


The First Provenance Challenge



Luc Moreau*, Bertram Ludäscher, Ilkay Altintas, Roger S. Barga, Shawn Bowers, Steven Callahan, George Chin Jr., Ben Clifford, Shirley Cohen, Sarah Cohen-Boulakia, Susan Davidson, Ewa Deelman, Luciano Digiampietri, Ian Foster, Juliana Freire, James Frew, Joe Futrelle, Tara Gibson, Yolanda Gil, Carole Goble, Jennifer Golbeck, Paul Groth, David A. Holland, Sheng Jiang, Jihie Kim, David Koop, Ales Krenek, Timothy McPhillips, Gaurang Mehta, Simon Miles, Dominic Metzger, Steve Munroe, Jim Myers, Beth Plale, Norbert Podhorszki, Varun Ratnakar, Emanuele Santos, Carlos Scheidegger, Karen Schuchardt, Margo Seltzer, Yogesh L. Simmhan, Claudio Silva, Peter Slaughter, Eric Stephan, Robert Stevens, Daniele Turi, Huy Vo, Mike Wilde, Jun Zhao, Yong Zhao

University of Southampton

SUMMARY

The first Provenance Challenge was set up in order to provide a forum for the community to understand the capabilities of different provenance systems and the expressiveness of their provenance representations. To this end, a Functional Magnetic Resonance Imaging workflow was defined, which participants had to either simulate or run in order to produce some provenance representation, from which a set of identified queries had to be implemented and executed. Sixteen teams responded to the challenge, and submitted their inputs. In this paper, we present the challenge workflow and queries, and summarise the participants contributions.

KEY WORDS: Provenance, representation, storing, recording, capabilities, queries

*Correspondence to: Electronics and Computer Science, University of Southampton Southampton, SO17 1BJ, U.K

†E-mail: contact author: l.moreau@ecs.soton.ac.uk



1. Introduction

The term *provenance* is commonly used in the context of art to denote the documented history or the chain of ownership of an art object. Provenance helps determine the authenticity and therefore the value of art objects. If the provenance of data produced by computer systems could be determined as it can for some works of art, then users, in their daily applications, would be able to interpret and judge the quality of data better [21]. In particular, the scientific and grid communities consider that, in order to support reproducibility, workflow management systems will be required to track and integrate provenance information as an integral product of the workflow [8]. Several surveys of provenance are available [3, 17, 25].

Against this background, the *International Provenance and Annotation Workshop* (IPAW'06), held on May 3-5, 2006 in Chicago, involved some 50 participants interested in the issues of data provenance, process documentation, data derivation, and data annotation [20, 2]. During a session on provenance standardisation, a consensus began to emerge, whereby the provenance research community needed to understand better the capabilities of the different systems, the representations they used for provenance, their similarities, their differences, and the rationale that motivated their designs. Hence, the first Provenance Challenge was born, and from the outset, the challenge was set up to be *informative* rather than *competitive*. In this editorial, we describe the challenge and provide a view on the contributions by the participating teams.

2. The Provenance Challenge

The challenge was defined by Simon Miles and Luc Moreau (U. of Southampton) and Mike Wilde and Ian Foster (U. of Chicago/Argonne Nat. Lab.) in May 2006; it was then reviewed by a larger group, including Juliana Freire (U. of Utah) and Jim Myers (NCSA), before a public review period by the IPAW'06 participants. It was published on June 19th and concluded with a two-day workshop, at GGF 18, in Washington, DC, on Sep. 13-14, 2006.

2.1. Instructions to Participants

The aim of the Provenance Challenge is to establish an understanding of the capabilities of available provenance-related systems and, in particular, examine the following.

- The representations that systems use to document details of processes that have occurred;
- The ability of each system to answer provenance-related queries;
- What each system considers to be within scope of the topic of provenance (regardless of whether the system can yet achieve all problems in that scope).

To help achieve the aim, a simple example workflow was defined to form the basis of the challenge. This workflow is inspired from a real experiment [28] in the area of Functional Magnetic Resonance Imaging (fMRI). Here, the term *workflow* [9] is used to denote a series of procedures performed in a system, each taking some data as input and producing other data as output. The procedures and workflows in the challenge problem are not defined in terms



of any particular technology (e.g., EXE files, Web Services, in the case of procedures; BPEL, compiled executable, batch file, in the case of workflows). Instead, participants could adopt their technology of choice.

