

# 1 New Grid Scheduling 2 and Rescheduling Methods 3 in the GrADS Project

4 F. Berman,<sup>1</sup> H. Casanova,<sup>1</sup> A Chien,<sup>1</sup> K. Cooper,<sup>2</sup> H. Dail,<sup>1</sup>  
5 A. Dasgupta,<sup>2</sup> W. Deng,<sup>3</sup> J. Dongarra,<sup>4</sup> L. Johnsson,<sup>5</sup>  
6 K. Kennedy,<sup>2,6</sup> C. Koelbel,<sup>2</sup> B. Liu,<sup>5</sup> X. Liu,<sup>1</sup> A. Mandal,<sup>2</sup>  
7 G. Marin,<sup>2</sup> M. Mazina,<sup>2</sup> J. Mellor-Crummey,<sup>2</sup> C. Mendes,<sup>3</sup>  
8 A. Olugbile,<sup>1</sup> M. Patel,<sup>5</sup> D. Reed,<sup>3</sup> Z. Shi,<sup>4</sup> O. Sievert,<sup>1</sup>  
9 H. Xia,<sup>1</sup> and A. YarKhan<sup>4</sup>

Received ■; revised ■; accepted ■

10  
11  
12  
13  
14  
15  
16

The goal of the Grid Application Development Software (GrADS) Project is to provide programming tools and an execution environment to ease program development for the Grid. This paper presents recent extensions to the GrADS software framework: a new approach to scheduling workflow computations, applied to a 3-D image reconstruction application; a simple stop/migrate/restart approach to rescheduling Grid applications, applied to a

<sup>1</sup>Department of Computer Science and Engineering, University of California at San Diego, San Diego, CA 92093. E-mail: {berman, casanova, achien, hdail, lxin, aoo, osievert, huaxia}@ucsd.edu

<sup>2</sup>Computer Science Dept., Rice University, Houston, TX 77005. E-mail: {keith, anshuman, ken, chk, anirban, mgabi, mmzn, johnmc}@rice.edu

<sup>3</sup>Department Computer Science, University of Illinois, Urbana, IL 61801. E-mail: {weideng, cmendes, reed}@uiuc.edu

<sup>4</sup>Innovative Computing Lab, University of Tennessee, Knoxville, TN 37996. E-mail: {dongarra,yarkhan,shi}@utk.edu

<sup>5</sup>Department Computer Science, University of Houston, Houston, TX 77204. E-mail: {johnsson,bliu2,mpate}@uh.edu

<sup>6</sup>To whom correspondence should be addressed.



17  
18  
19  
20  
21

QR factorization benchmark; and a process-swapping approach to rescheduling, applied to an N-body simulation. Experiments validating these methods were carried out on both the GrADS MacroGrid (a small but functional Grid) and the MicroGrid (a controlled emulation of the Grid).

---

**KEY WORDS:** Grid computing; scheduling; rescheduling.

23

## 1. INTRODUCTION

24  
25  
26  
27  
28  
29  
30

Since 1999, the Grid Application Development (GrADS) Project has worked to enable an integrated computation and information resource based on advanced networking technologies and distributed information sources. In other words, we have been attacking the problems inherent in Grid computing.<sup>(1)</sup> In theory, the Grid connects computers, databases, instruments, and people into a seamless web of advanced capabilities. In practice, its lack of usability has limited its application to specialists.

31  
32  
33  
34  
35  
36  
37  
38

Because the Grid is inherently more complex than stand-alone computer systems, Grid programs must reflect this complexity at some level. However, we believe that this complexity should *not* be embedded in the main algorithms of the application, as is often now the case. Instead, GrADS provides software tools that manage the Grid-specific details of execution with minimal effort by the scientists and engineers who write the programs. This increases usability and allows the system to perform substantial optimizations for Grid execution.

39  
40  
41  
42  
43  
44  
45  
46  
47  
48

Figure 1 shows the program development framework that GrADS pioneered in response to this need.<sup>(2)</sup> Two key concepts are central to this approach. First, applications are encapsulated as *configurable object programs (COPs)*, which can be optimized rapidly for execution on a specific collection of Grid resources. A COP includes *code* for the application, a *mapper* that determines how to map an application's tasks to a set of resources, and a *performance model* that estimates the application's performance on a set of resources. Second, the system relies upon *performance contracts* that specify the expected performance of modules as a function of available resources.

49  
50  
51  
52  
53  
54

The left side of Fig. 1 depicts tools used to construct COPs from either domain-specific components or low-level (e.g. MPI) programming. In either case, GrADS provides prototype tools that semi-automatically construct performance models and mappers. Although they are not the major focus of this paper, some of these tools are described in more detail in Section 3 below.

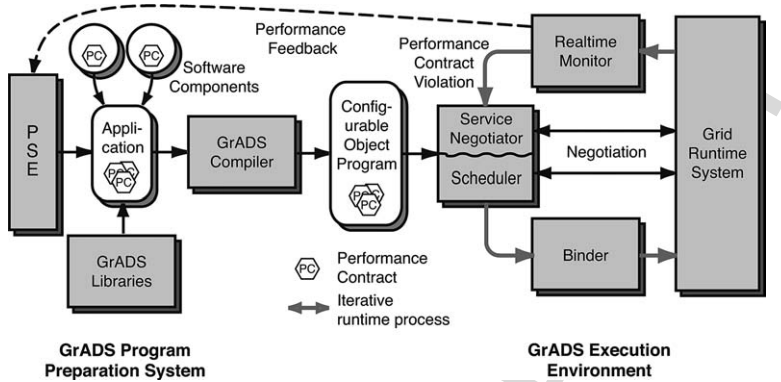


Fig. 1. GrADS Program Preparation and Execution Architecture.

55 The right side of Fig. 1 depicts actions when a COP is delivered  
 56 to the execution environment. The GrADS infrastructure first determines  
 57 which resources are available and, using the COP's mapper and perfor-  
 58 mance model, schedules the application components onto an appropriate  
 59 subset of these resources. Then the GrADS software invokes the *binder*  
 60 to tailor the COP to the chosen resources and the *launcher* (not shown) to  
 61 start the tailored COP on the Grid.

