

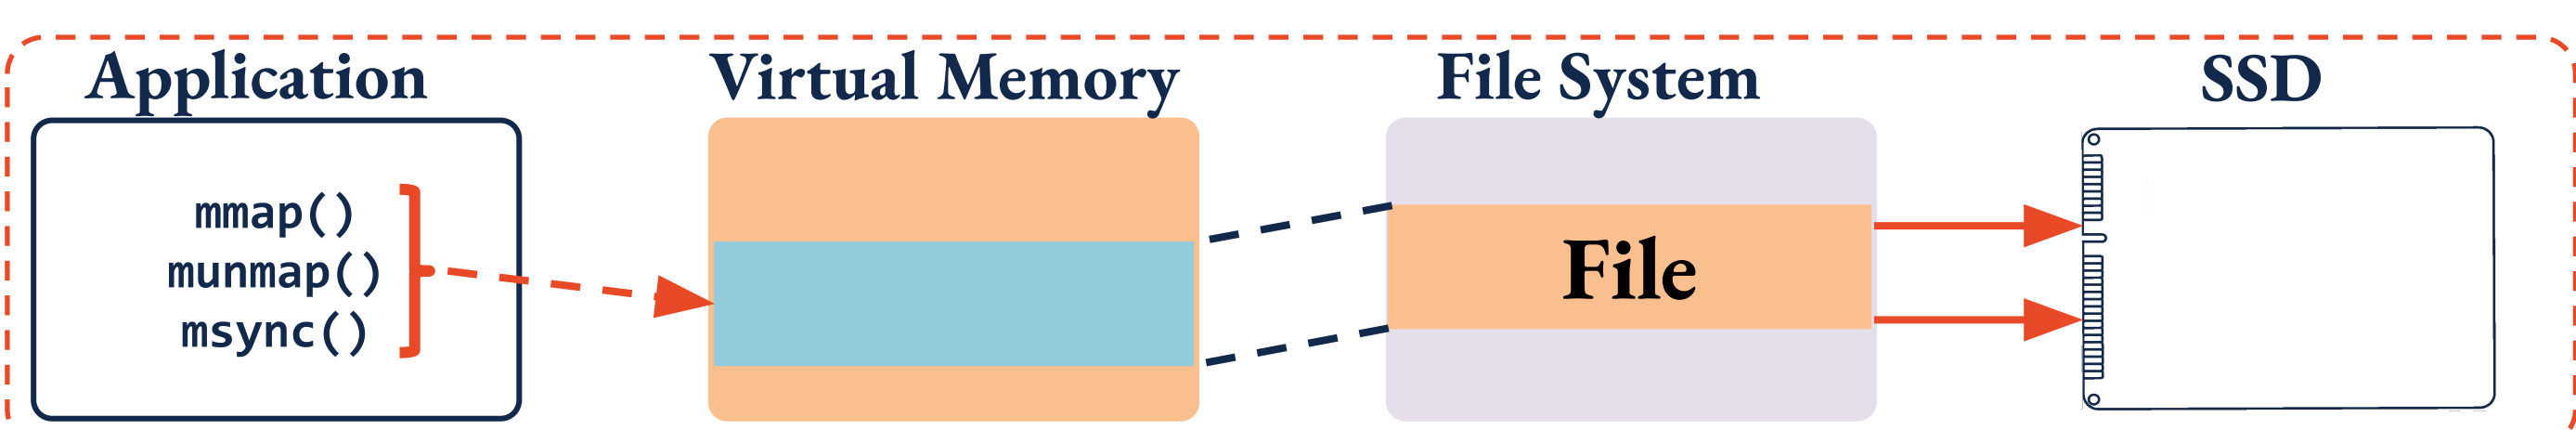
# FlatFlash: Exploiting the Byte-Accessibility of SSDs Within A Unified Memory-Storage Hierarchy

Ahmed Abulilla<sup>1</sup>, Vikram Sharma Mailthody<sup>1</sup>, Zaid Qureshi<sup>2</sup>, Jian Huang<sup>1</sup>, Nam Sung Kim<sup>1</sup>, Jinjun Xiong<sup>3</sup>, Wen-mei Hwu<sup>1</sup>

