

EMFSyncer Solution to the TTC 2020 Round-trip Migration Case

Artur Boronat

School of Informatics, University of Leicester, UK
artur.boronat@le.ac.uk

1 Introduction

The *Round-Trip Migration of Object-Oriented Data Model Instances* case [BKvPK20] for the TTC'2020 exposes an evolution problem in the context of web API development, where different APIs work on top of a common data model. To accommodate new or changing requirements in the system, API designers need to ensure backward-compatible changes in the underlying data model. Assuming an agile software development environment where data models may evolve rapidly, the case proposes the use of round-trip migration services, relying on Model-Driven Engineering (MDE) technology, for enabling the co-existence of different versions of the data model at run time. In particular, the Eclipse Modeling Framework (EMF) [SBPM09] is used to encode data models with the EMF meta-modeling language, namely Ecore, so that tools built atop EMF can be used to implement data migration services.

The solution presented in this article features the EMF-SYNCER [Bor19], a synchronization tool for bridging MDE-agnostic programs and MDE-aware programs, which may have some similarities in their object-oriented data models¹, at run time. EMF-SYNCER provides a generic synchronization strategy that exploits such similarities automatically and is, therefore, a suitable candidate for solving the proposed problems, since each problem relies on a small change and there is a large overlap between two versions of the same data model. The EMF-SYNCER assumes that MDE-agnostic programs are implemented in a JVM language and that MDE-aware programs represent their data model using EMF-generated code.

Given a source data model (represented in a package of classes in the MDE-agnostic program), a target data model (represented as an Ecore model in an EMF-based program), and a collection of objects representing the source program state, EMF-SYNCER translates the MDE-agnostic objects as an EMF model instance in the target EMF-based program, synchronizing the state of both programs at run time. EMF-SYNCER automatically infers structural similarities between object-oriented data models by mapping object structural features by name, when found, translating both attribute and reference values. This translation can be performed using a push-based model, where the entire source program state is migrated, or using a pull-based model, where only those feature values accessed in the target program are migrated. Once synchronization is established, changes that have been applied to target EMF model instances can be incrementally back-propagated to the source MDE-agnostic counterpart. Incrementality of back-migration entails that only changes in target model instances are propagated back and merged within the source instance.

The aforementioned generic mapping strategy that is built in EMF-SYNCER can be customized in order to allow for more complex data transformations between the data models involved. A custom mapping strategy is declared with a mapping specification that maps a source feature type to a target feature type, possibly including feature value transformations, either from source to target, or from target to source, or both. Two main custom mapping strategies can be declared:

- Renaming of feature types: such renamings may affect the name of the class, where the feature is declared, and the name of the feature. This mapping strategy allows for transformations that modify the data model

Copyright held by the author(s).

In: Proceedings of the 13th Transformation Tool Contest, published at <http://ceur-ws.org>

¹We refer to these data models as domain models in [Bor19].

either syntactically, when either a class or a feature is renamed in the target model, or structurally, when a feature type is moved to an unrelated object type. In this case, feature values are transformed using the generic synchronization strategy.

- Transformation of feature values, with or without renaming of feature types: feature value transformations are expressed using Xtend lambda expressions defined for a contextual object type, from the source data model, and result in a single target feature value. This mapping strategy allows for semantic transformations, where the feature value can be computed using navigation expressions from the contextual type. A feature value transformation is applied in a single direction, either from source to target or from target to source. A programmer can opt to provide feature value transformations in both directions (as in `Task_2` in §2.2), only in one (as in `Task_3` in §2.3), or none at all. The last case corresponds to a simple renaming of feature types as explained above.

For the TTC 2020 case, we are relying on the observation that an EMF-based program can be regarded as a plain Java program. Hence, the given Java program will refer to the source data model and the EMF-based program will refer to the target data model. In the following sections, we discuss the solution to the case in §2, and the evaluation of the solution in §3 using the evaluation criteria proposed in the case, wrapping up with some conclusions in §4.

2 Solution

The solutions to the different tasks of the case are explained below and are available at <https://github.com/emf-syncer/ttc20-roundtrip>.

