

Centralized versus Distributed Schedulers for Bag-of-Tasks Applications

Olivier Beaumont¹, Larry Carter², Jeanne Ferrante²,
Arnaud Legrand³, Loris Marchal⁴ and Yves Robert⁴

Abstract—Multiple applications that execute concurrently on heterogeneous platforms compete for CPU and network resources. In this paper we consider the problem of scheduling applications to ensure fair and efficient execution on a distributed network of processors. We limit our study to the case where communication is restricted to a tree embedded in the network, and the applications consist of a large number of independent tasks (Bags of Tasks) that originate at the tree’s root. The tasks of a given application all have the same computation and communication requirements, but these requirements can vary for different applications. The goal of scheduling is to maximize throughput of each application while ensuring a fair sharing of resources between applications.

We can find the optimal asymptotic rates by solving a linear programming problem that expresses all necessary problem constraints, and we show how to construct a periodic schedule from any linear program solution. For single-level trees, the solution is characterized by processing tasks with larger communication-to-computation ratios at children with larger bandwidths. For multi-level trees, this approach requires *global* knowledge of all application and platform parameters. For large-scale platforms, such global coordination by a centralized scheduler may be unrealistic. Thus, we also investigate decentralized schedulers that use *only local* information at each participating resource. We assess their performance via simulation, and compare to an optimal centralized solution obtained via linear programming. The best of our decentralized heuristics achieves the same performance on about two-thirds of our test cases, but is far worse in a few cases. While our results are based on simple assumptions and do not explore all parameters (such as the maximum number of tasks that can be held on a node), they provide insight into the important question of fairly and optimally scheduling heterogeneous applications on heterogeneous grids.

Index Terms—parallel computing, scheduling, multiple applications, bag of tasks, resource sharing, fairness, throughput.

1: Laboratoire LaBRI, CNRS-INRIA Bordeaux, France, Olivier.Beaumont@labri.fr

2: Dept. of Computer Science and Engineering, University of California, San Diego, USA

{carter, ferrante}@cs.ucsd.edu

3: CNRS, Laboratoire LIG, projet INRIA MESCAL, Grenoble, France, Arnaud.Legrand@imag.fr

4: Laboratoire LIP, CNRS-INRIA, École Normale Supérieure de Lyon, France. {Loris.Marchal, Yves.Robert}@ens-lyon.fr

I. INTRODUCTION

In this paper, we consider the problem of scheduling multiple applications that are executed concurrently, hence that compete for CPU and network resources. The target computing platform is a heterogeneous network of computers structured either as a *star network* (a one-level rooted tree) or a multi-level rooted tree. In both cases we assume full heterogeneity of the resources, both for CPU speeds and link bandwidths.

Each application consists of a large collection of independent equal-sized tasks, and all tasks originate at the tree’s root. This scenario is somewhat similar to that addressed by existing systems. For instance BOINC [1] is a centralized scheduler that distributes tasks for participating applications, such as SETI@home, ClimatePrediction.NET, and Einstein@Home. The applications can be very different in nature, e.g., files to be processed, images to be analyzed or matrices to be manipulated. Consequently, we assume each application has an associated *communication-to-computation ratio* for all of its tasks. This ratio proves to be an important parameter in the scheduling process.

The scheduling problem is to maintain a balanced execution of all applications while using the computational and communication resources of the system effectively to maximize throughput $\alpha^{(k)}$ of each application (the average number of tasks of application A_k , $1 \leq k \leq K$ processed per time-unit). For each application, the root node must decide which workers (i.e. which subtree) the tasks are sent to. For multi-level trees, each non-leaf worker must make similar decisions: which tasks to compute locally, and which to forward to workers further down in the tree. The scheduler must also ensure

a fair management of the resources. If all tasks were equally important, the scheduler should aim to process the same number of tasks for each application. We could generalize this by allowing each application A_k to be assigned a *priority weight* $w^{(k)}$ that quantifies its relative value. For instance, if $w^{(1)} = 3$ and $w^{(2)} = 1$, the scheduler should try to ensure that three tasks of A_1 are executed for each task of A_2 . However, we avoid using weights in the following to simplify the presentation¹.

We will consider both centralized and decentralized schedulers. For smaller platforms it may be realistic to assume a centralized scheduler, which makes its decisions based upon complete and reliable knowledge of all application and platform parameters. With such knowledge at our disposal, we are able to determine an *optimal* schedule, i.e. a schedule that maximizes the fair throughput asymptotically. This is done by formulating all constraints into a linear programming problem, and using the solution to construct a periodic schedule. Except during the (fixed length) start-up and clean-up periods no schedule can have higher throughput. For single-level rooted trees, we provide an interesting characterization of the optimal solution: applications with larger communication-to-computation ratio should be processed by the workers with larger bandwidths, independently of the communication-to-computation ratios of the workers.

For large-scale platforms, particularly ones in which resource availability changes over time, a centralized scheduler may be undesirable. Only local information, such as the current capacity (CPU speed and link bandwidth) of a processor's neighbors, is likely to be available. One major goal of this paper is to investigate whether decentralized scheduling algorithms can reach optimal throughput, or at least achieve a significant fraction of it. We provide several decentralized heuristics that rely exclusively on local information to make scheduling decisions. The key underlying principles of these heuristics come from our characterization of the optimal solution for star networks: give priority to high-bandwidth children, and assign them tasks of larger communication-to-computation ratios. We evaluate the decentralized heuristics through extensive simulations using SimGrid [3], and use a centralized algorithm (guided by the linear program solution) as a reference basis.

The rest of the paper is organized as follows. Section II is devoted to an overview of related work. In Section III, we state precisely the scheduling problem under con-

sideration, with all application and platform parameters, and discuss the objective function used afterwards. Section IV explains how to analytically compute the best solution, using a linear programming approach, and characterizes the solution for single-level trees. Section V is then a discussion on the design of several decentralized scheduling heuristics, while Section VI provides an experimental comparison of these heuristics. Finally, we state some concluding remarks in Section VII.

II. RELATED WORK

We classify related research in three main areas:

A. Bag of Tasks Scheduling on Computational Grids

Bag of Tasks applications are parallel applications whose tasks are all independent. This framework is motivated by problems that are addressed by collaborative computing efforts such as SETI@home [4], factoring large numbers [5], the Mersenne prime search [6], and those distributed computing problems organized by companies such as Entropia [7]. One can argue that Bag of Tasks applications are most suited for computational grids, because communication can easily become a bottleneck for tightly-coupled parallel applications.

Condor [8] and APST [9], [10] are one of the first projects providing specific support for such applications. Condor was initially conceived for campus-wide networks [8], but has been extended to run on grids [11]. While APST is user-centric and does not handle multiple-applications, Condor is system-centric. Those two projects are designed for standard grids but more recent and active projects like OurGrid [12] or BOINC [1] target more distributed architectures like desktop grids. BOINC [1] is a centralized scheduler that distributes tasks for participating applications, such as SETI@home, ClimatePrediction.NET, and Einstein@Home. The set of resources is thus very large while the set of applications is small and very controlled. OurGrid is a brazilian project that encourages people to donate their computing resources while maintaining the symmetry between consumers and providers. While APST, Condor and BOINC all rely on a centralized scheduler and may suffer from scalability issues, OurGrid is the only framework we know of where scheduling is done in a fully distributed way, following the Peer-to-Peer approach.

However, all these projects generally focus on designing and providing a working infrastructure, and they do not provide any analysis of scheduling techniques suited to such environments.

¹All results can be easily extended when adding weights: simply replace $\alpha^{(k)}$ by $\frac{\alpha^{(k)}}{w^{(k)}}$. See [2] for further details.

B. Steady-State Scheduling

Because the number of tasks to be executed on the computing platform is expected to be very large (otherwise why deploy the corresponding application on a distributed platform?), it makes sense to focus on *steady-state* optimization problems rather than on standard *makespan* minimization problems. Minimizing the makespan, i.e. the total execution time, is a NP-hard problem in most practical situations [13], [14], [15], while it turns out that the optimal steady-state can often be characterized very efficiently, with low-degree polynomial complexity.

