Treating SSDs as Memory

Anirudh Badam
Researcher

Hot Data Outpacing Memory Capacity

Zuckerberg's law: The number of photos accessed is doubling every 18 months

Cisco Mobile Traffic Survey: Mobile multimedia data set to increase 56% per year

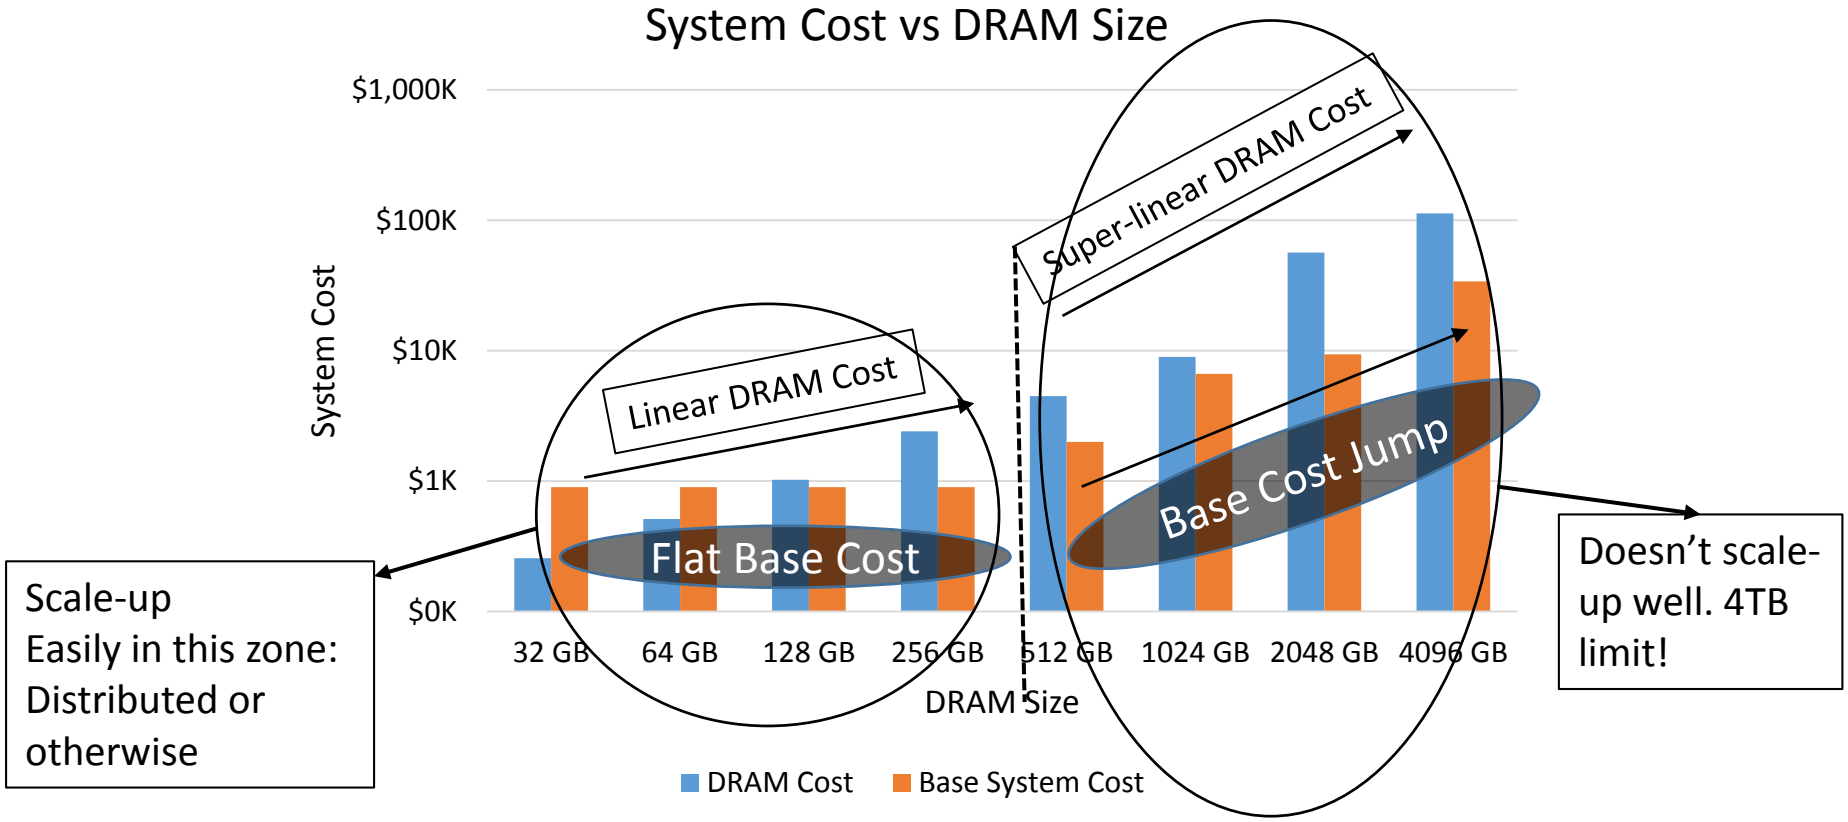

Use Flash to Augment DRAM

Flash is…	SSD form factor	PCIe Flash	Insight
Really Dense	1TB	20TB	20x denser than DRAM
Really Cheap	$0.75/GB	$5.00/GB	5x cheaper than DRAM
