

An NMF solution to the TTC2021 incremental recompilation of laboratory workflows case

Georg Hinkel¹

¹Tecan Software Competence Center GmbH, Peter-Sander-Straße 41a, 55252 Wiesbaden, Germany

Abstract

This paper presents a solution to the Incremental Recompilation Laboratory Workflows Case at the TTC 2021 using the .NET Modeling Framework (NMF). This solution is able to derive an incremental change propagation almost entirely in an implicit manner.

Keywords

Incremental, Model transformation, Laboratory automation, Workflows

1. Introduction

The transformation of high-level process models to low-level jobs actually executed on machines is a common problem not only in laboratory automation but also in other domains such as smart production. In these domains, it is desirable to adapt an executed process in case of errors or at least avoid wasting resources if it is clear that the complete workflow cannot be performed completely. As there are typically a lot of things that could go wrong, it is desirable to design a transformation system in such a way that an incremental change propagation can be inferred, i.e. does not have to be specified by the developer.

To assess to what degree current model transformation tools are able to infer an incremental change propagation in such scenarios, the Transformation Tool Contest¹ 2021 hosts a case for incremental recompilation of laboratory automation workflows. This paper presents a solution to this case using the .NET Modeling Framework (NMF) [1].

NMF is a framework built to support model-driven engineering, incremental model analyses and incremental model transformations. In particular, NMF Expressions [2] is an incrementalization system able to incrementalize arbitrary function expressions and NMF Synchronizations [3, 4] is an incremental model transformation approach. Using both tools in combination, it is possible to solve the incremental laboratory workflows case in a declarative manner such that the required change propagations can be derived mostly implicitly.

The remainder of this paper is structured as follows:

TTC'21: Transformation Tool Contest, Part of the Software Technologies: Applications and Foundations (STAF) federated conferences, Eds. A. Boronat, A. García-Domínguez, and G. Hinkel, 25 June 2021, Bergen, Norway (online).

georg.hinkel@tecan.com (G. Hinkel)

© 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)

¹<https://www.transformation-tool-contest.eu>

Section 2 gives a brief overview on how NMF Expressions and NMF Synchronizations work. Section 3 explains the actual solution. Section 4 evaluates the solution against the reference solution. Finally, Section 5 concludes the paper.

2. NMF Expressions and NMF Synchronizations

NMF Expressions [2] is an incrementalization system integrated into the C# language. It takes expressions of functions and automatically and implicitly derives an incremental change propagation algorithm. This works by setting up a dynamic dependency graph that keeps track of the models' state and is adapted when necessary. The incrementalization system is extensible and supports large parts of the Standard Query Operators (SQO²).

NMF Synchronizations is a model synchronization approach based on the algebraic theory of synchronization blocks. Synchronization blocks are a formal tool to run model transformations in an incremental (and bidirectional) way [4]. They combine a slightly modified notion of lenses [5] with incrementalization systems. Model properties and methods are considered morphisms between objects of a category that are set-theoretic products of a type (a set of instances) and a global state space Ω .

A (well-behaved) in-model lens $l : A \leftrightarrow B$ between types A and B consists of a side-effect free GET morphism $l \nearrow \in \text{Mor}(A, B)$ (that does not change the global state) and a morphism $l \searrow \in \text{Mor}(A \times B, A)$ called the PUT function that satisfy the following conditions for all $a \in A, b \in B$ and $\omega \in \Omega$:

$$\begin{aligned} l \searrow (a, l \nearrow (a)) &= (a, \omega) \\ l \nearrow (l \searrow (a, b, \omega)) &= (b, \tilde{\omega}) \quad \text{for some } \tilde{\omega} \in \Omega. \end{aligned}$$

²<http://msdn.microsoft.com/en-us/library/bb394939.aspx>; SQO is a set of language-independent standard APIs for queries, specifically defined for the .NET platform.

The first condition is a direct translation of the original PUTGET law. Meanwhile, the second line is a bit weaker than the original GETPUT law because the global state may have changed. In particular, we allow the PUT function to change the global state.