2.1 Task 1: Create/Delete Field

In this task, a new feature `age` is added to the class `Person` in the modified data model M_2 and when a `Person` instance is migrated, its feature value is set to `-1`. When the instance is migrated back, this information is lost.

The solution `Task1_M1_M2_M1`, shown in the listing below, uses the pattern used in all of the solutions. The EMF-SYNCER is used by instantiating the class `EMFSyncer` between a source Java package name list `#['scenario1_v1']` and a target Ecore model in the constructor. All of the solutions use the push-based model in `syncForward`, which is specified by indicating `syncer.syncingStrategy = SyncingStrategy.EAGER`, forcing the migration of the entire source instance.

In the method `migrate`, the statement `syncer.syncForward(person_v1)` migrates the source `Person` instance obtaining the target instance with the `age` feature, which is not initialized until the statement `person_v2.age = -1` is evaluated. In the method `migrateBack`, the statement `syncer.syncBack(person_v2)` migrates the instance back to the source Java program.

The method `modify` is used to implement changes to target model instances that need to be propagated back. This logic is hardcoded in test cases in the benchmark framework. EMF-SYNCER needs to track which changes are performed in the target model instance in order to enable incremental back migration. Therefore, this logic has been moved from test cases to task classes.

```
1 class Task_1_M1_M2_M1 extends AbstractTask {
2     val Syncer syncer
3
4     new (EPackage model1, EPackage model2) {
5         super(model1, model2);
6         syncer = new Syncer(#[ 'scenario1_v1' ], model2)
7         syncer.syncingStrategy = SyncingStrategy.EAGER
8     }
9
10    override migrate(EObject instance) {
11        val person_v1 = instance as scenario1_v1.Person
12        val person_v2 = syncer.syncForward(person_v1) as scenario1_v2.Person
13        person_v2.age = -1
14        return person_v2;
15    }
16
17    override migrateBack(EObject instance) {
18        val person_v2 = instance as scenario1_v2.Person
19        return syncer.syncBack(person_v2) as scenario1_v1.Person
20    }
21
22    override modify(EObject instance) { }
```

It is important to note that the bidirectional transformation between both data models is completely inferred by the EMF-SYNCER automatically in this task. In the back migration, as the `age` feature does not exist in the source program, it is not propagated.

The solution to the symmetric problem `Task1_M2_M1_M2` is achieved by inverting the direction in which the EMF-SYNCER is applied to the data models, as seen in the listing below. In this case, the data model M2 is regarded as the Java program and the data model M1 as the EMF program. In this case, the value of the `age` feature of a `Person` class is preserved on the back migration because this feature does not exist in M1 and, therefore, no changes can be applied to the feature `age`. In the following subsections, solutions to symmetric problems are achieved by flipping, as for this task, the direction in which the EMF-SYNCER is applied and are not included in the paper.

```

1 class Task_1_M2_M1_M2 extends AbstractTask {
2   val EMFSyncer syncer
3
4   new (EPackage model1, EPackage model2) {
5     super(model1, model2);
6     syncer = new Syncer(['scenario1_v2'], model1)
7     syncer.syncingStrategy = SyncingStrategy.EAGER
8   }
9
10  override migrate(EObject instance) {
11    val person_v2 = instance as scenario1_v2.Person
12    return syncer.syncForward(person_v2) as scenario1_v1.Person
13  }
14
15  override migrateBack(EObject instance) {
16    val person_v1 = instance as scenario1_v1.Person
17    return syncer.syncBack(person_v1) as scenario1_v2.Person
18  }
19
20  override modify(EObject instance) { }
21 }

```

2.2 Task 2: Rename Field

In this task, the feature `age` of the class `Person` is renamed to `ybirth`, and its semantics is changed by representing the age and the year of birth, respectively. This renaming involves domain-specific semantics that is not present in the Ecore model, requiring a mapping specification so that the EMF-SYNCER can perform the data transformation correctly.

