

Scheduling Divisible Loads on Star and Tree Networks: Results and Open Problems

O. Beaumont¹, H. Casanova², A. Legrand³, Y. Robert³, Y. Yang²

1: LaBRI, UMR CNRS 5800, Bordeaux, France

`Olivier.Beaumont@labri.fr`

2: Dept. of Computer Science and Engineering,

San Diego Supercomputer Center,

University of California, San Diego, USA

`{casanova,yang}@cs.ucsd.edu`

3: LIP, UMR CNRS-INRIA 5668, ENS Lyon, France

`{Arnaud.Legrand,Yves.Robert}@ens-lyon.fr`

Abstract

Many applications in scientific and engineering domains are structured as large numbers of independent tasks with low granularity. These applications are thus amenable to straightforward parallelization, typically in master-worker fashion, provided that efficient scheduling strategies are available. Such applications have been called *divisible loads* because a scheduler may *divide* the computation among worker processes arbitrarily, both in terms of number of tasks and of task sizes. Divisible load scheduling has been an active area of research for the last fifteen years. A vast literature offers results and scheduling algorithms for various models of the underlying distributed computing platform. Broad surveys are available that report on accomplishments in the field. By contrast, in this paper we propose a unified theoretical perspective that synthesizes previously published results, several novel results, and open questions, in a view to foster novel divisible load scheduling research. Specifically, we discuss both one-round and multi-round algorithms, and we restrict our scope to the popular star and tree network topologies, which we study with both linear and affine cost models for communication and computation.

Keywords: parallel computing, scheduling, divisible load.

I. INTRODUCTION

Scheduling the tasks of a parallel application on the resources of a distributed computing platform efficiently is critical for achieving high performance. The scheduling problem has

been studied for a variety of application models, such as the well-known directed acyclic task graph model for which many scheduling heuristics have been developed [1]. Another popular application model is that of independent tasks with no task synchronizations and no inter-task communications. Applications conforming to this admittedly simple model arise in most fields of science and engineering. A possible model for independent tasks is one for which the number of tasks and the task sizes, i.e. their computational costs, are set in advance. In this case, the scheduling problem is akin to bin-packing and a number of heuristics have been proposed in the literature. Another flavor of the independent tasks model is one in which the number of tasks and the task sizes can be chosen arbitrarily. This corresponds to the case when the application consists of an amount of computation, or *load*, that can be divided into any number of independent pieces. This corresponds to a perfectly parallel job: any sub-task can itself be processed in parallel, and on any number of workers. In practice, this model is an approximation of an application that consists of large numbers of identical, low-granularity computations. This *divisible load* model has been widely studied in the last several years, and *Divisible Load Theory* (DLT) has been popularized by the landmark book written in 1996 by Bharadwaj, Ghose, Mani and Robertazzi [2].

DLT provides a practical framework for the mapping on independent tasks onto heterogeneous platforms, and has been applied to a large spectrum of scientific problems, including linear algebra [3], image processing [4], [5], video and multimedia broadcasting [6], [7], database searching [8], [9], and the processing of large distributed files [10]. These applications are amenable to the simple master-worker programming model and can thus be easily implemented and deployed on computing platforms ranging from small commodity clusters to computational grids [11]. From a theoretical standpoint, the success of the divisible load model is mostly due to its analytical tractability. Optimal algorithms and closed-form formulas exist for the simplest instances of the divisible load problem. This is in sharp contrast with the theory of task graph scheduling, which abounds in NP-completeness theorems and in inapproximability results.

There exists a vast literature on DLT. In addition to the landmark book [2], two introductory surveys have been published recently [12], [13]. Furthermore, a special issue of the *Cluster Computing* journal is entirely devoted to divisible load scheduling [14], and a Web page collecting DLT-related papers is maintained [15]. Consequently, the goal of this paper

is not to present yet another survey of DLT theory and its various applications. Instead, we focus on relevant theoretical aspects: we aim at synthesizing some important results for realistic platform models. We give a new presentation of several previously published results, and we add a number of new contributions. The material in this paper provides the level of detail and, more importantly, the unifying perspective that are necessary for fostering new research in the field.

In this paper, we limit our discussion to star-shaped and tree-shaped logical network topologies, because they often represent the solution of choice to implement master-worker computations. Note that the star network encompasses the case of a bus, which is a homogeneous star network. We consider two types of model for communication and computation: linear or affine in the data size. Our major results include several optimality results in the case where each processor receives a single message (one-round strategies). We also provide new analytical formulations and performance characterizations in the more general case where each processor receives its load divided into several sub-messages (multi-round strategies).

The extended version of this paper [16] reviews works that study other network topologies. It also surveys constant-cost models, Numerous extensions to the original DLT framework have been proposed recently. Release times and buffer-size capacities are taken into account in [17], multiple applications competing for resources are dealt with in [18], and extensions to grid platforms are considered in [19], [20]. In contrast, we limit ourselves to the basic case of a single load to be distributed on a star-shaped or tree-shaped platform.

The rest of this paper is organized as follows. In Section II, we detail our platform and cost models. We also introduce the algorithmic techniques that have been proposed to schedule divisible loads: one-round and multi-round algorithms. One-round algorithms are described in detail in Section III and multi-round algorithms in Section IV. Finally, we conclude in Section V.

II. FRAMEWORK

A. Target architectures and cost models

We consider either star-graphs or tree-graphs, and either linear or affine costs, which leads to four different platform combinations.





Fig. 1. Heterogeneous star graph, with the linear cost model. Fig. 2. Heterogeneous tree graph.

As illustrated in Figure 1, a *star network* $\mathcal{S} = \{P_0, P_1, P_2, \dots, P_p\}$ is composed of a master P_0 and of p workers P_i , $1 \leq i \leq p$. There is a communication link from the master P_0 to each worker P_q . In the linear cost model, each worker P_q has a (relative) computing power w_q : it takes $X.w_q$ time units to execute X units of load on worker P_q . Similarly, it takes $X.g_q$ time units to send X units of load from P_0 to P_q . Without loss of generality we assume that the master has no processing capability (otherwise, add a fictitious extra worker paying no communication cost to simulate computation at the master).

In the affine cost model, a latency is added to computation and communication costs: it takes $W_q + X.w_q$ time units to execute X units of load on worker P_q , and $G_q + X.g_q$ time units to send X units of load from P_0 to P_q . It is acknowledged that these latencies make the model more realistic.

For communications, the one-port model is used: the master can only communicate with a single worker at a given time-step. We assume that communications can overlap computations on the workers: a worker can compute a load fraction while receiving the data necessary for the execution of the next load fraction. This corresponds to workers *equipped with a front end* as in [2]. A *bus network* is a star network such that all communication links have the same characteristics: $g_i = g$ and $G_i = G$ for each worker P_i , $1 \leq i \leq p$.

Essentially, the same one-port model, with overlap of communication with computation, is used for tree-graph networks. A tree-graph $\mathcal{T} = \{P_0, P_1, P_2, \dots, P_p\}$ (see Figure 2) simply is an arborescence rooted at the master P_0 . We still call the other resources *workers*, even

though non-leaf workers have other workers (their children in the tree) to which they can delegate work. In this model, it is assumed that a worker in the tree can simultaneously perform some computation, receive data from its parent, and communicate to at most one of its children (sending previously received data).

B. Algorithmic strategies: one-round versus multi-round

We denote by W_{total} the total load to be executed. The key assumption of DLT is that this load is perfectly divisible into an arbitrary number of pieces, or *chunks*. The master can distribute the chunks to the workers in a single *round* (also called “installment” in [2]), so that there is a single communication between the master and each worker. The problem is to determine the size of these chunks and the order in which they are sent to the workers.

