UT-CS-01-460

April 28, 2001

Numerical Libraries And The Grid: The GrADS Experiments With ScaLAPACK

Antoine Petitet, Susan Blackford, Jack Dongarra, Brett Ellis, Graham Fagg, Kenneth Roche, and Sathish Vadhiyar

University of Tennessee

Computer Science Department

Innovative Computing Laboratory

Numerical Libraries And The Grid: The GrADS Experiments With ScaLAPACK

Antoine Petitet[‡], Susan Blackford[†], Jack Dongarra[†], Brett Ellis[†], Graham Fagg[†], Kenneth Roche[†], and Sathish Vadhiyar[†]

Abstract:
This paper describes an overall framework for the design of numerical libraries on a computational Grid of processors where the processors may be geographically distributed and under the control of a Grid-based scheduling system. A set of experiments are presented in the context of solving systems of linear equations using routines from the ScaLAPACK software collection along with various grid service components, such as Globus, NWS, and Autopilot.

Motivation On The Grid

The goal of the *Grid Application Development Software (GrADS) project [1]* is to simplify distributed heterogeneous computing in the same way that the World Wide Web simplified information sharing over the Internet. The GrADS project is exploring the scientific and technical problems that must be solved to make Grid applications development and performance tuning for real applications an everyday practice. This requires research in four key areas; each validated in a prototype infrastructure that will make programming on grids a routine task:

1. Grid software architectures that facilitate information flow and resource negotiation among applications, libraries, compilers, linkers, and runtime systems;
2. Base software technologies, such as scheduling, resource discovery, and communication, to support development and execution of performance-efficient Grid applications;
3. Languages, compilers, environments, and tools to support creation of applications for the Grid and solution of problems on the Grid; and
4. Mathematical and data structure libraries for Grid applications, including numerical methods for control of accuracy and latency tolerance.

In this paper we will describe the issues of developing a prototype system designed specifically for the use of numerical libraries in the grid setting and the results of experiments with routines from the ScaLAPACK library [2] on the Grid.

Motivation On Numerical Libraries

[‡] Sun France Benchmark Center, Paris, France.
[†] Department of Computer Science, University of Tennessee, Knoxville TN 37996. The work is supported in part by the National Science Foundation contract GRANT #E81-9975020, SC R36505-29200099, R01-1030-09.

The primary goals of our effort in numerical libraries are to develop a new generation of algorithms and software libraries needed for the effective and reliable use of dynamic, distributed and parallel environments, and to validate the resulting libraries and algorithms on important scientific applications. To consistently obtain high performance in the Grid environment will require advances in both algorithms and supporting software.

Some of the challenges in this arena have already been encountered. For example, to make effective use of current high-end machines, the software must manage both communication and the memory hierarchy. This problem has been attacked with a combination of compile-time and run-time techniques. On the Grid, the increased scale of computation, depth of memory hierarchies, range of latencies, and increased variability in the run-time environment will make these problems much harder.

To address these issues, we must rethink the way that we build libraries. The issues to consider include software architecture, programming languages and environments, compile versus run-time functionality, data structure support, and fundamental algorithm design. Among the challenges that we must deal with are:

- The library software must support performance optimization and algorithm choice at run time.
- The architecture of software libraries must facilitate efficient interfaces to a number of languages, as well as effective support for the relevant data structures in those languages.
- New algorithmic techniques will be a prerequisite for latency tolerant applications.
- New scalable algorithms that expose massive numbers of concurrent threads will be needed to keep parallel resources busy and to hide memory latencies.
- Library mechanisms that will interact with the grid to dynamically deploy resources in solving the posed users' problems.

These considerations lead naturally to a number of important areas where research in algorithm design and library architecture is needed.

Grid-Aware Libraries To enable the use of the Grid as a seamless computing environment, we are developing parameterizable algorithms and software annotated with performance contracts. These annotations will help a dynamic optimizer tune performance for a wide range of architectures. This tuning will, in many cases, be accomplished by having the dynamic optimizer and run-time system provide input parameters to the library routines that will enable them to make a resource-efficient algorithm selection. We are also developing new algorithms that use adaptive strategies by interacting with other GrADS components. For example, libraries will incorporate performance contracts for dynamic negotiation of resources, as well as runtime support for adaptive strategies to allow the compiler, scheduler, and runtime system to influence the course of execution.

We are using the "Grid information service" and the "Grid event service" to obtain the information needed to guide adaptation in the library routines. These services will be discussed in more detail later in this paper.

Latency-Tolerant Algorithm Remote latency is one of the major obstacles in achieving high efficiency on today's high-performance computers. The growing gap between the speed of the microprocessors and memory coupled with a deep memory hierarchy implies that the memory subsystem has become a large performance factor in modern computer systems such as the Department of Energy's Advanced Strategic Computing Initiative (ASCI) computers. In the unpredictable and dynamic world of the Grid, this problem will be even worse. Research into latency tolerant algorithms that explore a wider portion of the latency and bandwidth/memory space is needed. Furthermore, tools are needed for measuring and managing latency. We are designing and constructing numerical libraries that are parameterized to allow their performance to be optimized over a range of current and future memory hierarchies, including the ones expected in computational Grids.

Compiler-Ready Libraries In the past, library development has typically focused on one architecture at a time with the result that much of the work must be repeated to migrate the code to a new architecture and its memory hierarchy. We are exploring the design and development of parameterized libraries that permit performance tuning across a wide spectrum of memory hierarchies. Some developers of portable libraries rely on tools like the HPC Compiler to analyze and parallelize their programs. These compilers are large and complex; they do not always discover parallelism when it is available. Automatic parallelization may be adequate for some simple types of scientific programming, but experts writing libraries find it frustrating, because they must often perform tedious optimizations by hand—optimizations that could be handled by the compiler if it had a bit of extra information from the programmer. Programmers may find it both easier and more effective to annotate their code with information that will help the compiler generate the desired code with the desired behavior. We have begun to identify opportunities where information about algorithms contained in library functions can help the compiler and the run–time environment and work with the GrADS compiler group to develop a new system of annotation to provide information about semantics and performance that aids in compilation. At the library interface level, this would include memory–hierarchy tuning parameters and semantic information, such as dependency information to make it possible to block the LU factorization of a matrix, and floating–point semantic information (for example, to indicate that it is acceptable to reorder certain floating point operations or to handle exceptions in particular ways). The goal is to make it possible to "build in" to the compiler knowledge about these libraries far beyond what can be derived by a compile-time analysis of the source.

