

GAUGE: Grid Automation and Generative Environment

Francisco Hernández, Purushotham Bangalore, Jeff Gray, Zhijie Guan, Kevin Reilly
{hernandf, puri, gray, zhijie, reilly}@cis.uab.edu
Department of Computer and Information Sciences
University of Alabama at Birmingham

Abstract

The Grid has proven to be a successful paradigm for distributed computing. However, constructing applications that exploit all the benefits that the Grid offers is still not optimal for both inexperienced and experienced users. Recent approaches to solving this problem employ a high-level abstract layer to ease the construction of applications for different Grid environments. These approaches help facilitate construction of Grid applications, but they are still tied to specific programming languages or platforms. A new approach is presented in this paper that uses concepts of domain-specific modeling (DSM) to build a high-level abstract layer. With DSM, the users are able to model Grid applications without being bound to specific programming languages or platforms. An additional benefit of DSM provides the capability to generate software artifacts for various Grid environments.

This paper presents the Grid Automation and Generative Environment (GAUGE). The goal of GAUGE is to automate the generation of Grid applications to allow inexperienced users to exploit the Grid fully. At the same time, GAUGE provides an open framework in which experienced users can build upon and extend to tailor their applications to particular Grid environments or specific platforms. GAUGE employs domain-specific modeling techniques to accomplish this challenging task.

1 Introduction

The Grid has proven to be a successful paradigm for distributed computing [13], [33]. It provides dependable, collaborative and secure access to remote computational, data, and instrumentation resources. Different Grid frameworks have emerged to provide support for this model of computation (e.g., Globus [15], Legion [28], and Unicore [38]). The Globus Toolkit [15] is perhaps the most widely used framework to construct Grid applications. Globus provides a common middleware that considers resources as entities of a virtual organization. The middleware is formed by different components that provide services to integrate distributed resources in a Grid computing environment. Creating an application that uses Globus requires the composition of several of these components. Unfortunately, constructing Grid applications is still difficult despite all the capabilities that Globus provides [21]. The Java Commodity Grid Toolkit (Java CoG Kit) was later created to assist in the development of applications using Globus [27] services. This was a step towards simplifying the construction of applications for the Globus Toolkit. The Java CoG Kit helps a user navigate the intricacies of the Globus components more easily

by introducing a new programming model for the Grid. Furthermore, the Java CoG Kit provides many utility components that enhance the functionality of Globus.

However, developing applications for the Grid remains difficult for many users. This is due in part to the complexity of the back-end systems, where even potentially inexperienced users are exposed to all the details of the underlying Grid technologies [22]. Traditionally, Problem Solving Environments (PSE), or portals [32], have been developed to ease the construction of Grid applications for the un-savvy Grid user. PSE's provide a high-level view for specifying Grid-enabled applications and rely on middleware to connect with the Grid component resources [14]. This kind of tool expedites simple tasks (i.e., simple job submissions, and checking the status of a previously submitted job), but it lacks the flexibility to define a complex sequence of tasks. Thus, an inexperienced user is not equipped with the capabilities for exploiting the Grid without having to decipher the intricacies of the Grid technologies. In such cases, users can become overwhelmed by the numerous low-level details involved in engineering a Grid application.

Experienced Grid users (developers) also struggle when developing Grid applications. We observe two principal reasons for this problem: 1) current software engineering practices (e.g., reusability, modeling and rapid prototyping) have not been fully explored for the Grid model, and 2) the Grid applications usually consist of different back-end implementations, but there is no standard for the application development process [5]. Thus, a developer has to write application-specific code for each back-end. A solution to improve this kind of development is to introduce a high-level abstraction layer for hiding the complex back-ends and provide functionality for different Grid architectures [5], [23], [39]. Yet, it can be observed that performing the abstraction at a programming-language level is still not optimal, and that the two problems enumerated above can be lessened if the abstraction is realized at the *domain level*. Domain-specific modeling (DSM) [18] enables the user to employ familiar concepts in the Grid domain to construct models of Grid applications. These models can then be synthesized (i.e., translated) into different representations (e.g., XML configuration files or source code). With this technology, a user focuses on higher levels of abstraction at the problem space and is able to avoid low-level details, such as low-level Grid services and their usage. Domain-specific modeling has been useful in automating different kinds of applications in which the environment is dynamic; including embedded systems [31], automotive manufacturing [30], and complex QoS applications [9].

This paper presents the Grid Automation and Generative Environment (GAUGE). GAUGE aims to enable inexperienced users to take full advantage of the Grid infrastructure and at the same time assists developers in modeling and prototyping Grid applications. Because the use of Globus and the Java CoG Kit has proven to be a successful combination, GAUGE relies on this combination for interfacing with Grid infrastructure. The approach employed by GAUGE provides a high-level abstract layer for the construction of Grid applications. This layer is composed of visual models that are constructed using concepts of domain-specific modeling. Programs that manage the application execution are generated from the corresponding visual models. Thus, users need not learn how to use the Globus Toolkit or the Java CoG Kit in order to develop Grid-enabled applications. GAUGE also provides an open framework in which

developers can reuse components and tailor the automation environment to work with a particular Grid Architecture.

The remainder of this paper is organized as follows. Section 2 provides an overview of GAUGE. Section 3 outlines the requirements considered when developing the automation environment. The architecture of GAUGE is introduced in Section 4. Section 5 presents the implementation of GAUGE, including the construction of the user interface layer, the domain transformation layer, and the code generation layer. Section 6 presents related work, and Section 7 outlines future directions for research in this area. Finally, Section 8 gives conclusions of our work.

2 GAUGE Tool

A *Grid application* is an executable software artifact that is suited to work in a Grid environment. Under current practice, a Grid application is created by composing different sub-tasks that generate an execution flow. Typically, these applications also require specification of the data flow needed to accomplish the execution flow. The data flow is required because it establishes pre- and post-conditions for task execution [3]. The combination of both flows defines a *Workflow*. The complexity of building a workflow system for the Grid arises from the ability of each of the constituent tasks to use different resources. Thus, an in-depth understanding of the underlying environment (hardware and software) is typically needed to define a Grid workflow. The Grid application consists of the control code that executes tasks on different resources (execution flow), and also moves and copies files between those resources (data flow).

