

GrADSolve—a grid-based RPC system for parallel computing with application-level scheduling[☆]

Sathish S. Vadhiyar^{a,*} and Jack J. Dongarra^b

^aDepartment of Computer Science, University of Tennessee, 107, Ayres Hall, Knoxville, TN 37996-1301, USA

^bDepartment of Computer Science, University of Tennessee, Oak Ridge National Laboratory, USA

Received 9 September 2003

Abstract

Although some existing Remote Procedure Call (RPC) systems provide support for remote invocation of parallel applications, these RPC systems lack powerful scheduling methodologies for the dynamic selection of resources for the execution of parallel applications. Some RPC systems support parallel execution of software routines with simple modes of parallelism. Some RPC systems statically choose the configuration of resources for parallel execution even before the parallel routines are invoked remotely by the end user. These policies of the existing systems prevent them from being used for remotely solving computationally intensive parallel applications over dynamic computational Grid environments. In this paper, we discuss a RPC system called GrADSolve that supports execution of parallel applications over Grid resources. In GrADSolve, the resources used for the execution of parallel application are chosen dynamically based on the load characteristics of the resources and the characteristics of the application. Application-level scheduling is employed for taking into account both the application and resource properties. GrADSolve also stages the user's data to the end resources based on the data distribution used by the end application. Finally, GrADSolve allows the users to store execution traces for problem solving and use the traces for subsequent solutions. Experiments are presented to prove that GrADSolve's data staging mechanisms can significantly reduce the overhead associated with data movement in current RPC systems. Results are also presented to demonstrate the usefulness of utilizing the execution traces maintained by GrADSolve for problem solving.

© 2003 Elsevier Inc. All rights reserved.

Keywords: RPC; Grid; GrADSolve; Application-level scheduling; Data staging; Execution traces

1. Introduction

Remote Procedure Call (RPC) mechanisms have been studied extensively and have been found to be powerful abstractions for distributed computing [12,13]. In RPC frameworks, the end user invokes a simple routine to solve problems over remote distributed resources. A number of RPC frameworks have been implemented and are widely used [1,15–17,27,34,38,41,45]. In addition to providing simple interfaces for uploading applications into the distributed systems and for remote invocation of the applications, some of the RPC systems

also provide service discovery, resource management, scheduling, security and information services.

The role of RPC in Computational Grids [23] has been the subject of many recent studies [19,20,33,37,39]. Computational Grids consist of large number of machines ranging from workstations to supercomputers and strive to provide transparency to end users and high performance for end applications. While high performance is achieved by the parallel execution of applications on a large number of Grid resources, user transparency can be achieving by employing RPC mechanisms. Hence, Grid-based RPC systems need to be built to provide the end users the capability to invoke remote parallel applications on Grid resources using a simple sequential procedure call.

Though there are a large number of RPC systems that support remote invocation of parallel applications [15,16,22,29,35,37,38], the selection of resources for the

[☆]This work is supported in part by the National Science Foundation contract GRANT #EIA-9975020, SC #R36505-29200099 and GRANT #EIA-9975015.

*Corresponding author.

E-mail addresses: vss@cs.utk.edu (S.S. Vadhiyar), dongarra@cs.utk.edu (J.J. Dongarra).

execution of parallel applications in these systems does not take into account the dynamic load aspects that are associated with Computational Grids. Some of the parallel RPC systems [8,18,21,22,26,28–30,35] are mainly concerned with providing robust and efficient interfaces for the service providers to integrate their parallel applications into the systems, and for the end users to remotely use these parallel services. In these systems the users or the service providers have to provide their own scheduling mechanisms if the resources for end application execution have to be dynamically chosen. In most cases, the users and the service providers lack the expertise to implement scheduling techniques. Some RPC systems [15,38] provide scheduling services in addition to the basic functionality of providing interfaces to the service providers and users. In these systems, scheduling methodologies are employed to choose between different parallel domains that implement the same parallel services. But within a parallel domain, the number and configuration of resources are fixed at the time when the services are uploaded into the RPC system and hence are not adaptive to the load dynamics of the Grid resources.

In this paper, we propose a Grid-based RPC system called GrADSolve¹ that enables the users to invoke MPI applications on remote Grid resources from a sequential environment. GrADSolve combines the easy-to-use RPC mechanisms of NetSolve [2,15] and powerful application-level scheduling mechanisms inherent in the GrADS [11] project. Application-level scheduling has been proven to be a powerful scheduling technique for providing high performance [9,10].

In addition to providing easy-to-use interfaces for the service providers to upload the parallel applications into the system and for the end users to remotely invoke the parallel applications, GrADSolve also provides interfaces for the service providers or library writers to upload execution models that provide information about the predicted execution costs of the applications. This information is used by GrADSolve to perform application-level scheduling and to dynamically choose the resources for the execution of the parallel applications based on the load dynamics of the Grid resources. GrADSolve also uses the data distribution information provided by the library writers to partition the users' data and stage the data to the different resources used for the application execution. Our experiments show that the data staging mechanisms in GrADSolve help reduce the data staging times in RPC systems by 20–50%. GrADSolve also uses the popular Grid computing tool, Globus [25] for transferring data between the user and the end resources and for launching the application

on the Grid resources. In addition to the above features, GrADSolve also enables the users to store execution traces for a problem run and use the execution traces for the subsequent problem runs. This feature helps in significantly reducing the overhead incurred due to the selection of the resources for application execution and the staging of input data to the end resources.

