

# WILEY

*Publishers Since 1807*

THE ATRIUM, SOUTHERN GATE, CHICHESTER, WEST SUSSEX P019 8SQ

**\*\*\*IMMEDIATE RESPONSE REQUIRED\*\*\***

Your article may be published online via Wiley's EarlyView® service ([www.interscience.wiley.com](http://www.interscience.wiley.com)) shortly after receipt of corrections. EarlyView® is Wiley's online publication of individual articles in full text HTML and/or PDF format before release of the compiled print issue of the journal. Articles posted online in EarlyView® are peer-reviewed, copyedited, author corrected, and fully citable via the article DOI (for further information, visit [www.doi.org](http://www.doi.org)). EarlyView® means you benefit from the best of two worlds--fast online availability as well as traditional, issue-based archiving.

Please follow these instructions to avoid delay of publication

**READ PROOFS CAREFULLY**

- This will be your only chance to review these proofs. **Please note that once your corrected article is posted online, it is considered legally published, and cannot be removed from the Web site for further corrections.**
- Please note that the volume and page numbers shown on the proofs are for position only.

**ANSWER ALL QUERIES ON PROOFS** (Queries for you to answer are attached as the last page of your proof.)

- List all corrections and send back via e-mail, or mark all corrections directly on the proofs and send the scanned copy via e-mail. Please do not send corrections by fax or in the post.
- E-mail corrections to: [cpeproofs@wiley.co.uk](mailto:cpeproofs@wiley.co.uk)

**CHECK FIGURES AND TABLES CAREFULLY**

- Check size, numbering, and orientation of figures.
- All images in the PDF are downsampled (reduced to lower resolution and file size) to facilitate Internet delivery. These images will appear at higher resolution and sharpness in the printed article.
- Review figure legends to ensure that they are complete.
- Check all tables. Review layout, title, and footnotes.

**COMPLETE CTA (if you have not already signed one)**

- Please send a scanned copy with your proofs and post your completed original form to the address below. **We cannot publish your paper until we receive the original signed form.**
- Post to: Tim Williams, CPE Project Manager, Sunrise Setting Ltd, 12A Fore Street, St Marychurch, Torquay, Devon TQ1 4NE, UK

**COMPLETE OFFPRINT ORDER FORM**

- Fill out the attached offprint order form. It is important to return the form even if you are not ordering offprints. You may, if you wish, pay for the offprints with a credit card. Offprints will be mailed only after your article appears in print. This is the most opportune time to order reprints. If you wait until after your article has been published, the offprints will be considerably more expensive.
- E-mail to: [offprints@wiley.co.uk](mailto:offprints@wiley.co.uk)

**RETURN**

- PROOFS**
- OFFPRINT ORDER FORM**
- CTA (If you have not already signed one)**

**RETURN IMMEDIATELY AS YOUR ARTICLE MAY BE POSTED ONLINE SHORTLY AFTER RECEIPT**

# WILEY AUTHOR DISCOUNT CARD

As a highly valued contributor to Wiley's publications, we would like to show our appreciation to you by offering a **unique 25% discount** off the published price of any of our books\*.

To take advantage of this offer, all you need to do is apply for the **Wiley Author Discount Card** by completing the attached form and returning it to us at the following address:

The Database Group  
John Wiley & Sons Ltd  
The Atrium  
Southern Gate  
Chichester  
West Sussex PO19 8SQ  
UK

In the meantime, whenever you order books direct from us, simply quote promotional code **SOO1W** to take advantage of the 25% discount.

The newest and quickest way to order your books from us is via our new European website at:

**<http://www.wileyeurope.com>**

Key benefits to using the site and ordering online include:

- Real-time SECURE on-line ordering
- The most up-to-date search functionality to make browsing the catalogue easier
- Dedicated Author resource centre
- E-mail a friend
- Easy to use navigation
- Regular special offers
- Sign up for subject orientated e-mail alerts

So take advantage of this great offer, return your completed form today to receive your discount card.

Yours sincerely,



Verity Leaver  
E-marketing and Database Manager

#### \*TERMS AND CONDITIONS

This offer is exclusive to Wiley Authors, Editors, Contributors and Editorial Board Members in acquiring books (excluding encyclopaedias and major reference works) for their personal use. There must be no resale through any channel. The offer is subject to stock availability and cannot be applied retrospectively. This entitlement cannot be used in conjunction with any other special offer. Wiley reserves the right to amend the terms of the offer at any time.

# REGISTRATION FORM FOR 25% BOOK DISCOUNT CARD

To enjoy your special discount, tell us your areas of interest and you will receive relevant catalogues or leaflets from which to select your books. Please indicate your specific subject areas below.

<p><b>Accounting</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Public <span style="float: right;">[ ]</span></li> <li>• Corporate <span style="float: right;">[ ]</span></li> </ul>	<p><b>Architecture</b> <span style="float: right;">[ ]</span></p>
<p><b>Chemistry</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Analytical <span style="float: right;">[ ]</span></li> <li>• Industrial/Safety <span style="float: right;">[ ]</span></li> <li>• Organic <span style="float: right;">[ ]</span></li> <li>• Inorganic <span style="float: right;">[ ]</span></li> <li>• Polymer <span style="float: right;">[ ]</span></li> <li>• Spectroscopy <span style="float: right;">[ ]</span></li> </ul>	<p><b>Business/Management</b> <span style="float: right;">[ ]</span></p>
<p><b>Encyclopedia/Reference</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Business/Finance <span style="float: right;">[ ]</span></li> <li>• Life Sciences <span style="float: right;">[ ]</span></li> <li>• Medical Sciences <span style="float: right;">[ ]</span></li> <li>• Physical Sciences <span style="float: right;">[ ]</span></li> <li>• Technology <span style="float: right;">[ ]</span></li> </ul>	<p><b>Computer Science</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Database/Data Warehouse <span style="float: right;">[ ]</span></li> <li>• Internet Business <span style="float: right;">[ ]</span></li> <li>• Networking <span style="float: right;">[ ]</span></li> <li>• Programming/Software Development <span style="float: right;">[ ]</span></li> <li>• Object Technology <span style="float: right;">[ ]</span></li> </ul>
<p><b>Earth &amp; Environmental Science</b> <span style="float: right;">[ ]</span></p>	<p><b>Engineering</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Civil <span style="float: right;">[ ]</span></li> <li>• Communications Technology <span style="float: right;">[ ]</span></li> <li>• Electronic <span style="float: right;">[ ]</span></li> <li>• Environmental <span style="float: right;">[ ]</span></li> <li>• Industrial <span style="float: right;">[ ]</span></li> <li>• Mechanical <span style="float: right;">[ ]</span></li> </ul>
<p><b>Hospitality</b> <span style="float: right;">[ ]</span></p>	<p><b>Finance/Investing</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Economics <span style="float: right;">[ ]</span></li> <li>• Institutional <span style="float: right;">[ ]</span></li> <li>• Personal Finance <span style="float: right;">[ ]</span></li> </ul>
<p><b>Genetics</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Bioinformatics/Computational Biology <span style="float: right;">[ ]</span></li> <li>• Proteomics <span style="float: right;">[ ]</span></li> <li>• Genomics <span style="float: right;">[ ]</span></li> <li>• Gene Mapping <span style="float: right;">[ ]</span></li> <li>• Clinical Genetics <span style="float: right;">[ ]</span></li> </ul>	<p><b>Life Science</b> <span style="float: right;">[ ]</span></p>
<p><b>Medical Science</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Cardiovascular <span style="float: right;">[ ]</span></li> <li>• Diabetes <span style="float: right;">[ ]</span></li> <li>• Endocrinology <span style="float: right;">[ ]</span></li> <li>• Imaging <span style="float: right;">[ ]</span></li> <li>• Obstetrics/Gynaecology <span style="float: right;">[ ]</span></li> <li>• Oncology <span style="float: right;">[ ]</span></li> <li>• Pharmacology <span style="float: right;">[ ]</span></li> <li>• Psychiatry <span style="float: right;">[ ]</span></li> </ul>	<p><b>Landscape Architecture</b> <span style="float: right;">[ ]</span></p>
<p><b>Non-Profit</b> <span style="float: right;">[ ]</span></p>	<p><b>Mathematics/Statistics</b> <span style="float: right;">[ ]</span></p>
	<p><b>Manufacturing</b> <span style="float: right;">[ ]</span></p>
	<p><b>Material Science</b> <span style="float: right;">[ ]</span></p>
	<p><b>Psychology</b> <span style="float: right;">[ ]</span></p> <ul style="list-style-type: none"> <li>• Clinical <span style="float: right;">[ ]</span></li> <li>• Forensic <span style="float: right;">[ ]</span></li> <li>• Social &amp; Personality <span style="float: right;">[ ]</span></li> <li>• Health &amp; Sport <span style="float: right;">[ ]</span></li> <li>• Cognitive <span style="float: right;">[ ]</span></li> <li>• Organizational <span style="float: right;">[ ]</span></li> <li>• Developmental and Special Ed <span style="float: right;">[ ]</span></li> <li>• Child Welfare <span style="float: right;">[ ]</span></li> <li>• Self-Help <span style="float: right;">[ ]</span></li> </ul>
	<p><b>Physics/Physical Science</b> <span style="float: right;">[ ]</span></p>

[ ] I confirm that I am a Wiley Author/Editor/Contributor/Editorial Board Member of the following publications:

SIGNATURE: .....

