ECHO: Decoupling Datacenter Studies from Access to Large-Scale Applications

Christina Delimitrou[1], Sriram Sankar[2], Kushagra Vaid[2], Christos Kozyrakis[1]
[1] Stanford University, [2] Microsoft

Accurate Modeling and Generation of DC Storage Workloads

Storage Modeling Overview

Workload **Modeling** & **Generation** increasingly critical for DC studies because:

- Replay original applications in all storage system configurations is **impractical**
- Datacenter workloads are **not publicly available**
- **Storage system ~20-30% of TCO and power consumption** of the total system
- No concise methodology for modeling and generation of DC applications

Features	IOMeter	DiskSpd
Inter-arrival Time (mean or distribution)	×	✓
Intensity Knob	×	✓
Spatial Locality	×	✓
Trace Replay	×	✓
Different Levels of Granularity	×	✓
File Accesses*	×	✓

Comparison: IOMeter – DiskSpd

Model

State Diagram-Based Markov Chain:

- **State**: Logical block range on disk (LBNs)
- **Transition**: Probability of switching between block ranges
- **Stats**: rd/wr, rnd/seq, block size, inter-arrival time

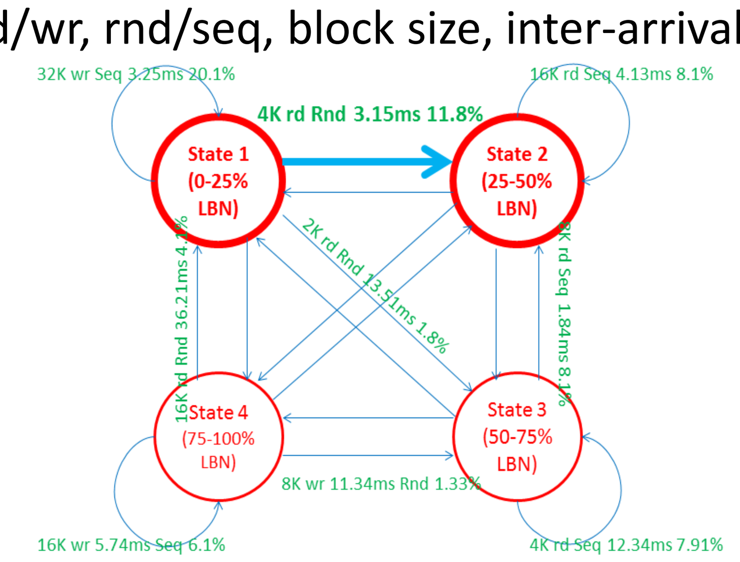

Extend to hierarchical representation → optimal number of levels per application

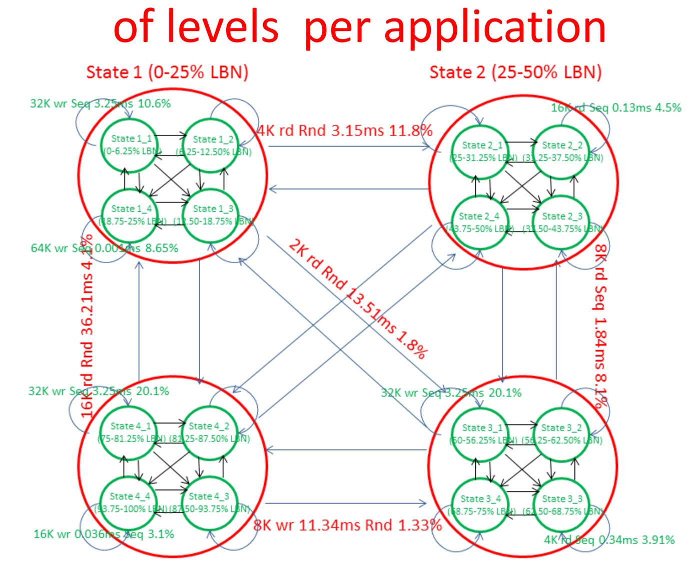

Reduce Model Complexity: Hierarchical rather than **Flat** representation → Spatial locality within rather than across states

Validation – Storage Model

Two Step Process:

1. Collect traces on production systems → create workload models
2. Generate synthetic workloads from models → validate I/O and performance metrics