Thus, the contributions of our research are:

- (1) Design and development of an RPC system that utilizes standard Grid Computing mechanisms for invocation of remote parallel applications from a sequential environment.
- (2) Selection of resources for parallel application execution based on load conditions of the resources and application characteristics.
- (3) Maintenance of execution traces for problem runs.

Section 2 describes in brief the GrADS and NetSolve projects. The architecture of GrADSolve, the various entities in the GrADSolve system and the support for the entities in the GrADSolve system are explained in Section 3. The support in the GrADSolve system for maintaining execution traces is explained in Section 4. In Section 5, the experiments conducted in GrADSolve are explained and results are presented to demonstrate the usefulness of the data staging mechanisms and execution traces in GrADSolve. Section 6 looks at the related efforts in the development of parallel RPC systems. Section 7 presents conclusions and future work.

2. Background of GrADSolve

GrADSolve evolved from two projects, GrADS [11] and NetSolve [15]. In this section, the overviews of GrADS and NetSolve are presented.

2.1. The GrADS project

The Grid Application Development Software (GrADS) project is a multi-university research project which works to simplify distributed heterogeneous computing in the same way that the World Wide Web simplified information sharing. The GrADS project intends to provide tools and technologies for the development and execution of applications in a Grid environment. In the GrADS vision, the end user simply presents their parallel application to the framework for execution. The framework is responsible for scheduling the application on an appropriate set of resources, launching and monitoring the execution, and, if necessary, rescheduling the application on a different set of resources. A high-level view of the GrADS

¹The system is called GrADSolve since it is derived from the experiences of the GrADS [11] and NetSolve [2,15] projects.

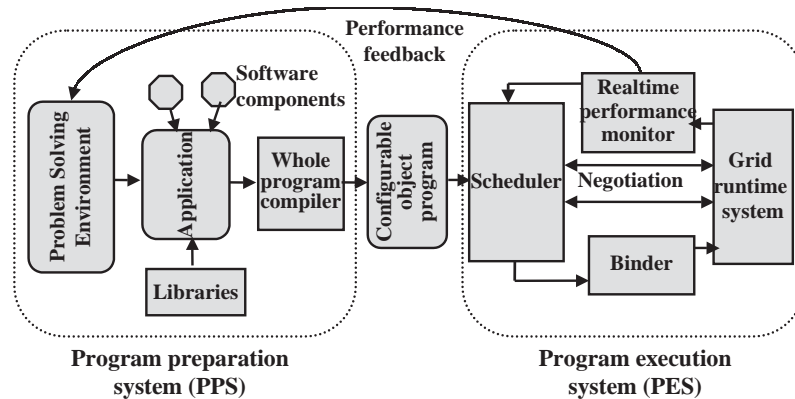


Fig. 1. The GrADS architecture.

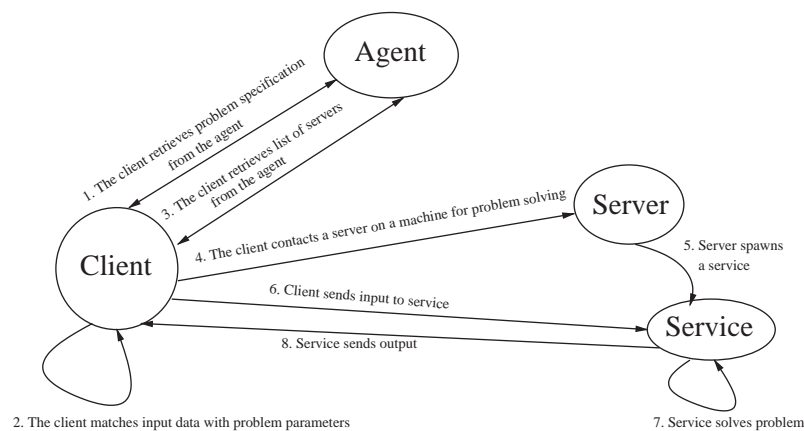


Fig. 2. Overview of NetSolve system.

architecture is shown in Fig. 1. For more details, the readers are referred to [11].

2.2. NetSolve—a brief overview

NetSolve [15] is a Grid computing system developed at University of Tennessee. It is a Remote Procedure Call (RPC)-based system used for solving numerical applications over remote machines. The NetSolve system consists of three main components—agent, server and client. The working of the NetSolve system is illustrated in Fig. 2. Although NetSolve supports remote execution of parallel applications, the amount of parallelism is fixed at the time the server daemons are started. For more details, the readers are referred to [15].

3. The GrADSolve system

The general architecture of GrADSolve is shown in Fig. 3. At the core of the GrADSolve system is a XML database implemented with Apache Xindice [3]. GrADSolve uses XML as a language for storing information

about different Grid entities. This database maintains four kinds of tables—*users*, *resources*, *applications* and *problems*. The Xindice implementation of the XML-RPC standard [45] was used for storing and retrieving information to and from the XML database. There are three human entities involved in GrADSolve—*administrators*, *library writers* and *end users*. The role of these entities in GrADSolve and the functions performed by the GrADSolve system for supporting these entities are explained in the following subsections.

