

Computer Systems Performance Analysis and Benchmarking (37-235)

Analytic Modeling
Simulation
Measurements / Benchmarking

Lecture by:

Prof. Thomas Stricker
Michela Taufer

Assignments/Projects:

Christian Kurmann

Textbook:

Raj Jain, "The Art of Computer Systems Performance Analysis", 1991 Wiley & Sons, New York

Topic of Today:

- **Memory Systems Benchmarks (memperf)**
- **Modelling of an application (OPAL)**

Memory System Performance Characterization for PCs, Workstations and Parallel Systems

Thomas M. Stricker



Laboratory for Computer Systems
Swiss Institute of Technology
CH-8092 Zürich


Intel Microprocessor Research Lab

February 10, 1997

Outline

- **Notion of a global address space**
- **Extended copy transfer (ECT) characterization**
 - Bandwidth, spatial locality, temporal locality
- **Performance of local/remote memory accesses and transfers**
 - DEC8400, T3D, T3E
 - Intel P6, uni- and twin- Processor
- **High speed communication and the memory system**
- **Evaluation with an application kernel**

Rechenleistung wächst explosionsartig!



Extended Copy Transfer Characterizations

Contribution: New method to characterize the performance of memory systems:

- **Categories**
 - **Access pattern, stride** (temporal locality)
 - **Working set** (spatial locality)
- **Value**
 - **Transfer bandwidth** (large amount of data)
- **Same chart resulting from one microbenchmark**
 - **Local** and **Remote** transfers

