

Supporting round-trip data migration for web APIs with Henshin

Daniel Strüber¹

¹Faculty of Science, Radboud University Nijmegen, Postbus 9010, 6500 GL Nijmegen, The Netherlands

Abstract

We present a solution to the Round-Trip Migration case of the Transformation Tool Contest 2020, based on the Henshin model transformation language. The task is to support four scenarios of transformations between two versions of the same data metamodel, a problem inspired by the application scenario of Web API migration, where such a round-trip migration methodology might mitigate drawbacks of the conventional “instant” migration style. Our solution relies on Henshin’s visual syntax, which seems well-suited to capture the problem on an intuitive level, since the syntax is already similar to the scenario illustrations in the case description. We discuss the five evaluation criteria *expressiveness*, *comprehensibility*, *bidirectionality*, *performance*, and *reusability*.

Keywords

data migration, model transformation, tool contest contribution

1. Introduction

During the evolution of web-based systems, provided APIs may change over time. Implementing API changes in an instant usually leads to severe complications, as developers of client software need to rapidly respond to such changes. A desirable alternative is a *round-trip migration system* that provides support for several API versions in parallel, while internally managing the data in such way that consistency is ensured.

To study how transformation tools may be an enabling technology for such a system, this problem is the subject of a case in the 2020 edition of TTC, called *Round-Trip Migration of Object-Oriented Data Model Instances* [1]. The case description includes four scenarios of changes that need to be supported by such a system: *create/delete field*, *rename field*, *declare field optional/mandatory*, and *multiple edits*. The task is to develop transformations capturing these changes between two versions M_1 , M_2 of the same metamodel. In each scenario, four transformations are required: the migration and subsequent back-migration for both possible round-trip directions (M_1 - M_2 - M_1 , M_2 - M_1 - M_2).

In this paper, we present a solution to the case based on the Henshin model transformation language [2, 3]. Henshin is a model transformation language that supports the declarative, graph-based specification of in-place transformations. The basic features of Henshin’s tool set are

a suite of editors and an interpreter kernel. It also offers a variety of advanced features, such code generation for parallel graph pattern matching and support for various transformation analyses.

Henshin provides an expressive visual syntax that aims to support usability during transformation development [3]. In fact, combining the object-diagram paradigm with change descriptions, Henshin’s syntax is similar to the informal notation used in the case description. Whereas the case description illustrates the change of specific instances, Henshin captures such changes in the form of reusable *rules*. Rules express basic match-and-change patterns. To specify control flow between rules where needed, Henshin provides *composite units* that orchestrate the execution of a number of sub-units and rules. From a variety of available units (including random execution and loops), our solution uses one: Sequential units, supporting the specification of rules in a given order as well as the included data flow.

2. Solution

The solution comprises a set of Henshin rules and units, as well some Java-based glue code. The reason for having the glue code is to establish the connection to the benchmark framework by implementing the provided task interface (constant over the different scenarios). The solution contains 8 modules (one per scenario and case), encapsulating rules and units. For scenarios 1–3, each module contains 2 rules, one for forward and one for back-ground migration. For scenarios 4, each module contains 2 units, each orchestrating 4 rules (one of them reused). The full solution and setup instructions are available at <https://github.com/dstrueber/ttc2020>.

TTC’20: Transformation Tool Contest, Part of the Software Technologies: Applications and Foundations (STAF) federated conferences, Eds. A. Boronat, A. García-Domínguez, G. Hinkel, and F. Krikava, 17 July 2020, Bergen, Norway (online).

✉ d.strueber@cs.ru.nl (D. Strüber)

🌐 <https://www.danielstrueber.de/> (D. Strüber)


🆔 0000-0002-5969-3521 (D. Strüber)

© 2021 Copyright for this paper by its authors. Use permitted under Creative


Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)





(a) Scenario 1, M_1 - M_2 - M_1 case.



(b) Scenario 1, M_2 - M_1 - M_2 case.

Figure 1: Full solution for Scenario 1: create/delete field.

2.1. Scenario 1: create/delete field.

The solution for scenario 1 comprises four rules, shown in Fig. 1: two for both directions, one for migration and one for back-migration. Each rule has a name and a parameter list; parameters have a name and one of the directions *in*, *out*, *var*. *In* and *out* parameters are passed into and out of the rule from the usage context (in our case, the Java-based glue code), *var* parameters (variables) are used internally to propagate information between different parts of the rule. Each rule comprises a number of *preserve* elements (shown in gray), specifying the context for creations, and *create* elements (shown in green), specifying newly created elements. For example, the top *migrate* rule specifies a person from metamodel 1 as context for the creation of a new person based on metamodel 2, obtaining the same name and the age value of -1. The *migrateBack* rule for the M_2 - M_1 - M_2 case requires a *trace* person as a source for the age value. The glue code for the previously applied *migrate* rule ensures that this object is available and can be passed into *migrateBack*.

