

# A Detailed Comparison of Two Transaction Processing Workloads

Robert Stets,

Kouros Gharachorloo, Luiz André Barroso



WWC-5  
November 25, 2002

# Online Transaction Processing (OLTP)

- Major workload for high performance servers
- Very challenging workload
  - lack of instruction-level parallelism
  - large memory overhead
- Much research on the standard benchmarks
  - TPC-B: Models a large bank with multiple branches
  - TPC-C: Models a warehouse distribution center
- *How does processor and memory system performance of TPC-B and TPC-C compare?*

# TPC-B and TPC-C

- Evaluated both workloads side-by-side
  - hardware monitoring
  - full-machine simulation

- Workloads differ in some basic metrics, *e.g.*,

|       | Cycles-per-instruction | Dirty Miss Fraction |
|-------|------------------------|---------------------|
| TPC-B | 8.5                    | 17.8%               |
| TPC-C | 5.3                    | 9.6%                |

- *Surprisingly, performance of workloads is similar across a range of architectural choices!*

## Outline

- Introduction
- Background
  - benchmark descriptions
- Results
- Conclusion

# TPC-B: Debit-Credit

- Models a banking system with frequent updates to arbitrary accounts
- 4 tables
  - branch, teller, account, history
- 1 transaction: balance update transaction
  - 3 row select and updates
  - 1 row insertion

# TPC-C: Order-Entry

- Models a whole-sale supplier
- 9 tables
  - warehouse, stock, district, customer, order...
- 5 transactions: New-order (45%), payment (43%)...
- New-order transaction (average 10 items)
  - 10 selects
  - 11 select and updates
  - 12 insertions

## Outline

- Introduction
- Background
- Results
  - methodology
  - hardware and simulation results
- Conclusion

# Software and Hardware


- Software
  - Oracle v8.0.4 (and v7.1.3)
  - IProbe (hardware counter collection)
- Hardware monitoring
  - AlphaServer 8400
    - 8- 612MHz 21164
    - 8K/1-way, 96K/3-way on-chip, 4M/1-way board caches
  - AlphaServer DS20
    - 2- 500MHz 21264
    - 64K/2-way on-chip, 4M/1-way board caches

# Simulation


- SimOS simulator
  - Full-system simulation
  - configurable
- Base system
  - 1GHz processor
  - 64K/2-way L1
  - 4MB/1-way Board-level cache
  - L2 Hit latency: 30ns
  - Memory latency: 100ns

