



Evaluation of Reallocation Heuristics for Moldable Tasks in Computational Dedicated and non Dedicated Grids

Ghislain Charrier, Yves Caniou, Frédéric Desprez

► To cite this version:

Ghislain Charrier, Yves Caniou, Frédéric Desprez. Evaluation of Reallocation Heuristics for Moldable Tasks in Computational Dedicated and non Dedicated Grids. [Research Report] RR-7365, INRIA. 2010, pp.38. inria-00509393

HAL Id: inria-00509393

<https://inria.hal.science/inria-00509393>

Submitted on 12 Aug 2010

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

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

***Evaluation of Reallocation Heuristics for Moldable
Tasks in Computational Dedicated and non
Dedicated Grids***

Yves Caniou — Ghislain Charrier — Frédéric Desprez

N° 7365

August 12, 2010

____ Domaine 3 ____

 ***apport
de recherche***

Evaluation of Reallocation Heuristics for Moldable Tasks in Computational Dedicated and non Dedicated Grids

Yves Caniou, Ghislain Charrier, Frédéric Desprez

Domaine : Réseaux, systèmes et services, calcul distribué
Équipe-Projet GRAAL

Rapport de recherche n° 7365 — August 12, 2010 — 35 pages

Abstract: Grid services often consist of remote sequential or rigid parallel application executions. However, moldable parallel applications, linear algebra solvers for example, are of great interest but requires dynamic tuning which has mostly to be done interactively if performances are needed. Thus, their grid execution depends on a *remote and transparent* submission to a possibly different batch scheduler on each site, and means an *automatic tuning* of the job according to the local load.

In this report we study the benefits of having a middleware able to automatically submit and reallocate requests from one site to another when it is also able to configure the services by tuning their number of processors and their walltime. In this context, we evaluate the benefits of such mechanisms on four multi-cluster Grid setups, where the platform is either composed of several heterogeneous or homogeneous, dedicated or non dedicated clusters. Different scenarios are explored using simulations of real cluster traces from different origins.

Results show that a simple scheduling heuristic is good and often the best. Indeed, it is faster and thus can take more jobs into account while having a small execution time. Moreover, users can expect more jobs finishing sooner and a gain on the average job response time between 10% and 40% in most cases if this reallocation mechanism combined to auto-tuning capabilities is implemented in a Grid framework. The implementation and

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

the maintenance of this heuristic coupled to the migration mechanism in a Grid middleware is also simpler because less transfers are involved.

Key-words: Reallocation, batch schedulers, meta-schedulers, computational grids, moldable tasks

Evaluation d'heuristiques de réallocation de tâches moldables dans les grilles de calculs dédiées et non dédiées

Résumé : L'appel à des services présents sur les grilles de calcul correspondent généralement à l'exécution d'une application séquentielle ou rigide. Cependant, il est possible d'avoir des applications parallèles moldables, telles que des solveurs linéaires, qui sont d'un grand intérêt, mais qui demandent une adaptation dynamique pour obtenir de bonnes performances. Leur exécution nécessite donc d'avoir un accès *distant et transparent* à différents gestionnaires de ressources, demandant donc une *adaptation automatique* de l'application en fonction de la charge locale.

Dans ce rapport, nous étudions les bénéfices découlant de l'utilisation d'un intergiciel de grille capable de soumettre et de réallouer des requêtes d'un site à l'autre tout en configurant automatiquement les services en choisissant le nombre de processeurs ainsi que la durée d'exécution estimée. Dans ce contexte, nous évaluons les gains apportés par de tels mécanismes sur quatre grilles de calcul différentes où la plate-forme est composée de plusieurs grappes, homogène ou hétérogènes, dédiées ou non. Nous explorons différents scénarios par la simulation de traces de tâches provenant de réelles exécutions.

Les résultats montrent que l'utilisation d'une heuristique d'ordonnancement simple est efficace, souvent amplement suffisante, voire la meilleure. En effet, elle est plus rapide à l'exécution et permet de prendre plus de requêtes en compte. Les utilisateurs peuvent espérer une majorité de requêtes terminant plus tôt si elle est utilisée, ainsi qu'une réduction du temps d'attente du résultat d'entre 10% et 40% dans la plupart des cas lorsque le mécanisme de réallocation couplé à l'adaptation automatique sont présents dans l'intergiciel. De plus, l'implantation et la maintenance de cette heuristique couplée au mécanisme de migration de tâches dans un intergiciel de grille est aussi plus facile car moins de transferts sont nécessaires.

Mots-clés : Réallocation, ordonnanceurs batch, méta-ordonnanceurs, grilles de calcul, tâches moldables

1 Introduction

In order to meet the evergrowing needs in computing capabilities of scientists of all horizons, new computing paradigms have been explored including the Grid. The Grid is the aggregation of heterogeneous computing resources connected through high speed wide area networks.

Computing resources are often parallel architectures managed by a local resource manager, called batch scheduler. In such a case, the local submission of a job requires at least a number of processors and a walltime. The walltime is the expected execution time for this job, given by the user or computed using data mining techniques. In most local resource management systems, when the walltime is reached, the job is killed, so users tend to over-evaluate the walltime to be sure that their job finishes its execution. Furthermore, giving an estimation of the execution time of a job is not an easy task and is influenced by the number of processors, which is generally chosen depending on external parameters such as the cluster load. Errors made at the local resource level may have a great impact on the global scheduling as shown in [BG09]. Errors can come from mistakes on the walltime as well as a burst of submission as shown in [SYIE09]. Thus, having a mechanism to accommodate bad scheduling decisions is important.

The context of this work, takes place in a multi-cluster Grid connected through a high bandwidth network. We propose a reallocation mechanism that takes into account scheduling errors by moving waiting jobs between clusters. The mechanism we propose can be used to connect different clusters together while each cluster keeps its local scheduling or resource allocation policies. Each job submitted onto the platform is executed automatically without any intervention from the user.

Two reallocation algorithms are studied with two heuristics each. We evaluate each couple (algorithm, heuristic) by comparing them on different metrics to an execution where reallocation is not performed. We extend the simulations realized in [CCD10b] by focusing on moldable tasks instead of parallel rigid tasks. The middleware is able to determine the number of processors and the walltime automatically for each task. Furthermore, we study the algorithms on dedicated platforms as well as non dedicated platforms. *We aim at showing the expectations in terms of performance with regard to the increased complexity of the jobs management done by the middleware.* We analyze the results on different metrics, and we show that obtained gains are very good in the majority of the simulations we perform. Gains are larger on dedicated platforms than on non dedicated platforms. We show that in most cases reallocating jobs will let jobs to finish sooner and diminish their average response time between 10% and 40%. Furthermore, results definitely confirm the counter intuitive fact that even for moldable jobs, whose number of processors varies if migrated, the simplest heuristic, both algorithmically and in implementation complexity, is the best to use.

