

LA-UR-00-6049

*Approved for public release;
distribution is unlimited.*

<i>Title:</i>	CartaBlanca— A Pure-Java, Component-based Systems Simulation Tool for Coupled Nonlinear Physics on Unstructured Grids—An Update
<i>Author(s):</i>	<i>W. B. VanderHeyden E. D. Dendy N. T. Padial-Collins</i>
<i>Submitted to:</i>	<i>Concurrency & Practice</i>

Los Alamos NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U. S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness. Form 836 (8/00)

CartaBlanca— A Pure-Java, Component-based Systems Simulation Tool for Coupled Nonlinear Physics on Unstructured Grids –An Update[©]

W. B. VanderHeyden

Los Alamos National Laboratory
Theoretical Division and
Los Alamos Computer Science
Institute
Los Alamos, NM 87545
wbv@lanl.gov

E. D. Dendy

Los Alamos National Laboratory
Theoretical Division and
Los Alamos Computer Science
Institute
Los Alamos, NM 87545
dendy@lanl.gov

N. T. Padial-Collins

Los Alamos National Laboratory
Theoretical Division and
Los Alamos Computer Science
Institute
Los Alamos, NM 87545
nelylanl@lanl.gov

ABSTRACT

This paper describes a component-based nonlinear physical system simulation prototyping package written entirely in Java using object-oriented design. The package provides scientists and engineers a “developer-friendly” software environment for large-scale computational algorithm and physical model development. The software design centers on the Jacobian-Free Newton-Krylov solution method surrounding a finite-volume treatment of conservation equations. This enables a clean component-like implementation. We first provide motivation for the development of the software and then discuss software structure. Discussion includes a description of the use of Java’s built-in thread facility that enables parallel, shared-memory computations on a wide variety of unstructured grids with triangular, quadrilateral, tetrahedral and hexahedral elements. We also discuss the use of Java’s inheritance mechanism in the construction of a hierarchy of physics systems objects and linear and nonlinear solver objects that simplify development and foster software re-use. We also provide a brief review of the Jacobian-Free Newton-Krylov nonlinear system solution method and discuss how it fits into our design. Following this, we show results from example calculations and then discuss plans including the extension of the software to distributed memory computer systems.

1. INTRODUCTION

Specialized simulation software for nonlinear physical systems is one of the central research products from many of the programs here at Los Alamos National Laboratory (LANL) and at other similar institutions. Such systems are often 3-dimensional and are solved on unstructured computational grids. Thus, simulation software for these systems can be quite complex. Very often, simulation projects involve the modification of existing software to produce new capabilities. Furthermore, program goals change frequently as funding priorities evolve or as research developments lead to new branches of investigation. As such, application development is often the bottleneck on these projects. There is, therefore, substantial incentive for software and software

environments that are “developer-friendly”—easy for scientific and engineering software developers to modify and extend. Fortunately, two technology trends address this need in a coordinated fashion.

First, object-oriented and component-based software has made an enormous impact in the commercial software arena for telecommunications and e-Business. Enhancement of software developer productivity has come from a variety of sources including new languages such as Java which support object orientation and component software and a host of software developer productivity tools including graphical design applications, graphical debuggers, etc., that work with these languages. At the same time, significant advances have occurred in nonlinear systems solutions methods. In particular, Jacobian-Free Newton-Krylov (JFNK), [3], methods have been employed and extended to provide robust and flexible solution methods for a wide variety of coupled nonlinear physics simulation capabilities. In addition, the combination of JFNK and the finite-volume technique, [9], fits well into an object-oriented or component-based software environment.

Thus, the goal of the CartaBlanca project is to produce a modern flexible software environment for prototyping physical models, discretization schemes and solution methods for nonlinear physics problems on unstructured grids. CartaBlanca employs an object-oriented, component-based design using the Java programming language. CartaBlanca uses unstructured grids composed of either, triangular or quadrilateral mesh elements in 2 dimensions, or tetrahedral or hexahedral mesh elements in 3 dimensions. CartaBlanca employs the finite-volume method, [9], to provide for flexibility with regard to both meshes and physical effects. Finally, CartaBlanca uses Java’s built-in thread facility for shared-memory parallelization.

1.1. Status

We currently have a great deal of infrastructure for CartaBlanca in place. Our updated list of features since first reporting on CartaBlanca, [28], includes:

- Accepts unstructured grids in 2 and 3 dimensions with triangular, tetrahedral, quadrilateral and hexahedral elements.

- Accepts Metis, [16], generated mesh partition files; computes in parallel on Metis sub domains.
- Has a Graphical User Interface (GUI) for problem specification.
- Uses abstract classes for state, physics and solvers objects; imposes a uniform component-like interface for developers.
- Pre-conditioned Conjugate gradient and GMRES Krylov linear solvers interact seamlessly with physics objects.
- Nonlinear quasi-Newton solver wraps around linear Krylov solvers to provide a JFNK solver.
- Physics objects available for high accuracy scalar advection, interface tracking, heat transfer and multiphase flow.
- Automatic generation of Tecplot (from Amtec Engineering) graphics files for arbitrary physical systems.
- Embedded software-testing facility based on the JUnit framework, [15].
- Direct remote access to CVS revision control server is enabled for efficient team code development.

1.2. Java Performance

A recent article by Schatzman and Donehower, [25], provides a useful discussion of the potential pitfalls involved and tips on how to program in Java to try to achieve performance comparable to that obtainable using C, C++ and FORTRAN. Similar information is also available in reference, [27]. More encouraging news regarding Java performance has been reported by Boisvert, et. Al., [4]. These authors provide an up-to-date review, which shows that Java performance has improved significantly over the past several versions of the Java Version Machine (JVM). Remarkably, they show that, for some cases, Java is slightly faster than C when tested back-to-back on a suite of scientific computing benchmark applications called Scimark from the US National Institute of Standards and Technology. In addition, Reinholtz, [23], writes that not only is it possible for Java to have better general performance than C++, but also that it is likely that this will actually occur. He argues that Java performance can exceed that of C++ because dynamic compilation gives the Java compiler access to runtime information not available to a C++ compiler. Furthermore, the rapidly growing market for embedded systems will drive such performance improvements to extend battery life. Considering the above information and the important benefits from Java's strong typing, clean design, object-orientation, it is our opinion that Java is now a serious alternative language for scientific and engineering computing applications.

