



HAL
open science

The Master-Slave Paradigm with Heterogeneous Processors

Olivier Beaumont, Arnaud Legrand, Yves Robert

► **To cite this version:**

Olivier Beaumont, Arnaud Legrand, Yves Robert. The Master-Slave Paradigm with Heterogeneous Processors. [Research Report] RR-4156, LIP RR-2001-13, INRIA, LIP. 2001. inria-00072467

HAL Id: inria-00072467

<https://inria.hal.science/inria-00072467>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

*The Master-Slave Paradigm with Heterogeneous
Processors*

Olivier Beaumont,
Arnaud Legrand and

Yves Robert

No 4156

April 2001

———— THÈME 1 ————

 *Rapport
de recherche*

The Master-Slave Paradigm with Heterogeneous Processors

Olivier Beaumont,
Arnaud Legrand and
Yves Robert

Thème 1 — Réseaux et systèmes
Projet ReMaP

Rapport de recherche n° 4156 — April 2001 — 21 pages

Abstract: In this paper, we revisit the master-slave tasking paradigm in the context of heterogeneous processors. We assume that communications take place in exclusive mode. We present a polynomial algorithm that gives the optimal solution when a single communication is needed before the execution of the tasks on the slave processors. When communications are required both before and after the task processing, we show that the problem is at least as difficult as a problem whose complexity is open. In this case, we present a guaranteed approximation algorithm. Finally, we present asymptotically optimal algorithms when communications are required before the processing of each task, or both before and after the processing of each task.

Key-words: heterogeneous processors, master-slave tasking, communication, matching, complexity.

(Résumé : tsvp)

This text is also available as a research report of the Laboratoire de l'Informatique du Parallélisme
<http://www.ens-lyon.fr/LIP>.

Paradigme Maître-Esclaves sur un ensemble de processeurs hétérogènes

Résumé : Dans ce rapport, nous nous intéressons au paradigme maître/esclaves pour des plateformes hétérogènes. Nous supposons que les communications ont lieu de façon exclusives. Nous donnons un algorithme polynomial qui donne une solution optimale au problème de l'allocation des tâches lorsqu'une seule communication est nécessaire avant le traitement des tâches sur les différents processeurs. Lorsqu'une communication avant et une communication après le traitement des tâches sont nécessaires, nous montrons que le problème est aussi difficile qu'un autre problème dont la complexité est ouverte. Dans ce cas, nous présentons un algorithme d'approximation polynomial garanti. Enfin, nous présentons des algorithmes asymptotiquement optimaux quand des communications sont nécessaires avant ou avant et après le traitement de chaque tâche.

Mots-clé : Ressources hétérogènes, maître/esclaves, communications, couplages, complexité.

1 Introduction

Master-slave tasking is a simple yet widely used technique to execute independent tasks under the centralized supervision of a control processor. In the standard implementation of master-slave, the tasks are executed by identical processors (the slaves). We revisit the master-slave paradigm in the framework of heterogeneous computing resources: slave processors have different computation speeds. We present several scenarios to model the communication pattern between the master and the slaves. In all cases, such communications will take place in exclusive mode on a dedicated hardware resource (such as a serial bus).

To give a single motivation, this framework applies to any Monte Carlo simulation where large numbers of identical and independent simulations are run for different values of the random number generator seed. Monte Carlo simulations are widely used in various areas such as cellular microphysiology [13], reactor simulations [14] or studies on conformations of proteins [12].

The rest of the paper is organized as follows. In Section 2 we state four different variants of the master-slave problem: (i) with communications only before the dispatching of the tasks, (ii) with communications both before and after the processing of the tasks, (iii) with communication before each task processing and (iv) with communication both before and after each task processing. We give in Section 3 a polynomial time algorithm that solves the first problem. The second problem seems intrinsically more difficult and we prove in Section 4 that it is at least as difficult as a problem whose complexity is open; we also prove a guaranteed approximation algorithm for the second problem in Section 4, and we report some simulation results. We present asymptotically optimal algorithms when communications are required before each task (third problem) in Section 5, and when communications are required both before and after each task (fourth problem) in Section 6. We briefly survey related work in Section 7. Finally, we give some remarks and conclusions in Section 8.

2 Problem statement

The target architectural platform is represented in Figure 1. The master M and the p slaves P_1, P_2, \dots, P_p communicate through a shared medium, typically a bus, that can be accessed only in exclusive mode. At a given time-step, at most one processor can communicate with the master, either to receive data from the master or to send results back to the master.




Figure 1: The target master-slave architecture.

We assume that there is a pool of independent tasks to be processed by the p slaves. All tasks are of same-size, i.e. they represent the same amount of processing. Tasks are considered to be atomic (execution cannot be preempted once initiated). Processors are heterogeneous; more precisely, slave P_i requires t_i units of time to process a single task. We say that t_i is the “cycle-time” of processor P_i . Each P_i will execute c_i tasks (where c_i is to be determined) from the pool. Regardless of the hypotheses concerning communication costs, there are two (related) optimization problems:

MinTime(C) Given a total number of tasks C , determine the best allocation of tasks to slaves, i.e. the allocation $\mathcal{C} = \{c_1, c_2, \dots, c_p\}$ s.t. $\sum_{i=1}^p c_i = C$ and which minimizes the total execution time.

MaxTasks(T) Given a time bound T , determine the best allocation of tasks to slaves, i.e. the allocation $\mathcal{C} = \{c_1, c_2, \dots, c_p\}$ s.t. all processors complete their execution within T units of time and $\sum_{i=1}^p c_i = C$ is maximized.

In the paper, we concentrate on solving the second problem **MaxTasks(T)**. Given the solution to this problem, we find a solution to **MinTime(C)** by using binary search on T , calling **MaxTasks(T)** several times, and returning the smallest value of T for which the answer is at least C .

We now state some specific hypotheses for the communication costs. For each modeling of these communication costs, we analytically formulate the **MaxTasks(T)** problem.

2.1 Without any communication cost

Assume first that there is no communication cost at all. It is not difficult to solve both previous problems using a greedy algorithm. The solution of problem **MaxTasks(T)** is straightforward: we let $c_i = \lfloor \frac{T}{t_i} \rfloor$ for all i , $1 \leq i \leq p$, which obviously defines the optimal solution.

2.2 With an initial scattering of data

The formulation of this problem is taken from Andonie et al. [1], who study the implementation of distributed backpropagation neural networks on heterogeneous networks of workstations, using the PVM library [6]. The training of the neural network is divided into computational phases. At each phase, the training pattern is distributed among the slaves, which are different-speed processors. Before executing any task, each slave must receive some data file from the master processor. Because the communication medium is exclusive, it is not relevant to distinguish whether the data file is the same for all slaves (then the master executes a broadcast operation) or whether it is different (then the master executes a scatter operation): we only assume that each slave must receive the same amount of data, and that each reception costs t_{com} units of time. In the model of Andonie et al. [1], there is no communication cost paid to send the results back to the master. In general, when the slaves compute “yes/no” results, the cost of returning the results may well be neglected in front of the cost of the initial scatter and/or of the computations. Note that we deal with another model, including communication costs both before and after the tasks, in Section 2.3.

