

Scalable Graph Query Evaluation and Benchmarking with Realistic Models

Gábor Szárnyas

Budapest University of Technology and Economics
Department of Measurement and Information Systems
MTA-BME Lendület Research Group on Cyber-Physical Systems
szarnyas@mit.bme.hu

Abstract

Model queries are widely used in model-driven engineering toolchains: models are checked for errors with validation queries, model simulations and transformations require complex pattern matching, while injective mappings for views are defined with model queries. Efficient and scalable evaluation of complex queries on large models is a challenging task. To achieve scalable graph query evaluation, I identified key challenges such as the lack of credible benchmarks and difficulties of obtaining real models for performance testing. To address these challenges, my contributions target (1) distributed incremental graph queries, (2) a cross-technology benchmark for model validation, (3) characterization of realistic models, and (4) realistic models generation.

Keywords distributed queries, model validation, model generation, benchmarking

1. Problem and Motivation

Model-Driven Engineering (MDE) is a development methodology used in many application domains such as critical applications (automotive, avionics and railway systems [7, 31, 58]). To increase the efficiency of development, MDE facilitates the use of models in various modelling languages targeting different levels of abstraction. Models can be used not only for presenting the structure and behaviour of the system, but also for synthesizing various design artifacts (such as source code, configuration files, documentation). To catch design flaws early, *model validation* techniques check the *well-formedness* of models. Design rules and *well-formedness constraints* are often captured in the form of *graph patterns* [9] to highlight invalid model ele-

ments to systems engineers. MDE tools check these patterns by evaluating *graph queries*.¹

1.1 Scalable Graph Queries

As models are rapidly increasing in size and complexity, efficient execution of model validation operations is challenging for the currently available toolchains, like ARTOP [2], Capella [38] or Papyrus [53].

The last decade brought considerable improvements in distributed storage and query technologies, known as NoSQL systems. These systems provide quick evaluation of simple retrieval operations and they are able to answer complex queries in a scalable manner, albeit not instantly. Providing quick response times for evaluating such queries over large and evolving data sets is still a challenging task.

Graph queries capturing validation constraints are often complex, including many join, antijoin and filtering operations. However, most query technologies cannot efficiently evaluate such operations for models with 10 million model elements [48], while models of critical systems, software and geospatial models are often 1–2 orders of magnitude larger [43]. A possible solution for scalable graph queries is to use distributed query processing techniques [16, 59]. This brings us to the first research question I investigated.

RQ 1. *How to incrementally evaluate graph queries over a distributed platform?*

1.2 Benchmarking

To assess the performance of a graph query engine, a benchmark framework is of high importance. According to the *Benchmark Handbook* [21], a useful benchmark is (1) relevant, (2) portable, (3) scalable, and (4) simple. To ensure *relevance*, the benchmark must use a *representative* workload and models similar to realistic ones. Providing relevant results, while also guaranteeing the other three properties (portability, scalability and simplicity) is a major challenge.

[Copyright notice will appear here once 'preprint' option is removed.]

¹ In this paper, I use the term *graph* as a synonym for *instance model*.

For real-world industrial systems, both metamodels and instance models are protected by intellectual property rights (IPR). For example, AUTOSAR [7] is not an open standard, but only available to members of the consortium, therefore it is not suitable for an open performance benchmark. Similarly, engineering models in the avionics and railway domains are also not available to the public.

These challenges confirm the need for a benchmark framework, which provides a real-world-like workload scenario and evaluates realistic queries on realistic models. Therefore, the second research question is the following.

RQ 2. *How to assess query technologies for a continuous model validation scenario?*

1.3 Characterization of Realistic Models

While existing generators may produce large models in increasing sizes, these models are usually simple and synthetic, which hinders their credibility for industrial and research benchmarking purposes. Up to my best knowledge, there are no existing techniques to characterize models used in MDE practice. To develop such a technique, first I had to address questions about model metrics, such as:

- Which metrics can be used for characterizing models?
- Is it possible to distinguish models of different domains, purely based on their metrics?

To answer these questions, I conducted a literature review in other disciplines, e.g. network theory and social network analysis. The high-level goal of the research is to answer the following question.

RQ 3. *What makes a model realistic?*

1.4 Generating Realistic Models

Custom generators of graph-based models are used in MDE for many purposes such as functional testing and performance benchmarking of modeling environments to ensure the correctness and scalability of tools. However, none is capable of generating realistic models scalable in size:

- Logic-based synthesis (like Alloy [25]) generate well-formed models but lack scalability.
- Rule-based approaches [48] are capable of generating large models by using transformation rules or random mutations to add new elements. However, they provide no guarantees that the resulting model is realistic. Some approaches do not even guarantee well-formedness, which is a prerequisite for realistic models.

