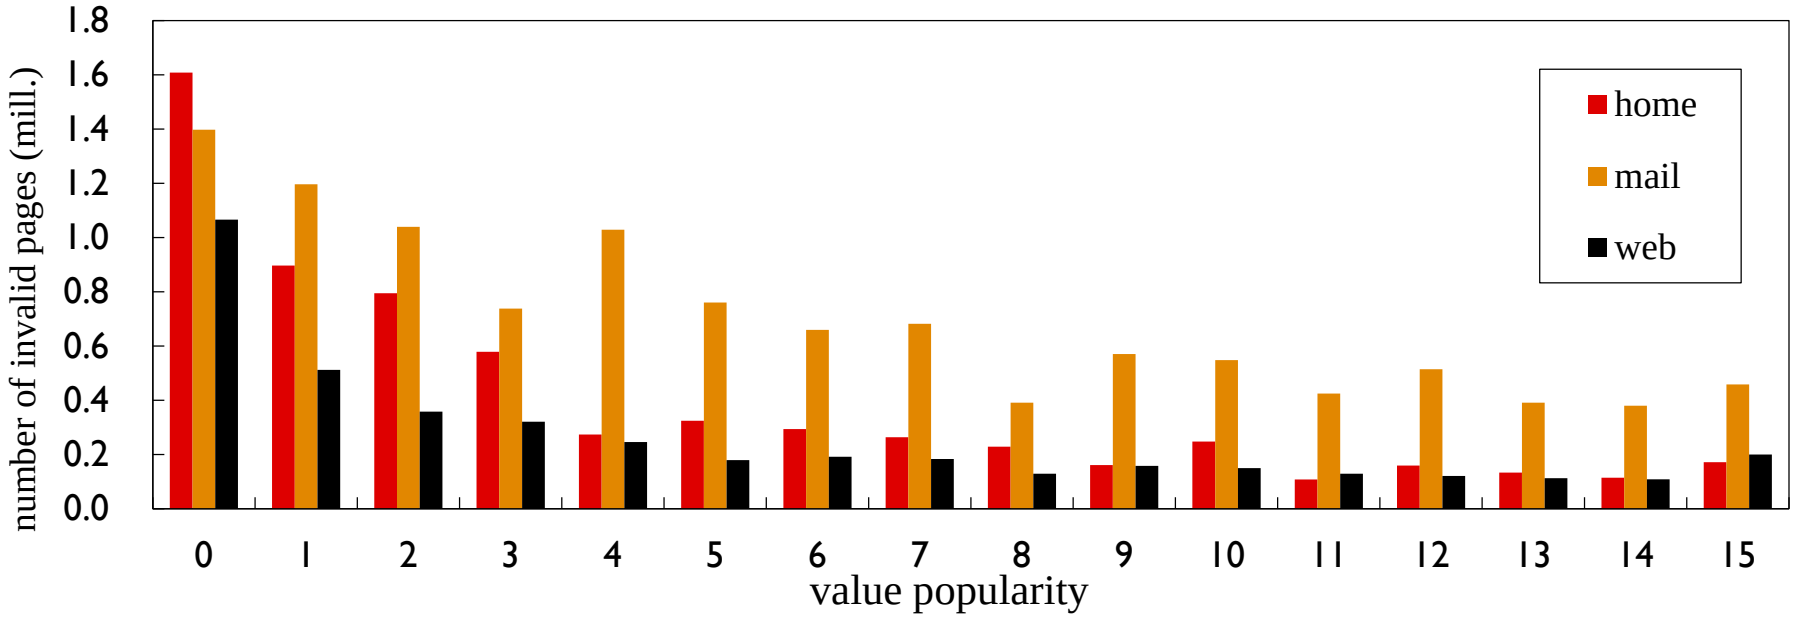

Motivations

- ***Locality Characterization***

We characterize two types of value locality to effectively assist us in dividing the invalid pages into highly recyclable and less recyclable:

  - value popularity

recyclable probability is closely related to the value popularity

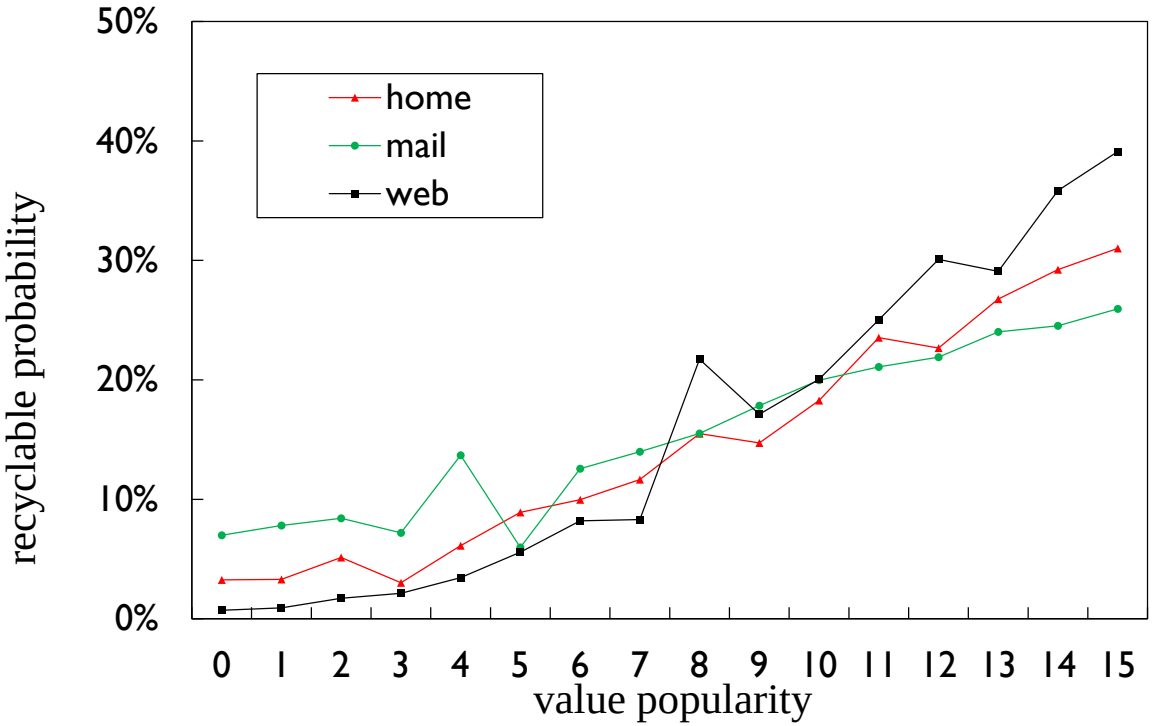

Motivations

- ***Locality Characterization***
  - Value version

a characterization on date representing how long an invalid page has been generated

value version also affects the recycle probability

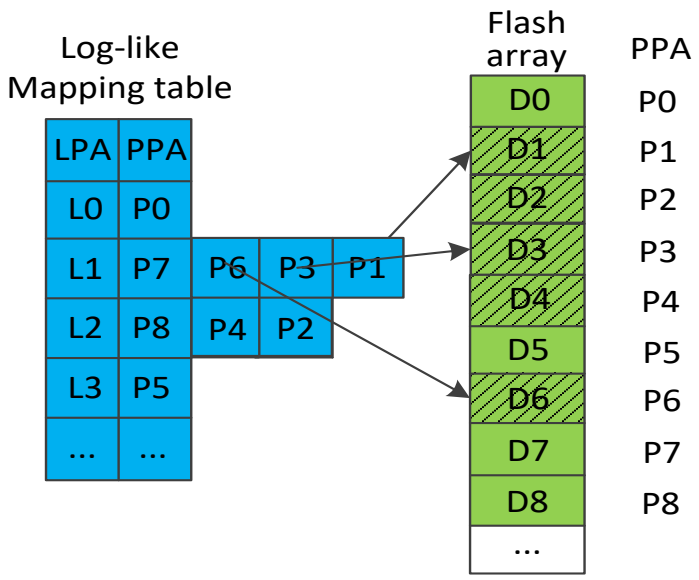

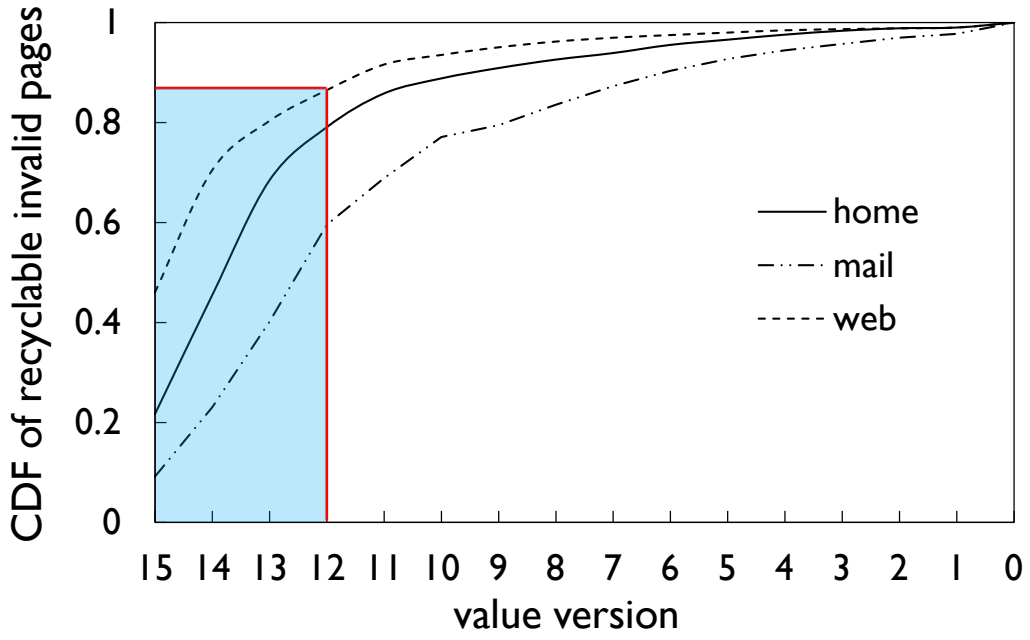

Challenges

- **Limited RAM capacity**

  Designing a content-aware FTL needs extra metadata which puts additional pressures on the limited RAM

- **Without semantic hints**

  more difficult to identify recyclable invalid pages, and requires more complicated design

- **Minimize the overhead**

  Recyclable SSD still need to retain high data access performance even when there is little locality among invalid pages.