A mapping specification consists of a collection of mappings between feature types by name, `'Person'`, `'age'`, `/* <-> */ 'Person'`, `'ybirth'`, and by adding optional feature value transformations as lambda expressions. For example, in the solution, the lambda expression

```

1 val person_v1 = it as scenario2_v1.Person
2 Integer.valueOf(Calendar.getInstance().get(Calendar.YEAR) - person_v1.age) as Object

```

gets a `Person` instance from the source data model and returns the `ybirth` value. When the EMF-SYNCER applies the transformation `syncForward`, the feature value for `Person::ybirth` will be obtained by applying this expression. The rest of the solution follows the same structure as the solution explained in §2.1.

```

1 class Task_2_M1_M2_M1 extends AbstractTask {
2   val EMFSyncer syncer
3
4   // custom mapping strategy: M1 <--> M2
5   val public static mapping = new EMFSyncerMapping(
6     'Person', 'age', /* <-> */ 'Person', 'ybirth',
7     // source to target feature value transformation
8     [
9       val person_v1 = it as scenario2_v1.Person
10      Integer.valueOf(Calendar.getInstance().get(Calendar.YEAR) - person_v1.age) as Object
11    ],
12  ],
13 }

```

```

14 // target to source feature value transformation
15 [
16     val person_v2 = it as scenario2_v2.Person
17     Integer.valueOf(Calendar.getInstance().get(Calendar.YEAR) - person_v2.ybirth) as Object
18 ]
19 )
20
21 new (EPackage model1, EPackage model2) {
22     super(model1, model2);
23     syncer = new EMFSyncer(['scenario2_v1'], model2, newArrayList(mapping))
24     syncer.syncingStrategy = SyncingStrategy.EAGER
25 }
26
27 override migrate(EObject instance) {
28     val person_v1 = instance as Person
29     return syncer.syncForward(person_v1) as scenario2_v2.Person
30 }
31
32 override migrateBack(EObject instance) {
33     val person_v2 = instance as scenario2_v2.Person
34     return syncer.syncBack(person_v2) as Person
35 }
36
37 override modify(EObject instance) { }
38 }

```

2.3 Task 3: Declare Field Optional/Mandatory

In this task, the multiplicity of a feature type is modified so that the feature is mandatory in one data model and optional in the modified version. This task exposes a difference between a MDE-agnostic program, where it is not possible to know whether a feature is optional in plain Java, and a MDE-aware one, where this information is encoded in the Ecore model. As the EMF-SYNCER treats the source program as a plain Java program, it disregards the information contained in the Ecore model, and the logic to transform null values in to empty Strings needs to be provided in a custom mapping. The expression `person_v2.name ? : ""` returns an empty String when the `name` of the `Person` instance is null, which is checked using the Elvis operator `?:`. The rest of the data transformation is fully inferred by the EMF-SYNCER .

```

1 class Task_3_M1_M2_M1 extends AbstractTask {
2     val EMFSyncer syncer
3
4     // custom mapping strategy: M1 <--> M2
5     val public static mapping = new EMFSyncerMapping(
6         'Person', 'name', /* <--> */ 'Person', 'name',
7         null,
8         // target to source feature value transformation
9         [
10            val person_v2 = it as scenario3_v2.Person
11            person_v2.name ? : ""
12        ]
13    )
14
15    new (EPackage model1, EPackage model2) {
16        super(model1, model2);
17        syncer = new EMFSyncer(['scenario3_v1'], model2, newArrayList(mapping))
18        syncer.syncingStrategy = SyncingStrategy.EAGER
19    }
20
21    override migrate(EObject instance) {
22        val person_v1 = instance as scenario3_v1.Person
23        return syncer.syncForward(person_v1) as Person
24    }
25
26    override migrateBack(EObject instance) {
27        val person_v2 = instance as Person
28        syncer.syncBack(person_v2) as scenario3_v1.Person
29    }
30
31    override modify(EObject instance) {
32        val person_v2 = instance as Person
33        syncer.track[ person_v2.name = null ]
34        return person_v2
35    }
36 }

```

This solution also contains an example of how to track changes performed in the target instance, in the `modify` method, that need to be migrated back. This change is encoded in the corresponding test case in the original test framework.

