DEVELOPING AN ARCHITECTURE TO SUPPORT THE IMPLEMENTATION AND DEVELOPMENT OF SCIENTIFIC COMPUTING APPLICATIONS

Dorian C. Arnold, Jack J. Dongarra

University of Tennessee

Knoxville, TN, USA

Abstract The increasing complexity and computational demands of scientific computing applications has led to a need for robust software architectures to facilitate the conceptualization, design, implementation, deployment and maintenance of these applications. This paper attempts to shed light on how the unique characteristics of scientific computing applications, as well as computational scientists, make impositions upon the framework used to support scientific research efforts. We use our experience with NetSolve, a toolkit designed just for such interactions, as a means to present the approach of one infrastructure to support scientific computing. We further discuss how NetSolve implements the unique model of using a single system to aggregate, manage and access distributed hardware *and* software resources.

Keywords: scientific computing, distributed systems, grid computing, heterogeneous network computing, NetSolve

1. INTRODUCTION

Today, scientists and engineers have become completely reliant on the computer as a tool for advanced modeling, simulation experimental analyses. Scientific simulations ranging from that of the electromagnetic field of a diffractor in a virtual human to warfare scenarios with tens of thousands of components interacting together are now possible thanks to significant advancements in computer technology. While the rapid rate of microprocessor performance growth has been constant, decades ago scientists realized that an intuitive way to increase the computational capacity of any single processor/machine is to use software systems that leverage computer networking infrastructures to connect multiple pro-

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35407-1_22](https://doi.org/10.1007/978-0-387-35407-1_22)

R. F. Boisvert et al. (eds.), *The Architecture of Scientific Software*

© IFIP International Federation for Information Processing 2001

cessors/machines together. Indeed, the scientific computing application (SCA) that runs solely on a single processor has become a rare breed. Supercomputers, with massively parallel processors (MPP) or shared memory multi-processors (SMP), are very common and today's fastest computers operate at speeds of over 2 TFlop/s. These speeds have been achieved by improved chip technology and computing architectures that support parallel processing. As computer scientists and engineers explore ways to increase computational capacities and capabilities, computational scientists now face highly complex architectures with each variation itself requiring a variation of the parallel programming model needed for code accuracy and performance. Furthermore, since every SCA demands interdisciplinary expertise (mathematics, computer science, and the domain science at the very least), there must be simplistic ways for the computational scientist to leverage the effort of algorithm developers, software system designers and hardware engineers and merge them with his own expertise.

The emergence of technologies like PVM [1] and standards like MPI [2] show strong efforts to develop a common parallel programming environment. Even more recently, there has been the emergence of Grid computing concepts [3] that envision software technologies exploiting today's high network connectivity to create a single, global virtual machine. However, little attention has been focused on the computational scientists and what is required for them to accomplish useful tasks without the colossal nightmare of becoming a computer scientist and a mathematician (and maybe even a magician) overnight. The goal of this article is to analyze the unique needs of SCAs and the scientists that develop them in an attempt to establish the issues that supporting software infrastructures should resolve. In the following, we also discuss the NetSolve distributed computing environment which attempts to address these issues in a practical and efficient way. Section 2 defines what we have found to be the fundamental characteristics of scientific computing that should help to mold the supporting scientific computing application infrastructure (SCAI). After a general overview of the NetSolve system in Sect. 3, Sect. 4 evaluates some of the features of NetSolve based on the gathered requirement of scientists and their applications. What we hope to present is not only an introduction to NetSolve, but a means by which to evaluate any infrastructure that claims to support scientific computing based on the needs of that community.

2. CHARACTERISTICS OF SCIENTIFIC COMPUTING

Scientific applications typically have four phases as depicted in Fig. 1. Regardless of whether the application is Graphical User Interface (GUI) based or text based, the first phase entails the gathering and preparation of input data. This may mean getting user input, allocating memory requirements, constructing specialized data structures and the like. After the input is made ready, it is passed to the computationally intensive phase of the process when complex algorithms are run on the prepared data. After the data processing, often there is an analysis phase which may be used to determine if further processing of the data is necessary, among other things. The data processing and analyses phases may undergo several iterations until finally some output results are made available. This may be in the form of textual output or graphical images that visualize the simulation and/or its results.

