

An NMF solution for the *Petri Nets to State Charts* case study at the TTC 2013

Georg Hinkel

Karlsruhe Institute of Technology
Karlsruhe, Germany

georg.hinkel@student.kit.edu

Thomas Goldschmidt

ABB Corporate Research
Ladenburg, Germany

thomas.goldschmidt@de.abb.com

Lucia Happe

Karlsruhe Institute of Technology
Karlsruhe, Germany

lucia.kapova@kit.edu

Software systems are getting more and more complex. Model-driven engineering (MDE) offers ways to handle such increased complexity by lifting development to a higher level of abstraction. A key part in MDE are transformations that transform any given model into another. These transformations are used to generate all kinds of software artifacts from models. However, there is little consensus about the transformation tools. Thus, the Transformation Tool Contest (TTC) 2013 aims to compare different transformation engines. This is achieved through three different cases that have to be tackled. One of these cases is the Petri Net to State Chart case. A solution has to transform a Petri Net to a State Chart and has to derive a hierarchical structure within the State Chart. This paper presents the solution for this case using NMF Transformations as transformation engine.

1 Introduction

The challenge of the Petri Nets to State Charts case of the TTC 2013 [2] is to transform a Petri Net into a State Chart. Furthermore, both the Petri Net and the State Chart are afterwards to be reduced. This paper presents a solution for these tasks using NMF¹, an open source project to support model-driven engineering on the .NET platform for the transformation part. The reduction of Petri Net and State Charts has been written in general purpose C# that is embedded in the transformation and makes use of the trace model created by NMF Transformations. The solution is available on SHARE².

2 .NET Modelling Framework (NMF)

The .NET Modelling Framework is an open source project that provides support for model-driven software development on the .NET platform. An essential part is the model transformation engine, NMF TRANSFORMATIONS, which allows to write rule-based transformations in arbitrary .NET languages using an internal DSL [1]. The reason to implement the transformation language as internal DSL is mainly that transformation languages ought to be Turing complete [3] and thus, many advantages of external DSLs attenuate. An internal DSL, however, can make use of features of its host language. Developers used to this language feel familiar with the DSL.

NMF TRANSFORMATIONS makes it possible to specify model transformations directly in C#. For this purpose, NMF TRANSFORMATIONS has a simple abstract syntax but hides the complexity in the attributes of the metaclasses which are representing functions. These functions can be specified with general purpose code that contain code as sophisticated as required. Although the transformation language

¹<http://nmf.codeplex.com>

²http://is.ieis.tue.nl/staff/pvgorp/share/?page=ConfigureNewSession&vdi=XP-TUe_TTC13::NMF_TTC13::NMF_updated_NMF_LiveContest.vdi

might seem quite verbose, especially when compared with external model transformation languages, C# has been chosen as host language to make it easier to write and thus maintain these transformations for C# developers.

Currently, NMF does not contain a metamodeling foundation, e.g. based on MOF. Instead, NMF TRANSFORMATIONS uses the concepts of the CLR (the virtual machine used on the .NET platform) to represent models and operates on plain objects (POCOs). Thus, we used an interop component to EMF, which generates classes from an Ecore metamodel. Furthermore, there exists a serializer component to load and store simple models (models that do not have references to other files).

Beside NMF, to the best of our knowledge, there is hardly any framework that supports model-driven engineering on the .NET platform. Microsoft offers a Visualization and Modeling SDK³ for Visual Studio, which is a tool for graphical editors, and T4 as Text-To-Text-Transformation engine. However, although T4 is included in Visual for many years now, there is still hardly support for editing T4 templates. There is an add-in providing syntax-highlighting and code-completion, but still the support is much less than for writing normal e.g. C# code. Furthermore, T4 has some restrictions like no inheritance is allowed within a T4-template. These restrictions and the lack of out-of-the-box tool support make model-driven development hard. NMF TRANSFORMATIONS makes it possible to use the full tool support for C# also for model transformations.

3 Solution

As the case description was divided into several parts of the transformation, this section is also divided in several subsections each presenting the solution of a subtask within the transformation. Thus, section 3.1 presents the initialization before section 3.2 explains the reduction.

3.1 Initialization

The task of the initialization is to create an initial structure for the statechart model. Although it is not mentioned in the description, the first rule for this initialization is that for a Petri Net, a Statechart has to be created with a topstate, which is an AND element.

A M2M-transformation in NMF TRANSFORMATIONS is specified through transformation rules, which are represented by classes. These transformation rules may only be called once per input arguments within a transformation context, i.e. a transformation pass. Furthermore, transformation rules may define dependencies to other rules. These dependencies are necessary to set, in order to let NMF TRANSFORMATIONS know which rules to call and to derive the inputs of these rules. NMF TRANSFORMATIONS operates on plain CLR objects and therefore does not know the structure of the metamodel. Thus the structure of the domain model has to be reflected in the dependencies of the transformation rules.

