WARCIP: Write Amplification Reduction by Clustering I/O Pages

Jing Yang, Shuyi Pei and Qing Yang
Dept. of Electrical, Computer, and Biomedical Engineering, University of Rhode Island
Kingston, RI, USA 02881
Shenzhen DAPU Microelectronics Co., Ltd
Shenzhen, China
{**jyang**,spei,qyang}@ele.uri.edu

Garbage Collection Overheads

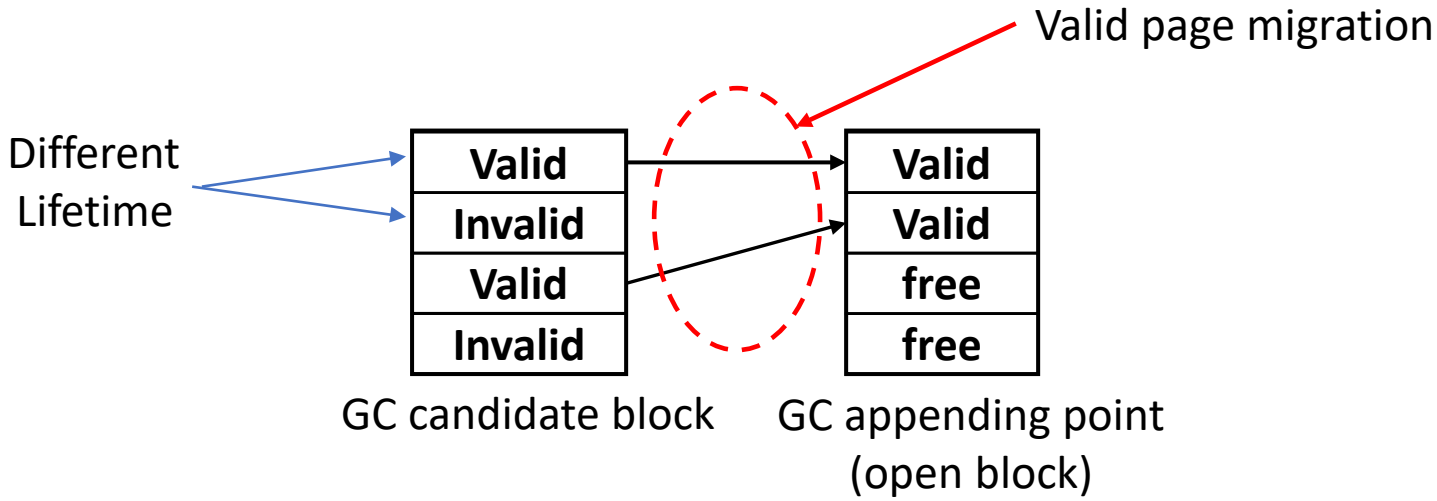

- Drags down I/O performance
  - Consumes a significant amount of back-end bandwidth of SSDs
- Reduces SSD's endurance
  - Valid page migration causes extra writes to NAND flashes

Garbage Collection Overheads (cont.)

