

# The Epsilon Solution to the OCL2PSQL Case

Antonio García-Domínguez  
a.garcia-dominguez@aston.ac.uk  
Aston University  
Birmingham, UK

## ABSTRACT

There have been several attempts to map Object Constraint Language queries to SQL: one of these is the OCL2PSQL mapping proposed by Nguyen and Clavel. In this paper, I describe an implementation of OCL2PSQL using two languages from the Eclipse Epsilon project: 740 lines of ETL code for the model-to-model transformation itself, and 96 lines of EGL code for a model-to-text transformation that produces more readable SQL than the reference version. The solution passes all correctness tests set out in the original framework: the transformation has a median time of 0.93s across all scenarios.

## KEYWORDS

OCL, SQL, model transformation, abstract syntax graphs

### ACM Reference Format:

Antonio García-Domínguez. 2021. The Epsilon Solution to the OCL2PSQL Case. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

The OMG Object Constraint Language is a well known, standardized language for specifying constraints and queries in models: while typically associated with UML, it has been widely used for MOF-based modelling languages, and particularly those based on the Ecore implementation of EMOF. Given the increasingly large size of the models used by Model-Driven Engineering practitioners, one common solution is to persist them in databases. This has motivated attempts to map the OCL queries (written in terms of the abstract syntax of the modeling language) to SQL queries that run directly on the underlying relational database used to store the models.

Some of these attempts have used imperative features such as loops and cursors to deal with iterators, possibly reducing their compatibility across vendors (due to the limited standardization of these features). OCL2PSQL (“OCL to pure SQL”) is an approach that provides a mapping of nested iterators while staying entirely within the broadly standardized and declarative parts of SQL [2]. The OCL2PSQL TTC case has selected a core subset of this mapping (with some erratum added since the original release of the mapping), and has invited tool authors to demonstrate the usability,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*Conference'17, July 2017, Washington, DC, USA*

© 2021 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

conciseness and ease of understanding of their model transformation languages through alternative implementations of this subset.

This paper presents an outline of a solution based on the Eclipse Epsilon family of model management languages. Since the original release in 2006 [1], Eclipse Epsilon has grown to include languages for model-to-model transformation (the Epsilon Transformation Language), model-to-text transformation (the Epsilon Generation Language), model validation, pattern matching, model migration, unit testing and other tasks. The solution passes all correctness checks, though some minor refinements of the proposed mappings were required.

The rest of the paper is structured as follows: Section 2 explains the overall structure of the solution. Section 3 describes the key features of the implemented model-to-model transformation from OCL to SQL. Section 4 describes the alternative model-to-text transformation that has been developed from the SQL models to textual SQL queries. Finally, Section 5 presents the preliminary performance results obtained by the solution author.

## 2 OVERALL STRUCTURE

The Epsilon solution to the OCL2PSQL case is a Java application, using Apache Maven for its dependency management. Epsilon is usable as a standalone Java library, with stable versions available through Maven Central and snapshot versions available through the OSS Sonatype repository. The solution should work in Epsilon 2.3.0, but uses the latest 2.4.0 interim versions to avoid a warning message when using the Epsilon Generation Language.

The solution reuses the basic scaffolding of the reference solution, including the generated code for the OCL and SQL metamodels, and the classes responsible for interpreting the environment variables, communicating with the MySQL database, and performing the correctness tests. The solution adds the following Java classes:

- **SAMPLELAUNCHER**, which transforms all OCL queries without using the environment variables of the benchmark framework. This is mostly for internal development.
- **OCL2SQL**, which encapsulates a model-to-model transformation written in the Epsilon Transformation Language into an easy-to-use Java class. More information on the ETL transformation is available in Section 3.
- **SQL2TEXT**, which replaces the model-to-text transformation in the reference solution with one based on the Epsilon Generation Language. More details on the EGL transformation are given in Section 4.

Finally, the solution changed the code of the **SOLUTION** class in the reference solution to use the **OCL2SQL** and **SQL2TEXT** classes from above.

