

Provenance and data differencing for workflow reproducibility analysis

Paolo Missier^{1*}, Simon Woodman¹, Hugo Hiden¹, Paul Watson¹

¹ *School of Computing Science, Newcastle University, Newcastle-Upon-Tyne, UK*

SUMMARY

One of the foundations of science is that researchers must publish the methodology used to achieve their results so that others can attempt to reproduce them. This has the added benefit of allowing methods to be adopted and adapted for other purposes. In the field of e-Science, services – often choreographed through workflow, process data to generate results. The reproduction of results is often not straightforward as the computational objects may not be made available or may have been updated since the results were generated. For example, services are often updated to fix bugs or improve algorithms. This paper addresses these problems in three ways. Firstly, it introduces a new framework to clarify the range of meanings of “reproducibility”. Secondly, it describes a new algorithm, PDIFF, that uses a comparison of workflow provenance traces to determine whether an experiment has been reproduced; the main innovation is that if this is not the case then the specific point(s) of divergence are identified through graph analysis, assisting any researcher wishing to understand those differences. One key feature is support for user-defined, semantic data comparison operators. Finally, the paper describes an implementation of PDIFF that leverages the power of the e-Science Central platform which enacts workflows in the cloud. As well as automatically generating a provenance trace for consumption by PDIFF, the platform supports the storage and re-use of old versions of workflows, data and services; the paper shows how this can be powerfully exploited in order to achieve reproduction and re-use.

Copyright © 0000 John Wiley & Sons, Ltd.

Received ...

KEY WORDS: e-science, reproducibility, provenance, scientific workflow

1. INTRODUCTION

1.1. Motivation

A key underpinning of scientific discourse is the ability independently to verify or refute experimental results that are presented in support of a claim. With e-science in particular, such a need translates into expectations that the computational machinery used by one research group to produce the results be made publicly available in sufficient detail as to be used by third parties to construct a scientific argument in support of, or against, those results. Whilst it has long been the practice of scientific scholarly communication to require experimental methods to be described alongside the experimental data, it is generally accepted that reproducing an experiment based on its paper description requires substantial effort from independent researchers, either in the lab or in form of new program implementation.

Expectations regarding reproducibility in science have, however, been evolving towards higher levels of automation of the validation and reproduction process. This is driven by the observation

*Correspondence to: E-mail: Paolo.Missier@ncl.ac.uk

that actual acceleration of the production of scientific results, enabled by technology infrastructure which provides elastic computational facilities for “big data” e-science, is critically dependent upon the large-scale availability of datasets themselves, their sharing, and their use in collaborative settings.

Following this observation, new social incentives are being tested with the goal of encouraging, and in some cases forcing, the publication of datasets associated with the results published in papers. Evidence of increasing sharing can be found mostly in the form of online repositories. These include, amongst others, the recently created Thomson’s Data Citation Index[†]; the Dryad repository for published data underlying articles in the biosciences (datadryad.org/); the databib.org registry of data repositories (databib.org/) (over 200 user contributed repository entries at time of writing), as well as Best Practice guidelines for data publication, issued by data preservation projects like DataONE[‡].

This dissemination phenomenon combines with the increasing rate at which publishable data sets are being produced, resulting in the need for automated validation and quality control of the experimental methods and of their outcomes, all of which imposes new reproducibility requirements. Here the term *reproducibility* denotes the ability for a third party who has access to the description of the original experiment and its results to reproduce those results using a possibly different setting, with the goal to confirm or dispute the original experimenter’s claims. Note that this subsumes, and is hopefully more interesting, than the simpler case of *repeatability*, i.e. the attempt to replicate an experiment and obtain the same exact results when no changes occur anywhere.

In this paper we focus on the specific setting of workflow-based experiments, where reproduction involves creating a new version of the original experiment, while possibly altering some of the conditions (technical environment, scientific assumptions, input datasets). In this setting, we address the need to support experimenters in comparing results that may diverge because of those changes, and to help them diagnose the causes of such divergence. Depending on the technical infrastructure used to carry out the original experiment, different technical hurdles may stand in the way of repeatability and reproducibility. The apparently simple task of repeating a completely automated process that was implemented for example using workflow technology, after a period of time since the experiment was originally carried out, may be complicated by the evolution of the technical environment that supports the implementation, as well as by changes in the state of surrounding databases, and other factors that are not within the control of the experimenter. This well-known phenomenon, known as “workflow decay”, is one indication that complete reproducibility may be too ambitious even in the most favourable of settings, and that weaker forms of partial reproducibility should be investigated, as suggested by De Roure *et al.* [1]. Yet, it is easy to see that this simple form of repeatability underpins the ability to reuse experiments designed by others as part of new experiments. Cohen-Boulakia *et al.* elaborate on reusability specifically for workflow-based programming [2]: a older workflow that has ceased to function is not a candidate for reuse as part of a new workflow that is being designed today.

Additional difficulties may be encountered when more computational infrastructure with specific dependencies (libraries, system configuration) is deployed to support data processing and analysis. For some of these, well-known techniques involving the use of virtual machines (VMs) can be used. VMs make complex applications self-contained by providing a runtime environment that fulfills all of the applications dependencies. One can, for instance, create a VM containing an application consisting of a variety of bespoke scripts, along with all the libraries, databases and other dependencies that ensure the scripts can run successfully regardless of where the VMs are deployed, including on public cloud nodes.

[†]wokinfo.com/products_tools/multidisciplinary/dci/

[‡]notebooks.dataone.org/bestpractices/

1.2. Goals

VM technology is useful to provide a self-contained runtime environment, and indeed Groth *et al.* show that a combination of VM technology for partial workflow re-run along with provenance can be useful in certain cases to promote reproducibility [3]. However, our contention is that it is *the capability to compare execution traces and datasets with each other* that sits at the core of any argumentation that involves the reproduction of scientific results. Simple as it sounds, this aspect of reproducibility seems to have received relatively little attention in the recent literature, with a few exceptions as noted later. Yet, without appropriate operational definitions of *similarity of data*, one cannot hope to answer the question of whether the outcome of an experiment is a valid confirmation of a previous version of the same, or of a similarly designed experiment. The work described in this paper is focused primarily on this aspect of reproducibility, namely on adding data-comparison capability to an existing e-science infrastructure, and to make it available to user scientists in addition to other reproducibility features, as a tool for supporting scientific discourse.

Our long term goal is to support experimenters in answering two related questions. Firstly, given the ability to carry out (variations of) the same experiments, to what extent can the outcomes of those experiments be considered *equivalent*? Secondly, if they are not equivalent, what are the possible causes for the divergence?

Answering the first question is hard in general, as it requires knowledge of the type of the data, their format, and more importantly, their semantics. Consider for example two workflows, which process data in a similar way but at some stage employ two different statistical learning techniques to induce a classification model. Presumably, their outputs will be different for any execution of each of the two, and yet they may still be sufficiently close to be comparable. Thus, in this case a legitimate question is, are there suitable measures of data similarity that help support high-level conclusions of the form “the second method confirms the results of the first on a specific set of input datasets”? Answering this question requires knowledge of machine learning theory, namely to establish whether two classifiers can be considered equivalent, although the results they produce are not identical. While the general problem of establishing data equivalence is knowledge-intensive, limiting the scope to simpler structural comparison may still lead to effective conclusions. For instance, for some applications that operate on XML data, it may be enough to just define measures of similarity between two XML documents based on computing their structural differences, a problem successfully addressed by Wang *et al.* [4]. These examples suggest a spectrum of complexity in data similarity, starting from the simplest (two datasets are the same if their files are byte-wise identical), to structurally more complex data structures (XML “diff” operators), to end with the harder problems of semantics-based similarity as in the earlier machine learning example.

Our second question, regarding the diagnosis of the possible causes for dissimilarity of results, requires the ability to capture and then analyse the provenance traces associated with each of the experiments. Structurally, the provenance trace of a workflow execution is an acyclic digraph (DAG) in which the nodes denote either a piece of data or an activity, and the edges denote relationships amongst data and activities, which reflect the production and consumption of data throughout the workflow execution steps (a simple but more formal definition is given in Sec. 2). While the problem of comparing two traces is in principle a classic one of finding subgraph isomorphisms in directed labelled graphs [5, 6], greatly simplifying assumptions can be made by observing that workflow provenance traces are a specific type of acyclic digraph with additional labelling, namely regarding the ports on which activities produce and consume data items, which make it possible to reduce the space of potential matches. Our assumption is that two provenance traces pertaining to two similar workflows are expected to be structurally similar, and to contain data nodes that are comparable, in the sense clarified above. When this is the case, then there is hope that one can detect points in the traces where results begin to diverge, and to diagnose the causes of such divergence by pairwise graph traversal.

