

# VRM: A Failure-Aware Grid Resource Management System

Lars-Olof Burchard\*, César A. F. De Rose†, Hans-Ulrich Heiss\*, Barry Linnert\*, Jörg Schneider\*

\*{baron,heiss,linnert,komm}@cs.tu-berlin.de  
Technische Universität Berlin, GERMANY

†derose@inf.pucrs.br  
PUCRS, Porto Alegre, BRASIL

**Abstract**—For resource management in Grid environments, advance reservations turned out to be very useful and hence are supported by a variety of Grid toolkits. However, failure recovery for such systems has not yet received the attention it deserves. In this paper, we address the problem of remapping reservations to other resources, when the originally selected resource fails. Instead of dealing with jobs already running, which usually means checkpointing and migration, our focus is on jobs that are scheduled on the failed resource for a specific future period of time but not started yet. The most critical factor when solving this problem is the estimation of the downtime. We avoid the drawbacks of under- or overestimating the downtime by a dynamic load-based approach that is evaluated by extensive simulations in a Grid environment and shows superior performance compared to estimation-based approaches.

## I. INTRODUCTION

Currently, Grid research moves its focus from the basic infrastructure that enables the allocation of resources in dynamic and distributed environments in a transparent way to more advanced management systems that accept and process complex jobs and workflows consisting of numerous sub-tasks and, e.g., provide guarantees for the completion of such jobs. In this context, the introduction of service level agreements (SLA) provides flexible mechanisms for agreeing on the quality-of-service (QoS) provided by various resources, including mechanisms for negotiating SLAs [1]. The introduction of SLAs involves prices for resource usage and also implies fines that must be paid when the assured QoS is not met. Depending on the scenario, this may be, e.g., a missed deadline for the completion of a sub-job in a workflow. Consequently, the definition of SLAs demands for control over each job and its required resources at any stage of the job's lifetime from the request negotiation to the completion. An example for a resource management framework covering these aspects is the *Virtual Resource Manager* (VRM) architecture described in [2].

A typical example for a complex workflow in a Grid is depicted in Fig. 1. The workflow processed in the distributed environment consists of five sub-tasks which are executed one after another in order to produce the final result, in this case the visualization of the data. This

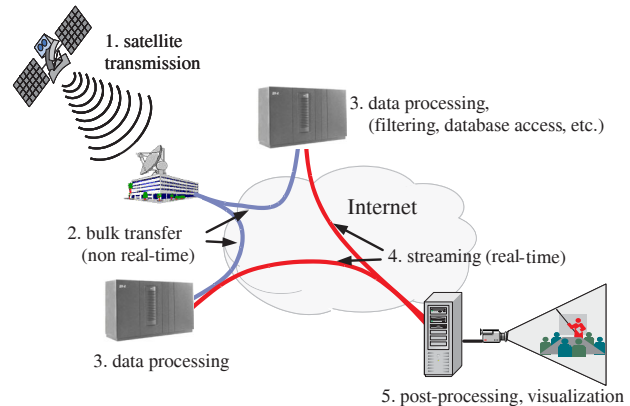


Fig. 1. Example: Grid application with time-dependent tasks.

includes network transmissions as well as computations on two cluster computers.

One important aspect in this context is the behavior of the management system in case of failures. While current research mainly focused on recovery mechanisms for those jobs that are already active, in advance reservation environments it is also necessary to examine the impact of failures to admitted but not yet started jobs or sub-jobs. In contrast to the sophisticated and difficult mechanisms needed to deal with failures for running jobs, e.g., checkpointing and migration mechanisms, jobs not yet started can be dealt with in a transparent manner by remapping those affected jobs to alternative resources.

As previous work in this area showed it is necessary to deal with these affected but not yet active jobs. Consequently we enhanced the VRM as our prototype for next generation Grid management systems with failure recovery functionalities.

In this paper, we present the novel, load-based approach that has been developed for our VRM framework and has the advantage of adapting to the actual load situation and remapping affected jobs accordingly. The adaptiveness is an important feature as estimations of the failure duration are unreliable.

The remainder of this document is organized as follows: firstly, the general problem is described, including the properties of the advance reservation environment and the necessary assumptions and conditions to apply our

approach. Based on these general remarks, the VRM framework is introduced as it is the target application for the failure recovery mechanism. Following that, the load-based approach for remapping jobs is presented. In Sec. V, the strategies are evaluated using extensive simulations, showing the superiority of our approach compared to estimations. Before the paper is concluded with some final remarks, related work important for this paper is outlined.

## II. PROBLEM DEFINITION

In this section the general properties of advance reservation systems are described as well as different aspects related to the problem of remapping jobs in case of failures.

