

# Automatic Generation of Workflow Provenance

Roger S. Barga<sup>1</sup> and Luciano A. Digiampietri<sup>2</sup>

<sup>1</sup> Microsoft Research, One Microsoft Way  
Redmond, WA 98052, USA

<sup>2</sup> Institute of Computing, University of Campinas,  
Sao Paolo, Brazil  
barga@microsoft.com

**Abstract.** While workflow is playing an increasingly important role in e-Science, current systems lack support for the collection of provenance data. We argue that workflow provenance data should be automatically generated by the enactment engine and managed over time by an underlying storage service. We briefly describe our layered model for workflow execution provenance, which allows navigation from the conceptual model of an experiment to instance data collected during a specific experiment run, and back.

## 1 Introduction

Many scientific disciplines today are data and information driven, and new scientific knowledge is often obtained by scientists scheduling data intensive computing tasks across an ever-changing set of computing resources. Scientific workflows represent the logical culmination of this trend. They provide the necessary abstractions that enable effective usage of computational resources, and the development of robust problem-solving environments that marshal both data and computing resources. Part of the established scientific method is to create a record of the origin of a result, how it was obtained, experimental methods used, the machines, calibrations and parameter settings, quantities, etc. It is the same with scientific workflows, except here the result provenance is a record of workflow activities invoked, services and databases used, their versions and parameter settings, data sets used or generated and so forth. Without this provenance data, the results of a scientific workflow are of limited value.

Though the value of result provenance is widely recognized, today most workflow management systems generate only limited provenance data. Consequently, this places the burden on the user to manually record provenance data in order to make an experiment or result explainable. Once the experiment is finished, the specific steps leading to the result are often stored away in a private file or notebook. However, weeks or months may pass before the scientist realizes a result was significant, by which time provenance has been forgotten or lost – it simply slips through the cracks.

We believe the collection of provenance data should be automatic, and the resulting provenance data should be managed by the underlying system. Moreover, a robust provenance trace should support multiple levels of representations. In some cases, it is suitable to provide an abstract description of the scientific process that took place, without describing specific codes executed, the data sets or remote services

invoked. In other cases, a precise description of the execution of the workflow, with all operational details such as parameters and machine configurations provided, is exactly what is required to explain a result. Given no single representation can satisfy all provenance queries, the system should generate multiple related representations simultaneously. The ideal model will allow users to navigate from the abstract levels into lower levels and map the latter backwards to higher level abstractions. This also allows scientists to specify exactly what they wish to share from their experiment.


In this paper, we suggest the workflow enactment engine should be responsible for generating workflow provenance *automatically* during runtime, and an underlying storage service should be responsible for managing workflow provenance data. Provenance collection and management should happen transparently. That is, users should not have to take any special actions or execute special commands to have provenance of their workflow execution collected and maintained. Indeed, unless users take extraordinary action, such as purging an experiment result, an accurate provenance record will be maintained for results residing on their system. By making the enactment system responsible for the creation and collection of execution provenance, we are able to generate provenance automatically, freeing users from having to manually track provenance during execution. Equally important, when properly collected and maintained, workflow provenance data can be indexed and queried as a *first class data product* using conventional data management tools and techniques. In this paper, we also outline a *layered* or *factored* model for representing workflow execution provenance, from an abstract description of the workflow (scientific process) to a precise description of execution level details for a single experiment, which enables provenance data to be linked and defined in a way for more effective discovery, integration and cooperation.

## 2 Workflow Execution Provenance Model

In this section we outline the machine-readable model we propose to represent provenance data captured during workflow execution. The attractions of creating a machine-readable model of an experiment, that is, a workflow, are not difficult to envisage. Three key benefits that a *multilayered model* can bring are explanation, validation, and insight. Having an abstract model of the science being undertaken, and being able to use the model to interpret behavior (services and data sets used), together make it possible to explain or reason about the science. For validation, knowing which individual activities were executed, identifying branches taken during execution, steps skipped or inserted, and specific parameters supplied at runtime is essential when trying to establish trust in a result. The ability to compare a set of related experiment executions against an abstract workflow model, to identify common patterns or options not yet explored, can be used to gain insight in authoring new experiment designs. A related benefit, not discussed in this short paper is optimization – a layered or *factored* model can expose opportunities to efficiently store provenances traces in the underlying storage manager.