We review one-round algorithms in Section III. For large loads, the single round approach is not efficient due to the idle time incurred by the last workers to receive chunks. To reduce the makespan, i.e. the total execution time, the master can send chunks to the workers in multiple rounds so that communication is pipelined and overlapped with computation. Additional questions in this case are: “How many rounds should be scheduled?”; and “What are the best chunk sizes at each round?” We discuss multi-round algorithms in Section IV.

In fact, we point out that there is no formal definition of a *round*, which intuitively is a sequence of communications from the master to different workers. The situation is clear for one-round algorithms, formally defined as computations where each participating worker receives a single message. This is not so clear for multi-round algorithms, which we define by contraposition: there is some worker that receives at least two messages. However, from the *multi-round* denomination, we expect some periodicity, i.e. some regular pattern reproduced repeatedly, in the sequence of communications emitted by the master. All published multi-round algorithms exhibit such a regular pattern.

III. ONE-ROUND ALGORITHMS

For one-round algorithms, the first problem is to determine in which order the chunks should be sent to the different workers (or equivalently to sort the workers), given that the master can perform only one communication at a time. Once the communication order has been determined, the second problem is to decide how much work should be allocated to

each worker P_i : each P_i receives α_i units of load, where $\sum_{i=1}^p \alpha_i = W_{\text{total}}$. The final objective is to minimize the makespan, i.e. the total execution time.




Fig. 3. Pattern of a solution for dispatching a divisible load, using a star network and the linear cost model. All workers complete execution at the same time-step T_f .

A. Star network and linear cost model

This is the simplest platform combination, denoted as STARLINEAR. Let α_i denote the number of units of load sent to worker P_i , such that $\sum_{i=1}^p \alpha_i = W_{\text{total}}$. Figure 3 depicts the execution, where T_i denotes idle time of P_i , i.e. the time elapsed before P_i begins its processing. The goal is to minimize the total execution time, $T_f = \max_{1 \leq i \leq p} (T_i + \alpha_i w_i)$, according to the linear model defined in Section II. In Figure 3, all the workers participate in the computation, and they all finish computing at the same time (i.e. $T_i + \alpha_i w_i = T_f, \forall i$). This is a general result:

Proposition 1: In any optimal solution of the STARLINEAR problem, all workers participate in the computation, and they all finish computing simultaneously.

Note that Proposition 1 has been proved for the case of a bus in [2]. To the best of our knowledge, this is a new result for the case of a heterogeneous star network. Furthermore, in the case of a bus, any load distribution order is optimal [2], but this property does not extend to the heterogeneous case (for which the optimal ordering will be characterized in Proposition 2 below).

Proof: We first prove that in an optimal solution all workers participate to the computation. Then, we prove that in any optimal solution, all workers finish computing simul-

taneously.

Lemma 1: In any optimal solution, all workers participate in the computation.

Proof: Consider an optimal solution: up to a renumbering of the processors, assume that the ordering is P_1, P_2, \dots, P_p . Suppose that at least one worker is kept fully idle. In this case, at least one of the α_i , $1 \leq i \leq p$, is zero. Let us denote by k the largest index such that $\alpha_k = 0$.

Case *k*: Consider the solution of STARLINEAR, where we add P_k at the end of the initial solution. (hence we use the ordering $P_1, \dots, P_{k-1}, P_{k+1}, \dots, P_p, P_k$). By construction, $\alpha_p \neq 0$, so that the communication medium is not used during at least the last $\alpha_p w_p$ time units. Therefore, it would be possible to process at least $\frac{\alpha_p w_p}{g_k + w_k} > 0$ additional units of load with worker P_k , which contradicts the assumption that the original solution was optimal: P_k does more work than in the original solution where it was kept idle, and all the other processors execute the same amount of load units.

Case *k'*: We can modify the optimal solution of STARLINEAR by giving some work to P_p without increasing the execution time. Let k' be the largest index such that $\alpha_{k'} \neq 0$. By construction, the communication medium is not used during at least the last $\alpha_{k'} w_{k'} > 0$ time units. Thus, as previously, it would be possible to process at least $\frac{\alpha_{k'} w_{k'}}{g_p + w_p} > 0$ additional units of load with worker P_p , which leads to a similar contradiction.

Therefore, in any optimal solution, all workers participate in the computation. ■

It is worth pointing out that the above property does not hold true if we consider solutions in which the communication ordering is fixed *a priori*. For instance, consider a platform comprising two workers: P_1 (with $g_1 = 4$ and $w_1 = 1$) and P_2 (with $g_2 = 1$ and $w_2 = 1$). If the first chunk has to be sent to P_1 and the second chunk to P_2 , the optimal number of units of load that can be processed within 10 time units is 5, and P_1 is kept fully idle in this solution. On the other hand, if the communication ordering is not fixed, then 6 units of load can be performed within 10 time units (5 units of load are sent to P_2 , and then 1 to P_1). In the optimal solution, both workers perform some computation, and both workers finish computing at the same time, which is stated in the following lemma.

Lemma 2: In the optimal schedule, all workers finish computing simultaneously.

Proof: Consider an optimal solution. All the α_i 's have strictly positive values (Lemma 1). Consider the following linear program:

$$\begin{aligned} & \text{MAXIMIZE } \sum \beta_i, \\ & \text{SUBJECT TO} \\ & \begin{cases} \text{LB}(i) \quad \forall i, & \beta_i \geq 0 \\ \text{UB}(i) \quad \forall i, & \sum_{k=1}^i \beta_k g_k + \beta_i w_i \leq T \end{cases} \end{aligned}$$

The α_i 's satisfy the set of constraints above, and from any set of β_i 's satisfying the set of inequalities, we can build a valid solution of the STARLINEAR problem that process exactly $\sum \beta_i$ units of load. Therefore, if we denote by $(\beta_1, \dots, \beta_p)$ an optimal solution of the linear program, then $\sum \beta_i = \sum \alpha_i$.

It is known that one of the extremal solutions \mathcal{S}_1 of the linear program is one of the convex polyhedron \mathcal{P} induced by the inequalities [21, chapter 11]: this means that in the solution \mathcal{S}_1 , there are at least p inequalities among the $2p$ equalities. Since we know that for any optimal solution of the STARLINEAR problem, all the β_i 's are strictly positive (Lemma 1), then this vertex is the solution of the following (full rank) linear system

$$\forall i, \quad \sum_{k=1}^i \beta_k g_k + \beta_i w_i = T.$$

Thus, we derive that there is an optimal solution where all workers finish their work at the same time.

Let us denote by $\mathcal{S}_2 = (\alpha_1, \dots, \alpha_p)$ another optimal solution, with $\mathcal{S}_1 \neq \mathcal{S}_2$. As already pointed out, \mathcal{S}_2 belongs to the polyhedron \mathcal{P} . Now, consider the following function f :

$$f : \begin{cases} \mathbb{R} & \rightarrow \mathbb{R}^n \\ x & \mapsto \mathcal{S}_1 + x(\mathcal{S}_2 - \mathcal{S}_1) \end{cases}$$

By construction, we know that $\sum \beta_i = \sum \alpha_i$. Thus, with the notation $f(x) = (\gamma_1(x), \dots, \gamma_p(x))$:

$$\forall i, \gamma_i(x) = \beta_i + x(\alpha_i - \beta_i),$$

and therefore

$$\forall x, \quad \sum \gamma_i(x) = \sum \beta_i = \sum \alpha_i.$$

Therefore, all the points $f(x)$ that belong to \mathcal{P} are extremal solutions of the linear program.

