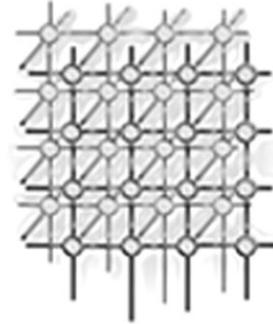


Programming environments for multidisciplinary Grid communities



N. Ramakrishnan^{*,†}, L. T. Watson, D. G. Kafura, C. J. Ribbens
and C. A. Shaffer

Department of Computer Science, Virginia Tech, Blacksburg, VA 24061, U.S.A.

SUMMARY

As the power of computational Grids increases, there is a corresponding need for better usability for large and diverse communities. The focus in this paper is on supporting multidisciplinary communities of scientists and engineers. We discuss requirements for Grid computing environments (GCEs) in this context, and describe several core support technologies developed to meet these requirements. Our work extends the notion of a programming environment beyond the compile–schedule–execute paradigm, to include functionality such as collaborative application composition, information services, and data and simulation management. Systems designed for five different applications communities are described. These systems illustrate common needs and characteristics arising in multidisciplinary communities and motivate a high-level design framework for building GCEs that meet those needs. Copyright © 2002 John Wiley & Sons, Ltd.

KEY WORDS: Grid computing environments; problem solving environments; multidisciplinary Grid communities; compositional modeling

1. INTRODUCTION

Grid computing environments (GCEs) have increasingly gained attention in the past few years. Advances in technological infrastructure as well as a better awareness of the needs of application scientists and engineers have been the primary motivating factors. In particular, the shift in emphasis

*Correspondence to: Naren Ramakrishnan, Department of Computer Science, Virginia Tech, Blacksburg, VA 24061, U.S.A.

†E-mail: naren@cs.vt.edu

Contract/grant sponsor: National Science Foundation; contract/grant number: DMS-9625968; EIA-9974956; EIA-9984317; EIA-0103660 and MCB-0083315

Contract/grant sponsor: NASA; contract/grant number: NAG-2-1180

Contract/grant sponsor: AFOSR; contract/grant number: F496320-99-1-0128

Contract/grant sponsor: Defense Advanced Research Projects Agency; contract/grant number: N00014-01-1-0852

Contract/grant sponsor: Virginia Tech ASPIRES Program



from low-level application scheduling and execution [1] to high-level problem-solving (the focus of this special issue) signals that Grid computing will become increasingly important as a way of doing science. We use the term GCE to broadly denote any facility by which a scientist or engineer utilizes Grid services and resources to solve computational problems. Our definition thus includes facilities from high-performance scientific software libraries [2] augmented with Grid access primitives to domain-specific problem-solving environments (PSEs) [3,4] that provide targeted access to applications software.

GCEs extend the notion of a programming environment beyond the compile–schedule–execute paradigm to include functionality such as networked access [5], information services, data management, and collaborative application composition. This is especially true when designing such systems for supporting *multidisciplinary Grid communities*, the focus of this paper. Our emphasis at Virginia Tech has been on exploiting the high-level problem-solving context of such ‘virtual organizations’ [6] and building on the Grid architectures, services, and toolkits (e.g. [7–10]) being developed by the Grid community. In working with large, concerted groups of scientists and engineers in various applications (aircraft design, watershed assessment, wireless communications system design, to name a few), we have identified several recurring themes important for supporting and sustaining such communities. Our goal in this paper is to document these themes, present a high-level design framework for building GCEs [11], and describe some solutions we are working on to address the concomitant needs.

In the remainder of this section, we present usage scenarios from multidisciplinary communities that will help us characterize requirements for GCEs. We also describe a high-level framework for building and organizing programming environments for such communities. In Section 2, we describe PSEs that we have built for five different multidisciplinary communities. Section 3 discusses a variety of issues pertaining to software systems support for GCEs. In particular, our semistructured data management facility plays a central role in exploiting the rich problem-solving context of multidisciplinary Grid communities. Two other elements of our GCE framework are described in Section 3: Sieve (a collaborative component composition workspace) and Symphony (a framework for managing remote legacy resources). We conclude with a brief discussion of future directions in Section 4.

1.1. Multidisciplinary Grid communities: scenarios

We begin by describing some scenarios to illustrate the needs typical of multidisciplinary Grid communities. We posit that there are fundamental differences in the usage patterns for a single researcher (or even a group of collaborators) working on a relatively homogeneous problem as compared to the usage patterns found in the communities we have in mind. For example, how does a Grid community for solving matrix eigenvalue problems differ from, say, one for aircraft design or wireless communications? We identify three scenarios that are suggestive of the distinctions we would like to make.

- *Scenario 1.* A specialist in ray tracing, a channel modeler, and a computer scientist are addressing the problem of determining the placement of wireless base stations in a square mile area of a large city such that the coverage is optimal [12]. In terms of execution, this problem involves a computation expressed as a digraph of components, written in multiple

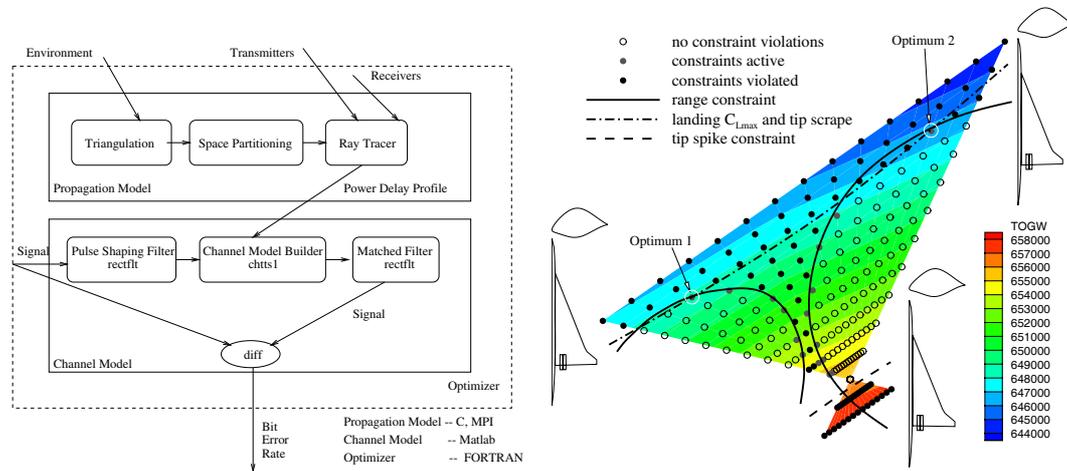


Figure 1. Left-hand side, compositional modeling for designing a wireless communications system. Right-hand side, a slice of an aircraft configuration design space through three design points.

languages (C, Matlab, and FORTRAN), and enclosed in an optimization loop (see Figure 1, left). Notice that information is exchanged between executions in three different languages and is streaming between the optimizer and the simulation. In addition, a variety of intermediate results are produced, not all of which are direct performance data. Such results are typically cached to improve performance, visualized at different stages of the execution, or simply saved for later inspection. Furthermore, the components (codes) are developed at different times by different researchers and many are still under active development. Their I/O specifications hence cannot be enumerated in advance to achieve matching of components. Further, the possibilities of how components could be cascaded and combined can itself evolve over time. How can a programming environment be designed that allows the binding of problem specifications to arbitrary codes and allows their arbitrary composition?

- *Scenario 2.* A team of aircraft design engineers and numerical analysts are attempting to minimize the take-off gross weight (TOGW) for an aircraft configuration design involving 29 design variables with 68 constraints [13] (see Figure 1, right). High-fidelity codes dealing with aerodynamics, mechanics, and geometry determine how changes in design variables affect the TOGW. This application domain is characterized not by an abundance of data, but rather by a scarcity of data (owing to the cost and time involved in conducting simulations). Consequently, the solution methodology involves a combination of high accuracy computations, surrogate modeling (to provide response surface approximations for unsampled design regions [14]) and a robust data management system to help focus data collection on the most promising regions. As a result, evaluating a design point might involve executing a high-fidelity computation, using low-fidelity approximations to obtain an estimate of the TOGW and/or querying a database to lookup previously conducted simulations. In addition, the resources for computing and data storage



could be geographically distributed. How can a single environment provide unified access to such diverse facilities and what programming abstractions are available that allow its efficient and effective use?

- *Scenario 3.* A group of computer scientists, nuclear physicists, and performance engineers are modeling Sweep3D [15,16], a complex ASCI benchmark for discrete ordinates neutron transport. They concur that efficient modeling of this application (over the 90% accuracy level) requires analytical modeling, simulation, and actual system execution paradigms simultaneously [17]! They use a metasytem infrastructure to combine these various models together in a unified manner. However, they are undecided over when to *switch* codes during the computation—do they use a low-level simulator for 80% of the available time, and then switch to analytic models or can they be confident of extrapolating using analytical models even earlier? What if excessive multi-threading on a machine leads to too many fluctuations in their estimates? What system architectures are available that enable compositional modeling when information about component choices is obtained during the computation (rather than before)?

1.2. Multidisciplinary grid communities: themes

The key dominant theme in these scenarios, and one that is well accepted as an integral aspect of Grid computing, is the ability to do *compositional modeling* [18,19]. In the context of problem-solving, Forbus [20] defines this term as ‘combining representations for different parts of a (computation) to create a representation of the (computation) as a whole’. In this paper, we employ this term to convey merely an approach to problem solving and its use is not meant to imply an implementation technology, such as distributed object components (although that is one of the common ways of providing the functionality). For instance, a scientist explicitly moving input and output files across multiple program executables can be viewed as performing compositional modeling (albeit in a very primitive manner). Thus, a component could be any piece of software, executable, model fragment, or even a set of equations that helps the scientist to formalize the process of modeling a computation.

A second aspect (again, one whose assertion will hardly be controversial) is *collaboration*. By definition, a GCE for a Grid community must support groups of scientists and engineers, not just single investigators. These users rely on each others’ codes and data, contribute results to the total effort, communicate in a variety of ways, and organize themselves around subproblems in ways that are hard to predict. They may need to collaborate in real time on a given simulation, but they are often at physically separate locations. Collaborative workspaces are fundamental to the way multidisciplinary research is conducted.

While the above two aspects are underscored in many Grid projects, GCEs for multidisciplinary communities have a unique responsibility (and opportunity) to exploit the larger *context* of the scientific or engineering application, defined by the ongoing activities of the pertinent community. Typical GCEs only deal with one simulation at a time. The larger context we allude to here may include previous scientific results which can be used to improve the efficiency of current simulations or avoid computation altogether if a desired result is already available. The context may denote the fact that a simulation is being run as part of a higher-level problem-solving strategy, e.g. involving optimization or recommendation. Context also implies previous computational experience or performance, e.g. Grid resources may be assigned more intelligently if the performance of previous similar simulations is known. A final example of context is the fact that a given simulation is often part of an *ensemble* of



simulations; recognizing this aspect can help in creating more sophisticated simulation management tools.

As we will show below, the synergy resulting from consideration of all of the above three aspects—compositional modeling, collaboration, context—poses a unique set of research issues pertinent for multidisciplinary communities. An important goal of our approach is to maximize the synergy between Grid computing on the one hand and multidisciplinary scientific problem-solving on the other. Thus, we are trying to answer two questions: (i) ‘How can a multidisciplinary community setting be exploited to better use a Grid?’; and (ii) ‘How can the Grid setting be exploited to better serve a scientific problem-solving community?’

1.3. GCEs for multidisciplinary Grid communities: characteristics

Abstracting from the scenarios described above, and reflecting on the three themes just discussed, where do we locate multidisciplinary Grid communities in the ‘space’ of computational Grid users? To answer that question, we find that the following three dimensions are useful in characterizing GCEs for multidisciplinary communities. These dimensions should not be viewed as a one-to-one translation of the above themes into features; rather, they are the most pertinent forms of distinctions that will help us identify requirements for GCEs for multidisciplinary communities.