Related work

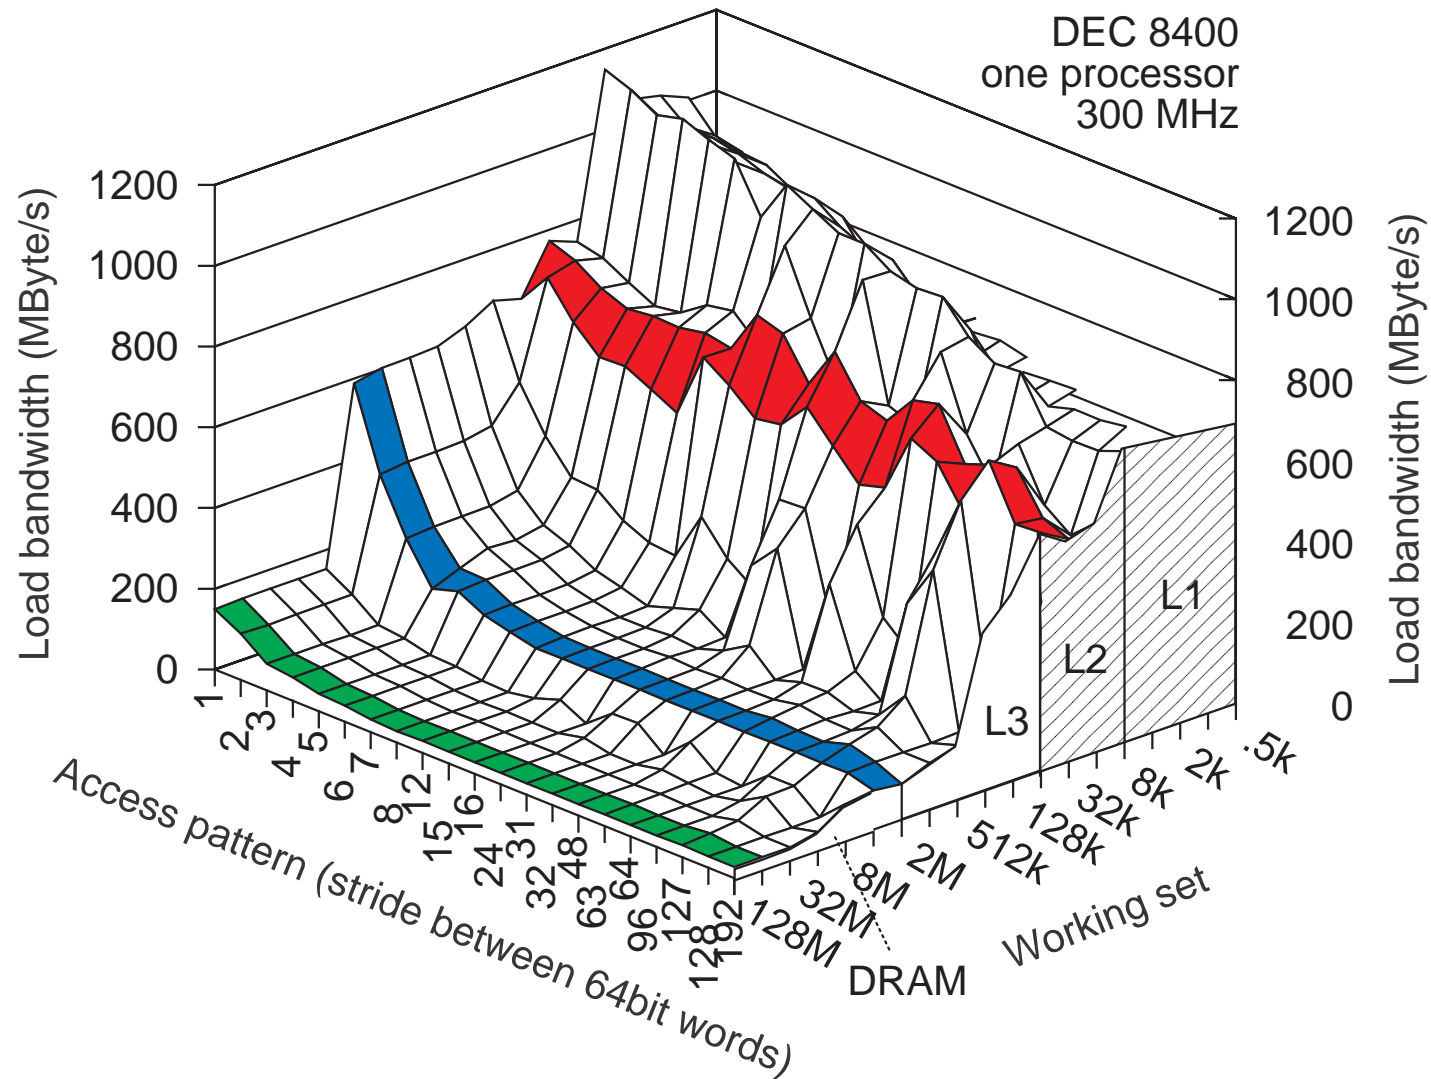
- **Trace based evaluation of a memory system**
- **Classic NUMA model**
- **A basic copy transfer model [ISCA95]**

Caveat


Some parameter combinations are hard to measure, even with carefully tuned C code:

- **Reduced performance for *large strides* and *small working-sets* is a measurement artifact and not architecture related.**
- **Compilers occasionally generate suboptimal instruction schedules for loads/stores.**


Local Load Access (DEC 8400)



