Int J Digit Libr (2007) 7:99–107
DOI 10.1007/s00799-007-0018-5

REGULAR PAPER

Provenance Explorer-a graphical interface for constructing scientific publication packages from provenance trails

Jane Hunter · **Kwok Cheung**

Published online: 13 July 2007

Abstract Scientific communities are under increasing pressure from funding organizations to publish their raw data, in addition to their traditional publications, in open archives. Many scientists would be willing to do this if they had tools that streamlined the process and exposed simple provenance information, i.e., enough to explain the methodology and validate the results without compromising the author's intellectual property or competitive advantage. This paper presents *Provenance Explorer*, a tool that enables the provenance trail associated with a scientific discovery process to be visualized and explored through a graphical user interface (GUI). Based on RDF graphs, it displays the sequence of data, states and events associated with a scientific workflow, illustrating the methodology that led to the published results. The GUI also allows permitted users to expand selected links between nodes to reveal more fine-grained information and sub-workflows. But more importantly, the system enables scientists to selectively construct "scientific publication packages" by choosing particular nodes from the visual provenance trail and dragging-and-dropping them into an RDF package which can be uploaded to an archive or repository for publication or e-learning. The provenance relationships between the individual components in the package are automatically inferred using a rules-based inferencing engine.

Keywords eScience · Provenance · Visualization · Inferencing · Publications

J. Hunter (✉)
ITEE, The University of Queensland, St. Lucia, QLD, Australia
e-mail: jane@itee.uq.edu.au

K. Cheung
AIBN, The University of Queensland, St. Lucia, QLD, Australia
e-mail: kwokc@itee.uq.edu.au

1 Introduction and objectives

Digital library researchers have tended to concentrate on technologies to support digital objects at the scholarly publishing and e-learning end of the research chain, rather than the raw data being produced at the beginning of the chain. However the emerging eScience infrastructure is laying the foundation for new forms of intellectual products that require new modes of curation, publication and collaborative interaction. Already, scientific communities and their funding bodies, are talking about the need for scientists to publish their raw data sets, experimental details, analytical methods and visualizations, in addition to the traditional scholarly publications. A record of the complete scientific discovery process enables peers to review the method of conducting the science as well as the final conclusions. It also enables greater sharing, re-use and comparison of scientific results, reduces duplication and insures against data loss because the additional contextual and provenance information enhances the repeatability and verifiability of the results.

However these new information formats present significant challenges to digital library researchers, who are used to dealing with file-based digital objects. Our aim is to provide a system that enables scientists to easily construct "scientific publication packages" that link the raw data to the derivative products and final publications within a single composite digital object which can be uploaded to an open archive or institutional repository. The associated provenance information will be provided through the RDF relationships which are either explicitly recorded or inferred using a rule-based inferencing engine. By enabling the scientists to interactively select the specific workflow components to be published, they can hide those components which add unnecessary complexity or which they want to keep private or confidential in order to protect their intellectual property.

In order to achieve our objectives, we assume the availability of a number of pre-existing components that underpin the system. In particular we require:

1. A provenance collection system capable of recording the complete set of processing or experimental steps in the physical world (i.e., the laboratories or processing plants) as well as the data post-processing, analysis, visualization and derivation steps in the digital domain.
2. Semantic descriptions of provenance information that has been recorded in RDF, using a common underlying data model or ontology. If the sequence of events, inputs and outputs is recorded using a machine-processable model (expressed in RDF/OWL) then many of the relationships between the individual components are either explicitly captured or can be inferred later, as required, using logic-based inferencing rules.
3. A common underlying data model expressed in RDF/OWL for capturing and modelling the provenance data.
4. A set of rules for inferring new implicit relationships between indirectly related entities in the provenance trail.

By basing the system on Semantic Web technologies, we can use the inferencing capabilities to generate alternative but still correct views of the data provenance. Alternative views of provenance are required for a number of reasons. Simplified views of highly complex workflows may be required for teaching or publication purposes. Restricted views which hide certain information or details are required to protect the intellectual property associated with particular scientific processes. This is particularly important within collaborating teams of scientists to protect individual IP but still enable controlled sharing and validation of the overall process.

Hence our objectives are to leverage existing RDF-based workflow tools and the captured provenance data and metadata in order to:

- generate visualizations of the lineage of the data and its products, i.e., the relationships between the different derivative products generated during the scientific process;
- streamline the construction of publication or e-learning packages (that link the raw data to its derivatives and traditional scholarly publications);
- enable interactive selection of the components of a scientific workflow to be included in the publication, keeping certain data or processing steps private in order to protect intellectual property and maintain competitive advantage;
- dynamically infer the provenance relationships between entities within the coarse-grained views to be published.

The remainder of this paper is structured as follows: Sect. 2 describes related work; Sect. 3 describes the case study we used for evaluation and testing; Sect. 4 describes the system architecture and components; Sect. 5 describes the implementation and user interface and Sect. 6 concludes with an evaluation, discussion and future work plans.