3.1. Administrators

The GrADSolve administrator is responsible for managing the users and resources of the GrADSolve system. The administrator initializes the XML database and creates entities for different users in the XML database by specifying a *user configuration file*. The user configuration file contains information for different users, namely the user account names for different resources and the location of the home directories on different resources in the GrADSolve system. The administrator also creates the *resources* table in the

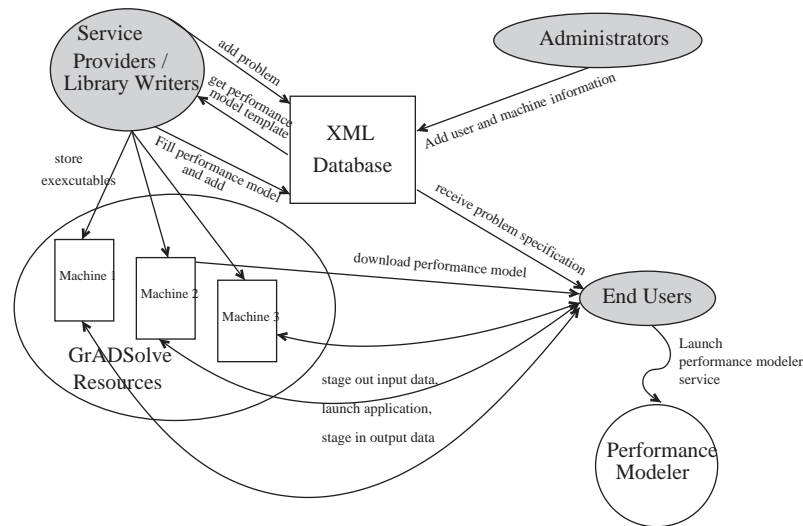


Fig. 3. Overview of GrADSolve system.

```

PROBLEM sparse_solve
C ROUTINE sparse_solve(IN string package,
                       IN int N,
                       IN double SM[N][N]<nnz,indices,pointer>,
                       IN int nnz,
                       IN int indices[nnz],
                       IN int pointer[N],
                       IN double rhs[N],
                       IN double pivot,
                       IN int permutation,
                       OUT double sol[N])
TYPE = parallel

```

Fig. 4. An example GrADSolve IDL for a sparse factorization problem.

Xindice database and adds entries for different resources in the GrADSolve system by specifying a *resource configuration file*. The various information in the configuration file, namely the names of the different machines, their computational capacities, the number of processors in the machines and other machine specifications, are stored as XML documents.

3.2. Library writers

The library writer uploads his application into the GrADSolve system by specifying an Interface Definition Language (IDL) file for the application. Even though there are a number of robust IDL systems for High Performance Computing including the OMG IDL [31] and the SIDL (Scientific IDL) from the BABEL project [40], GrADSolve implements and uses its own IDL system. In addition to the primitives for MPI applications and complex data types that are supported by the existing IDL systems, GrADSolve supports sparse matrix data types. Supporting sparse matrix data types

is essential for integrating popular and efficient high performance libraries including PETSC [5–7], AZTEC [4], SuperLU [42], etc. Fig. 4 illustrates an IDL with support for sparse matrices that is unique to GrADSolve IDL. In the IDL file, the 3rd parameter, *SM*, is a sparse matrix represented by a compressed-row format.

After the library writer submits the IDL file to the GrADSolve system, GrADSolve translates the IDL file to a XML document similar to the mechanisms in SIDL. The GrADSolve translation system also generates a wrapper program that acts as an entry point for remote execution of the actual function. The wrapper program performs initialization of the parallel environment, reads input data from files, invokes the actual parallel routine and stores output data to files. The GrADSolve system then compiles the wrapper program with the object files and the libraries specified in the IDL file and with the appropriate parallel libraries if the application is specified as a parallel application in the IDL file. The GrADSolve system then stages the executable to the different resources in the Grid using the Globus

```

int areResourcesSufficient(int N, int NB, double* A,
                          double* B,
                          RESOURCEINFO* resourceInfo,
                          SCHEDULESTRUCT* schedule){
}

int getExecutionTimeCost(int N, int NB, double* A,
                        double* B,
                        RESOURCEINFO* resourceInfo,
                        SCHEDULESTRUCT* schedule,
                        double* cost){
}

int mapper(int N, int NB, double* A, double* B,
           RESOURCEINFO* resourceInfo,
           SCHEDULESTRUCT* inputSchedule,
           SCHEDULESTRUCT* mapperSchedule){
}

```

Fig. 5. A performance model template generated by the GrADSolve system for the ScaLAPACK QR problem.

GridFTP mechanisms and stores the locations of the executables in the XML database.

The library writer also has the option of adding an execution model for the application. If the library writer wants to add an execution model, he executes the *getperfmodel_template* utility, specifying the name of the application. The utility retrieves the problem description of the application from the XML database and generates a performance model template file. The performance model template file contains definitions for three functions to help the library writers to convey information about his library to the GrADSolve system—*areResourcesSufficient* for conveying if a given set of resources are adequate for problem solving, *getExecutionTimeCost* for conveying the predicted execution cost of the application if executed on a given set of resources and an optional function *mapper* for specifying the data distribution of the different data used by the application. The performance model template file generated by the *getperfmodel_template* for a ScaLAPACK QR problem is shown in Fig. 5. The library writer uploads his execution model by executing the *add_perfmodel* utility. The *add_perfmodel* utility uploads the execution model for the application by storing the location of the execution model in the XML database corresponding to the entry for the application.