2.4 Task 4: Multiple Edits

This task combines `Task_1` and `Task_2` with the aim of analysing reuse mechanisms that can be employed. The solution below reuses the transformation logic of `Task_1`, as it is handled by the EMF-Syncer automatically, and it reuses the mapping specification of `Task_2`, which is defined as a static field. The rest of the solution is as in §2.1, after renaming the corresponding namespaces.

```
1 class Task_4_M1_M2_M1 extends AbstractTask {
2   val EMFSyncer syncer
3
4   new (EPackage model1, EPackage model2) {
5     super(model1, model2);
6     syncer = new EMFSyncer(['scenario4_v1'], model2, newArrayList(Task_2_M1_M2_M1.mapping))
7     syncer.syncingStrategy = SyncingStrategy.EAGER
8   }
9
10  override migrate(EObject instance) {
11    val container_v1 = instance as scenario4_v1.Container
12    return syncer.syncForward(container_v1) as scenario4_v2.Container
13  }
14
15  override migrateBack(EObject instance) {
16    val container_v2 = instance as scenario4_v2.Container
17    syncer.syncBack(container_v2) as scenario4_v1.Container
18  }
19
20  override modify(EObject instance) { }
21 }
```

3 Evaluation

In this section, the author provides an evaluation of the solution according to the evaluation criteria proposed in [BKvPK20].

3.1 Expressiveness

Two test cases are provided for checking the correctness of each task. A test case provides the input for a round-trip migration and the expected output. The test framework has been adapted in order to apply the EMF-Syncer, which works with code generated from Ecore models using EMF. The main properties (namely, `name`, `nsPrefix` and `nsUri`) of the `EPackage` in each Ecore model were updated in order to generate disjoint namespaces. In test cases, loading resources for each Ecore model was modified in order to use the corresponding generated factory for each model instance. Model instances, used as input/output data, were modified to refer to the corresponding `nsUri`. In some cases, the test case also performs a change in the target instance, as in `Task_3`, that needs to be propagated back to the source instance. Such modifications have been encoded in the task itself, in the method `modify`, as the EMF-Syncer needs to track those changes. Such changes refer to implementation details and do not alter the correctness properties checked² by the test suite.

Solutions for all of the tasks have been implemented and all of them pass the correctness check.

3.2 Comprehensibility

Task solutions have been implemented using Xtend. However, the EMF-Syncer can be used from Java programs as well. Given that the transformation in most of the solutions is inferred automatically and that solutions use generated code, instead of using the EMF reflection API for accessing/mutating values, the author argues that solutions are likely to be more comprehensible than the reference ones.

When custom mappings are required, e.g. in `Task_2` and in `Task_4`, these are defined by instantiating the class `EMFSyncerMapping`, where feature value transformations are defined as Xtend lambda expressions. Such mapping expressions could have been defined similarly in Java as well.

²The test case `task_3_M1_M2_M1_b` had to be updated in order to check that the `name` was migrated back correctly.

3.3 Bidirectionality

In the solution for most of the tasks, the EMF-SYNCER infers both transformations to be applied, `syncForward` and `syncBack`, automatically. In such cases, the programmer does not need to provide a transformation specification and a bidirectional data transformation is being used internally so that Java instances can be migrated to an EMF program, and back again, at run time. A custom mapping specification containing only feature type renamings is fully bidirectional. Feature value transformation expressions are unidirectional though.

Therefore, the EMF-SYNCER provides support for bidirectional transformations by default for all of the solutions and accommodates special cases with unidirectional feature value transformations, which correspond to a fraction of the data migration to be performed.

3.4 Re-usability

Re-usability is internalized in the EMF-SYNCER by inferring transformations automatically. That is, when a user does not need to provide a mapping specification, the transformation logic in the EMF-SYNCER is reused for any data model change. For example, reuse of the logic transformation in `Task_1` falls under this category.

On the other hand, as EMF-SYNCER is a JVM library, a programmer can rely on reuse mechanisms provided by the JVM language of choice. For example, reuse of the transformation logic in `Task_2` falls under this category. The mapping specification defined for `Task_2` is reused by calling a static field and by using a common interface for the class `Person`, which has been implemented using inheritance in `Task_4`.

3.5 Performance

We have run the performance tests both for the reference solution and for the EMF-SYNCER solution on a MacBookPro11,5 Core i7 2.5 GHz, with four cores and 16 GB of RAM. The runtime results (in ms.) obtained are displayed in Fig. 1. The EMF-SYNCER solution exhibits a linear growth with respect to the number of iterations. However, the EMF-SYNCER solution is more efficient than the reference one thanks to its support for incremental migration of changes. While the reference solution takes about 31 s. for 2 million iterations, the EMF-SYNCER solution takes about 9 s, with an improvement factor of 70%.