## Results

# Hardware Monitoring



**Debit-Credit (TPC-B)**  
CPI: 8.5



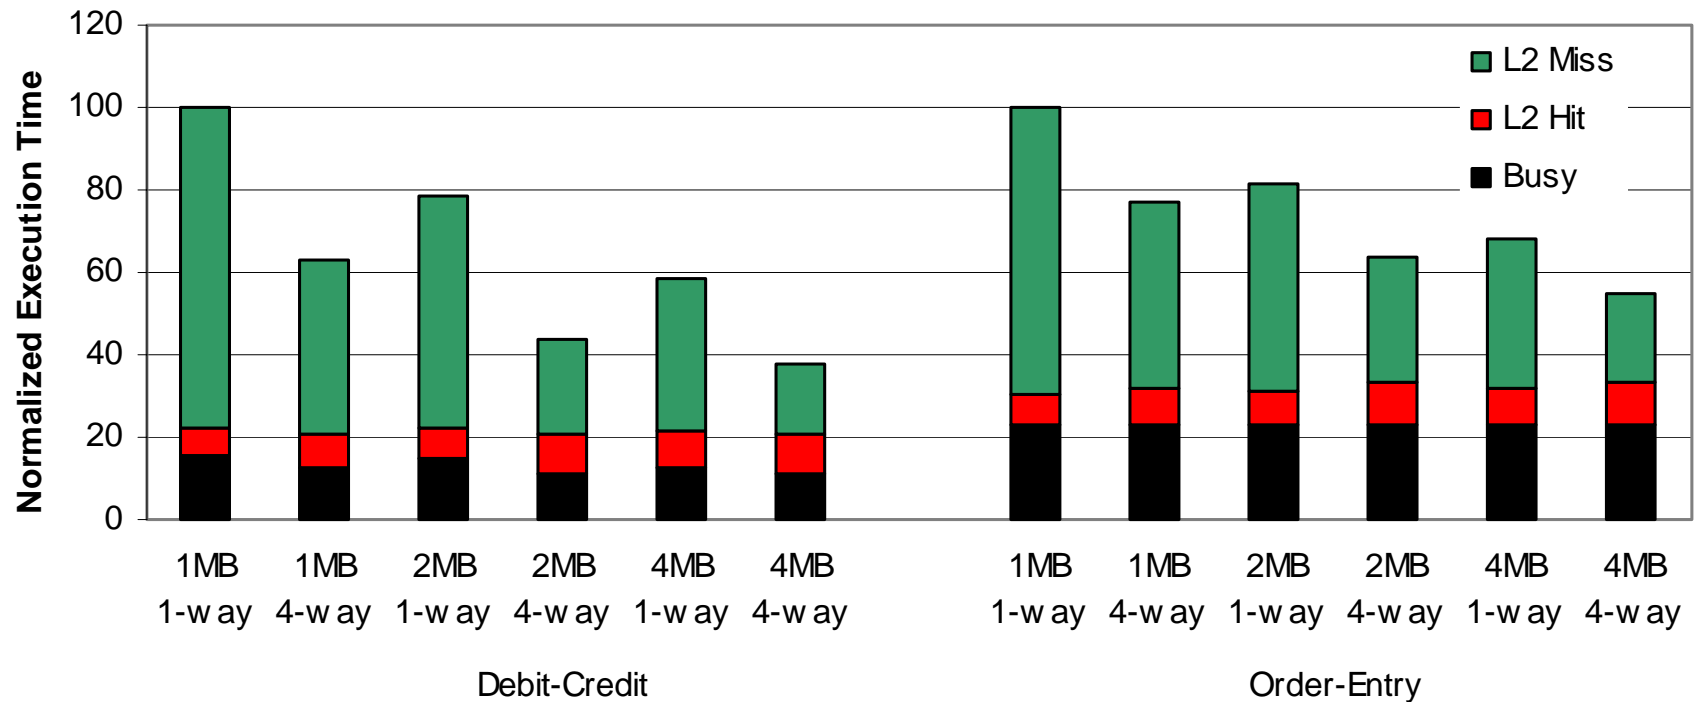
**Order-Entry (TPC-C)**  
CPI: 5.3

(AlphaServer 8400)

| Parameter           | Debit-Credit | Order-Entry |
|---------------------|--------------|-------------|
| Icache (global)     | 13.6%        | 13.4%       |
| Dcache (global)     | 22.2%        | 31.0%       |
| Scache (global)     | 7.5%         | 7.5%        |
| Bcache (global)     | 1.6%         | 1.0%        |
| Dirty miss fraction | 17.8%        | 9.6%        |

