NON-VOLATILE MEMORY DBMS

Joy Arulraj

15-799 : Final Presentation

Carnegie Mellon

TRADITIONAL DBMS

- Long Transactions
  - —Interactive workload
- Small Memory Capacity
  - —Disk latency

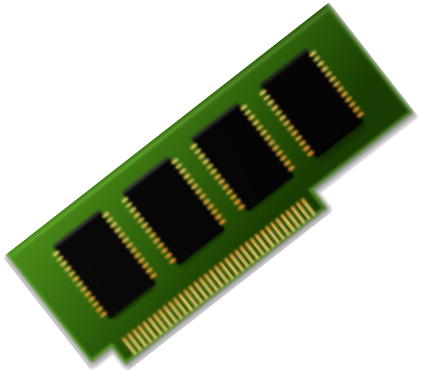

REALITY CHECK

- Short Transactions
- Repetitive Workloads
- Large Memory Capacity

MAIN-MEMORY DBMS

MAIN-MEMORY DBMS

- Disk used only for logging/recovery
- High-throughput OLTP

CHALLENGES

- DRAM SCALING LIMIT
  - Reliable sensing
- RECOVERY LATENCY
  - Throughput

NVM DBMS

PROPERTIES

	DRAM	NVM	SSD	DISK
READ LATENCY	1x	2-5x	500x	10^5x
WRITE LATENCY	1x	2-5x	5000x	10^5x
PERSISTENCE	✗	✓	✓	✓
SCALABLITY	✗	✓	✓	✓
BYTE-LEVEL ACCESS	✓	✓	✗	✗