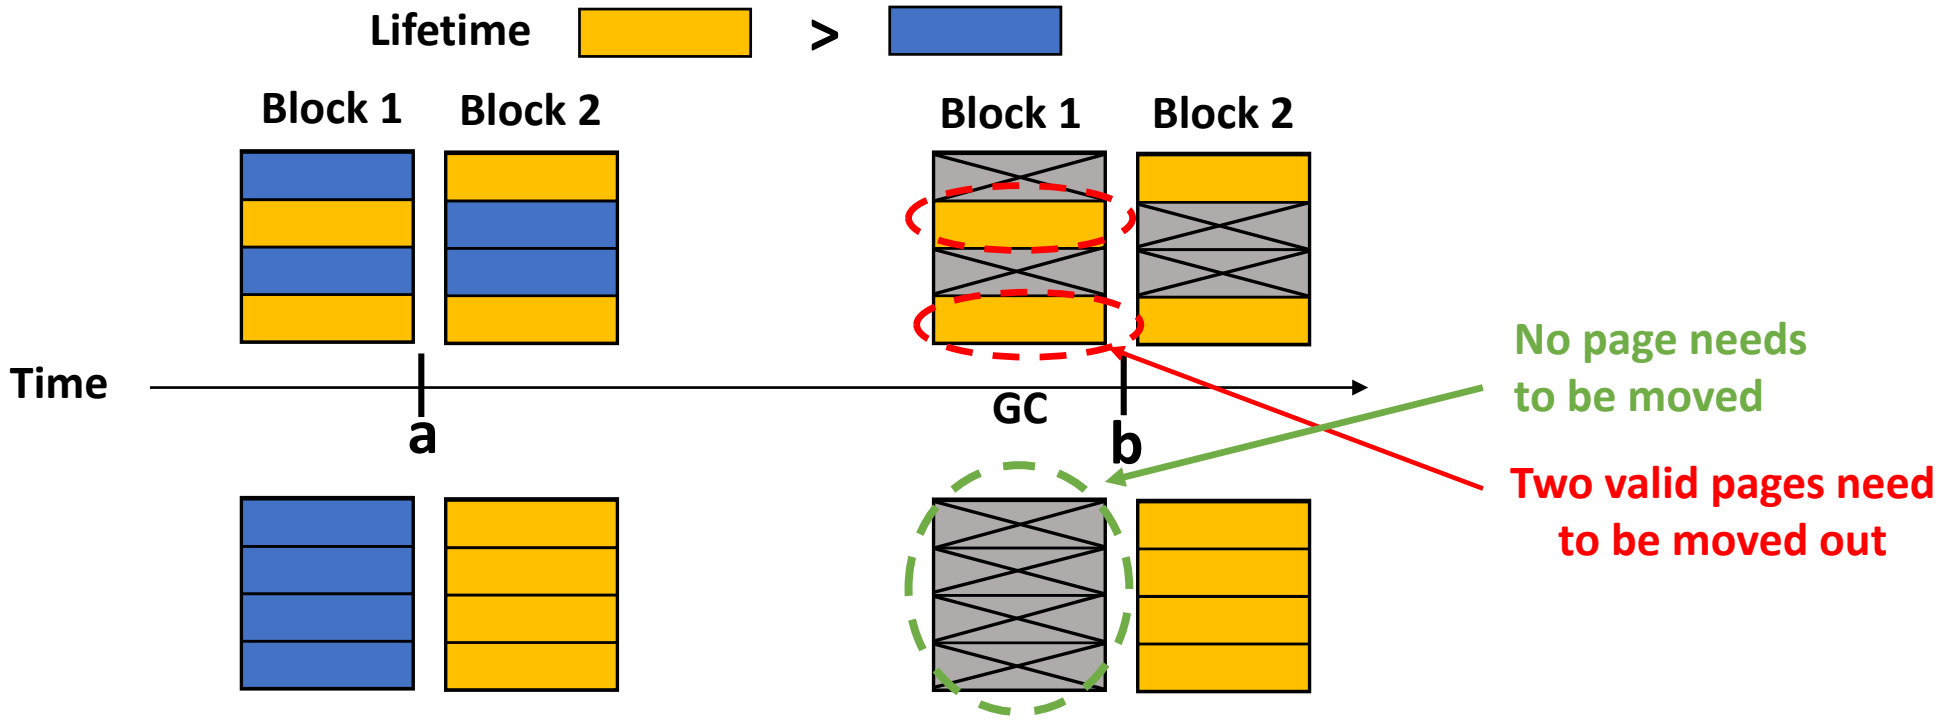

WA can be reduced by minimizing lifetime variance of pages in a flash block.

Existing Solutions

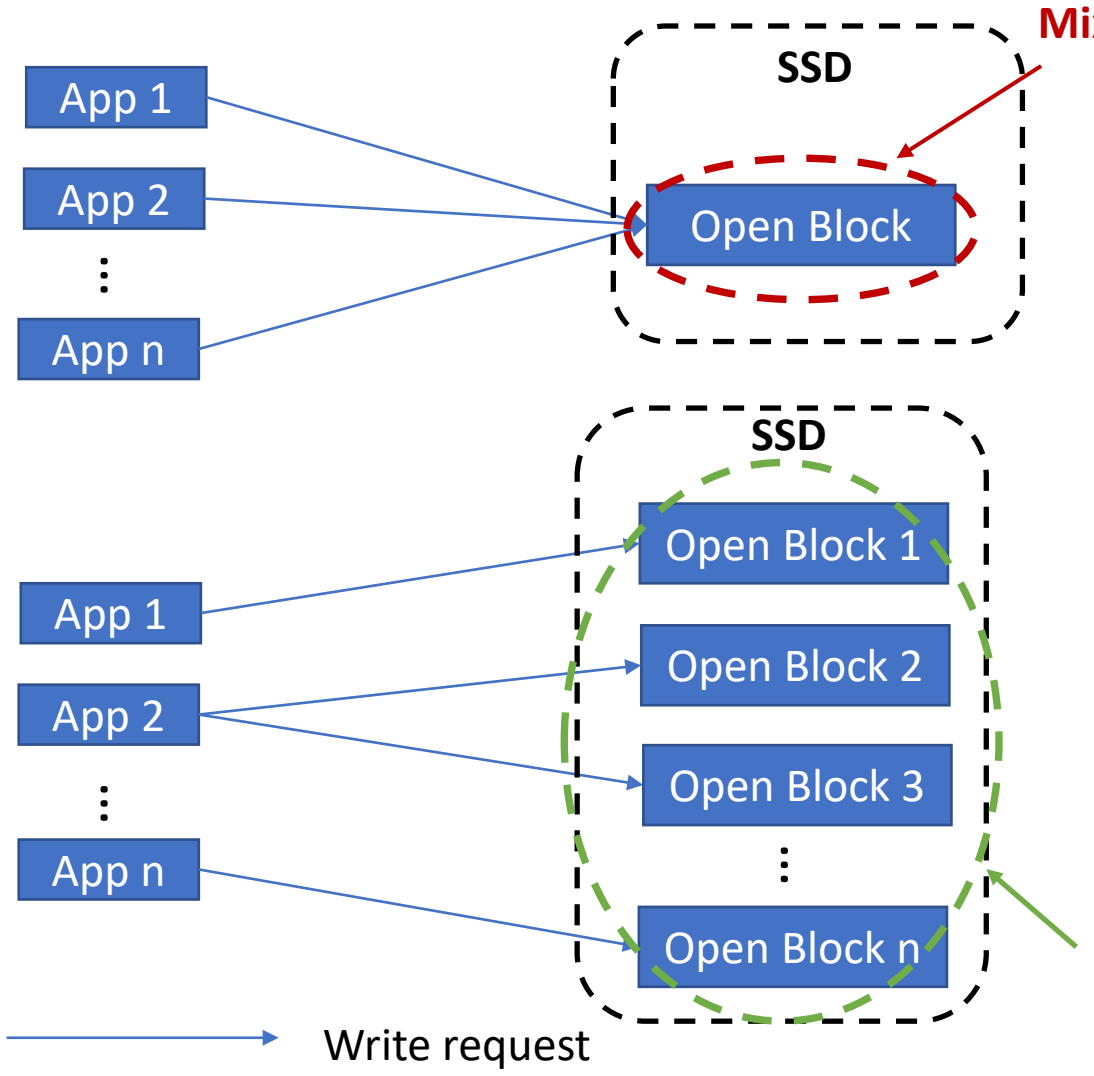

- **Traditional SSD** only has one or two appending points
  - Pages with various lifetimes are mixed in the same erase unit
- **Multi-stream SSD** has multiple appending points
  - Data from different streams are physically separated to different blocks
  - Less lifetime variance in a flash block
  - Reduce write amplification

