

CONF-9706144--1

Title:

Performance Characterization and Validation of ASCI
Applications: A Memory Centric View

Author(s):

Olaf M. Lubeck
Yong Luo
Harvey Wasserman
Federico Bassetti

RECEIVED

AUG 14 1997

OSTI

MASTER

Submitted to:

PAID '97
Denver, Colorado
June 1, 1997

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

**Los Alamos
National Laboratory**

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

hj

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

Performance Characterization and Validation of ASCI Applications: A Memory Centric View

Olaf M. Lubeck
Yong Luo
Harvey Wasserman
Federico Bassetti

Los Alamos National Laboratory

Extended Abstract

Performance and scalability of high performance scientific applications on large scale parallel machines are more dependent on the hierarchical memory subsystems of these machines than the peak instruction rate of the processors employed. The dependence is likely to increase in the future [1,4]. While single-processor performance may double every eighteen months, memory bandwidth increases by only 15% during the same period.

In addition, distributed shared memory (DSM) architectures are now being implemented which extend the concept of single-processor cache hierarchies across an entire physically-distributed multi-processor machine. Machines which will be available to the Department of Energy's Accelerated Strategic Computing Initiative (ASCI) can have as many as 128 processors in a single DSM. Scalability of these machines to large numbers of processors is ultimately tied to issues of memory hierarchy performance, which includes data migration policies and distributed cache coherence protocols. Investigations of the performance improvements of applications over time and across new generations of machines must explicitly account for the effects of memory performance.

In this paper, we characterize application performance with a "memory-centric" view. The applications are a representative part of the ASCI workload. Using a simple Mean Value Analysis (MVA) strategy and observed performance data, we infer the contribution of each level in the memory system to the application's overall performance in cycles per instruction (CPI). Our empirical model accounts for the overlap of processor execution with memory accesses.

Measurements of ASCI codes were obtained on the latest Origin 2000 and Power Challenge machines from Silicon Graphics, Inc. (SGI). These machines provide a very interesting perspective since they use the identical MIPS R10000 processor but differ significantly in the design of the memory subsystems [2,3].

The full paper describes the parts of the machine architecture relevant to this work, codes from the ASCI workload, the model and empirical methodology, validation of the model using a combination of measurement and simulation, results, analysis and major conclusions.

In this abstract, we will provide an abbreviated discussion of the model and present one characterization of the ASCI codes that was obtained using the model.

The application analysis in the paper uses a typical mean value parameterization [5] to separate CPU execution time from stall time due to memory loads/stores. The key issue in

the analysis is to determine the amount of memory access time that is overlapped by computation.

The model projects the overall CPI of an application as a function of CPU execution time and average memory access times:

$$CPI = CPI_0 + \sum_{i=1}^{nlevels} h_i * t_i \quad (1)$$

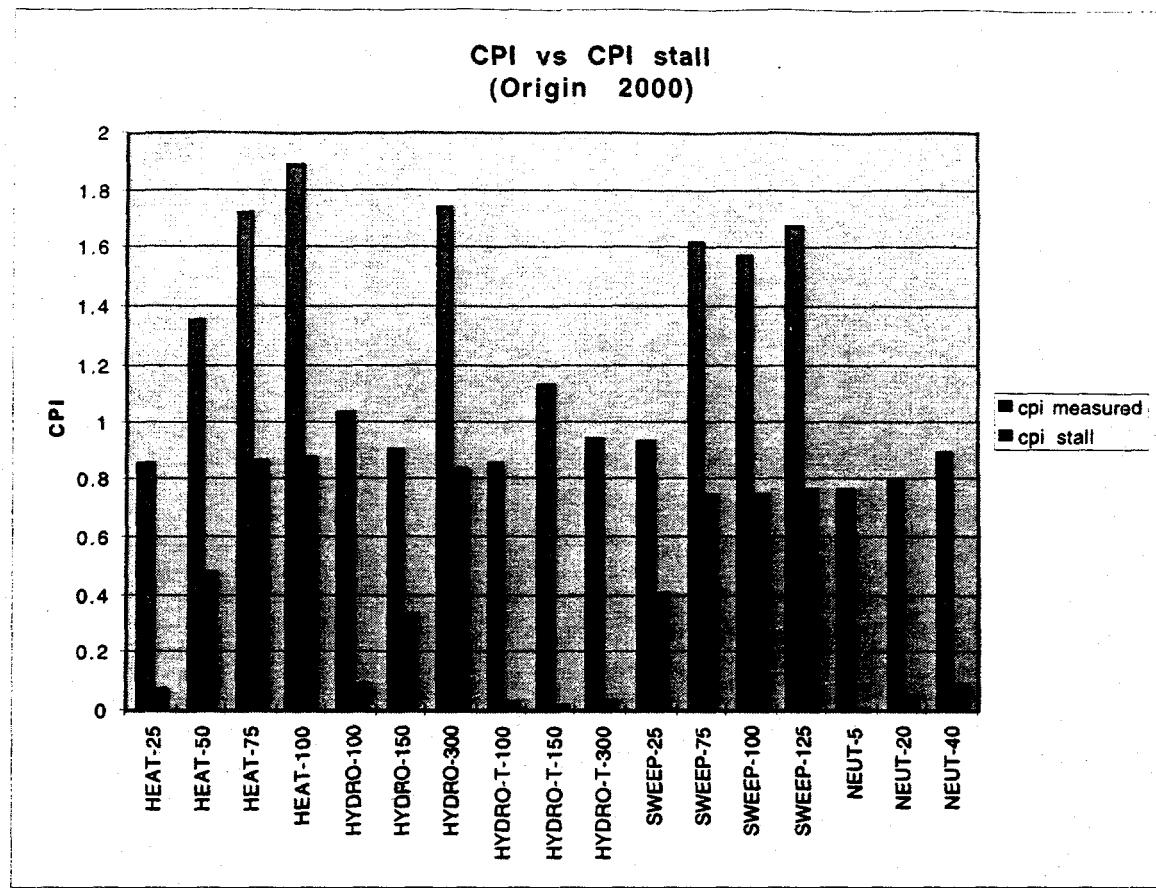
where CPI_0 is the CPI of the application assuming that all memory accesses take 1 clock period (CP), and h_i and t_i are, correspondingly, the hits per instruction and average non-overlapped access times for the i th level cache.

If no overlap of CPU execution and memory access occurs, every memory access to the i th level incurs the full round-trip latency, which we denote as T_i . We define a measure of the overlap of memory accesses with computation as m_0 , where

$$CPI = CPI_0 + (1-m_0) \sum_{i=1}^{nlevels} h_i * T_i \quad (2)$$

Although this overlap is not directly measurable, using the relationship in Eq. 1, we can infer the overlap for an individual application by performing a least-squares fit of the model parameters to observed execution data (CPI, h_i) for that code. CPI and hits are obtained from the MIPS R10000 hardware performance counters [6]. Validation of the inferred model parameters is accomplished by using the model to predict performance on a different machine configuration. In addition, confidence in the methodology is further established with an independent measurement of CPI_0 using an R10000 simulator made available from SGI [7] and from direct measurement of the overall CPI of a small problem that fits in the L1 cache.

The figures below show the results in terms of overall CPI and non-overlapped memory access stall cycles. The results suggest that the ASCI applications on the Power Challenge are indeed dominated by memory access time with two exceptions (HYDRO-T and NEUT). The Origin data suggest a significant improvement in the memory sub-system over the Power Challenge. Overall CPIs have decreased by more than a factor of two and stall time is proportionately less.



1. Wulf, W. A. and McKee, S. A. "Hitting the Memory Wall: Implications of the Obvious," University of Virginia Department of Computer Science Technical Report 19??, available from the author at wulf@virginia.edu.
2. MIPS Technologies, Inc., "R10000 Microprocessor Product Overview."
3. Yeager, K. C., "The MIPS R10000 Superscalar Microprocessor," IEEE Micro, April, 1996, pp 28-40.
4. Luo, Y., Lubeck, O.M., and Wasserman, H. J., "Preliminary Performance Study of the SGI Origin2000," Los Alamos National Laboratory Unclassified Release LA-UR -334, 1997.
5. Vernon, M.V, Lazowska, E. D., and Zahorjan, J., "An Accurate and Efficient Performance Analysis Technique for Multiprocessor Snooping Cache-consistency Protocols," in Proc. 15th Annu. Symp. Comput. Architecture, Honolulu, HI, June, 1988, pp 308-315.
6. Zagha, M., Larson, B., Turner, S., and Itzkowitz, M., "Performance Analysis Using the MIPS R10000 Performance Counters," Proc. Supercomputing '96, Pittsburgh, PA, December, 1996, IEEE Computer Society, Los Alamitos, California, pp -
7. Private communication, Steve Turner, Silicon Graphics, Inc., January, 1997.

**CPI vs CPI stall
(Power Challenge)**

