

# **Impact of Database Scaling on DSS Workload Characteristics on SMP Systems**

Ramendra K. Sahoo,  
Krishnan Sugavanam  
Ashwini Nanda

IBM T.J. Watson Research Center


# Motivation

- Question: How to profile and evaluate realistic commercial workloads?
- Are the workload characteristics independent of database scaling ?



# Workload Evaluation Options

- Simulation with scaled down workloads.
- Realistic workloads with hardware emulation.



# Setup Environment



System Under  
Test-



MemorIES



L3 PC  
Console

# Tools

- MemorIES -I.
- pmcount.
- Standard UNIX (AIX) tools.



# Experimental Platform

- 12-way IBM S7A SMP
  - MemorIES board replaces 4 processors.
  - 262 MHz Northstar processors.
  - 8MB L2 caches.
  - 24GB main memory.
  - More than 1TB disk space.
  - L2 size/associativity can be varied from 8MB 4-way to 1MB direct-mapped.

# Preliminary Results

- Impact of DSS workload on L3 behavior.
- Effect of DB size on L3 cache miss characteristics.
- Impact of DB size on L2 miss characteristics.


# Impact of DSS Workload on L3 Miss Behavior






# Impact of DSS Workload on L3 Miss Behavior (Contd.)

**Load Avg. for 30GB TPC-H Power Run**

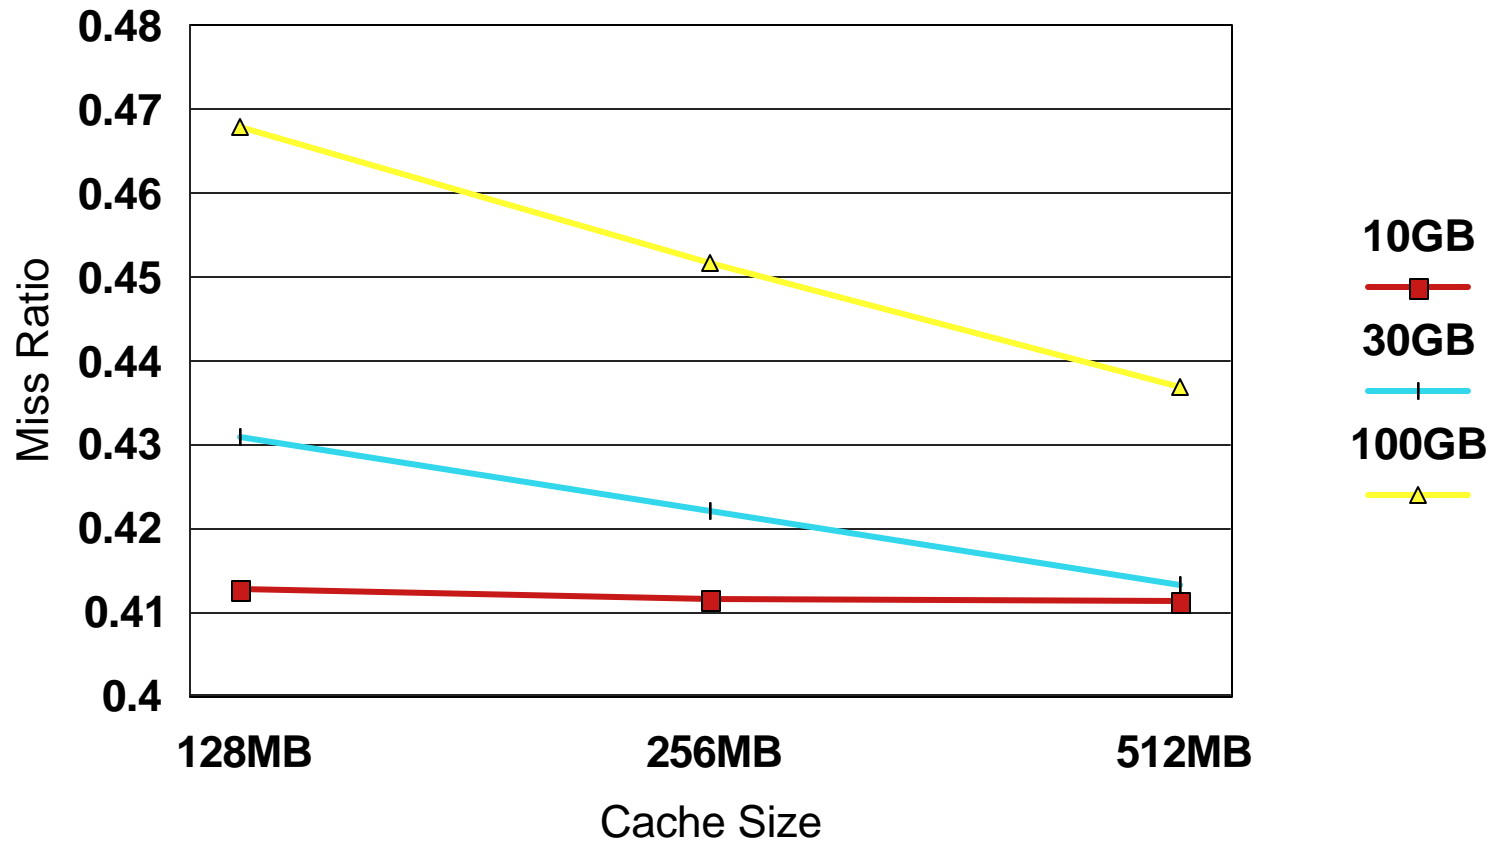


**Load Avg. for 30GB TPC-H Throughput Run**



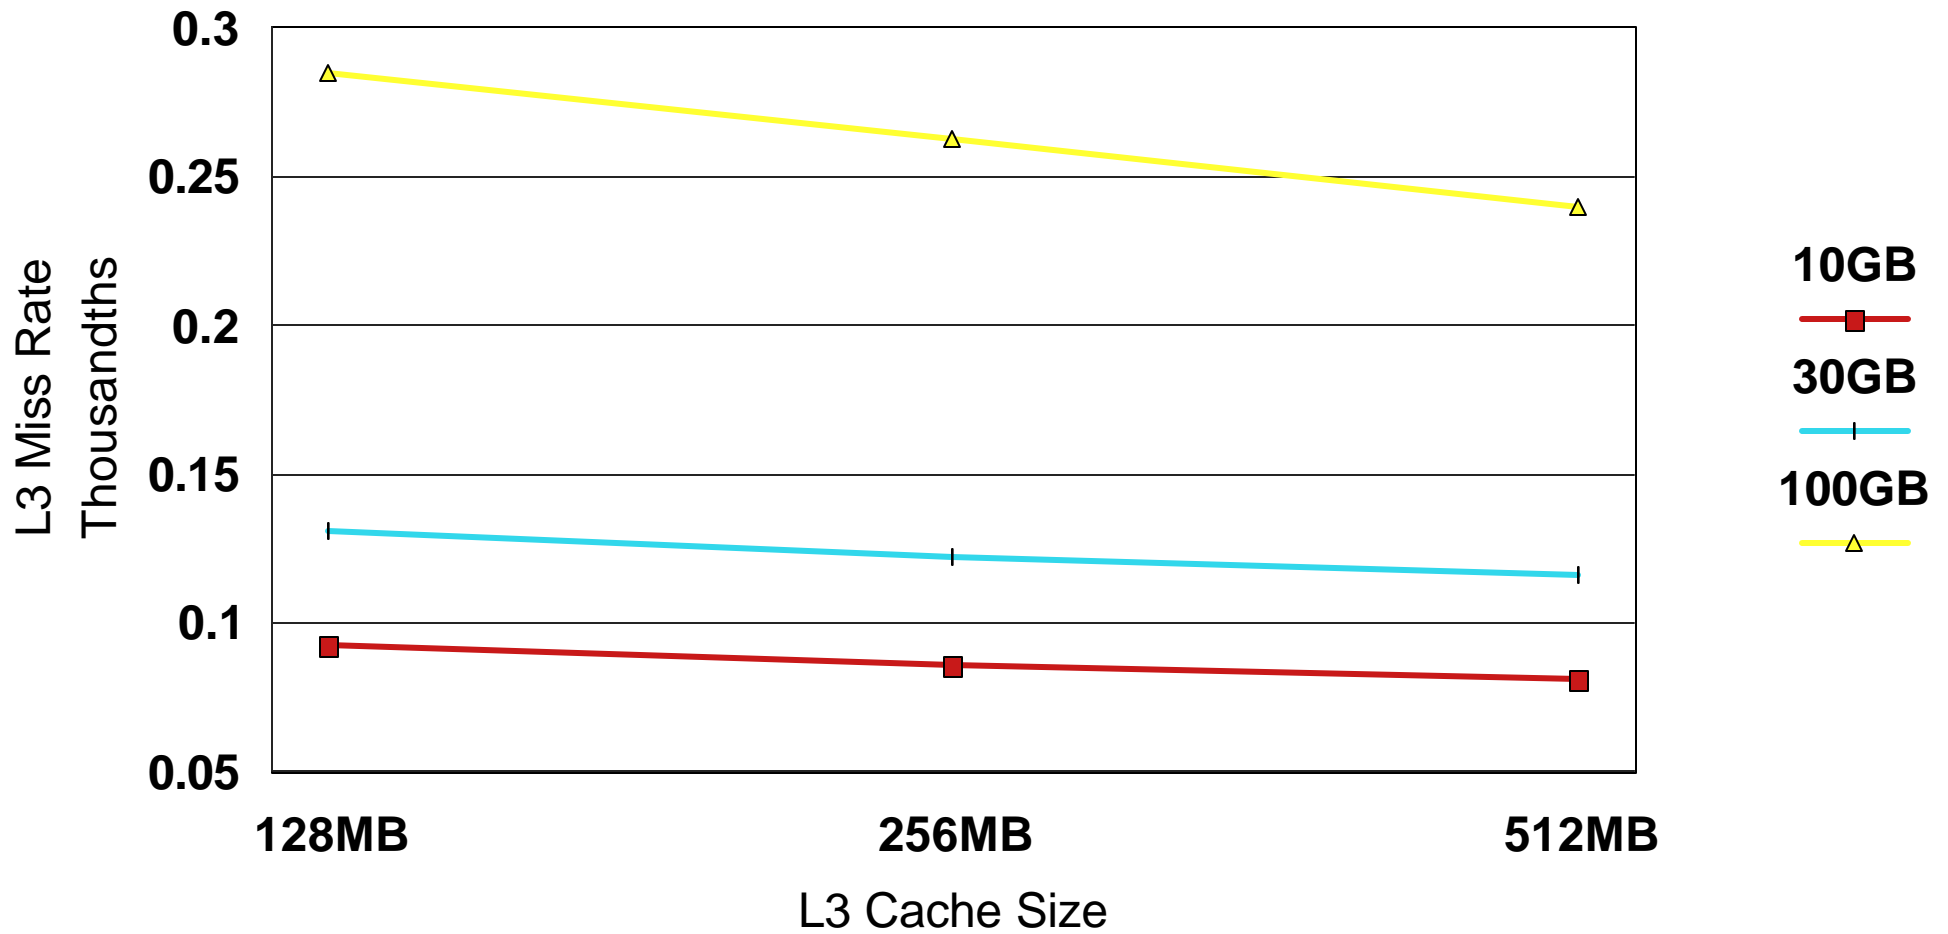
# Effect of DB Size on L3 Miss Characteristics