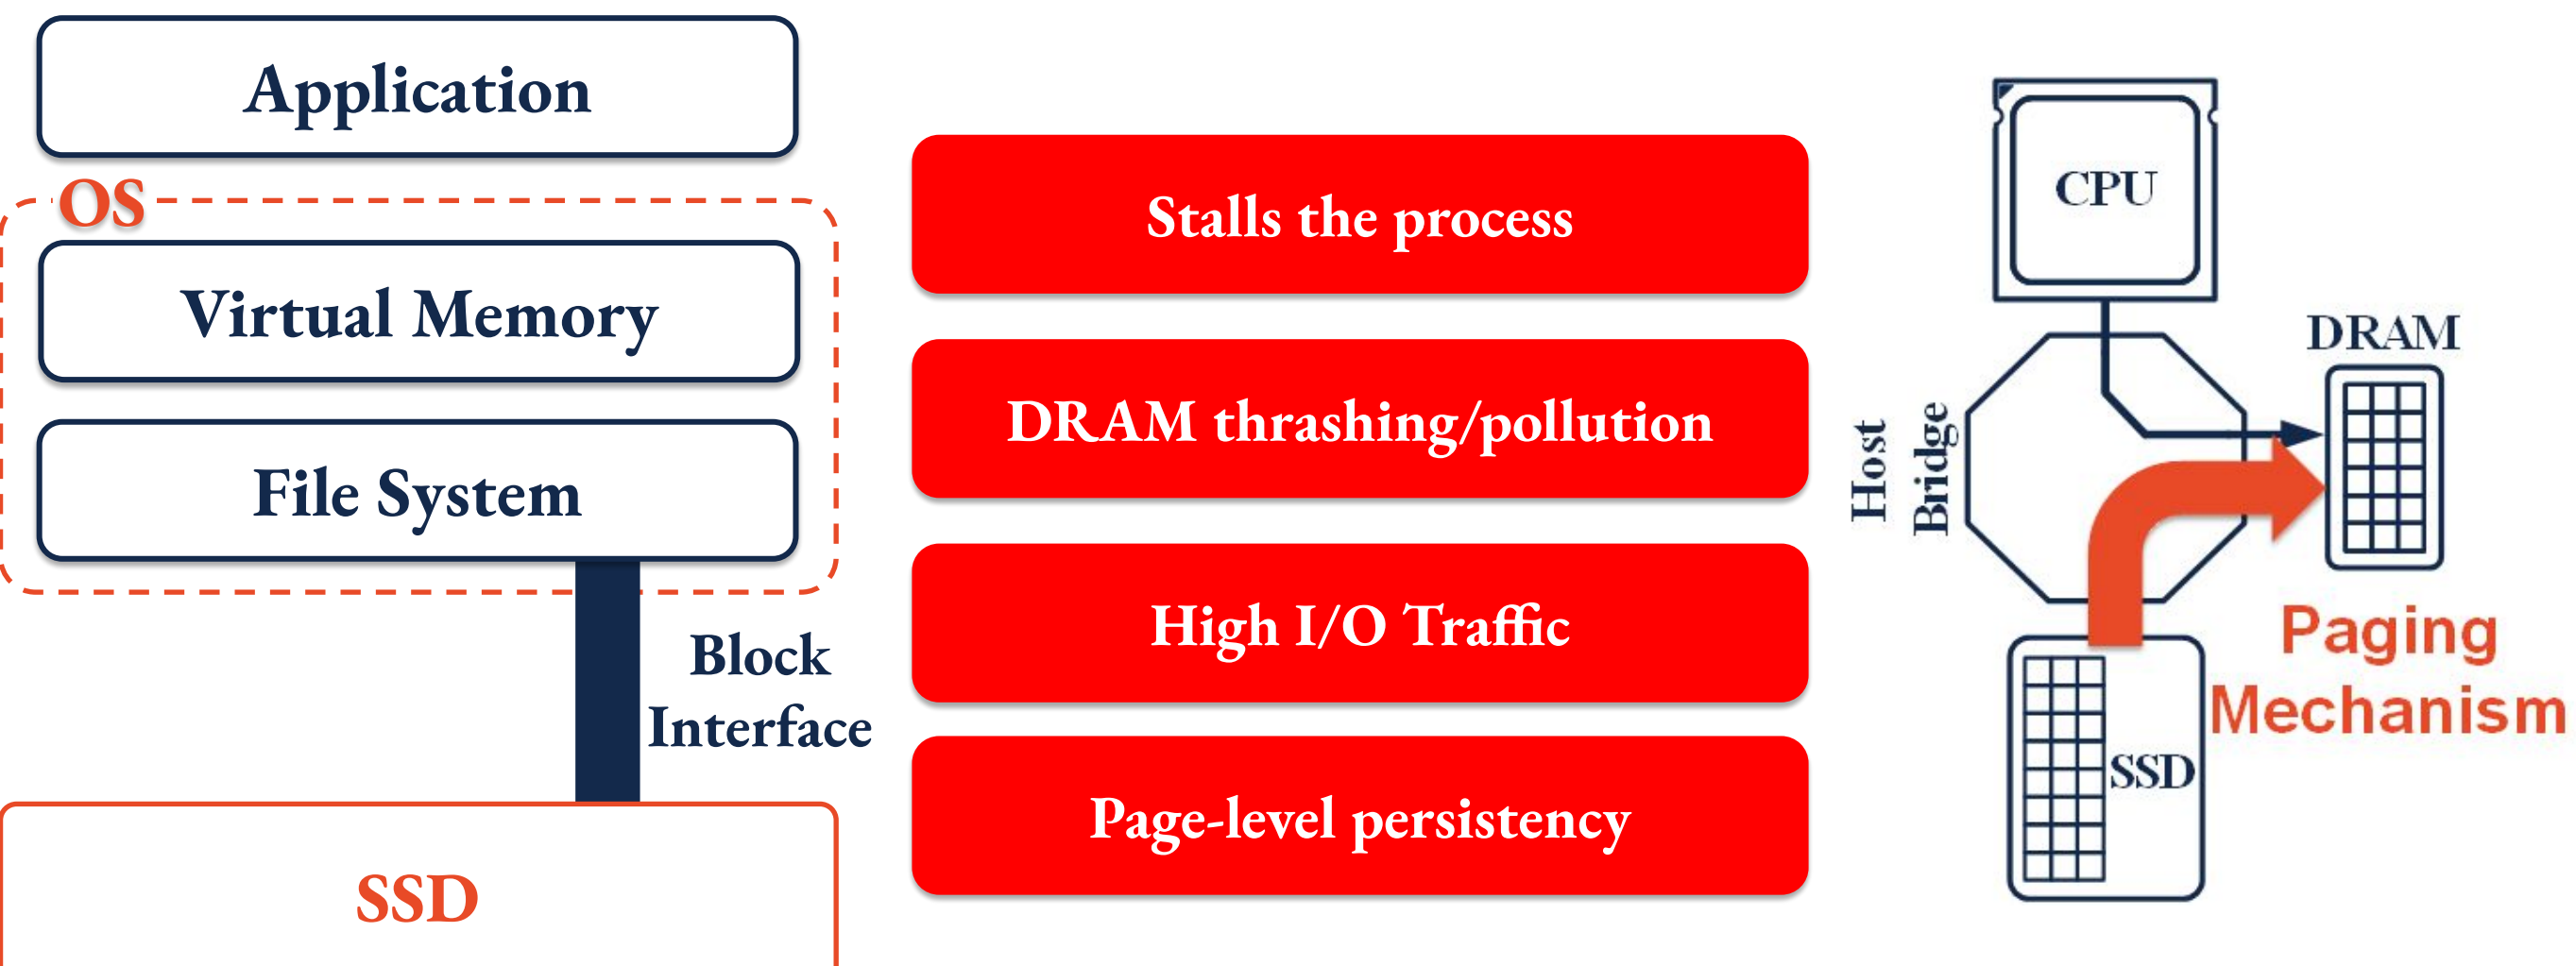
<sup>1</sup> ECE, <sup>2</sup> CS, University of Illinois at Urbana-Champaign, Urbana, IL 61801

<sup>3</sup> Cognitive Computing & University Partnership, IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598