It is an open research question if it is possible to ensure these properties.

RQ 4. *How to generate scalable and realistic models?*

2. Preliminaries

This section introduces an example used throughout the paper and presents the concept of *incremental queries*.

2.1 Running Example: Railway Network

As a running example, I use a small railway network, defined on the metamodel of the Train Benchmark [48], a model validation benchmark (the benchmark and my related contributions are discussed in Section 4.2).²

Figure 1 shows a schematic representation of the network, with routes (1–3), switches and segments. As the first switch is set to a straight position and the second switch is set to a diverging position, a train passing through this track would follow route #3, hence that route *active*.

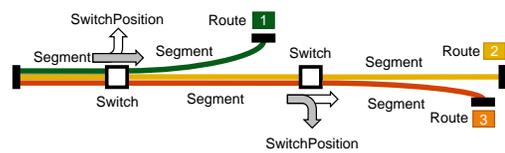


Figure 1. Railway example model. The positions of the switches designate route #3 as active.

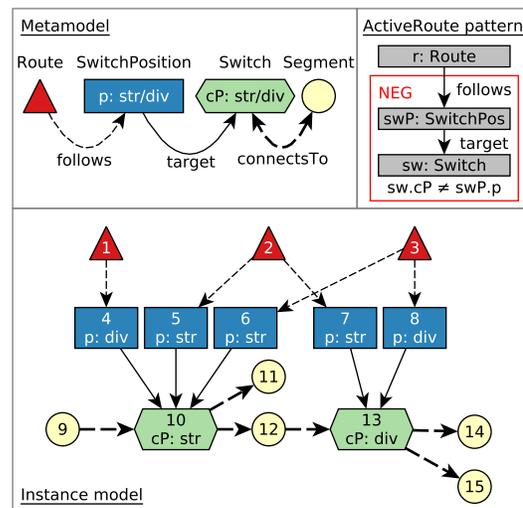


Figure 2. The graph of the railway example in Figure 1. The metamodel is shown in the top left corner of the figure. The graph pattern for finding the active routes is shown in the top right corner.

Modeling tools often represent their models as graphs. Figure 2 shows the example network as a labelled, attributed graph, along with the metamodel of the graph. Routes follow a set of *switch positions* that contain the prescribed *position* (straight or diverging) of the switch. The railway track consists of connected switches and segments.

²To guarantee that the example is concise and easy to understand, the example only uses a fraction of the Train Benchmark metamodel. The benchmark uses models that are significantly more complex: they contain more metamodel elements (types) and consist of more elements (objects).

The *active route* can be determined by evaluating a *graph query* (by *graph pattern matching*). A route is active if all its switches are in the position prescribed by the switch positions of the route. In other words, a route is active if none of its switches are set to a different position as the prescribed position. This results in the pattern shown in the upper right corner of Figure 2. In the example, the graph query selects route #3 as the active one, as both its switch positions (6 and 8) are satisfied by the corresponding switches (10 and 13).

2.2 Incremental Query Evaluation

In many use cases, queries are continuously evaluated, while changes affect only a restricted part of data. The queries and transformations for simulation and well-formedness validation in MDE are typical examples of such a workload. The goal of *incremental query evaluation* is to speed up such queries, utilizing the (partial) results obtained during the previous executions of the query to compute the latest set of changes. For example, if the *current position* of the second switch in Figure 1 changes from diverging to straight, the change only affects a small part of the graph (node 13 in Figure 2). This allows the incremental query engine to quickly reevaluate the query: in this case, the *active route* is changed from #3 to #2.

Incremental query evaluation algorithms use additional data structures for caching interim results, hence they consume more memory than search-based, non-incremental algorithms. In other words, they trade memory consumption for execution speed. While incremental query engines provide quick response times for various use cases [9, 48], their excessive memory consumption limits their scalability.

3. Related Work

To appropriately address all the research questions in the context of MDE, a wide range of multidisciplinary topics needs to be covered.

Distributed incremental graph queries. The Rete algorithm was originally created by Charles Forgy for rule-based expert systems [20]. Bunke et al. [14] were the first to propose the Rete algorithm in the context of graph transformations. Bergmann et al. adapted the algorithm for the Eclipse Modeling Framework in the EMF-INCQUERY project [9], now part of the VIATRA project [55].

Query languages and execution engines have been developed to support *incremental graph queries* on a single-machine environment. Drools [27] is an incremental business rule engine for Java-based systems. INSTANS [40] provides incremental queries over RDF [57].

Various distributed, but non-incremental graph query systems exist, including an approach based on SAP HANA [29], a graph transformation tool using the Bulk Synchronous Parallel graph processing model [30], and Trinity, an RDF-based query engine [45].

