

SimP: Secure Interoperable Multi-Granular Provenance Framework

Amani Abu Jabal, Elisa Bertino

Dept. of Computer Science, Purdue University, West Lafayette, USA

{aabujaba,bertino}@purdue.edu

Abstract— We propose a provenance framework which includes an expressive provenance model able to represent the provenance of any data object captured with various granularities. The model is represented according to relational and graph specifications. The framework is interoperable with two standard provenance models: OPM and PROV. In addition, the framework captures access control policies for data objects and secures the provenance storage itself. We have integrated our framework with CRIS - a real world system for managing scientific data.

Keywords—*Provenance Model, Provenance Framework, Relational, Graph*

I. INTRODUCTION

Data provenance, one kind of metadata, pertains to the derivation history of a data object starting from its original sources [3]. The term data object refers to data in any format (e.g., files, database records, or workflow templates). One example of provenance model for scientific applications is Chimera [1] that documents workflows generating data. Data provenance is critical for many purposes including: assessing data quality based on its ancestral data and derivations, detecting sources of errors and anomalies, providing an audit trail for regulatory purposes, and assessing data trustworthiness.

Building a comprehensive provenance infrastructure involves addressing the following requirements:

1. **multi-granular provenance model**: the provenance infrastructure should provide a rich provenance model capable to represent the provenance of data objects with different granularities (e.g., file, database record, data in a workflow);
2. **provenance queries**: the infrastructure should support queries for inspecting various aspects of the provenance;
3. **security**: the framework should be able to capture the access control policies, according to which users had been granted permissions to the data, at the time of data access; the infrastructure should also control access to provenance data storage as provenance may be sensitive;
4. **interoperability services**: the framework should provide services supporting provenance integration across different systems.

Despite a large number of research efforts devoted to provenance management, only a few provenance infrastructures have been proposed. Chimera [1], myGrid [2], and Karma [7] are examples of provenance systems. However, the provenance models of these systems are tailored to their

specific applications and therefore are not general enough. PASS [4] is a provenance management system for file systems; it provides a custom query tool but it does not support security and different granularity levels for provenance metadata. In addition, there are two standard provenance models: Open Provenance model (OPM) [5] and PROV [11]. These two models are interoperable and generic so that they are able to represent provenance for different systems and applications. However, their major limitation is that they are not able to represent metadata about access control policies. Ni et al. [8] has proposed a provenance model that focuses on access control policies for provenance. However, Ni's model is not able to support different granularity levels. Sultana and Bertino [14] have designed an initial comprehensive provenance infrastructure. However, this infrastructure has several limitations, including the lack of a query language and lack of interoperability services for provenance. In addition, such framework has not been implemented nor integrated with an actual data management system.

In this paper, we thus design and implement the first comprehensive provenance infrastructure addressing the four requirements discussed earlier. Our contributions include:

- A data provenance model extended from the provenance model proposed by Sultana and Bertino [14]. We provide specifications of this model according to a relational and graph model.
- A mapping ontology to support interoperation of our provenance model with both OPM and PROV.
- The integration of our provenance framework with the *Computational Research Infrastructure for Science* (CRIS) [13]. CRIS is widely used at Purdue University for managing scientific data from many different research areas, including biology, bio-chemistry, water management, and social sciences.

The rest of the paper is organized as follows: Section II introduces our provenance framework including the provenance model, mapping ontology. Section III discusses related work and Section IV outlines conclusions and future work.

II. PROVENANCE FRAMEWORK

Our provenance framework as shown in Fig. 1 is composed of several components that we discuss in what follows.




Figure 1: SimP Framework

A. Provenance Model

In our provenance model, called SimP, the provenance of a data object records the history of its input *data*, *processes*, *operations*, *communications*, *actors*, *environments*, and *access controls*.

