

The Epsilon Solution to the KMEHR to FHIR Case

Antonio Garcia-Dominguez¹

¹University of York, York, UK

Abstract

In recent years, there has been a push towards standardisation of interchange formats for health records: this case tackles the problem of mapping Patient Summarized Records between the Belgian KMEHR and international HL7 FHIR standards. The case proposed two evaluation dimensions: time/space efficiency, and understandability. This paper presents a port of the original ATL rules to the Epsilon Transformation Language (ETL), with a focus towards helping the user understand the transformation via transient visualisations using the Picto tool bundled with recent versions of Epsilon. While ETL is not as fast as ATL, the ETL script is 1,096 lines long, compared to the 1,319 lines of ATL. At the same time, during the development of this solution, a few improvements were identified and made to Epsilon itself.

Keywords

ATL, ETL, healthcare information systems, model transformation, model visualisation

1. Introduction

As described in the original case, allowing different healthcare systems to efficiently share information with each other is increasingly important to deliver rapid and effective care (e.g. when needing medical treatment in a different country). This has motivated the creation of a number of specifications to exchange the information that the system has about a patient. The Belgian KMEHR (Kindly Marked-Up Electronic Health Care Record)¹ is composed of an XML Schema describing a message grammar, a recognized set of medical transactions that use that grammar, and a set of reference tables with values to be used in those messages. Health Level 7 (HL7) is a set of international standards for exchanging health records, which includes the FHIR (Fast Healthcare Interoperability Resources) standards framework.

The original case description focused on transforming one specific type of healthcare record, known by HL7 as an International Patient Summary (IPS)². An IPS is supposed to follow a patient across borders and make it possible to quickly provide the information needed for unplanned, cross-border care. The case description included a reference implementation with over 1,300 lines of ATL code, and specified two evaluation dimensions for alternative solutions: 1) their efficiency in time and space (in system memory), and 2) the understandability of the transformations.

This paper presents a solution based on the Eclipse Epsilon family of model management languages, available from Github³. The transformation itself is a close port of the original ATL rules to the Epsilon Transformation Language (ETL) [1], but it is complemented by the use of the Epsilon Picto [2] tool to produce interactive visualisations of its transformation trace. While ETL does not outperform the reference solution, it does take up fewer lines of code (1,096 lines, compared to the 1,319 of ATL), and its interactive visualisations can help understand the transformation and correct issues. The solution passes all the correctness checks provided by the case authors, and automated tests check that the produced FHIR models are line-by-line equivalent to the expected ones (except for changes due to unique IDs which change between executions).

The rest of the paper is structured as follows. Section 2 explains the ETL model-to-model transformation. Section 3 describes how the internal ETL transformation traces are serialised into a transformation trace model and then visualised with Picto. Section 4 presents the performance results of the solution and compares them against those of the reference ATL solution. Finally, Section 5 provides general conclusions for this paper, and points out several areas in which this work can be refined.

2. Model transformation

The original case description did not provide a detailed specification of the transformation involved, and the author was unfamiliar with KMEHR and FHIR. For that reason, the approach taken was to translate the ATL rules to ETL, given their similar approaches.

Due to time limitations, the ETL script is contained in a single `.et1` file that is 1,096 lines long, with a **pre**

TTC'23: Transformation Tool Contest, Part of the Software Technologies: Applications and Foundations (STAF) federated conferences, Eds. A. Boronat, A. García-Domínguez, and G. Hinkel, 20th July 2023, Leicester, UK.

✉ a.garcia-domínguez@aston.ac.uk (A. García-Dominguez)

🆔 0000-0002-4744-9150 (A. García-Dominguez)

© 2023 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.ehealth.fgov.be/standards/kmehr/en>

²<https://build.fhir.org/ig/HL7/fhir-ips/>

³<https://github.com/agarciadom/ttc2023-kmehr2fhir>

Listing 1: Excerpt of the Posology rule in ETL

```

1 rule Posology
2   transform s: KMEHR!PosologyType
3   to t: FHIR!MedicationStatement, msid: FHIR!Id,
4     /* ... */
5   {
6     var i = s.eContainer();
7     var tx = i.eContainer();
8     var f = tx.eContainer();
9     // ...
10  }

```

(pre-execution) block being used to initialise the various predefined mappings from special values in KMEHR to FHIR, and instantiate the XMLTYPEFACTORY needed to deal with certain data types (i.e. dates).

Rules were mostly mapped 1-to-1 from the ATL sources, but some adaptations were needed due to subtle differences between ATL and ETL, which will be listed below.

