

A Model for User-Oriented Data Provenance in Pipelined Scientific Workflows

Shawn Bowers¹, Timothy McPhillips¹, Bertram Ludäscher^{1,2},
Shirley Cohen³, Susan B. Davidson³

¹UC Davis Genome Center, University of California, Davis

²Department of Computer Science, University of California, Davis

³Computer and Information Science, University of Pennsylvania

{sbowers, ludaesch, tmcphillips}@ucdavis.edu

{shirleyc, susan}@cis.upenn.edu

Abstract. Integrated provenance support promises to be a chief advantage of scientific workflow systems over script-based alternatives. While it is often recognized that information gathered during scientific workflow execution can be used automatically to increase fault tolerance (via checkpointing) and to optimize performance (by reusing intermediate data products in future runs), it is perhaps more significant that provenance information also may be used by scientists to reproduce results from earlier runs, to explain unexpected results, and to prepare results for publication. Current workflow systems offer little or no direct support for these “scientist-oriented” queries of provenance information. Indeed the use of advanced execution models in scientific workflows (*e.g.*, process networks, which exhibit pipeline parallelism over streaming data) and failure to record certain fundamental events such as *state resets* of processes, can render existing provenance schemas useless for scientific applications of provenance. We develop a simple provenance model that is capable of supporting a wide range of scientific use cases even for complex models of computation such as process networks. Our approach reduces these use cases to database queries over event logs, and is capable of reconstructing complete data and invocation dependency graphs for a workflow run.

1 Introduction

The importance of provenance information in scientific data and workflow management is widely recognized, as witnessed, *e.g.*, by specialized workshops [4,1], research projects [17], and surveys [3,20] dedicated to this topic, and by investigations on foundations of data provenance for queries and transformations [5,8,23]. However, current scientific workflow systems still offer little or no support for queries of interest to the end-users of these systems, *i.e.*, researchers in the natural and physical sciences. In this paper, we argue that concrete use cases, expressed in terms that are meaningful to the scientist, should drive the design of a provenance system. Moreover, such systems should be designed in terms of the models of computation (MoC) that govern the execution of scientific workflows to ensure that all pertinent events are recorded in the execution log.

Fig. 1 shows an example workflow for inferring phylogenetic trees approximating the evolutionary relationships between organisms. DNA sequences for homologous genes from a number of taxa are provided as input to the workflow. Actor A1 performs an initial alignment of the sequences (*e.g.*, using the program ClustalW [22]), and Actor A2 refines this initial alignment (*e.g.*, using Gblocks [6]).¹ Actor A3 infers a set of phylogenetic trees from the aligned se-

¹ Following KEPLER terminology, we call workflow components *actors*.

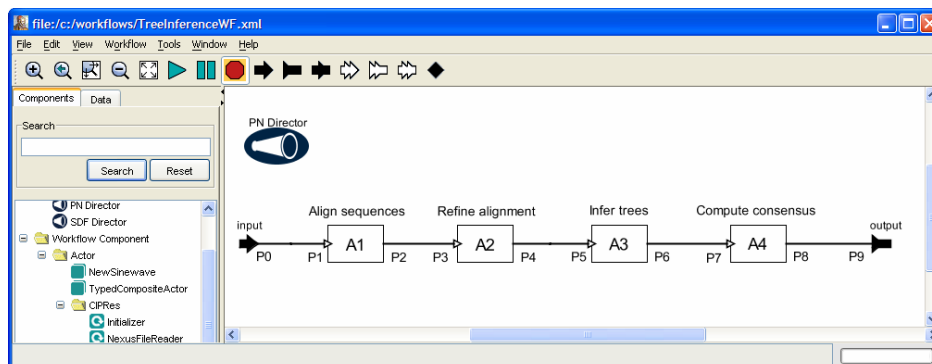


Fig. 1. A workflow for computing phylogenetic trees from input DNA sequences.

