2006 TAIWAN INTERNATIONAL SCIENCE FAIR

CATEGORY : Computer Science

PROJECT : A load-balancing strategy for coarse-grained tree searches as applied to fractal image compression

AWARDS : Computer Science Third Award

SCHOOL : Raffles Junior College

FINALISTS : Koh Pang Wei

COUNTRY : Singapore

Abstract of Exhibit
Taiwan International Science Fair

CATEGORY: **Computer Science**

TITLE: **A load-balancing strategy and depth-prediction heuristic for coarse-grained tree searches as applied to fractal image compression**

NAME: **Koh Pang Wei**

COUNTRY: **Singapore**

(This project received the first prize for the following local science competitions: the IBM-sponsored High Performance Computing Quest and the Institute of Electrical and Electronics Engineers' (IEEE) Science Symposium.)

Introduction and Purpose

An exact solution to many current computational problems, such as the famous Travelling Salesman Problem (TSP), require a complete tree traversal in order to determine. This is often unfeasible, as the time complexity of the tree traversal grows exponentially with the size of the input, thus leading to an essentially computationally intractable problem. The branch and bound technique is an approach commonly used to speed this process. It entails dynamically pruning off branches of the tree in which the answer is probably not found in, hence reducing the amount of data that is needed to be traversed and the total time and resources required to perform the computation.

In this paper, we introduce a new load-balancing strategy for the execution of such a branch and bound algorithm in parallel, using a three-tiered hierarchical approach, to perform fractal image compression, which is essentially a complete tree traversal problem. This novel heuristic is aimed at achieving optimal load-balancing and minimising unnecessary network traffic and bottlenecking, which functions by predicting the optimum search depth and hence controlling the coarseness of the input that is assigned to each worker node. Our scheme additionally enables us to tailor to the specifications of different clusters, as the heuristic is adjusted based on network speed and processor speed, which vary appreciably from cluster to cluster. We further discuss how to apply our method to other large tree search problems, such as the TSP and other NP-complete problems.

We have also enhanced an existing load-balancing strategy outlined in Crivelli *et. al.* (2004, IBM Journal of Research and Development), by prioritising the re-allocation of idle worker nodes such that supervisors who are in need of more help receive a larger share of the free workers.

Procedures

Using the Message Passing Interface protocol, we created a fractal image compressor with the aforementioned load-balancing strategies, and tested compression times on different images, with varying amounts of processors, as well as different constant factors that we multiplied the granularity of the input given to each worker node by, in order to test the efficacy of our heuristic. The running times of our algorithm were benchmarked against a static load-balancing system, and a two-tiered master-slave approach. Every image was decompressed after compression for quality testing.

Data and Discussion

Through empirical simulation, we verified that our new load-balancing and depth-prediction strategy improved running times of the compression algorithm significantly. Our approach was also more scalable, that is, able to support more processors on the cluster while retaining a decent amount of speedup per processor, since our reassignment of idle workers ensured that all the processors were being fully utilized, with a minimum of network overhead expended in the reassignment.

Furthermore, using a dynamic depth-prediction heuristic provided a substantial amount of speedup over static depth allocation, by enabling a more efficient partitioning of the search tree amongst worker nodes. The exact value of this speedup varies from picture to picture, but we were able to cut down execution times by approximately a third, as the dynamic system allows for greater flexibility with the depth of the search tree assigned to each worker node, depending on the situation.

Conclusion

Our results show that the new load-balancing and depth-prediction strategy do indeed provide a significant speedup over more conventional approaches. They are widely applicable to other computational problems, even to fields such as bioinformatics, e.g. protein folding, and we believe that as such, our load-balancing method will be useful in many industrial and research-based applications.

評語

Controlling the granularity of the problem input assigned to each computing wade is a feasible solution to speed up the exertion time of parallel problem . The result of this project can be imparked by comparing with the outcome of running in a PC cluster in which the granularity is controlled by the OS.