1.3. Related Efforts

CartaBlanca builds on some important existing software programs here at LANL. Here is a brief description of these programs. CFDLIB, [17], a FORTRAN 77 program, is an outgrowth of the Caveat program, [1], which provides a flexible finite-volume multiphase flow simulation capability. The Telluride program, [31], provides a multi-material simulation capability with interface tracking on unstructured grids. Telluride makes extensive use of the FORTRAN 90 module concept. The CHAD program, [21], is a flexible node-based finite-volume simulation program for flow simulation on unstructured grids. CHAD uses FORTRAN 90 and accommodates hybrid grids using an edge-based connectivity data structure, [26]. Finally, Kokopelli, is a C++ program for interfacial and polymer flow problems. The

authors have had extensive experience with these programs. The design and implementation of CartaBlanca builds on the lessons learned from these experiences.

It is also worthwhile to note two relevant examples from outside our laboratory. First, Hauser, et. Al., [14], have produced an object-oriented, pure-Java simulation code for aerospace applications. They employ a multi-block structured grid computation scheme and use Java's thread and RMI facilities for parallelization and to enable remote interaction between a graphical user interface and the numerical application.

Another effort worth noting here is that of Hatakeyama, et. Al., [13], who describe an object-oriented paradigm for flow simulation software in which the object-oriented concept is used at the computational node level to produce flexible abstractions. They demonstrate their concepts with a C++, structured grid simulation of a wind tunnel with a test object and show that they can easily insert and extract arbitrary-shaped flow obstacles.

1.4. Outline

In Section 2 we provide an overview of the finite-volume method for discretization of conservation equations. In Section 3 we describe the major features of the Jacobian-Free Newton-Krylov solution method. We then proceed in Section 4 to give an overview of the CartaBlanca software packages. We follow this in Section 5 with some results from example calculations. Finally, in Section 6 we provide a discussion of conclusions and plans.

2. FINITE VOLUME METHOD

CartaBlanca is based on the finite-volume method, [9], for conservation equations. CartaBlanca adopts the node-based version of this scheme with edge-based connectivity, [21], [26]. We provide here a very simplified outline of the method. For an arbitrary control volume V with bounding surface A the generic conservation statement is of the form

$$\frac{d}{dt} \int_V q dV + \int_A \vec{f} \cdot \vec{n} dS + \int_V s dV = 0, \quad (1)$$

where q is the density of some conserved quantity such as mass, momentum or energy, \vec{f} is the local flux of this conserved quantity due to a variety of mechanisms, \vec{n} is an outward normal vector defined on the surface of the control volume, and s is a generalized source density. The first and third integrals in Equation (1) are over the entire space of the control volume; the second integral is over the surface of the control volume. The derivative on the first integral quantity in Equation (1) is with respect to time. For numerical computations, Equation (1) is discretized in time and in space on a computational grid. On such a grid, conservation nodes are connected by edges as shown in Figure 1.




Figure 1. Control volume for i th node.

Each node is associated with a polyhedral control volume, V_i , as depicted in Figure 1. For each node, the averaged value of the conserved density is defined as

$$q_i \equiv \frac{1}{V_i} \int q dV. \quad (2)$$

The quantities q_i are, typically, the state variables for the numerical simulation. Similarly, the average source over each control volume is

$$s_i \equiv \frac{1}{V_i} \int s dV. \quad (3)$$

Let \bar{f}_e be the average flux on the control volume face associated with edge e . Then, if we integrate the Equation (1) over a time step, Δt , using, for example, a first-order difference approximation for the time derivative, we obtain the discretized form of the conservation equation

$$q_i^{n+1} V_i^{n+1} - q_i^n V_i^n + \Delta t \left\{ \sum_{edges} \bar{f}_e \cdot \bar{n}_e A_e + s_i V_i \right\} = 0, \quad (4)$$

where the superscripts n and $n+1$ denote the present and future time levels, respectively. Of course, the fluxes and source terms are generally functions of space, time and the state variables, q_i . Thus, the set of discretized conservation equations for all nodes and all types of conservation quantities forms a nonlinear algebraic system. The physics for a given application lies in the definition of the fluxes and sources in Equation (4). The aim of CartaBlanca is to provide scientists and engineers a friendly environment using object-oriented Java for the implementation of component-like physics and solver objects for the solution of the corresponding coupled nonlinear conservation equations.

3. JACOBIAN-FREE NEWTON-KRYLOV METHOD

We may write the set of conservation equations in the compact, abstract form

$$F_i(q^{n+1}) = 0 \quad (5)$$

where F_i denotes the left hand side of Equation (4) and q^{n+1} denotes the entire set of state variables at the advanced time. The quantity F_i is called the residual function. The system represented by Equation (5) is, in general, nonlinear. We employ the Jacobian-Free Newton-Krylov method, [3], in CartaBlanca to solve these systems. We provide here a brief outline in order to motivate our discussion of the software design. Newton's method for a nonlinear system begins with an initial guess of the solution, $q_j^{n+1(0)}$, where the superscript in parenthesis denotes the iterate level. This is, typically, the solution from time level n . Newton's method then proceeds through a series of iterations involving the solution of a sequence of linear systems

$$J_{ij}(q^{n+1(k)}) \delta_j^k = -F_i(q^{n+1(k)}), \quad (6)$$

along with the update

$$q^{n+1(k+1)} = q^{n+1(k)} + \delta^k, \quad (7)$$

where there is an implied summation in Equation (6) on the repeated index, j . The goal, of course, is to proceed until we find the solution to Equation (5). The matrix quantity, J_{ij} , is the Jacobian matrix defined as

$$J_{ij}(q) = \frac{\partial F_i(q)}{\partial q_j}. \quad (8)$$

Explicit formation of the Jacobian matrix is typically a very expensive computation. Fortunately, the JFNK method takes advantage of the fact that Krylov linear solution methods require only the evaluation of matrix-vector products, Jv (where v is a Krylov vector), and not the matrix J by itself, [3]. Furthermore, matrix-vector products can be approximated numerically using a directional difference formula,

$$Jv \approx \frac{F(q + \varepsilon v) - F(q)}{\varepsilon}, \quad (9)$$

where ε is some small scalar perturbation parameter, [3]. This approximation allows us to structure CartaBlanca in such a way that the physics developer can focus on providing residual functions inside physics objects or components. Using the abstraction embodied in Equation (5) we have genericized the rest of the infrastructure for solving and processing physics problems so that developers can work simultaneously on a variety of different problems using the same software.