Current Numerical Subroutine Libraries

Current numerical libraries for distributed memory machines are designed for heterogeneous computing, and are based on MPI [3] for communication between processes. Two such widely used libraries are ScaLAPACK and PETSc [4], designed for dense and sparse matrix calculations, respectively. ScaLAPACK assumes a two-dimensional block cyclic data distribution among the processes, and PETSc provides a block-row or application-specific data distribution. The user must select the number of processes associated with an MPI communicator, and also select the specific routine/algorithm to be invoked.

In the case of ScaLAPACK, the user has total control over the exact layout of the data, specifying the number of block rows and the number of block columns in each block to be distributed. These blocks are then distributed cyclically to maintain proper load balance of the application on the machine. It is the user's responsibility to distribute the data prior to invoking the ScaLAPACK routine. If the user makes a poor choice of data layout, then this can significantly affect his application's performance. All data that has been locally distributed can be explicitly accessed in the user's program.

For PETSc, the library has a variety of data distribution routines from which to choose. The user can select the default block row distribution where the size of the block is determined by the PETSc system as a function of the size of the global matrix to be distributed. PETSc will choose the block size such that the matrix is evenly distributed among the processes. The user can also select an application-specific block row distribution whereby the size of the block is a function of the application to be run on that process's data. In contrast to ScaLAPACK, the user does not have explicit access to individual elements in the data structure. The user must use specialized PETSc matrix manipulation routines to access the matrix data.

For both libraries, the user is responsible for making many decisions on how the data is decomposed, the number of processors used, and which software is to be chosen for the solution. Given the size of the problem to be solved, the number of processors available, and certain characteristics of the processors, such as CPU speed and the amount of available memory per processor, heuristics exist for selecting a good data distribution to achieve efficient performance. In addition, the ScaLAPACK Users' Guide [2] provides a performance model for estimating the computation time given the speed of the floating point operations, the problem size, and the bandwidth and latency associated with the specifics of the parallel computer. Equation 1 provides the model used by ScaLAPACK for solving a dense system of equations.

$$T(n,p) = C_f t_f + C_v t_v + C_m t_m \qquad \textbf{Eq 1}$$

where

$C_f = \frac{2n^3}{3p}$	Total number of floating-point operations per processor
$C_v = (3 + \frac{1}{4}\log_2 p)\frac{n^2}{\sqrt{p}}$	Total number of data items communicated per processor
$C_m = n(6 + \log_2 p)$	Total number of messages
t_f	Time per floating point operation
t_v	Time per data item communicated
t_m	Time per message
n	Matrix size
p	Number of processors
$T(n,p)$	Parallel execution time for a problem of size n run on p processors

The performance model assumes that the parallel computer is homogeneous with respect to both the processors and communication network. With the Grid both of these assumptions are incorrect and a performance model becomes much more complex. With the dynamic nature of the grid environment, the Grid "scheduler" must assume the task of deciding how many processors to use and the placement of data. This selection would be performed in a dynamic fashion by using the state of the processors and the communication behavior of the network within the grid in conjunction with a performance model for the application. The system would then determine the data layout, the number and location of the processors, and perhaps the algorithm selection for a given problem for the best "time to solution" on the Grid.

Adapting Current Libraries To The Grid Environment

Our goal is to adapt the existing distributed memory software libraries to fit into the Grid setting without too many changes to the basic numerical software. We want to free the user from having to allocate the processors, make decisions on which processors to use to solve the problem, how the data is to be decomposed to optimally solve the problem,

allocating the resources, starting the message passing system, monitoring the progress, migrating or restarting the application if a problem is encountered, collecting the results from the processors, and returning the processors to the pool of resources.

Implementation Outline Of The GrADS' ScaLAPACK Demo

The ScaLAPACK experiment demonstrates the feasibility of solving large-scale numerical problems over the Grid and analyzes the added cost of performing the calculation over machines spanning geographically distributed sites. We solve a simple linear system of equations over the Grid using Gaussian elimination with partial pivoting via the ScaLAPACK routine, PDGESV. We illustrate how users, without much knowledge of numerical libraries can seamlessly use numerical library routines over the Grid. We also outline the steps that are necessary for a library writer to integrate his library into the Grid System. (The Appendix contains a more detailed description of these parts.) While ease of use is an important aspect of the Grid, performance improvement and the ability to perform tasks that are too large to be performed on a single tightly-coupled cluster are important motivations behind the use of the Grid. Our experiments show that effective resources can be selected to solve the ScaLAPACK problem and, for our experiments, scalability (as the problem size and the number of processors increases) of the software is maintained.

Before the user can start his application, the Grid system is assumed to have initialized three components – *Globus MDS [5], NWS [6] sensors* on all machines in the *Globus MDS* repository, and the *Autopilot Manager/Contract Monitor[7]*. We assume that the user has already communicated with the Grid system (Globus) and has been authenticated to use the grid environment. The Globus MDS maintains a repository of all available machines in the Grid, and the NWS sensors monitor a variety of system parameters for all of the machines contained in the Globus MDS repository. This information is necessary for modeling the performance of the application, and for making scheduling decisions of the application on the Grid. Autopilot was designed and is maintained at the University of Illinois, Urbana-Champaign, (UIUC). It is a system for monitoring the application execution and enabling corrective measures, if needed, to improve the performance while the application is executing. . The Autopilot Manager must be running on one of the machines in the Grid prior to the start of the experiment. The library routine (in our experiment this is the ScaLAPACK code itself) must be instrumented with calls to Autopilot monitoring. It is hoped in the future that this instrumentation could be done by the compilation system. However, for our experiment, all instrumentation was inserted by hand. As the application executes, performance data is accumulated and communicated to the Autopilot manager via the *Contract Monitor*. This Contract Monitor must be started on the machine to which the user will be initiating the experiment. Work is underway to build a generic contract monitor that will be able to automatically maintain and detect the Autopilot managers that exist on the Grid. At present, however, the name of the machine running the Autopilot manager must be

