

The Fulib solution to the TTC 2020 migration case

Sebastian Copei¹, Albert Zuendorf¹

¹Kassel University, Germany

1. Introduction

At Kassel University we are working on a solution for bidirectional transformations based on event sourcing for about a year, now. It turned out, that the TTC 2020 migration case [1] is a special case of a bidirectional transformation and that our approach provides a reasonable solution for it.

2. Design

The idea and design for our solution stems from Domain Driven Design [2] and Event Sourcing [3]. Basically, we use two editors `M1Editor` and `M2Editor`, one for each model of the case study, cf. Figure 1. Each editor holds the current object model (based on the corresponding ecore model). In addition, each editor provides editing commands following the command design pattern of [4]. All operations on the object model are encapsulated within editor commands. Each editor keeps track of all executed commands (and the used command parameters) within its `event store`, cf. Figure 1. To enable collaboration of `M1Editor` and `M2Editor`, both editors provide the same set of commands: a `HavePerson` command with parameters `id`, `name`, and `age` and a `HaveDog` command with parameters `id`, `ownerId`, `name`, and `age`, cf. Figure 3 and Figure 2.

each editor implements the command execution differently according to its specific ecore model. As an example, Listing 1 shows the implementation of the `HavePerson` command within `M2Editor`. Line 22 of Listing 1 shows how the `age` parameter of the `HavePerson` command is turned into a `ybirth` value for model M2.

Generally, we consider `M1Editor` and `M2Editor` as independent programs that may run on different computers, concurrently. Therefore each editor is able to serialize its event store (in `yaml` format) and to send its commands to the other editor. Correspondingly, the receiving editor is able to deserialize and execute the commands, too.




Figure 1: Design




Figure 2: Objects

In [5] we have developed theoretical foundations for this command (or event) sharing between multiple editors: basically, we require that multiple commands that address the same object (have the same id) overwrite each other, i.e.: if `c1` and `c2` are commands with the same id than applying `c1` and then `c2` is similar to applying only `c2`. In addition, commands that work on different objects may be executed in any order (are commutative), i.e. if `c3` and `c4` are commands with different ids, applying first `c3` and then `c4` must result in the same object model as applying `c4` first and `c3` second. While these are pretty strong conditions, it turned out to be easy to implement the commands according to these rules. For the TTC2020 migration case we will discuss this in Section 3.

Provided with overwriting and commutative commands, our editors are able to merge commands executed by themselves or received from another editor: if a new command arrives that uses an id that is already used by some old command, the new command is executed and then the old command is replaced by the new command in the event store. If a new command arrives that uses a new id, the new command is executed and added to the

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).
EMAIL: sco@uni-kassel.de (S. Copei); zuendorf@uni-kassel.de (A. Zuendorf)

© 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)



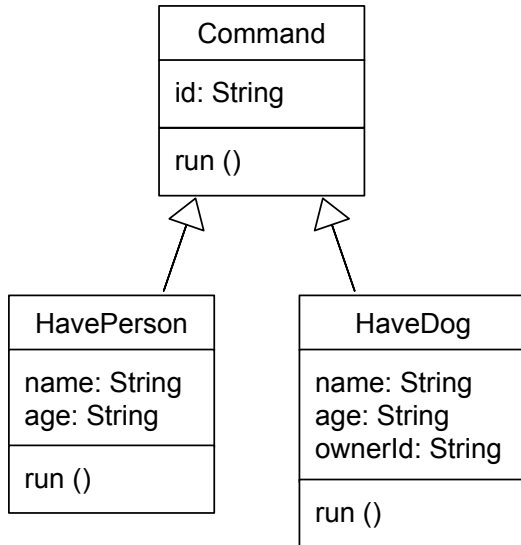


Figure 3: Command Classes

event store. The event store is treated as a set, i.e. the order of the commands (and the order of the command execution) does not matter due to our commutativity condition.

Overall this editor design enables us to implement the TTC2020 migration case as outlined in Figure 4. The migrate step is invoked by the test or benchmark Tasks. The loading of objects is done by loading the corresponding xml files via EMF mechanisms. The parsing steps are discussed in Section 4. The sending is done via our yaml serialization and deserialization. The execution of commands is discussed in Section 3. The modification step is again done by the test or benchmark tasks. Similarly, the test and benchmark tasks invoke the migrate back step. The migrate back step uses parsing and sending and execution similar to the migrate forward step. However, some details dealing with missing information are discussed in Section 4.