62 Once launched, execution is tracked by the *contract monitor*, which  
 63 detects anomalies and invokes, when necessary, the *rescheduler* to take  
 64 corrective action. Performance monitoring in GrADS is based on Auto-  
 65 pilot,<sup>(3)</sup> a toolkit for real-time application and resource monitoring and  
 66 closed-loop control. Autopilot provides sensors for performance data  
 67 acquisition, actuators for implementing optimization commands and a  
 68 decision-making mechanism based on fuzzy logic. Part of the tailoring  
 69 done by the binder is to insert the sensors needed for monitoring a par-  
 70 ticular application. Autopilot then assesses the application's progress using  
 71 performance contracts,<sup>(4)</sup> which specify an agreement between application  
 72 demands and resource capabilities. The contract monitor takes periodic  
 73 data from the sensors and uses Autopilot's decision mechanism to ver-  
 74 ify that the contract is being met. If a contract violation occurs, the  
 75 monitor takes corrective action, such as contacting a GrADS rescheduler.  
 76 GrADS incorporates a variety of utilities associated with contract mon-  
 77 itoring, including a Java-based Contract Viewer GUI to visualize the per-  
 78 formance contract validation activity in real-time.

79 To support research into and evaluation of GrADS capabilities,  
 80 GrADS has constructed two research testbeds. The *MacroGrid* consists  
 81 of Linux clusters with GrADS software installed at several participat-  
 82 ing GrADS sites, including clusters at University of California at San

83 Diego (UCSD, 10 machines), University of Tennessee at Knoxville (UTK,  
84 24 machines), University of Illinois at Urbana-Champaign (UIUC, 24  
85 machines), and University of Houston (UH, 24 machines). The experi-  
86 ments in Section 3 and Section 4.1 run on this testbed. The *MicroGrid* is  
87 a Grid emulation environment that runs on clusters and permits experi-  
88 mentation with extreme variations in network traffic and loads on com-  
89 pute nodes.<sup>(5)</sup> Section 4.2 describes experiments run on this platform. (We  
90 earlier ran very similar experiments on the MacroGrid, validating both the  
91 MicroGrid's emulation and the rescheduling method's practicality.<sup>(6)</sup>)

92 The experiments we describe exercise many parts of the GrADS envi-  
93 ronment. This paper closes with a brief discussion of what we learned  
94 from these experiences, and an outline of future work.

## 95 2. LAUNCHING COMPONENTS ON THE GRID

96 Once an application schedule has been chosen, the GrADS *applica-*  
97 *tion manager* must prepare the configurable object program and map it  
98 onto the selected resource configuration. In turn, the application man-  
99 ager invokes the binder, which is responsible for creating and configu-  
100 ring the application executable, instrumenting it, and then launching it on  
101 the Grid. The original GrADS binder did most of its work by editing the  
102 entire application binary, which limited its applicability to homogeneous  
103 collections of processors (such as our original testbed). It soon became  
104 clear that this approach would not suffice for a general system because  
105 most grids (including later generations of our own testbed) are heteroge-  
106 neous and because many grid programs require linking against libraries of  
107 components preinstalled on Grid resources.

108 To address these issues, we developed a new distributed GrADS binder  
109 that executes on all Grid resources specified in the schedule. The new binder  
110 receives three sets of inputs: resource specific information (such as hardware  
111 and software capabilities) via the GrADS Information Service (GIS), char-  
112 acteristics of the target architecture that can be used for machine-specific  
113 optimizations, and a *compilation package* that consists of the application's  
114 source code in an intermediate representation, a list of required libraries,  
115 and a script to configure the application for compilation.

116 A binder process executes on each machine chosen by the scheduler.  
117 For this to be possible, the global binder must know the locations of all  
118 software resources, including application-specific libraries, general libraries,  
119 and the binder itself. To that end, the global binder queries the GIS to  
120 locate necessary software on the scheduled node, starting with the local  
121 binder code. The global binder then launches the local binder process,  
122 which further queries GIS for the locations of application-specific libraries,

123 instruments the code with Autopilot sensors, configures, compiles, and  
 124 links the application. Finally, the global binder enables the launch of  
 125 the application. If the application is an MPI application, then a global  
 126 synchronization must be carried out as part of the MPI protocol at the  
 127 beginning of the execution. In this case, the binder returns control to the  
 128 application manager which launches the application after synchronization.  
 129 In non-MPI applications, the binder launches the application and notifies  
 130 the application manager when the program terminates.

131 Note that by using a high-level representation of the program and  
 132 configuring and compiling it only at the target machine, the binder nat-  
 133 urally deals with heterogeneous resources. This is important in any Grid  
 134 context. Moreover, preserving high-level program information until the  
 135 target machine is known also provides opportunities for architecture-  
 136 specific optimizations.

### 137 3. SCHEDULING WORKFLOW GRAPHS

138 Workflow applications are an important class of programs that can  
 139 take advantage of the power of Grid computing, such as the LIGO<sup>(7)</sup> pul-  
 140 sar search image processing applications.<sup>(8)</sup> As the name suggests, a work-  
 141 flow application consists of a collection of components that need to be  
 142 executed in a partial order determined by control and data dependences.

143 The previous version of the GrADS scheduler was designed to sup-  
 144 port tightly-coupled MPI applications<sup>(9-11)</sup> and was not well suited to  
 145 workflow applications. On the other hand, existing approaches to work-  
 146 flow scheduling, such as Condor DAGMan,<sup>(12)</sup> are not able to effectively  
 147 exploit the performance modeling available within GrADS to produce  
 148 better schedules. To address these shortcomings, we developed a new  
 149 GrADS workflow scheduler that resolves the application dependences and  
 150 schedules the components, including parallel components, onto available  
 151 resources using GrADS performance models as a guide.

#### 152 3.1. Workflow Scheduling

153 A Grid scheduler for a workflow application must be guided by an  
 154 objective function that it tries to optimize, such as minimizing communi-  
 155 cation time or maximizing throughput. For the GrADS Project, we have  
 156 chosen to minimize the overall job completion time, also known as the  
 157 *makespan*, of the application. The GrADS scheduler builds up a model of  
 158 Grid resources using services such as MDS<sup>(13)</sup> and NWS.<sup>(14)</sup> The sched-  
 159 uler also obtains performance models of the application using a scalable  
 160 technique developed for GrADS. Using these models, the scheduler then  
 161 provides a mapping from the workflow components to the Grid resources.