2 Related work

2.1 RDF-based Provenance capture tools

Our aim is to take the output from existing RDF-based provenance capture systems and to develop a visualization tool that dynamically generates customized views of the provenance trail. A number of such systems are available. For example, Kepler [1] is a scientific workflow system designed for multiple disciplines which enables scientists to design and execute workflows. Recently, Kepler embedded a new provenance recording component that collects data and workflow provenance at runtime. Other RDF-based workflow systems that capture provenance metadata in RDF include Taverna [2], Triana [3] and GridNexus [4].

Similarly, Recentris' Collaborative Electronic Research Framework (CERF)[1] and the SmartTea [5] and MyTea [6] systems are examples of RDF-based laboratory notebook systems. They provide a unified electronic record-keeping environment for scientists to capture, curate, annotate, and archive their data, and to integrate the data into electronic lab notebook-like pages.

These systems can be integrated relatively seamlessly into the Provenance Explorer system because they are Java-based applications that produce RDF graphs. Furthermore, the Protégé-OWL Plugin API can be used as the interface between the output from these systems and Provenance Explorer.

2.2 Common Provenance data model

One of the difficulties associated with managing provenance data associated with scientific workflows is that we are trying to store, describe and relate entities, data objects and events from both the real world as well as the digital domain and generated by a range of different organizations, individuals, disciplines, instruments and analytical tools. Integration of data originating from such a complex range of possible scenarios is not possible through conventional schema-mapping approaches. Semantic interoperability and semantic mediation is necessary to relate disparate data sources and workflow components to each other. Semantic mediation requires that the components of a scientific process are described semantically using terms and values defined within an ontology (expressed in a standard machine-processable language such as OWL). Additional reasoning can be performed across the

[1] http://www.rescentris.com/

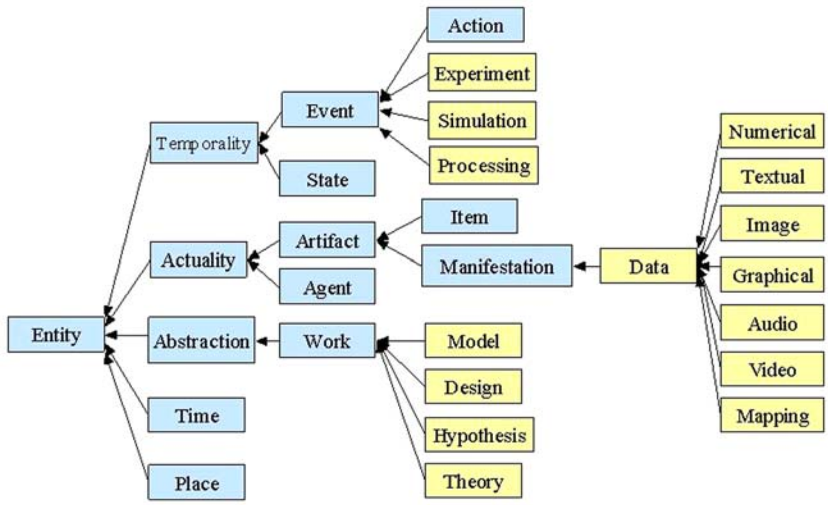

Fig. 1 Extensions to the ABC model to support eScience Provenance

provenance trails through the definition and application of logic-based rules expressed in SWRL (Semantic Web Rule Language).

The Harmony project's ABC model [7] is an event-aware, top level metadata model developed for the library, museum and archival domains to capture the events that a digital object undergoes during its lifecycle. However, it also provides an ideal top-level ontology that can be extended to define the entities and properties associated with scientific workflows. Figure 1 illustrates the upper class hierarchy for the extended ABC model that we have developed to support eScience provenance. The new classes are shaded. Associated with each of these new sub-classes are a set of properties specific to that sub-class. In addition we defined SWRL rules for inferring relationships between entities that were not directly related. For example:

IF (Event A precedes Event B) AND (Event B precedes Event C)
THEN (Event A precedes Event C)

2.3 Provenance visualization tools

A number of tools have previously been developed to enable the visualization of provenance trails but none provide the combination of features we require.

The *Prototype Lineage Server* [8] allows users to browse lineage information by navigating through the sets of metadata that provide useful details about the data products and transformations in a workflow invocation. Web server scripts on the lineage server query the lineage database, and provide a Web browser interface that allows navigation via HTML links. Views are restricted to parent and children metadata objects. Clicking on a parent object will move that link to the center of the screen and show that object's parents. Clicking on the metadata object link in the center of the screen will bring up the XML metadata for an object.

Pedigree Graph [9], one of tools in Multi-Scale Chemistry (MCS) portal from the Collaboratory for Multi-Scale Chemical Science (CMCS), is designed to enable users to view multi-scale data provenance. The portlet provides scientists with a two-dimensional visualization of a data object or file and all of its scientific pedigree relationships. The view is static, and rendered straight from GXL (Graphical eXchange Language) files but users are able to traverse the tree by clicking on links.