3.3. End users

The end users solve problems over remote GrADSolve resources by writing a client program in C or Fortran. The client program includes an invocation of a

routine called *gradsolve()* passing to the function, the name of the end application and the input and output parameters needed by the end application.

The invocation of the *gradsolve()* routine triggers the execution of the GrADSolve *Application Manager*. As a first step, the Application Manager verifies if the user has credentials to execute applications on the GrADSolve system. GrADSolve uses Globus Grid Security Infrastructure (GSI) [14] for the authentication of users. If the application exists in the GrADSolve system, the Application Manager registers the problem run in the *problems* table of the XML database. The Application Manager then retrieves the problem description from the XML database and matches the user's data with the input and output parameters required by the end application.

If an execution model exists for the end application, the Application Manager downloads the execution model from the remote location where the library writer had previously stored the execution model. The Application Manager compiles the execution model programs with algorithms for scheduling [32,46] and starts the application-specific *Performance Modeler* service. The Application Manager then retrieves the list of machines in the GrADSolve system from the *resources* table in the XML database, and retrieves various performance characteristics of the machines including the peak performance of the resources, the load on the machines, the latency and the bandwidth of the networks between the machines and the free memory available on the machines from the Network Weather Service (NWS) [44]. The Application Manager passes the list of machines, along with the resource characteristics to the Performance Modeler service to determine if the resources are sufficient to solve the problem. If the resources are sufficient, the Application Manager proceeds to the *Schedule Generation* phase.

In the Schedule Generation phase, the Application Manager first determines if the end application has an execution model. If an execution model exists, the Application Manager contacts the Performance Modeler service and passes the problem parameters and the list of machines with the machine capabilities. The Performance Modeler service uses the execution model supplied by the library writer along with certain scheduling heuristics [32,46] to determine a final schedule for application execution and returns the final list of machines to the Application Manager. Along with the final list of machines and the predicted execution cost for the final schedule, the Performance Modeling service also returns information about the data distribution for the different data in the end application. If an execution model does not exist for the end application, the Schedule Generation phase adopts default scheduling strategies to generate the final schedule for end application execution. At the end of the Schedule

Generation phase, the GrADSolve Application Manager receives a list of machines for final application execution. The Application Manager then stores the status of the problem run and the final schedule in the *problems* table of the XML database corresponding to the entry for the problem run.

The Application Manager then creates working directories on the remote machines of the final schedule for end application execution and enters the *Application Launching* phase. The Application Launching phase consists of several important functions. The Application Launcher stores the input data to files and stages these files to the corresponding remote machines chosen for application execution using the Globus GridFTP mechanisms. If data distribution information for an input data does not exist, the Application Launcher stages the entire input data to all the machines involved in the end application execution. If the information regarding data distribution for an input data exists, the Application Launcher stages only the appropriate portions of the data to the corresponding machines. This kind of selective data staging significantly reduces the time needed for the staging for entire data especially if a large amount of data is involved.

After the staging of input data, the Application Launcher launches the end application on the remote machines chosen for the final schedule using the Globus MPICH-G [24] mechanism. The end application reads the input data that were previously staged by the Application Launcher and solves the problem. The end application then stores the output data to the corresponding files on the machines in the final schedule. If the end application finished execution, the Application Launcher copies the output data from the remote machines to the user's address space. The staging in of the output data from the remote locations is a reverse operation of the staging out of the input data to the remote locations. The GrADSolve Application Manager finally returns success state to the user client program.

4. Execution traces in GrADSolve—storage, management and usage

One of the unique features in the GrADSolve system is the ability provided to the users to store and use execution traces of problem runs. There are many applications in which the outputs of the problem depend on the exact number and configuration of the machines used for problem solving. For example, considering the problem of adding a large number of double precision numbers, one of the parallel implementations of the problem is to partition the list of double precision numbers among all processes of the parallel application, compute local sums of the numbers in each process and

then compute the global sum of the local sums computed on each process. The final sum obtained for the same set of double precision numbers may vary from one problem run to another depending on the number of elements in each partition, the number of processes used in the parallel application and the actual processors used in the computation. This is due to the impact of the round off errors caused by the addition of double precision numbers. In general, ill-conditioned problems or unstable algorithms can give rise to vast changes in output results due to small changes in input conditions. For these kinds of applications, the user may desire to use the same input environment for all problem runs. Also, during the testing of new numerical algorithms over the Grid, different groups working on the algorithm may want to ensure that the same results are obtained when the algorithms are executed with same input data on the same configuration of resources.

To guarantee reproducibility of numerical results in the above situations, GrADSolve provides capability to the users to store *execution traces* of problem runs and use the execution traces during subsequent executions of the same problem with the same input data. For storing an execution trace of the current problem run, the user executes his GrADSolve program with a configuration file called *input.config* in the working directory containing the line, `TRACE_FLAG = 1`.