Challenges in Stream Management

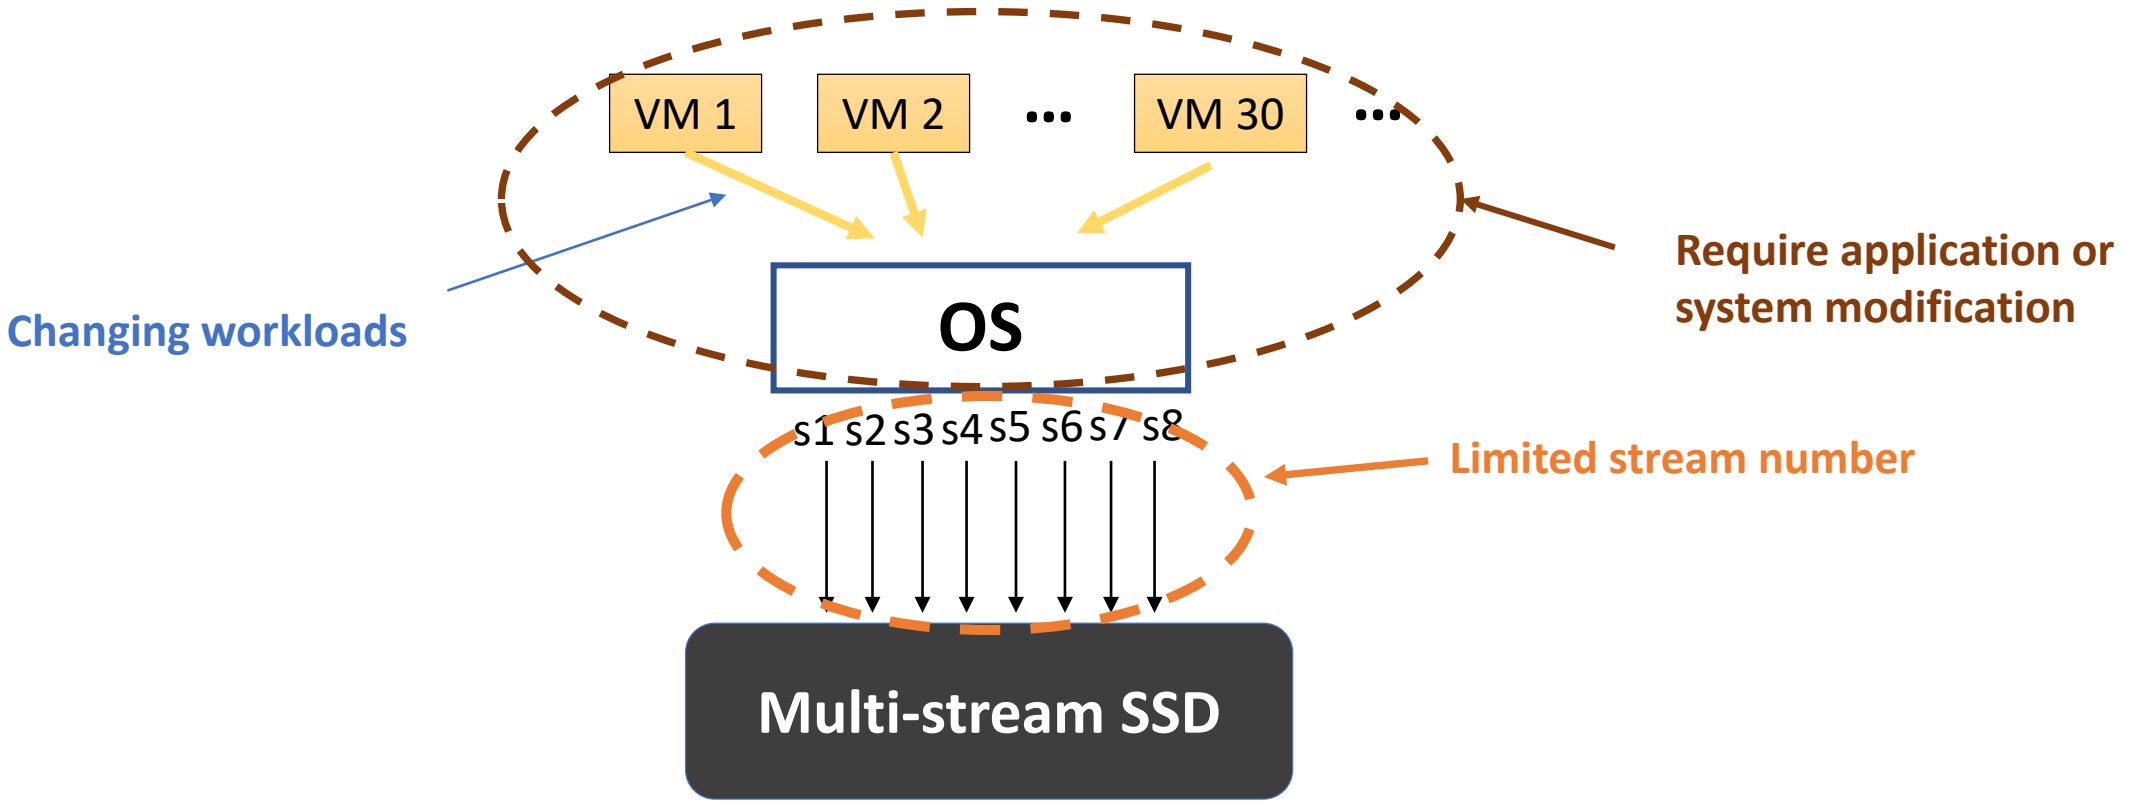

Challenges in stream management (cont.)