STORAGE CHOICES

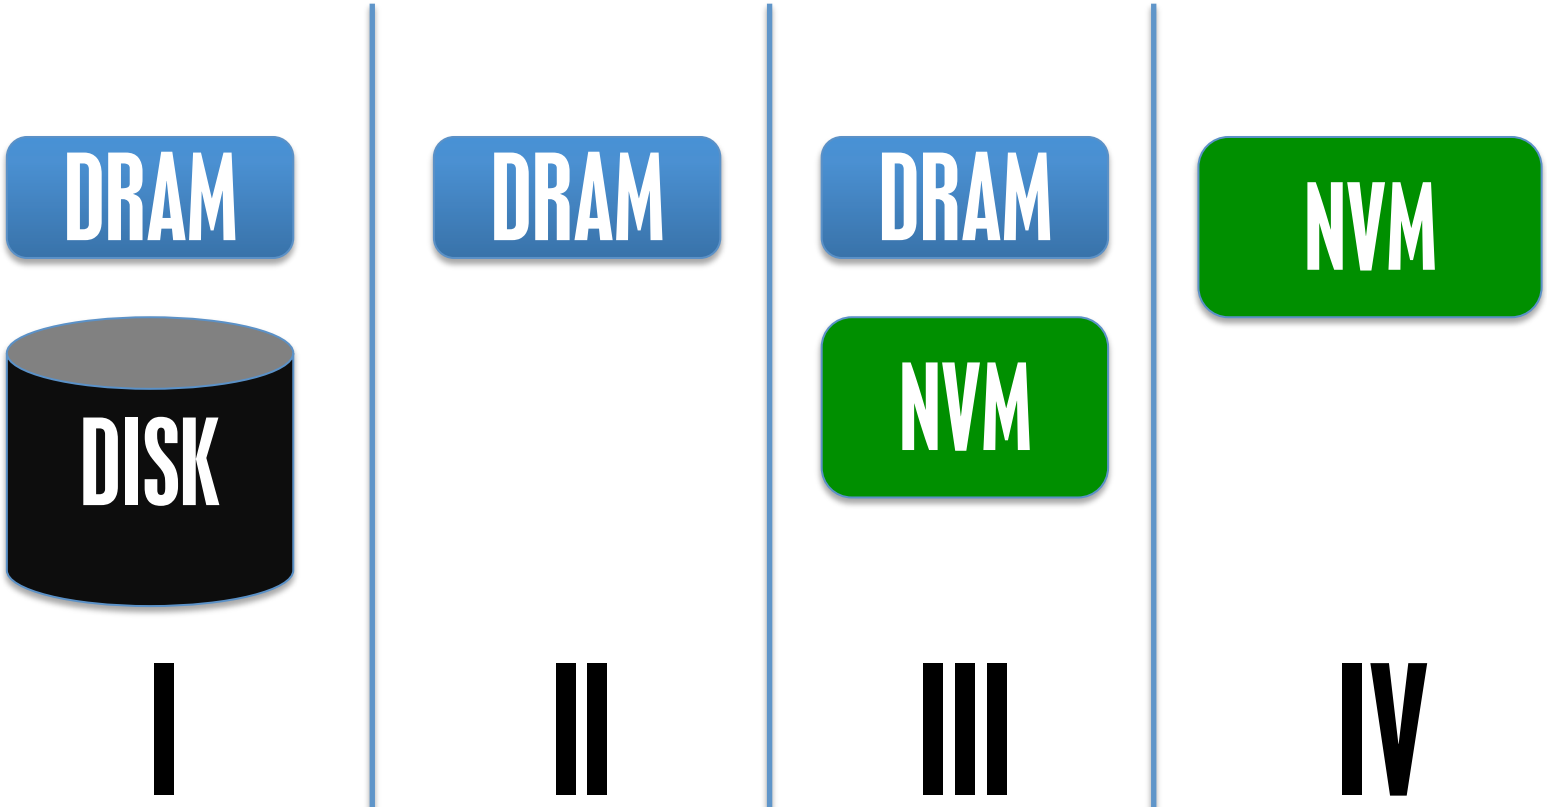

ENVIRONMENT

- INTEL NVM EMULATOR
  - –Instrumented motherboard
- PERSISTENT MEMORY FILE SYSTEM
  - –MMAP interface to PM

NVM HARDWARE EMULATOR

- READ LATENCY
  - –LLC Miss Stalls
- WRITE BANDWIDTH
  - –Throttling in memory controller