Since \mathcal{P} is a convex polyhedron and both \mathcal{S}_1 and \mathcal{S}_2 belong to \mathcal{P} , then $\forall 0 \leq x \leq 1$, $f(x) \in \mathcal{P}$. Let us denote by x_0 the largest value of $x \geq 1$ such that $f(x)$ still belongs to \mathcal{P} : at least one constraint of the linear program is an equality in $f(x_0)$, and this constraint is not satisfied for $x > x_0$. Could this constraint be one of the $\text{UB}(i)$'s? the answer is no, because otherwise this constraint would be an equality along the whole line $(\mathcal{S}_2 f(x_0))$, and would remain an equality for $x > x_0$. Hence, the constraint of interest is one of the $\text{LB}(i)$'s. In other terms, there exists an index i such that $\gamma_i(x_0) = 0$. This is a contradiction since we have proved that the γ_i 's correspond to an optimal solution of the STARLINEAR problem. Therefore $\mathcal{S}_1 = \mathcal{S}_2$, the optimal solution is unique, and in this solution, all workers finish computing simultaneously. ■

Altogether, this concludes the proof of Proposition 1. ■

To be able to characterize the optimal solution, there remains to determine the best ordering for the master P_0 to send work to the workers:

Proposition 2: An optimal ordering for the STARLINEAR problem is obtained by serving the workers in the ordering of non decreasing link capacities g_i .

We give a new, shorter proof of this proposition 2, which originally appeared in [2, chapter 7], but under the hypothesis that all workers participate and finish computing simultaneously. Note that we have shown that this latter property indeed holds for any optimal schedule.

Proof: The proof is based on the comparison of the amount of work that is performed by the first two workers, and then proceeds by induction. To simplify notations, assume that P_1 and P_2 have been selected as the first two workers. There are two possible orderings, as illustrated in Figure 4. For each ordering, we determine the total number of units of load $\alpha_1 + \alpha_2$ that are processed in T time-units, and the total occupation time, t , of the communication medium during this time interval. We denote with upper-script (A) (resp. (B)) all the quantities related to the first (resp. second) ordering.

Let us first determine the different quantities $\alpha_1^{(A)}$, $\alpha_2^{(A)}$, and $t^{(A)}$ for the upper ordering in Figure 4:

- From the equality $\alpha_1^{(A)}(g_1 + w_1) = T$, we get:

$$\alpha_1^{(A)} = \frac{T}{g_1 + w_1}. \quad (1)$$

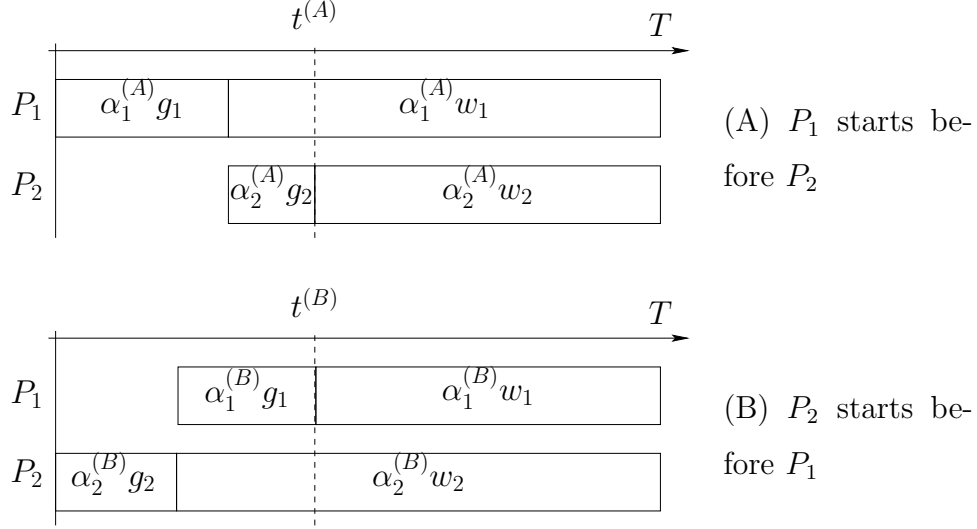


Fig. 4. Comparison of the two possible orderings.

- Using the equality $\alpha_1^{(A)} g_1 + \alpha_2^{(A)} (g_2 + w_2) = T$, we obtain (from equation (1)):

$$\alpha_2^{(A)} = \frac{T}{g_2 + w_2} - \frac{T g_1}{(g_1 + w_1)(g_2 + w_2)}. \quad (2)$$

Therefore, the overall number of processed units of load is equal to (by (1) and (2)):

$$\alpha_1^{(A)} + \alpha_2^{(A)} = \frac{T}{g_1 + w_1} + \frac{T}{g_2 + w_2} - \frac{T g_1}{(g_1 + w_1)(g_2 + w_2)}. \quad (3)$$

and the overall occupation time of the network medium is equal to (using the previous equalities and $t^{(A)} = \alpha_1^{(A)} g_1 + \alpha_2^{(A)} g_2$):

$$t^{(A)} = \frac{T g_1}{g_1 + w_1} + \frac{T g_2}{g_2 + w_2} - \frac{T g_1 g_2}{(g_1 + w_1)(g_2 + w_2)}. \quad (4)$$

A similar expression can be obtained for scenario (B) and we derive that:

$$(\alpha_1^{(A)} + \alpha_2^{(A)}) - (\alpha_1^{(B)} + \alpha_2^{(B)}) = \frac{T(g_2 - g_1)}{(g_1 + w_1)(g_2 + w_2)}, \quad (5)$$

and

$$t^{(A)} = t^{(B)}. \quad (6)$$

Thanks to these expressions, we know that the occupation of the communication medium does not depend on the communication ordering. Therefore, we only need to consider the number of processed units of load in both situations. Equation (5) indicates that one should send chunks to the worker with the smallest g_i first: if $g_2 > g_1$, more work is done in time T with the ordering (A) than with (B).

We now proceed to the general case. Suppose that the workers are already sorted so that $g_1 \leq g_2 \leq \dots \leq g_p$. Consider an optimal ordering of the communications σ , where chunks are sent successively to $P_{\sigma(1)}, P_{\sigma(2)}, \dots, P_{\sigma(p)}$. Let us assume that there exists an index i such that $\sigma(i) > \sigma(i+1)$. Furthermore, let us consider the smallest such index if multiple ones exist. Consider now the following ordering:

$$P_{\sigma(1)}, \dots, P_{\sigma(i-1)}, P_{\sigma(i+1)}, P_{\sigma(i)}, P_{\sigma(i+2)}, \dots, P_{\sigma(p)}.$$

Then, $P_{\sigma(1)}, \dots, P_{\sigma(i-1)}, P_{\sigma(i+2)}, \dots, P_{\sigma(p)}$ perform exactly the same number of units of load, since the exchange does not affect the overall communication time, but together, $P_{\sigma(i+1)}$ and $P_{\sigma(i)}$ perform $\frac{T(g_{\sigma(i)} - g_{\sigma(i+1)})}{(g_{\sigma(i+1)} + w_{\sigma(i+1)})(g_{\sigma(i)} + w_{\sigma(i)})}$ more units of load, where T denotes the remaining time after communications to $P_{\sigma(1)}, \dots, P_{\sigma(i-1)}$. Therefore, the initial ordering σ is not optimal, which is a contradiction. Therefore index i does not exist, which proves that in an optimal ordering the workers are sorted by non-decreasing values of the g_i 's. ■

According to Proposition 2, we now re-order the workers so that $g_1 \leq g_2 \leq \dots \leq g_p$. The following linear program aims at computing the optimal distribution of the load:

$$\begin{array}{ll} \text{MINIMIZE } T_f, & \\ \text{SUBJECT TO} & \\ \left\{ \begin{array}{ll} (1) \alpha_i \geq 0 & 1 \leq i \leq p \\ (2) \sum_{i=1}^p \alpha_i = W_{\text{total}} \\ (3) \alpha_1 g_1 + \alpha_1 w_1 \leq T_f & \text{(first communication)} \\ (4) \sum_{j=1}^i \alpha_j g_j + \alpha_i w_i \leq T_f & \text{(i-th communication)} \end{array} \right. & \end{array}$$

Theorem 1: The optimal solution for the STARLINEAR problem is given by the solution of the linear program above.

