

Automatic mapping tasks to cores. Evaluating AMTHA Algorithm in multicore architectures.

Laura De Giusti¹, Emilio Luque², Franco Chichizola¹, Marcelo Naiouf¹, Armando De Giusti¹

¹Instituto de Investigación en Informática (III-LIDI) – School of Computer Sciences–UNLP.
Argentina

²Universidad Autónoma de Barcelona (UAB) - Computer Architecture and Operating System
Department (CAOS) Spain

ldgiusti@lidi.info.unlp.edu.ar, emilio.luque@uab.es, {francoch, mnaiouf, degiusti}@lidi.info.unlp.edu.ar

Abstract. The AMTHA (Automatic Mapping Task on Heterogeneous Architectures) algorithm for task-to-processors assignment and the MPAHA (Model of Parallel Algorithms on Heterogeneous Architectures) model are presented.

The use of AMTHA is analyzed for multicore processor-based architectures, considering the communication model among processes in use.

The results obtained in the tests carried out are presented, comparing the real execution times on multicores of a set of synthetic applications with the predictions obtained with AMTHA.

Finally current lines of research are presented, focusing on clusters of multicores and hybrid programming paradigms.

Keywords: Parallel Algorithm Models – Task-to-Processor Mapping - Performance Prediction – Multicore Architectures –

1 Introduction

A cluster is a parallel processing system formed by a set of PCs interconnected over some kind of network and that cooperate as if they were an “only and integrated” resource, regardless of the physical distribution of its components. When two or more clusters are connected over a LAN or WAN network, we are in the presence of a multicluster [13].

The hardware and operating system of the participating machines may be different; each machine may even be a “multiprocessor”, as is the case in multicore architectures that are so relevant nowadays. Multicore processors include several processing elements within an integrated circuit. This type of architectures are considered as a solution to the limitation of one core machines to increase computing power due to the increase in temperature [13][14][15].

Figure 1 shows the typical design of a current multicore architecture, composed by two processors that share the main memory. Each of these processors in turn is formed by four cores that share one L₃ cache memory (this memory may not be present in some models). There is also an L₂ cache memory that is shared by pairs of cores. Finally, each core has its own L₁ cache memory.

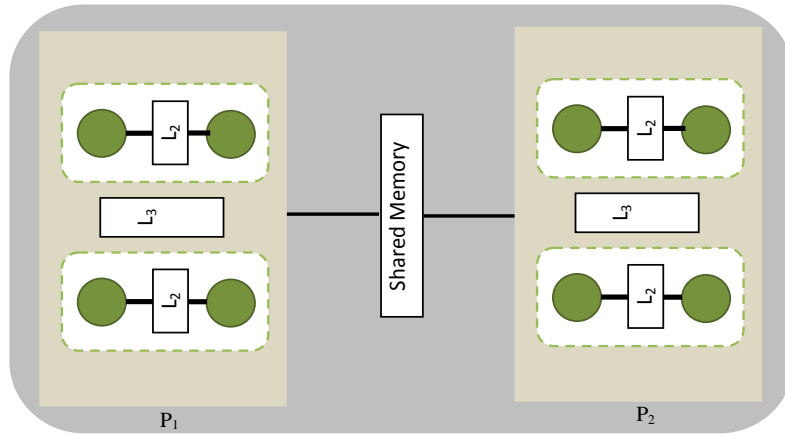


Fig. 1. Diagram of a multicore structure.

As it can be seen from Fig. 1, the communication between the various cores is done through the various memories of the architecture. Thus, the communication time between two cores is given by the time required to access the corresponding memory. In the case of Fig. 1, there are three levels of shared memory with their corresponding communication times. It is possible to build clusters using multicores; Fig. 2 shows a diagram of this type of architecture.

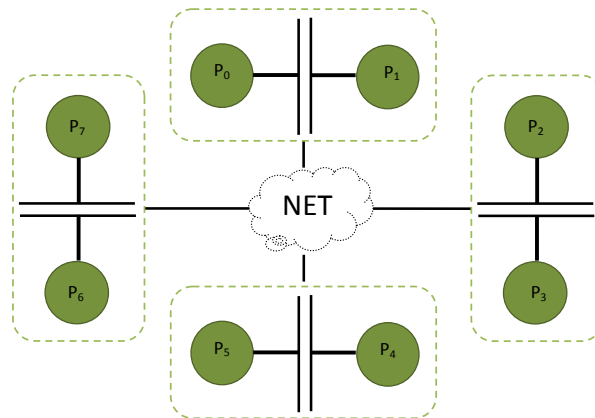


Fig. 2. Diagram of a multicore cluster.