During the registration of the problem run with the XML database, the value of the `TRACE_FLAG` variable is stored. The GrADSolve Application Manager proceeds to other stages of its execution. After the end application completes its execution and the output data are copied from the remote machines to the user's address space, the Application Manager, under default mode of operation, removes the remote working directories used for storing the files containing the input data for the end application. But when the user wants to store the execution trace of the problem run, i.e. when the *input.config* file contains "`TRACE_FLAG = 1`" line, the Application Manager retains the input data used for the problem run in the remote machines. At the end of the problem run, the Application Manager generates an output configuration file called *output.config* containing the line, `TRACE_KEY = <key>`. The value *key* in the *output.config* is a pointer to the execution trace stored for the problem run.

When the user wants to execute the problem with the execution trace previously stored, he executes his client program specifying the line, `TRACE_KEY = <key>` in the *input.config* file. The value *key* in the *input.config*, is the same value previously generated by the GrADSolve Application Manager when the execution trace was stored. The Application Manager first checks if the `TRACE_KEY` exists in the *problems* table of the XML database. If the `TRACE_KEY` does not exist, the Application Manager displays an error message to the

user and aborts operation. If the TRACE_KEY exists for an execution trace of a previous problem run, the Application Manager registers the current problem run with the XML database and proceeds to the other stages of its execution. During the Schedule Generation phase, the Application Manager, instead of generating a schedule for the execution of the end application, retrieves the schedule used for the previous problem run corresponding to the TRACE_KEY, from the *problems* table in the XML database. The Application Manager then checks if the capacities of the resources in the schedule at the time of trace generation are comparable to the current capacities of the resources. If the capacities are not comparable, the Application Manager displays an error message to the user and aborts the operation. If the capacities are comparable, the Application Manager proceeds to the rest of the phases of its execution. During the Application Launching phase, the Application Manager, instead of staging the input data to remote working directories, copies the input data and the data distribution information, used in the previous problem run corresponding to the TRACE_KEY, to the remote working directories. The use of the same number of machines and the same input data used in the previous schedule also guarantees the use of the same data distribution for the current problem run. Thus GrADSolve guarantees the use of the same execution environment used in the previous problem run for the current problem run, and hence guarantees reproducibility of numerical results.

To support the storage and use of execution traces in the GrADSolve system, two trigger functions are associated with the XML database. One trigger function called *trace_usage_trigger* updates the last usage time of an execution trace when the execution trace is used for a problem run. Another trigger function called *cleanup_trigger* is used for periodically deleting entries in the *problems* table of the XML database thereby maintaining the size of the *problems* table in the database. The *cleanup_trigger* is invoked whenever a new entry corresponding to a problem run is added to the *problems* table. The *cleanup_trigger* employs a longer duration for those problem runs for which execution traces were stored.

5. Experiments and results

The GrADS testbed consists of about 40 machines from University of Tennessee (UT), University of Illinois, Urbana-Champaign (UIUC) and University of California, San Diego (UCSD). For the sake of clarity, our experimental testbed consists of four machines:

- a 933 MHz Pentium III machine with 512 MBytes of memory located in UT,

- a 450 MHz Pentium II machine with 256 MBytes of memory located in UIUC and
- two 450 MHz Pentium III machines with 256 MBytes of memory located in UCSD.

The two UCSD machines are connected to each other by 100 MBytes switched Ethernet. Machines from different locations are connected by Internet. In the experiments, GrADSolve was used to remotely invoke the ScaLAPACK driver for solving a linear system of equations, $AX = B$. The driver invokes ScaLAPACK QR factorization for the factorization of matrix, A . Block-cyclic distribution was used for the matrix A . A GrADSolve IDL was written for the driver routine and an execution model that predicts the execution cost of the QR problem was uploaded into the GrADSolve system. The GrADSolve user invokes the remote parallel application by passing the size of the matrix A and the right-hand side vector, B to the *gradsolve()* call.

GrADSolve was operated in three modes. In the first mode, the execution model did not contain information about the data distribution used in the ScaLAPACK driver. In this case, GrADSolve transported the entire data to each of the locations used for the execution of the end application. This mode of operation is practiced in RPC systems that do not support the information regarding data distribution. In the second mode, the execution model contained information about the data distribution used in the end application. In this case, GrADSolve transported only the appropriate portions of the data to the locations used for the execution of end application. In the third mode, GrADSolve was used with an execution trace corresponding to a previous run of the same problem. In this case, data is not staged from the user's address space to the remote machines, but temporary copies of the input data used in the previous run are made for the current problem run.

Fig. 6 shows the times taken for data staging and other GrADSolve overhead for different matrix sizes




Fig. 6. Data staging and other GrADSolve overhead.

Table 1
Machines chosen for application execution

Matrix size	Machines
1000	1 UT machine
2000	1 UT machine
3000	1 UT machine
4000	1 UT machine
5000	1 UT, 1 UIUC machines
6000	1 UIUC, 1 UCSD machines
7000	1 UIUC, 1 UCSD machines
8000	1 UT, 1 UIUC, 2 UCSD machines

and for the three modes of GrADSolve operation. Since the times taken for the execution of the end application are same in all the three modes, we focus only on the times taken for data staging and possible Grid overheads. The machines that were chosen by the GrADSolve application-level scheduler for the execution of end application for different matrix sizes are shown in Table 1. The UT machine was used for smaller problem sizes since it had larger computing power than other machines. For matrix size, 5000, a UIUC machine was also used for the execution of parallel application. For matrix sizes, 6000 and 7000, the available memory in the UT machine at the time of the experiments was less than the memory needed for the problems. Hence UIUC and UCSD machines were used. For matrix size, 8000, all four machines were needed to accommodate the problem. All the above decisions were automatically made by the GrADSolve system taking into account the sizes of the problems and the resource characteristics at the time of the experiments.

