

ProvGen: generating synthetic PROV graphs with predictable structure

Hugo Firth and Paolo Missier

School of Computing Science, Newcastle University, UK
{h.firth, paolo.missier}@ncl.ac.uk

Abstract. This paper introduces *provGen*, a generator aimed at producing large synthetic provenance graphs with predictable properties and of arbitrary size. Synthetic provenance graphs serve two main purposes. Firstly, they provide a variety of controlled workloads that can be used to test storage and query capabilities of provenance management systems at scale. Secondly, they provide challenging testbeds for experimenting with graph algorithms for provenance analytics, an area of increasing research interest. *provGen* produces PROV graphs and stores them in a graph DBMS (Neo4J). A key feature is to let users control the relationship makeup and topological features of the graph, by providing a seed provenance pattern along with a set of constraints, expressed using a custom Domain Specific Language. We also propose a simple method for evaluating the quality of the generated graphs, by measuring how realistically they simulate the structure of real-world patterns.

1 Introduction

Every piece of data ever produced, either manually or automatically, has a provenance. This is metadata that provides an account of how the data was created. Examples include a blog’s author, the history of a piece of software along with its contributors, the instruments used to take a measurement, and their settings; or a description of an experimental process used to produce a scientific result. The PROV data model for provenance [MMB⁺12], endorsed in 2013 by the W3C, provides a formal and domain-agnostic grounding for provenance, in the form of UML and OWL models, and RDF, XML, and relational (PROV-N [MM-CSR12]) serializations. We refer to PROV instances as digraphs, where nodes are of three possible types: *Entities* (for data, documents, anything that has provenance), *Activities*, which model the execution of a data consumption and production process; and *Agents*, to whom Entities can be attributed, and who hold responsibility for carrying out Activities. The edges represent instances of relationships amongst the nodes, which are documented in the PROV-DM specification [MMB⁺12].

The provenance traces associated with a homogeneous data collection (a scientific data repository, all the blogs hosted on a particular site, all the artifacts associated with a complex software project) also naturally form a collection. Such collections grow in size both with the number of underlying data products,

and with the complexity of their production process. Fig.1 suggests how different collections can be placed into a space defined by volume, i.e., the number of traces in a collection, and by the typical size of a trace within a collection. For instance, many small traces (upper left) may be associated with a large repository of scientific data, while complex software with a long history may be represented by many large traces (upper right), as exemplified by the Git2Prov [DMV+13] tool.

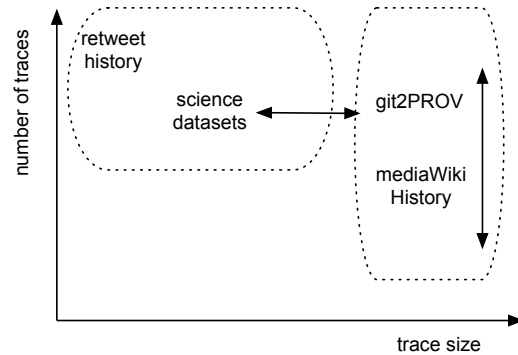


Fig. 1: A simple space for homogeneous provenance collections

Arguably, the value of provenance comes not only from querying the content of individual traces, but also from analytics, which can only be computed on whole collections. It is therefore important for practical applications to demonstrate the effectiveness of a data and service architecture to manage large bodies of provenance, with special focus on the upper quadrant of our size/volume space. Thus, we expect that the design of scalable repositories for provenance traces should be a natural concern in provenance management. A number of recent efforts have been documented on nascent provenance management infrastructure [CAB+13, CLFF10, LLCF11, MMW+12], and there is evidence of the emergence of applications that require provenance querying in a variety of settings (eg [MOnH+13, ddOOn+12]). However, unlike other “big data” domains such as Linked Data and more generally RDF triple stores, where performance benchmarking is established practice, to the best of our knowledge no community-made benchmarking and commonly accepted datasets that are specific to provenance are available.¹ This makes it difficult to benchmark and compare different implementations with regards to storage techniques, query models, and analysis algorithms.

¹ The use of community datasets for comparing the performance of predictive models has also long been commonplace within the data mining and KDD community, where challenge datasets are regularly used.