Really Fast (Random)	100K IOPS	1M IOPS	Not a TPS bottleneck
Really Fast (Sequential)	4Gbps	48Gbps	Not a BW bottleneck
Really Fast (Latency)	200µS	100µS	Intra-datacenter latency
Really Efficient	200mW	150W	10x better than DRAM

Augment DRAM with Flash

Why augment DRAM with Flash?

Unmodified Memory Intensive Applications

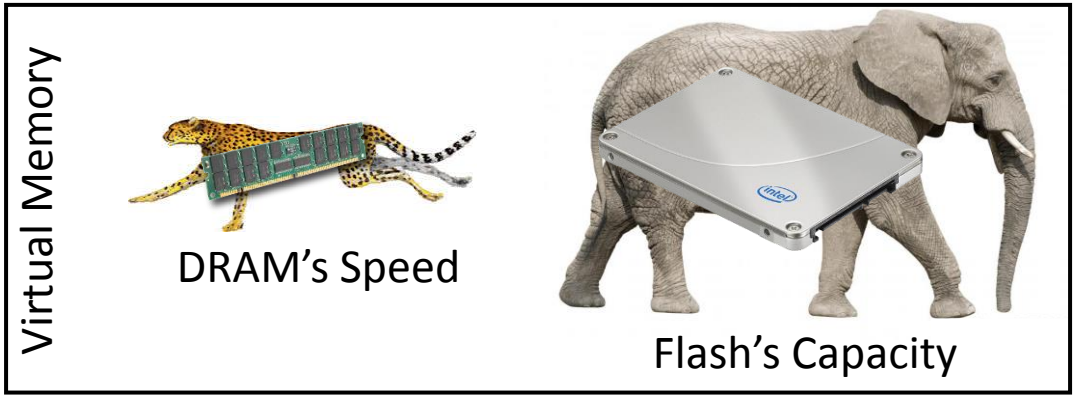

What's wrong with paging/memory mapped files?