Proof: Direct consequence of Propositions 1 and 2. Note that inequalities (3) and (4) will be in fact equalities in the solution of the linear program, so that we can easily derive a closed-form expression for T_f . ■

We point out that this is linear programming with rational numbers, hence of polynomial complexity. Finally, we consider the variant where the master is capable of processing chunks (with computing power w_0) while communicating to one of its children. It is easy to see that the master is kept busy at all times (otherwise more units of load could be processed). The

optimal solution is therefore given by the following linear program (where $g_1 \leq g_2 \leq \dots \leq g_p$ as before):

$$\begin{aligned} & \text{MINIMIZE } T_f, \\ & \text{SUBJECT TO} \\ & \left\{ \begin{array}{ll} (1) \alpha_i \geq 0 & 0 \leq i \leq p \\ (2) \sum_{i=0}^p \alpha_i = W_{\text{total}} & \\ (3) \alpha_0 w_0 \leq T_f & \text{(computation of the master)} \\ (4) \alpha_1 g_1 + \alpha_1 w_1 \leq T_f & \text{(first communication)} \\ (5) \sum_{j=1}^i \alpha_j g_j + \alpha_i w_i \leq T_f & \text{(} i\text{-th communication)} \end{array} \right. \end{aligned}$$

Closed-form expressions have been derived for the optimal solution of the STARLINEAR problem, both for homogeneous and heterogeneous star networks [2]. The formulation in terms of linear program is useful because it can be directly extended to deal with tree networks (Section III-B). It also serves as the basis for the mixed linear program formulation in the affine cost model (Section III-C), and to derive the asymptotically optimal multi-round algorithm (Section IV-C). Finally, we point out that the first use of linear programming in the DLT theory appeared in [22].

B. Tree network and linear cost model

All the results in the previous section can be extended to a tree-shaped network. There is however a key difference with the beginning of Section III-A: each worker now is capable of computing and communicating to one of its children simultaneously. However, because of the one-round hypothesis, no overlap can occur with the incoming communication from the node's parent.

We use a recursive approach, which replaces any set of leaves and their parent by a single worker of equivalent computing power. This idea of collapsing subnets into equivalent processors originates in [23], [24].

Lemma 3: A single-level tree network with parent P_0 (with input link of capacity g_0 and cycle-time w_0) and p children P_i , $1 \leq i \leq p$ (with input link of capacity g_i and cycle-time w_i), where $g_1 \leq g_2 \leq \dots \leq g_p$, is equivalent to a single node with same input link capacity




Fig. 5. Replacing a single-level tree by an equivalent node.

g_0 and cycle-time $w_{-1} = 1/W$ (see Figure 5), where W is the solution to the linear program:

$$\begin{array}{l}
 \text{MAXIMIZE } W, \\
 \text{SUBJECT TO} \\
 \left\{ \begin{array}{l}
 (1) \alpha_i \geq 0 \qquad \qquad \qquad 0 \leq i \leq p \\
 (2) \sum_{i=0}^p \alpha_i = W \\
 (3) Wg_0 + \alpha_0 w_0 \leq 1 \\
 (4) Wg_0 + \alpha_1 g_1 + \alpha_1 w_1 \leq 1 \\
 (5) Wg_0 + \sum_{j=1}^i \alpha_j g_j + \alpha_i w_i \leq 1
 \end{array} \right.
 \end{array}$$

Proof: Here, instead of minimizing the time T_f required to execute load W , we aim at determining the maximum amount of work W that can be processed in one time-unit. Obviously, after the end of the incoming communication, the parent should be constantly computing. We know that all children (i) participate in the computation and (ii) terminate execution at the same-time. Finally, the optimal ordering for the children is given by Proposition 2. This completes the proof. Note that inequalities (3), (4) and (5) will be in fact equalities in the solution of the linear program, so that we can easily derive a closed-form expression for $w_{-1} = 1/W$. ■

Lemma 3 provides a constructive way of solving the problem for a general tree. First we traverse it from bottom to top, replacing each single-level tree by the equivalent node. We do this until there remains a single star. We solve the problem for the star, using the results of Section III-A. Then we traverse the tree from top to bottom, and undo each transformation in the reverse ordering. Going back to a reduced node, we know which amount of time it is working. Knowing the ordering, we know which amount of time each of the children is

working. If one of this children is a leaf node, we have computed its load. If it is a reduced node, we apply the transformation recursively.

Instead of this pair of tree traversals, we could write down the linear program for the whole tree: when it receives something, a given node knows exactly what to do: compute itself all the remaining time, and feed its children in decreasing bandwidth order. However, the size of the linear program would grow proportionally to the size of the tree, hence the recursive solution is to be preferred.

C. Star network and affine cost model

To the best of our knowledge, the complexity of the STARAFFINE problem is open. The main difficulty arises from resource selection: contrarily to the linear case where all workers participate in the optimal solution, it seems difficult to decide which resources to use when latencies are introduced. However, the second property proved in Proposition 1, namely simultaneous termination, still holds true:

Proposition 3: In an optimal solution of the STARAFFINE problem, all participating workers finish computing at the same time.

Proof: The proof is very similar to the STARLINEAR case. Details can be found in the extended version [16]. ■

Proposition 4: If the load is large enough, then for any optimal solution (i) all workers participate and (ii) chunks must be sent in the order of non decreasing link capacities g_i .

Proof: Consider a valid solution of the STARAFFINE problem with time bound T . Suppose, without loss of generality, that $\alpha_{\sigma(1)}$ units of load are sent to $P_{\sigma(1)}$, then $\alpha_{\sigma(2)}$ to $P_{\sigma(2)}$, \dots and finally $\alpha_{\sigma(k)}$ to P_k , where $\mathcal{S} = \{P_{\sigma(1)}, \dots, P_{\sigma(k)}\}$ is the set of workers that participate to the computation. Here, σ represents the communication ordering and is a one-to-one mapping from $[1 \dots k]$ to $[1 \dots p]$. Moreover, let n^{TASK} denote the optimal number of units of load that can be processed using this set of workers and this ordering.

- Consider the following instance of the STARLINEAR problem, with k workers $P'_{\sigma(1)}, \dots, P'_{\sigma(k)}$, where $\forall i, G'_i = 0, W'_i = 0, g'_i = g_i, w'_i = w_i$ and $T' = T$. Since all computation and communication latencies have been taken out, the optimal number of units of load n_1^{TASK} processed by this instance is larger than the number of units of load n^{TASK} processed by

the initial platform. From Theorem 1, the value of n_1^{TASK} is given by a formula

$$n_1^{\text{TASK}} = f(\mathcal{S}, \sigma) \cdot T,$$

where $f(\mathcal{S}, \sigma)$ is either derived from the linear program, or explicitly given by a closed form expression in [25]. What matters here is that the value of n_1^{TASK} is proportional to T .