A *process* manipulates input data objects by performing a sequence of operations to generate other data objects. A process might be a service in a user application (e.g., a workflow in a scientific experiment application) or an operating system level process (e.g., executing a script by shell command in UNIX). At a more detailed level, an *operation* executes a sub-task which is part of a process execution. The operations may generate/modify persistent data or intermediate results. In provenance, it is crucial to capture the origin of the generated data. Such information is captured by a *data lineage* entity. In data lineage, a data item is described by the source data and the operations used to derive the data item from source data. Lineage helps in producing the data dependency graph of a data object and implicitly describing the process dependency.

The operations in the same process or in two different processes interact by real or virtual messages. In our model, such interaction is referred to as *communication*. There are two types of communication between operations. The first type refers to the completion of an operation followed by the start of another operation. The second type refers to an operation executed (i.e., started and completed) within the execution of another operation. We refer to the first type of communication as *sequential* and to the second as *composition*. A communication may involve data passing if the operation which initiated the communication generates data. Examples of communications include data flow and copy-paste in UNIX. On the other hand, a process may initialize another process to be executed. The process which invokes the initialization is the parent process and the newly created process is the child process.

Processes (including their operations) and data are manipulated by *actors* which can be human subjects or workflows. Capturing information about actors, who actuate the activities changing data objects, helps in detecting intrusion, data misuse, or system changes. A user may authorize other users to perform certain activities on his behalf. In this model, *data* objects are attributed to actors to identify users who inserted input data or generated output by executing a process. Processes also have a context that affects their execution and output. Such context is represented by the

environment which refers to a set of parameters and system configurations. Environment information helps in understanding the system context in which processes were executed and data output were generated.

Our model is security-aware. Following the model by Ni et al. [8], our model is able to represent the access control policies in place at the time of data manipulation by the actors. Information about access control policies include which actors are authorized to utilize which processes and operations on which data. Such information is modeled by the *Access Control Policy* entity which includes actor and policy information. The policy object might refer to processes, or operations specified by the policy subject.

All fundamental provenance entities contain a domain attribute which might be used to specify the scope of provenance information (e.g., where processes and data manipulations executed) especially when providing a provenance storage for different systems. In addition, the domain value might include more detailed scope information (e.g., a particular application or workflow). The domain attribute is essential for providing an abstract domain view of the provenance graph.

Our framework supports the specification of the provenance model according to two representations: relational and graph.

1) Relational Representation

The relational model representation of SimP is shown in Fig. 2. Based on the abstract description of our model, the fundamental entities are stored in six fundamental tables (i.e., *Data*, *Processes*, *Operations*, *Actors*, *Environments*, and *Access Control Policies*). In addition, there are tables maintaining many-to-many relationships among the fundamental tables (i.e., *Lineages*, *Communications*, *Process Input Data*, *Process Output Data*, *Operation Input Data*, *Operation Output Data*, and *Delegations*). These tables are connected through a set of referencing relations (i.e., foreign key constrains). Each table has a unique identifier and consists of several attributes.

Each data record contains description, value, and actor ID. An example of a data object is a file. The actual data object identifier is different from data record identifier in the provenance storage. Suppose that the data object is a script file named f_i . This file might be edited by different users at different times. Thus, the provenance storage contains multiple records for f_i and each record is identified by different data ID but the description attributes of these records are similar (i.e. f_i). In our model, every data object is attributed to an actor to trace who created the object or generated it by executing a process.

Each process record has a unique ID and is executed by an actor in a certain environment. A process manipulates certain data and may generate other data so process's input and output data are recorded in the Process Input Data and Process Output Data tables. Processes are categorized into two types: application process or system process. If the process belongs to an application (e.g., a scientific workflow), it contains a workflow ID. Otherwise, it is a system process. In the case of a

system process, the workflow ID attribute refers to the workflow hosting the application process which encapsulates the system process. Each process has a description which is the actual identifier of the process (e.g., executing the script file (*fi*)). A process might be executed multiple times. Therefore, in order to capture provenance for the same process multiple times we have an automatic generated ID to distinguish process records. A process might be initialized (i.e., forked) by another process so we store such information in the parent process attribute.

