A Detailed Comparison of TPC-C versus TPC-B

Robert Stets, Luiz André Barroso,
Kourosh Gharachorloo, Ben Verghese

Western Research Laboratory
Compaq Computer Corporation

Online Transaction Processing (OLTP)

- Major workload for high performance servers
- OLTP characterization
  - Large number of clients
  - Short transactions that update small amounts of data

- Transaction Processing Council OLTP benchmarks
  - TPC-B: Models a large bank with multiple branches
  - TPC-C: Models a company that distributes products

- *How does processor and memory system performance of TPC-B and TPC-C compare?*

TPC-B versus TPC-C

- Evaluated scaled-down models of these two benchmarks
  - Performance monitoring of cutting-edge hardware
  - Full-machine simulation of next generation hardware

- Similarities
  - Both show 1.3x speedup from out-of-order execution
  - Working sets of both can fit in reasonably sized L2
- Difference
  - TPC-B suffers from high sharing overheads

Outline

- Introduction
- Performance monitoring
- Simulation
- Conclusion

Our evaluation is ongoing, and we hope to include additional results in the final presentation. If slides change or are added, we will distribute updated handouts at the workshop.

Performance Monitoring: Platform

- Monitoring via IProbe and Digital Continuous Profile Infrastructure (DCPI)

- Eight processor AlphaServer (AS) 8400
  - Alpha 21164 processor at 612MHz
    - Three cache levels: 8K/1-way, 96K/3-way, 4MB/1-way
    - Memory latency: 357 nanoseconds

Performance Monitoring: AS 8100

TPC-B

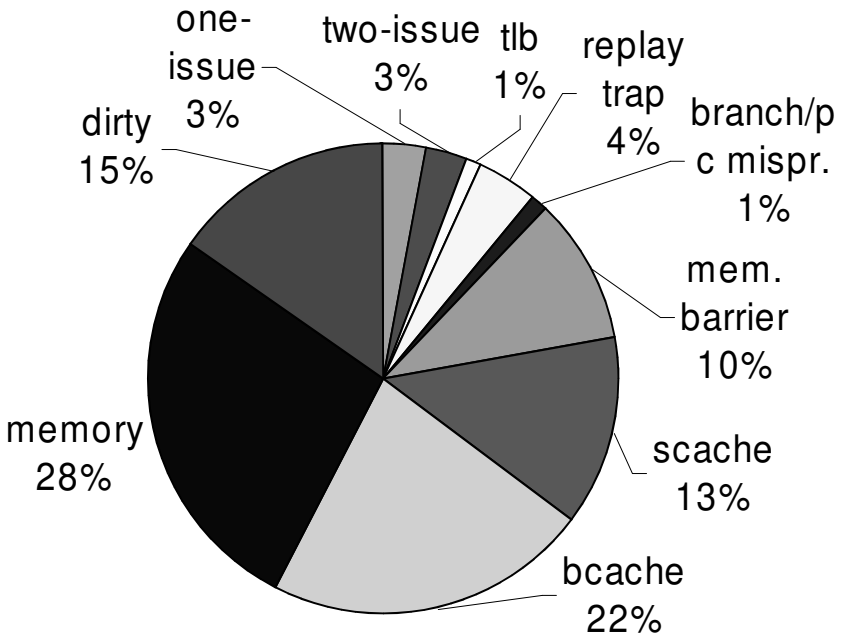

TPC-C

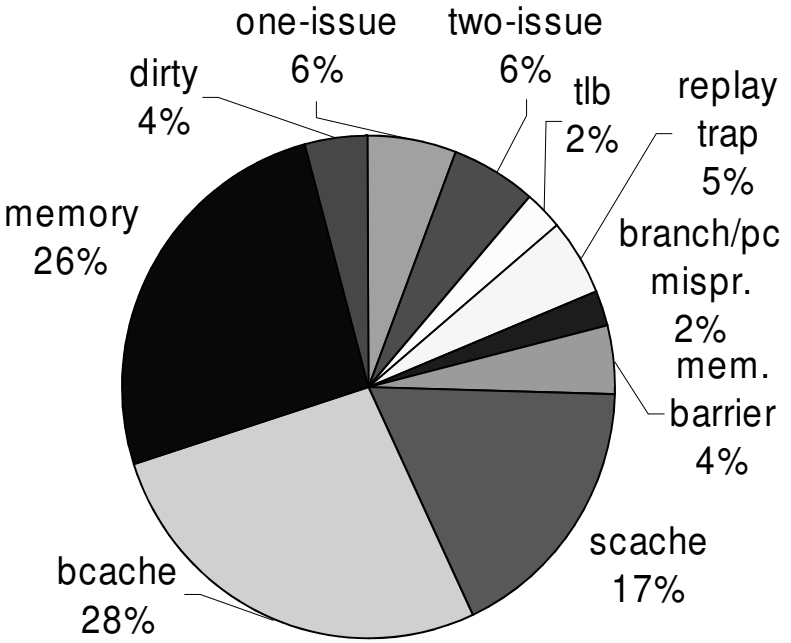

	TPC-B	TPC-C
icache miss rate	**12.4%**	**13.3%**
dcache miss rate	**31.7%**	**31.5%**
scache miss rate	**7.4%**	**7.1%**
bcache miss rate	**1.6%**	**1.0%**
CPI	**11.8**	**5.9**
% dirty misses	**26.8%**	**9.4%**