The *MyGrid* project renders graph-based views of RDF-coded provenances using *Haystack* [10]. Haystack is a Semantic Web browser that enables developers to provide tailored views over RDF-metadata and to visualize networks of semantic relationships among provenance resources associated with experiments. Haystack is highly resource-consumptive because its execution is based on Adenine, a high level programming language developed on top of Java. Hence the response time to user's instructions can be slow.

The *VisTrails* system [11] was developed by the University of Utah for building, storing, editing and visualizing workflows and interactively tracking workflow execution and evolution. Although it uses graphs to visualize workflows and provenance trails, it differs from the Provenance Explorer in that it is not designed to generate personalized views of provenance—adapted for publication or teaching purposes or to suit a user's interest or access permissions.

So although there are existing systems that enable visualization of RDF-encoded provenance graphs, the unique aspect of our Provenance Explorer system is its ability to dynamically generate tailored views of provenance trails (suitable for publication) using a combination of user input and semantic reasoning.

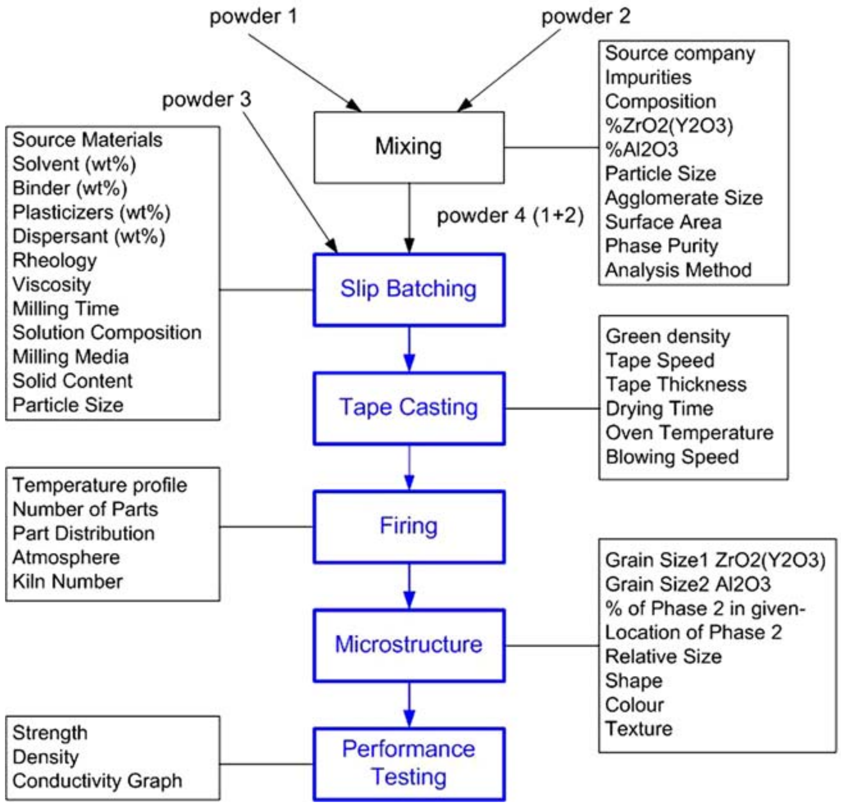

Fig. 2 A logical view of the manufacturing and testing process for fuel cell electrolytes

2.4 Complex digital object publishing tools

Within this paper, the focus is on the tools that enable the construction and publishing of scientific publication packages through their uploading to repositories such as Fedora [12], aDORe [13] or DSpace [27]. A recent special issue of the Journal of Digital Libraries on complex digital objects, includes several papers that focus on technologies to support the storage, management and dissemination of complex digital objects – not dissimilar to the Scientific Publication Packages that we are proposing in this paper. In Lagoze et al. [12], describe the Fedora open source digital repository service, that is designed to manage complex digital objects (and the relationships between their components). It uses an RDF-based relationship model to represent relationships among digital objects and their components, to support distributed information networks such as the National Science Digital Library (NSDL).

The aDORe system [13] developed at the Los Alamos National Laboratory research library also provides a standards-based repository for managing and accessing complex digital objects. Objects are encoded in XML using the MPEG-7 DIDL [14] and a limited set of object relationships can be expressed using RDF.

XML-based representations of composite objects such as METS[2] and the MPEG-21 DIDL [15] provide syntactic interoperability, but do not provide the necessary semantic interoperability or the ontology-based reasoning that can be applied to complex objects described using OWL. Our decision to model and represent the Scientific Publication Packages in RDF/OWL was based on these requirements.