Thus, we have to specify that whenever a Petri Net is transformed, also all places of that Petri Net have to be transformed, before or after the transformation of the Petri Net. The only rule that is actually called by NMF TRANSFORMATIONS automatically is the rule that matches the transformation request to transform a Petri Net into a state chart. This rule is depicted in Figure 1.

The first rule in Figure 1, *PetriNet2StateChart*, is a rule that transforms a Petri Net into a State Chart. The only thing that happens is that *PetriNet2TopState* is called and the resulting top state is applied to the State Chart. The second rule, *PetriNet2TopState* creates this top state for each Petri Net. This rule already contains information for the next two rules: For every Place within the PetriNet, a corresponding

³<http://archive.msdn.microsoft.com/vsvmsdk>

```

public class PetriNet2StateChart : TransformationRule<Net, Statechart>
{
    public override void RegisterDependencies()
    {
        Require(Rule<PetriNet2TopState>(), (chart, topState) => chart.TopState = topState);
    }
}

public class PetriNet2TopState : TransformationRule<Net, AND>
{
    public override void RegisterDependencies()
    {
        RequireMany(Rule<Place2Or>(),
            selector: net => net.Places,
            persistor: (and, places) => and.Contains.AddRange(places));

        RequireMany(Rule<Transition2HyperEdge>(),
            selector: net => net.Transitions,
            persistor: (and, transitions) => and.Contains.AddRange(transitions));
    }
}

```

Figure 1: The transformation rules to transform the Petri Net

OR element should be created and added to the initial AND top state. Furthermore, any transition should be transformed to a HyperEdge. Note that within this rule, there is no information about how the places are to be transformed, it is only required that they are transformed using the *Place2OR*-rule. The same applies for the transitions. The *Transition2HyperEdge* rule again requires the incoming and outgoing places to be transformed. However, the engine takes care that any place is only transformed once.

There can be multiple parameters specified to filter the dependencies, select the inputs for the dependent transformation rule or persist the output. The code example in Figure 1 already contains several cases. In the require dependency of the *PetriNet2StateChart* rule, there is only a persistor specified. This persists the output of the dependent rule back to the output of the current rule, i.e. the top state of the State Chart is set. The first *RequireMany* statement further has a selector specified. This selector chooses the input for the dependent transformation rule, i.e. it specifies the places that should serve as inputs for the *Place2Or* transformation rule. The last dependency works in the same way.

The initialization rules can directly be reflected in transformation rules. However, due to space limitations, they are not displayed here. In total, the initialization consists of just five transformation rules:

- *PetriNet2StateChart* that transforms a Petri Net into a State Chart
- *PetriNet2TopState* that creates the top state for a State Chart
- *Place2Basic* that creates the Basic element for each place in the Petri Net
- *Place2Or* that creates the OR element for each place in the Petri Net
- *Transition2HyperEdge* that creates the HyperEdge for a each transition

With these five transformation rules, the initialization task is completed. The demanded equivalence function is implicitly stored in the transformation context, as it contains a trace where we just need to resolve a place using the *Place2Basic* or *Place2OR* rule. However, this trace functionality is not serialized by default. If we wanted to serialize the trace into a file, we would have to do this on our own.

3.2 Reduction

As NMF TRANSFORMATIONS only checks application conditions for transformation rules once, we implemented the reduction in general purpose code in C#. However, this general purpose code is embedded in the transformation and makes use of the transformation engine, especially of the provided trace functionality.

Therefore, we extend the transformation with the given reduction rules and include the reduction code. As the rules apply on transitions, we write this code in the *Transform*-method of the *Transition2HyperEdge*-rule. We can just call the reduction rule from within the *Transform*-method, as both the reduction rules and the transformation is written in C#. We can even put the helper methods that we need into the class representing the *Transition2HyperEdge*-rule.

The implementations of the two main reduction rules are both placed within a single method. Both methods consist of the following steps: (1) At first, it is checked whether the reduction rule is applicable at all. If not, the method immediately returns. (2) Next, the rule is applied to the Petri Net and (3) to the State Chart model. (4) Finally, the rule checks whether another rule is applicable with new inputs. Besides step (3), all these steps are implemented using traditional straight-forward imperative code that operates on the in-memory model representations. The only point where the reduction needs an access to the correspondence established by the model transformation is when the reduction rules are applied to the State Chart model. To have an access to the trace functionality, both methods have an additional parameter that is the transformation context.

```
var or_transformed = context.CallTransformation<Place>(andRulePlace2Or, or)
    .Output as OR;
```

Figure 2: Calling a transformation rule within the reduction

As the reduction rules rely on the trace functionality, it is necessary to use update the trace when new OR compounds have to be created. This is done via the code snippet shown in Figure 2 where *andRulePlace2Or* is a cached reference to an additional transformation rule that transforms a place into an OR compound without creating a Basic element.

