

## The PBase Scientific Workflow Provenance Repository

Víctor Cuevas-Vicentín<sup>1</sup>, Parisa Kianmajd<sup>1</sup>,  
Bertram Ludäscher<sup>1</sup>, Paolo Missier<sup>2</sup>, Fernando  
Chirigati<sup>3</sup>, Yaxing Wei<sup>4</sup>, David Koop<sup>3</sup>, Saumen  
Dey<sup>1</sup>

<sup>1</sup>University of California at Davis, <sup>2</sup>Newcastle  
University, <sup>3</sup>Polytechnic Institute of NYU, <sup>4</sup>Oak  
Ridge National Laboratory

victorcuevasv@gmail.com,  
parisa.kianmajd@gmail.com, ludaesch@ucdavis.edu  
Paolo.Missier@ncl.ac.uk, fchirigati@nyu.edu,  
weiy@ornl.gov, dkoop@nyu.edu, scdey@ucdavis.edu

### Abstract

Scientific workflows and their supporting systems are becoming increasingly popular for compute-intensive and data-intensive scientific experiments. The advantages scientific workflows offer include rapid and easy workflow design, software and data reuse, scalable execution, sharing and collaboration, and other advantages that altogether facilitate “reproducible science”. In this context, provenance, information about the origin, context, derivation, ownership, or history of some artifact plays a key role, since scientists are interested in examining and auditing the results of scientific experiments.

However, in order to perform such analyses on scientific results as part of extended research collaborations, an adequate environment and tools are required. Concretely, the need arises for a repository that will facilitate sharing scientific workflows and their associated execution traces in an interoperable manner, also enabling querying and visualization. Furthermore, such functionality should be supported while taking performance and scalability into account.

With this purpose in mind we introduce PBase, a scientific workflow provenance repository implementing the PROV-ONE proposed standard, which extends the emerging W3C PROV standard for provenance data with workflow specific concepts. PBase is built on the Neo4j graph database, thus offering capabilities such as declarative and efficient querying. Our experiences demonstrate the power gained by supporting various types of queries for provenance data. In addition, PBase is equipped with a friendly user interface tailored for the visualization of scientific workflow provenance data, making also the specification of queries and the interpretation of their results easier and more effective.

*Submitted 28 October 2013*

Correspondence should be addressed to Víctor Cuevas-Vicentín, 1 Shields Avenue, Davis, California, 95616, United States. Email: victorcuevasv@gmail.com



The 9th International Digital Curation Conference will take place on 24–27 February 2014 in San Francisco. Please ensure you use the guidance in this template to produce your paper. Please submit your paper in one of the following formats: Microsoft Word (.doc, .docx), Open Document Format (.odt) or Rich Text (.rtf). <http://www.dcc.ac.uk/events/idcc14/submissions>

## Introduction

The origin and processing history of an artifact is known as its *provenance*. Data provenance is an important form of metadata that explains how a particular data product was generated, including the system and the steps in the computational process involved along with the user responsible for its execution, time, and resources used, such as parameter settings, input data, software tools, etc. Provenance information provides transparency and helps to audit processes and interpret data products. Common uses and applications of provenance include quality control, data curation, debugging, data discovery, and generally, the validation, attribution, and reproducibility of scientific results.

The state of the art of scientific workflow systems such as Kepler (Ludäscher et al., 2006), Taverna (Wolstencroft et al., 2013), VisTrails (Freire et al., 2006) provide controlled environments for specifying and enacting complex computational pipelines for which provenance information is automatically captured by the system in the form of traces (albeit often in proprietary formats). However, they lack the capabilities to enable the users to effectively query and visualize the provenance traces associated with a particular workflow. In addition, they do not address the need of scientists of sharing their workflow-based computational experiments for the scrutiny and benefit of the rest of the community. Therefore, the need arises for a repository allowing multiple users to store and query scientific workflow provenance in an interoperable manner. In this work we document the task of developing and putting into use such a repository, called PBase. Our ultimate goal is to incorporate such a repository into DataONE<sup>1</sup>, a large scale and federated data infrastructure serving the Earth sciences community.

