A Case for Efficient Hardware/Software Cooperative Management of Storage and Memory

Justin Meza[1], Yixin Luo[1], Samira Khan[1,3], Jishen Zhao[2], Yuan Xie[2,4], Onur Mutlu[1]
[1]Carnegie Mellon University, [2]Pennsylvania State University, [3]Intel Labs, [4]AMD Research

PROBLEM DEFINITION

Observation

- **The energy consumed executing operating system and file system code to access persistent data becomes an important contributor to total energy in future systems**

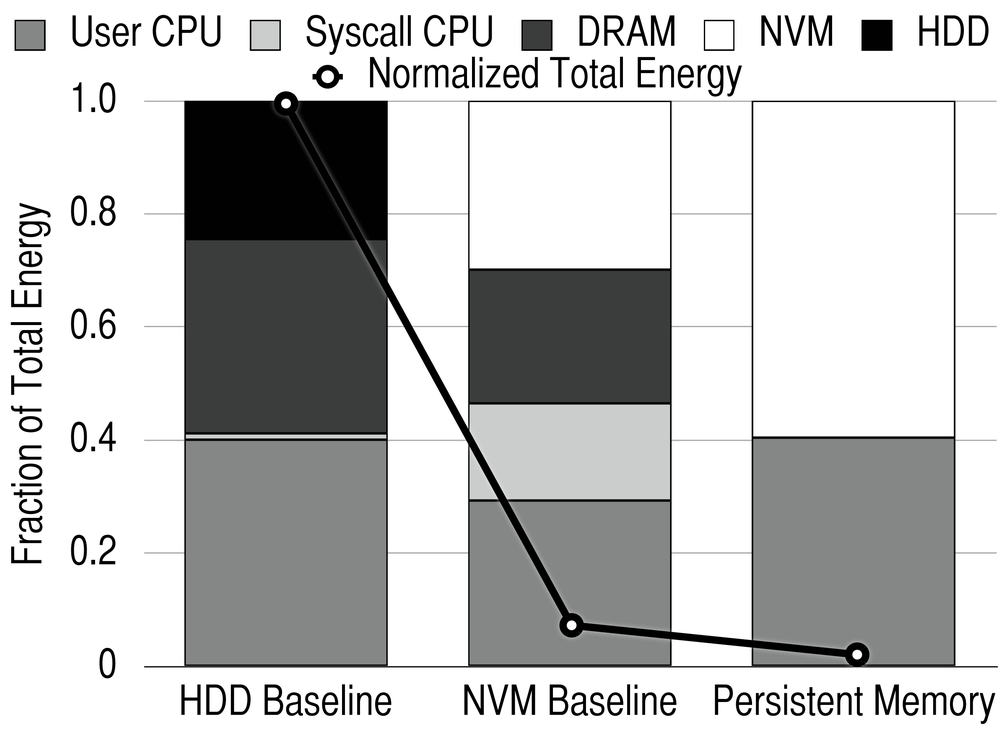

This is because in traditional two-level storage models:

- **Software manipulates data in volatile memory using a load-store interface**
- **Storage systems maintain persistence in non-volatile memory using a file-system interface**

KEY IDEA

- **Coordinate the management of memory and storage in a single address space using hardware-software collaborative designs**

PERSISTENT MEMORY MANAGER

```
1 int main(void) {
2   // data in file.dat is persistent
3   FILE myData = "file.dat";
4   myData = new int[64];
5 }
6 void updateValue(int n, int value) {
7   FILE myData = "file.dat";
8   myData[n] = value; // value is persistent
9 }
```

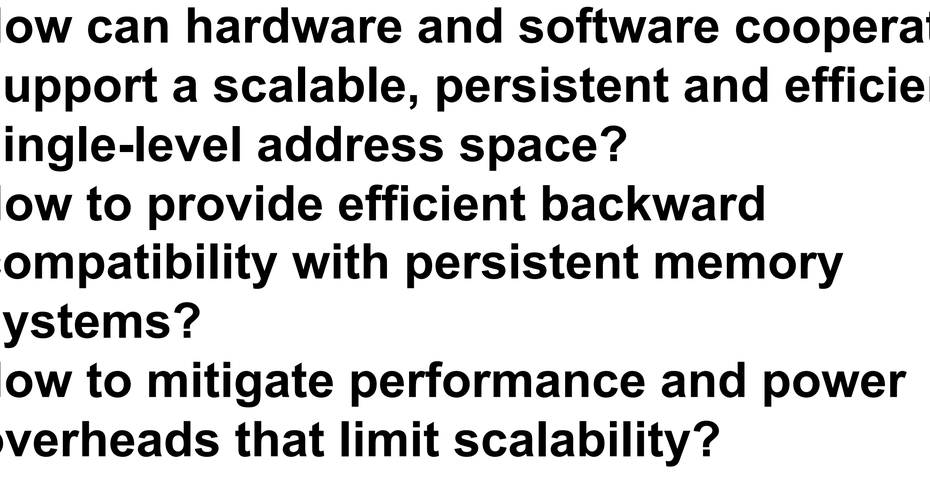

OPPORTUNITIES & BENEFITS

- **Eliminating system calls for file operations**
- **Eliminating file system operations**
- **Efficient data mapping to hybrid memory**
- **Providing security and reliability efficiently**

