



An RMS for Non-predictably Evolving Applications

Cristian Klein, Christian Pérez

► **To cite this version:**

Cristian Klein, Christian Pérez. An RMS for Non-predictably Evolving Applications. IEEE International Conference on Cluster Computing, Sep 2011, Austin, Texas, United States. 2011. <inria-00606908>

HAL Id: inria-00606908

<https://hal.inria.fr/inria-00606908>

Submitted on 7 Jul 2011

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.

An RMS for Non-predictably Evolving Applications

Cristian KLEIN, Christian PÉREZ

INRIA / LIP, ENS de Lyon
46 allée d’Italie, 69364 Lyon Cedex 7, France
{cristian.klein,christian.perez}@inria.fr

Abstract—Non-predictably evolving applications are applications that change their resource requirements during execution. These applications exist, for example, as a result of using adaptive numeric methods, such as adaptive mesh refinement and adaptive particle methods. Increasing interest is being shown to have such applications acquire resources on the fly. However, current HPC Resource Management Systems (RMSs) only allow a static allocation of resources, which cannot be changed after it started. Therefore, non-predictably evolving applications cannot make efficient use of HPC resources, being forced to make an allocation based on their maximum expected requirements.

This paper presents CooRMv2, an RMS which supports efficient scheduling of non-predictably evolving applications. An application can make “pre-allocations” to specify its peak resource usage. The application can then dynamically allocate resources as long as the pre-allocation is not outgrown. Resources which are pre-allocated but not used, can be filled by other applications. Results show that the approach is feasible and leads to a more efficient resource usage.

I. INTRODUCTION

High-Performance Computing (HPC) resources, such as clusters and supercomputers, are managed by a Resource Management System (RMS) which is responsible for multiplexing resources among multiple users. Generally, computing nodes are space-shared, which means that a user gets an exclusive access. Requesting a resource allocation is mostly done using *rigid* jobs: the user requests a number of nodes for a limited duration, then the RMS decides when the allocation is served.

Increased interest has been devoted to giving more flexibility to the RMS, as it has been shown to improve resource utilization [1]. If the RMS can reshape a request, but only before the allocation starts, the job is called *modal*. Otherwise, if the RMS can change the allocation while it is served, the job is called *malleable*. Taking advantage of modal jobs often comes for free, as commonly used programming models (e.g., MPI) are already suitable for creating modal applications. However, writing malleable applications is somewhat more challenging as the application has to be able to reconfigure itself during execution. How to write malleable applications [2], [3] and how to add RMS support for them [4], [5], [6] has been extensively studied.

Improving utilization can be done further by supporting *evolving* applications. These are applications which change

their resource requirements during execution as imposed by their internal state. Unlike malleability, the change of the resource allocation is not imposed by the RMS, but by the nature of the computations. Being able to provide applications with more nodes and memory on the fly is considered necessary for achieving exascale computing [7].

We define three types of evolving applications: (i) **fully-predictable**, whose evolution is known at start (e.g., static workflows), (ii) **marginally-predictable** (e.g., non-DAG workflows containing branches), whose evolution can only be predicted within a limited horizon of prediction, and (iii) **non-predictable**, whose evolution cannot be known in advance. Non-predictably evolving applications (NEA) are more-and-more widespread, due to the increased usage of adaptive numeric methods, such as Adaptive Mesh Refinement (AMR) [8] and Adaptive Particle Methods (APM) [9].

Unfortunately, support for evolving applications is insufficient in current RMSs. In batch schedulers, evolving applications are forced from the beginning to allocate enough resources to fulfill their peak requirements. In Clouds, which support dynamic provisioning, large-scale resource requests may be refused with an “out-of-capacity” error, thus, allocating resources on the fly may lead to the application reaching an out of memory condition.

This paper proposes CooRMv2 an RMS that enables efficient scheduling of evolving applications, especially non-predictable ones. It allows applications to inform the RMS about their maximum expected resource usage (using **pre-allocations**). Resources which are pre-allocated but unused can be filled by malleable applications.

The main contribution of the paper is three-fold: (i) to devise a model for NEAs based on AMR applications; (ii) to propose CooRMv2 an extension to our previous RMS, for efficiently supporting such applications; (iii) to show that the approach is feasible and that it can lead to considerable gains.

The remaining of this paper is organized as follows. Section II motivates the work and refines the problem statement by presenting a model of an AMR application. Section III introduces CooRMv2, a novel RMS which efficiently supports various types of applications as shown in Section IV. Section V evaluates the approach and highlights the gains that can be made. Section VI discusses related work. Finally, Section VII concludes the paper and opens up perspectives.

II. A MODEL FOR NON-PREDICTABLY EVOLVING APPLICATIONS

This section presents a model of a non-predictably evolving application based on observations of AMR applications. First, the evolution of the working set size and the speed-up of the application are modeled. Second, the model is analyzed and the problem statement is refined.

A. Working Set Evolution Model

Our goal is to devise a model for simulating how the data size evolves during an AMR computation. In order to have a model which is parametrizable, we first devise a “normalized” evolution profile, which is independent of the maximum size of the data and the structure of the computations. The model has to be simple enough to be easy to analyze and use, but at the same time it should resemble observed behaviors. To this end, we consider that the application is composed of 1000 steps. During step i the size of the data s_i is considered constant ($s_i \in [0, 1000]$).