4. SOFTWARE

CartaBlanca is composed, at present, of twelve separate packages. Each contains classes that perform distinct functions. We have tried to design these classes to serve as software components that can be interchanged in a "plug and play" mode by developers. We have also tried to write the utility classes in such a way that physics and solver class developers need not concern themselves with the implementation of the parallel features of the software.

In the following, we describe each of these packages and the classes they contain. We also describe the interactions and associations between the classes in the different packages. We choose to start the discussion with the mesh and input packages. These are low-level packages; they are used by many but make sparing use of other packages. We then work our way up through the remaining packages of increasing complexity until we finally describe the main package, which contains the main methods. Before proceeding to the discussion of the CartaBlanca software packages, we start by commenting on our general design approach and on our software engineering methods.

4.1. Approach

Our approach to the design of CartaBlanca includes the following general guiding principles. First, we have endeavored to make use of object-orientation at the highest levels from a physical point of view. Thus, our objects are things that exist over entire sections of the computational domain or grid, rather than at individual nodes. This choice was made based on the idea that this would yield higher numerical performance by avoiding excessive overhead at the node level. This choice also provides a smoother transition into object-oriented programming for developers more familiar with procedural scientific legacy codes. Nevertheless, this approach has allowed us to make substantial use of Java and its object-oriented features.

Another principle we have employed is to make the top levels of the program as generic as possible so that the developer can plug physics into the appropriate program locations and then have the rest of the program able to immediately interact. This was accomplished, in part, by the use of abstract classes, which provide general functionality and interfaces for things such as physics and solver objects. Thus, these objects are like components.

4.2. Software Engineering

Our approach to team programming follows the lightweight processes advocated in the recent article by Fowler, [10]. Iterative programming and component development has, for example, been very useful. The use of team coding has also proved helpful.

In order to foster the team software approach, we have incorporated the JUnit, [15], testing facility into CartaBlanca. This has been useful in that any developer can perform tests easily on their local computing platform to make sure his modifications have not corrupted the software. This is in contrast to a situation in which software testing is performed using specialized software available only on a certain computing platform.

We have found it very helpful to use a common integrated development environment (IDE) for our software development. We are currently using JBuilder 4.0 Professional by Borland Technologies. JBuilder gives us an identical programming environment on our Windows NT and Solaris workstations. JBuilder is also available for LINUX operating systems. The JBuilder environment, conveniently, recognizes the JavaDoc @todo functionality. We use this feature as a simple issues tracking mechanism.

In addition to the JBuilder IDE, we use the GNU CVS revision control software for our software repository. We run CVS as a 'pserver' on one of our Solaris workstations. Thus, we can check pieces of software in and out over the network directly. We currently run a simple implicit heat transfer, scalar advection and

multiphase flow problems (discussed in Section 5) as test problems before committing software modifications to our CVS repository.

4.3. I/O Package

CartaBlanca reads mesh files (see Section 4.6) and writes graphics files (see Section 4.11). In the initial stages of the project, an I/O package was imported from S. J. Chapman, [5], for these operations, [28]. This package contained classes with methods that enabled the developer to write C-language-syntax file print and read statements. We have since abandoned this package in favor of using Java's very effective Reader and Writer classes in the java.io package, [12]. We make use of the non-standard ExponentialFormat class provided in reference [12] for writing doubles to text files. At present, this is the only class contained in the I/O Package.

4.4. Input Package

The input package contains the basic input facilities for problem specification. The user specifies parameters such as solver tolerances, physical properties and boundary conditions using a graphical user interface (GUI). Problem data is written to a 'ProblemSpecifier' class object. The 'ProblemSpecifier' object contains all input information from the GUI. It can be queried by as needed for information in the rest of the program. The 'ProblemSpecifier' object is serializable. This feature is used to save ProblemSpecifier settings to disk. This eliminates the need for any text-based input files, other than the mesh files (see section 4.6).

The GUI is contained in three classes named 'TabbedInputClass', 'TabbedInputFrame' and 'TabbedInputFrame_AboutBox.' The GUI covers several categories of input separated into several tabbed input frames. The input categories are General Information, Physics, Linear Solver, Nonlinear Solver, Pre-conditioner, Initial Conditions, Boundary Conditions and Materials. A snapshot of the GUI interface is shown in Figure 2.

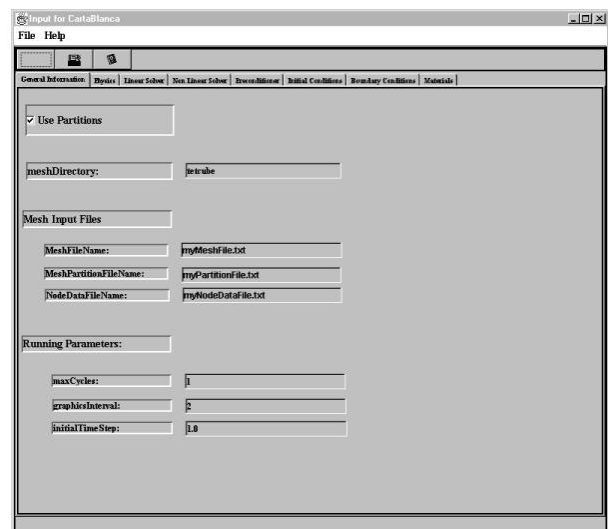


Figure 2. Snapshot of GUI. Tabs enable user to provide input on the various categories of input.

The user can click on the various tabs along the top of the GUI to access the different categories of input. As the user types in new information into the fields of the GUI, the information is written to the ProblemSpecifier object. When the user exits the GUI, the ProblemSpecifier is output to disk as a serialized object for future use and the rest of the program then begins executing based on the information in the ProblemSpecifier object.

4.5. Communications Package

The communications package contains classes of objects that provide functionality for inter-partition communication and for global mesh operations. The class CyclicBarrier provides a simple barrier that objects may invoke to synchronize calculations. The implementation was modeled on the barrier class provided in Chapter 5 of Oaks and Wong, [20]. The CyclicBarrier is used, for example, in discrete operations such as divergence field computations in which communication of flux quantities among mesh partitions are required.