162 A stricter definition of the problem can be formulated with the  
 163 help of two sets: the set  $C = \{c_1, c_2, \dots, c_m\}$  of available application  
 164 components from the application DAG, and the set  $G = \{r_1, r_2, \dots, r_n\}$  of  
 165 available Grid resources. The goal of the scheduler is to construct a map-  
 166 ping from elements of  $C$  onto elements of  $G$ .

167 For each application component, the GrADS workflow scheduler  
 168 ranks each eligible resource, reflecting the fit between the component and  
 169 the resource. Lower rank values, in our convention, indicate a better  
 170 match for the component. After ranking the components, the scheduler  
 171 collates this information into a performance matrix. Finally, it runs heu-  
 172 ristics on the performance matrix to schedule components onto resources.

*Computing rank values* The scheduler ensures that resources meet cer-  
 tain minimum requirements for a component. Resources that do not qual-  
 ify under these criteria are given a rank value of infinity. For all other  
 resources, the rank of the resource  $r_j$  is calculated by using a weighted  
 sum of the expected execution time on the resource and the expected cost  
 of data movement for the component  $c_i$ :

$$rank(c_i, r_j) = w_1 \times eCost(c_i, r_j) + w_2 \times dCost(c_i, r_j) \quad (1)$$

173 The expected execution time  $eCost$  is calculated using a performance  
 174 modeling technique that will be described in the next section. The cost of  
 175 data movement  $dCost$  is estimated by a product of the total volume of data  
 176 required by the component and the expected time to transfer data given  
 177 current network conditions. For this measurement, NWS is used to obtain  
 178 an estimate of the current network latency and bandwidth. The weights  $w_1$   
 179 and  $w_2$  can be customized to vary the relative importance of the two costs.

180 *Scheduling application components* Once ranks have been calculated, a  
 181 performance matrix is constructed. Each element of the matrix  $p_{ij}$  denotes  
 182 the rank value of executing the  $i$ th component on the  $j$ th resource. This  
 183 matrix is used by the scheduling heuristics to obtain a mapping of com-  
 184 ponents onto resources. Such a heuristic approach is necessary since the  
 185 mapping problem is NP-complete.<sup>(15)</sup> We apply three heuristics to obtain  
 186 three mappings and then select the schedule with the minimum makespan.  
 187 The heuristics that we apply are the min-min, the max-min, and the suffer-  
 188 age heuristics.<sup>(16,17)</sup>

### 189 3.2. Component Performance Modeling

190 As described in the previous section, estimating the performance  
 191 of a workflow component on a single node is crucial to construct-  
 192 ing a good overall workflow schedule. We model performance by build-  
 193 ing up an architecture-independent model of the workflow component

194 from individual component models. To obtain the component models, we  
195 consider both the number of floating point operations executed and the  
196 memory access pattern. We do not aim to predict an exact execution  
197 time, but rather provide an estimated resource usage that can be converted  
198 to a rough time estimate based on architectural parameters. Because the  
199 resources are architecture-independent, our models can be used on widely  
200 varying node types.

201 To understand the floating point computations performed by an  
202 application, we use hardware performance counters to collect operation  
203 counts from several executions of the program with different, small-size  
204 input problems. We then apply least squares curve-fitting on the collected  
205 data.

206 To understand an application's memory access pattern, we collect  
207 histograms of memory reuse distance (MRD)—the number of unique  
208 memory blocks accessed between a pair of references to the same block—  
209 observed by each load and store instruction.<sup>(18)</sup> Using MRD data col-  
210 lected on several small-size input problems to the application, we model  
211 the behavior of each memory instruction, and predict the fraction of hits  
212 and misses for a given problem size and cache configuration. To deter-  
213 mine the cache miss count for a different problem size and cache config-  
214 uration, we evaluate the MRD models for each reference at the specified  
215 problem size, and count the number of accesses with predicted reuse dis-  
216 tance greater than the target cache size.

### 217 3.3. Workflow Scheduling Test Case

218 In this section, we apply some of the strategies described in the previ-  
219 ous sections to the problem of adapting EMAN,<sup>(19)</sup> a bio-imaging applica-  
220 tion developed at Baylor College of Medicine, for execution on the Grid  
221 using the GrADS infrastructure. EMAN automates a portion producing  
222 3-D reconstructions of single particles from electron micro-graphs. Human  
223 intervention and expertise is needed to define a preliminary 3-D model  
224 from the electron micro-graphs, but the refinement from a preliminary  
225 model to the final model is fully automated. This refinement process is the  
226 most computationally intensive step and benefits the most from harness-  
227 ing the power of the Grid. Figure 2 shows the components in the EMAN  
228 refinement workflow, which forms a linear graph in which some compo-  
229 nents can be parallelized.

230 We have conducted experiments on workflow scheduling with two  
231 EMAN data-sets-GrOEL, a small data-set with 200 MB input data and  
232 rdv, a medium data-set with 2 GB input data. For these experiments, we  
233 used 6 nodes from the Itanium IA-64 cluster [i2-53 to i2-58] at UH and 7

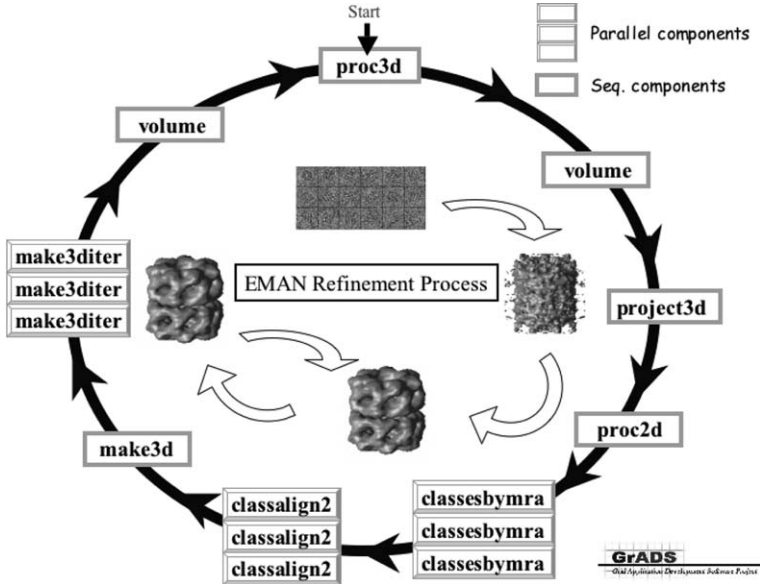


Fig. 2. EMAN refinement workflow.