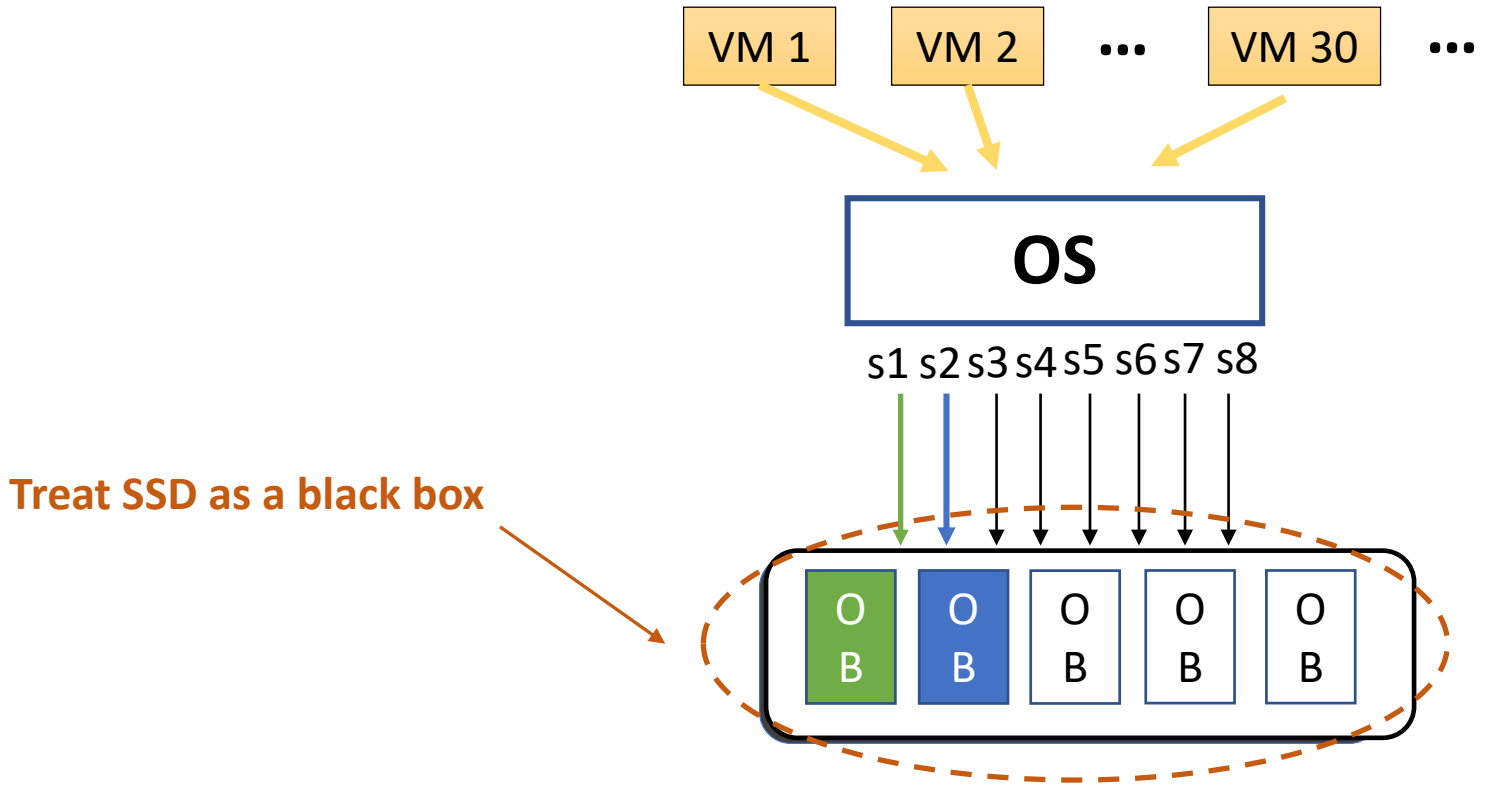

WARCIP: Write Amplification Reduction by Clustering I/O Pages

- Minimize rewrite interval (RWI) variance of pages **in a flash block**
- Learn current I/O patterns to fine-tune its clustering activities
- Device-side data clustering solution