This is somewhat counter-intuitive, given the amount of provenance that is generated, in domains such as those alluded to above. In fact, only a handful of real datasets are currently available through a community process, i.e., the first ProvBench initiative in 2013 (<http://bit.ly/1fB0swR>)², and even fewer conform to the recent PROV standard and are therefore interoperable. Existing benchmarking datasets which apply to RDF triple stores³ are not adequate, because they fail to account for the specific data model and semantics of PROV, as well as for the specific requirements of provenance query and analysis.

1.1 Contributions

Our assumption is that synthetic PROV graphs can be a valuable complement to emerging natural provenance collections, provided that their structural properties reflect specific provenance patterns, with control over their repetition and variability, and at varying scales. Such graphs can be used both for benchmarking emerging provenance management systems, as well as to test analytics algorithms that operate naturally on large provenance collections.

Our main contribution (Sec. 2) is the design and implementation of *provGen*, a PROV generator that is designed to help populate the space described in Fig. 1. *provGen* “grows” collections of synthetic PROV graphs in a way that conforms to real-life provenance patterns. These are currently user-defined and modelled after patterns found in specific domains, and which reflect the nature of the data generation process described by the provenance. For instance, the prevalent provenance pattern for a Media Wiki website, which we refer to as the “document revision” model, involves multiple revisions of articles, by multiple editors (Fig. 2). Git repositories exhibit similar patterns, which reflect the revision history of the code. These patterns are different, for instance, from those for the provenance of data generated using a workflow, which reflect the consumer-producer graph structure of the dataflow specification.

Users control the “shape” of the graph being generated by *provGen* by providing two main elements. The first is a *seed graph*, which determines the specific types of nodes and the relationships amongst them to be considered, in an otherwise random generation process. The second element is a set of constraints, expressed using a dedicated Domain Specific Language (DSL), which limit the possible ways in which nodes and relationships are added. These two elements ensure a predictable general shape for the generated graph, as well as its compliance to PROV.

As discussed later, *provGen* relies on a graph DBMS backend (Neo4J). In particular, the generation algorithm is based on graph rewrite rules that are implemented using a combination of Cypher queries and Create statements.

² Further contributions are expected from the second ProvBench in 2014 (<http://bit.ly/1c0q5rS>).

³ The W3C maintains a list of those <http://bit.ly/1lhjvvn>

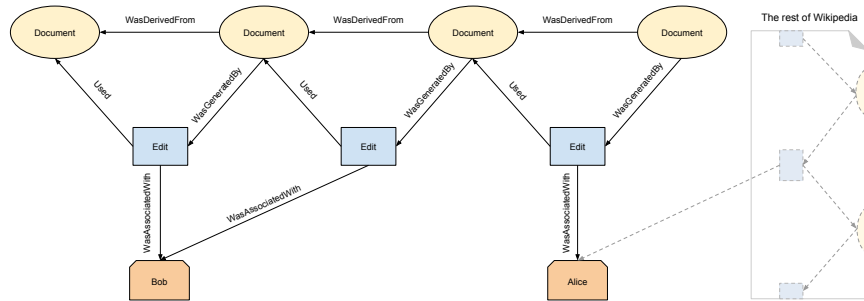


Fig. 2: The *document revision* provenance pattern in Wikipedia includes multiple derivation and editing activities by multiple user or bot agents.

1.2 Related work

A growing body of research is devoted to generating large bodies of synthetic graph data, either using purely random models [KN09, ER60], or by generating graphs that exhibit specific statistical properties [BA99, BB05, LCKF05]. One example is the *preferential attachment* model. Popularised by Barabasi and Albert [BA99], this model states that as new vertices are added to a graph, the probability of creating a relationship with node n is inversely proportional to the degree of n . This model generates a graph with a degree distribution which follows a power law.

An issue common to these models, emphasised for instance in a comprehensive survey on graph generators [CF06], is their focus on enforcing global properties of the generated graph, such as degree distribution, clustering coefficient, etc. A potential reason for this focus is that these generators are aimed at simulating social networks [PBE13, BB05], the statistical properties of which are based on large sets of examples, and thus are fairly well understood [MMG⁺07]. In contrast, our generation strategy relies on user-specified patterns, rather than a large set of pre-existing examples (in the future, we hope to be able to use patterns that have been automatically discovered from existing graphs, by means of standard graph mining techniques [KK04]). This has the advantage that the overall topology of the graph can be made to reflect desired semantic properties of the data, such as the average number of usages for a certain type of entity, the average number of association of an agent with activities, and so forth. Pham *et al.* [PBE13] are amongst the few to have addressed this problem. However, they focus on a loosely related issue, namely the correlation between node and relationship properties, such as an increased likelihood to be called “Joachim” if you live in Germany, and on generating realistic synthetic value dictionaries accordingly.