The steady-state approach has been pioneered by Bertsimas and Gamarnik [16]. It has been used successfully in many situations [17]. In particular, steady-state scheduling has been used to schedule independent tasks on heterogeneous tree-overlay networks [18], [19]. This is the same problem dealt with in the present paper, but restricted to a single application. Bandwidth-centric scheduling is introduced in [18], and extensive experiments are reported in [20]. Autonomous protocols for bandwidth-centric scheduling are investigated by Carter *et al.* [21]. Such distributed autonomous protocols have been obtained only on tree platforms. That is why in the current more complex context, we restrict our study to tree platforms. The steady-state approach has also been used by Hong *et al.* [22] who extend the work of [18] to deploy a divisible workload on a heterogeneous platform. However to the best of our knowledge, the steady-state scheduling approach has never been used in a multiple-application context.

C. Fairness

In a multi-user environment, resources have to be fairly shared between users. This issue becomes more and more critical as the size of the system increases. There is actually a large gap between what is known on the theoretical side in game theory and what is implemented in practical Bag of Tasks scheduling environments. In most practical environments, some fairness is ensured through the use of hand-tuned priorities or reciprocation-based economy [23]. Market-inspired economy and auction-based mechanism could also be used. But in most existing work related to fairness in grid environments, the mechanism is not based on clear definitions of fairness criteria.

Fairness is yet a classical criterion in network bandwidth allocation. Optimizing the sum of the throughputs is known as maximizing the throughput or *profit* of a network. Optimizing this kind of objective is natural for an access provider who receives an amount of

money proportional to the throughput that he/she is able to provide and who wants to maximize his/her profit. However, this criterion is known to be unfair and can lead to starvation. That is why in section III-D, we choose to maximize the minimum of $\alpha^{(k)}$ rather than the sum. This criterion is known in the literature as *max-min* and is intuitively fair since all throughputs are computed to be as close as possible from each other. Between these two extremes, other criteria can be found (e.g., *proportional fairness* that maximizes $\sum \log(\alpha^{(k)})$ or *minimum potential delay* that minimizes $\sum \frac{1}{\alpha^{(k)}}$). In fact all these criteria (profit, proportional fairness and minimum potential delay) amount to maximizing the arithmetic, geometric and harmonic mean of the throughput [24]. It is well-known in the networking community [25] that max-min fairness is generally achieved by explicit-rate calculation (e.g., in ATM networks) and rather hard to achieve in a fully-decentralized way. Yet, fully distributed algorithms are known to realize proportional fairness (such as some variants of TCP). Adapting such algorithms to Bag of Tasks scheduling environments seems challenging as both communications and computations are involved.

III. FRAMEWORK AND MODELS

In this section, we clarify the assumptions underlying our work; while they are overly simplistic, we believe nevertheless that they provide insight into the important question of how to optimally and fairly schedule heterogeneous applications on heterogeneous grids.

A. Platform Model

The target computing platform is either a single-level tree (also called a *star network*) or an arbitrary tree. The processor at the root of the tree is denoted P_0 . There are P additional “worker nodes” P_1, P_2, \dots, P_P ; each worker P_u has a single parent $P_{p(u)}$, and the link between P_u and its parent has bandwidth b_u . We assume a linear-cost communication model, hence it takes X/b_u time units to send a message of size X from $P_{p(u)}$ to P_u . For sake of simplicity, we ignore processor-task affinities; instead, we assume only the number of floating-point operations per second (c_u for processor P_u) determines the application execution speed.

There are several scenarios for the operation of the processors. In this paper, we concentrate on the *full overlap, single-port model* [26], [27]. In this model, a processor node can simultaneously receive data from one of its neighbors, perform some (independent) computation, and send data to one of its neighbors. At any given time, there are at most two communications involving a given processor, one sent and the other received.

B. Application Model

We consider K applications, A_k , $1 \leq k \leq K$. The root node P_0 initially holds all the input data necessary for each application A_k . Each application is composed of a set of independent, equal-sized tasks. We can think of each A_k as a bag of tasks, and the tasks are files that require some processing. A task of application A_k is called a task of *type* k . We assume one can express the computational requirements of tasks as a number of floating-point operations, and we let $c^{(k)}$ be the amount of computation (in floating point operations) required to process a task of type k . Similarly, $b^{(k)}$ is the size (in bytes) of (the file associated to) a task of type k . We assume that the only communication required is outwards from the root, i.e. that the amount of data returned by the worker is negligible. Our results are equally applicable to the scenario in which the input to each task is negligible but the output is large. Note that our notations use subscripts for platform resources (bandwidth b_u , CPU speed c_u) and superscripts for application parameters (bytes $b^{(k)}$, floating-point operations $c^{(k)}$).

C. Steady-State Scheduling

If each application had an unlimited supply of tasks, each application should aim at maximizing its average number of task processed per time-unit (the *throughput*). When the number of tasks is very large, optimizing the steady-state throughput enables to derive periodic asymptotically optimal schedules for the makespan [18], [19]. In our setting where each application has a very large number of tasks, we should thus try to optimize the steady-state throughput of each application.

More formally, for a given infinite schedule we can define $N^{(k)}(t)$ the number of tasks of type k processed in time interval $[0, t]$. The throughput for application k of such a schedule is defined as $\alpha^{(k)} = \liminf_{t \rightarrow \infty} \frac{N^{(k)}(t)}{t}$. Similarly we can define:

- $\alpha_u^{(k)}$, the average number of tasks of type k executed by P_u per time unit.
- $sent_{u \rightarrow v}^{(k)}$, average the number of tasks of type k received by P_v from P_u per time unit.

We will see in Section IV-A that the $\alpha_u^{(k)}$'s and the $sent_{u \rightarrow v}^{(k)}$'s are linked by linear equations and satisfy linear constraints, which enables to derive upper bounds on the throughput. It is possible to build a *periodic schedule*, i.e. a schedule that repeatedly begins and ends in the same state (see [18], [19] for more details), from values for $\alpha_u^{(k)}$ and $sent_{u \rightarrow v}^{(k)}$ returned by the linear program.

The throughput achieved by this periodic schedule for each application k is optimal. In other words, when the

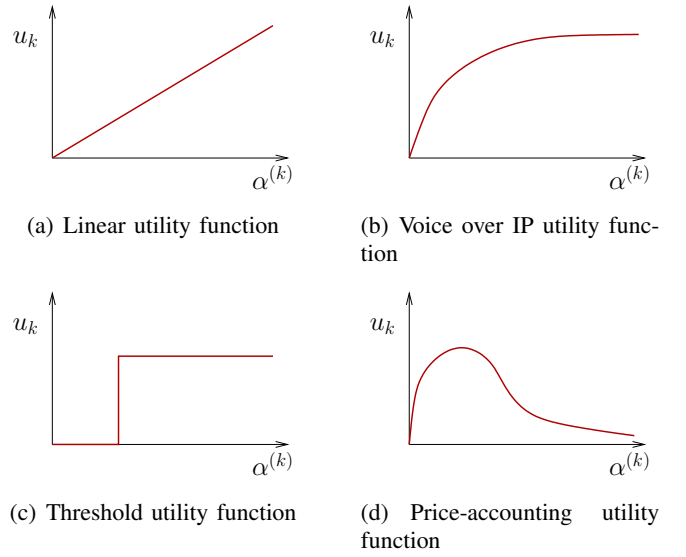


Fig. 1. Examples of utility functions

number of tasks per application is large, this approach enables to circumventing the NP-completeness of the makespan optimization problem while deriving efficient schedules.

