

An NMF Solution to the TTC 2021 OCL to SQL Case

Georg Hinkel

Am Rathaus 4b, 65207 Wiesbaden, Germany
georg.hinkel@gmail.com

Abstract

This paper presents a solution to the OCL to SQL translation case at the Transformation Tool Contest (TTC) 2021 using dynamic C# code, but without any dedicated model transformation language, using NMF for the model representation. The transformation tools of NMF are not used because the case does not fall under NMF's definition of a model transformation problem.

1 Introduction

The Object Constraint Language (OCL) is an important language to denote expressions based on models. If the models are stored in a database, it is desirable to translate these queries to SQL statements such that they can be processed directly by the database.

Using models for OCL and SQL, the OCL to SQL case at the Transformation Tool Contest (TTC) asks tool authors to transform models of OCL queries into models of SQL statements. For this purpose, metamodels are provided to understand OCL and SQL as models.

The .NET Modeling Framework [1] is a framework for model-driven engineering on the .NET platform and paper presents a solution of the OCL to SQL case using NMF.

In the remainder of the paper, I first briefly introduce the Dynamic Language Runtime that is heavily used for the solution in this paper in Section 2. Section 3 presents the solution and Section 4 shows the results achieved with the solution.

2 Dynamic C#

The solution makes use of the dynamic language runtime (DLR) that is part of the .NET Framework but perhaps not so widely known. The idea of the DLR is to allow elements of dynamic programming languages in the scope of the .NET runtime. These features are also available in C#, in particular the ability for late binding. That is, by converting variables to `dynamics`, the compiler sees that method calls are only resolved at runtime, based on the usual C# overload selection principles which the compiler attaches to make them available at runtime. However, especially when passing dynamic objects only as parameters, the compiler is able to calculate the set of methods that are candidates for a certain call already, which makes the actual call very efficient. Further, editors such as Visual Studio even show errors, if no suitable candidates could be found, the reference count counts all candidates and the "Go To Definition" feature lists all candidates.

3 Solution

To discuss the solution, I first give an overview in Section 3.1 before Sections 3.2, 3.3 and 3.4 go into details for the actual translation process, pruning and printing the SQL statement models to strings.

Copyright held by the author(s).

In: A. Garcia-Dominguez, G. Hinkel, and A. Boronat (eds.): Proceedings of the 14th Transformation Tool Contest, on-line, 25-06-2021

3.1 Overview

NMF does have a model transformation language (NTL, [2], [3]) but I decided not to use it for this case. Why? According to the philosophy of NTL, the biggest challenge of a model transformation is to establish an isomorphism between source and target models that provides a tracing functionality and that is used to ensure that certain input model elements are only transformed once and not once for every reference. This is because maintaining such a trace is difficult in general-purpose programming languages because it requires a lot of bookkeeping – one essentially requires a dedicated hashtable for each type and as soon as inheritance is in place, things start to become messy.

However, both the OCL and the SQL metamodels are essentially expression models that have a tree structure with very few cross-references, even none in the case of SQL. Because NMF takes containments very serious and model elements must always have exactly one parent, trying to add an existing model element to a containment reference of another model element removes it from its old container. Therefore, not only that a trace is not needed, it is even counter-productive.

Since the availability of a trace is not an argument in favor of NTL, the question is whether NTL still adds value against a pure general-purpose code solution and I believe the answer is plainly no. Especially using features like the DLR, the late binding can be implemented directly in C# with concepts known by a lot more developers and therefore easier to understand and better supported by tools.

Therefore, I decided to create a solution to the case using plain C# code making use of DLR features.

3.2 Translator

The general idea of the solution is to translate the OCL expressions in a (mutable) context to SQL expressions. This context includes a notion of open variables and their types as well as the body of the enclosing SQL statement and a counter of temporary tables created for a statement in order that they do not get confused. While simple expressions can be mapped to simple SQL expressions, other OCL expressions require to modify the context in which they are called.