The remainder of the report is as follows. In Section 2, we present related work. In Section 3 we describe mechanisms and the scheduling algorithms used in this work. Then we explain the experimental framework in Section 4, giving information about the simulator we developed, on the platforms simulated with real-world traces, scenarios of experiments that were conducted as well as the metrics on which results are compared in Section 5. Finally we conclude in Section 6.

2 Background

Parallel applications can be rigid, moldable or malleable [FRS⁺97]. A rigid application has a fixed number of processors. A moldable application can be executed with different number of processors, but once the execution started, this number can not change. Finally, the most permissive applications are malleable. The number of processors used can be modified “on the fly” during execution. This kind of applications is often an iterative application where the number of processors can be changed at the beginning of each iteration.

In [CB02] authors use moldable jobs to improve the performance in supercomputers. The user provides the scheduler *SA* with a set of possible requests that can be used to schedule a job. Such a request is represented by a number of processors and a walltime. *SA* chooses the request providing the earliest finish time. The evaluation of *SA* is done using real traces from the Parallel Workload Archive and their results show an average improvement on the response time of 44%, thus justifying the use of moldable jobs instead of rigid ones. In our work, we use the same kind of technique to choose the number of processors and the walltime of jobs. However, the user does not provide any information. The middleware is able to automatize everything thus facilitating the user’s actions and can choose to migrate jobs from on site to another one.

Another approach is given in [SR10]. ReSHAPE is a framework designed to execute malleable tasks, so it can also execute moldable tasks. It is a batch scheduler that can choose the number of processors of submitted tasks in order to minimize the queue length. This choice is done automatically. We also want to have a mechanism to choose the number of processors automatically on the clusters, but our mechanisms do not require any particular privilege on the different sites: we want to keep the underlying batch scheduler. Because ReSHAPE can change the number of processors in a queue once the job has been submitted, it can not provide a guaranty on when a job will finish in the worst case. Indeed, processors can be removed from a job so that other can have more, thus increasing the overall system efficiency. Batch schedulers using CBF can provide the completion time guaranty and we need it in our study.

Authors in [GC08] present a study of different meta-scheduling policies where each task uses its own meta-scheduler to be mapped on a parallel resource. Once submitted, a task is managed by the local scheduler and is never reallocated. In order to take advantage of the multi-site environment considered in our work, we use a central meta-scheduler to select a cluster for each incoming task because we place ourselves in the GridRPC context where clients do not know the computing resources. Also, once a task is submitted to the local scheduler, our approach let us cancel it and resubmit it elsewhere.

In order to migrate waiting jobs from one cluster to another, Yue presents the Grid-Backfilling in [Yue04]. Each cluster sends a snapshot of its state to a central scheduler at fixed intervals. Then the central scheduler tries to back-fill jobs in the queue of other clusters. The computation done by the central scheduler is enormous since it works with the Gantt chart of all sites. All clusters are homogeneous in power and size. In our work, the central scheduler is called upon arrival of each job in order to balance the load among clusters. During the reallocation phase, it gathers the list of all the waiting tasks and asks the local schedulers when a job would complete, but it does not perform complex computations. Furthermore, in our work, clusters are heterogeneous in size and power and we consider moldable jobs.

Author in [HSC09] present a study of the benefits of using moldable jobs in an heterogeneous computational grid. In this paper, the authors show that using a Grid meta-scheduler to choose on which site to execute a job coupled with local resource management schedulers able to cope with the moldability of jobs improves the average response time. In our work, instead of letting the local schedulers decide of the number of processors for a job, we keep existing infrastructure and software and we add a middleware layer that takes the moldability into account. Thus, our architecture can be deployed in existing Grids without modifications of the existing. Furthermore, this middleware layer renders reallocation between sites possible.

3 Task Reallocation

In this section, we describe the proposed tasks reallocation mechanism. First, we present the architecture of the Grid middleware (Section 3.1). Then we present the different algorithms used for the tasks reallocation (Section 3.2).

3.1 Architecture of the Middleware

We use a similar architecture than in [CCD10a], very close to the GridRPC [SLD⁺04] standard from the Open Grid Forum¹. Thus it can

¹<http://www.ogf.org>

be implemented in GridRPC compliant middleware such as DIET [CD06] or Ninf [SNS⁺97]. Because such a middleware is deployed on existing resources and has limited possibilities of action on the local resource managers, we developed a mechanism that only uses simple queries such as submission, cancellation, and estimation of the completion time.

The architecture relies on three main components: the **client** has computing requests to execute, and contacts an **agent** in order to obtain the reference of a **server** able to process the request. In our proposed architecture, one server is deployed on the front-end of each parallel resource, in which case it is in charge of interacting with the batch scheduler to perform the submission, cancellation or estimation of the completion date of a job. The server is also in charge of deciding how many processors should be used to execute the request, taking into account the load of the parallel resource. Benefiting from servers estimations, the *agent maps every incoming requests using a MCT strategy* (Minimum Completion Time [MAS⁺99]), and *decides of the reallocation with a second scheduling heuristic*.

3.2 Algorithms

This section presents the algorithm used to decide of the number of processors and walltime for each task (Section 3.2.1), the two versions of the reallocation mechanism (Section 3.2.2), and the scheduling heuristics used for reallocation (Section 3.2.3).

3.2.1 Tuning Parallel Jobs at Submission Time

The choice of the number of processors and walltime is done by the server each time a request arrives, either for the submission of a job or for an estimation of completion time. To determine the number of processors to allocate to the job, the server performs several estimations with different number of processors and returns the best size, *i.e.*, the one giving the earliest completion time. To estimate the completion time, the server can directly query the batch scheduler (but this capability is generally not present) or have it's own mechanism to compute the estimated completion time by simulating the batch algorithm for example.

The simplest idea to obtain the best size for the job is to perform an *exhaustive search*: For all possible number of processors (from one to the number of processors of the cluster), the estimation method provides a completion time as regard to the current load of the cluster. This method is simple and will choose the best size for jobs, however, it is time consuming. Indeed, each estimation is not instantaneous. Thus, for a large cluster, the estimation must be done a lot of times and the finding of the number of processors can require a long time.

In [SR10] the authors benchmark different sizes of the LU application from the NAS parallel benchmarks². Their study show a strictly increasing speedup up to 32 processors (adding processors always decreases execution time). But after this point, the execution time increases. It is due to the computation to communication ratio of the job becoming too small. This kind of job is not uncommon, thus we consider moldable jobs with *strictly increasing speedups until a known number of processors*.

Thus, in order to improve the speed in choosing the number of processors of a task, we can restrict the estimation from one processor to the limit of processors of the job. For jobs that don't scale very well, this will greatly reduce the number of calls to the estimation method thus reducing the time needed to find the most suitable number of processors.

Because of the hypothesis that speedup is strictly increasing until a maximum number of processors, we propose to perform a *binary search* on the number of processors to find how many of them to allocate to the job. Instead of estimating the completion time for each possible number of processors, we start by estimating the time for 1 processor and for the maximum number of processors. Then, we perform a classical binary search on the number of processors. This reduces the number of estimations from n to $\log_2 n$.