Greedy Clustering

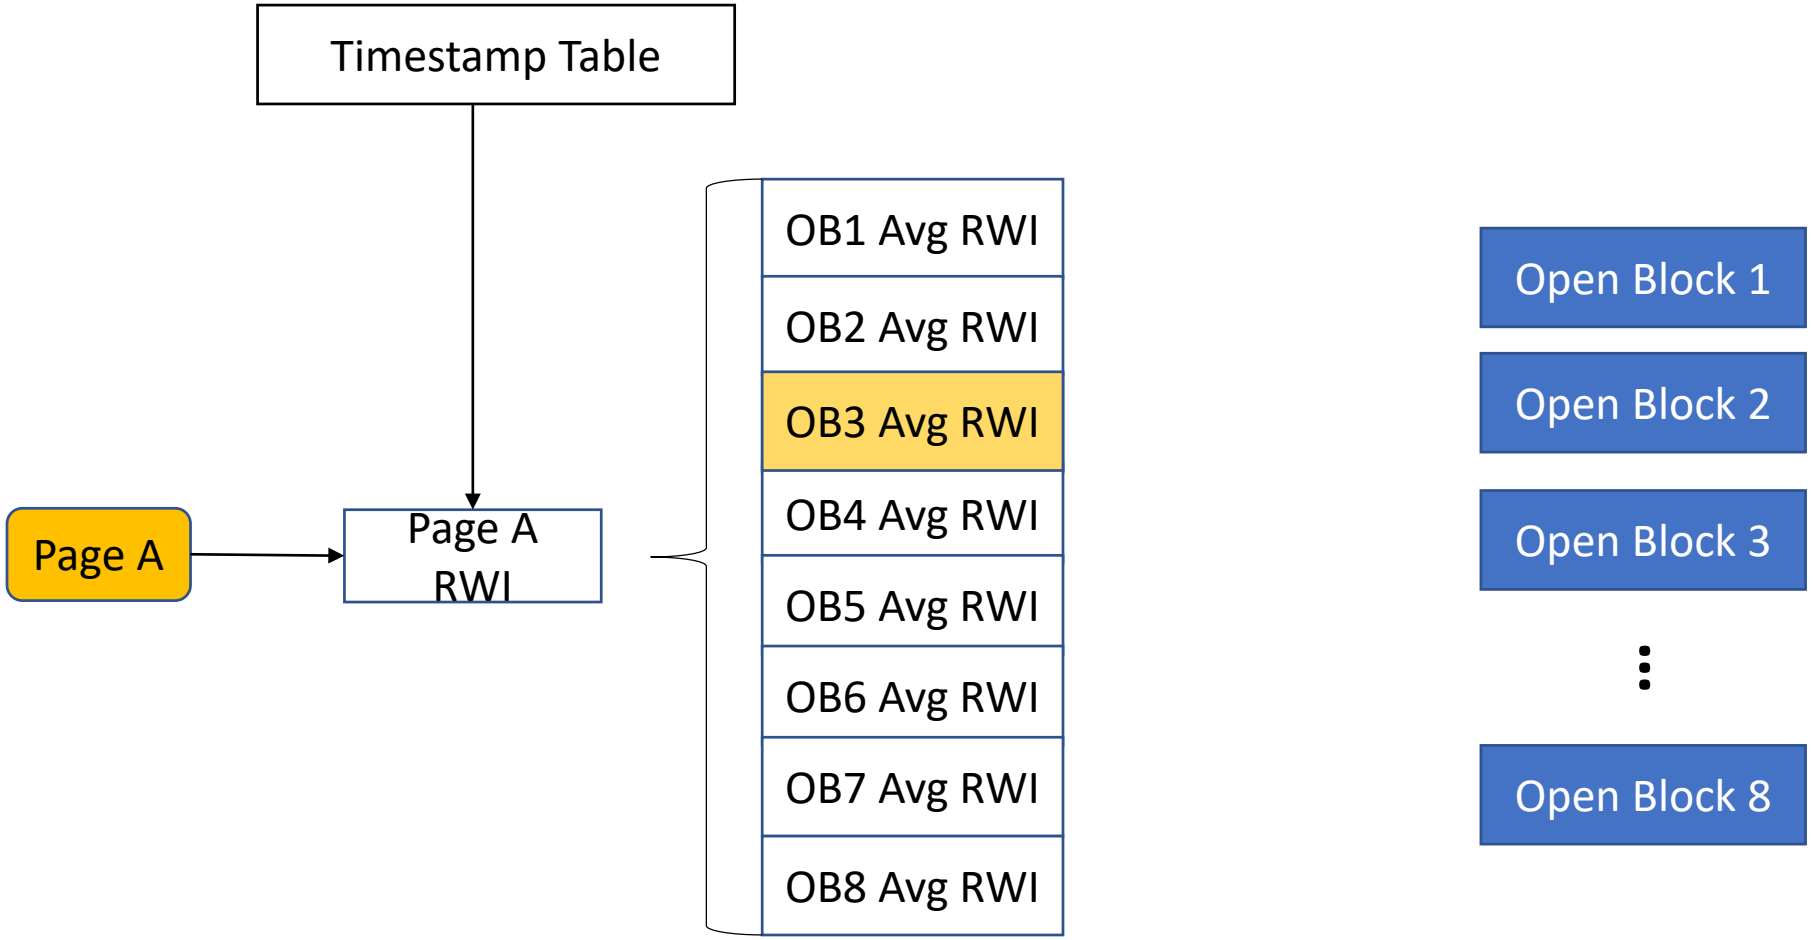

Dynamic Split-and-Merge

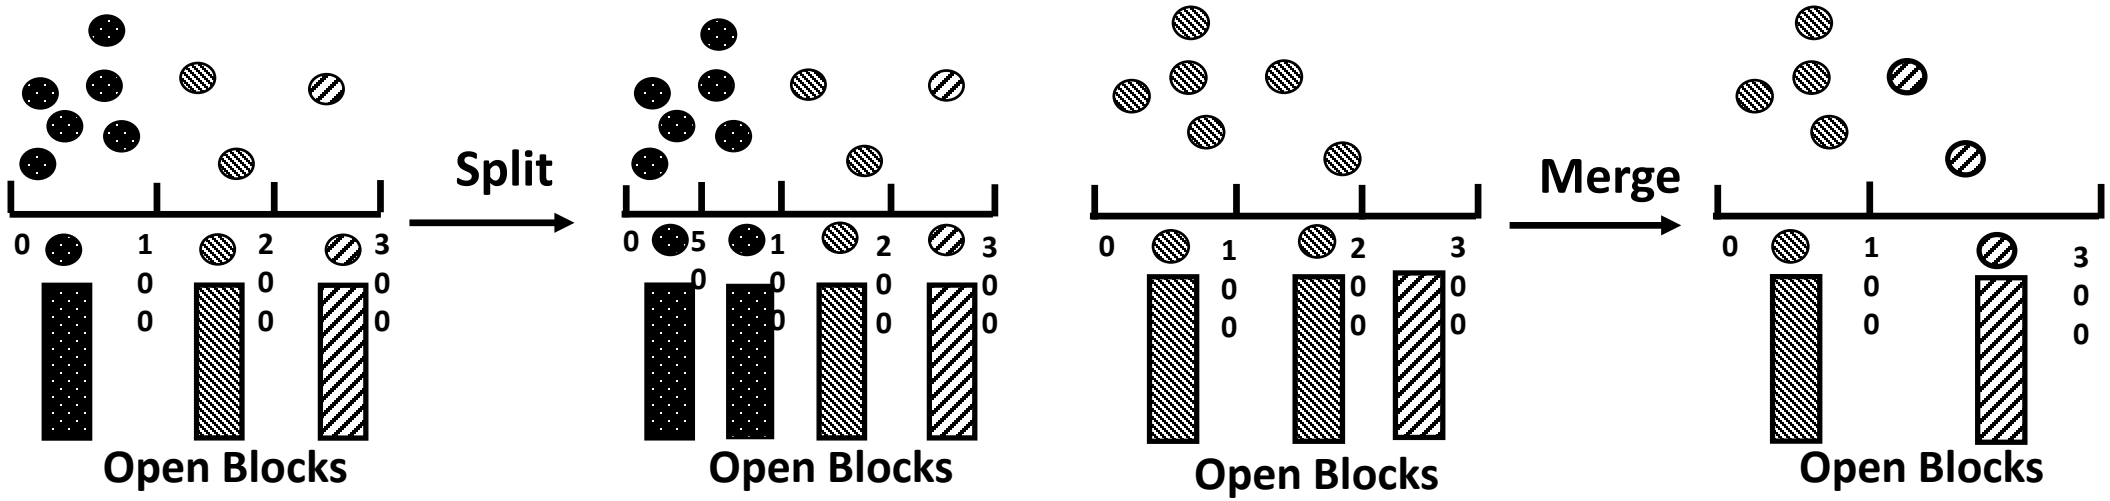

Feedback

- Collect valid pages' information during GC
- Double the recorded rewrite interval of each valid page