## Using Flash as Non-Volatile Memory



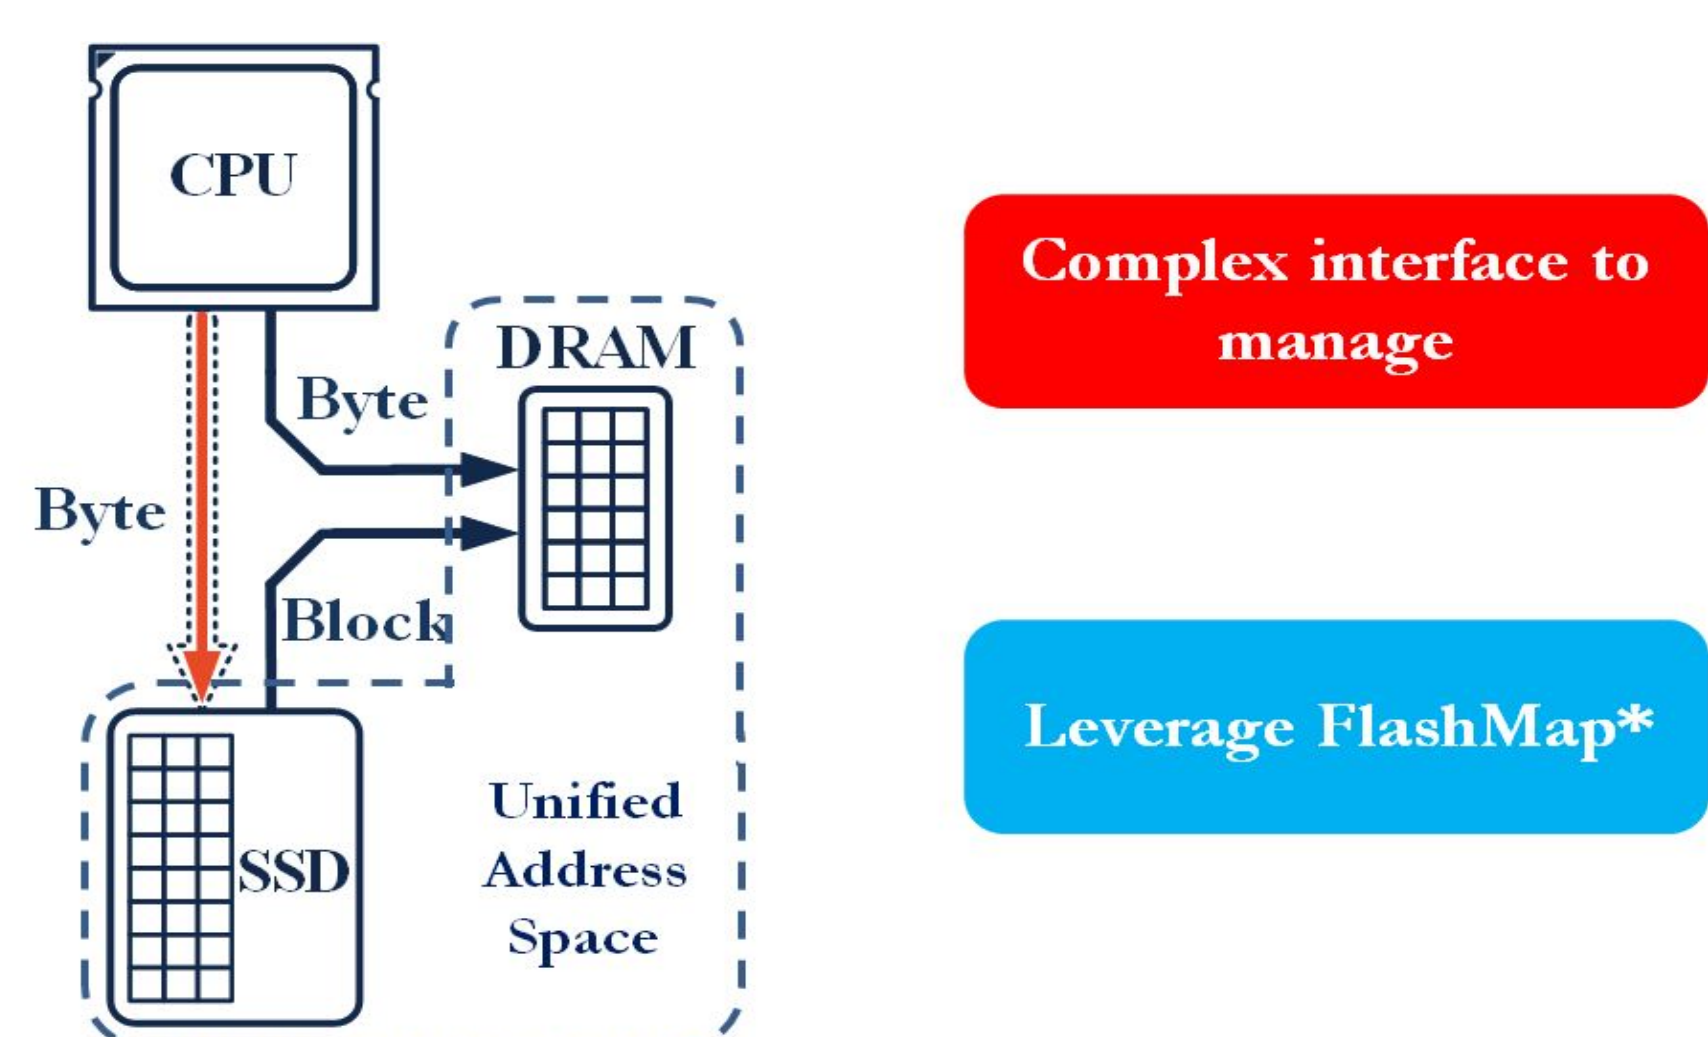
## Limitations of Block I/O Interface



## Enabling the Byte-Accessibility of SSDs




## Unifying the Memory and Storage



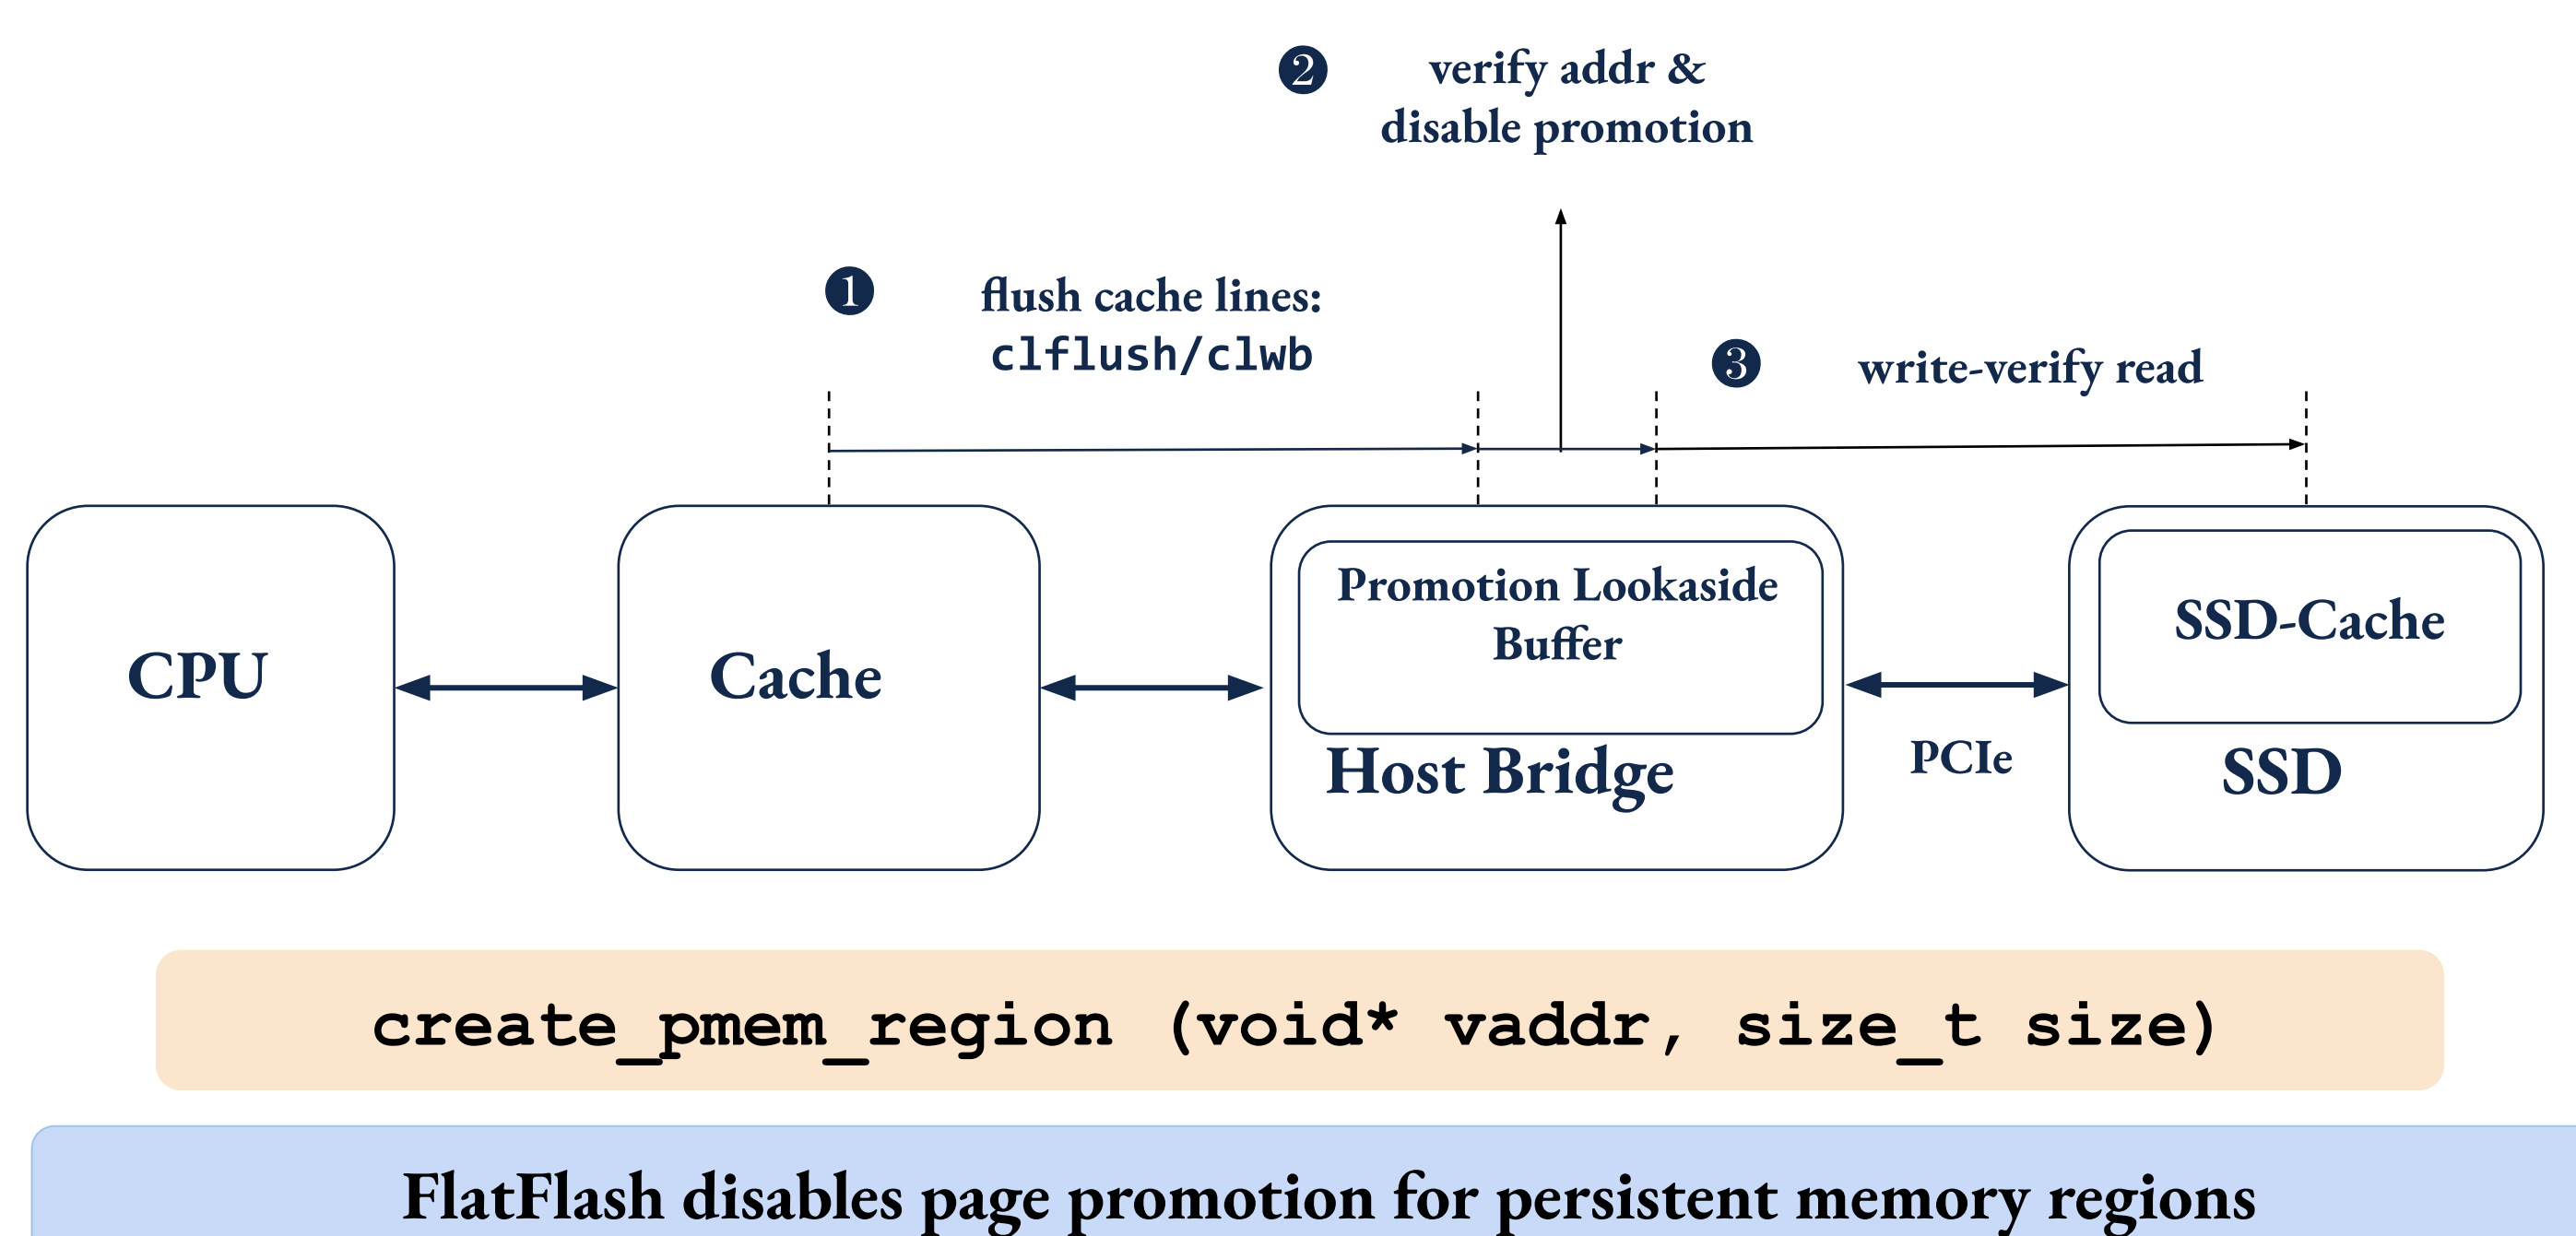
## Performance Challenges

Overhead Source	Average ( $\mu$ sec)	Category
Read a cache line in SSD-Cache via PCIe MMIO	4.8	Page Promotion to Host DRAM
Write a cache line in SSD-Cache via PCIe MMIO	0.6	
Promote a page from SSD-Cache to host DRAM	12.1	Off-Critical Path Page Promotion
Update PTE and TLB entry in host machine	1.4	
Page table walking to get the page location	0.7	