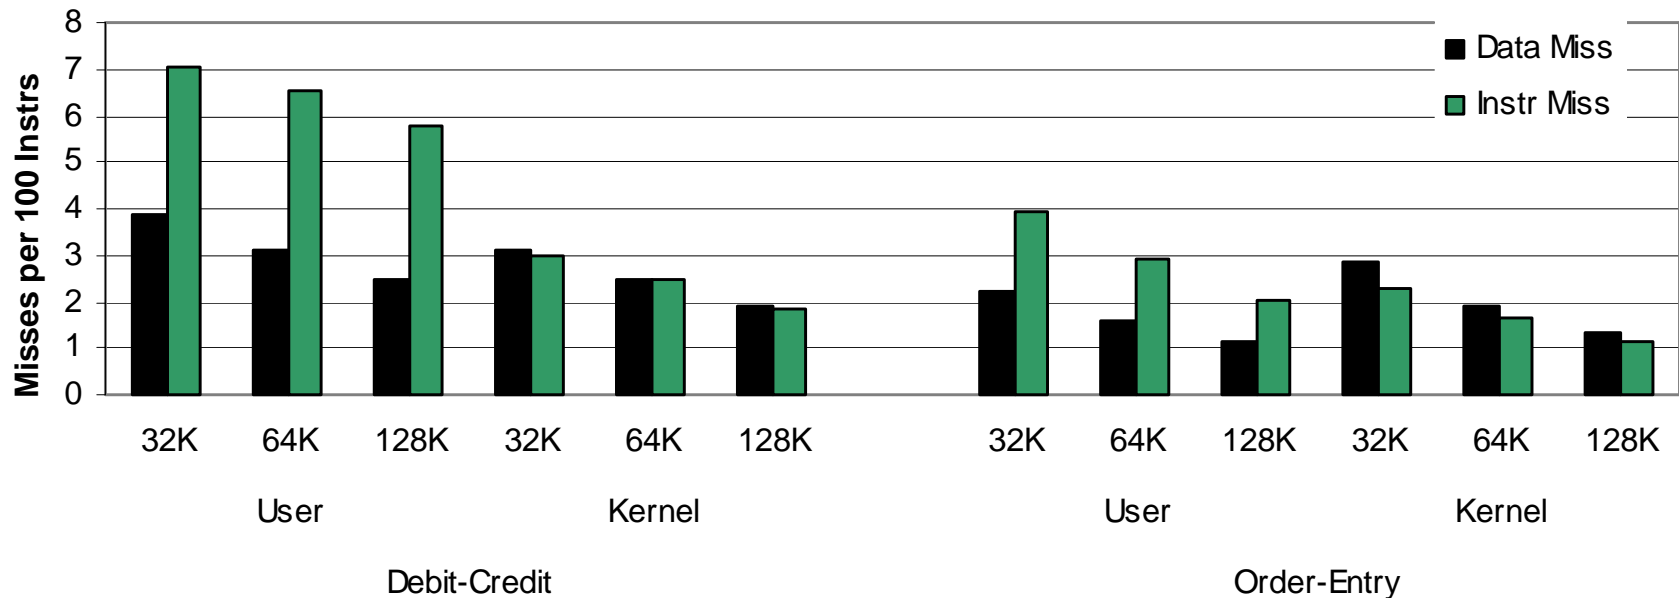


# L2 Cache Configurations



- Similar benefits from increased size
- Debit-Credit benefits more from associativity

