

AMSC/CMSC 662 Course Project Ideas

Ramani Duraiswami

November 3, 2009

1 Your own project

This should be a project on something you are doing research on, and have good understanding of the underlying problem. The goal is to use some or all of the ideas in the course, and your background knowledge to try various methods of improving the solution of the problem. For this to be an acceptable project, but the approach and objectives must be clearly stated from the outset. Your project proposal should specify the problem your code will address, what parts have already been parallelized, if any, and what further progress you expect to make as part of your class project. The project report consists of a report of the improvement.

2 Fast kernel sums (or N-body problems)

The kernel sum problem may be defined as evaluating

$$s(y_j) = \sum_{i=1}^m \alpha_i \phi(x_i, y_j), \quad j = 1, \dots, n.$$

Here the points x and y are points in d -dimensional space \mathbb{R}^d . This problem arises in many domains, including vortex methods, molecular dynamics, gravitation, boundary element methods, kernel density estimation in statistics, kernel methods in machine learning, etc. Let $r = \|x - y\|$. Some examples of these functions are

$$\begin{aligned} &1/r, r, \exp(-r^2/h^2), \exp(-r^2/h^2), (1 + \sqrt{3}r/h) \exp(-\sqrt{3}r/h), \\ &\exp(-2 \sin^2(\pi r/h)), (1 - r^2/h^2) \times 1(r/h < 1) \end{aligned} \tag{1}$$

The goal would be to distribute the sum on a machine with both local shared memory parallelism, and distributed processing.

The complexity of this algorithm is quadratic, $O(mn)$. Results should be compared against the best single processor linear algorithm.

3 Parallel Page Rank

This can be a joint project for up to two people. Several algorithms for page rank exist (the basis of web search). the goal would be to implement them on a heterogeneous architecture.

Some links: <http://spinner.cofc.edu/~langvillea/PRDataCode/index.html?referrer=webcluster>

<http://www.stanford.edu/~dgleich/projects/ppagerank/index.html>

4 Sparse Matrix vector products and heterogeneous CG

This can be a joint project for up to two people. Sparse matrices arise in many domains. The cost of a sparse matrix vector product is proportional to the number of non-zero entries in a matrix. The first task is to create a cache-aware shared memory multiprocessor sparse matrix vector product, and then extend it to a distributed architecture. Once that is done the fast matrix-vector product can be used with iterative algorithms for linear system solution. In case the matrix is symmetric the conjugate gradient algorithm can be used.

5 Mixed shared memory GPU implementations of Linear Algebra algorithms

Algorithms such as QR, LU, Cholesky etc. can be implemented in a mixed shared memory (OpenMP), GPU version. This would involve cooperative work with my graduate students.

6 Other suggestions

<http://www.cs.utk.edu/~dongarra/WEB-PAGES/SPRING-2008/additional-projects.pdf>