- Consider now the following instance of the STARLINEAR problem, with k workers $P'_{\sigma(1)}, \dots, P'_{\sigma(k)}$, where $\forall i, G'_i = 0, W'_i = 0, g'_i = g_i, w'_i = w_i$ and $T' = T - \sum_{i \in \mathcal{S}} (G_i + W_i)$. Clearly, the optimal number of units of load n_2^{TASK} processed by this instance of the STARLINEAR problem is lower than n^{TASK} , since it consists in adding all the communication and computation latencies before the beginning of the processing. Moreover, as previously n_2^{TASK} is given by the formula

$$n_2^{\text{TASK}} = f(\mathcal{S}, \sigma) \left(T - \sum_{i \in \mathcal{S}} (G_i + W_i) \right).$$

Therefore, we have

$$f(\mathcal{S}, \sigma) \left(1 - \frac{\sum_{i \in \mathcal{S}} (G_i + W_i)}{T} \right) \leq \frac{n^{\text{TASK}}}{T} \leq f(\mathcal{S}, \sigma).$$

Hence, when T becomes arbitrarily large, then the throughput of the platform, $\frac{n^{\text{TASK}}}{T}$, becomes arbitrarily close to $f(\mathcal{S}, \sigma)$, i.e. the optimal throughput if there were no communication and computation latencies. Moreover, we have proved that if there are no latencies, then $f(\mathcal{S}, \sigma)$ is maximal when \mathcal{S} is the set of all the workers, and when σ satisfies

$$g_j > g_i \implies \sigma(i) > \sigma(j).$$

Therefore, when T is sufficiently large, then all the workers should be used and the chunks should be sent to workers in the ordering of non decreasing link capacities g_i . In this case, if $g_1 \leq \dots \leq g_p$, then the following linear system provides an asymptotically optimal solution

$$\forall i, \quad \sum_{k=1}^i (G_k + g_k \alpha_k) + W_i + g_i w_i = T.$$

This solution is optimal if all g_i are different. Determining the best way to break ties among workers having the same bandwidth is an open question. ■

In the general case, we do not know whether there exists a polynomial-time algorithm to solve the STARAFFINE problem. However, we can provide the solution (with potentially

exponential cost) as follows: we start from the mixed linear programming formulation of the problem proposed by Drozdowski [8], and we extend it to include resource selection. In the following program, y_j is a boolean variable that equals 1 if P_j participates in the solution, and $x_{i,j}$ is a boolean variable that equals 1 if P_j is chosen for the i -th communication from the master:

MINIMIZE T_f ,

SUBJECT TO

$$\left\{ \begin{array}{ll} (1) \alpha_i \geq 0 & 1 \leq i \leq p \\ (2) \sum_{i=1}^p \alpha_i = W_{\text{total}} & (3) y_j \in \{0, 1\} \quad 1 \leq j \leq p \\ (4) x_{i,j} \in \{0, 1\} & 1 \leq i, j \leq p \\ (5) \sum_{i=1}^p x_{i,j} = y_j & 1 \leq j \leq p \\ (6) \sum_{j=1}^p x_{i,j} \leq 1 & 1 \leq i \leq p \\ (7) \alpha_j \leq W y_j & 1 \leq j \leq p \\ (8) \sum_{j=1}^p x_{1,j} (G_j + \alpha_j g_j + W_j + \alpha_j w_j) \leq T_f & \text{(first communication)} \\ (9) \sum_{k=1}^{i-1} \sum_{j=1}^p x_{k,j} (G_j + \alpha_j g_j) + \sum_{j=1}^p x_{i,j} (G_j + \alpha_j g_j + W_j + \alpha_j w_j) \leq T_f & 2 \leq i \leq p \quad \text{(i-th communication)} \end{array} \right.$$

Equation (5) implies that P_j is involved in exactly one communication if $y_j = 1$, and in no communication otherwise. Equation (6) states that at most one worker is activated for the i -th communication; if $\sum_{j=1}^p x_{i,j} = 0$, the i -th communication disappears. Equation (7) states that no work is given to non participating workers (those for which $y_j = 0$) but is automatically fulfilled by participating ones. Equation (8) is a particular case of equation (9), which expresses that the worker selected for the i -th communication (where $i = 1$ in equation (8) and $i \geq 2$ in equation (9)) must wait for the previous communications to complete before starting its own communication and computation, and that all this quantity is a lower bound of the makespan. Contrarily to the formulation of Drozdowski [8], this mixed linear program always has a solution, even if a strict subset of the resources are participating. We state this result formally:

Proposition 5: The optimal solution for the STARAFFINE problem is given by the solution of the mixed linear program above (with potentially exponential cost).

D. Tree network and affine cost model

This is the most difficult platform/model combination, and very few results are known. However, we point out that Proposition 4 can be extended to arbitrary tree networks: when

T becomes arbitrarily large, latencies become negligible, and an asymptotically optimal behavior is obtained by involving all resources and by having each parent communicate with its children in order of non decreasing link capacities.

IV. MULTI-ROUND ALGORITHMS

Under the one-port communication model described in Section II-A, one-round algorithms lead to poor utilization of the workers. As seen in Figure 3, worker P_i remains idle from time 0 to time T_i . To alleviate this problem, *multi-round* algorithms have been proposed. These algorithms dispatch the load in multiple rounds of work allocation and thus improve overlap of communication with computation. By comparison with one-round algorithms, work on multi-round algorithms has been scarce. The two main questions that must be answered are: (i) what should the chunk sizes be at each round? and (ii) how many rounds should be used? The majority of works on multi-round algorithms assume that the number of rounds is fixed and we review corresponding results and open questions in Section IV-A. In Section IV-B we describe recent work that attempts at answering question (ii). Finally, we deal with asymptotic results in Section IV-C, which of course are of particular interest when the total load W_{total} is very large.

A. Fixed number of rounds, homogeneous star network, affine Costs

As for one-round algorithms, a key question is that of the order in which chunks should be sent to the workers. However, to the best of our knowledge, all previous work on multi-round algorithms with fixed number of rounds only offer solution for homogeneous platforms, in which case worker ordering is not an issue. Given a fixed number of rounds M , the load is divided into $p \times M$ chunks, each corresponding to a α_j ($j = 0, \dots, pM - 1$) units of load such that $\sum_{j=0}^{pM-1} \alpha_j = W_{\text{total}}$. The objective is to determine the α_j values that minimize the overall makespan.

Intuitively, the chunk size should be small in the first rounds, so as to start all workers as early as possible and thus maximize overlap of communication with computation. It has been shown that the chunk sizes should then increase to optimize the usage of the total available bandwidth of the network and to amortize the potential overhead associated with each chunk. In the last round, chunk sizes should be decreasing so that all workers finish




Fig. 6. Pattern of a solution for dispatching the load of a divisible job, using a bus network ($g_i = g$), in multiple rounds, for 4 workers. All 4 workers complete execution at the same time. Chunk sizes increase during each of the first $M - 1$ rounds and decrease during the last round.

computing at the same time (following the same principle as in Section III). Such a schedule is depicted in Figure 6 for four workers.

Bharadwaj et al. were the first to address this problem with the multi-installment scheduling algorithm described in [26]. They reduce the problem of finding an optimal schedule to that of finding a schedule that has essentially the following three properties: (i) there is no idle time between consecutive communications on the bus; (ii) there is no idle time between consecutive computation on each worker; and (iii) all workers should finish computing at the same time. These properties guarantee that the network and compute resources are at maximum utilization.

In [26], the authors consider only linear costs for both communication and computation. The three conditions above make it possible to obtain a recursion on the α_j series. This recursion must then be solved to obtain a close form expression for the chunk sizes. One method to solve the recursion is to use generating functions and the rational expansion theorem [27].

We recently extended the multi-installment approach to account for affine costs [28]. This was achieved by rewriting the chunk size recursion in a way that is more amenable to the use of generating functions when fixed latencies are incurred for communications and computations. Since it is more general but similar in spirit, we only present the affine case

here.