Unfortunately, we have been able to find only a few papers on the evolution of the refined mesh. Most of the papers focus either on providing scalable implementations [10] or showing that applying AMR to a certain problem does not reduce its precision as compared to a uniform mesh [8]. From the few papers we found [11], [12], we extracted the main features of the evolution: (i) it is mostly increasing; (ii) it features regions of sudden increase and regions of constancy; (iii) it features some noise.

Using the above constraints, we derive the following model, that we call the **acceleration-deceleration model**. First, let the evolution of the size of the mesh be characterized by a velocity v_i , so that $s_i = s_{i-1} + v_i$. The application is assumed to be composed of multiple *phases*. Each phase consists of a number of steps uniformly distributed in [1, 200]. Each step i of an even phase increases v_i ($v_i = v_{i-1} + 0.01$), while each step i of an odd phase decreases v_i ($v_i = v_{i-1} \times 0.95$). Next, a Gaussian noise of $\mu = 0$ and $\sigma = 2$ is added. The above values have been chosen so that the profiles resemble the ones found in the cited articles. At the end, the profile is normalized, so that the maximum of the series s_i is 1000.

To obtain the actual, non-normalized data size profile S_i at step i , given the maximum data size S_{\max} , one can use the following formula: $S_i = s_i \times S_{\max}$. Figure 1 shows some examples of data size evolutions returned by the model, which are compatible with those presented in [11], [12].

B. A Speed-up Model

The next issue is to find a function $t(n, S)$ that returns the duration of a step, as a function of the number of allocated nodes n and the size of the data S . Both n and S are assumed constant during a step.

Instead of having a precise model (such as in [13]) which would be applicable to only one AMR application, we aimed at finding a simple formula, which could be fit




Fig. 1. Examples of obtained AMR working set evolutions.




Fig. 2. Fitting the AMR speed-up model against the actual data for different mesh sizes.

against existing data. To this end, we chose to use the data presented in [14] as it shows both strong and weak scalability of an AMR application for data sizes which stay constant during one experiment. After trying several functions and verifying how well they could be fit, we found a good one to be:

$$t(n, S) = A \times \frac{S}{n} + B \times n + C \times S + D$$

Figure 2 presents the result of logarithmically fitting this formula against actual data. The model fits within an error of less than 15% for any data point. The values of the parameters, used in the rest of the paper, are: $A = 7.26 \times 10^{-3} \text{ s} \times \text{node} \times \text{MiB}^{-1}$, $B = 1.23 \times 10^{-4} \text{ s} \times \text{node}^{-1}$, $C = 1.13 \times 10^{-6} \text{ s} \times \text{MiB}^{-1}$, $D = 1.38 \text{ s}$, $S_{\max} = 3.16 \text{ TiB}$.

C. Analysis of the Model

Let us analyze the model and refine the problem statement. Users want to adapt their resource allocation to a target criterion. For example, depending on the computational budget of a scientific project, an application has to run at a given **target efficiency**. This allows scientists to receive the results in a timely fashion, but at the same time control the amount of used resources.

To run the above modeled application at a target efficiency e_t throughout its whole execution, given the evolution of the data size S_i , one can compute the number of nodes n_i that have to be allocated during each step i of the computation. Let us denote the consumed resource area (node-count \times duration) in this case $A(e_t)$. In practice, while executing the application, one does not need any a priori knowledge of the size of the data, as n_i only depends on the current S_i .




Fig. 3. End-time increase when an equivalent static allocation is used instead of a dynamic allocation.

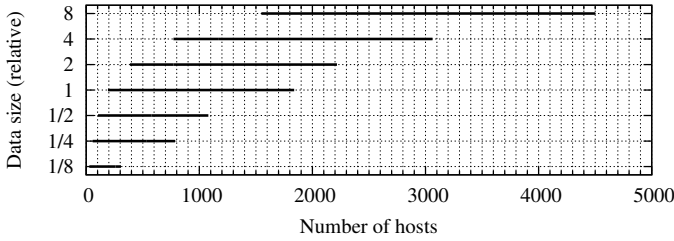


Fig. 4. Static allocation choices for a target efficiency of 75%.

However, most HPC resource managers do not support dynamic resource allocations, which forces a user to request a static number of nodes. Let us define the **equivalent static allocation** a number of nodes n_{eq} which, if assigned to the application during each step, leads to the consumed resource area $A(e_t)$. Computing n_{eq} requires to know all S_i a priori. Using simulations, we found that n_{eq} exists for $e_t < 0.8$. Of course, with a static allocation, some steps of the application run more efficiently, while others run less efficiently than e_t . However, interestingly, the end-time of the application increases with at most 2.5% (Figure 3). We deduce that, if users had an *a priori* knowledge of the evolution of the data size, they could allocate resources using a *rigid* job and their needs would be fairly well satisfied.

Unfortunately, the evolution of an AMR application is generally not known a priori, which makes choosing a static allocation difficult. Let us assume that a scientist wants to run her application efficiently. E.g., she wants her application not to run out of memory, but at the same time, she does not want to use 10% more resources than $A(75\%)$. Figure 4 shows the range of nodes the scientist should request, depending on the maximum data size S_{max} . We observe that taking such a decision without knowing the evolution of the application in advance is difficult, in particular, if the behavior is highly unpredictable.