Operation record attributes include ID, description, and process ID. Depending on applications, the operation description attribute might contain different definitions (e.g., a function name or a block of source code statements). The output of an operation might be intermediate or persistent. The output is intermediate if it is used as input for another operation while a persistent output is final. The persistent output of operations is considered also as the output for the container process. On the other hand, all input data of a process can be input data for all its contained operations.

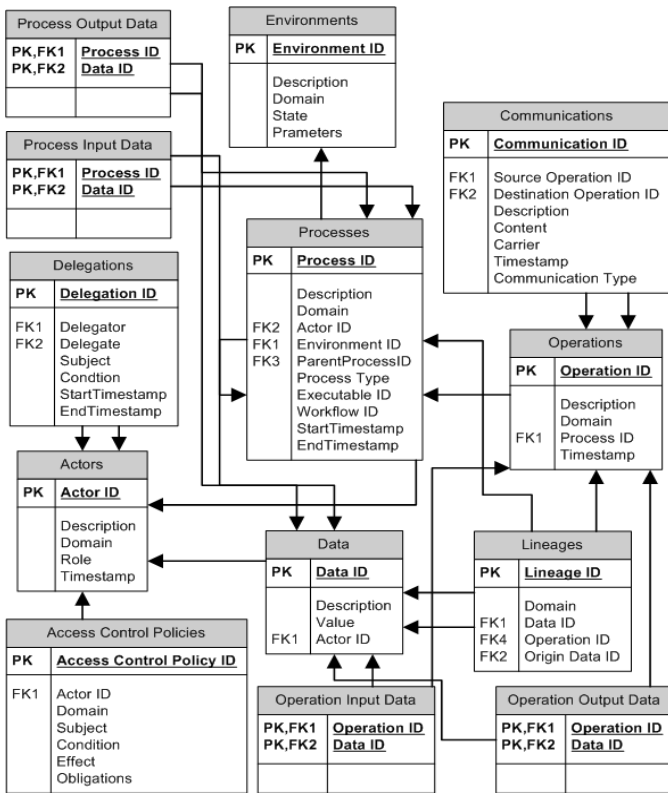


Figure 2: ER-Diagram of the SimP Model

Communication record attributes include description, carrier, the source operation ID, the destination operation ID, and type of communication. The detail level of the communication description depends on the applications. The communication channel (e.g., HTTP or SOAP) is described by the carrier attribute. The type of communication between two operations might be *sequential* or *composition*.

Lineage record attributes include lineage ID, data ID, and the ID of the operation that produced this data operating on the

input data identified by the data ID. A data item has multiple lineage records if it is generated by an operation which received multiple input data.

The attributes of an actor record include actor ID, description, and role. A role is a job of the actor. For an actor, the most recent record contains its current job while its previous records might include its previous jobs. An actor can delegate the execution of a process to another actor. Such information is stored in a Delegation record which contains delegator, delegate, subject, condition, start timestamp, and end timestamp. The subject might refer to processes or operations while the condition refers to the identifier of the subject.

The content of an environment record includes domain, state, and parameters attributes. In our model, each process belongs to one environment while an environment contains multiple processes. An environment timestamp is determined by its parent process record (i.e., the first process executed within the environment).

The attributes of an access control policy record include access policy ID, actor ID, subject, condition, effect, obligations. Such records capture the access control enforced at the time when the data of interest has been manipulated. The actor ID records the actor privileged to perform certain operations or processes. The policy subject might refer to processes, operations, or communications while the condition refers to the identifier (e.g., process or operation description) of the subject.

2) Graph Representation

The relational model is suitable for storing provenance in a relational database. However, as the standard provenance models (i.e., OPM and PROV) are modelled according to a graph-based format, developing an interoperability ontology mapping from these models onto our model requires a corresponding graph representation of our model. Thus, we also introduce a graph model specification of SimP.