Due to the constraint on the communication medium, the p messages will be sent one after the other. Obviously, it cannot hurt to send the messages as soon as possible, i.e. at time steps $0, t_{\text{com}}, 2t_{\text{com}}, \dots, (p-1)t_{\text{com}}$. The problem is then to determine the ordering of the p messages, i.e. the permutation σ of $\{1, 2, \dots, p\}$ such that slave P_i receives the message at time $\sigma(i)t_{\text{com}}$. We are ready to state the optimization problem analytically:

MaxTasks1(T) Given a time bound T , determine the best allocation of tasks to slaves, i.e. a permutation σ and an allocation $\mathcal{C} = \{c_1, c_2, \dots, c_p\}$ s.t. all processors complete their execution within T units of time and the total number of tasks is maximized:

$$\max \left(\sum_{i=1}^p c_i \mid \sigma \text{ permutation and } \forall i \in [1, p] : \sigma(i)t_{\text{com}} + c_i t_i \leq T \right)$$

2.3 With initial and final communications

As pointed out above, it is natural to assume that after the processing of their tasks, slave processors will send some data back to the master. Because this message may well have a different size than the message initially sent by the master, we model this situation by using two communication costs, t_{com}^1 for the messages sent by the master to the slaves, and t_{com}^2 for the messages sent by the slaves to the master.

As above, we look for a permutation σ_1 which determines the ordering of the initial messages from the host: the host sends data to slave P_i at time $\sigma_1(i)t_{\text{com}}^1$. But we also look for a second permutation σ_2 which determines the ordering of the final messages sent back to the host: given a time bound T , slave P_i sends data back to the host at time $T - \sigma_2(i)t_{\text{com}}^2$. This formulation is without any loss of generality: some slave processor P_i might send its message earlier than this bound, but we can always shift the communication




Figure 2: Delaying messages sent back to the host.

pattern as stated, i.e. delay some messages, as illustrated Figure 2. We are ready to state the optimization problem analytically:

MaxTasks2(T) Given a time bound T , determine the best allocation of tasks to slaves, i.e. two permutations σ_1 and σ_2 , and an allocation $\mathcal{C} = \{c_1, c_2, \dots, c_p\}$ s.t. all processors complete their execution within T units of time and the total number of tasks is maximized:

$$\max \left(\sum_{i=1}^p c_i \mid \sigma_1, \sigma_2 \text{ permutations and } \forall i \in [1, p] : \sigma_1(i)t_{\text{com}}^1 + c_i t_i + \sigma_2(i)t_{\text{com}}^2 \leq T \right)$$

2.4 With communications before each task processing

We also consider the case when communications are required between the master and the slave before the processing of each task. In this third model, we consider that the cost of such a communication is t_{com} . This model is quite natural: some specific input data may well have to be propagated from the master to the slave before computation can start.

We look for three functions $f_{\text{startcomm}}$, $f_{\text{startcomp}}$ and f_{proc} : $f_{\text{startcomm}}(i)$ represents the time-step at which the communication required by task i will begin; $f_{\text{startcomp}}(i)$ represents the time-step at which the processing of task i will begin on processor $f_{\text{proc}}(i)$. The functions $f_{\text{startcomm}}$, $f_{\text{startcomp}}$ and f_{proc} must fulfill the following conditions:

- $\forall i \geq 1, f_{\text{startcomm}}(i+1) - f_{\text{startcomm}}(i) \geq t_{\text{com}}$, which states that communications take place in exclusive mode.
- $\forall i \geq 1, f_{\text{startcomp}}(i) \geq f_{\text{startcomm}}(i) + t_{\text{com}}$, which states that the processing of task i cannot start before the end of the communication required by task i .
- $\forall 1 \leq i < j, \text{ if } f_{\text{proc}}(i) = f_{\text{proc}}(j) = k, \text{ then } f_{\text{startcomp}}(j) \geq f_{\text{startcomp}}(i) + t_k$, which states that tasks are processed sequentially on each processor k .
- $\forall 1 \leq i < j, \text{ if } f_{\text{proc}}(i) = f_{\text{proc}}(j) = k, \text{ then}$

$$[f_{\text{startcomm}}(j), f_{\text{startcomm}}(j) + t_{\text{com}}] \cap [f_{\text{startcomp}}(i), f_{\text{startcomp}}(i) + t_k] = \emptyset,$$

which states that communications and computations cannot be overlapped on processor k .




Figure 3: Grouping some messages sent by the host to a given processor.

This formulation is quite general. Note that each processor can perform several communications before processing the corresponding tasks, as illustrated on Figure 3. We are ready to state the optimization problem analytically:

MaxTasks3(T) Given a time bound T , determine the best allocation of tasks to slaves, i.e. three functions $f_{\text{startcomm}}$, $f_{\text{startcomp}}$ and f_{proc} satisfying all the conditions stated above, s.t. all processors complete their execution within T units of time and the total number of tasks is maximized:

$$\max (N \mid \forall i \leq N, f_{\text{startcomp}}(i) + t_{f_{\text{proc}}}(i) \leq T)$$

2.5 With communications both before and after each task processing

It is natural to assume that after the processing of each task, slave processors will send some data back to the master. As previously, we model this situation by using two different communication costs, t_{com}^1 for the messages sent by the master to the slaves, and t_{com}^2 for the messages sent by the slaves to the master.

We look for four functions $f_{\text{startcomm}}^1$, $f_{\text{startcomm}}^2$, $f_{\text{startcomp}}$ and f_{proc} : $f_{\text{startcomm}}^1(i)$ represents the time-step at which the communication from the host required before task i will begin (just as $f_{\text{startcomm}}(i)$ in the previous section); similarly, $f_{\text{startcomm}}^2(i)$ represents the time-step at which the communication back to the host after task i will begin; finally, $f_{\text{startcomp}}$ and $f_{\text{proc}}(i)$ are defined as before: $f_{\text{startcomp}}(i)$ represents the time-step at which the processing of task i will begin on processor $f_{\text{proc}}(i)$. The functions $f_{\text{startcomm}}^1$, $f_{\text{startcomm}}^2$, $f_{\text{startcomp}}$ and f_{proc} have to fulfill the following conditions:

- $\forall i \geq 1, \forall j \geq 1, \forall (k, l) \in \{1, 2\}$, if $k \neq l$ or $i \neq j$ then

$$[f_{\text{startcomm}}^k(i), f_{\text{startcomm}}^k(i) + t_{\text{com}}^k] \cap [f_{\text{startcomm}}^l(j), f_{\text{startcomm}}^l(j) + t_{\text{com}}^l] = \emptyset,$$

which states that communications take place in exclusive mode.

- $\forall i \geq 1, f_{\text{startcomp}}(i) \geq f_{\text{startcomm}}(i) + t_{\text{com}}^1$, which states that the processing of task i cannot start before the end of the communication from the host required by task i .
- $\forall i \geq 1, f_{\text{startcomp}}(i) + t_{f_{\text{proc}}}(i) \leq f_{\text{startcomm}}^2(i)$, which states that the communication back to the host required after task i cannot start before the end of the processing of task i .