For technical reasons, as in [26], we number the chunks in the reverse order in which they are allocated to workers: the last chunk is numbered 0 and the first chunk is numbered $Mp - 1$. Instead of developing a recursion on the α_j series directly, we define $\gamma_j = \alpha_j * w$, i.e. the time to compute a chunk of size α_j on a worker not including the W latency. Recall that in this section we only consider homogeneous platforms and thus $w_q = w$, $G_q = G$, $g_q = g$, and $G_q = G$ for all workers $q = 1, \dots, p$. The time to communicate a chunk of size α_j to a worker is $G + \gamma_i/R$, where R is the computation-communication ratio of the platform: w/g . We can now write the recursion on the γ_j series:

$$\forall j \geq P \quad W + \gamma_j = (\gamma_{j-1} + \gamma_{j-2} + \gamma_{j-3} + \dots + \gamma_{j-N})/R + P \times G \quad (7)$$

$$\forall 0 \leq j < P \quad W + \gamma_j = (\gamma_{j-1} + \gamma_{j-2} + \gamma_{j-3} + \dots + \gamma_{j-N})/R + j \times G + \gamma_0 \quad (8)$$

$$\forall j < 0 \quad \gamma_j = 0 \quad (9)$$

Eq. 7 ensures that there is no idle time on the bus and at each worker in the first $M - 1$ rounds. More specifically, Eq. 7 states that a worker must compute a chunk in exactly the time required for all the next P chunks to be communicated, including the G latencies. This equation is valid only for $j \geq P$. For $j < P$, i.e. the last round, the recursion must be modified to ensure that all workers finish computing at the same time, which is expressed in Eq. 8. Finally, Eq. 9 ensures that the two previous equations are correct by taking care of out-of-range α_j terms. This recursion describes an infinite α_j series, and the solution to the scheduling problems is given by the first pM values.

As in [26], we use generating functions as they are convenient tools for solving complex recursions elegantly. Let $\mathcal{G}(x)$ be the generating function for the series γ_j , that is $\mathcal{G}(x) = \sum_{j=0}^{\infty} \gamma_j x^j$. Multiplying Eq. 7 and Eq. 8, manipulating the indices, and summing the two gives:

$$\mathcal{G}(x) = \frac{(\gamma_0 - P \times G)(1 - x^P) + (P \times G - W) + G\left(\frac{x(1-x^{P-1})}{1-x}\right) - (P-1)x^P}{(1-x) - x(1-x^P)/R}.$$

The rational expansion method [27] can then be used to determine the coefficients of the above polynomial fraction, given the roots of the denominator polynomial, $Q(x)$. The values of the γ_j series, and thus of the α_j series, follow directly. If $Q(x)$ has only roots of degree 1 then the simple rational expansion theorem can be used directly. Otherwise the more complex

general rational expansion theorem must be used. In [28] we show that if $R \neq P$ then $Q(x)$ has only roots of degree one. If $R = P$, then the only root of degree higher than 1 is root $x = 1$ and it is of degree 2, which makes the application of the general theorem straightforward. Finally, the value of γ_0 can be computed by writing that $\sum_{j=0}^{Mp-1} \gamma_j = W_{\text{total}} \times w$. All technical details on the above derivations are available in a technical report [28]. We have thus obtained a closed-form expression for optimal multi-installment schedule on a homogeneous star network with affine costs.

B. Computed number of rounds, star network, affine costs

The work presented in the previous section assumes that the number of rounds is fixed and provided as input to the scheduling algorithm. In the case of linear costs, the authors in [2] recognize that infinitely small chunks would lead to an optimal multi-round schedule, which implies an infinite number of rounds. When considering more realistic affine costs there is a clear trade-off. While using more rounds leads to better overlap of communication with computation, using fewer rounds reduces the overhead due to the fixed latencies. Therefore, an key question is: What is the optimal number of rounds for multi-round scheduling on a star network with affine costs?

While this question is still open for the recursion described in Section IV-A, our work in [29] proposes a scheduling algorithm, Uniform Multi-Round (UMR), that uses a restriction on the chunk size: all chunks sent to workers during a round are identical. This restriction limits the ability to overlap communication with computation, but makes it possible to derive an optimal number of rounds due to a simpler recursion on chunk sizes. Furthermore, this approach is applicable to both homogeneous and heterogeneous platforms. We only describe here the algorithm in the homogeneous case. The heterogeneous case is similar but involves more technical derivations and we refer the reader to [30] for all details.

As seen in Figure 7, chunks of identical size are sent out to workers within each round. Because chunks are uniform it is not possible to obtain a schedule with no idle time in which each worker finishes receiving a chunk of load right when it can start executing it. Note in Figure 7 that workers can have received a chunk entirely while not having finished to compute the previous chunk. The condition that a worker finishes receiving a chunk right when it can start computing is only enforced for the worker P_p , which is also seen in




Fig. 7. Pattern of a solution for dispatching the load of a divisible job, using a bus network ($g_i = g$), in multiple **uniform** rounds, for 4 workers. All workers complete execution at the same time. Chunk sizes are fixed within the first $M - 1$ rounds but increase from round to round. Chunk sizes decrease during the last round.

the figure. Finally, the uniform round restriction is removed for the last round. As in the multi-installment approach described in Section IV-A, chunks of decreasing sizes are sent to workers in the last round so that they can all finish computing at the same time.

Let α_j be the chunk size at round j , which is used for all workers during that round. We derive a recursion on the chunk size. To maximize bandwidth utilization, the master must finish sending work for round $j + 1$ to all workers right when worker P finishes computing for round j . This can be written as

$$W + \alpha_j w = P(G + \alpha_{j+1} g), \quad (10)$$

which reduces to

$$\alpha_j = \left(\frac{g}{Pw}\right)^j (\alpha_0 - \gamma) + \gamma, \quad (11)$$

where $\gamma = \frac{1}{w - Pg} \times (PG - W)$. The case in which $w - Pg = 0$ leads to a simpler recursion and we do not consider it here for the sake of brevity.

Given this recursion on the chunk sizes, it is possible to express the scheduling problem as a constrained minimization problem. The total makespan, \mathcal{M} , is:

$$\mathcal{M}(M, \alpha_0) = \frac{W_{\text{total}}}{P} + MW + \frac{1}{2} \times P(G + g\alpha_0),$$

where the first term is the time for worker P to perform its computations, the second term the overhead incurred for each of these computations, and the third term is the time for the

master to dispatch all the chunks during the first round. Note that the $\frac{1}{2}$ factor in the above equation is due to the last round during which UMR does not keep chunk sizes uniform so that all workers finish computing at the same time (see [29] for details).

Since all chunks must satisfy the constraint that they add up to the entire load, one can write that:

$$\mathcal{G}(M, \alpha_0) = \sum_{j=0}^{M-1} P\alpha_j - W_{\text{total}} = 0. \quad (12)$$

The scheduling problem can now be expressed as the following constrained optimization problem: minimize $\mathcal{M}(M, \alpha_0)$ subject to $\mathcal{G}(M, \alpha_0) = 0$. An analytical solution using the Lagrange Multiplier method [31] is given in [29], which leads to a single equation for the optimal number of round, M^* . This equation cannot be solved analytically but is eminently amenable to a numerical solution, e.g. using a bisection method.

The UMR algorithm is a heuristic and has been evaluated in simulation for a large number of scenarios [30]. In particular, a comparison of UMR with the multi-installment algorithm discussed in Section IV-A demonstrates the following. The uniform chunk restriction minimally degrades performance compared to multi-installment when latencies are small (i.e. when costs are close to being linear). However, as soon as latencies become significant, this performance degradation is offset by the fact that an optimal number of rounds can be computed and UMR outperforms multi-installment consistently. Finally, note that a major benefit of UMR is that, unlike multi-installment, it is applicable to heterogeneous platforms. In this case the question of worker ordering arises and UMR uses the same criterion as that given in Proposition 2: workers are ordered by non-decreasing link capacities.

C. Asymptotic performance, star network, affine costs

In this section, we derive asymptotically optimal algorithms for the multi-round distribution of divisible loads. An algorithm is asymptotically optimal if the ratio of the time to execute a workload W_{total} over the optimal time to execute this workload tends to 1 as W_{total} tends to infinity. As in previous sections, we use a star network with affine costs.

The sketch of the algorithm that we propose is as follows: the overall processing time T is divided into k regular periods of duration T_p (hence $T = kT_p$, but k (and T_p) are yet to be determined). How to choose them? Intuitively, the constraints are the following.