quences (e.g., using DNAPARS [19]), and actor A4 computes the consensus of these trees (e.g., using CONSENSE [19]). For such scientific workflows we would like to (a) enable the scientist to ask scientific questions about a run by providing convenient queries against the execution log; and (b) have the system track the true data dependencies within a run so that answers to such scientific questions may be accurate as possible. For example, the system should recognize *independent* “sub-runs” as such. The workflow in Fig. 1 may process multiple sets s_1, s_2, \dots of independent DNA sequences (e.g., corresponding to distinct genes in the taxa of interest) within a single workflow run R . In such cases, the system should *not* infer that the data products resulting from the different s_i are interdependent. Rather, the system should answer accurately questions such as: *Which phylogenetic trees were used to produce this consensus tree? Which DNA sequences does this consensus tree depend on? Which input DNA sequences were **not** used to derive any output consensus tree?*

In this paper, we develop a provenance model designed to support such user-oriented queries even for pipelined models of computation where tracking of actual data dependencies can be complex. Consider, e.g., an actor A in an environmental monitoring workflow which computes a running average (of temperature, say) for each received measurement data token. Thus, upon each *firing* of A , the actor consumes a temperature token and emits a new running average token. To calculate the running average over multiple firings, A must maintain *state*. For the provenance system this means that every produced data token must be recorded as dependent not just on those input tokens received since the last time the actor fired, but on all tokens received since A was initialized. Conversely, if A is to limit the running averages to readings taken on a particular day, then A ’s state is *reset* once per day. There are no dependencies between tokens produced after a reset and tokens consumed prior to the reset. This observation naturally partitions token streams, as well as actor firings, into semantically meaningful firing *rounds*. Clearly, a provenance system should be able to observe and record new rounds of firing to avoid reporting false dependencies. Below we show how one can capture this essential information in a simple tabular event log. Our approach is easy to implement, e.g., in the KEPLER system, because some firing events can be automatically captured by a workflow scheduler (a *director* in KEPLER). Announcing a new round of firing (e.g., by signaling a reset event), on the other hand, can be seen as an actor’s responsibility, since an actor “knows” when it is beginning an independent task (such as a “sub-run” s_i above).

2 Preliminaries

Workflow Graphs, Actors, and Tokens vs. Data Objects. We adopt notions and terminology from KEPLER, a scientific workflow system extending PTOLEMY II. Workflows are composed by placing *actors* on a design canvas, and “wiring” them together to form the desired workflow graph (Fig. 1). Actors communicate through their input and output *ports*. In a *workflow graph* W , output ports can be connected to input ports, establishing a unidirectional dataflow *channel*. Actors communicate through these channels by passing *tokens*. By default tokens are immutable and “disposable”, *i.e.*, every token t is written only once [14] and thus lives only between its creation on an output port, and its consumption at subsequent input ports. Thus, even if an actor passes on a data object unchanged, a new token-id is created, facilitating tracking of token dependencies. A separate object-id is used to track object dependencies. By $object(t)$ we denote the data *object* represented by the token t . To support user-oriented queries, we associate with an object o one or more types $types(o)$. The *ports* of an actor A are denoted $ports(A)$. We assume that port-ids are globally unique, *i.e.*, include a unique actor-occurrence-id and a port-name which is unique to the actor occurrence. Ports are either *input* or *output*, so $ports(A) = in(A) \dot{\cup} out(A)$. Some ports $pars(A) \subseteq in(A)$ may be distinguished as *parameter* inputs for configuring A ’s behavior. The signature $\Sigma_W := in(W) \rightarrow out(W)$ of a workflow W is given by a set of distinguished inputs $in(W)$ and outputs $out(W)$. As shown in Figure 1, the distinguished workflow input and output ports are connected to a subset of the input and output ports of the workflow’s actors.

Directors. The model of computation (MoC) of a workflow is *not* defined by actors, but specified by a separate component called a *director*. Thus, KEPLER allows workflow designers to choose among different MoCs by choosing appropriate directors. A director specifies and (effectively) mediates all inter-actor communication, separating workflow scheduling and runtime orchestration (a director’s concern) from individual actor execution (an actor’s concern). This separation achieves a form of *behavioral polymorphism* [13], resulting in more reusable actor components and subworkflows. KEPLER provides a variety of directors that implement process network (PN and SDF), discrete event (DE), continuous time (CT), and finite state transducer (FST) semantics.