1.3. Contributions

We offer two main contributions, both set in the context of workflow-based e-science. Our reference implementation is based upon the e-Science Central workflow management system, developed by Hiden *et al.* in our group [7]. The first contribution (Sec. 2) is a framework to characterise the *reproducibility space*, defined by the change in time of one or more of the technical elements that make up a computational experiment, leading to various forms of “workflow decay”, i.e. disfunctional or broken workflows.

The second contribution, described in Sec. 3, is PDIFF, an algorithm and prototype implementation to help users analyse diverging outputs that occur when trying to reproduce workflows over time. This is complemented by an initial set of data comparison functions that address typical data types used in e-Science Central workflows. As these workflows are routinely used to produce, for example, machine learning models for chemical engineering research [8], our suite of data comparison functions includes the statistical similarity of predictive models produced using machine learning techniques (in addition to comparing plain text files and XML documents). This set of functions is easily extensible, as they are implemented using workflow blocks, which can be rapidly implemented and deployed on the middleware.

Finally, in Sec. 4 we describe how PDIFF and the data diff functions leverage the current e-Science Central middleware and complement the functionality of its existing provenance management sub-system. The technique described in the paper is not specific to the e-Science Central platform, however. Indeed, since e-Science Central provenance is encoded using the Open Provenance Model [9] (and will in the near future be aligned with the new PROV model [10]), it applies to any provenance system that adopts the standard.

1.4. Related work

The perceived importance of enabling transparent, reproducible science has been growing over the past few years. There is a misplaced belief, however, that adopting e-science infrastructures for computational science will automatically provide longevity to the programs that encode the experimental methods. Evidence to the contrary, along with some analysis of the causes of non-reproducibility, is beginning to emerge, mainly in the form of editorials [11], [12]. The argument in support of reproducible research, however, is not new, dating back from 2006 [13]. More recently, Dummond and Peng separately make a distinction between repeatability and reproducibility [14], [15]. Bechhofer *et al.* make a similar distinction concerning specifically the sharing of scientific data [16].

In high profile Data Management conferences like SIGMOD, established initiatives to encourage contributors to make their experiments reproducible and shareable have been ongoing for a number of years, with some success[§].

Technical solutions for ensuring reproducibility in computational sciences date back over 10 years [17], but have matured more recently with workflow management systems (WFMS) like Galaxy [18], [19] and VisTrails [20]. In both Galaxy and VisTrails, repeatability is facilitated by ensuring that the evolution of the software components used in the workflows is controlled by the organizations that design the workflows. This “controlled services” approach is shared by other WFMS such as Kepler [21] and Knime (www.knime.org/), an open-source workflow-based integration platform for data analytics. This is in contrast to WFMS like Taverna [22], which are able to orchestrate the invocation of operations over arbitrary web services, which makes it very general but also very vulnerable to version changes as well as maintenance issues of those services. A recent analysis of the decay problem, conducted by Zhao *et al.* on about 100 Taverna workflows authored between 2007 and 2012 that are published on the myExperiment repository, reveals that more than 80% of those workflows either fail to produce the expected results, or fail to run altogether [23]. Such a high failure rate is not surprising given that experimental workflows by their nature tend to make use, at least in part, of experimental services which are themselves subject to frequent changes or

[§]The SIGMOD reproducibility initiative: www.sigmod.org/2012/reproducibility.shtml.

may cease to be maintained. There are no obvious solutions to this problem, either, but the analysis technique described in this paper applies to Taverna provenance traces [24], which comply with the Open Provenance Model (and now with the W3C PROV model [10]). In this respect, e-Science Central sits in the “closed world” camp, with versioned services data and workflows stored and managed by the platform. By default a workflow will make use of the latest version of a service, but this can be overridden, to the extent that past versions are still available on the platform. In any case, the version history of each service can be queried, a useful feature for the application described in this work.

The use of provenance to enable reproducibility has been explored in several settings. Moreau showed that a provenance trace expressed using the Open Provenance Model can be given a denotation semantics (that is, an interpretation) that effectively amounts to “executing” a provenance graph and thus reproducing the computation it represents [25]. In the workflow setting, to the best of our knowledge, VisTrails and Pegasus/Wings [26] are the only WFMS where the evolution of workflows is recorded, leading to a complete provenance trace that accounts for all edit operations performed by a designer to create a new version of an existing workflow, as shown by Koop *et al.* [27]. A similar feature is available in e-Science Central, namely in the form of a record history of all versions of a workflow. Whilst edit scripts can potentially be constructed from such history, those are currently not explicit or in the form of a provenance trace. Thus, our divergence analysis is solely based upon the provenance of workflow executions, rather than the provenance of the workflow itself. Note that this makes the technique more widely applicable, i.e. to the majority of WFMS for which version history is not available.

Our technique falls under the general class of graph matching [28]. Specifically relevant is the heuristic polynomial algorithm for change detection in un-ordered trees, by Chawathe and Garcia-Molina [29]. Although matching provenance traces can be viewed as a change detection problem, it is less general because additional constraints on the data nodes, derived from the semantics of provenance graphs, reduce the space of possible matching candidates for any given node.

The idea of comparing workflow and provenance traces was explored by Altintas *et al.* [30] for the Kepler system, with the goal of exploiting a database of traces as a cache to be used in lieu of repeating potentially expensive workflow computation whenever possible. This “smart rerun” is based on matching provenance trace fragments with an upcoming computation fragment in a workflow run, and replace the latter with the former when a match is found. This serves a different purpose from PDIFF, however. Whilst in SRM success is declared when a match is found, in PDIFF the interesting points in the graphs are those where divergence, rather than matching nodes, are detected.

Perhaps the most prominent research result on computing the difference of two provenance graphs is that of Bao *et al.* [31]. The premise of the work is that general workflow specifications that include arbitrary loops lead to provenance traces for which “differencing” is as complex as general subgraph isomorphism. The main insight is that the differencing problem becomes computationally feasible under the additional assumption that the workflow graphs have a series-parallel (s-p) structure, with well-nested forking and looping. While the e-Science Central workflow model does allow the specification of non-s-p workflows, its loop facility is highly restricted, i.e. to a map operator that iterates over a list. Additionally, PDIFF relies on named input and output ports to steer the coordinated traversal of two provenance graphs, as explained in Sec. 3.1. The combination of these features leads to a provenance comparison algorithm where each node is visited once.

2. A FRAMEWORK FOR REPRODUCIBILITY IN COMPUTATIONAL EXPERIMENTS

The framework presented in this section applies to generic computer programs for which provenance traces can be generated. However, we restrict our attention to scientific workflows and workflow management systems, as the approach described in the rest of the paper has been implemented specifically for a workflow-based programming environment, namely eScience Central. Scientific workflows and their virtues as a programming paradigm for the rapid prototyping of science applications have been described at length by Deelman *et al.* [32]. For our purposes, we adopt

a minimal definition of workflow as a directed graph $W = \langle T, E \rangle$ consisting of a set T of computational units (*tasks*) which are capable of producing and consuming data items on *ports* P and a set $E \subset T \times T$ of edges representing data dependencies amongst processors, such that $\langle t_i.p_A, t_j.p_B \rangle \in E$ denotes that data produced by t_i on port $p_A \in P$ is to be routed to port $p_B \in P$ of task t_j during the course of the computation[¶]. We write:

$$tr = exec(W, d, ED, wfms) \quad (1)$$

to denote the execution of a workflow W on input data d , with the additional indication that W depends on external elements such as third-party services and database state, collectively denoted ED (for “External Dependencies”). W also depends on a runtime environment $wfms$, namely the workflow management system.

The outcome of the execution is a graph-structured *execution trace* tr . A trace is defined by a set D of *data items*, a set A of *activities*, and a set $\mathcal{R} = \{R_1 \dots R_n\}$ of relations $R_i \subset (D \cup A) \times (D \cup A)$. Different provenance models, including the OPM and PROV, introduce different types of relations. For the purpose of this work, we adopt a simple model $\mathcal{R} = \{used, genBy\}$ where $used \subset A \times P \times D$ and $genBy \subset D \times A \times P$ denote that an activity a has used data item d from its port p , written $used(a, p, d)$, and that d has been generated by a on port p , written $genBy(d, a, p)$, respectively. Thus, a trace is simply a set of *used* and *genBy* records which are created during one workflow run, from the observation of the inputs and outputs of each workflow block that is activated. In particular, activities that appear in a trace correspond to instances, or invocations, of tasks that appear in the workflow. Distinguished data items in a trace for a workflow run are the inputs and the outputs of the workflow, denoted $tr.I$ and $tr.O$ respectively. Formally, $tr.I = \{d \in D \mid a \in A \Rightarrow (d, a) \notin genBy\}$. Symmetrically, $tr.O = \{d \in D \mid a \in A \Rightarrow (a, d) \notin used\}$.

