

Solving the TTC'15 Train Benchmark Case Study with SIGMA

Filip Křikava

University Lille 1 - LIFL, France

INRIA Lille, Nord Europe

filip.krikava@inria.fr

In this paper we describe a solution for the *Transformation Tool Contest 2015* (TTC'15) Train Benchmark case study using SIGMA, a family of Scala internal *Domain-Specific Languages* (DSLs) that provides an expressive and efficient API for model consistency checking and model transformations. We solve both the core and the extension tasks.

1 Introduction

In this paper we describe our solution for the TTC'15 Train Benchmark case study [3] using the SIGMA [1]. SIGMA is a family of Scala¹ internal DSLs for model manipulation tasks such as model validation, model to model (M2M), and model to text (M2T) transformations. Scala is a statically typed production-ready *General-Purpose Language* (GPL) that supports both object-oriented and functional styles of programming. It uses type inference to combine static type safety with a “*look and feel*” close to dynamically typed languages. Furthermore, it is supported by the major integrated development environments bringing EMF modeling to other IDEs than traditionally Eclipse (the solution was developed in IntelliJ IDEA²).

SIGMA DSLs are embedded in Scala as a library allowing one to manipulate models using high-level constructs similar to the ones found in the external model manipulation DSLs. The intent is to provide an approach that developers can use to implement many of the practical model manipulations within a familiar environment, with a reduced learning overhead as well as improved usability and performance. The solution is based on the *Eclipse Modeling Framework* (EMF) [2], which is a popular meta-modeling framework widely used in both academia and industry, and which is directly supported by SIGMA.

The complete source code is available on Github³ in the fork of the original case study repository. In the Appendix A we provide a steps how to install it locally.

¹<http://scala-lang.org>

²<https://www.jetbrains.com/idea/>

³<https://github.com/fikovnik/trainbenchmark-ttc>

2 Solution Description

Essentially, the solution for this transformation case study consist of queries that check for violations of a number of constrains and repair transformations that fixes them. SIGMA provides a dedicated model consistency checking DSL with the ability to provide quick fixes repairing invariant validations. However, given the structure of the case study source, we realize the solution in the way it can be easily compared to the reference implementations in both Java and EMF-IncQuery. Another reason is that it shows that SIGMA can be easily integrated in existing modeling projects and other high-level abstractions can be easily developed. We therefore only rely on the SIGMA common infrastructure which provides a set of basic operations for model navigation (*i.e.* projecting information from models) and modification (*i.e.* changing model properties or elements).

The description of the solution is split in two parts: (1) the core part that describes the queries and repair transformations, (2) the integration part gives an overview how it has been integrated in the case study source code.

The solution has been implemented in a separate module `hu.bme.mit.trainbenchmark.ttc.benchmark.sigma` in a fork of the original code based provided by the case study.

2.1 Queries and Repair Transformations

Essentially, we have created a small internal DSL that builds on SIGMA that allows us to solve the given benchmark cases in an expressive and compact way—*i.e.* to express the model queries and repair transformations.

The top-level constructs in the DSL is a *constraint*. A constraint is composed of a model *query* that finds all model instances violating a certain model restriction and a *repair* transformation correcting the failed instances. Concretely, a query is a function that given a model element—*i.e.* a context of the constraint in the classical model consistency checking—returns a set of *matches*. A match can be either a single instance or a tuple of instances of model elements that are related to the violations.

2.2 Constraint DSL

A typical way of creating an internal DSL in Scala is by designing a library that allows one to write fragments of code with domain-specific syntax. These fragments are woven within Scala own syntax so that it appears different.

One way to represent the above concepts is using a Scala case class:

```

1 case class Constraint[A <: EObject, B <: AnyRef] (
2   query: (A) => Iterable[B],
3   repair: (B) => Unit
4 )

```

This defines a case class with a field for both query and repair. A case class in Scala is like a regular class which some additional properties out which, in our case, the important one is that it can be instantiated without the `new` keyword and thus limiting the language noise. The two

type parameters A , B specify the model context for the query and the types of matches the query produces. The input type is further constrained to be a subtype of an `EObject`.