3 Case study

Within the University of Queensland, materials scientists within the Australian Institute for Bioengineering and Nanotechnology are investigating the optimization of fuel cells—an alternative environment-friendly energy source to fossil fuels. Fuel cell efficiency depends on the internal structure of the fuel cell components and their interfaces. Electrolytes are one of the primary fuel-cell components. Figure 2 llustrates the complex set of steps involved in the process of manufacturing and testing electrolytes. Associated with each step in the workflow is a set of parameters, only some of which

[2] Metadata encoding and transmission standard <http://www.loc.gov/standards/mets/>

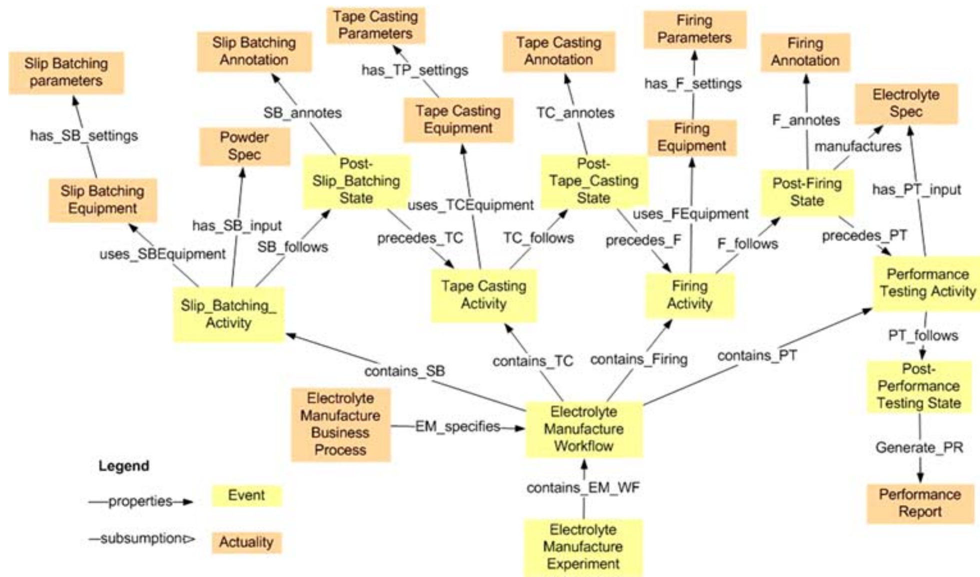

Fig. 3 Provenance model for the electrolyte manufacturing and testing process

are controllable. The objective of the fuel-cell scientists is to determine the optimum combination of controllable parameters in order to attain the maximum strength, efficiency and longevity of the fuel cell for the minimum cost [16].

Through the FUSION project [17] we have been collaborating with a team of fuel cell scientists on the development of an eScience workflow and provenance capture system that records the data associated with each of the steps in the electrolyte manufacturing and testing process and enables its statistical analysis in order to generate new workflows [16]. Through this work we have access to data records from a series of manufacturing and testing experiments. Hence we decided to use this application as a case study for evaluating and attaining user feedback on the Provenance Explorer system.

The first step involved modeling the workflow in Fig. 2 and representing it in OWL. We decided to use the extended ABC ontology [7] to track the life cycle of digital objects. We first had to extend the ABC ontology to describe processing, simulation and experimental events. Given this extended ontology, we were able to represent the workflow instances corresponding to Fig. 2 in OWL. This is illustrated in Fig. 3.

Given the OWL representations of the provenance data associated with the fuel cell manufacturing and testing process, the aim was to generate customized graphical visualizations of the data using the Provenance Explorer system.

4 System architecture

Figure 4 illustrates the overall system architecture and its key components. The three key components of the system are:

- The knowledge base which consists of SWRL.OWL files that contain the provenance instance data and metadata and the inference rules.
- the Provenance Visualizer and
- Algernon, a rule-inference engine.

The SWRL.OWL files are input to both the Provenance Visualizer and Algernon. Jena as the interface between the Provenance Visualizer and the SWRL.OWL files. The protégé-OWL plugin provides the interface acts between Algernon and the SWRL.OWL files. Jena [18], developed by HP Labs, provides the programmatic environment for RDF, RDFS and OWL. Jena supports SPARQL [19] which is used to query the SWRL.OWL files. The Protégé-OWL Plugin [20] was used to generate the SWRL.OWL files and to retrieve the rules from the SWRL.OWL files for Algernon to process at runtime. Algernon [21] is a rule-inference engine that supports both forward and backward chaining rules of inference, and implements Access-Limited Logic. However because Algernon does not support the inference of subsumption between properties or comply with the SWRL rule format, the rules retrieved from SWRL.OWL files by Protégé-OWL Plugin APIs had to be transformed to the Algernon-compliant rules before being imported to Algernon at runtime.