Figure 1 The four typical phases of a scientific computing application

So far, we have made no mention of possible parallelizations that may take place in the program. Typically, the only difference between parallel applications and serial ones is that in a parallel application, the data processing phase is distributed amongst multiple processors. From this point on, we assume the more common scenario of a parallel application. At this point, we distinguish two general categories of par-

allel applications; this categorization depends on whether or not there is communications amongst the computational nodes during the computation. Figure 2 depicts what we refer to as cooperative parallelism where there is interprocess communication amongst the computational nodes. Figure 3 shows the other kind of parallelism where there are concurrently running modules on multiple processors, but no communications amongst the nodes; we term this independently parallel execution.

Figure 2 The parallel scientific computing application with cooperative parallelism

Figure 3 The parallel scientific computing application with independent parallelism

The following characteristics of SCAs have been collected from a variety of sources, including [4] in an attempt to provide a complete picture of the demands of scientific computing. The format of the items that follow is a description of the characteristic followed by a single (*emphasized*) sentence that summarizes the impact that characteristic has on SCAs:

•Knowledge Base of the Scientist

Whether it be a computational chemistry or nuclear engineering application, the computational scientist will have a deep knowledge and understanding of the concepts of their scientific domain. They may also possess a thorough understanding of relevant mathematical concepts. However, their expertise in programming and other aspects of software engineering is typically limited. In larger collaborations, the application development is often carried out by a concert of domain scientists, computer scientists and, perhaps, mathematicians. Not every organization and research effort can afford such luxuries, nor do they wish to.

The SCAI must provide an easy and intuitive programming model that allows computational scientist to integrate complex supporting software

and highly-optimized mathematical modules without advanced levels of expertise.

•Composition of Existing Software Components

To help alleviate the previous problem, highly optimized numerical software libraries are often made available. Numerical solvers (linear/non-linear systems solvers, differential equation solvers, etc.) are the fundamental subprograms found in practically all SCAs; packages like LAPACK [5] and PETSc [6] are only two examples of libraries that implement some of these sophisticated techniques. As is the case with these examples, a wide array of this software is freely available, and even open-sourced. For obvious reasons, the domain scientist will want to leverage the years, if not decades, of research and development incorporated in the older packages or the novel concepts and capabilities of the newer ones.

The SCAI must provide ways to easily incorporate largely varied types of computational software modules.

•Granularity of Computational Modules

The granularity of the computational modules can be defined as the relation between the floating point performance of the module and the average communication bandwidth of a single processor (Flop/Byte) [7]. A higher granularity means that the computational modules take on large, atomic chunks of work at a time, while a lower granularity implies that not much processing takes place before data is transported. The granularity of the computational modules can depend upon the class to which that module belongs. Lower level classifications like a mathematical linear system solver (or even the matrix multiply routine upon which the solver depends) leads to low granularities. Higher level classifications like a data transformation module that incorporates linear systems solvers, or an image processing module that uses a series of data transformations yield higher granularity. Depending on the application, it may be warranted to use a large granularity, a small one, or both.

Software module granularity should bear little or no impact on the performance or ease-of-integration of the SCAI.

•Too Many Choices

As can be seen, there are many software packages available from a variety of sources. While this allows for much flexibility when deciding which packages to use, it also implies climbing a steep learning curve to determine the peculiarities of each package to make a fair evaluation of which suits your purposes. The veritable alphabet soup of packages available

can be daunting, yet discovering which package to use is the least of the troubles. As is the case of iterative methods, an algorithm is only as good as the manner in which you use it. Veterans of iterative methods agree that finding the right combination of solver, pre-conditioner, scaling and re-ordering is an art form developed only from experimentation within the application area [8].

The SCAI must hide the complexities and intricacies of the underlying computational modules while providing convenient ways for the scientist to discover and use the services he wants.

•Problem Capacity

Often, the memory or processing demands of an SCA exceed the capacity of workstations and requires multiple processors to accurately and efficiently solve the problem specifications. In other cases, it may be that the only appropriate algorithms discovered by the computational scientist involve code that only run on a particular server. Or, combining these two scenarios, the only feasible solvers are specialized parallel algorithms on a remote distributed memory MPP system [8].

The SCAI must be able to reliably and efficiently execute modules on remote servers.

•Interaction Levels

Although there definitely is a place for GUIs in rapid prototyping and quick experimentation with SCAs, generally, these high-performance computing applications can take hours, days and weeks to complete a single task. The user may not want to keep his GUI running and definitely does not want to be required to interact for that duration. And while GUIs at times provide a more convenient programming model, the day has not yet arrived when a graphical programming environment can implement a doubly nested for loop executing a module hundreds of times more conveniently than a scripting language.