Design of Recyclable SSD

Architecture Overview

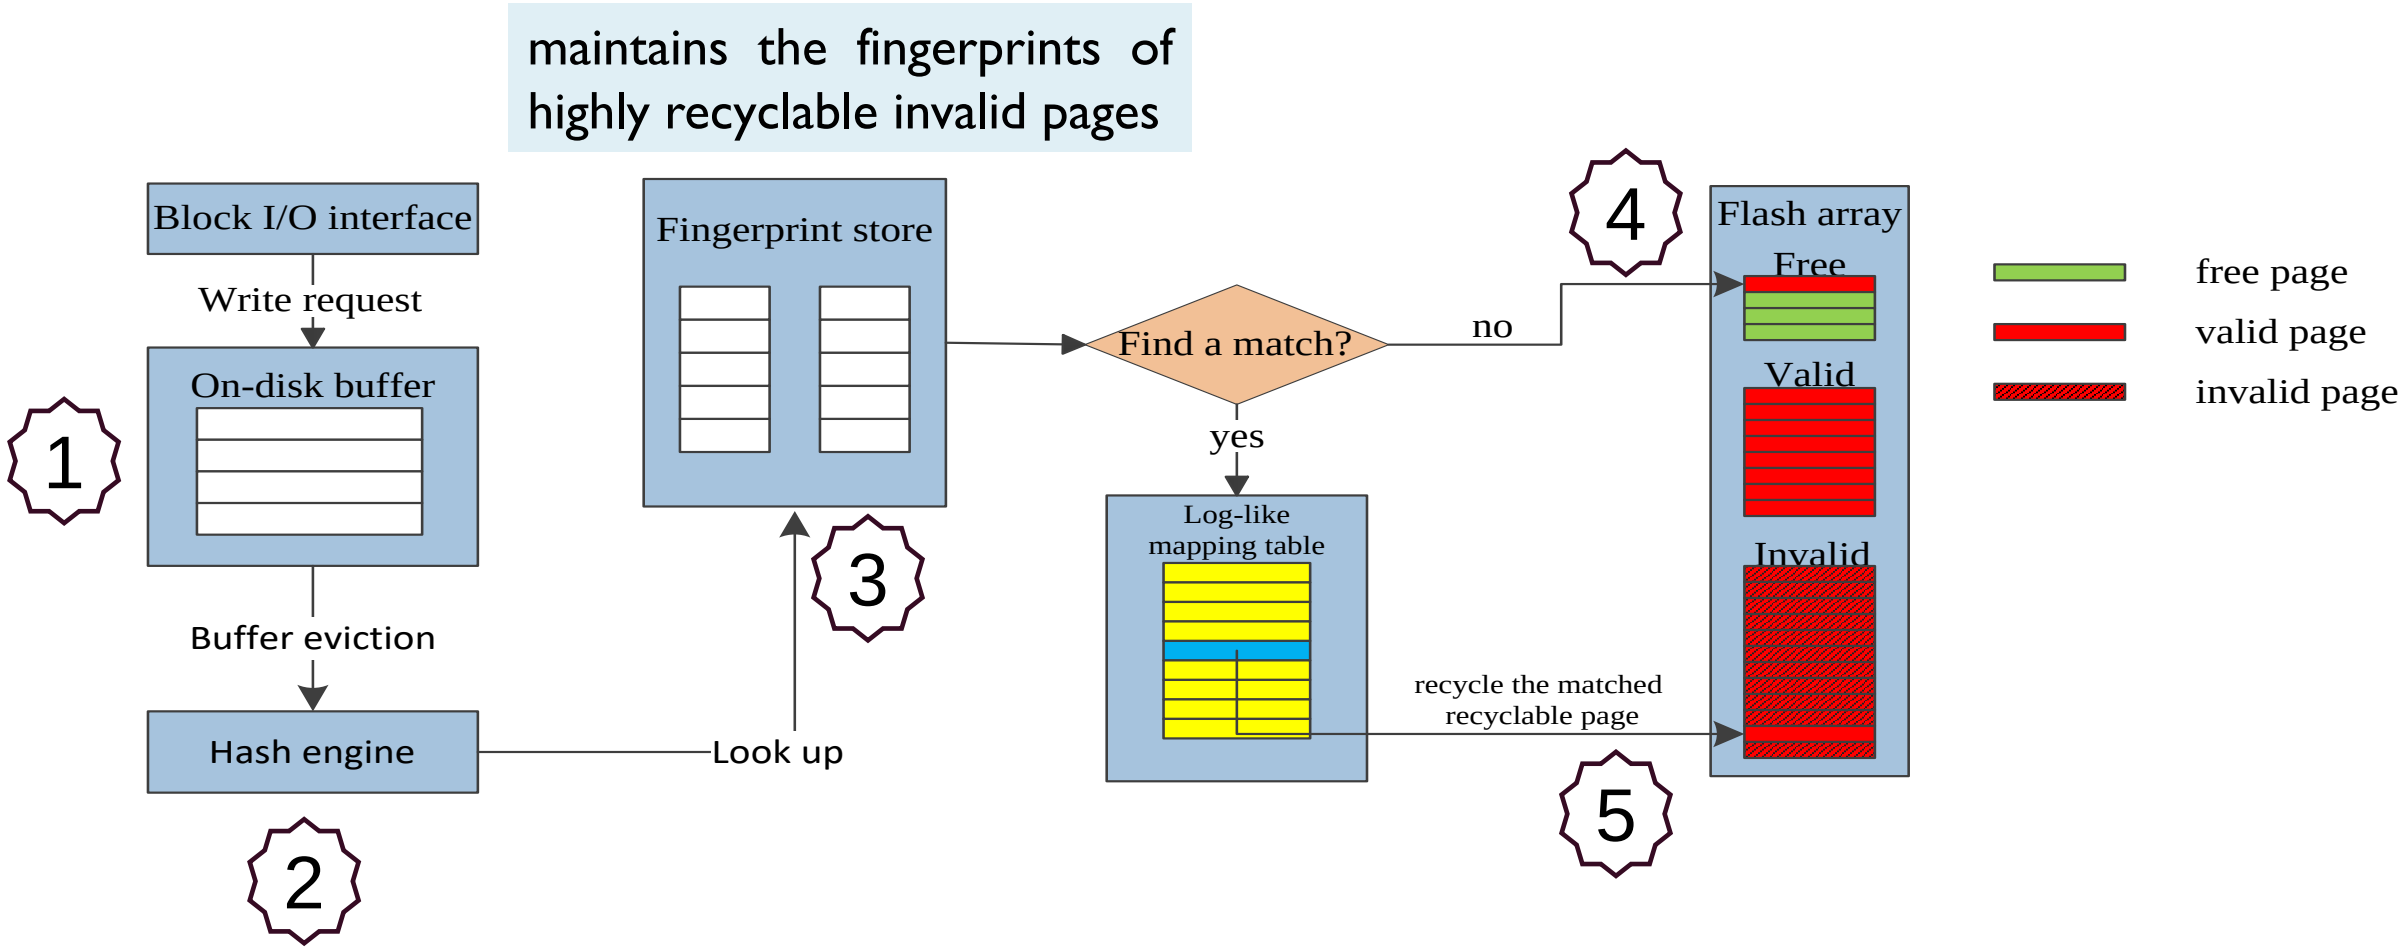

Design of Recyclable SSD

- **Extra components**

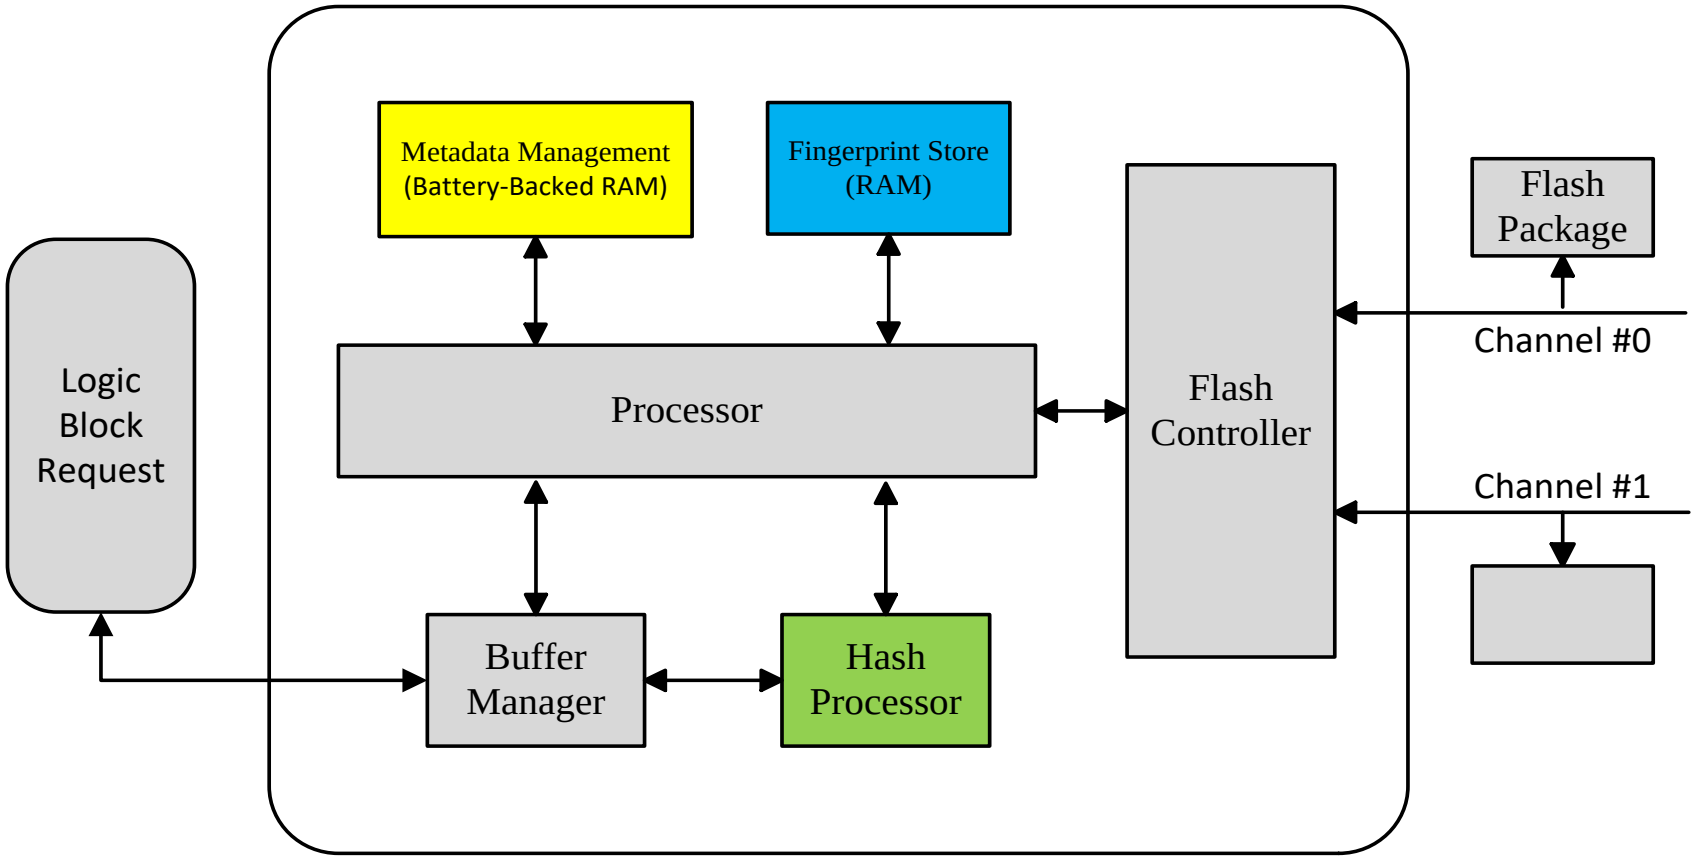

Design of Recyclable SSD

- **Metadata management**
  - Log-like mapping table (LMT)

The LMT uses a linked list to store all the physical page addresses that have ever been mapped to the same logical page address.

