



The Effects of System Options on Code Performance

Courtenay T. Vaughan
ctvaugh@sandia.gov

Sandia National Laboratories
May 2007



Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.





XT3 Available Options

- **small vs. large pages**
 - Controls memory page size
 - large pages are default
 - `yod -small_pages ...`
- **eager vs. rendezvous message sends**
 - Controls message protocol
 - rendezvous protocol is default
 - `export MPI_PTL_EAGER_LONG=1`
- **Catamount malloc vs. GNU malloc**
 - Controls which malloc is being used
 - Link time option (`cc ... -lgmalloc ...`)



Test Parameters


- Codes were run for all combinations of options
- Codes were run using only one core per socket
- All results on a given number of processors for a given code were run using the same nodes on the machine
- Most of the results are from Red Storm (2.4 GHz processors) while some are from our test system (2.0 GHz processors)
- Most tests were run once
 - Early experience with the test system indicated that this was sufficient



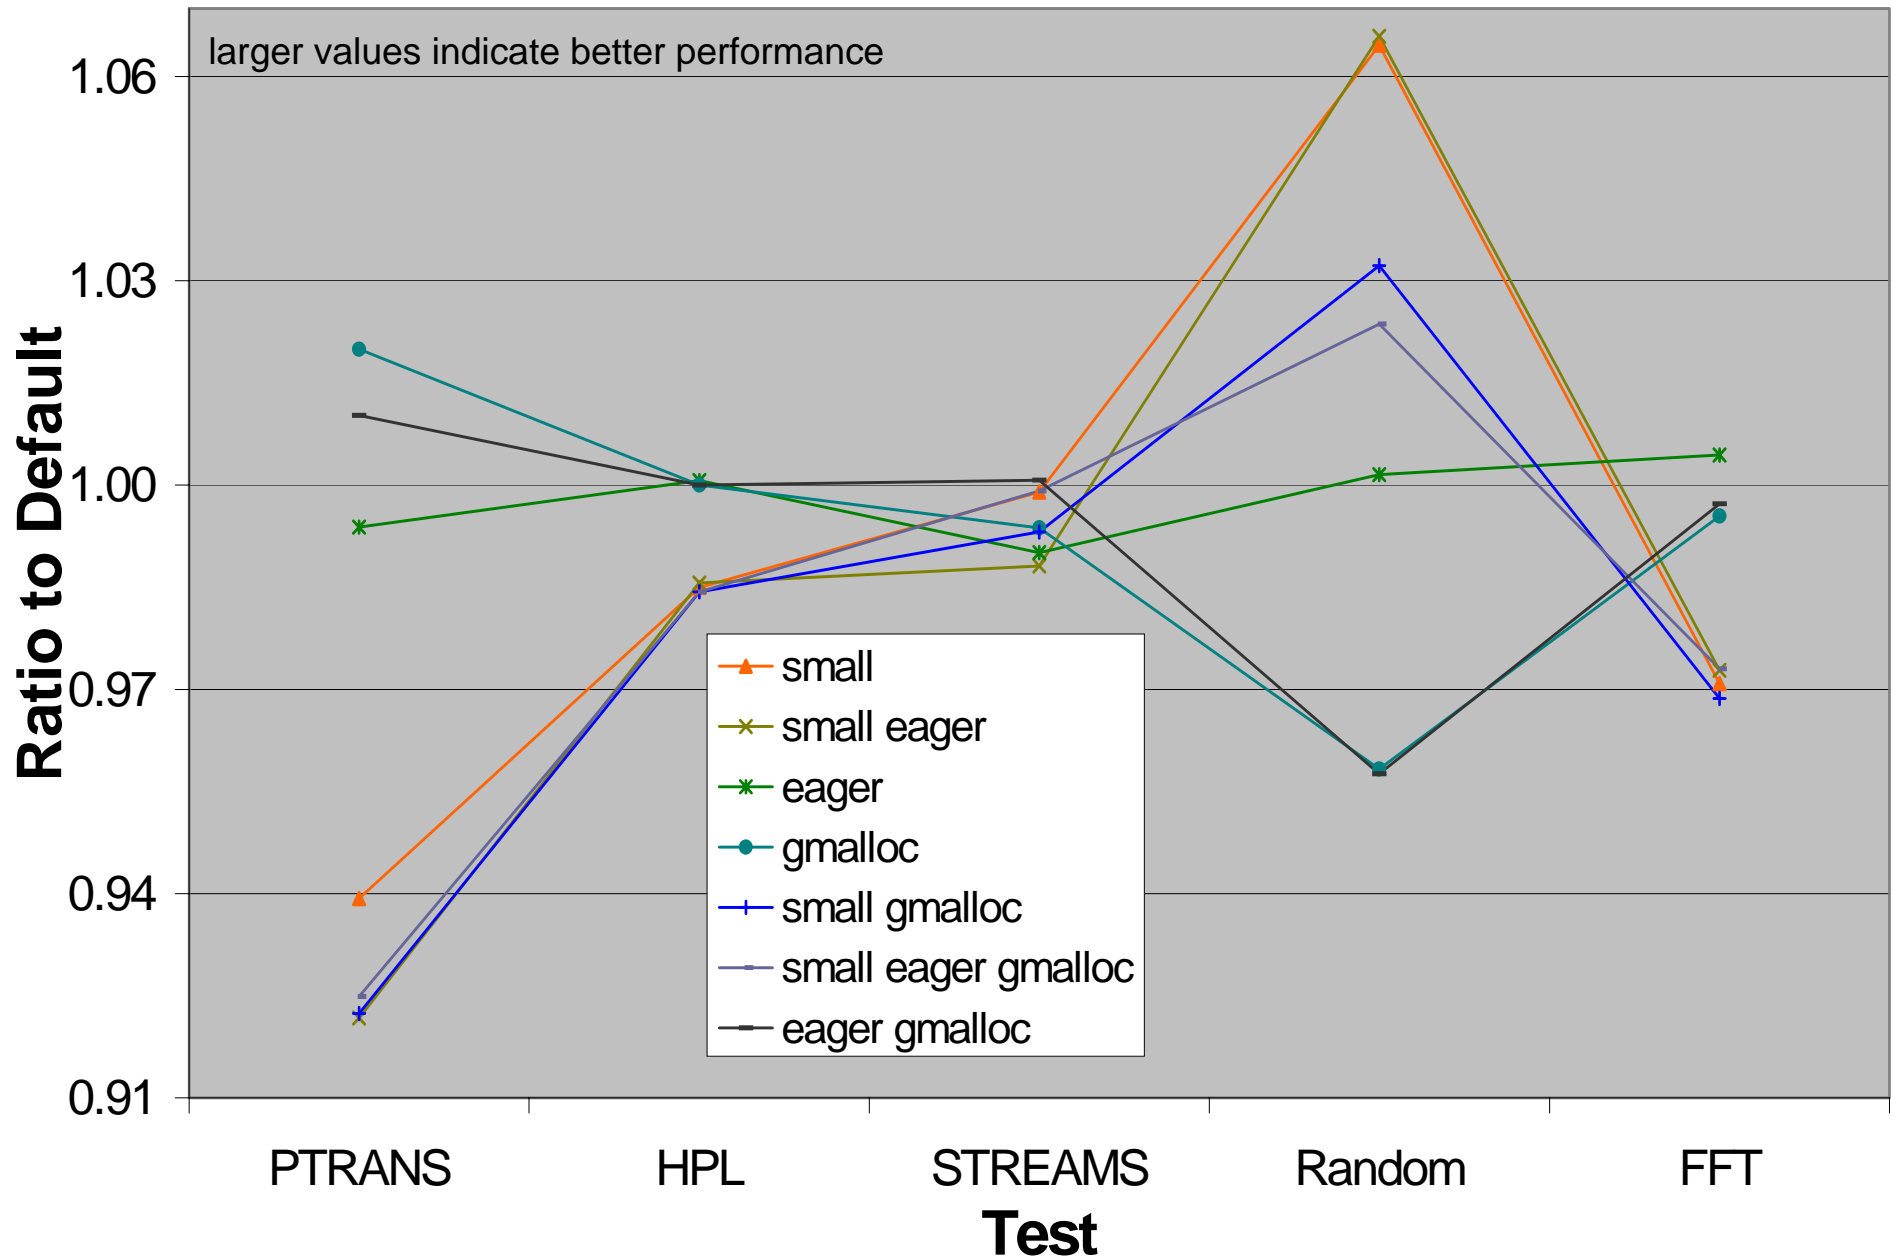
HPCC

- **Series of 7 benchmarks in one package. We are using 5 of them:**
 - **PTRANS - matrix transposition**
 - **HPL - Linpack direct dense system solve**
 - **STREAMS - Memory bandwidth**
 - **Random Access - Global random memory access**
 - **FFT - large 1-D FFT**
- **Code is C plus libraries**

HPCC - 64 processors, N = 80003




HPCC - 384 processors, N = 150035



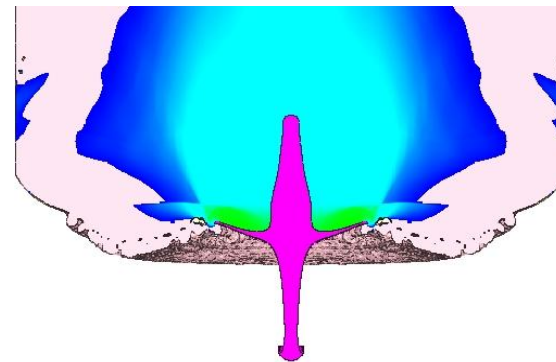


CTH

- Three-dimensional shock hydrodynamics code
- Ran in flat mesh mode - no AMR (Automatic Mesh Refinement)
- Shaped charge problem
- 90 x 216 x 90 cells per processor
- Code is mostly FORTRAN with a little C