Our focus in this challenge was on provenance and not on running the experiment. Hence, to facilitate take-up, we allowed challenge participants to implement procedures as “dummies”, i.e., as fake procedures that make use of the input, output, and intermediate data we provided, take the right input and produce the right output, but do not execute real code. Alternatively, participants could execute the real workflow after installing the necessary libraries.

Different systems use different representations for provenance information. In order to explore the capabilities of these different representations, we also defined a set of core queries, and asked participants to show how they addressed those queries.

Challenge participants were invited to upload the following information to the Provenance Challenge TWiki [5], to then allow comparison.

- Representation of the workflow in their system.
- Representation of provenance for the example workflow.
- Representation of the result of the core queries.
- Contributions to a matrix of queries vs. systems, indicating for each whether: (1) the query can be answered by the system, (2) the system cannot answer the query now but considers it relevant, (3) the query is not considered relevant to the project.

Each participant was also invited to optionally contribute the following.

- Additional queries (beyond the core queries) that illustrate the scope of their system;
- Extensions to the example workflow that the participant feels illustrates the unique aspects of their system;
- Any categorisation of queries that the project considers to have practical value.

2.2. The Provenance Challenge FMRI workflow

The purpose of the challenge workflow is to create population-based “brain atlases” from the fMRI Data Center’s archive of high resolution anatomical data. Specifically, for an input fMRI data set, it produces average images along the axis X, Y, and Z, after aligning each input sample with a reference image. The workflow comprises procedures and data items flowing between them, respectively shown as ovals and rectangles in Figure 1. The workflow can be seen as having five stages, where each stage is depicted as a horizontal row of the same procedure in the figure. Note that the term “stage” is introduced only to describe the workflow; we do not specify how “stages” should be realised in a concrete implementation.

Individual procedures employ the AIR (automated image registration) suite (bishopw.loni.ucla.edu/AIR5/index.html) to create an averaged brain from a collection of high resolution anatomical data, and the FSL suite (www.fmrib.ox.ac.uk/fsl) to create 2D images across each sliced dimension of the brain. In addition to the data items shown in the figure, there are other inputs to procedures (constant string options), details of which can be found on the Provenance Challenge TWiki [5].

The inputs to a workflow are a set [13] of new brain images (Anatomy Image 1 to 4) and a single reference brain image (Reference Image). All input images are 3D scans of a brain of

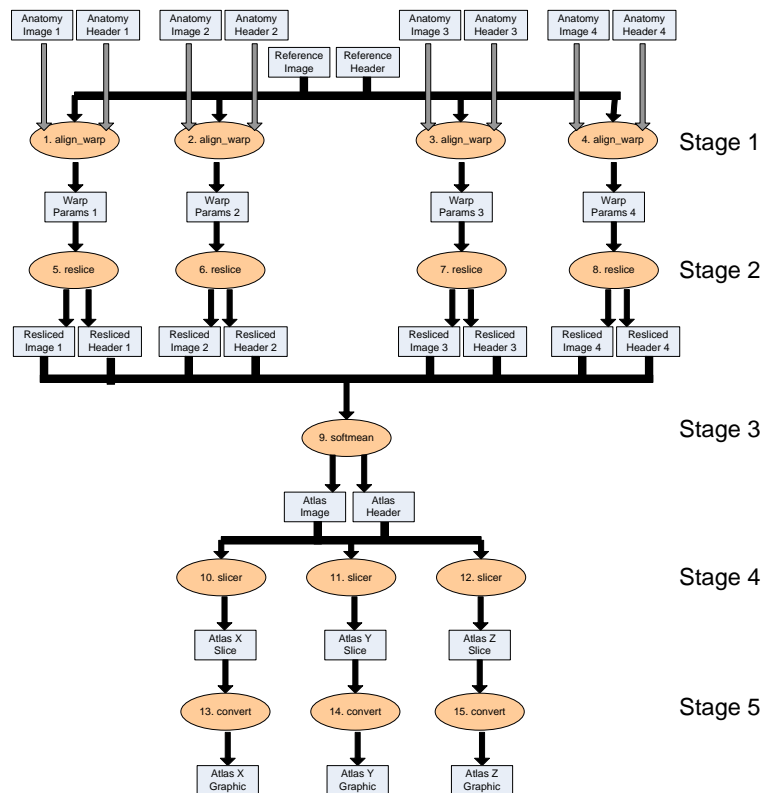


Figure 1. The Provenance Challenge Workflow

varying resolutions, so that different features are evident. For each image, there is the actual image and the metadata information for that image (Anatomy Header 1 to 4).