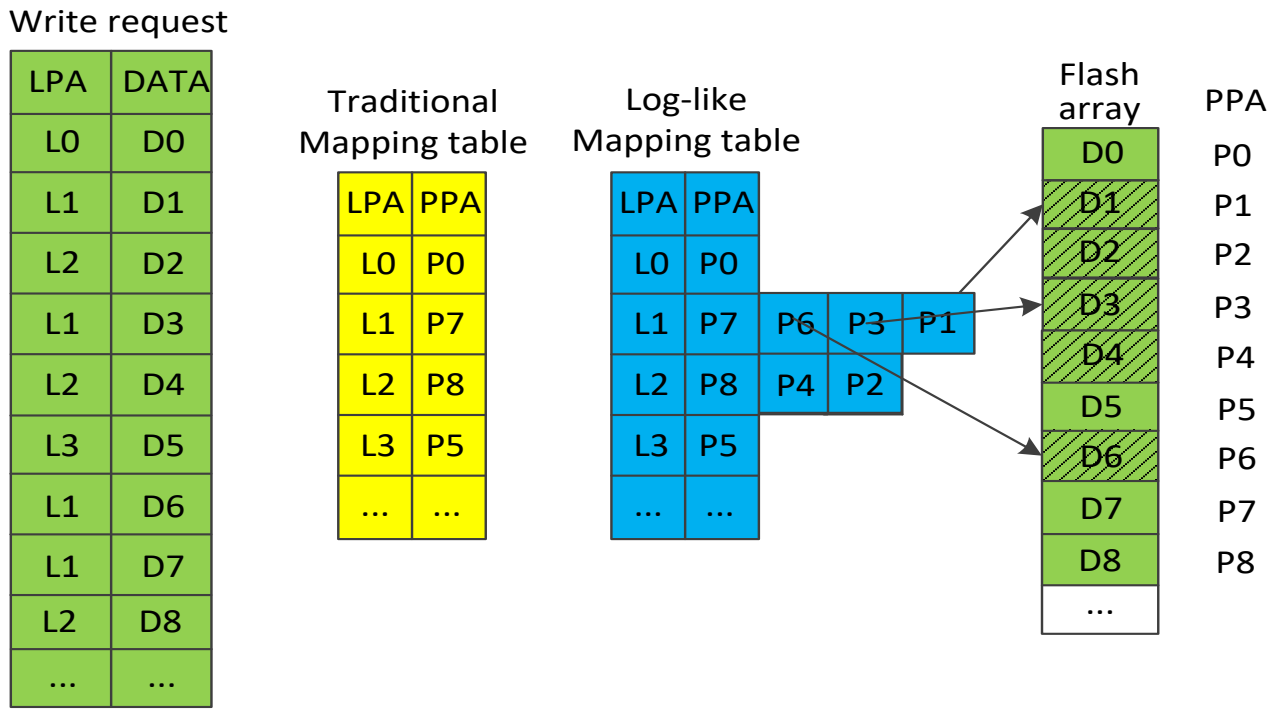

identify the value version

Design of Recyclable SSD

- **Metadata management**
  - Accessing table

record the number of times each physical page has been accessed

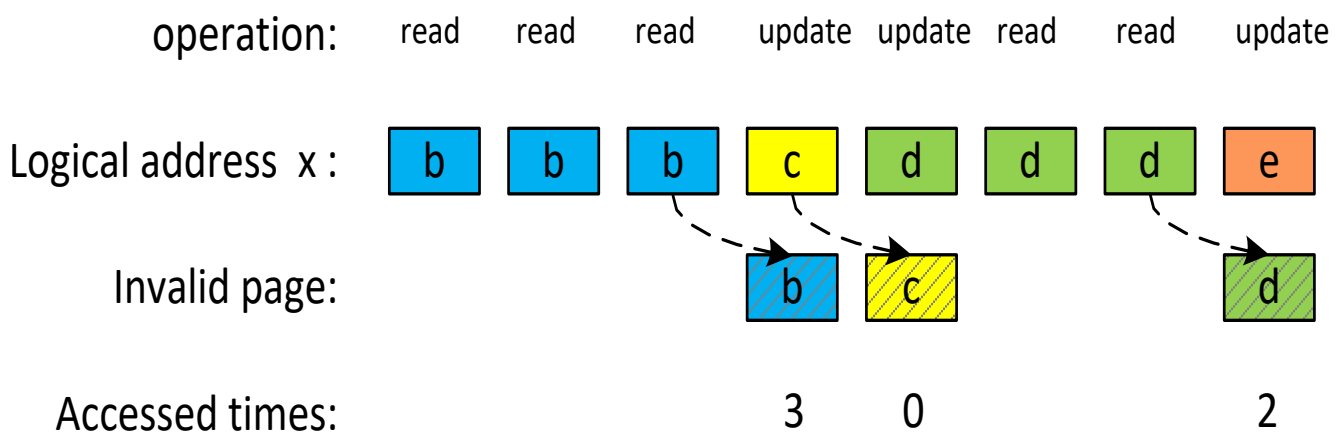

identify the value popularity

Design of Recyclable SSD

- **Fingerprint Store management**
  - Recyclable probability calculation

$$pop_{inv}(p) = (1+\alpha)^{ace_{inv}(p)}$$

α is a growth factor

in the range of 0.1 to 0.15

$$ReProb_{ini}(p) = \frac{pop_{inv}(p) - 1}{(1+\alpha)^{15} - 1}$$

$$ReProb_{fin}(p) = ReProb_{fin}(p) * (1-\beta)^{(15-ver_{inv}(p))}$$

β is a decay factor

in the range 0.15 to 0.2

Design of Recyclable SSD