Ideally, a distinction should be made between the resources that an evolving application expects to use in the worst case and the resources that it actually uses. The unused resources should be filled by malleable applications, that can free these resources if the evolving application requests them. However, since applications might belong to different users, this cannot be done without RMS support. The next section proposes such an RMS.

III. CoORMv2

This section introduces CoORMv2, an RMS which allows to efficiently schedule malleable and evolving applications, whether fully-, marginally- or non-predictable. First, it explains the concepts that are used in the system (Section III-A). Second, an RMS implementation is proposed (Section III-B). The section concludes by showing an example of interaction between the RMS, a malleable application and an evolving application (Section III-C).

A. Principles

The application interacts with the RMS through a specialized supervision process (e.g., its launcher) which runs on a dedicated machine, such as the front-end. At the core of the RMS-application interaction are four concepts: requests, request constraints, high-level operations and views. Let us detail each one of them.

1) *Requests*: A **request** is a description, sent by applications to the RMS, of the resources an application wants to have allocated. As in usual parallel batch schedulers, a request consists of the number of nodes (node-count), the duration of the allocation and the cluster on which the allocation should take place¹. It is the RMS that computes a start-time for each request. When the start-time of a request is the current time, node IDs are allocated to the request and the application is notified.

CoORMv2 supports the following three request types. **Non-preemptible** requests (also called run-to-completion, default in most RMSs [16]) ask for an allocation that once started cannot be interrupted by the RMS. **Preemptible** requests ask for an allocation that can be interrupted by the RMS, whenever it decides to, similar to OAR’s best-effort jobs [17].

In order to allow a non-predictably evolving application to inform the RMS about its maximum expected resource usage, CoORMv2 supports a third request type which we shall call **pre-allocation**. No node IDs are actually associated with the request, the goal of a pre-allocation being to allow the RMS to mark resources for possible future usage. In order to have node IDs allocated, an application has to submit non-preemptible requests inside the pre-allocation. Pre-allocated resources cannot be allocated non-preemptibly to another application, but they can be allocated preemptibly.

Section IV shows that these types of requests are sufficient to schedule evolving and malleable applications.

2) *Request Constraints*: Since CoORMv2 deals with applications whose resource allocation can be dynamic, each application is allowed to submit several requests, so as to allow it to express varying resource requirements. However, the application needs to be able to specify some constraints regarding the start-times of the requests relative to each other. To this purpose, each request r has two more attributes: **relatedTo**—points to an existing

¹In practice, separate batch queues are used for each cluster [15].




Fig. 5. Visual description of request constraints (Section III-A2).

request (noted r_p)—and **relatedHow**—specifies how r is constrained with respect to r_p .

CoORMv2 defines three possible values for **relatedHow** (Figure 5): **FREE**, **COALLOC**, and **NEXT**. **FREE** means that r is not constrained and **relatedTo** is ignored. **COALLOC** specifies that r has to start at the same time as r_p . **NEXT** is a special constraint that specifies that r has to start immediately after r_p , with r and r_p sharing common resources. If r requests more nodes than r_p , then the RMS will allow the application to keep the resources allocated as part of r_p and will send additional node IDs when r is started. If r requests fewer nodes than r_p , then the application has to release node IDs at the end of r_p .

3) *High-Level Operations*: Let us now examine how applications can use the request types and constraints to fulfill their need for dynamic resource allocation.

Having defined the above request types and request constraints, the targeted applications only need two low-level operations: **request()** adds a new request into the system, while **done()** immediately terminates a request (i.e., its duration is set to the current time minus the request’s start-time) and, for requests constrained with **NEXT**, specifies which node IDs have been released. Using these two operations, an application can perform high-level operations. We shall present two of which are the most relevant: spontaneous update and announced update.

A **spontaneous update** is the operation through which an application immediately requests the allocation of additional resources. It can be performed by calling **request()** with a different node-count, then calling **done()** on the current request. Given the initial request state in Fig. 6(a), Fig. 6(b) shows the outcome of a spontaneous update.

An **announced update** is the operation through which an application announces that it will require additional resources at a future moment of time (i.e., after an **announcement interval** has passed), thus allowing the system (and possible other applications) some time to prepare for the changes of resource allocation. It can be performed by first calling **request()** with a node-count equal to currently allocated number of nodes and a duration equal to the announcement interval, then calling **request()** with a new node-count and, finally, calling **done()** on the current request (see Figure 6(c)).

Note that, in order to guarantee that at update is successful (i.e., the RMS can allocate additionally requested resources), updating non-preemptible requests can only be guaranteed if they can be served from a pre-allocation.




Fig. 6. Performing an update.

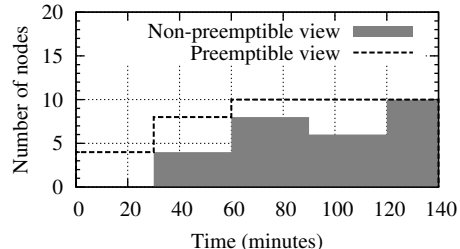


Fig. 7. Example of views for one cluster.