Although the VRM architecture is capable of handling various kinds of resources at the same time, this work takes only computing resources into account. Comparable kinds of resources, e.g., storage space or bandwidth, can be handled by this mechanism as well.

### A. Properties of the Advance Reservation Environment




Fig. 2. Advance reservations: status about future utilization

Advance reservations are requests for a certain amount of resources during a specified period of time. In general, a reservation can be made for a fixed period of time in the future, called *book-ahead interval* ( $t_{BA}$ , see Fig. 2). The book-ahead interval is divided into *slots* of fixed size, e.g., minutes, and reservations can be issued for a consecutive number of slots [3]. The time between issuing a request and the start time of the request is called *reservation time*  $r$ . The finishing time or duration for a given request must be specified to ensure reliable admission control, i.e., to determine whether or not sufficient resources can be guaranteed for the requested period. As depicted in Fig. 2, this approach requires to keep the status of each resource, i.e., information about future requests which are already admitted, for the whole book-ahead interval.

In this environment, failure recovery not only has to handle already active jobs, but also those which are admitted but not yet started, so-called *inactive jobs*. This means that the affected inactive jobs have to be *remapped in advance* to another matching resource. As the timing parameters start and stop time were specified during the admission, e.g., as part of a service level agreement (SLA), jobs can only be moved to another resource but not shifted in the temporal dimension.

The general benefits of remapping in advance are shown in [4] where estimations of the failure duration were used. As the end of the failure is usually unknown, it is not

easy to decide which active jobs on the resource have to be taken into account for remapping. This paper presents a downtime independent approach to decide this question basing on the actual load situation.

### B. Implications of the Environment

In order to implement failure recovery mechanisms in the advance reservation environment, it is necessary to consider the types of available resources, i.e., resource heterogeneity or homogeneity play an important role. Whenever identical hardware and software infrastructure is available, including a wide range of properties such as processor type, cache sizes, operating system version, or libraries, mapping an inactive job to another resource is relatively simple. It may be even possible to migrate a running job, e.g., with support from checkpointing mechanisms, some systems provide libraries for that purpose [5].

However, in a heterogeneous environment the same task is more complex and difficult, even mapping inactive jobs to a different resource, e.g., a cluster, may have consequences on the run-time of the respective processes and hence, must be considered. For example, mapping a job from a 2 GHz processor to a 1 GHz processor will increase the overall execution time. In such an environment, the migration of running jobs is even more difficult. Consequently, in the context of this study active jobs are considered to be not remappable. For many resource types, such as cluster systems or parallel computers, such functionality lacks completely or has to be implemented explicitly by the application. However, this assumption is not crucial for the usage of our approach or the success of the remapping strategy itself.

Furthermore, we consider a homogeneous environment with resources of at least very similar hardware and software infrastructure which allows jobs to be executed on any of the available resources. Because of the known problems with multisite applications, we assume in the following that jobs cannot be split among several resources. An example for an actual application environment of our approach is given in the following section.

## III. GRID ENVIRONMENT

The approach presented here has been developed in the context of the *Virtual Resource Manager* (VRM) described in [2]. Since the VRM supports QoS by means of SLAs, failure recovery is an essential feature. To meet SLAs in terms of, e.g., guaranteed completion times for complex compute jobs, the VRM needs to have complete control over the resources and the jobs during run-time (run-time responsibility). In the following, we briefly describe the architecture of the VRM as the environment of our failure recovery mechanism.

The *Administrative Domain Controller* (ADC) constitutes the central management component of the VRM architecture (see Fig. 3). The ADC is responsible for establishing so called *Administrative Domains* (AD) which

consist of a number of underlying local resources and their local *Resource Management Systems* (RMS). These management systems may control arbitrary types of resources, e.g., cluster systems, parallel computers, or networks and are connected to the ADC by *Active Interfaces*. Such an Active Interface is also available to connect to another ADC responsible for a subdomain.