Comparing the first two modes in Fig. 6, we find that for smaller problem sizes, the times taken for data staging in both the modes are the same. This is because only one machine was used for problem execution and the same amount of data is staged in both the modes when only one machine is involved for problem execution. For larger problem sizes, the times for data staging with distribution information is less than 20–55% of the times taken for staging the entire data to remote resources. Thus the use of data distribution information in GrADSolve can give significant performance benefits when compared to staging the entire data that is practiced in some of the RPC systems. Data staging in the third mode is basically the time taken for creating temporary copies of data used in the previous problem runs in remote resources. We find this time to be negligible when compared to the first two modes. Thus execution traces can be used as caching mechanisms to use the previously staged data for problem solving. The GrADSolve overheads for all the three modes are found to be the same. This is because of the small number of machines used in the experiments. For experiments when large number of machines are used,

we predict that the overheads will be higher in the first two modes than in the third mode. This is because in the first two modes, the application-level scheduling will explore large number of candidate schedules to determine the machines used for end application while in the third mode, a previous application-level schedule will be retrieved from the database and used.

6. Related work

A number of parallel RPC systems have been built in the context of Object Management Group (OMG) [30], Common Component Architecture Forum (CCA) [18] and Grid research efforts [15,38,39].

The Object Management Group (OMG) [30] has been dealing with specifying objects for both sequential and parallel applications. The Data Parallel CORBA specification describes parallel objects that enable the object implementer to take advantage of parallel resources for achieving high performance. The specification defines interfaces for both the implementers of the objects and the client to use the remote parallel services. The specification does not deal with dynamic selection of resources for parallel computing. The PaCO [35,36] and PaCO++ [19,18] systems from the PARIS project in France are implemented within the CORBA [17] framework to encapsulate MPI applications in RPC systems. The data distribution and redistribution mechanisms in PaCO are much more robust than in GrADSolve and support invocation of remote parallel applications either from sequential or parallel client programs. Recently, the PARIS project has been investigating coupling multiple applications of different types in Grid frameworks [20,33]. Similar to the Data Parallel CORBA specification, the parallel CORBA objects in the PaCO projects do not support dynamic selection of resources for application execution as in GrADSolve. The selection of resources for parallel execution taking into account the load aspects of the resources is a necessity in dynamic Computational Grids. Also, GrADSolve supports Grid-related security models by employing Globus mechanisms. And finally, GrADSolve is unique in maintaining execution traces that can help bypass the resource selection and data staging phases.

The Common Component Architecture Forum (CCA) [18] has been investigating the deployment and use of both parallel and sequential components. Their “MxN Redistribution” working group has been dealing with the issues of data redistribution when multiple parallel components are coupled together. The CUMULVS MxN interface [26] from Oak Ridge National Laboratory, the PAWS environment [8] from Los Alamos National Laboratory and the PARDIS SPMD objects [28] from Indiana University work within the

CCA to develop a parallel RPC standard. The main goals of these systems include providing interoperability between different components, building user interfaces for conveying information about the parallel data, developing communication schedules to communicate the data between different components and synchronizing data transfers. These projects delegate the responsibility of scheduling or choosing the end resources for parallel application to the implementers of parallel components. In most cases, the implementers of the components lack the expertise to include scheduling technologies. In GrADSolve, application-level scheduling is an integral component of the system and requires the implementors of the parallel components to only provide information about their parallel applications.

NetSolve [15] and Ninf [38] are Grid computing systems that support task parallelism by the asynchronous execution of a number of remote sequential applications. OmniRPC [37] is an extension of Ninf and supports asynchronous RPC calls to be made from OpenMP programs. But similar to the approaches in NetSolve and Ninf, OmniRPC supports only master-worker models of parallelism. NetSolve and Ninf also support remote invocation of MPI applications, but the amount of parallelism and the locations of the resources to be used for the execution are fixed at the time when the applications are uploaded to the systems and hence are not adaptive to dynamic loads in the Grid environments. Recently, Grid-RPC [39] has been proposed to standardize the efforts of NetSolve and Ninf. The current Grid-RPC standard does not specify scheduling methodologies to choose the resources for execution of remote parallel applications.

7. Conclusions and future work

In this paper, an RPC system for efficient execution of remote parallel software was discussed. The efficiency is achieved by dynamically choosing the machines used for parallel execution and staging the data to remote machines based on data distribution information. The GrADSolve RPC system also supports maintaining and utilizing execution traces for problem solving. Our experiments showed that the GrADSolve system is able to adapt to the problem sizes and the resource characteristics and yielded significant performance benefits with its data staging and execution trace mechanisms.