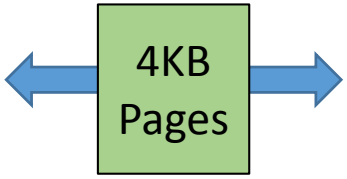

Too much data transfer

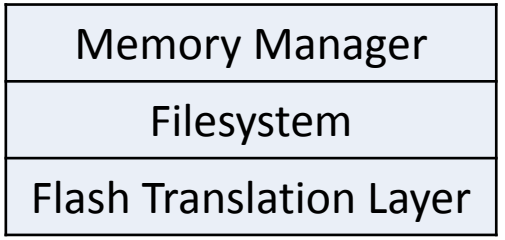

Too much software overhead

Reduce Traffic via Sparse Virtual Memory

Main idea: allocate sparse virtual pages, pack DRAM for efficiency, move as little data as possible between DRAM and the SSD.

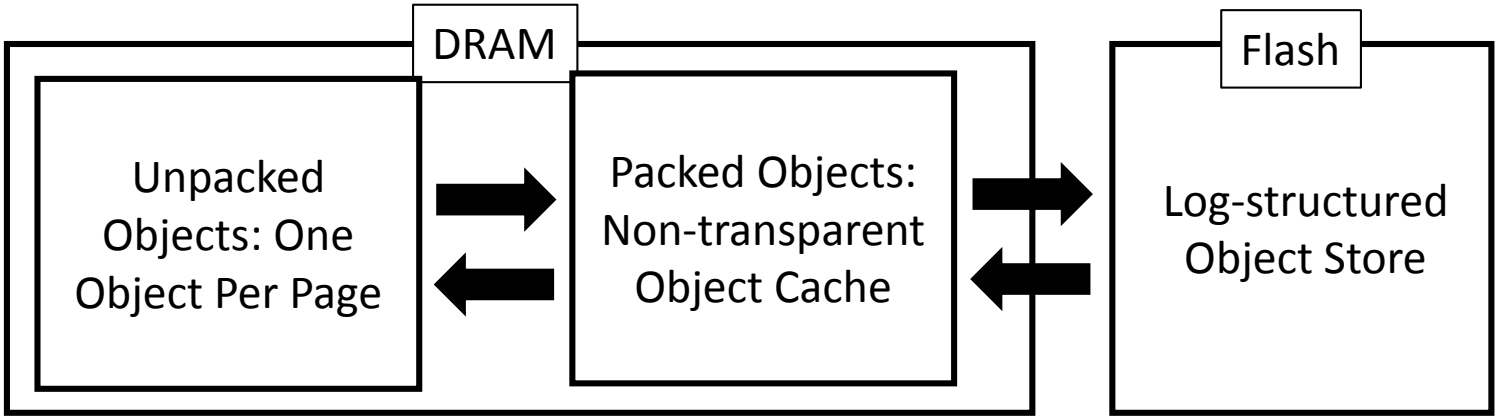

SSDAlloc NSDI'11

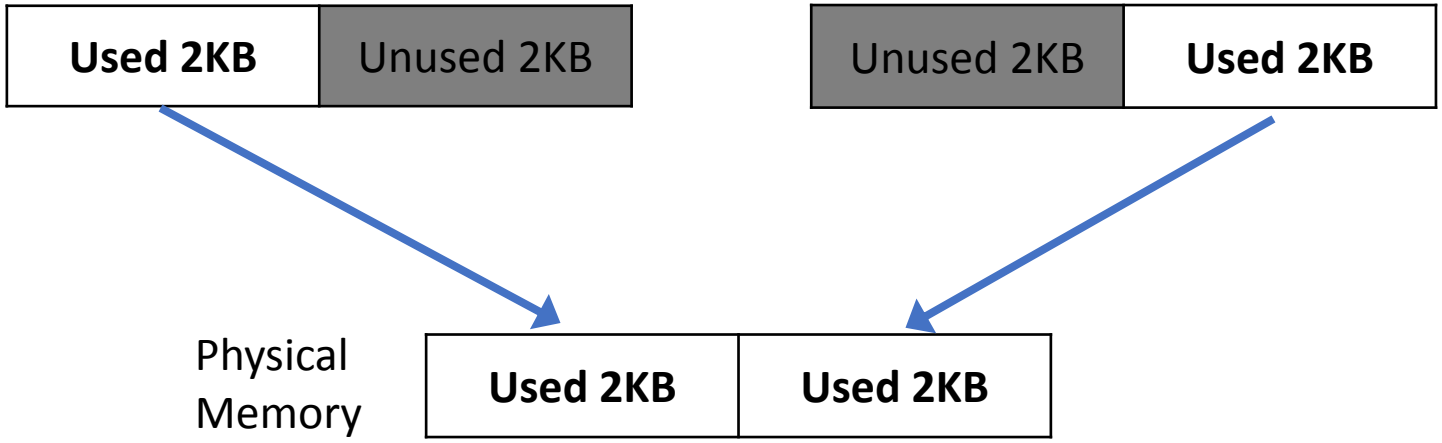

Chameleon TRIOS'11

Reducing the Software Overhead

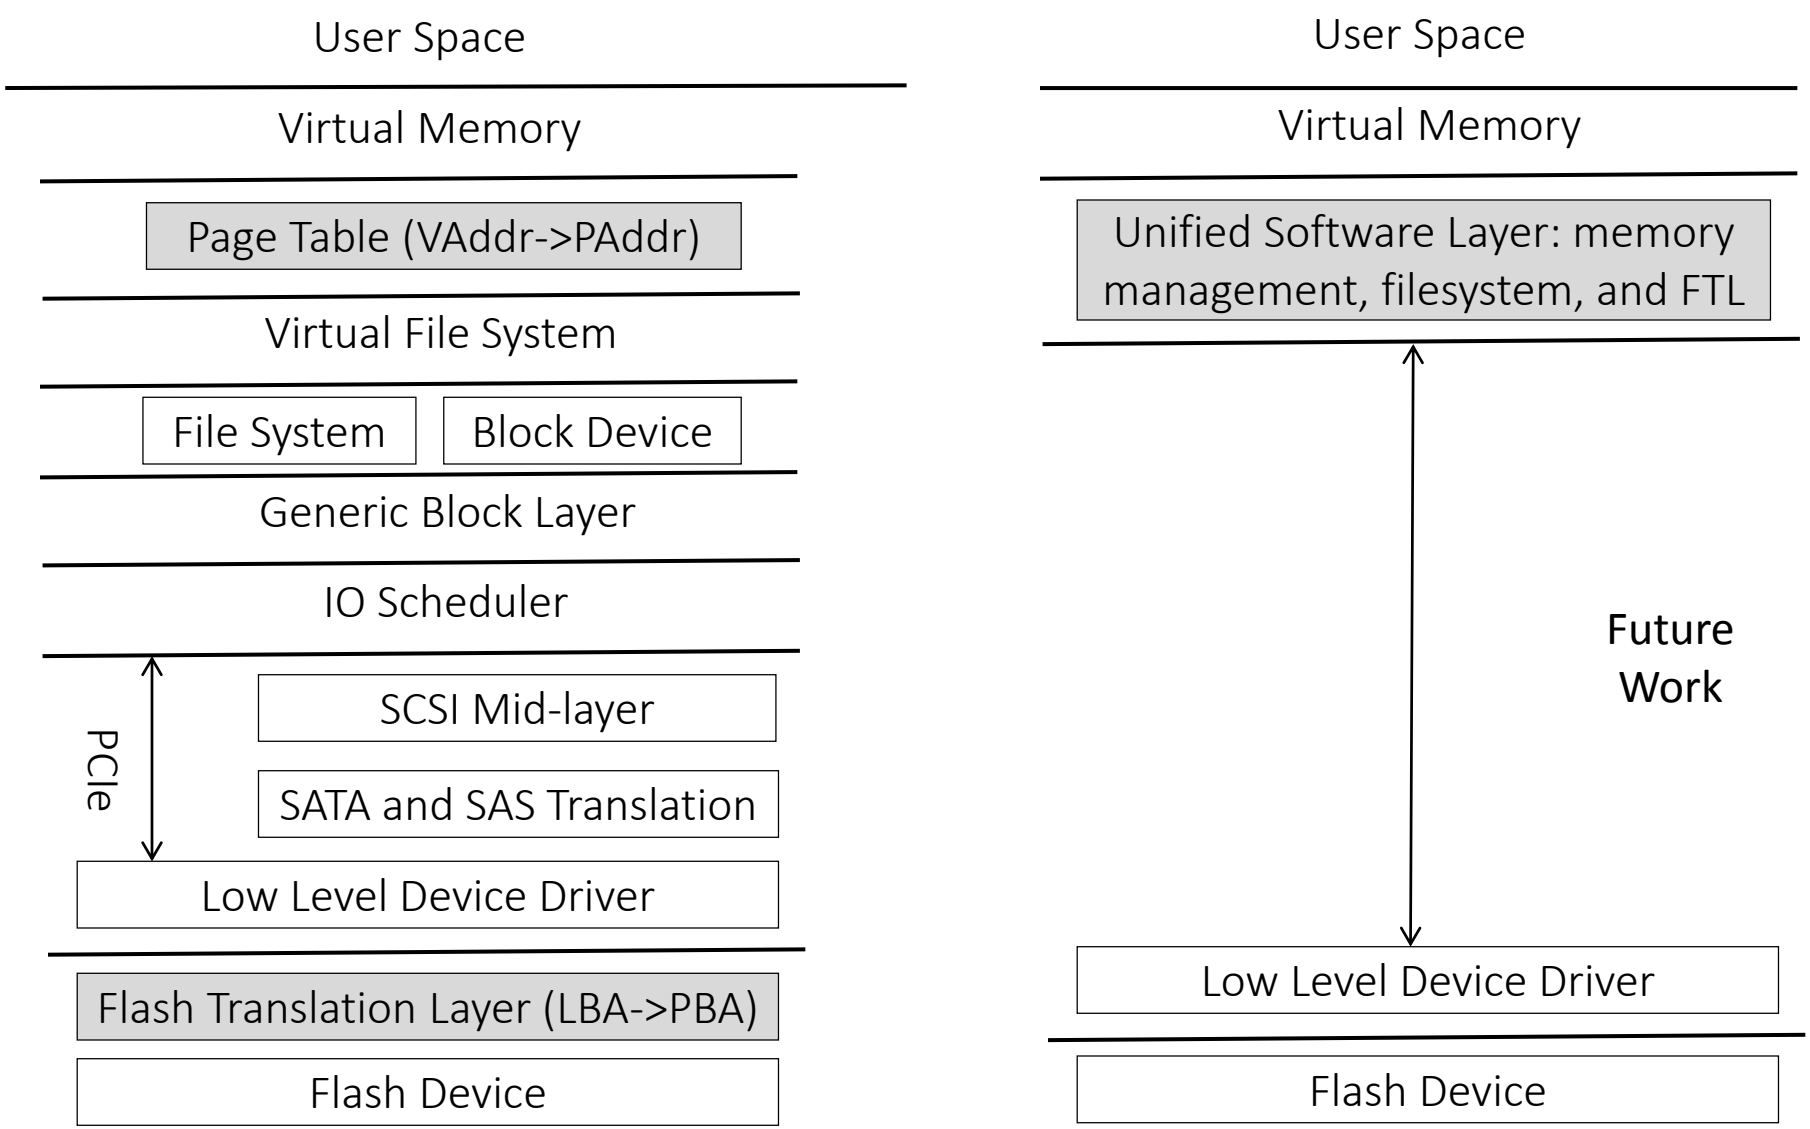

Questions?

- Flash to augment DRAM
  - Uses sparse virtual memory pages to obtain fine-granular usage information
  - Allows density beyond DRAM scaling, at low cost
- Reduce software overhead
  - Move data between DRAM and SSD quickly
  - Unification of software layers
  - Not lose the functionality of any

Thanks!

anirudh.badam@microsoft.com