In particular cases the binary search will not provide the optimal result because of the back-filling. Let us consider an example in order to illustrate this behavior. Consider a cluster of 5 processors and a job needing 7 minutes to be executed on a single processor. With a perfect parallelism, this jobs needs 3.5 minutes to run on 2 processors, 2.33 on 3, 1.75 on 4 and 1.4 on 5. Upon submission, the cluster has the load represented by hatched rectangles in Figure 1. First, the binary search evaluates the completion time for the job on 1 and 5 processors (top of the figure) and obtains completion times of 7 and 7.4 minutes respectively. Then, the number of processors is set to 3 (middle of 1 and 5). The evaluation returns a completion time of 7.33 (bottom left of the figure). The most promising completion time was obtained with 1 processor, thus the binary search looks between 1 and 3. Finally, the best completion for the tested values time is obtained for 2 processors: 6.5 minutes (bottom right). However, the best possible completion time the job could have is 1.75 minutes with 4 processors. Indeed, with 4 processors, the jobs can start as soon as submitted, but this value was disregarded by the binary search. During our tests to verify the behavior of the binary search on thousands jobs, the results were the same as the exhaustive search which means that the "bad" cases are rare.

If the maximum number of processors of a job is large, using the binary search reduces enormously the number of estimations to do, potentially by orders of magnitude. For example, if a job can be executed on 650 processors

²<http://www.nas.nasa.gov/Resources/Software/npb.html>




Figure 1: Estimations made by the binary search.

the exhaustive search performs 650 estimations of completion time and the binary search performs only 10. The binary search in this case is thus 65 times faster.

3.2.2 Reallocation Algorithm

The first algorithm, **regular**, presented in Algorithm 1 works as follows: It gathers the list of all jobs in the waiting queues of all clusters; it selects a job with a scheduling heuristic; if it is possible to submit the job somewhere else with a better estimated completion time (ECT) of at least a minute, it submits it on the other cluster and cancels the job at its current location; finally, it starts again with the remaining jobs. The one minute threshold is here to consider some small data transfer that can take place, and to diminish the number of reallocations bringing almost no improvement.

To have a better idea of what is done, consider an example of two batch systems with different loads (see Figure 2). At time t , task f finishes before its walltime, thus releasing resources. Task j is then scheduled earlier by the local batch scheduler. When a reallocation event is triggered by the meta-scheduler at t_1 , it reallocates tasks h and i to the second batch system because their expected completion time is better there. To reallocate the tasks, each one is sequentially submitted to the second batch and canceled on the first one. In this example, the two clusters are identical so the tasks have the same execution time on both clusters, and the tuning of the parallel jobs (choice of number of processors to allocate to task h and i) is the same due to the same load condition. In an heterogeneous context, the length and

Algorithm 1 Reallocation algorithm.

```

 $l \leftarrow$  waiting jobs on all clusters
while  $l \neq \emptyset$  do
  Select a job  $j$  in  $l$ 
  Get a newECT for  $j$ 
  if  $j.newECT + 60 < j.currentECT$  then
    Cancel  $j$  on its current cluster
    Submit  $j$  to the new cluster
  end if
   $l = l \setminus \{j\}$ 
end while

```

even the number of processors allocated to the tasks would change between the clusters. Note that a task starting earlier on a cluster does not imply that it will also finish earlier.




Figure 2: Example of reallocation between two clusters.

The second algorithm, **all-cancellation**, detailed in Algorithm 2 starts by canceling all waiting jobs of all clusters. The agent keeps a reference for all jobs. Then, it selects a job with a scheduling heuristic. Finally, it submits the job to the cluster giving the minimum estimated completion time and loops on each of the remaining jobs.

Note that it does not mean that all parallel jobs will be tuned in the maximum of their performance since platforms are not necessarily dedicated to the Grid middleware, each cluster has its own load. It may be better to use less resources, thus have a longer execution time, but start earlier.

Algorithm 2 Reallocation algorithm with cancellation.

```

 $l \leftarrow$  waiting jobs on all clusters
Cancel each job in  $l$ 
while  $l \neq \emptyset$  do
    Select a job  $j$  in  $l$ 
    Submit  $j$  to the cluster with MCT
     $l = l \setminus \{j\}$ 
end while

```

The reallocation event in both versions of the algorithm is triggered periodically every hour, based on previous works conducted in [CCCD09] where a smaller period did not change the results but required more network transfers and potentially more reallocations.

Because both reallocation algorithm use an estimation of the completion time, it is mandatory that clusters use a batch scheduling algorithm able to give some guaranties on the completion time to guaranty the results. The two main algorithms offering these guaranties, presented in [FRS04], are First-Come-First-Served (FCFS) and Conservative Back-Filling (CBF). Both algorithm make reservations for each job and jobs can never be delayed once the reservation done. However, jobs can be scheduled earlier if new resources become available. Batch schedulers using one of these algorithms are common. Other algorithms such as Easy Back-Filling (EBF) introduced in [Lif95] or the well-known Shortest Job First (SJF) [FRS⁺97] do not guaranty a completion time and thus should not be used without adding specialized prediction mechanisms to the servers.

3.2.3 Scheduling Heuristics for Reallocation

We focus on two heuristics to use to select a job at each iteration. With the first one, jobs are processed in their submission order. In the remainder of the report, we refer to this policy as MCT because jobs are submitted in their original submission order and the jobs are submitted to the cluster with the Minimum Completion Time (MCT).

The second policy executes the MinMin heuristic on a subset of the jobs. MinMin asks the estimated completion time of all jobs and selects the job with the minimum of the returned values. In this report, MinMin is executed on the 20 oldest jobs. We use this limit to avoid a too long reallocation time. Indeed, MinMin has to update the estimations of completion times of all the remaining jobs at each iteration to select the job with the minimum of the ECTs. Because the all-cancellation algorithm needs to resubmit all jobs, it executes MinMin on the 20 oldest jobs and then the remaining jobs are processed in their original submission order, leading to a MCT policy.

We have two scheduling heuristics, MCT and MinMin, as well as two reallocation algorithms, namely regular and all-cancellation. Thus, we have four couples of algorithm that we refer in the remainder of this report as *MCT-reg*, *MCT-can*, *MinMin-reg*, and *MinMin-can*.

4 Experimental Framework

In this section we depict the experimental framework by presenting the simulator we implemented to run our experiments (Section 4.1), the description of the jobs (Section 4.2), the simulated platforms (Section 4.3), and the metrics used to compare the heuristics (Section 4.4). Finally, the experiments are described (Section 4.5).