Interfaces for the library writers for expressing more capabilities of the end application are currently being designed. These capabilities include the ability of the application to be pre-empted and resumed later with different processor configuration. These capabilities will allow GrADSolve to adapt to changing Grid scenarios. The current GrADSolve system employs application-

level scheduling requiring the implementers to provide information about their libraries. In the future, we plan to employ methods for automatic determination of the information about the libraries similar to the efforts in the Prophecy [43] project. Though GrADSolve currently provides basic security by the authentication of users and service providers through Globus mechanisms, it does not provide privacy with regard to message transactions and also does not support validation of results. We plan to employ encryption of data to provide privacy and signed software mechanisms to assure integrity of results.

Acknowledgments

The authors would like to thank the managers of the GrADS project for providing valuable input during the development of GrADSolve. We acknowledge the use of machines in the GrADS testbed for the experiments conducted in the research. We also thank the research teams from different institutions, namely the Pablo research group from the University of Illinois, Urbana-Champaign, the Grid Research and Innovation Laboratory (GRAIL) from the University of California, San Diego and the Innovative Computing Laboratory (ICL) from University of Tennessee, for the support and maintenance of the machines in the GrADS testbed and for enabling of experiments needed for the research.

References

- [1] Apache Xindice, <http://xml.apache.org/xindice>.
- [2] P. Arbenz, W. Gander, M. Oettli, The remote computation system, *Parallel Comput.* 23 (1997) 1421–1428.
- [3] D. Arnold, S. Agrawal, S. Blackford, J. Dongarra, M. Miller, K. Seymour, K. Sagi, Z. Shi, S. Vadhiyar, Users' Guide to NetSolve V1.4.1, Innovative Computing Dept. Technical Report ICL-UT-02-05, University of Tennessee, Knoxville, TN, June 2002.
- [4] AZTEC, <http://www.cs.sandia.gov/CRF/aztec1.html>.
- [5] S. Balay, W.D. Gropp, L.C. McInnes, B.F. Smith, Efficient management of parallelism in object oriented numerical software libraries, in: E. Arge, A.M. Bruaset, H.P. Langtangen (Eds.), *Modern Software Tools in Scientific Computing*, Birkhauser Press, Basel, 1997, pp. 163–202.
- [6] S. Balay, W.D. Gropp, L.C. McInnes, B.F. Smith, PETSc home page, <http://www.mcs.anl.gov/petsc>, 1999.
- [7] S. Balay, W.D. Gropp, L.C. McInnes, B.F. Smith, PETSc 2.0 users manual, Technical Report ANL-95/11 - Revision 2.0.24, Argonne National Laboratory, 1999.
- [8] P.H. Beckman, P.K. Fasel, W.F. Humphrey, S.M. Mniszewski, Efficient Coupling of Parallel Applications Using PAWS, in: *Proceedings of the Seventh IEEE International Symposium on High Performance Distributed Computing*, IEEE, New York, 1998, pp. 215–223.
- [9] F. Berman, High-performance schedulers, in: I. Foster, C. Kesselman (Eds.), *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, Los Altos, CA, 1999, pp. 279–293, ISBN 1-55860-475-8.