Cross-technology benchmark for continuous validation.

Numerous benchmarks have been proposed to compare the performance of query and transformation engines, but no openly available cross-technology benchmarks exist for continuous model validation.

The first transformation benchmark was proposed in [56], which gave an overview of typical application scenarios of graph transformations together with their characteristic features. Many transformation challenges have been proposed as cases for graph and model transformation contests. However, only [22, 61] focus on query performance, while others measure the usability of the tools, the conciseness and readability of the query languages and tests various advanced features, including reflection, traceability, etc.

There are numerous benchmarks from the area of semantic databases. *SP²Bench* [44] features a synthetic DBLP-like dataset, the *Berlin SPARQL Benchmark (BSBM)* [11] simulates an e-commerce application, while the *DBpedia SPARQL benchmark* [35] features a real data set with queries based on real-world user queries. The Linked Data Benchmark Council (LDBC) recently developed the Social Network Benchmark [19], a cross-technology benchmark, which provides an interactive workload and focuses on navigational pattern matching (i.e. traversal operations). While some of these benchmarks feature update operations and hence measure incremental query performance, they provide workloads that significantly differ from MDE use cases.

Characterization of realistic models. Revealing essential structural similarities and differentiations among networks from different fields is a fundamental objective in network theory with a wide range of applications. The authors of [15] list 22 areas using network theory, including social network analysis, transportation, biomolecular networks and chemistry. Network theory is also studied in physics, e.g. in the context of statistical mechanics [5]. However, most of these applications use untyped (one-dimensional) networks. So far, existing multidimensional studies only focused on models of a single application domain, such as neighbourhood and centrality analysis of a social network [12], relevance and correlation analysis of different dimensions in Flickr [28], community detection in the network of YouTube [52].

The authors of [10] use graph metrics to capture the structure and evolution of software products and processes in order to detect significant structural changes, help estimate bug severity, prioritize debugging efforts, and predict defect-prone releases in software engineering. Metrics are also used for understanding the main characteristics of domain-specific metamodels, to study model transformations with respect to the corresponding metamodels, and search correlations between them via analytical measures [41].

Realistic model generation. The *SP²Bench* [44] benchmark uses a generator based on the statistics of the DBLP

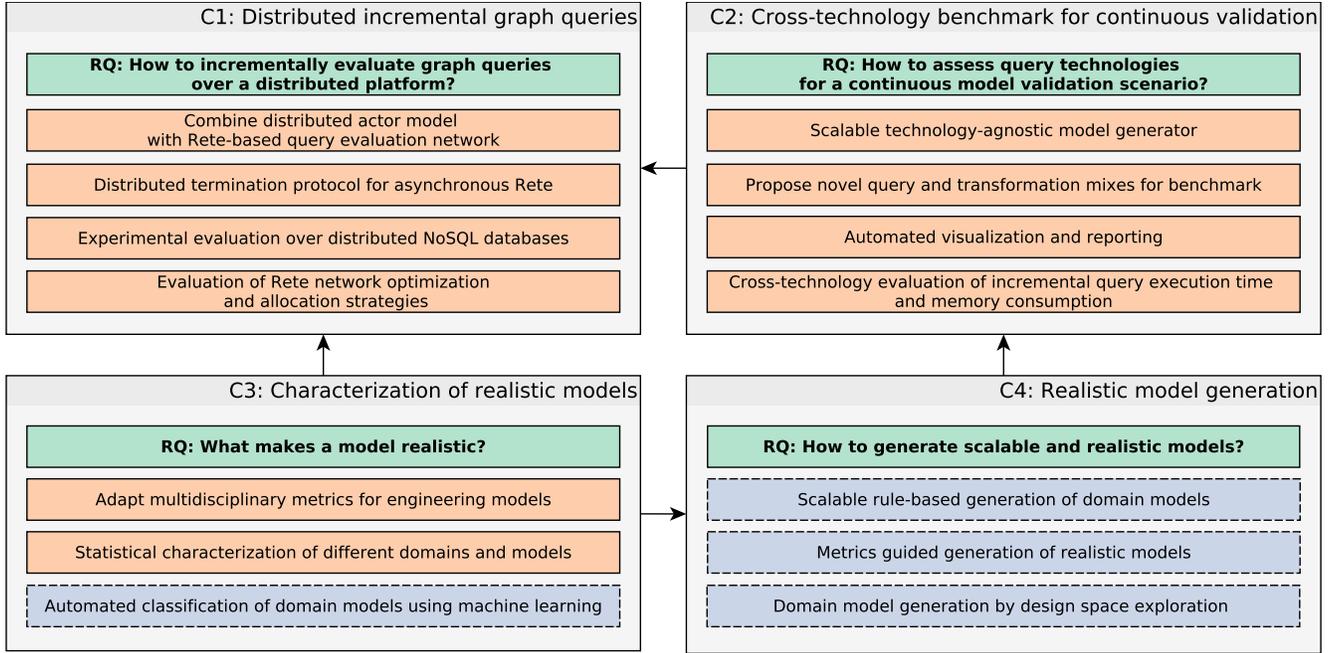


Figure 3. Graph of the contributions. Research questions are typeset in bold. Proposed contributions are noted with dashed border. The arrows indicate the relationship between contributions, e.g. the results for C3 could be directly used for C1 and C4.