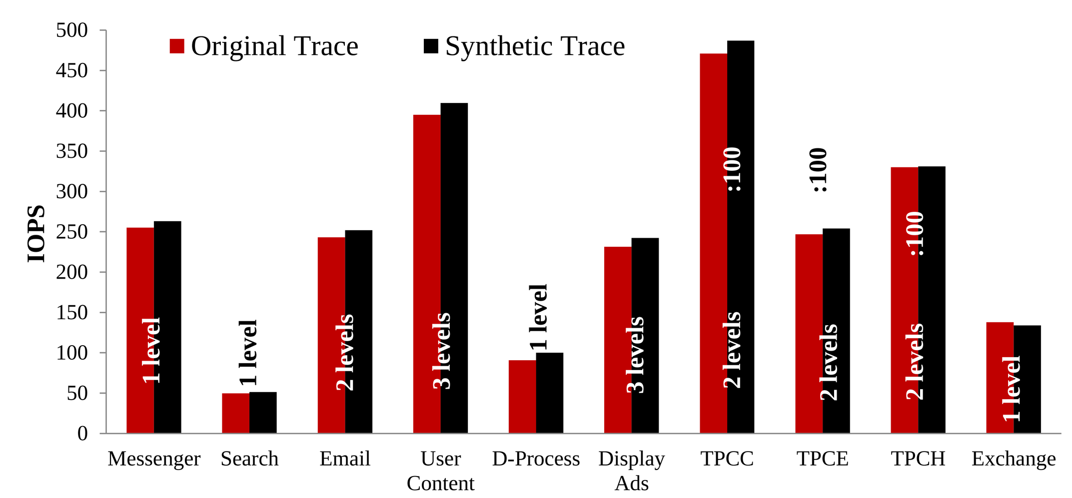

Less than 5% deviation between original and synthetic workloads

Use Cases

- **SSD Caching**
  1. Most I/O aggregated in space for DC apps
  2. Progressive SSD caching (0-32GB)
  3. Use models to quantify benefit in performance

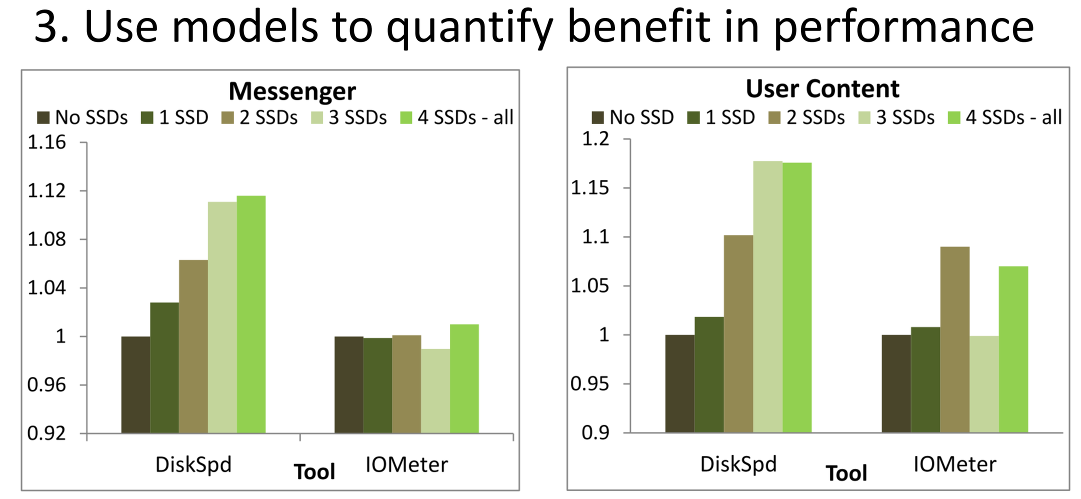

DiskSpd → Accurate performance predictions
IOMeter → No or inconsistent speedup

- **Defragmentation Benefits**
  1. Most (> 80%) disk I/O accesses are random
  2. Potential for performance/energy benefits in defragmentation
  3. Use models to predict benefits of defragmentation

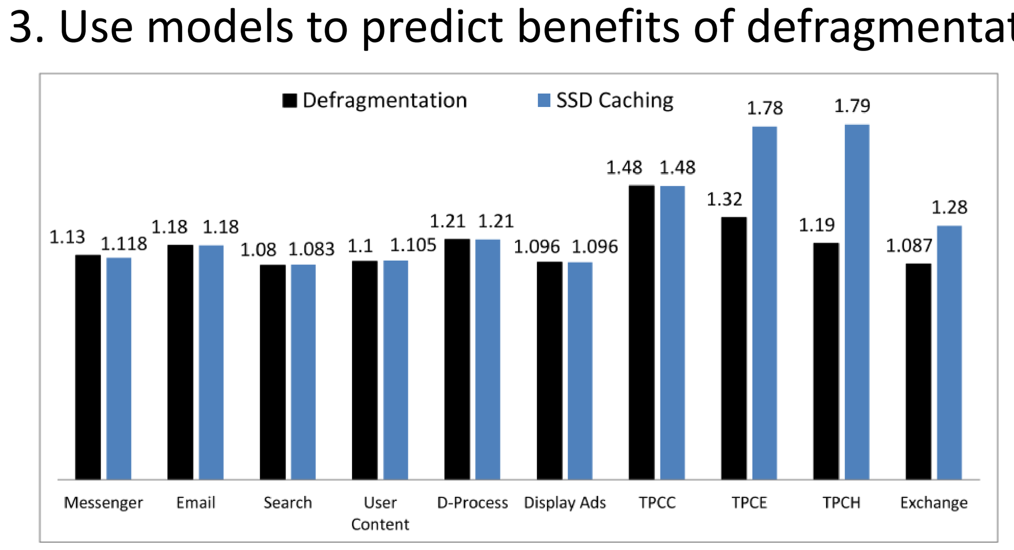

Modeling and Generation of DC Network Workloads in Time and Space

Network Modeling Overview

Network Modeling is critical for network-related DC studies.