Given this case class we can write a constraint using:

```
1 Constraint[A, (B, C)](
2   query = ...
3   repair = ...
4 )
```

The query and repair are defined as functions $A \rightarrow \text{Iterable}[B]$ and $B \rightarrow \text{Unit}$ where `Unit` is like `void` in Java.

In some cases the match returned by the query is of the same type as the query context. The query can be therefore simplified to a boolean expression selecting instances on which it evaluates to true. For these type of queries we provide a dedicated constructs called `BooleanConstraint`:

```
1 case class BooleanConstraint[A <: EObject : ClassTag](
2   query: (A) => Boolean,
3   repair: (A) => Unit
4 )
```

For example, the first query, *PosLength*, which finds all the segments with negative length can be written as:

```
1 BooleanConstraint[Segment](
2   query = segment => segment.length < 0,
3   repair = segment => segment.length += -segment.length + 1
4 )
```

In Scala, name `=> ...` is a function literal of one parameter function. We do not have to specify the types of the parameter nor the result as they will be inferred by the Scala compiler.

Next to constraint, we define a validator that allows one to check and consequently repair the incorrect model instances:

```
1 abstract class Validator[A <: EObject, B <: AnyRef] {
2   def check(container: EObject): Iterator[B]
3   def repair(matches: Iterator[B]): Unit
4 }
```

It is defined as an abstract class with two methods that correspond to the two operations we want to do—*i.e.* check and repair. Finally, we provide an implementation for the two types of queries that we have:

```
1 case class ConstraintValidator[A <: EObject, B <: AnyRef](constraint: Constraint[A, B])
2   extends Validator[A, B] {
3
4   override def check(container: EObject) =
5     container.eAllContents collect { case x: A => x } flatMap constraint.query
6
7   override def repair(matches: Iterator[B]) =
8     matches foreach constraint.repair
9 }
```

The implementation is straight forward. For all elements contained in a `container`, we first collect all instances of the required context type and then query them using the query function provided by the given constraint. The repair simply executes the constraint repair function on the matching element.

The DSL is implemented in a Scala file `hu/bme/mit/trainbenchmark/ttc/benchmark/sigma/Co`

2.2.1 Constraints

In the following we describe the individual constraints that were part of the case study (case study tasks) except the *PosLength* that has already been shown above.

— *SwitchSensor*

```

1 BooleanConstraint[Switch] {
2   query = switch => switch.sensor.isEmpty,
3   repair = switch => switch.sensor = Sensor()
4 }

```

The `isEmpty` is a method that is defined on an `Option` type (coming from the standard Scala library) representing a type which may or may not have a value. Since in the model, the *sensor* reference of the *Switch* class is defined as optional (with cardinality 0..1), in SIGMA we represent the reference using the `Option` class. Not only makes this the cardinality expressed in the type definition, but it also prevents some of the `NullPointerException`s caused by traversing unset references. Technically, this is realized by implicit conversions (*cf.* Krikava *et al.* [1]).

— *SwitchSet*

```

1 Constraint[SwitchPosition, (Semaphore, Route, SwitchPosition, Switch)] {
2   query = swP => {
3     for {
4       semaphore <- Option(swP.route.entry) if semaphore.signal == Signal.GO
5       sw = swP.switch if sw.currentPosition != swP.position
6     } yield (semaphore, swP.route, swP, sw)
7   },
8
9   repair = {
10    case (_, _, swP, sw) => sw.currentPosition = swP.position
11  }
12 }

```

This is a more complex constraint that matches a tuple of model elements. It is using a *for comprehension*, a lightweight notation for expressing sequence comprehensions. From the Scala documentation⁴: Comprehensions have the form `for` (enumerators) `yield` *e*, where *enumerators* refers to a semicolon- or new line-separated list of enumerators. An *enumerator* is either a generator which introduces new variables, or it is a filter. A comprehension evaluates the body *e* for each binding generated by the enumerators and returns a sequence of these values.

⁴<http://docs.scala-lang.org/tutorials/tour/sequence-comprehensions.html>

In this concrete example, the generator is the optional value of the `Route.entry` reference. It either generates a single value in the case the actual instance contains one or it does not produce anything. Because, there is a mistake in the model (the `Route.entry` should have the cardinality set to `0..1` instead of `1`), we need to explicitly convert the reference to an `Option`.