- $\forall 1 \leq i < j$, if $f_{\text{proc}}(i) = f_{\text{proc}}(j) = k$,
 $f_{\text{startcomp}}(j) \geq f_{\text{startcomp}}(i) + t_k$, which states that tasks are processed sequentially on each processor k .
- $\forall 1 \leq i < j$, if $f_{\text{proc}}(i) = f_{\text{proc}}(j) = k$, then

$$[f_{\text{startcomm}}^1(j), f_{\text{startcomm}}^1(j) + t_{\text{com}}] \cap [f_{\text{startcomp}}(i), f_{\text{startcomp}}(i) + t_k] = \emptyset$$

and

$$[f_{\text{startcomm}}^2(i), f_{\text{startcomm}}^2(i) + t_{\text{com}}] \cap [f_{\text{startcomp}}(j), f_{\text{startcomp}}(j) + t_k] = \emptyset,$$

which states that communications and computations cannot be overlapped on processor k .




Figure 4: Grouping some messages sent by and to the host.

Again, this formulation is quite general. Note that each processor can perform several communications from the host before processing the corresponding tasks, as well as delaying several communications back to the host: this is illustrated on Figure 4. We are ready to state the optimization problem analytically:

MaxTasks4(T) Given a time bound T , determine the best allocation of tasks to slaves, i.e. four functions $f_{\text{startcomm}}^1$, $f_{\text{startcomm}}^2$, $f_{\text{startcomp}}$ and f_{proc} satisfying all the conditions stated above, s.t. all processors complete their execution within T units of time and the total number of tasks is maximized:

$$\max(N \mid \forall i \leq N, f_{\text{startcomm}}^2(i) + t_{\text{com}}^2 \leq T)$$

3 Solution with an initial scattering of data

3.1 Restricted search

To (partially) solve the **MaxTasks1(T)** problem of Section 2.2, Andonie et al. [1] restrict the search to allocations where the fastest processors start computing first. They use a dynamic programming algorithm to solve the optimization problem **MinTime(C)**. With our setting for problem **MaxTasks1(T)**, this amounts to sort the cycle-times as $t_1 \leq t_2 \leq \dots \leq t_p$ and to let $\sigma(i) = i$ for $1 \leq i \leq p$. The intuition is that fastest processors execute tasks more rapidly than the others, hence they should work longer.

However, the intuition is misleading in some cases. Assume for instance two slave processors ($p = 2$) with $t_1 = 5$ and $t_2 = 9$ and let $t_{\text{com}} = 1$. For the time bound $T = 28$, it is better to start the slow processor P_2

first: P_2 can then execute three tasks: $t_{\text{com}} + 3t_2 = 28 \leq T$; the fast processor, although started at time-step $2t_{\text{com}} = 2$, can execute five tasks: $2t_{\text{com}} + 5t_1 = 27 \leq T$. If we start the fast processor first, it cannot execute more than 5 tasks, while the second processor can execute only 2.




Figure 5: Bipartite graph for MaxTasks1(T).

3.2 Matching techniques

The optimal solution to the MaxTasks1(T) problem can be found using a weighted-matching algorithm. The idea is to draw a complete bipartite graph with $2p$ vertices, as shown in Figure 5. Vertices on the left represent processors, while vertices on the right represent possible values for the permutation σ . The edge from vertex P_i to vertex S_j is weighted with the maximum number of tasks that P_i can execute if $\sigma(i) = j$, namely $\lfloor \frac{T - jt_{\text{com}}}{t_i} \rfloor$. Extracting a matching from the graph enables to assign a different value of σ for each processor, thereby guaranteeing that σ is indeed a permutation. In fact, there is a one-to-one correspondence between matchings and permutations. Because the total weight of a given matching is the total number of tasks that can be executed for the corresponding choice of the permutation, our problem reduces to finding the maximum weighted matching in the bipartite graph. Efficient (polynomial) algorithms exist to solve this problem, see [7, 15]. More precisely, the complexity of finding the maximum weighted matching in a bipartite graph with $2p$ vertices is of order $O(p^3)$.

We work out the following example: assume $p = 3$ processors of cycle-times $t_1 = 4$, $t_2 = 5$ and $t_3 = 9$. Let $t_{\text{com}} = 1$, and consider the time bound $T = 118$. The weighted bipartite graph is shown Figure 6. The maximum weighted matching is unique, it corresponds to the permutation $\sigma(1) = 2$, $\sigma(2) = 3$ and $\sigma(3) = 1$. The total number of tasks is $29 + 23 + 13 = 75$.

To conclude this section, we formally state our result:

Proposition 1 *The optimal solution to the MaxTasks1(T) problem with initial messages can be found in time of order $O(p^3)$ with p processors using the above weighted-matching algorithm.*

4 Solution with initial and final communications

The solution to the MaxTasks2(T) problem with initial and final messages turns out to be surprisingly difficult. In fact, we do not know of any polynomial algorithm for the general case. We present an efficient guaranteed approximation using matching techniques, as explained in Section 4.1. In Section 4.2, we give some remarks about the complexity of MaxTasks2(T).




Figure 6: Bipartite graph with $p = 3$, $t_1 = 4$, $t_2 = 5$, $t_3 = 9$, $t_{\text{com}} = 1$, and $T = 118$.

4.1 Matching techniques




Figure 7: Bipartite graph with initial and final communications.

To take both permutations σ_1 and σ_2 into account, we build a bipartite graph $G = (V, E)$ with $3p$ vertices (i.e. $|V| = 3p$), as shown Figure 7. The p leftmost vertices F_i correspond to the first permutation σ_1 , the p center vertices P_i correspond to processors, and the p rightmost vertices S_i correspond to the second permutation σ_2 . Rather than a matching, we extract a 2-factor from the graph [7, 15]: more precisely, we select a subset E' of $2p$ edges so that in the graph $G = (V, E')$ each vertex F_i or S_i is exactly of degree 1, and each vertex P_i is exactly of degree 2. The complexity of extracting 2-factor from the graph with $3p$ vertices is of order $O(p^3)$ again, since we can solve independently the maximum weighted matching in both bipartite graphs with $2p$ vertices (on the left-hand side and on the right-hand side in Figure 7) in time of order $O(p^3)$.

The problem is that edge weights cannot be determined fully accurately; the inequality $\sigma_1(i)t_{\text{com}}^1 + c_i t_i + \sigma_2(i)t_{\text{com}}^2 \leq T$ translates into

$$c_i \leq \left\lfloor \frac{T - \sigma_1(i)t_{\text{com}}^1 - \sigma_2(i)t_{\text{com}}^2}{t_i} \right\rfloor,$$

and we need to know *both* $\sigma_1(i)$ and $\sigma_2(i)$ to compute c_i . Instead, we use the approximation

$$\left\lfloor \frac{T/2 - \sigma_1(i)t_{\text{com}}^1}{t_i} \right\rfloor + \left\lfloor \frac{T/2 - \sigma_2(i)t_{\text{com}}^2}{t_i} \right\rfloor.$$

This approximation enables us to weight the edges as follows: the edge between F_j and P_i is weighted as $\left\lfloor \frac{T/2 - jt_{\text{com}}^1}{t_i} \right\rfloor$ while the edge between P_i and S_k is weighted as $\left\lfloor \frac{T/2 - kt_{\text{com}}^2}{t_i} \right\rfloor$.

