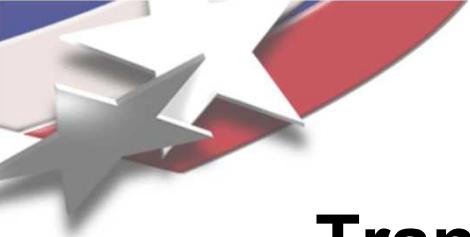


TLBs, DRAMs, and Other Scary Things

and Impact of Multi-core

April 8, 2008
Software Engineering Seminar Series

Kevin Pedretti
Scalable System Software, Dept. 1423
ktpedre@sandia.gov



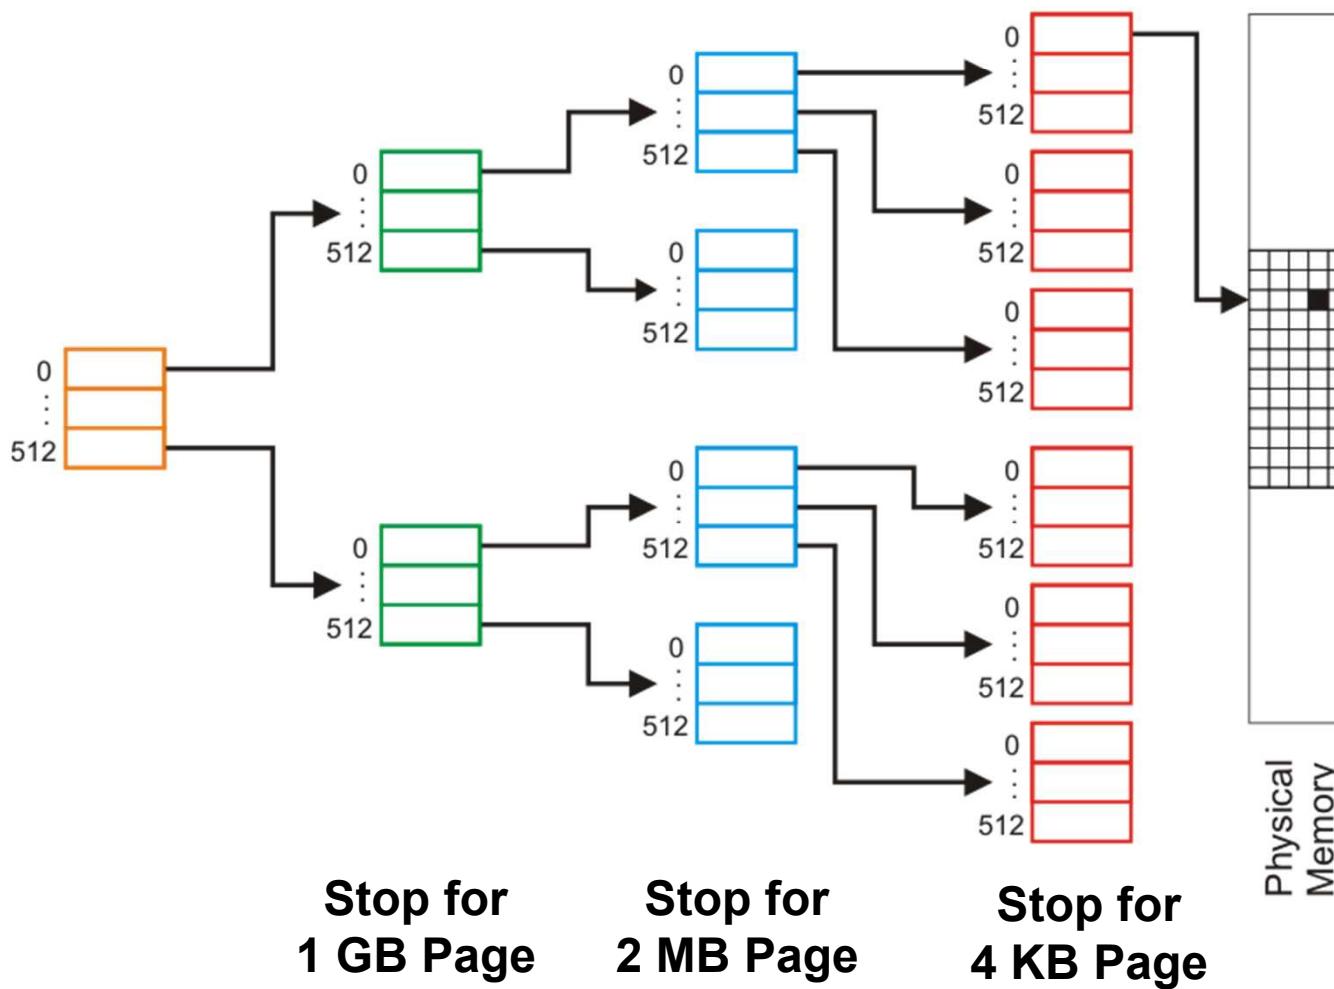
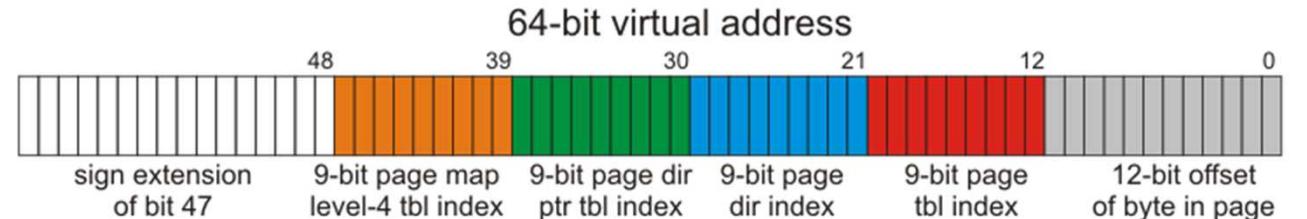
Translation Lookaside Buffers (TLB)