The Reduction class in the communications package provides for the computation of global quantities across the entire mesh such as a global maximum or a global sum. Global sums are required, for example, for mesh-wide dot products of vectors in the various Krylov solvers. The Reduction class accomplishes this by using static class variables for sums and extrema.

4.6. Mesh Package

The mesh package contains several classes that describe mesh elements, edges, interior boundary nodes, and partition and global meshes. These classes are built from mesh information read from mesh files. CartaBlanca requires three types of mesh information files, which follow the format used by the Metis mesh-partitioning program, [16]. The three files contain the mesh connectivity, the node coordinates and the partitioning of the mesh elements. Please see the Metis manual, [16], for a description of these files.

CartaBlanca requires mesh partitioning to be done in such a way that elements and not nodes are partitioned. Referring to Figure 3, the mesh partitioning for CartaBlanca must be done along node-edge connections. In the Figure, the heavier edge connections denote the boundary between partitions A and B. To implement this mode of partitioning in CartaBlanca, nodes on the partition boundaries are duplicated. In the example in the Figure, the three nodes along the partition boundary would be present in each partition as duplicates.




Figure 3. Partitioning in CartaBlanca. Meshes must be partitioned along node connections.

Figure 4 shows an example of an element-partitioned mesh for CartaBlanca.




Figure 4. Two-dimensional 3000 node partitioned mesh. Partitions were generated using Metis.

The mesh partitioning shown in Figure 4 was performed using the Metis program and the Metis output was then fed to CartaBlanca for computations. The actual plot was generated using the Tecplot program which operates on graphics output files from CartaBlanca (see Section 4.11) A further example mesh is shown in Figure 5 for the case of a 3-dimensional tetrahedral mesh.




Figure 5. Three-dimensional tetrahedral element mesh. The shading denotes the 4 partitions that were computed by Metis.

In general terms, the mesh package classes perform the following functions:

- Read and store data from mesh input files. This includes element connectivity, node coordinates and element partitioning,
- Compute all required element and node geometrical information such as cell face areas and normal vectors for the global mesh,
- Compute all edge connectivity and geometric information from the element information for the global mesh,
- Link all nodes via edge elements,
- Setup up partition meshes with links between global and partition mesh objects including nodes, elements and edges.
- Set up connectivity between duplicate nodes on different partitions.

4.7. Discrete Operations Package

The discrete operations package contains a class called Divergence which provides a variety of mesh-wide discrete operations including the computation of the divergence of a vector field, the gradient of a scalar at both mesh nodes and mesh faces as well as some more specialized operations. Some of the specialized operations include finding the maximum face-by-face inflow values for each node for advection calculations and finding the diagonal term of a mesh-wide matrix operator. All of these operations require communication and therefore use the duplicate node connectivity information from the mesh package classes and the barrier object from the communications package.

4.8. Physical Properties Package

Physical properties such as material densities and transport properties such as viscosities, mass diffusivities and thermal conductivities are required in simulations of physical systems. These quantities are often predicted using equations of state from system quantities such as temperature and pressure. The details of these predictions are kept separate from the solution of conservation equations by isolating the implementation in a separate package. This package contains, at present, classes for the prediction of material densities diffusivities, and inter-phase exchange parameters such as drag coefficients. The classes use information input by the user from the materials input pages of the GUI and provide methods to the physics package classes (see Section 4.9) for the materials properties predictions. Thus, the physical properties package classes also insulates the physics developers from changes in the details of the materials properties input specifications.

4.9. Physics Package

The physics package contains classes that allow a developer to encode the conservation equations that he or she would like to solve. The developer first must set up an AbsState class corresponding to his physical system. AbsState is a container class (see section 4.9.1). Once the AbsState class is set up, the user then can encode his conservation equations in an AbsProblemPhysics class. This is discussed in section 4.9.2. When specifying both the AbsState class and the AbsProblemPhysics class, the user must extend abstract classes

that provide the basic interface expected by the rest of CartaBlanca.

4.9.1. AbsState Class

The AbsState class is an abstract class that must be extended by the developer to provide a data container for state variables for specific physics problems. The state variables are fundamentally stored in a two dimensional array wherein the first dimension is the variable type and the second dimension is the node index. For example, if one is trying to solve a problem with state variables for pressure, and three components of velocity, then the first dimension of this array would be four. The two-dimensional representation is convenient for developers since they tend to work with the governing equations a field or state variable type at a time. The two-dimensional view is also a convenient format for the graphics package since it also processes the data a field at a time.

Krylov solvers, however, work in terms of a one-dimensional state vector. Thus, the AbsState class also provides a one-dimensional view of the same state data. Currently, the one-dimensional view is provided as a copy of the two dimensional data. The copy is performed using Java's System.arraycopy function for best performance.

4.9.2. AbsProblemPhysics Class

For linear physical systems, developers can specify their physical system behavior by extending the AbsProblemPhysics class. AbsProblemPhysics is an abstract class that lays out what CartaBlanca expects from physics objects. The most important feature of this class of objects is the methods to get the right and left hand side of the governing equations for the state variables. The solvers in CartaBlanca interact with these physics object methods to obtain the right-hand side of the linear equation system and the matrix-vector multiply. Another important behavior of AbsProblemPhysics objects is the pre-conditioning method. The Krylov solvers also interact with physics objects by invoking their pre-conditioning method. This method takes a Krylov vector from the Krylov solver and updates it according to some iterative improvement scheme. Currently, diagonal, Jacobi and symmetric successive over-relaxation pre-conditioning are available. Plans for a multigrid-like scheme are in place to obtain improved solver performance.

AbsProblemPhysics classes also inherit some methods for the base class for converting time n states to time $n+1$ states. These methods, of course, can be overridden in the derived classes to provide additional functionality.

4.9.3. NLAbsProblemPhysics Class

In the case of nonlinear physics problems, the matrix-vector multiply evaluation has to be provided in a generic fashion following Equation (9). The NLAbsProblemPhysics class of CartaBlanca extends the AbsProblemPhysics to provide this behavior. In this class of objects, the developer must encode the governing equations into methods that return the full nonlinear residual equation in the form of a left and right hand side. The left and right hand side correspond to the implicit and explicit parts of the governing equations. These objects invoke these nonlinear *get* methods from the overridden linear *get* methods from AbsProblemPhysics class using Equation (9) to produce a linear matrix-vector multiply evaluation. Since

NLAbsProblemPhysics inherits from AbsProblemPhysics, all other behavior, such as pre-conditioning is also available.