2 Graph generation model

Graph generation in *provGen* is an iterative process which starts from a single node. At each iteration, a collection of predefined *atomic rewrite* rules is used to

add a set of new nodes or relationships to the current graph. These rules account for all possible relation types that are defined in the PROV-DM specification. As an example, consider the definition of the $used(a, e)$ relation between an activity a and e an entity e . Three atomic graph rewrite rules are defined for this relation, namely (i) given an entity node e , add a new activity node a and an edge $used(a, e)$; (ii) given an activity node a , add a new entity node e and an edge $used(a, e)$; and (iii) given a pair of unrelated nodes (a, e) , add edge $used(a, e)$. Since each single PROV relation type induces three atomic rewrites, and we consider 13 types of relations from PROV, at each iteration *provGen* can potentially fire any of 39 different rules.

Users can control the execution of these rules and the overall effect of the generation process in three complementary ways, namely (i) by specifying a *seed graph*, (ii) by adding a set of constraints, and (iii) by specifying additional execution parameters. We now describe these in some detail.

1. *Seed graphs.* A seed graph specification restricts the set of rules to choose from, to only those corresponding to the relations that appear in the graph. As an example, the document revision pattern depicted in Fig. 2 may be expressed as follows, using PROV-N syntax:⁴

```
entity(e1, [ prov:type="Document" ])
entity(e2, [ prov:type="Document" ])
activity(a, 2013-11-16T16:00:00, 2013-11-16T16:05:00, [prov:type="edit"])
agent(ag, [ prov:type='prov:Person' ])
used(a, e1, 2013-11-16T16:00:00)
wasGeneratedBy(e2, a, -, [ ex:fct="save" ])
wasAssociatedWith(a, ag, -, [ prov:role="contributor" ])
wasDerivedFrom(e2, e1, a)
```

Using this graph, *provGen* determines that only *wasGeneratedBy*, *used*, *wasDerivedFrom* and *wasAssociatedWith* rules are to be used. Furthermore, it will associate the properties and values found in the seed graph, for instance `prov:type="edit"`, to the new nodes and relations.

2. *Constraints.* Even with this restriction, unconstrained generation would lead to a graph with arbitrarily high node degree and branching factor, which would bear little resemblance to the seed trace provided, except in its relationship makeup. To further control the generation process, the second user input consists of an additional set of constraints, specified using a natural and intuitive syntax. Constraints are syntactically similar to workflow *control-flow* patterns [VTKB03], expressing the required states of data being created.

Constraints consist of three structural components, as shown in the examples of Table 1, namely a *determiner*, an *imperative*, and a *condition*. The determiner is either variable (an **Agent**) or invariable (**the Agent**, `a1`) and determines the

⁴ Domain-specific properties have been added to nodes and relations to denote the role of entities, activities, and agents in the pattern.

Determiner	Imperative		Condition
	Requirement	Req. qualifier	
an Entity	has in degree	at most 1;	
an Agent	has relationship "WasAssociatedWith"	between 1, 1000 times, with distribution $\gamma(\dots, \dots)$,	unless it has relationship "ActedOnBehalfOf";
an Activity	has relationship "Used"	exactly 1 times,	unless it has property { "prov:type"="create" };
an Entity	has relationship "WasDerivedFrom",	at least 1 times,	unless it has relationship "WasGeneratedBy" with the Activity, a1, AND a1 has property { "prov:type"="create" };

Table 1: Examples of user-defined constraints for graph generation.

elements to which a constraint applies. Requirements on these elements are specified by means of the Imperative clause. For instance **has in degree** (the requirement) **at most 1** (a qualifier) allows a new incoming edge to be added to any Entity that has none. The qualifier may optionally include a probability distribution, as in the second example. This determines the likelihood that an action be taken in order to satisfy the requirement, namely the generation of a new *WasAssociatedWith* relation. Furthermore, a condition specifies the applicability of an imperative to a determined element, i.e. **when** (selective condition) or **unless** (greedy condition). Thus, the second constraint inhibits the creation of a new *WasAssociatedWith* relation for any Agent that already has a *ActedOnBehalfOf* relation associated to it. Conditions admit the use of logical connectives, as in the third and last constraint examples, and may predicate on properties that are mentioned in the seed graph, such as **prov:type** (pre-defined) or **ex:name** (user-defined). Finally, the last constraint shows an example of variable usage (**a1**).