**PLEASE COMPLETE THE FOLLOWING DETAILS IN BLOCK CAPITALS:**

TITLE AND NAME: (e.g. Mr, Mrs, Dr) .....

JOB TITLE: .....

DEPARTMENT: .....

COMPANY/INSTITUTION: .....

ADDRESS: .....

.....

.....

.....

01

# 02 Self adaptivity in 03 Grid computing 04

05

06

Sathish S. Vadhiyar<sup>1,\*,\dagger</sup> and Jack J. Dongarra<sup>2,3</sup>

08

09

<sup>1</sup>Supercomputer Education and Research Centre, Indian Institute of Science,  
Bangalore 560012, India

10

11

<sup>2</sup>Computer Science Department, University of Tennessee, Knoxville,  
TN 37996-3450, U.S.A.

12

13

<sup>3</sup>Computer Science and Mathematics, Oak Ridge National Laboratory, U.S.A.

14

15

16

## SUMMARY

17

18

19

20

21

22

23

24

25

26

27

Optimizing a given software system to exploit the features of the underlying system has been an area of research for many years. Recently, a number of self-adapting software systems have been designed and developed for various computing environments. In this paper, we discuss the design and implementation of a software system that dynamically adjusts the parallelism of applications executing on computational Grids in accordance with the changing load characteristics of the underlying resources. The migration framework implemented by our software system is aimed at performance-oriented Grid systems and implements tightly coupled policies for both suspension and migration of executing applications. The suspension and migration policies consider both the load changes on systems as well as the remaining execution times of the applications thereby taking into account both system load and application characteristics. The main goal of our migration framework is to improve the response times for individual applications. We also present some results that demonstrate the usefulness of our migration framework. Published in 2005 by John Wiley & Sons, Ltd.

28

29

KEY WORDS: self adaptivity; migration; GrADS; rescheduling; redistribution; checkpointing

30

31

32

## 1. INTRODUCTION

33

34

35

36

37

Optimization of software routines for achieving efficiency on a given computational environment has been an active area of research. Historically, the optimization was achieved by hand-tuning the software system to fit the needs of the computing environment. Although high optimization can be achieved,

38

39



40

41

42

\*Correspondence to: Sathish S. Vadhiyar, Supercomputer Education and Research Centre, Indian Institute of Science, Bangalore 560012, India

<sup>\dagger</sup>E-mail: vss@serc.iisc.ernet.in

43

44

Contract/grant sponsor: U.S. Department of Energy; contract/grant number: DE-AC05-00OR22725

Contract/grant sponsor: National Science Foundation; contract/grant number: ACI 0103759



01 this process was found to be tedious and needs considerable scientific expertise. Also, the hand-  
02 tuning process was not portable across different computing environments. Finally, hand customization  
03 does not take into account the run-time load dynamics of the system and the input parameters of the  
04 application.

05 The solution to the above-mentioned problems associated with hand-tuning software routines for the  
06 computing environment is to build *self-adaptive software system* that examines the characteristics of  
07 the computing environments and chooses the software parameters needed to achieve high efficiency  
08 on that environment. Recently, a number of self-adaptive software systems have been designed  
09 and implemented [1–6]. Some of the software systems apply adaptivity to the computational  
10 processors [1,2], some are tuned for communication networks [3], some are intended for workstation  
11 clusters [5] and some have been developed for computational Grids [6]. The various adaptive software  
12 systems also differ in the time when adaptivity is performed. Some perform adaptivity at installation  
13 time [2–4], while others perform adaptivity at run time [5,6].

14 There are very few self-adaptive software systems that dynamically adapt to changes in the load  
15 characteristics of the resources on computational Grids. Computational Grids [7] involve large resource  
16 dynamics, so the ability to migrate executing applications onto different sets of resources assumes great  
17 importance. Specifically, the main motivations for migrating applications in Grid systems are to provide  
18 fault tolerance and to adapt to load changes on the systems. In this paper, we focus on the migration  
19 of applications executing on distributed and Grid systems in order to adapt to the load dynamics of the  
20 resources.

21 There are at least two disadvantages in using the existing migration frameworks [8–13] for adapting  
22 to load dynamics. First, due to separate policies employed by these migration systems for suspension  
23 of executing applications and migration of the applications to different systems, applications can incur  
24 lengthy waiting times between when they are suspended and when they are restarted on new systems.  
25 Second, due to the use of predefined conditions for suspension and migration and due to the lack of  
26 knowledge of the remaining execution time of the applications, the applications can be suspended and  
27 migrated even when they are about to finish execution in a short period of time. This is certainly less  
28 desirable in performance-oriented Grid systems where the large load dynamics may lead to frequent  
29 satisfaction of the predefined conditions and hence could lead to frequent invocations of suspension  
30 and migration decisions.

31 In this paper, we describe a framework that defines and implements scheduling policies for migrating  
32 applications executing on distributed and Grid systems in response to varying resource load dynamics.  
33 In our framework, the migration of applications depends on

- 34 (1) the amount of increase or decrease in loads on the resources;
  - 35 (2) the point during the application execution lifetime when load is introduced into the system;
  - 36 (3) the performance benefits that can be obtained for the application due to migration.
- 37  
38  
39  
40  
41  
42  
43  
44



01 The framework has been implemented and tested in the GrADS system [6]. Our test results indicate  
02 that our migration framework can help improve the performance of executing applications by more  
03 than 30%. In this paper, we present some of the descriptions and results from our earlier work [14] and  
04 also present new experiments regarding dynamic determination of rescheduling cost.

05 In Section 2, we present a general overview of self-adaptive software systems by describing some  
06 systems that perform adaptivity. In Section 3, we describe the GrADS system and the life cycle of  
07 GrADS applications. In Section 4, we introduce our migration framework by describing the different  
08 components in the framework. In Section 5, we describe the API of the checkpointing library used  
09 in our migration framework. In Section 6, the various policies regarding rescheduling are dealt with.  
10 In Section 7, other issues relevant to migration are described in brief. In Section 8, we describe our  
11 experiments and provide various results. In Section 9, we present related work in the field of migration.  
12 We give concluding remarks and explain our future plans in Section 10.

13

14

## 15 2. SELF-ADAPTIVE SOFTWARE SYSTEMS—AN OVERVIEW

16

17 Recently, there have been a number of efforts in designing and developing self-adaptive software  
18 systems. These system differ in terms of the kind of computational environments, the kind of adaptive  
19 software system used and also the time when adaptivity is performed. The following sections describe  
20 some illustrative examples.

21

### 22 2.1. ATLAS

23

24 ATLAS [2] stands for Automatically Tuned Linear Algebra Software. ATLAS exploits cache locality  
25 to provide highly efficient implementations of BLAS (Basic Linear Algebra Subroutine) and few  
26 LAPACK routines. During installation, ATLAS studies various characteristics of the hardware  
27 including the size of the cache, the number of floating point units in the machine and the pipeline  
28 length to determine the optimal or near-optimal block size for the dense matrices, the number of  
29 loop unrollings to perform, the kind of instruction sets to use, etc. Thus, optimizations are performed  
30 for reducing the number of accesses to main memory and reduce loop overheads resulting in BLAS  
31 implementations that are competitive with the machine-specific versions of most known architectures.

32

### 33 2.2. ATCC

34

35 ATCC [3] (Automatically Tuned Collective Communications) is intended for optimizing MPI [15,16]  
36 collective communications for a given set of machines connected by networks of specific  
37 configurations. The collective communication routines form integral parts of most of the MPI-based  
38 parallel applications. During installation, ATCC conducts experiments for different algorithms and  
39 segment sizes for different collective communications, number of processors and message sizes.  
40 ATCC then gathers the times for individual experiments in a look-up table. When the user invokes a  
41 collective communication routine with a given message size and a given number of processors, ATCC  
42 looks up the table and chooses the best collective communication algorithm and segment size for  
43 communication. Recent versions of ATCC include performance models for collective communication  
44 algorithms to reduce the time taken for conducting actual experiments.



### 01 2.3. BeBOP

02  
03 The BeBOP project from Berkeley attempts to optimize sparse matrix kernels, namely, matrix-  
04 vector multiplication, triangular solve and matrix triple product for a given architecture. For each of  
05 the sparse matrix kernels, the BeBOP project considers a set of implementations and chooses the  
06 optimal or near-optimal implementation for a given architecture. Given a sparse matrix, machine,  
07 and kernel, the BeBOP approach in choosing an implementation consists of two steps. First, the  
08 possible implementations are benchmarked off-line in a matrix-independent, machine-dependent way.  
09 When the matrix structure is known during runtime, the matrix is sampled to extract relevant aspects  
10 of its structure, and performance models that combine the benchmark data and the estimated matrix  
11 properties are evaluated to obtain the near-optimal implementation. The BeBOP [4] approach has been  
12 successfully applied to optimize register blocking for sparse matrix-vector multiplication.