To query the correspondence established by the tracing functionality, we can simply use the trace component that is given as a reference of the transformation context and either resolve many places at once (as needed for the AND rule) or resolve single places (as needed for the OR rule). This tracing functionality is shown in the Figure 3.

```
var q_or = context.Trace.Resolve<Place, OR>(q);
var r_or = context.Trace.Resolve<Place, OR>(r);
```

Figure 3: Using the trace functionality to resolve single places in the OR rule

Figure 3 shows the tracing functionality without the specification of a transformation rule acting as key. Thus, the trace looks up all transformations from a Place to an OR, i.e. collects transformations using either *AndRulePlace2Or* or *Place2Basic*.

4 Analysis

4.1 Performance measurements

The average execution times on SHARE⁴ for the performance test cases are shown in table 1. Each test case was repeated five times. All the tests were created in a single batch mode to delete the influence that the code has to be compiled by the JIT compiler when the transformation is run for the first time. However, the transformation has not been reused, although NMF TRANSFORMATIONS allows such procedure. Instead, the transformation has been initialized every time from scratch.

The transformation used the original provided performance test cases with id-referencing scheme and thus, the load times are a bit slow.

4.2 Profiler results

As NMF TRANSFORMATIONS is used through an internal DSL for C#, the tool support integrated within Visual Studio and other tools can be applied to model transformations created with it. Besides a great debugging and refactoring support, this means that a profiler can be applied to the model transformations. This made it possible to unveil the biggest performance degrade in previous versions. These previous versions used ordered sets due to the index-based referencing scheme in the test models. As the profiler showed, the remove operation of these ordered sets was consuming large shares of the execution time. Indeed, by changing the underlying data structure for the collections, the performance increased by two magnitudes.

5 Conclusion

In this paper we have presented a solution to the TTC 2013 *Petri Nets to State Charts* case based on NMF TRANSFORMATIONS. It was not possible to support every bit of the transformation with NMF, as the reduction was entirely written in general purpose code interacting with the transformation engine. However, it was easy to integrate the general purpose code into the transformation.

We suggest the high points of our solution as

- **Good execution speed.** The biggest test model is transformed in a few seconds.
- **Easy integration of general purpose code.** The whole reduction part is written as general purpose code that interacts with the transformation.
- **Great tool support,** as NMF TRANSFORMATIONS can reuse for example debugging, profiling, refactoring, testing and continuous integration support for C#.

References

- [1] Martin Fowler (2010): *Domain-specific languages*. Addison-Wesley Professional.
- [2] Pieter Van Gorp & Louis M. Rose (2013): *The Petri-Nets to Statecharts Transformation Case*. <http://planet-s1.org/ttc2013/images/userdirs/122/ttc2013/pn2sc.pdf>.
- [3] Shane Sendall & Wojtek Kozaczynski (2003): *Model transformation: The heart and soul of model-driven software development*. *Software, IEEE* 20(5), pp. 42–45, doi:10.1109/MS.2003.1231150.

⁴http://is.ieis.tue.nl/staff/pvgorp/share/?page=ConfigureNewSession&vdi=XP-TUe_TTC13::NMF_TTC13::NMF_updated_NMF_LiveContest.vdi

A Appendix

Test case	Reading input	Transformation	Writing output
sp200	12.96ms	8.78ms	8.62ms
sp300	19.87ms	13.06ms	12.72ms
sp400	27.50ms	15.66ms	18.44ms
sp500	32.80ms	21.14ms	21.62ms
sp1000	72.09ms	55.31ms	41.10ms
sp2000	176.91ms	111.13ms	103.97ms
sp3000	295.12ms	175.32ms	142.97ms
sp4000	465.48ms	248.64ms	160.39ms
sp5000	618.60ms	299.40ms	241.14ms
sp10000	1,887ms	745.50ms	499.84ms
sp20000	6,273ms	1,549ms	1,600ms
sp40000	22,012ms	2,923ms	2,277ms
sp80000	83,417ms	5,639ms	6,453ms
sp100000	128,445ms	7,040ms	7,343ms
sp200000	488,709ms	13,532ms	18,317ms

Table 1: Execution times of the performance test cases