Pipelined Execution. In the process network MoC, the PN director executes each actor as a separate process (or thread). In this way, workflows that run using the PN director not only exhibit task parallelism, but also *pipeline parallelism*. For example, during a single workflow run, each actor in Figure 1 can execute multiple times, and different actors can execute simultaneously. A number of other MoCs models can be considered special cases of the basic Kahn-McQueen process network model [11]. In the synchronous dataflow (SDF) model [12], actors a priori define fixed token consumption and production rates. This allows the SDF director to statically schedule actors, while guaranteeing, *e.g.*, that (unlike in the general PN case) deadlocks cannot occur and that buffers have a fixed size. By DAG (directed acyclic graph) we denote a MoC that is common in job-centric grid workflows [21,9]: nodes represent jobs, and directed edges represent execution dependencies between jobs. Thus, a DAG director can simply execute the jobs in the partial order implied by the job dependency graph. Note that in this model, unlike in PN and SDF, each actor is invoked exactly once.

3 A Provenance Model for Pipelined Workflows

This section describes a provenance model that can handle the process network (PN) model of computation and thus specialized versions such as SDF and DAG as well. To execute a *workflow*

(graph) \mathcal{W} , we must “bind” (*i.e.*, select) *input data* \mathbf{i} on which \mathcal{W} will operate. Often \mathcal{W} is also parameterized using initial parameter settings \mathbf{p} . It is customary to record identifiers for \mathcal{W} , \mathbf{p} , and \mathbf{i} as part of the provenance information. Finally, a MoC M is needed (*e.g.*, PN, SDF, DAG) to determine how the workflow is executed.² Taken together, the equation $\mathbf{o} = M(\mathcal{W}_{\mathbf{p}}(\mathbf{i}))$ denotes a workflow execution in which the output \mathbf{o} is obtained by applying a suitable model of computation M to an appropriately instantiated workflow \mathcal{W} .

Run, Trace, Observables. Each MoC M formally defines the notion of legal *computations* or *runs*, such that one can determine whether a particular run R of a workflow \mathcal{W} is a legal representation (*wrt.* M) of an execution $\mathbf{o} = M(\mathcal{W}_{\mathbf{p}}(\mathbf{i}))$. A workflow *trace* \mathcal{T} is an approximation of a run R , according to a *model of provenance*. As recorded by a provenance model, a trace approximates a run by recording functional and non-functional *observables*. For example, an SDF director precomputes a static workflow schedule (based on actor consumption and production rates), and using this schedule signals each actor to fire in turn. Thus, actor firings are directly observed in SDF. In contrast, the *size* of a token (or rather the object it represents) and the *timestamp* when the token was created are non-functional observables: according to the MoC, the outcome does not depend on these. Non-functional observables can be useful to record, *e.g.*, to benchmark actor execution times or data transfer times between actors, but are not essential for determining data dependencies.³

The RWS Provenance Model. Here we consider a concrete model of provenance, called the RWS model, which records *read*, *write*, and (state) *reset* events for each actor in a workflow run. These events are stored in a relational *event log*, as shown in Fig. 2.⁴ During a workflow run, a read event is added to the event log each time an actor reads a token from a port. Similarly, a write event is added to the log each time an actor writes a token to a port. A series of reads followed by writes denotes an actor firing. Note that in a particular firing F_j , an actor may use data that it read in a previous firing F_i to generate output (*e.g.*, this is the typical behavior of a running-average actor, as described in Section 1). In this case, we say the actor maintains state across firings, and state reset events denote when the state is “flushed” (reset). The firings between reset events constitute a firing round.

As shown in Fig. 2, each row in the event log contains: (1) the “location” of the event E_{loc} , which is either a port (for read and write events) or an actor (for state reset events); (2) the event type E_{typ} , which is either ‘r’ for read events, ‘w’ for write events, or ‘s’ for state reset events; (3) the identifier of the token E_{tok} that was read or written at the port (null for state reset events); and (4) a firing count E_{fire} .