## L3 Miss Ratio vs. DB Size

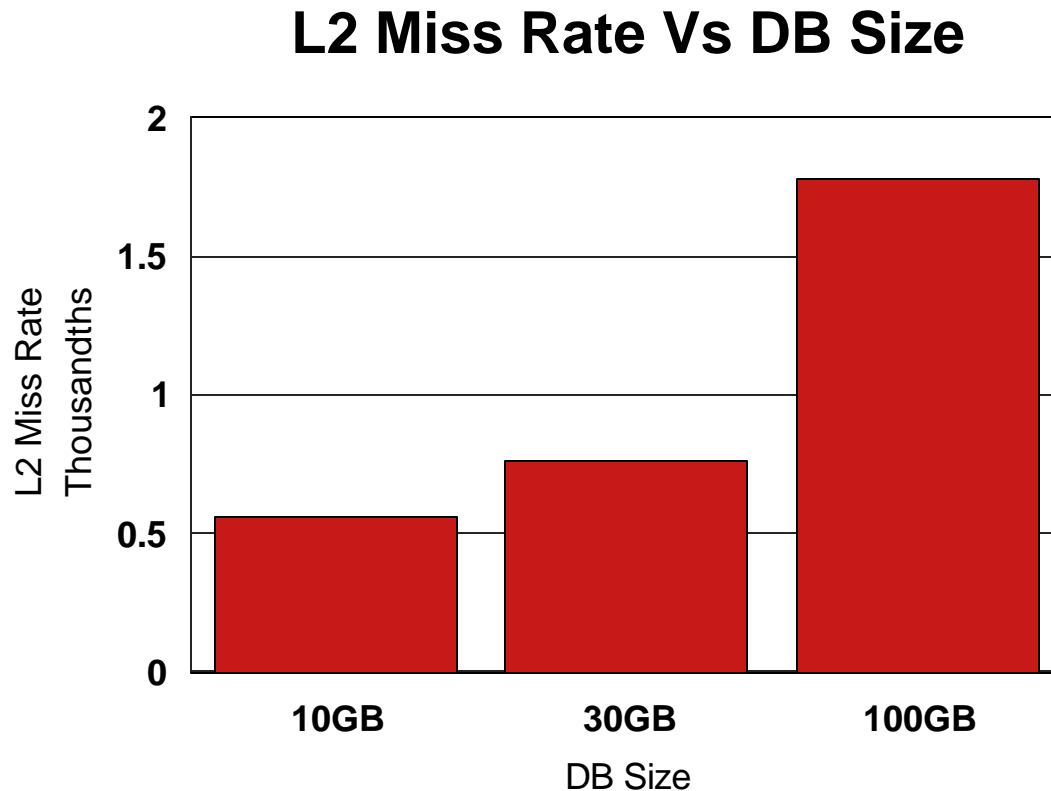


# Effect of DB Size on L3 Miss Characteristics (Contd.)

## L3 Miss Rate vs. DB Size



# Impact of DB Size on L2 Miss Characteristics




Increased working set => More misses.


Increase in miss rate is significant for 100GB => L3 will be useful.

# Impact of DB Size on L2 Miss Characteristics (Contd.)

### L2 Miss Breakup (128MB/L3)




### L2 Miss Breakup (256MB/L3)

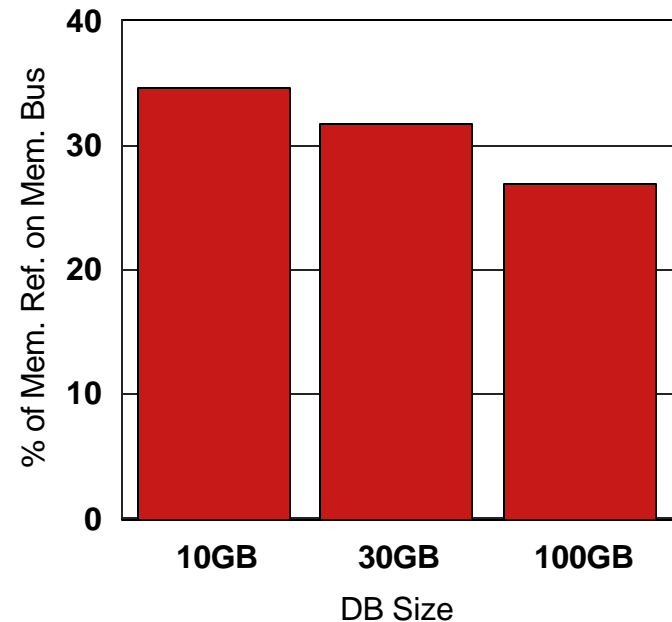


# Impact of DB Size on L2 Miss Characteristics (Contd.)

## DClaims (Invalidation Reqs.)



## Modified Interventions



As dataset size increases, misses increase => DClaims decrease.

Because of Database distribution on multiple disks, increased throughput (useful work) translates to increased cache traffic.

## Concluding Remarks

- DSS workload characteristics vary significantly with DB size.
- Misses per instruction can vary three fold between a 10GB and 100GB database.
- Invalidation requests and modified interventions on the memory bus decreases moderately with increase in database size.
- Due to the large variations in workload characteristics, it is important that realistic DB sizes are used for system evaluations.