Our graph model consists of six nodes and twelve types of edges. The graph nodes include *Data*, *Process*, *Operation*, *Actor*, *Environment*, and *Access Control policy*. Each graph node has a set of attributes similar to its corresponding entity table described in the relational model. Furthermore, an edge represents a relationship between the source of the edge and the destination of the edge. The kinds of the relationship (and thus the meaning of the edges) are specified by the types of the nodes that are the end-points of the edges. These types of relations are defined as follows:

- *used*: an edge that connects a source Process or Operation to a destination Data. It indicates that the process/operation required the availability of certain data to be able to complete its execution.
- *wasGeneratedBy*: an edge that connects a source Data to a destination Process or Operation. It indicates that the process/operation was required to initiate its execution for the data to have been generated.
- *wasDerivedFrom*: an edge that connects a source Data to a destination Data. It indicates that the destination

data needs to have been generated for the source data to be generated.

- *wasExecutedBy*: an edge that connects a source Process to a destination Actor. It indicates that a process with all its operations was executed by the actor.
- *wasInformedBy*: an edge that connects a source Operation to a destination Operation. It indicates that the execution of the destination operation was followed by the execution of the source operation.
- *wasEncapsulatedBy*: an edge that connects a source Operation to a destination Operation. It indicates that the execution of the source operation is part of the execution of the destination operation.
- *wasPartOf*: an edge that connects a source Operation to a destination Process. It indicates that the execution of the source operation is part of the execution of the destination process.
- *wasForkedBy*: an edge that connects a source Process to a destination Process. It indicates that the execution of the source process is initiated by the destination process.
- *wasInContext*: an edge that connects a source Process to a destination Environment. It indicates that the execution of the source process was in the context of the destination Environment.
- *wasGrantedTo*: an edge that connects a source Access Control Policy to a destination Actor. It indicates that the source access control policy was enforced on the destination actor at provenance capture time.
- *actedOnBehalfOf*: an edge that connects a source Actor to a destination Actor. It indicates that the execution of the action performed by the source actor was delegated to the source actor by the destination actor.
- *wasAttributedTo*: an edge that connects a source Data to a destination Actor. It indicates that the source data was manipulated by the destination actor.

Based on the proposed provenance specifications for relational and graph-based representations, our framework supports two types of storage: relational database (i.e., MySQL) or graph database (i.e., Neo4J) for storing provenance metadata. For this sake, our framework has an abstract storage interface. This abstract interface communicates with either MySQL adapter or Neo4J adaptor. The default storage is Neo4J and the framework can be configured to change to the second database. The administrator is able to change the target storage type at any time but he should first import the database from Neo4J to MySQL (or vice versa).

B. Interoperability With Other Provenance Models

Provenance interoperability refers to the ability to integrate and convert provenance information represented by different provenance models in order to facilitate provenance data interchange. Our model supports interoperability with two well-known standard provenance models: OPM [5] and PROV

[11]. Below we provide some background about OPM and PROV model ontologies and address the interoperability challenge by defining a mapping ontology that maps provenance information expressed in each of those two standard models into provenance information expressed in SimP.

1) Mapping from OPM Standard Model

OPM [5] represents provenance information as a directed graph which consists of nodes and edges. The nodes comprise three types of entities: *Artifacts* (i.e., data) which represent resources, *Processes* which represent actions or steps of action performed on artifacts, and *Agents* (i.e., actuators) which control the processes. The edges represent the dependencies and relations among the entities. There are five types of edges: *used*, *wasControlledBy*, *wasGeneratedBy*, *wasDerivedFrom*, and *wasTriggeredBy*. Each edge is distinguished by its source and destination: a process used an artifact, a process was controlled by an agent, an artifact was generated by a process, an artifact was derived from another artifact, and a process was triggered by another process. OPM also introduced the concept of *account* to represent a particular view of the provenance.