Because actor port identifiers are unique across a workflow, and tokens are written once, the port and token identifiers recorded for each read and write event enable the reconstruction of the flow of data through the workflow run. However, these events alone are not sufficient to reconstruct data dependencies. We use the state reset events (as described above) along with the firing count for this purpose. In particular, the firing count is incremented independently for each actor whenever (i) an actor switches from writing tokens to reading tokens, denoting a new firing of the actor, and (ii) whenever a state reset event occurs.

² Some MoCs might also be aware of resources such as cluster (or grid) nodes and transport protocols, and schedule a distributed workflow accordingly.

³ In existing systems, such timestamps are often the only information available and thus are also used to second-guess other properties such as token and object dependencies.

⁴ We focus only on a minimal set of observables that allow us to answer many science-oriented user questions (see next section), while ignoring non-functional observables such as timestamps, although such information can be easily added.

E_{loc}	E_{typ}	E_{tok}	E_{fire}	E_{loc}	E_{typ}	E_{tok}	E_{fire}	E_{loc}	E_{typ}	E_{tok}	E_{fire}	E_{loc}	E_{typ}	E_{tok}	E_{fire}
p_0	w	t_1	1	A_1	s	-	3	A_2	s	-	4	p_7	r	t_{25}	1
...				p_1	r	t_{17}	3	A_3	s	-	1	p_7	r	t_{26}	1
p_0	w	t_{18}	1	p_1	r	t_{18}	3	p_5	r	t_{22}	1	p_8	w	t_{29}	1
A_1	s	-	1	p_2	w	t_{21}	3	p_6	w	t_{24}	1	A_4	s	-	2
p_1	r	t_1	1	A_1	s	-	4	p_6	w	t_{25}	1	p_7	r	t_{27}	2
...				A_2	s	-	1	p_6	w	t_{26}	1	p_7	r	t_{28}	2
p_1	r	t_7	1	p_3	r	t_{19}	1	A_3	s	-	2	p_8	w	t_{30}	2
p_2	w	t_{19}	1	p_4	w	t_{22}	1	p_5	r	t_{23}	2	A_4	s	-	3
A_1	s	-	2	A_2	s	-	2	p_6	w	t_{27}	2	p_9	r	t_{29}	1
p_1	r	t_8	2	p_3	r	t_{20}	2	p_6	w	t_{28}	2	p_9	r	t_{30}	1
...				p_4	w	t_{23}	2	A_3	s	-	3				
p_1	r	t_{16}	2	A_2	s	-	3	A_4	s	-	1				
p_2	w	t_{20}	2	p_3	r	t_{21}	3	p_7	r	t_{24}	1				

Fig. 2. The event log for a run of the example workflow in Fig. 1.

Dependency Graphs. Using the RWS model, we are able to infer from the event log the *token dependency graph*. That is, for each token t , we can know which parent tokens $\{t_1, \dots, t_k\}$ directly contributed to the production of t (as the result of an actor firing). As an example, in the upper left of Fig. 3, $\{t_1, \dots, t_7\}$ are parent tokens of t_{19} . Conversely, t_{22} is the parent of t_{24}, t_{25}, t_{26} . The following Datalog program illustrates how the token dependency graph can be computed from the event log. The *event* relation corresponds to the event log and the *actor* relation contains a mapping from ports to their corresponding actors.

$$\begin{aligned}
\text{depends-on}(T_1, T_2) &:- \text{event}(P_1, w, T_1, C_1), \text{event}(P_2, r, T_2, C_2), \text{actor}(P_1, A), \text{actor}(P_2, A), \\
&\quad \text{reset}(A, C_b, C_e), C_b \leq C_2 \leq C_1 < C_e. \\
\text{reset}(A, C_b, C_e) &:- \text{event}(A, s, X, C_b), \text{event}(A, s, Y, C_e), C_b < C_e, \\
&\quad \neg \text{reset-between}(A, C_b, C_e). \\
\text{reset-between}(A, C_b, C_e) &:- \text{event}(A, s, X, C), C_b < C < C_e.
\end{aligned}$$

We say that T_1 depends on T_2 whenever $\text{depends-on}(T_1, T_2)$ is true.