### 14 2.4. LFC

15  
16 The LFC (LAPACK for Clusters) project [5] aims to simplify the use of parallel linear algebra software  
17 on computational clusters. Benchmark results are obtained for sequential kernels that are invoked by  
18 the parallel software. During run-time, adaptivity is performed by taking into account the resource  
19 characteristics of the computational machines and an optimal or near-optimal choice of a subset of  
20 resources for the execution of the parallel application is made by the employment of scheduling  
21 algorithms. LFC also optimizes the parameters of the problem, namely the block size of the matrix.  
22 LFC is intended for the remote invocation of parallel software from a sequential environment and  
23 hence employs data movement strategies. The LFC approach has been successfully used for solving  
24 ScaLAPACK LU, QR and Cholesky factorization routines.

## 27 3. THE GrADS SYSTEM

28  
29 GrADS (Grid Application Development Software) [6] is an ongoing research project involving a  
30 number of institutions and its goal is to simplify distributed heterogeneous computing in the same way  
31 that the World Wide Web simplified information sharing over the Internet. GrADS approach is similar  
32 to the LFC approach, but more suited to Grid computing due to the employment of Grid computing  
33 tools. The University of Tennessee investigates issues regarding integration of numerical libraries in the  
34 GrADS system. In our previous work [17], we demonstrated the ease with which numerical libraries  
35 such as ScaLAPACK can be integrated into the Grid system and the ease with which the libraries can  
36 be used over the Grid. We also showed some results to prove the usefulness of a Grid in solving large  
37 numerical problems.

38 In the architecture of GrADS, the user wanting to solve a numerical application over the Grid invokes  
39 the GrADS application manager. The life cycle of the GrADS application manager is shown in Figure 1.

40 As a first step, the application manager invokes a component called Resource Selector. The Resource  
41 Selector accesses the Globus Monitoring and Discovery Service (MDS) [18] to obtain a list of  
42 machines in the GrADS testbed that are available and then contacts the Network Weather Service  
43 (NWS) [19] to retrieve system information for the machines. The application manager then invokes a  
44 component called Performance Modeler with problem parameters, machines and machine information.





01  
02  
03  
04  
05  
06  
07  
08  
09  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44

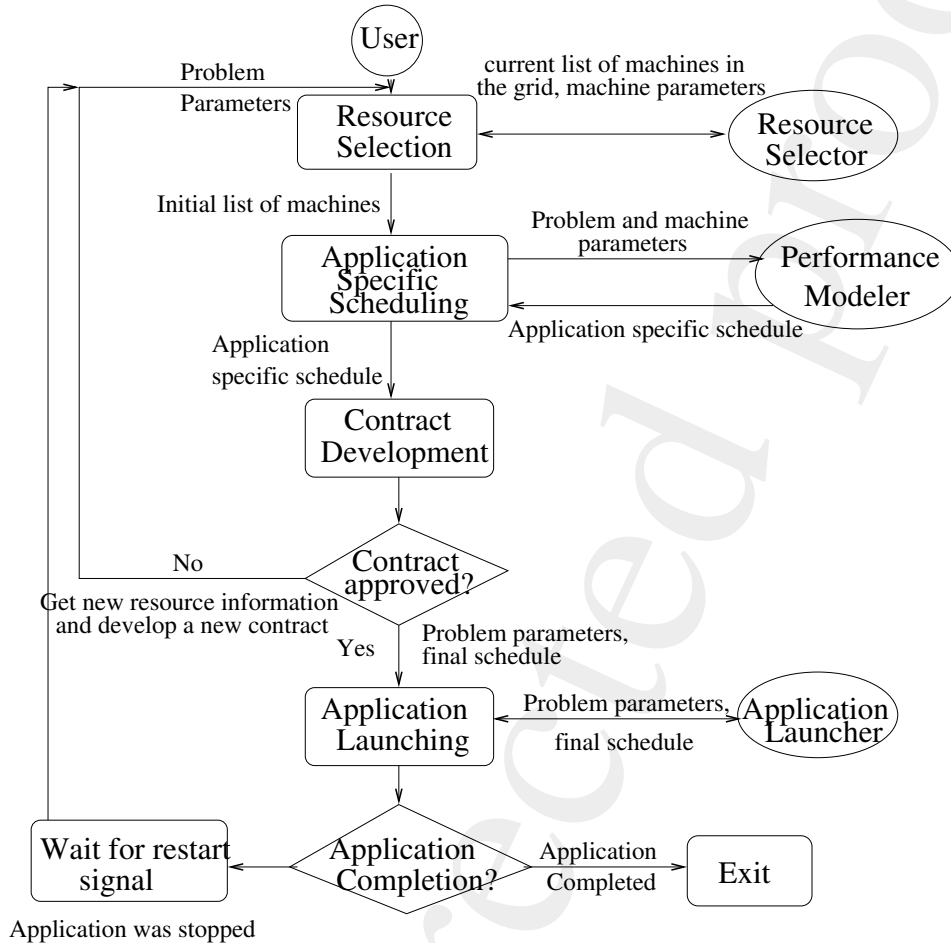


Figure 1. GrADS application manager.

The Performance Modeler, using an execution model built specifically for the application, determines the final list of machines for application execution. By employing an application-specific execution model, GrADS follows the AppLeS [20] approach to scheduling. The problem parameters and the final list of machines are passed as a contract to a component called Contract Developer. The Contract Developer may either approve or reject the contract. If the contract is rejected, the application manager develops a new contract by starting from the resource selection phase again. If the contract is approved, the application manager passes the problem, its parameters and the final list of machines to Application Launcher. The Application Launcher spawns the job on the given machines



01 using Globus job-management mechanism and also spawns a component called Contract Monitor.  
02 The Contract Monitor through an Autopilot mechanism [21] monitors the times taken for different  
03 parts of applications. The GrADS architecture also has a GrADS Information Repository (GIR) that  
04 maintains the different states of the application manager and the states of the numerical application.  
05 After spawning the numerical application through the Application Launcher, the application manager  
06 waits for the job to complete. The job can either complete or suspend its execution due to external  
07 intervention. These application states are passed to the application manager through the GIR. If the job  
08 has completed, the application manager exits, passing success values to the user. If the application is  
09 stopped, the application manager waits for the state of the end application to change to 'RESUME' and  
10 then collects new machine information by starting from the resource selection phase again.

11

12

#### 13 4. THE MIGRATION FRAMEWORK AND SELF ADAPTIVITY

14

15 Although the GrADS architecture explained in the previous section has provisions for continuing an  
16 end application after the application was stopped, it lacks components that perform the actual stopping  
17 of the executing end application and informing the application manager of the various states of the  
18 end application. Hence, the GrADS architecture as described in the previous section does not adapt the  
19 executing application to the changing resource characteristics once the application is committed to a set  
20 of resources. It is highly desirable to adapt and migrate the application to a different set of resources if  
21 the resources on which the application is executing do not meet the performance criteria. The ability to  
22 migrate applications in the GrADS system is implemented by adding a component called *Rescheduler*  
23 to the GrADS architecture. The migrating numerical application, *Migrator*, the *Contract Monitor* that  
24 monitors the application's progress and the *Rescheduler* that decides when to migrate, together form  
25 the core of the migrating framework. The interactions among the different components involved in  
26 the migration framework is illustrated in Figure 2. These components are described in detail in the  
27 following sections.

28

##### 29 4.1. Migrator

30

31 A user-level checkpointing library called SRS (stop restart software) is used to provide migration  
32 capability to the end application. The application, by making calls to the SRS API, achieves the  
33 ability to checkpoint data, to be stopped at a particular point in execution, to be restarted later on  
34 a different configuration of processors and to be continued from the previous point of execution.  
35 The SRS library is implemented on top of MPI and hence can be used only with MPI-based parallel  
36 programs. Since checkpointing in SRS is implemented at the application layer and not at the MPI layer,  
37 migration is achieved by clean exit of the entire application and restarting the application on a new  
38 configuration of resources. Although the method of rescheduling in SRS, by stopping and restarting  
39 executing applications, incurs more overhead than process migration techniques [22–24] where a single  
40 process or a set of processes of the application is either migrated to another processor or replaced by  
41 a set of processes, the approach followed by SRS allows reconfiguration of executing applications and  
42 achieves portability across different MPI implementations, particularly MPICH-G [25], a popular MPI  
43 implementation for Grid computing. The SRS library uses Internet Backplane Protocol (IBP) [26] for  
44 storage of the checkpoint data. IBP storage depots are started on all the machines in the GrADS testbed.