Theorem 1 *The previous approximation leads to tasks allocations that differs at most by p from the optimal solution.*

Proof First note that

$$\forall a, b : 0 \leq \lfloor a + b \rfloor - \lfloor a \rfloor - \lfloor b \rfloor \leq 1. \quad (4.1)$$

Therefore for any allocation (σ_1, σ_2, c) built using the previous approximation, we have

$$\forall i, c_i \leq \left\lfloor \frac{T/2 - \sigma_1(i)t_{\text{com}}^1}{t_i} \right\rfloor + \left\lfloor \frac{T/2 - \sigma_2(i)t_{\text{com}}^2}{t_i} \right\rfloor \leq \left\lfloor \frac{T - \sigma_1(i)t_{\text{com}}^1 - \sigma_2(i)t_{\text{com}}^2}{t_i} \right\rfloor,$$

and (σ_1, σ_2, c) is then a solution of the initial problem.

Let us note by

$$C_{app}(\sigma_1, \sigma_2) = \sum_{i=1}^p \left\lfloor \frac{T/2 - \sigma_1(i)t_{\text{com}}^1}{t_i} \right\rfloor + \left\lfloor \frac{T/2 - \sigma_2(i)t_{\text{com}}^2}{t_i} \right\rfloor$$

the approximative cost of two permutations and

$$\sigma_{app} = (\sigma_1^{(app)}, \sigma_2^{(app)}) = \operatorname{argmax}_{\sigma_1, \sigma_2} C_{app}(\sigma_1, \sigma_2)$$

the permutations built by our algorithm.

Let

$$C_{opt}(\sigma_1, \sigma_2) = \sum_{i=1}^p \left\lfloor \frac{T - \sigma_1(i)t_{\text{com}}^1 - \sigma_2(i)t_{\text{com}}^2}{t_i} \right\rfloor$$

denote the real cost of two permutations and let

$$\sigma_{opt}(\sigma_1^{(opt)}, \sigma_2^{(opt)}) = \operatorname{argmax}_{\sigma_1, \sigma_2} C_{opt}(\sigma_1, \sigma_2)$$

be some permutations that are optimal solutions to the original problem.

By definition,

$$C_{opt}(\sigma_{app}) \leq C_{opt}(\sigma_{opt}) \quad (4.2)$$

and

$$C_{app}(\sigma_{opt}) \leq C_{app}(\sigma_{app}). \quad (4.3)$$

Using equation (4.1) we find that

$$\forall \sigma : C_{opt}(\sigma) - p \leq C_{app}(\sigma) \leq C_{opt}(\sigma) \quad (4.4)$$

and therefore

$$\begin{aligned} C_{opt}(\sigma_{opt}) - p &\leq C_{app}(\sigma_{opt}) && \text{(using (4.4))} \\ &\leq C_{app}(\sigma_{app}) && \text{(using (4.3))} \\ &\leq C_{opt}(\sigma_{app}) && \text{(using (4.4))} \\ C_{opt}(\sigma_{opt}) - p &\leq C_{opt}(\sigma_{app}) \leq C_{opt}(\sigma_{opt}) && \text{(using (4.2)),} \end{aligned}$$

which means that our approximation leads to tasks allocations that differs at most by p from the optimal solution. ■

4.2 Some remarks about the complexity of MaxTasks2(T)

We have not found any polynomial algorithm for the general case, and we have not been able to prove the NP-completeness of MaxTasks2(T). Nevertheless, we can formulate a few remarks about the intrinsic difficulty of MaxTasks2(T). First, an exhaustive search of all possible permutations would have a complexity of order $O((p!)^2)$, which is impossible in practice as soon as $p \geq 9$ (about 44 minutes on a Pentium III 550 MHz for $p = 9$ and more than one day for $p = 10$). Moreover, the problem seems to be difficult even for very simple instances of MaxTasks2(T), as shown below. Indeed, let us consider the following open (polynomial vs. NP-complete) problem in combinatorial optimization (see [9]):

Permutation Sums:

Instance: Let $a_1 \leq a_2 \leq \dots \leq a_p$ be p positive integers satisfying $\sum_{i=1}^p a_i = p(p+1)$.

Question: Do there exist two permutations σ_1 and σ_2 of $\{1, 2, \dots, n\}$ such that

$$\forall i \in [1, p] : \sigma_1(i) + \sigma_2(i) = a_i.$$

Let us build the following instance of MaxTasks2(T):

Instance: Let $T = 3 \max_i a_i$, $t_{\text{com}} = 1$. Let $t_i = T - a_i$, $\forall i \in [1, p]$ denote the cycle time of processor P_i .

Question: Is it possible to perform p tasks within T units of time?

Let us consider the above instance of MaxTasks2(T). Since $t_i = T - a_i$ and $T = 3 \max_i a_i$, $2t_i > T$ and it is therefore impossible to execute more than one task with a given processor within T units of time. Thus, our instance of MaxTasks2(T) is very simple (with respect to the general formulation) since $\forall i$, $c_i \leq 1$. Indeed, the question can be stated as follows: is it possible to find two permutations σ_1 and σ_2 such that $\forall i \in [1, p] : \sigma_1(i) + \sigma_2(i) + t_i \leq T$.

Nevertheless, this instance of MaxTasks2(T) is as difficult as Permutation Sums. More precisely, if Permutation Sums is proved to be NP-complete then MaxTasks2(T) is also NP-complete and if MaxTasks2(T) can be solved in polynomial time, then it proves that Permutation Sums can also be solved in polynomial time. Indeed, let us suppose that

$$\exists \sigma_1, \sigma_2 \text{ such that } \forall 1 \leq i \leq p, \sigma_1(i) + \sigma_2(i) + t_i \leq T.$$

Then, $\forall 1 \leq i \leq p$, $\sigma_1(i) + \sigma_2(i) \leq a_i$. Moreover, let us suppose that

$$\exists i, \sigma_1(i) + \sigma_2(i) < a_i,$$

then

$$p(p+1) = \sum_{i=1}^p (\sigma_1(i) + \sigma_2(i)) < \sum_{i=1}^p a_i = p(p+1),$$

which is absurd. Thus, finally

$$\forall i \in [1, p] : \sigma_1(i) + \sigma_2(i) = a_i.$$


Thus, we can expect that the general instance of MaxTasks2(T) is intrinsically difficult.

4.3 Simulations


In this section, we compare solutions given by the approximation algorithm proposed in Section 4 with the optimal ones. For a given instance of the problem, we compute the difference between number of tasks processed using the optimal permutations (found with an exhaustive search) and the number of tasks processed using the approximation algorithm proposed in Section 4. Figure 8 depicts the mean value of this error for several tests when the number of processors varies :

1. In Figure 8(a), cycle times are not strongly heterogeneous (numbers at random between 15 and 25) and communication time can be large compared to cycle times (numbers at random between 10 and 90).