In addition to the token dependency graph, we are also able to infer the *object dependency graph* using the RWS model. Object dependencies describe user data lineage, and are crucial for our “user-oriented” queries. For example, the middle column of Fig. 3 shows the object dependencies for the workflow run of Fig. 1. Note that the object dependency graph differs slightly from the token dependency graph. Object dependency graphs can be computed from corresponding token dependency graphs and token-object mappings.

Finally, we can also infer the *actor-invocation dependency graph* from the RWS model. In particular, this graph can be built from state-reset events in the event log such that an actor invocation A_2^j depends on another actor invocation A_1^i whenever A_2^j reads a token that is written by A_1^i . Note that here, “invocation” refers to a firing round. The token, object, and actor-invocation dependencies for a workflow trace are directed acyclic graphs (DAGs) that in general may be unconnected (as in Fig. 3). It should be clear that all of the information stored in the event table is required to reconstruct these token, object, and invocation dependency graphs for a workflow trace. In particular, if state-reset events are not logged, each token written by an actor will appear to depend on all previous tokens read during prior firing rounds of the actor: *e.g.*, in the absence of state-reset events, t_{21} would be connected to tokens t_1 to t_{18} in the token dependency graph of Fig. 3.

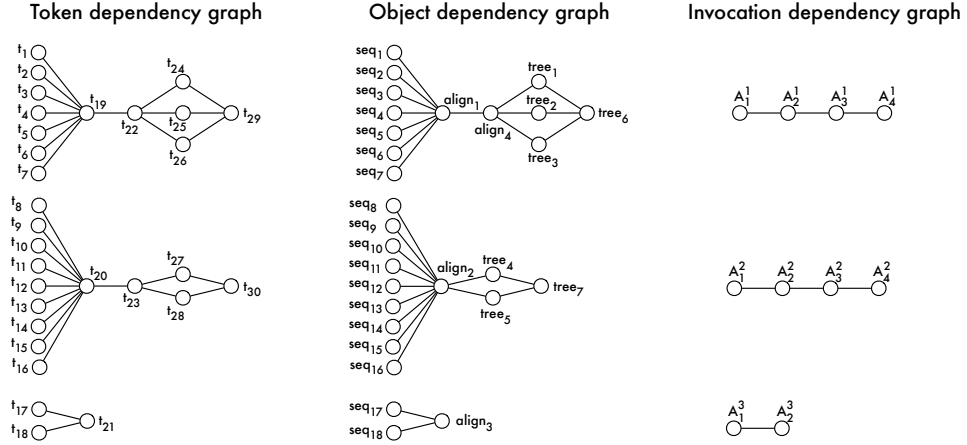


Fig. 3. Token, object, and actor invocation graphs for example phylogenetics workflow. Dependencies are shown from left to right. Note that all but one of the token-object mappings can be inferred from the graph structures; tokens t_{20} and t_{23} both map to the object align_2 .

4 Querying Workflow Traces

A wide range of scientifically relevant questions can be answered using the provenance model described above. To make access to event logs more convenient, we introduce the following primitive operations, which can be implemented, *e.g.*, as relational selections over the event log. The $\text{writer}(t)$ and $\text{reader}(t)$ operations return the ports that a token t was written to and read from, respectively (a token is written to a port exactly once, but can be read multiple times). The $\text{token-parents}(t)$ and $\text{token-children}(t)$ operations return the set of direct token dependencies for a token t , while $\text{token-ancestors}(t)$ and $\text{token-descendents}(t)$ are their transitive closures. The $\text{siblings}(t)$ operation returns the tokens with the same direct dependencies as t ; *e.g.*, because actor A3 can infer multiple trees from an alignment, given one of these trees, siblings returns the other trees computed from the same alignment. The $\text{origin}(o)$ and $\text{death}(o)$ operations return the first and last tokens in the trace that refer to the object o ; *e.g.*, the origin and death operations can be used to determine that the alignment object align_2 originated with token t_{20} (written by actor A1) and terminated with token t_{23} (written by actor A2).