A (single-valued) synchronization block S is an 8-tuple $(A, B, C, D, \Phi_{A-C}, \Phi_{B-D}, f, g)$ that declares a synchronization action given a pair $(a, c) \in \Phi_{A-C} : A \cong C$ of corresponding elements in a base isomorphism Φ_{A-C} . For each such tuple in states (ω_L, ω_R) , the synchronization block specifies that the elements $(f(a, \omega_L), g \nearrow (b, \omega_R)) \in B \times D$ gained by the lenses f and g are isomorphic with respect to Φ_{B-D} .




Figure 1: Schematic overview of unidirectional synchronization blocks

A schematic overview of a synchronization block is depicted in Figure 1. The usage of lenses allows these declarations to be enforced automatically and in both directions, if required. The engine computes the value that the right selector should have and enforces it using the PUT operation. Similarly, a multi-valued synchronization block is a synchronization block where the lenses f and g are typed with collections of B and D , for example $f : A \hookrightarrow B^*$ and $g : C \hookrightarrow D^*$ where stars denote Kleene closures.

Synchronization blocks have been implemented in NMF Synchronizations, an internal DSL hosted by C# [3, 4]. For the incrementalization, it uses the extensible incrementalization system NMF Expressions. This DSL is able to lift the specification of a model transformation/synchronization in three orthogonal dimensions:

- **Direction:** A client may choose between transformation from left to right, right to left or in check-only mode
- **Change Propagation:** A client may choose whether changes to the input model should be propagated to the output model, also vice versa or not at all
- **Synchronization:** A client may execute the transformation in synchronization mode between a left and a right model. In that case, the engine finds differences between the models and handles them according to the given strategy (only add missing elements to either side, also delete superfluous elements on the other or full duplex synchronization)

```

1 public void InitializeContext(IEnumerableExpression<ISample>
2     samples, ISynchronizationContext context) {
3     context.Data.Add( _platesKey, samples
4         .ChunkIndexed( 8, ( samples, column ) => new
5             ProcessColumn( column,
6                 samples.Select( tuple => new ProcessWell(
7                     tuple.Item2 % 96, tuple.Item1 ) ) ) )
8         .Chunk( 12, ( columns, plateIndex ) => new ProcessPlate(
9             $"Plate{plateIndex+1:00}", columns ) )
10        .AsNotifiable() );
11    context.Data.Add( _tubesKey, samples
12        .ChunkIndexed( 16, ( samples, tubeIndex ) => new Tubes( $
13            "Tube{tubeIndex+1:00}",
14            samples.Select( tuple => new ProcessWell(
15                tuple.Item2 % 16, tuple.Item1 ) ) ) )
16        .AsNotifiable() );
17 }

```

Listing 1: Setting up the mapping of samples to plates and wells

This flexibility makes it possible to reuse the specification of a transformation for a broad range of different use cases. Furthermore, the fact that NMF Synchronizations is an internal language means that a wide range of advantages from mainstream languages, most notably modularity and tool support, can be inherited [6].

3. Solution

This section describes key aspects of our solution. First, we describe how samples are chunked into plates and columns in Section 3.1. Next, we give a high-level overview of the actual synchronization in Section 3.2. Sections 3.3-?? describe how the different high-level protocol steps are synchronized. Section 3.6 explains the synchronization of the NEXT reference and finally, Section 3.7 shows how the synchronization is initiated.

3.1. Assignments of plates, columns and wells

As a first step, the samples to process are grouped into plates and columns. Each column consists of 8 wells that can be pipetted at the same time. This is done using the Chunk operation recently built into NMF. This comes in two versions, Chunk and ChunkIndexed where the latter also keeps the original index in the original collection. The code for calculating the assignments of samples to plates, columns and tubes is depicted in Listing 1. This listing shows how to register a collection of samples with a synchronization context. That is, because the assignment of plates is needed in many places throughout the synchronization, we put it as context.