- **Emphasis on component coding effort versus component composition effort.** Traditional programming environments emphasize either the coding of components (influenced by an implicit composition style) or the aspect of connecting them together (to prototype complex computations). For instance, when coding effort is paramount and composition is implemented in a distributed objects system (e.g. [21,22]), techniques such as inheritance and templates can be used to create new components. Other implementations involving parallel programming [23–25] or multi-agent coordination [26,27] provide comparable facilities (typically APIs) for creating new components. Component composition effort, on the other hand, emphasizes the modeling of a computation as a process of, say, graphically laying out a network of components (e.g. [28]). By providing a sufficiently rich vocabulary and database of primitive components, emphasis is shifted to composition rather than coding.

Design decisions made about component implementation and composition style indirectly influence the options available for composition and coding, respectively. This dimension distinguishes programming environments based on how they ‘carve up’ compositional modeling; which of these efforts do they emphasize more? By placing what forms of restrictions and assumptions on the other? In a multidisciplinary setting (e.g. *Scenario 1*), programming environments are required to emphasize both efforts in almost equal importance. The needs of the underlying application (in this example, wireless communications) render typical assumptions on both coding and composition style untenable.

- **Cognitive discordance among components.** An indirect consequence of typical compositional modeling solutions is that they commit the scientist to an implementation (and representation) vocabulary. For example, components in LSA [22] (and most object-based implementations) are required to be high-performance C++ objects, instantiated from class definitions. This is not a serious constraint for typical Grid communities since there is usually substantial agreement over



the methodology of computation. The only sources of discordance here involve format conversions and adherence to standards (e.g. matrices in CSR format versus matrices in CSC format). In multidisciplinary Grid communities (see *Scenarios 1 and 2*), there are huge differences in vocabulary (e.g. biologists, civil engineers, and economists using a watershed assessment PSE have almost no common terminology) and fundamental misunderstandings and disagreements about the way computations should be organized and modeled (e.g. aerodynamicists, control engineers, and structural engineers model an aircraft in qualitatively different ways). Furthermore composition in such a setting typically involves multiple legacy codes in native languages, and requires the ability to adjust to changing data formats, data sources (e.g. user-supplied, accessed through Grid information services, streamed from another module etc.). Cognitive discordance is a serious issue here, one that is impossible to address by committing to a standard vocabulary for implementing components. Such messiness should be viewed not as a limiting bottleneck, but a fundamental aspect of how multidisciplinary research is conducted.

- Sophistication of simulation management.** Traditional GCEs make a simple-minded distinction between the representation of a component and its implementation, suitable for execution on the Grid. Representation is usually intended to imply naming conventions and association of features (e.g. ‘is it gcc-2.7.2 compliant?’) to help in execution. Once again, this has not proved a serious constraint since Grid services have traditionally focused more on executing computations (single runs) and less on high-level problem solving. The sophistication of simulation management is directly related to the representational adequacy of components in the GCE. For situations such as described in *Scenarios 2 and 3*, the scientist would like to say ‘Conduct the same simulation as done on Friday, but update the response surface modeling to use the new numbers collected by Mark’. Or perhaps, ‘collect data from parameterized sweeps of all performance models of the Sweep3D code where the MPI simulation fragment occupies no more than 30% of the total time’. Simulation management can be viewed as a facility for both high-level specification of runs as well as a way to seamlessly *mix* computations and retrievals from a database of previously conducted simulations (see *Scenario 2*). This implies that data management facilities should be provided not as a separate layer of service, but as a fundamental mode by which the simulation environment in a GCE can be managed. The recent NSF-ITR-funded GriPhyN project [29] and the Sequoia data management system [30], both multidisciplinary endeavors, are motivated by similar goals. Simulation management also serves as a way of documenting the ‘history’ of computational runs and experiments. For example, in conducting parameterized sweeps [31], knowing that certain particular choices have been executed elsewhere on the Grid allows flexibility in load balancing and farming out computations to distributed resources.

1.4. GCEs for multidisciplinary Grid communities: a high-level architecture

Finally, by way of introduction, we present a high-level architecture or design framework for organizing and building GCEs for multidisciplinary Grid communities (see Figure 2). We believe that programming capabilities improve by recognizing modeling assumptions and explicitly factoring them out in a system design architecture. Figure 2 does not describe an architecture in the full sense of the word, e.g. with precisely defined interfaces between layers. However, it does separate out the various functions or modes that must be represented in a powerful and effective multidisciplinary community

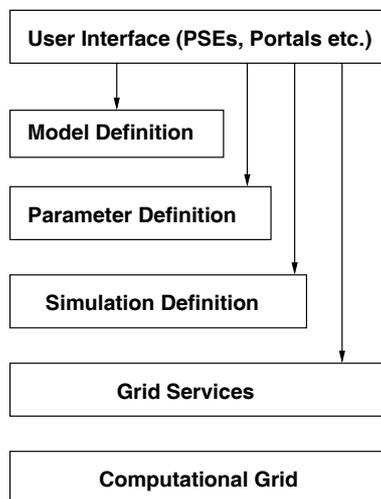


Figure 2. Layers of functionality needed to support multidisciplinary Grid communities.

GCE. The functional framework of the Grid summarized in Figure 2 is complementary to ones that are based on protocol layering (see [6]) and commodity computing (see [32]).

Model. A model is a directed graph of specific executable pieces defining the control-flow and data-flow in a computation, e.g. the digraph in Figure 1 (left). We distinguish between a model and its representation in a GCE; the representation might involve just the model's name or it might involve opening up the boxes (nodes in the digraph) and representing them in a more sophisticated fashion. Although models consist of ready-to-run pieces of code, these pieces may be parameterized.

Model instance. A model instance is a model with all parameters specified. Note that some of these parameters may not be specified until runtime. Thus, while there might not exist a static conversion from models to model instances, the distinction between model instances and models is still useful. For example, using two different input data sets with the same model corresponds to two different model instances and a parameter sweep tool can be used to generate such model instances.

Simulation. A simulation is a model instance assigned to and run on a particular computational resource on the Grid. It is useful to distinguish between a model instance and a simulation because, for example, a single model instance can be run (and re-run) many times using different computational resources or different random number sequences; each of these would be a new simulation by our conventions.



Given these definitions, the framework summarized in Figure 2 can be used to organize the various functions which should be supported in a GCE for a typical multidisciplinary Grid community. The *model definition* layer is where users who need to create or modify models find tools to support this activity. Users who simply use existing models require only trivial support from this layer. In the *parameter definition* layer we locate those activities that associate model instances with models. Examples include tools that generate parameter sweeps [31] or other types of model instance ensembles, as well as the use of problem-oriented scripting languages to generate multiple model instances. (Note that we are using ‘parameter’ in a very broad sense here, making no specific assumptions about exactly how these parameters are defined or what they include.) Another activity that is naturally found at the parameter definition level is a ‘database query’ mode, in which results from previous simulations are accessed, perhaps instead of doing new computations. The next layer, *simulation definition*, is where a model instance (or set of model instances) is assigned to Grid resources. In the simplest case, a user simply chooses some subset of available Grid resources to which the model instance should be mapped. More interesting, however, are the possibilities for simulation-management tools which take a set of model instances and assign them to the Grid, perhaps with sophisticated load balancing strategies or leveraging performance summaries from previous simulations. The lowest two levels appearing in Figure 2, *Grid services* and *computational Grid*, correspond to the software and hardware resources (e.g. Globus, networks, machines) that make computational Grids possible. As mentioned earlier, our emphasis has been on high-level, application-specific issues. We omit further discussion of the architecture, protocols, and services being developed elsewhere for these levels (e.g. see [6]).

Note that not all services or activities fit neatly into the categories shown. For example, in computational steering [28], model parameters may be modified and computational resources re-assigned at runtime; so the parameter and simulation definition services are interleaved with execution in this setting. Other important aspects of an effective GCE are not explicitly represented in Figure 2. For example, support for collaboration is implicit throughout. However, this high-level view of required layers of functionality helps organize and orthogonalize our efforts.

In keeping with the typical end-to-end design philosophy of the Grid [1], we have attempted to provide support for these new services as layers of abstraction over traditional low-level Grid scheduling and resource management facilities. In addition, our resulting high-level architecture ‘teases out’ typically blurred layers into distinct levels at which various services can be provided. Three of our specific contributions to this architecture include: (i) a lightweight data management system that supports compositional modeling (at the model definition level), helps view experiment evaluation as querying (at the parameter definition level), and provides bindings and semistructured representations (for all levels); (ii) a collaborative component composition workspace (*Sieve*) for model definition; and (iii) a framework for distributed resource control (*Symphony*) that provides core support for parameter and simulation definition and management. We describe these efforts in more detail in Section 3.

2. MOTIVATING APPLICATIONS

This section briefly describes five PSEs that are variously situated along the Grid community characteristic axes. These examples highlight the diversity of multidisciplinary communities that a unifying GCE architecture must support.



2.1. WBCSim

WBCSim is a prototype PSE that is intended to increase the productivity of wood scientists conducting research on wood-based composite materials, by making legacy file-based FORTRAN programs, which solve scientific problems in the wood-based composites domain, widely accessible and easy to use. WBCSim currently provides Internet access to command-line driven simulations developed by the Wood-Based Composites (WBC) Program at Virginia Tech. WBCSim leverages the accessibility of the Web to make the simulations with legacy code available to scientists and engineers away from their laboratories. WBCSim integrates simulation codes with a graphical front end, an optimization tool, and a visualization tool. The system converts output from the simulations to the virtual reality modeling language (VRML) for visualizing simulation results. WBCSim has two design objectives: (1) to increase the productivity of the WBC research group by improving their software environment; and (2) to serve as a prototype for the design, construction, and evaluation of larger scale PSEs. The simulation codes used as test cases are written in FORTRAN 77 and have limited user interaction. All the data communication is done with specially formatted files, which makes the codes difficult to use. WBCSim hides all this behind a server and allows users to supply the input data graphically, execute the simulation remotely, and view the results in both textual and graphical formats.

WBCSim contains four simulation models of interest to scientists studying wood-based composite materials manufacturing—rotary dryer simulation (RDS), radio-frequency pressing (RFP), composite material analysis (CMA), and particle mat formation (MAT). The rotary dryer simulation model was developed as a tool to assist in the design of drying systems for wood particles, such as used in the manufacture of particleboard and strandboard products. The rotary dryer is used in about 90% of these processes. The radio-frequency pressing model was developed to simulate the consolidation of wood veneer into a laminated composite, where the energy needed for cure of the adhesive is supplied by a high-frequency electric field. The composite material analysis model was developed to assess the strength properties of laminated fiber-reinforced materials, such as plywood. The mat formation model is used to calculate material properties of wood composites, modeling the mat formation process as wood flakes are deposited and then compressed into a mat. This model is crucial for all other manufacturing process models, as they require material properties as input.