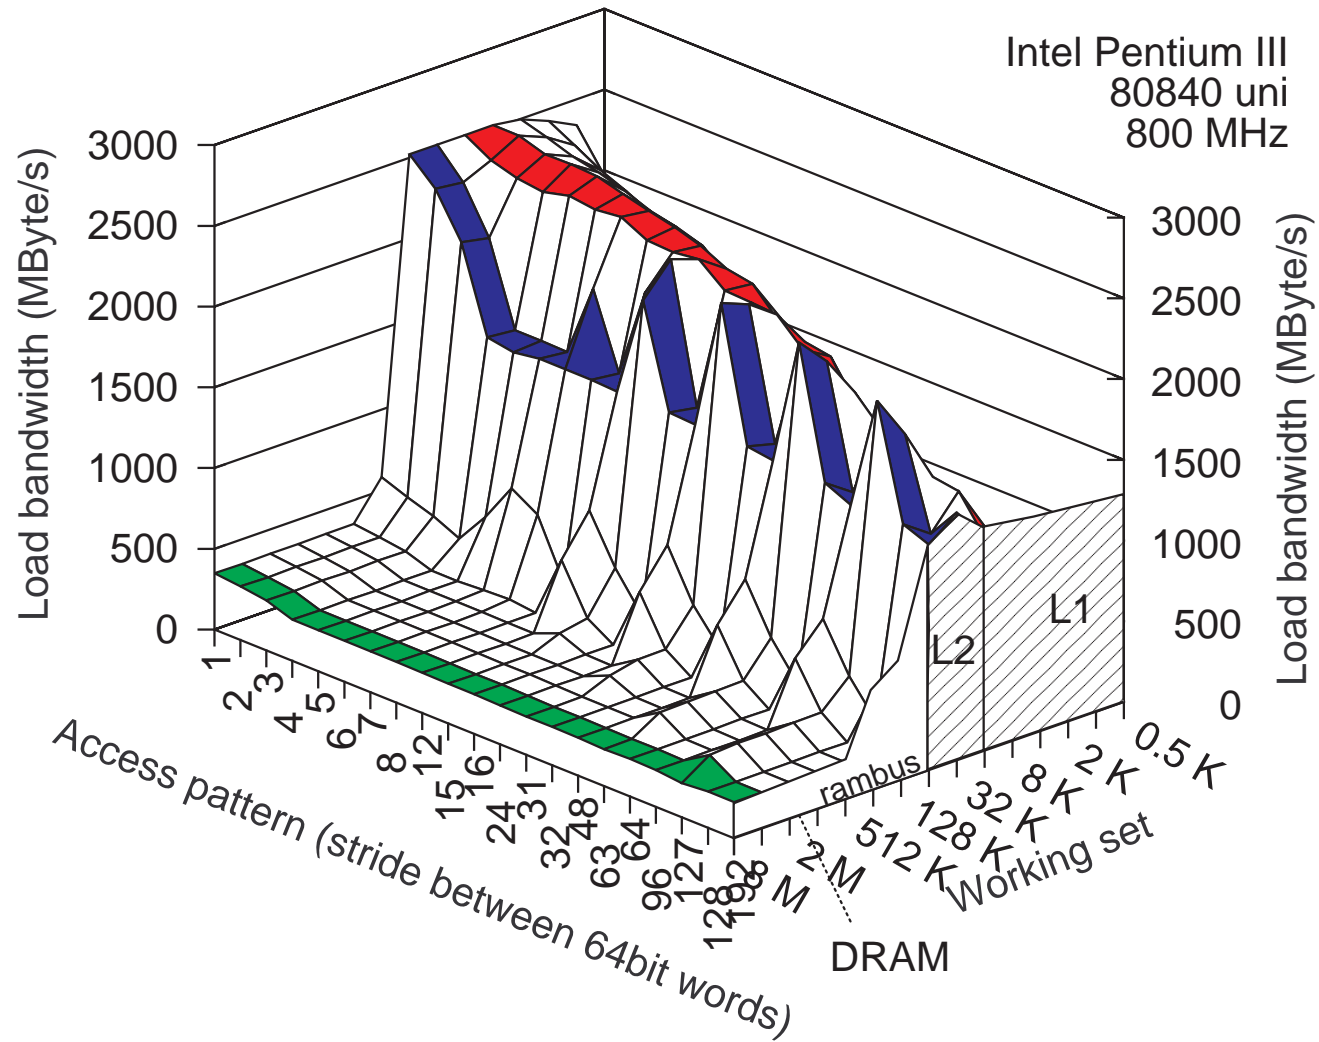
Local Load Access (Cray T3D)




Local Load Access (Cray T3E)




Load Datenstrom Intel 80840 (800MHz)




Comparison - Local (Working Set: DRAM)