Scripting capabilities should be available to allow complex chaining combinations of modules with no sacrifice of convenience and ease-of-use.

•Code Maintenance and Enhancement

The fact that the FORTRAN programming language is still popular after nearly 50 years implies two fundamental things: i) the laws of inertia apply to computer science as well (i.e. scientists and programmers are reluctant to diverge from a platform that can work for their purposes after tremendous time and efforts have been invested) and ii) it is extremely important for codes to be backward compatible and have a lifetime as long as possible. As a result, mixed-language programming

is very common in SCAs – hybrid applications now integrate C, FORTRAN and object-oriented languages like Java and C++ to exploit the different advantages of each platform be it performance, maintenance, security or programming methodologies (or, simply, code availability).

The SCAI must have multiple-language support and provide mechanisms which allow for easy replacement of components with newer, optimized, or simply corrected, versions without significant code modification.

•Performance

Last, but by no means least, SCAs demand optimal performance. It is the very nature of these high-performance applications to consume and dominate any and all resources they can access and still require days and weeks of computational time to complete. Data sets can easily extend into the hundreds of megabytes and gigabyte range for a single application, and this too needs to be considered, especially when data transfer to remote servers are involved.

The design of the SCAI must be such that it readily improves the performance of applications with minimal overhead due to component interactions.

3. THE NETSOLVE SYSTEM

We now briefly discuss the architecture and major components of the NetSolve system to provide background for the discussion to follow. NetSolve is being developed at the University of Tennessee's Innovative Computing Laboratory. Its original motivation was to alleviate the difficulties that domain scientists encounter when trying to locate/install/use numerical software, especially on multiple platforms. Today, the system has evolved into much more than a way to access numerical solver routines. NetSolve provides an environment that monitors and manages computational resources, allocating the services they provide to NetSolve-enabled client programs. Built upon standard Internet protocols, like TCP/IP sockets, it runs on popular variants of the UNIX operating system, and parts of the system are available for the Microsoft Windows platforms. Further documentation of NetSolve along with source code for the full product are available at icl.cs.utk.edu/netsolve. NetSolve has been recognized as a significant effort in research and development, and was named in R&D Magazine's top 100 list for 1999.

Figure 4 shows the infrastructure of the NetSolve system and its relation to the applications that use it. NetSolve and systems like it are often referred to as Grid middleware; this figure helps to show why. The shaded parts of the figure represent the NetSolve system. It can be seen

that NetSolve acts as glue layer that brings the application or user together with the hardware and/or software it needs to complete useful tasks.

Figure 4 Architectural Overview of the NetSolve System

At the top tier, the NetSolve client library is linked in with the user's application. Through NetSolve's application programming interface (API), client-users gain access to aggregate resources without needing knowledge of computer networking or distributed computing. In fact, the user does not even have to know remote resources are involved. NetSolve currently supports the C, FORTRAN, and Matlab programming languages as interfaces of implementation for client programs.

The NetSolve agent represents the gateway to the NetSolve system. It maintains a database of NetSolve servers along with their capabilities (hardware performance and allocated software) and dynamic usage statistics. It uses this information to allocate server resources for client requests. The agent uses this information to allocate server resources for client requests, finding the server that will service the request the quickest, while keeping a balanced load amongst its servers and keep track of failed ones. Requests are directed away from failed servers.

The NetSolve server is the computational backbone of the system. It is a daemon process that awaits client requests. The server can run on single workstations, clusters of workstations, symmetric multi-processors or machines with massively parallel processors. A key component of the NetSolve server is a source code generator which parses a NetSolve problem description file (PDF). This PDF contains information that allows the NetSolve system to create new service modules thereby enhancing the server's capabilities.

4. NETSOLVE AND SCIENTIFIC COMPUTING

The NetSolve system, as mentioned in the previous section, was designed specifically to accommodate the computational scientist. This section compares the system with the requirements of the scientific computing community as discussed in Sect. 2 by placing NetSolve's features alongside the corollaries of that discussion.

• **Corollary 1:** The SCAI should provide an easy programming model and API for novice non-computer scientists.

The NetSolve system allows for client users to embed functions from practically any software library into their applications without having to install, learn or maintain that package. Fig. 5 shows an example Matlab code, before and after the NetSolve API has been integrated. The code on the left hand-side is making a call to a Matlab native function to multiply two matrices A and B and store the result in C. The call to NetSolve, on the right hand-side of the figure, achieves the same result via the NetSolve framework. This example shows how NetSolve can be used to provide access to complicated software modules without requiring expert interactions of the user. Apart from module encapsulation, it allows one to create uniform interfaces of different packages of similar algorithms. For instance, the intricacies of sparse solvers like PETSc [6], AZTEC [9] and others can be hidden by a single common interface that takes an additional parameter defining the software package to use.