2.1. Evolution of workflow execution elements

At various points in time, each of the elements that participate in an execution may evolve, either independently or in lock-step. This is represented by subscripting each of the elements in (1), as follows:

$$tr_t = exec_t(W_i, ED_j, d_h, wfms_k), \text{ with } i, j, h, k < t \quad (2)$$

This indicates that a specific version of each of the elements that participate in the execution is used at time t . Suppose for instance that version 1 is initially defined at time t_1 for all elements: $\langle W_1, ED_1, d_1, wfms_1 \rangle$. One execution at this time produces tr_1 . Then, at time t_2 , a new version W_2 of the workflow is produced, everything else being unchanged. When executed, the workflow produces tr_2 . At time t_3 , further changes occur to the external dependencies and to the input dataset, leading to a new combination $\langle W_2, ED_3, d_3, wfms_1 \rangle$ and to the new result tr_3 . Fig. 1 shows this incremental evolution, with changes happening at t_4 to the WFMS itself, from $wfms_1$ to $wfms_4$, and finally once more to the dependencies, from ED_3 to ED_5 . The last column in the figure indicates “default” execution, that is, executions based on the latest available version of each element.

We use this simple model to give a precise meaning to the informal notions of repeatability and reproducibility introduced earlier, as follows.

- We say that a baseline execution of the form (2) is *repeatable* if it is possible to compute (2) at some time $t' > t$.
- We say that execution (2) is *reproducible* if it is possible to compute a new version of (2) at time $t' > t$, i.e.:

$$tr_{t'} = exec_{t'}(W_{i'}, ED_{j'}, d_{h'}, wfms_{k'})$$

in the case where at least one of the elements involved in the original execution at time t has changed, i.e., $W_i \neq W_{i'}$, $ED_j \neq ED_{j'}$, $d_h \neq d_{h'}$, or $wfms_k \neq wfms_{k'}$.

[¶]More complete formalizations of workflow models are available, but this level of abstraction will suffice for our purposes.




Figure 1. Example of evolution timeline

Notice that a situation where data changes: $d_h \rightarrow d_{h'}$ but W_i remains the same, or vice versa, may require additional adapters to ensure that the new input can be used in the original workflow, or that the new workflow can accommodate the original input, respectively.

These simple definitions entail a number of issues that make repeating and reproducing executions complex in practice. Firstly, execution repetition becomes complicated when at least one the elements involved in the computation has changed between t and t' , because in this case each of the original elements must have been saved and still made available. Secondly, reproducibility is not trivial either, because in general there is no guarantee that the same workflow can be run with new external dependencies, or using a new version of the WFMS, or using a new input dataset.

Furthermore, we interpret the notion of reproducibility as a means by which experimenters validate results using deliberately altered conditions ($W_{i'}$, $d_{j'}$). In this case, the goal may not be to obtain exactly the original result, but rather one that is similar enough to it to conclude that the original method and its results were indeed valid. Thus, we propose that reproducibility requires the ability to *compare the results* $tr_t.O$, $tr_{t'}.O$ of two executions, using dedicated similarity functions, i.e., $\Delta_D(tr_t.O, tr_{t'}.O)$ (for example, one such function may simply detect whether the two datasets are identical). However, a further complication arises as the criteria used to compare two data items in general depends on the type of the data, requiring the definition of a whole collection of Δ_D functions, one for each of possibly many data types of interest.

A final problem arises when one concludes, after having computed $exec_t(\dots)$, $exec_{t'}(\dots)$, and $\Delta_D(tr_t.O, tr_{t'}.O)$, that the results are in fact not similar enough to conclude that the execution has been successfully reproduced. In this case, one would like to identify the possible causes for such divergence.

We address these problems in the rest of the paper. In particular, after discussing general techniques for supporting repeatability and reproducibility, we focus on the last problem, namely the role of provenance traces in diagnosing divergence in non-reproducible results, and describe a prototype implementation of such a diagnostic tool built as part of the eScience Central WFMS, underpinned by simple but representative data similarity functions.

2.2. Mapping the reproducibility space

Our formulation can be used to summarize some of the repeatability/reproducibility scenarios mentioned earlier. The grid in Fig. 2 illustrates a distinction between elements whose evolution is under the control of the experimenter, namely W and d , and those that are beyond their control, namely ED and $wfms$. Ignoring the latter for a moment, one can interpret the co-evolution of methods and data in terms of objectives that are part of the scientific discourse, ranging from simply repeating the experiment (as defined above), to observing the effect of changes in the method, in the input data, or both. Our implementation of a provenance-based divergence analysis with associated data similarity functions is aimed mainly at supporting the latter three scenarios.

When external factors evolve as well, various forms of *decay* occur, leading to workflows that no longer deliver the intended functionality, or cannot be executed altogether. Our analysis tool may

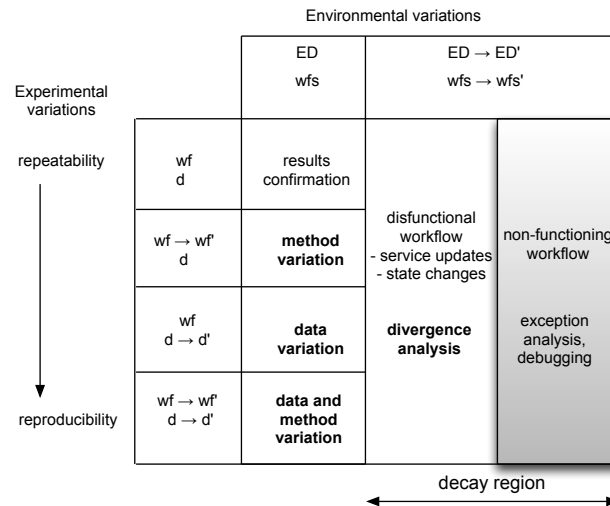


Figure 2. Mapping the reproducibility space

help diagnose the former case, where the results are not as expected and the cause is to be found in new versions of third-party services, updated database states, or the non-deterministic behaviour of some of the workflow blocks. In contrast, exceptions analysis is required to diagnose a broken workflow, and it is out of the scope of this work.

3. PROVENANCE AND DATA DIFFERENCING

3.1. Detecting divergence in execution traces

In this section we present PDIFF, an algorithm designed to compare two provenance traces, tr_t , $tr_{t'}$, that are generated from related executions, i.e., of the form (2) where at least one of the four contributing elements has changed between times t and t' . Specifically, the algorithm attempts to identify differences in the two traces, which are due to one of the causes portrayed in Fig. 2, namely workflow evolution, input data evolution, as well as non-destructive external changes such as service and database state evolution. Note that we only consider cases where both executions succeed, although they may produce different results. We do not consider the cases of workflow decay that lead to nonfunctional workflows, as those are best analysed using standard debugging techniques, and by looking at runtime exceptions.

The aim of the divergence detection algorithm is to identify whether two execution traces are identical, and if they are not, where do they diverge. Specifically, given two outputs, one from each of the two traces, the algorithm tries to match nodes from the two graphs by traversing them in lock-step as much as possible, from the bottom up, i.e., starting from the final workflow outputs. When mismatches are encountered, either because of differences in data content, or because of non-isomorphic graph structure, the algorithm records these mismatches in a “delta” data structure and then tries to continue by searching for new matching nodes further up in the traces. The delta structure is a graph whose nodes are labelled with pairs of trace nodes that are expected to match but do not, as well as with entire fragments of trace graphs that break the continuity of the lock-step traversal. Broadly speaking, the graph structure reflects the traversal order of the two graphs. A more detailed account of the construction can be found in Sec. 3.2. Our goal is to compute a final *delta tree* that can be used to explain observed differences in workflow outputs, in terms of differences throughout the two executions. We first illustrate the idea in the case of simple data evolution:

$$tr_t = exec_t(W, ED, d, wfms), tr_{t'} = exec_{t'}(W, ED, d', wfms)$$

The workflow for this example has three inputs and one output, and is portrayed in Fig. 3. The two execution traces in Fig. 4, denoted A and B, correspond to inputs $tr_t.I = \{d_1, d_2, d_3\}$ and $tr_{t'}.I = \{d_1', d_2', d_3\}$, respectively. The two outputs $d_F, d_{F'}$ differ, as expected, however note that not all of the inputs contribute to this mismatch. Indeed, observe that input d_1' on S_0 produces the same result as input d_1 . Running PDIFF on these two traces produces the delta graph, indeed a simple path in this case, shown in the rightmost part of Fig. 4. The path consists of pairs of mismatching nodes, that can be used to explain the mismatch $d_F \neq d_{F'}$. We can see for example that $d_2 \neq d_2'$ caused all of the intermediate data products downstream from S_1 to differ, eventually explaining the difference in the outputs.