library. The authors of [36] use *Boltzmann samplers* [17] to ensure efficient generation of uniform models.

OMOGEN [13] is a tool for automatically generating test models, used for testing model transformations. The tool takes a metamodel and a set of model fragments as its inputs and combines the fragments using several strategies to build valid instances.

gMark [8] is a domain-independent framework for synthesizing large graphs, allowing the user to specify parameters – size, types, degree distributions and other constraints – for the graphs to be generated. gMark is also able to generate query workloads with queries of different size, shape and selectivity.

4. Approach and Contributions

Figure 3 summarizes key the research questions and contributions of my research. This section presents my approach along with achieved and proposed contributions.

4.1 Distributed incremental graph queries

To achieve scalable incremental query evaluation, I adapted the Rete algorithm for distributed systems. I demonstrate the Rete algorithm works on the ActiveRoute query (Figure 2). As described in Section 2.1, the query collect Routes, where all Switches along the route are in the position prescribed by the corresponding SwitchPosition. In other words, without using the universal quantifier (\forall), it searches for routes that do *not* have a SwitchPosition which prescribes a position different from the current position of its target Switch [39].

Hence, the query can be formalized in relational algebra as:³

$$\text{route} \triangleright (\text{follows} \bowtie \sigma_{\text{currentPosition} \neq \text{position}} (\text{switch} \bowtie \text{target} \bowtie \text{switchPosition})) = \{\langle 3 \rangle\}$$

where \bowtie denotes the natural join operator that joins its operands based on their common attributes, and \triangleright denotes the antijoin operator (also known as the anti-semijoin [46]) that keeps the tuples from its left operand which do not have a matching tuples in its right operand.

Figure 4 shows a distributed Rete network implementing this relational algebra expression. The network is allocated to two machines, Server 1 and Server 2. This allows the query engine to scale for larger graphs, for which the Rete network would not fit in the memory of a single workstation. However, this approach still has a bottleneck limiting scalability: if a Rete node cannot fit to the memory of a single workstation, it will run out of memory.

³To formalize the query, the relations for the vertices and edges in Figure 2 can be defined as follows:

- $\text{route}(\text{route}) = \{\langle 1 \rangle, \langle 2 \rangle, \langle 3 \rangle\}$
- $\text{follows}(\text{route}, \text{switchPosition}) = \{\langle 1, 4 \rangle, \langle 2, 5 \rangle, \langle 2, 7 \rangle, \langle 3, 6 \rangle, \langle 3, 8 \rangle\}$
- $\text{switch}(\text{switch}, \text{currentPosition}) = \{\langle 4, \text{div} \rangle, \langle 5, \text{str} \rangle, \langle 6, \text{str} \rangle, \langle 7, \text{str} \rangle, \langle 8, \text{div} \rangle\}$
- $\text{target}(\text{switchPosition}, \text{switch}) = \{\langle 4, 10 \rangle, \langle 5, 10 \rangle, \langle 6, 10 \rangle, \langle 7, 13 \rangle, \langle 8, 13 \rangle\}$
- $\text{switchPosition}(\text{switchPosition}, \text{position}) = \{\langle 10, \text{str} \rangle, \langle 13, \text{div} \rangle\}$

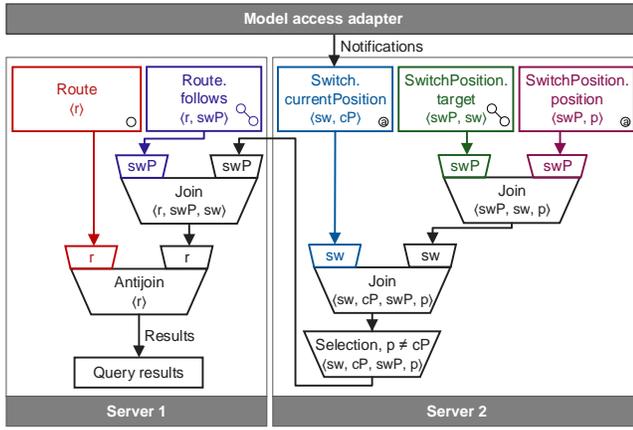


Figure 4. Rete network for the ActiveRoute pattern.

