SMIO: I/O Similarity Aware Virtual Machine Management in Virtual Desktop Environments

Min Li, Sushil Mantri, Pin Zhou, Ali R. Butt

{limin,sushil,butta}@cs.vt.edu, pinzhou@us.ibm.com

Introduction

- Highly scalable shared storage is a key component of Virtual desktop environments (VDEs)
- I/O reduction techniques identify and remove such duplicated I/O load from the shared storage system to improve efficiency
  - Virtual Machine (VM) images are usually created using the same golden image
  - VMs also deploy a similar set of applications such as anti-virus software
  - The effectiveness of these techniques depends on the amount of duplicated data accessed by the VMs running on a physical host
- Current VM placement solutions can lead to IO reduction inefficiency limiting storage scalability

Background

- I/O Reduction Techniques
  - Capo[FAST11']: maintains a bit-map to eliminate duplicate read requests and a host-side cache to reduce the number of I/O requests to the shared storage system
  - Seacache[MASCOTS12']: integrates host-side cache with storage-side deduplication, eliminates both duplicate read and write path traffic
- virtual machine management techniques
  - A centralized VM Manager maintains the global information such as VM resource allocation information, the VM location information, receive heart beats from PMs
  - It decides the VM placement and migration based on certain metrics. e.g. energy consumption, network traffic

SMIO Architecture

- Design goals:
  - scalability, low overhead, low bandwidth consumption, dynamicity