234 nodes from the IA-32 cluster [torc1 to torc7] at UTK. Note that the test-  
 235 bed is heterogeneous in terms of architecture, CPU-speeds, memory and  
 236 storage. Also, note that “classesbymra” is the most computationally intensive  
 237 step in the EMAN refinement and is a parameter sweep that can be  
 238 distributed across multiple clusters. “classalign2” on the other hand  
 239 cannot be distributed across multiple clusters.

240 Table-I shows the results of the run of the rrv data on unloaded  
 241 resources on the testbed. The first column represents the name of the component  
 242 in the linear DAG. The second column denotes the resources chosen by the  
 243 Workflow scheduler for the particular component. The third column denotes  
 244 the number of instances mapped by the Workflow scheduler to the selected  
 245 resources. The last column denotes the time it took for that component to  
 246 run on the selected set of resources.

247 For the sequential and single-cluster components, the scheduler chose  
 248 the best node or cluster for execution. The interesting case is the case of  
 249 the parameter sweep step called “classesbymra”. From the execution time  
 250 of the “classesbymra” step, the following can be inferred:

- 251 — The makespan of the “classesbymra” step was 84 h 30 min [the  
 252 time the instances finished on the UH cluster]. Since the instances  
 253 at the UTK machines finished in 81 h 41 min, it can be inferred



**Table I. Results of EMAN Workflow Execution with rdv Data**

Component	Resources Chosen	Num Instances	Component Exec Time
Proc3d	i2-58	1	<1 min
Project3d	i2-58	1	1 h 48 min
Proc2d	i2-58	1	<1 min
<i>Classesbymra</i>	<i>i2-53 to i2-58</i>	68 [ <i>i2-*</i> ]	84 h. 30 min
	<i>torc1 to torc7</i>	42 [ <i>torc*</i> ]	81 h. 41 min
Classalign2	i2-53 to i2-58	379	45 min
Make3d	i2-58	1	47 min
Proc3d	i2-58	1	<1 min
Proc3d	i2-58	1	<1 min

254 that the load was optimally balanced across the two clusters since  
 255 the granularity of a single instance is greater than 7 h.  
 256 — The optimal load balance is primarily due to accurate performance  
 257 models and efficient Work-flow scheduling. Rank of a “classes-  
 258 bymra” instance on a node in UH cluster was 5077.76 and on a  
 259 node in UTK cluster was 8844.91.

260 For the GroEL data-set, the makespan for the classesbymra step for  
 261 heuristic scheduling was compared with that obtained from random sched-  
 262 uling. Random scheduling picks a node randomly for the next available  
 263 instance. The results in Table-II use 2 nodes from the UH cluster and  
 264 7 nodes from the UTK cluster and all the resources are unloaded. The  
 265 number in the braces after execution times indicate the average number of  
 266 classesbymra instances mapped to the site. From these results, it can be  
 267 inferred that accurate relative performance models on heterogeneous plat-  
 268 forms combined with heuristic scheduling result in good load balance of  
 269 the classesbymra instances when the grid resources are unloaded. Heuris-  
 270 tic scheduling is better than random scheduling by 25 percent in terms of  
 271 makespan length.

272 The second set of results shows the effect of loaded machines on  
 273 the quality of schedule. Five loaded nodes from the UH cluster and 7  
 274 unloaded nodes from UTK cluster were used for these experiments. From  
 275 the results in Table-III, it is observed that there is uneven load balance due  
 276 to loading of the UH nodes. Random scheduling does better because the  
 277 random distribution maps more instances to the unloaded UTK cluster  
 which had more nodes in the universe of resources. So, it can be inferred

**Table II. Results for GrOEL Data with Unloaded Resources**

	Heuristic Run Average	Random Run Average
<i>Exectime(uh)</i>	12 min 42 sec [38]	6 min 3 sec [17]
<i>Exectime(utk)</i>	11 min 47 sec [60]	15 min 48 sec [81]
<i>Makespan</i>	12 min 42 sec	15 min 48 sec

**Table III. Results for GrOEL Data with Loaded Resources**

	Heuristic Run Average	Random Run Average
<i>Exectime(uh)</i>	16 min 41 sec [60]	6 min 38 sec [44]
<i>Exectime(utk)</i>	7 min 51 sec [38]	10 min 28 sec [54]
<i>Makespan</i>	16 min 41 sec	10 min 28 sec

278 that for performance model based scheduling to work, either the underlying  
 279 set of resources should be reliable [implying advanced reservation] or  
 280 the variability of resource performance can be predicted and taken into  
 281 account during scheduling.

282 The third set of results show the effect of inaccurate performance  
 283 models on the quality of schedule. A rank value of 4.57 instead of 7.60  
 284 was used for a *classesbymra* instance on a UH node. Rank value for  
 285 the UTK nodes was kept correct. Six nodes from the UH cluster and 7  
 286 nodes from the UTK cluster were used. From the results in Table-IV it  
 287 can be inferred that, inaccurate relative performance models on different  
 288 heterogeneous platforms result in poor load balance of the *classesbymra*  
 289 instances.

**Table IV. Results for GrOEL Data with Inaccurate Performance Models**

	Heuristic Run Average	Random Run Average
<i>Exectime(uh)</i>	21 min 37 sec [77]	5 min 24 sec [45]
<i>Exectime(utk)</i>	3 min 57 sec [21]	10 min 30 sec [53]
<i>Makespan</i>	21 min 37 sec	10 min 30 sec

## 290 4. RESCHEDULING

291 Normally, a contract violation activates the GrADS *rescheduler*. The  
 292 retcheduling process must determine whether rescheduling is profitable,  
 293 based on the sensor data, estimates of the remaining work in the appli-  
 294 cation, and the cost of moving to new resources. If rescheduling appears  
 295 profitable, the rescheduler computes a new schedule (using the COP's map-  
 296 per) and contacts *rescheduling actuators* located on each processor. These  
 297 actuators use some mechanism to initiate the actual migration or load bal-  
 298 ancing. Sections 4.1 and 4.2 describe two rescheduling mechanisms that  
 299 we have explored. Both rely on application-level migration, although we  
 300 designed both so that the required additional programming is minimal.  
 301 Whether a migration is done or not, the rescheduler may contact the con-  
 302 tract monitor to update the terms of the contract.