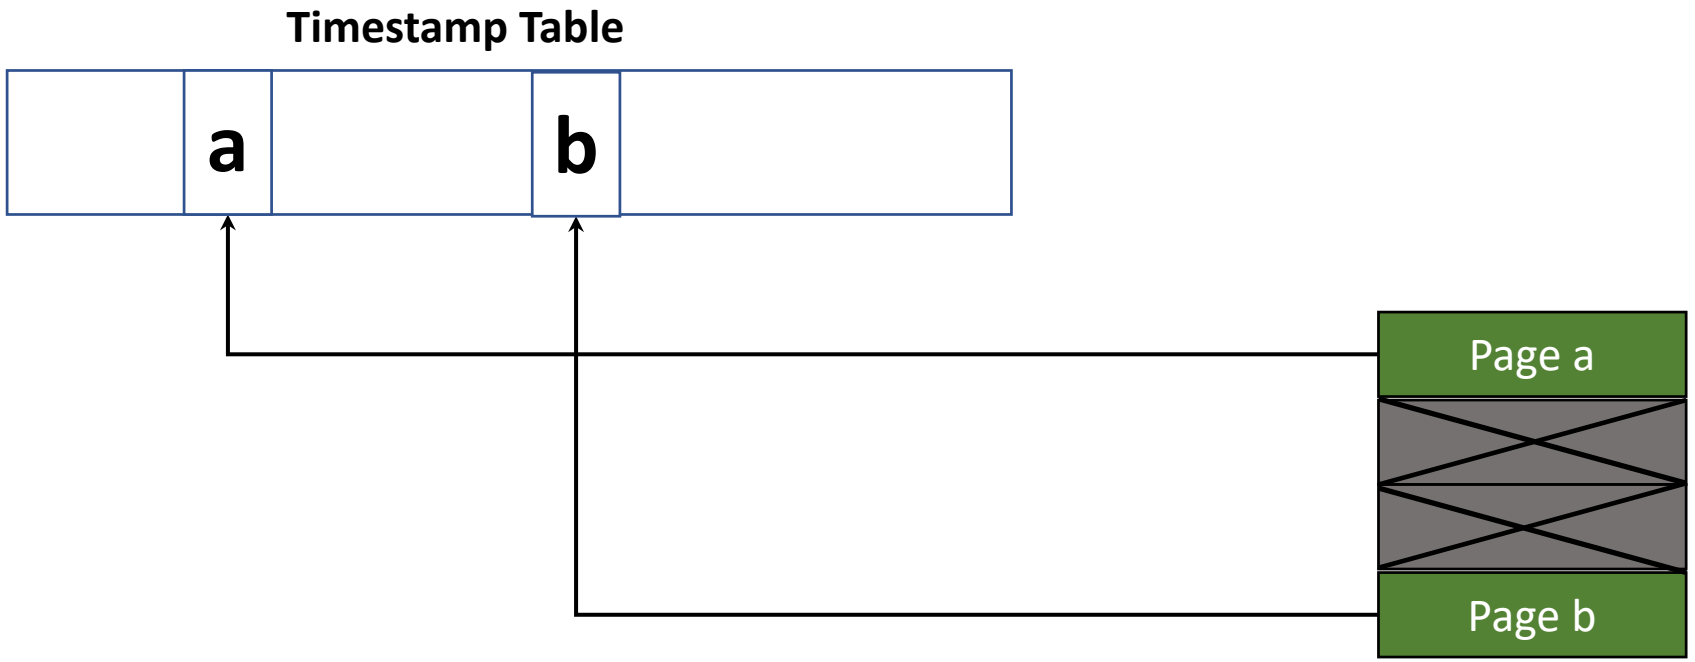

SSD Write Amplification Improvement

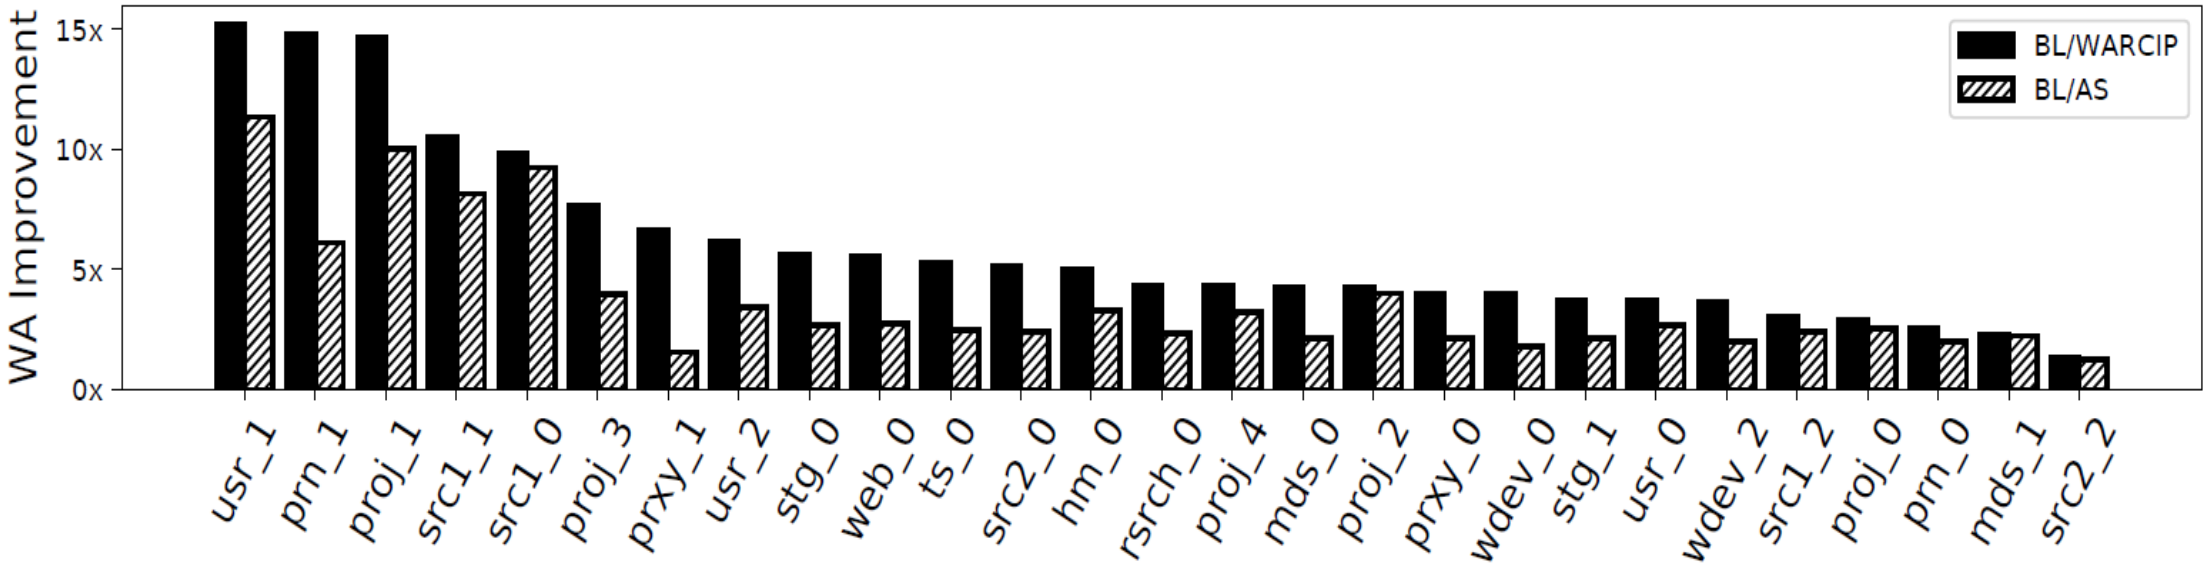

SSD Erasure Count Reduction

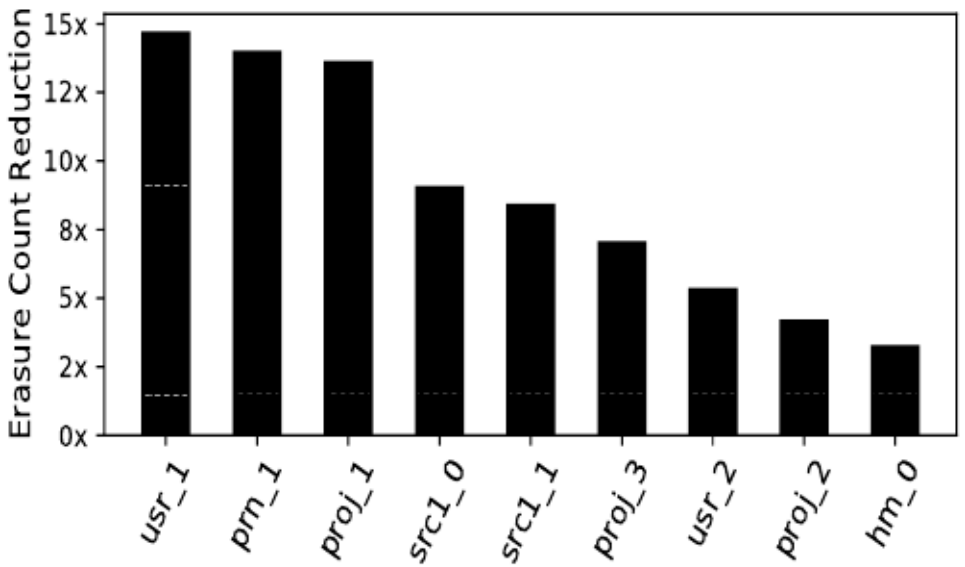