D. Objective Function

In this section, we present a few game theory notions and how they translate to our context. This help us to correctly define a metric to optimize in our setting. Game theory provides a general framework to model situations where many users compete for resources. Each user (in our context, each application) is characterized by a *utility* function u_k defined on $(\alpha_p^{(k)})_{1 \leq k \leq K, 1 \leq p \leq P}$, where K is the number of applications, and P is the number of computing resources. A variety of possible utility functions are shown in Figure 1; in the remaining of the article, we focus on linear utility functions:

$$u_k(\alpha) = \sum_p \alpha_p^{(k)} = \alpha^{(k)}$$

Our goal is to find scheduling strategies such that the utility of *each* user is maximized. However, as these users may compete for the same resources, it is generally not possible to simultaneously maximize the utility of each user. Instead, we employ a utility set U :

$$U = \{(u_1(\alpha), \dots, u_K(\alpha)) | \alpha \text{ is feasible}\}.$$

For tree-shaped platforms, the set of constraints on $\alpha_u^{(k)}$ is a set of linear inequalities (as seen later in Section IV-A), and we know that the utility set is thus a convex polyhedron. Using the same techniques as in [28], [17], [29], one can show for general platforms

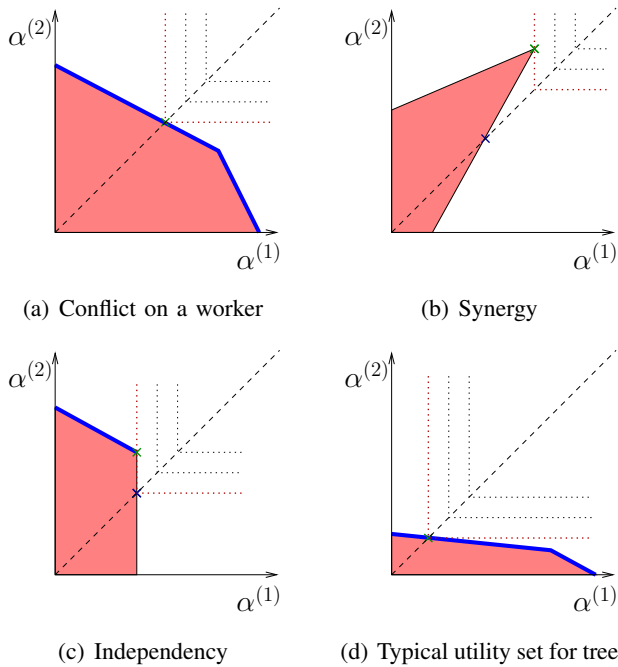


Fig. 2. A few examples of utility sets. Dotted lines are isolines of $(\alpha^{(1)}, \alpha^{(2)}) \rightarrow \min(\alpha^{(1)}, \alpha^{(2)})$ and bold lines represent Pareto optimal points.

that the utility set is also a convex polyhedron, as illustrated in Figure 2.

Figure 2(a) corresponds to the typical situation where two applications are competing on a single node:

$$\begin{cases} \alpha^{(1)} \cdot c^{(1)} + \alpha^{(2)} \cdot c^{(2)} \leq c_u & \alpha^{(1)} \geq 0 \\ \alpha^{(1)} \cdot b^{(1)} + \alpha^{(2)} \cdot b^{(2)} \leq b_u & \alpha^{(2)} \geq 0 \end{cases}$$

In a multi-user context, optimality is not defined as simply as in the single-user context, and it is common to use Pareto-optimality, defined as follows:

Definition 1 (Pareto-optimality): $\tilde{\alpha}$ is Pareto-optimal if and only if:

$$\forall \alpha, \exists i, u_i(\alpha) > u_i(\tilde{\alpha}) \Rightarrow \exists j, u_j(\alpha) < u_j(\tilde{\alpha})$$

In other words, $\tilde{\alpha}$ is Pareto optimal if it is impossible to strictly increase the utility of a user without strictly decreasing that of another. For example, in Figure 2(a), all points on the right-most boundaries are Pareto-optimal whereas in Figure 2(b), only the right-most point is Pareto-optimal. Any Pareto optimal point is thus *a priori* as worth of interest as any other Pareto optimal point. Defining fairness can be seen as defining a criterion for choosing amongst Pareto optimal points.

One of the most common fairness criterion is the well-known max-min fairness strategy [30], [25]. For a given allocation α , there is an application k whose utility is smaller than the other ones. A max-min fair allocation is such that this smaller utility is as large as possible. Such a criterion is thus more reasonably fair than for example trying to maximizing the sum of utilities (also known

as social optimum). Indeed with such a criterion, some applications may receive nothing at all (as for example in Figure 2(d)). Moreover, by using weighted throughput, any Pareto-optimal point is a max-min fair point for a given set of weights. We will thus focus in the following on max-min fairness:

$$\text{Maximize } \min_k \alpha^{(k)}$$

On Figure 2(a), the minimum of the $\alpha^{(k)}$ is maximized at a point where all the $\alpha^{(k)}$ have the same value. One can easily check that the only Pareto optimal point of Figure 2(b) is also the point such that the minimum of the $\alpha^{(k)}$ is maximized. However, one can also check that the points such that the $\alpha^{(k)}$'s all have the same value are not efficient, which corresponds to the well-known fact that giving the same thing to each user is not always a good option. In fact the shape of this utility set is rather uncommon and corresponds to a situation where there is a synergy between both users. Such situations may occur with caching mechanisms for example but cannot occur in our framework because all coefficients of the linear constraints are positive.

One may then wonder whether in our context max-min optimal solutions are always such that the throughput of all applications are the same or not. We will see now that this is true on trees but does not hold on general platforms. The utility set Figure 2(c) is typical of the case where two applications originate from different locations and where one of them can only use a limited area of the network (due to a very high communication to computation ratio and a small connectivity to the network for example). In such case, it may be possible to increase the throughput of the application with lower ratio ($\alpha^{(2)}$ here) without decreasing the throughput of the higher one ($\alpha^{(1)}$ here). However, if both applications start using the same resources, the throughput of one application can only increase at the expense of the throughput of another application. It is important to note that many different points maximize the minimum of the throughputs (all points belonging to U and to the lowest isoline of $\min(\alpha^{(1)}, \alpha^{(2)})$). However, only one of them is of interest (i.e. Pareto Optimal). It is well-known in the network community (see for example [30], [25]) that max-min fairness should be *recursively* defined in this case: the first minimum should be maximized, then the second should be maximized, and so on.

Such situations cannot occur on tree-shaped platforms as applications originate from the same location and thus always compete on the same set of resources. Note that this result does not only hold for the mentioned full overlap, single-port model but applies to any situation

where applications originate from the same location. That is why in the remaining of this article we can search for solutions where all application throughputs are equal. But in a more general situation, we should look for Pareto-optimal allocations and the previous stopping condition could not be used anymore.

IV. COMPUTING THE OPTIMAL SOLUTION

In this section, we show how to describe the optimal throughput using a linear programming formulation. For star networks we give a nice characterization of the solution, which will guide the design of some heuristics in Section V.

A. Linear Programming Solution

A summary of our notation follows:

- P_0 is the root processor and $P_{p(u)}$ is the parent of node P_u for $u \neq 0$.
- $\Gamma(u)$ is the set of indices of the children of node P_u .
- Node P_u can compute c_u floating-point operations per time unit, and, if $u \neq 0$, can receive b_u bytes from its parent $P_{p(u)}$.
- Each task of type k involves $b^{(k)}$ bytes and $c^{(k)}$ floating-point operations.

The linear programming formulation in (1) allows us to solve for the following variables:

- $\alpha_u^{(k)}$, the average number of tasks of type k executed by P_u per time unit on the whole platform.
- $\alpha^{(k)}$, the average number of tasks of type k executed per time unit.
- $sent_{u \rightarrow v}^{(k)}$, average the number of tasks of type k received by P_v from P_u per time unit.

Any valid schedule must satisfy the linear constraints of (1); we seek a schedule also satisfying the optimal value of the objective function.

$$\begin{aligned} & \text{MAXIMIZE } \min_k \{ \alpha^{(k)} \} \text{ UNDER THE CONSTRAINTS} \\ & \left\{ \begin{array}{l} \forall k, \quad \sum_u \alpha_u^{(k)} = \alpha^{(k)} \quad (\text{definition of } \alpha^{(k)}) \\ \forall k, \forall u \neq 0, \quad sent_{p(u) \rightarrow u}^{(k)} = \alpha_u^{(k)} + \sum_{v \in \Gamma(u)} sent_{u \rightarrow v}^{(k)} \quad (\text{data conservation}) \\ \forall u, \quad \sum_k \alpha_u^{(k)} \cdot c^{(k)} \leq c_u \quad (\text{computation limit at node } P_u) \\ \forall u, \quad \sum_{v \in \Gamma(u)} \frac{\sum_k sent_{u \rightarrow v}^{(k)} \cdot b^{(k)}}{b_v} \leq 1 \quad (\text{communication limit out of } P_u) \\ \forall k, u \quad \alpha_u^{(k)} \geq 0 \text{ and } sent_{u \rightarrow v}^{(k)} \geq 0 (\text{non-negativity}) \end{array} \right. \quad (1) \end{aligned}$$

All the input parameters to the linear programming problem are rational numbers, and the solution will be rational also (hence computed in polynomial time).