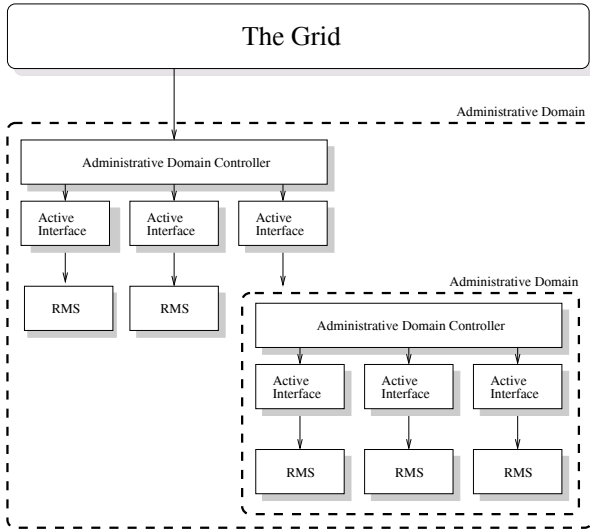


Fig. 3. Hierarchical Administrative Domain Structure

The failure recovery mechanism proposed in this paper will be situated in the ADC component. Once any of the underlying compute resources fails partly or entirely, the jobs allocated to the failed resource are mapped onto alternative resources within the same domain according to our strategy. Partial failures, e.g., of one or more nodes within a cluster computer, may also require the recovery mechanism to act, as the total capacity of a resource may be exhausted. Within the VRM, inactive jobs can be transparently mapped without further notification to the users which is a major advantage compared to other Grid resource management systems such as Globus [3].

The framework of the VRM provides control over not only compute resources, but also any other resource required for the remapping of jobs, i.e., also interconnection networks. For example, remapping a compute job with large amounts of input data requires a appreciable amount of time which must be considered and network transmissions must be planned accordingly which may include the reservation of network bandwidth. Allocation of network bandwidth is often available in dedicated networks for high performance Grid environments, e.g., *LambdaGrids* [6].

#### A. Using Batch Jobs to Improve the Performance

Batch jobs are frequently used also in planning based systems [7]. These jobs are not provided with fixed start and stop times but will be placed in a queue and are processed as soon as sufficient capacities are available. In our environment, such batch jobs can be placed on a failed resource behind the computed remapping interval as long

as the failure persists. The resource management system then places batch jobs onto the currently failed resource and once the remapping interval is extended or the failure persists beyond their anticipated start time, the batch jobs are simply postponed. Otherwise, they can be started. The rationale behind this approach is that no timing guarantees are given for batch jobs and thus, postponing these jobs is less costly compared to terminating planned jobs.

#### B. Discovery of Additional Resources

In case no sufficient resources are available or the accumulation of free resources does not satisfy the additional requirements arising with the remapping of jobs to different computing systems, the VRM is capable of searching for free resources in a peer-to-peer like network of other VRMs. For this purpose, knowledge about the job properties and requirements is essential. The problem of discovering suitable resources in another administrative domain or computing environment can be handled using, e.g., ontology-driven resource discovery as presented in [8]. Interfaces to other resource management systems, e.g., Globus, are also conceivable. As the VRM starts the search for alternative resources within its local domain, time-intensive search for suitable resources is only necessary when the local capacities are exhausted.

### IV. LOAD-BASED REMAPPING ALGORITHM

In this section, our load-based remapping approach used in the VRM is presented.

#### A. Remapping Strategies

In case a failure of a specific resource, e.g., a cluster, is noticed, the management system has to face two different tasks to minimize the impact of the failure. First, the management system has to determine all jobs that have to be taken into account for remapping and, as a second step, these jobs have to be remapped to other resources.

In order to remap the set of chosen jobs, the task is to maximize the success of the remapping according to some optimization criterion, e.g., the amount of successfully remapped jobs. Other optimization criteria are conceivable as well although not targeted in this paper, e.g., minimizing fines to be paid for terminated jobs.

Because the reservations are fixed in time it is not possible to shift the jobs to the future on the local system or alternative resources.

Finding feasible alternative resources for a specified set of jobs is a classical bin packing problem, but in the spatial dimension not in the temporal one [9]. Therefore, a simplified algorithm has to be used to solve the generally NP-hard problem online during the failure recovery.

Different strategies for remapping a set of affected jobs were evaluated in [4]. As it was shown, the difference among the individual strategies is rather small. This means, it may be possible to select a strategy that matches additional constraints, such as preferring long book-ahead times leading to FCFS strategy as used in this paper.

## B. Downtime-Independent Remapping

Former work showed, that predictions of the actual downtime are critical, as the downtime cannot be accurately anticipated [4]. In the case of underestimating the downtime the number of terminated jobs increases drastically. The mechanism integrated in the VRM does not rely on predictions and is independent of the actual downtime. The general approach used in the VRM is to identify and remap jobs which are unlikely to be safely remapped at any later point in time. For that purpose, in each time slot throughout the duration of the failure a *remapping interval* is calculated. The length of this remapping interval, computed as described in the following sections, is *independent* of the actual downtime, which is unknown to the Grid resource management system.

As described before, we do not deal with the mechanisms to recover already active jobs on the broken resource. So we assume, that either an adequate migration mechanism is available that handles all jobs that are already running, or—which is more likely—thus running jobs are simply terminated.