We will describe a steady-state operation, divided into periods of duration T_p . Initialization and clean-up will lead to “loosing” a few periods at the beginning and at the end of the execution, so there should be enough periods to render this sacrifice negligible. However, within each period, we want the impact of start-up costs to become negligible too. Hence the period length should be large enough. Altogether, a good trade-off will be to use periods whose length T_p is proportional to the square-root of the optimal execution time T^* , hence the number of periods k will also be proportional to T^* . The major difficulty is to perform the resource-selection. To decide which resource to involve, we come back to a linear programming formulation, as detailed below.

During a period of duration T_p , the master sends α_i units of load to worker P_i . It may well be the case that not all the workers are involved in the computation. Let $\mathcal{I} \subset \{1, \dots, p\}$ represent the subset of indices of participating workers. For all $i \in \mathcal{I}$, the α_i 's must satisfy the following inequality, stating that communication resources are not exceeded:

$$\sum_{i \in \mathcal{I}} (G_i + \alpha_i g_i) \leq T_p. \quad (13)$$

Since the workers can overlap communications and processing, the following inequalities also hold true:

$$\forall i \in \mathcal{I}, \quad W_i + \alpha_i w_i \leq T_p.$$

Let us denote by $\frac{\alpha_i}{T_p}$ the average number of units of load that worker P_i processes during one time unit, then the system becomes

$$\left\{ \begin{array}{l} \forall i \in \mathcal{I}, \quad \frac{\alpha_i}{T_p} w_i \leq 1 - \frac{W_i}{T_p} \quad (\text{overlap}) \\ \sum_{i \in \mathcal{I}} \frac{\alpha_i}{T_p} g_i \leq 1 - \frac{\sum_{i \in \mathcal{I}} G_i}{T_p} \quad (\text{1-port model}) \end{array} \right. ,$$

and our aim is to maximize the overall number of units of load processed during one time unit, i.e. $\rho = \sum_{i \in \mathcal{I}} \frac{\alpha_i}{T_p}$.

Consider the following linear program:

$$\begin{aligned} & \text{MAXIMIZE } \sum_{i=1}^p \frac{\alpha_i}{T_p}, \\ & \text{SUBJECT TO} \\ & \left\{ \begin{array}{l} \forall 1 \leq i \leq p, \quad \frac{\alpha_i}{T_p} w_i \leq 1 - \frac{\sum_{i=1}^p G_i + W_i}{T_p} \\ \sum_{i=1}^p \frac{\alpha_i}{T_p} g_i \leq 1 - \frac{\sum_{i=1}^p G_i + W_i}{T_p} \end{array} \right. \end{aligned}$$

This linear program is more constrained than the previous one, since $1 - \frac{W_i}{T_p}$ and $1 - \frac{\sum_{i \in \mathcal{I}} G_i}{T_p}$ have been replaced by $1 - \frac{\sum_{i=1}^p G_i + W_i}{T_p}$ in p inequalities. The linear program can be solved using a package similar to Maple (we have rational numbers), but it turns out that the technique developed in [32] enables us to obtain the solution in closed form. We refer the reader to [32] for the complete proof. Let us sort the g_i 's so that $g_1 \leq g_2 \leq \dots \leq g_p$, and let q be the largest index so that $\sum_{i=1}^q \frac{g_i}{w_i} \leq 1$. If $q < p$, let ε denote the quantity $1 - \sum_{i=1}^q \frac{g_i}{w_i}$. If $p = q$, we set $\varepsilon = g_{q+1} = 0$, in order to keep homogeneous notations. This corresponds to the case where the full use of all the workers does not saturate the 1-port assumption for out-going communications from the master. The optimal solution to the linear program is obtained with

$$\forall 1 \leq i \leq q, \quad \frac{\alpha_i}{T_p} = \frac{1 - \frac{\sum_{i=1}^p G_i + W_i}{T_p}}{g_i}$$

and (if $q < p$):

$$\frac{\alpha_{q+1}}{T_p} = \left(1 - \frac{\sum_{i=1}^p G_i + W_i}{T_p} \right) \left(\frac{\varepsilon}{g_{q+1}} \right),$$

and $\alpha_{q+2} = \alpha_{q+3} = \dots = \alpha_p = 0$.

With these values, we obtain a distribution whose throughput is :

$$\rho = \sum_{i=1}^p \frac{\alpha_i}{T_p} = \left(1 - \frac{\sum_{i=1}^p G_i + W_i}{T_p} \right) \rho_{\text{opt}} \quad \text{with } \rho_{\text{opt}} = \left(\sum_{i=1}^q \frac{1}{w_i} + \frac{\varepsilon}{g_{p+1}} \right).$$

Consider now an optimal multi-round distribution of B load units, and denote by ρ^* the average number of load units that is processed within one unit of time. If we denote by β_i^* the average number of units of load that is processed by worker P_i within one unit of time, the β_i^* 's satisfy the following set of inequalities, in which the G_i 's have been removed:

$$\left\{ \begin{array}{l} \forall 1 \leq i \leq p, \quad \beta_i^* w_i \leq 1 \\ \sum_{i=1}^p \beta_i^* g_i \leq 1 \end{array} \right.$$

Here, because we have no latencies, we can safely assume that all the workers are involved (and let $\beta_i^* = 0$ for some of them). We derive that:

$$\rho^* \leq \rho_{\text{opt}}.$$

If we denote by T^* the optimal time necessary to process B units of load, then

$$T^* \geq \frac{B}{\rho^*} \geq \frac{B}{\rho_{\text{opt}}}.$$

Let us denote by T the time necessary to process all B units of load with the algorithm that we propose. Since the first period is lost for processing, then the number k of necessary periods satisfies $\rho T_p(k-1) \geq B$ so that we choose

$$k = \left\lceil \frac{B}{\rho T_p} \right\rceil + 1.$$

Therefore,

$$T \leq \frac{B}{\rho} + 2T_p \leq \frac{B}{\rho_{\text{opt}}} \left(\frac{1}{1 - \sum_{i=1}^p \frac{G_i + W_i}{T_p}} \right) + 2T_p,$$

and therefore, if $T_p \geq 2 \sum_{i=1}^p G_i + W_i$,

$$T \leq \frac{B}{\rho_{\text{opt}}} + 2 \frac{B}{\rho_{\text{opt}}} \sum_{i=1}^p \frac{(G_i + W_i)}{T_p} + 2T_p.$$

Finally, if we set $T_p = \sqrt{\frac{B}{\rho_{\text{opt}}}}$, we check that

$$T \leq T^* + 2 \left(\sum_{i=1}^p (G_i + W_i) + 1 \right) \sqrt{T^*} = T^* + O(\sqrt{T^*}),$$

and thus that

$$\frac{T}{T^*} \leq 1 + 2 \left(\sum_{i=1}^p (G_i + W_i) + 1 \right) \frac{1}{\sqrt{T^*}} = 1 + O\left(\frac{1}{\sqrt{T^*}}\right),$$

which completes the proof of the asymptotic optimality of our algorithm.

Note that resource selection is part of our explicit solution to the linear program: to give an intuitive explanation of the analytical solution, workers are greedily selected, fast-communicating workers first, as long as the communication to communication-added-to-computation ratio is not exceeded.

We formally state our main result, for which a detailed proof can be found in [33]:

Theorem 2: For arbitrary values of G_i , g_i , W_i and w_i and assuming communication-computation overlap, the previous periodic multi-round algorithm is asymptotically optimal. Closed-form expressions for resource selection and task assignment are provided by the algorithm, whose complexity does not depend upon the total amount of work to execute.

V. CONCLUSION

The goal of this paper was to present a unified discussion of divisible load scheduling results for star and tree networks. In Section III we have discussed one-round algorithms for which the two main issues are: (i) selection and ordering of the workers, (ii) computation of the chunk sizes. Section IV focused on multi-round algorithms, with the two main issues: (i) computation of chunk sizes at each round, and (ii) choice of the number of rounds. Section IV also discussed multi-round scheduling for maximizing asymptotic application performance. For both classes of algorithms, we have revisited previously published results, presented novel results, and clearly identified open questions. Our overall goal was to identify promising research directions and foster that research thanks to our unified and synthesized framework.