As explained in Section III-D, on tree-shaped platforms, the previous solution is the max-min fair solution. On general platform, similar constraints can be written but solving this linear program would not give the max-min fair solution as the first minimum should be maximized, then the second should be maximized, and so on. This can easily be done in our setting by identifying which applications correspond to the first minimum by looking at saturated constraints (those inequalities that are in fact equalities at optimum point). One can then rerun the linear program, with the throughput of these applications fixed, to maximize the smallest throughput of the remaining applications. This process can be repeated until all applications are saturated. Max-min solutions can thus easily be computed in polynomial time, even on complex platforms.

B. Reconstructing a Periodic Schedule from a Linear Programming Solution

The linear inequalities in the linear programming problem (1) describe steady state behavior, but it is not immediately obvious that there exists a valid schedule satisfying these constraints that also achieves the desired throughput. Nevertheless, let us suppose that such a solution exists, and we have determined all the values $\alpha_u^{(k)}$, and $sent_{u \rightarrow v}^{(k)}$. We define a periodic schedule as follows.

Define the time period T_{period} to be the least common multiple of the denominators of these rational values. Thus, in one time period, there will be an integral number of tasks sent over each link and executed by each node. We give each node sufficient buffer space to hold twice the number of tasks it receives per time period. Each task received in period i will, in period $i+1$, either be computed locally or sent to a child. Since each node receives tasks from only one other node (its parent), there is no concern with scheduling the incoming communications to avoid conflicts. Further, each node is free to schedule its sends arbitrarily within a time period. Note that this schedule is substantially simpler than what is required when processors were connected as an arbitrary graph (cf. [19]).

A node at depth d does not receive any tasks during the first $d-1$ time periods, so will only enter “steady state mode” in time period d . Similarly, the root will eventually run out of tasks to send, so the final time periods will also be different from steady state. It is often possible to improve the schedule in the start-up and

clean-up time periods, which is the concern of the NP-complete makespan minimization problem. However, the periodic schedule described above is asymptotically optimal. More precisely, let z be the number of tasks executed by the periodic schedule in steady state during d time periods, where d is the maximum depth of any node that executes a positive number of tasks. Then our schedule will execute up to z fewer tasks than any possible (not necessarily periodic) schedule. More precisely, given a time-bound B for the execution, it can be shown that the periodic schedule computes as many tasks of each type as the optimal, up to a constant (independent of B) number of tasks. This result is an easy generalization of the same result with a single application [18], [19]. Note that as the applications we consider consist of a large number of independent tasks, z is generally much smaller than the total number of tasks of an application.

One final comment is that the time period T_{period} , and the amount of buffer space used, can be extraordinarily large, making this schedule impractical. We will revisit this issue later.

C. The Optimal Solution for Star Networks

When the computer platform is a star network, we can prove that the optimal solution has a very particular structure: applications with higher communication-to-computation ratio are scheduled on processors with higher bandwidth. Thus, if we order the processors according to their bandwidths, then each application is executed by a set of consecutive nodes, which we refer to as a *slice*. The application with the highest communication-to-computation ratio is executed by a first slice of processors, those with largest bandwidths. Then the next most communication-intensive application is executed by the next slice of processors, and so on. There is a possible overlap at the slice boundaries. For instance P_{a_1} , the processor at the boundary of the first two slices, may execute tasks for both applications A_1 and A_2 .

To simplify matters, we consider the root P_0 to be a worker with infinite bandwidth ($b_0 = +\infty$). The following proposition proves that the optimal solution has the structure described above.

Proposition 1: Sort nodes by bandwidth so that $b_0 \geq b_1 \dots \geq b_p$, and sort the applications by communication-to-computation ratio so that $\frac{b^{(1)}}{c^{(1)}} \geq \frac{b^{(2)}}{c^{(2)}} \dots \geq \frac{b^{(K)}}{c^{(K)}}$. Then there exist indices $a_0 \leq a_1 \dots \leq a_K$ such that only processors P_u , $u \in [a_{k-1}, a_k]$, execute tasks of type k in the optimal solution.

Proof: The key idea is to show that if a node P_i is assigned a task with a lower communication-to-

computation ratio than a task assigned to P_{i+1} , then these two nodes could swap an equal amount of computational work. This would reduce the communication time required by the schedule without changing any throughput. Thus, by a sequence of such swaps, any schedule can be transformed to one of the desired structure, without changing the fair throughput. See [2] for a detailed proof. ■

This characterization does not enable to determine the boundaries of the slices nor the $\alpha_u^{(k)}$ through analytical formulas. We did not succeed in deriving a counterpart of Proposition 1 for tree-shaped platforms. Intuitively, the problem is that a high-bandwidth child of node P_i can itself have a low-bandwidth, high-compute-rate child, so there is no *a priori* reason to give P_i only communication-intensive tasks. Still, we use the intuition provided by Proposition 1 and its proof to design the heuristic of Section V-E.

V. DEMAND-DRIVEN AND DECENTRALIZED HEURISTICS

As shown in Section IV-A, given a tree-shaped platform and the set of all application parameters, we are able to compute an optimal periodic schedule from any linear programming solution. However, this approach suffers from several serious drawbacks. First, the period of the schedule is the least common multiple of the denominators of the solution of the linear program (1). This period may be huge, requiring the nodes to have unreasonably large buffers to ensure uninterrupted steady-state behavior. The problem of buffer size has already been pointed out in [21], [31], where it is shown that no finite amount of buffer space is sufficient for every tree. It is also known that finding the optimal throughput when buffer sizes are bounded is a strongly NP-hard problem even in very simple situations [31].

Since unlimited buffer space is unrealistic, we will only consider *demand-driven* algorithms, which operate as follows. Each node has a local *worker* thread and a *scheduler* thread. The worker thread is an infinite loop that requests a task from the same node's scheduler thread and then, upon receiving a task, executes it. Figure 3 shows the pseudo-code for the scheduler thread. The "select" choices in line 5 depend on the particular heuristic used, and can be based on, for instance, the history of requests, task types it has received, and the communication times it has observed for its children.

A second problem that some schedulers (including those generated as in Section IV-B) encounter is that centralized coordination becomes an issue as the size of the platform grows. It may be costly to collect up-to-date information and disseminate it to all nodes in

- 1: **Loop**
- 2: If there will be room in your buffer, request work from your parent.
- 3: Get incoming requests from your local worker and children, if any.
- 4: Move incoming tasks from your parent, if any, into your buffer.
- 5: Select which thread (your local worker or a child's scheduler) to assign work to, and the type of application that will be assigned.
- 6: **If** you have a request and a task that match your choice **Then**
- 7: Send the task to the chosen thread (when the send port is free)
- 8: **Else**
- 9: Wait for a request or a task

Fig. 3. Demand-driven scheduler thread, run in each node

the system. Consequently, a decentralized scheduling algorithm, where all choices are based exclusively on *locally* available information, is desirable.

In the following we consider one demand-driven algorithm that is based on global information (derived from a solution to the linear programming problem), and four algorithms that are fully decentralized.

A. Centralized LP-based (LP)

If we know the computation power and communication speeds of all nodes in the distributed system, we can solve the linear programming problem (IV-A), obtaining values for the number of tasks of each type it should assign to each of its children during each time period. Thereafter, no further global communication is required.

Each scheduler thread uses a 1D load-balancing mechanism [32] to select a requesting thread and an application type. The 1D load-balancing mechanism works as follows: if task k should be chosen with frequency $f(k)$, and has already been chosen $g(k)$ times, then the next task to be sent will be of type ℓ , where $\frac{g(\ell)+1}{f(\ell)} = \min_k \frac{g(k)+1}{f(k)}$.

We might hope the LP heuristic would always converge to the optimal throughput, but we will see in Section VI-B.1 that this is not always the case, primarily because of insufficient buffer space.

B. First Come First Served (FCFS)

The FCFS heuristic is a very simple and common decentralized heuristic. Each scheduler thread simply fulfills its requests on a *First Come First Served* basis, using the tasks it receives in order from its parent. The root ensures fairness by selecting task types using a round-robin selection. This simple heuristic has the disadvantage, not shared by the other methods we consider, that an extremely slow communication link cannot be avoided. Thus, optimal performance should not be expected.