2. In Figure 8(b), cycle times are still not strongly heterogeneous (numbers at random between 15 and 25) but communication time are rather small compared to cycle times (between 1 and 9).
3. In Figure 8(c), cycle times are strongly heterogeneous (numbers at random between 5 and 35) and communication time can be large compared to cycle times (two numbers at random between 10 and 90).
4. In Figure 8(d), cycle times are still strongly heterogeneous (numbers at random between 5 and 35) and communication time are rather small compared to cycle times (two numbers at random between 1 and 9).




(a) comp 20 ± 5 ; comm 50 ± 40



(b) comp 20 ± 5 ; comm 5 ± 4



(c) comp 20 ± 15 ; comm 50 ± 40



(d) comp 20 ± 15 ; comm 5 ± 4

Figure 8: A few simulations

In all cases, the matching gives very satisfying results, since the difference between the optimal number of tasks that may be processed and the number of tasks processed using the matching algorithm is small. As we could expect, this difference grows with p and doesn't depend on the degree of heterogeneity.

5 Solution with communications before each task

In this section, we present an asymptotically optimal algorithm for $\text{MaxTasks3}(T)$: when T becomes large, the ratio of the number of tasks processed by this algorithm over the number of tasks executed by the optimal solution tends to one.

5.1 Theoretical bounds

In order to prove the asymptotic optimality of our algorithm, we need to determine the optimal number of tasks that can be performed if the cost of a communication between the master and the slave is t_{com} and the cycle times of slaves processors are $t_1 \leq t_2 \leq \dots \leq t_p$. Consider a valid communication and computation scheme, i.e. three functions $f_{\text{startcomm}}$, $f_{\text{startcomp}}$ and f_{proc} satisfying the conditions given in Section 2.4. Let T be the time bound and let N denote the maximal number of tasks that can be processed within T units of time:

$$N = \max\{n, \forall i \leq n, f_{\text{startcomp}}(i) + T_{f_{\text{proc}}(i)} \leq T\}.$$

Moreover, let

$$\forall i, 1 \leq i \leq p, \alpha_i = \frac{\text{Card}\{j, f_{\text{proc}}(j) = i\} t_{\text{com}}}{T}$$

be the ratio of the time spent by the master to perform communications with slave P_i over the time bound T .

Lemma 1 *With the above notations, we have*

$$N \leq \sum_i \frac{T}{t_{\text{com}} + t_i} \quad (5.1)$$

$$N \leq \frac{T}{t_{\text{com}}} \quad (5.2)$$

Proof

1. Since $\forall 1 \leq i \leq p, \alpha_i = \frac{\text{Card}\{j, f_{\text{proc}}(j)=i\} t_{\text{com}}}{T}$, then $\frac{T\alpha_i}{t_{\text{com}}} = \text{Card}\{j, f_{\text{proc}}(j) = i\}$ and

$$\frac{T \sum_{i=1}^p \alpha_i}{t_{\text{com}}} = \sum_{i=1}^p \text{Card}\{j, f_{\text{proc}}(j) = i\} = N. \quad (5.3)$$

2. Let us determine the maximal number of tasks that can be performed by slave $P_i, \forall 1 \leq i \leq p$. Since the overall cost of a task on slave P_i is $t_{\text{com}} + t_i$, we have

$$\text{Card}\{j, f_{\text{proc}}(j) = i\} \leq \frac{T}{t_{\text{com}} + t_i} \text{ i.e. } \alpha_i \leq \frac{t_{\text{com}}}{t_{\text{com}} + t_i}. \quad (5.4)$$

3. Let us determine the number of communications that can be performed by the master. The cost of a communication is t_{com} so that $Nt_{\text{com}} \leq T$ and then, using equation 5.3, we deduce

$$\sum_{i=1}^p \alpha_i \leq 1. \quad (5.5)$$

4. Using equations 5.3 and 5.4, we have

$$N \leq \sum_i \frac{T}{t_{\text{com}} + t_i}$$

5. Using equations 5.3 and 5.5, we deduce

$$N \leq \frac{T}{t_{\text{com}}}$$

■

In order to determine the optimal number of tasks that can be performed during T time steps, we need to distinguish two different cases, according to the value of $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}} + t_i}$. Indeed, it turns out that the communication network is not the limiting resource if $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}} + t_i} \leq 1$, but it becomes the limiting resource otherwise.

5.2 Solution of MaxTasks3(T) if $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}}+t_i} \leq 1$

To solve MaxTasks3(T), we propose an algorithm that consists in determining a pattern for communications and computations, that will be reproduced periodically throughout the execution.

Let $tc_i = t_{\text{com}} + t_i$, for $1 \leq i \leq p$, denote the overall cost of the processing of a task on slave P_i , since we cannot overlap communications and computations. Let T^{pattern} be the least common multiple of these p values tc_i : T^{pattern} determines the length of the pattern. Let $nb_i^{\text{pattern}} = \frac{T^{\text{pattern}}}{tc_i}$ be the number of tasks processed by processor P_i during the execution of the pattern.

Consider the following example: $t_{\text{com}} = 1$ and $p = 4$ slave processors with $t_1 = 2$, $t_2 = 3$, $t_3 = 3$ and $t_4 = 5$. We have $tc_1 = 3$, $tc_2 = 4$, $tc_3 = 4$ and $tc_4 = 6$. In this case $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}}+t_i} = 1$, $T^{\text{pattern}} = 12$, and $nb_1^{\text{pattern}} = 4$, $nb_2^{\text{pattern}} = 3$, $nb_3^{\text{pattern}} = 3$ and $nb_4^{\text{pattern}} = 2$.

To formally build the pattern, we need some complicated notations. We advise the reader to follow the construction for the example, using Figure 9. First we define time-steps and processors within the pattern, using three new functions $f_{\text{startcomm}}^{\text{pattern}}$, $f_{\text{startcomp}}^{\text{pattern}}$ and $f_{\text{proc}}^{\text{pattern}}$ which we define as follows (initially $nb_0^{\text{pattern}} = 0$):

- Determine which processor executes task number i :

$$\forall 1 \leq i \leq \sum_{k=1}^p nb_k^{\text{pattern}} : f_{\text{proc}}^{\text{pattern}}(i) = \min \left\{ j \mid i > \sum_{k=0}^{j-1} nb_k^{\text{pattern}} \right\}$$

- Determine the beginning of the communication and of the computation for task number i :

$$f_{\text{startcomm}}^{\text{pattern}}(i) = 1 + \sum_{k=0}^{j-1} nb_k^{\text{pattern}}(t_{\text{com}} + t_k) + (i - 1 - \sum_{k=0}^{j-1} nb_k^{\text{pattern}})t_{\text{com}},$$

$$f_{\text{startcomp}}^{\text{pattern}}(i) = 1 + \sum_{k=0}^j nb_k^{\text{pattern}}t_{\text{com}} + \sum_{k=0}^{j-1} nb_k^{\text{pattern}}t_k + (i - 1 - \sum_{k=0}^{j-1} nb_k^{\text{pattern}})t_j.$$




Figure 9: Example when $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}}+t_i} \leq 1$.