Listing 1 shows how this applies to boolean expressions where the literal is simply converted to an EQUALSTOEXPRESSION, either that $1 = 1$ for `true` or $1 = 0$ for `false`.

```
1 private IExpression GetExpression(SelectContext context, BooleanLiteralExp booleanLiteral) {
2     return new EqualsToExpression {
3         LeftExp = new LongValue { Value = 1 },
4         RightExp = new LongValue {
5             Value = booleanLiteral.BooleanValue.GetValueOrDefault() ? 1 : 0
6         }
7     };
8 }
```

Listing 1: Translating simple boolean expressions

Calls to `GetExpression` can be nested as denoted in Listing 2 that depicts how to translate AND call expressions.

```
1 return new AndExpression {
2     LeftExp = GetExpression(context, (dynamic)callExpression.Source),
3     RightExp = GetExpression(context, (dynamic)callExpression.Argument[0])
4 };
```

Listing 2: Nesting translation calls to translate an AND call expression

More interesting is the handling of the `AllInstances` method as depicted in Listing 3. Because it does not directly have an impact on the result, we return a null reference, but this time change the context and set it to the table with the name of the referred type.

```
1 private IExpression GetAllInstances(SelectContext context, IEntity referredType) {
2     var table = new Table { Name = referredType.Name };
3     if (context.Body.FromItem == null) {
4         context.Body.FromItem = table;
5     } else {
6         context.Body.Joins.Add(new Join { RightItem = table });
7     }
8     return null;
9 }
```

Listing 3: Handling the `AllInstances` method

To handle iterators, we need to determine how to bind the variable. For this, the considered subset of the OCL language knows to collections that can be iterated: A collection returned by the AllInstances method or an association of a different variable. In both cases, we add an open variable to the select context while calculating the expression for the iterator body and remove it afterwards.

With the iterators in place, we can implement the PROPERTYCALLEXP expressions as depicted in Listing 4.

```

1 private IExpression GetExpression(SelectContext context, PropertyCallExp propertyCall) {
2     switch (propertyCall.Source) {
3         case VariableExp variableRef:
4             var table = context.Variables[variableRef.ReferredVariable.Name];
5             return new Column {
6                 Table = new Table {
7                     Name = table,
8                     Alias = new Alias {
9                         Name = variableRef.ReferredVariable.Name
10                    }
11                },
12                Name = propertyCall.ReferredProperty.Name
13            };
14         default:
15             throw new NotSupportedException();
16     }
17 }

```

Listing 4: Transformation of a PROPERTYCALLEXP

The (syntactically allowed) case that a property of a property is queried would require adding more joins, which is ignored in the current solution, particularly given that this was not required for the reference inputs.

In case of an ASSOCIATIONCALLEXP, we register the join as last join in the context and add the join to the current select context as depicted in Listing 5.

```

1 private IExpression GetExpression(SelectContext context, AssociationClassCallExp association) {
2     switch (association.Source) {
3         case VariableExp variableRef:
4             var variable = variableRef.ReferredVariable.Name;
5             var associationEnd = association.ReferredAssociationEnds;
6             var alias = variable + "_" + associationEnd.Association;
7             context.Body.Joins.Add(new Join {
8                 Left = false,
9                 RightItem = new Table {
10                    Name = associationEnd.Association,
11                    Alias = new Alias { Name = alias }
12                },
13                OnExp = new EqualsToExpression {
14                    LeftExp = new Column {
15                        Table = new Table {
16                            Name = context.Variables[variable],
17                            Alias = new Alias { Name = variable }
18                        },
19                        Name = associationEnd.Name,
20                    },
21                    RightExp = new Column {
22                        Table = new Table {
23                            Name = associationEnd.Association,
24                            Alias = new Alias { Name = alias }
25                        },
26                        Name = context.Variables[variable] + "_id",
27                    }
28                }
29            });
30             context.LastJoin = Tuple.Create(variable, associationEnd);
31             return null;
32         default:
33             throw new NotSupportedException();
34     }
35 }

```

Listing 5: Transformation of an ASSOCIATIONCALLEXP

Perhaps the most interesting expression is the method to return the sizes. This is because the the aggregate drastically changes the execution of the query and we need to return rows for actually empty combinations. To do this, we create a temporary sub-select model with the current context query inside, group that query by all context variables and return a column of the temporary table. However, because this eliminates the open variables that might be needed elsewhere, we group the result by all open variables and add these variables to

the result. To make them available in the sub-select, which is the new context select statement, we add joins for each open variable from their original table.