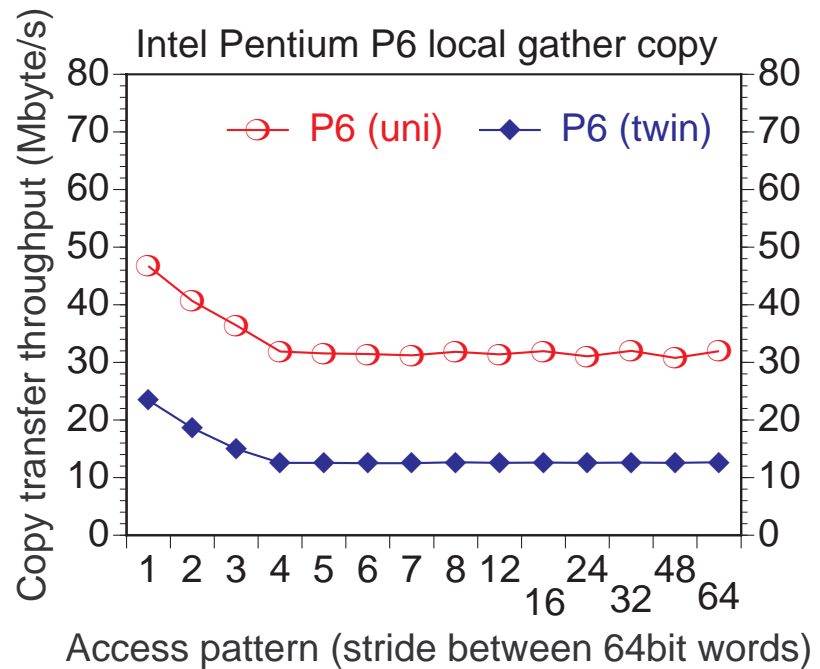

Local Load Access (P6 single)



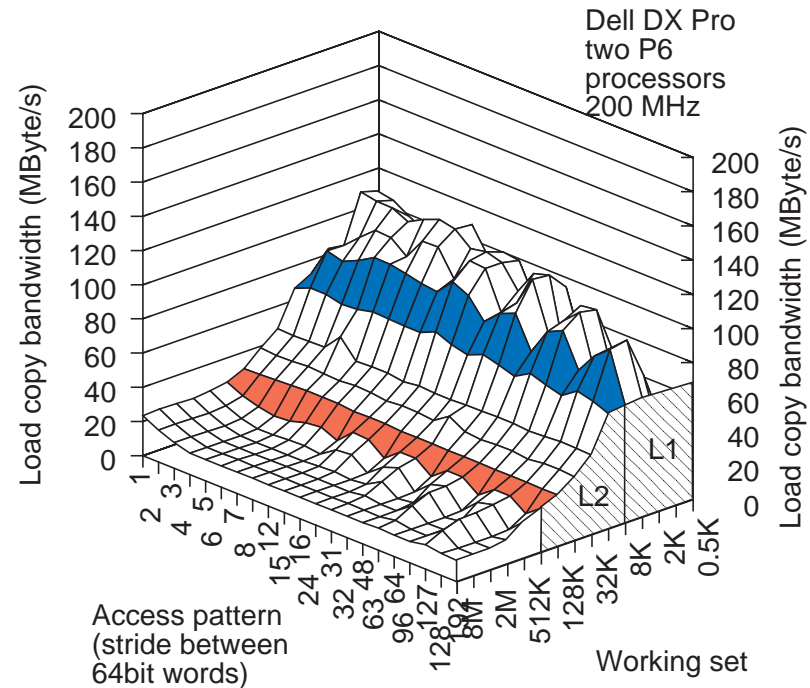

Local Load Access (P6 twin)