Using these techniques and algorithms, I made following contributions.

Combine distributed actor model with Rete-based query evaluation network. I designed a distributed architecture and prototyped INCQUERY-D, a Rete-based query engine using actors for distributed scalability. I presented a detailed performance evaluation in the context of model incremental well-formedness validation. The results showed nearly instantaneous complex query reevaluation well beyond 10M+ model elements, [47]. To further extend the scalability of the system, I proposed sharding individual Rete nodes in [32].

Distributed termination protocol for asynchronous Rete. As Rete is an asynchronous algorithm, determining if the network is in a consistent state w.r.t. the latest change set requires a distributed *termination protocol*. The protocol was also presented in [47] and [32].

Experimental evaluation over distributed NoSQL databases. The proposed architecture and algorithms are *representation-agnostic*. They have been integrated with the Neo4j graph database [23], the Titan distributed graph database and 4store, a semantic database [47].

Evaluation of Rete network optimization and allocation strategies. Allocating the Rete nodes in the cloud is a complex optimization problem, where the goal is to minimize the cost of communication between the nodes. I presented a solver-based approach for allocating Rete nodes in [33]. I also proposed optimization techniques used in relational query optimization for enhancing the performance of graph queries [50].

Uniqueness. Up to my best knowledge, existing technologies are either distributed [30, 45] or incremental [55], but there is no system that provides scalable, distributed incremental graph queries.

4.2 Cross-technology benchmark for continuous validation

In Section 2.1, I used a running example from the Train Benchmark framework. The Train Benchmark is an incremental model validation benchmark, continuously developed by the Fault-Tolerant Systems Research Group since 2010. I have significantly extended the Train Benchmark, both conceptually and implementation-wise. Figure 5 shows the inputs of the benchmark process, the benchmark phases and the benchmark results.

The Train Benchmark is a macro benchmark that aims to measure the performance of continuous model validation with graph-based models and constraints captured as queries. The benchmark is *cross-technology*, i.e. it is implemented on a range technologies. The serialization formats include Eclipse-based model-driven engineering toolchains (EMF), graph databases [42], relational databases (SQL) and semantic technologies (RDF [57]). The query engines include relational engines (SQLite, MySQL), graph transformation frameworks (VIATRA [55]), rule engines (Drools [27]), graph query engines (Neo4j [37]) and SPARQL engines (Sesame [3], Jena [1]). Also, the framework is extensible which allows users of the benchmark to incorporate new technologies.

Earlier versions of the benchmark have been continuously used for performance measurements since 2012 [47, 54]. The benchmark is also part of the benchmark suite used by the MONDO EU FP7 [34] project and was selected as a case for the 2015 Transformation Tool Contest [51] as well. The benchmark framework is available as an open-source project.⁴

Scalable technology-agnostic model generator. While the original benchmark framework included a model generator, its scalability was limited. I redesigned the model generator focusing on two aspects: (1) ensuring scalability for large models, and (2) allowing the framework users to easily adapt new representations.

Propose novel query and transformation mixes for benchmark. The workload profile of the benchmark simulates *real-world model validation scenarios* of users loading, validating and transforming their models. The transformations capture user edits and quick-fix like automated refactoring operations. Some queries in the benchmark are structurally similar to AUTOSAR [7] validation queries (presented in [9]), while other aim to test various features of graph query engines (such as efficient filtering and evaluation of negative conditions).

Automated visualization and reporting. The framework features end-to-end automation [24] to (1) set up configurations of benchmark runs, (2) generate large model instances

⁴<https://github.com/FTRSrg/trainbenchmark>

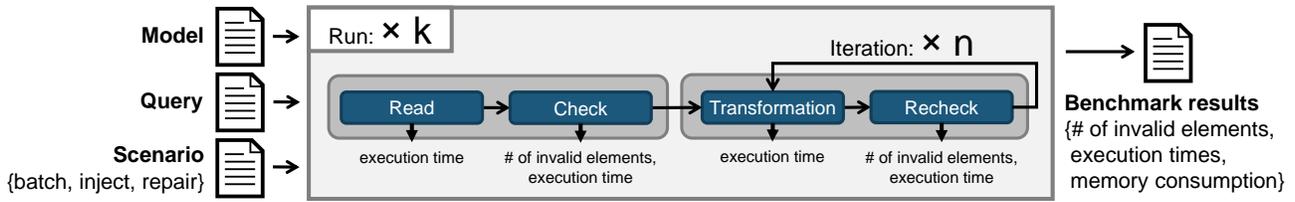


Figure 5. Phases of the Train Benchmark.