Simulation: Platform

- Simulation using SimOS-Alpha

- Eight processor "next-generation" AlphaServer
  - Alpha processor (1GHz)
    - single-issue, in-order and four-issue, out-of-order pipelines
    - Integrated on-chip L1, L2, memory, coherence, and network controllers
      - L1: various sizes, 2-way associative
      - L2: various sizes and associativities
  - NUMA architecture
    - Local memory latency: 75 nanoseconds
    - Remote memory latency: 225 nanoseconds

Simulation: In-order versus Out-of-order

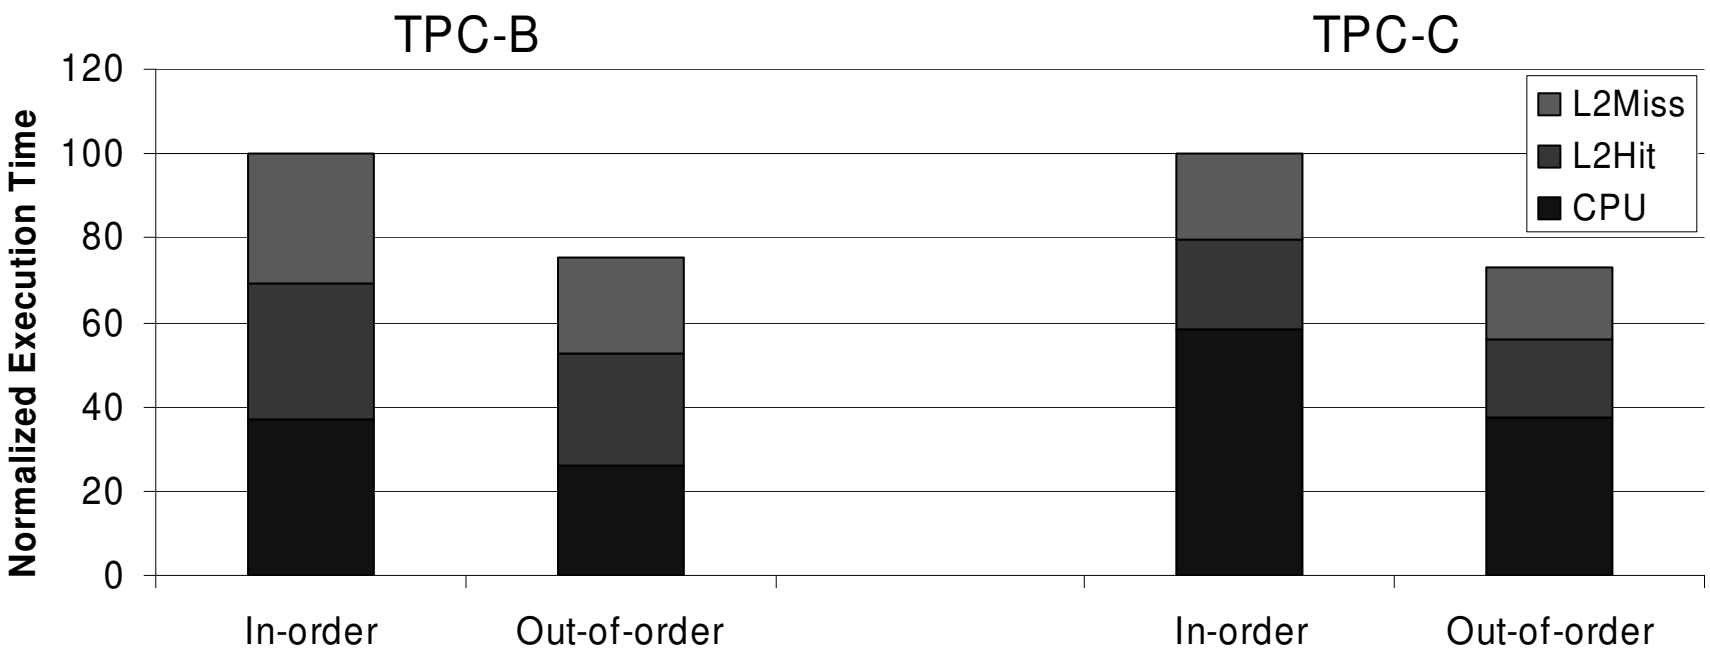