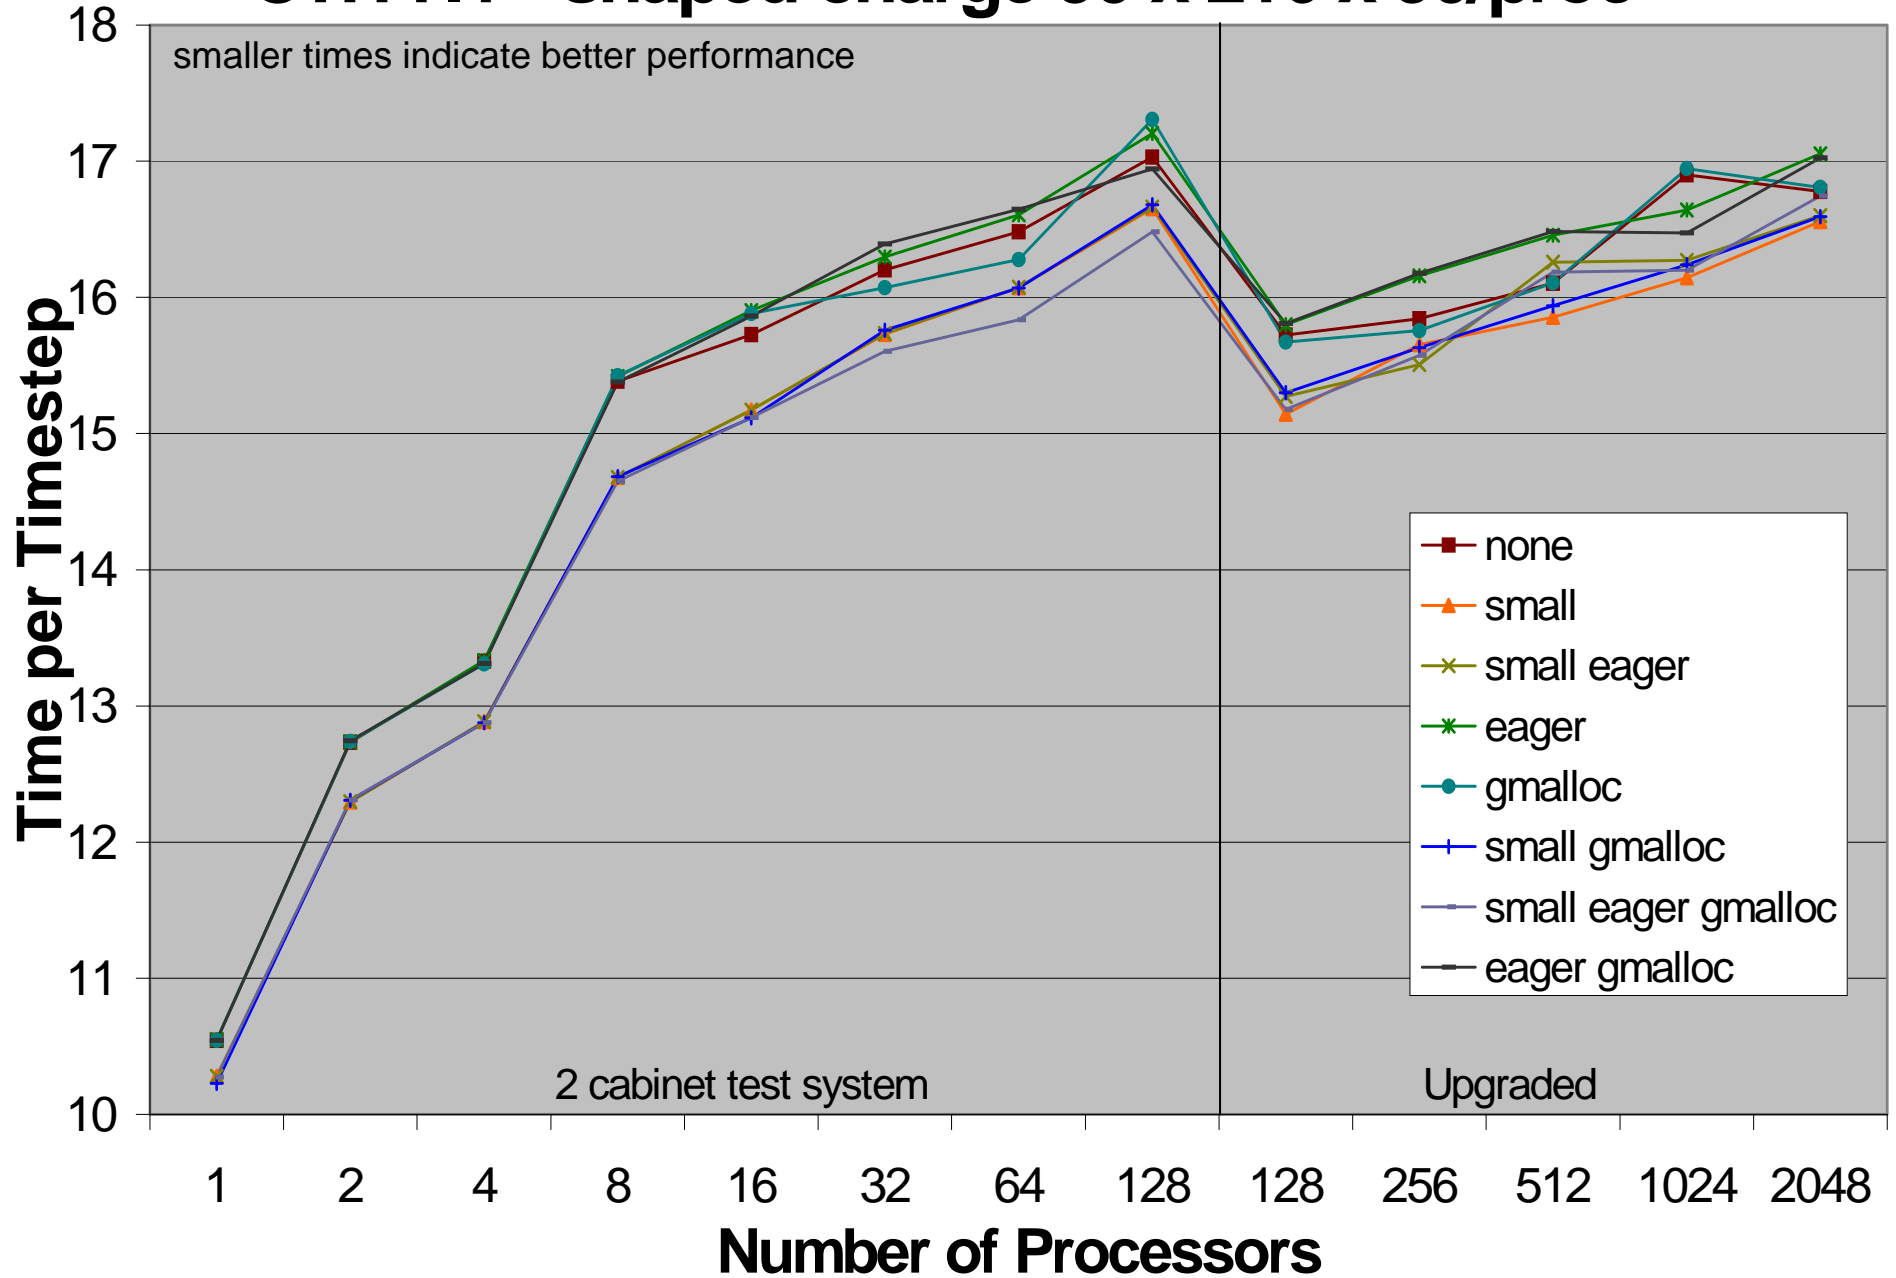


time = 0.0 ms



time = 0.3 ms

CTH 7.1 - shaped charge 90 x 216 x 90/proc