supplied by the user (see details below). There is also an ongoing effort in the compiler group of GrADS to produce a generic Configurable Object Program (COP) [1] and hence the user will not be required to maintain separate executables for each machine in the Grid. However, at present, this feature not available. After the user has compiled the executable, the user is responsible for copying this executable to every machine in the Grid.
After these preliminary steps have been completed, the user is now ready to execute his application on the Grid. The user interface to the ScaLAPACK experiment in the GrADS system is the routine, Grads_Lib_Linear_Solve. This routine is the main numerical library driver of the GrADS system and calls other components in GrADS. It accepts the following inputs: the matrix size, the block size, and an input file. This input file contains information such as the machine on which the Autopilot manager is running, the path to the contract monitor, the subset of machines in the Grid on which the user could run his application, the path to the executable on each machine, etc. In the future, this file will not be necessary since most of the parameters in the file will be maintained by the Grid system. For purposes of this experiment the input matrices are either read from a URL or randomly generated. A more flexible user interface for the generation and distribution of matrices is being developed. Future development will also encompass the automatic determination of the value of the "block size" that will yield the "best performance" of the application on the Grid. This "block size" determines the data distribution of the matrices to the processors in the Grid, and likewise the size of the computational kernel to be used in the block-partitioned algorithm.

The Grads_Lib_Linear_Solve routine performs the following operations:

1. Sets up the problem data
2. Creates the "coarse grid" of processors and their NWS statistics by calling the resource selector.
3. Refines the "coarse grid" into a "fine grid" by calling the performance modeler.
4. Invokes the contract developer to commit the resources in the "fine grid" for the problem.

Repeat Steps 2-4 until the "fine grid" is committed for the problem.

5. Launches the application to execute on the committed "fine grid".

The Appendix contains additional details of these steps, as well as pseudo-code for Grads_Lib_Linear_Solve() and APIs for the GrADS components.

1. Setting up the Problem

This step supplies the GrADS system data structures with the user-passed parameters. In the future, this step will involve contacting the GrADS name server to get information about the Autopilot managers and contract monitor. This step will also involve the building of the Configurable Object Program (COP), which will be an extension to the normal object code, encapsulating annotations about the runtime system and user preferences. These annotations will later be used for rescheduling the applications when the Grid parameters change. And finally, future enhancements will also entail the automatic determination of the best "block size" for this DGEMM-based application on each of the machines in the Grid.

2. Resource Selector

The resource selector contacts the MDS server, maintained by the Information Sciences Institute (ISI) as part of the Globus system, to check the status of the machines needed by the user for the application. If MDS does not detect failures with the machines, the resource selector then contacts the Network Weather Service (NWS), maintained at the University of Tennessee, to obtain machine-specific information pertaining to available CPU, available memory, and latency and bandwidth between machines. At the end of the resource selection step, a "coarse grid" is formed. This "coarse grid" is essentially all of the machines available along with the statistics returned by NWS.

3. Performance Modeler

The performance modeler calculates the "speed" of the machine as a percentage of the peak Mflop/s on the machine. The user currently supplies the peak Mflop/s rate of the machine. In the future, the GrADS system will be able to determine the peak Mflop/s rate of a machine. The percentage used in the calculation of "speed" is heuristically chosen by observing previous ScaLAPACK performance, in this case routine PDGESV, on the given architecture. Typically, PDGESV achieves approximately 75% of the local DGEMM (matrix multiply) performance per processor, and as the percentage of peak performance attained by DGEMM is approximately 75%, we use a heuristic measure of 50% of the theoretical peak performance for routine PDGESV from ScaLAPACK.

The performance modeler then performs the following steps in determining the collection of machines to use for the problem:

i. Determine the amount of physical memory that is needed for the current problem.
ii. If possible, find a fastest machine in the coarse grid that has the memory needed to solve the problem.
iii. Find a machine in the coarse grid that has the maximum average bandwidth relative to the other machines in the grid. Add this to the fine grid.
iv. Do the following:
    a. Find the next machine in the coarse grid that has maximum average bandwidth relative to the machines that are already in the fine grid.
    b. Calculate the new time estimate for the application using the machines in the fine grid. This time estimate is calculated by executing a performance model for the application. We are assuming that the library writer has provided a performance model for the library routine. This performance model takes into account the speed of the machines as well as the latency and bandwidth as returned by NWS.
    c. Repeat Step iv until the time estimate for the application run time increases.
v. If a single machine is found in step ii, the time estimate for the problem using the single machine is compared with the time estimate for the problem using the machines found in steps iii. If the time estimate for the machine found in Step ii is less than the time estimate for the machines found in Steps iii and iv, then use the single machine in Step ii for the fine grid. Else use the machines found in Steps iii-iv for the fine grid. If Step ii was not able to find a single machine that is able

to solve the problem, then the machines found in Steps iii and iv are used for the fine grid.

At the end of performance modeling, the fine grid, which consists of a subset of machines that can solve the problem in the fastest possible time, given the grid parameters, is committed to the run.

The performance model, in the current case of solving a system of linear equations on a heterogeneous computational grid, relies specifically on modeling the time to solution for the PDGESV kernel. This is the routine that is responsible for solving the linear system. A full description of how PDGESV solves the system may be found in the ScaLAPACK Users' Guide. It should be recalled that the process is dominated by a LU factorization of the coefficient matrix A. This is to say that solving the equations is a Level 3 BLAS process. As such, the time to do this factorization is arguably the major time constraint in the time to solution for solving a system of linear equations, in particular, as the problem size grows larger.