The stages of the workflow are as follows.

1. For each new brain image, `align_warp` compares the reference image to determine how the new image should be warped, i.e., the position and shape of the image adjusted, to match the reference brain. The output of each procedure in the stage is a *warp parameter set* defining the spatial transformation to be performed (Warp Params 1 to 4).
2. For each warp parameter set, the actual transformation of the image is done by `reslice`, which creates a new version of the original new brain image with the configuration defined in the warp parameter set. The output is a *resliced* image.
3. All the *resliced* images are averaged into one single image using `softmean`.



4. For each dimension (x, y and z), the averaged image is sliced, with the utility `slicer` , to give an atlas data set, i.e., a 2D atlas along a plane in that dimension, taken through the centre of the 3D image.
5. Each atlas data set is converted into a graphical atlas image using (the ImageMagick utility) `convert` .

2.3. Core Provenance Queries

In addition to the workflow, the challenge specified an initial set of provenance-related queries. These queries, based on the authors' experience [17], identify typical patterns of querying found in provenance systems.

- Q1.** Find the process that led to Atlas X Graphic / everything that caused Atlas X Graphic to be as it is. This should tell us the new brain images from which the averaged atlas was generated, the warping performed etc.
- Q2.** Find the process that led to Atlas X Graphic, excluding everything prior to the averaging of images with `softmean` .
- Q3.** Find the Stage 3, 4 and 5 details of the process that led to Atlas X Graphic.
- Q4.** Find all invocations of procedure `align_warp` using a twelfth order nonlinear 1365 parameter model (see model menu describing possible values of parameter “-m 12” of `align_warp`) that ran on a Monday.
- Q5.** Find all atlas graphic images outputted from workflows where at least one of the input Anatomy Headers had an entry `global_maximum=4095` . The contents of a header file can be extracted as text using the scanheader AIR utility.
- Q6.** Find all output averaged images of `softmean` (average) procedures, where the warped images taken as input were `align_warped` using a twelfth order nonlinear 1365 parameter model, i.e. “where `softmean` was preceded in the workflow, directly or indirectly, by an `align_warp` procedure with argument -m 12.”
- Q7.** A user has run the workflow twice, in the second instance replacing each procedures (convert) in the final stage with two procedures: `pgmtoppm` , then `pnmtojpeg` . Find the differences between the two workflow runs. The exact level of detail in the difference that is detected by a system is up to each participant.
- Q8.** A user has annotated some anatomy images with a key-value pair `center=UChicago` . Find the outputs of `align_warp` where the inputs are annotated with `center=UChicago` .
- Q9.** A user has annotated some atlas graphic images with a key-value pair where the key is `studyModality` . Find all the graphical atlas sets that have metadata annotation `studyModality` with values `speech` , `visual` or `audio` , and return all other annotations to these files.



3. An Analysis of Contributions to the Provenance Challenge

Following its publication, 16 teams responded to the challenge and submitted an entry to the Provenance Challenge TWiki [5]. This special issue contains each participating team's contribution to the challenge. Teams are referred to by the name of their system: Redux [1], Mindswap [12], Karma [26], JP [15], myGrid [27], VisTrails [22], ES3 [10], ZOOM [7], RWS [16], COMAD [4], PASS [24], SDG [23], NCSD2K and NCSCI [11], VDL [6], OPA [18], Wings/Pegasus [14].

In this section, we introduce a classification of the different approaches to help the reader gain a better understanding of provenance systems and their differences. To contrast the different approaches, we have identified a set of criteria, which have been grouped according to two categorisations.

Categorisation 1 is concerned with the broad characteristics of provenance systems, such as the environment in which they are embedded and the technologies they use. Such systems are usually developed in the context of research projects that have specific foci: understanding research motivations is also useful to appreciate some design decisions. Given that the purpose of provenance systems is to build a computer-based representation of provenance that can be queried and reasoned over, Categorisation 2 groups criteria pertaining to such representations; these criteria allow the reader to extract some of the fundamental concepts underpinning representations, and therefore, capabilities of systems. All findings are summarised in Figure 2.

Categorisation 1: Characteristics of Provenance Systems

C1.1 Execution Environment In many cases (but not all), provenance systems are embedded in a specific execution environment. The most common environments are workflow systems and operating systems. When embedded in a single execution environment, provenance representation may become (though does not have to be) dependent on the execution technology. On the one hand, such approaches may offer opportunities for optimisation, which indeed were exploited by some teams. On the other hand, it makes representations technology specific, and brings difficulties if applications are composed of several execution environments.