The first level in our model, illustrated in Figure 1 below, represents an abstract description of the scientific process. This layer captures *abstract activities* in the workflow and *links/relationships* among the activities or ports of activities, where an

abstract activity is a *class* of activities. These classes can be described by the function of the activity (e.g., `InvokeWebService`, or `Sequential`, `Parallel Branch` and `While` activities, etc) or through semantic description – for example, a class of *alignment* activities, where `Blastx` and `FASTA` are instances that belong to this class.




**Fig. 1.** Level L0, represents abstract service descriptions

The second level in our model, illustrated in Figure 2 below, represents an instance of the abstract model. This layer captures bindings or *instances of activities* and additional *relationships*, as classes of activities are instantiated. This level refines abstract activities specified in the previous level with a specific instance (for example, `InvokeWebService` is bound to a specific activity instance `InvokeBlastWebService` and `Alignment tool` is bound to `Blastx tool`). The additional information captured at this level is required to assign *fields of the classes* of the previous level. For example, the `InvokeWebService` has fields `ProxyClass` and `MethodName` that should be filled, with a proxy where the web service can be found and name of the method to be invoked.



**Fig. 2.** Level L1, represents service instantiation descriptions

In the next level of the model L2, illustrated in Figure 3, represents the execution of the workflow instance specified in level L1. From level L1 to level L2 the model captures information provided at runtime, such as *parameters* and properties that complete the specification of activities, including input data, runtime parameter supplied, activities inserted or skipped during execution, etc. Information captured at



**Fig. 3.** Level L2, represents data instantiation of the executable workflow

this level is sufficient to trace the execution of the workflow, from input parameters to individual activities executed and runtime parameters supplied to each activity.

The final level of our model L3, illustrated in Figure 4 below, represents runtime specific information. Here we capture operational specific details, such as the start and end time of the workflow execution, the start and end time of the individual activity execution, status codes and intermediary results, information about the internal state of each activity, along with information about the machines to which each activity was assigned for execution.




Fig. 4. Level L3, run-time provenance for a workflow execution

At each level of our model, we extend the amount of detail and amount of information of the previous level. The linkage between levels is facilitated by three types of operations: *extend* (class extension); *instance* (class instantiation) and *import* (reference to a specific model). Resuming, in level 0 the user designs a class from a model that is extended in level 1. In the level 2, the class of level 1 is instantiated. Level 3 imports the model of level 2 to make references to its entities.

### 3 Basic Workflow Execution Provenance Sample

In this section we illustrate the information that can be captured in our model using a simple example. Our implementation is based on the Windows Workflow Foundation