The input type `IEnumerable<ISample>` used in Line 1 of Listing 1 denotes an incrementalizable collection of samples. NMF essentially implements the Standard Query Operators of C# and a few more operators on

```

1 public class JobRequestToJobCollection :
2     SynchronizationRule<IJobRequest, IJobCollection> {
3     public override void DeclareSynchronization() {
4         SynchronizeManyLeftToRightOnly( SyncRule<ReagentToTrough
5             >(),
6         request => request.Assay.Reagents, jobCollection =>
7             jobCollection.Labware.OfType<ILabware, Trough>()
8         );
9         SynchronizeManyLeftToRightOnly( SyncRule<
10             ProcessPlateToMicroplate>(),
11         (request, context) => GetPlates(context),
12         (jobCollection, _) => jobCollection.Labware.OfType<
13             ILabware, Microplate>() );
14         SynchronizeManyLeftToRightOnly( SyncRule<
15             SamplesToTubeRunner>(),
16         (request, context) => GetTubes(context),
17         (jobCollection, _) => jobCollection.Labware.OfType<
18             ILabware, TubeRunner>() );
19         SynchronizeManyLeftToRightOnly( SyncRule<
20             ProtocolStepToJobsRule>(),
21         (request, _) => request.Assay.Steps,
22         (jobCollection, context) => new
23             CollectionOfJobCollections( jobCollection,
24             context ) );
25     }
26 }

```

Listing 2: The entry point synchronization rule

top of this interface in order to derive an incremental change propagation for a given query. That is, the system allows developers to obtain incremental updates of the results upon changes of the input models, such as adding a sample.

Lines 3–5 in Listing 1 calculate columns as chunks of samples. These columns are then chunked into microplates. Line 6 forces the incrementalization of this collection. The idea of this order as opposed to chunking the samples into plates and then further into columns is to allow NMF to rebalance samples between columns and then rebalance columns between plates. However, we did not specify a balancing strategy and thus, NMF will not try to rebalance the chunks. Similar, lines 8–9 calculate the collection of tube runners from different chunks of the input samples.

3.2. The model synchronization

The actual model synchronization is split into several synchronization blocks that act as isomorphisms. Each synchronization rule defines a list of synchronization blocks that define what data should be synchronized. The entry point synchronization rule, depicted in Listing 2, synchronizes an overall high-level job request with a low-level job collection. In this listing, Lines 3–4 denote that the reagents are mapped to troughs and lines 5–10 denote that that tube runners should be created to host the samples as well as microplates for processing. For the tubes and the microplates, we consume a second parameter in the lens to access the plate collections stored in the context (cf. Listing 1).

In Lines 11–13 of Listing 2, we define that the steps

```

1 public class AddReagentToJobsRule : SynchronizationRule<
2     AddReagent, JobsOfProtocolStep> {
3     public override void DeclareSynchronization() {
4         MarkInstantiatingFor( SyncRule<ProtocolStepToJobsRule>()
5         );
6         SynchronizeManyLeftToRightOnly(
7             SyncRule<AddReagentLiquidTransferToLiquidTransfer>(),
8         ( step, context ) => GetPlates( context )
9         .SelectMany(p => p.Columns, (plate, column) => new
10             AddReagentLiquidTransfer(column, plate, step))
11         .Where(transfer => transfer.Column.AnyValidSample.
12             Value),
13         ( jobsOfStep, _ ) => jobsOfStep.Jobs.OfType<IJob,
14             LiquidTransferJob>() );
15     }
16 }

```

Listing 3: Synchronizing the jobs for an ADDREAGENT protocol step

of the requested assay should be synchronized with the job collections in the low-level model. For this, we use a custom collection implementation that essentially groups the low-level jobs of the resulting job collection by name. This needs access to the transformation context as it will store information such as the affected samples of a job. The four calls to `SynchronizeManyLeftToRightOnly` basically define collection-valued unidirectional synchronization blocks that are only enforced from the left to the right.