Grid Automation and Generative Environment (GAUGE) is a tool that automates the development of *Grid applications*. The goal of GAUGE is to automate the generation of the control code that manages a particular Grid application in a manner that can function in any system that uses a Grid framework, such as the *Globus Toolkit* [15]. GAUGE supports the following Grid tasks: job submissions, explicit file transfers, Open Grid Service Architecture (OGSA) [16] compliant Grid services, and inner workflows. Two types of users are considered: operators and developers. Operators are knowledgeable about the domain of a specific application's tasks, as well as the manner in which the relevant tasks are combined to form the final Grid application. However, operators typically are not familiar with the underlying Grid technologies. In order for them to fully utilize the Grid, complex programs often must be developed to manage the control of the applications. Developers, the other type of users, are assumed to be Grid savvy. A developer possesses knowledge about the intricacies of the underlying Grid technologies, but often lacks ability in handling the specifics of the different application constituents for a specific domain. Maximizing the use of the Grid infrastructure requires a synergistic coordination between operator and developer roles. One of the objectives of GAUGE is to facilitate the interaction between these two types of users. GAUGE focuses on the domain experts (operators) and helps them by providing an environment in which they can graphically specify the workflow for their application and automatically generate the code that manages the execution of the workflow (execution and data flow). For the developers, this tool provides an open framework in which they can reuse components (e.g., the user interface and the domain transformation layer), and tailor the automation environment to work within a particular Grid environment.

Interaction with the system is realized through a graphical user interface that is provided by the Generic Modeling Environment (GME) [26]. GME is a domain-specific modeling environment that can be configured and adapted from meta-level specifications that describe the domain [24]. The domain-specific models are developed by first creating a meta-model that specifies the syntax and static semantics of a visual language; the dynamic semantics are introduced by an interpreter that synthesizes the models into different representations [24]. The Globus Domain Model (GDM) [21], [20] is the visual language that configures and adapts GME to function in the Grid domain. The visual language shields the underlying complexity from the user by incorporating concepts from the Grid domain, and associates those concepts with graphical elements (GDM constructs).

In addition to the Grid domain constructs, GDM supports a set of constructs that allows the specification of workflows from simple tasks that can be defined within GAUGE. According to [3], functionality of workflow systems can be expressed by a set of recurring abstractions in the specification of workflows (*workflow patterns*). These abstractions are based on a set of routing primitives [2] that control the workflow's execution: sequence, splits, joins, and iteration. GDM supports a basic set of control routing primitives for implementing the workflow's functionality: implicit AND-Split, implicit AND-Join, explicit XOR-Split, and explicit XOR-Join. With these basic constructs, the user controls the execution flow of the workflows, and is also able to define basic workflow patterns [4] such as: sequence, parallel split, synchronization, exclusive choice, simple merge, and arbitrary cycles. This set of patterns is important because it creates a workflow language capable of executing tasks in sequence and in parallel, executing a task an arbitrary number of times, and steering the execution flow according to particular conditions encountered during the execution of the workflow. All of these capabilities are required [25] when implementing a Grid workflow. The manner in which GDM is constructed is explained in detail in Section 5.1.

The user's interaction with GDM generates graphical models of the workflows that visually express the workflow's tasks and the execution and data flow between those tasks. GME provides an API that traverses the internal representation of the models. A model interpreter uses this API to translate the models into an application that manages the execution of the workflow. This is a typical example of software translation because a new artifact is created from a description at a different level of abstraction [19]. A key benefit of model-driven techniques is that different output representations can be generated based on the same model properties [26]. Currently, GAUGE is able to generate a Java application that interfaces with Globus using the Java CoG Kit. GAUGE can also generate an initial OGSA [16] Grid service representation that must be submitted to a Grid services environment for its execution.

Although the user interface of GAUGE is aimed at users who are not experts in the different Grid technologies, developers can also benefit by reusing GAUGE's infrastructure. Thus applications for different Grid environments can be generated from the same graphical models. This reduces the time spent in developing Grid architecture's prototypes, allowing developers to verify that a specific Grid architecture is feasible before investing resources in developing it for production purposes.

3 Requirements

Two requirements guided the design of GAUGE: (1) the simplicity in the specification of the workflows, and (2) the generality in the specifics of the workflow tasks. The challenge of the first requirement is hiding the underlying complexity (i.e., heterogeneous and distributed back-end resources, Grid technologies) from the specification of the workflows. The solution employed is to use concepts that are familiar to the users (i.e., location of hosts, job specifications, file transfers) and associate those concepts with graphical elements. This association defines a visual language. Thus, the users graphically combine these elements and generate the corresponding workflows. The second requirement permits the creation of a general tool that is not tied to a specific Grid environment. To fulfill this requirement, the workflow tasks are considered external with respect to GAUGE.

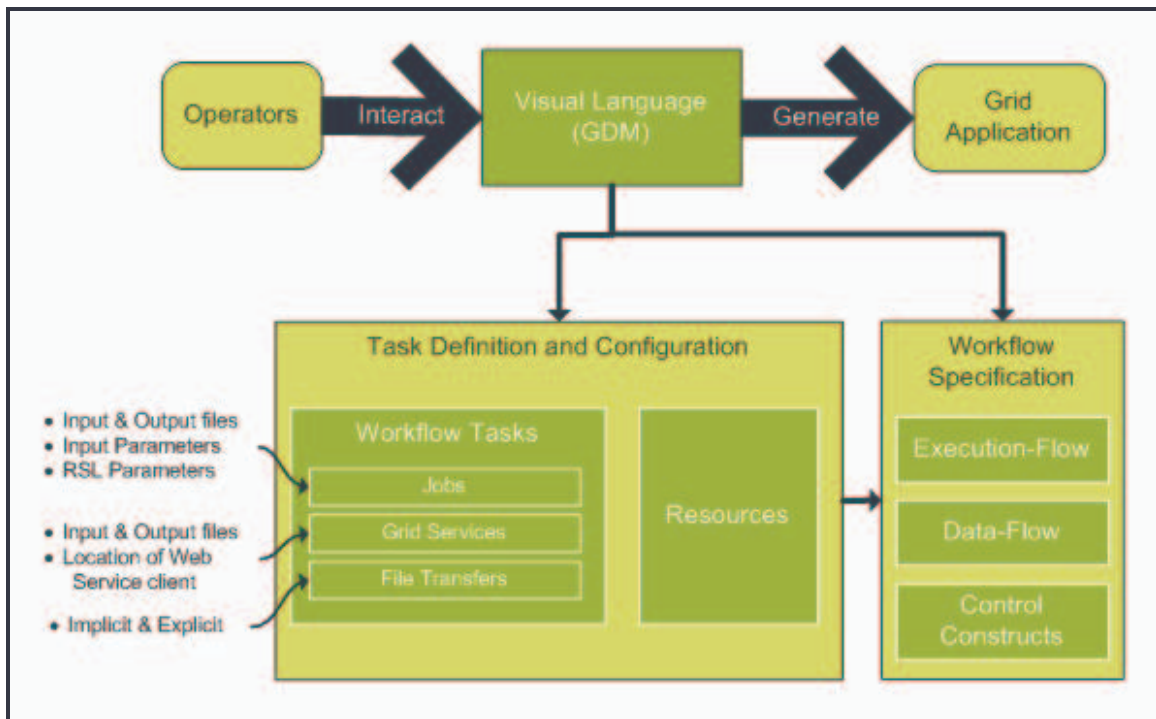


Figure 1. Function of the GAUGE Modeling Process. The operators interact with a visual language that provides definition and configuration of tasks, and specification of workflows, from which the code that forms the Grid application is generated.