Figure 6 provides a graphical overview of the physics class inheritance hierarchy that was generated directly from the Java source code using GPro from Embarcadero Technologies.

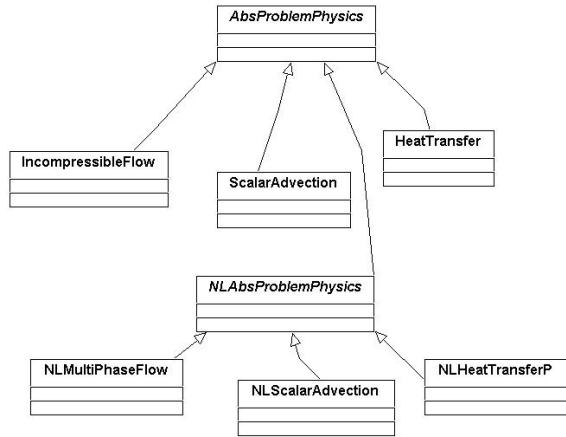


Figure 6. UML Class hierarchy diagram of the Physics package.

4.10. Solver Package

The solver package contains classes for linear and nonlinear solvers. As was the case for the classes in the physics package, an abstract solver class, AbsSolver, is provided as a parent for all solvers. Currently, this class has been extended to provide users Krylov solver classes based on Conjugate Gradient and both the standard and flexible variant of Gmres, [24]. In addition, an ‘explicit’ solver is provided for fully explicit calculations which essentially bypasses any solution method at all and simply returns the right hand side as the solution. Finally, a Newton-Krylov (JFNK) solver is provided for nonlinear problems. Each of these solvers communicates directly with physics objects through method invocations. The solver class inheritance hierarchy is shown in Figure 7.

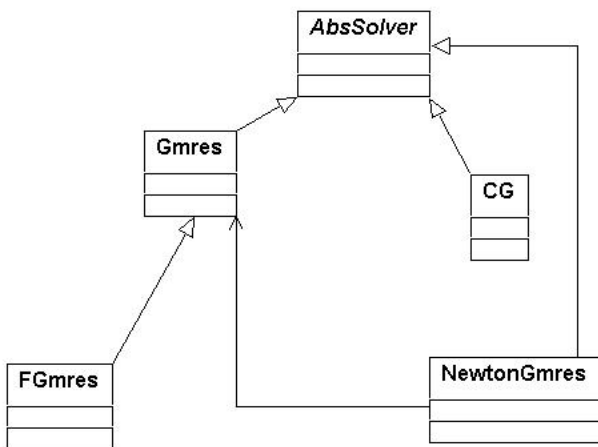


Figure 7. Solver package class hierarchy.

4.11. Graphics Package

The Graphics package, at present, contains only one class that can be used to produce Tecplot format output text files. The class interacts with the abstract state class so that it automatically knows about new state variables, etc. Eventually, this class will be extended to allow for additional plot file output formats. We also envision direct use of Java graphics.

4.12. Problem Driver Package

The ProblemDriver package contains the Driver class, a top-level driver for solving physics problems on each mesh partition. The Driver class implements Java’s Thread-class Runnable interface. This enables data-parallel computation in CartaBlanca with each thread corresponding to a particular mesh partition. Figure 8 shows a UML association diagram for the Driver class.

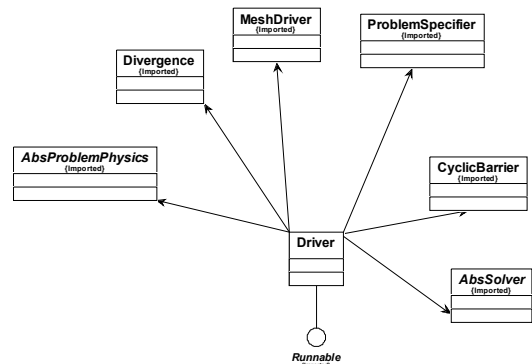


Figure 8. Association diagram for driver class.

As can be seen, the Driver class interacts with all the major CartaBlanca objects from an AbsProblemPhysics object to an AbsSolver object.

4.13. Boundary Conditions Package

Boundary conditions are required for the complete specification of all but the simplest physical problems. In order to build a layer of abstraction between the core physics classes and the user interface for boundary conditions, we have introduced a separate boundary conditions package. At present, this package contains only one class, which implements boundary conditions. This class takes user input data from the ProblemSpecifier object and provides methods for setting boundary fluxes for use in the conservations equations in the physics classes.

4.14. Main Package

The main package consists of several classes that contain the public static main method that drives the entire simulation. The class PhysMain contains a main method that instantiates all high-level objects and invokes the start method for all of the Driver objects for each mesh partition.

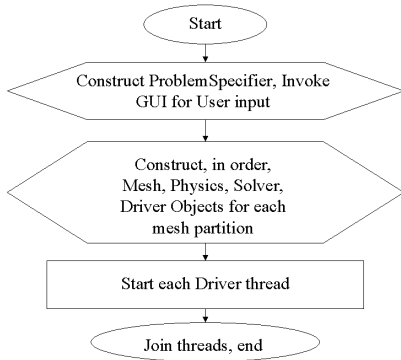


Figure 9. Flow chart for CartaBlanca main method.

5. RESULTS

Here we provide information on compilation, run-time performance and some preliminary simulation results from CartaBlanca. Computations were performed on a dual 500 Mhz Pentium III processor SGI 320 PC running a Windows NT operating system, on a single processor Sun Ultra 60 workstation running Solaris 2.7 operating system and on a 4-processor SGI Origin 200 workstation running the IRIX operating system. The entire software package compiles to byte code in a matter of a few seconds.

5.1. Scalar Advection

The first two examples are of explicit scalar advection calculations done on the triangular element grid shown in Figure 4. The results of the calculations are shown in Figure 10 and Figure 11. These calculations were performed using 2 threads in parallel on our dual-processor SGI PC. Each thread operated on a separate mesh partition.