PDIFF works by traversing the two traces from the bottom up. As long as the two graphs are isomorphic, as is the case in this simple example, the algorithm proceeds in a breadth-first fashion up the traces, simply testing for differences in data and activity nodes, alternately. A new node is added to the delta graph whenever mismatches are encountered. The structure is a graph, rather than a tree, because different branches can sometimes be merged, as shown in the next example. Note that mismatches may occur even when the input datasets are identical, for example when one or more of the services exhibits non-deterministic behaviour, or depends on external state that has changed between executions.




Figure 3. Simple workflow (and its representation in the e-Science Central GUI) to illustrate divergence analysis in the presence of data evolution




Figure 4. Example traces for two executions of the workflow of Fig. 3, and simple delta

The case of non-isomorphic traces, which arise in the case of workflow evolution or of evolution in the version of some of the services, present additional challenges to the matching algorithm. As a second, more elaborate example, consider the two workflows in Fig. 5 (not that these are two different evolutions of the workflow of Fig. 3). The main differences between them are as follows:

- S_0 is followed by $S_{0'}$ in W_A but not in W_B ;
- S_3 is preceded by $S_{3'}$ in W_B but not in W_A ;
- S_2 in W_A is replaced by a new version, S_{2v2} , in W_B ;
- S_1 in W_A is replaced by S_5 in W_B .

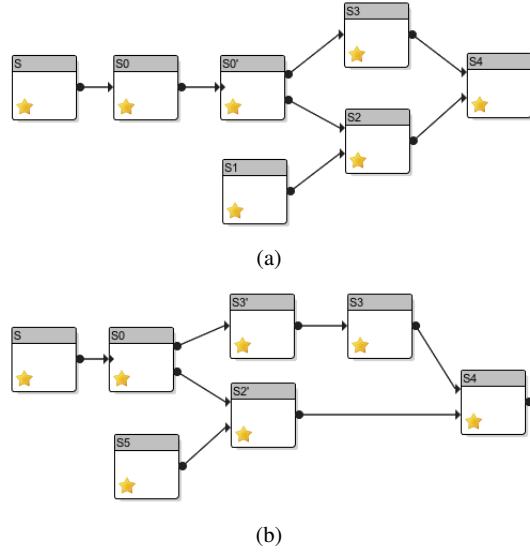


Figure 5. Example workflows to illustrate general divergence analysis

Let us assume that the input data is the same for both workflows, resulting in the following scenario:

$$tr_t = exec_t(W, ED, d, wfms), tr_{t'} = exec_t(W', ED', d, wfms)$$

where ED' consists of the version change in service S_2 . PDIFF operates on the assumption that the two traces are not too drastically different, and on the expectation that workflow evolution can be described using a combination of the patterns listed above. Although the algorithm provides no formal guarantees that all possible cases of graph structure mismatches will be correctly identified, this informal assumption is true for the two traces portrayed in Fig. 6, corresponding to one execution of each of the two example workflows. On these traces, PDIFF produces the delta graph shown in Fig. 7. The algorithm operates as follows.

Initially, the outputs $tr_t.O = x$ and $tr_{t'}.O = x'$ are compared, and the initial mismatch $x \neq x'$ is recorded in the delta graph. Note that the outputs are uniquely determined by a combination of the unique workflow run, the ID of the output block that produces them, and the name of the output port. Moving up, the services that produce x and x' , S_4 , are found to be the same, and no action is taken. As S_4 has consumed two inputs^{||}, the traversal branches out to follow each of the two ports, p_0, p_1 . Considering the left branch, the mismatching transient data $y \neq y'$ is detected and added to the delta graph. The traversal then continues past w, w' , when a service mismatch $S_{0'}, S_{3'}$ is detected. This is recorded in the delta graph, but it also triggers an exploration upwards on both traces, to find the nearest matching service node, if it exists, which would allow the pairwise comparison to resume. In this case, S_0 is such a “sync” node (it happens to be at the same level in both graphs in this example, but this is not necessarily the case). The graph fragments that were skipped during the search on each trace represent relevant structural differences in the two graphs, and are therefore recorded as part of the delta graph as well. These represent structural variations that occurred in each of the two

^{||}The algorithm assumes that the same service will always have the same number of ports.

graphs. In the next step, a version mismatch is detected between S and S_{v2} , and again added to the graph. The final step on this branch finds identical original input data, d_1 .




Figure 6. Example traces for two executions of the workflow of Fig. 5

A similar traversal occurs on the right side of the traces, starting at port p_1 of S_4 . The mismatch $z \neq z'$ is first recorded, followed by $S_2 \neq S_{2v2}$. In this case, however, PDIFF is able to detect that the mismatch is due to a simple service version upgrade, as service versions are available as part of the provenance trace, and thus it does not trigger any further upwards search. One more branching occurs at this point: on the left, $h \neq h'$, followed by a further service mismatch, $S_{0'} \neq S_0$. As we saw earlier, this causes PDIFF to search for the nearest matching service node, which is again S_0 . Note however that in this instance, the “diff” graph fragment on the right trace is null, because $S_{0'}$ has been added to one of the two workflows. The traversal resumes from S_0 on both sides. Because this node was used earlier as a sync point on the left branches of the two traces, any further traversal upwards in the graph would follow exactly the same steps, and therefore it can be avoided. Furthermore, the delta path starting with pair $\langle S, S_{v2} \rangle$ can be shared with this branch of the delta graph, resulting in a join in the graph as shown at the top of Fig. 7. Finally, we come to the p_1 branch of S_{2v2} , where $k \neq k'$ and then $S_1 \neq S_5$. Although this does trigger a new upwards search, this ends immediately upon finding the input data.

We have mentioned that PDIFF assumes that the two traces are similar enough for this comparison to progress to the end. Large differences amongst the traces are manifested as divergence that cannot be reconciled later in the graph traversal. In this case, the earlier divergence point is reported. For instance, if the same output is produced by drastically different workflows, then PDIFF reports that there are no sync points at all.

3.2. Support functions

The pseudo code for the algorithm appears in Alg. 1, using a conceptual recursive formulation for clarity. The support functions used by the algorithm are described informally, as follows.

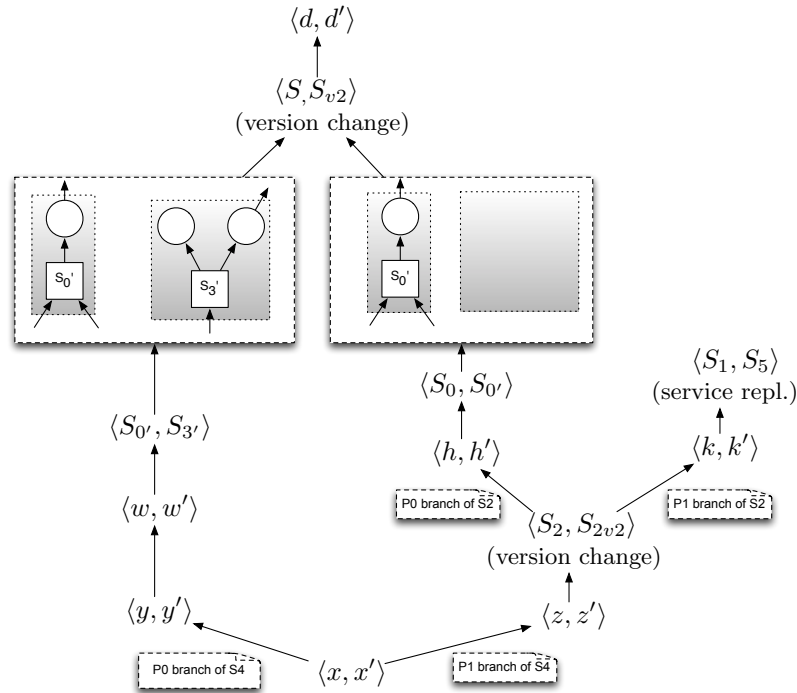


Figure 7. Delta graph from PDIFF on the traces of Fig. 6

- Functions $UPD(d)$ and $UPS(a, p)$ perform one step of the upwards graph traversal. Specifically, $UPD(d)$ traverses a $genBy(d, a)$ relation where d is a data node and returns the node for the activity a that generated d . $UPS(a, p)$ traverses a $used(a, d, p)$ relation where a is an activity node, and returns the node for the data d used by a through port p .
- Function $DDIFF(nL, nR)$ defines datatype-specific comparison of data. This is used when matching the inputs and outputs of two workflow executions, i.e., $tr_t.I$ and $tr_t.O$. Examples of such data comparison functions are presented in the next section. Note that intermediate results, represented as internal data nodes in the traces above, are transient data which are not explicitly stored in the provenance database. To save space, only their MD-5 hash is recorded in the node. This results in a crude but very efficient data mismatch test for internal nodes with boolean result. In contrast, in the case of actual data the function may produce a similarity value between 0 and 1. For simplicity, we assume that a threshold is applied, so that data diff functions all return TRUE if significant differences are detected, and FALSE otherwise**. Similarly, function $ADIFF(nL, nR)$ compares two service nodes using the service identifiers, and returns TRUE if the services are not the same. It can further flag differences that are due solely to mismatches in service versions.
- $FINDNODE(nL, nR)$ performs the search for a new matching node following a service node mismatch. It returns a 4-tuple consisting of the graph fragments that were skipped during the search (either or both can be null), as well as references to the matching service nodes in each of the two traces, if they exist (note that either both of these are valid references, or neither is). More formally, consider the subgraphs upwards from nL and nR , i.e., upL

