ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

M. Troyer
ETH Zürich, HIT G 31.8
CH-8093 Zürich

Set 7 - Hybrid MPI+OpenMP and Autocorrelations

Issued: April 27, 2015
Hand in: May 4, 2015

Question 1: Hybrid MPI+OpenMP

In previous exercises we parallelized the 2D diffusion code for distributed memory machines and for shared memory machines. In this exercise we will parallelize the same code using hybrid parallelism.

a) Include shared memory parallelism in the 2D diffusion code parallelized with MPI (you can start from `skeleton/diffusion2d_mpi.cpp`). You can use any of the shared memory programming models learned in the lecture.

A straightforward way to include threading into the code is to use OpenMP loops for the computation of the local domain.

Hybrid code: `solution/diffusion2d_hybrid.cpp`

b) Compare the performances of the pure MPI code and the hybrid code, Use different configurations to run the hybrid code, but use always a total of 48 cores (two nodes on Euler). Comment on the results.

Figure 1 presents the comparison of different runs with a total of 48 cores of the hybrid application. The horizontal green line is the elapsed time of the pure MPI program.

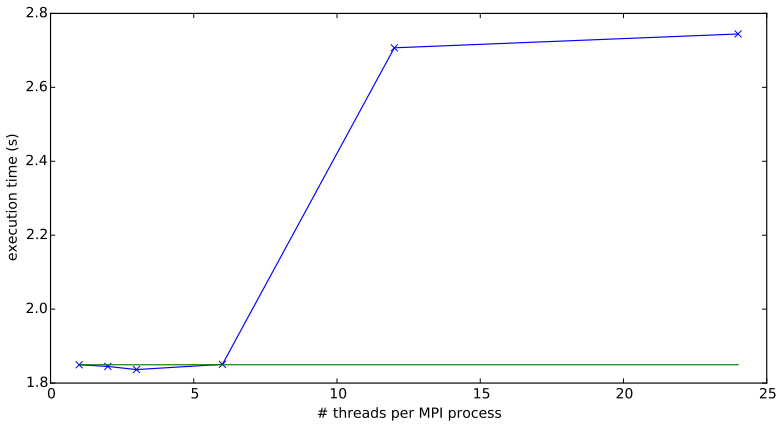

Figure 1: Comparison of different runs of the hybrid application (domain size $2^{14} \times 2^{14}$). Each run use a total of 48 cores.

The main effect shown in the figure is due to cache. Since we store the data in a row-wise format and the structure of the stencil operator requires reading data from the lower and upper rows, we effectively profit from cache only when two rows (of the local domain) fit in cache. By increasing the number of threads per process, we reduce the number of MPI processes, which directly affects the domain decomposition by making the blocks larger and we loose cache optimization.

Question 2: Dogs and Fleas

The dogs and fleas model[1] is a nice toy model to show the problems of autocorrelations in Markov-Chain Monte Carlo simulations.

a) Implement the dogs and fleas as described in section III B) of the paper using naive error estimates.

b) Implement the binning analysis method to obtain more reliable error estimates and compare the results.

   Note: For simplicity, you may assume that the total number of measurements is a power of two.

Summary

Summarize your answers, results and plots into a short PDF document. Furthermore, elucidate the main structure of the code and report possible code details that are relevant in terms of accuracy or performance. Send the PDF document and source code to your assigned teaching assistant.

[1] `http://dx.doi.org/10.1119/1.3247985` (accessible from within the ETH network)