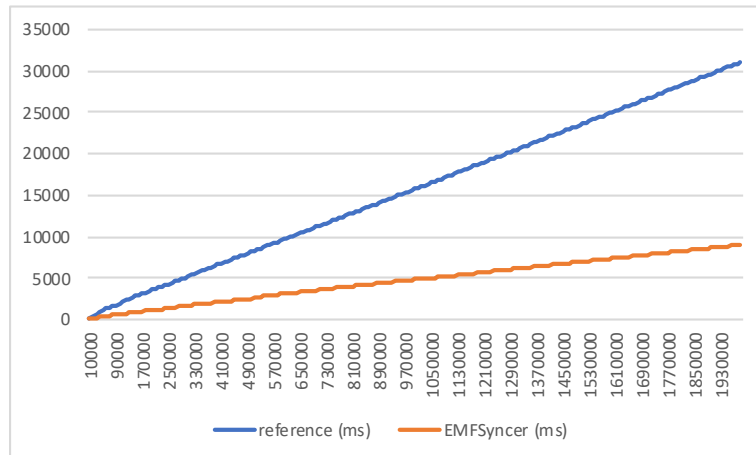


Figure 1: Performance results (ms.) along iterations

4 Conclusions

Following from the justification of the evaluation criteria in the section above, the evaluation results are summarized in the table below. Overall, this solution shows that the EMF-SYNCER helps in achieving a competitive trade-off between data migration specification and performance. On the one hand, the automatic inference of data transformations between different object-oriented data models reduces the need for specifying data transformations. When these have to be specified, a programmer can rely on their programming skills, using a JVM programming language, for reusing transformation logic. On the other hand, EMF-SYNCER can be used as an efficient data migration service at run time. For example, the EMF-SYNCER solution is faster than the reference

solution developed in raw Java. EMF-SYNCEER can be used for building more scalable solutions, involving very large models, by using the pull-based model for propagating changes in `syncForward` only when they are required in the target program.

	Expressiveness	Comprehensibility	Bidirectionality	Re-usability
Task 1: "Create/Delete Field"				
Task_1_M1_M2_M1	2 (2)	2 (2)	1 (1)	n.a.
Task_1_M2_M1_M2	2 (2)	2 (2)	1 (1)	n.a.
Task 2: "Create/Delete Field"				
Task_2_M1_M2_M1	2 (2)	2 (2)	1 (1)	n.a.
Task_2_M2_M1_M2	2 (2)	2 (2)	1 (1)	n.a.
Task 3: "Create/Delete Field"				
Task_3_M1_M2_M1	2 (2)	2 (2)	1 (1)	n.a.
Task_3_M2_M1_M2	2 (2)	2 (2)	1 (1)	n.a.
Task 4: "Create/Delete Field"				
Task_4_M1_M2_M1	2 (2)	2 (2)	1 (1)	4 (4)
Task_4_M2_M1_M2	2 (2)	2 (2)	1 (1)	4 (4)
	\sum : 16 (16)	\sum : 16 (16)	\sum : 8 (8)	\sum : 8 (8)

References

- [BKvPK20] Luca Beurer-Kellner, Jens von Pilgrim, and Timo Kehrer. Round-Trip Migration of Object-Oriented Data Model Instances. In *Proceedings of the 13th Transformation Tool Contest, a part of the Software Technologies: Applications and Foundations (STAF 2018) federation of conferences*, CEUR Workshop Proceedings. CEUR-WS.org, 2020.
- [Bor19] Artur Boronat. Code-first model-driven engineering: On the agile adoption of mde tooling. In *Proceedings of the 34th IEEE/ACM International Conference on Automated Software Engineering (ASE 2019), San Diego, CA, November 11-15*. ACM, 2019.
- [SBPM09] David Steinberg, Frank Budinsky, Marcelo Paternostro, and Ed Merks. *EMF: Eclipse Modeling Framework 2.0*. Addison-Wesley Professional, 2nd edition, 2009.