3.3. Synchronization of ADDREAGENT

The actual high-level process steps are translated using separate synchronization rules. That is, we synchronize a protocol step with the jobs implementing this protocol step. The approach to transform the other types of high-level jobs is conceptually similar, although the different complexity of the job types leads to a different complexity of the synchronization rules required. The synchronization rule for the synchronization of ADDREAGENT is depicted in Listing 3.

Line 3 marks the synchronization rule as instantiating for `ProtocolStepToJobsRule`, which means that the synchronization rule is used when the `ProtocolStepToJobsRule` is executed with an ADDREAGENT protocol step. Lines 4–9 denote the synchronization block that computes the elements from which to create the jobs, using a dedicated class to represent the request for a liquid transfer. The query calculates all columns of all plates that have at least any valid (i.e., not failed) sample.

Because the latter needs to be calculated incrementally for each `ProcessColumn`, the calculation (and its incrementalization) is separated into a static function (see Listing 4).

The reason to separate the logic into an `ObservingFunc` instance here is that the incre-

```

1 private static ObservingFunc<ProcessColumn, bool>
  _anyNonErrorSample = new ObservingFunc<ProcessColumn,
  bool>( c => c.AllSamples.Any( s => s.State !=
  SampleState.Error ) );
2 ...
3 AnyValidSample = _anyNonErrorSample.Observe( this );

```

Listing 4: Calculating whether a column has any sample that is not in the error state.

```

1 public class AddReagentLiquidTransferToLiquidTransfer :
  SynchronizationRule<AddReagentLiquidTransfer,
  LiquidTransferJob> {
2 public override void DeclareSynchronization() {
3 SynchronizeLeftToRightOnly( SyncRule<ReagentToTrough>(),
  step => step.AddReagent.Reagent, liquidTransfer =>
  liquidTransfer.Source as Trough );
4 SynchronizeLeftToRightOnly( SyncRule<
  ProcessPlateToMicroplate>(), step => step.Plate,
  liquidTransfer => liquidTransfer.Target as
  Microplate );
5 SynchronizeManyLeftToRightOnly( SyncRule<
  AddReagentTipToTipTransfer>(),
  step => step.Column.Samples
  .Where( s => s.Sample.State != SampleState.Error )
  .Select( s => new AddReagentTip( step, s ) ),
  liquidTransfer => new TipCollection( liquidTransfer.
  Tips ) );
6 SynchronizeManyLeftToRightOnly(
  ( step, _ ) => step.Column.AllSamples,
  ( liquidTransfer, context ) => GetAffectedSamples(
  context, liquidTransfer ) );
7 }
8 }
9 }
10 }
11 }
12 }
13 }
14 }

```

Listing 5: The synchronization of add reagent elements to actual LIQUIDTRANSFERJOB elements.

mentalization of a method in NMF involves some reflection and takes a bit of time while applying it to a particular element is rather cheap. Using a static instance essentially caches the incrementalization and applies it to multiple instances. For this reason, although supported by NMF, nested queries are currently rather slow and hence we refrain from using the C# query syntax in the mappings such as Listing 3.

The child synchronization rule `AddReagentLiquidTransferToLiquidTransfer` defines how the instances of this intermediate class are transformed into a low-level job as depicted in Listing 5. Line 3 defines that the source of the liquid transfer should be synchronized with the trough created for the reagent. Line 4 specifies that the reagent should be pipetted into the microplate created for the processing request. In Lines 5–9, the synchronization block denotes which tips exactly need to be created. We use an intermediate class and a custom collection again in Line 9 in order to control that a tip liquid transfer is only removed when it is still planned. Lines 10–12 specify that the samples created for this liquid transfer are stored inside the transformation context.