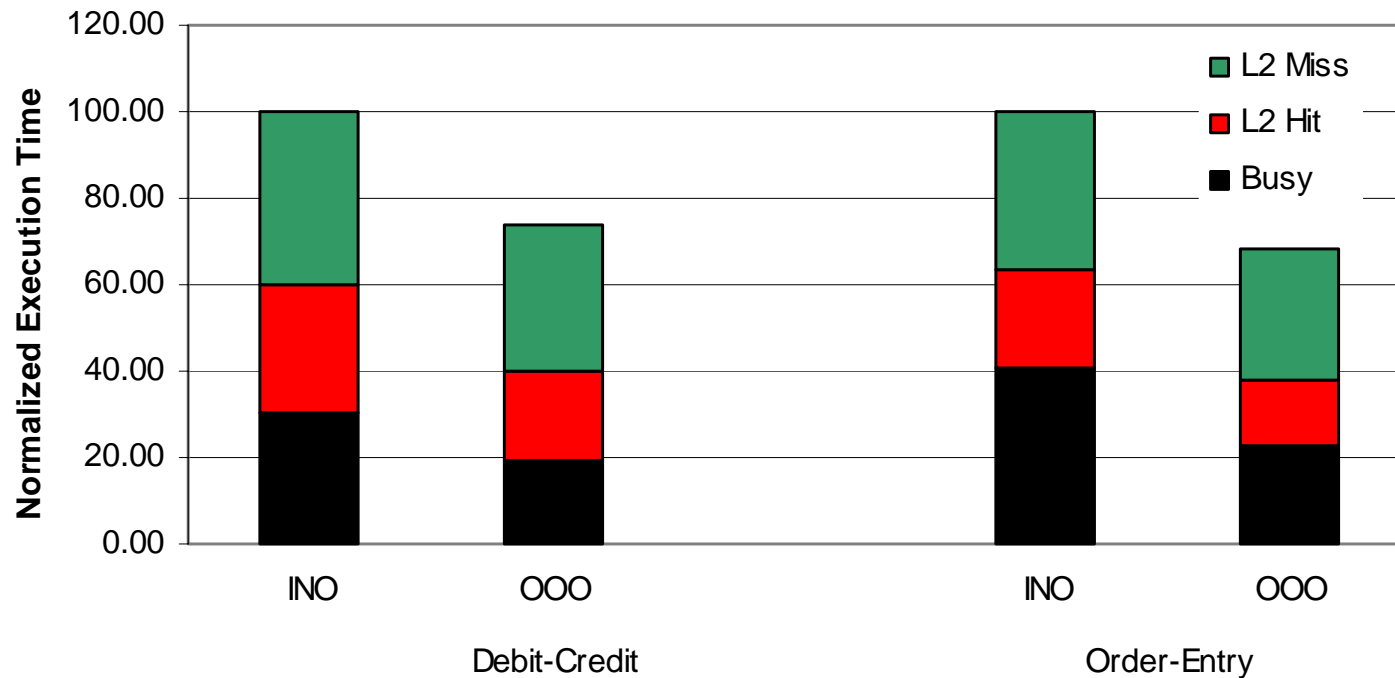
# L1 Cache Size



(2-way associative)

- Debit-Credit show significant misses even at 128KB size

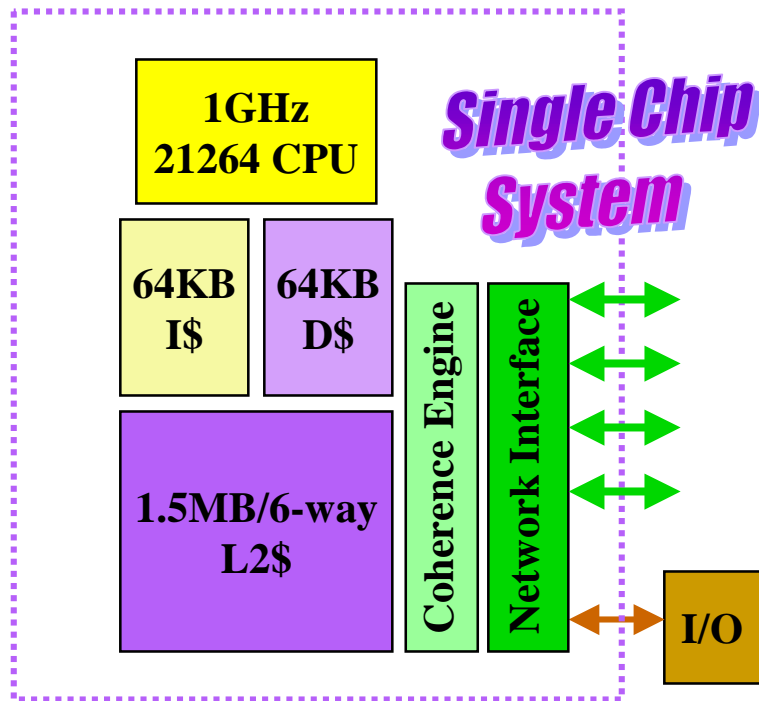
# Out-of-order Execution



- Similar performance gains from on hardware
  - AlphaServer DS20 (OOO) vs AlphaServer 8400 (INO)