- **Fingerprint Store management**
  - Fast loading

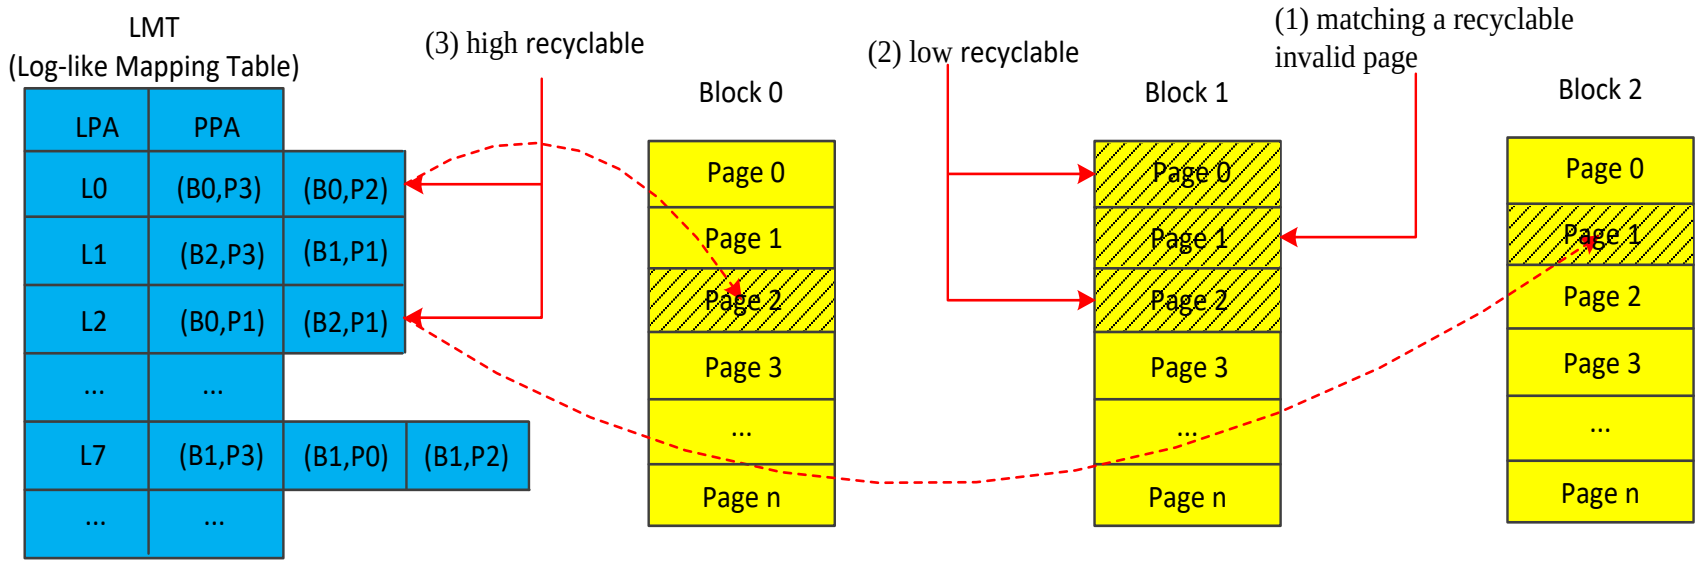

we use LMT to quickly locate those invalid pages which have spatial locality

Design of Recyclable SSD

- **Garbage collection**

  ReCyclable-aware GC

$$BlkStats(b) = \sum_{i=1}^{n} PageStats(i) * \left(1 - ReProb_{fin}(i)\right)$$

Evaluation

- **Metrics**

  Write reduction

  Number of block erasures

  Average response time

- **Setup**

  We set simulated SSD with the default configurations from the SSD extension

- **Datasets**

  home, *mail*, and *web*, from FIU

Evaluation

- **Write reduction**

With a moderate 3MB RAM size, the Recyclable SSD can achieve 10.5%-18.4% write reduction in three workloads