(3) execute benchmark measurements, (4) synthesize diagrams for measurements using R scripts⁵.

Cross-technology evaluation of incremental query execution time and memory consumption. This *cross-technology benchmark* can be adapted to different model representation formats and query technologies. This is demonstrated by 12+ reference implementations over four different technological spaces (EMF, graph databases, RDF and SQL) presented in [48].

Uniqueness. Compared to other benchmarks, the Train Benchmark has the following set of distinguishing features:

- The workload profile follows a real-world model validation scenario by updating the model with changes derived by simulated user edits or transformations.
- The benchmark measures the performance of both initial validation and incremental revalidation.
- This benchmark was designed with cross-technology adaptations in mind. It can be implemented with different model representation formats and query technologies.

4.3 Characterization of realistic models

In [49], I presented multidisciplinary graph metrics and evaluated them on instance models from different domains. As a result, I proposed some metrics which turned out to be useful for characterizing the structure of models.

Adapt multidisciplinary metrics for engineering models.

I performed a literature review and identified several graph metrics from other disciplines. For evaluating these metrics, I gathered instance models from software and systems engineering domains:

- AutoFOCUS system models [6],
- Building Information Models (BIM) [18],
- Capella system models [38],
- JaMoPP code models [26],
- railway models from the Train Benchmark [48],
- Yakindu [60] statecharts.

Statistical characterization of different domains and models. I used both exploratory and confirmatory data analysis techniques in order to determine the “usefulness” of metrics.

⁵<https://www.r-project.org/>

I considered a metric *useful* if it separates models of different domains from each other, while provides similar values for models within the same domain. I also investigated whether some of these metrics can *distinguish real models from auto-generated synthetic ones*.

Exploratory analysis relied on data visualization, while confirmatory analysis used statistical methods (such as performing Kolmogorov–Smirnov tests on the derived metrics distributions). My initial finding is that different versions of clustering coefficients (i.e. how tightly connected the model elements are) were particularly useful for such classifications. But, unsurprisingly, no single metric was able to sufficiently handle all the domains. The analysis also provides some insights that needs to be considered in future model generators to synthesize realistic models.

Automated classification of domain models using machine learning. As a *future research objective*, I plan to use *machine learning* techniques for automated classification of domain models.

Uniqueness. Up to my best knowledge, this is the first investigation for using multidimensional graph metrics for both *characterizing the realism* of models and *distinguishing different domain models* from each other.

4.4 Realistic model generation

As a *proposed contribution*, I plan to design and develop a generator that is capable of producing realistic models scalable in size. While there are solutions for generating either scalable or realistic models, there are no known approaches for the combination of both, rendering this a high-risk research task. The long-term research objective of generating scalable and realistic models breaks down to the following steps:

1. metrics guided generation of realistic models,
2. domain model generation by design space exploration [4],
3. scalable rule-based generation of domain models.

Acknowledgments

I would like to thank my advisor, Dániel Varró for his guidance during my research. I would also like to express my gratitude to István Ráth, Gábor Bergmann and numerous colleagues an co-authors for sharing their suggestions and ideas.