The advent of this type of distributed architectures that can be accessed by any user has promoted the growth of parallel processing as a technique to increase architecture

exploitation. Application programmers should implement this technique by describing its component tasks and the interaction between them. Once the application has been developed, the programmer or user of the application on a parallel architecture will have to decide how to do it. That is, they should select how many of the processors/cores in the architecture will be used and how the application tasks will be assigned to them, in order to achieve the best possible throughput of the architecture with the lowest response time. This problem of solving the distribution of tasks between processors is called scheduling.

The problem of the automated scheduling of tasks to processing elements (processors in conventional machines and cores in multicore computers) is highly complex [1]. This complexity can be briefly represented considering the two main elements relating the parallel application to the supporting architecture: each node's processing capacity and the cost in time of communicating two processing elements [2].

The goal of modeling processing architectures is to obtain an "abstract" or simplified version of the physical machine, capturing crucial characteristics and disregarding minor details of the implementation [3].

A model does not necessarily represent a given real computer, but allows studying classes of problems over classes of architectures represented by their essential components. In this way, a real application can be studied over the architecture model, allowing us to get a significant description of the algorithm, draw a detailed analysis of its execution, and even predict the performance [4].

In the case of parallel systems, the most currently used architectures – due to their cost/performance relation - are clusters and multiclusters of multicores; for this reason, it is really important to develop models fitting the characteristics of these platforms. Essential elements to be considered are the potential heterogeneity of processors/cores and the different communication resources (shared memory, buses, networks), which add complexity to the modeling [5][6].

At present, there are different graph-based models to characterize the behavior of parallel applications in distributed architectures [7]. Among these models, we can mention TIG (Task Interaction Graph), TPG (Task Precedence Graph), TTIG (Task Temporal Interaction Graph) [8], TTIGHA (Task Temporal Interaction Graph on Heterogeneous Architectures) [9] and MPAHA (Model on Parallel Algorithms on Heterogeneous Architectures) [10].

Once the graph modeling the application has been defined, the *scheduling* problem is solved by an algorithm that establishes an automatic mechanism to carry out the assignment of tasks to processing elements, searching for the optimization of some running parameter (usually, time) [11][12]. Among the known mapping/scheduling algorithms, we consider AMTHA (Automatic Mapping Task on Heterogeneous Architectures), a mapping algorithm to carry out the assignment of tasks, making up the application to the processors of the architecture [10]. This algorithm considers the heterogeneous characteristics of the architecture taken into account in MPAHA (Model on Parallel Algorithms on Heterogeneous Architectures) model [10].

The AMTHA algorithm was developed to carry out the scheduling of applications executed over cluster and multicluster architectures using conventional machines. The

focus of this paper is analyzing the operation and adaptability of the AMTHA algorithm to multicore cluster architectures.

In Section 2, the scheduling algorithm AMTHA and the MPAHA model are described. Section 3 deals with the possible use of AMTHA and MPAHA for multicore clusters. In Section 4, the experimental work carried out with a multicore machine is presented, and the results obtained are detailed in Section 5. Finally, Section 6 presents the conclusions and the future lines of work.

2 AMTHA mapping algorithm

AMTHA is a static mapping algorithm that is applied to the graph generated by the MPAHA model. It allows determining the assignment of tasks to the processors of the architecture to be used, minimizing the execution time of the application. This algorithm must also provide the order in which subtasks (forming the task) assigned to each processor should be executed (task scheduling) [10].

The MPAHA model is based on the construction of a directed graph $G(V,E)$, where:

- V is the set of nodes representing each of the tasks T_i of the parallel program. Each node represents a task T_i of the parallel program, including its subtasks (S_{t_j}) and the order in which they should be executed to perform the task. If there is a heterogeneous architecture, the computation times for each processor should be taken into account ($V_i(s,p)$ = execution time of subtask s in processor type p).
- E is the set of edges representing each of the communications between the nodes of the graph. An edge A between two tasks T_i and T_j contains the communication volume (in bytes), the source subtask ($\in T_i$) and a target subtask ($\in T_j$).

It should be noted that, given the heterogeneity of the interconnecting network, instead of representing the time required for the communication, the corresponding communication volume between two subtasks is represented.