...	...
A = load(input_matrix1);	A = load(input_matrix1);
B = load(input_matrix2);	B = load(input_matrix2);
C = matmul(A, B);	C = netsolve('matmul', A, B);
...	...

Figure 5 Matlab code: Left side before NetSolve, right side after NetSolve integration

• **Corollary 2:** The SCAI must provide ways to easily incorporate largely varied types of computational modules.

NetSolve provides a code generator that parses a NetSolve PDF in order to extend the servers' functional capabilities. Figure 6 shows a segment of a PDF that was used to integrate a module from a sub-surface fluid simulator. The **PROBLEM** parameter of this file defines the name we want client applications to use when referring to this module. The **INCLUDE** and **LIB** directives are used in the compilation of the module. Among other things, this PDF eventually describes the code that determines how to call the simulation code with the inputs given from a client program. The PDF facility provides a way for NetSolve to seamlessly integrate any type of computational modules. Furthermore, we have implemented a Java GUI that makes this process even easier.

```

@PROBLEM ipars
@INCLUDE "ipars.h"
@LIB /home/user/lib/libipars.a
@DESCRIPTION
Parallel Sub-Surface Flow Simulator
@INPUT 2
@OBJECT STRING CHAR model
IPARS physical model to use
@OBJECT FILE CHAR infile
Input data file
....

```

Figure 6 Sample PDF used to integrate new modules into the NetSolve system.

- **Corollary 3:** The system should adapt to varying granularities of computational modules.

While the PDF facility does not concern itself with the granularity of the computational module being integrated, an issue in any distributed environment is how the amounts and sizes of data transport affects performance. The first way NetSolve attempts to deal with this issue is by analyzing network bandwidths and latencies to choose the most conveniently located server resources to solve client requests. The second way we have optimized data communications is by creating an interface and infrastructure that allows a user to group or sequence a collection of NetSolve requests [10]. The system then analyzes the input and output parameters amongst all requests and caches common data near the relevant servers. Figure 7 illustrates the typical transactions that take place during a series of NetSolve requests by a single client. The important points to note are that parameter A is shared as an input for the first and second requests. Also, output parameters C and

D serve as inputs for subsequent requests. Figure 8 shows the reduction in data flow that occurs when the sequencing mode is employed.

Figure 7 Client-server interactions during a typical request scenario.

Figure 8 Client-server interactions during a "request sequence".

- **Corollary 4:** Complexities and intricacies of the underlying computational modules must be hidden, yet the modules must be conveniently accessible.

Through NetSolve, users are given access to complex algorithms that solve a variety of types of problems, one instance being linear systems solvers. All solvers, however, are not built alike; depending on the characteristics of the system being solved some perform poorly and others not at all. NetSolve has incorporated a large number of both direct and iterative solver algorithms for sparse/dense, symmetric/non-symmetric systems. To allow non-expert users to properly and efficiently use these algorithms without climbing the steep learning curve that would otherwise be involved, we have created an interface that allows them to generically call a "LinearSolve" routine which transparently analyzes the input matrix and determines which algorithm to use based on input characteristics. [11] further discusses this interface and the heuristics and decisions that are involved in the algorithm selection process. We envision using similar heuristics for classes of problems other than linear system solvers in the future.

- **Corollary 5:** The SCAI should reliably execute modules on remote servers.

By its very definition, NetSolve is a distributed computing environment that allows for remote problem execution. Its failure detection and fault-tolerant mechanisms allow the system to detect servers

that have failed to solve particular problems or server hosts that are non-responsive and direct new requests away from these resources. During a computation, the system attempts to use every appropriate and capable server host (from best to worse (see Corollary 8)) until a problem has been solved or the list of servers has been exhausted. Other investigations are leading to the development of heuristics to checkpoint NetSolve services so that a mid-service failure does not result in a lost of all previous computation. These checkpoints will be used to migrate the state of the interrupted service to other NetSolve servers where computation will resume.

- **Corollary 6:** Scripting capabilities must be available to allow users to chain complex combinations of modules.