In order to build PBase we face three main challenges. i) Defining a standard model for workflow provenance representation that is compatible with the main scientific workflow systems. ii) Characterizing the required functionality in terms of specific queries. iii) Developing the infrastructure that enables to evaluate such queries efficiently, while also being easy to use for the scientific community. We briefly describe next how we addressed those challenges. First, describing PROV-ONE, our proposed standard for scientific workflow provenance data, after which we present representative provenance queries. Subsequently, we describe the PBase architecture and how users interact with the system. Finally, we discuss related works, present our conclusions and discuss future work.

## The PROV-ONE Model for Workflow Provenance

Scientific workflow systems provide a user-friendly graphical interface to specify a computational process in the form of a directed graph of interconnected tasks. Such a graph is an abstraction that can be regarded as *prospective* provenance, since it details the steps to follow in order to generate the desired result. In addition, as mentioned earlier, the workflow system automatically captures the various events associated with the workflow execution, which is usually referred to as *retrospective* provenance. The full potential of provenance information is achieved when combining both. To this end,

---

<sup>1</sup> <http://www.dataone.org/>

in the context of the DataONE project, an extension of the W3C PROV standard called PROV-ONE has been developed.

PROV (PROV-WG, 2013) provides a basic and standardized model for the representation and exchange of provenance information. PROV-ONE uses and extends the information elements of PROV to describe both workflow-level and trace-level information. The adoption of PROV-ONE establishes a common model for PBase and brings the advantages in interoperability of the emerging PROV standard. Support for PROV-ONE, serialized as XML, was recently added to VisTrails and other workflow systems can be supported through corresponding wrappers as well.

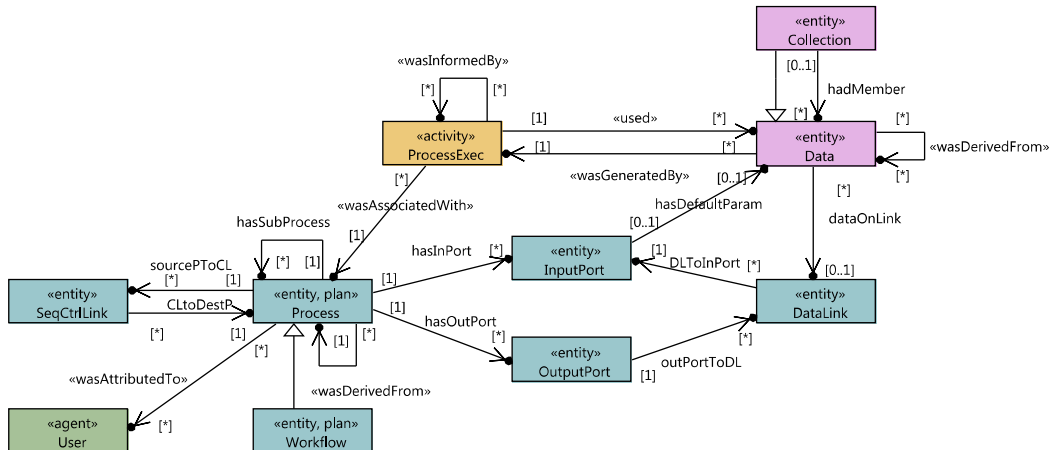


Figure 1. PROV-ONE conceptual model UML class diagram.

## PROV-ONE Conceptual Model

The PROV-ONE conceptual model is illustrated by the UML diagram of Figure 1. All classes have a correspondant PROV type denoted by a UML stereotype (e.g. «entity»), whereas this is the case for only a subset of the associations (e.g. «used»).

The various tasks that form part of a workflow are represented by the *Process* class. *Processes* can be either atomic or composite, the latter case specified through the *hasSubProcess* self association. A given *Process* can be distinguished as a *Workflow*. Each *Process* has a series of *InputPorts* and *OutputPorts*, while the ports from the various *Processes* are connected through *DataLinks*. Note that both input and output ports can be associated with multiple *DataLinks*, thus allowing workflow models in which a single output is copied and sent to multiple destinations as well as in which tasks take inputs from different sources through a single input port.