The software architecture for WBCSim is three-tiered: (i) the legacy simulations and various visualization and optimization tools, perhaps running on remote computers; (ii) the user interface; and (iii) the middleware that coordinates requests from the user to the legacy simulations and tools, and the resulting output. These three tiers are referred to as the developer layer, the client layer, and the server layer, respectively. The developer layer consists primarily of the legacy codes on which WBCSim is based. The server layer expects a program in the developer layer to communicate its data (input and output) in a certain format. Thus, legacy programs are ‘wrapped’ with custom Perl scripts, and each legacy program must have its own wrapper. The client layer consists of Java applets and is responsible for the user interface (see Figure 3). It also handles communication with the server layer, is the only layer that is visible to end-users, and typically will be the only layer running on the user’s local machine. The server layer is the core of WBCSim as a system distinct from its legacy code simulations and associated data viewers. The server layer is responsible for managing execution of the simulations and for communicating with the user interface contained in the client layer. WBCSim applications require sophisticated management of the execution environment; the server layer, written in Java, directs execution of multiple simulations, accepts multiple requests from clients concurrently,

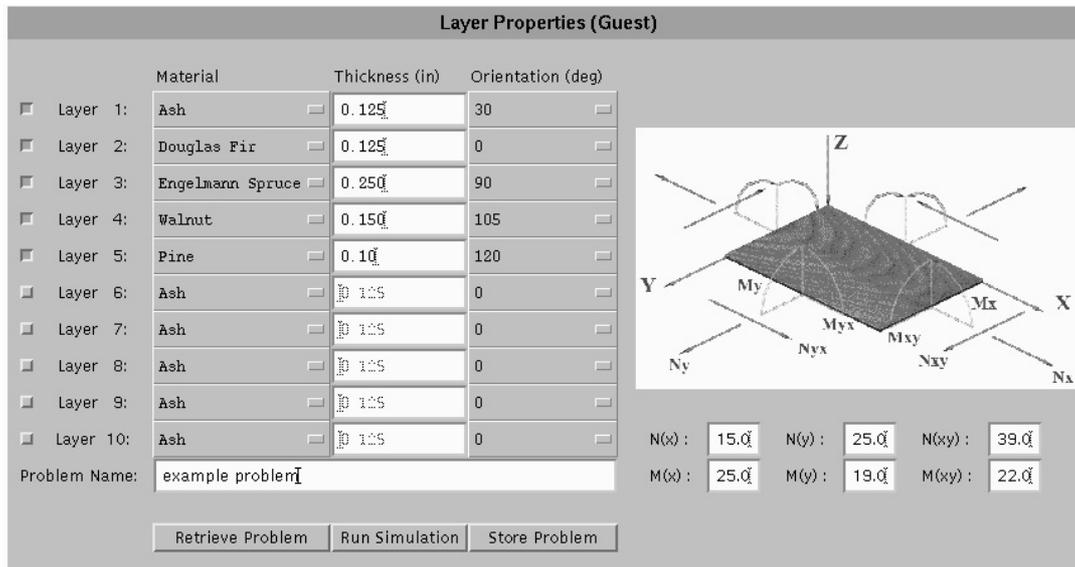


Figure 3. Interface to the CMA model in the WBCSim PSE [33].

and captures and processes messages that signify major milestones in the execution (such as the computation of an intermediate value). Graphical results from the simulations are communicated to the clients using an HTTP server.

2.2. VizCraft

VizCraft [13] is a PSE that aids aircraft designers during the conceptual design stage. At this stage, an aircraft design is defined by a vector of 10 to 30 parameters. The goal is to find a vector that minimizes a performance-based objective function while meeting a series of constraints. VizCraft integrates simulation codes to evaluate a design with visualizations for analyzing a design individually or in contrast to other designs. VizCraft allows the designer to easily switch between the view of a design in the form of a parameter set, and a visualization of the corresponding aircraft geometry. The user can easily see which, if any, constraints are violated. VizCraft also allows the user to view a database of designs using the parallel coordinates visualization technique. VizCraft is a design tool for the conceptual phase of aircraft design whose goal is to provide an environment in which visualization and computation are combined. The designer is encouraged to think in terms of the overall task of solving a problem, not simply using the visualization to view the results of the computation.

VizCraft provides a menu-driven graphical user interface to the high speed civil transport (HSCT) design code that uses 29 variables and 68 realistic constraints. This code is a large (million line) collection of C and FORTRAN routines that calculate the aircraft geometry in 3D, the design constraint

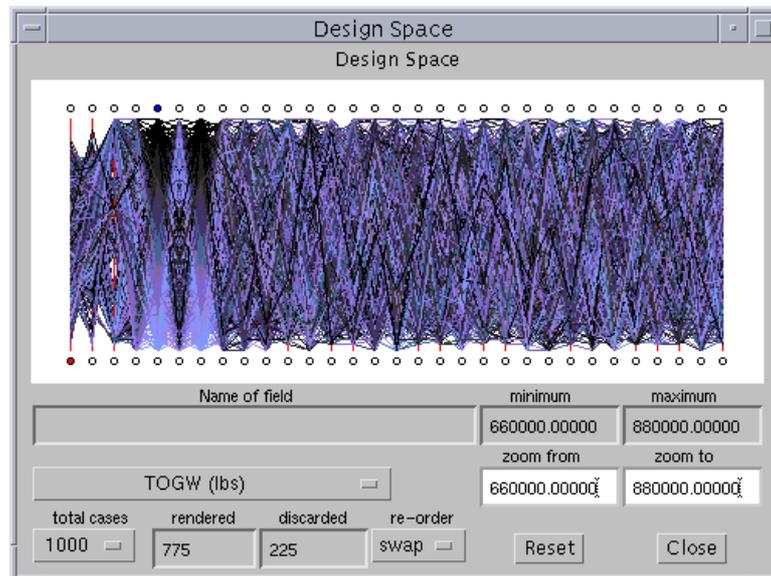


Figure 4. Visualizing 156 aircraft design points in 29 dimensions in a VizCraft [13] design view window. A careful assignment of variables to color drivers reveals an interesting association.

values, and the take-off gross weight (TOGW) value, among other things. VizCraft displays the HSCT platform (a top view), cross sections of the airfoil at the root, leading edge break, and tip of the wing, and color coded (red, yellow, green) constraint violation information. To help manage the large number of constraints, they are grouped conceptually as aerodynamic, geometric, and performance constraints. Design points, and their corresponding TOGW, are displayed via active parallel coordinates. The parallel coordinates are also color coded, and they can be individually scaled, reordered, brushed, zoomed, and colored. A parallel coordinate display for the constraints can be similarly manipulated. While the integration of the legacy multidisciplinary HSCT code into a PSE is non-trivial, the strength and uniqueness of VizCraft lie in its support for visualization of high dimensional data (see Figure 4).

2.3. L2W

Landscapes to Waterscapes (L2W) is a PSE for landuse change analysis and watershed management. L2W organizes and unifies the diverse collection of software typically associated with ecosystem models (hydrological, economic, and biological), providing a Web-based interface for potential watershed managers and other users to explore meaningful alternative land development and management scenarios and view their hydrological, ecological, and economic impacts. Watershed management is a broad concept entailing the plans, policies, and activities used to control water and related resources and processes in a given watershed. The fundamental drivers of change are



modifications to landuse and settlement patterns, which affect surface and ground waterflows, water quality, wildlife habitat, economic value of the land and infrastructure (directly due to the change itself such as building a housing development, and indirectly due to the effects of the change, such as increased flooding), and cause economic effects on municipalities (taxes raised versus services provided). The ambitious goal of L2W is to model the effects of landuse and settlement changes by, at a minimum, integrating codes/procedures related to surface and subsurface hydrology, economics, and biology. The development of L2W raises issues far beyond the technical software details, since the cognitive discordance between computer scientists (developing the PSE), civil engineers (surface and subsurface hydrology), economists (land value, taxes, public services), and biologists (water quality, wildlife habitat, species survival) is enormous. The disparity between scientific paradigms in a multidisciplinary engineering design project involving, say, fluid dynamicists, structural, and control engineers is not nearly as significant as that between computer scientists, civil engineers, economists, and biologists. A further compounding factor is that L2W should also be usable by governmental planners and public officials, yet another different set of users.

The architecture of the L2W PSE is based on leveraging existing software tools for hydrology, economic, and biological models into one integrated system. Geographic information system (GIS) data and techniques merge both the hydrologic and economic models with an intuitive Web-based user interface. Incorporation of the GIS techniques into the PSE produces a more realistic, site-specific application where a user can create a landuse change scenario based on local spatial characteristics (see Figure 5). Another advantage of using a GIS with the PSE is that the GIS can obtain necessary parameters for hydrologic and other modeling processes through analysis of terrain, land cover, and other features. Of all the PSEs described here, L2W is unique in that it is centered around a GIS. Currently, L2W integrates surface hydrology codes and economic models for assessing the effect of introducing settlement patterns. Wildlife and fisheries biologists were involved in the L2W project, but their data and models are not fully integrated as of this writing. The biological models include the effect of development on riparian vegetation, water quality, and fish and wildlife species.

2.4. S⁴W

Site-Specific System Simulator for Wireless communications (S⁴W) is a collaborative PSE for the design and analysis of wideband wireless communications systems. In contrast to the above described projects, the S⁴W project is occurring in parallel with the development of high-fidelity propagation and channel models; this poses a unique set of requirements for software system design and implementation (ref. *Scenario 1* in the introduction) [35]. S⁴W has the ability to import a 3D database representing a specific site (see Figure 6) and permits a wide range of radio propagation models to be used for practical communications scenarios [12]. For example, in a commercial wireless deployment, there is a need to budget resources, such as radio channel assignments and the number of transmitters. S⁴W allows wireless engineers to automatically drive the simulation models to maximize coverage or capacity, or to minimize cost. Furthermore, unlike existing tools, S⁴W permits the user to import measured radio data from the field, and to use this data to improve the models used in the simulation. A knowledge-based recommender system [36] provides improved modeling capability as the software corrects the environment model and the parameters in the propagation model, based on measured data. Finally, the ability to optimize the location of particular wireless portals in an arbitrary environment is

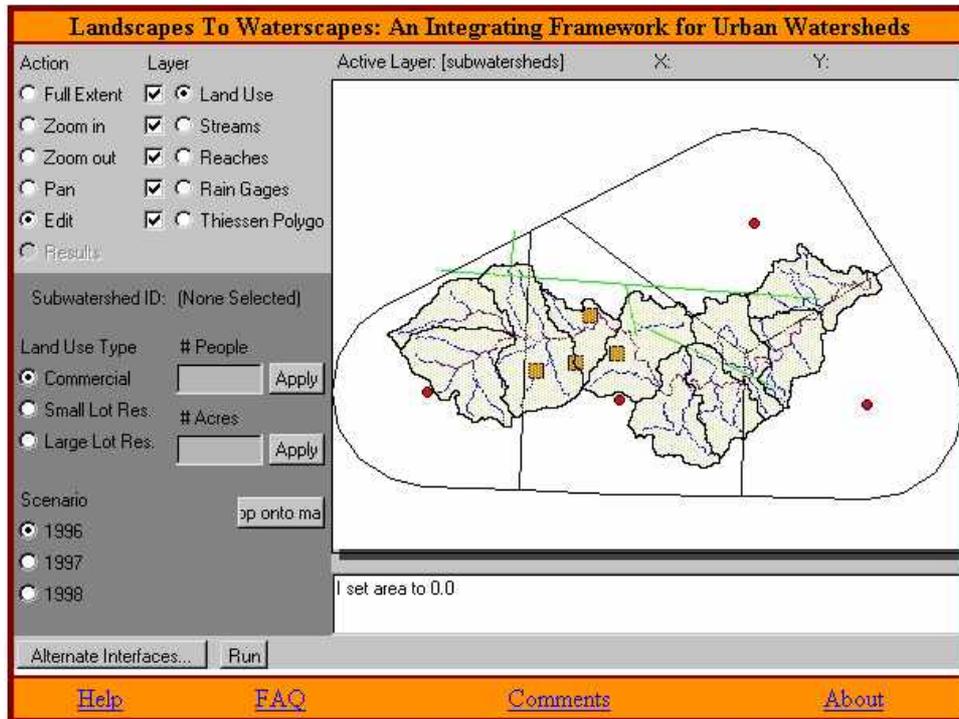


Figure 5. Front-end decision maker interface to the L2W PSE [34], depicting landuse segmentation of the Upper Roanoke River Watershed in Southwest Virginia.

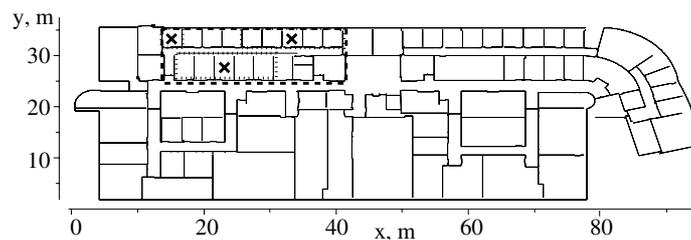


Figure 6. An example indoor environment in S⁴W showing locations of three transmitters optimized to cover eighteen rooms and a corridor.



a fundamental breakthrough for wireless deployment, and S⁴W has the ability to perform optimization based on a criterion of coverage, QoS, or cost.