Note that these constraints are in addition to those defined in the PROV-CONSTR document [CMM12]. For instance, *provGen* will not create a graph where entities are generated by multiple activities. The sketch in Fig.3 shows the different patterns obtained when generating the graph with and without enforcing the constraints.

A more complete account of the constraint DSL can be found as part of the *provGen* documentation⁵.

3. Execution Parameters. Finally, users may specify additional execution parameters to control the number of distinct (unconnected) graphs to be generated, as well as the average number of nodes and edges per graph. More advanced parameters can be used to control the average height (maximum depth) and width (maximum breadth) for each graph generated.

⁵ The *provGen* website: <http://bit.ly/1sPn0zg>




Fig. 3: Sketch of PROV graphs generated with and without enforcing user constraints

The combination of seed graph, constraints, and execution parameters leads to collections of PROV graphs that approximate real traces from different domains, and which can be used to populate selected areas of our provenance state (Fig.1). In Sec. 4 we briefly sketch the evaluation method we are using to test the quality of generated graphs, with respect to a large testbed of provenance graphs with known topological properties.

Overall, *provGen*'s generation process consists of a nested iteration loop. In the inner loop, *provGen* iterates over the set of active atomic rewrite rules. When a rule fires, any constraint that applies to the elements that it is operating upon is checked, and if any of those constraints is violated, the rule has no effect. This process is repeated in the outer loop, until a halting condition is satisfied, i.e., the desired size is reached, and the DSL constraints are satisfied.

3 Mapping the model to graph DBMS queries

provGen is implemented using the Neo4J graph DBMS⁶ as a back end. In particular, both atomic rewrite rules and user constraints are transparently compiled into **CREATE** and **MATCH** statements expressed in Cypher, Neo's declarative graph pattern language⁷. Queries (in addition to **CREATE** statements) are required at each iteration to test the requirements and conditions associated with user constraints (Table 1). This compilation step provides isolation from the data layer, delegating graph traversal to the underlying DBMS, and also provides flexibility for retargeting the graph generator to a different back end. A native graph DBMS also offers a more natural data model for PROV than a more traditional RDBMS solution.

The *provGen* architecture is shown in Fig. 4. Components are deployed on a server, which is reachable from a web based client application through a REST API. In the following sections, we focus on the steps involved in generating Cypher queries from rewrite rules and user constraints.

⁶ The Neo4j project: <http://bit.ly/Pwux7U>

⁷ Cypher documentation: <http://bit.ly/1k1IIMK>

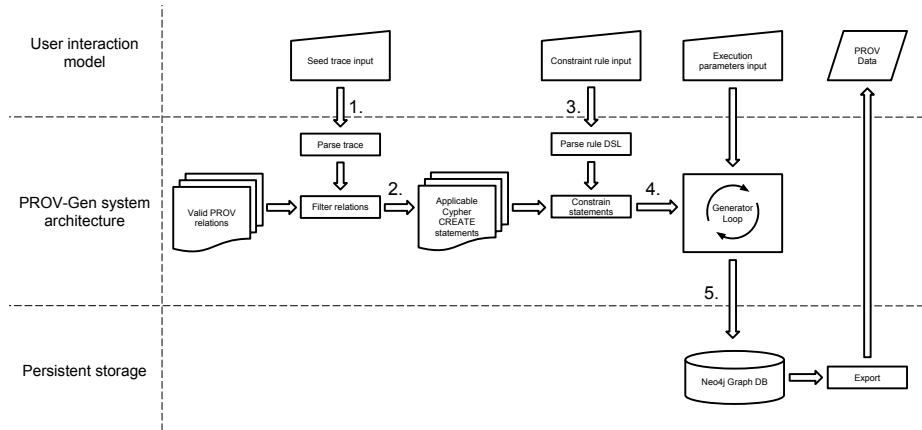


Fig. 4: *provGen* system architecture.