AMTHA considers an architecture with a limited number of heterogeneous processors. As regards the interconnecting network, the algorithm also considers that bandwidth and transmission speed can be heterogeneous.

The AMTHA algorithm uses the values of graph G generated by the MPAHA model; these values are the time required to compute a subtask in each type of processor, the communication volume with adjacent processors, and the task to which each subtask belongs.

Calculate **rank** for each task.
Whereas (not all tasks have been assigned)

1. Select the next task t to assign.
2. Chose the processor p to which task t should be assigned.
3. Assign task t (selected in step 1) to processor p (selected in step 2).
4. Update the **rank** of the tasks involved in step 3.

Fig. 3. Pseudo-code with the basic steps of the AMTHA algorithm.

The AMTHA algorithm assigns one task at a time until all tasks have been assigned. Figure 3 shows the pseudo-code with the main steps of the algorithm.

When the execution of the AMTHA algorithm ends, all the tasks have been assigned to one of the processors and the order in which the subtasks forming the tasks assigned to these processors will be executed has also been determined.

The following paragraphs describe each of the three steps followed during the execution of the AMTHA algorithm.

2.1 Calculating the rank of a task

Given a graph G , the rank of a task $Rk(T)$ is defined as the sum of the average times of the subtasks forming it and that are ready for execution (all predecessors have already been assigned to a processor and are already there). Equation 1 expresses this definition:

$$Rk(T) = \sum_{i \in L(T)} W_{avg}(St_i) \quad (1)$$

where:

$L(T)$ is the set of subtasks that are ready for task T .

$W_{avg}(St)$ is the average time of subtask St . The average time is calculated as shown in Equation 2.

$$W_{avg}(St_i) = \frac{\sum_{p \in P} V_{St_i}(\text{type of processor } p)}{\#P} \quad (2)$$

where P is the set of processors present in the architecture and $\#P$ is the number of processors forming this set.

2.2 Selecting the task to execute

After obtaining the *rank* of each application task, the task that maximizes it is selected. If there are two or more tasks that have the same maximum value, the algorithm breaks this tie by selecting the one that minimizes the total execution time average for the task. Equation 3 shows this calculation:

$$Tavg(T) = \sum_{i \in T} W_{avg}(St_i) \quad (3)$$

2.3 Selecting the processor

Selecting the processor involves choosing the computer within the architecture that minimizes the execution time when the selected task is assigned to that processor.

In order to understand how the time corresponding to processor p is calculated, the fact that each processor keeps a list of subtasks LU_p that were already assigned to it and that can be executed (all its predecessors are already placed), and another list that contains those subtasks that were assigned to p but whose execution is still pending LNU_p (some of their predecessors have not been placed yet) must be taken into account.

Therefore, to calculate which processor p will be selected, two possible situations are considered:

1. All subtasks of task t can be placed in p (that is, all its predecessors have been placed).
2. Some of the subtasks of t cannot be placed in p (this happens when some predecessor of this subtask has not been placed).

In the first case, the time T_p corresponding to processor p is given by the moment in which p finishes the execution of the last subtask of t . However, in the second case, the time T_p corresponding to processor p is given by the time when the last subtask of LU_p will finish plus the addition of all execution times in p for each of the subtasks on LNU_p .

2.4 Assigning the task to the selected processor

When assigning a task t to a processor p , there is an attempt to place each subtask St_k belonging to t to the processor at a moment in time when all the adjacent subtasks have already finished (including the predecessor subtask within t , if there is one) and its communications have been completed. The assignment can be a free interval between two subtasks that have already been placed in p , or an interval after them. If subtask St_k cannot be placed, it is added to the LNU_p list. Each time a subtask St_k is added to the LU list of one of the processors, an attempt is made to place all the predecessors belonging to already assigned tasks.

2.5 Updating the rank value of the pending tasks.

The first action within this step consists in assigning -1 as *rank* value to the task t that was assigned to processor p . The reason for this is to prevent task t from being re-selected for assignment.

Also, the following situation is considered in this step: for each subtask St_k placed in step 3.4, the need to update the *rank* of the tasks to which successor subtasks St_{succ} of St_k belong is analyzed; that is, if all predecessors of St_{succ} are already placed, then the *rank* of the task to which St_{succ} belongs is updated by increasing it by $W_{avg}(St_{succ})$.

3 MPAHA and AMTHA in multicore clusters

3.1 MPAHA model

The MPAHA model described in Section 2 does not require any modification to be used with multicore processors or multicore clusters. The directed graph $G(V,E)$ representing the tasks T_i and the communication among them do not change, if the parallel program is the same; regardless of the physical architecture.