While primitive software tools exist for cellular and PCS system design, none of these tools include models adequate to simulate broadband wireless systems, nor do they model the multipath effects due to buildings and other man-made objects. Furthermore, currently available tools do not adequately allow the inclusion of new models into the system, visualization of results produced by the models, integration of optimization loops around the models, validation of models by comparison with field measurements, and management of the results produced by a large series of experiments. One of the major contributions of S⁴W is a lightweight data management subsystem [35] that supports the experiment definition, data acquisition, data analysis, and inference processes in wireless system design. In particular, this facility helps manage the execution environment, binds representations to appropriate implementations in a scientific computing language, and aids in reasoning about models and model instances.

S⁴W is designed to enhance three different kinds of performance—software, product, and designer. Superior software performance is addressed in this project by: (i) developing fundamentally better wireless communication models; (ii) constructing better simulation systems composed from the component wireless models via the recommender; and (iii) the transparent use of parallel high-performance computing hardware via the composition environment's access to distributed resources. Superior product performance (the actual deployed wireless systems) is addressed by using optimization to design *optimal* rather than merely *feasible* systems. Superior designer performance is directly addressed by the synergy resulting from the integrated PSE, whose purpose is to improve designer performance and productivity.

2.5. Espresso

The Espresso project [37] addresses the entire lifecycle of microarray bioinformatics, an area where 'computing tools coupled with sophisticated engineering devices [can] facilitate discovery in specialized areas [such as genetics, environment, and drug design]' [38]. Microarrays (sometimes referred to as *DNA chips*) are a relatively new technique in bioinformatics, inspired by miniaturization trends in micro-electronics. Microarray technology is an experimental approach to study all the genes in a given organism simultaneously; it has rapidly emerged as a major tool of investigation in experimental biology. The basic idea is to 'print' DNA templates (*targets*), for all available genes that can be expressed in a given organism, onto a high-density 2D array in a very small area on a solid surface. The goal then is to determine the genes that are expressed when cells are exposed to experimental conditions, such as drought, stress, or toxic chemicals. To accomplish this, RNA molecules (*probes*) are extracted from the exposed cells and 'transcribed' to form complementary DNA (cDNA) molecules. These molecules are then allowed to bind (*hybridize*) with the targets on the microarray and will adhere only with the locations on the array corresponding to their DNA templates. Typically such cDNA molecules are tagged with fluorescent dyes, so the expression pattern can be readily visualized as an image (see Figure 7). Intensity differences in spots will then correspond to differences in expression levels for particular genes. Using this approach, one can 'measure transcripts from thousands of genes in a single afternoon' [38]. Microarrays thus constitute an approach of great economic and scientific importance, one whose methodologies are continually evolving to achieve higher value and to fit new uses.

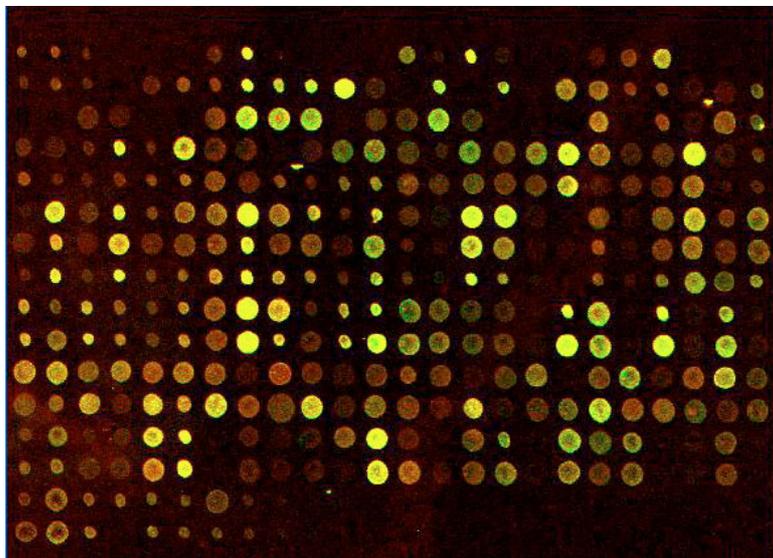


Figure 7. An example microarray image in Expresso [37] to study gene expression in Loblolly pine clones. The microarray is printed in four 24×16 sub-quadrants, one of which is shown here. Figure courtesy Y.-H. Sun (NCSU).

The Expresso PSE [37] is designed to support all microarray activities including experiment design, data acquisition, image processing, statistical analysis, and data mining. Expresso's design incorporates models of biophysical and biochemical processes (to drive experiment management). Sophisticated codes from robotics, physical chemistry, and molecular biology are 'pushed' deeper into the computational pipeline. Once designs for experiments are configured, Expresso continually adapts the various stages of a microarray experiment, monitoring their progress, and using runtime information to make recommendations about the continued execution of various stages. Currently, prototypes of the latter three stages of image processing, statistical analysis, and data mining are completely automated and integrated within our implementation.

Expresso's design underscores the importance of modeling both physical and computational flows through a pipeline to aid in biological model refinement and hypothesis generation. It provides for a constantly changing scenario (in terms of data, schema, and the nature of experiments conducted). The ability to provide expressive and high performance access to objects and streams (for experiment management) with minimal overhead (in terms of traditional database functionality such as transaction processing and integrity maintenance) [39] is thus paramount in Expresso.

The design, analysis, and data mining activities in microarray analysis are strongly interactive and iterative. Expresso thus utilizes a lightweight data model to intelligently 'close the loop' and address both experiment design and data analysis. The system organizes a database of problem instances and simulations dynamically, and uses data mining to better *focus* future experimental runs based on results from similar situations. Expresso also uses inductive logic programming (ILP), a relational data mining



technique, to model interactions among genes and to evaluate and refine hypothesized gene regulatory networks. One complete instance of the many stages in Espresso has been utilized to study gene expression patterns in Loblolly pine [40], in a joint project with the Forest Biotechnology group of North Carolina State University.

3. SYSTEMS SUPPORT FOR MULTIDISCIPLINARY GCE APPLICATIONS

This section describes several core systems support technologies useful for developing GCEs (and currently employed in the applications outlined so far). Many of these tools and frameworks rely on the notion of representations of components; we begin by motivating this idea.

3.1. Representations in a GCE

One of the main research issues in GCEs is modeling the fundamental processes by which knowledge about scientific models is created, validated, and communicated. As mentioned in Section 1 and illustrated in the many example systems of Section 2, the expressiveness with which a scientist could interact with a GCE is directly related to the adequacy of representation provided by the system. While it is true that there is no universal representation that is ideal for all purposes, traditional approaches employed in Grid projects (for representing models, model instances, and simulations) are not sufficient to support high-level problem solving.

Recall that we defined a model to denote a directed graph of specific computational codes or executables. The notion of the ‘representation of a model’ is open to many interpretations and intensely debated in the modeling literature (see for instance [41]); we will not attempt to settle this debate here. Instead, we adopt an operational definition for the representation of a model, namely that it is an abstraction of the model that permits useful problem-solving capabilities that would not be possible with the model alone. The abstraction could refer to the functional behavior of the model (e.g. a signature), the structural constituents of the model (e.g. a digraph of model fragments), a profile of its performance (to aid in design and analysis), its relationships to other models, and/or information about how it fits within the larger computational context of the GCE and the activities conducted within it.

Consider two extremes of representing a single computational component (the simplest model) in a GCE. A black-box representation would be one where just the name of the component serves as its representation. Such a solution will definitely aid in referring to the component over the Grid (e.g. ‘run the XYZ software’), but does not help in any further sophisticated reasoning about the component (e.g. ‘is XYZ an iterative algorithm?’). At the other extreme, a white-box representation is one where the component itself serves as its representation (for example, mathematical descriptions of scientific phenomena). Usually in such cases, the representation is accompanied by a domain theory and algorithms for reasoning in that representation. For instance, the QSIM system [42] is a facility where once a component (e.g. one for the cooling of liquid in a cup [20]) is defined, it is possible to reason about the properties and performance of that component (e.g. when are the contents of the cup drinkable?). While extremely powerful, such designs work well only within restrictive (and sometimes artificial) applications. An intermediate solution is to annotate components with (feature, value) pairs describing attribute-based properties. For instance, annotations might involve directives, flags, and hints for compiling the code on a specific platform.



These issues are amplified when we consider the model to be a digraph of computational components. While many projects distinguish between models and representations, two main approaches can be distinguished here. In the first category (projects such as [28,43–46]), representations are motivated by the need to manage the execution environment (e.g. ‘schedule this graph on the Grid, taking care to ensure that data files are properly matched’). Projects in the second category are based on AI research in compositional modeling [47–49] and are motivated by formal methods for reasoning with (approximate and qualitative) representations. The modeling and performance requirements in a multidisciplinary GCE setting mean that both approaches are too restrictive (see [35]).

With the advent of XML and the emergence of the Web as a large-scale semistructured data source, interest in semistructured representations has expanded into the GCE/PSE community. A plethora of markup languages, XML-based formats, and OO coding templates have been proposed for representing aspects of domain-specific scientific codes (e.g. SIDL [50]). In addition, a variety of formats have been proposed recently (e.g. SOX [51]) for defining metadata associated with various stages of computational experiments. A major advantage of such solutions is that the ensuing representations can be cataloged and managed by database technology.

Our goal is to investigate representations that: (i) allow the binding of problem specifications to models (and model instances), without insisting on an implementation vocabulary (for the models); (ii) can help us to reason both about the models being executed as well as data produced from such simulations; and (iii) help design facilities such as change management, high-level specification of simulations, recommendation, optimization, and reasoning (about models and model instances).

3.2. BSML: a binding schema markup language

Akin to other GCE projects, our emphasis here will be on semistructured representations for models. However we view markup languages such as XML as less of a data format, programming convention, or even a high-level abstraction of a programming environment. Rather, we view them as a vehicle to define *bindings* from representations to models in a GCE. Binding refers to the process of converting XML data to an appropriate encoding in a scientific computing language (the reverse process is fairly straightforward). There are several forms of bindings in a GCE—binding of values to language variables, converting an XML format to some native format that can be read directly by the model, and/or generating source code for a stub that contains embedded data and a call to the appropriate language function using these data as parameters. Notice that we do not make a distinction between invoking a component procedurally in a scientific computing language, generating code that invokes a component, or executing a program with command line arguments. All of these are bindings from one representation to various assumptions on the execution environment (which is presumably being handled by the existing computational setup). Our lack of any stringent assumptions on the computational codes or method of invoking models is fundamental to multidisciplinary research. From the viewpoint of the GCE, a single representation could be stored but which can allow all these forms of bindings to be performed.

A full description of our binding schema markup language (BSML) is beyond the scope of this article (for more details, see [35]). We briefly mention that BSML associates user-specified blocks of code with user-specified blocks of an XML file. ‘Blocks’ can be primitive datatypes, sequences, selections, and/or repetitions. Intuitively, primitive datatypes denote single values, such as double



```

<element name='pdp' >
  <sequence>
    <element name='rmsDelaySpread'
              type='double' />
    <element name='meanExcessDelay'
              type='double' />
    <element name='peakPower'
              type='double' />
    <code component="optimizer">
      <bind>print "$peakPower\n"</bind>
    </code>
  </sequence>
  <repetition>
    <sequence>
      <element name='time' type='double' />
      <element name='power' type='double' />
      <code component="chttt1|chttm">
        <bind>print " $time $power\n"</bind>
      </code>
    </sequence>
  </repetition>
  <code component="chttt1|chttm">
    <begin>print "M = [\n"</begin>
    <end>print "];\n"</end>
  </code>
</element>

```

Figure 8. BSML descriptions for a class of XML documents pertaining to power delay profiles (PDPs) in the S⁴W PSE. Sample bindings for MATLAB are shown by the bind tags.

precision numbers; sequences denote structures; selections denote multiple choices of conveying the same information; and repetitions denote lists. While not particularly expressive, this notation is meaningful to GCE component developers, simple and efficient to implement, and general enough to allow the building of more complex data representations.