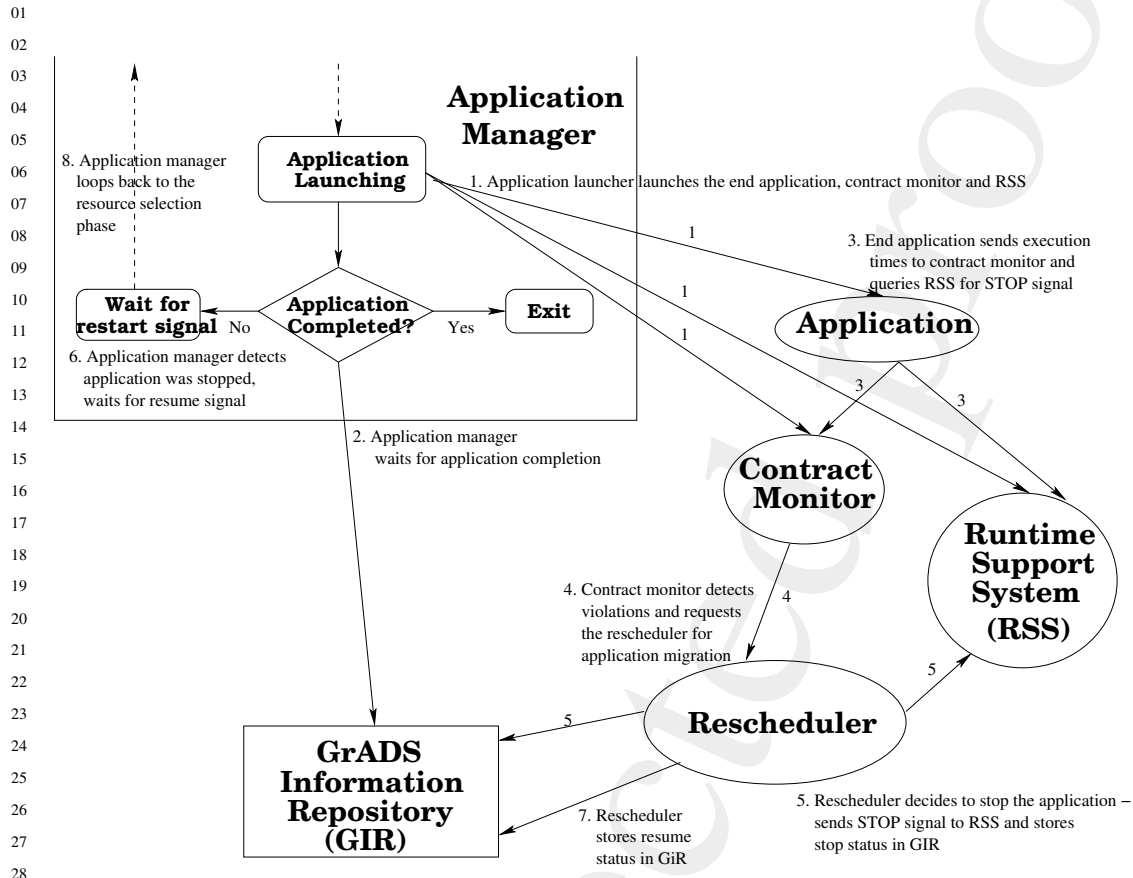


Figure 2. Interactions in migration framework.

The application launcher, apart from launching the end application and the contract monitor, also launches a component called RSS (run-time support system). RSS is included as part of the SRS checkpointing package. An external component (e.g. the rescheduler) wanting to stop an executing end application interacts with the RSS daemon. RSS exists for the entire duration of the application and spans across multiple migrations of the application. Before the actual parallel application is started, the RSS daemon is launched by the application launcher on the machine where the user invokes the GrADS application manager. The actual application through the SRS library knows the location of the RSS from the GIR and interacts with RSS to perform various functions. These functions include initialization of certain data structures in the library, whether the application needs to be stopped and storing and retrieving various information including pointers to the checkpointed data, processor configuration and data distribution used by the application. RSS is implemented as a threaded service that receives asynchronous requests from external components and the application.



---

## 01 4.2. Contract Monitor

02  
03 The Contract Monitor is a component that uses the Autopilot infrastructure to monitor the progress of  
04 applications in GrADS. Autopilot [21] is a real-time adaptive control infrastructure built by the Pablo  
05 group at University of Illinois, Urbana-Champaign. An autopilot manager is started before the launch of  
06 the numerical application. The numerical application is instrumented with calls to send the execution  
07 times taken for the different phases of the application to the contract monitor. The contract monitor  
08 compares the actual execution times with the predicted execution times. When the contract monitor  
09 detects large differences between the actual and the predicted performance of the end application, it  
10 contacts the rescheduler and requests it to migrate the application.

## 12 4.3. Rescheduler

13  
14 Rescheduler is the component that evaluates the performance benefits that can be obtained due to  
15 the migration of an application and initiates the migration of the application. The rescheduler is a  
16 daemon that operates in two modes: *migration on request* and *opportunistic migration*. When the  
17 contract monitor detects intolerable performance loss for an application, it contacts the rescheduler  
18 requesting it to migrate the application. This is called migration on request. In other cases when any  
19 contract monitor has not contacted the rescheduler for migration, the rescheduler periodically queries  
20 the GrADS Information Repository (GIR) for recently completed applications. If a GrADS application  
21 was recently completed, the rescheduler determines whether performance benefits can be obtained for  
22 an currently executing application by migrating it to use the resources that were freed by the completed  
23 application. This is called opportunistic rescheduling.

## 26 5. THE SRS API

27  
28 The application interfaces for SRS look similar to CUMULVS [27], but unlike CUMULVS, SRS  
29 does not require a PVM virtual machine to be set up on the hosts. The SRS library consists of six  
30 main functions: SRS\_Init(), SRS\_Finish(), SRS\_Restart\_Value(), SRS\_Check\_Stop(), SRS\_Register()  
31 and SRS\_Read(). The user calls SRS\_Init() and SRS\_Finish() in their application after MPI\_Init() and  
32 before MPI\_Finalize(), respectively. Since SRS is a user-level checkpointing library, the application  
33 may contain conditional statements to execute certain parts of the application in the start mode and  
34 certain other parts in the restart mode. In order to know whether the application is executed in the  
35 start or restart mode, the user calls SRS\_Restart\_Value() that returns zero and one on start and restart  
36 modes, respectively. The user also calls SRS\_Check\_Stop() at different phases of the application to  
37 check whether an external component wants the application to be stopped. If the SRS\_Check\_Stop()  
38 returns one, then the application has received a stop signal from an external component and hence  
39 should perform application-specific stop actions. There is no relationship between the locations of  
40 the SRS\_Check\_Stop() calls and the calls to extract the execution times of the different phases of  
41 application.

42 The user calls SRS\_Register() in his application to register the variables that will be checkpointed  
43 by the SRS library. When an external component stops the application, the SRS library checkpoints  
44 only those variables that were registered through SRS\_Register(). The user reads in the checkpointed



01 data in the restart mode using `SRS_Read()`. The user, through `SRS_Read()`, also specifies the previous  
02 and current data distributions. By knowing the number of processors and the data distributions used  
03 in the previous and current execution of the application, the SRS library automatically performs the  
04 appropriate data redistribution. For example, the user can start their application on four processors  
05 with block distribution of data, stop the application and restart it on eight processors with block-cyclic  
06 distribution. The details of the SRS API for accomplishing the automatic redistribution of data are  
07 beyond the scope of the current discussion. For the current discussion, it is sufficient that the SRS  
08 library is generic and has been tested with numerical libraries like ScaLAPACK and PETSC.

09

10

## 11 6. RESCHEDULING POLICIES

12

### 13 6.1. Policies for contacting the rescheduler

14

15 The contract monitor calculates the ratios between the actual execution times and the predicted  
16 execution times of the application. The tolerance limits of the ratios are specified as inputs to the  
17 contract monitor. When a given ratio is greater than the upper tolerance limit, the contract monitor  
18 calculates the average of the computed ratios. If the average is greater than the upper tolerance limit, it  
19 contacts the rescheduler, requesting that the application be migrated. The average of the ratios is used  
20 by the contract monitor to contact the rescheduler due to the following reasons.

21

22 (1) A competing application of short duration on one of the machines may have increased the  
23 load temporarily on the machine and hence caused the loss in performance of the application.  
24 Contacting the rescheduler for migration on noticing few losses in performance will result  
25 in unnecessary migration in this case, since the competing application will end soon and the  
26 application's performance will be back to normal.

27 (2) The average of the ratios also captures the history of the behavior of the machines on which  
28 the application is running. If the application's performance on most of the iterations has been  
29 satisfactory, then few losses of performance may be due to sparse occurrences of load changes  
30 on the machines.

31 (3) The average of the ratios also takes into account the percentage completed time of application's  
32 execution.