All remaining jobs using the broken resource within the remapping interval are considered for remapping. The management system searches for another resource with sufficient free capacities for every job. If the remapping is not successful, the job remains on the failed system until its start time since the failure may have ended until then, otherwise the job will be terminated. All other jobs are not remapped, even if they are assigned to the currently broken resource.

The failed resource is blocked for the remapping interval only. This means, new reservations for time slots after the remapping interval may be booked on the currently broken resource and are remapped later if necessary.

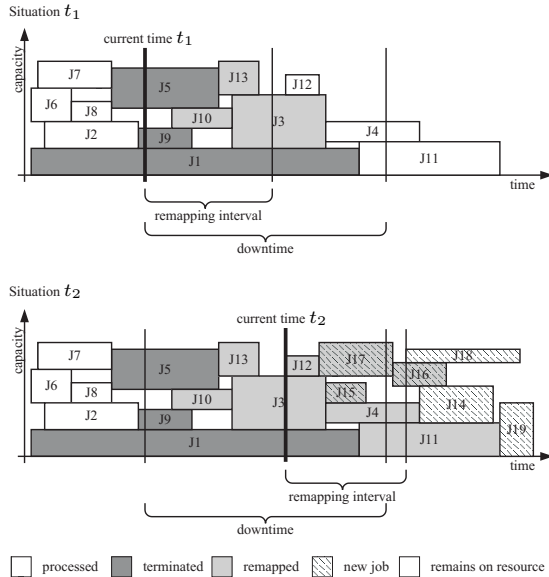


Fig. 4. Examples for the usage of the remapping interval during a resource downtime for two different time slots  $t_1$  (above) and  $t_2$  (below).

Fig. 4 is an illustration of the approach based on the remapping interval. For the first situation at  $t_1$ , this means that the dark gray jobs are handled by the failure recovery for active jobs (in the simplest case just terminated) as they were active when the failure occurred. The light gray jobs are considered for remapping and will be assigned to other resources if possible. The white jobs stay on the broken resource also if they reside within the—currently unknown—downtime. The time slots after the remapping interval are even available for new jobs.

Due to the recalculation of the remapping interval in each of the following time slots, more and more jobs mapped to the broken resource will be remapped. Jobs submitted after the failure occurs are also remapped if they are within the remapping interval, as can be seen in the second situation at  $t_2$  in Fig. 4. Since the downtime is unknown to the management system, also those jobs are remapped that will start after the downtime (J16). This remapping and the recalculation of the remapping interval is done until the resource is recovered.

## C. Objectives for the Remapping Interval

The main requirement is that the remapping interval must be sufficiently long, such that any job affected by a failure can be safely remapped with high probability.

On the other hand, all jobs that would start after the downtime but had been remapped, block free capacity on other resources. As by a longer remapping interval the broken resource is also blocked for a longer period no new jobs are accepted and mapped on the meanwhile recovered resource for this duration. Both effects combined reduce the number of accepted jobs, especially, for the time after the breakdown of the resource has ended. Therefore, it is necessary to determine the remapping interval as short as possible.

## D. Definitions

Before we give a formal description of the calculation of the remapping interval some notations are introduced.

As we assume that all resources consist of comparable nodes, the resource usage of jobs, the load and the capacity of the resources is measured in number of nodes. Most values will be normalized by the *total number of nodes*  $\bar{c}$  of all resources within the Grid. For other kinds of resources our algorithm is also applicable if an adequate measurement is given.

The load situation in the Grid of a time slot  $t_0$  is described by the *load profile*  $l_{t_0}(t)$ , with  $t \in \mathbb{N}$ , which is defined as the normalized total number of allocated nodes on all resources for each future time slot  $t_0 + t$ .

The capacity lost due to the breakdown of a resource is denoted by  $c^*$  while the load profile consisting only of jobs allocated to the broken resource is denoted by  $l_{t_0}^*(t)$ . Using this definition, the load profile of the not affected jobs can be defined as  $\hat{l}_{t_0}(t) = l_{t_0}(t) - l_{t_0}^*(t)$ . Accordingly, the set of jobs admitted to the system is denoted by  $J$ , whereas the

set of jobs allocated to the broken resource is denoted by  $J^*$ , and the set of unaffected jobs by  $\tilde{J} = J \setminus J^*$ .

The booking behavior is described by the average booking profile. The *booking profile*  $b_{t_x}(t)$ , with  $t \in \mathbb{N}$ , of a time slot  $t_x$  denotes the normalized number of booked nodes per future time slot  $t_x + t$  of all incoming reservations during this time slot. The average computed over all previous booking profiles is denoted by the *average booking profile*  $\bar{b}(t)$ . These profiles are illustrated in Fig. 5.