Consider, for example, representing a power delay profile (PDP) from the S⁴W project in XML. A PDP is a two-column table that describes the power received at a particular location during a specified time interval. Statistical aggregates derived from power delay profiles are used, for example, to optimize transmitter placement in S⁴W. We can use BSML to define bindings between PDPs and all applicable models in S⁴W. Applying a parser generated from such a BSML document (see Figure 8 for an example) to a PDP will yield native code in the underlying execution environment (in this case, an executable Matlab script that contains embedded data). For a different native execution environment, a different binding could be defined (for the same data type). Hence, our representation is truly semistructured.

Notice that we can rapidly prototype new model instances with this technique. Similarly, we can use the same BSML source to provide bindings for an optimizer (for locating base stations, see *Scenario 1* in Section 1). The feedback will be a sequence of peak powers, one number per line. Some twenty

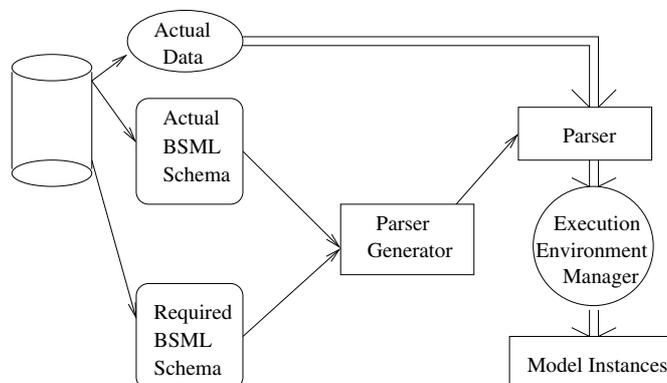


Figure 9. A facility for creating model instances from specifications represented in a binding schema markup language (BSML).

five lines of BSML source can therefore take care of data interchange problems for three components. Storing these PDPs in a database is also facilitated. From a system point of view, the schemas are the metadata and the software that translates and converts schemas is the parser generator. Figure 9 shows a typical configuration.

Both the data and the metadata are stored in a database. A parser is lazily generated for each combination of model's input port and the actual BSML schema whose data instance is connected to that port. Model descriptions can also be stored in the database. They consist of model id, description, schemas of its input and output ports, various execution parameters, and annotations, such as relations to other models. We do not provide any tools for checking the consistency of the output data with the output schemas because, unlike in Web or business domains, this is rarely an issue.

The execution environment manager (see Figure 9) glues the generated parsers to the components. For full-featured languages like FORTRAN, it will simply link the parser with the model's code. Prototyping for languages like Matlab requires more attention. The output of the parser for such languages is typically the source code that needs to be supplied to the interpreter. The exact way the parsers are linked to the model is specified by the model's execution parameters. From this point, each model together with a set of generated parsers looks like a program that takes a number of XML streams as inputs and produces a number of XML streams as outputs. This is an appropriate representation for the management of the execution environment. Our goal is similar to those in [52–54] in that the common motivation is management of the execution environment; at the same time, our concern with high-level problem-solving (see next three sections) shifts the emphasis from a unifying programming environment to one that allows a data-driven approach to managing large-scale multidisciplinary codes.

Finally, it is relatively straightforward to store any resulting data from such processes in a database system. If an RDBMS is used, we can reuse BSML to generate a number of SQL update statements in the same manner we used it to generate a Matlab script in Figure 8. One of these 'models' will then



connect to the database and execute these statements. This is no different from other format conversions happening in the GCE.

3.3. Format conversions and change management

One of the benefits of semistructured data representations is automatic format conversion. This feature is useful in the following situations: (i) a model is changed over time, but data corresponding to the older versions has already been recorded in the database system. An example from S⁴W is the evolution of the space partitioning parameters in the ray tracer. After we have realized that placing polygons at the internal nodes of the octree can improve space usage by an order of magnitude, more parameters have been added to space partitioning; (ii) several components need essentially the same parameters, but are not truly plug-and-play interchangeable. Minor massaging is necessary in order to make their I/O specifications match.

We model the following changes: insertions, deletions, replacements, and unit conversions. Insertions and deletions correspond to additions and removals of parameters. For example, a moving channel builder takes the same inputs as a static one, plus the velocity of the receiver. Thus, any input to a moving channel builder can be converted to the input to a static one by projecting out the receiver's velocity. Replacements represent changes in parameter representation, such as a conversion between spherical and rectangular coordinates. Unit conversions are a special case of conversions that are quite common and can be easily automated, such as conversions between watts and decibel milliwatts. Unit conversion can be performed by equation and constraint solvers [20].

In our XML representation, insertions can be handled by requiring default values for new parameters. Removals amount to deleting the old values. Replacements and unit conversions require user-supplied or automatically generated conversion filters. The modeling literature abounds in such conversions, but it is important to realize that conversion facilities are *ad hoc* by nature, and therefore only work for small changes in the schema. Typically, it is not necessary to find a globally optimal conversion sequence. A thorough treatment of change detection can be found in [55].

Change management can also be used to realize any problem-solving feature that involves transforming between semistructured representations. For example, consider the possibility that two students configure a GCE independently with different choices for various stages in a computational pipeline and arrive at contradictory results. They could then query the database for 'What is different between the experiments that produced data in directory A from the ones in directory B'?—providing responses such as 'A calibration threshold of 0.84 was used in B instead of 0.96 for A,' which are obtained by automatically analyzing the XML descriptions [56]. Change detection and processing is crucial in GCE projects such as Espresso, where objects of interest change formats, stations, and schema rapidly.

3.4. Executing simulations = querying

Recall that we defined a simulation as a model instantiated with inputs, along with an assigned computational resource. This captures the notion of applying multiple models to multiple inputs to generate a database of simulation results and performance data. In this section, we describe how our semistructured representations of models and bindings can aid in even higher level problem-solving facilities. In particular, we concentrate on facilities such as the 'parameter sweep' tool [31] and the



```
<experiment id='diff. prop.'>
  WHERE <environment id='$id'>
    <meta><type>urban</type></meta>
  </environment> CONTENT_AS $env IN "envs"
  CONSTRUCT <experiment id='diff. prop.: $id'>
    <model>...</model>
    <inputs>
      <input>$env</input>
      ...
    </inputs>
    <outputs>...</outputs>
  </experiment>
</experiment>
```

Figure 10. Constructing new XML data (to recalculate PDPs in S⁴W) using the XML-QL query notation. WHERE...CONSTRUCT is the format for expressing queries in this language. Notice that the query is parameterized by the \$env variable whose type is restricted to be urban.

‘database query’ mode (found in the parameter definition layer of Figure 2). The facilities described in this section: (i) produce model instances; and also (ii) associate data generated from simulations back with the corresponding model instances. We gloss over the aspect of how simulations corresponding to a model instance are actually executed, since they are addressed later. In particular, when we refer to executing a simulation we imply that some assignment of computational resources to model instances has been done by the simulation definition layer (see Figure 2).

In the database paradigm, a model instance can be represented as a view. Executing the simulation corresponds to materializing the view. The query behind the view is a join over models and data. To be meaningful, a simulation must further satisfy some *syntactic* and/or *semantic* constraints. Syntactic constraints ensure that the simulation can indeed be executed. Each simulation run must be given enough data and the data must conform to the appropriate schemas. Semantic constraints ensure that the models are meaningful in the specific problem domain. We will describe semantic constraints in the next section. In our framework, users can impose custom constraints, such as ‘use only the datasets from last week’. Specifying a model instance therefore maps naturally into a database query. This feature also supports the provision of iteration, aggregation, and composition operators by introducing minimal overhead in implementation. For example, compositions can be achieved by relational joins and aggregation by user-defined VIEWS.

Consider the following scenario. An S⁴W developer of ray tracing propagation models has added a model that takes diffraction into account. She now wants to recalculate the PDPs for the environments where diffraction is most significant, e.g. for urban outdoor environments. Her request for new simulations can be specified in a query notation such as XML-QL (see Figure 10 for how this can be done). The result of this query is a sequence of model instances, which can be associated with corresponding simulations and scheduled for execution. Not only is this form of specification concise, it also enables us to use well-known query optimization techniques to push costly operations ‘deeper’ into the computational pipeline [57]. In particular, query-based representations of model instances lose



the distinction between *conducting a simulation to collect data* and *looking up (already simulated) data from a database*. Coupled with Grid information services, such a representation can help determine if specified simulations have been conducted elsewhere on the Grid, to avoid duplication of effort. Furthermore, ‘nearby’ simulations can be retrieved to construct a surrogate function for the entire simulation or parts of it, thereby replacing costly executions with cheap surrogate function evaluations. In large scale, multidisciplinary, engineering design the construction and use of such surrogates has become standard practice. The query approach facilitates the automatic construction and validation of surrogates or functional approximations.

In our current implementation the above two modes—simulation by executing a code and simulation by querying—are provided as distinct services. Before we can seamlessly mix computations with retrievals from a database, a logic for such a facility has to be defined. For example, do we adopt a policy that looks for cached results before spawning out simulations? Or is it a best-search effort for doing costly operations on external clusters and conducting smaller simulations locally? Or does it use different fidelity approximations for different ranges of the parameter sweep? The design of these policies (along with their associated business and organizational ramifications) will influence the acceptability of the ‘execution as querying’ viewpoint.

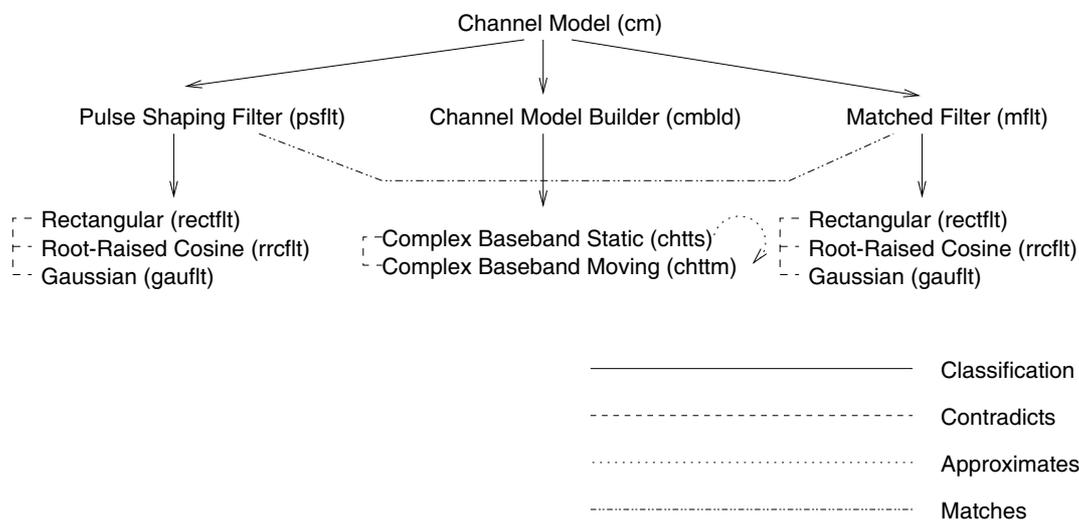
3.5. Reasoning and problem solving

What constitutes a good model? GCEs should provide the facility to reason about a model and its constituent parts in terms of the features of the problem being solved and the desired performance constraints. For instance, this might help in recommending models for new problem instances, based on representational goals and performance criteria. A lot of domain-specific knowledge [48] is required to arrive at promising model choices, but a few general rules can be outlined.