The function of the GAUGE modeling process is presented in Figure 1. Operators interact with the visual language in two ways: (1) defining and configuring the workflow tasks, and (2) specifying the execution flow that generates the workflow. Defining and configuring tasks involves specifying the resources (hardware) that are used for the tasks. Communication between tasks is conducted via files, thus each job and Grid service definition contains a list of the required input/output files. Jobs require basic input parameters (e.g., executable, directory, standard output), but the complete specification of Resource Specification Language (RSL) [36] is also permitted. The RSL parameters specify the tasks in more detail, thus maximizing the use of the resources. In addition to

the list of input and output files, Grid services require the path to the client that controls the execution of the Grid service. File transfers can be implicit or explicit, but the operator can only define the explicit file transfers. The execution flow defines the workflow by specifying the order of the tasks. The data flow is based on the execution flow, because it is guided by the list of input and output files that are defined for each workflow task. The data flow is implemented by moving files between tasks, which require adding implicit file transfers to the execution flow. Control flow constructs (XOR-split, AND-split, XOR-join, AND-join), are also required when specifying a workflow. With these basic constructs, the user is able to execute tasks in sequence and in parallel, execute a task an arbitrary number of times, and steer the execution flow according to particular conditions encountered during the execution of the workflow. After the workflows specifications are completed, the user is able to generate the code that manages the application's execution.

4 Architecture

The left side of Figure 2 shows the architecture of GAUGE. The major components consist of a front-end, a domain transformation tier, and a code-generation tier. The front-end is realized by the Generic Modeling Environment (GME) [26], which is configured by the visual language GDM. GDM uses concepts of workflow control mixed with concepts of the Grid domain. The middle tier (middle left side of the diagram) is comprised by the domain transformations layer. This layer provides a Java API that wraps the API provided by GME to access the GDM models. The lower left-hand side of the diagram shows the code generation layer. This layer generates code for a particular Grid environment, in our case it generates Java code that manages the execution of the workflows defined in GDM. A developer can extend GAUGE by replacing the code generation layer and by interacting with the middle layer generated code suited for a different Grid environment (i.e., not Globus), or the developer can generate code written in another programming language.

The structures of the generated applications currently supported are shown on the right side of Figure 2. The generated application consists of a Java program that controls and manages the correct execution of the workflows. This application performs the function of what is typically known as the workflow engine [25]. Work is underway to generate corresponding Grid services (conforming to the OGSA [16] specification) requests from the generated Java programs. Both types of applications interact with the back-end systems through the Java CoG Kit [27], and the Globus Toolkit [15]. Currently, we are working on generating two more output representations within the Globus environment: a Python [35] script that interacts with the Globus Toolkit using PyGlobus [34], and a GridAnt [25] specification file.

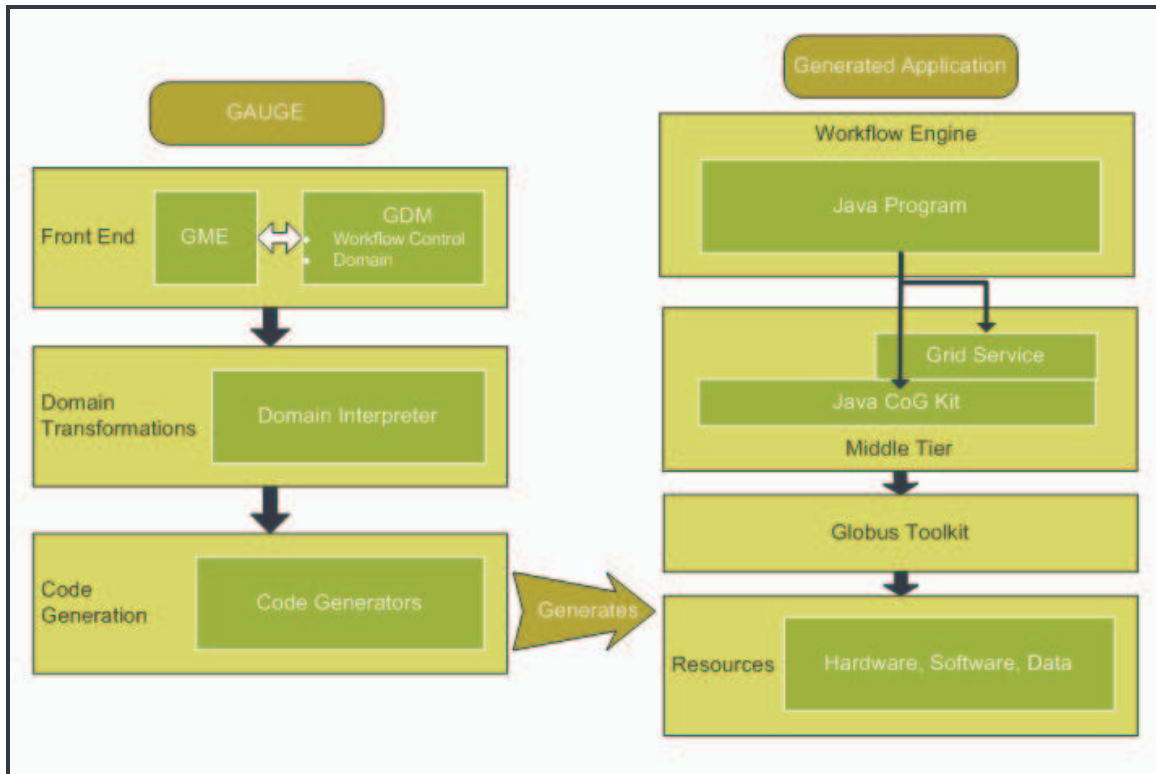


Figure 2. Architecture of the Modeling Environment. The left side of the figure presents the architecture for GAUGE. The right side presents the organization of the programs generated by GAUGE.

5 Implementation

In this section, the implementation of GAUGE is described. For the sake of clarity, the system is split into three components: user interface, domain transformation, and code generation.

5.1 User interface

The user interface presents an environment in which the users can graphically manage the workflows of their Grid applications. The front-end hides the intricacies of the underlying Grid environment from the end user. Therefore, it is not required to be proficient in Grid technologies to write complex Grid applications. The graphical interface is provided by the Generic Modeling Environment (GME), which is configured using the visual language GDM [21], [20]. This combination defines the interaction of the user with the system (section 5.1.1). Aside from the Grid domain constructs, the visual language also employs control constructs (section 5.1.2) that are generally used to generate workflow patterns; this facilitates control of the execution of the workflow.