The Provenance Visualizer, uses the graphical user interface (GUI) powered by JGraph [22] (an extension of Java Swing GUI Component to support directed graphs). The Provenance Visualizer GUI is divided into three panels horizontally:

1. The upper panel presents a graphical view of the provenance process modeled using RDF graphs.

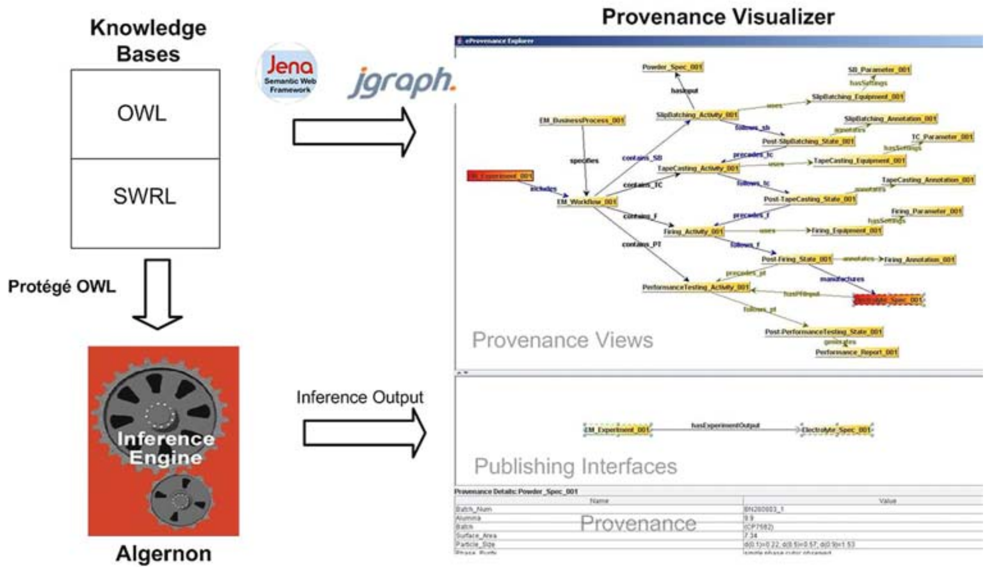

Fig. 4 Overview of system architecture and graphical user interface

2. The central panel is for dragging and dropping selected components from the upper panel into an RDF package. Any two can be linked manually with the relationships inferred automatically by Algernon.
3. Finally, the lower panel displays the provenance details (metadata) for the object highlighted in the upper panel.

Initially, authenticated users of Provenance Explorer are presented with a relatively coarse view of provenance. Finer-grained views may be accessed by clicking on specific links between entities—this expands the link to reveal the details of sub-workflows that have been hidden for reasons of either simplicity or confidentiality.

5 Demonstration and user interface

In the following section we describe the system from the point of view of a scientist in the FUSION project team investigating fuel cell modeling and optimization.

Firstly the researcher logs onto the Shibboleth Service provider where the Provenance Explorer service is installed. Initially, the user is redirected to Shibboleth's Identity Provider for authentication and authorization. Once authenticated, the user's attributes are returned back to the Service Provider and the user is granted access to the Provenance Explorer. The researcher searches for provenance of Batch Number 280818. Initially the researcher is presented with the basic view of the experiment provenance. This is the default view for all users with access privileges to the FUSION project's Provenance Explorer service. Figure 5 demonstrates the default expandable view. The darker arrows indicate relationships that can be expanded to reveal further fine-grained information about the sub-activities.

When the researcher clicks on a dark arrow, a request for additional information is generated and converted to a SPARQL query. The additional retrieved data is displayed in the top panel. Eventually by interactively drilling down via the links, the researcher is presented with the complete fine-grained view of the experiment or workflow. Figure 6 illustrates the complete view in the upper panel. The light arrows indicate parts of the expanded view and can be collapsed manually back to the original view i.e., the dark expandable links. If an individual node on the upper panel is selected, the complete provenance metadata for this node is displayed in the bottom panel. Figure 6 demonstrates this feature. Node *Powder_Spec_001* is highlighted in a circle on the upper panel, and the associated provenance information is displayed in the bottom panel.

Furthermore, using this interface, the researcher is able to manually construct a package of related components for publication or dissemination. This is performed by selecting nodes in the top panel and dragging and dropping them into the middle panel. By linking them manually, the relationship between the selected nodes is inferred by the Algernon rule-inference engine. For example, Fig. 6 demonstrates that the relationship inferred between the two selected nodes in the middle panel, *Experiment_001* and *Electrolyte_Spec_001* is *hasExperimentOutput*. The path used to infer this relationship is highlighted in blue (with the beginning and end nodes highlighted in red) in the upper panel. The extended ABC ontology that we are using (as described in Section 2.2) defines an experiment as comprising of a sequence of activities with particular pre-event and post-event states. One of the inferencing rules states that any output generated by one of the activities in the sequence is an output of the experiment. Hence the inferred relationship between