3.2 AMTHA algorithm

When the AMTHA algorithm is run over a multicore cluster architecture, the following issues should be considered:

- The tasks that are part of the applications to execute will now be placed in some of the cores in the architecture; these cores are grouped in the various processors.
- The heterogeneity of the architecture as regards communications is not only given by the existence of different interconnecting networks within the architecture, but also by the different memory levels (main or cache) shared by the cores within each multicore machine. That is, two cores of the global architecture may communicate through different levels of shared memory, or by means of messages sent through an interconnecting network, which can be seen in Fig 1.
- When the algorithm assigns a task, it must consider the communication costs with its predecessor tasks. To this end, data related to the communication types that occur through the interconnecting network used when working with conventional clusters are required, as well as additional information regarding average access times for each of the memory levels in the multicores, together with information about core distribution in the machine.

4 Experimental work

In order to analyze the applicability of the AMTHA algorithm over multicore architectures, a set of synthetic applications with various characteristics was generated (as indicated in Section 4.1). For each of these, task assignment to the different cores in the architecture using the AMTHA algorithm was determined, and the execution time of using such distribution was estimated (T_{est}).

Based on the distribution done with AMTHA, the application was executed over the architecture described in Section 4.2 in order to obtain the real execution time (T_{exec}).

Both times (T_{est} and T_{exec}) were compared to determine how well the AMTHA algorithm estimates the execution time.

4.1 Choosing the set of applications to evaluate the AMTHA algorithm

A set of applications was selected, in which each of them varied in terms of: number of tasks (15-25), task size (5-50 seconds), number of subtasks making up a task (3-6), communication volume among subtasks, and communication probability between two subtasks (5-35 %). In all the applications, the total computing time exceeds that of communications (coarse grained application).

4.2 Choosing the architecture for the tests

The most widely used multicore architecture is Dell Poweredge 1950. It has the following characteristics: 2 quad core, 2.33 GHz Intel Ceon e5410 processors; 4 Gb

of RAM (shared between both processors); 6 MB L2 cache for each pair of processor cores.

5 Results

To analyze the results of the tests carried out, the difference between the execution times over the real architecture (T_{exec}) and the estimated execution times obtained when assigning tasks with the AMTHA algorithm (T_{est}) is calculated.

In Fig. 4, the relative percentage ($\%Dif_{rel}$) of this difference in T_{exec} is shown. To calculate this value, Equation 4 is used.

$$\%Dif_{rel} = \frac{T_{exec} - T_{est}}{T_{exec}} * 100 \quad (4)$$

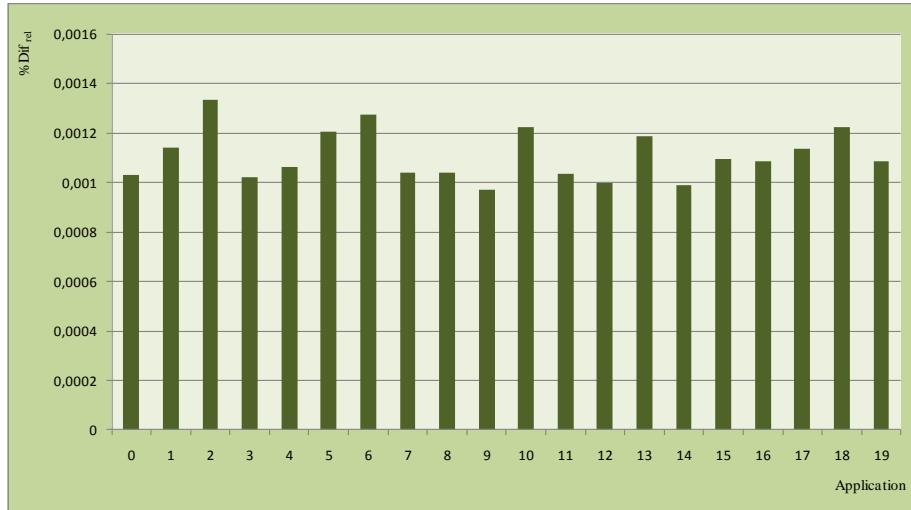


Fig. 4. $\%Dif_{rel}$ of the different tests carried out.

As the volume of communications (or the size of the transmitted packages) between tasks increases, so does the error as a function of the available cache in each core. However, in the tests carried out this value was never above 4% (always in applications with a much greater processing load than communication load).