First, a model must not contain any components that make incompatible assumptions about the phenomena being modeled. Following Nayak [48], we call such components *contradictory*. An example of contradictory components in S^4W is a class of model builders (static and moving). Second, some modeling choices may constrain the form of the rest of the model. For example, the signal filters of the transmitter and the receiver must *match*. And finally, models in a given class, say filters, often support similar forms of reasoning. We use the term *classification* to describe this aspect.

An example of these relations in S^4W is given in Figure 11. The labels represent a small model library and the links represent the relations. Note that these relations are domain specific and cannot be derived from the source in any general-purpose language. They must be supplied by the user (in this case, wireless system designer) as annotations to components. Such relations can then be used to prune the search space for recommendation and problem solving.

An alternate, data-driven approach to reasoning is described in [58–60], where we have explored the knowledge-based selection of solution components for individual application domains. Such recommender systems’ [61] can help in the natural process of a scientist/engineer making selections among various choices of algorithms/software in a GCE. They are typically designed off-line by organizing a battery of benchmark problems and algorithm executions [58], and mining it to obtain high-level rules that can form the basis of a recommendation. Visualization of performance data from scientific applications is also related to this aspect [62]. The importance of data mining in providing decision support in large-scale simulations [63] has also been recognized in other projects [64–69]. Our framework is novel in that it captures the entire problem-solving process prior to the scheduling

Figure 11. Relations between channel modeling components in S⁴W.

of a simulation and the semistructured format for model instances allows the embedding of mining functions as primitives into standard query languages (akin to [70]). It also facilitates the incorporation of performance models of the various parts of a computation [71] in reasoning and recommendation.

3.6. Sieve: a collaborative component composition workspace

Readers familiar with components and distributed internet-based applications will recognize that many goals of GCEs described earlier apply to other distributed applications. While the details differ, describing a computation as a network of components and providing access to a database (in this case the database of experimental runs) are not unique to computational science. While supporting legacy code is often central to computational science applications, this need is also not novel.

However, the combination of issues embodied in multidisciplinary GCE settings presents novel problems. These include the fact that individual runs of a simulation can take hours; the extensive and integrated use of visualization; the inherently distributed nature of the computation (i.e. certain sub-models may need to run on differing systems for reasons related to resource needs, or simply because they are legacy codes written for differing systems); the desire for synchronous collaboration; and the needs of multidisciplinary users, no one of whom is an expert in all aspects of the larger system (recall the L2W PSE).

We embody the model definition activities of a GCE in a (collaborative) visual workspace, in which the user places various objects. These objects are components that represent individual codes, optimization tools, visualization software, etc. These components are linked together by the user to form networks that indicate the flow of data or control. The links between components are often



represented by arrows. For example, a component representing an input file on some computer (with the name of the file and computer specified as parameters to the component) might be linked by an arrow to another component representing a model/optimizer combination. Another arrow links the model/optimizer combination to a visualizer. The intent is that the GCE will cause the input file to be moved to the machine storing the model and optimizer, and the model/optimizer will then be invoked. The output of this process will then be passed to the visualization (perhaps on another machine), with the results displayed on the user's screen. The fundamental interface design is similar to that of a modular visualization environment (MVE) or the Khoros image processing system [72].

Our primary tool for creating interfaces for model building is named Sieve. Sieve provides a collaborative workspace that supports the deployment of JavaBeans. At its heart, Sieve is a collaborative 'Java Beanbox' whose library APIs provide support for creating a variety of component-oriented applications. In particular, Sieve provides to programmers convenient facilities for defining not only the component themselves, but also for defining the actions performed by the links between components, along with specifying the format of the data that flows along those links. Users enter a collaborative session simply by specifying at the start which collaborative session to join (or initiate). Together with Symphony (see next section), Sieve provides the foundation for constructing PSEs, as their combination creates a collaborative environment for controlling distributed, legacy resources.

Sieve provides an environment for collaborative component composition that supports the following:

- a Java-based system compatible with standard WWW browsers;
- a convenient environment for generating visualizations through linking of data-producing modules with data-visualization modules;
- collaboration between users through a shared workspace, permitting all users to see the same visualizations at the same time;
- support for annotating the common workspace, visible to all users;
- a convenient mechanism for linking in new types of components.

Sieve presents the user with a large, scrollable workspace onto which data sources, processing modules, and visualization components may be dropped, linked, and edited. Figure 12 shows a Sieve workspace containing a simple data-flow network. This example illustrates a particular collection of beans that support various statistical visualizations.

Our design for Sieve allows processing and visualization modules to be generic, with all data-source-specific details hidden by the source modules. In addition, data-flow semantics can be tailored to the individual application characteristics and available Grid services. For instance, the components of Figure 12 conform to an API which allows data to be viewed by adjacent modules in the network as a two-dimensional table containing objects of any type supported by the Java language. For other applications, the components could invoke Grid services and file transfer utilities to effect the data-flow. Source modules simply convert raw data into the assumed data-flow representation. Processing modules can manipulate these data and present an altered or extended view. Visualization modules can then produce visual representations of the data in the system and also serve as interfaces for data selection.

Sieve's support for defining link types permits a great deal of flexibility in combining various collections of components. For example, if a component had been created that did not support the table format just described for our current visualization components, all that would be needed is for the programmer to create a link class that converts between the two data formats in question. The Sieve

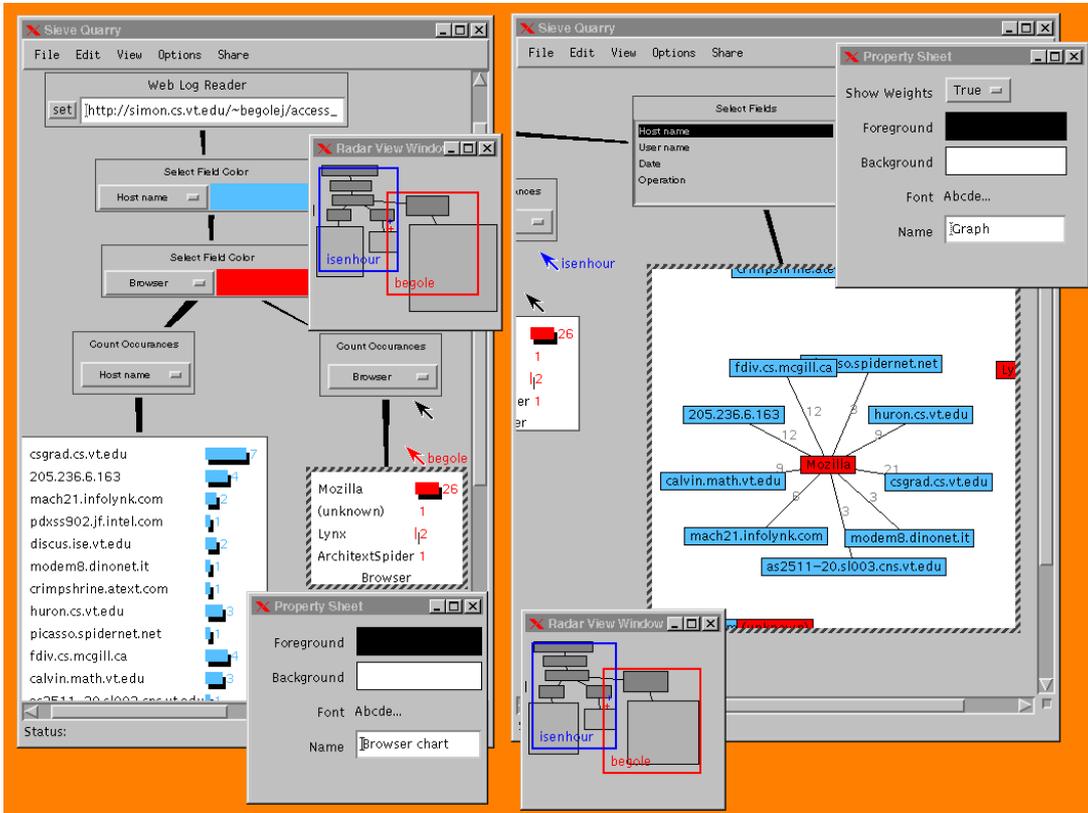


Figure 12. Example of a Sieve workspace with dataflow and annotations.

runtime environment will automatically deduce the appropriate link class to use when two components are linked together. Thus, if the new component (with its unique data format) were the source for the link, and one of the existing (table driven) visualization components were the target, then Sieve would determine that the link class to instantiate is the one that performs the appropriate data conversion between the two components.

3.7. Symphony: managing remote legacy resources

Symphony is a component-based framework for composing, saving, sharing, and executing simulations [73]. A simulation is a collection of programs and data resources intended to be executed as a unit. A simulation is composed by assembling individual Symphony components into an acyclic directed graph and thus corresponds to a representation of a model or model instance. The directed

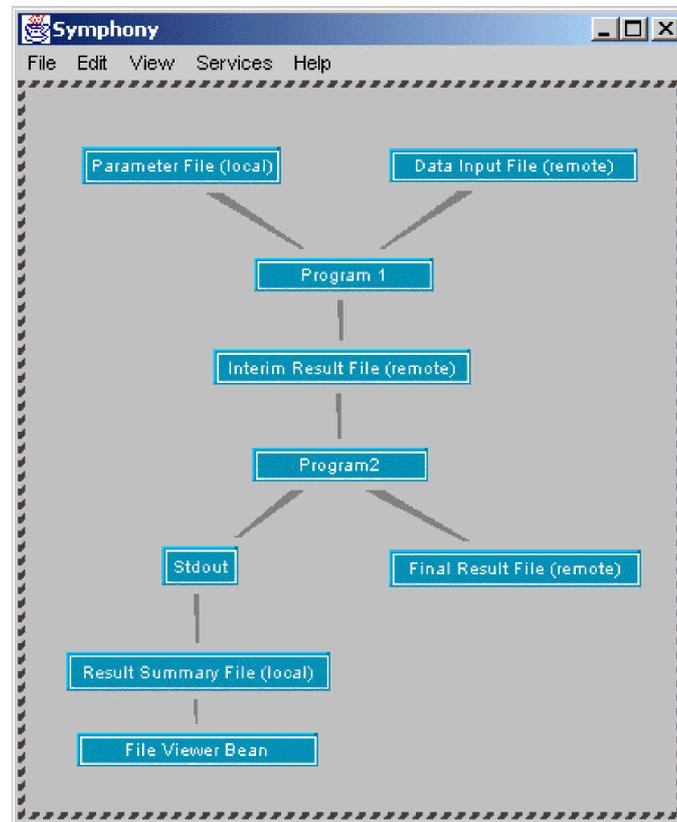


Figure 13. Design of a Symphony Bean Network.

graph models the data flow relationships among the individual components. Such a graph is shown in Figure 13 and described later. Each Symphony component is a surrogate for some actual program or data resource. The actual resources may, and usually do, reside on different machines that are able to communicate through an underlying network. In addition to modeling data-flow relationships, the connections between two Symphony components allow the components to synchronize the behavior of the computational resources they represent. Simulations may be saved by their creators for repeated reuse and shared with other users in a problem-solving community. Simulations are transparently executed in a manner that respects the data-flow requirements of individual programs in the composition. Execution transparency in this context means that all system level operations of program execution and of data movement across geographically distributed locations must be largely, if not totally, transparent to the user.



The Symphony framework, based on Sun's JavaBeans component architecture [74], has two principal elements: a composition environment in which a simulation is constructed and a back-end execution environment in which the described computation is performed. Symphony beans can be placed in any runtime environment that supports JavaBeans, such as Sieve. Figure 13 shows Symphony running in Sun's BeanBox. The BeanBox interface allows generic Symphony components to be selected from a palette (not shown in Figure 13) and dropped into the BeanBox's workspace. The properties of components in the workspace can be changed through customization to contain information about a specific computation resource for which the bean is a surrogate. Symphony includes Program beans that represent executable entities on some machine, several beans for representing sources or destinations of data including a File bean, and beans that model input and output data streams.