The important fact about this pattern is that $\forall 2 \leq i \leq p$,

$$\begin{aligned} & \left(f_{\text{startcomp}}^{\text{pattern}} \left(\sum_{k=0}^i nb_k^{\text{pattern}} \right) + t_i \right) - \left(f_{\text{startcomp}}^{\text{pattern}} \left(\sum_{k=0}^{i-1} nb_k^{\text{pattern}} \right) + t_{i-1} \right) \\ &= \left(f_{\text{startcomm}}^{\text{pattern}} \left(\sum_{k=0}^{i-1} nb_k^{\text{pattern}} + 1 \right) \right) - \left(f_{\text{startcomm}}^{\text{pattern}} \left(\sum_{k=0}^{i-2} nb_k^{\text{pattern}} + 1 \right) \right). \end{aligned}$$

This condition states that the difference between the date of the end of the processing of the last task on slave P_i and the date of the end of the processing of the last task on slave P_{i-1} is equal to the difference between the beginning of the communication required by the first task processed by slave P_i and the beginning of the communication required by the first task processed by processor P_{i-1} . This is the key-condition that ensures the periodicity of the whole computation and communication scheme from one pattern to the next one: it is possible to execute the pattern defined above every T^{pattern} time steps, as illustrated Figure 9. During the execution of one pattern, we process $\sum_{i=1}^p nb_i^{\text{pattern}}$ tasks.

We are now able to define the functions $f_{\text{startcomm}}(n)$, $f_{\text{startcomp}}(n)$ and $f_{\text{proc}}(n)$ corresponding to our algorithm. Let k s.t. $k \sum_{i=1}^p nb_i^{\text{pattern}} < n \leq (k+1) \sum_{i=1}^p nb_i^{\text{pattern}}$ and let $i = n - k \sum_{i=1}^p nb_i^{\text{pattern}}$. Then

$$\begin{aligned} f_{\text{startcomm}}(n) &= kT^{\text{pattern}} + f_{\text{startcomm}}^{\text{pattern}}(i), \\ f_{\text{startcomp}}(n) &= kT^{\text{pattern}} + f_{\text{startcomp}}^{\text{pattern}}(i), \\ f_{\text{proc}}(n) &= f_{\text{proc}}^{\text{pattern}}(i). \end{aligned}$$

One can easily check that the functions $f_{\text{startcomm}}$, $f_{\text{startcomp}}$ and f_{proc} satisfy all the conditions stated in Section 2.4.

5.3 Solution of MaxTasks3(T) if $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}} + t_i} > 1$

To solve MaxTasks3(T) when $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}} + t_i} > 1$, we slightly modify the algorithm proposed in previous section. Indeed, in this case, the network is the limiting resource. The algorithm consists in determining a communication and computation pattern so that the communication network is always in use. Some slower processors will be kept idle at some periods, or even will never be used.

First of all we sort the cycle-times of the slave processors and assume that

$$t_1 \leq t_2 \leq \dots \leq t_p.$$

Let tc_i be defined as previously and let

$$p_{\text{max}} = \max \left\{ k \mid \sum_{i=1}^k \frac{t_{\text{com}}}{t_{\text{com}} + t_i} \leq 1 \right\}.$$

p_{max} is the index of the last processor whose computation power will be fully used in the pattern.

Let T^{pattern} be the least common multiple of t_{com} and of the tc_i , $1 \leq i \leq p_{\text{max}}$. Moreover, define nb_i^{pattern} as follows:

$$\begin{aligned} \forall 1 \leq i \leq p_{\text{max}}, \quad nb_i^{\text{pattern}} &= \frac{T^{\text{pattern}}}{tc_i}, \\ nb_{p_{\text{max}}+1}^{\text{pattern}} &= \frac{T^{\text{pattern}} - t_{\text{com}} \sum_{i=1}^{p_{\text{max}}} nb_i^{\text{pattern}}}{t_{\text{com}}} \end{aligned}$$

and let

$$nb_i^{\text{pattern}} = 0, \quad \forall i > p_{\text{max}} + 1.$$

We see that processor number $p_{\text{max}} + 1$ is not used fully, while following processors are not used at all.

With these notations, we define $f_{\text{startcomm}}^{\text{pattern}}$, $f_{\text{startcomp}}^{\text{pattern}}$ and $f_{\text{proc}}^{\text{pattern}}$ as in the previous section. Again, the only difference is that slaves P_i , $i > p_{\text{max}} + 1$ are kept idle all the time, while slave $P_{p_{\text{max}}+1}$ is kept idle during the last $(T^{\text{pattern}} - nb_{p_{\text{max}}+1}^{\text{pattern}} tc_{p_{\text{max}}+1})$ time steps.

The construction of the pattern is illustrated in Figure 10, with $t_{\text{com}} = 1$, and $p = 5$ processors s.t. $t_1 = 2$, $t_2 = 3$, $t_3 = 3$, $t_4 = 4$ and $t_5 = 6$. In this case $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}} + t_i} > 1$, $p_{\text{max}} = 3$, $T^{\text{pattern}} = 60$, and $nb_1^{\text{pattern}} = 20$, $nb_2^{\text{pattern}} = 15$, $nb_3^{\text{pattern}} = 15$, $nb_4^{\text{pattern}} = \frac{60 - (20+15+15)*1}{1} = 10$ and $nb_5^{\text{pattern}} = 0$. Slave P_5 is not used at all, while slave P_4 stays idle the last $(60 - 10 * 5) = 10$ time-steps of the pattern.

We extend the definition of $f_{\text{startcomm}}^{\text{pattern}}$, $f_{\text{startcomp}}^{\text{pattern}}$ and $f_{\text{proc}}^{\text{pattern}}$ to $f_{\text{startcomm}}$, $f_{\text{startcomp}}$ and f_{proc} exactly as before. Again, one can easily check that the functions $f_{\text{startcomm}}$, $f_{\text{startcomp}}$ and f_{proc} satisfy all the conditions stated in Section 2.4.




Figure 10: Example when $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}} + t_i} > 1$.

5.4 Asymptotic optimality

In this section, we show that the algorithm presented in Sections 5.2 and 5.3 is asymptotically optimal.

- Consider first the case $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}} + t_i} \leq 1$. Let N be the number of tasks that can be performed using our algorithm. We define k such that

$$kT^{\text{pattern}} \leq T - t_{\text{com}} \sum_{i=1}^{p-1} nb_i^{\text{pattern}} < (k+1)T^{\text{pattern}}.$$

In this expression, $t_{\text{com}} \sum_{i=1}^{p-1} nb_i^{\text{pattern}}$ represents the delay before slave P_p begins to receive its first message, as illustrated in Figure 11, and k represents the number of patterns that can be entirely completed before time bound T . Then,

$$k \geq \frac{T}{T^{\text{pattern}}} - \frac{\sum_{i=1}^{p-1} nb_i^{\text{pattern}}}{T^{\text{pattern}}} - 1.$$

Therefore,

$$\begin{aligned} N &\geq k \sum_{i=1}^p nb_i^{\text{pattern}} \\ &\geq k \sum_{i=1}^p \frac{T^{\text{pattern}}}{t_{\text{com}} + t_i} \\ &\geq \left(\frac{T}{T^{\text{pattern}}} - \frac{\sum_{i=1}^{p-1} nb_i^{\text{pattern}}}{T^{\text{pattern}}} - 1 \right) \sum_{i=1}^p \frac{T^{\text{pattern}}}{t_{\text{com}} + t_i} \\ &\geq \left(\sum_{i=1}^p \frac{T}{t_{\text{com}} + t_i} \right) - \left(\sum_{i=1}^{p-1} nb_i^{\text{pattern}} + T^{\text{pattern}} \right) \left(\sum_{i=1}^p \frac{1}{t_{\text{com}} + t_i} \right) \end{aligned}$$

Lemma 1 (equation 5.1) states that the optimal number of tasks N_{opt} that can be performed within T units of time satisfies

$$N_{\text{opt}} \leq \sum_{i=1}^p \frac{T}{t_{\text{com}} + t_i}.$$




Figure 11: Several consecutive patterns.