6 Conclusions and future lines of work

As shown in Fig. 4, the $\%Dif_{rel}$ parameter did not increase beyond 0.0014 in the tests carried out. This confirms that the AMTHA algorithm is capable of successfully estimating the execution time of the application over a multicore architecture.

As regards related future lines of work, there are three areas to explore. First, tests should be extended to a multicore cluster (a Blade HP BL260c G5 with 64 cores in 8 blades with 2 INTEL E5405 processors with a quad core configuration and 2Gb of RAM will be incorporated soon). A second line of interest is the representation of communication time in multicore architectures when the messages sent exceed the capacity of the shared memories (main and cache). Finally, the necessary modifications to the model (MPAHA) and the scheduling algorithm (AMTHA) should be determined to consider hybrid programming models (integrating message passing and shared memory).

Reference

1. Grama A., Gupta A., Karypis G., Kumar V., "An Introduction to Parallel Computing. Design and Analysis of Algorithms. 2nd Edition". Pearson Addison Wesley (2003).
2. Kalinov A., Klimov S., "Optimal Mapping of a Parallel Application Processes onto Heterogeneous Platform". Proceeding of 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05), pp 123. IEEE CS Press (2005).
3. Leopold C., "Parallel and Distributed Computing. A survey of Models, Paradigms, and Approaches". Wiley, New York (2001).
4. Attiya H., Welch J., "Distributed Computing: Fundamentals, Simulations, and Advanced Topics. 2nd Edition". Wiley-IEEE, New Jersey (2004).
5. Topcuoglu H., Hariri S., Wu M., "Performance-Effective and Low-Complexity Task Scheduling for Heterogeneous Computing". IEEE Transactions on Parallel and Distributed Systems, vol. 13, pp. 260-274 (2002)
6. Goldman, "Scalable Algorithms for Complete Exchange on Multi-Cluster Networks". CCGRID'02, pp. 286 – 287. IEEE/ACM, Berlin (2002).
7. Roig C., Ripoll A., Senar M.A., Guirado F., Luque E., "Modelling Message-Passing Programas for Static Mapping". Euromicro Workshop on Parallel and Distributed Processing (PDP'00), pp. 229--236, IEEE CS Press, USA (1999).
8. Roig C., Ripoll A., Senar M., Guirado F., Luque E., "Exploiting knowledge of temporal behavior in parallel programs for improving distributed mapping". EuroPar 2000. LNCS, vol. 1900, pp. 262--71. Springer, Heidelberg (2000).
9. De Giusti L., Chichizola F., Naiouf M., De Giusti A. "Mapping Tasks to Processors in Heterogeneous Multiprocessor Architectures: The MATEHa Algorithm". Accepted for publication in the proceedings of the International Conference of the Chilean Computer Science Society (SCCC 2008). IEEE CS Press.
10. De Giusti L. "Mapping sobre Arquitecturas Heterogéneas". PhD Dissertation, Universidad Nacional de La Plata (2008).
11. Cuenca J., Gimenez D., Martinez J., "Heuristics for Work Distribution of a Homogeneous Parallel Dynamic Programming Scheme on Heterogeneous Systems". Proceeding of the 3rd International Workshop on Algorithms, Models

- and Tools for Parallel Computing on Heterogeneous Networks (HeteroPar'04), pp 354-361. IEEE CS Press (2004).
12. Cunha J.C., Kacsuk P., Winter S., "Parallel Program development for cluster computing: methodology, tools and integrated environments". Nova Science Pub., New York (2001).
 13. Thomas W. Burger. Intel Multi-Core Processors:Quick Reference Guide. http://cachewww.intel.com/cd/00/00/23/19/231912_231912.pdf
 14. Mc. Cool Michael "Programming models for scalable multicore programming" <http://www.hpcwire.com/features/17902939.html> (2007)
 15. Lei Chai, Qi Gao, Dhabaleswar K. Panda. "Understanding the Impact of Multi-Core Architecture in Cluster Computing: A Case Study with Intel Dual-Core System". IEEE International Symposium on Cluster Computing and the Grid 2007 (CCGRID 2007), pp. 471-478 (May 2007).
 16. Suresh Siddha, Venkatesh Pallipadi, Asit Mallick. "Process Scheduling Challenges in the Era of Multicore Processors"Intel Technology Journal, Vol. 11, Issue 04, November 2007.