5.1.1 GDM

GDM is a visual language that was constructed using concepts of domain-specific modeling. In domain-specific modeling, a design engineer creates models for a specific domain using concepts and terminology from that domain [18]. The domain-specific models are developed by first creating a meta-model that specifies the ontology of the domain. The meta-model serves as a paradigm, or language, that defines the syntax and

static semantics of the visual language; the dynamic semantics are introduced by an interpreter that synthesizes the models into different representations [24]. Domain-specific modeling has been useful in automating different kinds of applications in which the environment is dynamic; including embedded systems [31], automotive manufacturing [30], and complex QoS applications [9].

GDM was written using the general paradigm provided by GME [26]. This general paradigm consists of meta-language elements specified in the form of a UML class diagram [8]. The domain-specific concepts are expressed as stereotypes of specific classes. A design engineer is able to define the type of models that can be constructed by specifying the relationships between the stereotypes. The three basic classes supported by GME are atom, model, and connection. Atoms are classes of objects that do not contain other objects. Models are container classes, and connections associate models and atoms.

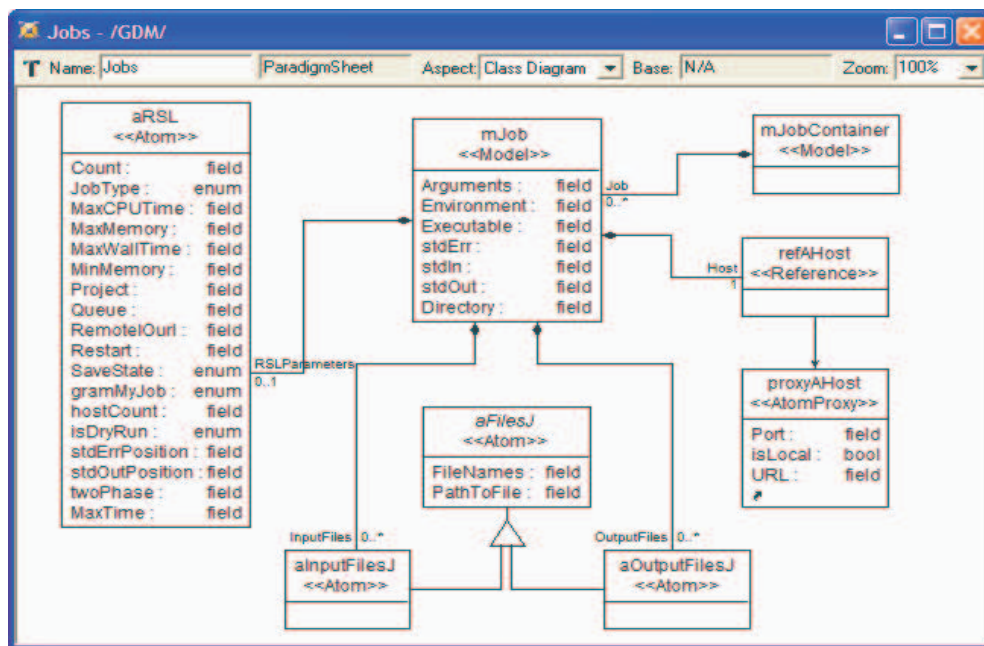


Figure 3. GDM language specification for job tasks. This figure presents the manner in which the relation of the different (GME) classes specify a construct in GDM.

Figure 3 illustrates the use of these basic classes that generate the job specification aspect of GDM. To fully define a job submission (see section 3), the user should specify the machine to which the job will be submitted, input parameters and the lists of input and output files. RSL parameters are optional. The central concept in the figure is mJob, which represents a job submission that is defined as a model to contain the rest of the specifications. The attributes that mJob contains are also specified in the class diagram. These attributes represent the input parameters to be specified when defining the job tasks. mJobContainer represents a container for job specifications that allows the user to group jobs according to the specific preferences. refAHost, defined below mJobContainer, represents the machine for the job. Notice that refAHost is declared as a reference class. This means that when declaring a job in GDM, a user must first define the host to use and then declare a reference to it. Any changes in the original host will be

reflected in all its references. Because refAHost is defined in other parts of the meta-model, refAHost is associated with a Proxy class that points to the original host definition. The line between mJob and refAHost indicates that mJob contains refAHost. The cardinality of this association indicates that there must be exactly one host defined for each job. aRSL permits the specification of all the RSL parameters for the job. This information is optional, as can be seen in the cardinality of the association (0..1). Finally, mJob also contains the list of input and output files. As typical in UML models, the triangle indicates that both kinds of lists inherit attributes and connections from a basic aFilesJ entity.

In addition to assisting in the construction of the visual language, GME provides an environment in which the configurations and the definitions of resources, tasks and workflows can be persistent and shared between users. All these features make GME an ideal choice for the front-end of GAUGE.

5.1.2 Control constructs

van der Aalst et al., [3] introduces the concept of workflow patterns to identify the functionality of workflow systems. From our observation, patterns are abstractions that keep recurring in the specification of workflows. These abstractions are based on a set of routing primitives that control the workflow's execution. GAUGE supports a set of basic control constructs that in turn allow the specification of workflow patterns.

GAUGE supports four constructs to manage the execution workflow. These constructs provide basic functionality: sequence, splits (XOR and AND), joins (XOR and AND), and iteration. Complex workflows can be specified through composition of these primitive functions.

Figure 4.a presents the manner in which the sequential construct is specified in GAUGE. The workflow example consists of a file transfer (FT), a job submission (A) and a Grid service execution (B). Each one of these three tasks requires the previous one to be complete before they can start processing. Notice the manner in which the data flow is conducted: B requires a file from A, "preproc". By labeling the connection, an implicit file transfer is conducted between the output files of A and the input files of B. However, there is no named connection between FT and A. The reason is that FT is an explicit file transfer, so the user is explicitly moving the input files of A. Although not shown in the figure, a task may require both types of transfers to collect its input files. This is also supported by GAUGE.

AND-Splits and AND-Joins are declared implicitly within the workflow tasks and do not require a specific construct. In AND-Splits, the sequential execution is broken into many parallel branches. The AND-Join must wait until all the branches have finished before it can begin its processing. The manner in which these constructs are enacted in GAUGE is depicted in figure 4.b. Job A splits the execution into Job B and Job C. Grid service D waits for the completion of B and C to begin its processing. The connections named *all* (between A-C, B-D, and C-D) indicate that all the output files from the source (A,B,C) need to be transferred to the destination (C,B,D). The star in the connection between A and B indicates that there are several files to be transferred between A and B, but that not all the output files from A need to be transferred.