- TLB purpose is to minimize performance impact of virtual memory
 - Given a virtual address, returns physical address
 - Size limited by speed requirements
- Yet another level of locality
 - Best thought of as a cache, like L1 & L2
 - Hit rates range from 0-100% depending on application



X86-64 Long PAE Page Table Lookup

(52-bit physical space, still 64-bit virtual space per process)



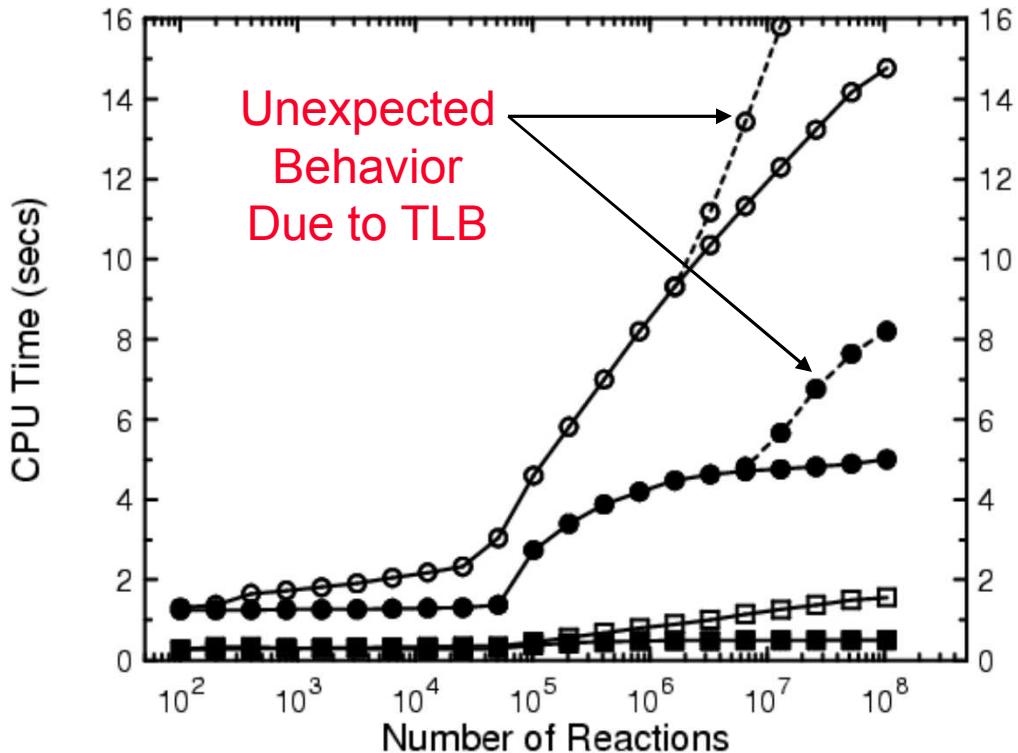


Red Storm Upgrade

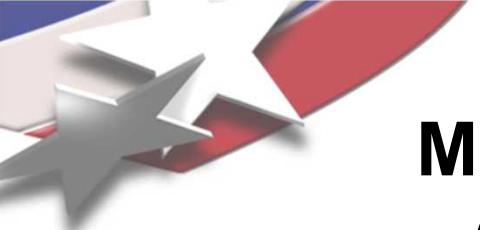
- In past, small pages have usually been better
 - `yod -small_pages -sz 8192 app`