3.1 From seed traces to MATCH query clauses

The first step involves parsing the seed traces. Since these user-supplied samples of PROV data may be serialized into multiple formats, parsing relies upon several third party libraries, including the OWLAPI⁸ and ProvToolbox.⁹ This step results in a subset of the 39 pre-defined atomic graph rewrite rules, mentioned in Sec. 2, to be selected for the generation step.

Rewrite rules are statically mapped to Cypher queries. As an example, below we show the queries responsible for creating the PROV *used* relationship. Note that multiple queries are required in order to account for the directed nature of PROV relationships and the ability to create a edge between two pre-existing nodes.

- (1) MATCH (a:Activity {}) CREATE (a)-[:USED {}]->(:Entity {})
- (2) MATCH (a:Entity {}) CREATE (a)<-[:USED {}]-(:Activity {})
- (3) MATCH (a:Activity {}), (b:Entity {}) CREATE (a)-[:USED {}]->(b)

Query fragment (1) matches any node *a* of type Activity, it creates a new Entity node, and it connects it to *a* using a *used* relationship. Symmetrically, (2) adds a new Activity node to any existing Entity node. Finally, (3) takes a pair of *existing* nodes *a* (Activity), *b* (Entity) and again creates a *used* relationship between them.

The examples above show empty brackets, to indicate that no properties are associated to the nodes and relationships. However, all properties associated to the elements of the seed trace are also associated to corresponding elements of the new graph. Thus, for example activities would have a property `prov:type`, inherited from the activity node in the seed graph above.

⁸ The OWLAPI project: <http://bit.ly/N9hsPM>

⁹ The ProvToolbox project: <http://bit.ly/1fV95nN>

3.2 Constraints as WHERE clauses

The DSL parser¹⁰ separates the component elements of each constraint, namely *determiner*, *imperative* and *condition*. Requirements may be expressed on various graph features, i.e., nodes in/out degree, relationship, property, etc. . . . Each type of requirement is compiled into a Cypher query **WHERE** clause. These clauses are then added to the **MATCH** statements that represent the atomic rewrite rules, to form complete queries. Consider the following example:

```
an Activity has relationship ‘‘Used’’ exactly 1 times,  
    unless it has property {‘‘ex:name’’:‘‘create’’};  
an Activity has degree at most 5;
```

These constraints are easily interpreted in the context of a document revision pattern, where activities are edits of document versions, which produce a new version. For these activities, we stipulate that they use only one entity (the original document). Activities that create new documents are exceptions, noted by the `ex:name=create` property, and these activities are allowed to use zero or more input documents. Additionally, we add an upper bound to an Activity node’s degree to illustrate a more complex constraint.

The constraint is compiled into query fragments (4) and (5) in the Cypher query below, where they are merged with the **MATCH** and **CREATE** clauses of atomic query (1) from the example above:

```
(1) MATCH (a:Activity {})  
(4) MATCH (a)-[r]-()  
(5) WHERE NOT a.ex_name = ‘‘create’’ AND NOT count(r) >= 5  
(1) CREATE (a)-[:USED {}]->(:Entity {})
```

The query specifies at the same time the node and relationship generation, and the constraint. The **MATCH** clauses bind variables `a` and `r` to an Activity and to the set of its edges, respectively (either incoming or outgoing, as no direction is specified). The **WHERE** clause ensures that the **CREATE** statement (which creates a new *used* relationship) is only executed on `a` if the ‘‘`ex:name`’’ property is not ‘‘create’’, and the number of edges in set `r` is at most 4.

3.3 Generator Loop

The generator loop (Fig. 4) accepts a collection of atomic create operations, selected and constrained as described above, and repeatedly iterates over it, executing each associated Cypher query against the underlying graph database.

The generator loop has several halting conditions: both explicit, where execution parameters, detailed in Sec. 2, halt generation as the order $|V|$ and size $|E|$ of the graph reach their specified maxima; and implicit, where constraint rules may prevent the execution of individual operations in order to avoid violating specified range requirements. Note that limits in cardinality imposed by execution parameters may be met before the minimum requirements of a constraint rule are satisfied. When this is the case, *provGen* gives priority to the user constraints, to ensure that those are not violated.