To see this, consider an extension of stage 8 where we reuse the open variable `c` as depicted in Listing 6. We refer to this query later on as stage 9.

```
1 Car.allInstances()->exists(c|c.owners->exists(p|p.name = 'Peter') and c.color='black')
```

Listing 6: Slight extension of the stage 8 query that reuses the open variable `c`

Note, the `exists` method is treated as a filter condition and an additional size aggregate. We need to keep the variable `c` in order to be able to check whether the color is black.

3.3 Pruning

The resulting SQL statement may join tables that are not actually needed, e.g. when joined tables are not actually needed. This gets apparent in challenge 8, where the open variable `c` is only used to calculate the size, but given that we are not interested in any of its properties, we do not actually need to join the `CAR` table once again after the initial context is gone.

```
1 if (selectBody.SelItems.Select(s => s.Exp).OfType<CountAllFunction>().Any()) {
2     return;
3 }
4 var expressionsToCheck = selectBody.SelItems.Select(s => s.Exp).ToList();
5 if (selectBody.WhereExp != null) {
6     expressionsToCheck.Add(selectBody.WhereExp);
7 }
8 var usedAliases = (from selectExp in expressionsToCheck
9                     from column in selectExp.Descendants().OfType<Column>()
10                    select column.Table.Alias.Name).Distinct();
11 for (int i = selectBody.Joins.Count - 1; i >= 0; i--) {
12     var join = selectBody.Joins[i];
13     if (join.RightItem is Table table && !usedAliases.Contains(table.Alias.Name)) {
14         selectBody.Joins.RemoveAt(i);
15     }
16 }
17 if (selectBody.FromItem is SubSelect subSelect) {
18     Prune(subSelect.SelectBody);
19 }
```

Listing 7: Pruning the joins of the resulting SQL statement

The implementation of the pruning is depicted in Listing 7. Aggregate (sub-)queries are not pruned because removing joins changes the number of result elements and thus the result get incorrect. Otherwise, we select all table aliases that appear either in the selection or in the where clause and remove all joins that join tables that are not actually needed. Lastly, we recurse in case the source is a sub-query.

3.4 Printer

The solution to print the SQL statement models to strings works similar by using the DLR to dispatch the different object types and then print them to strings.

```
1 public static string Print(IPlainSelect selectBody) {
2     var resultBuilder = new StringBuilder();
3     resultBuilder.Append($"SELECT_{string.Join(", ", selectBody.SelItems.Select(Print))}");
4     if (selectBody.FromItem != null) {
5         resultBuilder.Append($"FROM_{PrintFrom((dynamic)selectBody.FromItem)}");
6     }
7     foreach (var join in selectBody.Joins) {
8         resultBuilder.Append($"_{(join.Left.GetValueOrDefault()_?"LEFT":_?"INNER")}_{JOIN}_{PrintFrom((dynamic)join.RightItem)}_{ON}_{PrintExpression((dynamic)join.OnExp)}");
9     }
10    if (selectBody.WhereExp != null) {
11        resultBuilder.Append($"_WHERE_{PrintExpression((dynamic)selectBody.WhereExp)}");
12    }
13    if (selectBody.GroupBy != null) {
14        resultBuilder.Append($"_GROUP_BY_{string.Join(", ", selectBody.GroupBy.GroupByExps.Select(exp=>_PrintExpression((dynamic)exp))}_");
15    }
16    return resultBuilder.ToString();
17 }
```

Listing 8: Printing the resulting SQL statement using the DLR

As an example, the method to print the actual SQL statement is depicted in Listing 8. The query printer makes intensive use of the string interpolation available in C#.