In order to specify executable instances of a *Workflow*, default parameters can be defined for some of its constituting *Processes*. The default parameters are represented by *Data* elements, which can be on various types (e.g. XML, JSON, files, etc.) including collections thereof. In addition, sequential control links can be specified between *Processes* as denoted by the *SeqCtrlLink* class, which enforce that a given *Process* can only be executed after the other *Process* that shares the link has sent the required signal. Finally, a particular *Process* that specifies a *Workflow* or part of it can be associated with a *User* that assumes responsibility for its creation. A detailed presentation of the PROV-ONE model based on an OWL-2 ontology can be found in

(DataONE-ProvWG, 2014). An example of a concrete workflow and its corresponding trace is presented in the next section.

## Scientific Workflow Provenance Queries

PROV-ONE data is intrinsically graph-oriented, since workflows are represented as graphs having processes as nodes and traces in part as graphs whose nodes correspond to data entities (inputs and outputs) and computational activities (process executions). Therefore, the queries of interest over scientific workflow provenance correspond in many cases to queries over graphs, which have been extensively studied and employed in various domains.

Figure 2 a) presents an example workflow whose goal is to summarize ecological spatiotemporal data. The left branch calculates the standard deviation of the regional net ecosystem exchange data obtained from its input, and generates a map with 1 degree resolution illustrating the standard deviation of the data across the region. The right branch uses the same data to calculate monthly averages and presents the results in a plot. A trace generated by running this workflow is depicted in Figure 2 b), with data items as ellipses and process executions as rectangles (some data and process execution nodes have been omitted for clarity).

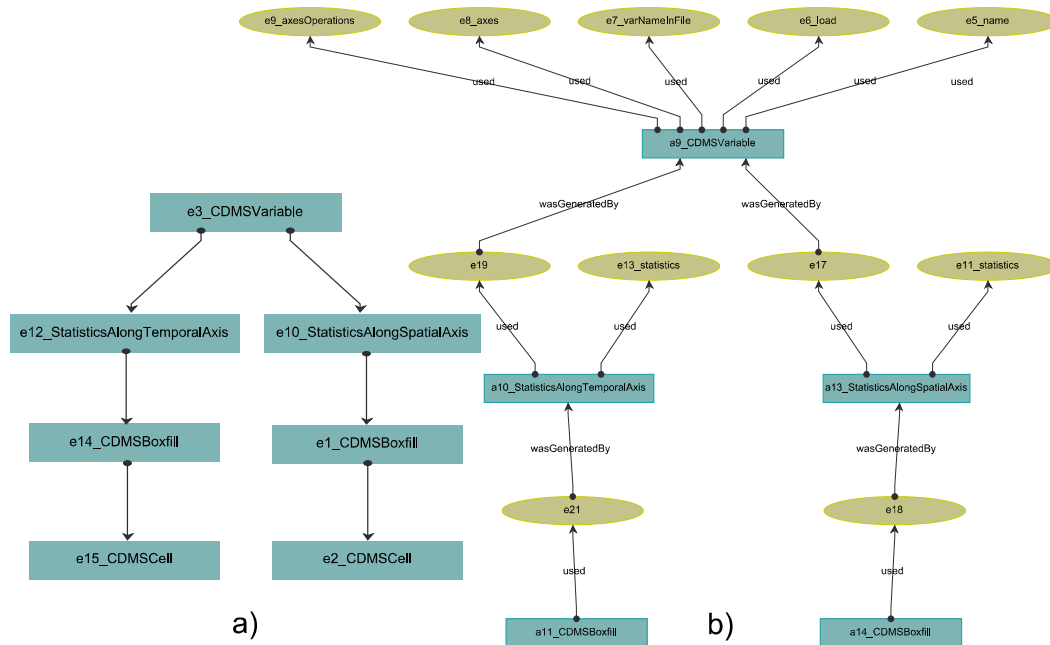


Figure 2. Example workflow (a) and corresponding trace (b).

In cooperation with climate scientists that employ scientific workflows such as the one presented in Figure 2 a) in their research, we identified a suite of queries that are relevant for scientific workflow provenance. The queries in the suite can be classified in four main types. i) *Lineage* queries, i.e., dealing with the derivation of data products. For example, “what are the datasets involved in the generation of this plot?”. ii) *Execution analysis* queries, for example, “find the processes in a workflow that were not

completed”. iii) *Search* queries, for example, “find all the workflows that used a bilinear regrid module”. iv) *Statistical* queries, such as “list the most used modules across all workflows”.