4.1 Simulator

In order to simulate task reallocation in a distributed environment composed of several clusters, we use SimGrid [CLQ08], a discrete events simulation toolkit designed to simulate distributed environments, and Simbatch [CG09], a batch systems simulator built on top of SimGrid. Simbatch, which has been tested against real life experiments, can simulate the main algorithms used in batch schedulers described by [FRS04]. In this study, we use the *Conservative Back-Filling* (CBF) [MF01] algorithm for the batch schedulers. CBF tries to find a slot in the queue (Back-filling) where the job can fit without delaying already scheduled jobs (Conservative). If it does not, the job is added at the end of the queue. CBF is available in batch systems such as Maui [JSC01], Loadleveler [KRM⁺01], and OAR [CDCG⁺05] among others.

The simulator is divided using the same components as the ones in the GridRPC standard introduced in Section 3.1:

The *client* requests the system for a service execution. It contacts the meta-scheduler that will answer with the reference of a server providing the desired service.

The *meta-scheduler* matches incoming requests to a server according to a scheduling heuristic (we use MCT in this paper) and periodically reallocates jobs in waiting queues on the platform using one of the reallocation scheduling heuristic described in Section 3.2.3.

The *server* is running on the front-end of a cluster and interacts with the batch system. It receives requests from the client and can submit jobs to the batch scheduler to execute the requests. It can also cancel a waiting job, return an estimation of the completion time of a request and return the list of jobs in the waiting state. For submission and estimation, the server uses an estimation function that automatically chooses the number of processors and the walltime of the request using the technique described in Section 3.2.1.

4.2 Jobs

We built seven scenarios of jobs submission, where for six of them, jobs come from traces of different clusters on GRID’5000 for the first six months of 2008. Table 1 gives the number of jobs per month on each cluster. The seventh scenario is a six month long simulation using two traces from the parallel workload archive (CTC and SDSC) and the trace of Bordeaux on GRID’5000. The trace from Bordeaux contains 74647 jobs, CTC has 42873 jobs and SDSC contains 15615 jobs. Thus, there is a total of *133135 jobs*. In the remainder of the paper, we refer at the different scenarios by the name of the month of the trace for the jobs from GRID’5000 and we refer to the jobs coming from CTC, SDSC, and GRID’5000 as “PWA-G5K”.

Month/Cluster	Bordeaux	Lyon	Toulouse	Total
January	13084	583	488	14155
February	5822	2695	1123	9640
March	11673	8315	949	20937
April	33250	1330	1461	36041
May	6765	2179	1573	10517
June	4094	3540	1548	9182

Table 1: Number of jobs per month and in total for each site trace.

In our simulations, we do not consider advance reservations (present in GRID’5000 traces). They are considered as simple submissions so the batch scheduler can start them when it decides to. To evaluate the heuristics, we compare simulations together so this modification does not impact the results. However, we can not compare ourselves with what happened in reality. Furthermore, note that we add a meta-scheduler to map the jobs onto clusters at submission time, as if a grid middleware is used. On the real platform, users submit the cluster of their choice (usually they submit to the site closest to them) so the simulations already diverge from reality.

The traces taken from the Parallel Workload Archive were taken in their standard original format, *i.e.*, they also contain “bad” jobs described by [FT06]. We want to reproduce the execution of jobs on clusters, so we need to keep all the “bad” jobs removed in the clean version of the logs because these jobs were submitted in reality.

4.2.1 Moldable Jobs

The jobs contained in the trace files are parallel rigid jobs. So, in order to simulate the moldable jobs, we defined 4 types of jobs using Amdahl’s law ($speedup = \frac{1}{(1-P) + \frac{P}{N}}$ with P the fraction of parallel code and N the number of processors). The law states that the expected speedup of an application is strictly increasing but the increase rate diminishes. The execution time of

an application tends to the execution time of the sequential portion of the application when adding more processors.

To obtain the 4 types of moldable jobs, we vary the parallel portion of the jobs that is sequential as well as the limit of processors until the execution time decrease. The different values for the parallel portion of code are 0.8, 0.9, 0.99 and 0.999. Figure 3 plots the speedups for the different values of parallel code for different number of processors. Note that the y-axis is log-scaled. The figure shows that there is some point where the speedup increase becomes negligible. For the limits, we chose to use 32, 96, 256, and 650 processors. These values were chosen in accordance to the gain on the execution time of adding one processor. When the gain becomes really small, chances are that the internal communications of the job will take most of the time and slow down the task. Furthermore, the 650 limit is given by the size of the largest cluster of our simulations. So, the 4 types of jobs we consider are 4 couples (parallel portion, limit of processors): $t1:(0.8, 32)$, $t2:(0.9, 96)$, $t3:(0.99, 256)$ and $t4:(0.999, 650)$.




Figure 3: Speedups for the Amdahl's law for different parallelism portions.

In the traces, there are more tasks using a small number of processors than tasks using a lot of processors. Thus, each job from the trace files was given a moldable type. In each simulation we present, there are 50% of jobs of type $t1$, 30% of type $t2$, 15% of type $t3$ and 5% of type $t4$. The type of a job is chosen randomly. In order to keep a more realistic set of jobs, we decided to keep the sequential jobs of the traces sequential. All the characteristics of the jobs are summarized in Table 2.

Job type	Parallelism	Limit	Repartition
$t1$	0.8	32	50%
$t2$	0.9	96	30%
$t3$	0.99	256	15%
$t4$	0.999	650	5%

Table 2: The different kind of jobs.

4.2.2 Simulating Realistic Parallel Jobs Traces

During the simulations, the server uses information from both the traces and the type of the job to choose a suitable number of processors and a walltime for the job. In order to do so, the server uses the binary search described in Section 3.2.1 to choose a number of processors and follows the following process to choose the walltime: First, it computes the speedup of the job in the trace file using Amdahl's law, the type of the job and the number of processors: $spd = amdahl(p, n_t)$ with p the parallel portion of the code and n_t the number of processors used in the trace file. Second, the server computes the walltime of the job on one processor: $w_1 = w_{n_t} * spd$. Third, the server computes the speedup of the job for the current number of processors chosen by the binary search: $spd_b = amdahl(p, n_b)$. Finally, the server computes the walltime for the job: $w_b = \frac{w_1}{spd_b}$.

To obtain the actual execution time for the moldable jobs, we keep the same difference ratio as the one in the trace file. Thus if the runtime of a job was twice smaller than walltime in the trace file, it will also be twice smaller than the walltime in the simulations, independently of the number of processors chosen for the job. Furthermore, the runtime and walltime are modified *in accordance with the speed of the cluster* given in Section 4.3.1.

4.3 Platform Characteristics

4.3.1 Computing Resources

We consider two platforms with different numbers of cores distributed among three sites. Each platform is used in an homogeneous case (all clusters are similar in processor speed, but not in number of processors) and in an heterogeneous case (clusters differs in terms of CPU speed and number of processors). In all cases, the batch schedulers use a CBF policy to schedule the jobs.