The following examples illustrate how the provenance operations can be combined to answer concrete questions of interest to a scientist using the workflow in Figure 1. For each high-level question below, we define a corresponding parameterized query using set-comprehension syntax⁵, along with the actual results for the event log given in Figure 2. Below, we use \mathcal{W} to denote the workflow graph (in Figure 1) and \mathcal{T} for the corresponding trace.

- **What DNA sequences were input to the workflow?** This is one of the first questions a scientist might ask about the workflow run. Given an object type $\$c$, the parameterized query

$$q_1(\$c) := \{o \mid t \in \text{tokens}(\mathcal{T}) \wedge \text{writer}(t) \in \text{in}(\mathcal{W}) \wedge \text{object}(t) = o \wedge \$c \in \text{types}(o)\},$$

returns the set of objects of type $\$c$ that were input to the workflow run. For our example trace, $q_1(\text{SEQUENCE})$ returns the objects seq_1 to seq_{18} . The expression $t \in \text{tokens}(\mathcal{T})$ selects a token

⁵ Queries could also be defined in Datalog or in query languages for graphs or semistructured data.

from the trace, the expression $writer(t) \in in(W)$ checks that the token was written by an input port of the workflow W , the expression $object(t) = o$ obtains the object associated with t , and the expression $\$c \in types(o)$ verifies that o has $\$c$ as a type.

- **What phylogenetic trees were output by the workflow?** This is another basic question that a scientist might initially ask after a run. Given the query

$$q_2(\$c) := \{o \mid t \in tokens(\mathcal{T}) \wedge reader(t) \in out(W) \wedge object(t) = o \wedge \$c \in types(o)\},$$

the expression $q_2(TREE)$ returns the objects $tree_6$ and $tree_7$.

- **What phylogenetic trees (intermediate or final) were created by the workflow?** This question requests both intermediate as well as final data products of a run. Given the query

$$q_3(\$c) := \{o \mid t \in tokens(\mathcal{T}) \wedge writer(t) \notin in(W) \wedge object(t) = o \wedge \$c \in types(o)\},$$

the expression $q_3(TREE)$ returns all $tree$ objects of Figure 3. Note that the expression $writer(t) \notin in(W)$ ensures that the returned trees were not given as input to the workflow.

- **What actor created this phylogenetic tree?** The following query returns the actors that first wrote the given object $\$o$:

$$q_4(\$o) := \{a \mid t \in origin(\$o) \wedge actor(writer(t)) = a\}.$$

The query returns A_3 for $tree_1$ to $tree_5$, and A_4 for $tree_6$ and $tree_7$. This question is of particular interest for workflows that employ multiple approaches for inferring phylogenetic trees.

- **Which phylogenetic trees were directly used to compute this consensus tree?** This question (*i.e.*, what is this tree the "consensus" of?) asks for the intermediate data products supplied to the actor producing a particular workflow output. Given the query

$$q_5(\$c, \$o) := \{o' \mid t \in origin(\$o) \wedge t' \in token-parents(t) \wedge object(t) = o' \wedge \$c \in types(o')\},$$

the expression $q_5(TREE, tree_6)$ returns $tree_1$ to $tree_3$; and $q_5(TREE, tree_7)$ returns $tree_4$ to $tree_5$.

- **What sequences input to the workflow does this consensus tree depend on?** This question illustrates how a workflow output can be related to the particular workflow inputs from which it was derived. Given the query

$$q_6(\$c, \$o) := \{o' \mid t \in origin(\$o) \wedge t' \in token-ancestors(t) \wedge writer(t') \in in(W) \wedge object(t') = o' \wedge \$c \in types(o')\},$$

the expression $q_6(SEQUENCE, tree_6)$ returns seq_1 to seq_7 , and the expression $q_6(SEQUENCE, tree_7)$ returns seq_8 to seq_{16} .