### 3 MODEL-TO-MODEL TRANSFORMATION WITH ETL

The Epsilon project includes several languages for performing model transformations:

- The Epsilon Object Language (EOL) is an OCL-inspired imperative language which is well suited for small in-place (endogenous) model transformations, though it can be used for purely imperative exogenous transformations as well.
- The Epsilon Transformation Language (ETL) builds on top of EOL by adding the concept of a *rule*, which transforms a certain type of source model element (possibly with some *guards* limiting its applicability) to a target model element. The rule scheduling can be kept entirely declarative, or can be controlled to some extent through the use of “greedy” or “lazy” rules. ETL is well-suited to exogenous transformations, where a new model is produced from the source model.
- Flock can be thought of as ETL with an automated “copy-unless-otherwise-stated” default strategy. It is well suited for *model migration* tasks, where a model has to undergo small changes from one version of a metamodel to the next.
- The Epsilon Wizard Language (EWL) is a variant of EOL which allows users to define “wizards” that users can manually trigger on specific model elements, performing small in-place transformations (perhaps with some simple user interaction).

Out of these languages, ETL was chosen since the original formulation of OCL2PSQL as a collection of recursive mappings by source type was a close match to the rules used by ETL. For the most part, each of those mappings was translated into an ETL rule, and calls to  $\text{map}_e(e)$  were translated into  $\text{e.equivalent}()$  calls in ETL. The  $\text{e.equivalent}()$  built-in operation retrieves the SQL subgraph produced from the OCL  $e$  subtree, allowing the different SQL subtrees to be linked together. The OCL2PSQL case did not require any manual rule scheduling or the use of greedy/lazy rules: the automated rule scheduling done by ETL based on source element types and guards was sufficient.

ETL allows for breaking up the transformation across several modules. This made it relatively easy to iteratively implement the various stages in OCL2PSQL and test out how it behaved for the various challenges. Listing 1 shows the main script of the ETL transformation: it is composed of a number of **import** statements that bring in the rules needed for each stage, and a **post** rule which places the one `PLAINSELECT` without a container into a `SELECT-STATEMENT` element as mandated by the SQL metamodel.

Some of the OCL2PSQL mappings had to produce significantly large SQL subtrees: to keep the code concise, a library of EOL operations (`utilities.eol`) was created and reused from the ETL rules. This library largely contained a set of functions for simple creation of SQL model elements, an implementation of the OCL2PSQL  $\text{SVars}_e(e)$  and  $\text{FVars}_e(e)$  functions, and several other miscellaneous functions. As a simple example, Listing 2 shows the code needed to transform OCL integer literals to the target SQL metamodel.

The ETL rules are for the most part a direct one-to-one translation from the descriptions at the end of the OCL2PSQL case, except for two changes.

Listing 1: Main ETL script

```

1 import 's0_literals.etl';
2 import 's1_equals_and.etl';
3 import 's2_allInstances.etl';
4 import 's3_size.etl';
5 import 's4_collect_variable.etl';
6 import 's5_attributes.etl';
7 import 's6_associationEnds.etl';
8 import 's7_exists.etl';
9 import 's8_existsWithFree.etl';
10
11 post {
12   var firstRootSelect = SQL!PlainSelect.all.selectOne(ps|ps.
13     eContainer.isUndefined());
14
15   var stmt = new SQL!SelectStatement;
16   stmt.selectBody = firstRootSelect;
17 }

```

Listing 2: Excerpt of ETL for stage 0 (integer literals)

```

1 import 'utilities.eol';
2
3 /*
4  * All these boil down to:
5  *
6  * mape(l) =
7  * SELECT l as res, 1 as val
8  */
9
10 rule IntLiteral
11 transform e:OCL!IntegerLiteralExp
12 to ps:SQL!PlainSelect {
13   ps.sellItems.add(longSelectItem('res', e.integerValue.asLong()));
14   ps.sellItems.add(longSelectItem('val', 1));
15 }

```

The first change was considering one special case listed in the original OCL2PSQL paper [2] but not in the case paper. The mapping of `collect` and `exists` in the original OCL2PSQL paper covered the case when  $v \notin \text{FVars}(b)$ , but this mapping had been omitted from the OCL2PSQL case description. It turned out that this special case was needed for some of the queries, e.g. challenge 0 in stage 4 (`Car.allInstances()->collect(c5)`).

The second change was due to an unexpected interaction between the recursive approach used to define OCL2PSQL, the definition of `SubSelect.selectBody` as a containment reference in the SQL metamodel, and how the ETL  $\text{e.equivalent}()$  operation works. ETL will only apply a certain rule once to each matching source model element, and from them on  $\text{e.equivalent}()$  will always return the same counterpart in the target model (e.g. the exact same object). This resulted in some SQL queries “losing” the body of their `SUBSELECT` objects to other subtrees of the SQL model, as they also needed the mapping for that part of the OCL expression.