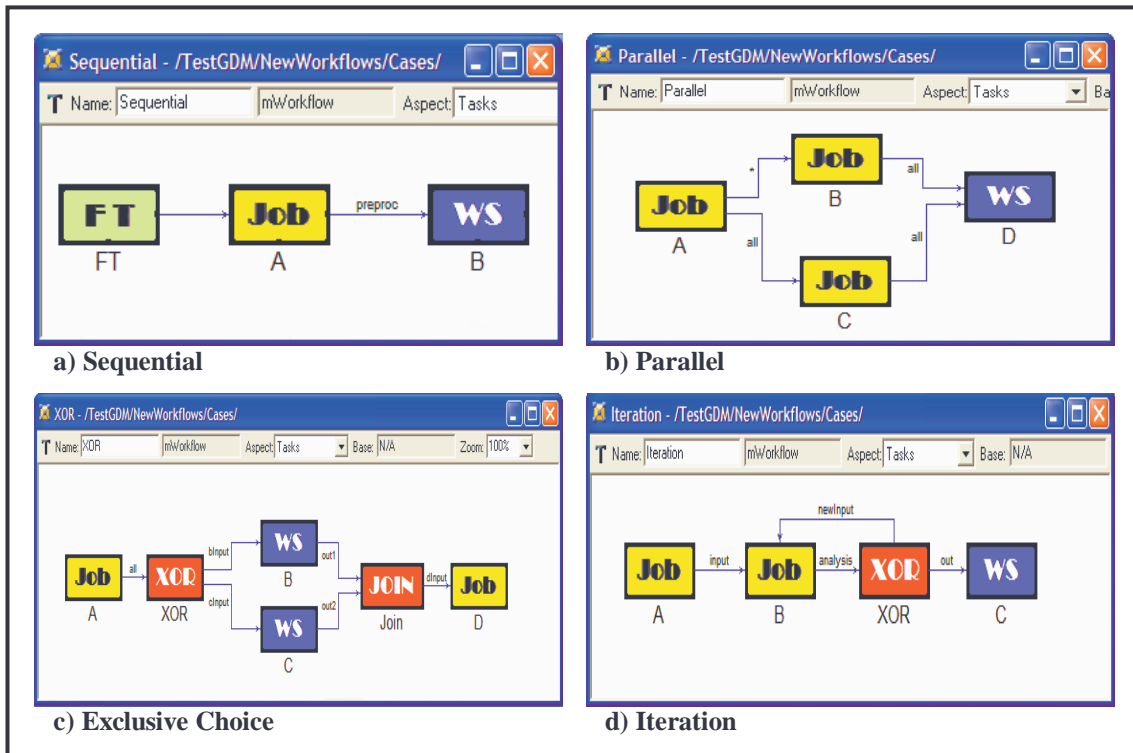


Figure 4. Control constructs supported by GAUGE. a) Sequential execution of tasks. b) AND-Split & Join, both B and C are executed after A. D is executed after B and C have completed their execution. c) XOR-Split & Join. After completion of A, either B or C is executed (but not both). D is executed after the completion of B or C. d) B can be executed many times until the condition in the XOR routes the flow to C.

The XOR-Split is the conditional construct supported in GAUGE. The function of this construct is to select one of several branches based on a decision. Different systems implement this construct differently. In some systems there is a trigger that indicates which branch to choose [4]. The challenge in implementing this construct in GAUGE is that tasks' executions are considered external to GAUGE. Because of this, GAUGE does not have knowledge about the data involved in the execution of the task. The approach taken is to ask the user to analyze the data and trigger the correct path to the workflow engine. The user does this by writing an analyzer program whose output file indicates the correct branch. Figure 4.c shows how to use the XOR constructs in GAUGE. After the execution of Job A, the task inside the XOR construct decides if the next task to execute is Grid service B or C. After the task inside the XOR makes the decision, it signals its choice to the XOR construct which in turn starts the execution of either B or C. Finally, the XOR indicates the branch taken to the Join construct. In this manner, the Join construct knows which of the branches was taken and can start the execution of D when either B or C finishes executing.

Another capability of XOR is that it allows the specification of iterations within a workflow. For example, in Figure 4.d after Job B has completed its execution, the XOR task analyzes the output and decides whether the next task to execute is Grid service C or Job B. In the case that B needs to be repeated, the XOR also provides it with a new input file ("newInput").

The last of the basic control constructs supported by GAUGE is inner workflows, which allow the user to reuse any previously defined workflows. Inner workflows are the basis for constructing hierarchical workflows. Figure 5 presents the way that inner workflows work in GAUGE. *Parallel* is defined as a normal workflow and is later used inside the sequential workflow *Hierarchical*. Any changes to *Parallel* are reflected in each workflow that uses *Parallel* as an inner workflow. This capability is based on the use of prototypes and clones, which is a feature of GME [24].