C. Coarse-Grain Bandwidth-Centric (CGBC)

This heuristic (CGBC) builds upon our previous work for scheduling a single application on a tree shaped platform [18], [19]. In bandwidth-centric scheduling, each node only needs to know the bandwidth to each of its children. The node's own worker thread is considered to be a child with infinite bandwidth. The scheduler thread prioritizes its children in order of bandwidth, so the greatest bandwidth has highest priority. The scheduler always assigns tasks to the highest-priority requester. Bandwidth-centric scheduling has been shown to have optimal steady-state throughput for single application, both theoretically and, when the buffers are sufficiently large, in extensive simulations.

Our *coarse-grain* heuristic assembles several tasks into a large one. More precisely, we build a macro-task out of one task of type k , for each k , and the macro-tasks are scheduled using the bandwidth-centric method. Thus, fairness is assured.

Unfortunately, even though bandwidth-centric scheduling can give optimal throughput of macro-tasks, the CGBC heuristic does not reach the optimal fair throughput. Indeed, Proposition 1 asserts that in star networks, nodes with faster incoming links should process only tasks with larger communication-to-computation ratios. But since a macro-task includes tasks of *all* types, CGBC will send communication-intensive tasks to some low-bandwidth nodes in a star network.

D. Parallel Bandwidth-Centric (PBC)

The *parallel bandwidth-centric* heuristic (PBC) supposes bandwidth-centric trees for each type of task, running all of them in parallel. More precisely, each node has K scheduler and K worker threads that run concurrently, corresponding to the K application types. Threads only communicate with other threads of their own type.

In all our other simulations, we enforce the one-port constraint for each scheduler thread. But for this *PBC* heuristic, we have not enforced this constraint globally across the schedulers, and a node may send as many as K tasks concurrently, one of each type. Instead, we model the contention on the port so the aggregate bandwidth does not exceed the port's limit (similarly, the node's processor can multitask between multiple tasks.) This gives the *PBC* strategy an unfair advantage over the other heuristics. In fact, it has been shown [21] that allowing *interruptible communication* (which is similar to concurrent communication) dramatically reduces the amount of buffer space needed to achieve optimal throughput.

E. Data-Centric Scheduling (DATA-CENTRIC)

This heuristic is our best attempt to design a decentralized, demand-driven algorithm that converges to a solution of the linear program (1). The idea is to start from the bandwidth-centric solution for the most communication-intensive application and to progressively trade some of these tasks for more computation-intensive ones. Doing so yields better values for the expected $\alpha_u^{(k)}$ and the expected $sent_{u \rightarrow v}^{(k)}$, which can in turn be used in the demand-driven algorithm of Figure 3. These frequencies are continuously recomputed so as to cope with potential availability variations. The rest of this subsection is devoted to details of the *trading* operations. As we have explained in Section III-D, in the optimal solution on trees, all applications have the same throughput. So this heuristic starts from an initial solution and updates the $\alpha_u^{(k)}$ until all throughputs are close to each other.

We sort the task types by non-increasing communication-to-computation ratios. We start the algorithm using the pure bandwidth-centric approach for tasks of type 1, but as the computation proceeds, a node will find itself receiving a mix of different types of tasks. To reduce the imbalance, the root iteratively applies the four operations described below, in the listed order of precedence. In the following, H (resp. L) denotes the application that currently has the highest (resp. lowest) throughput. As the root distributes all tasks, H and L are easy to identify. Those operations attempt to increase the number of tasks of type L that are assigned, sometimes by reducing the number of H 's.

Communication Trading Suppose H has a higher communication-to-computation ratio than L (which is the common case since we start with only tasks of type 1). Then if a child reports that it is not fully utilized (either because its CPU is idle or because it cannot keep

up with the requests it receives from its own children i.e. $\sum_k \alpha_u^{(k)} c^{(k)} < c_u$ or $\sum_{v \in \Gamma(u)} \frac{\sum_k sent_{u \rightarrow v}^{(k)} \cdot b^{(k)}}{b_v} = 1$) then the parent can *substitute* some tasks of type H by sending them in place of some tasks of type L to the under-utilized child. It should make the substitution in a way that keeps the communication time the same (i.e. in the ratio of $b^{(L)}$ H 's for $b^{(H)}$ L 's), and limited by the number that would make the throughputs equal. Last, let CPU denote the CPU computation time to execute all tasks currently assigned to processor P_u . Then we have $CPU = \sum_k \frac{\alpha_u^{(k)} \cdot c^{(k)}}{c_u}$, and we should not exceed the CPU capacity after the update. Therefore $\alpha_u^{(L)}$ is increased by ε_L and $\alpha_u^{(H)}$ is decreased by ε_H with the following constraints:

$$\begin{cases} \varepsilon_L b^{(L)} = \varepsilon_H b^{(H)} \\ 0 \leq \alpha^{(L)} + \varepsilon_L \leq \alpha^{(H)} - \varepsilon_H \\ CPU + \varepsilon_L \frac{c^{(L)}}{c_u} - \varepsilon_H \frac{c^{(H)}}{c_u} \leq 1 \end{cases}$$

Hence, we get:

$$\varepsilon_H = \min \left(\frac{\alpha^{(H)} - \alpha^{(L)}}{1 + \frac{b^{(H)}}{b^{(L)}}}, \frac{1 - CPU}{\frac{c^{(H)} c^{(L)}}{b^{(L)} c_u} \left(\frac{b^{(H)}}{c^{(H)}} - \frac{b^{(L)}}{c^{(L)}} \right)} \right).$$

Gap filling Suppose that some bandwidth from the root is not used, and that a remote processor P_u could receive more tasks of a low throughput application. This step calculates ε_L , the possible additional number of tasks of type L that this processor could handle. Let CPU denote the CPU computation time to execute all tasks currently assigned to processor P_u . Then we have $CPU = \sum_k \frac{\alpha_u^{(k)} \cdot c^{(k)}}{c_u}$, and the following computation limit on ε_L has to hold: $CPU + \varepsilon_L \frac{c^{(L)}}{c_u} \leq 1$. In addition, there must be enough free bandwidth along the path from the root node to P_u to send the additional tasks; therefore, for any node i along this path, we can define $bus_occupation(p(i)) = \sum_k \sum_j \frac{sent_{p(i) \rightarrow j}^{(k)} \cdot b^{(k)}}{b_j}$ and we need the following condition on ε_L to hold true: $bus_occupation(p(i)) + \varepsilon_L \frac{b^{(L)}}{b_i} \leq 1$. Lastly, to avoid over-reducing the imbalance between $\alpha^{(H)}$ and $\alpha^{(L)}$, we add the following constraint: $\alpha^{(H)} \geq \alpha^{(L)} + \varepsilon_L$. Therefore, we have:

$$\varepsilon_L = \min \left(\frac{c_u(1 - CPU)}{c^{(L)}}, \alpha^{(H)} - \alpha^{(L)}, \min_{i \in \text{path from the root to } P_u} \left(\frac{b_i(1 - bus_occupation(p(i)))}{b^{(L)}} \right) \right).$$

Bus de-saturation If the bus is saturated by tasks with a high communication-to-computation ratio, we may still be using only workers with high communication capacity. However, the workload might still be increased by using additional idle subtrees, i.e., the current tree must be *widened*. Thus, we need to reduce the amount of tasks of type H that are processed by the currently




Fig. 4. Rebalancing operations

used subtrees. This heuristic takes the branch with the smallest bandwidth that processes tasks of type H , and scales down the $\alpha_i^{(H)}$ and $sent_{i \rightarrow j}^{(H)}$ values of nodes on the branch by 10%. This operation allows us to decrease the communication resource utilization and precedes the next round of “Gap filling” operations.

Task trading on the root At some point it may be the case that application H is processed only on the root node. This heuristic will try to substitute ε_H tasks of type H for ε_L tasks of type L at the root. To do so, we need the following constraints: $\varepsilon_H \leq \alpha_{root}^{(H)}$, $\alpha^{(H)} - \varepsilon_H \geq \alpha^{(L)} + \varepsilon_L$ and $\varepsilon_L \cdot \frac{c^{(L)}}{c_{root}^{(L)}} = \varepsilon_H \cdot \frac{c^{(H)}}{c_{root}^{(H)}}$. Therefore, we have

$$\varepsilon_H = \min \left(\alpha_{root}^{(H)}, \frac{\alpha^{(H)} - \alpha^{(L)}}{1 + \frac{c^{(H)}}{c^{(L)}}} \right) \text{ and } \varepsilon_L = \frac{c^{(H)}}{c^{(L)}} \varepsilon_H.$$

The preceding operations are iteratively performed (with the listed order of precedence) until we reach a (tunable) balance, for example,

$$\frac{\max_k \{\alpha^{(k)}\} - \min_k \{\alpha^{(k)}\}}{\min_k \{\alpha^{(k)}\}} < 0.05.$$

The above operations may appear to need global knowledge about the tree. For example, it may seem at first sight that when performing a “Gap filling” operation, the master needs to know the path to its remote descendant P_u . However, this operation in fact simply amounts to computing a minimum along this path, which can be done via a classical and efficient distributed propagation mechanism. The same distributed technique can be used for all other operations as they only require information from immediate descendants in a single subtree.