The synchronization rule `AddReagentTipToTipTransfer` depicted in Listing 6 specifies the transfor-

```

1 public class AddReagentTipToTipTransfer :
  SynchronizationRule<AddReagentTip, ITipLiquidTransfer
  > {
2 public override void DeclareSynchronization() {
3 SynchronizeLeftToRightOnly( well => well.AddReagent.
  Volume, transfer => transfer.Volume );
4 SynchronizeLeftToRightOnly( well => well.TargetWell.Well,
  transfer => transfer.TargetCavityIndex );
5 SynchronizeRightToLeftOnly( well => IsSampleFailed( well.
  TargetWell.Sample ), transfer => transfer.Status
  == JobStatus.Failed );
6 }
7 }

```

Listing 6: Synchronization rule `AddReagentTipToTipTransfer`

mation of tip liquid transfers. Lines 3–4 synchronize the volumes and the target cavity (the source cavity is always 0 for a trough).

The last synchronization block in line 5 specifies that the failure of the tip transfer should be synchronized back to the high-level job request model.

3.4. Synchronization of DISTRIBUTESAMPLE

The synchronization of `DISTRIBUTESAMPLE` elements works exactly like the synchronization of `ADDREAGENT` with one important exception: While the source labware of an `ADDREAGENT` is accessible easily via the transformation trace from the reagent, this is unfortunately not as easy for `DISTRIBUTESAMPLE`.

As a reason, the current design of the solution has no direct connection between a column of a processing microplate and the tube runner that holds the samples. First, we created an approach that would calculate the mapping incrementally, but this turned out to be very resource-intensive both in terms of time and memory.

```

1 private static Tubes GetSourceTube( ITransformationContext
  context, ProcessColumn column ) {
2 return GetTubes( context )
3 .AsEnumerable()
4 .FirstOrDefault( t => t.Samples
5 .AsEnumerable()
6 .Any( s => column.Samples
7 .AsEnumerable()
8 .Any( s2 => s.Sample == s2.Sample ) ) );
9 }

```

Listing 7: Calculating the tube runner for a given column of a processing plate

The solution now is to break out of the incrementalization monad explicitly and calculate the source tube runner only once as depicted in Listing 7: We explicitly call the `AsEnumerable` method here in order to instruct the compiler to actually compile the lambda expressions used to calculate the tube runner. This, however, breaks the support of rebalancing the chunks making up the columns and plates.

```

1 public abstract class MicroplateProtocolStepRule<TProtocol,
  TJobRule, TJob> : SynchronizationRule<TProtocol,
  JobsOfProtocolStep>
2 where TProtocol : IProtocolStep
3 where TJob : class, IJob
4 where TJobRule : MicroplateJobRule<TProtocol, TJob>
5 {
6     public override void DeclareSynchronization() {
7         MarkInstantiatingFor( SyncRule<ProtocolStepToJobsRule>(
8             );
9         SynchronizeManyLeftToRightOnly(
10            SyncRule<TJobRule>( ),
11            ( step, context ) => GetPlates( context )
12            .Where( plate => plate.AnyValidSample.Value )
13            .Select( plate => Tuple.Create( step, plate ) ),
14            ( jobsOfStep, _ ) => jobsOfStep.Jobs.OfType<IJob, TJob
15            >( ) );
16    }
17 }
18 public abstract class MicroplateJobRule<TProtocol, TJob> :
  SynchronizationRule<Tuple<TProtocol, ProcessPlate>,
  TJob>
19 where TProtocol : IProtocolStep
20 where TJob : IJob
21 {
22     public override void DeclareSynchronization() {
23         SynchronizeManyLeftToRightOnly(
24             ( step, _ ) => step.Item2.AllSamples,
25             ( job, context ) => GetAffectedSamples( context, job )
26             );
27         SynchronizeRightToLeftOnly(
28             step => AreAllFailed( step.Item2.AllSamples ),
29             job => job.State == JobStatus.Failed );
30         SynchronizeLeftToRightOnly( SyncRule<
31             ProcessPlateToMicroplate>( ),
32             tuple => tuple.Item2.MicroplateProperty );
33     }
34     protected abstract Expression<Func<TJob, IMicroplate>>
35         MicroplateProperty { get; }
36 }

```

Listing 8: Template for synchronization of microplate processing protocol steps

3.5. Synchronization of WASH and INCUBATE

The synchronization of WASH steps and INCUBATE steps is very similar, because both steps (as many in lab automation) operate on entire microplates. The protocol step needs to be instantiated for each microplate used for sample processing.