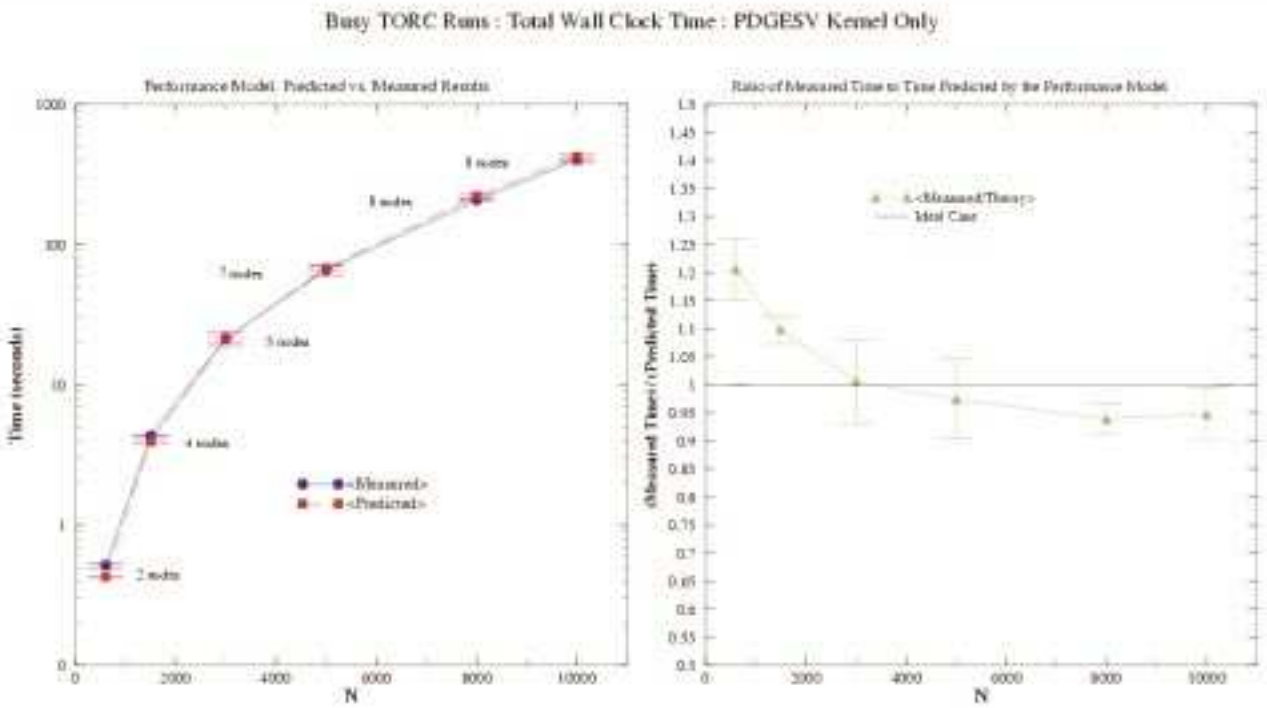

Figure 1: Performance Model vs. Measured Performance

Figure 1 provides a check on how precisely the current performance model is making predictions. The plots are a statistical study of grid-based runs on the TORC cluster at the University of Tennessee. The runs represent a homogeneous, non-dedicated cluster case.

This is the simplest realistic case possible since the communication lines, available memory, and CPU architectures are the same for each computing node. The way the model makes decisions is based upon the grid conditions returned at the time of request as described above. More specifically, in the current model the problem is distributed over the fine grid in a one-dimensional block cyclic fashion with N/NB sized panels (which are determined from the size of A, N, and the chosen block size, NB). In the LU factorization there are three major steps; a factorization phase, a broadcast phase, and an update phase. The current model predicts times for each of these phases as would be reflected on the root-computing node of the fine grid. The dominant phase is the update, which is a call to the Level 3 PBLAS [2] routine PDGEMM.

The first plot of Figure 1 is the total wall clock time measured for doing the linear solve for problem sizes ranging from N = 600 to N = 10,000. The corresponding predicted values are also shown. Five runs were made for each data point on this linear-log plot. The second plot gives a precise comparison of the ratios of these numbers. Naturally, in plot two, we see the outlying data points for the smaller problem sizes. If one disregards those points, the current model is better than 95% accurate on average for this simplistic, homogeneous cluster study.

4. Contract developer

Currently, the contract developer commits all of the machines that were returned by the performance modeler for the application. In the future, the contract developer will be more robust and be able to build contracts for the given user problem and the grid parameters.

5. Application Launcher

The application launcher spawns the parallel application over the machines in the fine grid, launches the contract monitor, and starts the application sensors. These sensors register themselves with the Autopilot manager and periodically send instrumentation data to the Autopilot manager. The contract monitor, through the Autopilot manager, contacts the sensors, obtains the data from the sensors, determines if the application behaves as predicted by the application model, and prints the output. In the future, the contract monitor will send its output to a scheduler, which may then reschedule, perhaps migrate, the application if the contracts are violated. As soon as the parallel application has been spawned, the input matrices are randomly generated (or read from a URL), and block-cyclically distributed among the processors in the fine grid. After the input matrices are distributed, the ScaLAPACK routine PDGESV is invoked to solve the system of linear equations, the validity of the solution is checked, and the results are returned to the user's application.

Experiments

Two sets of experiments were conducted. The first set of experiments compares the Grid version of ScaLAPACK (using MPICH-G) and the native ScaLAPACK (using MPICH-P4) as implemented on a cluster. These measurements give an estimate of the overhead cost associated with performing numerical computations using the GrADS system. The second set of experiments illustrate the functionality of GrADS by running the application on a dynamically chosen number of processors that exist and are available on the Grid.

In viewing the following graphs, it is important to note that the GrADS strategy for building the "fine grid" tries to optimize the time to solution. This model does not optimize for best efficiency. However, the model could be changed to alternatively optimize for best efficiency. There are many possible grid resources -- dedicated, shared, etc -- and the needs of the user changes depending upon the types of resources available to him. Similarly, his definition of "best performance" could be measured in terms of Mflop/s to most effectively utilize the CPUs in his machine, instead of minimum time to solution. Keeping these issues in mind, in Figure 2 the number of processors selected by the GrADS strategy for the "fine grid" is close to the best choice for "Raw ScaLAPACK" with respect to the "time to solution" criterion, and the one-dimensional mapping of processors constraint.

Comparison Of Grid And Raw ScaLAPACK

In the following experiments, ScaLAPACK runs using the GrADS system were compared with the native ScaLAPACK runs without the GrADS system on a local Linux cluster (TORC, http://icl.cs.utk.edu/projects/torc/) at the University of Tennessee (UT). Each machine is a dual processor 550 MHz Pentium III running Linux, and only one process was spawned per node. The comparison is between the (ScaLAPACK+MPICH-G) [8] performance over the Grid and the (ScaLAPACK+MPICH-P4) [9] [10] performance without the Grid. Although the goal of the GrADS system is to solve large problems across multiple wide-area clusters, the following experiments reveal the costs of Grid-related overhead in the most ideal setting, a local cluster. The results yield an approximate 30% overhead in running MPICH-G versus MPICH-P4. In a best-case scenario, this is a lower bound on the overhead to be incurred when running across geographically-distributed clusters. Experiments were run using the TORC cluster in dedicated and non-dedicated mode.