3. Commands

Our approach relies on commands that are shared between the two editors. In the simple TTC20 Migration case we need only two commands, cf. Figure 3 and Figure 2. To facilitate the merging of concurrent edits, we require that commands are overwriting and commutative. Thus, if we e.g. execute a HavePerson command with a certain id on an empty model, the first time we shall create a Person object and fill its attributes. However the second time, we shall not create a second Person

1. migrate M1 to M2
 - a) load M1 objects
 - b) parse M1 objects into M1 event store
 - c) send M1 commands to M2Editor
 - d) execute commands by M2Editor creating M2 objects and filling M2 event store
2. modify M2 objects to M2'
3. migrate M2' back to M1'
 - a) parse (modified) M2' objects and detect new commands and merge the new commands into the old M2 event store
 - b) send the updated M2 commands to the M1Editor
 - c) execute (modified and new) commands on M1 objects and update M1 event store

Figure 4: The Fulib Migration Approach

object but we shall lookup the already existing object and just adjust its attributes. In Line 5 of Listing 1 our HavePerson command achieves this behavior by using the getOrCreatePerson method of its editor. Our editors have hash tables for each model class (i.e. for Person and Dog) and the getOrCreatePerson method looks up this hash table. If the hash table already has a Person with the given id, this Person is returned. Otherwise, getOrCreatePerson creates a Person object, initializes its id, adds it to the hash table, and then returns the new Person. Thus, you may call getOrCreatePerson with a certain id as often as you like, it will always return the same Person object.

This getOrCreate mechanism also helps us to achieve commutativity for our commands. Commutativity for commands requires e.g. that you can execute a HaveDog command before you execute the HavePerson command for the dog owner. To allow this, our HaveDog commands uses getOrCreatePerson with the ownerId to retrieve the corresponding Person object. If the owner already exists, we just use it, otherwise, the owner object is created on the fly. In the latter case, a subsequent execution of the HavePerson command with the same id will retrieve the already existing Person object and then fill the Person's name and age (or ybirth).

Once the targeted Person object has been retrieved, Line 6 to Line 22 of Listing 1 fill the attributes of that Person. In this solution, we use EMF dynamic editing features to set the name and age or ybirth attributes. The different migration tasks use different ecore models with different properties. We store the current ecore model within the editors and thus the commands just query the current ecore model e.g. whether the current Person has either an age attribute or an ybirth attribute. Thus one implementation of our M2 .HavePerson com-

```

1 public class HavePerson extends Command {
2     @Override
3     public Object run(M2Editor editor) {
4        EObject person =
5         editor.getOrCreatePerson(getId());
6         EClass personClass =
7         (EClass) editor
8         .getEmfModel()
9         .getEClassifier("Person");
10        person.eSet(
11        personClass.getEStructuralFeature("name"),
12        name);
13        EStructuralFeature ageFeature =
14        personClass.getEStructuralFeature("age");
15        if (ageFeature != null) {
16            person.eSet(ageFeature, this.age);
17        }
18        EStructuralFeature ybirthFeature =
19        personClass
20        .getEStructuralFeature("ybirth");
21        if (ybirthFeature != null) {
22            person.eSet(ybirthFeature,
23            2020 - this.age);
24        }
25        return person;
26    }
27    ...

```

Listing 1: M2.HavePerson::run()

mand suffices to address all different migration tasks of the TTC2020 migration case. Note, the `M1.HavePerson` command also has only one implementation for all tasks, however this implementation is slightly simpler as the one of `M2.HavePerson` as no variant of model `M1` deals with `ybirth` attributes.

4. Parsing

The implementation of the different migration tasks provided by the TTC2020 migration case resources [6] load the different start models from XML files. After the first migration the task implementation may retrieve the migrated model and it may modify the migrated model directly via dynamic EMF means. This means, the task implementation does not use our editor commands to load or to modify the models. Thus, our first task is to parse the provided model and to derive the editing commands that correspond to it.

For this purpose, our editors provide a parse method that basically uses a visitor to travel through the current model. For each model object our visitor then calls special `parsePerson` or `parseDog` methods, respectively.