VI. SIMULATION RESULTS

A. Evaluation Methodology

1) *Throughput Evaluation*: It is not at all obvious how to determine that a computation has entered steady-state, and measuring throughput becomes even trickier when the schedule is not periodic. We took a pragmatic, heuristic approach for our experiments. Let T denote the earliest time that all tasks of some application were completed. Let $N^{(k)}(t)$ denote the number of tasks of type k that were finished in time period $[0, t]$. We define the achieved throughput ρ_k for application k by:

$$\rho_k = \frac{N^{(k)}((1 - \varepsilon)T) - N^{(k)}(\varepsilon T)}{(1 - 2\varepsilon)T}, \text{ where } \varepsilon \in [0, 0.5].$$

The ε factor allows us to ignore the initial and final instabilities (in practice, we set ε to be equal to 0.1). In the following, we will refer to ρ_k as the *experimental throughput* of application k as opposed to the expected throughput that can be computed by solving the linear program (1). Likewise, the minimum of the experimental throughputs is called the *experimental fair throughput*.

2) *Platform Generation*: The platforms used in our experiments are random trees described by two parameters: the number of nodes n and the maximum degree $degree_{max}$. To generate the interconnection network topology, we use a breadth-first algorithm (see [2] for more details) in order to have wide (rather than deep and narrow) trees. In our experiments, we generated a total of 150 trees, ten trees each with 5, 10, 20, 50 and 100 nodes, and maximum degree 2, 5, or 15.

We assigned capacity, latency and CPU floating point rate values on edges and nodes at random. Those values come from real measurements (performed using tools like pathchar) on machines spread across the Internet. CPU rates ranged from 22.151 Mflops (an old Pentium Pro 200MHz) to 171.667 Mflops (an Athlon 1800). Bandwidth ranged from 110 kb/s to 7 Mb/s and latency from 6 ms to 10 s. Note that in the SimGrid simulator [3] which we are using, latency and link capacity are limiting factors for determining the effective bandwidth of a connection.

3) *Application Generation*: An application is mainly characterized by its *communication-to-computation ratio (CCR)*. We used $CCR_{min} = 0.001$, which corresponds to the computation-intensive task of multiplying two 3500×3500 matrices. We also set an upper bound for CCR of 4.6, corresponding to the addition of two such matrices. In choosing application types, we picked CCR_{max} between 0.002 and 4.6, and then chose the applications’ CCR ’s to be evenly spaced in the range $[CCR_{min}, CCR_{max}]$. For simplicity we take K to be 2 or 3.

4) *Heuristic Implementation*: The experiments were performed using the SimGrid simulator [3], and the

values of c_i and b_i values were measured from within the simulator (based on the node values input), and used to make the decisions in the algorithms of Section V.

As explained in section V-E, the demand-driven algorithms send requests (involving a few bytes) from children to parents. Our simulations included the request mechanism, and we ensured that no deadlock occurred in our thousands of experiments, even when some load-variations occurred. Except where otherwise noted, throughput evaluations were performed using 200 tasks per application. Note, that we carefully checked by hand using a larger number of tasks (ten times more tasks) that the smaller number we used for our experiments was always sufficient to reach the steady-state (the experimental throughput was within 1% for all configurations when adding more tasks).




Fig. 5. Deviation of experimental fair throughput from expected theoretical throughput ($1 - \frac{\text{Experimental fair throughput}}{\text{Theoretical fair throughput}}$) for *LP*, *DATA-CENTRIC* and *CGBC*

B. Case study

1) *Theoretical versus Observed Throughput*: For the heuristics *LP*, *DATA-CENTRIC* and *CGBC*, we can easily compute the expected *theoretical* fair throughput as they all rely on explicit rate computation. This allowed us to explore how implementation issues result in the experimentally obtained fair throughput differing from the theoretical fair throughput. There are many reasons that these quantities might differ, such as the overhead of the request mechanism or a startup period longer than the 10% we allowed for. It turned out that the major cause of inefficiency was the limit on the buffer size.

Our experiments assumed enough buffer space to hold 10 tasks of any type. For this case, Figure 5 depicts the experimental fair throughput deviation from the expected theoretical throughput for heuristics *CGBC*, *LP* and *DATA-CENTRIC*. This deviation is computed as $1 - \frac{\text{Experimental fair throughput}}{\text{Theoretical fair throughput}}$. All three heuristics exhibited a similar distribution, so they were combined in this figure. The average deviation is equal to 9.426%. However,

when we increased the buffer size by a factor ten (and increased the number of tasks per application to 2000), the mean average deviation dropped to 0.334%. Even though the larger buffer size led to much better throughput, we considered it unrealistic, and used size 10 in all other experiments.

2) *Performance of Heuristics*: Let us first compare the relative performances of our five heuristics (*FCFS*, *PBC*, *CGBC*, *LP* and *DATA-CENTRIC*). More precisely, for each experimental setting (i.e. a given platform and a given *CCR* interval), we compute the (natural) logarithm of the ratio of the experimental fair throughput of *LP* with the experimental fair throughput of a given heuristic (applying a logarithm enables us to have a symmetrical value). This value is called in the following *logarithmic deviation*. Therefore, a positive value means that *LP* performed better than the other heuristic. Figure 6 depicts the histogram plots of these values.

First of all, we observe that most values are positive, which illustrates the superiority of *LP*. Next, observe on Figure 6(a) that *DATA-CENTRIC* is very close to *LP* most of the time, despite the distributed computation of the expected $\alpha_u^{(k)}$ and $sent_{u \rightarrow v}^{(k)}$ values. However, the geometric average² of these ratios is equal to 1.164, which is slightly larger than the geometric average for *CGBC* (1.156). The reason is that even though in most settings *DATA-CENTRIC* ends up with a very good solution, in a few instances it performed very badly (up to 16 times worse than *LP*). In contrast, *CGBC* (see Figure 6(d)) is much more stable since its worst performance is only twice that of *LP*. Note that these failures happen on any type of tree (small or large, narrow or wide) and that the geometric averages of these two heuristics are always very close to each other. We have also checked that these failures are not due to an artifact of the decentralized control of the scheduling by ensuring that the theoretical throughput has the same behavior (i.e. the bad behavior actually comes from the computation of the expected $\alpha_u^{(k)}$ and $sent_{u \rightarrow v}^{(k)}$). We are still investigating the reasons that *DATA-CENTRIC* fails on some instances, and suspect that it is due to the use of the (sometimes misleading) intuition of Proposition 1. Indeed, in this heuristic, applications with a high communication-to-computation ratio are performed mainly on the subtrees with the best bandwidths at the root while applications with a low communication-to-computation are performed primarily on the subtrees with the worst bandwidths at the root, which is definitely not optimal on particular

²It is a well-known fact [33] that arithmetic average of ratios can lead to contradictory conclusions when changing the reference point. Therefore, we use a geometric average of ratios which is known to be closer to the common idea of *average ratio*.