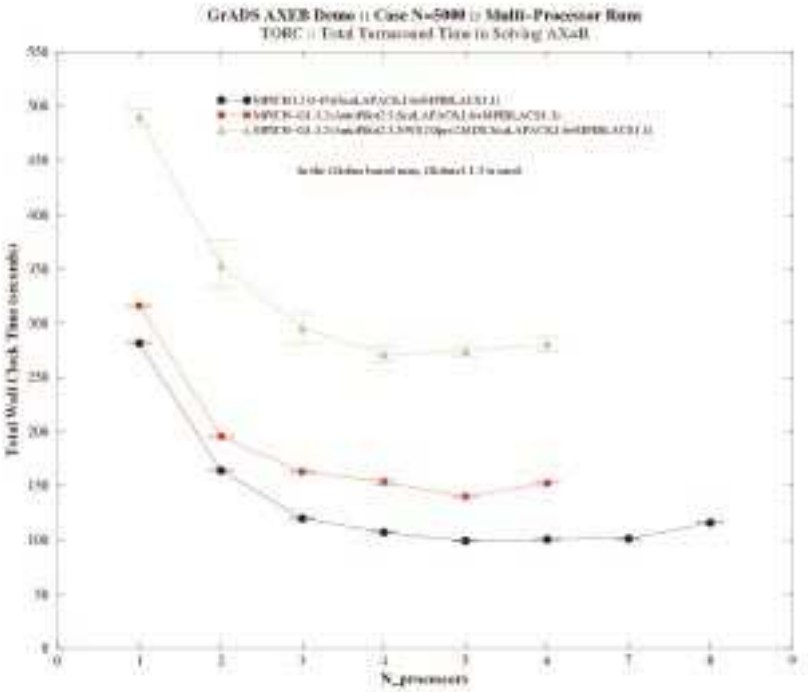

Figure 2: Ax = b, n = 5000, Multiprocessor Runs

Figure 2 depicts a sample case from the non-dedicated runs on TORC. It is a statistical set of measurements (ten runs for each point on the graph) for the total turnaround time involved in solving a linear system of equations with 5000 unknowns. In other words the time interval from the time the user submits a request to the time the system has been solved and the application has been cleanly removed from the computing resource is compared.

As mentioned, the GrADS linear system solver is built with ScaLAPACK embedded in grid-based software. Thus, it is important to understand what overhead is introduced when going from the raw ScaLAPACK runs to the full grid-based application. A fair question to ask for the grid-based runs for a given problem size is: "How can one guarantee that the fine grid will consist of exactly X processors, where X is varying for a set problem size?" Actually, if you allow the entire set of grid-based computing resources to be considered when solving the N = 5000 problem, the answer is you cannot. However, one can constrain the study to the TORC cluster and to a preset number of processors through a well-defined configuration file used at runtime. Further, although restriction to a cluster is easy to impose, the number of nodes requested versus the number that is actually selected for the fine grid cannot be imposed with certainty. This is due to the performance model. The point is, once the maximum number of compute nodes requested becomes larger than the optimal number predicted by the model, the request is ignored and the model chooses the number it thinks would best solve the problem in a timely fashion. Naturally, since the grid resources are dynamic, this optimal number of machines in the fine grid will vary. Thus, the number of nodes for the grid-based runs is seen to terminate at six processors in the plot.

The plot in Figure 2 compares three sets of measurements:

1) The MPICH-P4 1.2.0 + ScaLAPACK1.6 + MPIBLACS1.1 numbers -which are referred to as the "Raw ScaLAPACK" runs,
2) The MPICH-G1.1.2 (Globus-based MPI) + AutoPilot2.3 + ScaLAPACK1.6 + MPIBLACS1.1 runs, and
3) The MPICH-G1.1.2 (Globus-based MPI) + NWS2.0pre2 + MDS + AutoPilot2.3 + ScaLAPACK1.6 + MPIBLACS1.1 runs - the actual grid runs.

In this example, the raw ScaLAPACK had a minimal runtime, on average, when run on five processors (as did the type (2) runs). The grid-based runs were solved the fastest on four of the TORC computing nodes.

Clearly, going from the type 1) to the type 3) runs incurs additional overhead in time associated with the gathering of information for the Grid run. A detailed breakdown of where this extra time comes from in the grid application is provided in Figure 3.