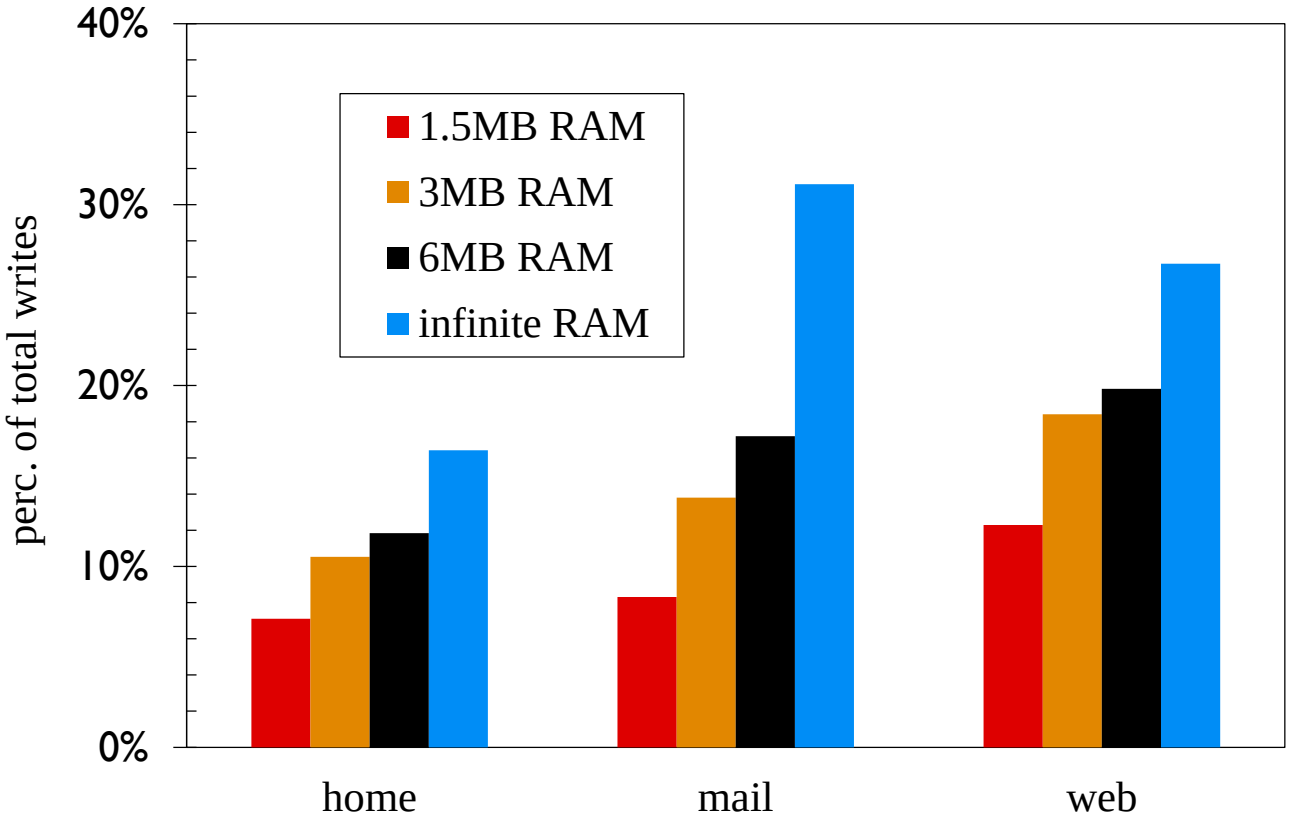

RAM sizes of 1.5MB, 2MB, and 6MB maintaining 64K, 128K, and 256K entries

Evaluation

- **Number of block erasures**

the *web* has exhibited the best results: the normalized number of block erasures under FS of 64K, 128K, and 256K are 45.2%, 41.1%, and 39.6% respectively

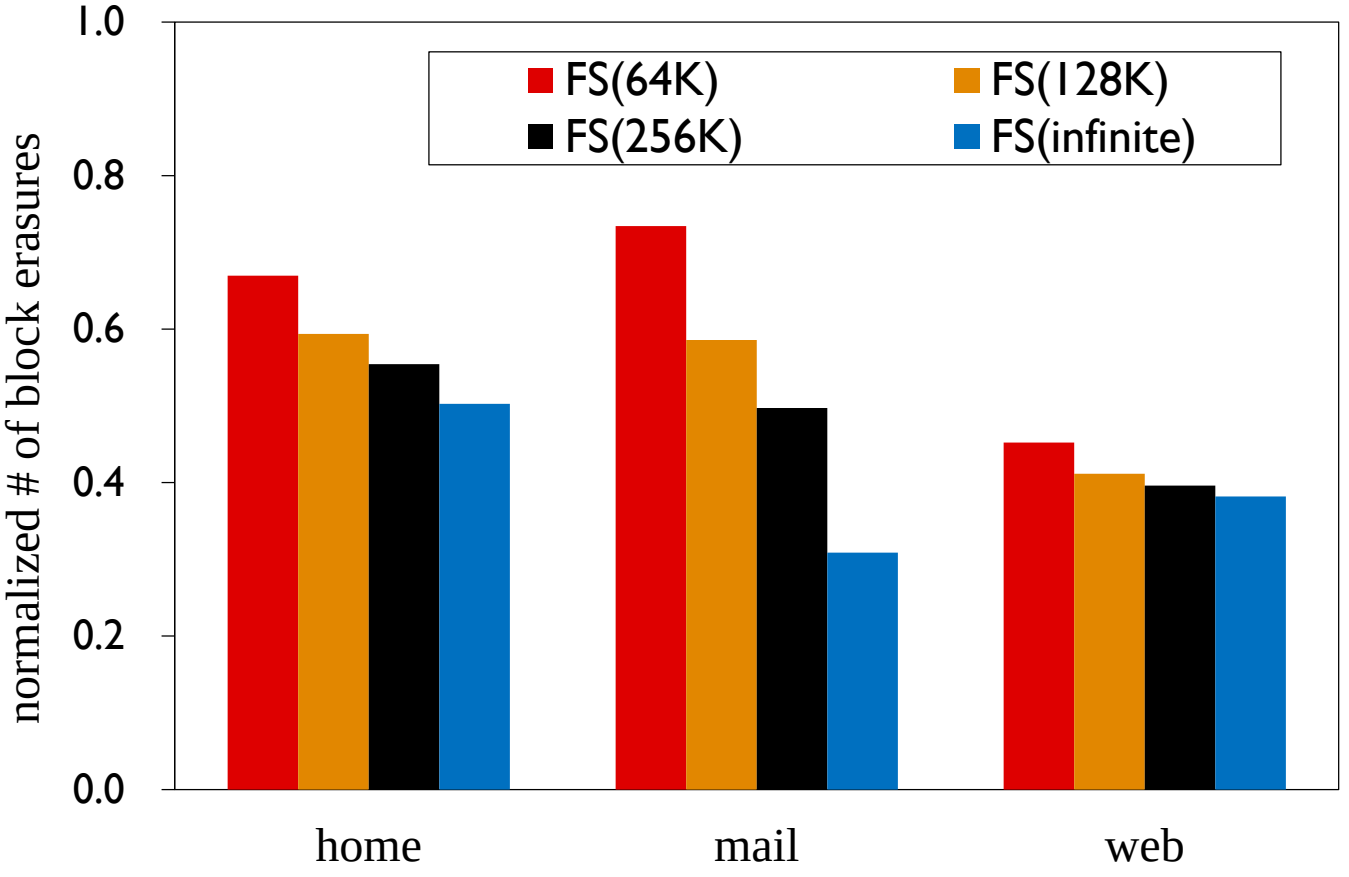

Evaluation

Average response time

the average response time of the Recyclable SSD also benefit a lot from write and block erasure reductions