Table 1 presents six queries that are incorporated into our implementation and their query type. Queries 1-3 can also be characterized as focusing on workflows, while queries 4-6 focus on traces. This distinction becomes relevant for user interaction as discussed in the next section.

#	Query	Type
1	Compute the number of invocations of process <code>e7_Regrid</code>	Statistical
2	Find all inputs of a workflow across executions	Lineage
3	Find the modules of processes that were not executed	Execution
4	Find the process executions that were not completed	Execution
5	Find the processes that used data item <code>e10_regrid_method</code>	Search
6	Find the data products influenced by module <code>e7_Regrid</code>	Search

Table 1. Example queries for scientific workflow provenance data.

## The PBase Provenance Repository

### Provenance Data Storage and Querying

The challenges involved in developing the infrastructure capable of evaluating the aforementioned queries are twofold. First, we need an expressive query language covering diverse types of queries. Second, query evaluation should be efficient considering that provenance traces can be significantly large. We opted to base our infrastructure on the NoSQL graph database Neo4j<sup>2</sup>, following the rationale that it is customized for graph data and that it has recently incorporated Cypher, a declarative graph query language.

1	START n = node:node_auto_index(name="e7_Regrid") MATCH m-[:wasAssociatedWith]-n RETURN count(m)
2	START n=node(*) MATCH (n)-[:used]-() WHERE not ((n)-[:wasGeneratedBy]->()) AND HAS(n.wfID) AND (n.wfID="wf1") RETURN DISTINCT n
3	START n=node(*) MATCH n<-[:wasAssociatedWith]-m WHERE HAS(n.vtType) AND HAS(n.wfID) AND HAS(n.runID) AND n.vtType="vt:module_exec" AND n.completed=-1 AND n.wfID="wf1" AND n.runID="ex1" RETURN m
4	START n=node(*) WHERE HAS(n.completed) AND n.completed=-1 AND HAS(n.wfID) AND n.wfID="wf1" RETURN DISTINCT n
5	START n = node:node_auto_index(name="e10_regrid_method") MATCH n<-[:used]-a RETURN distinct a
6	START n=node(*) MATCH n-[*]->a WHERE n-[:wasGeneratedBy]->() AND HAS(n.wfID) AND n.wfID="wf1" AND HAS(a.module) AND a.module="e7_Regrid" RETURN DISTINCT n

Table 2. Cypher specification of the queries in Table 1.

Table 2 presents the Cypher specification of the queries presented in the previous section. The basic structure of a Cypher query comprises a `START` clause that defines a

<sup>2</sup> <http://www.neo4j.org/>

series of starting points for graph exploration, a `MATCH` clause that specifies a pattern to be matched in the graph, a `WHERE` clause specifying additional filter conditions, and a `RETURN` clause that outputs the results, which can consist of complete nodes and edges or only of attributes thereof. A detailed discussion of the language is presented in (Neo Technology, 2014).

In some cases, however, the performance of Neo4j was found to be unsatisfactory when querying large traces. An alternative to address this issue is to incorporate new indexing and encoding techniques into Neo4j to achieve better performance. For reachability queries, a particular case of lineage that determines whether there is a directed path between a pair of nodes, we achieved major improvements by incorporating the tree cover encoding introduced in (Agrawal et al., 1989) to the stored workflows and traces. This form of encoding can be generated in quadratic time on the number of nodes and in the worst case can be of a linear size for some nodes. In practice, however, its size is significantly lower and enables to answer reachability queries in nearly a constant time. For a test graph of 10,000 nodes the encoding enabled us to evaluate 8,000 random queries in under a minute on a laptop, while it took more than two hours without it.

### System Architecture

The architecture of our current prototype is presented in Figure 3. The system adopts a Java platform three tier architecture with a client application running on a Web browser, while the server runs on the Tomcat application server. The server in turn is separated into the application logic and the database, the latter supported by Neo4j 1.9.1 running on embedded mode. Currently, our prototype only enables the creation of a database from a VisTrails exported XML trace file, which contains the workflow specification and the various traces corresponding to the workflow runs. This file is submitted to PBase through the *TraceUpload* component of the Web client as an HTTP POST request, which includes a user supplied identifier for the database to be created.