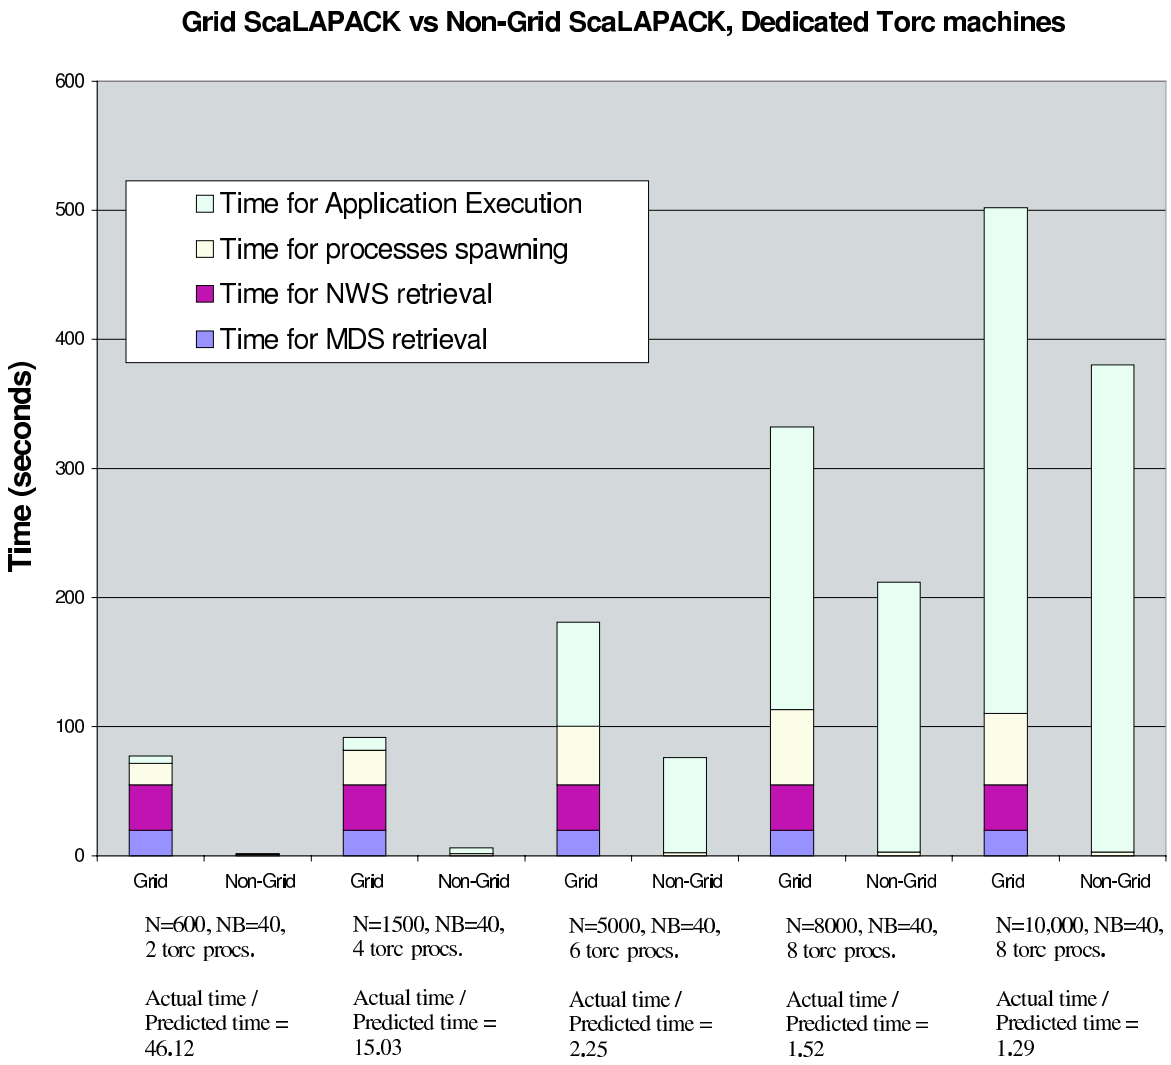

Figure 3: Overhead in Grid Runs

With the exception of process spawning, the Grid overhead remains more or less constant as the size of the problem and the number of processors chosen increase. The time to spawn processes in the Grid is noticeably more expensive than in the "non-Grid" case, and increases with the number of processes. As the complexity of the problem being solved is $O(n^3)$, for very large problems this overhead will become negligible.

Experiment 2: ScaLAPACK Across Clusters

In this experiment, two separate clusters from UT, named TORC and CYPHER (http://www.cs.utk.edu/sinrg/), and a cluster from the University of Illinois Urbana-Champaign (UIUC) called OPUS were used. For this experiment a total of four TORC machines, six CYPHER machines and eight OPUS machines were available. The following table shows some of the system parameters of the machines.

	TORC	**CYPHER**	**OPUS**
Type	Dual Pentium III	Dual Pentium III	Pentium II
OS	Red Hat Linux 2.2.15 SMP	Debian Linux 2.2.17 SMP	Red Hat Linux 2.2.16
Memory	512 MB	512 MB	128 or 256 MB
CPU speed	550 MHz	500 MHz	265 – 448 MHz
Network	Fast Ethernet (100 Mbit/s) (3Com 3C905B) and switch (BayStack 350T) with 16 ports	Gigabit Ethernet (SK-9843) and switch (Foundry FastIron II) with 24 ports	IP over Myrinet (LANai 4.3) + Fast Ethernet (3Com 3C905B) and switch (M2M-SW16 & Cisco Catalyst 2924 XL) with 16 ports each

The following graph shows the total time taken for GrADS for each experiment as the matrix size is increased.

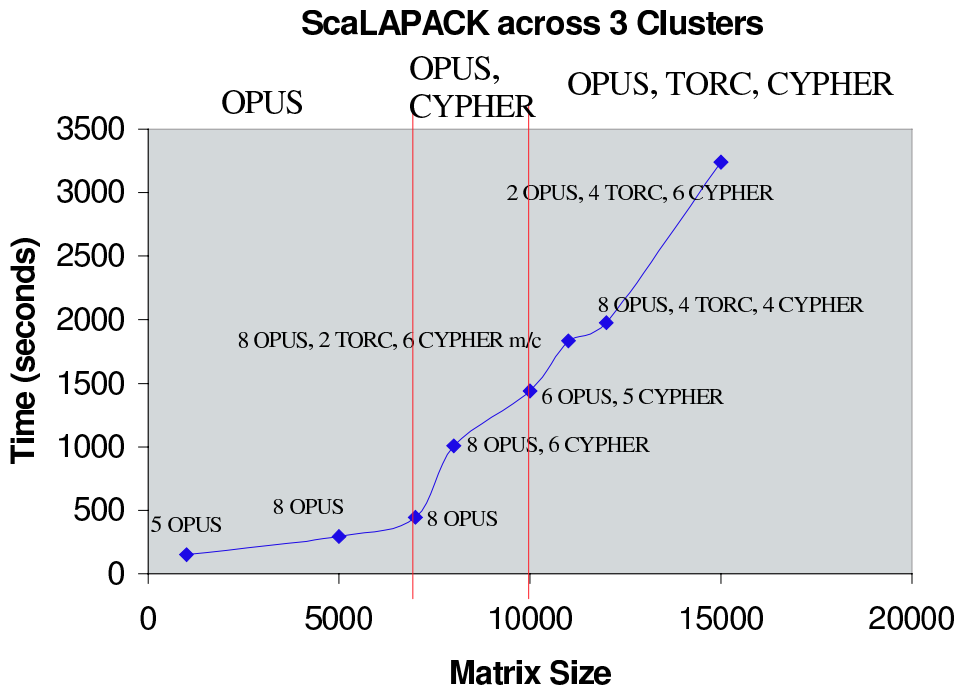

In the above experiments, GrADS chose only OPUS machines at UIUC for matrix sizes up to 8000. In fact, for this problem size, the system can be solved on one cluster. The OPUS and the CYPHER clusters are comparable in terms of their network parameters. At the time the experiments were conducted, the OPUS cluster was found to have better network bandwidth than the CYPHER cluster due to the network load on CYPHER. Hence, the OPUS cluster was the best choice among the pool of resources. For a matrix size of 8000, the amount of memory needed per node is of the order of 512 MB. Since none of the UIUC machines has this much memory, GrADS utilized both the UIUC and CYPHER machines at UT for matrix sizes greater than 8000. The GrADS framework gave preference to the CYPHER machines over the TORC machines because of the superior network speed of the CYPHER network. Hence, we find a steep increase in execution time from matrix size 7000 to 8000.