The conversion from OPM to our provenance model is straightforward because of the rich vocabularies in our model. Table 1 shows the mapping from the OPM graph to our provenance model graph. Besides the intuitive mapping provided in Table 1, we need to consider the following:

- OPM does not have a finer granularity level of description for a process as a set of operations. When converting a process from OPM to SimP, we create a process and an operation and link them by a *wasPartOf* edge. Mapping an OPM process into a SimP operation assumes that the relations connected with the OPM process are mapped onto relations with SimP operation. So the purpose of creating a dummy SimP process is only to host the created SimP operation.
- In OPM, the intercommunication between a process and another process is identified by one edge type, that is, *wasTriggeredBy*, which does not capture the exact meaning of process-process relations (i.e., communication relation or parent-child relation). In our model, these relations are represented by three types of edges (*wasForkedBy*, *wasEncapsulatedBy*, and *wasInformedBy*). *wasForkedBy* represents the parent-child relation between two processes. On the other hand, *wasInformedBy* and *wasEncapsulatedBy* represent communications between two operations (in the same process or different processes). More specifically, a *wasInformedBy* represents a sequential communication and a *wasEncapsulatedBy* represents a composition communication. When performing a conversion, we map the OPM *wasTriggeredBy* edge to *wasInformedBy* in our model under the assumption that all OPM processes are related with sequential communication.
- OPM supports a particular view of provenance by including the “account” attribute of the graph nodes. By contrast, in our model, we support such view of provenance using the environment node. So each node in OPM that has an

account attribute is mapped onto a *wasInContext* edge connecting the mapped node with the environment node.

TABLE 1: MAPPING FROM OPM TO SIMP

	<i>OPM</i>	<i>SimP</i>
Nodes	Process	Process, Operation, WasPartOf
	Artifact	Data
	Agent	Actor
Edges	Used	Used
	WasGeneratedBy	WasGeneratedBy
	WasDerivedFrom	WasDerivedFrom
	WasControlledBy	WasExecutedBy
	WasTriggeredBy	WasInformedBy

2) Mapping from PROV Standard Model

A W3C provenance standard model called PROV [11] was published in 2013 based on a revision of OPM. A PROV graph contains three types of nodes: *Entities* (i.e., data) to represent resources, *Activities* to represent to actions performed on entities and *Agents* to model parties responsible for activities. Additionally, PROV includes seven types of edges: *used* (some entity was used by an activity), *wasAssociatedWith* (an agent was engaged in some activity), *wasGeneratedBy* (an entity was generated by an activity), *wasDerivedFrom* (an entity derived another entity), *wasAttributedTo* (an agent used an entity), *actedOnBehalfOf* (an agent acted on behalf of another agent) and *wasInformedBy* (an activity sent its result data to another activity).

Table 2 shows the mapping from a PROV graph to a SimP graph. In addition, we address the same issues discussed for the mapping from OPM.

TABLE 2: MAPPING FROM PROV TO SIMP

	<i>PROV</i>	<i>SimP</i>
Nodes	Activity	Process, Operation, WasPartOf
	Entity	Data
	Agent	Actor
Edges	Used	Used
	WasGeneratedBy	WasGeneratedBy
	WasDerivedFrom	WasDerivedFrom
	WasAssociatedWith	WasExecutedBy
	WasInformedBy	WasInformedBy
	WasAttributedTo	WasAttributedTo
	ActedOnBehalfOf	ActedOnBehalfOf

By following the mapping ontology described earlier, we implemented a conversion tool that facilitates the conversion from the standard models (OPM, and PROV) to SimP model. The input of the tool is an XML -formatted file containing data provenance encoded according to the OPM model or the PROV model. The conversion tool stores the converted provenance data into our provenance storage, that is, MySQL or Neo4J.

