Studying the impact of synchronization frequency on scheduling tasks with dependencies in heterogeneous systems

F. M. Ciorba, I. Riakiotakis, G. Papakonstantinou
School of Electrical & Computer Engineering, National Technical University of Athens, Greece

T. Andronikos
Department of Informatics Ionian University Greece

A. T. Chronopoulos
Department of Computer Science, University of Texas at San Antonio, TX USA

Many scheduling algorithms have been devised for nested loops with and without dependencies on general heterogeneous distributed systems ([1] and references therein). However, none addressed the case of dynamically computing and allocating chunks of non-independent tasks to processors. We propose a theoretical model that results in a function that estimates the parallel time of tasks in loops with dependencies on heterogeneous systems. We show that the minimum parallel time is obtained with a synchronization frequency that minimizes the function giving the parallel time. The accuracy of the model is validated through experiments from a practical application. For more details refer to [2].

To find the optimal synchronization frequency, we build a theoretical model for heterogeneous dedicated systems, in which workers have different computational powers. Loops with dependencies are efficiently scheduled on heterogeneous systems with self-scheduling algorithms [1]. The self-scheduling algorithms are based on the *master-worker* model. The master assigns work to workers upon request. Due to the data dependencies, applying self-scheduling algorithms to loops with dependencies yields a *pipelined* parallel execution. In the case of one master and NP workers, each assignment round corresponds to a pipeline with NP stages. Our approach assumes that the nested loop is represented in Cartesian space with at least 2 dimensions. One dimension is partitioned by the master into chunks according to a self-scheduling algorithm. In a pipeline organization, each worker synchronizes with its neighbors. Thus, synchronization points are inserted along the other dimension. A synchronization interval, denoted by h, represents the number of elements in the index space along the synchronization dimension. Data produced at the end of one pipeline are fed to the next pipeline. It is obvious that the synchronization frequency plays an important role in the total parallel time. Frequent synchronization implies excessive communication, whereas infrequent synchronization may limit the parallelism.

In order to estimate the theoretical parallel time on a heterogeneous system for the case of multiple assignment rounds (pipelines), i.e., the number of processors is smaller than the total number of chunks, we assume that a problem of the original index space size can be decomposed into p subproblems (pipelines) of (equal) size in which each processor is assigned one chunk. Thus, one subproblem corresponds to one assignment round. These subproblems are inter-dependent in the sense that (part of) the data produced by one subproblem are consumed by the next subproblem. Upon completion of one subproblem, the processor assigned the last chunk of the subproblem transmits (in a single message) all necessary data to the processor assigned the first chunk of the next subproblem. The time to complete this data transfer, represents the time to send and receive a data packet of size equal to the size of the synchronization dimension.

Hence, the theoretical parallel time is given by the sum of the parallel time of each subproblem, plus the data transfer time from one subproblem to the next, and the time necessary for for the master to assign work to every worker. The optimal value of h, that minimizes the parallel time is found by differentiating the estimated parallel time function with respect to h. Extensive experimental tests of our model show that using the optimal value of h as determined above, one can obtain an actual parallel time that if not minimum, is guaranteed to be very close to the actual minimum. The proposed model offers a feasible practical alternative to exhaustive testing of an application for estimating the optimal synchronization interval. More results are given in [2].

References

[1]   F.M. Ciorba, T. Andronikos, I. Riakiotakis, A.T. Chronopoulos and G. Papakonstantinou. Dynamic Multi Phase Scheduling for Heterogeneous Clusters, *Proc. of the 20th IEEE Int'l Par. & Dist. Proc. Symp. (IPDPS'06)*, Greece, 2006.
[2]   Studying the impact of synchronization frequency on scheduling tasks with dependencies in heterogeneous systems, http://www.cslab.ntua.gr/~cflorina/publications.

IEEE
computer
society