C1.2 Execution Environment (for the challenge) When systems allow for multiple execution environments, we indicate which one was actually used for the challenge.

C1.3 Representation Technology Provenance is represented and stored using a range of technologies, including relational databases (RDBMS), semantic web technologies (RDF, OWL), and internal private formats. Several systems also expose provenance according to an XML view.

C1.4 Query Language Systems offer query interfaces that operate over the stored representation of provenance. In some systems, the supported language is standard, whereas for others, it is purpose-built.



	Redux	Mindswap	Karma	JP	myGrid	VisTrails	ES3	ZOOM	RWS	COMAD	PASS	SDG	NCSD2K	NCSCI	VDL	OFA	Wings/Pegasus
1. Characteristics of Provenance Systems	Workflow system	Workflow system	Workflow system	Workflow system and job submission	Workflow system	Workflow system	Operating system	Workflow system	Workflow system	Workflow system	Operating system	Workflow system	Visual Program. Env	Workflow system	Workflow system	Workflow system	Workflow system
1.1 Execution Environment (actual system)	Web	Web	Web	EGEE gLite	Workflow system	Workflow system	Operating system	Workflow system	Workflow system	Workflow system	Operating system	Workflow system	Visual Program. Env	Workflow system	Workflow system	Workflow system	Workflow system
1.2 Execution Environment (for the challenge)	WWF	XBay and BPTEL	XBay and BPTEL	EGEE VO	VisTrails system and shell	VisTrails system and shell	IDL and Bash	Technology independent	Kepler, Ptolemy	Kepler, Ptolemy	Linux	Kepler, Ptolemy	DZK	Cyber Integrator	VDS		Wings and Pegasus
1.3 Provenance Representation	RDBMS	XML View and RDBMS	XML View and RDBMS	key-value pairs in RDBMS	XML View and RDBMS	XML View and RDBMS	XML view	RDBMS	Internal view	Internal and XML view	Internal	RDF (external and SAMWeb internal)	RDF	RDBMS	SQL	Internal and XML view	OWL and RDBMS
1.4 Query Language	SQL	Karma query API	Karma query API	JPS + JPPS	ProQA + JPS + JPPS	ProQA + JPS + JPPS	ESS	SQL + Transitive Closure	Internal Graph	Internal Graph	Custom	Semantic Querying	Extended (SQL)	SQL	+XML	Querying	XQuery, XQuery, XQuery
1.5 Research Emphasis	ER/SQ	ER	ER	RS/Q	ER/Q	ER/Q	RS/Q	S/Q	ER/Q	ER/Q	ER/S	ER/S	ER/S	ER/S	ER/S	ER/S	ER/S
1.6 Challenge Implementation	Run	Run	Run	Run	Simulated	Run	Run	Simulated	Partial	Partial	Run	Partial	Run	Run	Run	Run	Run
2. Properties of Provenance Representation	yes	no	no	yes	yes	yes	no	yes	yes	yes	no	yes	no	no	yes	no	yes
2.1 Includes workflow representation	E	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D
2.2 Distinguishes Causal Flow of Events	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)	(+AS/+AI)
2.3 Arbitrary annotations in scope/implemented	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)	(+TS/+TR)
2.4 Time supported/required	keys for data	GUIDs for data, processes, workflows	GUIDs for data, processes, workflows	no	keys for all data	keys for all data	assigned to every provenance object	assigned to every provenance object	yes for tokens	collections of members	yes, unique	ids (urls or lids) for data	URIs	URIs	logical names	logical names and file domain	logical file names and file domain
2.5 Naming required (if yes, then what)	no	port level	port level	any GUID assignable	changes to workflow	changes to workflow	I/O of any type	I/O of any type	collections of tokens	collections of tokens	collections of tokens	collections of tokens	any relationship	any relationship	any relationship	any relationship	any relationship
2.6 Tracked data, and granularity	contents	grouping on data, processes, and nested workflows	grouping on data, processes, and nested workflows	arbitrary	contents	contents	file or process	file or process	tokens	tokens	tokens	tokens	tokens	tokens	tokens	tokens	tokens
2.7 Abstraction mechanisms	layered provenance model	layered provenance model	layered provenance model	aspects	layered provenance model	layered provenance model	script or job steps	user view of composite steps	view of composite steps	view of composite steps	view of composite steps	view of composite steps	view of composite steps	view of composite steps	view of composite steps	view of composite steps	view of composite steps

Figure 2. Summary of Contributions



C1.5 Research Emphasis Teams have different research objectives when investigating provenance concepts. Their research may focus variously on techniques for *executing* (E) workflows such as the one defined in the challenge; *recording* (R) a description of a process being executed; *storing* (S) descriptions of process in persistent storage; and/or *querying* (Q) stored descriptions, in a way that captures the user’s interest.