# Chip-level Integration


Alpha 21364



- “Glue-less” multiprocessing
- Lower latency
- Higher bandwidth

|                | Latencies (ns) |       |        |
|----------------|----------------|-------|--------|
|                | L2 Hit         | Local | Remote |
| Base           | 30             | 100   | 175    |
| All Integrated | 15             | 75    | 150    |


# Chip-level Integration



- Debit-Credit benefits more from integration

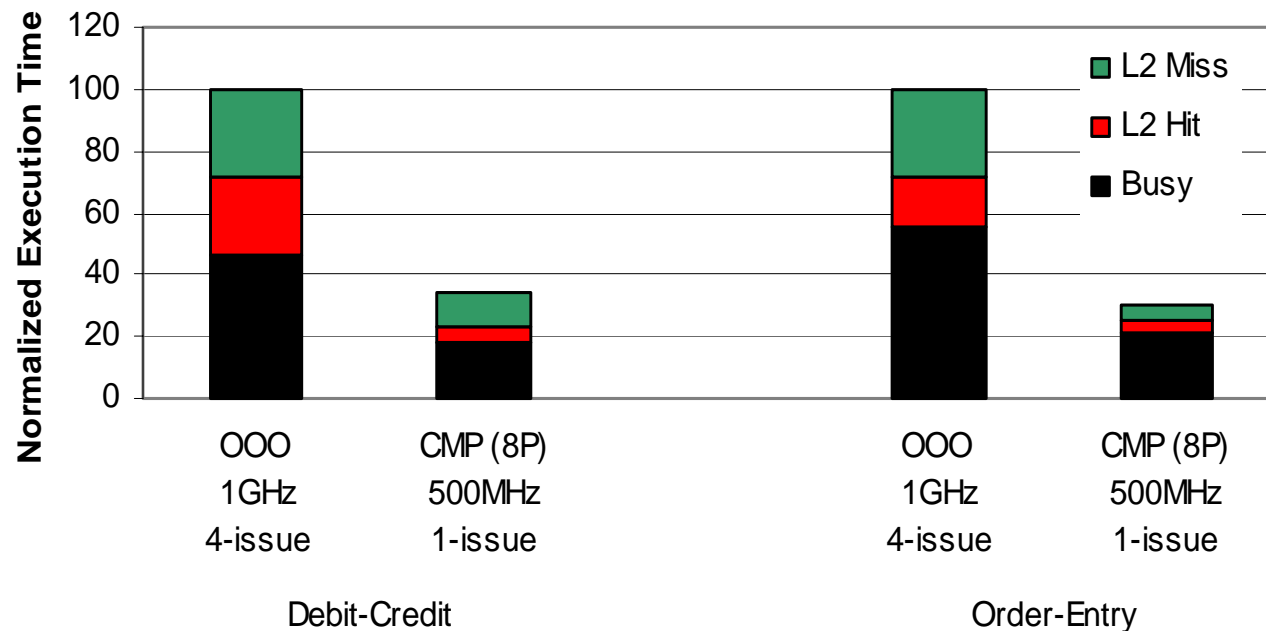
# Chip-multiprocessing

- Based on Piranha [ISCA, 2000]
  - 8 single-issue, in-order cores
  - 32K/1-way L1s
  - 1MB/8-way shared L2



- Assumes an ASIC design process
  - target 500MHz

# Chip-multiprocessing



- Debit-Credit
  - larger relative memory stall time
  - smaller improvement from CMP
  - penalized more by smaller L2 cache

## Conclusions

- Both workloads show similar performance across a variety of architectural choices
  - cache organizations
  - out-of-order execution
  - chip-level integration
  - chip-multiprocessing
- TPC-B may be more interesting for architectural studies
  - simpler to set up
  - more memory intensive