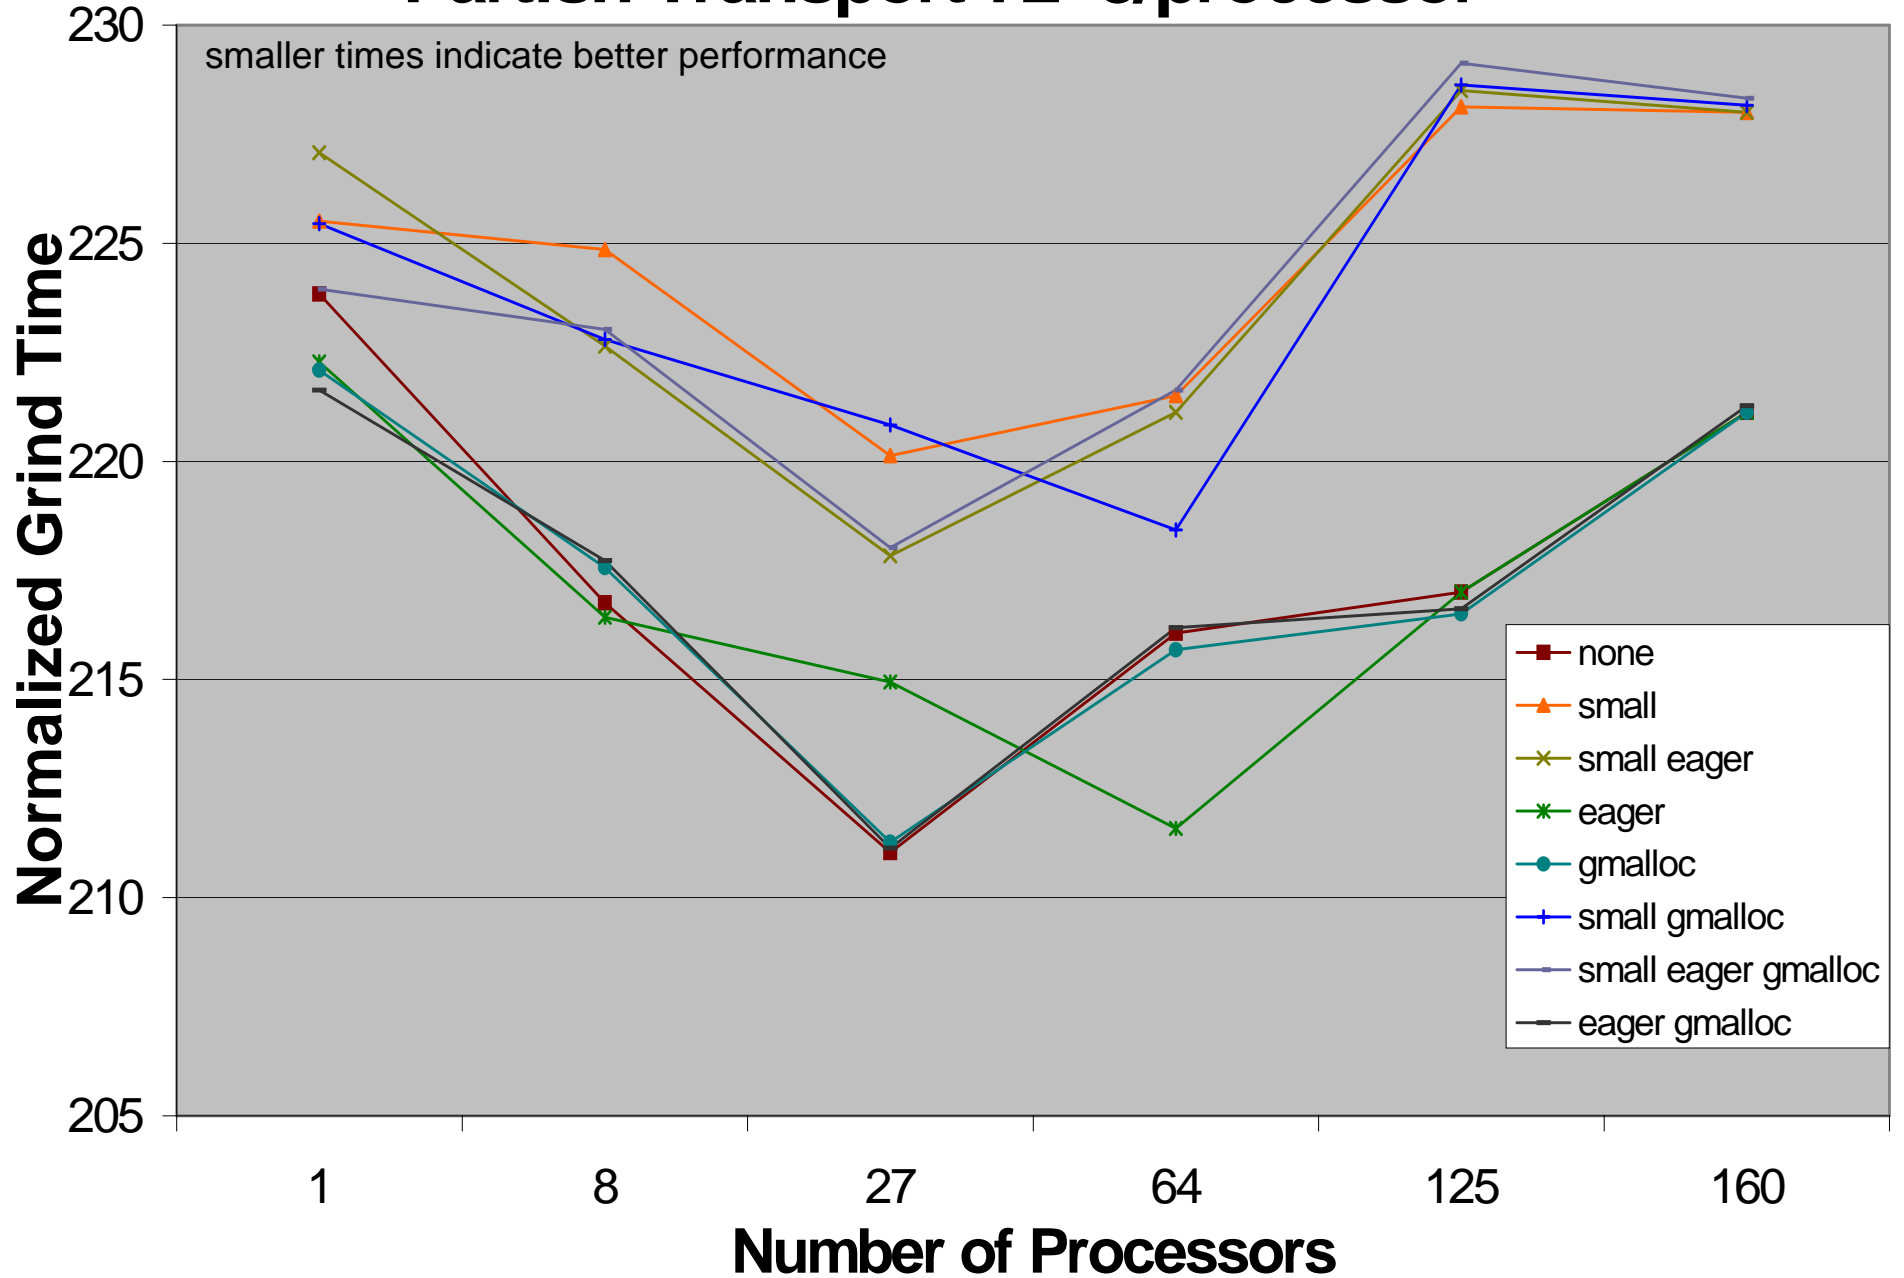




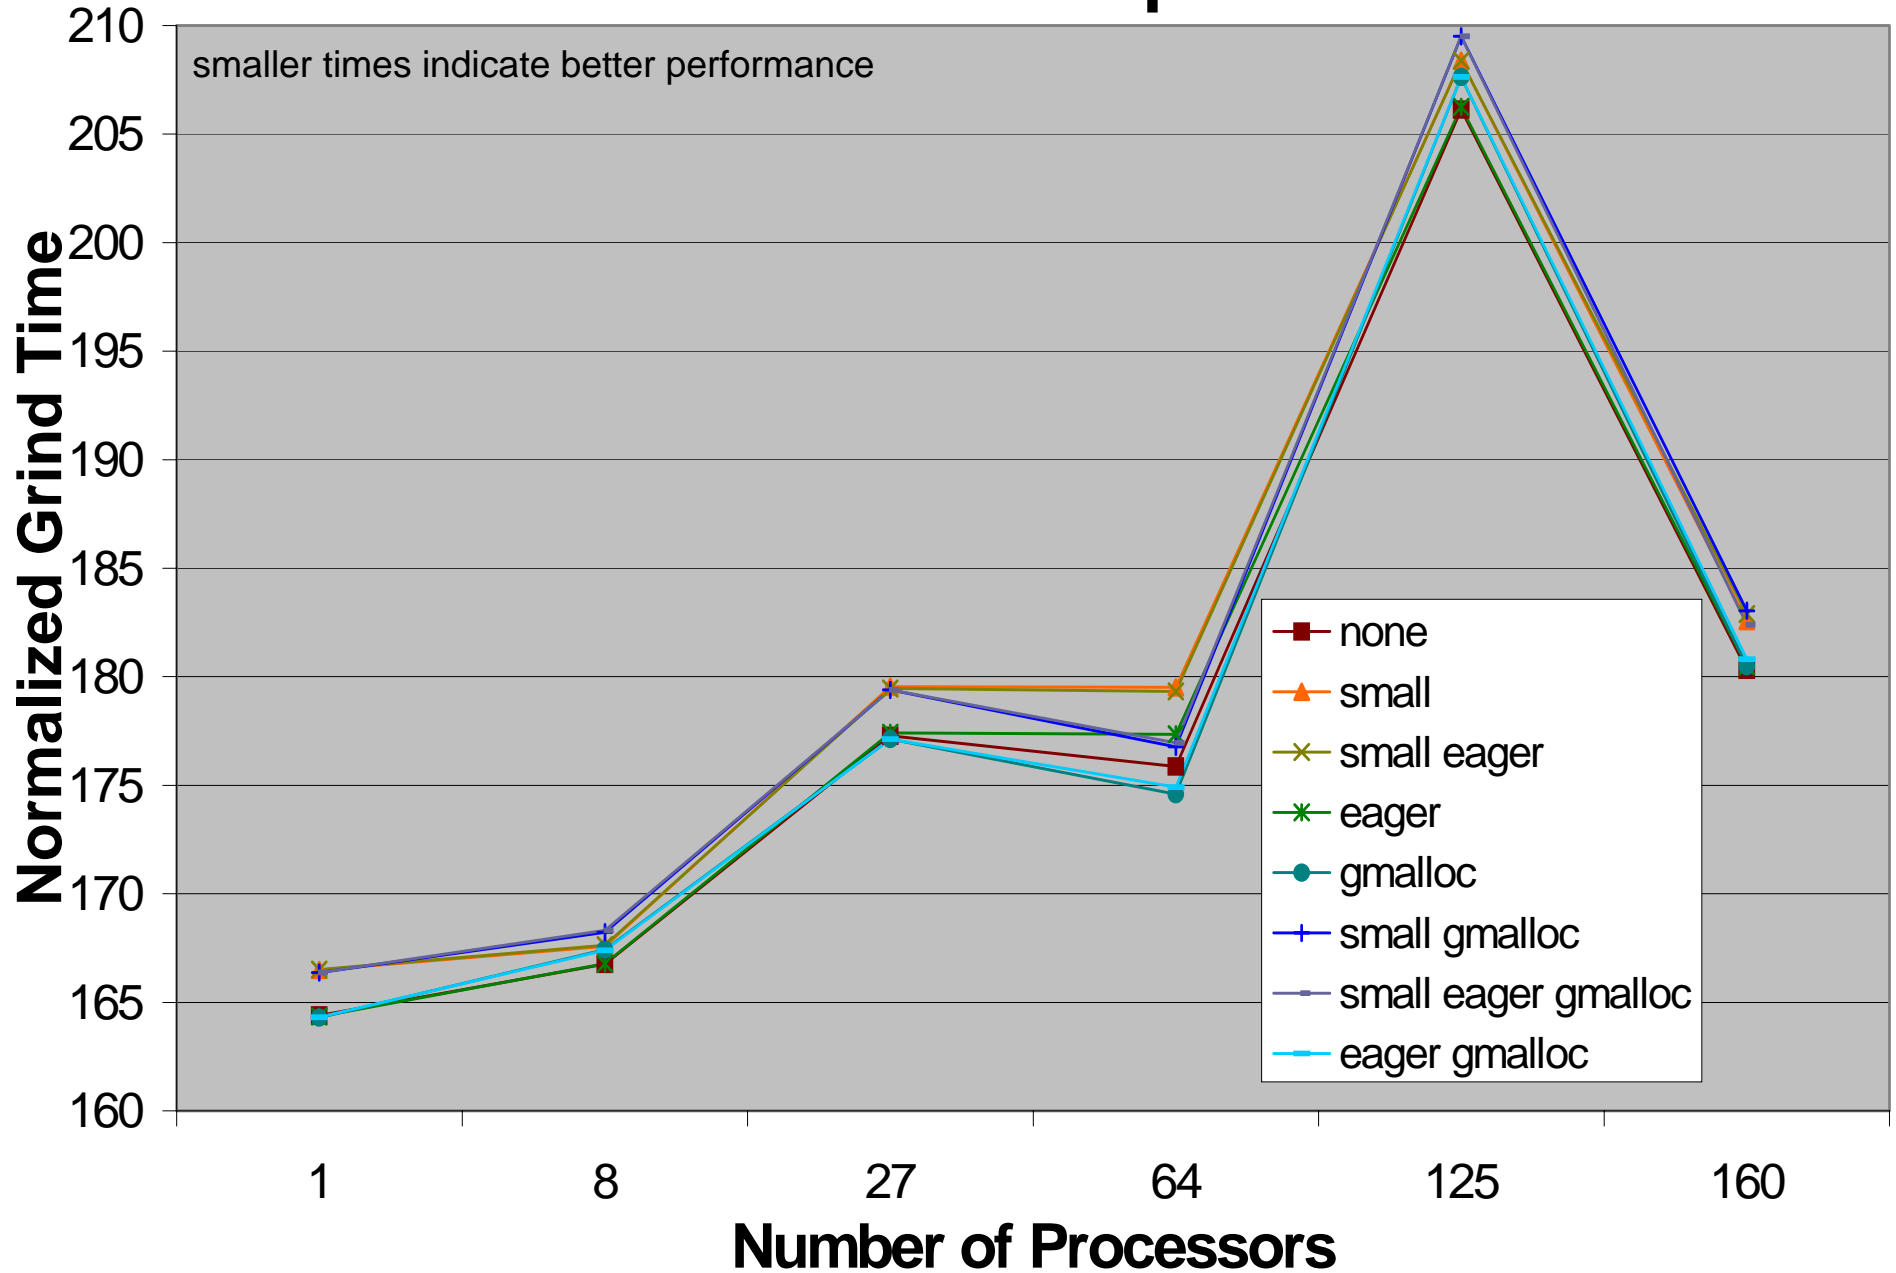
Partisn

- **LANL code that solves the Boltzmann transport equation**
- **Has a transport and diffusion phase**
- **SN timing problem with 72^3 cells per processor**
- **Run only on 2 cabinet test system**
- **Code is mostly FORTRAN**