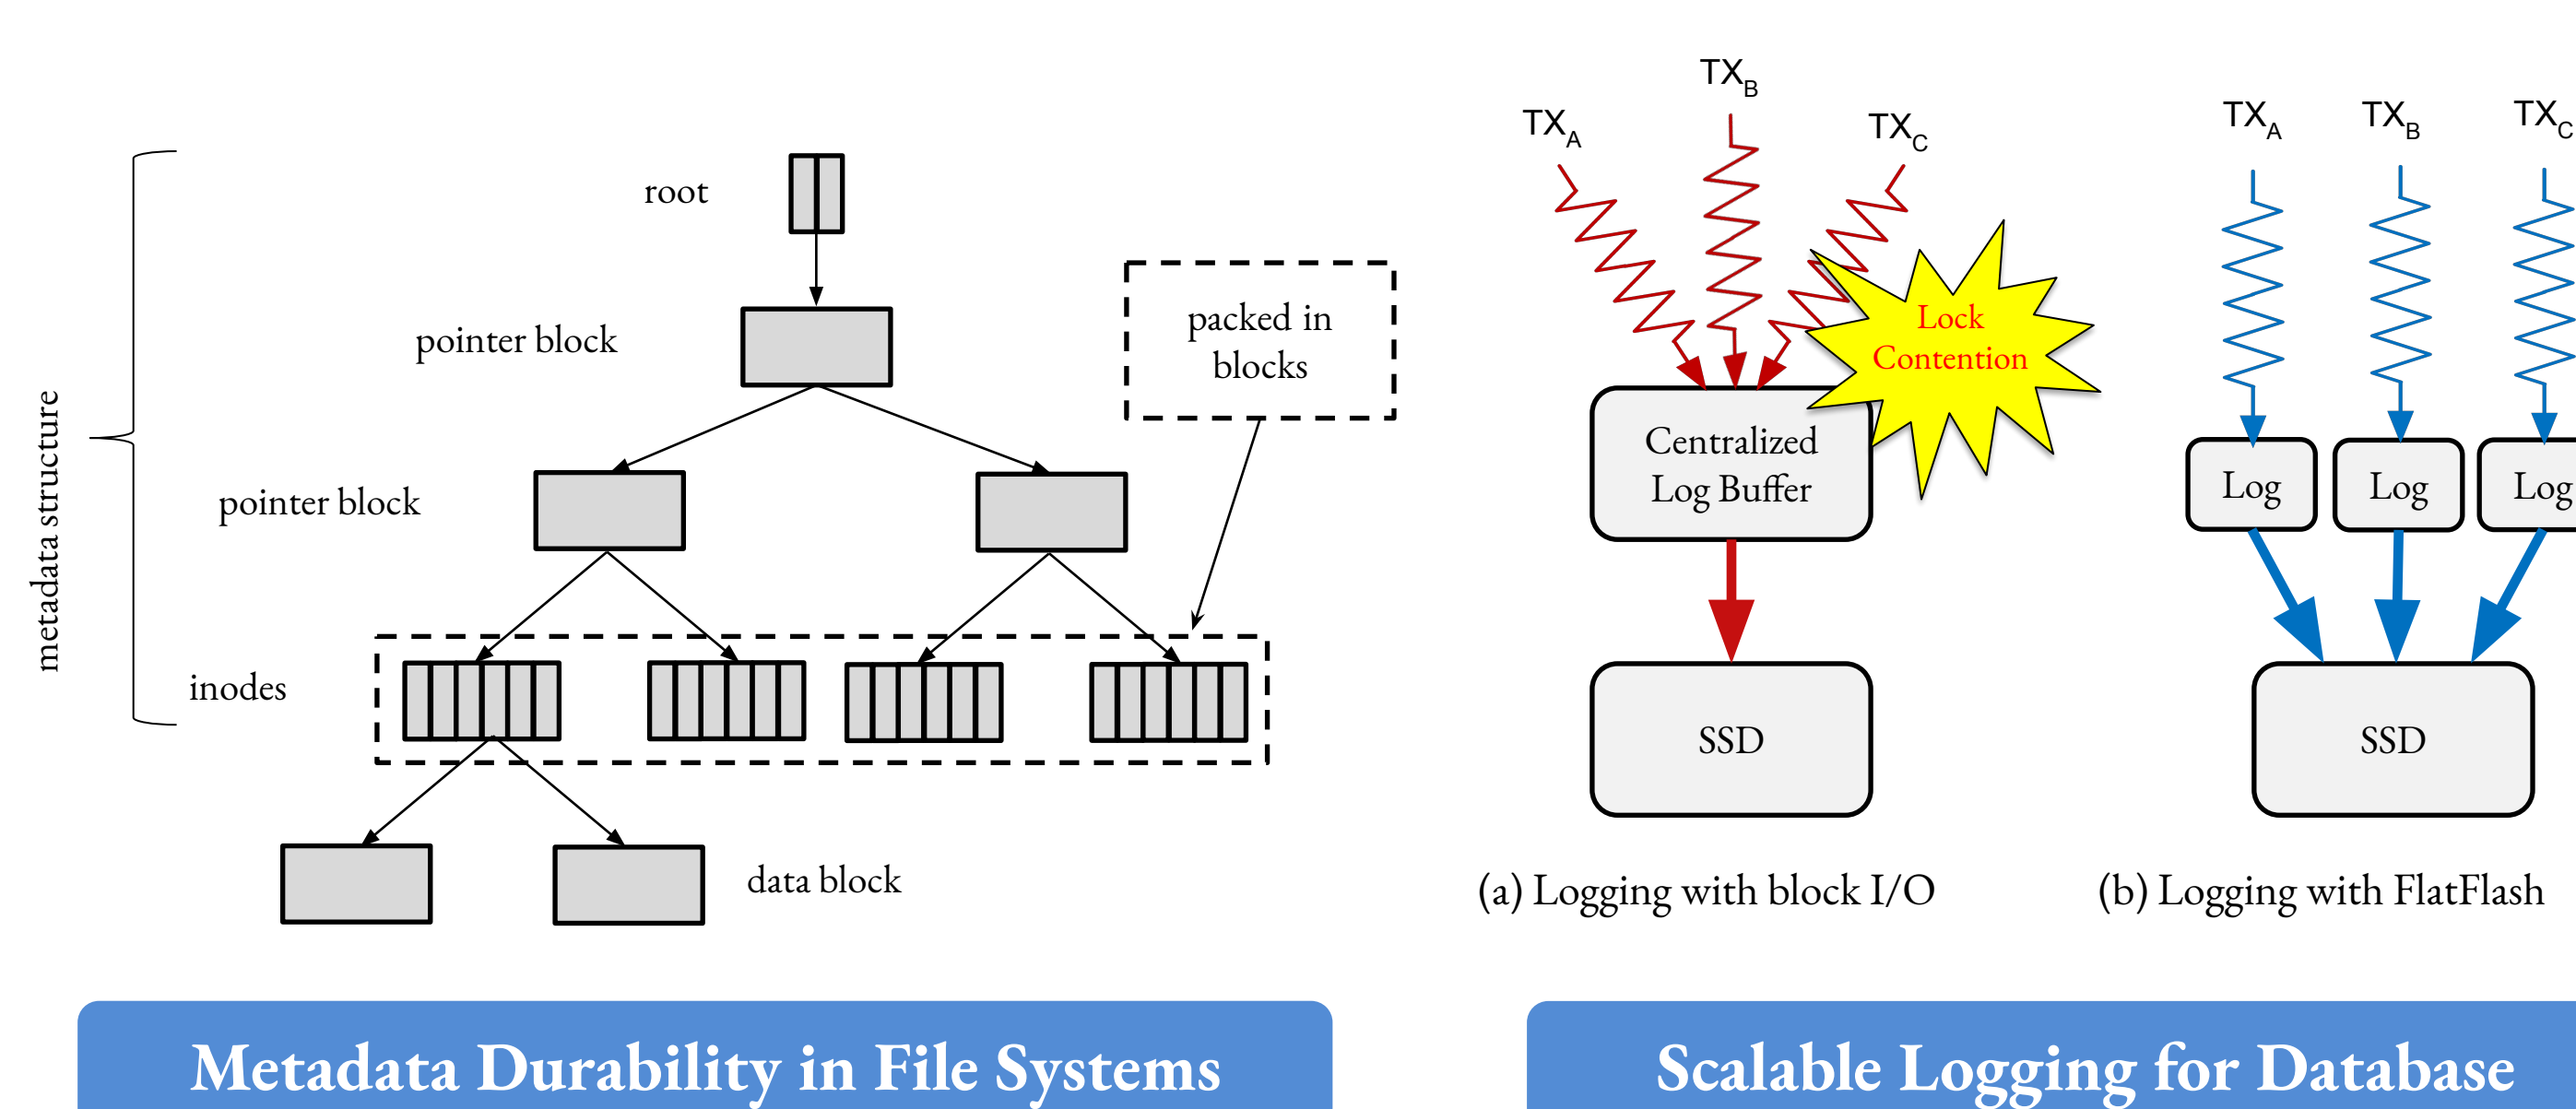
## Ensuring Data-Consistency During Page Promotion