**Whilst the transient nature of intermediate data does preclude the use of content-based comparison functions, note that one can control the granularity of provenance, and thus the points where actual data is stored with it, by breaking down a complex workflow into a composition of smaller ones. Using small workflows is indeed typical for existing e-Science Central applications.

Algorithm 1 PDIFF: Provenance trace divergence detection

```

function PDIFF( $\Delta, nL, nR$ )
  if  $nL$  is null or  $nR$  is null then
    return  $\Delta$ 
  end if
  if  $nL.type == DATA\_NODE$  then ▷ when comparing two data nodes:
    if DDIFF( $nL, nR$ ) then ▷ data mismatch
       $\Delta \leftarrow \text{ADDDelta}(\Delta, \langle nL, nR \rangle)$  ▷ add this mismatch to  $\Delta$ 
    end if
    return PDIFF( $\Delta, \text{UPD}(nL), \text{UPD}(nR)$ ) ▷ and continue up the graph
  end if
  if  $nL.type == ACTIVITY\_NODE$  then ▷ when comparing two activity nodes:
    if ADIFF( $nL, nR$ ) then ▷ activity mismatch
      if (different version of the same service) then
         $\Delta \leftarrow \text{ADDDelta}(\Delta, \langle nL, nR \rangle)$  ▷ add this version mismatch to  $\Delta$  and continue
      else
         $\Delta \leftarrow \text{ADDDelta}(\Delta, \langle nL, nR \rangle)$  ▷ add this service mismatch to  $\Delta$ 
         $\langle GL, GR, nL', nR' \rangle \leftarrow \text{FINDNODE}(nL, nR)$  ▷ resync traversal if possible
        if  $nR'$  is null and  $nL'$  is null then ▷ resync failed
          return  $\Delta$  ▷ return
        end if
         $\Delta \leftarrow \text{ADDDelta}(\Delta, \langle GL, GR \rangle)$  ▷ resync succeeded, record mismatching graphs
         $\text{stopFlag} \leftarrow \text{ISDELTAStop}(\Delta)$  ▷ has a branch been joined in the delta graph?
        if  $\text{stopFlag}$  then
          return  $\Delta$  ▷ sync'ed to a previous sync point, delta graphs joined. this traversal terminates
        end if
        if  $nR$  is not null then
           $nR \leftarrow nR'$ 
        end if
        if  $nL$  is not null then
           $nL \leftarrow nL'$ 
        end if
      end if
    end if
  end if
  return MERGEDELTA( $\{ \text{PDIFF}(\Delta, \text{UPS}(nL, p), \text{UPS}(nR, p)) \mid p \in nL.ports \}$ ) ▷ traverse each input port
end if
end function

```

$= \{ \text{UPS}(nL, p_i) \}$ for all input ports p_i of nL , and $upR = \{ \text{UPS}(nR, p_j) \}$ for all input ports p_j of nR . FINDNODE finds the closest matching activity nodes, i.e., the earliest encountered activity nodes nL' in upL and nR' in upR such that ADIFF(nL', nR') is false. It records the mismatching nodes encountered in the new graph fragments GL, GR , respectively, or it returns null node identifiers if the search fails.

- Finally, function ADDDELTA($\Delta, \text{deltaElement}$) adds *deltaElement* to Δ and returns the new delta graph. This function also needs to be able to join tree branches on sync points as discussed in our previous examples. Suppose a delta node g containing two graph trace fragments is added to the delta, and suppose n is next sync node found by FINDNODE. In this case, g is annotated with n . If n is found again at a later time following a new graph mismatch g' , then this is an indication that the traversal can stop, and that the graph above g can be shared with g' . Fig. 7 shows one example. Auxiliary function ISDELTAStop(Δ) returns true if a merge has occurred, indicating that the current traversal may terminate.

This function is generalized by MERGEDELTA(*deltaset*), which combines a set of Δ graphs into a new graph by making them siblings of a new common root.

We now briefly analyse the complexity of PDIFF. Let N, M be the total number of nodes in the left and right graph, respectively. The trace traversal is driven by functions UPD() and UPS(), which perform one step upwards in each of the two graphs. As noted, port names ensure that only matching *genBy()* (resp. *used()*) relations are traversed. Thus, if the two traces were identical, PDIFF would complete the traversal in $N = M$ steps. On the other hand, a mismatch amongst two activity nodes may trigger a call to the FINDNODE() operator, which involves searching for pairs of matching activity nodes across the two traces. Here port names cannot be used to assist the matching, which

therefore requires one activity node lookup in one of the graph, for each activity node in the other graph. In the worst case, this results in $N.M$ comparisons for a call to `FINDNODE()` that occurs near the bottom of the graph (i.e., at the first activity node from the bottom). As the traversal progresses upwards, more calls to `FINDNODE()` may occur, but they operate on progressively smaller fragments of the graphs, and any new search will stop at the first sync node that was found in a previous search (i.e., when function `ISDELTA STOP()` returns true). Considering the extreme cases, suppose that divergence is detected at the bottom of the graphs, and the only sync node is at the very top. This involves only one call to `FINDNODE()`, with cost approximately $N.M$, and then `PDIFF` terminates. In this case, the algorithm involves at most $\min(N, M) + N.M$ steps. The other extreme case occurs when the sync nodes are one step away from the diverging nodes. Now `FINDNODE()` is called repeatedly, at most $\min(N, M)$ times, but each time it completes after only one comparison, resulting in $2.\min(N, M)$ steps.

To summarize, `PDIFF` is designed to assist reproducibility analysis by using provenance traces from different executions to identify the following possible causes of diverging output values for those executions:

- data and workflow evolution;
- non-deterministic behaviour in some of the services, including the apparent non-determinism due to changes in external service (or database) state;
- service version upgrades.

3.3. Differencing data (*DDiff*)

Determining the similarity between two data sets is highly dependant not only on the type and representation but also the semantics of the data. Not all data sets can be compared using a byte-wise comparator to determine if they are equivalent. For example, two models produced from the same data set may well be equivalent but are unlikely to be byte-wise identical. In addition, a Boolean response of equivalent or not equivalent is a somewhat naïve approach and may provide false negatives and be less amenable to statistical analysis.

Within e-Science Central we have built a flexible architecture in which it is possible to embed different algorithms for different types and representations of data. Following the standard extensibility model of e-Science Central, new data comparison functions can be implemented as workflow blocks (typically as Java classes that interact with the system's API).

We provide algorithms which calculate the equivalent of three classes of data. The algorithms are implemented as e-Science Central workflows (described in Section 4) which implement a common interface allowing the system to determine dynamically which one to use for the relevant data set.

Text or CSV Comparing textual files is straightforward, and tools such as GNU `diff` or `cmp`^{††} are commonly used and well understood. By default GNU `diff` works by comparing each file, line by line, and returning the differences, known as *hunks*. It is possible to configure `diff` to ignore whitespace or text case differences which is useful for scientific data sets where, for example, the header row may have different capitalisation but the data sets are otherwise identical. If no *hunks* are returned by `diff` the files are identical. However, if *hunks* are returned, the files differ in some way. In order to calculate the similarity, s , we return the percentage of lines that are unchanged, converted to a number in the range $0 \leq s \leq 1$

XML Comparing XML files is more complex than comparing textual data due to the semantic rules governing its structure. For example, the ordering of attributes is not important, generally the order of nodes is not important (although this is schema dependent) and where namespaces are used the prefix is not important. Existing work on comparing XML documents often tries to minimise the transformation necessary to turn one document into another (the so called, 'edit script') or maximise the common structure from one document to another [33, 4, 34].