Partisan Transport 72³/processor




Partisan Diffusion 72³/processor






Presto

- **Structural mechanics with contact algorithm**
- **Walls problem**
 - Two sets of two brick walls colliding
 - 10240 elements/proc
- **Each brick contained on a processor**
 - All communication for contact
- **Code is C++**



**32 processor problem
colored by processor number**

PRESTO - Walls, 10240 elem/proc, t = 5.0e-3

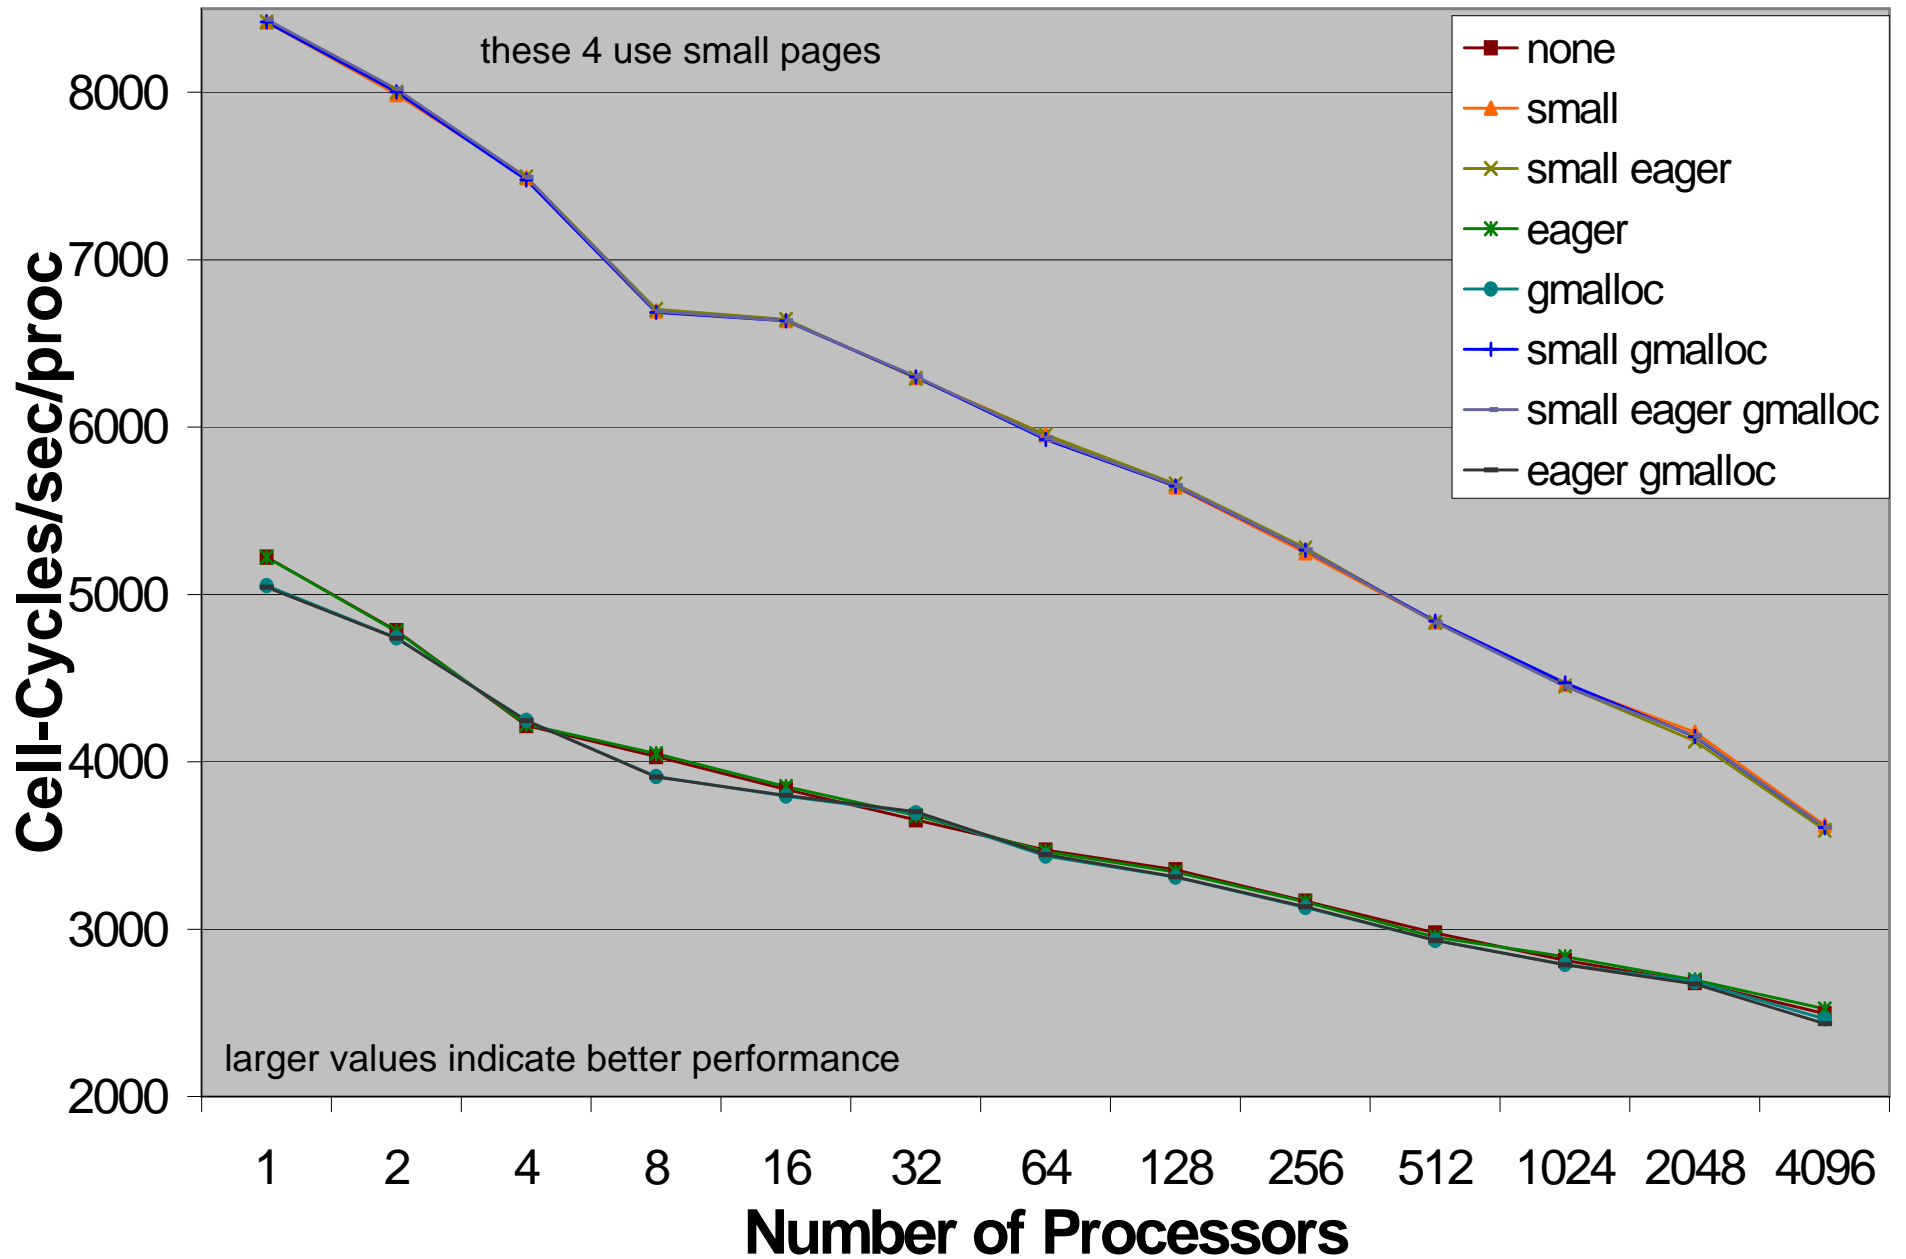




SAGE

- **LANL Eulerian Hydrocode with AMR (Adaptive Mesh Refinement)**
- **timing_c problem**
 - has adaptation and heat conduction
 - 250000 cells per processor
- **Code is mostly FORTRAN 90**

SAGE - timing_c - 250000 cells/proc





PAPI Results for Different Page Sizes

App	Large pages					Small pages				
	TLB miss	Cache access	Access /miss	Cache hit rate	FP intensity	TLB miss	Cache access	Access /miss	Cache hit rate	FP intensity
PTRANS	3.30e ⁷	6.37e ¹⁰	1934	98.2%	0.150	1.03e ⁹	6.54e ¹⁰	63.8	98.2%	0.144
HPL	5.57e ⁸	3.16e ¹²	5670	99.1%	1.693	2.98e ⁸	3.17e ¹²	10644	99.1%	1.686
PARTISN	1.48e ⁹	6.14e ¹¹	414	92.8%	0.558	1.38e ⁹	6.18e ¹¹	449	92.7%	0.554