¹⁰ The parser is implemented using Scala parser combinators: <http://bit.ly/1cURrAo>

4 Evaluation methodology

The main purpose of *provGen* is to fulfill the need to generate a possibly large number of provenance graphs for data domains where provenance is not yet routinely collected, or is not abundant. Yet, our evaluation of the system’s effectiveness relies on precisely those domains where large provenance collections are available. Specifically, we evaluate *provGen* by comparing selected properties of existing “real-world” provenance graphs, which we call *control* set, to those of generated graphs (the *test* set) intended to emulate them. Using this approach, we aim to empirically demonstrate that *provGen* may be configured to generate datasets that are “similar” to those produced by multiple different sources of provenance.

Our evaluation is ongoing. Here we illustrate the approach using one single control set, namely a set of Wikipedia provenance traces, representative of the *document revision* pattern, taken from the ProvBench repository and compliant with PROV.¹¹ The control graphs include about 4,000 nodes and 6,000 relationships. Our test set consists of two synthetic datasets of roughly the same size as the control, produced using *provGen* with a user-created seed trace for the document revision pattern, along with constraints and parameters.

In this initial evaluation we have considered three simple criteria. Firstly, we note that in the control set, which follows the linear Wikipedia pattern (Fig. 2), each Entity is used exactly once. Thanks to our user constraints, this is easily replicated exactly in the test set. Secondly, as example criteria we additionally consider the *number of associations per Agent*, and the *average number of entities with distinct titles contributed to, per Agent*. In the control, each Agent has 2.4 associations on average (std dev. 6.2), while in our test set it has 2.9. The average number of contributions per Agent is 1.1 in the control (std dev 0.8), while in the test is 1.8. Encouraged by these preliminary results, we are now in the process of more extensively testing *provGen* using a variety of criteria that can be easily measured both on control and on test graph.

5 Conclusion

In this paper we have presented *provGen*, a PROV-specific graph generator driven by user-defined seed graphs, which represent provenance patterns, and additional user-defined constraints designed to enforce semantics properties of the generated graph. Constraints are expressed in a dedicated “plain english” constraints language.

One feature that sets *provGen* apart from existing approaches to graph generation is that it provides users with local control over topological features and statistical characteristics of the graph. Constraints are evaluated locally for each node created, thus avoiding the complexity of verifying them globally. *provGen* is implemented using a Neo4J graph database back end. Graph rewrite rules and

¹¹ ProvBench’2013 CFP: <http://bit.ly/1fB0swR>

user constraints are both mapped to Cypher queries. Rewrite rules are mapped to `CREATE` clauses, while constraints are compiled into `WHERE` clauses. The two are blended together into complete Cypher queries, so that graph generation relies entirely on Neo4J’s native query engine.

We have also briefly discussed our approach to evaluating the effectiveness of *provGen* in generating “real-world” provenance, i.e., by comparing some of its key statistical properties with those of real graphs within the same class. We are currently experimenting with a variety of seed graph patterns, and more extensively evaluating *provGen*’s capability to mimick real provenance. Currently seed patterns must be manually designed or discerned. In future, an attempt to collate a collection of patterns common to provenance data, as has been done with workflow specifications [VTKB03], could prove useful.

Graph generation performance is another concern we are currently addressing. Generating large scale graphs requires efficient execution of the `MATCH-CREATE-WHERE` queries shown above, on graphs of increasing size. We are finding that Neo4J may not be an optimal choice, as it is geared for OLTP workloads with consequent transaction management overhead. However, our architecture is flexible and allows for experimentation, as changing the back end simply requires retargeting the mapping of rules and constraints to a different query language.

References

- BA99. AL Barabási and R Albert. Emergence of scaling in random networks. *science*, pages 1–11, 1999.
- BB05. Vladimir Batagelj and Ulrik Brandes. Efficient generation of large random networks. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 71(3 Pt 2A):036113, March 2005.
- CAB⁺13. Artem Chebotko, John Abraham, Pearl Brazier, Anthony Piazza, Andrey Kashlev, and Shiyong Lu. Storing, Indexing and Querying Large Provenance Data Sets as RDF Graphs in Apache HBase. In *2013 IEEE Ninth World Congress on Services*, pages 1–8. IEEE, June 2013.
- CF06. Deepayan Chakrabarti and Christos Faloutsos. Graph mining. *ACM Computing Surveys*, 38(1):2–es, June 2006.
- CLFF10. Artem Chebotko, Shiyong Lu, Xubo Fei, and Farshad Fotouhi. RDFProv: A relational RDF store for querying and managing scientific workflow provenance. *Data & Knowledge Engineering*, 69(8):836–865, August 2010.
- CMM12. James Cheney, Paolo Missier, and Luc Moreau. Constraints of the Provenance Data Model. Technical report, 2012.
- ddOOn⁺12. João Carlos de A. R. Gonçalves, Daniel de Oliveira, Kary A C S Ocaña, Eduardo S Ogasawara, and Marta Mattoso. Using Domain-Specific Data to Enhance Scientific Workflow Steering Queries. In *Procs. IPAW’12*, pages 152–167, Santa Barbara, California, 2012. Springer.
- DMV⁺13. Tom De Nies, Sara Magliacane, Ruben Verborgh, Sam Coppens, Paul T Groth, Erik Mannens, and Rik Van de Walle. Git2PROV: Exposing Version Control System Content as W3C PROV. In *Poster and Demo Proceedings of the 12th International Semantic Web Conference*, pages 125–128, 2013.