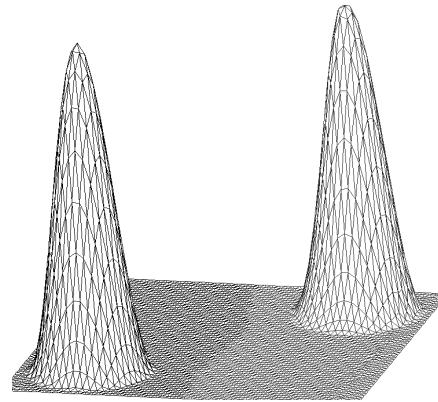


Figure 10. Results from a continuum scalar advection simulation on a two-dimensional triangular mesh.




Figure 11. Results from an interface-tracking scalar advection simulation on a two-dimensional triangular mesh.

Each calculation started with a pulse of concentration in the lower left hand corner of the domain. In the first example, the concentration pulse has a Gaussian spatial distribution while in the second example; the pulse was a spatial step function with circular shape. This material was advected towards the top right according to the conservation equation

$$\frac{\partial c}{\partial t} + \frac{\partial c u_i}{\partial x_i} = 0 \quad (10)$$

where c is concentration, u_i is the i th component of velocity, t is time and x_i is the i th spatial coordinate. There is an implied summation over the repeated index in the second term. In these example calculations, the velocity was a constant with component values of one in the horizontal and vertical directions. The final conditions of the concentration pulses are shown in the top-right of Figure 10 and Figure 11. Note that in the case of Figure 10 we used the continuous field version of CartaBlanca's advection facility while in Figure 11 we used the interface-tracking version of the advection facility, [6].

The calculations used CartaBlanca's explicit solver to advance the solution of Equation (10) over 120 time cycles. Although Equation (10) is relatively simple, it is commonly known that it is

difficult to obtain accurate numerical solutions for this equation that minimize the artificial numerical smearing of such concentration pulses and also avoid the creation of artificial extrema, especially on triangular grid meshes. Also, doing such calculations using interface tracking is also a very difficult task. These examples demonstrate that CartaBlanca has advanced advection algorithms that can be used by developers for a wide variety of applications.

5.2. Implicit Heat Transfer

The second example is a simple implicit heat transfer calculation performed on similar triangular element meshes. CartaBlanca's flexible Gmres solver was used without preconditioning to compute the temperature field as a function of time according to the equation

$$\frac{\partial T}{\partial t} = \frac{\partial}{\partial x_i} \left(\alpha \frac{\partial T}{\partial x_i} \right) \quad (11)$$

where T is temperature and α is the thermal diffusivity, which was set to one for the example problem. Zero-gradient boundary conditions were used which correspond physically to a perfectly insulated box. The final state is, therefore, a uniform temperature of one-half. CartaBlanca computed the correct solution on a 46000-node triangular element mesh, using both a single mesh partition and multiple mesh partitions running on separate threads in parallel.

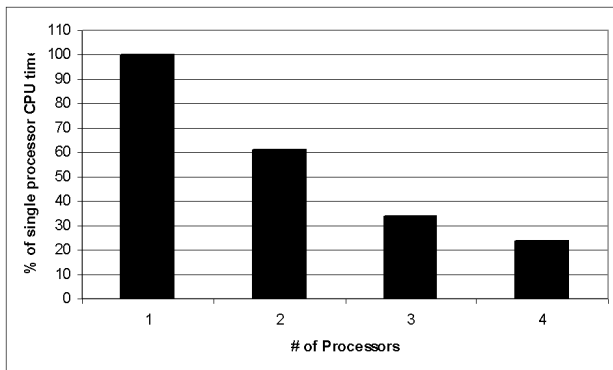


Figure 12. Preliminary example of scaling behavior of CartaBlanca on linear heat transfer problem. Values are wall clock time for problem solution scaled by that for a single processor calculation.

Figure 12 shows the scaling behavior for this heat transfer problem on a 4-processor SGI Origin 200 workstation. These are very preliminary results and are likely to change as we perform more tests on the parallel performance of CartaBlanca. They do show, however, that CartaBlanca appears to scale well for this problem.

5.3. Multiphase Flow

We also can compute multiphase flow using CartaBlanca. The details of the governing equations and the numerical algorithm are beyond the scope of this publication. Those interested in the

governing equations for multiphase flow should consult reference, [7], for example. Our numerical algorithm is closely related to the one used in CFDLIB, [17]. We plan to publish a detailed exposition of our algorithm in the near future. In this paper we demonstrate our ability to compute multiphase flows on complex geometries and to show the performance of CartaBlanca compared to the CFDLIB FORTRAN code. Our first example calculation is of the flow of sand grains and water in an hourglass. Our second example is a computation of a so-called broken dam experiment due to Martin and Moyce, [19].

5.3.1. Hourglass flow

We show here, some results from computations performed on a triangular element mesh in the shape of an hourglass. Figure 13 shows both the mesh and the initial conditions for the calculation. The mesh contains 1018 elements and 578 nodes. The overall dimensions of the problem are 10 cm by 10 cm. The width of the neck of the hourglass is 1 cm. The contours indicate the initial distribution of materials. The dark region in the middle of the hourglass is filled with 60% (by volume) sand and 40% water. The remainder of the space is filled with water. At time zero, gravity is "turned on" and the sand begins to fall through the hourglass throat. The computation of this example problem was performed with CartaBlanca on our SGI PC using a single mesh partition. Figure 14 and Figure 15 show the time evolution of the as calculated by CartaBlanca. The computation took about 15 minutes on the PC.

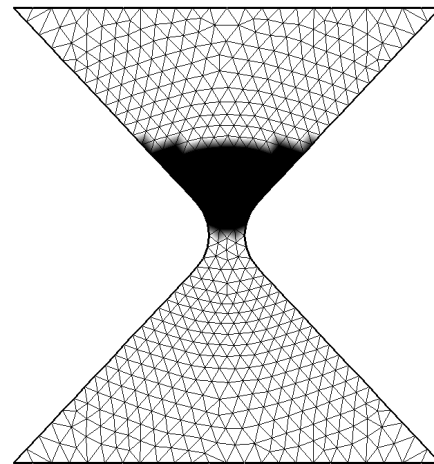


Figure 13. Initial conditions and mesh for hourglass problem. Dark region is 60% sand and 40% water by volume.




Figure 14. Hourglass flow problem at 0.2 seconds, real time.




Figure 15 Hourglass flow problem at 0.36 seconds, real time.