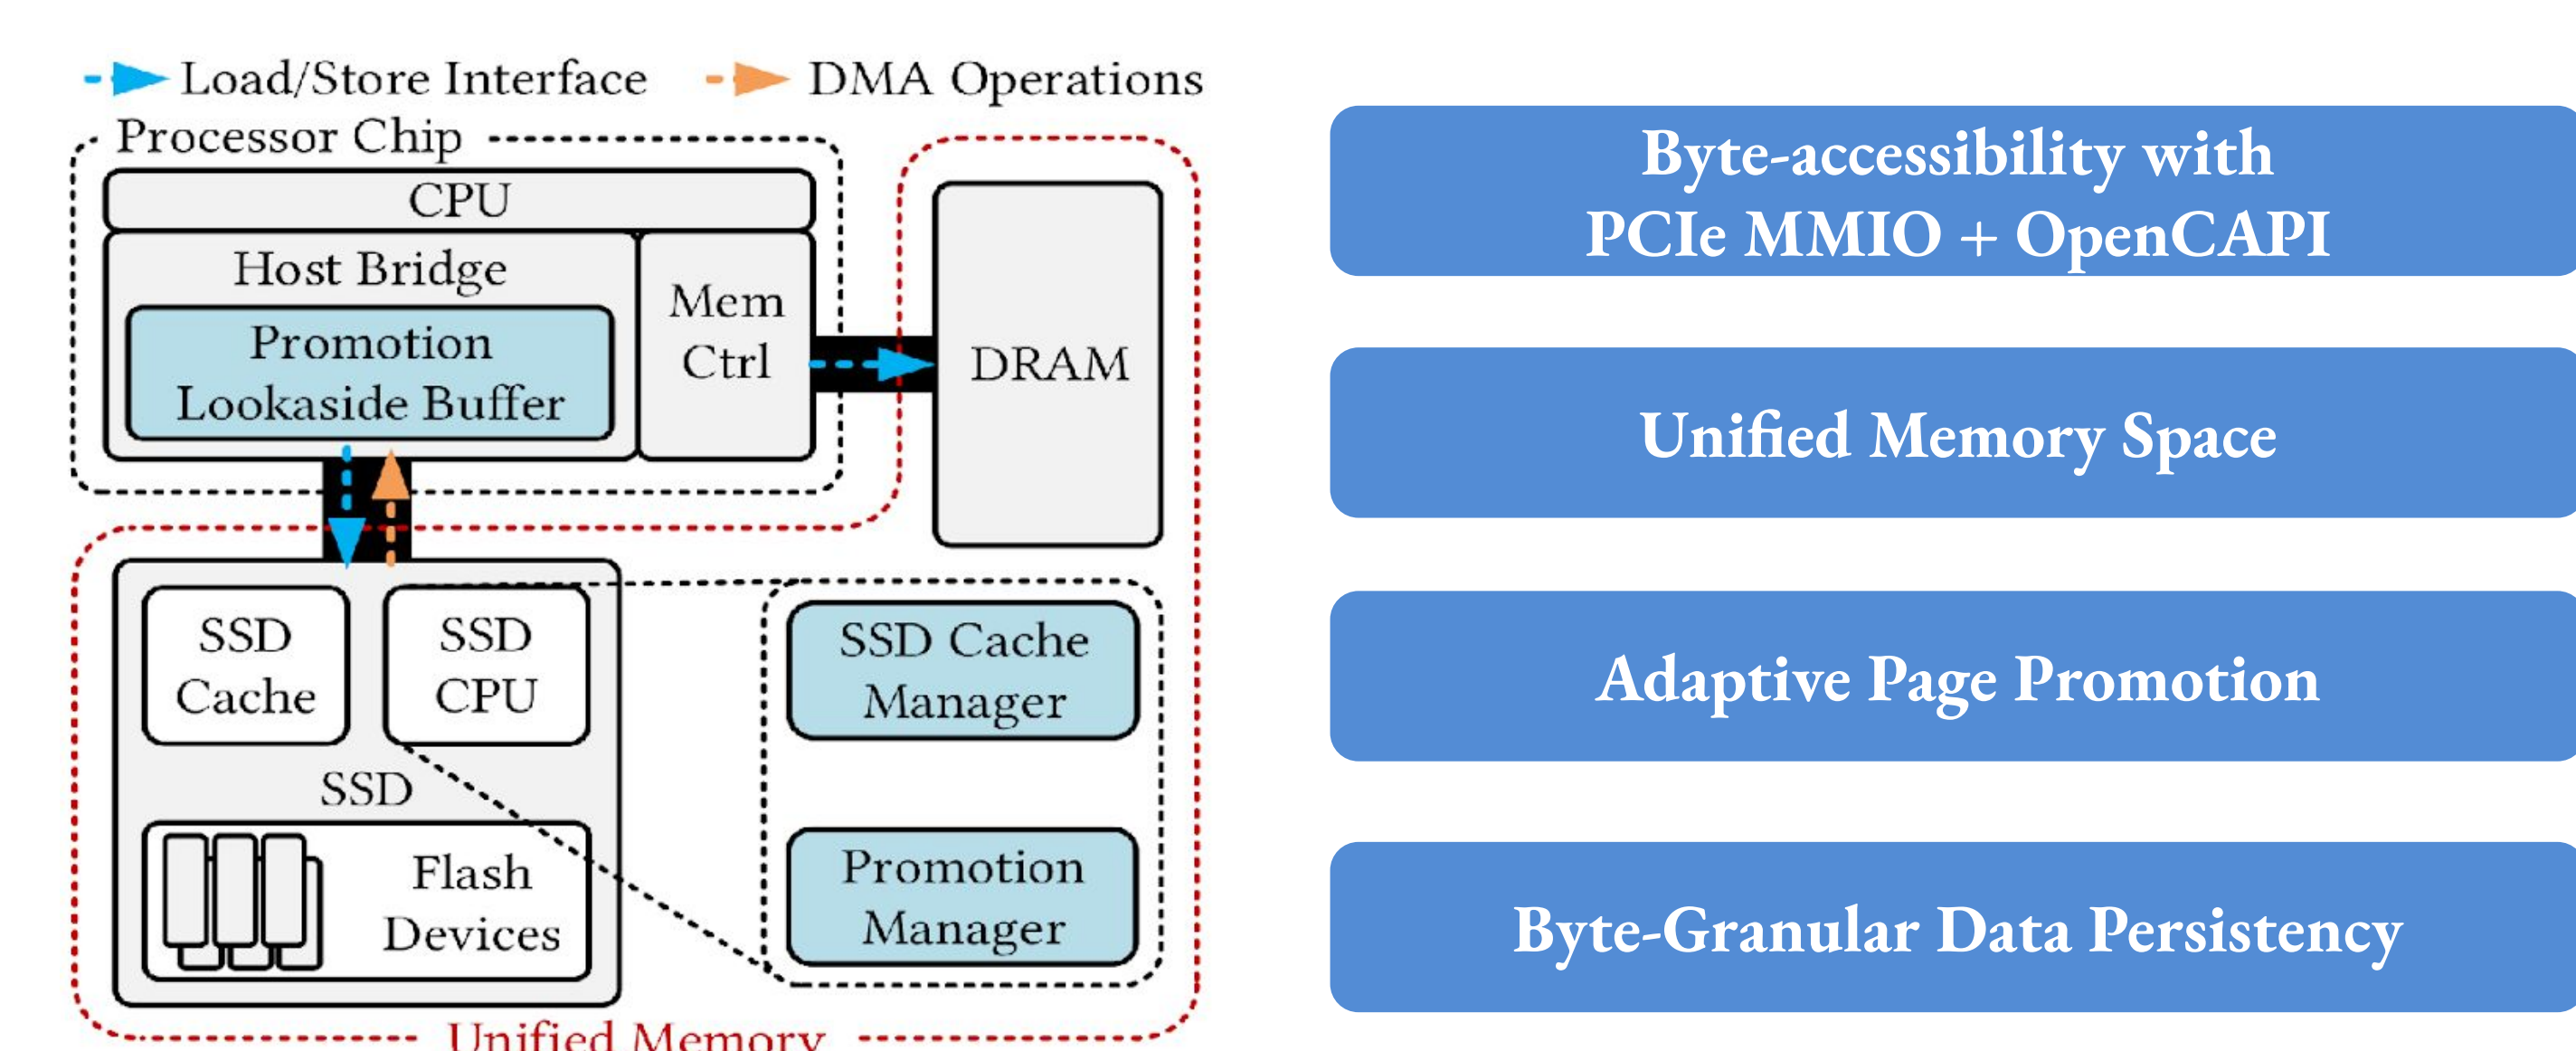
## Enabling Byte-Granular Data Persistence