References

- [1] Apache Jena. <http://jena.apache.org/>.
- [2] Artop: The AUTOSAR Tool Platform. <https://www.artop.org/>.
- [3] Sesame: RDF API and query engine. <http://www.openrdf.org/>.
- [4] H. Abdeen, D. Varró, H. A. Sahraoui, A. S. Nagy, C. Debreceni, Á. Hegedűs, and Á. Horváth. Multi-objective optimization in rule-based design space exploration. In *ASE*, pages 289–300, 2014.
- [5] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74(1):47–97, January 2002.
- [6] V. Aravantinos et al. AutoFOCUS 3: Tooling concepts for seamless, model-based development of embedded systems. In *Joint Proceedings of ACES-MB & WUCOR co-located with MODELS*, pages 19–26, 2015.
- [7] AUTOSAR Consortium. *The AUTOSAR Standard*. <http://www.autosar.org/>.
- [8] G. Bagan, A. Bonifati, R. Ciucanu, G. H. L. Fletcher, A. Lemay, and N. Advokaat. Generating flexible workloads for graph databases. *Proc. VLDB Endow.*, 9(13):1457–1460, Sept. 2016.
- [9] G. Bergmann, Á. Horváth, I. Ráth, D. Varró, A. Balogh, Z. Balogh, and A. Ökrös. Incremental evaluation of model queries over EMF models. In *MODELS*, pages 76–90, 2010.
- [10] P. Bhattacharya, M. Iliofotou, I. Neamtii, and M. Faloutsos. Graph-based analysis and prediction for software evolution. In *ICSE*, pages 419–429, 2012.
- [11] C. Bizer and A. Schultz. The Berlin SPARQL benchmark. *International Journal on Semantic Web & Information Systems*, 5(2):1–24, 2009.
- [12] P. Bródka, P. Kazienko, K. Musial, and K. Skibicki. Analysis of neighbourhoods in multi-layered dynamic social networks. *Int. J. Computational Intelligence Systems*, 2012.
- [13] E. Brottier, F. Fleurey, J. Steel, B. Baudry, and Y. L. Traon. Metamodel-based test generation for model transformations: an algorithm and a tool. In *ISSRE*, pages 85–94, 2006.
- [14] H. Bunke, T. Glauser, and T. Tran. An efficient implementation of graph grammars based on the RETE matching algorithm. In *Graph-Grammars and Their Application to Computer Science, 4th International Workshop, Bremen, Germany, March 5-9, 1990, Proceedings*, pages 174–189, 1990.
- [15] L. d. F. Costa, O. N. Oliveira, G. Travieso, F. A. Rodrigues, P. R. Villas Boas, L. Antiquiera, M. P. Viana, and L. E. Correa Rocha. Analyzing and modeling real-world phenomena with complex networks: a survey of applications. *Advances in Physics*, 60(3):329–412, 2011.
- [16] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. In *OSDI*, pages 137–150, 2004.
- [17] P. Duchon, P. Flajolet, G. Louchard, and G. Schaeffer. Boltzmann samplers for the random generation of combinatorial structures. *Combinatorics, Probability & Computing*, 13(4-5):577–625, 2004.
- [18] C. Eastman, P. Teicholz, R. Sacks, and K. Liston. *BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Designers, Engineers and Contractors*. Wiley Publishing, 2008.
- [19] O. Erling, A. Averbuch, J. Larriba-Pey, H. Chafi, A. Gubichev, A. Prat-Pérez, M. Pham, and P. A. Boncz. The LDBC social network benchmark: Interactive workload. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, pages 619–630, 2015.
- [20] C. Forgy. Rete: A fast algorithm for the many patterns/many objects match problem. *Artif. Intell.*, 19(1):17–37, 1982.
- [21] J. Gray, editor. *The Benchmark Handbook for Database and Transaction Systems (2nd Edition)*. Morgan Kaufmann, 1993.
- [22] T. Horn, C. Krause, and M. Tichy. The TTC 2014 movie database case. *TTC 2014*, page 93, 2014.
- [23] B. Izsó, G. Szárnyas, I. Ráth, and D. Varró. IncQuery-D: Incremental graph search in the cloud. In *BigMDE*, 2013.
- [24] B. Izsó, G. Szárnyas, I. Ráth, and D. Varró. MONDO-SAM: A framework to systematically assess MDE scalability. In *BigMDE@STAF*, pages 40–43, 2014.
- [25] D. Jackson. Alloy: a lightweight object modelling notation. *ACM Trans. Softw. Eng. Methodol.*, 11(2):256–290, 2002.
- [26] JaMoPP. *The Java Model Parser and Printer*, 2016. <http://www.jamopp.org/index.php/JaMoPP>.
- [27] JBoss. Drools. <http://www.jboss.org/drools>.
- [28] P. Kazienko, K. Musial, and T. Kajdanowicz. Multidimensional social network in the social recommender system. *IEEE Trans. Systems, Man, and Cybernetics*, 41(4):746–759, 2011.
- [29] C. Krause, D. Johannsen, R. Deeb, K. Sattler, D. Knacker, and A. Niadzelka. An SQL-based query language and engine for graph pattern matching. In *ICGT*, 2016.
- [30] C. Krause, M. Tichy, and H. Giese. Implementing graph transformations in the bulk synchronous parallel model. In *FASE*. 2014.
- [31] B. Luteberget, C. Johansen, and M. Steffen. Rule-based consistency checking of railway infrastructure designs. In *Integrated Formal Methods - 12th International Conference, IFM 2016, Reykjavik, Iceland, June 1-5, 2016, Proceedings*, pages 491–507, 2016.
- [32] J. Maginecz and G. Szárnyas. Sharded joins for scalable incremental graph queries. In *23rd PhD Mini-Symposium*, Budapest University of Technology and Economics, 2016.
- [33] J. Makai, G. Szárnyas, I. Ráth, Á. Horváth, and D. Varró. Optimization of incremental queries in the cloud. In *3rd International Workshop on Model-Driven Engineering on and for the Cloud (CloudMDE) at MODELS*, 2015.
- [34] MONDO project. *Scalable Modeling and Model Management on the Cloud Project, 7th EU Framework Programme*, 2016.
- [35] M. Morsey, J. Lehmann, S. Auer, and A.-C. N. Ngomo. DBpedia SPARQL benchmark: Performance assessment with real queries on real data. In *ISWC*, 2011.
- [36] A. Mougnot, A. Darrasse, X. Blanc, and M. Soria. Uniform random generation of huge metamodel instances. In *ECMDA-FA*, pages 130–145, 2009.
- [37] Neo Technology. Neo4j. <http://neo4j.org/>.