33 If the rescheduler refuses to migrate the application, the contract monitor adjusts its tolerance  
34 limits to new values. Similarly when a given ratio is less than the lower tolerance limit, the contract  
35 monitor calculates the average of the ratios and adjusts the tolerance limits if the average is less than  
36 the lower tolerance limit. The dynamic adjusting of tolerance limits not only reduces the amount of  
37 communication between the contract monitor and the rescheduler, but also hides the deficiencies in the  
38 application-specific execution time model.

39

### 40 6.2. Policies for migration

41

42 For both *migration on request* and *opportunistic migration* modes, the rescheduler first contacts the  
43 NWS to obtain the updated information for the machines in the Grid. It then contacts the application-  
44 specific performance modeler to evolve a new schedule for the application. Based on the current



Table I. Times for rescheduling phases.

Rescheduling phase	Time (s)
Writing checkpoints	40
Waiting for NWS to update information	90
Time for application manager to get new resource information from NWS	120
Evolving new application-level schedule	80
Other Grid overhead	10
Starting application	60
Reading checkpoints and data redistribution	500
Total	900

total percentage completion time for the application and the predicted total execution time for the application with the new schedule, the rescheduler calculates the remaining execution time,  $ret\_new$ , of the application if it were to execute on the machines in the new schedule. The rescheduler also calculates  $ret\_current$ , the remaining execution time of the application if it were to continue executing on the original set of machines. The rescheduler then calculates the rescheduling gain as

$$rescheduling\_gain = \frac{(ret\_current - (ret\_new + 900))}{ret\_current}$$

The number 900 in the numerator of the fraction is the worst case time in seconds needed to reschedule the application. The various times involved in rescheduling are given in Table I. The times shown in Table I were obtained by conducting several experiments with different problem sizes and obtaining the maximum times for each phases of rescheduling. Thus, the rescheduling strategy adopts a pessimistic approach for rescheduling, with the result that migration of applications will be avoided in certain cases where the migration could yield performance benefits.

If the rescheduling gain is greater than 30%, the rescheduler sends stop signal to the RSS and hence to the executing application, and stores the 'STOP' status in GIR. The application manager then waits for the state of the end application to change to 'RESUME'. After the application has stopped, the rescheduler stores 'RESUME' as the state of the application in the GIR thus prompting the application manager to evolve a new schedule and restart the application on the new schedule. If the rescheduling gain is less than 30% and if the rescheduler is operating in the *migration on request* mode, the rescheduler contacts the contract monitor prompting the contract monitor to adjust its tolerance limits.

The rescheduling threshold [28] that the performance gain due to rescheduling must cross for rescheduling to yield significant performance benefits depends on the load dynamics of the system resources, the accuracy of the measurements of resource information and may also depend on the particular application for which rescheduling is made. Since the measurements made by NWS are fairly accurate, the rescheduling threshold for our experiments depended only on the load dynamics of the system resources. By means of trial-and-error experiments using a range of different problem



01 sizes for the different applications that were considered and for different configurations of the available  
02 resources, we determined the rescheduling threshold for our testbed to be 30%. Rescheduling decisions  
03 made below this threshold may not yield performance benefits in all cases.

04

05

## 06 **7. OTHER MIGRATION ISSUES**

07

08 The calculation of remaining execution and percentage completion times of the application forms the  
09 backbone of our rescheduling architecture and lends uniqueness to our approach when compared with  
10 other migration research efforts. The contract monitor, based on the actual and predicted execution  
11 times of the different phases of the executing application and the predicted execution time from the  
12 execution model, calculates the refined expected execution time of the application. Based on the current  
13 elapsed time and the refined expected time of the executing application, the total percentage completion  
14 time and the remaining execution time of the application are calculated by the rescheduler. When  
15 calculating the remaining execution time of the application on a new set of resources, the total predicted  
16 execution time from the execution model for the new set of resources is also taken into account.  
17 Although our approach of calculating the remaining execution and percentage completion times is  
18 most suitable for iterative applications, it can also be applied to other kinds of applications.

19 Also, in order to prevent possible conflicts between different applications due to rescheduling, the  
20 rescheduler is implemented as a single GrADS service that is contacted by the contract monitors  
21 of different applications. The rescheduler implements a queuing system and at any point in time  
22 services the request for a single application by contacting the corresponding application manager of  
23 that application. The stopping of the application by the rescheduler occurs in two steps. First, the  
24 external component contacts the RSS and sends a signal to stop the application. This stop signal occurs  
25 concurrently with application execution. When the application executes the next `SRS_Check_Stop()`  
26 call, it contacts the RSS, obtains the stop information from RSS and proceeds to stop.

27

28

## 29 **8. EXPERIMENTS AND RESULTS**

30

31 The GrADS experimental testbed consists of about 40 machines that reside in institutions across  
32 United States including the University of Tennessee, the University of Illinois, the University of  
33 California at San Diego, Rice University, etc. For the sake of clarity, our experimental testbed consists  
34 of two clusters, one in the University of Tennessee and another in the University of Illinois, Urbana-  
35 Champaign. The characteristics of the machines are given in Table II. The two clusters are connected by  
36 means of the Internet. Although the Tennessee machines are dual-processor machines, the applications  
37 in the GrADS experiments use only one processor per machine.

38 About five applications, namely, ScaLAPACK LU and QR factorizations, ScaLAPACK eigenvalue  
39 problems, PETSC CG solver and the heat equation solver, have been integrated into the migration  
40 framework by instrumenting the applications with SRS calls and developing performance models  
41 for the applications. In general, our migration framework is suitable for iterative MPI-based parallel  
42 applications for which performance models predicting the execution costs can be written. In the  
43 experiments shown in this paper, ScaLAPACK QR factorization was used as the end application.  
44 Similar encouraging results were also obtained for other applications.



Table II. Resource characteristics.

Cluster name	Location	Nodes	Processor type	Speed (MHz)	Memory (MB)	Network	Operating system	Globus version
<i>msc</i>	Tennessee	8	Pentium III	933	512	100 Mb switched Ethernet	Redhat Linux 7.3	2.2
<i>opus</i>	Illinois	16	Pentium II	450	256	1.28 Gbit s <sup>-1</sup> full duplex Myrinet	Redhat Linux 7.2 (2.4.18 kernel)	2.2

The performance model of ScaLAPACK QR factorization was derived by simulating the routine PDGEQRF. The simulation was based on benchmark performance of matrix multiplication and other basic linear algebra routines on the resource testbed and a prediction of communication costs for a given set of network links and for given message sizes. A more detailed description of the QR performance model is beyond the scope of this paper. For a general idea of the methodology, the reader is referred to earlier work [17]. The application was instrumented with calls to SRS library such that the application can be stopped by the rescheduler at any point of time and can be continued on a different configuration of machines. The data that was checkpointed by the SRS library for the application included the matrix, A, and the right-hand side vector, B. Only the PDGEQRF routine and the driver routine for PDGEQRF were modified for instrumentation with SRS calls. The percentage increase in size of the code due to the modifications was less than 4%. Lower tolerance limit of 0.7 and upper tolerance limit of two were used as thresholds for the contract monitor. These thresholds were derived by conducting preliminary performance model validation tests on the testbed.

### 8.1. Migration on request

In all of the experiments in this section, four Tennessee and eight Illinois machines were used. A given matrix size for the QR factorization problem was input to the application manager. For large problem sizes, the computation time dominates the communication time for the ScaLAPACK application. Since the Tennessee machines have higher computing power than the Illinois machines, the application manager by means of the performance modeler chose the four Tennessee machines for the end application run. A few minutes after the start of the end application, artificial load is introduced into the four Tennessee machines. This artificial load is achieved by executing a certain number of loading programs on each of the Tennessee machines. The loading program used was a sequential C code that consists of a single looping statement that loops forever. This program was compiled without any optimization in order to achieve the loading effect.

Due to the loss in predicted performance caused by the artificial load, the contract monitor requested the rescheduler to migrate the application. The rescheduler evaluated the potential performance benefits that can be obtained by migrating the application to the eight Illinois machines and either migrated the application or allowed the application to continue on the four Tennessee machines. The rescheduler was





01 operated in two modes—a default and a non-default mode. The normal operation of the rescheduler  
02 is its default mode, and the non-default mode is to force the opposite decision of whether or not to  
03 migrate. Thus, in cases when the default mode of the rescheduler was to migrate the application, the  
04 non-default mode was to continue the application on the same set of resources, and in cases when  
05 the default mode of the rescheduler was to not migrate the application, the non-default mode was to  
06 force the rescheduler to migrate the application by adjusting the rescheduling cost parameters. For each  
07 experimental run, results were obtained for both when rescheduler was operated in the default and non-  
08 default mode. This allowed us to compare both scenarios and to verify whether the rescheduler made  
09 the right decisions.

10 Three parameters were involved in each set of experiments—the size of the matrices, the amount of  
11 load on the resources and the time after the start of the application when the load was introduced into  
12 the system. The following three sets of experiments were obtained by fixing two of the parameters and  
13 varying the other parameter.