STREAM	1.73e ⁴	4.25e ⁹	2.45e ⁵	93.5%	0.344	6.97e ⁶	4.34e ⁹	623	93.6%	0.333
FFT	2.03e ⁷	3.34e ⁹	164	97.6%	0.466	1.52e ⁷	3.30e ⁹	217	97.6%	0.441
CTH	1.08e ¹⁰	1.28e ¹²	119	96.9%	0.570	3.30e ⁹	1.32e ¹²	398	97.0%	0.588
Random	8.50e ⁸	4.32e ¹⁰	50.8	99.5%	0.0	1.01e ⁸	4.57e ¹⁰	452	99.5%	0.0
SAGE	2.90e ⁹	1.29e ¹¹	44.5	98.7%	0.226	2.29e ⁷	1.02e ¹¹	4457	98.3%	0.231

codes ranked by large page preference



Observations from PAPI Data

- The L1 data cache hit rate does not seem to change with different page sizes
- The floating point intensity of a calculation has no bearing on which page size is better
- The number of cache accesses per TLB miss can indicate something about which page size to use
- The number of L1 data cache accesses can vary with the page size



Summary

- **No set of options is always best**
- **Results from benchmarks do not necessarily translate to codes**
- **Small pages generally helps and can help significantly**
- **GNU malloc helps C++ codes**
- **The eager message protocol has small effects for the codes that were tested**