The *web* workload, which has the most value locality, shows a 29.7% improvements in response times

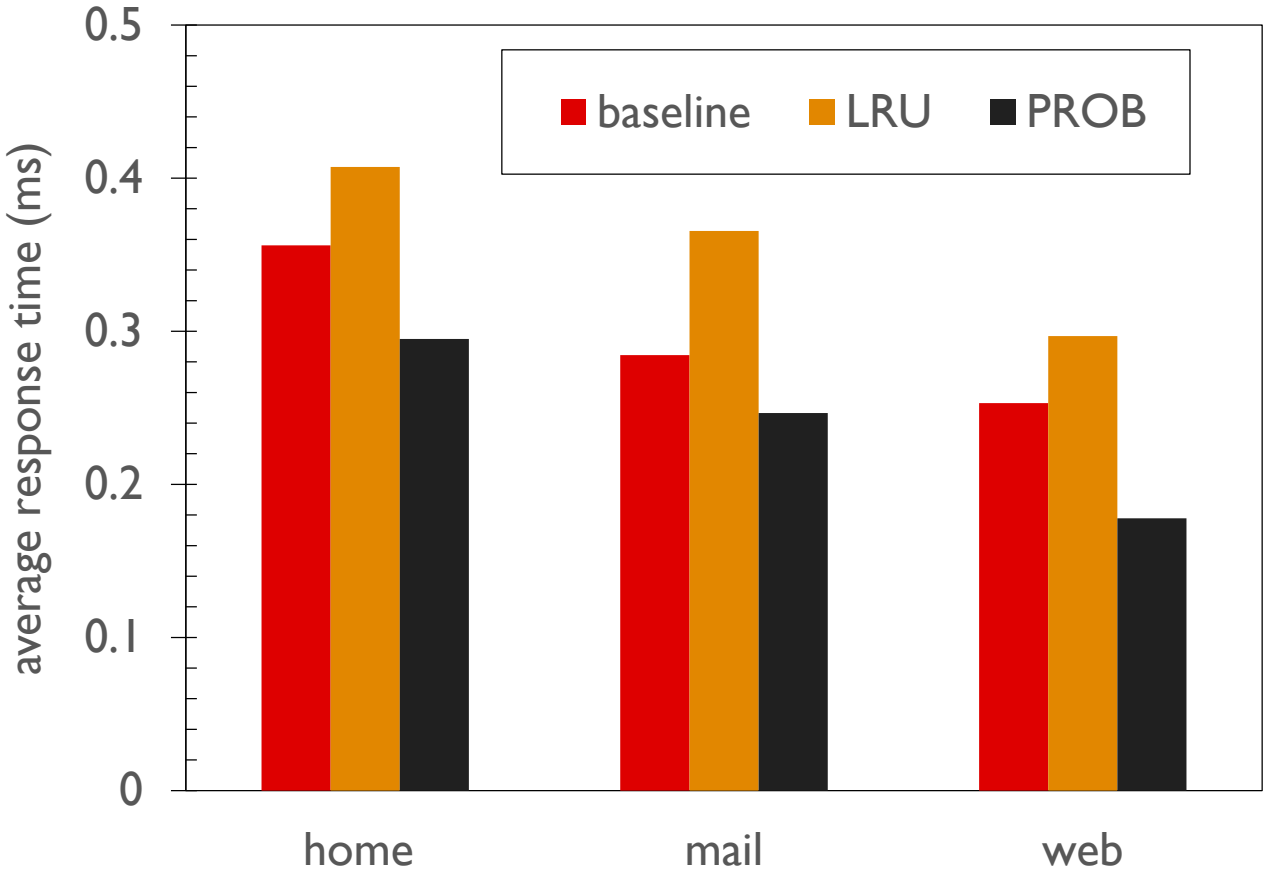

LRU
manages the FS as a queue with LRU eviction policy
PROB
maintains the FS as an ordered queue sorted by the recycle probability

Conclusion

- **Our study is the first attempt to address the benefit from recycling invalid pages without reprogramming or cleaning**
  - Recyclable SSD can extend SSD lifetime while slightly improving performance
  - The characterized localities ensuring that the implementation of the Recyclable SSD is feasible
  - the number of block erasures can be decreased resulting in significant lifetime extension

- **Most related works for extending SSD lifetime are orthogonal to the design of the Recyclable SSD and can be augmented with recycling invalid pages.**

Thank you!

Questions?

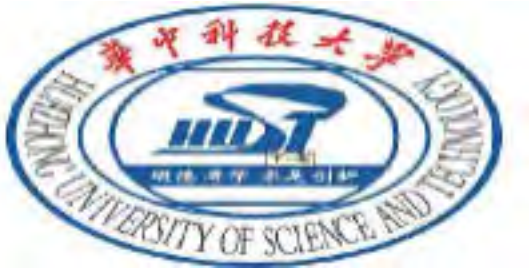

Email: hushaofu@hust.edu.cn