The first platform corresponds to the simulation of three clusters of GRID'5000 [BCC⁺06]. The three clusters are Bordeaux, Lyon, and Toulouse. Bordeaux is composed of 640 cores and is the slowest cluster (if clusters are heterogeneous). Lyon has 270 cores and is 20% faster than Bordeaux (in the heterogeneous case). Finally, Toulouse has 434 cores and is 40% faster than Bordeaux (still in the heterogeneous case).

The second platform corresponds to experiments mixing the trace of Bordeaux from GRID'5000 and two traces from the Parallel Workload Archive³. The three clusters are Bordeaux, CTC, and SDSC. Bordeaux has 640 cores and is the slowest cluster. CTC has 430 cores and is 20% faster than Bordeaux. Finally, SDSC has 128 cores and is 40% faster than Bordeaux.

4.3.2 Dedicated Vs. Non Dedicated

On real life sites, tasks can be either submitted by a Grid middleware or by local users. Thus, we investigate the differences in behavior of our mechanism depending on heuristics: on **dedicated platforms**, where all tasks have been submitted through our middleware; on **non dedicated platforms** where two third of the jobs issued from the traces are directly submitted through batch schedulers by simulated local users. Both setups will be investigated on both homogeneous and heterogeneous platforms.

4.4 Evaluation metrics

In order to evaluate the reallocation algorithm and the behavior of the scheduling heuristics, we use different metrics. The first type of metrics is system centered metrics. The second type is user centered metrics.

- System metrics

Jobs impacted by reallocation: The percentage of jobs whose completion time is changed compared to an execution without reallocations. Only the jobs whose completion time changes are interesting in our study.

Number of reallocations relative to the total number of jobs:

We give the percentage of reallocations in comparison of the number of jobs. A job can be counted several times if it migrated several times so it is theoretically possible to have more than 100% reallocations. A small value is better because it means less transfers.

- User metrics

Jobs finishing earlier: Percentage of jobs that finished earlier with reallocation than without. This percentage is taken only from the jobs whose completion time changed with reallocation. A value higher than 50% means that there were more jobs finishing early than late.

³<http://www.cs.huji.ac.il/labs/parallel/workload/>

Gain on average job response time: Authors in [FR98] present the notion of response time. It corresponds to the duration between submission and completion. Complementary to the previous one, **the average job response time** of the jobs impacted by reallocation relatively to the scenario without reallocation defines the average ratio that the duration of a job can issue. A ratio of 0.8 means that on average, jobs spent 20% less time in the system, thus giving the results faster to the users.

Figure 4 illustrates why jobs can be delayed and others finishing earlier onto a platform composed of two clusters. At time 0 a reallocation event is triggered. A task is reallocated from cluster 2 to cluster 1 with a greater number of processors allocated to it according to our algorithm. Thus, some tasks of cluster 2 are advanced in the schedule. On cluster 1, as expected, the task is back-filled. However, assume the task finishing at time 6 finishes at time 2 because the walltime was wrongly defined (see the task with the dashed line). Thus, because of the newly inserted task, the large task on cluster 1 is delayed. Note that, even with FCFS, reallocation can also cause delay. If a job is sent to a cluster, all the jobs submitted after may be delayed. Inversely, the job that was reallocated to another cluster now leaves some free space and it may be used by other jobs to diminish their completion time.




Figure 4: Side effects of a reallocation.

4.5 Experiment

An experiment is a tuple (reallocation algorithm, heuristic, platform-trace, platform heterogeneity, dedicated, seed) where the seed is used to draw the type of a job in the trace, and concerning non dedicated platform, to draw if a job is submitted to the middleware or directly to the local scheduler. We used 10 different random seeds, hence, in addition to the reference experiment

using MCT without reallocation, we conducted $28+2*2*2*7*2*10$, *i.e.*, **1158 experiments** in total.

5 Results

This section presents the results for the four scenarios presented. First, we present the results on an homogeneous dedicated platform in Section 5.1. Then, we add heterogeneity to this platform in Section 5.2. Section 5.3 presents the results for a non dedicated homogeneous platform. Finally, the results for a non dedicated heterogeneous platform are presented in Section 5.4 before some concluding remarks on the results in Section 5.5.

Figures in this Section show the minimum, the maximum, the median, the lower, and higher quartiles, and the average of the 10 runs of each experience. Concerning the figures in a non dedicated environment, results only take into account the jobs we control. Jobs of external loads are not represented in the plots.

5.1 Homogeneous/Dedicated Platform

This section presents the results for an homogeneous dedicated platform. All clusters have the same speed (but a different number of processors), and all the jobs are managed by the middleware, *i.e.*, there is no external load on the clusters.

Figure 5 plots the percentage of jobs that were impacted when reallocation is performed on homogeneous dedicated platform. It shows that there are almost always more than half of the jobs impacted by reallocation. Except February, all experiments have an average close to 80% or 90%. These numbers show that most jobs are impacted. Most of the jobs that are not impacted are submitted when part of the platform is empty and can start execution as soon as they arrive so reallocation has no impact on them. Both MCT-reg and MinMin-reg have a very similar impact on the jobs. As for the all-cancellation versions of the algorithms, they impact more jobs than the regular versions on the three months where the impact is less than 80%. However, when the impact is around 90%, it is very close. This high number of impacted jobs comes from the moldability. Indeed, moldability increases the number of opportunities to move jobs between clusters or to resubmit a job on the same cluster with a different number of processors.

The number of reallocations relative to the total number of jobs performed in each experiment is plotted in Figure 6. The all-cancellation algorithms make a lot more reallocations than the regular algorithms. They always produce more than twice reallocations, and often a lot more. Because all jobs are canceled, there are more possibilities for jobs to change location during re-submission. MinMin-can produces the most reallocations but only a few more than MCT-can. Concerning the regular algorithms,




Figure 5: Jobs impacted on homogeneous dedicated platform.

MCT-reg performs more reallocations. Because of the limit on the number of jobs taken into account for MinMin-reg, there are less possibilities to re-allocate. The number of reallocations for MCT-reg corresponds to 17.34% of the total number of jobs on average. The maximum percentage attained in an experiment is 29.24%. With MCT-can, the average number of reallocation corresponds to 59.7% and the maximum is 136.93% (in the June experiment). Some jobs were moved several times, thus the values higher than 100%. The number of reallocations done by MinMin-reg correspond to 5.45% of the jobs on average and 11.13% maximum. MinMin-can has an average of 66.11% and the maximum is 148.99%. On the number of reallocations, using the regular algorithm is better because it limits the transfers between clusters.

The percentage of jobs finishing earlier when reallocation is performed on an homogeneous dedicated platform is shown in Figure 7. It only takes into account the jobs that were impacted when reallocation is performed. Thus a value higher than 50 means that more than half of the impacted jobs finished early. In almost all cases, reallocation brings an improvement. The all-cancellation versions of the algorithms are better than the regular algorithms (except in January for MCT-reg). The best result is given by MCT-can where there are 75% jobs early for only 25% late in March. MCT-can is the best most of the time, followed by MinMin-can. MCT-reg is third and finally, MinMin-reg gives the worst results.