A usability limitation in this particular scenario (involving two metamodels that, from the user's perspective, are actually versions of the same metamodel) is that there is no visual distinction between elements from different metamodels. This issue is mitigated by the use of names (e.g., *instance1* and *instance2*, referring to instances from M_1 and M_2) as well as the coloring of separate actions (preserve, create).

2.2. Scenario 2: rename field.

Figure 2 shows two out of the four rules for scenario 2, capturing the M_1 - M_2 - M_1 case (for symmetry reasons, the other two rules are identical). In attribute calculations, Henshin supports the use of JavaScript expressions, including standard library objects such as *Date*, as we use to specify the conversion between birth dates and age. During execution, JavaScript expressions are evaluated using the Nashorn JavaScript engine.

2.3. Scenario 3: declare field optional/mandatory.

Figure 3 shows two out of four rules for scenario 3, capturing the M_1 - M_2 - M_1 case. Again, we rely on Henshin's capability of using JavaScript calculations in attribute values. In the *migrateBack* rule, we use the ternary operator to check if the name of the provided person object is null, and specify different calculation outcomes for both cases. The M_2 - M_1 - M_2 case (not shown) is largely similar; the main difference is that the *migrate* rule includes a handling of empty strings, in the same way that the shown *migraterules* handles the *null* case.

2.4. Scenario 4: multiple edits.

While scenarios 1–3 are straightforward, scenario 4 has a design space of possible solutions. The rationale of our solution is to illustrate the reuse capabilities of Henshin, as requested in the case description. Still, we notice an inherent trade-off between reuse, simplicity and performance. The granularity of reuse in this case is so fine




Figure 2: Scenario 2: rename field (M_1 - M_2 - M_1 case. The rules for M_2 - M_1 - M_2 look the same).




Figure 3: Scenario 3: Declare field optional/mandatory (M_1 - M_2 - M_1 case. The rules for M_2 - M_1 - M_2 look similar).

that the price of reuse appears high: Compared to a solution without reuse (capturing the edits in a single rule), the specification is much larger and therefore harder to understand and less efficient to execute.

Figure 4 shows the implementation of the *migrate* step in the M_1 - M_2 - M_1 direction, comprising a sequential unit that invokes four rules – three of them presented with the unit, one of them reused from Scenario 2. (For reading convenience, the rule from Scenario 2 was actually copied and included into the Scenario 4 transformation file, but it would be possible to maintain an actual reference as well—using EMF’s concept of remote references.) The unit produces an instance of a M_2 container object from an instance of a M_1 container object, and also has four internal *var* parameters. The order of the rule invocations is specified in an activity-diagram-like notation. Data flow is specified by passing parameters between rule invocations (question marks refer to variables, which do not have to be set from the context). Rule *getObject1* fetches the person and dog objects from the container. Rules *task2_migrate* and *dog_migrate* produce the migrated objects. Rule *connectMigrated2* connects the migrated objects with a reference and encapsulates them in a container, which is yielded as output.

The solution for *migrateBack* and both steps in the M_2 - M_1 - M_2 direction look largely similar. A small additional sophistication is that *migrateBack* for M_1 - M_2 - M_1 requires a copy (a.k.a. trace) of the previously migrated dog instance, which we implement in a similar manner as in the M_2 - M_1 - M_2 case of scenario 1.

3. Evaluation

We discuss our solution in the light of the five evaluation criteria from the case description [1].

Expressiveness. The assessment of expressiveness is based on the number of passing test cases: the more test cases a solution covers, the more expressive it is. According to this criterion, the present solution is maximally expressive, since it passes all test cases.

Comprehensibility. With our solution, we aimed at providing a primarily declarative solution. We achieved this goal by specifying all parts of the change logic using Henshin’s declarative rule and control flow concepts. In addition, our solution includes some glue code, written in Java, to implement the provided *Task* interface.

Bidirectionality. Our solution (like any solution passing all test cases) is inherently bidirectional, since it supports transformations in both directions. Still, Henshin does not provide dedicated support for automatically generating the transformation in either direction, and hence, leaves the effort for bidirectionality implementation to the user. An extension of Henshin with dedicated support for bidirectionality exists in an early stage [4].