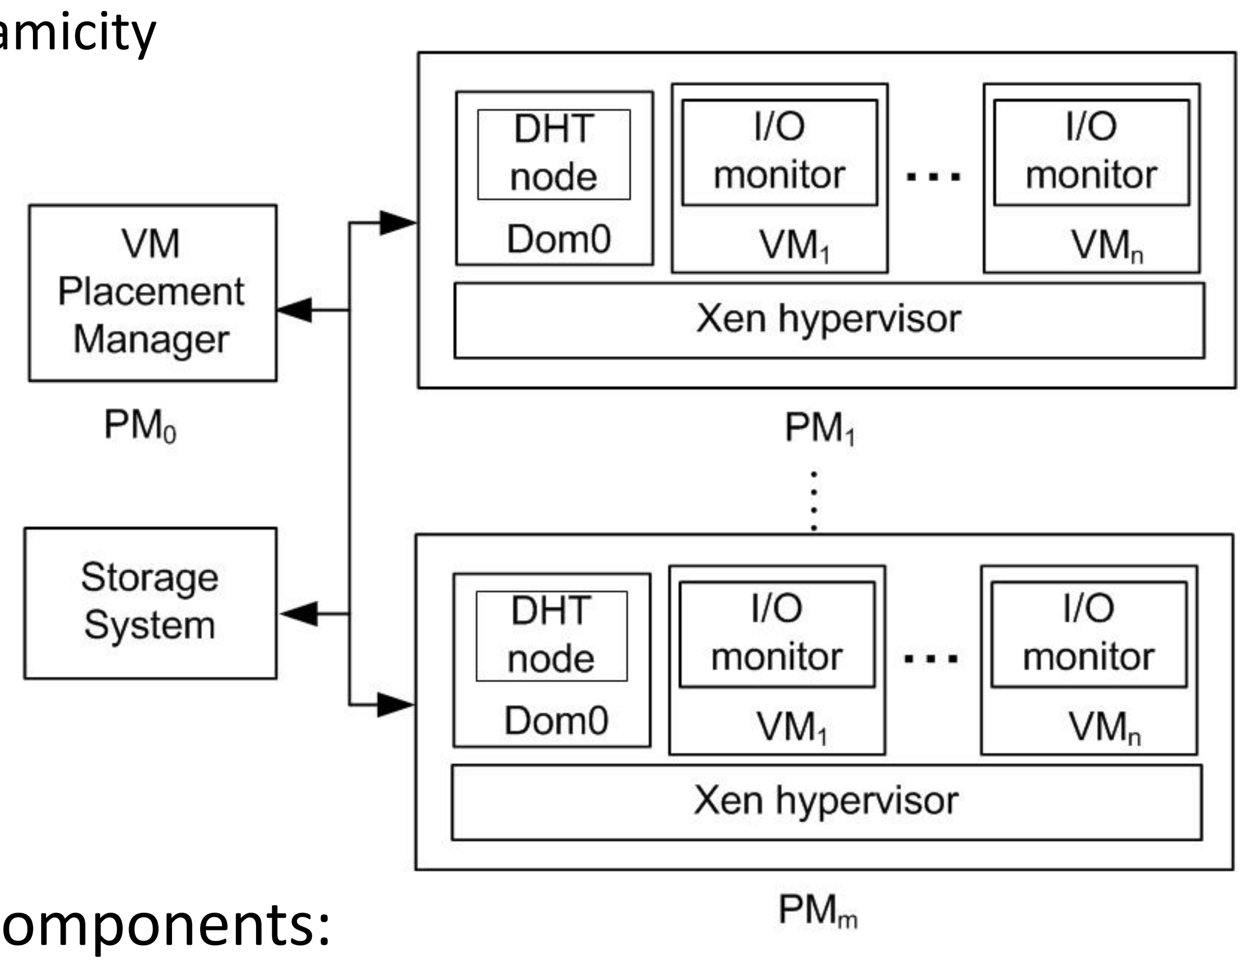

- System components:
  - IO monitor detects I/O similarity among different VMs
  - VM manager utilizes hierarchical clustering to produce a new I/O-similarity-aware VM placement scheme periodically;
  - VM Manager migrates the VMs when benefits of similar VM consolidation outweigh migration cost.

This work was sponsored in part by the NSF under Grants CCF-0746832, CNS-1016793, and CNS-1016408

Hierarchical Clustering

- Data sharing matrix:
  - common unique blocks accessed by both clusters (α_{ij}), the number of total unique blocks accessed by both the clusters (β_{ij})

$$\begin{bmatrix} - & (\alpha_{12},\beta_{12}) & \cdots & (\alpha_{1n},\beta_{1n}) \\ \vdots & & \ddots & \vdots \\ & \cdots & & (\alpha_{(n-1)n},\beta_{(n-1)n}) \end{bmatrix}$$

- Global benefit-cost matrix: *mcost* is the migration cost

$$\begin{bmatrix} - & \gamma_{12} & \cdots & \gamma_{1n} \\ \vdots & & \ddots & \vdots \\ & \cdots & & \gamma_{(n-1)n} \end{bmatrix}$$

$$\gamma_{ij} = \sum_k \alpha_{ij} / \sum_k \beta_{ij} - mcost_{ij} / mcost_{max}.$$

- An example of hierarchical clustering

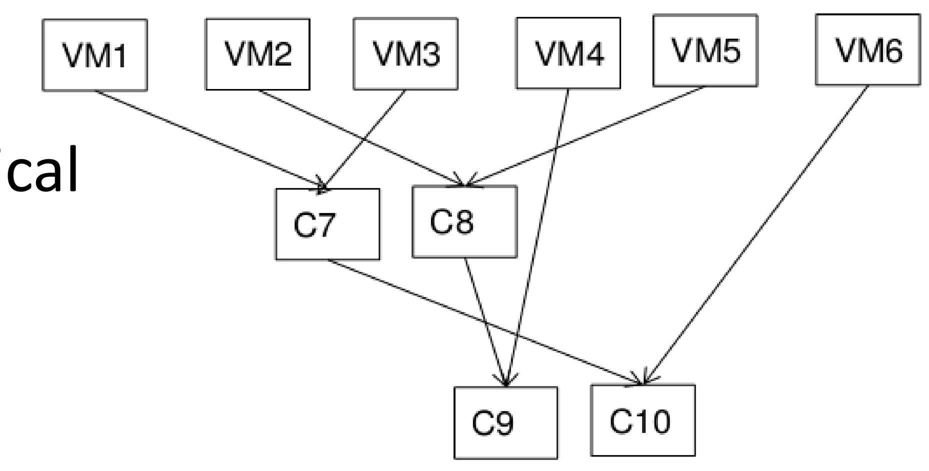

Evaluation

- Impact of SMIO on I/O reduction efficiency
  - Trace driven simulation
  - SMIO can effectively improve the IO reduction efficiency by up to 4.9X.
- The monitoring overhead of SMIO is negligible.