The PBase server application implements a series of Web services to handle user requests. In the case of submitted traces, these are processed by the *ProvTraceUploadService* that stores the file temporarily on the server, after which it is processed by the *PROVBuilder* component in the database tier, which relies on the *GraphDAO* component to store the data in an actual Neo4j database. The graph data for Neo4j is complemented by the tree cover encoding discussed earlier that is generated by its corresponding component. One database is created for each original XML trace file with its supplied identifier.

The stored workflows and traces can be accessed and queried through the *PBaseGUI*, whose generated requests are handled by a series of restful Web services (characterized as resources) in the PBase application. Requests are sent through HTTP GET, and are served through Cypher queries issued to the Neo4j database by the *GraphDAO* component, which also serializes the results in JSON to be sent back to the client.

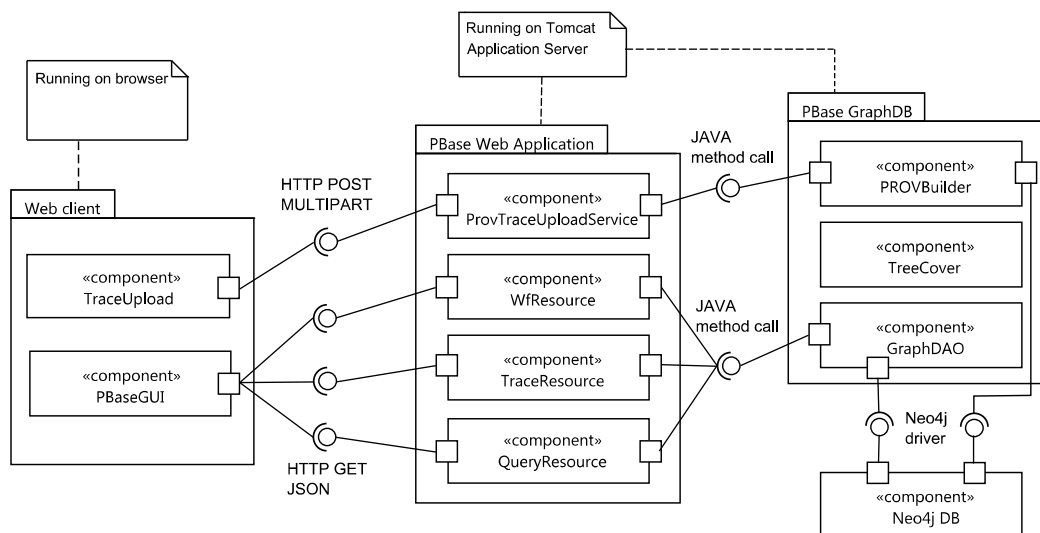


Figure 3. PBase System Architecture.

### User Interaction with PBase

As a proof of concept user interface aimed at end users, we developed a Web-based GUI that enables users to upload a provenance trace, visualize the workflow alongside its various traces, issue queries and obtain visualizations of their results. Our experience with climate scientists shows that this makes PBase much easier to use in comparison to the generic tools included in Neo4j, particularly for users that are not familiar with scientific workflow provenance in the first place. The Web-based GUI was developed with the YUI 3.14.1 framework<sup>3</sup> and the mxGraph 2.3.0.5 library<sup>4</sup>, a snapshot of it is depicted in Figure 4. Next we describe the main aspects of its functionality.

#### *Workflow and Trace Loading and Visualization.*

In order to load a workflow and its corresponding traces from the database, the user supplies the identifier specified when the trace was uploaded. Once the workflow and traces have been loaded, a visualization of the workflow is generated on the left-hand side and of the trace on the right-hand side. If multiple traces are associated with the workflow, these can be changed from a menu on the load tab on the right-hand side panel.

The size of the visualization panels of the workflow and the trace can be changed in order to focus on one or the other if desired. Their respective graph visualizations can also be zoomed in or out and dragged independently. Initially, only the identifiers of the various process, process execution and data nodes are displayed, but the additional properties associated with the nodes can be visualized by overlaying text boxes that can be shown or hidden as required. The tree cover encoding can also be visualized in an analogous manner.