PM FILE SYSTEM

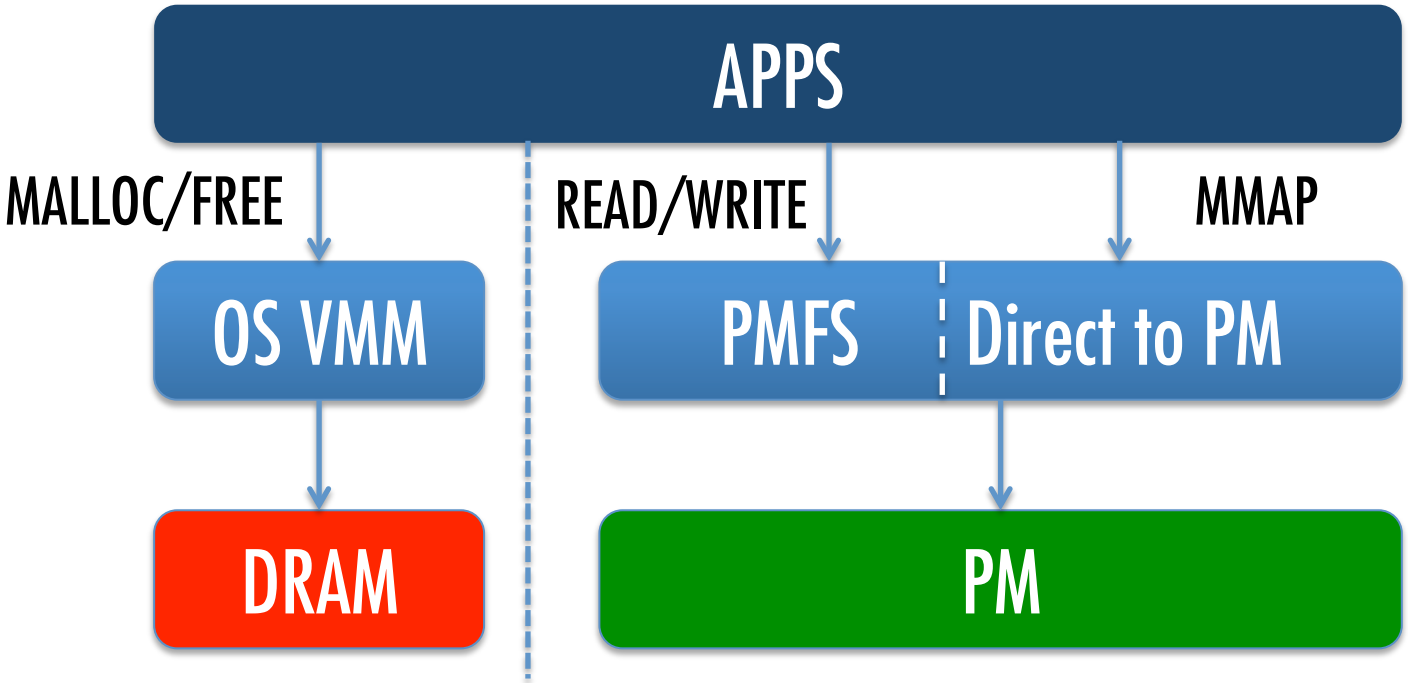

GOALS

- MMAP-BASED STORAGE MANAGER
- EVALUATION ON NVM EMULATOR
- MOVE INDEX STORAGE TO NVM

IMPLEMENTATION

- STORAGE MANAGER
  - —H-Store Table
  - —Per-table memory mapped file
  - —Metadata for recovery

IMPLEMENTATION

- STORAGE MANAGER
  - Pool Storage
  - String Pool (VARCHAR)

IMPLEMENTATION

- STL ALLOCATOR
  - –Index Storage
  - –On top of Storage Manager
  - –Ordered and Unordered map

EXPERIMENTS

SETUP

- INTEL NVM EMULATOR
  - 62 GB DRAM
- YCSB BENCHMARK
  - Zipfian distribution
  - Read Only (100% Reads)
  - Update Heavy (50% Updates, 50% Reads)