Then we have,

$$1 \geq \frac{N}{N_{opt}} \geq 1 + O\left(\frac{1}{T}\right),$$

which achieves the proof of the asymptotic optimality of our algorithm in the case $\sum_{i=1}^p \frac{t_{com}}{t_{com}+t_i} \leq 1$.

- Consider now the case $\sum_{i=1}^p \frac{t_{com}}{t_{com}+t_i} > 1$. Let N be the number of tasks that can be performed using our algorithm. We define k such that

$$kT^{\text{pattern}} \leq T - t_{com} \sum_{i=1}^{p_{max}} nb_i^{\text{pattern}} < (k+1)T^{\text{pattern}}.$$

In this expression, $t_{com} \sum_{i=1}^{p_{max}} nb_i^{\text{pattern}}$ represents the delay before slave $P_{p_{max}+1}$ begins to receive its first message and k represents the number of patterns that can be entirely completed before time bound T . Then,

$$k \geq \frac{T}{T^{\text{pattern}}} - \frac{t_{com} \sum_{i=1}^{p_{max}} nb_i^{\text{pattern}}}{T^{\text{pattern}}} - 1.$$

Therefore,

$$\begin{aligned} N &\geq k \left(\sum_{i=1}^{p_{max}} nb_i^{\text{pattern}} + nb_{p_{max}+1} \right) \\ &\geq k \frac{T^{\text{pattern}}}{t_{com}} \\ &\geq \left(\frac{T}{t_{com}} \right) - \left(\sum_{i=1}^{p_{max}} nb_i^{\text{pattern}} + \frac{T^{\text{pattern}}}{t_{com}} \right). \end{aligned}$$

Lemma 1 (equation 5.2) states that the optimal number of tasks N_{opt} that can be performed within T units of time satisfies

$$N_{opt} \leq \frac{T}{t_{com}}.$$

Then we have

$$1 \geq \frac{N}{N_{opt}} \geq 1 + O\left(\frac{1}{T}\right),$$

which achieves the proof of the asymptotic optimality of our algorithm in the case $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}}+t_i} > 1$.

We formally state this important result:

Theorem 2 *Let $N_{\text{opt}}(T)$ be the optimal number of tasks that can be executed within T time-steps. Let $N(T)$ be the number of tasks executed by the algorithm of Section 5.2 if $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}}+t_i} \leq 1$, and by the algorithm of Section 5.2 if $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}}+t_i} > 1$. Then*

$$\lim_{T \rightarrow \infty} \frac{N(T)}{N_{\text{opt}}(T)} = 1.$$

5.5 Comparison with a greedy algorithm

In this section, we compare the results obtained with the algorithm presented in Sections 5.2 and 5.3 against the results obtained with a greedy algorithm, which works as follows: at each time step, if k slaves with respective cycle times $t_{i_1} \leq t_{i_2} \leq \dots \leq t_{i_k}$ are waiting for a communication from the master, and if the communication network is free during the next t_{com} time steps, then a communication is performed between the master and the fastest slave P_{i_1} . In Figure 12, we display the solution obtained with $t_{\text{com}} = 1$ and $p = 4$ slave processors with $t_1 = 2$, $t_2 = 3$, $t_3 = 3$ and $t_4 = 5$. These results are to be compared with those obtained by our algorithm (here, we have $\sum_{i=1}^p \frac{t_{\text{com}}}{t_{\text{com}}+t_i} = 1$), and displayed in Figure 9.

The greedy algorithm also leads to a periodic allocation (the time period is 9); it is able to process 8 tasks every 9 time steps but neither the computing resources nor the communication medium are saturated. Our algorithm is able to process 12 tasks every 12 time steps, thus leading to an improvement of order $\frac{9}{8}$.

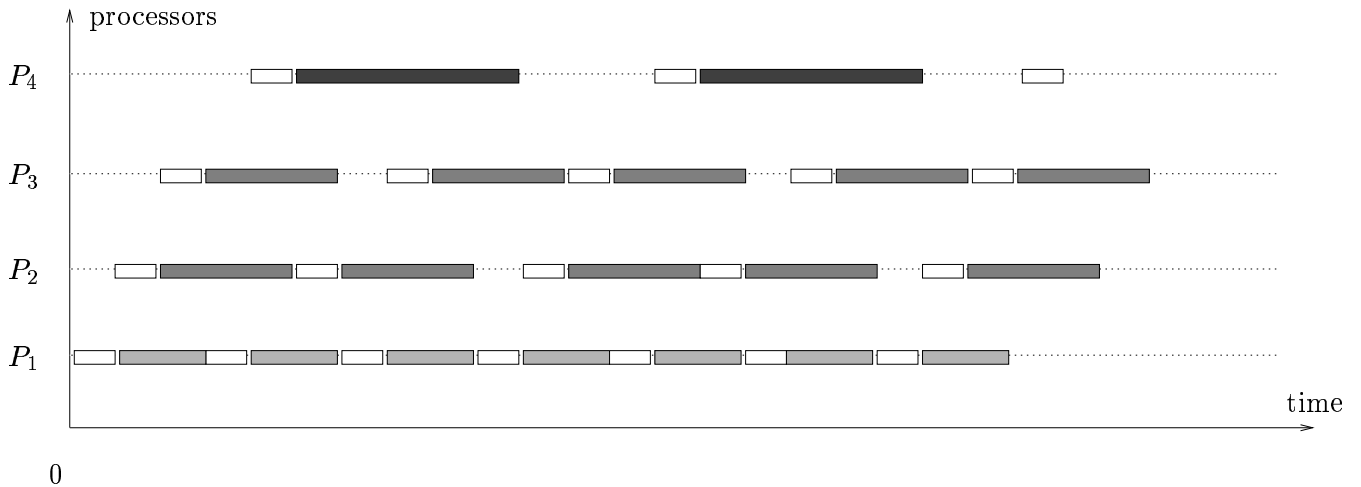


Figure 12: Execution with the greedy algorithm.

6 Solution with communications both before and after each task

In this section, we present an asymptotically optimal algorithm for $\text{MaxTasks4}(T)$. The algorithm is very similar to the one presented in Section 5, so we only outline the sketch of the algorithm, and describe it through an example, without detailing the proofs.