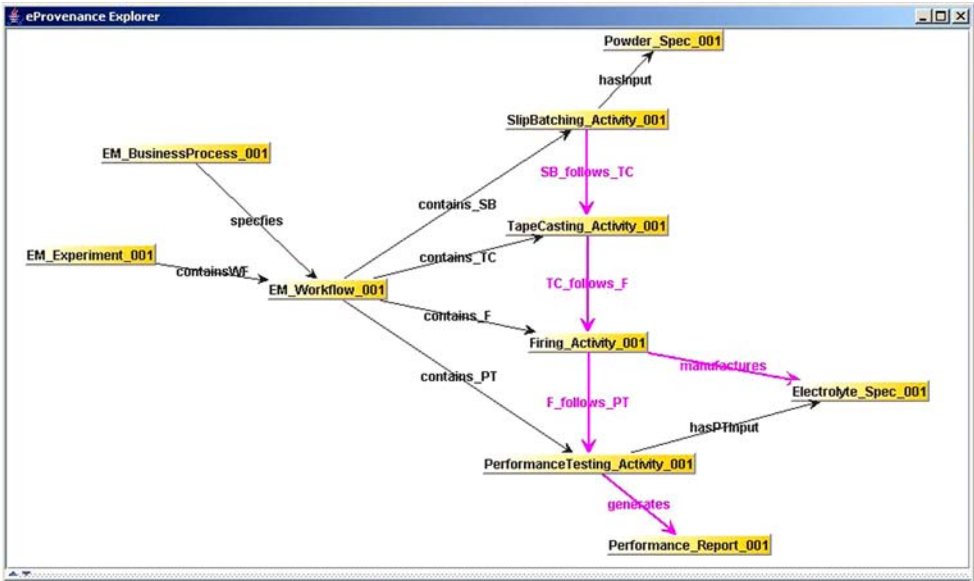

Fig. 5 A default expandable view

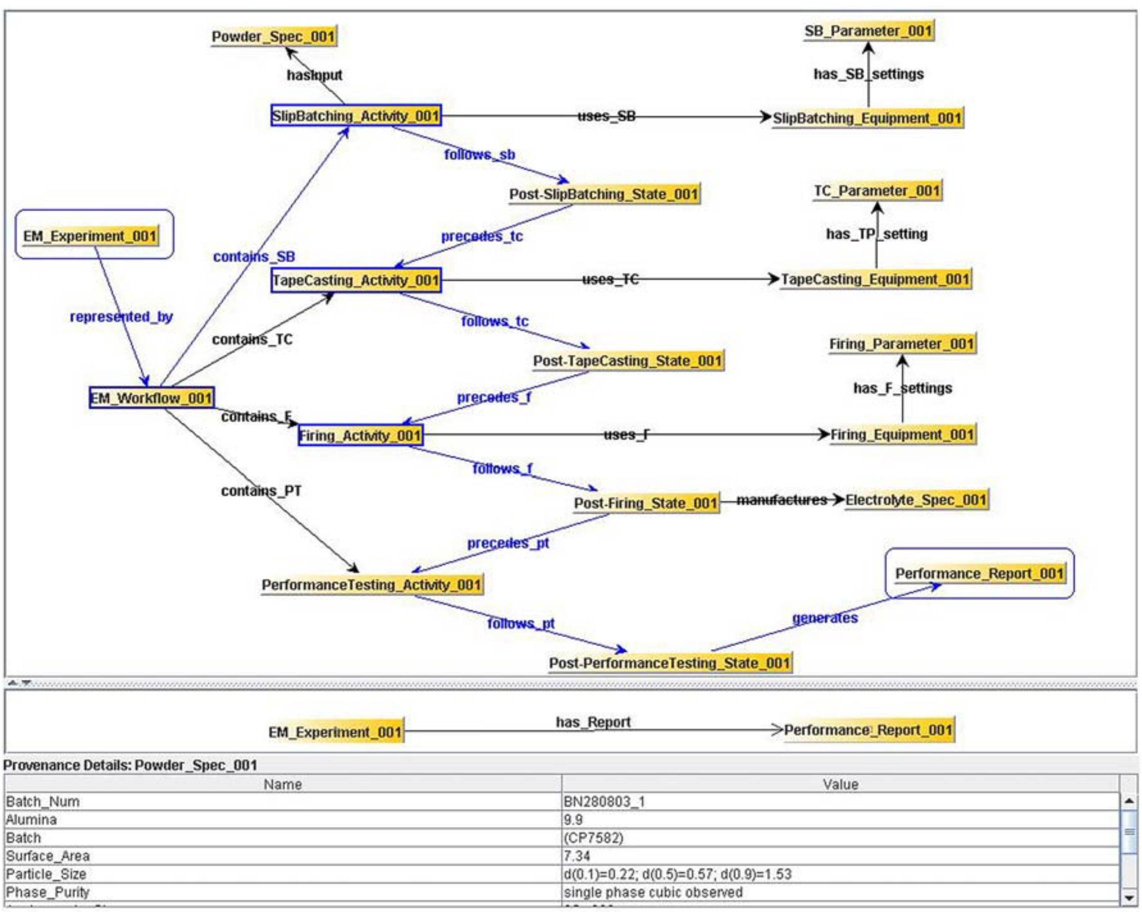

Fig. 6 An expanded fine-grained provenance view and illustration of provenance inferencing

Experiment_001 and *Electrolyte_Spec_001* is *hasExperimentOutput*.