DC network patterns are very hard to recreate empirically → need for validated analytical models with provable accuracy and error bounds

No existing robust method to model and generate representative DC network workloads

Two-fold proposal:

1. **Distribution Fitting Model**: Capture temporal variations in load → represent with known mathematical distributions

2. **Markov Chain Spatial Model**: Capture patterns in spatial locality of network load → recreate accurate network traffic in DC

Validation – Temporal Model

Distribution fitting model:

Input → Network per server bandwidth traces
Output → Superposition of known distributions

Model = Gaussian + Exponential + Gaussian + Gaussian + Constant

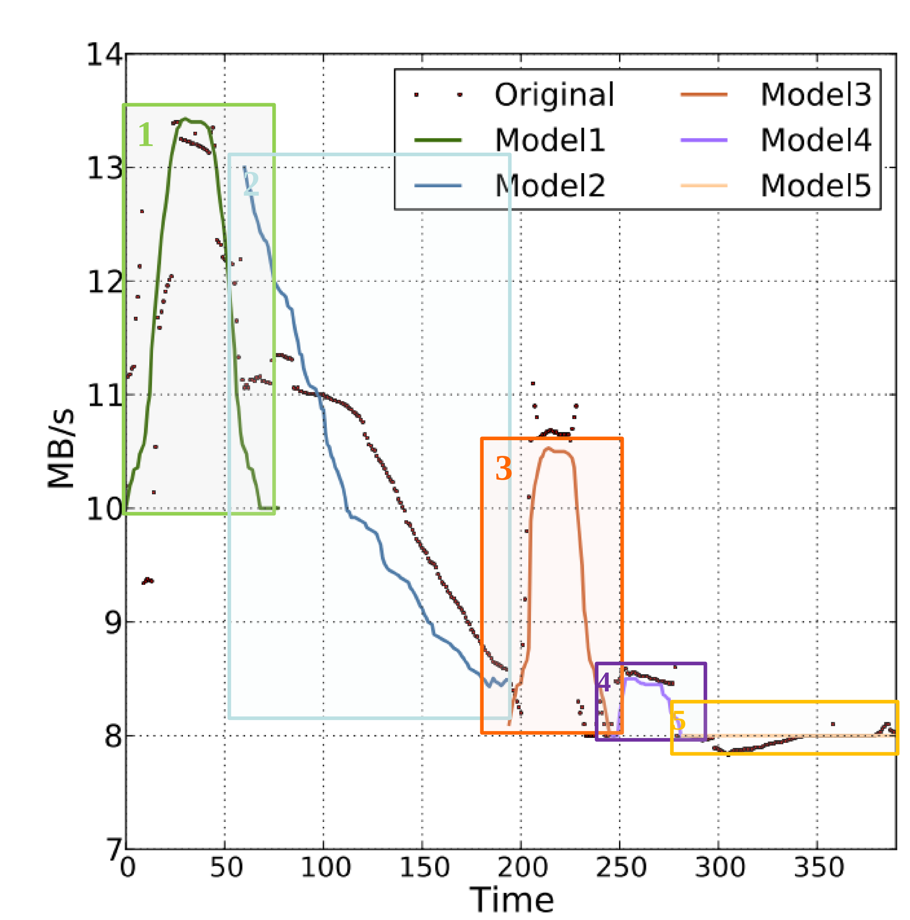

Less than 4.8% deviation between original workload and model

Validation – Spatial Model

Markov Chain Spatial Model:

1. Hierarchical representation of network traffic pattern in DC
2. Capture spatial patterns in network activity

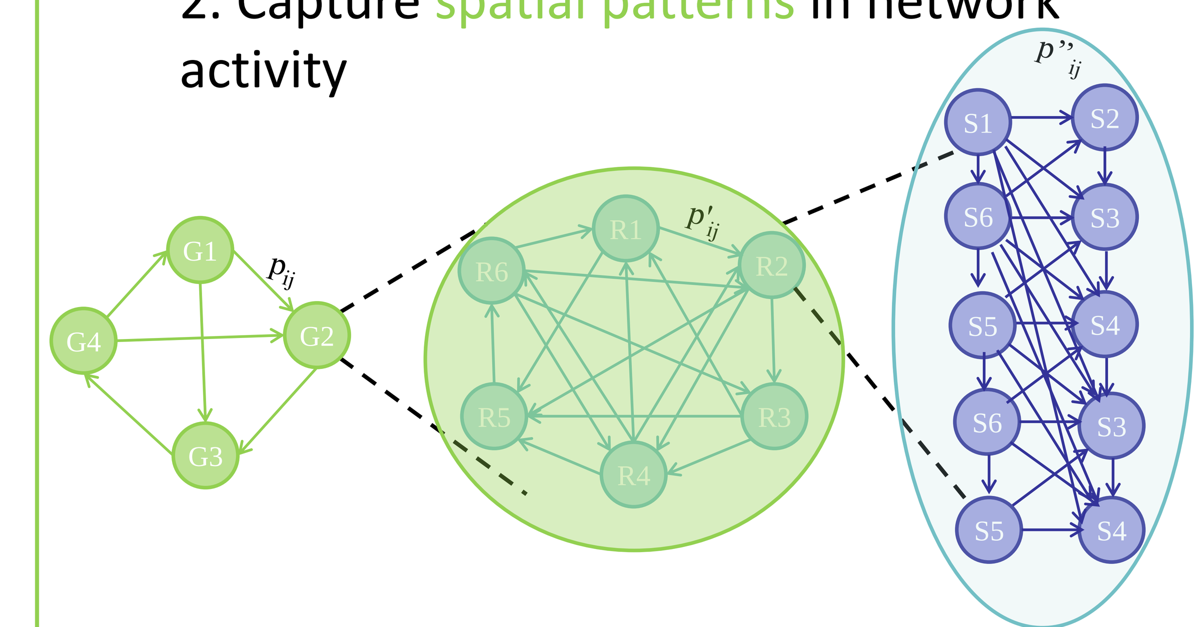

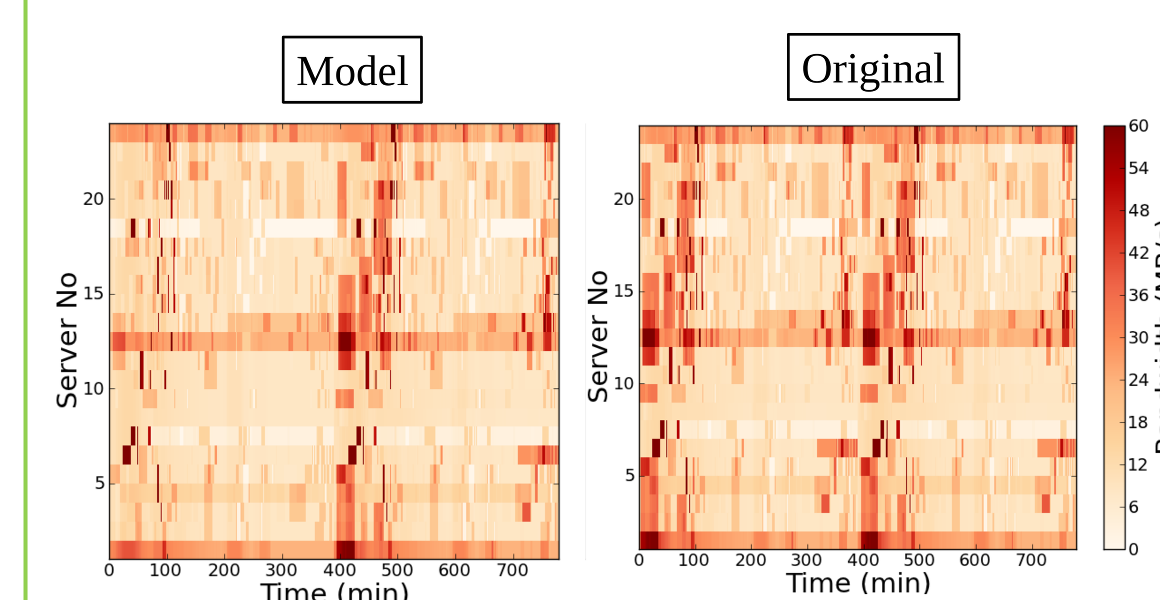

Less than 6.8% deviation between original network traffic and model

Conclusions and Future Work

- **Storage Model:**
  1. Accurate **representation** of DC access patterns
  2. **Generation** of representative DC storage workloads
  3. **Motivated two important challenges** in DC design
- **Network Model:**
  1. Two concise & accurate modeling approaches (temporal and spatial model)
  2. **High accuracy** in generation process
  3. Recreation of **actual DC network traffic**

Future Work:

1. Expand similar methodology to other system parts
2. Integrate the different models
3. Perform additional system studies