^{††}gnu.org/software/diffutils/

For the purposes of comparing XML documents we use XOM (xom.nu/) which is able to canonicalise XML documents so that they can be compared. Again, similarity, s is the percentage of the document which has changed where $0 \leq s \leq 1$.

Mathematical Models Some modelling algorithms (for example, Neural Networks [35]) have an element of random initialisation and can be expected to give different results each time they are applied to the same dataset. Whilst these models may fail a byte-wise comparison, their predictive capabilities might be statistically the same. To calculate the similarity of models, we use the ANCOVA test (Analysis of Covariance - [36]) which analyses the predictive performance of two models. This is demonstrated below in Figure 8, which shows a plot of two models' estimates of a variable against the actual observed values for that variable. A good model in this case would have a regression line at 45 degrees with very little spread in the actual and observed points. Different models applied to the same data set could be expected to have slightly altered regression line slopes, and the differencing tool needs to be able to take account of this.

The ANCOVA test is divided into two parts: the first determines whether the slope of the regression lines is significantly different (A and B); whilst the second checks whether the y-axis intercepts of the two regression lines (C and D) are equivalent. It should be noted that whilst there are many more tests that should be done in order to compare two models (analysis of the model residuals, comparison of the model error measurements, etc.), for the purposes of demonstrating the expandability of the data diff if two models pass the ANCOVA tests, they will be considered equivalent. It is also important to note that the comparison is implemented as a workflow within e-Science Central and, as such, can be expanded to include additional tests if necessary.




Figure 8. Comparison of Model fit

4. INFRASTRUCTURE FOR REPRODUCIBILITY IN E-SCIENCE CENTRAL

The e-Science Central system [37] is a portable cloud 'platform-as-a-service' that can be deployed on either private clusters or public clouds, including Amazon EC2 and Windows Azure. Cloud

computing has the potential to give scientists the computational resources they need, when they need them. However, cloud computing does not make it easier to build the often complex, scalable secure applications needed to support science. e-Science Central was designed to overcome these obstacles by providing a platform on which users can carry out their research, and build high-level applications. Figure 9 shows the key components of the e-Science Central ‘Platform as a Service’ sitting on an ‘Infrastructure as a Service’ cloud. It combines three technologies – Software as a Service (so users only need a web browser to do their research), Social Networking (to support sharing and community interaction) and Cloud Computing (to provide storage and computational power). Using only a browser, users can upload data, share it in a controlled way with colleagues, and analyse the data using either a set of pre-defined services, or their own, which they can upload for execution and sharing. A range of data analysis and programming languages is supported, including Java, Javascript, R and Octave. From the point of view of users, this gives them the power of cloud computing without them actually having to manage the complexity of developing cloud-specific software – they can create services in a variety of languages, upload them into e-Science Central, and have them run transparently on the cloud. These services can be composed into workflows to automate the analysis.

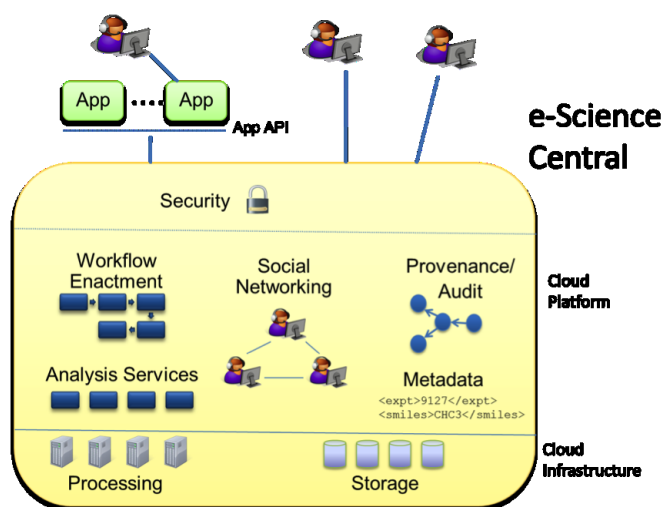


Figure 9. e-Science Central Components

Versioning is an integral storage feature in e-Science Central, allowing users to work with old versions of data, services and workflows. All objects (data sets, services and workflows) are stored in files through a virtual filestore driver than can be mapped onto a range of actual storage systems including standard local and distributed filesystems, and Amazon S3. When a file is stored, if a previous version exists then a new one is automatically created.

Whilst three of the elements, workflows, services and libraries (the latter two being part of *ED*), presented in Section 2 are within the control of e-Science Central and are versioned this alone is not sufficient to detect workflow decay. By default, workflow W will use the latest version of a service, S_{nV_i} when it is executed at time t . However at time, t' a service version may have been updated and so S_{nV_j} where $j > i$ would be used. It is possible to override this behaviour and always use a particular version of a service if desired. In addition, even if a workflow evolves $W_i \rightarrow W_j$ the evolution may only effect branches which do not appear in a particular execution trace.

4.1. Capturing provenance traces in e-Science Central

The provenance system within e-Science Central is used to capture the history and life cycle of every piece of data within the system. For instance: who created the data? who has downloaded the data? what version of which services (in what workflow) accessed the data? and who has the data been shared with?

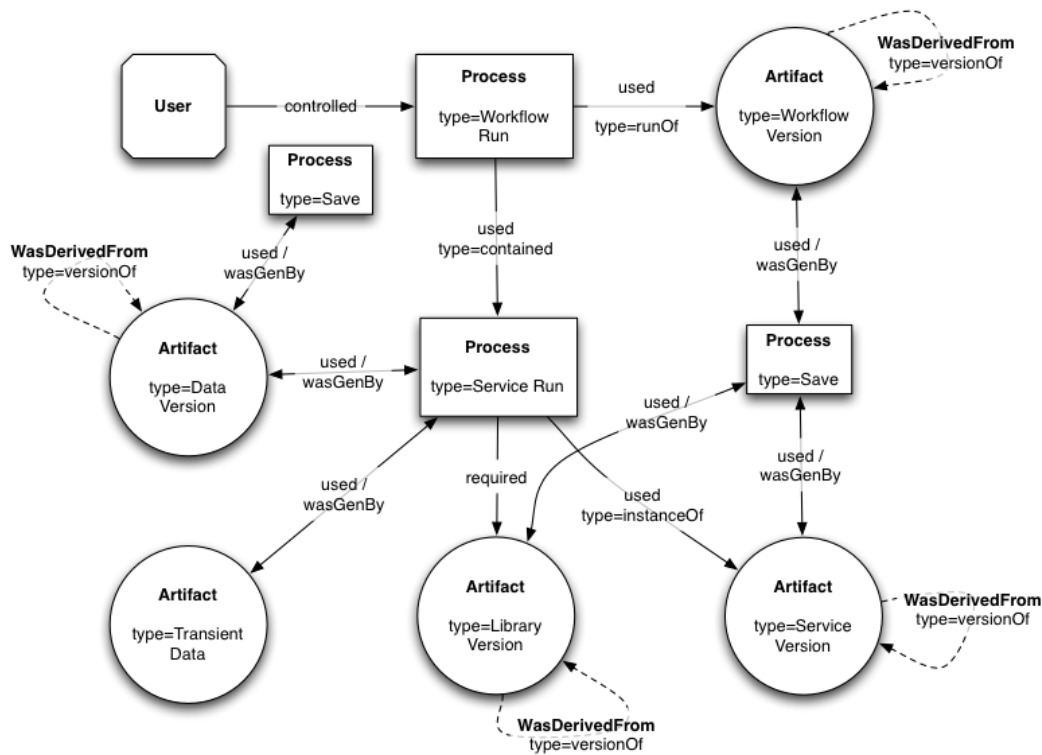


Figure 10. e-Science Central Provenance Model

The provenance data model, shown in Figure 10, is based on the Open Provenance Model (OPM) Version 1.1 [9], and can be used to produce a directed acyclic graph of the history of an object. Objects in OPM are categorised as either artifacts, processes or actors which correspond to nodes within the graph. Vertices in the graph represent relationships between two objects and are of types such as *wasGeneratedBy*, *used*, *wasControlledBy* and *wasDerivedFrom*. The OPM Core model has been extended with subclasses to identify the different types of processes and artifacts we are concerned with. For example, execution of a workflow and a service within a workflow have been differentiated. The relationship between workflow execution and service execution is of type *contained* and is not strictly required by our model but its inclusion makes it easier to generate the different views of the process. The artifacts described in the model are also subclassed in order to differentiate between versions of data, services, libraries and workflows. Without this subclassing, it would be necessary to either encode the object type in the identifier (not desirable as it adds an unnecessary layer of obfuscated information) or perform many lookups to answer the question ‘what type of object has a particular identifier’.

The relationships between processes and artifacts is specified in the OPM standard with the *Save* process taking one version of an artifact and generating a new version.