**Listing 3: Fix for “subtree stealing” in ETL**

```

233
234
235 1 operation copyIfContained(value) {
236 2   var emfTool = new Native(
237 3     "org.eclipse.epsilon.emc.emf.tools.EmfTool");
238 4   if (value.eContainer.isDefined()) {
239 5     return emfTool.ecoreUtil.copy(value);
240 6   }
241 7   return value;
242 8 }

```

For instance, consider the final challenge:

```

243
244
245
246
247 Car.allInstances()->exists(c|c.owners->exists(p|p.name = 'Peter'))
248

```

In this query, the SQL mapping of `c.owners->exists(...)` and the SQL mapping of `p.name` both require using the mapping of `c.owners` as a subquery. ETL successfully maps `c.owners` to SQL, but EMF will not allow a single model element to be contained from more than one place. Since the mapping of `c.owners->exists(...)` will complete last, it will effectively “steal” the subquery representing `c.owners` from the mapping of `p.name`.

The fix for this issue turned out to be simple, as shown in Listing 3. All uses of the `e.equivalent()` operation were wrapped into a new EOL operation: the operation tested if this “stealing” was about to take place (i.e. if the `PLAINSELECT` was already contained in another `SUBSELECT`), and if so it performed a deep cloning of the SQL subtree.

A better fix (which unfortunately would have required a rewriting of the input files for this case) would be to change the SQL metamodel so that `SubSelect.selectBody` is no longer a containment reference, and the same `PLAINSELECT` can be reused from multiple `SUBSELECT` model elements.

Even further, this suggests that the OCL2PSQL mapping really produces a SQL expression graph (where some subexpressions are reused) rather than a SQL abstract syntax tree. Instead of running the same subquery from several places, it may be advisable to redefine OCL2PSQL so it produces a sequence of SQL queries rather than a single large SQL query: it would run these reused subqueries first, and then provide their results to the higher-level queries. Otherwise, there may be a risk that the SQL query could grow exponentially if sufficiently large subqueries have to be duplicated across several locations.

Overall, the transformation required writing 14 rules across 528 lines of ETL code, with a support library of EOL operations that was 212 lines long. These line counts included whitespace and comments.

## 4 MODEL-TO-TEXT TRANSFORMATION TO SQL WITH EGL

The reference solution included a model-to-text transformation that produced the SQL query to be run in MySQL from the SQL model. During the development of this solution, it was found that the generated SQL was difficult to read in the presence of multiple levels of subqueries, as it was entirely on one line.

In order to improve the readability of the SQL queries and help with the debugging, an alternative implementation was written in 96 lines of EGL. The EGL template traverses the SQL model recursively from the root `SELECTSTATEMENT`, breaking up `SELECT` statements, `CASE` expressions, joins, and subqueries across multiple lines.

**Listing 4: Excerpt of EGL to generate SQL query text**

```

291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324

```

```

1 [%= SelectStatement.all.first.generate() %][%
2
3 @template
4 operation SelectStatement generate() { %]
5 [%=self.selectBody.generate(%);
6 [% }
7
8 @template
9 operation PlainSelect generate() { %]
10 SELECT
11 [% for (si in self.sellItems) { %]
12 [%=si.generate() + (hasMore ? ", " : "")%]
13 [% }
14 if (self.fromItem.isDefined()) { %]
15 FROM [%=self.fromItem.generate() %]
16 [% }
17 for (join in self.joins) { %]
18 [%=join.generate(%)]
19 [% }
20 if (self.whereExp.isDefined()) { %]
21 WHERE [%=self.whereExp.generate() %]
22 [% }
23 if (self.groupBy.isDefined()) { %]
24 [%= self.groupBy.generate() %]
25 [%
26 }

```

Listing 4 shows an excerpt of the EGL script: the first line is the entry point of the entire script, kicking off the recursive descent of the SQL model from the `SELECTSTATEMENT`. The EGL script makes heavy use of *template* operations, which allow EGL templates in their body and return strings which can be used within expressions (e.g. for concatenating separators, as in line 12).

The script also uses some of the built-in Epsilon variables: *hasMore* is a built-in Epsilon variable available in loops which is true if and only if there are more values after the current one.