4) *Views*: For the targeted applications, simply sending requests to the RMS is not enough. Applications need to be able to adapt their requests to the availability of the resources. For example, instead of requesting a large number of nodes which are available only at a future time, a moldable application might want to request fewer nodes to reduce its waiting time, thus reducing its end-time.

To this end, the RMS provides each application a **view** containing the current and future availability of resources. Views should be regarded as the *current information* that the RMS has, but should not be taken as a guarantee. Applications can scan their view and estimate when a request would be served by the RMS. This helps moldable applications as they do not need to emulate the scheduling algorithm of the RMS. When the state of the resources changes, the RMS will push new views to the applications, so that they can update their requests, if necessary.

On a more technical note, a view stores for each cluster a step function, that stores the **cluster availability profile**: the x-axis is the absolute time and the y-axis is the number of available nodes as illustrated in Figure 7.

In CoORMv2, two types of views are sent to each application i : **non-preemptive** ($V_{-P}^{(i)}$) and **preemptive view** ($V_P^{(i)}$). The former allows applications to estimate when pre-allocations and non-preemptive requests will be served, while the latter allows it to estimate when preemptive requests will be served.

The preemptive view is also used to signal an application, that it has to release some preemptively allocated resources, either immediately or at a future time. In CoORMv2, each application is supposed to cooperate and immediately release node IDs when it is asked to, using the above presented update operation. Otherwise, if a protocol violation is detected, the RMS kills the application’s processes and terminates the session with it.

B. RMS Implementation

This section presents an RMS implementation that we have realized. A COORMv2 RMS has three tasks: (i) for each application i compute a non-preemptible ($V_{-P}^{(i)}$) and a preemptible view ($V_P^{(i)}$), (ii) compute the start-time of each request and (iii) start requests and allocate node IDs.

For brevity, only a high-level description of the implementation is given. The RMS runs a scheduling algorithm whenever it receives a `request` or `done` message. In order to coalesce messages coming from multiple applications at the same time and reduce system load, the algorithm is run at most once every **re-scheduling interval**, which is an administrator-chosen parameter. A value for this parameter is given in Section V.

The scheduling algorithm goes as follows. Applications are sorted in a list based on the time the applications connected to the RMS. First, pre-allocation requests are scheduled using Conservative Back-Filling (CBF) [18]. Next, inside these pre-allocations, non-preemptible requests are scheduled. Non-preemptible requests which cannot be served from pre-allocations are implicitly wrapped in pre-allocations of the same size. The remaining resources are used to schedule preemptible requests.

Regarding preemptive views and the start-time of preemptive requests, they are computed so as to achieve equi-partitioning. However, should an application not use its partition, other applications are allowed to fill it in.

The proposed scheduling algorithm has a linear complexity with respect to the number of requests in the system. We have developed a Python implementation (with more compute-intensive parts in C++) which is able to handle approximately 500 requests/second on a single core of an Intel® Core™2 Duo CPU @ 2.53 GHz.

The above implementation is just one possible implementation and can easily be adapted to other needs. For example, the amount of resources that an application can pre-allocate can be limited, by clipping its non-preemptible view.

C. Example Interaction

Let us describe the RMS-Application protocol using an example. Assume there is one NEA and one malleable application in the system (see Figure 8). First, the NEA connects to the RMS (step 1). In response, the RMS sends the corresponding views to the application (step 2). The NEA sends a pre-allocation (step 3) and a first non-preemptible request (step 4), to which the RMS will immediately allocate resources (step 5). Similar communication happens between the RMS and the malleable application, except that a preemptible request is sent (step 6–9). As the computation of the NEA progresses, it requires more resources, therefore, it performs a spontaneous update (step 10–11). As a result, the RMS updates the view of the malleable application (step 12), which immediately frees some nodes (step 13–14). Then, the RMS allocates these nodes to the NEA (step 15).

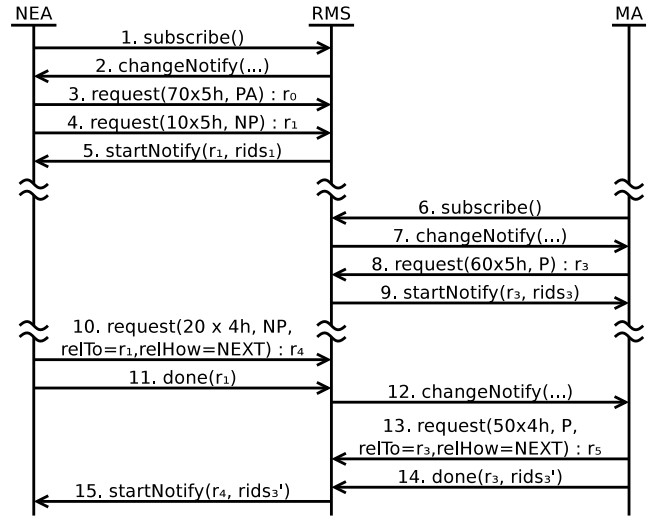


Fig. 8. Example of an interaction between the RMS, a Non-predictably Evolving Application and a Malleable Application.

IV. APPLICATION SUPPORT

This section evaluates how the COORMv2 architecture supports the types of application presented in Section I.

A **rigid application** sends a single non-preemptible request of the user-submitted node-count and duration.