6 Discussion and conclusion

6.1 User feedback

Initial feedback from the fuel-cell scientists involved in the FUSION project has been very positive. The system enables them to quickly and intuitively understand quite complex workflows and to compare different workflows. They are able to pinpoint problems within a particular workflow and to generate new experimental workflows accordingly. Users can understand the system very quickly because of its close analogy to the web—the use of hyperlinks for information exploration and navigation. Furthermore, with regard to the data's validity, the scientists can intuitively track the data's provenance via the complete graphical view of visualized scientific processes and the detailed metadata associated with any node. The users were also very positive about the security framework—in particular the use of Shibboleth to enable single sign-on and the advantages of being able to hide certain steps or the details associated with the process, in the publication.

However, users did raise concerns with regard to scalability and searching. At this stage, our demonstration involves multiple instances of a single workflow. In reality, the scientists may need to search, retrieve and compare multiple experiments simultaneously and the experimental workflows may be very different. Moreover, the current methods by which scientists can discover and retrieve experimental workflows are limited. Currently the system only permits search and retrieval of experiments via a unique ID. Scientists would like to be able to search for experiments via particular attributes, e.g., particular parameter values. The optimum methods for describing, indexing and discovering workflows require further investigation and direct input from the end-users.

6.2 Limitations and future work

The provenance metadata, graphical views and inferencing rules of the Provenance Explorer were all based on the provenance model in Fig. 1. This model is an extension of the ABC model developed within the Harmony project — extended to support experiments in laboratories. This model provides the semantic underpinning of the system, and the ontology's robustness may become a significant issue if/when the system is expanded across domains and organizations. Colomb argues that formal ontologies, such as DOLCE [23] and BWW [24], provide a rich meta-vocabulary and abstract data types, and well-understood structural organizational principles, thereby technically enhancing the reliability of material ontologies [25] such as our ontology. Thus, it may be worth carrying out further investigation on formal ontologies to determine how they can make the provenance model more reliable and rational in terms of the data structures.

To date the workflows that we have considered have really only focused on the provenance data/metadata and inferencing rules associated with processing events in a laboratory or manufacturing/processing plant. We need to extend the underlying model and the inferencing rules to support the data processing activities in the digital domain, e.g., reformatting, segmentation, normalization, compression etc.

We have begun investigating the use of XACML access policies to dynamically restrict access to particular workflow components within customized views. We require a streamlined mechanism for defining access policies and associating them with provenance data. For example, the individual or type of participant who is responsible for a particular activity or set of activities should have access to all of the provenance data associated with those activities and all sub-activities.

Another limitation of the current system is that it currently only supports expansion down one level of detail. Ideally users would be able to incrementally drill down to multiple levels of detail. For example one link can be expanded to two links, each of which can be further expanded. This may prove quite complex to implement because it involves multiple levels of inferencing rules and the specification of access policies associated with relationships at multiple levels.

Finally the packages of components that are able to be constructed provide a very efficient mechanism: for publishing and sharing scientific results; for teaching complex scientific concepts; and for the selective archival, curation and preservation of scientific data. Although we currently enable these packages to be saved, they are not indexed or able to be searched and retrieved. Tools are required to enable these RDF packages to be described, uploaded to institutional repositories (such as Fedora) and discovered and retrieved for reuse. In addition, we are continuing to track the outcomes of the Science Commons Initiative[3]. The Science Commons Licensing sub-project is exploring standard open agreements to facilitate licensing of intellectual property and the exchange of research materials. Our aim is to provide tools to enable scientists easily to attach Science Commons licences to SPPs and their components when they want to share them—without sacrificing intellectual property rights.

6.3 Conclusions

A recent OECD report on the scientific publishing industry [26] recommends that governments make publicly funded

[3] http://science.creativecommons.org/

research findings more widely available in order to boost innovation and get a better return on their investment. Consequently scientists are under increasing pressure to publish their experimental and evidential data together with the related traditional scholarly publication(s). But the infrastructure required to support these new forms of scientific publishing is still immature and currently relies on an ad hoc assemblage of software that is inadequate for the task. In this paper, we have described the Provenance Explorer system that we have developed to help fill this gap. It is a provenance visualization system that dynamically generates a graphical view of a provenance trail using RDF graphs. It enables users to intuitively explore the data provenance associated with scientific workflows or experiments by expanding or collapsing sub-worfklows through a hypermedia GUI. The interface also enables scientists to quickly and easily select and wrap related outputs into a single package for publication, peer-review or e-learning—and to have the provenance trail between the components automatically inferred. By basing the system on underlying semantic web technologies (including RDF, OWL, SWRL and the Algernon inferencing engine) it has been quick to implement, has greater flexibility and adaptability and the data being produced can more easily be shared, re-used, compared and integrated with other data sources.