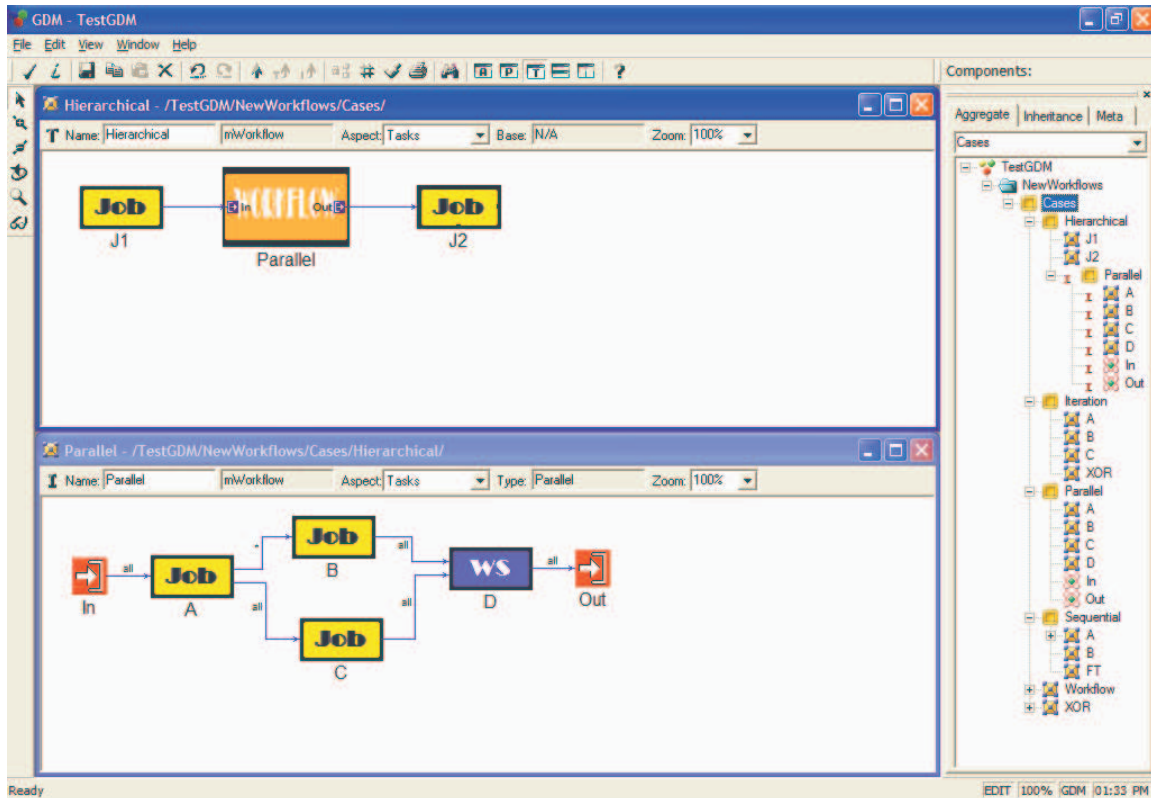


Figure 5. Inner Workflow. Parallel is an inner workflow that is embedded into a sequential workflow.

With these basic constructs, GAUGE is able to generate the following workflow patterns [4]: Sequence, Parallel Split, Synchronization, Exclusive Choice, Simple Merge, and Arbitrary Cycles. These patterns are important because they provide the capabilities required to implement a Grid workflow managing system.

5.2 Domain transformation

The domain transformation layer is responsible for parsing the models created with GDM, and converting the low-level representation of the models (meta-model) into a high-level representation (Grid domain) from which code can be generated. GME provides an API that traverses the internal representation of the GDM models. A model interpreter uses this API to translate the models into an application that manages the execution of the workflow. In our previous work [21], once the workflows are specified, a model interpreter traverses the internal representation of the models and generates the control code that manages the workflow execution. The problem with this approach is

that when working with different output representations (e.g., Java and Python), the process of parsing the models must be repeated for each representation.

The solution to this problem is to provide an intermediate representation from which any code generator can interpret the models. The intermediate layer is implemented by providing a Java API that wraps the API given by GME to access the GDM models. This new intermediate representation is presented at a higher level of abstraction than the one provided by GDM. While interacting with the front-end, the user constructs the models using GDM, which in turn is built using the GME general paradigm (e.g., models, atoms, connections). This layer's API consists of a collection of Java objects that parse the models and translate them from their GME domain representation to a Grid domain representation. Thus, instead of considering models, atoms, and connections, a developer accessing these objects would work with jobs, Grid services, and file transfers. This implementation eases the development of software artifacts for different Grid environments. Thus, a developer can reuse the graphical models and build code generators for specific Grid environments.

5.3 Code generation

The final layer of GAUGE generates code from the visual models. Using the intermediate layer, GAUGE accesses the models and generates the corresponding code that controls the execution of the Grid application. Different code generators can be constructed. We have implemented a Java-based representation and a Java Grid service representation (Initial tests have been constructed using the Systinet WASP server for Java version 4.6.1.; please see <http://www.systinet.com> for more information on the services environment.), both of which use the Java CoG kit to interface with the Globus Toolkit. Currently, we are working on generating two more output representations: a Python [35] script that interacts with the Globus Toolkit using PyGlobus [34], and a GridAnt [25] specification file. It is worthwhile to notice that all of these output representations are based in the same Grid domain models. Thus any other implementation that a developer wishes to construct would be transparent to the end-user as long as the domain-model language (GDM) is maintained.

In the Java applications that are currently supported by GAUGE, data dependence is defined by the sequence of tasks (e.g., if B is defined after A, B waits until A finishes to start execution). In the case of parallel constructs, data dependence is defined at the origin of the split. When the parallel paths merge the AND-Join construct depends on the completion of all the parallel branches (synchronization). Furthermore, when the execution flow splits (AND-Split, or XOR-Split), the merging task is blocked until the parallel tasks finish their execution. At present, the generated applications communicate directly with the Java CoG Kit. This causes scalability problems due to the generation of specific code for each workflow task. A solution to this problem is current under investigation and it consists of developing a reusable workflow engine and generating appropriate configurations from the graphical models.

The infrastructure presented in this section allows developers to augment and extend GAUGE by providing different code generators to tailor the generated code to a particular Grid environment or platform. At the same time it provides an easy to use

automation and generative environment in which an inexperienced user can exploit the Grid without familiarity with the intricacies of the underlying technologies.

6 Related Work

The foundation of GAUGE is to abstract the Grid environment into a high-level layer such that the essence of the workflow is not bound to a specific Grid environment. This high-level layer presents the Grid as a set of familiar concepts so that a user can easily specify Grid applications, i.e., workflows. The specification of tasks and workflows is performed by interacting with GME, which is configured by the GDM visual language. In addition to the Grid constructs, GDM also includes workflow specific constructs (e.g., Splits and Joins). Because the focus of this paper encompasses different areas, we present work on the areas that are relevant to this project.

6.1 Grid Abstraction

Although there are several studies related to this topic (e.g., [23], [39]), the work by Hategan et al., [5] presents a similar work to ours. Hategan proposes a technology and architecture-independent abstraction layer to provide interoperability across multiple Grid implementations, resulting in the Open Grid Computing Environment (OGCE). The main function of OGCE is to serve as a technology-independent, open, and extensible framework for client-side Grid development. The abstractions provided by OGCE are comparable to those introduced by GDM in this paper. For example, the task concepts presented in [5] contain concepts similar to those involved in the GDM job specification (Figure 3). However, the main difference between the studies is in the level of abstraction. While in this paper the abstraction layer is realized at a domain-model level, in [5] the abstraction layer is at a programming language level (Java). There are several advantages of working at the model level, instead of the programming language level:

1. Because the essence of a problem is captured, models are much more amenable to analysis than the lower-level code representation (where the accidental complexities of programming language syntax and semantics often get in the way of higher-level analysis of the problem). A model can be translated easily into input to tools performing a range of checks and analyses.
2. Because the representation is more abstract, it is often the case that a change can be made to a model that would have required many more adaptations at the code level. For those concerns that are crosscutting in nature, a simple change to a model can affect literally hundreds of places in the corresponding source code [19]. Thus, models better facilitate the rapid exploration of design decisions and changes to system configuration.
3. By focusing on the problem at the modeling level, rather than the specific technology used for a solution, the developer and user can protect and isolate themselves from technology obsolescence. That is, the models are not tied to any one solution and can easily adapt to the next generation's hot new technology.