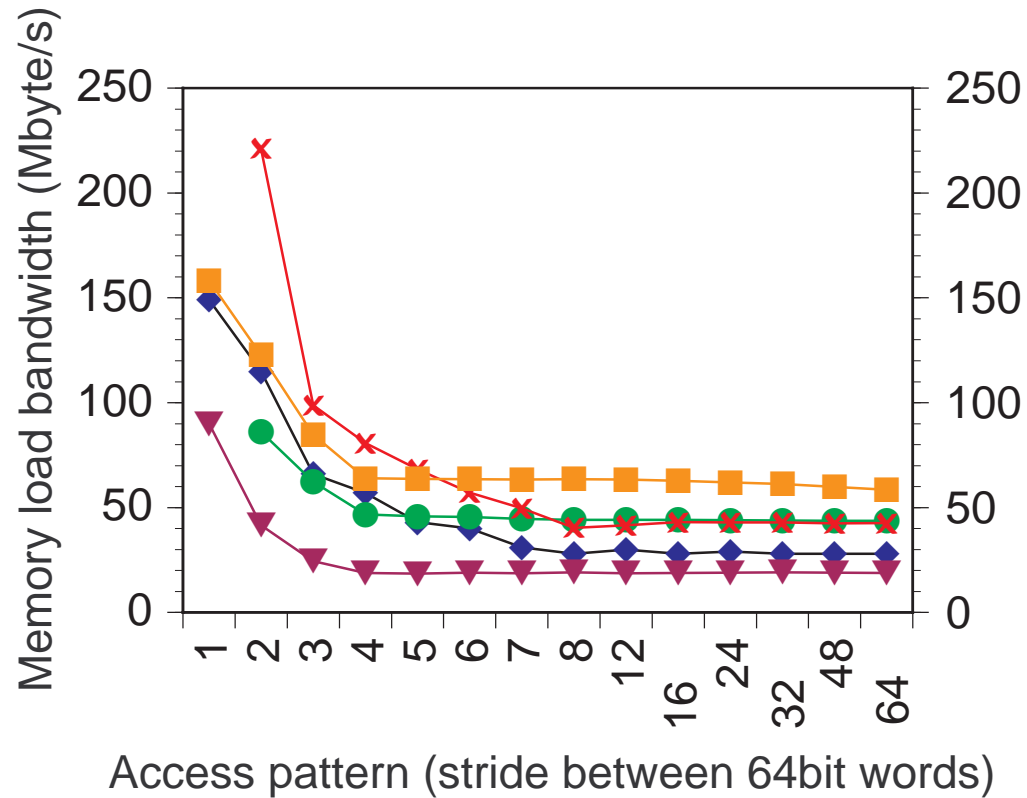
Comparison - Local Copy (Working Set: DRAM)



Local Copy (P6 single/twin)




Comparison - Local Access (Working Set: DRAM)




- ◆— DEC8400
- Cray T3D
- ×— Cray T3E
- P6 (uni)
- ▼— P6 (twin)

Entwicklung 1996 to 2000




80440FX (200MHz)



80840 (800MHz)