C. Integration with the CRIS System

We integrated our provenance model in the *Computational Research Infrastructure for Science* (CRIS) [13]. CRIS is a

scientific data management workflow cyberinfrastructure for scientists lacking extensive computational expertise. The application is currently used by a community of users in Agronomy, Biochemistry, Bioinformatics, and Health Care Engineering at Purdue University. Previously, CRIS had its own provenance model mainly based on versioning mechanism. Such mechanism is not able to capture provenance at different granularities since it is data centric. Our model is data and operational centric in that we maintain metadata about the derivation of every data object (i.e., what is the origin data object and what is the deriving operation).

Within the CRIS system, we have a provenance logging component based on aspect oriented programming to instrument the application code. The logging component collects a set of appropriate provenance logs while the CRIS is running. The logs use the XML-based representation of provenance records to facilitate parsing them into the SimP model. Periodically, the CRIS provenance logs are fetched and converted into another XML-format file. The new XML file contains a data dependency provenance graph following the SimP representation ontology.

D. Security

Due to the sensitivity of data provenance information collected and stored in provenance storage, security is an essential factor and requirement. Hence, we incorporated an additional entity to specify the privileges granted to actors for accessing the provenance storage. Such entity, referred to as Provenance Query Authorization, records information about which actors has which authorizations. An authorization is expressed as subject, condition, and effect fields. The subject refers to the type of provenance storage entity (e.g., Processes entity) while the condition field identifies the provenance entity (e.g. Process Description value). The effect field indicates the authorization status (e.g., granted or revoked).

The security requirement in our framework is thus addressed by the following features: a) the access control policy which is a main entity in the SimP model; and b) the provenance query authorization to secure the access to the sensitive information in the provenance storage.

E. Granularity

Another key feature of our provenance framework is the ability to specify and modify the granularity level needed to capture provenance for specific records or entities. For this purpose, in our provenance database, we include an additional entity, referred to as *Granularity Policy*, which is not part of the provenance model but it is part of our provenance framework. A granularity policy enables users to specify the desired level of provenance details to be captured and stored (e.g., capturing provenance at the activity level in scientific provenance workflow, or the operating system level to capture system configurations). A Granularity Policy record includes the following fields: granularity policy ID, actor ID, subject, condition, granularity type. The subject may refer to the targeted process, operations, or communications, whereas the condition refers to the identifier of the subject. The granularity type is the detail level of the required provenance.

The multi-granularity requirement in our framework is thus addressed by the following features: a) a provenance model able to represent provenance metadata for data objects at various granularity levels; b) support for the integration with systems utilizing different provenance capturing mechanisms (i.e., provenance for workflow data or file system); and c) support for granularity preferences based on granularity policies.

III. RELATED WORK

Chimera [1], myGrid [2], and Karma [7] are examples of workflow-based provenance systems in which the captured provenance is about data-centric workflows. In Chimera, workflow data is represented by a special language, referred to as Virtual Data Language (VDL) which describes provenance as relationships among datasets, procedures, invocations of procedures or tasks. In myGrid, the provenance model represents service invocations and their information (i.e., parameters, start and end times, data products used and derived). Karma collects provenance at three levels (i.e., Workflow, Service, and Application) and it uses the concept of activities (e.g., Workflow-Started or Workflow-Finished) that take place at these three levels to collect provenance

On the other hand, PreServ [12] is an example of process-based provenance systems. A process-based provenance model represents the relationships between services and data. These relationships have three categories: service-service (i.e., source and sink service), data-data (i.e., data derivation), service-data (a service consumes data or a service produces data).

An example of operating system based provenance system is PASS [4]. PASS works at the level of shared storage system to capture provenance about executed programs, their inputs, and outputs. Another example is ES3 [6] which records provenance metadata including input/output data objects, and domain names.

As already mentioned, models of existing provenance systems apply only to specific applications/domains and do not support security. OPM and PROV provide standard provenance representations. However, they are not able to represent information on access control policies. OPM Toolbox [9] allows one to create OPM graphs while ProvToolbox [15] allows one to create PROV graphs and convert between PROV data model representations (e.g., XML or JSON).