We also find that the number of machines chosen for a matrix size of 10,000 is smaller than the number of machines chosen for matrix size 8000. This is because certain UIUC machines have small memory size and these machines were found to be not optimal by the GrADS system for matrix size of 10,000. Also, since the experiment was conducted on non-dedicated systems, the GrADS scheduling system did not choose some machines from the collection when the system and network loads corresponding to those machines significantly increased.

For matrix sizes greater than 10,000, machines from all three of the clusters were chosen. We find that the transition from 10,000 to 11,000 is not as steep as the transition from

7000 to 8000. This is because the transition from 7000 to 8000 involved Internet connections between UIUC and UT machines and the transition from 10,000 to 11,000 involved UT campus interconnections between the CYPHER and the TORC machines. As can be seen from these experiments, the GrADS infrastructure is making intelligent decisions based on the application and the dynamics of the system parameters.

We also ran the experiment utilizing all the machines in the system including six TORC machines, twelve CYPHER machines and eleven UIUC machines. Due to memory limitations (and CPU load on the machines), the maximum problem size that was solvable by the collection of machines was a matrix of size 30,000. In this case, GrADS chose seventeen processors to solve the problem. These seventeen machines consisted of eight TORC machines and nine CYPHER machines. The total time taken for the problem was 81.7 minutes. Of this, 55 seconds for retrieving information from MDS and NWS and the remaining time was taken for launching and executing the application. Thus, for the problem size of 30,000, GrADS was able to achieve 213.4 Mflop/s. The theoretical peak performance achievable is 500 Mflop/s on CYPHER and 550 Mflop/s on TORC. Thus GrADS was able to achieve 42.6% of the peak performance, while the raw SCALAPACK can achieve about 50% of the peak performance. Thus, we find that the performance of GrADS over the Grid is not far from the performance of the native numerical application over a local cluster. The GrADS system as configured was not able to solve problem sizes larger than 30,000 due to the memory limitations (and CPU load) of the available machines.

The timings for "Grid ScaLAPACK" need to more precisely reflect the "total overhead" cost for performing ScaLAPACK on the Grid. "Grid ScaLAPACK" timings were performed with the NWS clique leader not included, as one of the computational nodes, so the complete "overhead" associated with NWS is not totally reflected in the timings. Also, the TORC nodes are dual processors, and the timings were performed on one processor of each node, but the kernel was configured as a dual processor. As the communication is done over IP, it is not known the amount of "overhead communication" cost that is not being captured by the timings, as this cost is being offset on the second processor per node. For future work, timings will be performed with the TORC nodes configured as uniprocessors and the clique leader included as one of the compute nodes, to more accurately reflect the "total overhead" associated with performing ScaLAPACK over the Grid.

The timings reported for "Grid ScaLAPACK" and "Raw ScaLAPACK" include the time to spawn the MPI processes, generate the matrices, solve the system, and perform an error check to validate the computed solution. The cost of the error check is negligible. In terms of performance, on the TORC cluster, for example, each machine's theoretical peak performance is 550 Mflop/s (each processor is 550 MHz and one flop per cycle), and ATLAS's [11] DGEMM achieves 400 Mflop/s, 73% of the theoretical peak performance. As a general rule of thumb, when optimizing for best efficiency per processor, ScaLAPACK achieves 75% of DGEMM performance, approximately 300 Mflop/s per processor. This measure greatly depends upon the network and the mapping of the processors (one-dimensional versus two-dimensional). Thus, in its best case, "Raw

ScaLAPACK" performs at 55% of the theoretical peak performance of the machine. As can be seen from the "Grid ScaLAPACK" timings, and the N=30,000 case, the code achieves approximately 210 Mflop/s per processor, which is 40% of the theoretical peak performance. This performance is quite good considering the fact that we have a heterogeneous group of machines connected across the Internet, most of which are slower than 550MHz.

Conclusions

The set of experiments that are reported here was more challenging than originally anticipated. Part of this stems from the fact that we had to coordinate a number of machines across different administrative domains and also in part because of the varying degrees of maturity in the software and the sheer amount of software involved in getting the experiments in place and maintaining a workable configuration over a long period of time. Hopefully this situation will improve as the software matures, more sites engage in grid based computing, and the software infrastructure is more widely used.

Part of the point of conducting these experiments was to show that using geographically distributed resources under a grid framework through the control of the library routine can lead to an improved time to solution for a user. As such the results, for this modest number of experiments, show that performing a Grid-based computation can be a reasonable undertaking. In the case of solving dense matrix problems we have the situation where there are $O(n^2)$ data to move and $O(n^3)$ operations to perform. So the fact that we are dealing with geographically distributed systems is not a major factor in performance when the data has to be moved across slow networks. If the problem characteristics were different the situation may not be the same in terms of Grid-feasibility.

Future work will involve the development of a system that implements a migration system if the time to solution violates the performance contract and a mechanism to deal with fault tolerance.