### E. Calculation of the Remapping Interval

For the calculation of the remapping interval  $i$ , the load situation in the Grid is taken into account.

For each time slot  $t_0$  this calculation has to be repeated. Initially,  $t_0$  is the moment the failure occurs, but later on the calculation is repeated for each time slot until the resource recovers.




Fig. 5. Determination of the remapping interval based on the combination of the weighted current load profile and the average booking profile.

First, a *weighted combined profile*  $\hat{l}_{t_0, \zeta}(t)$  is created for the current time slot  $t_0$  based on the current load profile of the unaffected jobs  $\tilde{l}_{t_0}(t)$  combined with the weighted load profile of the affected jobs  $\zeta \cdot l_{t_0}^*(t)$  (with  $\zeta \geq 1$ ) and with the average booking profile  $\bar{b}(t)$ . This profile is some kind of expectation of the load after this time slot, as it sums up the currently booked load and the average incoming load during one time slot. Weighting the load of the affected jobs higher than the load of the unaffected jobs results in a narrower profile for small or underloaded resources and in a higher profile in case a heavily loaded, large resource is broken. A *threshold*  $\eta$  with  $\eta \in [0, 1]$  is used to determine the remapping interval  $[t_0, t_0 + i_{\eta, \zeta}(t_0)]$  on the base of the weighted combined profile. The remapping interval is defined by the time after which all values of the weighted combined profile are lower than  $\eta$  (see Fig. 5):

$$i_{\eta, \zeta}(t_0) := \begin{cases} \min \{i \mid \forall t > i : \eta > \hat{l}_{t_0, \zeta}(t)\} & \text{if } > 1 \\ 1 & \text{otherwise} \end{cases}$$

$$\hat{l}_{t_0, \zeta}(t) := \tilde{l}_{t_0}(t) + \zeta l_{t_0}^*(t) + \bar{b}(t)$$

The length of the remapping interval has the lower bound of 1 to ensure, that at least all jobs which are supposed to start within the current time slot are remapped; otherwise they will fail. An algorithmic overview of this calculation is given in Fig. 6.

```

resource failed at  $t_0$ 
// terminate all active jobs
for each job  $j \in J^*$  do
  if  $j.start < t_0$  do
    terminate  $j$ 
while resource down do
  // calculate the profiles
  init  $\hat{l}_{t_0}$ 
  for each job  $j \in J^*$  do
    for each  $t \in [j.start, j.stop]$  do
       $\hat{l}_{t_0}(t) += \zeta j.numberOfNodes$ 
  for each job  $j \in \tilde{J}$  do
    for each  $t \in [j.start, j.stop]$  do
       $\hat{l}_{t_0}(t) += j.numberOfNodes$ 
  add average booking profile  $\bar{b}$  to profile  $\hat{l}_{t_0}$ 
  // calculate  $i$ 
  for  $i = t_{ba}$  to 1 do
    if not  $\hat{l}_{t_0}(i) < \eta$  break
  // remap jobs within remapping interval
  for each  $j \in J^*$  do
    if  $j.start < t_0 + i$  do
      remap  $j$  if possible
  wait for next time slot:  $t_0 = t_0 + 1$ 

```

Fig. 6. Algorithmic overview on the calculation of the remapping interval and the downtime independent remapping.

## V. EVALUATION

In this section, the results of our load-based approach are outlined. In particular, the impact of various choices for  $\eta$  and  $\zeta$  on the terminated jobs as well as on the overall resource utilization is examined in a simulation environment.

### A. Simulation Environment

The simulations were made assuming an infrastructure of several cluster and parallel computers with homogeneous nodes, i.e., each job is capable of running on any of the machines involved with exactly the same speed. If an adequate metric to compare the machines is given our algorithm could cover inhomogeneous Grids as well.

The simulations only serve the purpose of showing the general impact of failures and since according to [10] the actual distribution of job sizes, job durations etc. do not affect the general quality of the results generated even when using simple models, the simulations were made using a simple synthetic job and failure model.

The simulations were done using the simulation mode of the VRM framework. A simulated user reserved the resources in advance with the reservation time being exponentially distributed with a mean of 100 slots. The duration of the jobs were uniformly distributed in the interval [250, 750] and each job demanded for  $2^k$  nodes with  $k$  uniformly distributed in the interval [1, 8].

Furthermore, a central administrative domain controller (ADC) instance was simulated controlling an infrastructure that consisted of different parallel computers with

varying number of compute nodes. In total there were eight machines with 512, 256, 256, 128, 128, 96, 32, and 32 nodes. Obviously, some jobs can only be executed on the larger machines. For each resource a local resource management system capable to support advance reservations was simulated together with an active interface (AI) communicating with the simulated ADC. In our model all jobs were submitted to the ADC and no additional jobs were submitted to the local resource management systems.