All elements of the NetSolve system are accessible via the API of the client libraries. Using the functions of this API, users can embed calls to NetSolve in compiled languages, like C or FORTRAN, or interpreted languages, like Matlab and Mathematica. The nature of these environments make it possible for the user to invoke NetSolve in as simple or as complex a way as possible. The asynchronous interface further allows the user to make non-blocking request to NetSolve. The call returns immediately with a handle to the request that the user can use to probe to see if the request has completed and retrieve the results. This interface allows users to invoke multiple calls to NetSolve that would then run on different hosts, further improving application turnaround time.

To aid in the development of SCAs that do Monte Carlo simulations, parameter sweeps and other applications with simple task-parallel structures, i.e. independently parallel programming, we have created a task farming interface. The task farming interface allows the user to make a single call to netsolve, requesting multiple instances of the same problem. Rather than single parameters, the user passes arrays of parameters as input/output and designates how NetSolve should iterate across the arrays for the task farm. The main challenge in this effort is scheduling. Indeed, for long running farming applications it is to be expected that the availability and workload of resources within the server pool will change dynamically. [12] discusses the design of this infrastructure and also presents an adaptive scheduling algorithm used by the task farming interface to assign tasks to the server resources.

- **Corollary 7:** The SCAI must be have multiple-language support and its components should be completely pluggable.

As already mentioned, the NetSolve API has been implemented in many programming environments. Further motivated by the need to support various platforms (among other things), we have implemented client proxies to act on behalf of the Netsolve client. The proxy, a separate process that resides on the client host, handles (almost) all interactions with the underlying meta-computing resources. With a standard interface between the client and all proxies, it is possible, especially for third party developers, to easily add new language support to the NetSolve system. They would simply write libraries that interface the NetSolve proxies from their language of choice, allowing programs of that language to become NetSolve-enabled. Figure 9 depicts the main idea behind the proxy. The client libraries interact with the proxy thanks to a standard API and the proxy interacts with the meta-computing system using system-specific mechanisms. This allows NetSolve-enabled clients to leverage other meta-computing resources apart from those provided by NetSolve. The NetSolve proxy, for instance, uses the agent to discover services, contacts the appropriate server and establishes a session with that server who then receives input data from the client, executes his service and return output data.

Figure 9 Proxy Architecture

- **Corollary 8:** The design of the SCAI must optimize performance with minimal overhead to the resources it occupies.

The NetSolve agent uses both static and dynamic information from the servers to assign requests to the best server at that point in time. The algorithms for each service is configured with a complexity that describes the computational time of the algorithm based on input sizes.

Performance is measured by the LINPACK benchmark upon server initiation, and the server monitors its host reporting workload information. All these parameters, along with network bandwidth and latency information, are used whenever a request is received to rank the appropriate servers from best to worst. This list is sent to the client who then uses the fastest server to solve his problem in order to optimize performance.

The major drawback in distributed systems is often data transfers. Though not yet implemented, under consideration are heuristics to consider if performance might be improved by using the client's host as a server by uploading the necessary software once, rather than transferring data to a server (on possibly several occasions).

5. CONCLUSION

Scientific Computing Applications are a prominent part of the field of Computer Science. These applications and the computational scientists that implement them have unique characteristics that should forge the software infrastructures needed to support their development. We have discussed the relevant characteristics they possess and present NetSolve, an environment for remotely solving computational problems, as a system that hopes to address these issues. The system is a work-in-progress, and at the heart of this effort lies the philosophy that convenient interfaces and ease of administration are most important; every effort is made not to sacrifice these elements as the system evolves to meet the needs of the scientific computing community.

References

- [1] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, and V. Sunderam (1994). *PVM: Parallel Virtual Machine. A Users' Guide and Tutorial for Networked Parallel Computing*. The MIT Press, Cambridge, Massachusetts.
- [2] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra (1996). *MPI: The Complete Reference*. The MIT Press, Cambridge, Massachusetts.
- [3] I. Foster and C. Kesselman, eds. (1998). *The Grid, Blueprint for a New computing Infrastructure*. Morgan Kaufmann Publishers, Inc.
- [4] International Federation for Information Processing (2000). *Distinctive Characteristics of Scientific Applications*. http://www.nsc.liu.se/boein/ifip/woco8_dc.html.
- [5] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney,