14 In the first set of experiments, the artificial load consisting of 10 loading programs was introduced  
15 into the system five minutes after the start of the end application. The bar chart in Figure 3 was obtained  
16 by varying the size of the matrices, i.e. the problem size on the  $x$ -axis. The  $y$ -axis represents the  
17 execution time in seconds of the entire problem including the Grid overhead. For each problem size,  
18 the bar on the left represents the execution time when the application was not migrated and the bar on  
19 the right represents the execution time when the application was migrated.

20 Several points can be observed from Figure 3. The time for reading checkpoints occupied most of  
21 the rescheduling cost since it involves moving data across the Internet from Tennessee to Illinois and  
22 redistribution of data from four to eight processors. On the other hand, the time for writing checkpoints  
23 is insignificant since the checkpoints are written to local storage. The rescheduling benefits are more  
24 for large problem sizes since the remaining lifetime of the end application when load is introduced  
25 is larger. There is a particular size of the problem below which the migrating cost overshadows the  
26 performance benefit due to rescheduling. Except for matrix size 8000, the rescheduler made correct  
27 decisions for all matrix sizes. For matrix size 8000, the rescheduler assumed a worst-case rescheduling  
28 cost of 900 seconds while the actual rescheduling cost was close to about 420 seconds. Thus, the  
29 rescheduler evaluated the performance benefit to be negligible while the actual scenario points to the  
30 contrary. Thus, the pessimistic approach by using a worst-case rescheduling cost in the rescheduler will  
31 lead to underestimating the performance benefits due to rescheduling in some cases. We also observe  
32 from the figure that the times for reading checkpoints and data distribution do not necessarily increase  
33 linearly with increasing matrix sizes. For example, the time for data distribution is more for matrix  
34 size 11 000 than for matrix size 12 000. This is due to the transient loads associated with the Internet  
35 between Tennessee and Illinois.

36 In the second set of experiments, matrix size 12 000 was chosen for the end application and artificial  
37 load was introduced 20 min into the execution of the application. In this set of experiments, the amount  
38 of artificial load was varied by varying the number of loading programs that were executed. In Figure 4,  
39 the  $x$ -axis represents the number of loading programs and the  $y$ -axis represents the execution time in  
40 seconds. For each amount of load, the bar on the left represents the case when the application was  
41 continued on four Tennessee machines and the bar on the right represents the case when the application  
42 was migrated to eight Illinois machines.

43 Similar to the first set of experiments, we find only one case when the rescheduler made an incorrect  
44 decision for rescheduling. This case, when the number of loading programs was five, was due to the




Figure 3. Problem sizes and migration.

insignificant performance gain that can be obtained due to rescheduling. When the number of loading programs was three, we were not able to force the rescheduler to migrate the application because the application completed during the time for rescheduling decision. Also, the greater the load, the higher the performance benefit due to rescheduling because of larger performance losses for the application in the presence of heavier loads. However, the most significant result in Figure 4 was that the execution times when the application was rescheduled remained almost constant irrespective of the amount of load. This is because, as can be observed from the results when the number of loading programs was 10 and when the number was 20, the more the amount of load, the earlier the application was rescheduled.




Figure 4. Load amount and migration.

Hence our rescheduling framework was able to adapt to the external load. As with Figure 3, we see that the times for checkpoint reading show variance for the same matrix size in Figure 4. Again, this is due to the variance in network loads on the Internet connection between Tennessee and Illinois.

In the third set of experiments, shown in Figure 5, equal amounts of load consisting of seven loading programs was introduced at different points of execution of the end application for the same problem of matrix size 12 000. The  $x$ -axis represents the elapsed execution time in minutes of the end application when the load was introduced. The  $y$ -axis represents the total execution time in seconds. Similar to the previous experiments, the bars on the left denote the cases when the application was not rescheduled and the bars on the right represent the cases when the application was rescheduled.




Figure 5. Load introduction time and migration.

As can be observed from Figure 5, there are diminishing returns due to rescheduling as the load is introduced later into the program execution. The rescheduler made wrong decisions in two cases—when the load introduction times are 15 and 20 min after the start of end application execution. While the wrong decision for 20 min can be attributed to the pessimistic approach of rescheduling, the wrong decision of the rescheduler for 15 min was determined to be due to the faulty functioning of the performance model for the ScaLAPACK QR problem for Illinois machines. The most startling result in Figure 5 is when the load was introduced 23 min after the start of the end application. At this point, the program almost completed and hence rescheduling will not yield performance benefits for



01 the application. The rescheduler was able to evaluate the scenario correctly and avoid unnecessary  
02 rescheduling of the application. Most rescheduling frameworks will not be capable of achieving this  
03 since they do not possess the knowledge regarding remaining execution time of the application.

## 05 8.2. Opportunistic migration

07 In this set of experiments, we illustrate opportunistic migration in which the rescheduler tries to  
08 migrate an executing application when some other application completes. For these experiments, two  
09 problems were involved. For the first problem, matrix size of 14 000 was used and six Tennessee  
10 machines were made available. The application manager, through the performance modeler, chose the  
11 six machines for the end application run. Two minutes after the start of the end application for the  
12 first problem, a second problem of a given matrix size was input to the application manager. For the  
13 second problem, the six Tennessee machines on which the first problem was executing and two Illinois  
14 machines were made available. Due to the presence of the first problem, the six Tennessee machines  
15 alone were insufficient to accommodate the second problem. Hence, the performance model chose the  
16 six Tennessee machines and two Illinois machines for the end application and the actual application  
17 run involved communication across the Internet.

18 In the middle of the execution of the second application, the first application completed and hence the  
19 second application can be potentially migrated to use only the six Tennessee machines. Although this  
20 involved constricting the number of processors for the second application from eight to six, there can  
21 be potential performance benefits due to the non-involvement of Internet. The rescheduler evaluated  
22 the potential performance benefits due to migration and made an appropriate decision.

24 Figure 6 shows the results for two illustrative cases when matrix sizes of the second application  
25 were 13 000 and 14 000. The  $x$ -axis represents the matrix sizes and the  $y$ -axis represents the execution  
26 time in seconds. For each application run, three bars are shown. The bar on the left represents the  
27 execution time for the first application that was executed on six Tennessee machines. The middle bar  
28 represents the execution time of the second application when the entire application was executed on six  
29 Tennessee and two Illinois machines. The bar on the right represents the execution time of the second  
30 application, when the application was initially executed on six Tennessee and two Illinois machines  
31 and later migrated to execute on only six Tennessee machines when the first application completed.

32 For the second problem, for both matrix sizes 13 000 and 14 000, for the second problem, the  
33 rescheduler made the correct decision of migrating the application. We also find that for both problem  
34 cases, the second application was almost immediately rescheduled after the completion of the first  
35 application.

## 37 8.3. Predicting redistribution cost

39 As observed in Figures 3–5, the rescheduler can make wrong decisions for rescheduling in  
40 certain cases. In cases where the rescheduler made the wrong decision, the rescheduler decided  
41 that rescheduling the executing application will not yield significant performance benefits for the  
42 application, while the actual results point to the contrary. This is because the rescheduler used the  
43 worst-case times shown in Table I for different phases of rescheduling while the actual rescheduling  
44 cost was less than the worst-case rescheduling cost.




Figure 6. Opportunistic migration.

As shown in Table I, the cost for reading and redistribution of checkpoint data is the highest of the various costs involved in rescheduling. The checkpoint reading and redistribution are performed in a single operation where the processes determine the portions and locations of data needed by them and read the checkpoints directly from the IBP [26] depots. The data redistribution cost depends on a number of factors including the number and amount of checkpointed data, the data distributions used for the data, the current and future processors sets for the application used before and after rescheduling, the network characteristics, particularly the latency and bandwidth of the links between the current and future processor sets, etc. The rescheduling framework was extended to predict the



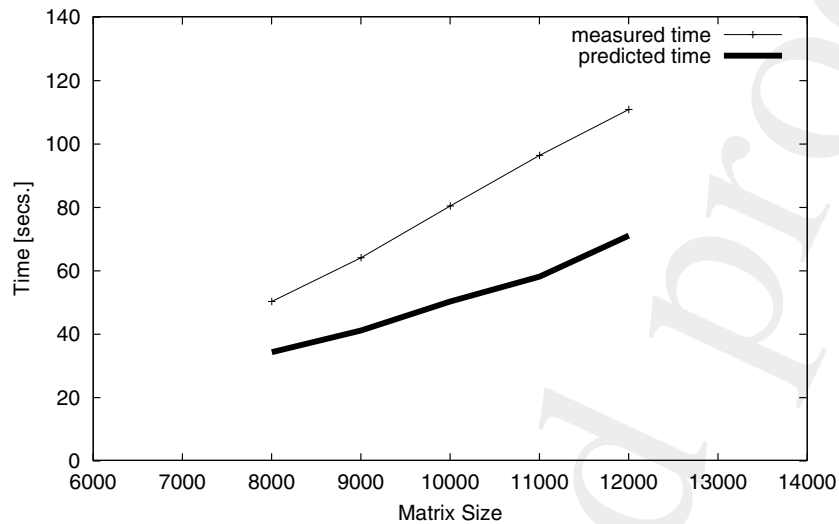
01 redistribution cost and use the predicted redistribution cost for calculating the gain due to rescheduling  
02 the executing application. Although the time for writing the checkpoints also depends on the size of  
03 the checkpoints, the checkpoint writing time is insignificant because the processes write checkpoint  
04 data to the local storage. Hence, the time for checkpoint writing is not predicted in the rescheduling  
05 framework.