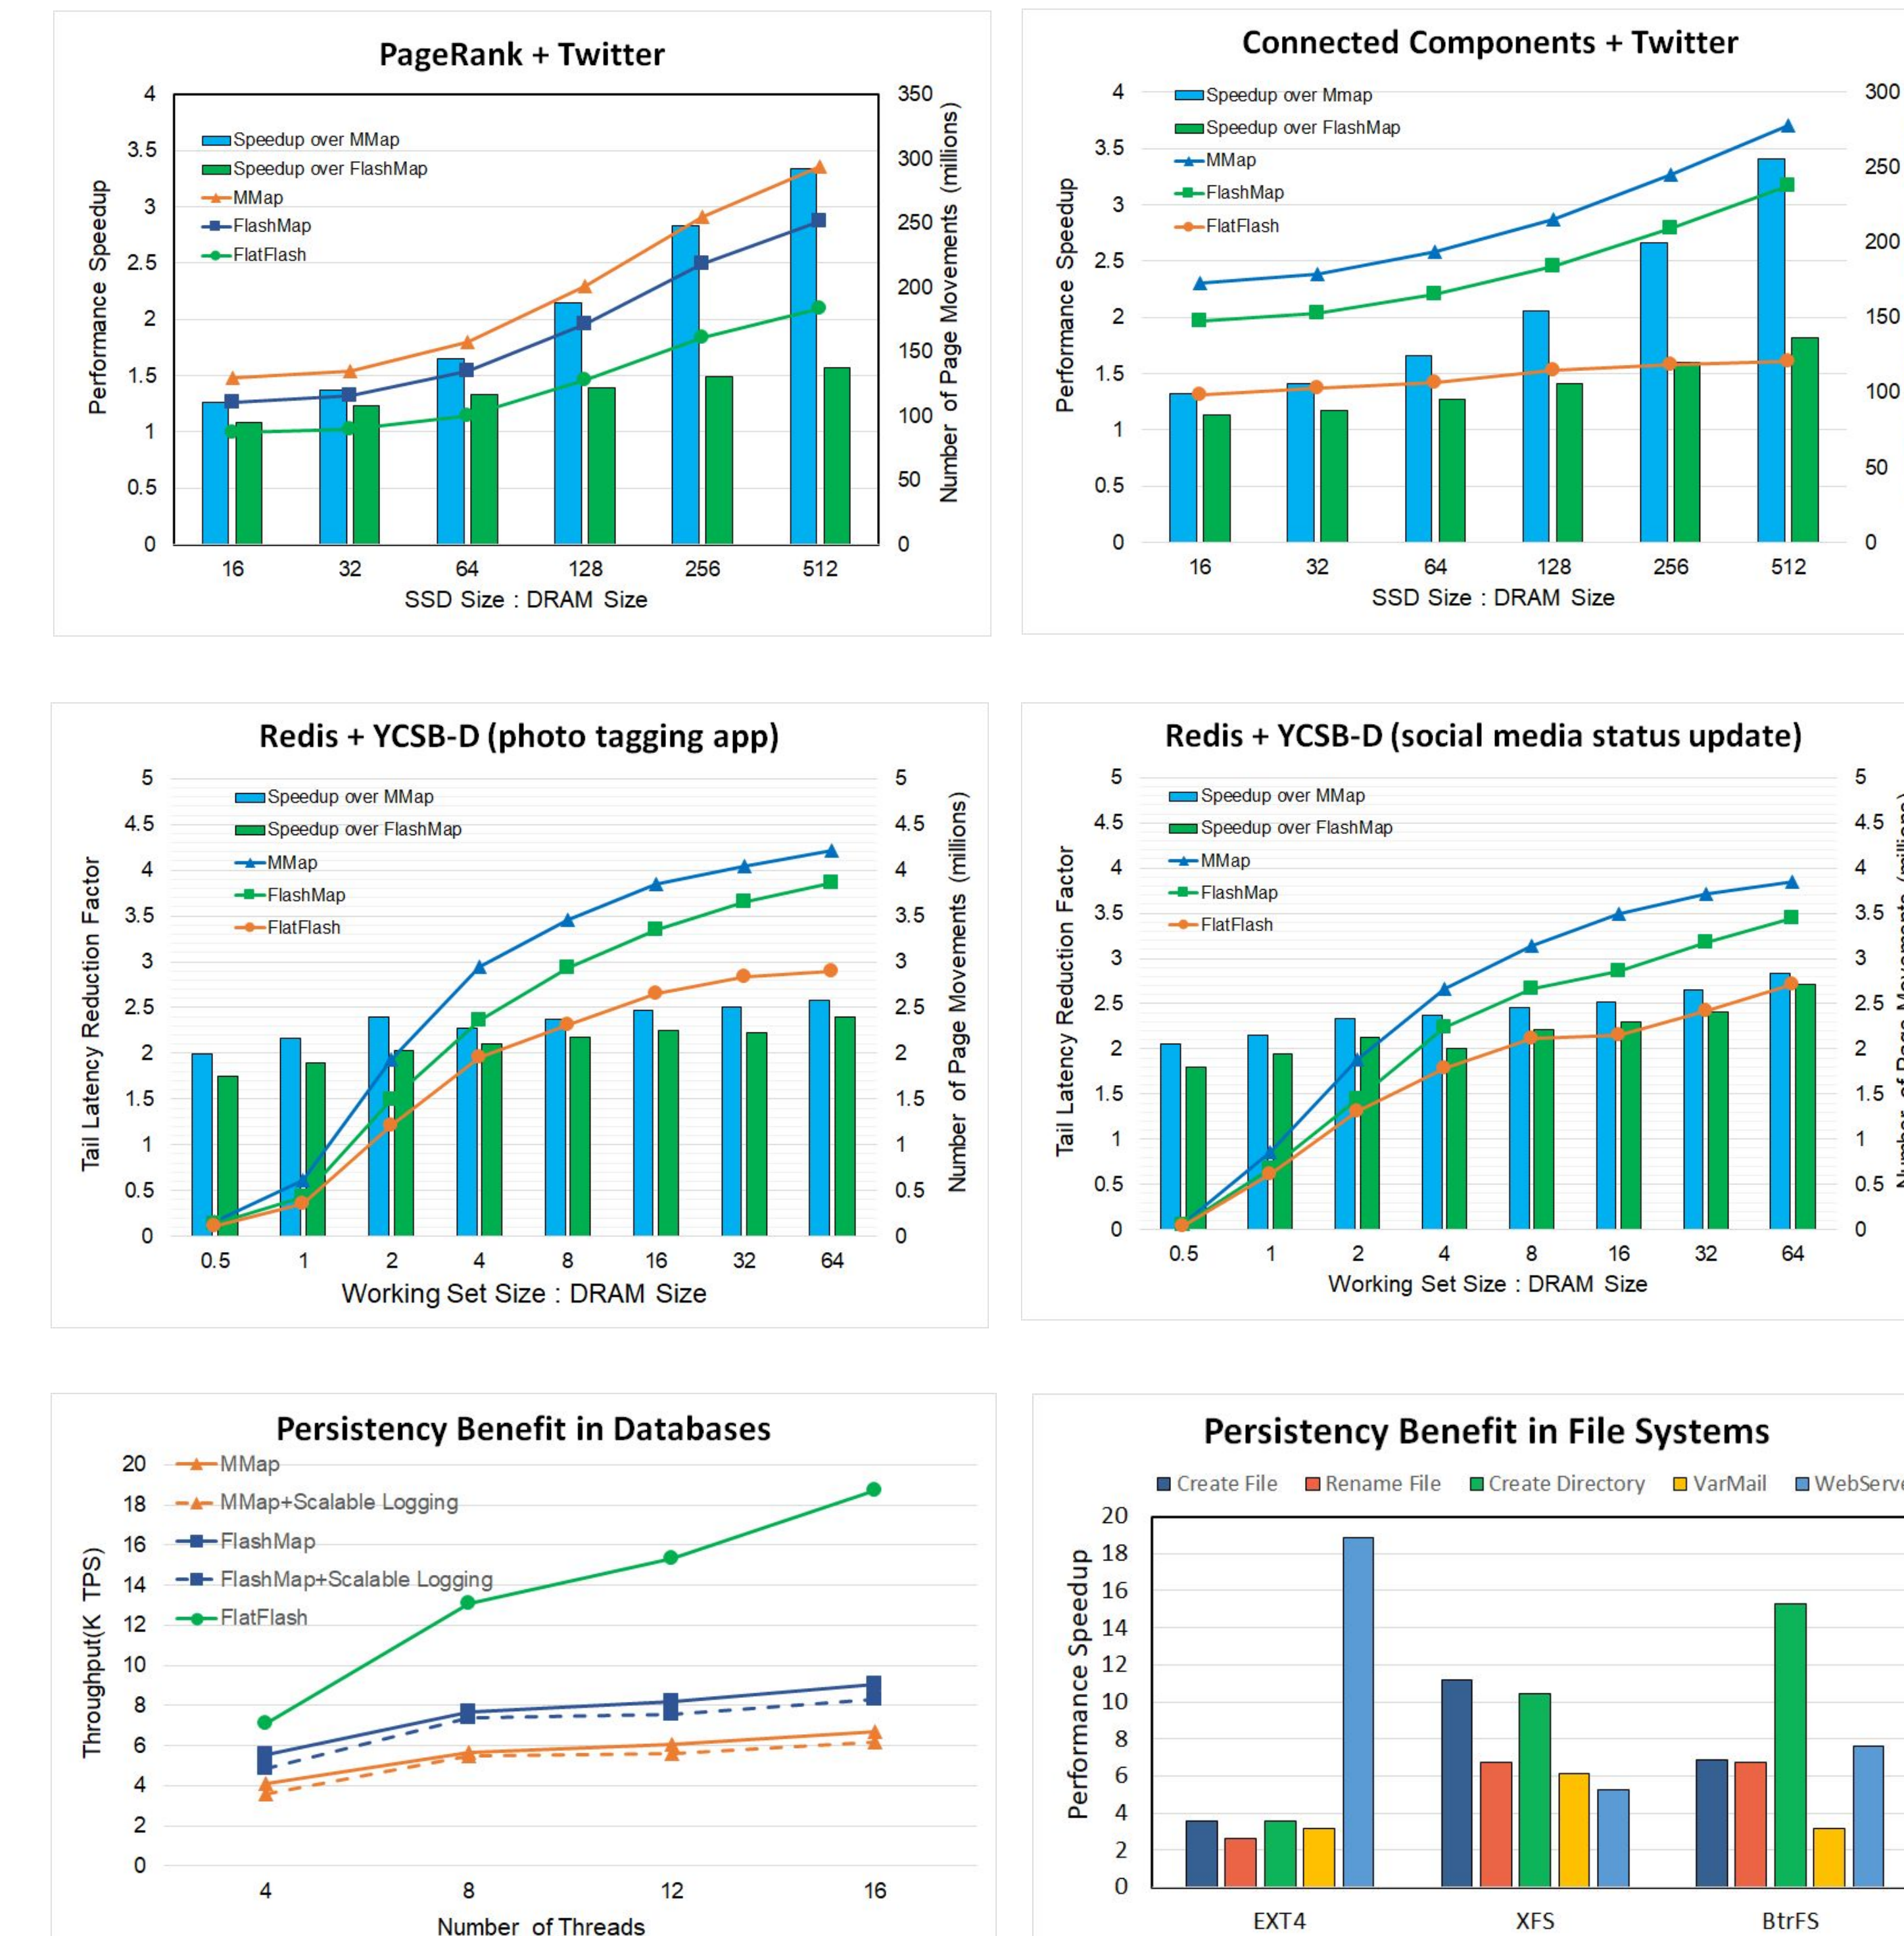
## System Implications of FlatFlash



## Put It All Together



## Results



Paper Download