6.2 Workflow Systems

The idea of composing applications from reusable components is not new. For example, WebFlow [1] introduces a platform-independent system that dynamically composes new applications from reusable components by clicking and dragging icons. The Job model of UNICORE [38] uses a set of directed acyclic graphs, and also permits the use of conditional and iterative execution of job groups or tasks. DAGMan [11] also maps a direct acyclic graph specification onto a physical environment. The Symphony framework [29] uses a graphical user interface for rapid collaborative development of grid applications following a data flow paradigm. Triana [37] also offers a visual programming model for the dynamic composition of predefined software components.

Other works propose languages to specify Grid workflows. For example Grid Workflow [7] focuses on proposing a standard for the sequence of complex high-performance computational tasks within a Grid. GridAnt [25] uses an XML base language to specify client-side workflows. GridAnt is also able to submit the executions of tasks or file transfers by using a workflow engine based on the Apache ANT tool [6]. Another set of tools use artificial intelligence to handle the automatic creation of workflows [12], [40]. The workflow based on GAUGE has been influenced by these projects, as well as van der Aalst's studies on workflows [2] and workflow patterns [4], [3].

7 Future Work

The GAUGE project is in its initial phase. The current implementation of GAUGE can handle only a limited number of tasks in the workflow due to the Java output representation. Having created the visual language and the basic three tier infrastructure, we plan to extend the functionality in four different areas: (1) improve the user interaction with the front-end, (2) integrate GAUGE with GIS [10] services, (3) write generators for different output representations (e.g., PyGlobus and GridAnt), and (4) improve the scalability of the generated applications. In addition to this work in progress, future directions that will be considered involve the following aspects:

1. Improve GDM in three ways: (1) extend the language to allow a complete mapping of the GRID, which will allow a better integration with different Grid environments, (2) introduce more workflow control constructs (e.g., OR-Split, and OR-Join), which would permit the implementation of more workflow patterns making GDM a more robust language, and (3) mapping the system to the new OGCE [5] specification.
2. Improve the design of the conditional constructs. The solution that is currently under investigation is the use of AI techniques to automatically execute these kinds of decisions. In particular the use of rule-based systems is being considered. Several tools exist that help to construct rule-based systems. However, because of the particular constraints imposed by the environment (GME), the tool that is being analyzed for this purpose is the Java Expert System Shell (JESS) [17]. JESS's easy integration with programs written in the Java language is the key reason for choosing this tool.

8 Conclusions

This paper presents the architecture and implementation of the Grid Automation and Generative Environment (GAUGE), which enables inexperienced users to take full advantage of the Grid infrastructure. The architecture of GAUGE provides a high-level view for the construction of Grid applications using the Globus Toolkit that shields the intricacies and accidental complexities of the Grid environment.

GAUGE helps the inexperienced Grid users by providing an environment in which they can graphically specify the workflow for their application and automatically generate the code that manages the execution of the workflow (execution and data flow). GAUGE also helps developers by providing an open framework in which they can reuse components and tailor the automation environment to work with different Grid environments.

The novelty of the manner in which GAUGE was designed is that it builds an abstract layer that hides the complexities of the Grid environments. This abstract layer is represented with the visual language GDM. Users manipulate visual models that represent Grid concepts, instead of having to write code for a specific infrastructure and in a specific programming language. GDM was created using domain-specific modeling techniques. The benefits of using domain-specific modeling techniques are:

1. Domain modeling removes the accidental complexities of creating workflows in a Grid by focusing on higher-levels of abstraction at the problem space rather than solution space, such as specific low-level Grid services and their usage.
2. When exploring various workflow scenarios, modeling tools and their interpreters facilitate the more rapid ability to change the workflow details. That is, it is easier to manipulate and change domain models rather than the associated code. This is an important benefit when rapid prototyping of Grid applications is needed.
3. Model-driven techniques possess the ability to generate multiple artifacts from the same model. Thus, with the same domain knowledge different output representations can be generated, which is essential when the Grid application consists of different back-end implementations.

The potential impact of this research and initial implementation is the reduction of the development time involved in generating applications for the Grid. Through GAUGE, users are not required to learn how to use specific Grid technologies to develop Grid-enabled applications. Rather, they construct graphical models that are at a more appropriate level of abstraction for describing the essence of the problem for a specific domain.

GAUGE's use of domain-specific modeling also addresses challenges faced by developers in improving the use of modern software engineering practices for the Grid programming model, and developing applications that interface with different Grid environments. The benefit offered by this work is an abstract high-level Grid model that can serve potentially as a basis for an extended model to completely map the Grid and facilitate the interaction between different Grid technologies and models that are not tied to specific programming languages or Grid frameworks.

References

1. Akarsu, E., Fox, F., Furmanski, W., and Haupt, T., "WebFlow – High-Level Programming Environment and Visual Authoring Toolkit for High Performance Distributed Computing," *Proceedings of Supercomputing 1998*.
2. van der Aalst, W.M.P., and van Hee, K., *Workflow Management: Models, Methods, and Systems*, MIT Press, 2002.
3. van der Aalst, W.M.P, ter Hofstede, A.H.M., Kiepuszewski, B., and Barros, A.P., "Advanced Workflow Patterns," In O. Etzion and P. Scheuermann, editors, *7th International Conference on Cooperative Information Systems (CoopIS 2000)*, volume 1901 of Lecture Notes in Computer Science, pages 18-29. Springer-Verlag, Berlin, 2000.
4. van der Aalst, W.M.P, ter Hofstede, A.H.M., Kiepuszewski, B., and Barros, A.P. "Workflow Patterns," *Distributed and Parallel Databases*, 14(3), pp 5-51, July 2003.
5. Amin, K., Hategan, M., von Laszewski, G., and Zulezec, N., "Abstracting the Grid," *Proceedings of the 12th Euromicro Conference on Parallel, Distributed and Network-Based Processing (PDP 2004)*, 11-13 February 2004, La Coruña, Spain.
6. "The Apache ANT Project." <http://ant.apache.org/>
7. Bivens, H.P., *Grid Workflow*, Grid Computing Environments Working Group Document, 2001.
8. Booch, G., Jacobson, I., and Rumbaugh, J., *The Unified Modeling Language User Guide*, Addison-Wesley, 1998.
9. Bapty, T., Neema, S., and Gray, J., "Model-Integrated Computing For Composition of Complex QoS Applications Using The Generic Modeling Environment (GME)," *OMG Workshop on Real-Time and Embedded Distributed Object Computing*, Washington, DC, July 15-18, 2002.
10. Czajkowski, K., Fitzgerald, S., Foster, I., C., and Kesselman, C., "Grid Information Services for Distributed Resource Sharing," In *Proceedings of the Tenth IEEE International Symposium on High-Performance Distributed Computing (HPDC-10)*, 181-184. IEEE Press, August 2001.
11. "DAGMan (Directed Acyclic Graph Manager)," <http://www.cs.wisc.edu/condor/dagman/>
12. Deelman, E., Blythe, J., Gil, Y., Kesselman, C., Mehta, G., Vahi, K., Blackburn, K., Lazzarini, A., Arbree, A., Cavanaugh, R., and Koranda, S., "Mapping Abstract Complex Workflows onto Grid Environments," *Journal of Grid Computing*, Vol.1, no. 1, 2003, pp. 25-39.
13. Foster, I., Gieraltowski, J., Gose, S., et al., "The Grid2003 Production Grid: Principles and Practice," *Proceedings 13th IEEE International Symposium on High Performance Distributed Computing (HPDC-13)*, 2004.
14. Fox, G., Gannon, D., and Thomas, M., "A Summary of Grid Computing Environments," chapter in Berman, F., Fox, G., and Hey, T., *Grid Computing: Making the Global Infrastructure a Reality*, John Wiley & Sons Ltd, Chichester, 2003.
15. Foster, I., and Kesselman, C., "Globus: A Metacomputing Infrastructure Toolkit," *International Journal of Supercomputer Applications*. 1997; **11**:115-128. <http://www.globus.org>