Figure 8 plots the average response time of the impacted jobs relative to the runs without reallocation. With MCT-reg, the minimum gain is 10% and the maximum is a little more than 50% (response time divided by a factor of two). The MCT-can version of reallocation gives even better results, especially on the PWA-G5K trace. The MinMin-can heuristic is




Figure 6: Reallocations on homogeneous dedicated platform.




Figure 7: Percentage of jobs early on homogeneous dedicated platform.

close to MCT-can, but a little less good except in January. The MinMin-reg heuristic is the worst. Indeed, it is always the less good heuristic and it gives negative results in several experiments. It confirms the results of the percentage of jobs early on this trace. These results show that when reallocation is performed, we can expect a gain between 20% and 40% on the average response time of the jobs (except with MinMin-reg).




Figure 8: Relative average response time on homogeneous dedicated platform.

5.2 Heterogeneous/Dedicated Platform

In this section, clusters are heterogeneous in number of processors and in speed (cf. Section 4.3). All requests are done to our Grid middleware, thus there are no local jobs submitted.

The percentage of jobs impacted is shown in Figure 9. In six experiments for the two traces March and June, extreme cases where almost 100% of the jobs that were impacted by reallocation appear. This happens when the platform has a few phases with no job. If there are always jobs waiting, the reallocation is able to move jobs more often thus impacting a bigger portion of the jobs. Apart from these cases, the number of jobs impacted varies between the traces from 25 to 95%. All-cancellation algorithms usually impacts more jobs. MinMin-can impacts more jobs on average than the other heuristics. MCT-reg and MinMin-reg have close results. There are less jobs impacted by reallocation of heterogeneous platforms than on the homogeneous one because some cluster are faster, so the global speed increased. Waiting queues are thus smaller.




Figure 9: Jobs impacted on heterogeneous dedicated platform.

The number of reallocations relative to the total number of jobs is plotted in Figure 10. The pattern in this figure is the same as the one on the homogeneous platform but with less reallocations. All-cancellation algorithms always produce more reallocations. The regular algorithms give results inferior than 15% so the number of reallocations is quite small compared to the total number of jobs. However, with the all-cancellation algorithms, it is possible to go to a value as high as 50%. Because all-cancellation empties the waiting queues, more jobs have the opportunity to be reallocated. With the regular algorithms, jobs close to execution have a very small chance of being reallocated. The regular version of the reallocation algorithm is better on this metric. Compared to the experiment on homogeneous platform, the

number of reallocations is smaller by a factor close to 2. Because the overall execution of the jobs is faster, there are less possibilities to reallocate jobs.




Figure 10: Reallocations on heterogeneous dedicated platform.

Figure 11 plots the percentage of jobs early on heterogeneous platform. In this case, 3 experiments produce more jobs late than early. In April without all-cancellation there are always more jobs late (less than 4%) when reallocation is performed. However in most cases, it is better to reallocate. MinMin-reg gives the worst results. It is followed by MCT-reg, then MinMin-can and finally MCT-can is the best with up to 64% of tasks early! This ranking is the same as the one on homogeneous platform.




Figure 11: Percentage of jobs early on heterogeneous dedicated platform.

Concerning the average relative response time, the plot in Figure 12 shows a clear improvement in most cases. Excluding MinMin-reg, most

gains are comprised between 10% and 40%. On average, MCT-can is the best heuristic. The reallocation without all-cancellation can worsen the average response time. It happened in 6 experiments (3 with MCT-reg and 3 with MinMin-reg). The loss is small for MCT-reg (less than 5%) thus it is not a problem. The all-cancellation versions are always better than their corresponding regular algorithm except in February for MCT-reg. Some experiments present a gain on the average response time while there were more jobs late than early (MCT-reg in April for example): The gains were high enough to compensate for the late jobs.




Figure 12: Relative average response time on heterogeneous dedicated platform.

5.3 Homogeneous/Non Dedicated Platform

In this section, we present the results where the platform is homogeneous but non dedicated. More precisely, 33% of the jobs are selected randomly as moldable jobs managed by the middleware. The remaining jobs are used as external load (for the grid middleware) on the clusters. These jobs are submitted as they were in the traces and the middleware can not act on them.

Figure 13 plots the percentage of the moldable jobs that are impacted by reallocation. The external load jobs on the clusters are not taken into account. The percentage of impacted jobs usually is below 60%. Because of the external load, the queues are more static (some jobs are in the queue and cannot be moved to another cluster) than on a dedicated platform, thus reallocations are harder to make. Consequently, the percentage of jobs impacted is lower than on a dedicated platform. In most experiments, the impacted jobs vary very little. However, there are extreme cases where the

impacted jobs are almost 100% with MinMin-reg. The difference of results between the heuristics is often very small.




Figure 13: Jobs impacted on homogeneous non-dedicated platform.

The number of reallocations relative to the total number of jobs is shown in Figure 14. As for the previous platforms, the all-cancellation algorithms reallocate more jobs. However, the difference with the regular algorithms is smaller. Both algorithms using all-cancellation give similar results. The MinMin-reg part at the beginning of the reallocations for MinMin-can has almost no influence here. Due to the limit of jobs to take into account, MinMin-reg is the heuristic that makes the less reallocations. Indeed, it is either almost equal to MCT-reg or a little lower. MCT-reg makes a number of reallocations corresponding to 2.49% of the moldable jobs with a maximum value of 9.51%. For the all-cancellation version, the number of reallocations corresponds to 5.86% of the moldable jobs on average and has a maximum of 22.45%. Concerning MinMin-reg it has an average of 1.89% and the maximum is 5.89%. Finally, MinMin-can is the algorithm with the highest values with an average of 6.18% and a maximum of 23.08%. The external load on the clusters prevents a lot of reallocations, thus their small number.

Figure 15 presents the percentage of jobs early on homogeneous non dedicated platform. Only the impacted moldable jobs are taken into account. In the most favorable case there are more than 90% of the jobs that finish earlier with reallocation. However in March, two experiments executed with algorithms based on MinMin produce more than 80% of jobs late in March. Without the extreme cases, the percentage of jobs early are between 50% and 80%, with a mean higher than 65%. So, on this platform, it is good to reallocate. Algorithms using all-cancellation are behind the regular versions by a few percents. The all-cancellation algorithms gives better results than




Figure 14: Reallocations on homogeneous non-dedicated platform.

expected. Indeed, by canceling all the moldable jobs, the external jobs in the queues are scheduled earlier. However, the algorithms can compensate this. On traces where the external loads are really high, this would not be true. Indeed, without free space, the jobs would be cancelled and resubmitted behind the jobs of the external load. This behavior could cause starvation until free nodes are available.




Figure 15: Percentage of jobs early on homogeneous non-dedicated platform.