 - Not always (next slide)
- Red Storm being partially upgraded to quad-core
- For quad-core, large pages almost always better
 - Need '`-best_page_size`' option?

	Small Page TLB (4 KB)	Large Page TLB (2 MB)	Super Jumbo TLB (1 GB)
Dual-core	32+512	8+0	N/A
Quad-core	48+512	48+128	48+0

TLB Gets in Way of Algorithm Research

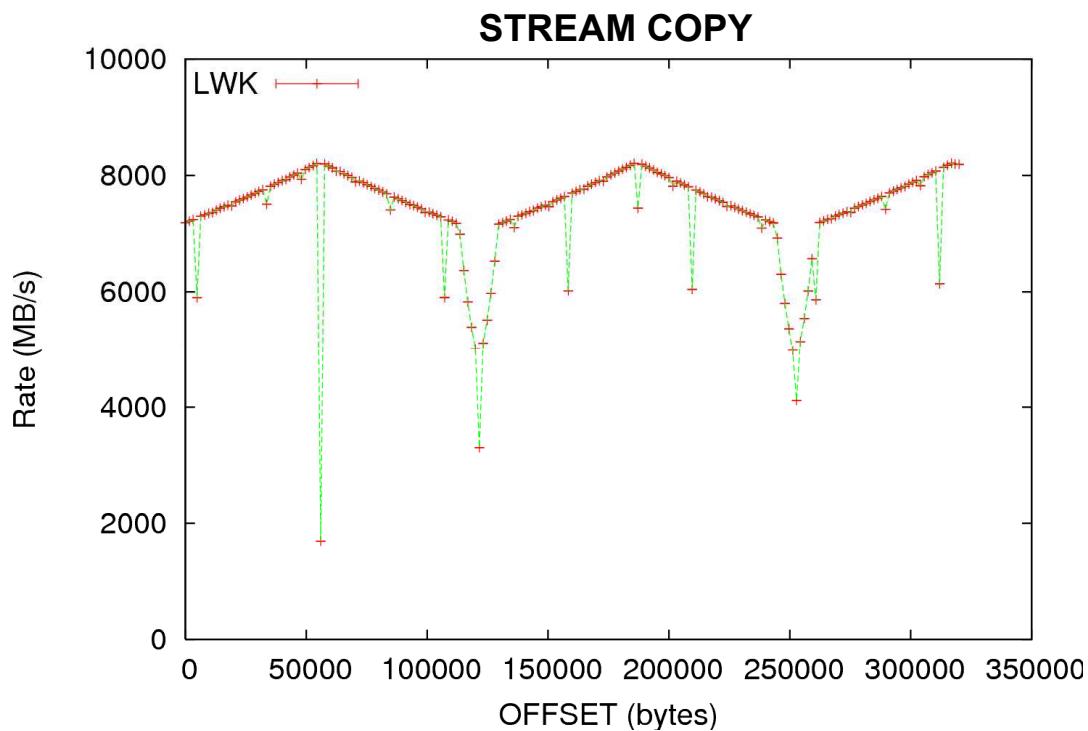


TLB misses increased with large pages,
but time to service miss decreased dramatically (10x).
Page table fits in L1! (vs. 2MB per GB with small pages)



Mysterious STREAM Sawtooth (N=2000000, ~16MB Arrays)