An additional component simulated the failures to occur periodically every 1500 slots with a failure duration of always 500 slots. The resource to break down completely was chosen randomly among the 8 machines. To simulate varying load situations we chose different values for the average reservation time  $r$ . Each simulation run had a duration of 20,000 slots and the simulations were repeated until a sufficiently small confidence interval (within  $\pm 5\%$  of the mean with 95% confidence) was reached.

In order to assess the performance of the load-based approach, two metrics were chosen that reflect both the amount of jobs that were affected but could not be successfully remapped onto alternative resources and the number of jobs that had to be rejected because resources are blocked due to the failure recovery. The first metric is the *termination ratio*, which is defined as follows:

$$\text{termination ratio} := \frac{|A^*|}{|A|},$$

with  $A$  being the set of affected jobs and  $A^* \subset A \subset J^*$  (see Sec. IV-D) being the set of unsuccessfully terminated jobs. The second metric is called *request blocking ratio*, defined as

$$\text{request blocking ratio} := \frac{|R|}{|S|},$$

with  $S$  denoting the set of all submitted requests and  $R$  denoting the rejected requests.

As described in previous sections, for the simulations we assume that running jobs cannot be migrated which is the usual case in a high performance computing scenario. However, our model is general enough to cover also the migration of running jobs, in particular this is discussed in Sec. V-D.

### B. Performance of the Load-Based Remapping Approach

In Fig. 7, the termination ratio is depicted for different choices of  $\eta$  and  $\zeta$ . In general, it can be seen that the impact of  $\eta$  is relatively high, whereas  $\zeta$  can only be used for fine-tuning. In this setting, the best results are obtained for  $\zeta = 2$  with  $\eta = 0.775$ . As  $\zeta$  is related to the fragmentation caused by resource boundaries, using larger values of  $\zeta$  does not improve the overall performance. Consequently, the following examinations were conducted using  $\zeta = 2$ .

In Fig. 8, the termination ratio of the load-based remapping approach is depicted only in dependence of  $\eta$ . The different choices of  $r$  were selected to create load situations where the average reservation time is shorter ( $r = 300$ )




Fig. 7. Termination ratio for different choices of  $\eta$  and  $\zeta$  ( $r = 300$ ).

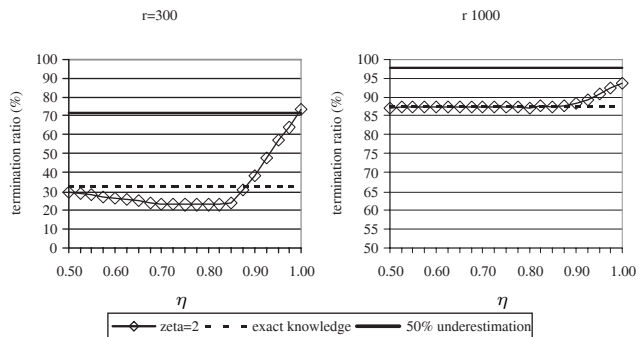


Fig. 8. Termination ratio for different average reservation times.

or longer ( $r = 1000$ ) than the failure duration (500). To compare our novel approach with the estimations done in [4], in Fig. 8 the termination ratio using the exact failure duration to identify the jobs to be remapped (which is a priori unknown) and an underestimation of the failure duration by 50% are depicted.

It can be observed, that for  $r = 300$ , the termination ratio shows a minimum at  $\eta \approx 0.8$  and from thereon rises. The curve for  $r = 1000$  reaches its minimum earlier and then remains stable. As already described in [4], underestimations of the failure duration result in considerable increase of the termination ratio. The reason that the exact knowledge performs worse than the load-based approach results from the fact that the failed resource cannot be locked for the whole duration of the failure using the load-based approach. Hence, additional jobs which must be remapped may be placed onto the resource. When this remapping is successful, it contributes positively to the termination ratio. In contrast, when a resource is locked at the occurrence of the failure for the entire failure time, such jobs will never be placed onto the failed resource.

The self-adaptiveness of our approach can be seen more clearly when investigating the development of the actual remapping interval lengths during a single failure. This is depicted in Fig. 9. The diagram shows the length of the computed remapping intervals during a single failure for a sample period of 32 slots using  $\eta = 0.8$ .




Fig. 9. Snapshot of the remapping interval lengths computed during a single failure ( $\eta = 0.8$ ,  $r = 300$ ).