The synchronization rule templates for protocol steps operating on a single microplate is depicted in Listing 8. There are two rule templates, one for synchronizing a protocol step with a collection of low-level jobs, the other for actually synchronizing the protocol step for a given microplate into a given job. The `MicroplateProtocolStepRule` class marks the rule as instantiating and registers the calls to the child rule. The template for the latter, `MicroplateJobRule`, registers affected samples, sets the samples to failed (using another lens called `AreAllFailed` in Line 25) and synchronizes the target microplate. As the target metamodel does not use a shared base class for jobs operating on microplates, the rule template uses an abstract property

```

1 public class WashToJobsRule : MicroplateProtocolStepRule<
  Wash, WashToWashJob, WashJob> { }
2
3 public class WashToWashJob : MicroplateJobRule<Wash,
  WashJob> {
4     protected override Expression<Func<WashJob, IMicroplate>>
  MicroplateProperty => wash => wash.Microplate;
5     public override void DeclareSynchronization() {
6         base.DeclareSynchronization();
7         SynchronizeManyLeftToRightOnly(
8             tuple => tuple.Item2.Columns.SelectMany( c => c.Samples.
9                 Where( s => s.Sample.State != SampleState.Error ).
10                Select( s => s.Well ) ),
11             wash => wash.Cavities );
12    }
13 }

```

Listing 9: Synchronization of WASH steps

such that instance rules have to specify the property used to store the microplate.

The instantiation of the rule templates for WASH elements is depicted in Listing 9. Since the rule to synchronize WASH protocol steps is sufficiently described using the synchronization template, we do not need to provide any further specification other than the type parameters to be used, including a reference to the child rule. Unfortunately, the C# compiler is not (yet?) able to infer the type parameters `TProtocol` and `TJob`, so they must be specified explicitly.

For the synchronization of a WASH in conjunction with a specific processing plate, we need to specify the property holding the microplate (in Line 4) and handle the additional reference to the cavities that should be washed. For this, we need to override the declaration of the synchronization rule. Because we do want to inherit the declaration of the template, we need to call the base declaration in line 6. Then, we add the synchronization of the cavities in lines 7–9.

The synchronization of INCUBATE protocol steps works in the same way, except that the child rule extends the template with synchronization blocks for temperature and duration.

3.6. Synchronization of the Next reference

In order for the scheduler to be able to actually schedule the low-level jobs, the base class for jobs keeps a reference to the next and previous jobs. That is, the scheduler may only schedule a job if all previous jobs are completed and in the opposite direction, the job is a prerequisite for all next jobs.

To aid this situation, we use a utility class called `CollectionBinding` that essentially enforces the synchronization of elements between an incrementalizable source collection (typically a query) and a target collection that should be adapted. The implementation is

```

1 CollectionBinding.Create(
2     _nextJobs.Jobs.Where( j => ProtocolSynchronization.
3         GetAffectedSamples( _context, j ).Intersect( samples
4             ).Any() ),
5     item.Next )

```

Listing 10: Binding the next low-level jobs to the jobs of the next job collection that affect the same samples

depicted in Listing 10. The query calculates the jobs for which the set of affected samples intersects the affected samples of the current job. The return value is an instance of the `IDisposable` interface, the typical interface in .NET to dispose objects. In this case, the binding is stopped when disposed. Since NMF supports bidirectional references, only one direction of the the association has to be set manually, the other is set automatically by NMF.

Unfortunately, the management of the collection binding currently has to be done manually by handling change events of the jobs created for a job collection.