The repair function is defined using a pattern matching construct allowing us to concisely assign variables from the matching tuple.

— *RouteSensor*

```

1  Constraint[Route, (Route, Sensor, SwitchPosition, Switch)](
2    query = route => {
3      for {
4        swP <- route.follows
5        sw = swP.switch
6        sensor <- sw.sensor if !(route.definedBy contains sensor)
7      } yield (route, sensor, swP, sw)
8    },
9
10   repair = {
11     case (route, sensor, _, _) => route.definedBy += sensor
12   }
13 )

```

The implementation is similar to the the previous case. It is also based on a for comprehension and closely follows the description of the query.

— *SemaphoreNeighbor*

```

1  Constraint[Route, (Semaphore, Route, Route, Sensor, Sensor, TrackElement, TrackElement)](
2    query = route1 => {
3      for {
4        sensor1 <- route1.definedBy if route1.exit != null
5        te1 <- sensor1.elements
6        te2 <- te1.connectsTo
7        sensor2 <- te2.sensor
8        route2 <- sensor2.sContainer[Route] if route1 != route2
9        semaphore = route1.exit if semaphore != route2.entry
10     } yield (semaphore, route1, route2, sensor1, sensor2, te1, te2)
11   },
12
13   repair = {
14     case (semaphore, _, route2, _, _, _, _) => route2.entry = semaphore
15   }
16 )

```

Again based on the for comprehension. Additionally, we provide a shortcut using the `route1.exit != null` so immediately skip the route instances that do not have an exit semaphore set.

The complete solution is implemented in a Scala object⁵ `hu.bme.mit.trainbenchmark.ttc.benchmark.sigma.Solution`.

2.3 Integration

The integration consists in making our solution work within the provided benchmark framework. For this we simply follow what has been provided for the reference implementation in Java. The

⁵A Scala object defines a single instance of a class.

following files were created:

- `SigmaBenchmarkMain.scala` contains the main application running the benchmark.
- `SigmaBenchmarkLogic.scala` contains the logic selecting which benchmark shall be run.
- `SigmaBenchmarkCase.scala` contains the adaptors bridging the framework with our constraint language. This file also defines `SigmaBenchmarkComparator` that is used to compare the matches as required by the case study. It is a general comparator that either compares single instances (results from boolean constraints violations) or tuples (regular constraints violations).

3 Evaluation

In this section we provide an evaluation of our solution following the categories given by the case study.

3.1 Correctness and Completeness of Model Queries and Transformations

We developed a solution for all of the tasks required by the case study and the solution passes the provided tests.

3.2 Conciseness

The solution itself consists of 52 lines of Scala code the internal DSL developed for this case study. The DSL itself has been implemented using 20 lines of Scala code using SIGMA. The integration part consists of three files with the total of 65 lines. All measures are source lines only excluding comments and new lines. Given these measures, we believe that the code is quite concise.

3.3 Readability

Next to being concise, the solution is also quite expressive. This means that the given problem (queries and repair transformations) naturally maps into the implementation. The higher-level abstraction provided by both SIGMA and the internal DSLs helps to facilitate it making a significant improvement over the Java reference implementation. Next to being concise and expressive, the code is also type-safe as Scala is statically typed language. A notable consequence is that it is very easy to use the DSL with an IDE like Eclipse or IntelliJ that provides a robust code completing functionalities, outline views and other features increasing one's productivity.

In summary, while readability is a subjective matter and largely depends on the background and experience of users, we believe that SIGMA scores well. Thanks to the syntax of Scala which is close to one of Java/C++ and hence shall be familiar to many developers. The expressiveness of the first-order logic collection operation should be familiar to anyone knowing OCL or any other function language.

3.4 Performance on Large Models

The tests have been performed on an 2.3 GHz Intel Core i7 machine with 16 GB of RAM being dedicated to the JVM process. We ran our solution together with the reference implementation in Java. We used the model instances from size 1 to 8192 and set 8GB memory to be dedicated to the JVM. The corresponding results are shown in the figures 1 and 2. We compare them to the Java solution which is shown in the figures 3 and 4. We get a similar performance which has been expected due to the fact that Scala compiles directly to Java bytecode and we use the same underlying libraries for accessing EMF models. This shows that we can leverage from concise and expressive queries without sacrificing performance.