We have discussed affine cost models and have seen that they often lead to significantly more complex scheduling problems than when linear models are assumed. These models are generally considered more realistic, and we even contend that, given current trends, linear models are quickly becoming increasingly inappropriate. In terms of communication, technology trends indicate that available network bandwidth is rapidly augmenting. Therefore, latencies account for an increasingly large fraction of communication costs. A similar observation can be made in terms of computation. Due to the absence of stringent synchronization requirements, divisible workload applications are amenable to deployment on widely distributed platforms. For instance, computational grids [11] are attractive for deploying large divisible workloads. However, initiating computation on these platforms incurs potentially large latencies (i.e., due to resource discovery, authentication, creation of new processes, etc.). Consequently, it is clear that divisible workload research should focus on affine cost models for both communication and computation.

Acknowledgments. We thank the reviewers for their helpful comments and suggestions, which greatly improved the final version of the paper.

REFERENCES

- [1] B. A. Shirazi, A. R. Hurson, and K. M. Kavi, *Scheduling and load balancing in parallel and distributed systems*. IEEE Computer Science Press, 1995.

- [2] V. Bharadwaj, D. Ghose, V. Mani, and T. Robertazzi, *Scheduling Divisible Loads in Parallel and Distributed Systems*. IEEE Computer Society Press, 1996.
- [3] S. Chan, V. Bharadwaj, and D. Ghose, "Large matrix-vector products on distributed bus networks with communication delays using the divisible load paradigm: performance and simulation," *Mathematics and Computers in Simulation*, vol. 58, pp. 71–92, 2001.
- [4] C. Lee and M. Hamdi, "Parallel image processing applications on a network of workstations," *Parallel Computing*, vol. 21, pp. 137–160, 1995.
- [5] X. Li, V. Bharadwaj, and C. Ko, "Distributed image processing on a network of workstations," *Int. J. Computers and Applications (ACTA Press)*, vol. 25, no. 2, pp. 1–10, 2003.
- [6] D. Altılar and Y. Paker, "An optimal scheduling algorithm for parallel video processing," in *IEEE Int. Conference on Multimedia Computing and Systems*. IEEE Computer Society Press, 1998.
- [7] —, "Optimal scheduling algorithms for communication constrained parallel processing," in *Euro-Par 2002*, ser. LNCS 2400. Springer Verlag, 2002, pp. 197–206.
- [8] M. Drozdowski, "Selected problems of scheduling tasks in multiprocessor computing systems," Ph.D. dissertation, Instytut Informatyki Politechnika Poznańska, Poznań, 1997.
- [9] J. Blazewicz, M. Drozdowski, and M. Markiewicz, "Divisible task scheduling - concept and verification," *Parallel Computing*, vol. 25, pp. 87–98, 1999.
- [10] R. Wang, A. Krishnamurthy, R. Martin, T. Anderson, and D. Culler, "Modeling communication pipeline latency," in *Measurement and Modeling of Computer Systems (SIGMETRICS'98)*. ACM Press, 1998, pp. 22–32.
- [11] I. Foster and C. Kesselman, Eds., *The Grid: Blueprint for a New Computing Infrastructure*. San Francisco, USA: Morgan Kaufmann Publishers, Inc., 1999.
- [12] V. Bharadwaj, D. Ghose, and T. Robertazzi, "Divisible load theory: a new paradigm for load scheduling in distributed systems," *Cluster Computing*, vol. 6, no. 1, pp. 7–17, 2003.
- [13] T. Robertazzi, "Ten reasons to use divisible load theory," *IEEE Computer*, vol. 36, no. 5, pp. 63–68, 2003.
- [14] D. Ghose and T. Robertazzi, Eds., *Special issue on Divisible Load Scheduling*. Cluster Computing, 6, 1, 2003.
- [15] T. Robertazzi, "Divisible Load Scheduling," URL: <http://www.ece.sunysb.edu/~tom/dlt.html>.
- [16] O. Beaumont, H. Casanova, A. Legrand, Y. Robert, and Y. Yang, "Scheduling divisible loads for star and tree networks: main results and open problems," LIP, ENS Lyon, France, Tech. Rep. RR-2003-41, Sept. 2003.
- [17] V. Bharadwaj and G. Barlas, "Scheduling divisible loads with processor release times and finite size buffer capacity constraints in bus networks," *Cluster Computing*, vol. 6, no. 1, pp. 63–74, 2003.
- [18] —, "Efficient scheduling strategies for processing multiple divisible loads on bus networks," *J. Parallel and Distributed Computing*, vol. 62, pp. 132–151, 2002.
- [19] D. Yu and T. Robertazzi, "Divisible load scheduling for grid computing," in *PDCS'2003, 15th Int'l Conf. Parallel and Distributed Computing and Systems*. IASTED Press, 2003.
- [20] H. Wong, D. Yu, V. Bharadwaj, and T. Robertazzi, "Data intensive grid scheduling: multiple sources with capacity constraints," in *PDCS'2003, 15th Int'l Conf. Parallel and Distributed Computing and Systems*. IASTED Press, 2003.
- [21] A. Schrijver, *Theory of Linear and Integer Programming*. New York: John Wiley & Sons, 1986.
- [22] R. Agrawal and H. Jagadish, "Partitioning techniques for large-grained parallelism," *IEEE Trans. Computers*, vol. 37, no. 12, pp. 1627–1634, 1988.

- [23] T. Robertazzi, "Processor equivalence for a linear daisy chain of load sharing processors," *IEEE Trans. Aerospace and Electronic Systems*, vol. 29, pp. 1216–1221, 1993.
- [24] S. Bataineh, T. Hsiung, and T.G.Robertazzi, "Closed form solutions for bus and tree networks of processors load sharing a divisible job," *IEEE Transactions on computers*, vol. 43, no. 10, pp. 1184–1196, Oct. 1994.
- [25] S. Charcranon, T. Robertazzi, and S. Luryi, "Optimizing computing costs using divisible load analysis," *IEEE Transactions on computers*, vol. 49, no. 9, pp. 987–991, Sept. 2000.
- [26] V. Bharadwaj, D. Ghose, and V. Mani, "Multi-installment load distribution in tree networks with delays," *IEEE Trans. on Aerospace and Electronc Systems*, vol. 31, no. 2, pp. 555–567, 1995.
- [27] R. Graham, D. Knuth, and O. Patashnik, *Concrete Mathematics*. Wiley, 1994.
- [28] Y. Yang and H. Casanova, "Extensions to the multi-installment algorithm: affine costs and output data transfers," Dept. of Computer Science and Engineering, University of California, San Diego, Tech. Rep. CS2003-0754, July 2003.
- [29] —, "UMR: A multi-round algorithm for scheduling divisible workloads," in *International Parallel and Distributed Processing Symposium (IPDPS'2003), Nice, France*. IEEE Computer Society Press, April 2003.
- [30] —, "Multi-round algorithm for scheduling divisible workload applications: analysis and experimental evaluation," Dept. of Computer Science and Engineering, University of California, San Diego, Tech. Rep. CS2002-0721, 2002.
- [31] D. Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods*. Athena Scientific, Belmont, Mass., 1996.
- [32] O. Beaumont, L. Carter, J. Ferrante, A. Legrand, and Y. Robert, "Bandwidth-centric allocation of independent tasks on heterogeneous platforms," in *International Parallel and Distributed Processing Symposium (IPDPS'2002)*. IEEE Computer Society Press, 2002.
- [33] O. Beaumont, A. Legrand, and Y. Robert, "Optimal algorithms for scheduling divisible workloads on heterogeneous systems," LIP, ENS Lyon, France, Tech. Rep. 2002-36, Oct. 2002.