Listing 2 shows the `parsePerson` method of our `M2Editor`. Basically, Line 5 to Line 20 of Listing 2 use dynamic EMF features to retrieve the parameters needed for the corresponding `HavePerson` command. Line 20 to Line 25 then create the desired `HavePerson` command and provide its parameters. Finally, Line 26 executes the

```

1 public class M2Editor {
2     ...
3     private void parsePerson(EObject instance) {
4         EClass eClass = instance.eClass();
5         String name =
6         (String) instance
7         .eGet(eClass
8         .getEStructuralFeature("name"));
9         int age = -1;
10        EStructuralFeature ageFeature =
11        eClass.getEStructuralFeature("age");
12        if (ageFeature != null) {
13            age = (Integer) instance.eGet(ageFeature);
14        }
15        EStructuralFeature ybirthFeature =
16        eClass.getEStructuralFeature("ybirth");
17        if (ybirthFeature != null) {
18            age = 2020 -
19            (Integer) instance.eGet(ybirthFeature);
20        }
21        HavePerson havePerson = new HavePerson()
22        .setName(name)
23        .setAge(age);
24        String id = getPersonId(instance);
25        havePerson.setId(id);
26        execute(havePerson);
27    }
28    ...

```

Listing 2: M2Editor::parse()

command. Thereby, the editor adds the new command to the event store.

One crucial step during parsing is the retrieval of ids for the model objects, cf. Line 15 of Listing 2. Usually, we require that the model objects have an `id` attribute. In the TTC2020 migration case this is not true. We solve this by storing the ids for model objects within the hash tables that are used by our `getOrCreate` methods. To look up the id of a model object that is already stored in such a hash table, we search e.g. the hash table for Persons for the given instance. If an entry exists, we return the corresponding id. If there is no entry yet (i.e. after loading the initial object model) we just create a new id and add the instance to the hash table under this new id and then return the new id. Within the backward migration step the hash table will already contain the instance and the old id is retrieved.

There is one special case when merging the commands created by the parse methods into the editor's event store: In migration Task4 of the TTC2020 migration case [1], in model `M2` the `Dog` has no `age` attribute. Thus, when we parse the `Dog` object during the backward migration step, the `parseDog` method will not find any age information. However, this age information is still contained within the original `HaveDog` command that has been received from `M1Editor` and that has been executed and added to the event store of `M2Editor` during forward transformation. To keep the age information alive, if the `parseDog` method cannot find the age attribute in

the parsed model object, it tries to retrieve the old command from the M2Editor event store and copies the age information from there into the new command.

Altogether, our parsing approach utilizes that our command execution is overriding and commutative. Due to the commutativity, the parsing visitor may visit the objects of the current input model in any order and thus create the editor commands in any order. Due to the overriding property, the backward migration may just (re)execute the detected commands and this will overwrite the old commands and blend into the event store, easily.¹

5. Results

Overall, our editor and command approach was very well suited for the TTC2020 migration case. We were able to address all migration task with the same implementation. While our design might appear a little bit over engineered we took benefit from the `FuLibServiceGenerator`, a tool we are just building and that generates a lot of boiler plate code for editors and commands. After all we just had to implement the run methods of the 2 commands for the two editors and the parsing methods for the two editors as described above.

We believe that the use of overwriting and commutative commands provides a great leverage for the parsing and merging of command sets. Overall, our design is able to handle much more complicated migration cases which we address in our current work.

Unfortunately, our performance is very poor, on a test run the Fulib solution took 9.6 seconds for 10000 iterations while the original solution used only 0.3 seconds. We had no time to go into the details of this.

You find our solution on:

Github:

<https://github.com/fujaba/ttc2020MigrationCaseByFulib>

Docker: [zuendorf/fulib-solution-ttc2019](https://github.com/fujaba/ttc2020MigrationCaseByFulib)

References

- [1] ttc2020migration, Ttc2020 case: Round-trip migration of object-oriented data model instances, https://www.transformation-tool-contest.eu/2020_roundtrip.pdf, 2020. Last viewed 13.06.2020.
- [2] E. Evans, Domain-driven design: tackling complexity in the heart of software, Addison-Wesley Professional, 2004.
- [3] V. Vernon, Implementing domain-driven design, Addison-Wesley, 2013.
- [4] E. Gamma, R. Helm, R. Johnson, J. Vlissides, Design patterns: Abstraction and reuse of object-oriented design, in: European Conference on Object-Oriented Programming, Springer, 1993, pp. 406–431.
- [5] S. Copei, A. Zündorf, Mx for microservices, in: Proc. Dagstuhl Seminar, volume 18491, 2019.
- [6] ttc2020resources, Ttc2020 case: Round-trip migration of object-oriented data model instances, github resources, <https://github.com/lbeurkellner/ttc2020>, 2020. Last viewed 13.06.2020.

¹Fortunately, the TTC2020 migration case does not include any delete operations during model modification. To handle deletion of model objects one has to remove the corresponding old commands from the event store and one has to propagate this command removal to the other editor and the other editor needs to remove the command, too, and it needs to undo the command in order to roll back the corresponding model changes. This rollback needs some careful dealing with our `getOrCreate` operations.