

# Building Scientific Workflows for Earth System Modelling with Windows Workflow Foundation

Matthew J. Fairman<sup>1</sup>, Andrew R. Price<sup>1</sup>, Gang Xue<sup>1</sup>, Marc Molinari<sup>1</sup>,  
Denis A. Nicole<sup>2</sup>, Timothy M. Lenton<sup>3</sup>, Robert Marsh<sup>4</sup>,  
Kenji Takeda<sup>1</sup>, and Simon J. Cox<sup>1</sup>

<sup>1</sup>Microsoft Institute for High Performance Computing,  
School of Engineering Sciences, University of Southampton, Southampton, UK  
{mjf, andrew, gx, mm, ktakeda, sjc}@soton.ac.uk

<sup>2</sup>School of Electronics and Computer Science, University of Southampton, UK  
dan@ecs.soton.ac.uk

<sup>3</sup>School of Environmental Sciences, University of East Anglia, Norwich, UK  
t.lenton@uea.ac.uk

<sup>4</sup>National Oceanography Centre, University of Southampton, Southampton, UK  
rma@noc.soton.ac.uk

**Abstract.** The GENIE project has built a Grid enabled framework that facilitates the integration, execution and management of component models for the study of the Earth system over millennial timescales. The existing framework supports collaborative study of GENIE models across heterogeneous compute grids through scripted workflows in the Matlab environment. In this paper, Windows Workflow Foundation technologies are applied to demonstrate the benefits of an environment that provides rapid composition, event driven logic and reliable hosting of scientific workflows. These improvements are demonstrated through a parametric study of bi-stability of the oceanic thermohaline circulation in a GENIE model.

**Keywords:** Workflow, GENIE, Windows Workflow Foundation.

## 1 Introduction

E-Science applications are usually composite systems that involve collaborations among various computational resources and software components in a distributed, heterogeneous environment. Workflow technologies have become a key element to e-Science systems, which help to orchestrate interactions with services so that they are seamlessly knitted together to implement the desired behaviour of the system.

A typical example of such an e-Science system can be found in GENIE [1]. The GENIE project has developed a framework for the composition, execution and management of integrated Earth system models. Component codes (e.g. ocean, atmosphere, land surface, sea-ice, ice-sheets, biogeochemistry, etc.) of varying resolution and complexity can be flexibly coupled together to form a suite of efficient climate models capable of simulation over millennial timescales. The project brings

together a distributed group of environmental scientists with a common interest in developing and using GENIE models to understand the Earth system. Grid computing technology supports the virtual organization to collaboratively study the models from the framework.

Earth system simulations are both computationally and data intensive processes, often consisting of a discrete series of steps. The GENIE framework has been designed to support running of such simulations across multiple distributed data and computing resources over a lengthy period of time. It has therefore put rigorous requirements on the selection of workflow technologies for GENIE.

The recently released Microsoft Windows Workflow Foundation (WF) [2] is a general, extensible framework for developing workflow solutions. As an integral part of the .NET Framework 3.0 [3], WF provides a common development model for creating workflows, a set of tools that facilitate workflow design, and a robust workflow hosting environment. It also allows seamless integration with other .NET technologies to provide support for distributed communication and rich user experiences. We have therefore applied WF to the GENIE framework to provide a solid base for the design and running of Earth system simulations.

In this paper we discuss and demonstrate how the WF technologies are applied to build efficient and flexible workflows for running Earth system simulations based on the GENIE framework. In section 2 we describe the functions of the GENIE workflows and explain why the attempt has been made to adopt WF technologies. Section 3 gives a detailed summary of the features of WF. Section 4 describes the design and running of our WF based GENIE workflow. Results from a scientific study of bi-stability in the thermohaline circulation (THC) in a GENIE model, undertaken using our WF based system, are presented in section 5. We draw our conclusions in section 6.

## 2 Workflow in the GENIE Framework

The GENIE framework has been designed to facilitate collaboration between a distributed team of environmental scientists to perform large ensemble studies of Earth system models that would be impossible for a single user. In contrast to system likes GridSolve [4] and Nimrod/G [5], in which a central manager controls all available resources and makes decisions about where to assign work units, the GENIE framework devolves responsibility for the management of work units to the client side, and pools compute resources contributed by members of the project to carry out large parametric studies. A central shared database accessible from all clients has been set up to provide the means to upload model configurations and experiment definitions, and be queried against to find available work units.