- and D. Sorensen (1999). *LAPACK Users' Guide, third edition*. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- [6] S. Balay, W. D. Gropp, and B. F. Smith (1997). *Modern Software Tools in Scientific Computing*. Birkhauser Press, 163-202.
 - [7] B. Monien, R. Diekmann, R. Feldmann, R. Klasing, R. Luling, K. Menzel, T. Romke, and U. Schroeder (1998). Efficient Use of Parallel & Distributed Systems: From Theory to Practice in *Lecture Notes in Computer Science*. Springer-Verlag, Berlin.
 - [8] R. Bramley, D. Gannon, T. Stuckey, J. Villacis, J. Balsabrumanian, E. Ackman, F. Breg, S. Diwan, and M. Govindaraju. *Component Architectures for Distributed Scientific Problem Solving*.
 - [9] S. A Hutchinson, Shadid J. N., and Tuminaro R. S (1995). *Aztec user's guide: Version 1.1*. Technical Report SAND95-1559, Sandia National Laboratories.
 - [10] D. C Arnold, D. Bachmann, and J. Dongarra (2000). Request Sequencing: Optimizing Communication for the Grid, in *Euro-Par 2000 - Parallel Processing* (A. Bode, T. Ludwig, W. Karl, and R. Wismuller, eds.), Springer-Verlag, Berlin, 1213-1222.
 - [11] D. C. Arnold, S. Blackford, J. Dongarra, V. Eijkhout, and T. Xu (2000). Seamless Access to Adaptive Solver Algorithms, in *16th IMACS WORLD CONGRESS 2000 on Scientific Computation, Applied Mathematics and Simulation*, <http://imacs2000.epfl.ch/Abstracts/RechAbstSess.asp>, Session: 172, Paper: 6.
 - [12] H. Casanova, M. Kim, J. S. Plank, and J. Dongarra (1999). Adaptive Scheduling for Task Farming with Grid Middleware, in *The International Journal of Supercomputer Applications and High Performance Computing*.

DISCUSSION

Speaker: Jack Dongarra

Morven Gentleman : The description of NetSolve for naive users sounds attractive when it works. What have you done about feedback to naive users when something goes wrong? For example, if the computing grid is international as described, the error handling system must be internationalized in the classical sense so that the error information is presented in the local language of the reader.

Jack Dongarra : At this point NetSolve only returns to the user the error condition or flags raised by the library call.

Brian Smith : What effort have you made to support robustness with respect to hardware, network, and software failures to solve the requested problem?

Jack Dongarra : The NetSolve system ensures that a user request will be completed unless every single resource capable of servicing the request has failed. When a client sends a request to a NetSolve agent, it receives a sorted list of computational servers to try. When the client has succeeded in contacting one of these servers, the numerical computation starts. If the contacted server fails during the computation, then another server is contacted and the computation restarts. In the current implementation, all input data resides on the user's machine during the computation and can then travel to another server after a failure. Future implementations will hand off the input data to NetSolve storage servers where it will stay until the computation has finished (after possible failures and restarts). The whole restart process is transparent to the user. This fault-tolerance mechanism is rather primitive (even though effective), and we are investigating techniques to allow task-checkpointing among servers.

David W. Walker : This question relates to data persistence. In a sequence of NetSolve calls in which the output from one call is used as the input to a subsequent call, can the data be manipulated without sending it back to the client after every call?

Jack Dongarra : We have a mechanism called request sequencing that provides for data persistence on the servers. With request sequencing a user can group together requests for grid services through NetSolve to exploit some common characteristics of these requests and minimize network traffic. Request sequencing can be used to implement affective scheduling policies and enable more expedient resource allocation methods.

Ivor Philips : Different computers may have put up different versions/releases of the same software. How can I, as a user, know that I

have used the most current version? This seems to point to a serious grid management issue.

Jack Dongarra : With NetSolve there is an audit trail generated for a call so a user can see what routine/version/machine was used in the invocation. The system does not provide the ability to have the user select a specific version.

Anne Trefethen : How much help does NetSolve offer the user in choosing an algorithm? For example, the Matlab interface shown, i.e., `NetSolve ('itmeth', 'PetSc', A, b)`, doesn't appear to indicate which solver and preconditioner should be used. Does NetSolve make that choice?

Jack Dongarra : We are developing a client side system that will probe the data and try to figure out what library is the "best" to use. This, together with what is available and expected time to solution, will determine the actual execution.

Milind Bhandarkar : What work has been done in the context of NetSolve about cost models and licensing models?

Jack Dongarra : The NetSolve system does not have a cost model today (although it does have hooks to do accounting). Clearly some cost model is needed to account for the cycles and software that is used, but today our experiments have focused on building mechanisms to use software and hardware in the grid. Licensing is a bit more tricky since most existing agreements have not taken into account the grid environment.