3.7. Starting the synchronization

To run the solution, we create a new context for the model synchronization, initialize the samples and start the model synchronization in the direction *LeftToRight* with change propagation in both directions.

```

1 _context = new SynchronizationContext( _synchronization,
2     SynchronizationDirection.LeftToRight,
3     ChangePropagationMode.TwoWay );
4 _synchronization.InitializeContext( _jobRequest.Samples,
5     _context );
6 _synchronization.Synchronize( ref _jobRequest, ref
7     _jobCollection, _context );

```

Listing 11: Starting the model synchronization

4. Evaluation

The strongest point of the presented solution is that the change propagation can be inherited mostly from a declarative specification. As a consequence, essentially all types of changes are supported, not just the change types executed by the benchmark framework. This means that new types of error handling do not necessarily have an effect on the transformation but they are supported by default. The declarative specification, however, keeps the understandability of the solution at a good level. The attendees of TTC will judge on the understandability compared to the reference solution.

In addition to the inherited change propagation, the solution also means that no changes to the metamodel code are necessary and the model representation can be reused independently of the transformation.

Before we describe the results in terms of performance, keep in mind that the reference solution is a solution tailored manually and explicitly for the given types of changes, without any incrementalization system or alike. Therefore, it is hard to beat it in terms of performance and the strengths of our solution are rather in the declarativeness and large variety of supported change types.

To evaluate the solution in terms of performance, we have run the benchmark on a system equipped with an Intel Core i7-8850H CPU clocked at 2.6Ghz and 32GB RAM, running Windows 10. The results are discussed in the remainder of this section.

4.1. Scaling Samples

The results in terms of time to execute the initial transformation for the scaling samples scenario are depicted in Figure 2. In this scenario, the different models represent loads of 8 samples (size 1) to 256 samples (size 32), applied to a simple ELISA assay model. The results show that whereas the time for the reference solution is essentially constant at around 50ms, the initial time for the NMF solution grows worse than linear, it is more like quadratic.

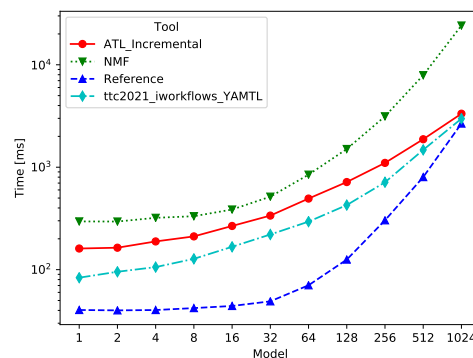


Figure 2: Time for the initial transformation in the scale samples scenario

We profiled the NMF solution. The results show that much of the time is lost because NMF Synchronizations executes the incrementalization of the queries used to specify the synchronizations over and over again instead of reusing it. Furthermore, the collection binding depicted in Listing 10 also requires the system to be incrementalized over and over again. We expect that the performance gap could be reduced, if the frameworks can be adapted to cache the incrementalization properly.

The results for propagating the state changes of low-level elements are depicted in Figure 3. This includes both changing the sample state and potentially removing low-




Figure 3: Results for the average time for an update in the scale samples scenario

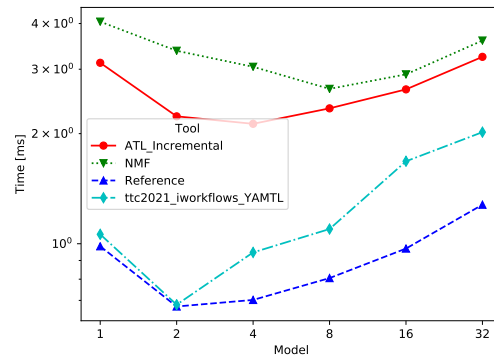


Figure 5: Results for the average time for an update in the scale assay scenario

level job elements as they have become obsolete (since all processed samples failed). Although the NMF solution is slower, it is still within few milliseconds even for the largest models considered.

4.2. Scaling the Assay

In the scale assay scenario, all model sizes use 96 samples, but the number of protocol steps varies: Whereas the smallest model (size 1) uses a simplified ELISA assay model with 8 steps, the largest model uses 32 repetitions (256 protocol steps in total). The execution time for the initial transformation is depicted in Figure 4. The runtime of the NMF solution is again quadratic in the size of the model while the reference solution remains fast.

takes longer for smaller models up to some point. This is because the benchmark framework uses Python's `subprocess.Popen` to spawn the processes for the model sizes and they can make use of JIT optimizations of earlier runs. Again, the propagation of the changes happens in a few milliseconds, both for NMF and the reference solution, the reference solution being slightly faster.