READ-ONLY WORKLOAD

2X DRAM LATENCY

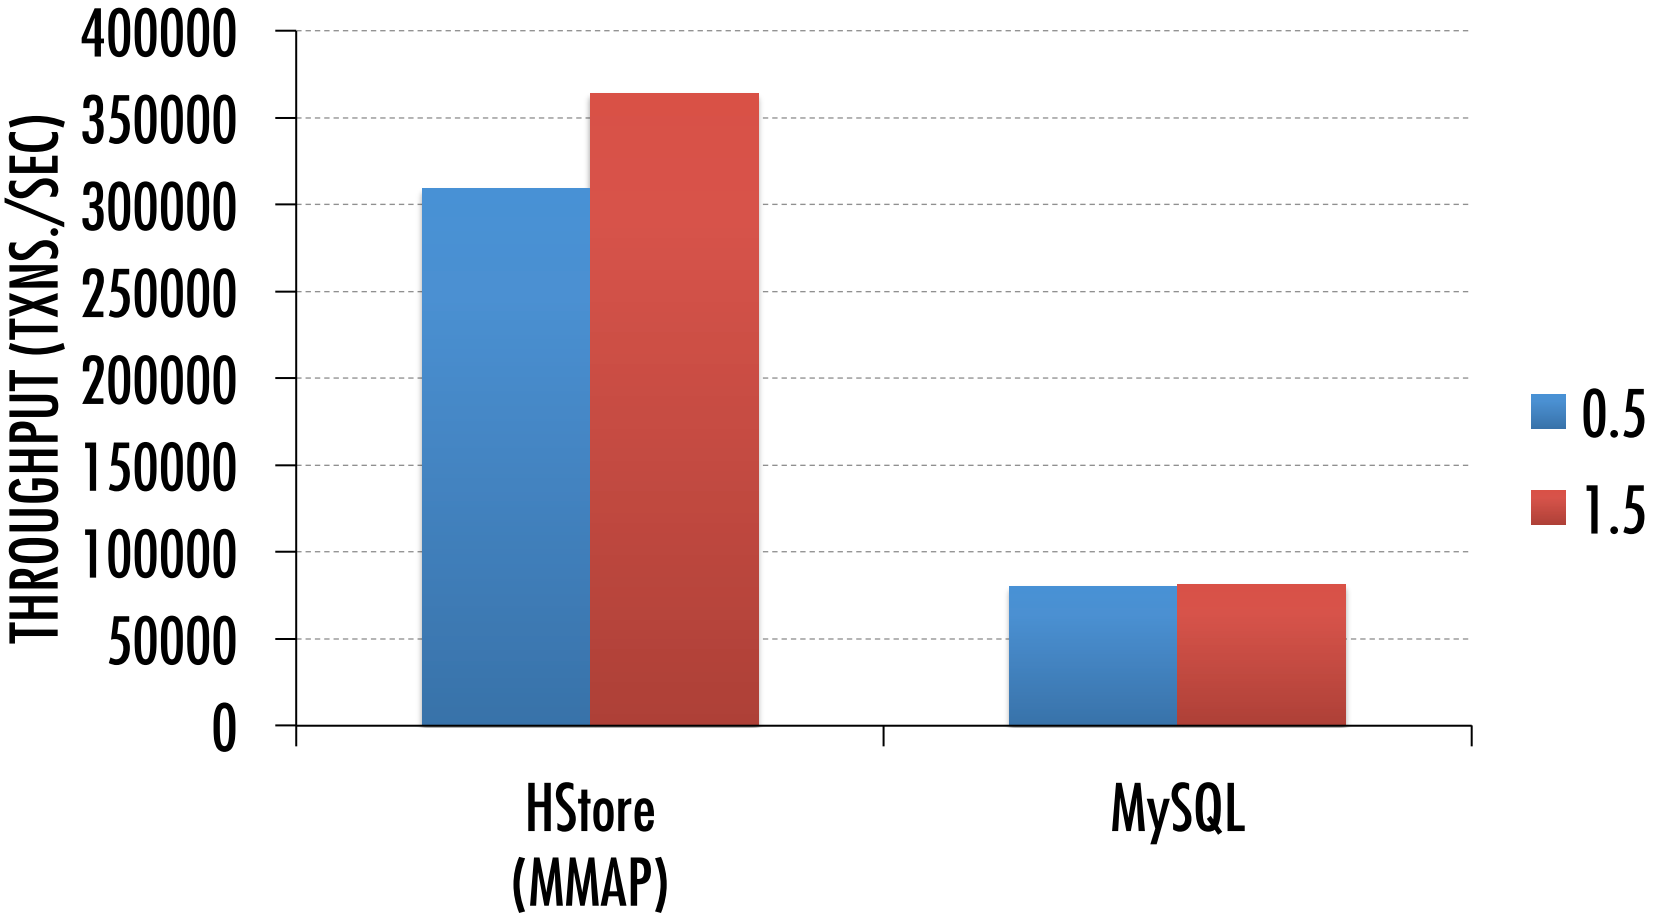

16X DRAM LATENCY

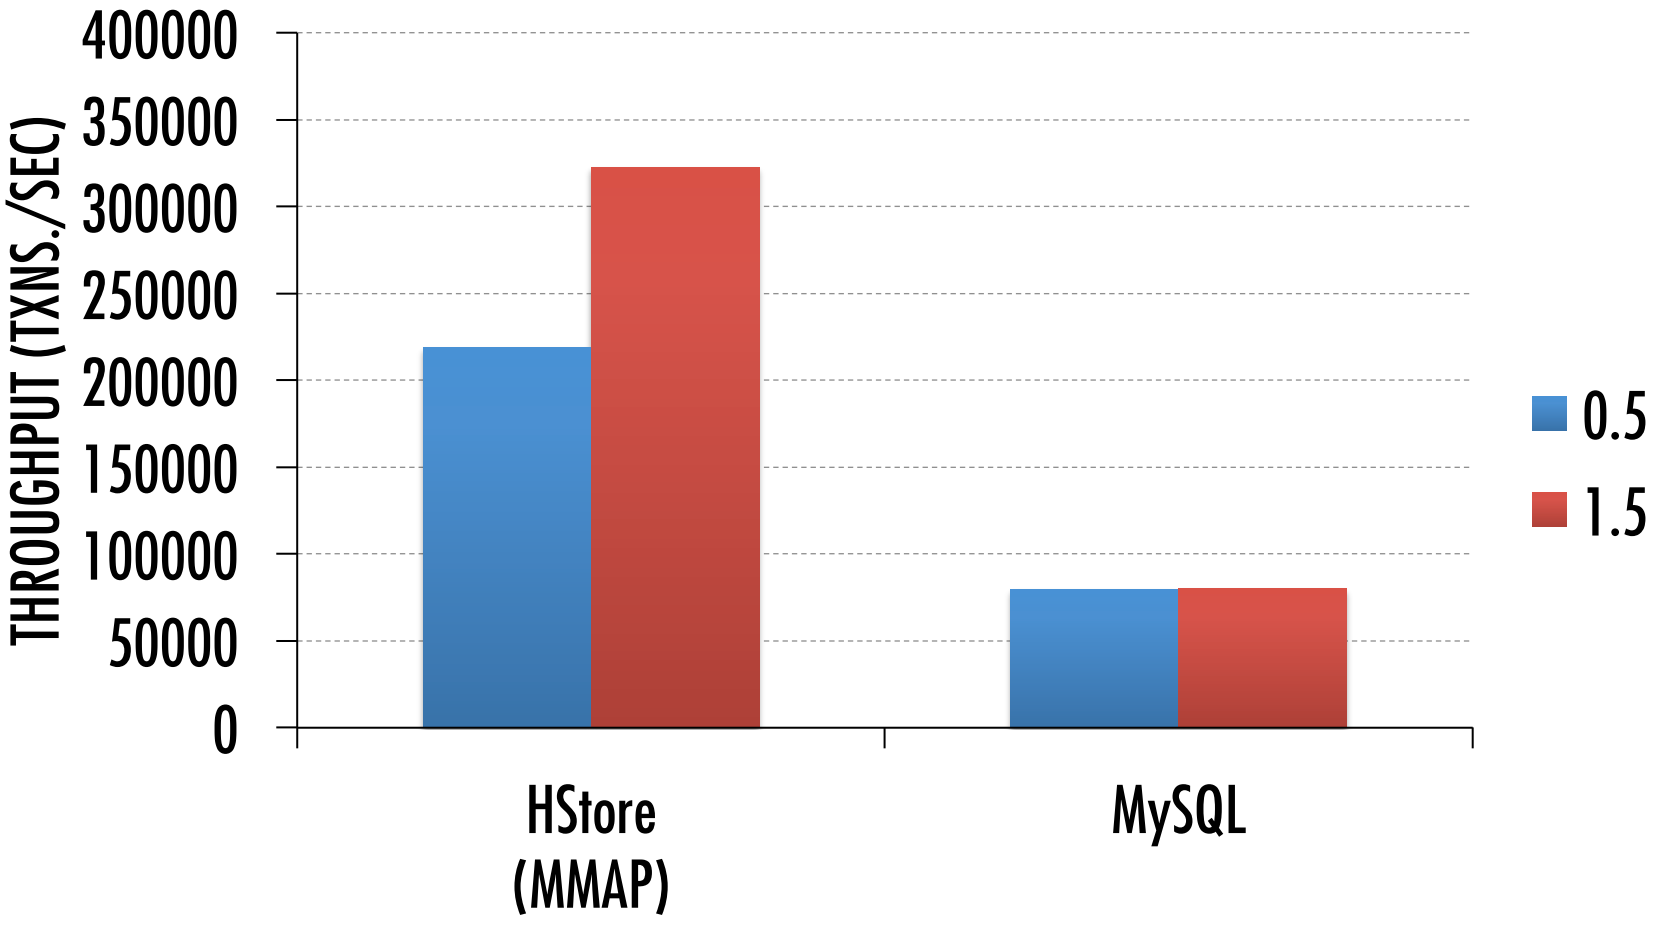