A GENIE workflow running at the client side is therefore responsible for the following tasks when contributing to a particular study:

- Querying the central database for available work units.
- Submitting the work units to the compute resources specified by the user.

- Monitoring the status of the submitted jobs; retrieving and checking the results once jobs are finished.
- Post-processing the results and saving them in the database.

The workflow runs until there is no work available, or it is terminated by the user.

Presently, the GENIE workflow is implemented using the Matlab scripting language and is hosted in the Matlab environment. While simplicity and flexibility are achieved through scripting, a few problems remain. The scripting approach for workflow tends to be less efficient when dealing with external events. When running a GENIE simulation, the user needs to keep running the script all the time and query the central database periodically to find work items, submit and monitor jobs, and post-process data from completed work items. Time intervals between two scheduled tasks can only be set in a speculative mode. Events such as new work units being added and compute tasks completing are unlikely to be handled efficiently. Callback mechanisms or publish/subscribe mechanisms would be an optimal choice for improvement. For example, the database could publish new work as it becomes available and the clients would subscribe for these events. However, it is difficult to make such modifications in the current scripting environment.

Another problem with hosting workflows in the scripting environment is the lack of support for persisting state of the workflow. While it is possible for the users to manually save some of the state information, restoring the running of a workflow in the cases of a system or network failure is always difficult. The GENIE system is robust to such failures since the experiment status is always defined in the shared repository but resumption of a study can only be achieved by restarting the workflow.

Finally, implementing the GENIE workflow in the Matlab script ties the system to a single user environment. It is very difficult to migrate existing workflows to other systems. Workflows in scripting languages are also less well-structured and might require individual human knowledge to interpret or modify. This makes it hard to maintain or to try to reuse the workflow activities.

As an attempt to address these problems, we have exploited the WF technologies to develop a dynamic and robust workflow implementation which actively manages the operations of the GENIE simulations with improved efficiency and stability.

### 3 Windows Workflow Foundation

Windows Workflow Foundation (WF) [2] is a Microsoft technology for quickly building workflow enabled applications. It consists of a programming model, a host engine and tools for defining, executing, and managing workflows. WF significantly enhances a developer's ability to model and support sophisticated science and engineering processes. Being part of the recently released .NET Framework 3.0, WF can be seamlessly integrated with other .NET technologies to provide a solid base for the design and running of Earth system simulations.

#### 3.1 Workflow Design with WF

WF supports two built-in workflow styles, i.e. sequential workflows, which carry out a series of operations in a pre-defined pattern, and state machine workflows, which

are driven by external events. This enables WF to accommodate both the well structured, automated system workflows and the loosely defined, flexible human workflows. WF adopts a component like approach, where each step of a workflow can be implemented by a specific piece of software called ‘activity’. A base activity library is supplied which contains a group of pre-defined activities most commonly used in both sequential and state machine workflows. Users are also allowed to build their own customised activities, build new activities by aggregating existing ones, and reuse activities in different workflows.

WF is a general framework not limited to a single language or a single tool. Workflows can be created using any programming language that is compliant with the .NET Framework’s Common Language Specification, such as C#, Visual Basic and Visual C++. In addition, WF supports defining workflows in XAML (eXtensible Application Markup Language) [6], an XML based object initialization language, so that different workflow designers can be used to generate or parse WF workflows. WF also provides its own graphical designer, usually hosted in Visual Studio 2005, to support visual design of workflows.

### **3.2 Workflow Hosting with WF**

WF workflows can be hosted in any Windows process ranging from simple console applications to large, sophisticated enterprise services. The workflow execution is carried out within the WF runtime engine, which provides important services, such as persisting workflow state, tracking the execution, and supporting transactions. The WF runtime can make decisions to unload long-running workflows that have been inactive for a period of time, and load them back when necessary. It also works as a proxy for all interactions between the workflow and external software, including Web services, in an asynchronous style.

Being an integrated part of the .NET Framework, it is possible for WF workflows to take advantage of other Microsoft technologies including Windows Communication Foundation (WCF) [7], Windows CardSpace [8] and MS SQL Server 2005 [9] to address problems such as communication in a distributed environment, federated security control and data access. Application of some of these technologies together with WF can be found in our work on the Earth system modelling workflows, which is discussed in the following section.