- [10] F. Berman, R. Wolski, The AppLeS Project: a status report, Proceedings of the Eighth NEC Research Symposium, Berlin, Germany (1997).
- [11] F. Berman, A. Chien, K. Cooper, J. Dongarra, I. Foster, D. Gannon, L. Johnsson, K. Kennedy, C. Kesselman, J. Mellor-Crummey, D. Reed, L. Torczon, R. Wolski, The GrADS project: software support for high-level grid application development, *Internat. J. High Performance Appl. Supercomput.* 15 (4) (2001) 327–344.
- [12] B. Bershad, T. Anderson, E. Lazowska, H. Levy, Lightweight remote procedure call, *ACM Trans. Comput. Systems* 8 (1) (1990) 37–55.
- [13] A. Birrell, B. Nelson, Implementing remote procedure calls, *ACM Trans. Comput. Systems* 2 (1) (1984) 39–59.
- [14] R. Butler, D. Engert, I. Foster, C. Kesselman, S. Tuecke, J. Volmer, V. Welch, A national-scale authentication infrastructure, *IEEE Comput.* 33 (12) (2000) 60–66.
- [15] H. Casanova, J. Dongarra, NetSolve: a network server for solving computational science problems, *Internat. J. Supercomputer Appl. High Performance Comput.* 11 (3) (1997) 212–223.
- [16] C.-C. Chang, G. Czajkowski, T. von Eicken, MRPC: a high performance RPC system for MPMD, *Parallel Comput.* 29 (1) (1999) 43–66.
- [17] CORBA, <http://www.corba.org>.
- [18] Common Component Architecture, <http://www.cca-forum.org>.
- [19] A. Denis, C. Prez, T. Priol, Portable parallel CORBA objects: an approach to combine parallel and distributed programming for grid computing, in: Proceedings of the Seventh International Euro-Par'01 Conference (EuroPar'01), Springer, Berlin, 2001, pp. 835–844.
- [20] A. Denis, C. Prez, T. Priol, Towards high performance CORBA and MPI middlewares for grid computing, in: C.A. Lee (Ed.), Proceedings of the Second International Workshop on Grid Computing, Lecture Notes in Computer Science, Vol. 2242, Springer, Berlin, 2001, pp. 14–25.
- [21] A. Denis, C. Pérez, T. Priol, PadicoTM: an open integration framework for communication middleware and runtimes, *Future Generation Comput. Systems* 19 (2003) 575–585.
- [22] A. Denis, C. Pérez, T. Priol, Achieving portable and efficient parallel CORBA objects, *Concurrency and Computation: Practice and Experience* 15 (10) (2002) 891–909.
- [23] I. Foster, C.K. (Eds.), *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, Los Altos, CA, 1999, ISBN 1-55860-475-8.
- [24] I. Foster, N. Karonis, A grid-enabled MPI: message passing in heterogeneous distributed computing systems, in: Proceedings of SuperComputing'98 (SC98), Orlando, FL (1998).
- [25] I. Foster, C. Kesselman, Globus: a metacomputing infrastructure toolkit, *Internat. J. Supercomput. Appl.* 11 (2) (1997) 115–128.
- [26] G.A. Geist, J.A. Kohl, P.M. Papadopoulos, CUMULVS: providing fault-tolerance, visualization and steering of parallel applications, *Internat. J. High Performance Comput. Appl.* 11 (3) (1997) 224–236.
- [27] Java Remote Method Invocation (Java RMI), java.sun.com/products/jdk/rmi.
- [28] K. Keahey, D. Gannon, PARDIS: a parallel approach to CORBA, in: Proceedings of the Sixth IEEE International Symposium on High Performance Distributed Computing, IEEE, New York, 1997, pp. 31–39.
- [29] J. Maassen, R. van Nieuwpoort, R. Veldema, H. Bal, T. Kielmann, C. Jacobs, R. Hofman, Efficient Java RMI for parallel programming, *ACM Trans. Programming Languages Systems* 23 (6) (2001) 747–775.
- [30] Object Management Group, <http://www.omg.org>.
- [31] OMG IDL, http://www.omg.org/gettingstarted/omg_idl.htm.
- [32] A. Petitet, S. Blackford, J. Dongarra, B. Ellis, G. Fagg, K. Roche, S. Vadhiyar, Numerical Libraries and the Grid: The GrADS Experiments with ScaLAPACK, *J. High Performance Appl. Supercomput.* 15 (4) (2001) 359–374.
- [33] C. Prez, T. Priol, A. Ribes, A parallel CORBA component model for numerical code coupling, in: C.A. Lee (Ed.), Proceedings of the Third International Workshop on Grid Computing, Lecture Notes in Computer Science, Springer, Berlin, 2002.
- [34] R. Rabenseifner, The DFN remote procedure call tool for parallel and distributed applications, in: K. Franke, U. Hübner, W. Kalfa (Eds.), Proceedings of Kommunikation in Verteilten Systemen—KiVS'95, Chemnitz-Zwickau, 1995, pp. 415–419.
- [35] C. René, T. Priol, MPI code encapsulation using parallel CORBA object, in: Proceedings of the Eighth IEEE International Symposium on High Performance Distributed Computing, IEEE, New York, 1999, pp. 3–10.
- [36] C. René, T. Priol, MPI code encapsulating using parallel CORBA object, *Cluster Comput.* 3 (4) (2000) 255–263.
- [37] M. Sato, M. Hirano, Y. Tanaka, S. Sekiguchi, OmniRPC: A grid RPC facility for cluster and global computing in OpenMP, in: Workshop on OpenMP Applications and Tools (WOMPAT2001), 2001.
- [38] H.N.M. Sato, S. Sekiguchi, Design and implementations of ninf: towards a global computing infrastructure, *Future Generation Comput. Systems, Metascomputing Issue* 15 (5–6) (1999) 649–658.
- [39] K. Seymour, H. Nakada, S. Matsuoka, J. Dongarra, C. Lee, H. Casanova, Overview of GridRPC: a remote procedure call API for grid computing, in: M. Parashar (Ed.), Lecture Notes in Computer Science, Grid Computing—GRID 2002, Vol. 2536, Third International Workshop, Springer, Baltimore, MD, USA, 2002, pp. 274–278.
- [40] SIDL from BABEL project, <http://www.llnl.gov/CASC/components/babel.html>.
- [41] Simple Object Access Protocol (SOAP), <http://www.w3.org/TR/SOAP>.
- [42] SuperLU, <http://crd.lbl.gov/~xiaoye/SuperLU>.
- [43] V.E. Taylor, X. Wu, J. Geisler, R. Stevens, Using kernel couplings to predict parallel application performance, in: Proceedings of the 11th IEEE International Symposium on High Performance Distributed Computing, IEEE, New York, 2002, pp. 125–135.
- [44] R. Wolski, N. Spring, J. Hayes, The network weather service: a distributed resource performance forecasting service for meta-computing, *J. Future Generation Comput. Systems* 15 (5–6) (1999) 757–768.
- [45] XML-RPC, <http://www.xmlrpc.com>.
- [46] A. Yarkhan, J. Dongarra, Experiments with scheduling using simulated annealing in a grid environment, in: M. Parashar (Ed.), Lecture Notes in Computer Science, Grid Computing—GRID 2002, Vol. 2536, Third International Workshop, Springer, Baltimore, MD, USA, 2002, pp. 232–242.