8 processors (in-order)
64K/2-way L1
2M/8-way L2

- Out-of-order execution provides a 1.3x speedup for both benchmarks

Simulation: Impact of L2 Configuration

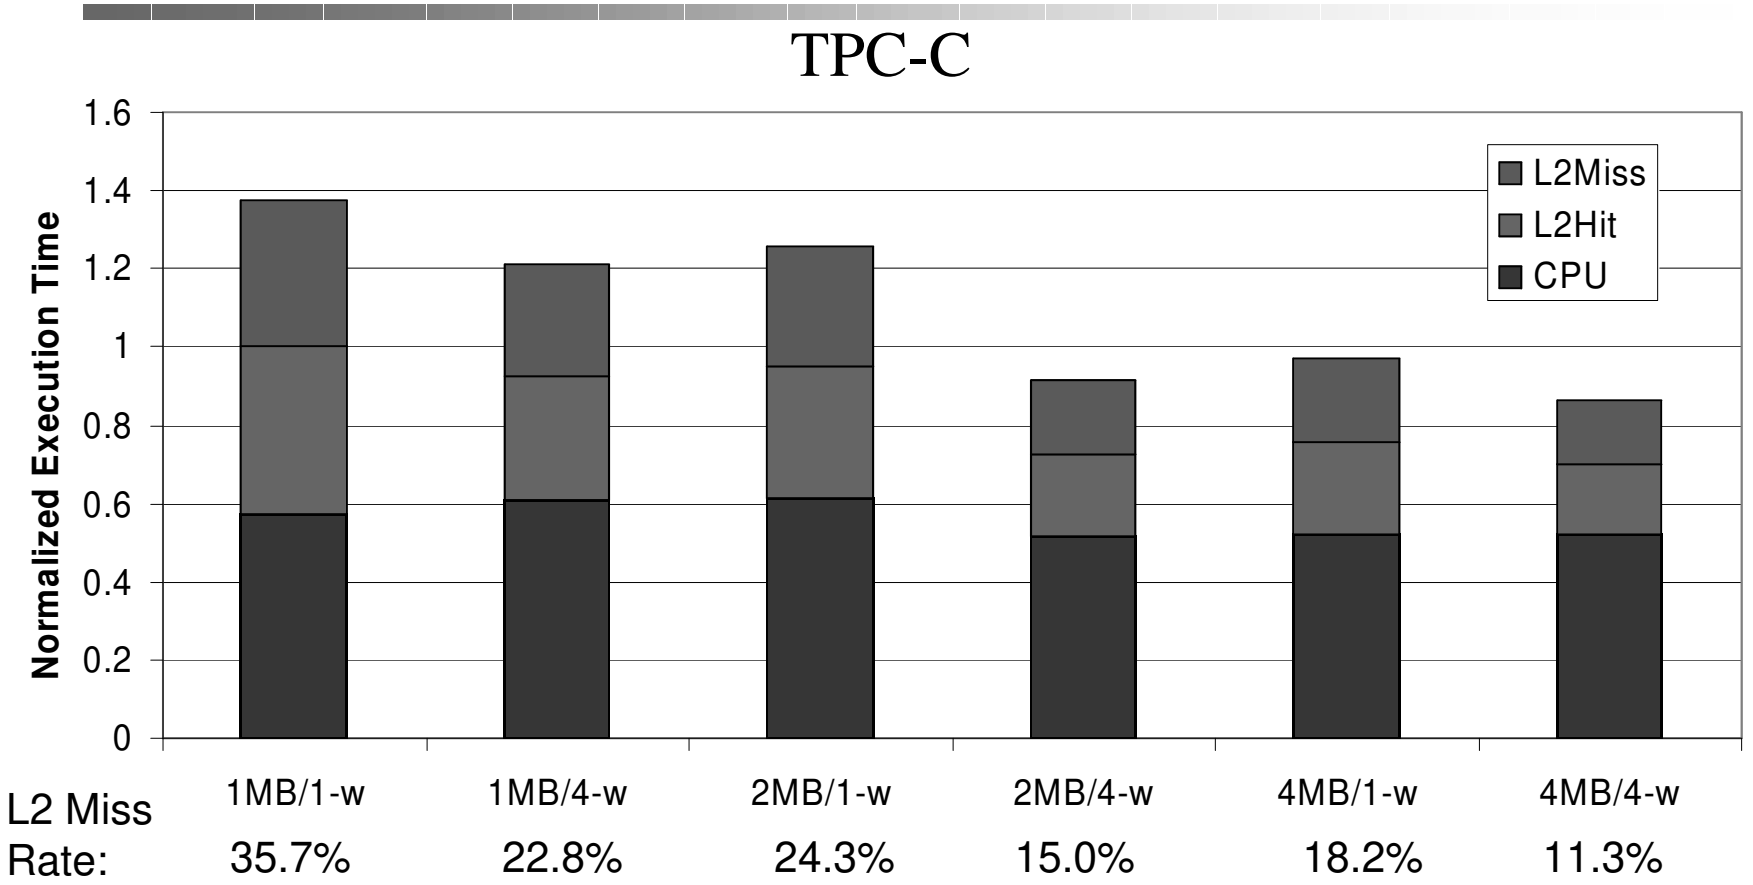

- TPC-B shows similar trends [Barroso, 1998]

Conclusions: TPC-B versus TPC-C

- Both benchmarks spend the majority of time in the memory system
  - Out-of-order execution is not effective in hiding memory stalls
  - Working sets fit in reasonably sized L2

- TPC-B has higher sharing and memory system overhead
  - More dirty misses
  - Higher memory barrier stall time