### 303 4.1. Rescheduling by Stop and Restart

304 Our first approach to rescheduling relied on application migration  
 305 based on a stop/restart approach. The application is suspended and  
 306 migrated only when better resources are found for application execution.  
 307 When a running application is signaled to migrate, all application pro-  
 308 cesses checkpoint user specified data and terminate. The rescheduled exe-  
 309 cution is then launched by restarting the application on the new set  
 310 of resources, which then read the checkpointed data and continue the  
 311 execution.

#### 312 4.1.1. Implementation

313 We implemented a user-level checkpointing library called SRS (Stop  
 314 Restart Software)<sup>(20)</sup> to provide application migration support. Via calls  
 315 to SRS, the application can checkpoint data, be stopped at a particular  
 316 execution point, be restarted later on a different processor configuration  
 317 and be continued from the previous point of execution. SRS can trans-  
 318 parently handle the redistribution of certain data distributions (e.g., block  
 319 cyclic) between different numbers of processors (i.e., N to M processors).  
 320 The SRS library is implemented atop MPI and is hence limited to MPI-  
 321 based parallel programs. Because checkpointing in SRS is implemented at  
 322 the application rather than the MPI layer, migration is achieved by exiting  
 323 of the application and restarting it on a new system configuration.

324 The SRS library uses the Internet Backplane Protocol (IBP)<sup>(21)</sup> for  
 325 checkpoint data storage. An external component (e.g., the rescheduler)  
 326 interacts with a daemon called Runtime Support System (RSS). RSS

327 exists for the duration of the application execution and can span multiple  
328 migrations: Before the application is started, the launcher initiates  
329 the RSS daemon on the machine where the user invokes the GrADS  
330 application manager. The actual application, through the SRS, interacts  
331 with RSS to perform some initialization, to check if the application needs  
332 to be checkpointed and stopped, and to store and retrieve checkpointed  
333 data.

334 The contract monitor retrieves the application's registration through  
335 the Autopilot<sup>(3)</sup> infrastructure. The applications are instrumented with sensors  
336 that report the times taken for the different phases of the execution  
337 to the contract monitor.

338 The contract monitor compares the actual execution times with predicted  
339 ones and calculates the ratio. The tolerance limits of the ratio are  
340 specified as inputs to the contract monitor. When a given ratio is greater  
341 than the upper tolerance limit, the contract monitor calculates the average of  
342 the computed ratios. If the average is greater than the upper tolerance limit,  
343 it contacts the rescheduler, requesting that the application be migrated. If  
344 the rescheduler chooses not to migrate the application, the contract monitor  
345 adjusts its tolerance limits to new values. Similarly, when a given ratio is less  
346 than the lower tolerance limit, the contract monitor calculates the average  
347 of the ratios and lowers the tolerance limits, if necessary.

348 The rescheduler component evaluates the performance benefits that  
349 might accrue by migrating an application and initiates the migration.  
350 The rescheduler daemon operates in two modes: *migration on request* and  
351 *opportunistic migration*. When the contract monitor detects unacceptable  
352 performance loss for an application, it contacts the rescheduler to request  
353 application migration. This is called migration on request. Additionally,  
354 the rescheduler periodically checks for a GrADS application that has  
355 recently completed. If it finds one, the rescheduler determines if another  
356 application can obtain performance benefits if it is migrated to the newly  
357 freed resources. This is called opportunistic rescheduling. In both cases,  
358 the rescheduler contacts the Network Weather Service (NWS) for updated  
359 Grid resource information. The rescheduler uses the COP's performance  
360 model to predict remaining execution time on the new resources, remaining  
361 execution time on the current resources, and the overhead for migration  
362 and determines if migration is desirable.

#### 363 4.1.2. Evaluation

364 We have evaluated stop/restart rescheduling based on application  
365 migration for a ScaLAPACK<sup>(22)</sup> QR factorization application. The  
366 application was instrumented with calls to the SRS library that

367 checkpointed application data including the matrix A and the right-hand  
368 side vector B.

369 In the experiments, 4 UTK machines and 8 UIUC machines were  
370 used. The UTK cluster consists of 933 MHz dual-processor Pentium III  
371 machines running Linux and connected to each other by 100 Mb switched  
372 Ethernet. The UIUC cluster consists of 450 MHz single-processor Pentium  
373 II machines running Linux and connected to each other by 1.28 Gbit/sec-  
374 ond full-duplex Myrinet. The two clusters are connected via the Internet.

375 A given matrix size for the QR factorization problem was input to the  
376 application manager. Initially, the scheduler used the more powerful UTK  
377 cluster. However, five minutes after the start of the application, an artificial  
378 load was introduced on a UTK node, which could make it more efficient  
379 to execute the application the UIUC cluster.

380 The contract monitor requested the rescheduler to migrate the appli-  
381 cation due to the loss in predicted performance caused by the artificial  
382 load. The rescheduler evaluated the potential performance benefits due to  
383 migration and either migrated the application or allowed the application  
384 to continue on the original machines.

385 The rescheduler was operated in two modes — default and forced.  
386 In normal operation, the rescheduler works under default mode, while the  
387 forced mode allows the rescheduler to require the application to either  
388 migrate or continue on the same set of resources. Thus, if the default  
389 mode is to migrate the application, the forced mode will continue the  
390 application on the same set of resources and vice versa. For the experi-  
391 ments, results were obtained for both modes, allowing comparison of the  
392 scenarios and verification that the rescheduler made the right decision.

393 Figure 3 was obtained by varying the size of the matrices (i.e., the  
394 problem size) on the  $x$ -axis. The  $y$ -axis represents the execution time in  
395 seconds of the entire problem including the Grid overhead. For each prob-  
396 lem size, the left bar represents the running time when the application was  
397 not migrated and the right bar represents the time when the application  
398 was migrated.

399 Several observations can be made from Fig. 3. First, the time for  
400 reading checkpoints dominated the rescheduling cost, as it involves mov-  
401 ing data across the Internet and redistributing data to more processors. On  
402 the other hand, the time for writing checkpoints is insignificant since the  
403 checkpoints are written to IBP storage on local disks.