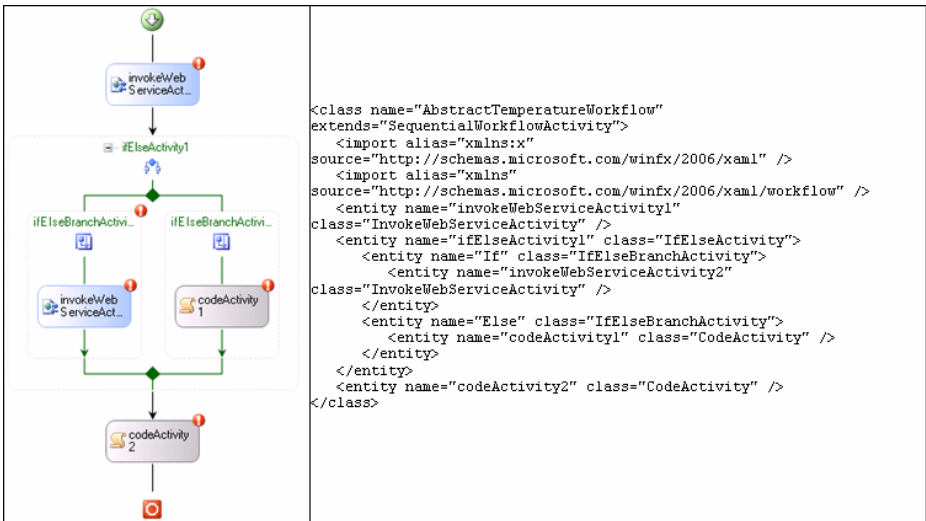


Fig. 5. Workflow level L0

(WinFX) [1], an extensible framework for developing workflow solutions. WinFX has a predefined set of activities (If-Else, While, Parallel, InvokeWebService, etc) and allows for user-defined custom activities by object inheritance from base classes.

In level L0 the user defines type of workflow (*Sequential*, *StateMachine* or *Rule Driven*) she wants to construct and identifies classes of activities available for use. Fig. 5 shows the graphical representation of the workflow and corresponding code. The data stored in the provenance for this level supports general queries, such as: what experiments in my collection use web services? (*invokeWebServiceActivities*) or what experiments utilize alignment tools? This level, in general, describes the design space for all workflows that are an instance of this abstract model.

In level L1 the user replaces abstract activities with concrete activities, and enters fields pre-defined for the individual abstract activities. For example, the activity *invokeWebServiceActivity1* of Figure 5 was specialized to *captureTemperatureWS* in Fig. 6 and the property *ProxyClass* was filled with the desired proxy. Another property that was filled in was *MethodName*, with the name of the method from the Web Service the user wants to invoke. The provenance model of L2 models a specific instance of the abstract workflow model that is being prepared for execution.

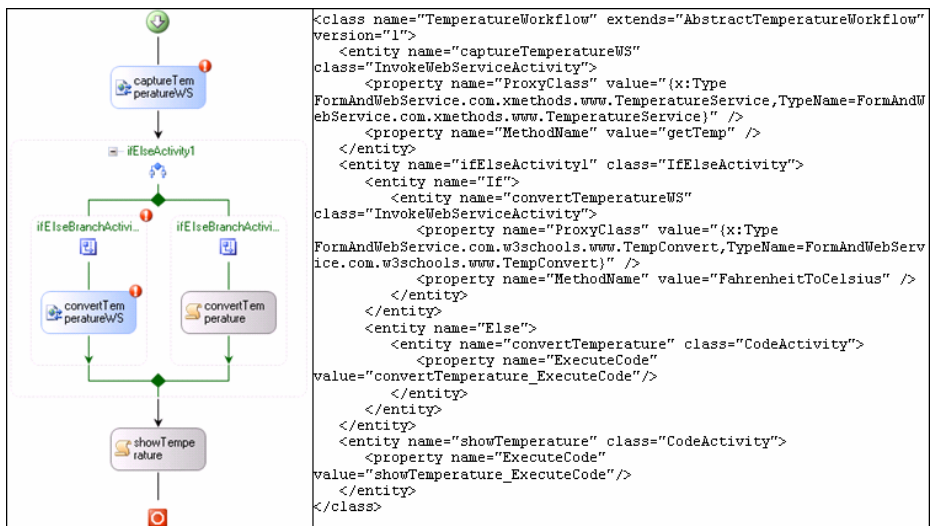


Fig. 6. Workflow level L1

From the provenance in Level L1 the user can pose queries about specific activities: What experiments (workflows) used the web service *TemperatureService*? Or, what experiments invoke the following method of a service, or what experiments use the same alignment tool?

