

Modeling the Performance of the Gaussian Computational Chemistry Code on the x86 Architecture

J. Antony¹, M. J. Frisch², and A. P. Rendell³

Abstract: Gaussian is a widely used scientific code with application areas in chemistry, biochemistry and material sciences. Computational methods implemented by Gaussian require the evaluation and manipulation of many integrals, representing the interaction between and among electrons and nuclei. For this rate-limiting step to proceed efficiently, cache blocking of batches of integrals is utilised. Larger blocks generate fewer operations, but result in greater cache misses. At present, cache blocking is statically determined.

In this study, hardware performance counters and cache simulation is used to characterise the cache behaviour of Gaussian. A simple linear performance model (LPM) is proposed

$$Cycles = \alpha * (I_{Count}) + \beta * (L1_{Misses}) + \gamma * (L2_{Misses})$$

where parameters α , β , γ are determined using a least squares fit of cycles to instruction counts (I_{Count}) and total L1, L2 cache miss counts ($L1_{Misses}$, $L2_{Misses}$). Intuitively α corresponds to a measure of how well the code is using the superscalar nature of the processor, β to the average cost in cycles for a L1 cache miss and γ to the average cost in cycles for a L2 cache miss.

Results are presented for three x86 processors – the Intel Pentium M, P4 and the AMD Opteron. Parameters obtained from a smaller calculation, when used to estimate cycle counts of larger ones, yield good estimates for the Pentium M and Opteron, but not for the P4. Possible reasons for the P4 discrepancy are discussed.

Using the LPM model the performance of Gaussian on x86 architectures with different cache sizes is explored by combining instruction and cache miss counts obtained from cache simulation (using Valgrind/Cachegrind) with relevant α , β and γ values. These results show that larger cache line sizes are beneficial.

^{1,3} Department of Computer Science, Australian National University
CSIT Building, # 108, North Road, Canberra, ACT, Australia
joseph.antony@anu.edu.au, alistair.rendell@anu.edu.au

² Gaussian Inc.
340 Quinnipac Street, Bldg. 40, Wallingford, CT 06492, USA