404 In addition, the rescheduling benefits are greater for large problem  
405 sizes because the remaining lifetime of the application is larger. For matrix  
406 sizes of 7000 and below, the migration cost overshadows the performance  
407 benefit due to rescheduling, while for larger sizes the opposite is true. Our  
408 rescheduler actually kept the computation on the original processors for




Fig. 3. Problem size and migration.

409 matrix sizes up to 8000. So, except for matrix size 8000, the rescheduler  
410 made the correct decision.

411 For matrix size 8000, the rescheduler assumed an experimen-  
412 tally-determined worst-case rescheduling cost of 900s while the actual  
413 rescheduling cost was about 420s. Thus, the rescheduler evaluated the  
414 performance benefit to be negligible. Hence, in some cases, the pessimistic  
415 approach of assuming a worst-case rescheduling cost will lead to underesti-  
416 mating the performance benefits due to rescheduling.

417 In another paper,<sup>(23)</sup> we examine the effects of other parameters (e.g.,  
418 the load and the time after the start of the application when the load was  
419 introduced) and the use of opportunistic rescheduling.

## 420 4.2. Rescheduling by Processor Swapping

421 Although very flexible, the natural stop, migrate and restart approach  
422 to rescheduling can be expensive: each migration event can involve large

423 data transfers. Moreover, restarting the application can incur expensive  
424 startup costs, and significant application modifications may be required for  
425 specialized restart code. Our process swapping approach, which was initially  
426 described in,<sup>(24)</sup> provides an alternative that is lightweight and easy  
427 to use, but less flexible than our migration approach.

#### 428 4.2.1. Basic Approach

429 To enable swapping, the MPI application is launched with more  
430 machines than will actually be used for the computation; some of these  
431 machines become part of the computation (the *active* set) while some do  
432 nothing initially (the *inactive* set). The user's application sees only the  
433 active processes in the main communicator (MPI\_Comm\_World); communication  
434 calls are hijacked, and user communication calls to the active  
435 set are converted to communication calls to a subset of the full process  
436 set.

437 During execution, the contract monitor periodically checks the performance  
438 of the machines and swaps slower machines in the active set with  
439 faster machines in the inactive set. This approach requires little application  
440 modification (as described in<sup>(24)</sup>) and provides an inexpensive fix for  
441 many performance problems. On the other hand, the approach is less flexible  
442 than migration—the processor pool is limited to the original set of  
443 machines, and the data allocation can not be modified.

444 MPI Swapping was implemented in the GrADS rescheduling architecture  
445 in which performance contract violations trigger rescheduling. The  
446 swapping rescheduler gathers information from sensors, analyzes performance  
447 information and determines whether and where to swap processes.  
448 We have designed and evaluated several policies<sup>(6)</sup> and we have experimentally  
449 evaluated our process swapping implementation using an N-body  
450 solver.<sup>(6,24)</sup>

#### 451 4.2.2. Evaluation

452 This section describes how we used the MicroGrid to evaluate the  
453 GrADS rescheduling implementation.

454 *The MicroGrid* Understanding the dynamic behavior of rescheduling  
455 approaches for Grids requires experiments under a wide range of resource  
456 network configurations and dynamic conditions. Historically, this has been  
457 difficult, and simplistic experiments with either a few resource configurations  
458 or simple models of applications have been used. We use a  
459 general tool, the MicroGrid, which supports systematic, repeatable, scalable,

460 and observable study of dynamic Grid behavior, to study the behavior  
461 of the process swapping rescheduling system on a range of network  
462 topologies. We show data from a run of an N-body simulation, under  
463 the N-N rescheduling system, running on the MicroGrid emulation of a  
464 distributed Grid resource infrastructure.

465 The MicroGrid allows complete Grid applications to execute on a  
466 set of virtual Grid resources. It exploits scalable parallel machines as  
467 compute platforms for the study of applications, network, compute, and  
468 storage resources with high fidelity. For more information on the Micro-  
469 Grid see.<sup>(5,25,26)</sup>

470 *Experiments with process-swapping rescheduling* The first step in using  
471 the MicroGrid is to define the virtual resource and network infrastruc-  
472 ture to be emulated. For our demonstration, we created a virtual Grid  
473 which is a subset of the GrADS testbed, consisting of machines at UCSD,  
474 UIUC, and UTK. The virtual Grid includes two clusters at UTK and  
475 UIUC and a single compute node at UCSD. The UTK cluster includes  
476 three 550 MHz Pentium II nodes. The UIUC cluster consists of three  
477 450 MHz Pentium II machines. Both clusters are internally connected by  
478 Gigabit Ethernet. The single UCSD machine is a 1.7 GHz Athlon node.  
479 The latency between UCSD and the other two sites is 30 ms and between  
480 UTK and UIUC the latency is 11 ms. These configurations are described  
481 for MicroGrid in standard Domain Modeling Language (DML) and a  
482 simple resource description for the processor nodes.

483 The MicroGrid uses a Linux cluster at UCSD to implement its  
484 Grid emulation. We allocated two 2.4 GHz dual-processor Xeon machines  
485 for network simulation, and seven 450 MHz dual-processor Pentium II  
486 machines to model the compute nodes in the above virtual Grid.

487 To perform the process swapping rescheduling experiment on the vir-  
488 tual Grid, we first launched the MicroGrid daemons (instantiating the vir-  
489 tual Grid). From this point on, all processes launched on UCSD, UTK,  
490 or UIUC machines ran on the virtual Grid nodes. Second, we launched  
491 the contract monitor infrastructure (the Autopilot manager and contract  
492 monitor processes) and rescheduler process on the UCSD node. Third, we  
493 launched the N-body simulation application to the UTK and UIUC clus-  
494 ters which then connected to the contract monitor and rescheduler. All  
495 three of the initial active application processes started on the UTK nodes.  
496 At (virtual) time 80 s, we added two competitive processes to consume  
497 CPU time on one UTK machine. The rescheduling infrastructure detected  
498 poor performance and migrated all three working application processes to  
499 the UIUC cluster by time 150 s. Figure 4 shows the resulting application  
500 progress, first slowed by the competitive load, then increased by the migra-  
501 tion to free resources.






Fig. 4. Emulated application progress during N-body demonstration run.