References

1. Altintas, I., Barney, O., Jaeger-Frank, E.: Provenance collection support in the Kepler scientific workflow system. In: International provenance and annotation workship (IPAW'06). Chicago (2006)
2. Oinn, T., et al.: Taverna: A tool for the composition and enactment of bioinformatics workflows. Bioinformatics J. **20**(17), 3045–3054 (2004)
3. Majithia, S., et al.: Triana: A Graphical Web Service Composition and Execution Toolkit. in IEEE International Conference on Web Services (ICWS'04). IEEE Computer Society (2004)
4. Brown, J.L., et al.: GridNexus: a grid services scientific workflow system. Int. J. Comput. Info. Sci. **6** (2), (2005)
5. Schraefel, M.C., et al.: Breaking the Book: Translating the chemistry lab book into a pervasive computing lab environment. in CHI04. Vienna (2004)
6. Gibson, A., et al.: myTea: connecting the web to Digital Science on the Desktop. in World Wide Web Conference. 2006. Edinburgh
7. Lagoze, C., Hunter, J.: The ABC Ontology and Model. J. Digi. Info. **2** (2), (2001)
8. Bose, R., Frew, J.: Composing lineage metadata with XML for custom satellite-derived data products. in Scientific and Statistical Database Management, 2004. In: Proceedings of 16th International Conference (2004)
9. Myers, J.D., Pancerella, C., Lansing, C., Schuchardt, K.L., Didier, B.: Multi-scale science: supporting emerging practice with semantically derived provenance. In: ISWC 2003 Workshop: semantic web technologies for searching and retrieving scientific data. Sanibel Island (2003)
10. Zhao, J., et al.: Using Semantic Web technologies for representing e-science provenance. In: 3rd international semantic web conference. Hiroshima (2004)
11. Freire, J., et al.: Managing rapidly-evolving scientific workflows. In: International provenance and annotation workship (IPAW'06). Chicago (2006)
12. Lagoze, C., et al.: Fedora: An architecture for complex objects and their relationships. J. of Digit. Libr. (2005)
13. Van de Sompel, H., et al.: aDORe: A modular, standards-based digital object repository. Comput. J. **48**(5), 514–535 (2005)
14. Hunter, J., Nack, F.: An Overview of the MPEG-7 desiption definition language (DDL) proposals. Signal Process. Image Communi. J. Special Issue on MPEG-7 **16**, 271–293 (2000)
15. Bekaert, J., Hochstenbach, P., Van-de-Sompel, H.: Using MPEG-21 DIDL to represent complex digital objects in the Los Alamos National Laboratory Digital Library. D-Lib Mag. **9** (11), (2003)
16. Hunter, J., Cheung, K.: Generating eScience Workflows from Statistical Analysis of Prior Data. in APAC'05. Royal Pines Resort, Gold Coast (2005)
17. Hunter, J., Drennan, J., Little, S.: Realizing the Hydrogen Economy through Semantic Web Technologies. IEEE Intell. Syst. J.—Special Issue on eScience 40–47, (2004)
18. Carroll, J.J., et al.: Jena: implementing the semantic web recommendations. In: Proceedings of the 13th international World Wide Web conference on Alternate track papers & posters. ACM Press, New York pp. 74–83 (2004)
19. McCarthy, P.: Search RDF data with SPARQL: SPARQL and the Jena Toolkit open up the semantic Web, in developerWorks. IBM (2005)
20. O'Connor, M., et al.: Writing Rules for the Semantic Web Using SWRL and Jess. In: 8th International Protege Conference, Protege with Rules Workshop. Madrid (2005)
21. Crawford, J.M., Kuipers, B.J.: Algernon — a tractable system for knowledge-representation. SIGART Bull. **2**(3), 35–44 (1991)
22. Alder, G.: The JGraph Swing Component, in Department of Computer Science. Federal Institute of Technology ETH, Zurich (2002)
23. Gangemi, A., et al.: Sweetening Ontologies with DOLCE. In: 13th International conference on knowledge engineering and knowledge management. Siguenza (2002)
24. Weber, R.: Ontological foundations of information systems. Monograph No. 4. Coopers& Lybrand Accounting Research Methodology, Melbourne (1997)
25. Colomb, R.M.: Formal versus Material Ontologies for information Systems interoperation in the Semantic Web. Comput. J. **49** (1), (2006)
26. Digital Broadband Content: Scientific Publishing, in Source-OECD Science & Information Technology. OECD - Organisation for Economic Co-operation and Development (2005)
27. Mackenzic, S.: D Span for E-print archives. High energy Phys Libr webzine (9), (2004)