### C. Impact on the Amount of Accepted Jobs and Utilization




Fig. 10. Request blocking ratio for  $r = 300$  and  $r = 1000$ .

The impact of our approach on the amount of rejected jobs can be observed in Fig. 10: The amount of rejected jobs decreases with growing  $\eta$ . Therefore, the request blocking ratio behaves opposite to the termination ratio. Failure situations are considered as exceptions from the normal operation and hence the impact on the overall request blocking ratio is limited. In particular, selecting a small value for  $\eta$  does not impact the request blocking ratio and hence, selecting  $\eta \leq 0.6$ , as also indicated in the previous section, can be considered as a reasonable choice.

### D. Comparison with Job Migrations

Although not assumed for the previous examinations, migration of running jobs is an interesting feature, allowing the management system to provide an almost completely transparent failure recovery. In the most extreme case, jobs may be repeatedly migrated throughout their life-time. One possibility to support job migration is checkpointing. In this case, the failure recovery also needs to take care of the transfer of the checkpointed data to the backup resource and of the correct restart of the job at the new location. Moreover, since checkpointing is only performed at certain intervals, the required rollback to the last checkpoint increases the run-time of the job.

In Fig. 11, the possible benefit of checkpointing and job migration is outlined. The diagram shows the common approach, i.e., jobs can only be transferred as a whole and active jobs cannot be remapped. In contrast, providing

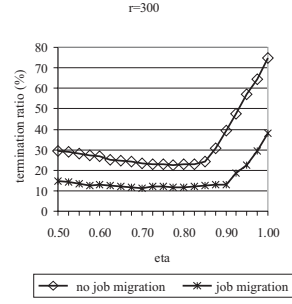


Fig. 11. Performance benefit using checkpointing and migration

means for migration of running jobs has a significant benefit as the termination ratio can be significantly reduced, i.e., around 50% more jobs can be successfully remapped to other resources. The diagram shows an idealized and unrealistic situation where data transfer over the network and rollbacks to previous checkpoints do not require time. However, the potential of using job migrations is apparent.

## VI. RELATED WORK

Advance reservations are an important allocation strategy, widely used, e.g., in Grid toolkits such as Globus [11], as they provide simple means for planning of resources and in particular co-allocations of different resources. Besides flexible and easy support for co-allocations, e.g., in case complex workflows need to be processed, advance reservations also have other advantages such as an increased admission probability when reserving sufficiently early, and reliable planning for users and operators. Support for advance reservations has been integrated into several management systems for distributed and parallel computing [3], [12]. In [2], advance reservations have been identified as essential for a number of higher level services, such as SLAs.

In the context of Grid computing, failure recovery mechanisms are particularly important as the distributed nature of the environment requires more sophisticated mechanisms than needed in a setting with only few resources that can be handled by a central management system.

In general, failure detection and recovery mechanisms focus on the requirements to deal with applications that are already active. The Globus heartbeat monitor HBM [3] provides mechanisms to notify applications or users of failures occurring on the used resources. The recovery mechanisms described in this paper can be initiated by the failure detection of the HBM. In [13], a framework for handling failures in Grid environments was presented, based on the workflow structure. The framework allows users to select different failure recovery mechanisms, such as simply restarting jobs, or - more sophisticated - checkpointing and migration to other resources if supported by the application to be recovered. The different recovery mechanisms are discussed and compared. However, the framework can only be used for the recovery of active ap-

plications, inactive applications that are already assigned to resources but not yet started are not taken into account.

In [4], the basic requirements and opportunities for failure recovery in planning based resource management systems have been examined. In particular, it was shown that remapping of admitted but not yet active jobs is essential in order to reduce the number of terminated jobs. It was also shown that the best results in terms of termination probability and overall resource utilization are achieved, when exact knowledge of the actual duration of a failure is available and any jobs commencing during this interval are remapped. However, estimations of the actual downtime are a questionable approach as these estimations are inherently unreliable and underestimations lead to a significantly higher termination ratio than possible with exact knowledge. In [14], we extend the approach of remapping in advance from [4] and introduce the concept to repeatedly calculate the remapping interval. This approach is extended here to be applicable in the Virtual Resource Manager (VRM)[2].

## VII. CONCLUSION AND FUTURE WORK

In this paper, the novel load-based failure recovery strategy used in the VRM framework was presented. The mechanism is applicable in any environment where distributed resources must be managed and failures of the system are critical, e.g., SLAs are given for the correct and complete execution of a job. In particular, co-allocation environments such as Grids are target environments for our strategy.