Vergleich - Local Copy



- ◆ DEC8400
- ✕ Intel PPro
- ▼ CrayT3E
- SGI Origin
- Intel PII 840er

Local/Remote: Parallel Computers


Parallel & Network Computers



Nectar (Classic, Gigabit, Credit), SHRIMP
Cray T3D/T3E, Intel Paragon, NEC Cenju

P: Processor **M: Memory**

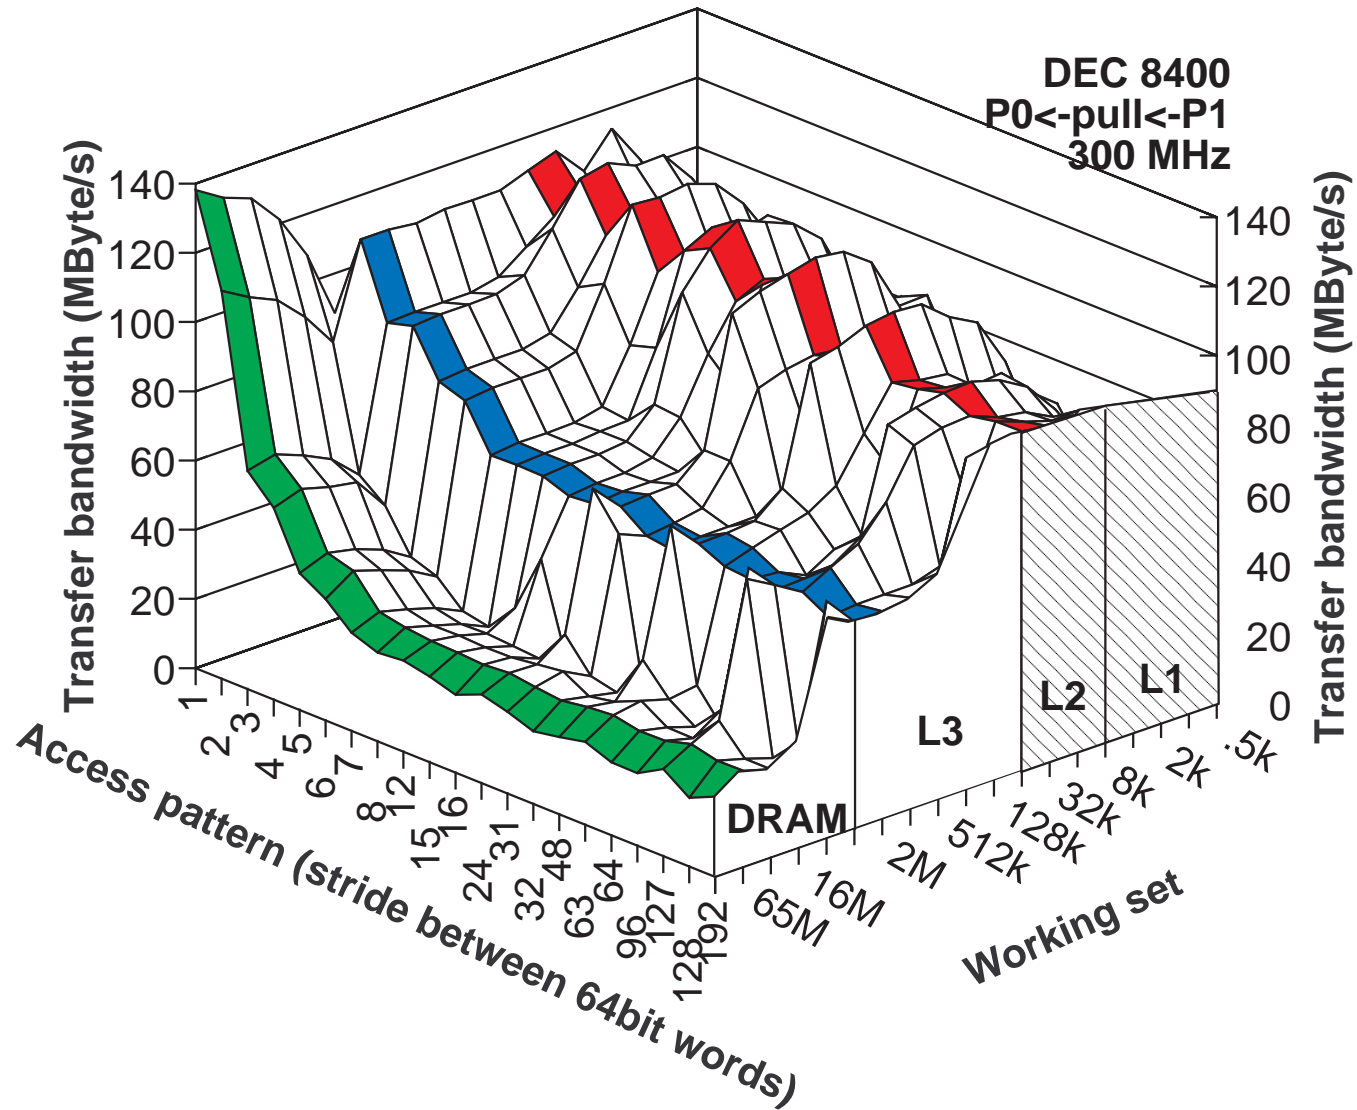
Symmetric Multiprocessors





SGI Power Challenge, DEC 8400
Cray J90, Pentium SMPs

C: on-chip/on-board Caches



Remote Copy Transfer (DEC 8400)