In the level L2, which is the executable workflow, all data sets are identified and data parameters are supplied so that the workflow is ready to be executed. One example of parameter that was filled is *zipcode* from *captureTemperatureWS* that has the value “98007” (see more details in Fig. 7). In this level of the model, we can

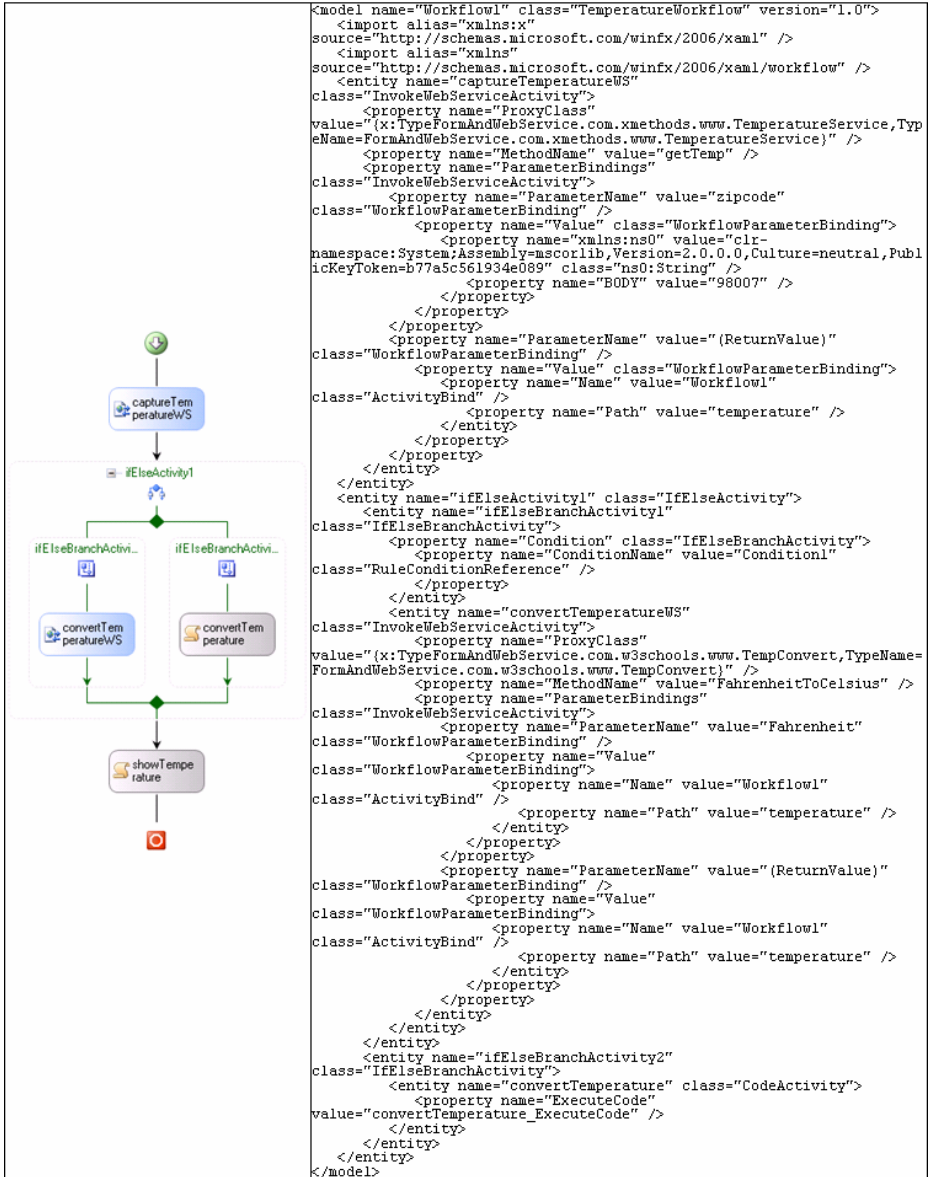


Fig. 7. Workflow level L2

pose a number of queries which do not involve runtime information. These queries can involve parameters, activities and abstract activities/structures. What workflows used the value “98007” in *zipcode* parameters? What was the precondition for the workflow follow the path *If* in the *IfThenElseActivity1*?