The load-based algorithm is based on previous work on this field which showed, that estimations of the actual downtime of a resource have a particularly negative impact on the termination ratio. Consequently, our approach adapts to the actual load situation and determines a remapping interval accordingly, which diminishes the danger of underestimating failure durations as any job is remapped before it is actually endangered of being terminated. Simulations showed how to select appropriate values for the used parameters to gain the best performance of the failure recovery.

The strategy presented in this paper is generic, i.e., it can easily be applied to almost any resource type and any resource management system. This is particularly important for next generation Grid systems, which essentially need to support higher level quality-of-service guarantees, e.g., specified by SLAs, as in the context of the VRM [2].

Future work will deal with the possibility to integrate checkpointing and migration mechanisms into the load-based approach, which has the potential to dramatically increase the performance of the failure recovery mechanisms as presented before. This approach will allow to more efficiently utilize the available resources as gaps due to temporal fragmentation can be filled. Important issues in this context are the time required for job migration over a network. Moreover, a completely automated selection of both parameters will be developed simplifying the deployment of the system. However, as the simulations conducted

for this paper indicate, a static choice is also reasonable and leads to good results.

## REFERENCES

- [1] Czajkowski, K., I. Foster, C. Kesselman, V. Sander, and S. Tuecke, "SNAP: A Protocol for Negotiating Service Level Agreements and Coordinating Resource Management in Distributed Systems," in *8th Intl. Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP)*, Edinburgh, Scotland, UK, ser. Lecture Notes in Computer Science (LNCS), vol. 2537. Springer, January 2002, pp. 153–183.
- [2] Burchard, L.-O., M. Hovestadt, O. Kao, A. Keller, and B. Linnert, "The Virtual Resource Manager: An Architecture for SLA-aware Resource Management," in *4th Intl. IEEE/ACM Intl. Symposium on Cluster Computing and the Grid (CCGrid)*, Chicago, USA, 2004.
- [3] "The Globus Project," <http://www.globus.org/>.
- [4] Burchard, L.-O. and B. Linnert, "Failure Recovery in Distributed Environments with Advance Reservation Management Systems," in *15th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM)*, Davis, USA, ser. Lecture Notes in Computer Science (LNCS), vol. 3278. Springer, 2004, to appear.
- [5] Litzkow, M., T. Tannenbaum, J. Basney, and M. Livny, "Checkpoint and Migration of UNIX Processes in the Condor Distributed Processing System," University of Wisconsin - Madison Computer Sciences Department, Tech. Rep. UW-CS-TR-1346, April 1997.
- [6] DeFanti, T., C. de Laat, J. Mambretti, K. Neggers, and B. S. Arnaud, "TransLight: A Global-Scale LambdaGrid for E-Science," *Communications of the ACM*, vol. 46, no. 11, pp. 34–41, November 2003.
- [7] Keller, A., and A. Reinefeld, "Anatomy of a Resource Management System for HPC Clusters," in *Annual Review of Scalable Computing, vol. 3*, Singapore University Press, 2001, pp. 1–31.
- [8] Heine, F., M. Hovestadt, and O. Kao, "Towards Ontology-Driven P2P Grid Resource Discovery," in *5th IEEE/ACM International Workshop on Grid Computing*, to appear, 2004.
- [9] Garey, M. and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.
- [10] Lo, V., J. Mache, and K. Windisch, "A Comparative Study of Real Workload Traces and Synthetic Workload Models for Parallel Job Scheduling," in *4th Workshop on Job Scheduling Strategies for Parallel Processing*, Orlando, USA, ser. Lecture Notes in Computer Science (LNCS), vol. 1459. Springer, 1998, pp. 25–46.
- [11] Foster, I., C. Kesselman, C. Lee, R. Lindell, K. Nahrstedt, and A. Roy, "A Distributed Resource Management Architecture that Supports Advance Reservations and Co-Allocation," in *7th International Workshop on Quality of Service (IWQoS)*, London, UK, 1999, pp. 27–36.
- [12] Snell, D., M. Clement, D. Jackson, and C. Gregory, "The Performance Impact of Advance Reservation Meta-scheduling," in *6th Workshop on Job Scheduling Strategies for Parallel Processing*, Cancun, Mexico, ser. Lecture Notes in Computer Science (LNCS), vol. 1911. Springer, 2000, pp. 137–153.
- [13] Hwang, S. and C. Kesselman, "Grid Workflow: A Flexible Failure Handling Framework for the Grid," in *12th Intl. Symposium on High Performance Distributed computing (HPDC)*, Seattle, USA. IEEE, 2003, pp. 126–138.
- [14] L.-O. Burchard, B. Linnert, and J. Schneider, "A distributed load-based failure recovery mechanism for advance reservation environments," in *5th ACM/IEEE Intl. Symposium on Cluster Computing and the Grid (CCGrid)*, 2005.