8th. World Congress on Computational Mechanics (WCCM8)
5th. European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS 2008)
June 30 – July 5, 2008
Venice, Italy

ENHANCING DATA LOCALITY AND INTERFACE COMMUNICATION IN EDGE-BASED COMPUTATIONS

***Marcos A.D. Martins[1], Renato N. Elias[2], José J. Camata[2], Alvaro L.G.A. Coutinho[2]**

[1] Petrobras Research Center
Av. Horacio Macedo 850,
Rio de Janeiro, Brazil
marcos.martins@petrobras.com.br

[2] Center for Parallel Computing and
Dept. of Civil Engineering,
Federal University of Rio de Janeiro
PO Box 68506, Rio de Janeiro, RJ 21945, Brazil
{renato, camata, alvaro}@nacad.ufrj.br

Key Words: *Data reordering, Unstructured grids, Edge-based finite elements, Collective Communication.*

ABSTRACT

The performance optimization of implicit unstructured grids codes based on Krylov solvers has many issues. Among them, enhancing data locality and interface communication are two important factors. Data locality can be improved by node, edge and element reordering techniques. The combination of data reordering techniques with edge-based data structures generates lots of algorithmic possibilities. Determining which one yields the best performance solution for a given hardware and software platform is not a trivial task [1]. Code performance depends on many factors: computational architecture, compiler options, data layout, number of degrees of freedom per node (i.e. algorithm complexity) and the algorithm structure itself. According to previous experiments [1], there is no ultimate data layout known prior to execution, unless a probe is performed, since the many possible combinations may produce unexpected results.

This work presents a simple method to determine which data layout is the best one in terms of processing time for a given computer platform. This method is simply based on the choice of the best results after probing the several possible data layouts. The probing is performed in matrix-vector product loops at execution time. Considering a particular parallel platform, as clusters (heterogeneous or not), the best data layout can be different from each other and thus, different data structures can be used together for the same model. The various data layouts and matrix-vector product modes available can be set by a combination of the following concepts: nodal ordering (improving pipelining processing, bandwidth optimization and indirect addressing optimization), edge ordering (targeting indirect addressing access and memory hierarchy optimization) and loop configuration (chunk and lists length and loop unrolling) [6]. Preliminary probing results in the simulation of a sedimentary basin at northeast of Brazil, shows 256 possible data layouts. Results point towards non-standard data layouts coexisting among different processors during the analysis.

Concerning message passing, the standard approach is to partition data before

performing data reordering on each processor. Therefore, another aspect to be considered in parallel computations is interface (subdomain) communication costs. The performance of these operations are critical to high performance computing and depends on many factors, including for example, physical topology of the system, number of processes involved and message sizes [3]. Applications with overlapped computations and communications are less latency sensitive and thus can achieve a good parallel performance and scalability. In this context, non-blocking collective operations have been employed to avoid unnecessary synchronization [2].

Recently, techniques of floating point compression also have been studied as a mechanism to improve the communication performance. The aim is to reduce the message latencies reducing the message size [4]. Santhosh Kumar *et. al*. [5] shows the impact the both lossless and loss compression schemes on the scalability of an atmospheric application on clusters. In this work, we conjugate the non-blocking operations and compression data schemes to improve communication in our edge-based incompressible flow solver [7].

We are now undertaking computational experiments in cluster environments to study the combined effects of interface non-blocking overlapping communication and computation, floating point compression/decompression and probing in matrix-vector products and residual computations in edge-based unstructured grid computations. The results of these studies will be communicated at the time of the conference.

REFERENCES

[1] A.L.G.A. Coutinho, M.A.D. Martins, R.M. Sydenstricker, R.N. Elias, "Performance comparison of data reordering algorithms for sparse matrix-vector multiplication in edge-based unstructured grid computations", *Int. J. Numer. Meth. Engng*; Vol. **66**, pp. 431–460, (2006).

[2] T. Hoefler, P. Gottschilng, A. Lumsdaine, W. Rehm, "Optimizing a conjugated gradient solver with non-blocking collective operations", *Parallel Computing*, Vol. **33**, pp. 627-633, (2007).

[3] J. Pjesivac-Grbovic, T. Angskun, G. Bosilca, G.E. Fagg, E. Gabriel, J.J. Dongarra, "Performance analysis of MPI collective operations", *Cluster Computing*, Vol. **10**, pp. 127-143, (2007).

[4] M. Burtscher, P. Ratanaworabhan, "High throughput compression of double-precision floating-point data", 2007 *Data Compression Conference*, pp. 293-302. March 2007.

[5] V. Santhosh Kumar, R. Nanjundiah, M.J. Thazhuthaveetil, R. Govindarajan, " Impact of message compress on the scalability of an atmospheric modelling application on clusters" *Parallel Computing*, 2007 (in press).

[6] M. Martins, R. Elias, A. Coutinho, "EdgePack: A Parallel Vertex and Node Reordering Package for Optimizing Edge-Based Computations in Unstructured Grids", *Lecture Notes in Computer Science*, Vol **4395**, pp. 292 – 304, (2007).

[7] R.N. Elias, A.L.G A. Coutinho, "Stabilized edge-based finite element simulation of free-surface flows", *Int. J. Numer. Meth. Fluids*, Vol. **54**, pp. 965–993, (2007).