<sup>3</sup> <http://yuilibrary.com/>

<sup>4</sup> <http://jgraph.github.io/mxgraph/>

### *Workflow and trace querying and results visualization.*

The user can select from a suite of example queries based on their textual description, for instance, Queries 1-3 from Table 1. Once the query is selected from the menu, its corresponding Cypher specification appears on the accompanying text area in the querying tab panel, in our example these would be queries 1-3 from Table 2. The Cypher specifications are generated automatically to include the identifiers of the loaded workflow and trace. The user can then execute the query directly or edit it as required. Alternatively, she can also specify an entirely new query from the basic query structure. This functionality is present for both the workflow and the trace in its corresponding panels on the GUI.

After the query is executed its results are presented in tabular form on a pop-up window, as shown in the left-hand side of Figure 4. If node attributes form part of the result these are presented directly in the table, whereas if entire nodes are part of the result their identifier is shown while the actual nodes can be highlighted on the corresponding graph if desired. Since lineage is particularly important for provenance data, PBase offers the capability to visualize the reachability of all the nodes in a workflow or a trace with respect to a particular node. Concretely, the user can select a node on the workflow or trace graph and the descendants and ancestors of that node will be highlighted using different colors (no change occurs if a node is neither an ancestor nor a descendant). Such highlighting is performed entirely on the client side via the GUI by taking advantage of the encoding pre-computed on the server, thus resulting in a more fluid user interaction.

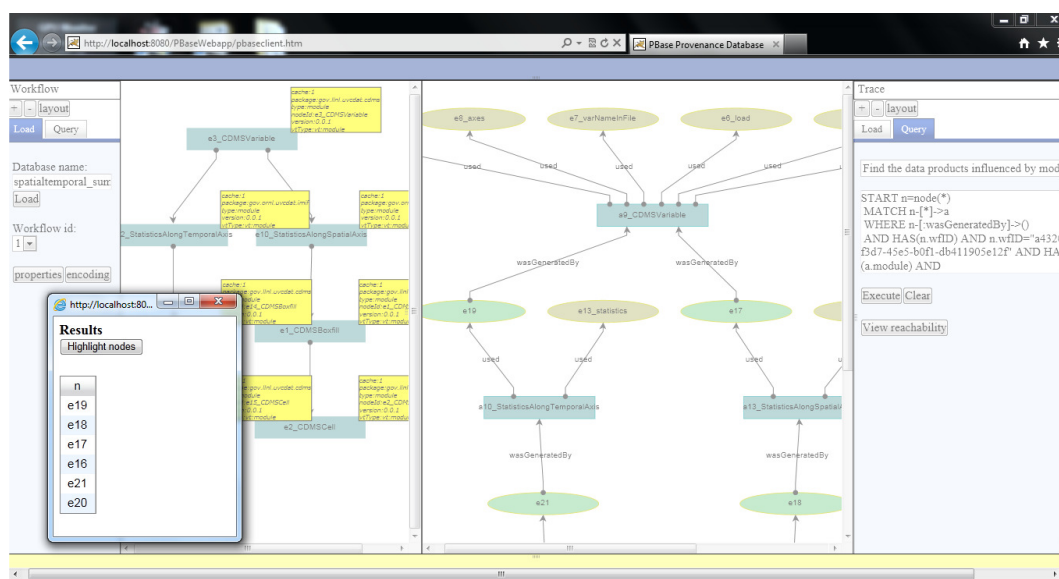


Figure 4. PBase Web-based graphical user interface.

## Related Work

Regarding the PROV-ONE model, independent extensions of PROV for scientific workflows are presented in (Costa et al., 2013) as well as in (Belhajjame et al., 2013) (focusing on workflow preservation). Although PROV-ONE may be in its core very



similar to these models, we additionally developed a repository with querying and visualization capabilities representing a more complete and flexible infrastructure. Our work in part complements previous results presented in (Missier et al., 2012), which focused on enabling the specification of mappings between data produced and used in different traces. This paper focuses instead on enabling declarative querying and effective visualization. In particular, we believe that the adoption of a graph query language can make the system easier to use as opposed to the approach presented in (Marinho et al., 2012) using Prolog, for example.

## Future Work

Future work involves enhancing the capabilities of PBase with queries that cannot be evaluated in Cypher or other graph query languages, such as finding workflows and traces by graph similarity or performing sophisticated keyword searches that take ranking into account. Integration with the DataONE Cyberinfrastructure is also envisaged to enable scientists to publish, search, explore, and query scientific data provenance in DataONE. Support for other workflow systems is also planned to make the system more practical to the scientific community. Although our experience with Neo4j was satisfactory overall, in particular regarding the expressivity of the Cypher query language, we are also exploring RDF/SPARQL graph databases as an alternative. PROV is inclined towards these Semantic Web standards and consequently, RDF/SPARQL are being adopted in the most recent PROV-ONE specification.