C1.6 Challenge Implementation Some teams executed the challenge workflow (run); others executed the challenge with fake image processing components (partial), making use of the data and intermediary results published with the challenge definition; finally, others fully simulated its execution (simulated).

Categorisation 2: Properties of Provenance Representation

At some level of abstraction, provenance captures a notion of a causal graph, explaining how a data product or event came to be produced in an execution. However, there are variations on this theme, as indicated by the following criteria.

C2.1 Includes Workflow Representation Some systems assume that an explicit representation of a workflow is part of the provenance representation, whereas others do not have such an assumption, and hence rely on other means to describe executions.

C2.2 Data Derivation vs. Causal Flow of Events Some systems describe derivation of data (e.g., conversion was applied to “.pgm” input to produce “.gif” output), whereas others document causal flow of events (e.g., writing of a file is followed by its opening for reading). Some are capable of characterising both data and event oriented views.

C2.3 Annotations Annotations entered by users may provide valuable information pertaining to data products or executions. While most systems were able to support the challenge queries related to annotations, not all systems considered annotations to be in the scope of provenance. In the matrix **+AS** (resp. **-AS**) denotes that annotations are in scope of provenance (resp. not in scope), whereas **+AI** (resp. **-AI**) indicates annotations were implemented (resp. not implemented) for the challenge queries.

C2.4 Time A representation of provenance does not have to include time, but it is perceived that it is practical for users to be able to refer to time. Therefore, most systems support a notion of time, so that users can refer to executions or data products according to the time they took place or were produced. However, this requirement brings the challenge of identifying which clock to use, given that distributed clocks may return different times. In the matrix, **+TS** (resp. **-TS**) indicates that time is supported for challenge queries (resp. not supported), whereas **+TR** (resp. **-TR**) time denotes that time is required (resp. not required) for capturing a correct representation of provenance.

C2.5 Naming In order to be able to identify data products, some systems require each product to be identified by a unique name, typically created during workflow execution; such a name



can then be used to query about the provenance of data products. Other systems do not require names to be assigned, but see the identification of data items as a query in itself.

C2.6 Tracked data, Granularity Systems are capable of tracking the provenance of different kinds of data; some introduce restrictions on the granularity of data they can track the provenance of. For instance, systems may or may not deal with collections, files, bytes, or bits.

C2.7 Abstraction mechanisms When processes or data products are complex, it is useful to describe them with different levels of abstractions, sometimes hiding details of execution or representation, and at other times providing them. Some provenance systems provide support for this, by introducing new concepts in their provenance representation.

4. Conclusions

The rest of this special issue consists of papers describing the different systems summarised in Figure 2. We judge that the Provenance Challenge was highly successful, as measured by the number of participating teams, the quality of their submissions, and the discussions that resulted during the workshop. A number of lessons were learned from the challenge.

- At times, provenance queries were considered ambiguous. In the future, it would be interesting to specify them better, more precisely and unambiguously, and to characterise the performance implications of the queries.
- While most participating teams could tackle all queries, it is unclear yet whether they all obtained the same or equivalent answers.
- The community lacks consistent and coherent terminology for provenance-related concepts [19]. A consistent terminology would help outsiders to easily grasp issues and compare systems.

Following discussions at the two-day workshop, the provenance research community has decided to organise a second Provenance Challenge to address some of these issues in a systematic manner.

REFERENCES

1. Roger S. Barga and Luciano A. Digiampietri. Automatic capture and efficient storage of escience experiment provenance. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
2. Raj Bose, Ian Foster, and Luc Moreau. Report on the International Provenance and Annotation Workshop (IPAW06). *Sigmod Records*, pages 51–53, September 2006.
3. Rajendra Bose and James Frew. Lineage retrieval for scientific data processing: A survey. *ACM Computing Surveys*, 37(1):1–28, March 2005.
4. Shawn Bowers, Timothy M. McPhillips, and Bertram Ludäscher. Provenance in collection-oriented scientific workflows. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
5. <http://twiki.ipaw.info/bin/view/Challenge/FirstProvenanceChallenge>, June 2006.