16. Foster, I., Kesselman, C., Nick, J., and Tuecke, S., "The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration," Open Grid Service Infrastructure WG, Global Grid Forum, June 22, 2002.
<http://www.globus.org/research/papers/ogsa.pdf>
17. Friedman-Hill, E., *Jess in Action: Rule-Based Systems in Java*, Manning Publications Company, 2003.
18. Gray, J., Bapty, T., Neema, S., and Tuck, J., "Handling Crosscutting Constraints in Domain-Specific Modeling," *Communications of the ACM*, October 2001, pp. 87-93.
19. Gray, J., Zhang, J., Lin, Y., Wu, H., Roychoudhury, S., Sudarsan, R., Gokhale, A., Neema, S., Shi, F., and Bapty, T., "Model-Driven Program Transformation of a Large Avionics Product Line Architecture," *Generative Programming and Component Engineering (GPCE 2004)*, Springer-Verlag LNCS, Vancouver, BC, October 2004.
20. Hernández, F., "Domain-Specific Models and the Globus Toolkit," Department of Computer and Information Sciences, University of Alabama at Birmingham, Tech. Rep. UABCIS-TR-2004-0504-1, 2004.
21. Hernández, F., Bangalore, P., Gray, J., and Reilly, K., "A Graphical Modeling Environment for the Generation of Workflows for the Globus Toolkit," to appear in *Workshop on Component Models and Systems for Grid Applications*, 18th Annual ACM International Conference on Supercomputing Saint-Malo, France June 26-July 1, 2004.
22. Haupt, T., Bangalore, P., and Henley, G., "Mississippi Computational Web Portal," in *Concurrency and Computation: Practice and Experience, Special Issue: Grid Computing Environments*, vol. 14, no. 13-15, November-December 2002.
23. Jhoney, A., Kuchhal, M., and Venkatakrisnan, "Grid Application Framework for Java (GAF4J)," IBM Software Labs, India, Tech Rep., 2003.
<https://secure/alphaworks.ibm.com/tech/GAF4J>
24. Karsai, G., Maroti, M., Lédeczi, A., Gray, J., and Sztipanovits, J., "Composition and Cloning in Modeling and Meta-Modeling," *IEEE Transactions on Control System Technology* (special issue on *Computer Automated Multi-Paradigm Modeling*), March 2004, pp. 263-278.
25. von Laszewski, G., Amin, K., Hategan, M., Zaluzec, N. J., Hampton, S., and Rossi, A., "GridAnt: A Client-Controllable Grid Workflow System," *Proceedings of the 37th Hawaii International Conference on System Science*, Jan 5-8, 2004, Island of Hawaii, Big Island.
26. Lédeczi, A., Bakay, A., Maroti, M., Volgyesi, P., Nordstrom, G., Sprinkle, J., and Karsai, G., "Composing Domain-Specific Design Environments," *IEEE Computer*, November 2001, pp. 44-51.
27. von Laszewski, G., Foster, I., Gawor, J., Lane, P., "A Java Commodity Grid Toolkit," *Concurrency and Computation: Practice and Experience*, Volume 13, no. 8-9, pp 643-662, 2001, <http://www.cogkits.org/>
28. Legion Worldwide Virtual Computer. <http://legion.virginia.edu>.
29. Lorch, M., and Kafura, D., "Symphony – A Java-based Composition and Manipulation Framework for Computational Grids," *Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid2002)*, Pages 136-143, May 21-24, 2002, Berlin, Germany.
30. Long, E., Misra, A., and Sztipanovits, J., "Increasing Productivity at Saturn," *IEEE Computer*, August 1998, pp. 35-43.

31. Neema, S., Bapty, T., Gray, J., and Gokhale, A., "Generators for Synthesis of QoS Adaptation in Distributed Real-Time Embedded Systems," *First ACM SIGPLAN/SIGSOFT Conference on Generative Programming and Component Engineering (GPCE '02)*, Springer-Verlag LNCS 2487, Pittsburgh, PA, October 6-8, 2002, pp. 236-251.
32. Novotny, J., "The Grid Portal Development Kit," *Concurrency and Computation: Practice and Experience, Special Issue: Grid Computing Environments*, vol. 14, no. 13-15, November-December 2002.
33. Pearlman, L., Kesselman, C., Gullapalli, S., et al., "Distributed Hybrid Earthquake Engineering Experiments: Experiences with a Ground-Shaking Grid Application" *Proceedings 13th IEEE International Symposium on High Performance Distributed Computing (HPDC-13)*, 2004.
34. "Python Globus (pyGlobus)," <http://dsd.lbl.gov/gtg/projects/pyGlobus/>
35. van Rossum, G., and Drake, F., *The Python Language Reference Manual*, Network Theory Ltd, 2003.
36. "The Globus Resource Specification Language RSL v1.0," http://www.globus.org/gram/rsl_spec1.html.
37. "Triana Workflow," <http://www.triana.co.uk>
38. "Unicore," <http://www.unicore.de/>
39. Wolski, R., Brevik J., Obertelli, G., Sprong, N., and Su, A., "Writing Programs that Run EveryWare on the Computational Grid," *IEEE Transactions on Parallel and Distributed Systems*, vol. 12, no. 10, pp. 1066-1080, October 2001.
40. Yeo, B., Hung, T., and Khoo, B., "An Agent-based Grid Flow Management Framework for the Problem Solving Environment (PSE)," in *Scientific Workflow Management Mini-Symposium*, GlobusWorld, 2004.