OPEN RESEARCH QUESTIONS

1. **How to tailor applications for systems with persistent memory?**
2. **How can hardware and software cooperate to support a scalable, persistent and efficient single-level address space?**
3. **How to provide efficient backward compatibility with persistent memory systems?**
4. **How to mitigate performance and power overheads that limit scalability?**

INITIAL EXPLORATION: PERFORMANCE & ENERGY BENEFITS

CPU	16 cores 1.6 GHz Avg 1.41W Peak 149 W
L1	32 KB private
L2	4 MB shared
MEM	DRAM 100 cycles 100mW static NVM 160(480) cycle 45mW static Disk 4ms 6Gbps 1W average

HB: HDD Baseline **NB: NVM Baseline**
HW: HDD without OS/FS overhead **PM: Persistent Memory**

Performance

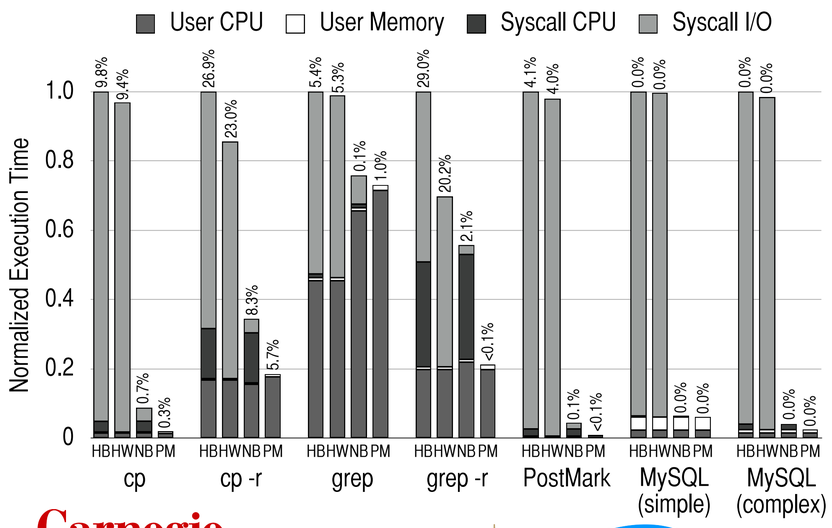

Energy

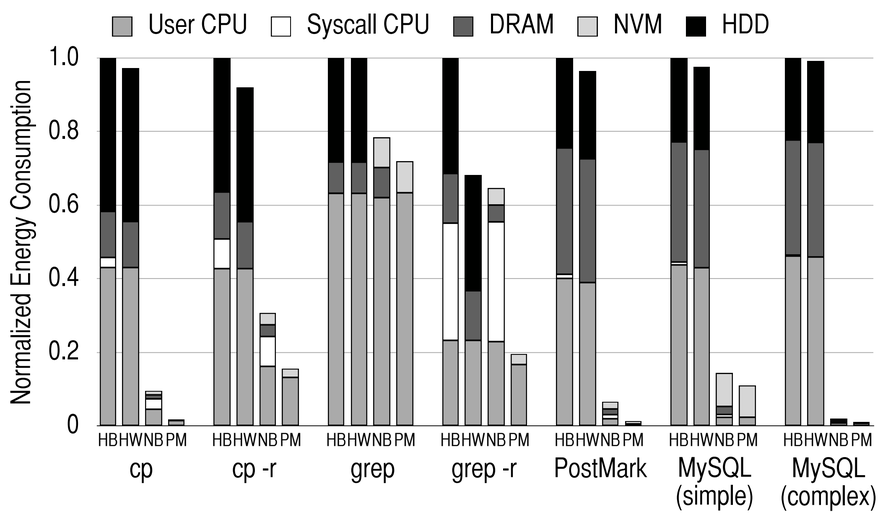

Intel Science & Technology Center for Cloud Computing

Carnegie Mellon University · Georgia Tech · intel · Princeton University · UC Berkeley · University of Washington