As previously, we define a pattern for communications and computations, that will be reproduced periodically. The pattern consists in two main phases:

- The first phase consists in both backward and forward communications between the master and the slaves,
- The second phase consists in task processing by the slaves.

Let t_{com}^1 be the communication cost for the messages from the master to the slaves, t_{com}^2 the communication cost for the messages from the slaves to the master, t_i , $1 \leq i \leq p$ the cycle times of the slaves, and T the time bound. Basically, the pattern of communications and computations is the same as those defined in Section 5, with $t_{\text{com}} = t_{\text{com}}^1 + t_{\text{com}}^2$.

The construction of the pattern is illustrated in Figure 13, with $t_{\text{com}}^1 = 2$, $t_{\text{com}}^2 = 1$, and $p = 3$ processors: $t_1 = 8$, $t_2 = 8$, and $t_3 = 9$. In this case, $t_{\text{com}} = 3$, $\sum_i \frac{t_{\text{com}}}{t_{\text{com}} + t_i} \leq 1$, $T^{\text{pattern}} = 132$, and $nb_1^{\text{pattern}} = 12$, $nb_2^{\text{pattern}} = 12$ and $nb_3^{\text{pattern}} = 11$.




Figure 13: Pattern when communications are required both before and after each task

Of course, the first pattern is not executed entirely, since no backward communication is required between the slaves and the master at the beginning of the execution. Similarly, the processing of tasks during the last pattern may be useless, since corresponding backward messages from the slaves to the master may well not have been completed.

Nevertheless, this does not impact the asymptotic optimality of this algorithm:

Theorem 3 *Let $N_{\text{opt}}(T)$ be the optimal number of tasks that can be executed within T time-steps, and let $N(T)$ be the number of tasks executed by our algorithm. Then*

$$\lim_{T \rightarrow \infty} \frac{N(T)}{N_{\text{opt}}(T)} = 1.$$

7 Related work

To the best of our knowledge, the most related work is presented in the paper of Andonie et al [1] which we have already quoted in Section 3.

Several theoretical papers deal with complexity results for parallel machine problems with a server, establishing complexity (NP-completeness) results [8, 10, 4] and providing guaranteed approximations [11]. Before processing, each job must be loaded on a machine, which takes a certain setup time. All these setups have to be done by a single server which can handle at most one job at a time. Our first problem (with initial messages only) is a very special instance of this class of server problems.

Our second problem (with initial and final messages) is a special instance of the job-shop scheduling problem (see problem SS18 in [5]) where each job consists of only three tasks, the first and last of which having to be executed by the two machines dedicated to communications. Because this instance is very specific, we do not know its complexity (polynomial versus NP-complete).

Generally speaking, note that our four problems differ from those studied in the literature with a server and start-up times in that (i) all tasks are identical and independent, and (ii) communication times (the counterpart of the set-up times) are identical too. The difficulty lies solely in the heterogeneity of the computing resources.

8 Conclusion

In this paper, we have shown that deriving efficient algorithms for the master-slave paradigm, in the framework of heterogeneous computing resources and communication links used in exclusive mode, turns out to be surprisingly difficult. More precisely, we have designed an optimal polynomial algorithm in the case of an initial scattering of data and provided a guaranteed polynomial approximation algorithm in the case of initial and final communications. We conjecture this last problem to be intrinsically difficult even on (intuitively) simple instances. Finally, we have presented asymptotically optimal algorithms for the case where each task processing must be preceded (and possibly followed) by a communication from (back to) the master.

The different variants of the master-slave problem that we have addressed in this paper seem quite representative of a large class of regular problems that exhibit a simple solution in the context of homogeneous processors but turn out to raise several algorithmic difficulties in the context of heterogeneous resources [2, 3]. Data decomposition, scheduling heuristics, load balancing, were known to be hard problems in the context of classical parallel architectures. They become extremely difficult in the context of heterogeneous clusters, not to speak about metacomputing platforms. This is a nice challenge to investigate for adventurous algorithmicians ...

References

- [1] R. Andonie, A.T. Chronopoulos, D. Grosu, and H. Galmeanu. Distributed backpropagation neural networks on a PVM heterogeneous system. In *Parallel and Distributed Computing and Systems Conference (PDCS'98)*, pages 555–560. IASTED Press, 1998.
- [2] Olivier Beaumont, Vincent Boudet, Fabrice Rastello, and Yves Robert. Matrix-matrix multiplication on heterogeneous platforms. In *2000 International Conference on Parallel Processing (ICPP'2000)*. IEEE Computer Society Press, 2000.
- [3] Pierre Boulet, Jack Dongarra, Fabrice Rastello, Yves Robert, and Frédéric Vivien. Algorithmic issues on heterogeneous computing platforms. *Parallel Processing Letters*, 9(2):197–213, 1999.
- [4] P. Brucker, C. Dhaenens-Flipo, S. Knust, S.A. Kravchenko, and F. Werner. Complexity results for parallel machine problems with a single server. Technical Report Reihe P, No. 219, Fachbereich Mathematik Informatik, Universität Osnabrück, 2000.
- [5] Michael R. Garey and Davis S. Johnson. *Computers and Intractability, a Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1991.
- [6] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam. *PVM Parallel Virtual Machine: A Users' Guide and Tutorial for Networked Parallel Computing*. The MIT Press, 1996.
- [7] M. Gondran and M. Minoux. *Graphs and Algorithms*. John Wiley & Sons, 1984.
- [8] N. Hall, C.N. Potts, and C. Sriskandarajah. Parallel machine scheduling with a common server. *Discrete Applied Mathematics*, 102:223–243, 2000.
- [9] Steve Hedetniemi. Open Problems in Combinatorial Optimization. World Wide Web document, URL: <http://www.cs.clemson.edu/~hedet/algorithms.html>.
- [10] S.A. Kravchenko and F. Werner. Parallel machine scheduling problems with a single server. *Mathematical Computational Modelling*, 26:1–11, 1997.
- [11] H. Lee and M. Guignard. A hybrid bounding procedure for the workload allocation problem on parallel unrelated machines with setups. *Journal of the Operational Research Society*, 47:1247–1261, 1996.
- [12] Kizhake Soman, Robert Fraczekiewicz, Christian Mumenthaler, Berthold von Freyberg, and Thomas Schaumann and Werner Braun. FANTOM - (Fast Newton - Raphson Torsion Angle Minimizer). World Wide Web document, URL: http://www.scsb.utmb.edu/fantom/fm_home.html. a program for "the

calculation of conformations of linear and cyclic polypeptides and proteins with low conformational energies including distance and dihedral angle constraints from nuclear magnetic resonance experiments or for modeling purposes."

- [13] J.R. Stiles, T.M. Bartol, M.M. Salpeter, and M.M. Salpeter. Monte Carlo simulation of neuromuscular transmitter release using MCell, a general simulator of cellular physiological processes. *Computational Neuroscience*, pages 279–284, 1998.
- [14] J. Watts and S. Taylor. A practical approach to dynamic load balancing. *IEEE Transactions on Parallel and Distributed Systems*, 9(3):235–248, 1998.
- [15] D.B. West. *Introduction to Graph Theory*. Prentice Hall, 1996.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399