The average response times plotted in Figure 16 have most of the time an improvement with reallocation. The average response time is worse with reallocation than without only with the June trace and only by a few percents. In the other cases, the gains are usually between 5% and 50%. The

best gain obtained provides an improvement of more than a factor two. The gains obtained are good for the few number of reallocation performed as were presented in Figure 14. In this non dedicated context, the all-cancellation algorithms work as well as the regular version of the algorithm. Canceling the jobs enable the jobs from the external load to jump in front of the moldable jobs. The all-cancellation versions performs more reallocations, so it is preferable to use the regular version. Furthermore, the all-cancellation algorithms can cause starvation by letting jobs from the external jump in front of the moldable jobs. There is no notable difference between MinMin-reg and MCT-reg, thus MCT-reg should be used. It has a lower complexity and takes all the jobs into account which gives it a better potential to make better decisions.




Figure 16: Relative average response time on homogeneous non-dedicated platform.

5.4 Heterogeneous/Non Dedicated Platform

In this section, we present the results on a non dedicated heterogeneous (in processors number as well as speed) platform where 33% of the jobs are moldable and submitted to the middleware.

The percentage of jobs impacted by reallocation is plotted in Figure 17. The two all-cancellation heuristics impact more jobs than the regular ones, but the difference between them is really small. There is one experiment in March where MinMin-reg impacts almost all jobs: a scheduling decision taken at the beginning of the experiment impacts all the following job completion dates. For a given trace, the number of impacted jobs usually does not vary a lot.




Figure 17: Jobs impacted on heterogeneous non-dedicated platform.

Figure 18 plots the number of reallocations relative to the total number of moldable jobs. The number of reallocations is very small. In most cases, there are only a few dozens reallocations. The all-cancellation algorithms always reallocates more than the regular versions, *but not by far*. In a lot of cases, the number of reallocations corresponds to less than 1% of the number of jobs. Thus, on a non dedicated platform, the reallocation mechanism does not produce many transfers.

Most experiments except the worst case for PWA-G5K and March with MinMin-reg result *in more than half of the jobs early* as we can see in Figure 19. The 90% jobs late in March with MinMin-reg are from the same experiment where almost all jobs were impacted in Figure 17. Most experiments exhibit a percentage of jobs early close to 70%. All-cancellation again produces less jobs early than regular. MCT-reg and MinMin-reg are the two heuristics of choice, but MinMin-reg gives mitigate results for PWA-G5K so MCT-reg is a better choice.




Figure 18: Reallocations on heterogeneous non-dedicated platform.




Figure 19: Percentage of jobs early on heterogeneous non-dedicated platform.

Figure 20 shows that the different heuristics give results close to one another on the relative average response time. All-cancellation heuristics usually have a smaller difference between the minimum and the maximum gains. Depending on the experiment, results vary a lot. In some experiments, the average response time is divided by more than two, but in other it is augmented with a maximum of 40%. However on all experiments, the average gain is positive. Thus reallocation is expected to provide a gain.




Figure 20: Relative average response time on heterogeneous non-dedicated platform.

5.5 Remarks on Results

Patterns in some of the results in the previous sections keep repeating on the different platforms. Indeed, MCT-reg and MinMin-reg usually give similar results on non dedicated platforms, often in favor of MCT-reg. On dedicated platforms however, MCT-reg is clearly better than MinMin-reg. MinMin-reg may give better results if it is able to take more jobs into account during reallocation. But, if it takes more jobs into account, its execution time grows exponentially. Furthermore, the two algorithms with all-cancellation also give very similar results with a small advantage for the MCT-reg on dedicated platforms.

The all-cancellation algorithms can cause starvation. In a non dedicated platform, it is obvious that starvation can happen. Indeed, when canceling jobs, jobs from the external load (for the Grid middleware) will be rescheduled in front of the moldable jobs managed by the middleware system. This may explain the worst cases peaks in Figure 20. Even in a dedicated environment with MinMin-can, it is possible for a job to be delayed indefinitely. If the job is long, it will always be resubmitted after others and may never start execution. However, such cases did not happen in our simulations because there are always phases of low load where the queues can be emptied.

The results presented in this paper show that the heuristic of choice is MCT with or without all-cancellation whether the platform is dedicated or not. Indeed, MinMin is too complex in time to react in a decent time regarding the submission rate of jobs onto the platform. In a previous study [CCD10b], we used several other selection heuristics such as MaxMin, Sufferage, MaxGain, and MaxRelGain but these heuristic did not prove better than MCT or MinMin. Because these algorithms have the same complex-

ity than MinMin, we argue that they may also give poor results, especially because of worst cases.

6 Conclusion and Perspectives

In this report, we presented a reallocation mechanism that can be implemented in a GridRPC middleware and used on any multi-cluster Grid without modifying the underlying infrastructure. Parallel jobs are tuned by the Grid middleware each time they are submitted to the local resource manager (which implies also each time a job is migrated). We achieve this goal by only querying batch schedulers with simple submission or cancellation requests. Users ask the middleware to execute some service and the middleware manages the job automatically.

We have investigated two reallocations algorithms, the key difference between them being that one, regular, cancels a task once it is sure that the expected completion time is better on another cluster, and the other, all-cancellation, cancels all waiting jobs before testing reallocation. We also considered two scheduling heuristics to make the decision of migrating a job to another site. We conducted 1158 experiments and analyzed them on 4 different metrics.

On dedicated clusters, the cancellation of all the waiting jobs proves to be very efficient to improve the average job response time. On the other hand in a non dedicated environment, the algorithm that does not cancel waiting jobs behaves better. On both platforms, surprisingly, there is not a great number of migrating tasks, but all tasks take benefit those migration since the percentage of impacted tasks is high. In term of performances, users can expect more jobs finishing sooner, and an improvement of the jobs response time from a few percents to more than 50%! Only a few cases give bad results leading to an increase of the average job response time.

The next step of this work is the implementation of the reallocation mechanism in the DIET GridRPC middleware. DIET already provides most of the needed features. The missing features are the cancellation of a job in batch schedulers (numerous are supported) which is easy to implement and the reallocation mechanism itself. This last point should be quite straightforward because all communications are already handled by the middleware. We intend to implement both reallocation mechanisms with MCT. Indeed, we need the regular algorithm to work on non dedicated platforms. We plan also to implement the all-cancellation mechanism because DIET can be used in a dedicated environment. Furthermore, we could use this in the SPADES⁴ project where we plan to maintain a set of reserved resources on a site which are managed by our own embedded batch scheduler.