4.3. New samples

The results for the initial transformation in the new samples scenario are depicted in Figure 6. Not very surprisingly, they are similar to the scaling samples case because the parameters for the initial model are exactly the same.

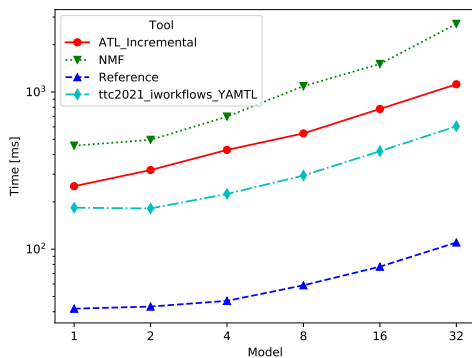


Figure 4: Time for the initial transformation in the scale assay scenario




Figure 6: Time for the initial transformation in the new samples scenario

The results for updates are depicted in Figure 5. The results look awkward because the change propagation

The difference to the scaling samples scenario is that new samples are introduced during the runtime of the




Figure 7: Results for the average time for an update in the new samples scenario

benchmark. The results for propagating these changes are depicted in Figure 7. Still, the changes are propagated within a few milliseconds, but this time this is much slower than in the other two scenarios, in line with the performance issues in the initial transformation.

5. Conclusion

The solution has shown that it is possible to derive an incremental change propagation for most of the transformation. Only for smaller parts such as the synchronization of the next reference, manual code is necessary. The evaluation shows that the performance is good, although it cannot keep up with the reference solution. In exchange, the solution supports much more types of changes. However, the solution also shows a performance problem caused by the current inability of NMF to cache the query incrementalization properly.

References

- [1] G. Hinkel, NMF: A multi-platform modeling framework, in: A. Rensink, J. Sánchez Cuadrado (Eds.), *Theory and Practice of Model Transformation*, Springer International Publishing, Cham, 2018, pp. 184–194.
- [2] G. Hinkel, R. Heinrich, R. Reussner, An extensible approach to implicit incremental model analyses, *Software & Systems Modeling* (2019). URL: <https://doi.org/10.1007/s10270-019-00719-y>. doi:10.1007/s10270-019-00719-y.
- [3] G. Hinkel, Change Propagation in an Internal Model Transformation Language, in: D. Kolovos, M. Wimmer (Eds.), *Theory and Practice of Model Transformations: 8th International Conference, ICMT*

2015, Held as Part of STAF 2015, L’Aquila, Italy, July 20–21, 2015. Proceedings, Springer International Publishing, Cham, 2015, pp. 3–17. URL: http://dx.doi.org/10.1007/978-3-319-21155-8_1. doi:10.1007/978-3-319-21155-8_1.

- [4] G. Hinkel, E. Burger, Change propagation and bidirectionality in internal transformation DSLs, *Softw. Syst. Model.* 18 (2019) 249–278. URL: <https://doi.org/10.1007/s10270-017-0617-6>. doi:10.1007/s10270-017-0617-6.
- [5] J. N. Foster, M. B. Greenwald, J. T. Moore, B. C. Pierce, A. Schmitt, Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem, *ACM Transactions on Programming Languages and Systems (TOPLAS)* 29 (2007). URL: <http://doi.acm.org/10.1145/1232420.1232424>. doi:10.1145/1232420.1232424.
- [6] G. Hinkel, T. Goldschmidt, E. Burger, R. Reussner, Using Internal Domain-Specific Languages to Inherit Tool Support and Modularity for Model Transformations, *Software & Systems Modeling* (2017) 1–27. URL: <http://rdcu.be/oTED>. doi:10.1007/s10270-017-0578-9.