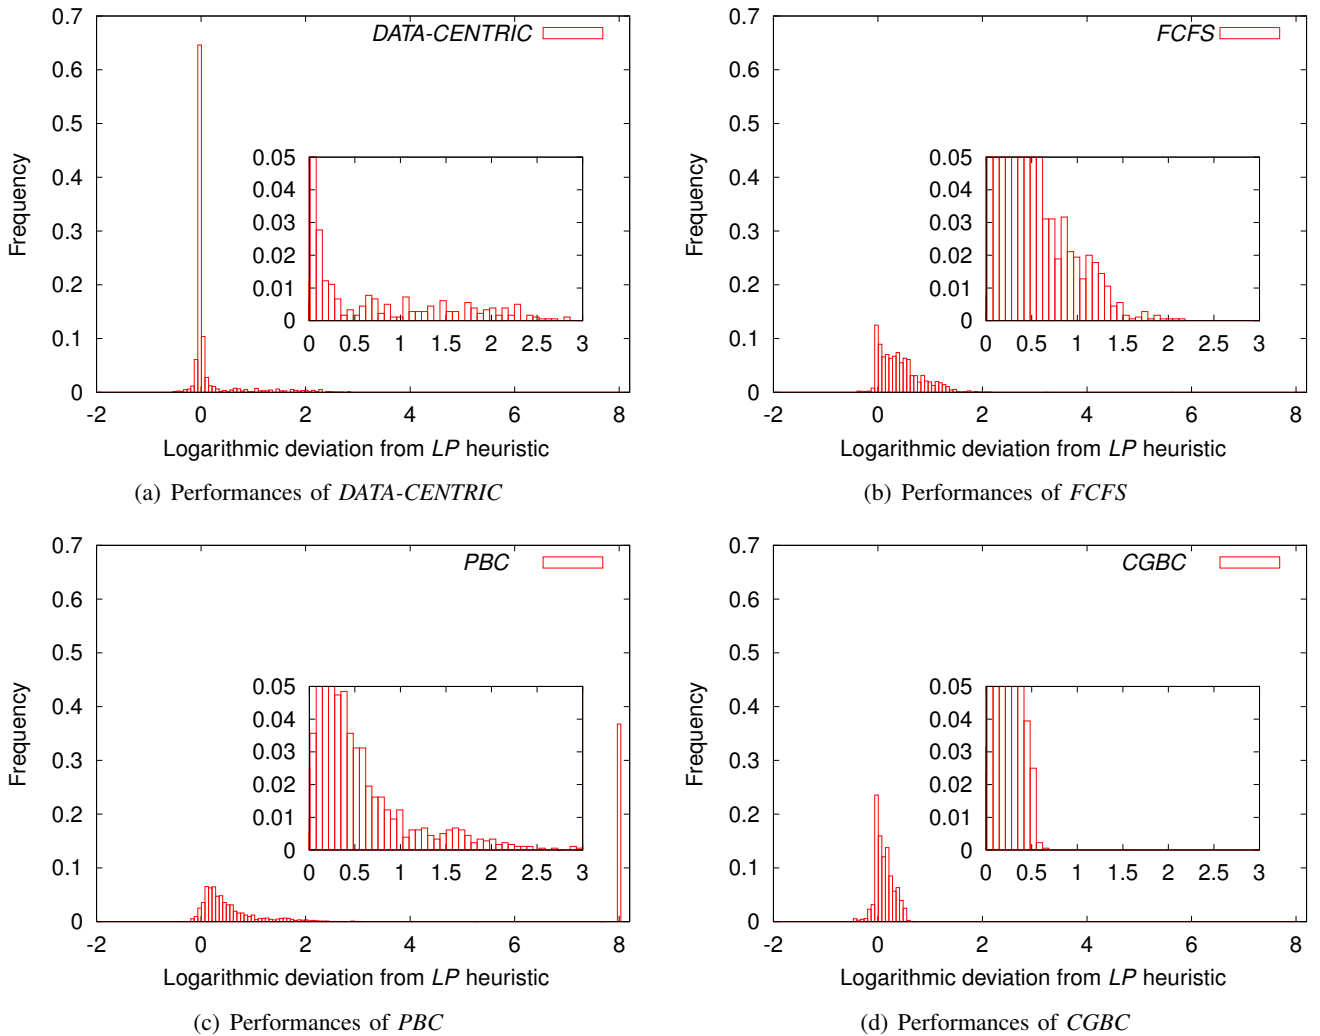


Fig. 6. Logarithmic deviation from *LP* performances ($\ln \left(\frac{\min_k \alpha^{(k)} \text{ for } LP}{\min_k \alpha^{(k)} \text{ for other heuristic}} \right)$).

instances.

Unsurprisingly, *PBC* leads to very bad results. In many situations (more than 35%), an application has been particularly unfavored and the fair experimental throughput was close to 0. The logarithm of the deviation for these situations has been normalized to 8. These poor results advocate the need for fairness guarantees in distributed computing environments like the ones we consider. As a matter of fact, this is somehow similar to the *Bushel of AppLeS* problem [34]: all applications identify the same resources as “best” for their tasks and seek to use them at the same time, thus making them no longer optimal or available.

Lastly, the geometrical average of *FCFS* is 1.564 and in the worst case, its performance is more than 8 times worse than *LP*. On the average, *FCFS* is therefore much worse than *LP*. On small platforms, the performances for *FCFS* and *CGBC* have the same order of magnitude. However, on larger ones (50 and 100 workers), *CGBC*

performs much better (geometrical average equal to 1.243) than *FCFS* (geometrical average equal to 2.0399).

VII. CONCLUSION

In this paper, we presented several heuristics for scheduling multiple Bag of Tasks applications on a tree-connected platform composed of heterogeneous processing and communication resources. Our contributions to this problem are the following:

- We outlined a theoretical scheduling model for Bag of Tasks applications scheduling environments with a particular emphasis on the fairness issue.
- We presented a centralized algorithm which, given the performance of all resources, computes an optimal schedule with respect to throughput maximization for tree networks. We also have characterized optimal solution on single-level trees.
- Since centralized algorithms may be unrealistic on large-scale platforms, we then presented sev-

eral distributed algorithms based on decentralized heuristics.

- We evaluated the efficacy of these heuristics using a wide range of realistic simulation scenarios, using limited buffer space. The results obtained by the most sophisticated heuristics are quite reasonable compared to the optimal centralized algorithm.

The experimental analysis provided in this article is an average-case analysis on small data-sets. It would be certainly be instructive to perform a worst-case analysis for the previous heuristics. In particular, we conjecture that it should be possible to establish a performance guarantee (an approximation factor) for *CGBC* and *FCFS*. However, we believe that *DATA-CENTRIC* and *PBC* are much harder to analyze.

We have seen with the *PBC* heuristic that non-cooperative approaches, where each application optimizes its own throughput, lead to a particularly unfair Nash equilibrium [35]. This Nash Equilibrium has been recently analytically studied in a multi-port context [36]. Extending this work to a single-port setting seems very challenging.

Another approach could be a cooperative approach where several decision makers (each of them being responsible for a given application) cooperate in making the decisions such that each of them will operate at its optimum. This situation can be modeled as a cooperative game like in [37] where this approach is successfully used for simpler problems such as bandwidth sharing in networks and another fairness criteria (proportional fairness). In particular, fully distributed protocols achieving proportionally fair allocations have been proposed. However in our situation, hierarchical resource sharing is rather hard to model, which renders such an approach quite challenging.

Acknowledgements

The authors would like to thank the reviewers for their detailed comments and suggestions which greatly helped us to improve the paper.