⁴ANRProject08-ANR-SEGI-025

References

- [BCC⁺06] Raphaël Bolze, Franck Cappello, Eddy Caron, Michel Daydé, Frederic Desprez, Emmanuel Jeannot, Yvon Jégou, Stéphane Lanteri, Julien Leduc, Noredine Melab, Guillaume Mornet, Raymond Namyst, Pascale Primet, Benjamin Quetier, Olivier Richard, El-Ghazali Talbi, and Irena Touché. Grid'5000: A Large Scale and Highly Reconfigurable Experimental Grid Testbed. *International Journal of High Performance Computing Applications*, 20(4):481–494, 2006.
- [BG09] Marta Beltrán and Antonio Guzmán. The Impact of Workload Variability on Load Balancing Algorithms. *Scalable Computing: Practice and Experience*, 10(2):131–146, June 2009.
- [CB02] Walfredo Cirne and Francine Berman. Using Moldability to Improve the Performance of Supercomputer Jobs. *Journal of Parallel and Distributed Computing*, 62(10):1571–1601, 2002.
- [CCCD09] Yves Caniou, Eddy Caron, Ghislain Charrier, and Frédéric Desprez. Meta-Scheduling and Task Reallocation in a Grid Environment. In *The Third International Conference on Advanced Engineering Computing and Applications in Sciences (ADV-COMP'09)*, pages 181–186, Sliema, Malta, October 11-16 2009.
- [CCD10a] Yves Caniou, Ghislain Charrier, and Frédéric Desprez. Analysis of Tasks Reallocation in a Dedicated Grid Environment. Technical Report RR-7226, Institut National de Recherche en Informatique et en Automatique (INRIA), March 2010.
- [CCD10b] Yves Caniou, Ghislain Charrier, and Frédéric Desprez. Analysis of Tasks Reallocation in a Dedicated Grid Environment. In *IEEE International Conference on Cluster Computing 2010 (Cluster 2010)*, Heraklion, Crete, Greece, September 20-24 2010. To appear.
- [CD06] Eddy Caron and Frédéric Desprez. DIET: A Scalable Toolbox to Build Network Enabled Servers on the Grid. *International Journal of High Performance Computing Applications*, 20(3):335–352, 2006.
- [CDCG⁺05] Nicolas Capit, Georges Da Costa, Yiannis Georgiou, Guillaume Huard, Cyrille Martin, Grégory Mounié, Pierre Neyron, and Olivier Richard. A Batch Scheduler with High Level Components. *CoRR*, abs/cs/0506006:9, 2005.

- [CG09] Yves Caniou and Jean-Sébastien Gay. Simbatch: An API for simulating and predicting the performance of parallel resources managed by batch systems. In *Workshop on Secure, Trusted, Manageable and Controllable Grid Services (SGS), held in conjunction with EuroPar'08*, volume 5415 of *LNCS*, pages 217–228, 2009.
- [CLQ08] Henri Casanova, Arnaud Legrand, and Martin Quinson. Sim-Grid: a Generic Framework for Large-Scale Distributed Experiments. In *10th IEEE International Conference on Computer Modeling and Simulation*, March 2008.
- [FR98] Dror G. Feitelson and Larry Rudolph. Metrics and Benchmarking for Parallel Job Scheduling. In *Proceedings of the Workshop on Job Scheduling Strategies for Parallel Processing (IPPS/JSSPP '98)*, pages 1–24, London, UK, 1998. Springer-Verlag.
- [FRS⁺97] Dror G. Feitelson, Larry Rudolph, Uwe Schwiegelshohn, Kenneth C. Sevcik, and Parkson Wong. Theory and Practice in Parallel Job Scheduling. In *IPPS '97: Proceedings of the Job Scheduling Strategies for Parallel Processing*, pages 1–34, London, UK, 1997. Springer-Verlag.
- [FRS04] Dror G. Feitelson, Larry Rudolph, and Uwe Schwiegelshohn. Parallel job scheduling - a status report. In *Job Scheduling Strategies for Parallel Processing*, 2004.
- [FT06] Dror G. Feitelson and Dan Tsafrir. Workload Sanitation for Performance Evaluation. In *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 221–230, Austin, Texas, March 2006.
- [GC08] Francesc Guim and Julita Corbalán. A Job Self-scheduling Policy for HPC Infrastructures. In *Job Scheduling Strategies for Parallel Processing (JSSPP)*, 2008.
- [HSC09] Kuo-Chan Huang, Po-Chi Shih, and Yeh-Ching Chung. Adaptive Processor Allocation for Moldable Jobs in Computational Grid. *International Journal of Grid and High Performance Computing*, 1(1):10–21, January-March 2009.
- [JSC01] David Jackson, Quinn Snell, and Mark Clement. Core Algorithms of the Maui Scheduler. In *JSSPP '01: Revised Papers from the 7th International Workshop on Job Scheduling Strategies for Parallel Processing*, pages 87–102, London, UK, 2001. Springer-Verlag.

-
- [KRM⁺01] Subramanian Kannan, Mark Roberts, Peter Mayes, Dave Brelsford, and Joseph Skovira. *Workload Management with LoadLeveler*. IBM Press, 2001.
- [Lif95] David A. Lifka. The ANL/IBM SP scheduling system. In *In Job Scheduling Strategies for Parallel Processing*, pages 295–303. Springer-Verlag, 1995.
- [MAS⁺99] Muthucumaru Maheswaran, Shoukat Ali, Howard Jay Siegel, Debra Hensgen, and Richard F. Freund. Dynamic Matching and Scheduling of a Class of Independent Tasks onto Heterogeneous Computing Systems. *Journal of Parallel and Distributed Computing*, 59:107–131, 1999.
- [MF01] Ahuva W. Mu’alem and Dror G. Feitelson. Utilization, Predictability, Workloads, and User Runtime Estimates in Scheduling the IBM SP2 with Backfilling. *IEEE Transactions on Parallel and Distributed Systems*, 12(6):529–543, 2001.
- [SLD⁺04] K. Seymour, C. Lee, F. Desprez, H. Nakada, and Y. Tanaka. The End-User and Middleware APIs for GridRPC. In *Workshop on Grid Application Programming Interfaces, In conjunction with GGF12*, Brussels, Belgium, September 2004.
- [SNS⁺97] Mitsuhsa Sato, Hidemoto Nakada, Satoshi Sekiguchi, Satoshi Matsuoka, Umpei Nagashima, and Hiromitsu Takagi. Ninf: A Network Based Information Library for Global World-Wide Computing Infrastructure. In *HPCN Europe*, pages 491–502, 1997.
- [SR10] Rajesh Sudarsan and Calvin J. Ribbens. Design and performance of a scheduling framework for resizable parallel applications. *Parallel Computing*, 36:48–64, 2010.
- [SYIE09] Ozan Sonmez, Nezih Yigitbasi, Alexandru Iosup, and Dick Epema. Trace-Based Evaluation of Job Runtime and Queue Wait Time Predictions in Grids. In *International Symposium on High Performance Distributed Computing (HPDC’09)*, June 11-13 2009.
- [Yue04] Jianhui Yue. Global Backfilling Scheduling in Multiclusters. *Lecture notes in Computer Science*, 3285:232–239, 2004.



Centre de recherche INRIA Grenoble – Rhône-Alpes
655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

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