The self referential relationship between artifacts of type *WasDerivedFrom*, shown with a dashed line, indicates a second level, inferred relationship omitting the *Save* process. This implies that $D_{i,v}$ is in some way derived from $D_{i,v-1}$ when $v > 0$.

Our model deals with two different types of data artifact: *Data Version* and *Transient Data*. This is due to the semantics of workflows in e-Science Central: data generated by a workflow must be explicitly saved using a service which ‘exports’ the data back into the e-Science Central repository. Any data which is not explicitly exported will be discarded when the workflow completes.

4.2. Provenance database implementation

Given that the provenance structure being stored is a directed acyclic graph, we chose to store it in the non-relational graph database, Neo4j (www.neo4j.org). Neo4j differs from traditional relational databases as its structure is not in terms of tables and rows and columns, but in nodes, relationships and properties. This provides a much more natural fit to our model, and allows us to store the provenance graph directly instead of encoding it in a relational model. Neo4j has also been shown to scale well, and most importantly, to perform well in terms of queries even when storing a very large graph structure [38]. Libraries are provided to allow users to work with Neo4j directly in either Java or Ruby.

Various components within e-Science Central all log provenance information. From the outset the system was designed such that the order in which provenance events are received by the server is not important. For instance, we could have an arbitrary interleaving of the events which signify that a ‘Service has run’ and that the ‘Service accessed some data’. Each of these events contains a subset of information which must be added to the provenance database: the first event contains the service name, start and completion time (and some other information) whereas the latter defines the identity of the data which was read. Instead of SQL queries, Neo4j supports an operation known as a *traversal* whereby the user defines a starting node and rules about what types of relationship/node to ‘traverse’. The result is a set of paths from the starting node to the acceptable ending nodes. The provenance server is decoupled from e-Science Central with a durable JMS queue which allows e-Science Central to minimise the number of synchronous write operations which must be performed in a synchronous request.

4.3. Implementing PDIFF using Neo4j

The Neo4j database used by e-Science Central contains not just the execution traces which we are interested in traversing for PDIFF but also all the provenance for other items within e-Science Central. For example, there may be multiple executions of each workflow but we are only interested in two specific executions. Also, the provenance may indicate a data set may have been created through a *Save* process which does not concern us.

The Neo4j traversal framework allows us to traverse the graph but restrict the traversal to relationships of a particular *type* and *direction*. During most of our traversals we restrict ourselves to the *used* and *genBy* types with direction *outgoing*. This allows us to navigate back from the output data set in each execution trace back to the original input data sets. However, the semantics of Neo4j traversals are that you are guaranteed to visit each node in the graph but not traverse every edge. This would have the potential to miss branches in the execution trace which we should be comparing. To mitigate this we make additional calls to the database to ensure we retrieve all the relationships for each node.

As we traverse through the graph we build up a report containing the tree of ‘Delta’ descriptions making use of the functions described below.

UPD(*d*) Navigates from either a Transient Data or Data Version node along an outgoing *genBy* relationship to the Service Run node which created it and returns this node. If there is no such relationship null is returned indicating the end of the graph. In this case, *d* is a Transient Data node and an exception is thrown as this is an invalid starting point in the execution trace.

UPS(*a*, *p*) Navigates along an outgoing *used* relationship to a Transient Data or Data Version node and returns it. The Neo4j relationship which represents the edge contains an attribute for the port name *p* and so we filter based on this. If no such node is found an exception is thrown as this is an invalid starting point to the execution trace.

DDIFF(*NL*, *NR*) If *nL* and *nR* are Transient Data nodes then this returns whether or not the MD-5 hash values of the data are the same. If *nL* and *nR* are Data Version nodes then it will execute and e-Science Central workflow to determine whether or not the data sets are equivalent. The workflow executed is dependent on the MIME type of the data set as recorded as an attribute of the node in Neo4j.

ADIFF(NL, NR) This returns a 2-tuple containing whether or not the service identifier and version identifier for nL and nR are equal. In order to get the service and version identifier we must navigate the *instanceOf* relationship to the Service Version node in the provenance database. If either of the Service Version nodes is null then an exception is thrown as the graph is invalid.

FINDNODE(NL, NR) This performs a breadth first search of each of the execution traces in order to try to find two nodes that enable the analysis to be re-synchronised. For brevity we will assume in the following discussion that the Service Run nodes can be compared directly – in practice we must navigate along the *instanceOf* relationship and perform the comparison there. Initially we select nL' as a match candidate which is the next Service Run in the trace traversing along *used* and *genBy* with direction *outgoing*. We then perform a breadth first traversal of the same relationship types beginning from nR searching for a match with nL' . If a match is found it is recorded as *candidateL* along with the depth at which it was found. If no match was found, we traverse back one step further from nL' and repeat the search from nR .

If no candidate is found, the entire graph fragments from nL and nR are returned as there is no matching node that enables the graphs to be re-synchronised. However, if a match is found, the search is re-performed but searching for nR' nodes in the nL graph. When a second match is found, namely *candidateR*, the search terminates. The returned nodes (and associated graph fragments) are those with the lowest combined depths from nL and nR . The second search deals with the case where a closer node match is found had we searched the other graph first. We can terminate on finding a match as we know that this will either be the same match as the *candidateL* or a match with a lower combined depth. The nodes which are returned as a synchronisation point are also stored so that we are able to join future branches in the results.

ADDELTA(Δ , DATAELEMENT) Builds up the result graph. The results are constructed using a simple graph implementation where each node is able to contain one of the following objects: a Transient Data mismatch (the actual MD-5 hash values are of little use); the service and version identifiers in case of a service or version change; or two graphs representing the graph fragments which appear in one of the graphs but not the other. In the case of the graph fragments, the previous node nL and nR are included along with the subsequent nodes nL' and nR' for clarity. The result is stored in e-Science Central and can be exported to GraphML for visualisation or use in other applications.

At this time, the PDIFF implementation is not yet integrated into e-Science Central. In the next stage, the diagnosis functionality offered by the algorithm will be exposed to users through a simple interface. Users will be able to select two workflow runs and obtain a “diff” report on their divergence. Thanks to the relative small size of typical workflows, the diff graph can be visualized as part of the e-Science Central web-based graphical interface.

5. CONCLUSIONS

The ability to reproduce experimental results underpins scientific discourse and is key to establishing credibility in scholarly communication. Yet, several factors complicate the task of reproducing other scientists’ results in settings other than the original one. Such factors include differences in the experimental environment (different lab, research group, equipment technology) and insufficient detail in the description of the methods. In computational science, the operating environment typically includes e-science infrastructure facilities, and the methods are formally encoded in the form of executable scripts, workflows, or general computer programs. This level of automation creates the perception that the limitations in reproducibility of results can be easily overcome. In reality, however, the traditional limiting factors are simply replaced by others, including an unstable and evolving operational environment, and lack of maintenance of related programs with mutual dependencies, which become unable to function together. As a result, scientific applications suffer from forms of decay that limit their longevity, and thus their reuse and evolution in time.

In this paper we have focused specifically on workflow-based e-science, and on a scenario where attempts to reproduce earlier results translates into new runs of the same workflow at a later time. We assume that the workflow is still executable, but it may produce unexpected results or may have been made dysfunctional by changes in its system environment, or by uncontrolled evolution of the workflow itself or of its input data.

We have presented two main contributions. Firstly, we have proposed a framework meant to clarify the range of meanings of the often ambiguous term “reproducibility”. We have then used the framework to situate our second contribution, namely the PDIFF algorithm for diagnosing the divergence between two executions by comparing their associated provenance traces. We have also described a number of data comparison function, which underpin provenance comparison. Finally, we have discussed the current implementation of PDIFF on the e-Science Central workflow management middleware. PDIFF will be integrated in the public-facing version of e-Science Central in the near future, with attention to usability testing, which has not being addressed in the current prototype version.