First observed by Courtenay Vaughan running HPCC



Double arrays
For $i = 0$ to 2000000 (out of cache)

COPY

$$a[i] = b[i]$$

SCALE

$$a[i] = \text{scalar} * b[i]$$

ADD

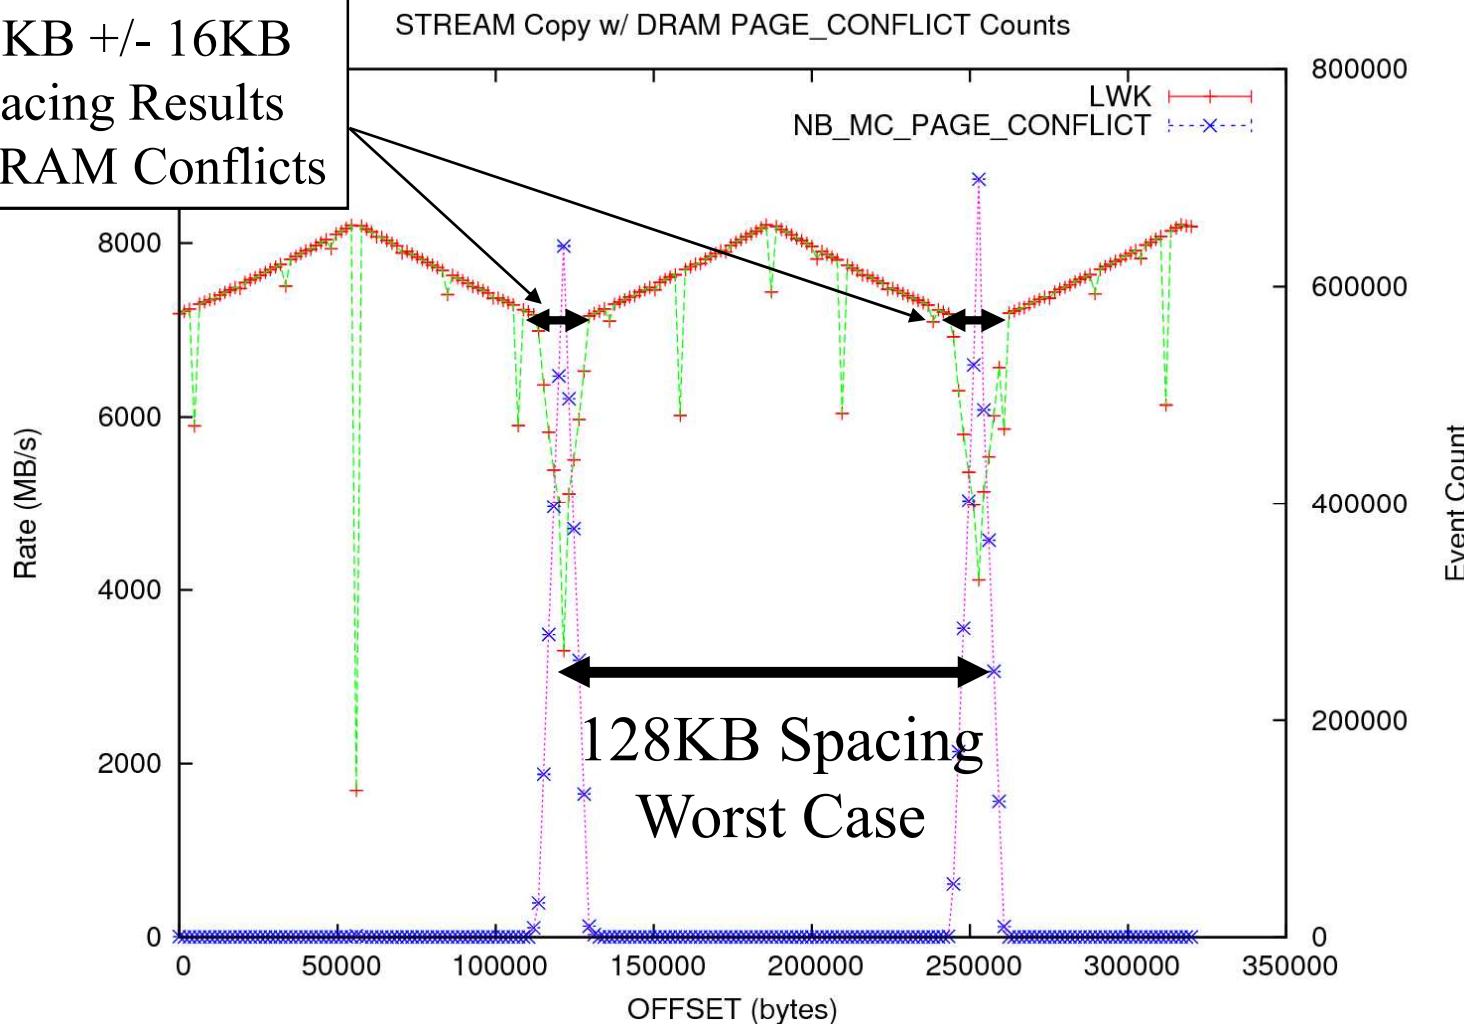
$$a[i] = b[i] + c[i]$$

TRIAD

$$a[i] = b[i] + \text{scalar} * c[i]$$

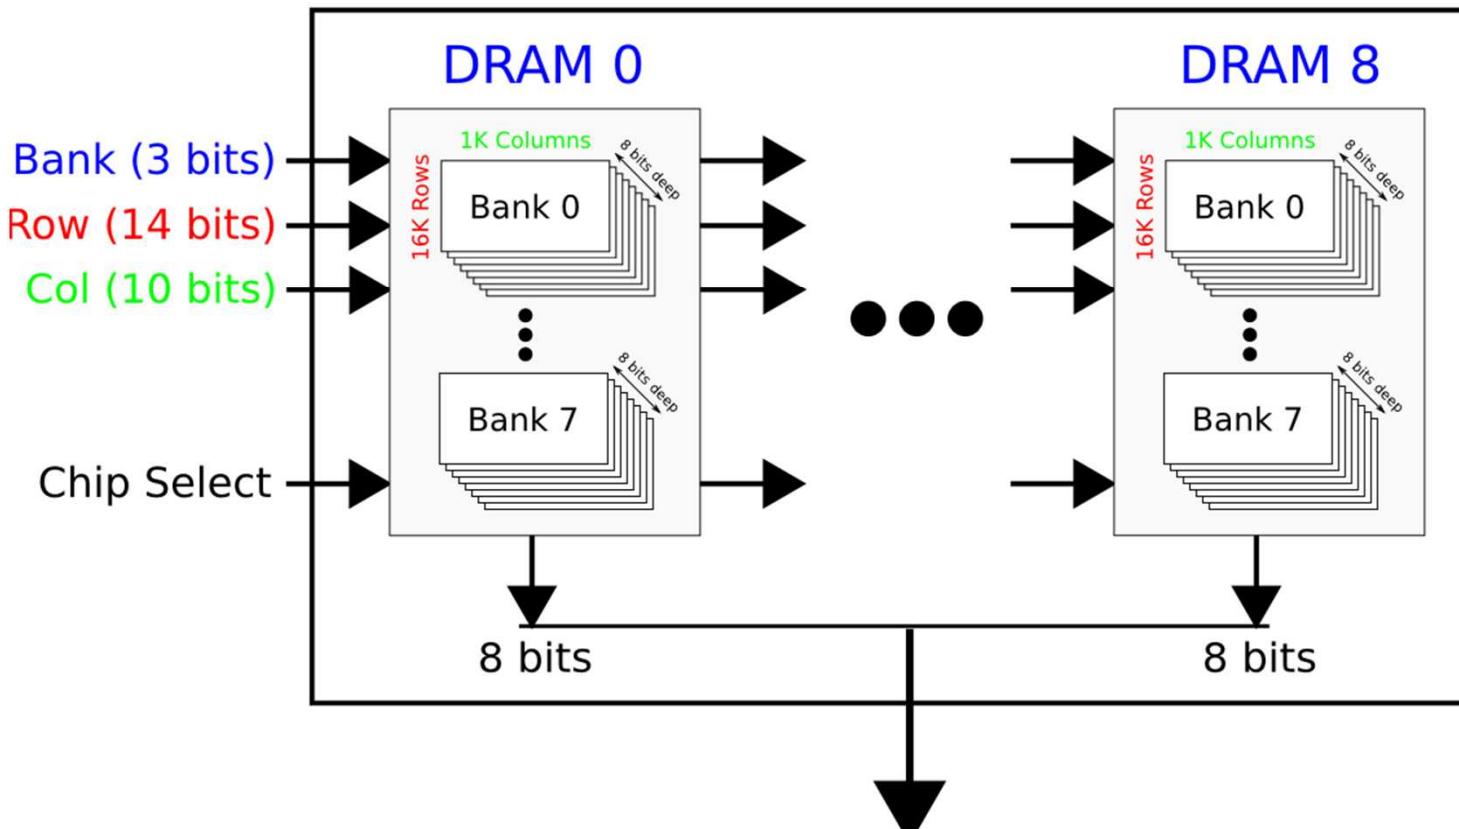
Mystery Solved

128KB +/- 16KB
Spacing Results
In DRAM Conflicts



DDR2 DIMM Architecture Example

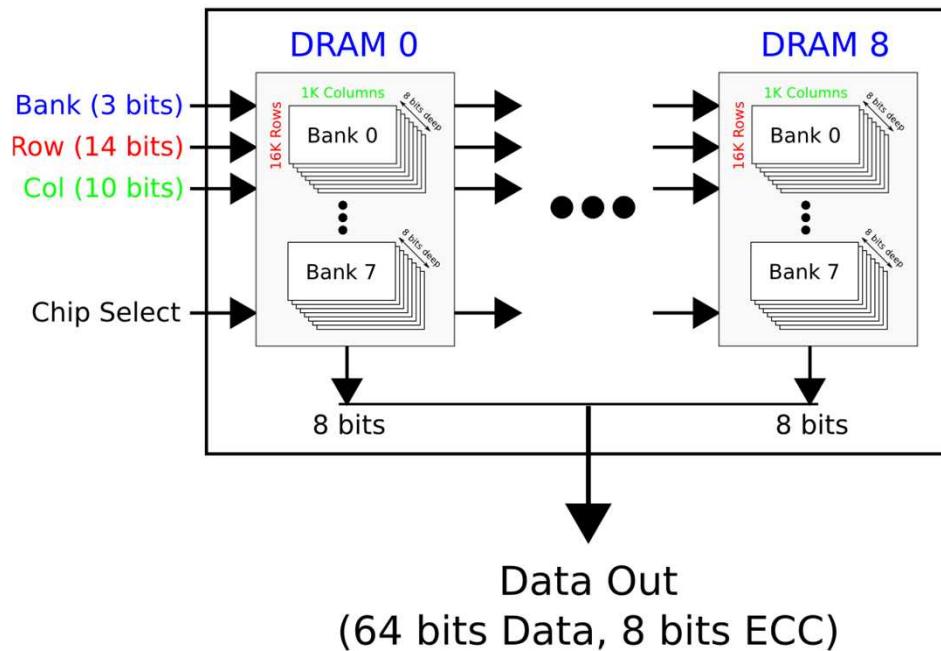
72-bit Wide DIMM (64-bit Data, 8-bit ECC)



Data Out
(64 bits Data, 8 bits ECC)

DDR2 DIMM Architecture Example

72-bit Wide DIMM (64-bit Data, 8-bit ECC)



Each DRAM row is
 $1\text{K columns} * 8 \text{ bits} = 1\text{K bytes}$

Each DIMM Row is
 $1\text{K bytes} * 8 \text{ chips} = 8\text{K bytes}$

Each Memory “Page” is
 $8\text{K bytes} * 2 \text{ DIMMs} = 16\text{K bytes}$

Addresses that are
 $16\text{K bytes} * 8 \text{ banks} = 128\text{K bytes}$
apart will result in a **Bank Conflict**
(Simultaneous accesses to
different rows in same
bank, aka Row Conflict)



Impact of Multi-core

- Main question: Does this matter for any apps?
- Lots of cores accessing same DRAM bank haphazardly decreases locality
 - HW memory controller can do better scheduling
 - OS can do something smarter
 - Application can do something smarter
 - Blocking, bringing in large unit-stride chunks into cache
 - Stream programming models