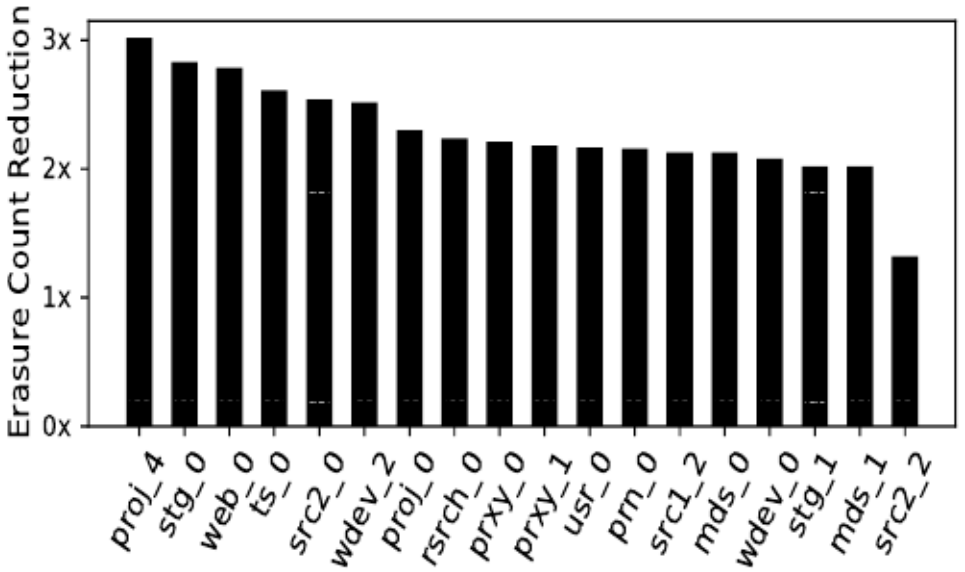

SSD Throughput

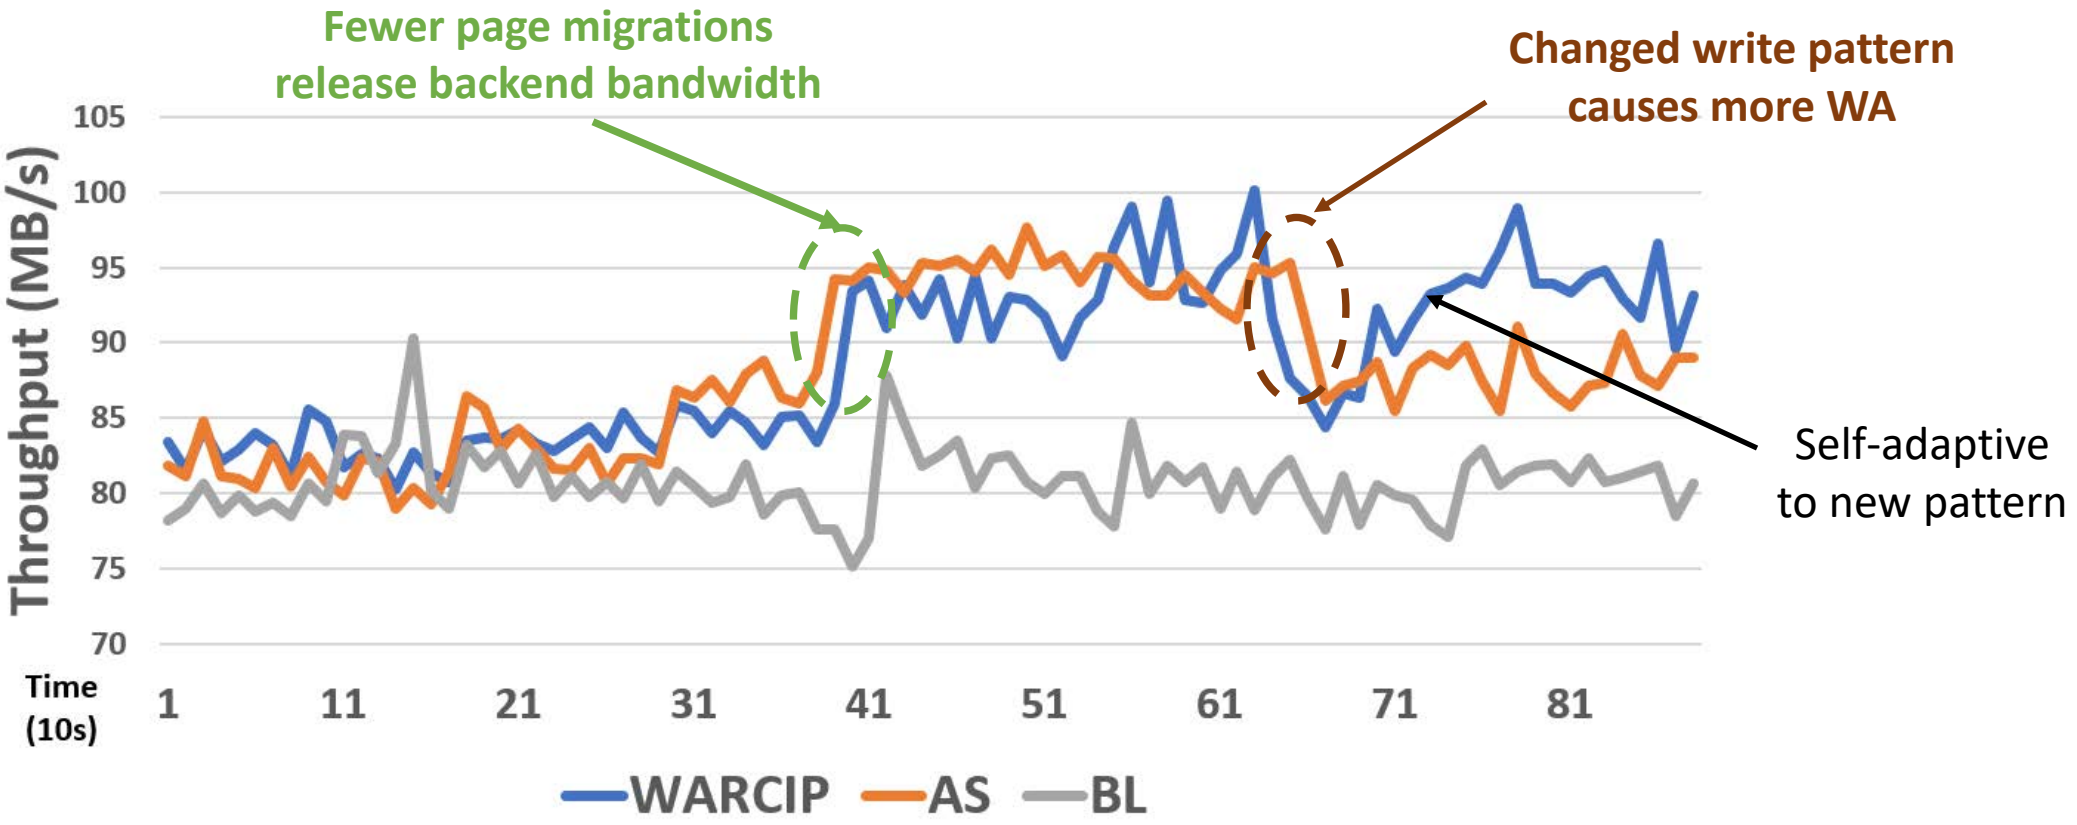

Conclusions

- We proposed WARCIP, a new approach to minimize the negative impact of garbage collection on SSD
  - WARCIP groups pages into a block according to their rewrite intervals to reduce write amplification caused by page migrations
  - Self-optimizing and adaptive to new I/O patterns
- Future work
  - Wear leveling