The last level (runtime level) can contain a great diversity of information. The kind of information to be stored will depend of the goals of the user. For example, in a system that utilizes a high performance Grid, the user may wish to know what activity (job) was executed in each processor of the Grid and when this execution occurred. In other systems, the user may wish to store information about the internal status of each activity. These two examples are allowed by our provenance model, but require changes in the applications to work (in the Grid example, the scheduler must send information about the process to the provenance system, and in the second example each activity must send information to the provenance system).

Fig. 8 shows a simple runtime record from the Executable Workflow of Fig. 7. In this example, only information about start and end time of the workflow and of the activities were stored and information about produced data.



Fig. 8. Workflow level 3

Note that in workflows designed using the Windows Workflow Foundation system there are no ports/channels to data exchange. This exchange is made via workflow global variables. Our provenance model allows for this information (and information about the code of each *CodeActivity*) be stored inside the workflow model (in level L2). The rules of the workflow also can be stored in the *ExecutableWorkflow* level. Fig. 9 shows the Code and Rules for the workflow *ExecutableWorkflow* in Fig. 7. Together, these are sufficient to recreate the original workflow specification.





required to validate a result. A multilevel representation would permit scientists to specify exactly what they wish to share, or retain, from their experiment record.

The work presented here is very much an initial investigation in an area of growing importance for e-Science. Our layered provenance model for workflow execution is a prototype to be used as a tool to explore user requirements, and to consider system requirements for managing provenance data. We envisage this provenance data to be a *first class* data resource in its own right, allowing users to both query experiment holdings to establish trust in a result and to drive the creation of future experiments.

## 5 Related Work

The challenge of workflow execution provenance is growing in both importance and as a research topic. Most existing provenance systems available today can be classified into one of two broad categories: i) Flow based provenance systems and ii) Annotation based provenance systems.

**Flow Based Provenance Systems:** In many existing Grid/Service workflow based systems, the provenance recorded is generally a variant of the following: when a service is invoked, the workflow manager stores the input, output and the service name in a provenance repository. Also, logically related service invocations are grouped by using a common ID. One disadvantage with this schema is that it can be sub-optimal for running provenance based queries, and precludes the ability to compress or reuse provenance data records. And, if not properly designed it can double the depth of the provenance trace, as traversals first have to lookup the process, then lookup inputs to the process to get to the parent file. Systems that have a workflow based provenance include: PASOA [2], Chimera [3], and myGrid [4].

**Annotation Based Provenance Systems:** Systems such as the storage resource broker (SRB) [5] and GeoDise [6] store provenance as name-value attribute pairs. While this schema can be used to build provenance trace, the construction again can be very inefficient for the same reasons as in workflow based provenance systems, and it does not lend itself to a layered model.

## References

- [1] Windows Workflow Foundation (WinFX), <http://msdn.microsoft.com/workflow/>.
- [2] PASOA, <http://twiki.pasoa.ecs.soton.ac.uk/bin/view/PASOA/WebHome>.
- [3] I. Foster, J. Voeckler, M. Wilde, and Y. Zhao. The Virtual Data Grid: A New Model and Architecture for Data-Intensive Collaboration. In CIDR, January. 2003.
- [4] J. Zhao, M. Goble, C. Greenwood, C. Wroe, and R. Stevens. Annotating, linking and browsing provenance logs for e-science.
- [5] M. Wan, A. Rajasekar, and W. Schroeder. Overview of the SRB 3.0: the Federated MCAT. <http://www.npaci.edu/DICE/SRB/FedMcat.html>, September 2003.
- [6] S. Cox, Z. Jiao, and J. Wason. Data management services for engineering. 2002.