The provenance model by Ni et al. [8] (extended by Cadenhead et al. [10]) is a general model supporting security. This model represents provenance data at a granularity of operation and thus is unable to distinguish between different granularity levels. Sultana and Bertino [14] provide an initial comprehensive provenance infrastructure which we extend in this paper.

IV. CONCLUSION AND FUTURE WORK

In this paper, we introduce SimP - a comprehensive provenance framework integrated with the scientific data

management system CRIS [13]. The framework includes a comprehensive provenance model provided with relational and graph specifications. Our provenance model is interoperable with the OPM and PROV provenance models.

As future work, we plan to design and implement specialized query languages for our framework and investigate efficient compression techniques for our provenance model.

ACKNOWLEDGMENT

This work is supported in part by NSF under award CICI-1547358. The opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any of the sponsors.

REFERENCES

- [1] Foster, I., Vöckler, J., Wilde, M. and Zhao, Y., 2002. Chimera: A virtual data system for representing, querying, and automating data derivation. In Scientific and Statistical Database Management, 2002. Proceedings. 14th International Conference on (pp. 37-46). IEEE.
- [2] Zhao, J., Goble, C., Stevens, R. and Bechhofer, S., 2004. Semantically linking and browsing provenance logs for e-science. In Semantics of a Networked World. Semantics for Grid Databases (pp. 158-176). Springer Berlin Heidelberg.
- [3] Simmhan, Y.L., Plale, B. and Gannon, D., 2005. A survey of data provenance in e-science. ACM Sigmod Record, 34(3), pp.31-36.
- [4] Muniswamy-Reddy, K.K., Holland, D.A., Braun, U. and Seltzer, M.I., 2006, June. Provenance-Aware Storage Systems. In USENIX Annual Technical Conference, General Track (pp. 43-56).
- [5] Moreau, L., Freire, J., JFutrelle, J., McGrath, R.E., Myers, J., aulson, P. The Open Provenance Model (v1.00), Tech. Rep., University of Southampton, URL <http://eprints.ecs.soton.ac.uk/14979/1/opm.pdf>, 2007.
- [6] Bowers, S., McPhillips, T., Riddle, S., Anand, M.K. and Ludäscher, B., 2008. Kepler/pPOD: Scientific workflow and provenance support for assembling the tree of life. In Provenance and Annotation of Data and Processes (pp. 70-77). Springer Berlin Heidelberg
- [7] Simmhan, Y.L., Plale, B. and Gannon, D., 2008. Query capabilities of the Karma provenance framework. Concurrency and Computation: Practice and Experience, 20(5), pp.441-451.
- [8] Ni, Q., Xu, S., Bertino, E., Sandhu, R. and Han, W., 2009. An access control language for a general provenance model. In Secure Data Management (pp. 68-88). Springer Berlin Heidelberg.
- [9] OPM Toolbox, <https://github.com/lucmoreau/OpenProvenanceModel>
- [10] Cadenhead, T., V. Khadilkar, M. Kantarcioglu, and B. Thuraisingham. "A language for provenance access control." In CODASPY, pp. 133-144. ACM, 2011.
- [11] PROV-Overview: <http://www.w3.org/TR/2013/NOTE-prov-overview-20130430/>
- [12] Hoekstra, R. and Groth, P., 2013. Linkitup: link discovery for research data. In AAAI Fall Symposium Series Technical Reports (No. FS-13-01, pp. 28-35). AAAI Publications.
- [13] Dragut, Eduard C., et al. "CRIS—Computational research infrastructure for science." Information Reuse and Integration (IRI), 2013 IEEE 14th International Conference on. IEEE, 2013.
- [14] Sultana, S. and Bertino, E., 2015. A Distributed System for The Management of Fine-grained Provenance. Journal of Database Management (JDM), 26(2), pp.32-47.
- [15] ProvToolbox, <https://github.com/lucmoreau/ProvToolbox>