The standardized architecture of JavaBeans provides for a large range of possible composition environments and operating systems that Symphony can work with. One such composition environment is Sieve (presented earlier), which adds the ability for several users to remotely collaborate on the same simulation. While Symphony was originally developed to exploit the graphical user interface of composition environments, we also envision that a simulation can be constructed and executed based on other representations (e.g. an XML description derived from a query-based specification of model instances). Languages for Grid workflow specification [75] are also relevant here.

When a simulation is executed, the components initiate, monitor, and synchronize the operations performed in a back-end execution environment. The back-end execution environment and the components collaborate to transport files between machines, execute programs, and connect data streams as needed to realize the computation specified in the network of beans. The operations are performed so as to respect the dataflow relationships defined by the simulation. The components are programmed to detect the completion of operations which they have initiated in the back-end environment. In the initial version of Symphony, user feedback on the progress of the computation was provided through graphical means (i.e. a component's icon changed colors to indicate its progress from not initiated, to in progress, to completed). The most recent version of Symphony includes a more sophisticated logging mechanism. The generated log file contains user-selectable levels of details on each operation. The back-end execution environment created for Symphony is an RMI-based server that provided access to resources local to the machine on which the server is running. This server offers a lightweight mechanism for systems that could not or did not want to use more sophisticated but more expensive Grid software. Our more recent work is to interface Symphony with one or more computational Grids and networks of workstations. At the moment, an interface to Globus has been developed. In general, we envision Symphony as an abstraction layer on top of popular Grid middleware systems that will ease and unify the access to Grid resources and provide a foundation for collaboration among Grid users through the exchange of pre-configured simulation components or even whole simulations. We believe that such a framework would be valuable to the community and would foster the more widespread application of computational Grids significantly.

A simple simulation is depicted in Figure 13. This simulation represents a remote program (Program 1) which takes a parameter file and a data file as input and generates an intermediate result file. This file is then used by another program (Program 2) to generate the final result which is also stored in a file and a summary of the computation which is sent to the standard output. The standard output in turn is redirected to a local file that is used as the input for a local file viewer application (a helper bean).



In terms of the functional layers presented in Figure 2, Symphony provides both simulation definition and model definition services. Symphony provides the runtime structure to synchronize the execution of a collection of programs without user intervention. As noted above, a configured and connected set of Symphony beans could be generated by programatic means or from some syntactic (e.g. XML) description. This description might itself be produced by the parameter definition layer in Figure 2. However, Symphony currently lacks a number of important features that should be provided by run management. For example, it should be possible to build cyclic simulations that would execute repeatedly. Such simulation would include a component to determine termination based on convergence or other optimization criteria. Additionally, support for parameter sweep style execution should be provided. Exploiting the concurrent execution of programs unrelated through data-flow constraints is also desirable. Symphony currently provides some facilities at the model definition layer because it, in cooperation with the composition environment, provides a user interface through which a simulation (a model) can be specified. More sophisticated model definition features are clearly desired. For example, it should be possible to aggregate part or all of a simulation as a single unit. This would allow for easier organization of larger simulations and also provide a better basis for sharing complex assemblies of components among users.

Some of the issues we are currently addressing in our work with Symphony are: defining a general syntax for resource declaration and job specification, support for the Globus middleware [8] and its security infrastructure, support for other middleware and security architectures, providing means for more effective data routing between resources, and support for resource discovery mechanisms and automatic resource allocation through super schedulers. A brief discussion follows.

A general syntax for resource declaration and job specification is needed to denote computational resources in a Grid-independent manner. We have had initial success with implementing a resource description that lets us create program beans which abstract programs accessible through our local Globus Grid by using the Globus Commodity Toolkit for Java. Our program beans can be customized with information collected from a Grid Information Index Service (GIIS [76]), from a local resource configuration file, or through manually entered parameters. We also added a Globus authentication service which provides an interface for all Symphony beans to access a default Globus user proxy; by this means, beans can authenticate themselves to Grid resources.

We are currently investigating the use of more efficient methods than used in the original version of Symphony to transport data files as well as executables from one location to another. For the Globus interface we are looking at employing third-party GSI-FTP and standard FTP services as well as globus access to secondary storage (GASS [77]) services which are based on the HTTPS protocol.

We are examining resource discovery and automatic resource allocation through super schedulers. We propose to interface with a resource brokerage system which could automate the resource selection process for compute resources by comparing the constraints stated by the software resource, user preferences and user rights to properties of the available resources in the Grid.

The Symphony framework requires a security architecture that will provide for finer grained and more flexible control of rights than is currently available on computational Grids. The requirement for finer grained control arises from Symphony's ability to exchange simulation components and whole simulations between users. Current Grid security infrastructures do not support these requirements. We believe that elements of the CRISIS [78] architecture may prove to be very useful to the aims of Symphony and should be implemented in future Grid security systems.



4. DISCUSSION

Our approach to GCEs can be characterized by an emphasis on high-level problem-solving facilities, and their implementation using traditional Grid information and lower-level services, for multidisciplinary Grid communities. Our unifying framework for multidisciplinary applications leverages concepts from data management, model representation, and compositional modeling into a cohesive methodology.

By viewing descriptions of model instances (and simulations) as entries in a database, programmatic descriptions of new model instances (and simulations) can be automatically created by writing queries. By writing BSMML specifications and using Symphony to associate the resulting simulations with the underlying computational environment, scientists and engineers are able to interact with GCEs in the vernacular of the domain. The distinctions made by our framework mirror other projects such as CACTUS applications [79], where the emphasis is on creating portals that integrate large-scale simulation and visualization.

The presented techniques also allow us to store descriptions, ‘run’ the descriptions to obtain data, record the data back in the database, and associate the data with the description that corresponds to its experimental setup. This emphasis on the entire problem-solving context facilitates sophisticated services such as change management.

In contrast to the variety of standards (many, XML-based) available for scientific data, our data model is aimed at capturing representations of simulations, not just simulation data. We posit that the description of a simulation is a more persistent representation of the data (it produces) than the data itself. As technology matures and evolves, recording how specific data was obtained is important for the purposes of ensuring repeatability and reliability. For example, if gridding technology for microarrays improves, then ‘running’ the same (stored) description with the new setup can be used to arrive at new, current, results. Since there is nothing in our design that commits us to a rigid schema, our data model can elegantly adapt to changes over time.

These requirements point to the future directions in the development of GCEs. We are now extending our ideas to apply to runtime scenarios, such as computational steering and closing-the-loop between design and analysis. Runtime recommendation of models will become pertinent in heterogeneous and distributed scenarios, where information about application characteristics is acquired only during the computation. Connections to Grid information services have to be established to enable GCEs to participate in large-scope projects that span multiple institutions and data sources.

The ‘simulation as querying’ viewpoint provides a useful conceptual abstraction for the effective utilization of computational Grids. A long-term goal is to design GCEs that help unify modeling, simulation, analysis, and design activities—this is especially pertinent in multidisciplinary applications. Our architecture and systems support technologies are uniquely designed to support such an integrated mode of investigation.

ACKNOWLEDGEMENTS

We acknowledge the numerous collaborators, post-doctoral research associates, and students who have been involved in the Virginia Tech PSEs Research Group over the past several years. The work presented in this paper is supported in part by National Science Foundation grants DMS-9625968, EIA-9974956, EIA-9984317, EIA-0103660, and MCB-0083315, NASA grant NAG-2-1180, AFOSR grant F496320-99-1-0128, the Defense



Advanced Research Projects Agency through the Office of Naval Research (N00014-01-1-0852), and the Virginia Tech ASPIRES program.

REFERENCES

1. Foster I, Kesselman C. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann: San Mateo, CA, July 1998.
2. Kennedy K. Telescoping languages: A compiler strategy for implementation of high-level domain-specific programming systems. *Proceedings of the 14th International Parallel and Distributed Processing Symposium (IPDPS'00)*. Springer, 2000.
3. Decker KM, Wylie BJ. Software tools for scalable multilevel application engineering. *International Journal of High Performance Computing Applications* 1997; **11**(3):236–250.
4. Saltz J, Sussman A, Graham S, Demmel J, Baden S, Dongarra J. Programming tools and environments. *Communications of the ACM* 1998; **41**(11):64–73.
5. Casanova H, Dongarra J. NetSolve: A network server for solving computational science problems. *International Journal of Supercomputer Applications and High Performance Computing* 1997; **11**(3):212–223.
6. Foster I, Kesselman C, Tuecke S. The anatomy of the Grid: Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications* 2001; **15**(3):200–222.
7. Baker M, Buyya R, Laforenza D. The Grid: International efforts in global computing. *Proceedings of the International Conference on Advances in Infrastructure for Electronic Business, Science, and Education on the Internet (SSGRR 2000)*, Italy, 2000.
8. Foster I, Kesselman C. Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications and High Performance Computing* 1997; **11**(2):115–128.
9. von Laszewski G. The Grid Portal Development Kit. *Grid Portal Collaboration Effort Document*, 2001.
10. Verma S, Parashar M, von Laszewski G, Gawor J. Corba Community Grid Toolkit (CoG). *Grid Computing Environments Community Practice (CP) Document*, 2001.
11. Fox G, Gannon D, Thomas M. Grid computing environments working group specification. *Global Grid Forum*, 2001.
12. Jiang J, Bae KK, Tranter WH, Verstak A, Ramakrishnan N, He J, Watson LT, Rappaport TS, Shaffer CA. S⁴W: A collaborative PSE for modeling of broadband wireless communication systems. *Proceedings of the Eleventh Annual MPRG Symposium on Wireless Personal Communications*. Mobile and Portable Radio Research Group Publications, 2001.
13. Goel A, Baker CA, Shaffer CA, Grossman B, Mason WH, Watson LT, Haftka RT. VizCraft: A problem-solving environment for aircraft configuration design. *IEEE/AIP Computing in Science and Engineering* 2001; **3**(1):56–66.
14. Knill DL, Giunta AA, Baker CA, Grossman B, Mason WH, Haftka RT, Watson LT. Response surface models combining linear and Euler aerodynamics for supersonic transport design. *Journal of Aircraft* 1999; **36**(1):75–86.
15. Koch KR, Baker RS, Alcouffe RE. Solution of the first-order form of the 3-D discrete ordinates equation on a massively parallel processor. *Transactions of the American Nuclear Society* 1992; **65**:198–199.
16. Lawrence Livermore National Laboratory. The ASCI Sweep3D Benchmark Code. http://www.llnl.gov/asci_benchmarks [1995].
17. Adve VS, Bagrodia R, Browne JS, Deelman E, Dube A, Houstis EN, Rice JR, Sakellariou R, Sundaram-Stukel DJ, Teller PJ, Vernon MK. POEMS: End-to-end performance design of large parallel adaptive computational systems. *IEEE Transactions on Software Engineering* 2000; **26**(11):1027–1048.
18. Browne JC, Berger E, Dube A. Compositional development of performance models in POEMS. *International Journal of High Performance Computing Applications* 2000; **14**(4):283–291.
19. Keahey K, Beckman P, Ahrens J. Ligation: Component architecture for high performance applications. *International Journal of High Performance Computing Applications* 2000; **14**(4):347–356.
20. Forbus KD. Qualitative reasoning. *The Computer Science and Engineering Handbook*, Tucker AB (ed.). CRC Press: Boca Raton, FL, 1996; 715–733.
21. Cummings JC, Crottinger JA, Haney SW, Humphrey WF, Karmesin SR, Reynders JW, Smith SA, Williams TJ. Rapid application development and enhanced code interoperability using the POOMA framework. *Proceedings of the SIAM Workshop on Object-Oriented Methods and Code Interoperability in Scientific and Engineering Computing (OO'98)*. SIAM Press: Philadelphia, PA, 1998.
22. Gannon D, Bramley B, Stuckey T, Villacis J, Balasubramanian J, Akman E, Breg F, Diwan S, Govindaraju M. Component architectures for distributed scientific problem solving. *IEEE Computational Science and Engineering* 1998; **5**(2):50–63.
23. Bodin, F, Beckman P, Gannon D, Narayana S, Yang SX. Distributed pC++: Basic ideas for an object parallel language. *Scientific Programming* 1993; **2**:7–22.