## **4 A Composite Workflow for GENIE Simulations**

A number of key components have been developed for the GENIE project to facilitate the running of Earth system simulations. These include a shared database hosted in Oracle 10g which contains the definitions, states and results of the experiments, a number of gateway systems to integrate compute resources from different HPC platforms including Microsoft Cluster Computing Server [10], Condor [11] and Globus [12], and a shared file store managed by using the Geodise Toolkit [13]. In order to bring these components together in a single system that manages the Earth system simulations, a composite workflow has been implemented using the WF technologies. Figure 1 shows the structure of the WF based Genie system.

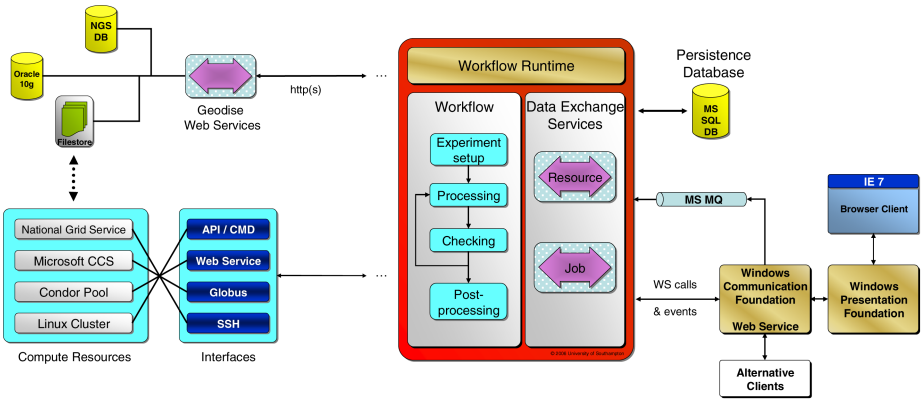


Fig. 1. The Complete Structure of the WF based GENIE Simulation System

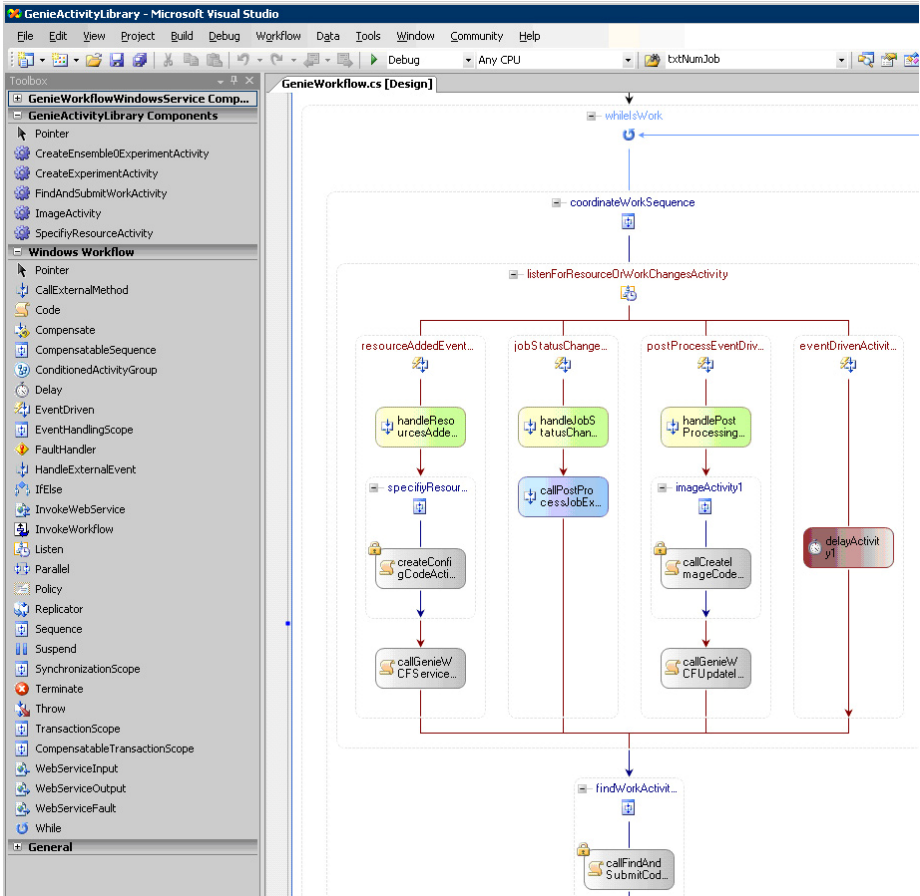


Fig. 2. Design of the GENIE workflow in the WF Designer with Visual Studio

The Genie WF workflow implementation automates the operations and interactions on the system components. It is a sequential workflow that consists of a while loop activity which contains a single parallel asynchronous event listening activity that waits for the following changes in the state of an experiment: newly specified resources, changes of job states and completion of work item post-processing. This not only provides a more responsive solution than the scripting based approach of polling, but also achieves better efficiency and robustness of the system by allowing the runtime to persist workflow instances in a database and unload them during periods of idleness. All of the asynchronous events are indications that either compute resources have become available or subsequent work items are now available. Consequently, a find work activity is performed after any one of these events has occurred. Figure 2 shows the main part of the design of the GENIE workflow.