## Conclusions

Provenance data is of foremost importance in eScience, where scientific workflow management systems already play a key role. In order to take advantage of scientific workflow provenance data it is essential to be able to share this data, the PROV-ONE model makes this feasible by extending PROV with workflow-specific concepts. In addition, it is necessary to be able to store and query efficiently such data, while also being able to do so in a user friendly manner and accompanied by adequate visualizations. PBase fulfills these requirements by adopting a powerful graph database and offering a user interface tailored for scientific workflow provenance. Our experiences show that the functionality offered by PBase represents an important first step for the wider adoption of scientific workflow provenance as part of the computer-based scientific process.

## Acknowledgements

The authors would like to thank the members of the DataONE EVA Working Group for their collaboration in the specification of the climate data analysis workflows, and of the Provenance Working Group in the specification of the PROV-ONE model and the functionality of PBase. This work was supported by NSF Award OCI-0830944 (DataONE).

## References

- Agrawal, R., Borgida, A., and Jagadish, H. V. (1989). Efficient management of transitive relationships in large data and knowledge bases. In Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data, SIGMOD '89, pages 253-262, New York, NY, USA. ACM.
- Belhajjame, K., Klyne, G., Garijo, D., Corcho, O., García-Cuesta, E., and Palma, R. (2013). Wf4ever research object model. URL: <http://wf4ever.github.io/ro/>
- Costa, F., Silva, V., de Oliveira, D., Ocaña, K., Ogasawara, E., Dias, J., and Mattoso, M. (2013). Capturing and querying workflow runtime provenance with PROV: A practical approach. In Proceedings of the Joint EDBT/ICDT 2013 Workshops, EDBT '13, pages 282-289, New York, NY, USA. ACM.
- DataONE-ProvWG (2014). PROV-ONE: A PROV extension data model for scientific workflow provenance. URL: <http://purl.org/provone/>
- Freire, J., Silva, C. T., Callahan, S. P., Santos, E., Scheidegger, C. E., and Vo, H. T. (2006). Managing rapidly-evolving scientific workflows. In Proceedings of the 2006 international conference on Provenance and Annotation of Data, IPAW'06, pages 10-18, Berlin, Heidelberg. Springer-Verlag.
- Ludäscher, B., Altintas, I., Berkley, C., Higgins, D., Jaeger, E., Jones, M., Lee, E. A., Tao, J., and Zhao, Y. (2006). Scientific workflow management and the Kepler system: Research articles. *Concurr. Comput. : Pract. Exper.*, 18(10):1039-1065.
- Marinho, A., Murta, L., Werner, C., Braganholo, V., Cruz, S. M. S. d., Ogasawara, E., and Mattoso, M. (2012). Provmanager: A provenance management system for scientific workflows. *Concurr. Comput. : Pract. Exper.*, 24(13):1513-1530.
- Missier, P., Ludäscher, B., Dey, S. C., Wang, M., McPhillips, T. M., Bowers, S., Agun, M., and Altintas, I. (2012). Goldentrail: Retrieving the data history that matters from a comprehensive provenance repository. *IJDC*, 7(1):139-150.
- Neo Technology, Inc. (2014). Neo4j Cypher Refcard 1.9. URL: <http://docs.neo4j.org/refcard/1.9/>
- PROV-WG, W3C. (2013). PROV-Overview: An Overview of the PROV Family of Documents. URL: <http://www.w3.org/TR/prov-overview/>
- Wolstencroft, K., Haines, R., Fellows, D., Williams, A., Withers, D., Owen, S., Soiland Reyes, S., Dunlop, I., Nenadic, A., Fisher, P., Bhagat, J., Belhajjame, K., Bacall, F., Hardisty, A., de la Hidalga, A. N., Vargas, M. P. B., Su, S., and Goble, C. (2013). The taverna workflow suite: designing and executing workflows of web services on the desktop, web or in the cloud. *Nucleic Acids Research*, First published online May 2, 2013, doi:10.1093/nar/gkt328.