24. Chandy KM, Kesselman C. *CC++: A Declarative, Concurrent, Object-Oriented Programming Notation (Research Directions in Concurrent Object-Oriented Programming)*. MIT Press: Cambridge, MA, 1993.
25. Foster I. Compositional parallel programming languages. *ACM Transactions on Programming Languages and Systems* 1996; **18**(4):454–476.
26. Drashansky TT. An agent based approach to multidisciplinary problem solving environments. *PhD Thesis*, Department of Computer Sciences, Purdue University, 1996.
27. Drashansky TT, Houstis EN, Ramakrishnan N, Rice JR. Networked agents for scientific computing. *Communications of the ACM* 1999; **42**(3):48–54.
28. Johnson C, Parker SG, Hansen C, Kindlmann GL, Livnat Y. Interactive simulation and visualization. *IEEE Computer* 1999; **32**(12):59–65.
29. Avery P, Foster I. Petascale virtual-data grids for data intensive science. White Paper. <http://www.griphyn.org> [April 2000].
30. Stonebraker M. Sequoia 2000: A reflection on the first three years. *IEEE Computational Science and Engineering* 1994; **1**(4):63–72.
31. Casanova H, Obertelli G, Berman F, Wolski R. The AppLeS parameter sweep template: User-level middleware for the Grid. *Proceedings Supercomputing Conference (SC'2000)*. IEEE Computer Society Press, 2000.
32. Fox GC, Furmanski W. High-performance commodity computing. *The Grid: Blueprint for a New Computing Infrastructure*, Kesselman C, Foster I (eds.). Morgan Kaufmann: San Mateo, CA, 1998; 237–255.
33. Goel A, Phanouriou C, Kamke FA, Ribbens CJ, Shaffer CA, Watson LT. WBCSim: A prototype problem solving environment for wood-based composites simulation. *Engineering with Computers* 1999; **15**:198–210.
34. Rubin EJ, Dietz R, Chanat J, Speir C, Dymond R, Lohani V, Kibler D, Bosch D, Shaffer CA, Ramakrishnan N, Watson LT. From landscapes to waterscapes: A PSE for landuse change analysis. *Proceedings of the 16th IMACS World Congress*, August 2000, Deville M, Owens R (eds.).
35. Verstak A, Vass M, Ramakrishnan N, Shaffer CA, Watson LT, Bae KK, Jiang J, Tranter WH, Rappaport TS. Lightweight data management for compositional modeling in problem solving environments. *Proceedings of the High Performance Computing Symposium, Advanced Simulation Technologies Conference*, 2001, Tentner A (ed.); 148–153.
36. Ramakrishnan N, Houstis EN, Rice JR. Recommender systems for problem solving environments. *Technical Report WS-98-08 (Working Notes of the AAAI-98 Workshop on Recommender Systems)*, Kautz H (ed.). AAAI/MIT Press, 1998; 91–95.
37. Alscher RG, Chevone BI, Heath LS, Ramakrishnan N. Espresso—A PSE for bioinformatics: Finding answers with microarray technology. *Proceedings of the High Performance Computing Symposium, Advanced Simulation Technologies Conference*, April 2001, Tentner A (ed.). Society for Computer Simulation International, 2001; 64–69.
38. Hamadeh H, Afshari CA. Gene chips and functional genomics. *American Scientist* 2000; **88**:508–515.
39. Grossman RL, Hanley D, Bailey S. High performance data management: A case for using lightweight, high performance persistent object managers in scientific computing. *Technical Report 95-R18*, Laboratory for Advanced Computing, University of Illinois, Chicago, IL, 1995.
40. Heath LS, Ramakrishnan N, Sederoff RR, Whetten RW, Chevone BI, Struble CA, Jouenne VY, Chen D, van Zyl L, Alscher RG. The Espresso Microarray Experiment Management System with application to the functional genomics of stress responses in loblolly pine. *Comparative and Functional Genomics* 2001; **3**(3):226–243.
41. Fishwick PA. *Simulation Model Design and Execution: Building Digital Worlds*. Prentice-Hall: Englewood Cliffs, NJ, 1995.
42. Kuipers B. Qualitative simulation. *Artificial Intelligence* 1986; **29**:289–338.
43. Bramley R, Chiu K, Diwan S, Gannon D, Govindaraju M, Mukhi N, Temko B, Yochuri M. A component based services architecture for building distributed applications. *Proceedings 9th IEEE International Symposium on High Performance Distributed Computing (HPDC'00)*. IEEE Computer Society Press: Los Alamitos, CA, 2000.
44. Gannon D, Bramley R, Stuckey T, Villacis J, Balasubramanian J, Akman E, Breg F, Diwan S, Govindaraju M. The linear system analyzer. *Enabling Technologies for Computational Science*, Houstis EN, Rice JR, Gallopoulos E, Bramley R (eds.). Kluwer: Dordrecht, 2000; 123–134.
45. Ioannidis Y, Livny M, Gupta S, Ponnekanti N. ZOO: A desktop experiment management environment. *Proceedings of the 22nd International Conference on Very Large Databases (VLDB'96)*. Morgan Kaufmann, 1996; 274–285.
46. Upson C, Faulhaber T, Kamins D, Schlegel D, Laidlaw D, Vroom F, Gurwitz R, van Dam A. The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications* 1989; **9**(4):30–42.
47. Falkenhainer B, Forbus K. Compositional modeling: Finding the right model for the right job. *Artificial Intelligence* 1991; **51**:95–143.
48. Nayak PP. Automated modeling of physical systems. *PhD Thesis*, Stanford University, 1992.



49. Rickel J, Porter B. Automated modeling of complex systems for answering prediction questions. *Artificial Intelligence* 1997; **93**:201–260.
50. Cleary A, Kohn S, Smith SG, Smolinski B. Language interoperability mechanisms for high-performance scientific applications. *Technical Report UCRL-JC-131823*, LLNL, 1998.
51. Gannon D. Grid Information Services and XML Schemata. White Paper. <http://www.extreme.indiana.edu/~gannon/schemata.html> [2000].
52. Allen G, Bengler W, Goodale T, Hege H-C, Lanfermann G, Merzky A, Radke T, Seidel E, Shalf J. Cactus tools for grid applications. *Cluster Computing* 2001; **4**:179–188.
53. Eidson TM. A component-based programming model for composite, distributed applications. *Technical Report ICASE Report No. 2001-15; NASA/CR-2001-21087*, Institute for Computer Applications in Science and Engineering (ICASE), May 2001.
54. Gannon D, Grimshaw A. Object-based approaches. *The Grid: Blueprint for a New Computing Infrastructure*, Kesselman C, Foster I (eds.). Morgan Kaufmann: San Mateo, CA, 1998; 205–236.
55. Chawathe S, Garcia-Molina H. Meaningful change detection in structured data. *Proceedings of the ACM-SIGMOD Conference on Management of Data*, Tucson, AZ. ACM Press, 1997; 26–37.
56. Abiteboul S, Buneman P, Suciu D. *Data on the Web: From Relations to Semistructured Data and XML*. Morgan Kaufmann: San Francisco, CA, 2000.
57. Hellerstein JM. Optimization techniques for queries with expensive methods. *ACM Transactions on Database Systems* 1998; **23**(2):113–157.
58. Houstis EN, Catlin AC, Rice JR, Verykios VS, Ramakrishnan N, Houstis CE. PYTHIA-II: A knowledge/database system for managing performance data and recommending scientific software. *ACM Transactions on Mathematical Software* 2000; **26**(2):227–253.
59. Ramakrishnan N, Ribbens CJ. Mining and visualizing recommendation spaces for elliptic PDEs with continuous attributes. *ACM Transactions on Mathematical Software* 2000; **26**(2):254–273.
60. Ramakrishnan N, Valdés-Pérez RE. Note on generalization in experimental algorithmics. *ACM Transactions on Mathematical Software* 2000; **26**(4):568–580.
61. Ramakrishnan N. Experiences with an algorithm recommender system. *Working Notes of the CHI'99 Workshop on Interacting with Recommender Systems*, Baudisch P (ed.). ACM SIGHI Press, 1999.
62. Shaffer E, Reed DA, Whitmore S, Schaeffer B. Virtue: Performance visualization of parallel and distributed applications. *IEEE Computer* 1999; **32**(12):44–51.
63. Chandy KM, Bramley R, Char BW, Reynders JVW. Report of the NSF workshop on problem solving environments and scientific IDEs for knowledge, information and computing (SIDEKIC'98). *Technical Report*, Los Alamos National Laboratory, 1998.
64. Foster I, Inslay J, Von Laszewski G, Kesselman C, Thiebaux M. Distance visualization: Data exploration on the Grid. *IEEE Computer* 1999; **32**(12):36–43.
65. Karin S, Graham S. The high performance computing continuum. *Communications of the ACM* 1998; **41**(11):32–35.
66. Moore RW, Baru C, Marciano R, Rajasekar A, Wan M. Data-intensive computing. *The Grid: Blueprint for a New Computing Infrastructure*, Kesselman C, Foster I (eds.). Morgan Kaufmann: San Mateo, CA, 1998; 105–129.
67. Moore RW, Prince TA, Ellisman M. Data-intensive computing and digital libraries. *Communications of the ACM* 1998; **41**(11):56–62.
68. Ramakrishnan N, Grama AY. Mining scientific data. *Advances in Computers* 2001; **55**:119–169.
69. Rice JR, Boisvert RF. From scientific software libraries to problem-solving environments. *IEEE Computational Science and Engineering* 1996; **3**(3):44–53.
70. Szalay A, Kunszt P, Thakar A, Gray J. Designing and mining multi-terabyte astronomy archives: The Sloan Digital Sky Survey. *Proceedings of the ACM-SIGMOD Conference on Management of Data*. ACM Press, 2000; 451–462.
71. Adve VS, Sakellariou R. Application representations for multiparadigm performance modeling of large-scale parallel scientific codes. *International Journal of High Performance Computing Applications* 2000; **14**(4):304–316.
72. Young M, Argiro D, Kubica S. Cantata: Visual programming environment for the Khoros system. *Computer Graphics* 1995; **29**(2):25–28.
73. Shah A, Kafura DG. Symphony: A Java-based composition and manipulation framework for distributed legacy resources. *Proceedings of the International Symposium on Software Engineering for Parallel and Distributed Systems*, Los Angeles, CA, May 1999. Springer, 1999; 2–12.
74. Sun Microsystems Inc. *The JavaBeans Component Architecture*. White Paper. <http://java.sun.com/products/javabeans/> [1998].
75. Bivens HP. Grid workflow. *Grid Computing Environments Working Group Document*, 2001.
76. Czajkowski S, Fitzgerald K, Foster I, Kesselman C. Grid information services for distributed resource sharing. *Proceedings 10th IEEE International Symposium on High-Performance Distributed Computing*. IEEE Computer Society Press: Los Alamitos, CA, 2001.



-
77. Bester J, Foster I, Kesselman C, Tedesco J, Tuecke S. GASS: A data movement and access service for wide area computing systems. *Proceedings of the Sixth Workshop on I/O in Parallel and Distributed Systems*, May 1999. ACM Press, 1999.
 78. Belani E, Vahdat A, Anderson T, Dahlin M. The CRISIS wide area security architecture. *Proceedings of the USENIX Security Symposium*, San Antonio, TX, January 1998. USENIX, 1998; 15–30.
 79. Allen G, Goodale T, Lanfermann G, Radke T, Seidel E, Bengler W, Hege H-C, Merzky A, Massó J, Shalf J. Solving Einstein's equations on supercomputers. *IEEE Computer* 1999; **32**(12):52–58.