The Genie workflow employs a set of external data exchange services that run within the workflow runtime. These services provide the means by which events occurring within the experiment system are passed into the workflow. When compute jobs are submitted, the workflow registers their identifiers with the Job Service which is responsible for monitoring the jobs' running statuses. The Job Service notifies the workflow instance of changes to the job status by firing a batched 'jobStatusChanged' event. In addition, the Job Service also performs asynchronous post-processing of results from completed jobs when the workflow receives a 'jobStatusChanged' event of either 'Completed' or 'Failed'. Another external data exchange service, the Resource Service, allows the users to instruct the workflow to create new experiment work units by registering additional compute resources with it.

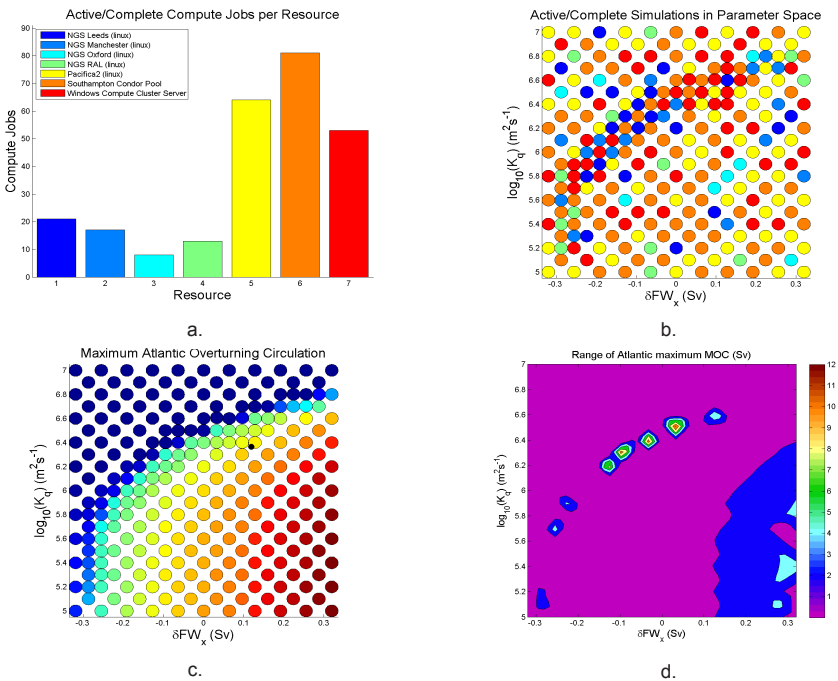
## 5 Earth System Simulation with WF

To demonstrate the WF based GENIE system, we repeat a systematic 2-parameter study of bi-stability in the Atlantic thermohaline circulation [14] using a higher resolution version of the GENIE model. The study investigates the C-GOLDSTEIN composition referred to as *genie\_ea\_go\_gs* in the GENIE framework nomenclature, which consists of a 3D frictional geostrophic ocean coupled to a simple 2D energy-moisture balance model and a 2D sea-ice component. The study seeks to investigate the properties of the Atlantic Thermohaline Circulation as a function of two of the key parameters that dominate atmospheric moisture transport in the model:

- $FW_x$ : Zonal moisture transport from the Atlantic to the Pacific.
- $K_q$ : Atmospheric moisture diffusivity - principally controls meridional moisture transport between the equator and poles.

A systematic sweep over a range of values for these parameters was undertaken. Each simulation was integrated over 4000 years to ensure an equilibrium state was achieved. The individual model runs each required approximately 2 hours compute time on a typical desktop. With the benefit of knowledge of the behaviour of the model at a lower resolution we chose to perform a coarser sampling of the parameter space (11x11 as opposed to 31x31) than the original study. The initial parameter sweep therefore required approximately 242 hours of compute time.