6. Ben Clifford, Ian Foster, Mihael Hategan, Tiberiu Stef-Praun, Michael Wilde, and Yong Zhao. Tracking provenance in a virtual data grid. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
7. Sarah Cohen-Boulakia, Olivier Biton, Shirley Cohen, and Susan Davidson. Addressing the provenance challenge using zoom. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
8. Ewa Deelman and Yolanda Gil (Eds.). Workshop on the challenges of scientific workflows. Technical report, Information Sciences Institute, University of Southern California, May 2006.
9. Geoffrey C. Fox and Dennis Gannon. Special issue: Workflow in grid systems. *Concurrency and Computation: Practice and Experience*, 18(10), 2006.
10. James Frew, Dominic Metzger, and Peter Slaughter. Automatic capture and reconstruction of computational provenance. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
11. Joe Futrelle and Jim Myers. Tracking provenance semantics in heterogeneous execution systems. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
12. Jennifer Golbeck and James Hendler. A semantic web approach to tracking provenance in scientific workflows. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
13. Denise Head, Abraham Z. Snyder, Laura E. Gitton, John C. Morris, and Randy L. Buckner. Frontal-hippocampal double dissociation between normal aging and alzheimer's disease. *Cerebral Cortex*, (doi:10.1093/cercor/bhh174), September 2004. fMRI Data Center Accession Number: 2-2004-1168X.
14. Jihie Kim, Ewa Deelman, Yolanda Gil, Gaurang Mehta, and Varun Ratnakar. Provenance trails in the wings/pegasus system. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
15. Ales Krenek, Jiri Sitera, Ludek Matyska, Frantisek Dvorak, Milos Mulac, Miroslav Ruda, and Zdenek Salvat. glite job provenance – a job-centric view. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
16. Bertram Ludäscher, Norbert Podhorszki, Ilkay Altintas, Shawn Bowers, and Timothy M. McPhillips. Models of computation and provenance, and the RWS approach. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
17. Simon Miles, Paul Groth, Miguel Branco, and Luc Moreau. The requirements of recording and using provenance in e-science experiments. *Journal of Grid Computing*, 2006.
18. Simon Miles, Paul Groth, Steve Munroe, Sheng Jiang, Thibaut Assandri, and Luc Moreau. Extracting Causal Graphs from an Open Provenance Data Model. *Concurrency and Computation: Practice and Experience*, 2007.
19. Luc Moreau. Usage of 'provenance': A Tower of Babel. Towards a concept map — Position paper for the Microsoft Life Cycle Seminar, Mountain View, July 10, 2006. Technical report, University of Southampton, June 2006.
20. Luc Moreau and Ian Foster, editors. *Provenance and Annotation of Data — International Provenance and Annotation Workshop, IPAW 2006*, volume 4145 of *Lecture Notes in Computer Science*. Springer-Verlag, May 2006.
21. Luc Moreau, Paul Groth, Simon Miles, Javier Vazquez, John Ibbotson, Sheng Jiang, Steve Munroe, Omer Rana, Andreas Schreiber, Victor Tan, and Laszlo Varga. The Provenance of Electronic Data. *Communications of the ACM*, 2007.
22. Carlos Scheidegger, David Koop, Emanuele Santos, Huy Vo, Steven Callahan, Juliana Freire, and Claudio Silva. Tackling the provenance challenge one layer at a time. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
23. Karen Schuchardt, Tara Gibson, Eric Stephan, and George Chin, Jr. Applying content management to automated provenance capture. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
24. Margo Seltzer, David A. Holland, Uri Braun, and Kiran-Kumar Muniswamy-Reddy. Pass-ing the provenance challenge. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
25. Y. Simmhan, B. Plale, and D. Gannon. A survey of data provenance in e-science. *SIGMOD Record*, 34(3):31–36, September 2005.
26. Yogesh L. Simmhan, Beth Plale, and Dennis Gannon. Querying capabilities of the karma provenance framework. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
27. Jun Zhao, Carole Goble, Robert Stevens, and Daniele Turi. Mining taverna's semantic web of provenance. *Concurrency and Computation: Practice and Experience*, in this issue, 2007.
28. Yong Zhao, Jed Dobson, Ian Foster, Luc Moreau, and Michael Wilde. A Notation and System for Expressing and Executing Cleanly Typed Workflows on Messy Scientific Data. *Sigmod Record*, 34(3):37–43, September 2005.