5.3.2. Broken Dam Flow

The final example calculation is of a broken dam experiment performed by Martin and Moyce, [19]. Martin and Moyce constructed a small box with dimensions of 28.57cm by 12.57cm by 2.3 cm out of plexiglass. Inside the box, they installed a hinged door that served as a dam for water at one end. The door was hinged so that it could be opened quickly to allow water to fall freely under the action of gravity. The experimental set up is shown in Figure 16

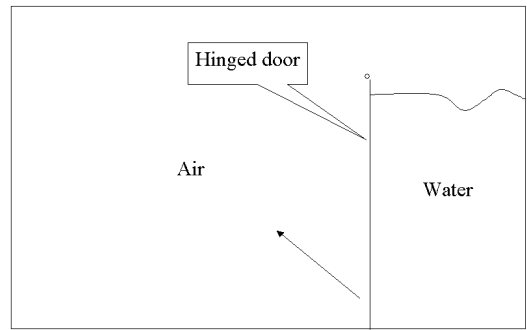


Figure 16. Schematic side-view diagram of broken dam experiment.

Martin and Moyce recorded the time and position of the front of water that moved across the bottom of the apparatus. We have used CartaBlanca to successfully simulate the Martin and Moyce broken dam experiment on a variety of two- and three-dimensional meshes composed of quadrilateral, triangular, hexahedral and tetrahedral elements. Figure 17 shows the tetrahedral element mesh we used, for example. This mesh consisted of 10076 elements and 2205 nodes.




Figure 17. Tetrahedral element mesh for the broken dam problem.

Figure 18 shows the initial condition of the interface for the calculation on the tetrahedral element mesh. The 'bumpiness' of the interface is an artifact of the irregular tetrahedral mesh and our, at present, simplified way of initializing materials in CartaBlanca. Currently, we simply assign to each node a volume fraction of either 0 or 1 according to its position relative to the air-water interface. In the future, we will implement the ability to assign fractional volume fractions to nodes that intersect the interface using a particle interpolation technique.

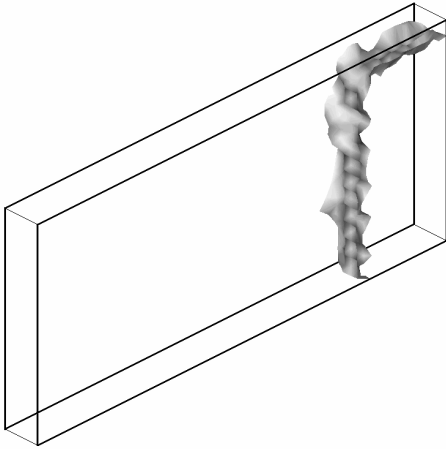


Figure 18. Initial condition for broken dam experiment simulation. Isosurface shows interface between water and air. Water is to the right and below the interface.

In Figure 19 through Figure 23 we show the progression of the air-water interface after opening the dam door to the point in time at 0.2 seconds after the opening of the dam door where the water reaches the far end of the box.




Figure 19 Interface at 0.04seconds.




Figure 20. Interface at 0.08seconds.




Figure 21. Interface at 0.12 seconds.




Figure 22. Interface at 0.16 seconds.




Figure 23. Interface at 0.2 seconds.

In Figure 24 we show a quantitative comparison of the water front position between the experiments of Martin and Moyce and the calculations using CartaBlanca. The CartaBlanca results are a little ahead of the Martin and Moyce data and may signal the need for improvements in the physical models incorporated into CartaBlanca. (The CartaBlanca simulation neglected viscous and surface tension effects.)




Figure 24. Comparison of CartaBlanca computation and data of Martin and Moyce for water front position as a function of time.

5.4. Performance

Our last task here is to give some measure of the performance of CartaBlanca. To provide a basis of comparison, we used CFDLIB, [17], and CartaBlanca to perform simulations of the Martin and Moyce broken dam problem on both two- and three-dimensional Cartesian meshes. We furthermore ran both codes using the same treatment of advection mass fluxes, namely, 1st order accurate donor cell advection. Finally, we used the Conjugate Gradient method for the solution of the pressure equation in CFDLIB and for the preconditioning of the non-linear pressure residual in CartaBlanca. Thus, the two sets of calculations, while not identical, were as close to the same as possible in terms of the numerical methods used. Table 1

provides a comparison of the speed of the codes in terms of the so-called grind time--average wall-clock processing time per time step per node.

Case\Code	Elements	Time Steps	CartaBlanca	CFDLIB
2D quadrilateral element mesh	1100	261	321	147
3D hexahedral element mesh	4400	279	779	400

Table 1 Performance comparisons between CartaBlanca and CFDLIB on the broken dam problem using 1st order (donor cell) fluxes for advection. Table entries are ‘grind times’ defined as average wall-clock processing time in microseconds per time step per node.

All calculations were performed on our Sun Ultra 60 workstation running Solaris 2.7. For the CartaBlanca calculations, we used Sun JDK 1.3.1 with the HotSpot JIT. For the CFDLIB calculations, we used Sun FORTRAN 77 compiler version 5.0 with optimization. As can be seen from the Table, CartaBlanca achieved 46% of the speed of CFDLIB in the two-dimensional case and 51% of CFDLIB in the three-dimensional case. While this is not a perfect side-by-side comparison of Java and FORTRAN it is a reasonably close comparison. The results are quite pleasing to us when we consider that CFDLIB is a highly optimized FORTRAN code, which has many man-years of effort behind it and a worldwide user base. Furthermore, CFDLIB is recognized as a fast multiphase flow code by our users. Finally, CFDLIB was written for structured grids and does not use indirect addressing, as does CartaBlanca.

6. CONCLUSIONS AND FUTURE PLANS

Since we are still in the development phase of this project, it is not appropriate to provide final assessments of the design and performance of CartaBlanca. In a qualitative sense, we can say that CartaBlanca has already been a useful tool for algorithm development since we have been able to do the research on our new advection scheme exclusively within CartaBlanca, [6]. That is, CartaBlanca has provided us with sufficient usability and performance that we did not choose to go “off-line” to some other program or development tool to do one of our main jobs, algorithm research and development. Therefore, in this sense, we can already claim some “bottom-line” success.

The performance comparison of CartaBlanca with CFDLIB is quite encouraging. We hope that with further improvements to the CartaBlanca software and with improved Java compilers and VM’s we will see the gap close.