A **moldable application** waits for the RMS to send a non-preemptible view, then runs a resource selection algorithm, which chooses a non-preemptible request. Should the state of the system change before the application starts, the RMS shall keep it informed by sending new views. This allows the application to re-run its selection algorithm and update its request, similarly to COORM [19].

A **malleable application** first sends a non-preemptible request r_{min} with its minimum requirements. Next, for the extra resources (i.e., the malleable part), the application scans its preemptive view V_P and sends a preemptible request r_{extra} , which is COALLOCATED with r_{min} . r_{extra} can either request a node-count equal to the node-count in V_P , or it can request fewer nodes. This allows an application to select only the resources it can actually take advantage of. For example, if the malleable application requires a power-of-two node-count, but 36 nodes are available in its preemptive view, it can request 32 nodes, leaving the other 4 to be filled by another application. During execution, the application monitors V_P and updates r_{extra} if necessary.

A **fully-predictably evolving application** sends several non-preemptible requests linked using the NEXT constraint. During its execution, if from one request to another the node-count decreases, it has to call `done` with the node IDs it chooses to free. Otherwise, if the node-count increases, the RMS sends it the new node IDs.

A **non-predictably evolving application** (NEA) first sends a pre-allocation request (whose node-count is discussed below), then sends an initial non-preemptible request. During execution, the application updates the non-

preemptible request as the resource requirements change. Since such updates can only be guaranteed as long as they happen inside a pre-allocation, such an application may adopt two strategies: sure execution or probable execution.

The application can opt for a **sure execution**, i.e., uninterrupted run-to-completion, if it knows the maximum node-count (n_{max}) it may require. To adopt this strategy, the size of the pre-allocation must be equal to n_{max} . In the worst case, n_{max} is the whole machine.

Otherwise, if the application only wants **probable execution**, e.g., because pre-allocating the maximum amount of required resources is impractical, then the application sends a “good-enough” pre-allocation and optimistically assumes never to outgrow it. If at some point the pre-allocation is insufficient, the RMS cannot guarantee that updates will succeed. Therefore, the application has to be able to checkpoint. It can later resume its computations by submitting a new, larger pre-allocation. Note that, in this case, the application might further be delayed as the RMS might have placed it at the end of the waiting queue.

If multiple NEAs enter the system two situations may occur. Either their pre-allocations are small enough to fit inside the system simultaneously, in this case the NEAs are launched at the same time, or their pre-allocations are too large to fit simultaneously, in which case the one that arrived later will be queued after the other. In both cases, the RMS is able to guarantee that whenever one of the NEAs requests an update inside its pre-allocation, it can actually be served. Resources that are pre-allocated to a NEA, but not allocated, cannot be non-preemptively allocated to another application. However, these resources can be filled by the malleable part of other applications. This approach is evaluated in the next section.

V. EVALUATION

This section evaluates COORMv2. First, the application and resource models are presented. Next, the impact on evolving and malleable applications is analyzed.

The evaluation is based on a discrete-event simulator, which simulates both the COORMv2 RMS and the applications. To write the simulator, we have first written a real-life prototype RMS and synthetic applications. Then, we have replaced remote calls with direct function calls and calls to `sleep()` with simulator events.

A. Application and Resource Model

As other works which propose scheduling new types of application [6], we shall only focus on the malleable and evolving applications, which can fully take advantage of COORMv2’s features. Therefore, we shall not evaluate our system against a trace of rigid jobs [20] as is commonly done in the community. Nevertheless, as shown in Section IV, the proposed system does support such a usage.

For the evaluation, two applications are used: a non-predictably evolving AMR and a malleable parameter-sweep application. Let us present how the two applications behave with respect to COORMv2.

1) *AMR Application*: A synthetic AMR application is considered, which behaves as a non-predictably evolving application with a sure execution (see Section IV). The application knows its speed-up model, but cannot predict how the working set will evolve (see Section II). Knowing only the current working set size, the application tries to target an efficiency of 75% using spontaneous updates.

The user submits the application by trying to “guess” the node-count n_{eq} of the equivalent static allocation (defined in Section II). The application uses this value for its pre-allocation. Inside this pre-allocation, non-preemptible allocations are sent/updated, so as to keep the application running at the target efficiency.

Since the user cannot determine n_{eq} a priori, her guesses will be different from this optimum. Therefore, we introduce as a simulation parameter the **overcommit factor**, which is defined as the ratio between the node-count that the user has chosen and the best node-count assuming a posteriori knowledge of the application’s behavior.

2) *Parameter-Sweep Application (PSA)*: Inspired by [21], we consider that this malleable applications is composed of an infinite number of single-node tasks, each of duration d_{task} . As described in Section IV, the PSA monitors its preemptive view. If more resources are available to it than it has currently allocated, it updates its preemptive request and spawn new processes. If the RMS requires it to release resources immediately, it kills a few tasks then updates its request. The computations done so far are lost. We measure the **PSA waste** which is the number of node-seconds wasted in such a case.

If the RMS is able to inform the PSA in a timely manner that resources will become unavailable, then the PSA waits for some tasks to complete, afterwards it updates its request to release the resources on which the completed tasks ran. No waste occurs in this situation.

Resource Model: We assume that resources consist of one single, large homogeneous cluster of n nodes. The re-scheduling interval of the RMS is set to 1 second, to obtain a very reactive system.