- ER60. P Erdős and A Rényi. On the Evolution of Random Graphs. In *Publication of the Mathematical Institute of the Hungarian academy of sciences*, pages 17–61. 1960.
- KK04. M. Kuramochi and G. Karypis. An efficient algorithm for discovering frequent subgraphs. *IEEE Transactions on Knowledge and Data Engineering*, 16(9):1038–1051, September 2004.
- KN09. Brian Karrer and M. E. J. Newman. Random graph models for directed acyclic networks. *Physical Review E*, 80(4):046110, October 2009.
- LCKF05. Jurij Leskovec, Deepayan Chakrabarti, Jon Kleinberg, and Christos Faloutsos. Realistic, mathematically tractable graph generation and evolution, using Kronecker multiplication. In *Proceedings of the 9th European conference on Principles and Practice of Knowledge Discovery in Databases*, pages 133–145. Springer-Verlag Berlin, Heidelberg, 2005.
- LLCF11. Chunhyeok Lim, Shiyong Lu, Artem Chebotko, and Farshad Fotouhi. OPQL: A First OPM-Level Query Language for Scientific Workflow Provenance. In *2011 IEEE International Conference on Services Computing*, pages 136–143. IEEE, July 2011.
- MMB⁺12. Luc Moreau, Paolo Missier, Khalid Belhajjame, Reza B’Far, James Cheney, Sam Coppens, Stephen Cresswell, Yolanda Gil, Paul Groth, Graham Klyne, Timothy Lebo, Jim McCusker, Simon Miles, James Myers, Satya Sahoo, and Curt Tilmes. PROV-DM: The PROV Data Model. Technical report, World Wide Web Consortium, 2012.
- MMCSR12. Luc Moreau, Paolo Missier, James Cheney, and Stian Soiland-Reyes. PROV-N: The Provenance Notation. Technical report, 2012.
- MMG⁺07. Alan Mislove, Massimiliano Marcon, Krishna P. Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement - IMC ’07*, page 29, New York, New York, USA, 2007. ACM Press.
- MMW⁺12. Anderson Marinho, Leonardo Murta, Cláudia Werner, Vanessa Braganholo, Sérgio Manuel Serra da Cruz, Eduardo Ogasawara, and Marta Mattoso. ProvManager: a provenance management system for scientific workflows. *Concurrency and Computation: Practice and Experience*, 24(13):1513–1530, 2012.
- MOnH⁺13. Marta Mattoso, Kary Ocaña, Felipe Horta, Jonas Dias, Eduardo Ogasawara, Vitor Silva, Daniel de Oliveira, Flavio Costa, and Igor Araújo. User-steering of HPC Workflows: State-of-the-art and Future Directions. In *Proceedings of the 2Nd ACM SIGMOD Workshop on Scalable Workflow Execution Engines and Technologies, SWEET ’13*, pages 4:1—4:6, New York, NY, USA, 2013. ACM.
- PBE13. Minh-Duc Pham, Peter Boncz, and Orri Erling. S3g2: A scalable structure-correlated social graph generator. In Raghunath Nambiar and Meikel Poess, editors, *Selected Topics in Performance Evaluation and Benchmarking*, pages 156–172. Springer Berlin Heidelberg, 2013.
- VTKB03. W.M.P Van der Aalst, A.H.M Ter Hofstede, B Kiepuszewski, and A.P Barros. Workflow patterns. *Distributed and Parallel Databases*, 14(1):5–51, 2003.