- [38] PolarSys. *Capella*. <https://www.polarsys.org/capella/>.
- [39] A. Rensink. Representing first-order logic using graphs. In *ICGT*, pages 319–335, 2004.
- [40] M. Rinne, E. Nuutila, and S. Törmä. INSTANS: high-performance event processing with standard RDF and SPARQL. In *ISWC*, 2012.
- [41] J. D. Rocco, D. D. Ruscio, L. Iovino, and A. Pierantonio. Mining correlations of ATL model transformation and metamodel metrics. In *MiSE*, pages 54–59, 2015.
- [42] M. A. Rodriguez and P. Neubauer. Constructions from Dots and Lines. *Bulletin of American Society for Information Science & Technology*, August/September, 2010.
- [43] M. Scheidgen, A. Zubow, J. Fischer, and T. H. Kolbe. Automated and transparent model fragmentation for persisting large models. In *Model Driven Engineering Languages and Systems - 15th International Conference, MODELS 2012, Innsbruck, Austria, September 30-October 5, 2012. Proceedings*, pages 102–118, 2012.
- [44] M. Schmidt, T. Hornung, G. Lausen, and C. Pinkel. SP2Bench: A SPARQL performance benchmark. Shanghai, China, 2009. IEEE.
- [45] B. Shao, H. Wang, and Y. Li. Trinity: a distributed graph engine on a memory cloud. In *SIGMOD*, 2013.
- [46] A. Silberschatz, H. F. Korth, and S. Sudarshan. *Database System Concepts, 5th Edition*. McGraw-Hill Book Company, 2005.
- [47] G. Szárnyas, B. Izsó, I. Ráth, D. Harmath, G. Bergmann, and D. Varró. IncQuery-D: A distributed incremental model query framework in the cloud. In *MODELS*, pages 653–669, 2014.
- [48] G. Szárnyas, B. Izsó, I. Ráth, and D. Varró. The Train Benchmark: Cross-technology performance evaluation of continuous model validation. *Software and Systems Modeling*, 2017. Accepted.
- [49] G. Szárnyas, Z. Kóvári, A. Salánki, and D. Varró. Towards the characterization of realistic models: Evaluation of multidisciplinary graph metrics. In *MODELS*, pages 87–94, New York, NY, USA, 2016. ACM.
- [50] G. Szárnyas, J. Maginecz, and D. Varró. Evaluation of optimization strategies for incremental graph queries. *Periodica Polytechnica, EECS*, 2017. Accepted.
- [51] G. Szárnyas, O. Semeráth, I. Ráth, and D. Varró. The TTC 2015 Train Benchmark case for incremental model validation. In *Proceedings of the 8th TTC, a part of STAF*, 2015.
- [52] L. Tang, X. Wang, and H. Liu. Community detection via heterogeneous interaction analysis. *Data Min. Knowl. Discov.*, 25(1):1–33, 2012.
- [53] The Eclipse Foundation. *Papyrus*, 2015. <https://eclipse.org/papyrus/>.
- [54] Z. Ujhelyi, G. Bergmann, Á. Hegedüs, Á. Horváth, B. Izsó, I. Ráth, Z. Szatmári, and D. Varró. EMF-IncQuery: An integrated development environment for live model queries. *Sci. Comput. Program.*, 98:80–99, 2015.
- [55] D. Varró, G. Bergmann, Á. Hegedüs, Á. Horváth, I. Ráth, and Z. Ujhelyi. Road to a reactive and incremental model transformation platform: three generations of the VIATRA framework. *SOSYM*, 15(3):609–629, 2016.
- [56] G. Varró, A. Schürr, and D. Varró. Benchmarking for graph transformation. In *VL/HCC*. IEEE Press, 2005.
- [57] W3C. *Resource Description Framework (RDF)*. <http://www.w3.org/standards/techs/rdf/>.
- [58] J. Whittle, J. E. Hutchinson, and M. Rouncefield. The state of practice in model-driven engineering. *IEEE Software*, 31(3):79–85, 2014.
- [59] R. S. Xin, J. E. Gonzalez, M. J. Franklin, and I. Stoica. Graphx: a resilient distributed graph system on spark. In *GRADES co-located with SIGMOD/PODS*, page 2, 2013.
- [60] Yakindu. *Statechart Tools*. <http://statecharts.org/>.
- [61] A. Zündorf. AntWorld benchmark specification, GraBaTs, 2008.