Performance. Figure 5 shows the execution time data from running the performance test on a dual-core i7-6600U CPU Ubuntu Linux system with 16 GB RAM. We observe a large slowdown, by two orders of magnitude, compared to the Java reference implementation. This

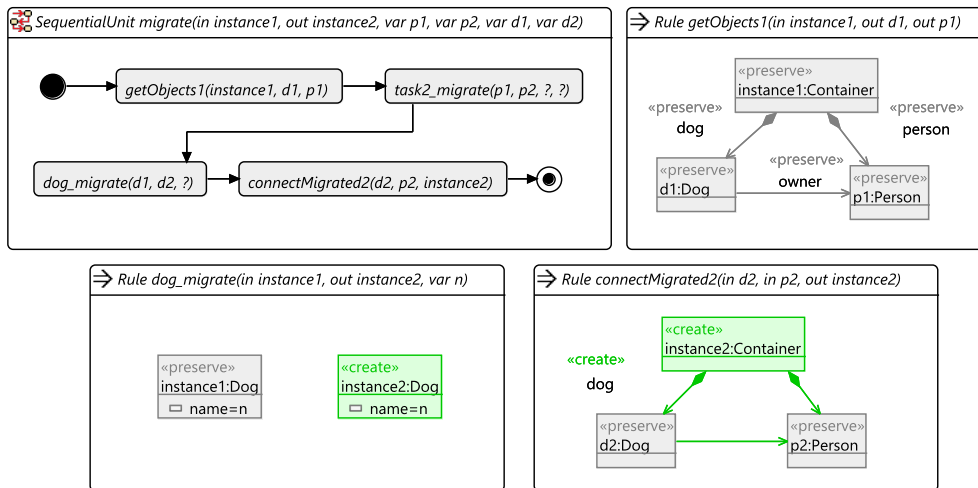


Figure 4: Scenario 4: Multiple edits; M1-M2-M1 migrate step (*migrate-back* and M2-M1-M2 look similar).




Figure 5: Runtime measurements

slowdown is to be expected due to the overhead of Henshin's generic matching engine, and the decision to use a deliberately inefficient implementation (in favor of reuse) in the most complicated scenario (scenario 4).

Re-usability. Henshin's composite units support the specification of a control flow. In scenario 4, we use this feature to reuse rules originally specified for the implementation of scenario 2. Doing so, we enable reuse by following a suggestion from the case description: "*delegating the migration of the referenced Person instances to migration functions which have been already defined.*" [1] Henshin has another reuse concept for specifying rule variants [5], but this concept was not applicable to the given case, since it does not support variations on the level of node types (dog vs. person). This experience emphasizes the need for more expressive reuse concepts for model transformations [6].

Furthermore, it can be emphasized that the Java-based glue code for implementing the task interface is constant over the different scenarios. Even though scenario 4 is the most complex scenario by far, we were able to fully reuse the glue code developed for scenarios 1–3 without further modifications.

4. Outlook

The trade-off related to reusability vs. simplicity/performance motivates further work on composition of transformation systems: The complicated specification of units and rules in scenario 4 could be translated into a single rule, which might both increase readability and execution performance. While there are various works on composition of rules from smaller parts [7, 8, 9], no existing work seems to consider control flow in the input transformation specification. The considered round-trip migration scenario is also relevant during the model-driven migration between different backend versions of content management systems [10].

Acknowledgement. Thanks to Antonio Garcia-Dominguez for providing an initial version of the performance comparison figure.

References

- [1] L. Beurer-Kellner, J. von Pilgrim, T. Kehrer, Round-trip migration of object-oriented data model instances, in: Transformation Tool Contest, 2020.

- [2] T. Arendt, E. Biermann, S. Jurack, C. Krause, G. Taentzer, Henshin: advanced concepts and tools for in-place emf model transformations, *Model Driven Engineering Languages and Systems (2010)* 121–135.
- [3] D. Strüber, K. Born, K. D. Gill, R. Groner, T. Kehrer, M. Ohrndorf, M. Tichy, Henshin: A usability-focused framework for emf model transformation development, in: *International Conference on Graph Transformation, 2017*, pp. 196–208. URL: <https://rgse.uni-koblenz.de/web/pages/research/papers/SBGGKOT17.pdf>.
- [4] C. Ermel, F. Hermann, J. Gall, D. Banzner, Visual modeling and analysis of emf model transformations based on triple graph grammars, *Electronic Communications of the EASST 54 (2012)*.
- [5] D. Strüber, J. Rubin, T. Arendt, M. Chechik, G. Taentzer, J. Plöger, Variability-based model transformation: formal foundation and application, *Formal Asp. Comput.* 30 (2018) 133–162.
- [6] M. Chechik, M. Famelis, R. Salay, D. Strüber, Perspectives of model transformation reuse, in: *International Conference on Integrated Formal Methods, 2016*, pp. 28–44.
- [7] A. Rensink, Compositionality in graph transformation, in: *International Colloquium on Automata, Languages and Programming, 2010*, pp. 309–320.
- [8] A. Anjorin, K. Saller, M. Lochau, A. Schürr, Modularizing triple graph grammars using rule refinement, in: *Fundamental Approaches to Software Engineering, 2014*, pp. 340–354.
- [9] L. Fritsche, J. Kosiol, A. Schürr, G. Taentzer, Shortcut rules, in: *STAF workshops, 2018*, pp. 415–430.
- [10] D. Priefer, P. Kneisel, D. Strüber, Iterative model-driven development of software extensions for web content management systems, in: *European Conference on Modelling Foundations and Applications, 2017*, pp. 142–157.