It is important to note that we do not developed any extra functionality for these benchmarks—*i.e.* no caching or incremental validations. On the other hand, functional approach we have selected makes it perfect for further parallelization. Moreover, for even larger models we could easily port the solution into big data analysis platforms such as Hadoop⁶. This would be rather difficult given the imperative Java implementation.




Figure 1: SIGMA fixed validation batch

⁶<https://hadoop.apache.org/>




Figure 2: SIGMA fixed revalidation

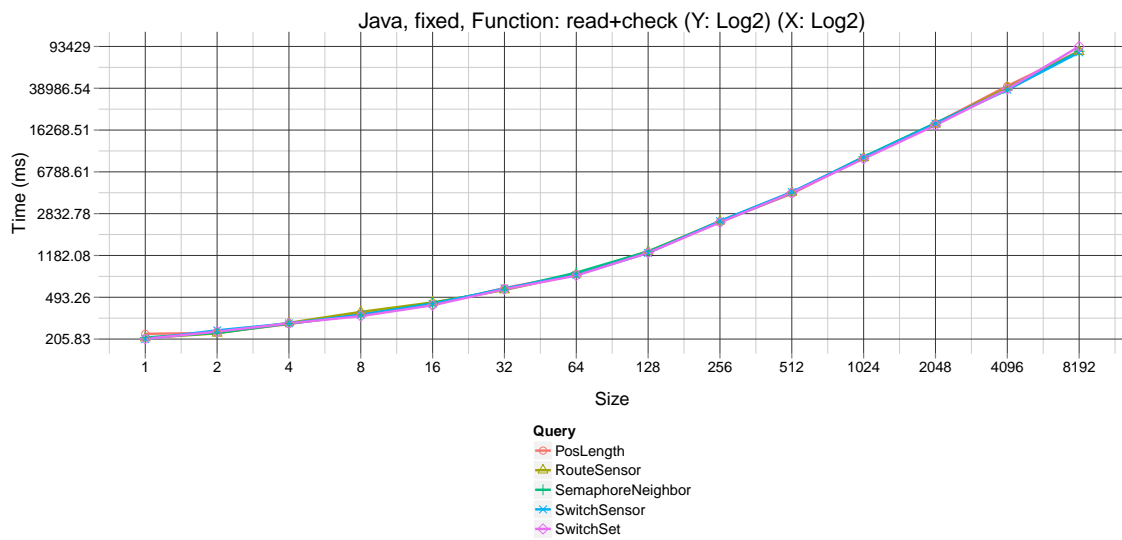


Figure 3: Java fixed validation batch




Figure 4: Java fixed revalidation

Acknowledgments This work is partially supported by the Datalyse project (| www.datalyse.fr/).

References

- [1] Filip Krikava, Philippe Collet & Robert France (2014): *SIGMA: Scala Internal Domain-Specific Languages for Model Manipulations*. In Juergen Dingel, Wolfram Schulte, Isidro Ramos, Silvia Abrahão & Emilio Insfran, editors: *Proceedings of 17th International Conference on Model-Driven Engineering Languages and Systems, MODELS 8767*, Springer International Publishing, Valencia, doi:10.1007/978-3-319-11653-2. Available at <http://link.springer.com/10.1007/978-3-319-11653-2>.
- [2] Dave Steinberg, Frank Budinsky, Marcelo Paternostro & Ed Merks (2008): *EMF: Eclipse Modeling Framework (2nd Edition)*. Addison-Wesley Professional. Available at <http://www.eclipse.org/emf>.
- [3] Gábor Szárnyas, Oszkár Semeráth, István Ráth & Dániel Varró (2015): *The TTC 2015 Train Benchmark Case for Incremental Model Validation**. In: *Transformation Tool Contest 2015*.

A Install

The only requirements for running the solution is git and maven. To reproduce the benchmark simply execute these steps in a command line:

```
1 $ git clone https://github.com/fikovnik/trainbenchmark-ttc
2 $ cd trainbenchmark-ttc
3 $ ./script/run.py
```