B. Scheduling with Spontaneous Updates

This experiment evaluates whether COORMv2 solves the initial problem of efficiently scheduling NEAs.

Two applications enter the system: one AMR and one PSA (PSA_1 with $d_{task} = 600$ s). Since we want to highlight the advantages that supporting evolving applications in the RMS brings, we shall schedule the AMR application in two ways: **dynamic**, in which the application behaves as described above, and **static**, in which the application is forced to use all the resources it has pre-allocated.

Regarding the resources, the number of nodes n is chosen so that the pre-allocation of the AMR application can succeed. By observing the behaviour of the application during some preliminary simulations, we concluded that for an overcommit factor of ϕ , having $n = 1400 \times \phi$ is sufficient.

Figure 9 shows the results of the simulations. Regarding the amount of resources effectively allocated to the AMR (AMR used resources), the figure shows that as the overcommit factor grows, a static allocation forces the application to consume more resources. This happens because, as the user chooses more nodes for the application than optimal, the application runs less efficiently. Moreover, the application is unable to release the resources it uses inefficiently to another application.

In contrast, a dynamic allocation allows the AMR application to cease resources it cannot use efficiently. Therefore, as the overcommit factor grows, a dynamic allocation allows the application to maintain an efficient execution. Note however, that thanks to the pre-allocation, if the resource requirements of the application increase, the application can request new resources and the RMS guarantees their timely allocation. Since, COORMv2 mandates that preemptible resources be freed immediately, the impact on the AMR application is negligible.

In this experiment, the RMS allocates the resources that are not used by the AMR application to the malleable PSA. However, since the AMR application does spontaneous updates and the RMS cannot inform the PSA in a timely manner that it has to release resources, the PSA has to kill some tasks, thus, computations are wasted (Figure 9 bottom). Nevertheless, we observe that the amount of resources wasted is smaller than the resources which would be used by an inefficiently running AMR application.

The PSA waste first increases as the overcommit factor increases, then stays constant after the overcommit factor is greater than 1. This happens because once the AMR application is provided with a big enough pre-allocation, it does not change the way it allocates resources inside it.

To sum up, this experiment shows that, with proper RMS support, resources can be more efficiently used, if the system contains non-predictably evolving applications. Nevertheless, PSA waste is non-negligible, therefore, the next section evaluates whether using announced updates makes sense to reduce this waste.

C. Scheduling with Announced Updates

This section evaluates whether using announced updates can improve resource utilization. To this end, the behavior of the AMR application has been modified as follows. A new parameter has been added, the **announce interval**: instead of sending spontaneous updates, the application shall send announced updates (see Section III-A3). The node-count in the update is equal to the node-count required at the moment the update was initiated. This means that the AMR receives new nodes later than it would require to maintain its target efficiency. If the announce interval is zero, the AMR behaves as with spontaneous updates. Otherwise, the application behaves like a marginally-predictably evolving application.

For this set of experiments, we have set the overcommit factor to 1. The influence of this parameter has already

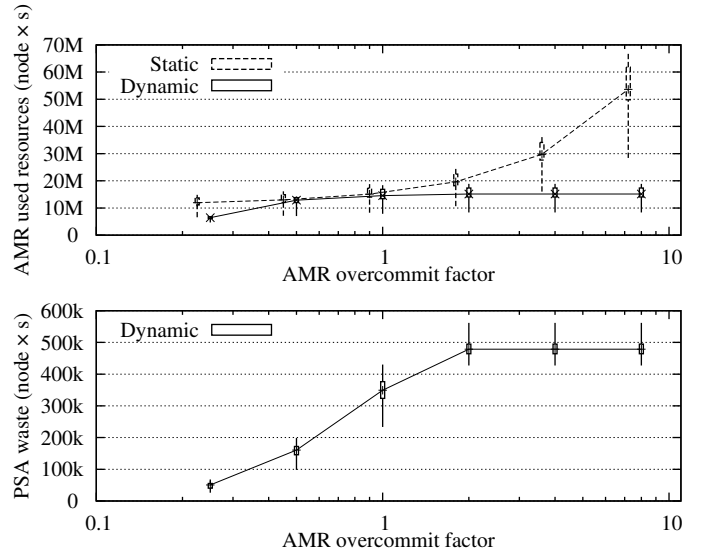


Fig. 9. Simulation results with spontaneous updates.

been studied in the previous section.

Using announced updates instead of spontaneous ones has several consequences. On the negative side, the AMR application is allocated fewer resources than required, thus, its end-time increases (see Figure 10). On the positive side, this informs the system some time ahead that new resources will need to be allocated, thus, it allows other applications more time to adapt.

Figure 10 shows that the PSA waste is decreasing as the announce interval increases. This happens because, in order to free resources, the PSA can gracefully *stop* tasks when they complete, instead of having to *kill* them. Once the announce interval is greater than the task duration d_{task} , no PSA waste occurs.

Let us discuss the **percent of used resources**, which we define as the resources allocated to applications minus the PSA waste as a percentage of the total resources. Figure 10 shows that announced updates do not generally improve resource utilization. This happens because the PSA cannot make use of resources if the duration of the task is greater than the time the resources are available to it. In other words, the PSA cannot fill resources unused by the AMR if the “holes” are not big enough: these “holes” are either wasted on killed tasks or released to the system. However, the higher the announce interval, the more resources the PSA frees, so that the RMS can either allocate them to another application (see next section), or put them in an energy saving mode.