- **Which input sequences were not used to derive any output consensus trees?** Here we are interested in whether there are any workflow inputs without corresponding workflow outputs. Such inputs may be considered the workflow equivalent of "phantom lineages" [23]. Given an input type $\$c_{in}$ and output type $\$c_{out}$, the query

$$q_7(\$c_{in}, \$c_{out}) := \{o \mid t \in tokens(\mathcal{T}) \wedge writer(t) \in in(W) \wedge object(t) = o \wedge \$c_{in} \in types(o) \wedge \{t' \mid t' \in token-descendants(t) \wedge reader(t') \in out(W) \wedge c_{out} \in types(object(t'))\} = \emptyset\},$$

returns the objects input to the workflow that do not produce any workflow outputs; *e.g.*, the expression $q_7(SEQUENCE, TREE)$ returns the sequences seq_{17} and seq_{18} . The query first finds workflow input tokens t that refer to objects of type $\$c_{in}$, and then checks (via a subquery) to make sure that t has no output tokens with objects of the type $\$c_{out}$.

- **What was the sequence alignment used in the process of inferring this tree?** This question requests the key intermediate data object used in producing a workflow result. A researcher may wish to examine the alignment to assess the reliability of the results, or reuse the alignment in another workflow. Given the query

$$q_8(\$c, \$o) := \{o' \mid t \in \text{origin}(\$o) \wedge t' \in \text{token-ancestors}(t) \wedge \text{object}(t') = o' \wedge \$c \in \text{types}(o') \wedge \{t'' \mid t'' \in \text{token-descendents}(t') \wedge \$c \in \text{types}(\text{object}(t''))\} = \emptyset\},$$

the expression $q_8(\text{ALIGNMENT}_4, \text{tree}_6)$ returns the sequence alignment align_4 , and the expression $q_8(\text{ALIGNMENT}_4, \text{tree}_7)$ returns the sequence alignment align_2 . The subquery above ensures that the object o' is the alignment directly used to infer the tree.

- **What actors were involved in creating this tree?** This question may be used, *e.g.*, when writing the methods section of a publication to cite the employed methods and implementations. Given the query

$$q_9(\$o) := \{a \mid t \in \text{origin}(\$o) \wedge \text{actor}(\text{writer}(t)) = a\} \cup \{a \mid t \in \text{origin}(\$o) \wedge t' \in \text{token-ancestors}(t) \wedge \text{actor}(\text{writer}(t')) = a\},$$

the expression $q_9(\text{tree}_6)$ returns actors A1 to A4.

- **Which actors did not produce any output for input derived from this input sequence?** This question provides an explanation for the phantom lineages revealed by q_7 above:

$$q_{10}(\$o) := \{a \mid t \in \text{origin}(o) \wedge t' \in \text{token-descendents}(t) \wedge \text{token-children}(t') = \emptyset \wedge \text{actor}(\text{reader}(t')) = a\},$$

The expressions $q_{10}(\text{seq}_{17})$ and $q_{10}(\text{seq}_{17})$ both return actor A2, indicating that this actor did not forward a refined sequence alignment of these two sequences to actor A3. This result is reasonable since no informative phylogenetic trees may be inferred from only two taxa.

5 Conclusion

Tracking provenance is an important aspect of scientific workflow systems. In this paper, we have focussed primarily on the problem of tracking data lineage within scientific workflow runs. The problem of data lineage has been widely studied in the database community [5,7,2,23], however, the primary focus there has been on transformations of data items expressed as database queries. As noted in [10], current provenance approaches for workflow systems (*e.g.*, [24,25,18]) record various kinds of metadata related to provenance. Despite these developments, however, little support exists in current systems to allow end-users to query provenance information in scientifically meaningful ways, in particular when advanced workflow execution models go beyond simple DAGs (*e.g.*, as in process networks). We have shown that a simple read/write/reset event-based provenance model is expressive enough to capture many relevant science-oriented provenance use cases. These use cases become queries against suitable views on top of the event log. This approach also marks the beginnings of a use-case and computation-model driven approach to provenance schema design. Using our framework, it is now meaningful to ask whether a provenance schema can answer specific use cases, since the latter become queries over the former. As future work we intend to extend our approach to support a wider array of operations, including “smart re-runs” (a workflow system requirement in [15]) and recovery, and to extend our current Prolog-based prototype to provide direct support (including query user interfaces) for our provenance model within KEPLER.