Reference

1. Berman, F., et al., *The GrADS Project: Software Support for High-Level Grid Application Development*. 2000, Rice University: Houston, Texas.
2. Blackford, L.S., et al., *ScaLAPACK Users' Guide*. 1997, Philadelphia, PA: Society for Industrial and Applied Mathematics.
3. Snir, M., et al., *MPI: The Complete Reference, Volume 1, The MPI Core, Second edition*. 1998, Boston: MIT Press.
4. Balay, S., et al., *PETSc 2.0 Users' Manual*. 1996, Argonne National Laboratory: Argonne, IL.
5. Foster, I. and C. Kesselman, *Globus: A Metacomputing Infrastructure Toolkit*. Intl J. Supercomputer Applications, 1997. **11**(2): p. 115-128.
6. Wolski, R., N. Spring, and J. Hayes, *The Network Weather Service: A Distributed Resource Performance Forecasting Service for Metacomputing*. Future Generation Computer Systems (to appear), 1999.
7. Ribler, R.L., et al. *Autopilot: Adaptive Control of Distributed Applications*. in *Proceedings of the 7th IEEE Symposium on High-Performance Distributed Computing*. 1998. Chicago, IL.
8. Foster, I. and N. Karonis. *A Grid-Enabled MPI: Message Passing in Heterogeneous Distributed Computing Systems*. in *Proc. SuperComputing 98 (SC98)*. 1998. Orlando, FL: IEEE.
9. Gropp, W., et al., *A High Performance, PortableImplementation of the MPI Message Passing Interface Standard*. Parallel Computing, 1996. **22**(6): p. 789-828.
10. Gropp, W. and W. Lusk, *Users' Guide for MPICH, A Portable Implementation of MPI*. 1996, Mathematics and Computer Science Division, Argonne National Laboratory.
11. Whaley, R., A. Petitet, and J. Dongarra, *Automated Empirical Optimization of Software and the ATLAS Project*. Parallel Computing, 2001. **27**(1-2): p. 3-25.

Appendix

Grads Numerical Library Interface:

The following example is for a user running an application on a sequential machine that is connected to the network. The user will make a grid-enabled library call and the computation will be done on a set of processors provided by the system. In this example the user wants to solve a system of linear equations using Gaussian elimination with partial pivoting. The framework provided here can be expanded to include other mathematical software.

We assume the user has already communicated with the system (Globus) and has been authenticated (we assume using grid-proxy-init).

The user must include "grads.h" in his program and invokes Grads_Lib_Linear_Solve() as follows:

ierr = Grads_Lib_Linear_Solve(USER_ARGS);

The USER_ARGS data structure contains all the parameters passed by the user including the matrix size, the block size, the list of machines on which the user wants to run his application, the machine which is running the Autopilot manager etc..

Below is the pseudo-code for Grads_Lib_Linear_Solve, as well as the APIs for the GrADS components detailed in this paper.

```
int Grads_Lib_Linear_Solve(  demo_args_T *USER_ARGS  )
{
        Grads_Resource_Coarse_Grid_Handle_T     coarse_grid;
        Grads_Lib_Fine_Grid_Handle_T            fine_grid;
        Grads_Lib_Problem_Handle_T              problem;
        int                                     n, match;
                                /* Create a problem of  type "Linear_Solve" */
        Grads_Lib_Problem_Create( Linear_Solve, &problem );
                                /* Set problem attributes */
        Grads_Lib_Problem_Set_Attr( Problem_Matrix_Size, &Matrix_Size, problem );
        Grads_Lib_Problem_Set_Attr( Problem_Block_Size, &Block_Size, problem );

        do
        {
                                        /* For a given problem, retrieve a grid */
                ierr = Grads_Resource_Selector( problem, &coarse_grid );
                                        /* Extract sub-grid to work with */
                Grads_Lib_Performance_Modeler( problem, coarse_grid, USER_ARGS,
                                        &fine_grid );
                /* we are done with coarse_grid; free memory  coarse_grid points to. */
```

```
            Grads_Resource_Coarse_Grid_Remove( problem, coarse_grid );
                                    /* Try to commit the fine grid for the problem */
            match = Grads_Contract_Developer( problem, fine_grid );
                                    /* If this list of machines is not good – release it */
            if( match != Grads_SUCCESS ) Grads_Lib_Grid_Free( problem, fine_grid );

        } while( match != Grads_SUCCESS );

        ierr = Grads_Application_Launcher( problem, fine_grid, USER_ARGS );

        /* Release the Resources */
        Grads_Lib_Fine_Grid_Remove( fine_grid );
        Grads_Lib_Problem_Remove ( problem );
        return( ierr );
}
```

The APIs for the GrADS components detailed in this paper are listed below. For complete information, please refer to the "grads.h" include file.

1. Resource Selector

Grads_Resource_Selector(problem, coarse_grid)

IN	problem	the problem handle
OUT	coarse_grid	handle to a structure with the following information
	int no_coarse_grid	number of processors potentially available
	int array name(no_coarse_grid)	names of the available processors (perhaps ip addresses)
	int array memory(no_coarse_grid)	amount of memory available on each of the processors
	int array communication(no_coarse_grid^2)	a 2-d array containing bandwidth and latency information on the link between available processors
	int array speed(no_coarse_grid)	peak speed for each processor according to some metric.
	int array load(no_coarse_grid)	load on each processor at the time the call was made.
OUT	ierr	error flag from the Resource_Selector

```
int Grads_Resource_Selector( Grads_Lib_Problem_Handle_T problem,
                             Grads_Resource_Coarse_Grid_Handle_T *coarse_grid )
```

2. Performance Modeler

Grads_Lib_Performance_Modeler(problem, coarse_grid, USER_ARGS, fine_grid)

IN	problem	a unique problem identifier for this library call
IN	coarse_grid	struct (see above call)
IN	USER_ARGS	struct
OUT	fine_grid	handle to a structure specifying the machine configuration to use
OUT	ierr	error code returned by the function

```
int Grads_Lib_Performance_Modeler(  Grads_Lib_Problem_Handle_T problem,
                                    Grads_Resource_Coarse_Grid_Handle_T coarse_grid,
                                    demo_args_T *USER_ARGS,
                                    Grads_Lib_Fine_Grid_Handle_T *fine_grid )
```

3. Contract developer

match = Grads_Contract_Developer(problem, fine_grid)

IN	problem	a unique problem identifier for this library call
OUT	fine_grid	handle to machine configuration
OUT	match	will be 0 if the processors are available for this run.

```
int Grads_Contract_Developer( Grads_Lib_Problem_Handle_T problem,
                              Grads_Lib_Fine_Grid_Handle_T fine grid )
```

4. Application Launcher

ierr = Grads_Application_Launcher(problem, fine_grid, USER_ARGS)

IN	problem	a unique problem identifier for this library call
IN	fine_grid	handle to machine configuration
IN	USER_ARGS	user arguments
OUT	ierr	error flag from the Application_Launcher

```
int Grads_Application_Launcher( Grads_Lib_Problem_Handle_T problem,
                                Grads_Lib_Fine_Grid_Handle_T fine_grid,
                                demo_args_T *USER_ARGS )
```