We observe that the percent of used resources has two peaks, when the announce interval is 300s and 600s. In these two cases, a “resonance” occurs between the AMR and the PSA. The “holes” left by the AMR have exactly the size required to fit a PSA task. On the contrary, when the announce interval is just below the duration of the task (i.e., 550s) the percent of used resources is the lowest, as




Fig. 10. Simulation results with announced updates.

the PSA has to kill tasks just before they end.

D. Efficient Resource Filling

In the previous section, we have observed that while announced updates reduce PSA waste, they do not improve resource utilization. This happens because the PSA’s task duration is greater than the “holes” left by the AMR. In this section, we evaluate whether COORMv2 is able to allocate such resources to another PSA, with a smaller task duration, in order to improve resource utilization.

To this end, we add another PSA in the system (PSA_2 with $d_{task} = 60$ s). By default, the RMS does equi-partitioning between the two PSA applications, however, when one of them signals that it cannot use some resources, the other PSA can request them.

To highlight the gain, we compare this experimental setup with an RMS which implements “**strict**” **equi-partitioning**: instead of allowing one PSA to fill the resources that are not requested by the other PSA, the RMS presents both PSAs a preemptible view, which makes them request the same node-count.

Figure 11 shows the percent of used resources. For readability, only medians are shown. Experiments show that COORMv2 is able to signal PSA_2 that some resources are left unused by PSA_1 . Resource utilization is improved, because PSA_2 is able to request these resources and do useful computations on them.




Fig. 11. Simulation results for two PSAs.

To sum up, COORMv2 is able to efficiently deal with the applications it has been designed for: evolving and malleable. On one hand, evolving applications can efficiently deal with their unpredictability by separately specifying what their maximum and what their current resource requirements are. Should evolving applications be able to predict some of their evolution, they can easily export this information into the system. On the other hand, malleable applications are provided with this information and are able to improve resource usage.

VI. RELATED WORK

To our knowledge, Feitelson [22] was the first to distinguish applications with dynamic resource usage into those that initiate the change themselves (*evolving*) and those that adapt to RMS’s constraints (*malleable*). Unfortunately, they are often ambiguously called “dynamic” applications in literature [6], [23]. Let us review RMSs which support these types of applications.

A. Malleable

ReSHAPE [6] is a framework for efficiently resizing iterative applications. After each iteration, applications report performance data to the RMS and the RMS attempts to divide resources, so as to optimize global speed-up. However, ReSHAPE assumes that all iterations of an application are identical and caches performance data. Therefore, applications whose iterations are irregular, such as AMR or APM applications, cannot be efficiently scheduled. Even if performance history were regularly flushed, ReSHAPE targets a specific type of applications, whereas COORMv2 also deals with other types of malleable application.

More general support for malleability can be found in KOALA [21]. Applications declare their minimum resource requirements, then the RMS divides the “extra” resources equally (*equi-partitioning*). The approach is appealing and has been successfully applied to scheduling Parameter-Sweep Applications (PSA) on idle resources. However, the minimum requirements are fixed at submittal and cannot be changed during the execution. Therefore, this solution cannot be used to schedule evolving applications.

The OAR batch scheduler supports low-priority, preemptible jobs called *best-effort*. If there are insufficient resources to serve a normal job, best-effort jobs are killed after a grace period. Best-effort jobs can be used to

support malleability [5], therefore, this type of job has inspired COORMv2's preemptible requests. However, best-effort jobs are killed even if only part of the allocated resources need to be freed. In contrast, COORMv2 allows an application to update its preemptible request and to free only as many resources as needed.

B. Evolving

In the Cloud Computing context, where consumed resources are paid for, great interest has been shown for applications that dynamically adapt their resource allocation to internal constraints [24]. However, most papers in this field only study how applications should adapt their resource requests, without studying how the RMS should guarantee that these resources can actually be allocated. Indeed, large-scale Cloud deployments can refuse an allocation request with an "out of capacity" error [25].

The Moab Workload Manager supports so-called "dynamic" jobs [23], which are both evolving and malleable: the application can grow/shrink both because its internal load changes and because the RMS asks it to. Moab regularly queries each application what its current load is, then it decides how resources are allocated. This feature has been mainly thought for interactive workloads and is not suitable for batch workloads. For example, if two non-predictably evolving applications arrive at the same time, instead of running them simultaneously, it might be better to run one after the other, so as to guarantee that peak requirements can be met. Also, Moab's interface does not allow marginally- or fully-predictable applications to present their estimated evolution to the system.

In contrast, COORMv2 allow an application to pre-allocate resources, so that the RMS can guarantee their availability when the application needs them. If two applications potentially require a large amount of resources, they shall be executed one after the other, to guarantee their completion. An application which can predict its evolution can export this information into the system, so as to optimize resource usage.

VII. CONCLUSION

This paper presented COORMv2, an RMS architecture which offers a simple interface to support various types of application, among others, malleable and fully-/marginally-/non-predictably evolving. Experiments showed that, if applications take advantage of the RMS's features, resource utilization can be improved a lot, while guaranteeing that resource allocations are always satisfied.