REFERENCES

- [1] "Berkeley Open Infrastructure for Network Computing," <http://boinc.berkeley.edu>.
- [2] O. Beaumont, L. Carter, J. Ferrante, A. Legrand, L. Marchal, and Y. Robert, "Scheduling multiple bags of tasks on heterogeneous master-worker platforms: centralized versus distributed solutions," LIP, Tech. Rep. 2005-45, Sept. 2005.
- [3] A. Legrand, L. Marchal, and H. Casanova, "Scheduling Distributed Applications: The SIMGRID Simulation Framework," in *Proceedings of the Third IEEE International Symposium on Cluster Computing and the Grid (CCGrid'03)*, May 2003.
- [4] SETI, URL: <http://setiathome.ssl.berkeley.edu>.
- [5] J. Cowie, B. Dodson, R.-M. Elkenbracht-Huizing, A. K. Lenstra, P. L. Montgomery, and J. Zayer, "A world wide number field sieve factoring record: on to 512 bits," in *Advances in Cryptology - Asiacrypt '96*, ser. LNCS, K. Kim and T. Matsumoto, Eds., vol. 1163. Springer Verlag, 1996, pp. 382–394.
- [6] Prime, URL: <http://www.mersenne.org>.
- [7] Entropia, URL: <http://www.entropia.com>.
- [8] M. Litzkow, M. Livny, and M. W. Mutka, "Condor - A hunter of idle workstations," in *Proceedings of the 8th International Conference of Distributed Computing Systems*. IEEE Computer Society Press, 1988, pp. 104–111.
- [9] F. Berman, R. Wolski, H. Casanova, W. Cirne, H. Dail, M. Faerman, S. Figueira, J. Hayes, G. Obertelli, J. Schopf, G. Shao, S. Smallen, N. Spring, A. Su, and D. Zagorodnov, "Adaptive computing on the grid using AppLeS," *IEEE Trans. Parallel Distributed Systems*, vol. 14, no. 4, pp. 369–382, 2003.
- [10] H. Casanova and F. Berman, "Grid'2002," in *Parameter sweeps on the grid with APST*, F. Berman, G. Fox, and T. Hey, Eds. Wiley, 2002.
- [11] J. Frey, T. Tannenbaum, I. Foster, M. Livny, , and S. Tuecke, "Condor-G: A Computation Management Agent for Multi-Institutional Grids," in *Proceedings of the Tenth IEEE Symposium on High Performance Distributed Computing (HPDC10)*. IEEE Computer Society Press, Aug. 2001.
- [12] W. Cirne, D. Paranhos, L. Costa, E. Santos-Neto, F. Brasileiro, J. Sauv e, F. A. B. da Silva, C. O. Barros, and C. Silveira, "Running Bag-of-Tasks Applications on Computational Grids: The MyGrid Approach," in *Proceedings of the ICCP'2003 - International Conference on Parallel Processing*, Oct. 2003.
- [13] M. R. Garey and D. S. Johnson, *Computers and Intractability, a Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
- [14] B. A. Shirazi, A. R. Hurson, and K. M. Kavi, *Scheduling and load balancing in parallel and distributed systems*. IEEE Computer Science Press, 1995.
- [15] P. Chr tienne, E. G. Coffman Jr., J. K. Lenstra, and Z. Liu, Eds., *Scheduling Theory and its Applications*. John Wiley and Sons, 1995.
- [16] D. Bertsimas and D. Gamarnik, "Asymptotically optimal algorithm for job shop scheduling and packet routing," *Journal of Algorithms*, vol. 33, no. 2, pp. 296–318, 1999.
- [17] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert, "Steady-state scheduling on heterogeneous clusters," *Int. J. of Foundations of Computer Science*, vol. 16, no. 2, pp. 163–194, 2005.
- [18] 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.
- [19] C. Banino, O. Beaumont, L. Carter, J. Ferrante, A. Legrand, and Y. Robert, "Scheduling strategies for master-slave tasking on heterogeneous processor platforms," *IEEE Trans. Parallel Distributed Systems*, vol. 15, no. 4, pp. 319–330, 2004.
- [20] B. Kreaseck, "Dynamic autonomous scheduling on heterogeneous systems," Ph.D. dissertation, University of California, San Diego, 2003.
- [21] L. Carter, H. Casanova, J. Ferrante, and B. Kreaseck, "Autonomous protocols for bandwidth-centric scheduling of independent-task applications," in *International Parallel and Distributed Processing Symposium IPDPS'2003*. IEEE Computer Society Press, 2003.
- [22] B. Hong and V. Prasanna, "Distributed adaptive task allocation in heterogeneous computing environments to maximize throughput," in *International Parallel and Distributed Processing Symposium IPDPS'2004*. IEEE Computer Society Press, 2004.

- [23] M. Mowbray, F. Brasileiroz, N. Andradez, J. Santanaz, and W. Cirnez, "Reciprocation-based economy for multiple services in peer-to-peer grids," in *Sixth IEEE International Conference on Peer-to-Peer Computing*. IEEE Computer Society Press, Sept. 2006.
- [24] T. Bonald and L. Massoulié, "Impact of fairness on internet performance," in *SIGMETRICS/Performance*, 2001, pp. 82–91.
- [25] L. Massoulié and J. Roberts, "Bandwidth sharing: Objectives and algorithms," *Transactions on Networking*, vol. 10, no. 3, pp. 320–328, June 2002.
- [26] P. Bhat, C. Raghavendra, and V. Prasanna, "Efficient collective communication in distributed heterogeneous systems," in *ICDCS'99 19th International Conference on Distributed Computing Systems*. IEEE Computer Society Press, 1999, pp. 15–24.
- [27] —, "Efficient collective communication in distributed heterogeneous systems," *Journal of Parallel and Distributed Computing*, vol. 63, pp. 251–263, 2003.
- [28] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert, "Pipelining broadcasts on heterogeneous platforms," *IEEE Trans. Parallel Distributed Systems*, vol. 16, no. 4, pp. 300–313, 2005.
- [29] O. Beaumont and L. Marchal, "Pipelining broadcasts on heterogeneous platforms under the one-port model," LIP, ENS Lyon, France, Research Report RR-2004-36, July 2004.
- [30] D. Bertsekas and R. Gallager, *Data Networks*. Prentice Hall, 1987.
- [31] O. Beaumont, A. Legrand, L. Marchal, and Y. Robert, "Independent and divisible tasks scheduling on heterogeneous star-shaped platforms with limited memory," in *PDP'2005, 13th Euromicro Workshop on Parallel, Distributed and Network-based Processing*. IEEE Computer Society Press, 2005, pp. 179–186.
- [32] O. Beaumont, V. Boudet, A. Petitot, F. Rastello, and Y. Robert, "A proposal for a heterogeneous cluster ScaLAPACK (dense linear solvers)," *IEEE Trans. Computers*, vol. 50, no. 10, pp. 1052–1070, 2001.
- [33] R. Jay, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. John Wiley and Sons, Inc., 1991.
- [34] F. Berman and R. Wolski, "TheAppLeS project: A status report," in *Proceedings of the 8th NEC Research Symposium, 1997*, available at <http://www.gcl.ucsd.edu/hetpubs.html#AppLeS>.
- [35] J. F. Nash, "Equilibrium points in n-person games," *Proceedings of the National Academy of Sciences USA*, vol. 36, pp. 48–49, 1950.
- [36] A. Legrand and C. Touati, "Non-cooperative scheduling of multiple bag-of-task applications," in *Proceedings of the 25th Conference on Computer Communications (INFOCOM'07)*, 2007.
- [37] H. Yaïche, R. R. Mazumdar, and C. Rosenberg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 667–678, 2000.



Olivier Beaumont received the PhD degree from the Université de Rennes in January 1999. He is currently an associate professor in the LaBRI laboratory in Bordeaux. His main research interests are parallel algorithms on distributed memory architectures.



Larry Carter received the AB degree from Dartmouth College in 1969 and the PhD degree in mathematics from the University of California at Berkeley in 1974. He worked at as a research staff member and manager at IBM's T.J. Watson Research Center for nearly 20 years in the areas of probabilistic algorithms, compilers, VLSI testing, and high-performance computation. Since 1994, Dr. Carter has been a professor in the Computer Science and Engineering Department of the University of California at San Diego. Between 1996 and 2000, he served as vice chair and then chair of the department. His current research interests include scientific computation, performance programming, parallel computation, and computer architecture. Prof. Carter is a senior fellow at the San Diego Supercomputing Center and a fellow of the IEEE.



Jeanne Ferrante received the PhD degree in mathematics from MIT in 1974. She was a research staff member at the IBM T.J. Watson Research Center from 1978 to 1994, and currently is a professor of computer science and associate dean of the Jacobs School of Engineering at the University of California, San Diego. Her work has included the development of intermediate representations for optimizing and parallelizing compilers, most notably the Program Dependence Graph and Static Single Assignment form. Her interests also include optimizing for parallelism and memory hierarchy and scheduling on large distributed platforms. She is a fellow of the ACM and a senior member of the IEEE.



Arnaud Legrand received the PhD degree from Ecole normale supérieure de Lyon in 2003. He is currently a CNRS researcher at the LIG Laboratory. He is mainly interested in parallel algorithm design for heterogeneous platforms and in scheduling techniques.



Loris Marchal received the PhD degree from Ecole normale supérieure de Lyon in 2006. He is currently a CNRS researcher at the LIP Laboratory at ENS Lyon. He is mainly interested in communication optimization on heterogeneous platforms and in scheduling techniques.



Yves Robert received the PhD degree from Institut National Polytechnique de Grenoble in January 1986. He is currently a full professor in the Computer Science Laboratory LIP at ENS Lyon. He is the author of four books, 80 papers published in international journals, and 100 papers published in international conferences. His main research interests are scheduling techniques and parallel algorithms for clusters and grids. He is a member of ACM and IEEE, and serves as an associate editor of *IEEE Transactions on Parallel and Distributed Systems*.