**Listing 5: SQL-specific LanguageFormatter used to indent the SQL query text**

```

325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348

```

```

1 private static class SQLFormatter extends LanguageFormatter {
2   private static final String increasePattern = "\\(\\s*$";
3   private static final String decreasePattern = "^\\)";
4
5   public SQLFormatter() {
6     super(Pattern.compile(increasePattern, Pattern.MULTILINE),
7           Pattern.compile(decreasePattern, Pattern.MULTILINE));
8   }
9 }

```

**Listing 6: SQL query for challenge 0 in stage 1, as generated by EGL**

```

351 1 SELECT
352 2 TEMP_left.res = TEMP_right.res AS res,
353 3 1 AS val
354 4 FROM (
355 5     SELECT
356 6         2 AS res,
357 7         1 AS val
358 8 ) AS TEMP_left
359 9 JOIN (
36010     SELECT
36111         3 AS res,
36212         1 AS val
36313 ) AS TEMP_right;
    
```

One useful feature in EGL is its ability to integrate *formatters* that postprocess the generated text. In particular, the EGL LANGUAGE-FORMATTER was customised for SQL (as shown in Listing 5) to automatically indent the lines of the SQL script to improve readability, while keeping the EGL script as simple as possible. This class only requires the regular expressions that should increase and decrease the indentation level after a match. Using this script, queries are generated in the more readable form shown in Listing 6.

**5 PRELIMINARY RESULTS**

After implementing the ETL model-to-model transformation and the EGL model-to-text transformation, Java code to encapsulate these transformations and integrate them with the TTC benchmark framework was added. The transformations passed all correctness cases for all scenarios across all stages and challenges.

In terms of execution time, the transformations were run in a Lenovo X1 laptop with an i7-6600U CPU running at 2.60GHz with 16GiB of physical RAM, running Ubuntu Linux 20.04.2 LTS with Linux 5.4.0-74-generic and the Oracle JDK 11.0.8. The default

Java memory allocation settings were used (no -Xmx or other JVM options were given). The Docker image provided by the OCL2PSQL case authors was used to run MySQL, using Docker Engine 20.10.7.

The execution and test times are shown in Figure 1. It must be noted that the execution times only cover the ETL model-to-model transformation: the testing times include the EGL model-to-text transformation and the execution of the SQL queries in MySQL. This is the same approach followed by the reference solution.

In general, it can be seen that the transformation generally completes within 1 second, and testing times are within 0.5 seconds. More specifically, transformation times ranged between 0.83s and 1.06s, with a median time of 0.93s. Testing times ranged between 0.33s and 0.44s, with a median time of 0.36s.

**REFERENCES**

[1] Dimitrios S. Kolovos, Richard F. Paige, and Fiona Polack. 2006. The Epsilon Object Language (EOL). In *Model Driven Architecture - Foundations and Applications, Second European Conference, ECMDA-FA 2006, Bilbao, Spain, July 10-13, 2006, Proceedings*. 128–142. [https://doi.org/10.1007/11787044\\_11](https://doi.org/10.1007/11787044_11)

[2] Hoang Nguyen Phuoc Bao and Manuel Clavel. 2019. OCL2PSQL: An OCL-to-SQL Code-Generator for Model-Driven Engineering. In *Future Data and Security Engineering*, Tran Khanh Dang, Josef Küng, Makoto Takizawa, and Son Ha Bui (Eds.). Springer International Publishing, Cham, 185–203. [https://doi.org/10.1007/978-3-030-35653-8\\_13](https://doi.org/10.1007/978-3-030-35653-8_13)

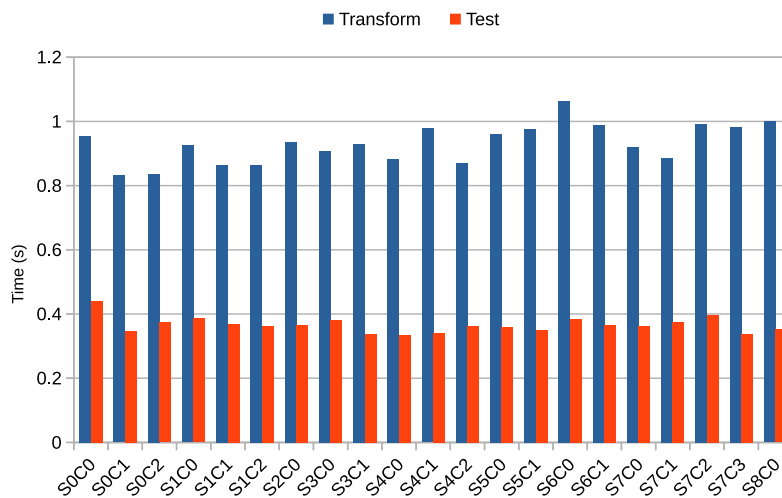


Figure 1: Transformation and test execution times in seconds per stage and challenge.