4 Evaluation

The solution has been integrated into the benchmark framework. The resulting queries are depicted in Listing 9.

```

1 ***** Stage#0 ***
2 +++ challenge#0: SQL: SELECT 2 res
3 +++ challenge#1: SQL: SELECT 'Peter' res
4 +++ challenge#2: SQL: SELECT 1 = 1 res
5 ***** Stage#1 ***
6 +++ challenge#0: SQL: SELECT 2 = 3 res
7 +++ challenge#1: SQL: SELECT 'Peter' = 'Peter' res
8 +++ challenge#2: SQL: SELECT 1 = 1 and 1 = 1 res
9 ***** Stage#2 ***
10 +++ challenge#0: SQL: SELECT Car_id res FROM Car
11 ***** Stage#3 ***
12 +++ challenge#0: SQL: SELECT tmp1.res res FROM (SELECT COUNT(*) res FROM Car) AS tmp1
13 +++ challenge#1: SQL: SELECT tmp1.res = 1 res FROM (SELECT COUNT(*) res FROM Car) AS tmp1
14 ***** Stage#4 ***
15 +++ challenge#0: SQL: SELECT 5 res FROM Car AS c
16 +++ challenge#1: SQL: SELECT c.Car_id res FROM Car AS c
17 +++ challenge#2: SQL: SELECT 1 = 0 res FROM Car AS c
18 ***** Stage#5 ***
19 +++ challenge#0: SQL: SELECT c.color res FROM Car AS c
20 +++ challenge#1: SQL: SELECT c.color = 'black' res FROM Car AS c
21 ***** Stage#6 ***
22 +++ challenge#0: SQL: SELECT tmp1.res res FROM (SELECT c.Car_id, COUNT(c_Ownership.ownedCars) res FROM Car AS c LEFT JOIN Ownership
    AS c_Ownership ON c.Car_id = c_Ownership.ownedCars GROUP BY c.Car_id) AS tmp1
23 +++ challenge#1: SQL: SELECT tmp1.res = 0 res FROM (SELECT c.Car_id, COUNT(c_Ownership.ownedCars) res FROM Car AS c LEFT JOIN
    Ownership AS c_Ownership ON c.Car_id = c_Ownership.ownedCars GROUP BY c.Car_id) AS tmp1
24 ***** Stage#7 ***
25 +++ challenge#0: SQL: SELECT tmp1.res > 0 res FROM (SELECT COUNT(*) res FROM Car AS c WHERE 1 = 1) AS tmp1
26 +++ challenge#1: SQL: SELECT tmp1.res > 0 res FROM (SELECT COUNT(*) res FROM Car AS c WHERE 1 = 0) AS tmp1
27 +++ challenge#2: SQL: SELECT tmp1.res > 0 res FROM (SELECT COUNT(*) res FROM Car AS c WHERE c.color = 'black') AS tmp1
28 +++ challenge#3: SQL: SELECT tmp2.res > 0 res FROM (SELECT COUNT(*) res FROM (SELECT c.Car_id, COUNT(c_Ownership.ownedCars) res
    FROM Car AS c LEFT JOIN Ownership AS c_Ownership ON c.Car_id = c_Ownership.ownedCars GROUP BY c.Car_id) AS tmp1 WHERE tmp1.
    res = 1) AS tmp2
29 ***** Stage#8 ***
30 +++ challenge#0: SQL: SELECT tmp2.res > 0 res FROM (SELECT COUNT(*) res FROM (SELECT c.Car_id, COUNT(p.Person_id) res FROM Car AS c
    LEFT JOIN Ownership AS c_Ownership ON c.Car_id = c_Ownership.ownedCars LEFT JOIN Person AS p ON c_Ownership.ownedCars = p.
    Person_id WHERE p.name = 'Peter' GROUP BY c.Car_id) AS tmp1 WHERE tmp1.res > 0) AS tmp2

```

Listing 9: Resulting SQL Statements

Notably, to reduce the influence of just-in-time compilation, we actually run the solution 100 times and divide the result by 100¹. The resulting transformation times then are in the range of up to 1.4ms for the stage 8 query and in the sub-millisecond area for most of the other queries. The time for the test lies around 20ms but that certainly gets more interesting once the solution is tested with larger databases.

References

- [1] G. Hinkel, “NMF: A multi-platform Modeling Framework,” in *Theory and Practice of Model Transformations: 11th International Conference, ICMT 2018, Held as Part of STAF 2018, Toulouse, France, June 25-29, 2018. Proceedings*, accepted, to appear, Springer International Publishing, 2018.
- [2] G. Hinkel, “An approach to maintainable model transformations using an internal DSL,” Master’s thesis, Karlsruhe Institute of Technology, 2013.
- [3] G. Hinkel, T. Goldschmidt, E. Burger, and R. Reussner, “Using Internal Domain-Specific Languages to Inherit Tool Support and Modularity for Model Transformations,” *Software & Systems Modeling*, pp. 1–27, 2017.

¹Actually, we do not because the smallest time unit in .NET happens to be 100ns, so we merge the division by 100 with the multiplication by 100.