UPDATE-HEAVY WORKLOAD

IMPACT OF NVM LATENCY

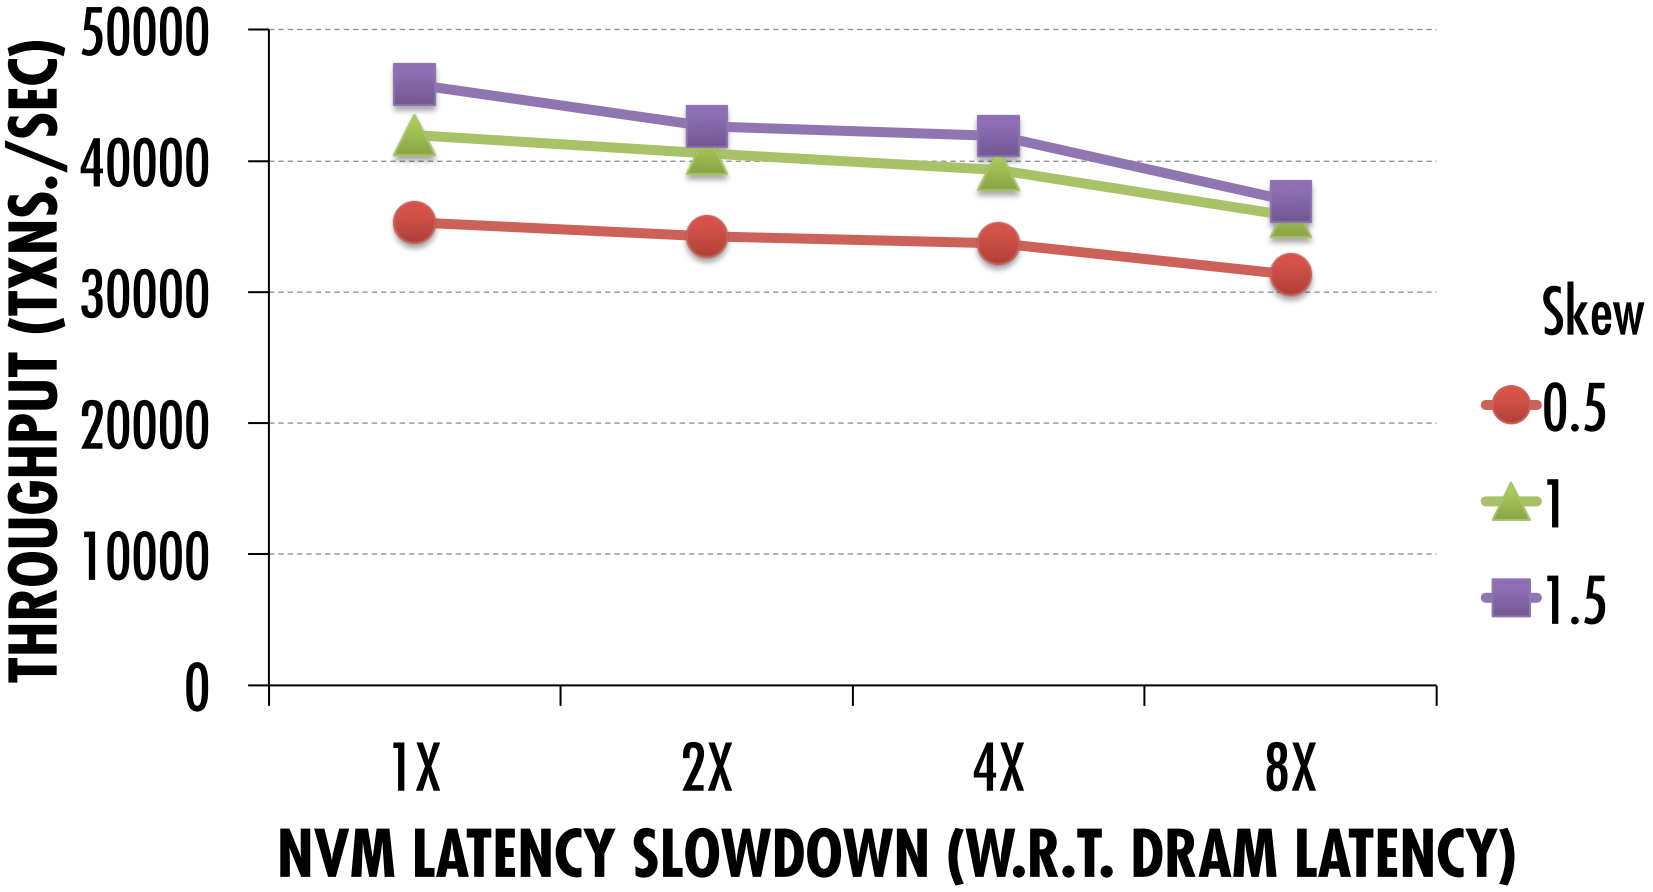

COMPARISON WITH DISK DBMS