We plan to pursue a number of promising leads in the near future. The following is a partial list covering some of the major possible directions:

- Fully investigate the use of Java native-code compilers. The GNU compiler, GCJ, is particularly attractive in that it provides the option of turning off array-bounds checking. This could provide significant speed-up, [11].

- Incorporate fluid-structure interactions to the nonlinear multiphase flow class for solidifying flow and elastic-plastic flow simulations.
- Extend CartaBlanca to distributed memory architecture for cluster-based computing. We are considering the use of JavaParty, [22], or Jackal, [29], [30], as a means to extend CartaBlanca to distributed memory systems.

7. ACKNOWLEDGMENTS

We gratefully acknowledge the support for this work from the Department of Energy and the Los Alamos Computer Science Institute (LACSI).

8. REFERENCES

- [1] F. L. Addessio, J. R. Baumgardner, J. K. Dukowicz, N. L. Johnson, B. A. Kashiwa, R. M. Rauenzahn, C. Zemach, *CAVEAT: A Computer Code for Fluid Dynamics Problems with Large Distortion and Internal Slip*, Los Alamos National Laboratory Report LA-10613-MS, Rev. 1, May, 1992.
- [2] R. M. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, H. Van der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition*, SIAM, Philadelphia, 1994.
- [3] P. N. Brown and Y. Saad, Hybrid Krylov Methods for nonlinear systems of equations, *SIAM J. Sci. Stat. Comput.*, 11(3):450-81, 1990.
- [4] R. F. Boisvert, J. Moriera, M. Phillipsen, R. Pozo, Java and Numerical Computing, *Computing in Science & Engineering*, 3(2):18-24, Mar.-Apr., 2001.
- [5] S. J. Chapman, *Java for Engineers and Scientists*, Prentice Hall, Upper Saddle River, NJ, 2000.
- [6] E. D. Dendy, N. T. Padiyal-Collins and W. B. VanderHeyden, *A General Purpose, Finite-Volume Advection Scheme for Continuous and Discontinuous Fields on Unstructured Grids*, in review, *J. Comp. Phys.*
- [7] D. A. Drew and S. L. Passman, *Theory of Multicomponent Fluids*, Springer, New York, 1999.
- [8] R. M. Eckstein, M. Loy and D. Wood, *Java Swing*, O'Reilly, Cambridge, 1998.
- [9] J. H. Ferziger and M. Peric, *Computational Methods for Fluid Dynamics*, Springer, New York, 1999.
- [10] M. Fowler, Put Your Process on a Diet, *Software Development Magazine*, 2(12), December 2000.
- [11] <http://gcc.gnu.org/java/index.html>
- [12] E. R. Harold, *Java I/O*, O'Reilly, Cambridge, 1998.
- [13] M. Hatakeyama, M. Watanabe and T. Suzuki, Object-Oriented Fluid Flow Simulation System, *Computers & Fluids*, 27(5):581-597, 1998.
- [14] J. Hauser, T. Ludewig, T. Gollnick, R. Winkelmann, R. Williams, J. Muylaert and M. Spel, A Pure Java Parallel Flow Solver, *AIAA paper 99-0549*.
- [15] JUnit, <http://www.junit.org/>.
- [16] G. Karypis, and V. Kumar, *METIS A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, Version 4.0*, University of Minnesota, Department of Computer Science / Army HPC Research Center, Minneapolis, MN (<http://www-users.cs.umn.edu/~karypis/metis/index.html>)
- [17] B. Kashiwa, N. T. Padiyal, R. M. Rauenzahn and W. B. VanderHeyden, A Cell-Centered ICE Method for Multiphase Flow Simulations, *FED-Vol. 185, Numerical Methods in Multiphase Flows, ASME*, 185:159-176, 1994.
- [18] C. T. Kelly, *Iterative Methods for Linear and Nonlinear Equations*, SIAM, Philadelphia, 1995.
- [19] J. C. Martin and W. J. Moyce, *An experimental study of the collapse of liquid columns on a rigid horizontal plane*, *Philos. Trans. R. Soc.*, A244:312-324, 1952.
- [20] S. Oaks, and H. Wong, *Java Threads*, O'Reilly, Cambridge, 1999.
- [21] P. J. O'Rourke, and M. S. Sahota, *CHAD: A Parallel, 3-D. Implicit, Unstructured-Grid, Multimatierial, Hydrodynamics Code for All Flow Speeds*, Los Alamos National Laboratory Report LA-UR-98-5663, October, 1998.
- [22] M. Philippsen, and M. Zenger, JavaParty: transparent remote objects in Java, *Concurrency-Practice and Experience*, 9:1225-1242, November 1997.
- [23] K. Reinholtz, Java will be faster than C++, *ACM Sigplan Notices*, 35(2):25-28, Feb 2000
- [24] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing, San Fransisco, 1995.
- [25] J. Schatzman and R. Donehower, High-Performance Java Software Development, *Java Report*, 6(2):24-41, 2001.
- [26] V. Selmin, The Node-Centered Finite Volume Approach: Bridge between Finite Differences and Finite Elements, *Comput. Methods in Appl. Mech. Engrg.*, 102(1):107-138, January, 1993.
- [27] J. Shirazi, *Java Performance Tuning*, O'Reilly, Cambridge, 2000.
- [28] W. B. VanderHeyden, E. D. Dendy and N. T. Padiyal-Collins, CartaBlanca- A Pure-Java, Component-based Systems Simulation Tool for Coupled Nonlinear Physics on Unstructured Grids, *Proceedings for the ACM 2001 Java Grande/ISCOPE Conference*, ACM Press, New York, 2001.
- [29] R. Veldema, R. F. H. Hofman, C. Jacobs, R. A. F. Bhoedjang, and H. E. Bal, Jackal, A Compiler-Supported Distributed Shared Memory Implementation of Java, submitted for publication. <http://www.cs.vu.nl/rveldema/jackal-2001.ps>, 2001.
- [30] R. Veldema, R. F. H. Hofman, R. A. F. Bhoedjang, and H. E. Bal, Runtime Optimizations for a Java DSM Implementation, *Proceedings for the ACM 2001 Java Grande/ISCOPE Conference*, ACM Press, New York, 2001.
- [31] Telluride, <http://public.lanl.gov/mww/HomePage.html>, Los Alamos National Laboratory Report LA-UR-99-1664, 1999.