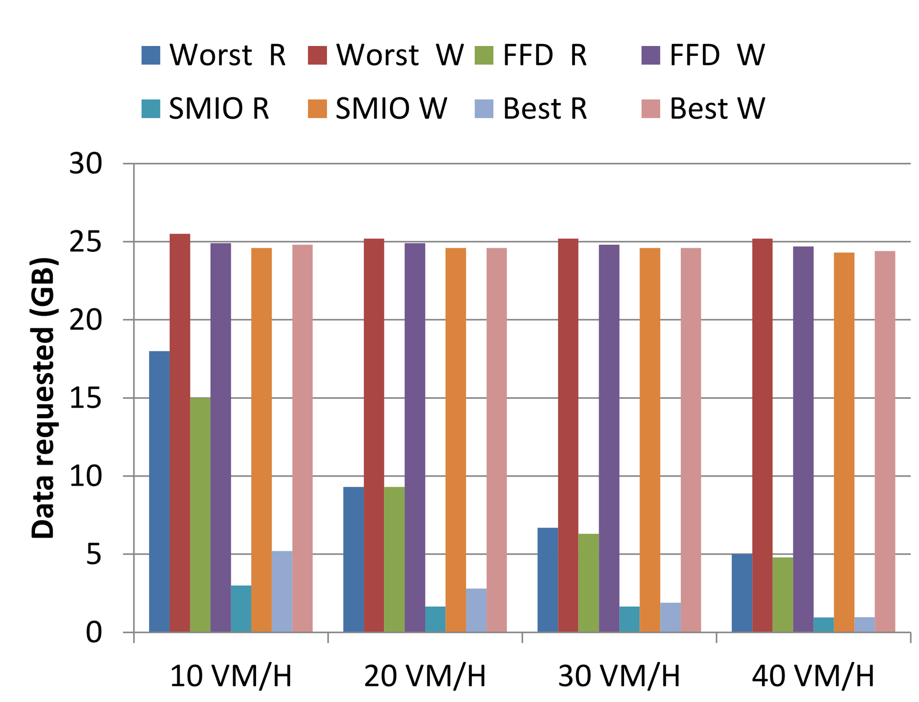

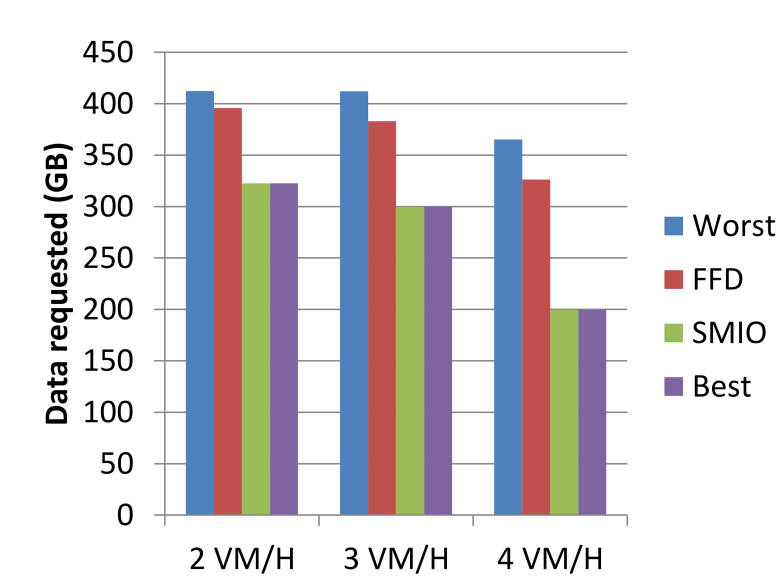

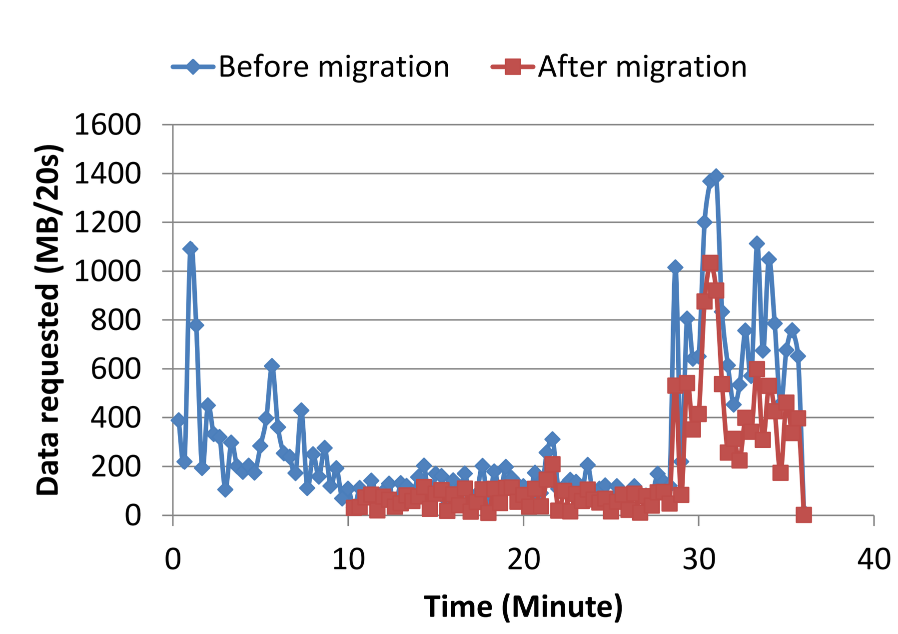

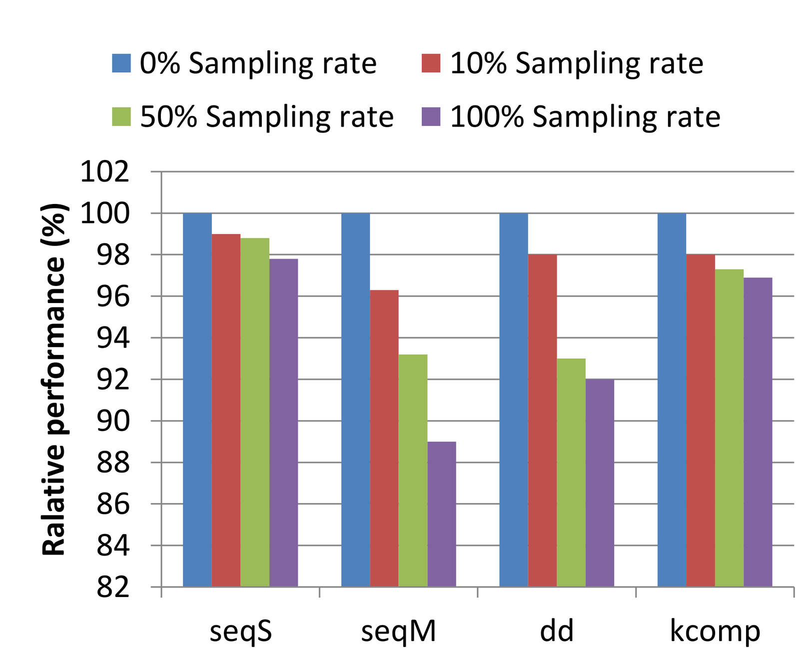

Conclusion and Future Work

- SMIO can improve the effectiveness of I/O reduction techniques by incorporating I/O similarity information
- Investigate the impact of sampling on the clustering effectiveness
- Investigate the scalability of SMIO in terms of increasing number of VMS