References

1. D. Berry, P. Buneman, M. Wilde, and Y. Ioannidis, editors. *e-Science Workshop on Data Provenance and Annotation*, National e-Science Centre, Edinburgh, December 2003.
2. D. Bhagwat, L. Chiticariu, W. C. Tan, and G. Vijayvargiya. An annotation management system for relational databases. In *Proc. of VLDB*, 2004.
3. R. Bose and J. Frew. Lineage retrieval for scientific data processing: A survey. *ACM Computing Surveys*, 37(1):1–28, 2005.
4. P. Buneman and I. Foster, editors. *Workshop on Data Derivation and Provenance*, Chicago, October 2002.
5. P. Buneman, S. Khanna, and W. C. Tan. Why and where: A characterization of data provenance. In *Proc. of ICDT*, volume 1973 of *LNCS*, 2001.
6. J. Castresana. Selection of conserved blocks from multiple alignments for their use in phylogenetic analysis. *Mol. Biol. Evol.*, 17:540–552, 2000.
7. Y. Cui and J. Widom. Lineage tracing for general data warehouse transformations. In *VLDB*, 2001.
8. Y. Cui, J. Widom, and J. Wiender. Tracing the lineage of view data in a warehousing environment. *ACM TODS*, 25(2), 2000.
9. E. Deelman, J. Blythe, Y. Gil, C. Kesselman, G. Mehta, S. Patil, M.-H. Su, K. Vahi, and M. Livny. Pegasus: Mapping scientific workflows onto the grid. In *Proc. of the European Across Grids Conference*, 2004.
10. C. Goble. Position statement: Musings on provenance, workflow and (semantic web) annotations for bioinformatics. In Buneman and Foster [4].
11. G. Kahn and D. B. MacQueen. Coroutines and networks of parallel processes. In *Proc. of the IFIP Congress*, 1977.
12. E. A. Lee and D. G. Messerschmitt. Static scheduling of synchronous data flow programs for digital signal processing. *IEEE Trans. Comput.*, C-36, 1987.
13. E. A. Lee and S. Neuendorffer. Actor-oriented models for codesign: Balancing re-use and performance. In *Formal Methods and Models for System Design*. Kluwer, 2004.
14. E. A. Lee and T. M. Parks. Dataflow process networks. *Proc. of the IEEE*, 83(5), 1995.
15. B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. A. Lee, J. Tao, and Y. Zhao. Scientific workflow management and the kepler system. *Concurrency and Computation: Practice & Experience*, 2005.
16. B. Ludäscher and C. A. Goble, editors. *ACM SIGMOD Record, Special Section on Scientific Workflows*, volume 34(3), September 2005.
17. L. Moreau, O. Rana, and D. Walker. Provenance aware service-oriented architecture (pasoa). pasoa.org, 2006.
18. T. M. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, R. M. Greenwood, T. Carver, K. Glover, M. R. Pocock, A. Wipat, and P. Li. Taverna: A tool for the composition and enactment of bioinformatics workflows. *Bioinformatics*, 20(17), 2004.
19. PHYLIP Phylogeny Inference Package. <http://evolution.gs.washington.edu/phylip.html>.
20. Y. Simmhan, B. Plale, and D. Gannon. A survey of data provenance in e-science. In Ludäscher and Goble [16].
21. D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: The Condor experience. *Concurrency – Practice and Experience*, 17(2–4), 2005.
22. J. D. Thompson, D. G. Higgins, and T. J. Gibson. CLUSTAL W: Improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position specific gap penalties and weight matrix choice. *Nucleic Acids Res.*, 22:4673–80, 1994.
23. J. Widom. Trio: A system for integrated management of data, accuracy, and lineage. In *Conference on Innovative Data Systems Research (CIDR)*, 2005.
24. S. Wong, S. Miles, W. Fang, P. Groth, and L. Moreau. Provenance-based validation of e-science experiments. In *ISWC*, 2005.
25. J. Zhao, C. Goble, R. Stephens, and S. Bechhofer. Linking and browsing provenance logs for e-science. In *ICSNW*, 2004.