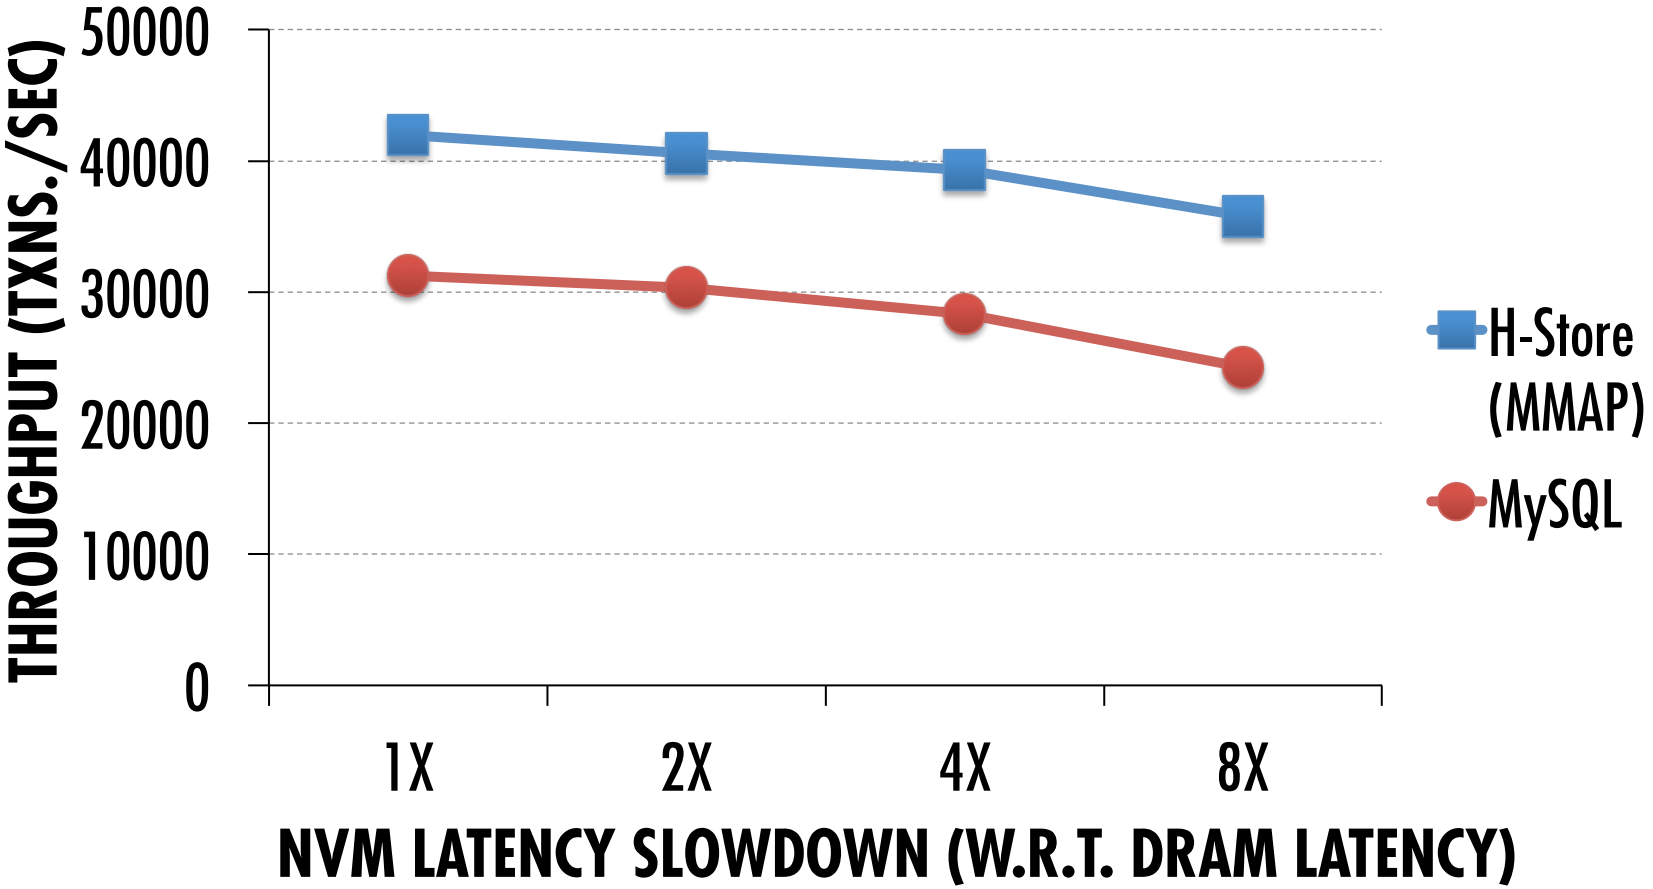

CONCLUSION

- Throughput comparison with MySQL
  - —4.5X on read-only workloads
  - —1.5X on update-heavy workloads
- Update-heavy workload
  - —msync overhead

CONCLUSION

- A new design ?
  - Recovery
  - Concurrency Control

IV

THANKS !