## 502 5. FUTURE DIRECTIONS: VIRTUAL GRIDS

503 GrADS provided a foundation for an evolving compilation and exe-  
 504 cution infrastructure, *GrAD-Soft*, which we and others have used to con-  
 505 duct a range of application experiments<sup>(27–31)</sup> such as those described in  
 506 this paper. These application experiments have not only validated the basic  
 507 GrADS approach, but have also informed our focus on the most critical  
 508 remaining challenges. These efforts are the focus of our new *Virtual Grid*  
 509 *Application Development Software (VGrADS)* project.

510 One of the key lessons of the GrADS project is that the complexity of  
 511 grid resource environments induces complexity in both application devel-  
 512 opment and execution. First, execution on a shared grid of heterogeneous  
 513 resources such as the TeraGrid forces an application developer to explic-  
 514 itly consider resource heterogeneity, dynamically fluctuating loads, and the  
 515 interaction between local users and resource policies. There is little ques-  
 516 tion that this complicates grid application programs, increasing program-  
 517 ming difficulty and discouraging grid applications. Second, a rigid view  
 518 by applications that prescribe a “perfect” set of resources, complicates  
 519 resource management requiring search of a great expanse of resources  
 520 and rapid, detailed matching of applications to resources. This too is a  
 521 major technological challenge. Finally, it is our observation from work-  
 522 ing with many leading grid application teams that when faced with com-  
 523 plex application performance structure and complex resource environments  
 524 compounded with poor predictive information, expert programmers are  
 525 reduced to use of *ad hoc* heuristics (albeit sophisticated ones) that require  
 526 much tuning and debugging to achieve acceptable resource utilization and  
 527 application performance.

528 Building on the knowledge and infrastructure of the GrADS project,  
529 our new approach adopts the concept of a *Virtual Grid (VGrid)* as a fun-  
530 damental element of the software architecture which supports a separation  
531 of concerns for VGrADS.

532 Vgrids cleanly separate high-level programming tools, applications,  
533 and services from the complexity of dynamic grid scheduling and resource  
534 management. This approach is analogous to one that has proven effective  
535 in sequential and parallel computing contexts, where optimizations tar-  
536 get abstract uniprocessors and multiprocessors rather than the phys-  
537 ical resources themselves. The same concept will form the basis of our  
538 approach to simplifying the task of Grid application development.

539 Virtual grids support simpler high-level program preparation tools by  
540 providing simplified resource management and simple monitored perfor-  
541 mance guarantees. This supports the development and use of more power-  
542 ful programming abstractions. We believe that virtual grids will enable the  
543 execution system to quickly and scalably identifying appropriate resources  
544 for applications, simplifying both application and system-level resource  
545 management. Finally, the virtual grid approach simplifies performance  
546 monitoring and resource adaptation by making explicit (and application  
547 neutral) the performance expectations and guarantees. In short, virtual  
548 grids provide a cleaner separation of responsibilities across the program  
549 preparation, execution system, and monitoring and adaptation systems,  
550 allowing each to be simplified and as a result more effective.

551 VGrADS research focuses on two major areas: execution environ-  
552 ments and programming tools. *Execution environment* research explores the  
553 synthesis, coordination, and measurement of grid resources. The goals of  
554 this work are to explore (1) aggregation and virtualization of resource and  
555 Grid service aggregates; (2) intelligent, rapid resource selection and man-  
556 agement in complex, heterogeneous environments; (3) performance mea-  
557 surement and tuning to achieve high individual application performance;  
558 and (4) fault-resilience through replication and intermediate, program state  
559 management. The resulting system will enable the nimble adaptation of  
560 applications to changing Grid conditions.

561 *Programming tools* research explores the mapping of two distinct,  
562 high-level programming models to VGrids. The *abstract parallel machine*  
563 model treats a computation as a collection of parallel tasks without  
564 concern for mapping that computation to the actual hardware. The  
565 *abstract component machine* model, on the other hand, represents a  
566 computation as a (possibly dynamic) graph of component invocations  
567 with specific data dependencies. In this model, applications and services  
568 might be high-level scripts that invoke operations from a component  
569 integration frame-work. The VGrADS execution system, working on

570 behalf of the application, will use VGrids to instantiate both of these  
571 programming models.

## 572 ACKNOWLEDGMENTS

573 The research reported here was supported under National Science  
574 Foundation awards 9975020, 0103759, and 0331645.

## 575 REFERENCES

- 576 1. I. Foster and C. Kesselman (eds.), *The Grid: Blueprint for a New Computing Infra-*  
577 *structure*, 2nd Ed., Morgan Kaufmann (2003).
- 578 2. K. Kennedy, M. Mazina, J. Mellor-Crummey, K. Cooper, L. Torczon, F. Berman,  
579 A. Chien, H. Dail, O. Sievert, D. Angulo, I. Foster, D. Gannon, S. L. Johnsson,  
580 C. Kesselman, R. Aydt, D. Reed, J. Dongarra, S. Vadhiyar, and R. Wolski, Towards  
581 a Framework for Preparing and Executing Adaptive Grid Programs, *Proceedings of*  
582 *NSF Next Generation Systems Program Workshop (International Parallel and Distrib-*  
583 *uted Processing Symposium)*, Fort Lauderdale, Florida (April 2002).
- 584 3. R.L. Ribler, H. Simitci, and D.A. Reed, The Autopilot Performance-directed Adaptive  
585 Control System, *Future Generation Computer Systems*, **18**(1):175–187 (September 2001).
- 586 4. F. Vraalsen, R.A. Aydt, C.L. Mendes, and D.A. Reed, Performance Contracts: Pre-  
587 dicting and Monitoring Grid Application Behavior, *Lecture Notes in Computer Science*,  
588 Vol. 2242, pp. 154–165, Springer Verlag (November 2001).
- 589 5. H. Song, X. Liu, D. Jakobsen, R. Bhagwan, X. Zhang, K. Taura, and A. Chien,  
590 The MicroGrid: A Scientific Tool for Modeling Computational Grids, *Proceedings of*  
591 *SC2000* (November 2000).
- 592 6. O. Sievert and H. Casanova, Policies for Swapping MPI Processes, *Proceedings of*  
593 *HPDC-12, the Symposium on High Performance and Distributed Computing* (June 2003).
- 594 7. B. Barish and R. Weiss, Ligo and detection of gravitational waves, *Physics Today*,  
595 **52**(10) (1999).
- 596 8. S. Hastings, T. Kurc, S. Langella, U. Catalyurek, T. Pan, and J. Saltz, Image Process-  
597 ing on the Grid: A Toolkit or Building Grid-enabled Image Processing Applications,  
598 *3rd International Symposium on Cluster Computing and the Grid* (2003).
- 599 9. K. Taura and A. Chien, A Heuristic Algorithm for Mapping Communicating Tasks  
600 on Heterogeneous Resources, *Heterogeneous Computing Workshop* (May 2000).
- 601 10. S. Vadhiyar and J. Dongarra, A Metascheduler for the Grid, *Proceedings of the High*  
602 *Performance Distributed Computing Conference* (July 2002).
- 603 11. R. Wolski, J. Plank, J. Brevik, and T. Bryan, G-commerce: Market Formulations Con-  
604 trolling Resource Allocation on the Computational Grid, *Proceedings of 2001 Interna-*  
605 *tional Parallel and Distributed Processing Symposium (IPDPS)* (March 2001).
- 606 12. Condor Team, Condor Version 6.4.7 Manual, [//www.cs.wisc.edu/condor/manual/v6.4/](http://www.cs.wisc.edu/condor/manual/v6.4/).
- 607 13. S. Fitzgerald, I. Foster, C. Kesselman, G. von Laszewski, W. Smith, and S. Tuecke,  
608 A Directory Service for Configuring High-Performance Distributed Computations, *Pro-*  
609 *ceedings of the 6th IEEE Symposium on High-Performance Distributed Computing*, pp.  
610 365–375 (August 1997), URL [papers/fitzgerald-hpdc97-mds.pdf](http://papers/fitzgerald-hpdc97-mds.pdf).