06 Similar to the SRS library, the rescheduling framework has also been extended to support common  
07 data distributions such as block, cyclic and block-cyclic distributions. When the end application calls  
08 `SRS_Register` to register data to be checkpointed, it also specifies the data distribution used for that  
09 data. If the data distribution is one of the common data distributions, the input parameter used for the  
10 distribution is stored in an internal data structure of the SRS library. For example, if a block-cyclic data  
11 distribution is specified for the data, the block size used for the distribution is stored in the internal  
12 data structure. When the application calls `SRS_StoreMap`, the data distributions used for the different  
13 data along with the parameters used for the distribution are sent to the RSS.

14 When the rescheduler wants to calculate the rescheduling cost of an executing application, it contacts  
15 the RSS of the application, and retrieves various information about the data that were marked for  
16 checkpointing including the total size and data types of the data, the data distributions used for the  
17 data and the parameters used for the data distributions. For each data that uses one of the common data  
18 distributions supported by the rescheduler, the rescheduler determines the data maps for the current  
19 processor configuration on which the application is executing and the future processor configuration  
20 where the application can be potentially rescheduled. A data map indicates the total number of panels  
21 of the data and the size and location of each of the data panel. The rescheduler calculates the data  
22 map using the data distribution and the parameters used for data distribution it collected from RSS.  
23 Based on the data maps for the current and future processor configuration and the properties of  
24 the networks between the current and future processor configuration it collected from NWS, the  
25 rescheduler simulates the redistribution behavior. The end result of the simulation is the predicted  
26 cost for reading and redistribution of checkpointed data if the application was rescheduled to the new  
27 processor configuration. The rescheduler uses this predicted redistribution cost for calculation of the  
28 potential rescheduling gain that can be obtained due to rescheduling the application.

29 An experiment was conducted in which the simulation model for predicting the redistribution  
30 cost was validated. In this experiment, four Tennessee and eight Illinois machines were used.  
31 A ScaLAPACK QR factorization problem was submitted to the GrADS Application Manager. Since  
32 the Tennessee machines were faster than the Illinois machines, the four Tennessee machines were  
33 chosen by the Performance Modeler for the execution of the end application. Five minutes after the  
34 start of the execution of the end application, artificial loads were introduced in the Tennessee machines  
35 by the execution of 10 loading programs on each of the Tennessee machines. When the Contract  
36 Monitor contacted the rescheduler requesting that the application be rescheduled, the rescheduler  
37 dynamically predicted the redistribution cost involved in rescheduling the application. Figure 7  
38 compares the predicted and the actual cost for redistribution of the data for different problem sizes.  
39 The  $x$ -axis denoted the matrix sizes used for the QR factorization problem and the  $y$ -axis represents  
40 the redistribution time.

41 From Figure 7, we find that the rescheduler was able to perform a reasonably accurate simulation  
42 of the redistribution of data. The actual redistribution cost was greater than the predicted redistribution  
43 cost by only 30–40 s. The difference is mainly due to the unpredictable behavior in the network  
44 characteristics of the Internet connection between Tennessee and Illinois. By employing the predicted



20 Figure 7. Redistribution cost prediction versus actual performance (four Tennessee and eight Illinois machines).

21  
22  
23  
24 redistribution cost, the rescheduler was able to make the right decisions for rescheduling for cases in  
25 Figures 3–5 when it previously made wrong decisions.

## 28 9. RELATED WORK

29  
30 Different systems have been implemented to migrate executing applications onto different  
31 sets of resources. These systems migrate applications either to efficiently use under-utilized  
32 resources [11,22,23,29,30], or to provide fault resilience [31], or to reduce the obtrusiveness  
33 to workstation owner [10,31]. The particular projects that are closely related to our work are  
34 Dynamite [11], MARS [13], LSF [12], Condor [10] and Cactus [32].

35 The Dynamite system [11] based on Dynamic PVM [30] migrates applications when certain  
36 machines in the system get under-utilized or over-utilized as defined by application-specified  
37 thresholds. Although this method takes into account application-specific characteristics it does not  
38 necessarily evaluate the remaining execution time of the application and the resulting performance  
39 benefits due to migration.

40 In LSF [12], jobs can be submitted to queues which have pre-defined migration thresholds. A job  
41 can be suspended when the load of the resource increases beyond a particular limit. When the time  
42 since the suspension becomes higher than the migration threshold for the queue, the job is migrated  
43 and submitted to a new queue. Thus LSF suspends jobs to maintain the load level of the resources while  
44 our migration framework suspends jobs only when it is able to find better resources where the jobs can





01 be migrated. By adopting a strict approach to suspending jobs based on pre-defined system limits, LSF  
02 gives less priority to the stage of the application execution whereas our migration framework suspends  
03 an application only when the application has enough remaining execution time so that performance  
04 benefits can be obtained by migration. And, lastly, due to the separation of the suspension and migration  
05 decisions, a suspended application in LSF can wait for a long time before it restarts executing on  
06 a suitable resource. In our migration framework, a suspended application is immediately restarted  
07 because of the tight coupling of suspension and migration decisions.

08 Of the Grid computing systems, only Condor [10] seems to migrate applications under workload  
09 changes. Condor provides powerful and flexible ClassAd mechanism by means of which the  
10 administrator of resources can define policies for allowing jobs to execute on the resources, suspending  
11 the jobs, and vacating the jobs from the resources. The fundamental philosophy of Condor is to increase  
12 the throughput of long running jobs and also respect the ownership of the resource administrators.  
13 The main goal of our migration framework is to increase the response times of individual applications.  
14 Similar to LSF, Condor also separates the suspension and migration decisions and hence has the same  
15 problems mentioned for LSF in taking into account the performance benefits of migrating applications.  
16 Unlike our rescheduler framework, the Condor system does not possess knowledge about the remaining  
17 execution time of the applications. Thus suspension and migrating decisions can be invoked frequently  
18 in Condor based on system load changes. This may be less desirable in Grid systems where system  
19 load dynamics are fairly high.

20 The Cactus [32] migration framework was also developed in the context of the GrADS project and  
21 hence follows most of the design principles of our migration framework. Their migration thorn is  
22 similar to our migrator and their performance detection thorn also performs contract monitoring and  
23 detects contract violation similar to our contract monitor. Their migration logic manager is similar in  
24 principle to our rescheduler. The differences lie in the decisions made to contact the rescheduler service  
25 for migration, and decisions made in the rescheduler regarding when to migrate. While our migration  
26 framework makes decisions using a threshold for the average performance ratio, the Cactus framework  
27 uses a maximum number of consecutive contract violations as the threshold for migration. Although  
28 Cactus allows the thresholds to be changed dynamically by the user using a HTTP interface, often  
29 the user does not possess adequate expertise in determining the threshold and hence a more automatic  
30 mechanism such as that followed in our approach is desirable for Grid systems. Also, the Cactus  
31 migration framework only uses the resource characteristics to discover better systems for migrating,  
32 whereas our system uses predicted application performance on the new systems. Also, similar to other  
33 approaches, Cactus does not take into account the remaining execution time of the application.

34 The GridWay framework [33] has a number of similarities with the GrADS framework both in  
35 terms of concepts and the design of the architecture. Hence, GridWay's job migration framework  
36 by Montero *et al.* [34] performs most of the functionalities of our migration framework. Their job  
37 migration framework takes into account the proximity of the execution hosts to the checkpoint and  
38 restart files. Their job migration framework also performs opportunistic migration and migration under  
39 performance degradation. However, their work does not mention about the migration of parallel MPI  
40 jobs and the possible reconfiguration of hosts and the redistribution of data. By considering dynamic  
41 redistribution costs based on network bandwidths, our migration framework indirectly takes into  
42 account the proximity of the new hosts to the checkpoint files. Lastly, the execution models used by our  
43 migration framework simulate the actual application and hence are more robust than their mathematical  
44 models.



## 10. CONCLUSIONS AND FUTURE WORK

Many existing migration systems that migrate applications under resource load conditions implement simple policies that cannot be applied to Grid systems. We have implemented a migration framework that takes into account both the system load and application characteristics. The migrating decisions are based on factors including the amount of resource load, the point during the application lifetime when the load is introduced, and the size of the applications. We have also implemented a framework that opportunistically migrates executing applications to make use of additional free resources. Experiments were conducted and results were presented to demonstrate the capabilities of the migration framework.