REFERENCES

1. Roure DD, Belhajjame K, Missier P, Al E. Towards the preservation of scientific workflows. *Procs. of the 8th International Conference on Preservation of Digital Objects (iPRES 2011)*, Singapore, 2011; 228–231.
2. Cohen-Boulakia S, Leser U. Search, adapt, and reuse: the future of scientific workflows. *SIGMOD Rec. Sep 2011*; **40**(2):6–16, doi:<http://doi.acm.org/10.1145/2034863.2034865>. URL <http://doi.acm.org/10.1145/2034863.2034865>.
3. Groth P, Deelman E, Juve G, Mehta G, Berriman B. A Pipeline-Centric Provenance Model. *The 4th Workshop on Workflows in Support of Large-Scale Science*, 2009.
4. Wang Y, DeWitt D, Cai JY. X-diff: an effective change detection algorithm for XML documents. *Data Engineering, 2003. Proceedings. 19th International Conference on*, 2003; 519 – 530, doi:10.1109/ICDE.2003.1260818.
5. Bertiss AT. A Backtrack Procedure for Isomorphism of Directed Graphs. *J. ACM Jul 1973*; **20**(3):365–377, doi:10.1145/321765.321766. URL <http://doi.acm.org/10.1145/321765.321766>.
6. Ullmann JR. An Algorithm for Subgraph Isomorphism. *J. ACM Jan 1976*; **23**(1):31–42, doi:10.1145/321921.321925. URL <http://doi.acm.org/10.1145/321921.321925>.
7. Hidden H, Watson P, Woodman S, Leahy D. e-Science Central: Cloud-based e-Science and its application to chemical property modelling. *Technical report cs-tr-1227*, School of Computing Science, Newcastle University 2011.
8. Cala J, Watson P, Woodman S. Cloud Computing for Fast Prediction of Chemical Activity. *Procs. 2nd International Workshop on Cloud Computing and Scientific Applications (CCSA)*, Ottawa, Canada, 2012.
9. Moreau L, Clifford B, Freire J, Futrelle J, Gil Y, Groth P, Kwasnikowska N, Miles S, Missier P, Myers J, *et al.*. The Open Provenance Model — Core Specification (v1.1). *Future Generation Computer Systems* 2011; **7**(21):743–756, doi:<http://dx.doi.org/10.1016/j.future.2010.07.005>.
10. Moreau L, Missier P, Belhajjame K, B'Far R, Cheney J, Coppens S, Cresswell S, Gil Y, Groth P, Klyne G, *et al.*. PROV-DM: The PROV Data Model. *Technical Report*, World Wide Web Consortium 2012. URL <http://www.w3.org/TR/prov-dm/>.
11. Hanson B, Sugden A, Alberts B. Making Data Maximally Available. *Science* 2011; **331**(6018):649, doi:10.1126/science.1203354.
12. Merall Z. Computational science: ...Error. *Nature* Oct 2010; **467**:775–777, doi:10.1038/467775a.
13. Peng RD, Dominici F, Zeger SL. Reproducible Epidemiologic Research. *American Journal of Epidemiology* 2006; **163**(9):783–789, doi:10.1093/aje/kwj093. URL <http://aje.oxfordjournals.org/content/163/9/783.abstract>.
14. Drummond C. Science, Replicability is not Reproducibility: Nor is it Good Science. *Procs. 4th workshop on Evaluation Methods for Machine Learning In conjunction with ICML 2009*, Montreal, Canada, 2009. URL <http://cogprints.org/7691/7/ICMLws09.pdf>.
15. Peng R. Reproducible Research in Computational Science. *Science* Dec 2011; **334**(6060):1226–1127.
16. Bechhofer S, De Roure D, Gamble M, Goble C, Buchan I. Research Objects: Towards Exchange and Reuse of Digital Knowledge. *Nature Precedings* 2010; URL <http://dx.doi.org/10.1038/npre.2010.4626.1>.
17. Schwab M, Karrenbach M, Claerbout J. Making Scientific Computations Reproducible. *Computing in Science Engineering* 2000; **2**(6):61–67, doi:10.1109/5992.881708. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=881708>.
18. Mesirov J. Accessible Reproducible Research. *Science* 2010; **327**. URL www.sciencemag.org.
19. Nekrutenko A. Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biology* 2010; **11**(8):R86. URL <http://dx.doi.org/10.1186/gb-2010-11-8-r86>.
20. Scheidegger C, Vo H, Koop D, Freire J. Querying and Re-Using Workflows with VisTrails. *Procs. SIGMOD*, 2008; 1251–1254. URL <http://portal.acm.org/citation.cfm?id=1376747>.
21. Ludäscher B, Altintas I, Berkley C. Scientific Workflow Management and the Kepler System. *Concurrency and Computation: Practice and Experience* 2005; **18**:1039–1065. URL <http://onlinelibrary.wiley.com/doi/10.1002/cpe.994/abstract;jsessionid=B39961C2FDAF407AB30BCB118612D171.d02t02>.

22. Missier P, Soiland-Reyes S, Owen S, Tan W, Nenadic A, Dunlop I, Williams A, Oinn T, Goble C. Taverna, reloaded. *Procs. SSDBM 2010*, Gertz M, Hey T, Ludaescher B (eds.), Heidelberg, Germany, 2010. URL <http://www.ssdbm2010.org/>.
23. Zhao J, Gomez-Perez J, Belhajjame K, Klyne G, Al E. Why Workflows Break - Understanding and Combating Decay in Taverna Workflows. *Procs. e-science conference*, Chicago, 2012.
24. Missier P, Paton N, Belhajjame K. Fine-grained and efficient lineage querying of collection-based workflow provenance. *Procs. EDBT*, Lausanne, Switzerland, 2010.
25. Moreau L. Provenance-based reproducibility in the Semantic Web. *Web Semantics: Science, Services and Agents on the World Wide Web* 2011; **9**(2):202–221, doi:10.1016/j.websem.2011.03.001. URL <http://www.sciencedirect.com/science/article/pii/S1570826811000163>.
26. Kim J, Deelman E, Gil Y, Mehta G, Ratnakar V. Provenance trails in the Wings-Pegasus system. *Concurrency and Computation: Practice and Experience* 2008; **20**:587–597, doi:<http://dx.doi.org/10.1002/cpe.1228>. URL <http://www3.interscience.wiley.com/journal/115805532/abstract>.
27. Koop D, Scheidegger C, Freire J, Silva C. The Provenance of Workflow Upgrades. *Provenance and Annotation of Data and Processes, Lecture Notes in Computer Science*, vol. 6378, McGuinness D, Michaelis J, Moreau L (eds.). Springer Berlin / Heidelberg, 2010; 2–16. URL http://dx.doi.org/10.1007/978-3-642-17819-1_2.
28. Bunke H. Graph Matching: Theoretical Foundations, Algorithms, and Applications. *Procs. Vision Interface*, 2000.
29. Chawathe SS, Garcia-Molina H. Meaningful change detection in structured data. *SIGMOD Rec.* Jun 1997; **26**(2):26–37, doi:10.1145/253262.253266. URL <http://doi.acm.org/10.1145/253262.253266>.
30. Altintas I, Barney O, Jaeger-Frank E. Provenance Collection Support in the {K}epler Scientific Workflow System. *IPAW*, 2006; 118–132, doi:http://dx.doi.org/10.1007/11890850_14.
31. Bao Z, Cohen-Boulakia S, Davidson S, Eyal A, Khanna S. Differencing Provenance in Scientific Workflows. *Procs. ICDE*, 2009, doi:<http://dx.doi.org/10.1109/ICDE.2009.103>. URL http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?tp=&arnumber=4812456&isnumber=4812372.
32. Deelman E, Gannon D, Shields M, Taylor I. Workflows and e-Science: An overview of workflow system features and capabilities. *Future Generation Computer Systems* 2009; **25**(5):528–540, doi:DOI:10.1016/j.future.2008.06.012. URL <http://www.sciencedirect.com/science/article/B6V06-4SYCPKX-2/2/bb631979e3dd7071ddede90bbff65a91>.
33. Schubert E, Schaffert S, Bry F. Structure-preserving difference search for XML documents. *Extreme Markup Languages®*, 2005.
34. Cobena G, Abiteboul S, Marian A. Detecting changes in XML documents. *Data Engineering, 2002. Proceedings. 18th International Conference on*, 2002; 41–52, doi:10.1109/ICDE.2002.994696.
35. Cybenko G. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems (MCSS)* 1989; **2**:303–314. URL <http://dx.doi.org/10.1007/BF02551274>, 10.1007/BF02551274.
36. Rutherford A. *Introducing ANOVA and ANCOVA : a GLM approach*. SAGE, 2001. URL <http://www.worldcat.org/isbn/9780761951605>.
37. Watson P, Hiden H, Woodman S. e-Science Central for CARMEN: science as a service. *Concurrency and Computation: Practice and Experience* December 2010; **22**:2369–2380, doi:<http://dx.doi.org/10.1002/cpe.v22:17>. URL <http://dx.doi.org/10.1002/cpe.v22:17>.
38. Dominguez-Sal D, Urbón-Bayes P, Giménez-Vañó A, Gómez-Villamor S, Martínez-Bazán N, Larriba-Pey JL. Survey of graph database performance on the HPC scalable graph analysis benchmark. *Proceedings of the 2010 international conference on Web-age information management, WAIM'10*, Springer-Verlag: Berlin, Heidelberg, 2010; 37–48. URL <http://dl.acm.org/citation.cfm?id=1927585.1927590>.