- 611 14. R. Wolski, N.T. Spring, and J. Hayes, The network weather service: a distributed  
612 resource performance forecasting service for metacomputing, *Future Generation Com-*  
613 *puter Systems*, **15**(5-6):757-768 (1999), URL [citeseer.nj.nec.com/wolski98network.html](http://citeseer.nj.nec.com/wolski98network.html).
- 614 15. M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory*  
615 *of Np-Completeness*, MIT Press (1979).
- 616 16. H. Casanova, A. Legrand, D. Zagorodnov, and F. Berman, Heuristics for Scheduling  
617 Parameter Sweep applications in Grid environments, *9th Heterogeneous Computing*  
618 *workshop (HCW'2000)* (2000).
- 619 17. Tracy D. Braun et al. A Comparison of eleven Static Heuristics for Mapping a Class  
620 of Independent Tasks onto Heterogeneous Distributed Computing Systems, *Journal of*  
621 *Parallel and Distributed Computing*, **61**:810-837 (2001).
- 622 18. G. Marin, *Semi-Automatic Synthesis of Parameterized Performance Models for Scientific*  
623 *Programs*, Master's thesis, Department of Computer Science, Rice University (April  
624 2003).
- 625 19. S. Ludtke, P. Baldwin, and W. Chiu, EMAN: Semiautomated Software for High-  
626 Resolution Single-Particle Reconstructions, *J. Struct. Biol.*, **128**:82-97 (1999), URL  
627 <http://ncmi.bcm.tmc.edu/homes/stevel/EMAN/doc>.
- 628 20. S. Vadhiyar and J. Dongarra, SRS A Framework for Developing Malleable and  
629 Migratable Parallel Applications for Distributed Systems, *Parallel Processing Letters*,  
630 **13**(2):291-312 (June 2003), iISSN 0129-6264.
- 631 21. J.S. Plank, M. Beck, W. Elwasif, T. Moore, M. Swany, and R. Wolski, The Internet  
632 Backplane Protocol: Storage in the Network, *NetStore99: The Network Storage Sym-*  
633 *posium* (1999).
- 634 22. L.S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra,  
635 S. Hammerling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R.C. Whaley, *Scal-*  
636 *APACK User's Guide* (1997).
- 637 23. S. Vadhiyar and J. Dongarra, A Performance Oriented Migration Framework for the  
638 Grid, *IEEE Computing Clusters and the Grid (CCGrid, <http://www.ccgriid.org>)* (May  
639 12-15 2003).
- 640 24. O. Sievert and H. Casanova, A Simple MPI Process Swapping Architecture for Iter-  
641 ative Applications, *The International Journal of High Performance Computing Applica-*  
642 *tions* (2004), to appear.
- 643 25. X. Liu and A. Chien, Traffic-based Load Balance for Scalable Network Emulation,  
644 *Proceedings of SC2003* (November 2003).
- 645 26. H. Xia, H. Dail, H. Casanova, F. Berman, and A. Chien, Evaluating the GrADS  
646 Scheduler in Diverse Grid Environments Using the MicroGrid (May 2003), submitted for  
647 publication.
- 648 27. A. Petitet, S. Blackford, J. Dongarra, B. Ellis, G. Fagg, K. Roche, and S. Vadhiyar,  
649 Numerical Libraries and the Grid, *Proceedings of SC'01* (November 2001).
- 650 28. M. Ripeanu, A. Iamnitchi, and I. Foster, Cactus Application: Performance Predictions  
651 in a Grid Environment, *Proceedings of European Conference on Parallel Computing*  
652 *(EuroPar)2001* (August 2001).
- 653 29. W. Chrabakh and R. Wolski, *GrADSAT: A Parallel SAT Solver for the Grid*, Techni-  
654 cal Report CS-2003-05, University of California, Santa Barbara (2003), available from  
655 <http://www.cs.ucsb.edu/research/trcs/index.shtml>.
- 656 30. H. Dail, *A Modular Framework for Adaptive Scheduling in Grid Application Develop-*  
657 *ment Environments*, Master's thesis, University of California, San Diego, Department of  
658 Computer Science and Engineering (Mardh 2002), available as UCSD Technical Report  
659 CS2002-0698.

- 660 31. A. Mandal, *Mapping HPF onto the Grid*, Technical report TR03-417, Department of  
661 Computer Science, Rice University, Houston (November 2002), URL <http://www.cs.rice.edu/~anirban/MStthesis.ps.gz>.  
662

Uncorrected Proof