The WF based GENIE system has been used to repeat the study of Marsh et al. using a 16-level seasonal version of the *genie\_ea\_go\_gs* model. Having provided a coarse mapping of the behaviour of the Atlantic THC in the model, the study proceeded by performing two further identical parameter sweeps using models with different boundary conditions. These secondary ensembles performed identical model runs initializing the models using particular end-states from the initial ensemble. Such runs are designed to investigate whether there are points in the parameter space of the atmospheric water transport properties for which the equilibrium THC state of the model is dependent upon its initial conditions. In Marsh et al. nine end-states were taken from the extremes of the parameter space to search for bi-stability. In this study we perform just two secondary ensembles using the end-states from the “top-left” and “bottom-right” points (Figure 3 (c)) which represent the extreme “off” and “on” states in the THC respectively. Following the coarse initial mapping the parameter space was further refined through the addition of another 100 (10x10) members to each ensemble. An interactive plot of the results was then exploited to add further ensemble members in the emerging regions of bi-stability in the THC to focus further work in the most appropriate locations.



**Fig. 3.** Study of bi-stability in the Atlantic Thermohaline Circulation. a) Summary of distribution of compute jobs across the computational grid, b) distribution of compute jobs in the parameter space, c) Strength of maximum Atlantic Overturning Circulation in the initial parameter sweep and d) narrow region of bi-stability in the parameter space.



The results of the bi-stability study are presented in Figure 3. Figure 3 (a) summarises the resource usage in performing the initial ensemble. The distribution of compute jobs broadly reflects the relative numbers of single concurrent compute tasks that the schedulers on each system allow to a single user. The distribution of simulations across the resource pool is plotted in the parameter space in Figure 3 (b). The user additions of model runs are focused along the “cliff-edge” of THC collapse (Figure 3 (c)) where a narrow region of bi-stability emerges (Figure 3 (d)). It is evident that further runs will be required to fully resolve the extent of the bi-stable region. First indications are that the newly tuned 16-level model has a narrower region of bi-stability than the original 8-level model [14]. A more detailed study of the parameter space and analysis of the results will be the subject of future work.

## 6 Summary

In this paper we discussed and demonstrated how the Windows Workflow Foundation technologies can be exploited to implement scientific workflows for e-Science systems such as GENIE. We reviewed features of WF technologies and explained how they have been deployed in the GENIE framework to facilitate coordinated operations among various distributed components. A collaborative study of bi-stability in the thermohaline circulation of an efficient climate model has been performed to demonstrate the running of the WF based GENIE system.

## References

1. Gulamali, M.Y., et al., GENIE: Delivering e-Science to the environmental scientist. UK e-Science All Hands Meeting, pp. 145-152, Nottingham, UK, 2003. ISBN 1-904425-11-9.
2. Shukla, D., Schmidt, B., Essential Windows Workflow Foundation. Addison Wesley Professional. ISBN-10: 0-321-39983-8
3. Microsoft .NET Framework 3.0. <http://msdn2.microsoft.com/en-gb/netframework/>
4. YarKhan, A., Seymour, K., Sagi, K., Shi, Z. and Dongarra, J., Recent Developments in Gridsolve. Int. J. of High Performance Computing Applications, 20 (2006), 131-141
5. Buyya, R., Abramson, D. and Giddy, J. Nimrod/G: An Architecture of a Resource Management and Scheduling System in a Global Computational Grid, HPC Asia 2000, May 14-17, 2000, pp 283 289, Beijing, China.
6. Microsoft, XAML Overview. <http://msdn2.microsoft.com/en-us/library/ms752059.aspx>
7. Microsoft, Windows Communication Foundation Architecture Overview.
8. <http://msdn.microsoft.com/library/en-us/dnlong/html/wcfarch.asp>
9. Chappell, D., Introducing Windows CardSpace. <http://msdn.microsoft.com/library/en-us/dnlong/html/introinfocard.asp>
10. Microsoft, SQL Sever 2005 Books Online. <http://www.microsoft.com/technet/prodtechnol/sql/2005/downloads/books.msp>
11. Microsoft, Windows Compute Cluster Server 2003 Product Overview.
12. <http://www.microsoft.com/windowsserver2003/ccs/overview.msp>
13. The Condor Project. <http://www.cs.wisc.edu/condor/>
14. The Globus Project. <http://www.globus.org>
15. The Geodise project. <http://www.geodise.org>
16. Marsh, R., *et al.*, Bistability of the thermohaline circulation identified through comprehensive 2-parameter sweeps of an efficient climate model. Clim. Dyn. 23 (2004) 761-777