Remote Copy Transfer on T3D (push/pull)

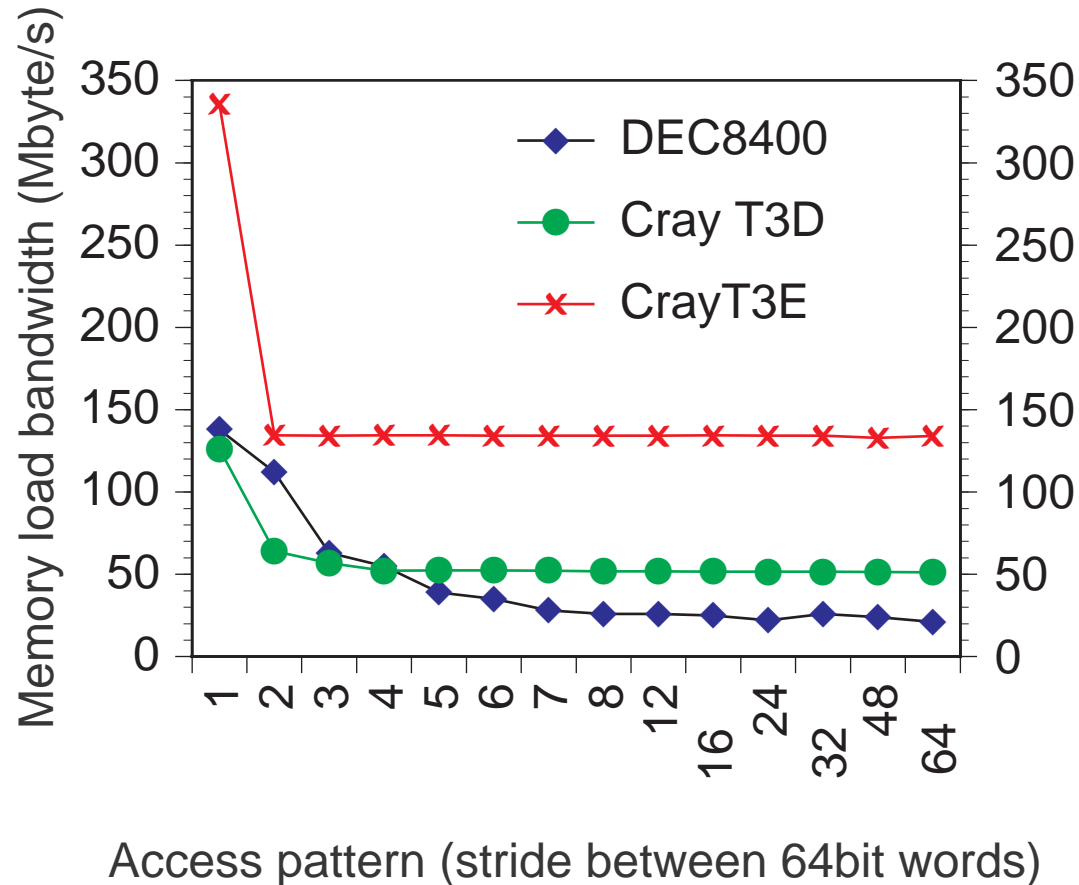


Remote Copy on Cray T3E (push, pull)





***Performance numbers on T3E are likely to improve
New machine software not as mature as T3D***

Comparison - Remote (Working Set: DRAM)



Remote Transfer im Cluster (Working Set: DRAM)



Conclusions

- The introduction of a **global address space** asks for a new memory system **performance model**.
- Streaming support (MPPs) instead of L3 caches (SMPs) can result in **better memory systems** and **faster computation**.
- **Fast communication** puts a high demand on the memory system. Current PC platforms need some improvements.