Future work can be divided in two directions. First, it would be interesting to study how accounting should be done in COORMv2, so as to determine users to efficiently use resources. Second, COORMv2 was designed for managing homogeneous clusters and assumes that the choice of node IDs can be left to the RMS. However, supercomputers feature a non-homogeneous network [26], therefore, an extension to COORMv2 should allow network-sensitive applications to optimize their placements on such platforms.

ACKNOWLEDGMENT

This work was supported by the French ANR COOP project, n°ANR-09-COSI-001-02.

Experiments presented in this paper were carried out using the Grid'5000 experimental testbed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (<https://www.grid5000.fr>).

REFERENCES

- [1] J. Hungershofer, "On the combined scheduling of malleable and rigid jobs," in *SBAC-PAD*, 2004, pp. 206–213.
- [2] K. El Maghraoui, T. J. Desell *et al.*, "Dynamic malleability in iterative MPI applications," in *CCGRID*, 2007, pp. 591–598.
- [3] J. Buisson, F. André *et al.*, "A framework for dynamic adaptation of parallel components," in *PARCO*, 2005, pp. 65–72.
- [4] J. Buisson, O. Sonmez *et al.*, "Scheduling malleable applications in multicloud systems," *CoreGRID*, Tech. Rep. TR-0092, 2007.
- [5] M. C. Cera, Y. Georgiou, O. Richard, N. Maillard *et al.*, "Supporting malleability in parallel architectures with dynamic CPUSETS mapping and dynamic MPI," in *ICDCN*, 2010.
- [6] R. Sudarsan and C. J. Ribbens, "ReSHAPE: A framework for dynamic resizing and scheduling of homogeneous applications in a parallel environment," *CoRR*, vol. abs/cs/0703137, 2007.
- [7] J. Dongarra *et al.*, "The international exascale software project roadmap," *IJHPCA*, vol. 25, no. 1, pp. 3–60, 2011.
- [8] T. Plewa, T. Linde, and V. G. Weirs, Eds., *Adaptive Mesh Refinement: Theory and Applications*. Springer, 2003.
- [9] D. R. Welch, T. C. Genoni, R. E. Clark, and D. V. Rose, "Adaptive particle management in a particle-in-cell code," *Journal of Computational Physics*, vol. 227, no. 1, pp. 143–155, 2007.
- [10] C. Burstedde *et al.*, "Extreme-scale AMR," in *SC*, 2010.
- [11] D. F. Martin, P. Colella *et al.*, "Adaptive mesh refinement for multiscale nonequilibrium physics," *Computing in Science and Eng.*, vol. 7, pp. 24–31, 2005.
- [12] G. L. Bryan, T. Abel, and M. L. Norman, "Achieving extreme resolution in numerical cosmology using adaptive mesh refinement: Resolving primordial star formation," in *SC*, 2001.
- [13] D. J. Kerbyson, J. H. Alme *et al.*, "Predictive performance and scalability modeling of a large-scale application," in *SC*, 2001.
- [14] J. Luitjens and M. Berzins, "Improving the performance of Uintah: A large-scale adaptive meshing computational framework," in *IPDPS*, 2010.
- [15] H. Casanova, "Benefits and drawbacks of redundant batch requests," *Journal of Grid Computing*, vol. 5, no. 2, 2007.
- [16] D. G. Feitelson, L. Rudolph, and U. Schwiigelshohn, "Parallel job scheduling – a status report," in *JSSPP*, 2004.
- [17] N. Capit, G. Da Costa *et al.*, "A batch scheduler with high level components," *CoRR*, vol. abs/cs/0506006, 2005.
- [18] A. W. Mu'alem and D. G. Feitelson, "Utilization, Predictability, Workloads, and User Runtime Estimates in Scheduling the IBM SP2 with Backfilling," *TPDS*, vol. 12, no. 6, pp. 529–543, 2001.
- [19] C. Klein and C. Pérez, "Untying RMS from Application Scheduling," INRIA, Research Report RR-7389, 2010.
- [20] "The parallel workloads archive." [Online]. Available: <http://www.cs.huji.ac.il/labs/parallel/workload>
- [21] O. O. Sonmez *et al.*, "Scheduling strategies for cycle scavenging in multicloud grid systems," in *CCGRID*, 2009, pp. 12–19.
- [22] D. G. Feitelson and L. Rudolph, "Towards convergence in job schedulers for parallel supercomputers," in *JSSPP*, 1996.
- [23] Adaptive Computing Enterprises, Inc., "Moab workload manager administrator guide, version 6.0.2." [Online]. Available: <http://www.adaptivecomputing.com/resources/docs/mwm>
- [24] S. Vijayakumar, Q. Zhu, and G. Agrawal, "Automated and dynamic application accuracy management and resource provisioning in a cloud environment," in *GRID*, 2010, pp. 33–40.
- [25] "Lessons learned building a 4096-core cloud HPC supercomputer." [Online]. Available: <http://blog.cyclecomputing.com/2011/03/cyclecloud-4096-core-cluster.html>
- [26] N. R. Adiga, M. A. Blumrich *et al.*, "Blue Gene/L torus interconnection network," *IBM J. Res. Dev.*, vol. 49, 2005.