We intend to provide more robust frameworks in the SRS system and in the rescheduler to efficiently predict the cost for the redistribution of data. Also, instead of fixing the rescheduler threshold at 30%, our future work will involve determining the rescheduling threshold dynamically based on the dynamic observation of load behavior on the system resources. We propose to investigate the usefulness of our approach for complex applications involving multiple components and/or written in multi-programming languages similar to the efforts of Mayes *et al.* [35]. Currently, the average of performance ratios is used to determine when a contract monitor will contact the rescheduler for migration. In the future, we plan to investigate more robust policies for contacting the rescheduler. Mechanisms for quantifying the deficiencies of the execution model detected during contract monitoring and communicating the information to the application developer also need to be investigated.

## ACKNOWLEDGEMENTS

The authors wish to thank the reviewers for their very helpful comments toward improving the quality of the paper. This research was supported in part by the Applied Mathematical Sciences Research Program of the Office of Mathematical, Information, and Computational Sciences, U.S. Department of Energy under contract DE-AC05-00OR22725 with UT-Battelle, LLC, and in part based on work supported by the National Science Foundation under Grant No. ACI 0103759.

## REFERENCES

1. Frigo M. FFTW: An adaptive software architecture for the FFT. *Proceedings of the ICASSP Conference*, vol. 3, 1998; 1381.
2. Whaley RC, Dongarra J. Automatically tuned linear algebra software. *SC98: High Performance Networking and Computing*, 1998.
3. Vadhiyar S, Fagg G, Dongarra J. Automatically tuned collective communications. *Proceedings of SuperComputing2000*, November 2000.
4. Vuduc R, Demmel JW, Yelick KA, Kamil S, Nishtala R, Lee B. Performance optimizations and bounds for sparse matrix-vector multiply. *Proceedings of Supercomputing*, Baltimore, MD, November 2002.
5. Chen Z, Dongarra J, Luszczek P, Roche K. Self adapting software for numerical linear algebra and LAPACK for clusters. *Parallel Computing* 2003; (submitted).
6. Berman F *et al.* The GrADS project: Software support for high-level Grid application development. *International Journal of High Performance Applications and Supercomputing* 2001; **15**(4):327–344.
7. Foster I, Kesselman C (eds.). *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann: San Mateo, CA, 1999.
8. Mirchandaney R, Towsley D, Stankovic JA. Adaptive load sharing in heterogeneous distributed systems. *Journal of Parallel and Distributed Computing* 1990; **9**:331–346.



- 01 9. Douglis F, Ousterhout JK. Transparent process migration: Design alternatives and the sprite implementation. *Software: Practice and Experience* 1991; **21**(8):757–785.
- 02 10. Litzkow M, Livney M, Mutka M. Condor—a hunter for idle workstations. *Proceedings of the 8th International Conference on Distributed Computing Systems*, 1988; 104–111.
- 03 11. van Albada GD, Clinckemallie J, Emmen AHL, Gehring J, Heinz O, van der Linden F, Overeinder BJ., Reinefeld A, Sloot PMA. Dynamite—blasting obstacles to parallel cluster computing. *High-Performance Computing and Networking (HPCN Europe '99)*, Amsterdam, The Netherlands, April 1995, Sloot PMA, Bubak M, Hoekstra AG, Hertzberger LO (eds.) (*Lecture Notes in Computer Science*, vol. 1593). Springer: Berlin, 1995; 300–310.
- 04 12. Zhou S, Zheng X, Wang J, Delisle P. Utopia: A load sharing facility for large, heterogeneous distributed computer systems. *Software: Practice and Experience* 1993; **23**(12):1305–1336.
- 05 13. Gehring J, Reinefeld A. MARS—A framework for minimizing the job execution time in a metacomputing environment. *Future Generation Computer Systems* 1996; **12**(1):87–99.
- 06 14. Vadhiyar S, Dongarra J. Performance oriented migration framework for the Grid. *Proceedings of the 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2003)*, Tokyo, Japan, May 2003; 130–137.
- 07 15. Snir M, Otto S, Huss-Lederman S, Walker D, Dongarra J. *MPI: The Complete Reference—The MPI Core* (2nd edn), vol. 1. MIT Press: Boston, MA, 1998.
- 08 16. MPI. <http://www-unix.mcs.anl.gov/mpi>.
- 09 17. Petitet A, Blackford S, Dongarra J, Ellis B, Fagg G, Roche K, Vadhiyar S. Numerical libraries and the Grid: The GrADS experiments with scalapack. *Journal of High Performance Applications and Supercomputing* 2001; **15**(4):359–374.
- 10 18. Fitzgerald S, Foster I, Kesselman C, von Laszewski G, Smith W, Tuecke S. A directory service for configuring high-performance distributed computations. *Proceedings of the 6th IEEE Symposium on High-Performance Distributed Computing*, 1997; 365–375.
- 11 19. Wolski R, Spring N, Hayes J. The network weather service: A distributed resource performance forecasting service for metacomputing. *Journal of Future Generation Computing Systems* 1999; **15**(5–6):757–768.
- 12 20. Berman F, Wolski R. The AppLeS project: A status report. *Proceedings of the 8th NEC Research Symposium*, May 1997.
- 13 21. Ribler RL, Vetter JS, Simitci H, Reed DA. Autopilot: Adaptive control of distributed applications. *Proceedings of the 7th IEEE Symposium on High-Performance Distributed Computing*, July 1998.
- 14 22. Casas J, Clark D, Galbiati P, Konuru R, Otto S, Prouty R, Walpole J. *MIST: PVM with Transparent Migration and Checkpointing*, 1995.
- 15 23. Casas J, Clark D, Konuru R, Otto S, Prouty R, Walpole J. MPVM: A migration transparent version of PVM. *Technical Report CSE-95-002*, 1995.
- 16 24. Stellner G. CoCheck: Checkpointing and process migration for MPI. *Proceedings of the 10th International Parallel Processing Symposium (IPPS '96)*, Honolulu, HI, 1996; 526–531.
- 17 25. Foster I, Karonis N. A Grid-enabled MPI: Message passing in heterogeneous distributed computing systems. *Proceedings of SuperComputing 98 (SC98)*, 1998.
- 18 26. Plank JS, Beck M, Elwasif WR, Moore T, Swamy M, Wolski R. The Internet backplane protocol: Storage in the network. *NetStore99: The Network Storage Symposium*, 1999.
- 19 27. Geist GA, Kohl JA, Papadopoulos PM. CUMULVS: Providing fault-tolerance, visualization and steering of parallel applications. *International Journal of High Performance Computing Applications* 1997; **11**(3):224–236.
- 20 28. Wolski R, Shao G, Berman F. Predicting the cost of redistribution in scheduling. *Proceedings of 8th SIAM Conference on Parallel Processing for Scientific Computing*, March 1997.
- 21 29. Saqabi KA, Otto SW, Walpole J. Gang scheduling in heterogeneous distributed systems. *Technical Report, OGI*, 1994.
- 22 30. Dikken L, van der Linden F, Vesseur JJJ, Sloot PMA. DynamicPVM: Dynamic load balancing on parallel systems. *Proceedings of High Performance Computing and Networking, Volume II, Networking and Tools*, Munich, Germany, April 1994, Gentsch W, Harms U (eds.) (*Lecture Notes in Computer Science*, vol. 797). Springer: Berlin, 1994; 273–277.
- 23 31. Arabe JNC, Lowekamp ABB, Seligman E, Starkey M, Stephan P. Dome: Parallel programming in a heterogeneous multi-user environment. *Proceedings of Supercomputing*, 1995.
- 24 32. Allen G, Angulo D, Foster I, Lanfermann G, Liu C, Radke T, Seidel E, Shalf J. The Cactus Worm: Experiments with dynamic resource discovery and allocation in a Grid environment. *The International Journal of High Performance Computing Applications* 2001; **15**(4):345–358.
- 25 33. Huedo E, Montero RS, Llorente IM. An experimental framework for executing applications in dynamic Grid environments. *Technical Report 2002-43, NASA/CR-2002-211960*, ICASE, November 2002.
- 26 34. Montero RS, Huedo E, Llorente IM. Grid resource selection for opportunistic job migration. *Proceedings of the 9th International Euro-Par Conference*, Klagenfurt, Austria, August 2003 (*Lecture Notes in Computer Science*, vol. 2790). Springer: Berlin, 2003; 366–373.
- 27 35. Mayes K, Riley GD, Ford RW, Luján M, Freeman L, Addison C. The design of a performance steering system for component-based Grid applications. *Performance Analysis and Grid Computing*, Getov V, Gerndt M, Hoisie A, Maloney A, Miller B, (eds.). Kluwer Academic: Dordrecht, 2003; 111–127.

## Annotations from cpe927.pdf

### Page 1

---

*Annotation 1*

Au:

Correspondence author OK as set or should it be Dongarra?

### Page 22

---

*Annotation 1*

Au:

Please supply publisher and publisher location details for refs [1-4,10,14,18,20-22,24-26,28,31]?

Please update ref [5]?

Please supply last date website accessed for ref [16]?

What is ref [21], Technical report, etc?

Any Institution for ref [23]?