2.1. Assignment and equivalent mapping

The ATL `<-` operator is similar to the ETL `:=` operator, in that it will assign to the l-value the results of transforming the r-value. However, `:=` cannot be used with objects that are not to be transformed (e.g. strings), and therefore users must be careful to use `=` or `:=` depending on the situation.

ETL does not have an equivalent to the ATL `mapsTo` keyword: in an ETL rule, every object in the `to` clause is considered equivalent to the source object. This implies that any calls to `equivalent()` or uses of the `:=` operator must be careful to filter the r-value so only objects compatible with the l-value are left.

2.2. Number of source objects per rule

ETL only allows one source object per rule, unlike ATL which allows for multiple source objects. While ATL required using a local search compiler to avoid multiple scans of the model to obtain pattern matches, ETL rules were simply written to apply to the innermost object, and use `x.eContainer()` to access its containers as needed. Comparing the Posology rules in ETL (Listing 1) and ATL (Listing 2), the ETL approach is simpler to specify and compute.

2.3. Helper operations and lazy rules

ATL helper operations were translated to Epsilon Object Language operations, in generally a very straightforward

Listing 2: Excerpt of the Posology rule in ATL

```

rule Posology {
  from
    f : KMEHR!FolderType,
    tx : KMEHR!TransactionType,
    i : KMEHR!ItemType,
    s : KMEHR!PosologyType (
      i.posology = s and
      tx.item->includes(i) and
      f.transaction->includes(tx) and
      i.isMedication
    )
  to
    t : FHIR!MedicationStatement mapsTo s (
      // ...
    ),
    // ...
}

```

way (with some simplifications, as the Epsilon languages do not have the `.` and `->` distinction that ATL has for accessing members, which is awkward to use for most new users). In addition, ETL allows for manually specifying whether the operation results should be cached for future calls (to save time or produce consistent results, like `uuid()` which generates random Universally Unique Identifiers or UUIDs for a given FHIR object), or not (to save memory): this level of control is not apparent in ATL.

Some ATL lazy rules were translated to ETL lazy rules (e.g. the rule hierarchy dedicated to `COMPOSITION-SECTIONS`), whereas others were turned into EOL operations as it wasn't deemed useful to pollute transformation traces with them (e.g. the creation of `FHIRSTRINGS`). Note that some basic FHIR types did use lazy rules, precisely for populating the transformation trace, e.g. `FHIRDATE` and `FHIRDATETIME`. It is important to note that ETL suffers a major performance hit as soon as any lazy rule is used, as mentioned in its official documentation⁴: due to time limitations, it was not possible to re-engineer the transformation to completely avoid them (e.g. by adding the appropriate guards).

2.4. Rule inheritance

It was noted that rule inheritance was somewhat different in ETL and ATL, when both the base and extending rule are non-abstract. Initial inspection of results suggests that ATL will execute the base and extending rule (if both applicable) on the same objects, whereas ETL will end up producing a different set of objects.

⁴<https://eclipse.dev/epsilon/doc/etl/>

Listing 3: Rule inheritance in ATL with multiple non-abstract rules

```

1 rule SumEHRTransaction {
2   from
3     f : KMEHR!FolderType,
4     s : KMEHR!TransactionType (
5       f.transaction->includes(s) and
6       s.cd->exists(cd | cd.value = 'sumehr')
7     )
8   to /* ... */
9
10  rule SumEHRTransactionWithAuthor
11    extends SumEHRTransaction {
12      from
13        f : KMEHR!FolderType,
14        s : KMEHR!TransactionType (
15          not s.txAuthor.oclIsUndefined()
16        )
17      to /* ... */
18
19    rule SumEHRTransactionWithCustodian
20      extends SumEHRTransaction {
21        from
22          f : KMEHR!FolderType,
23          s : KMEHR!TransactionType,
24          i : KMEHR!ItemType (
25            s.item->includes(i) and
26            i.cd->exists(cd | cd.value = 'gmdmanager')
27          )
28        to /* ... */

```

Listing 4: Adapting multiple non-abstract rule inheritance of ATL to ETL

```

1 rule SumEHRTransaction
2   transform s: KMEHR!TransactionType
3   to t: FHIR!Composition,
4     cid: FHIR!Id,
5     patRef: FHIR!Reference,
6     cStatus: FHIR!CompositionStatus,
7     dateTime: FHIR!fhir::DateTime
8   {
9     guard: s.cd.exists(cd | cd.value = 'sumehr')
10    /* ... */
11    if (s.txAuthor().isDefined()) { /* ... */
12    if (s.custodianItem().isDefined()) { /* ... */
13  }

```

Listing 5: Mapping ATL enumeration literals

```

helper def : genderMap : Map(
KMEHR!CDSEXvalues,
FHIR!AdministrativeGenderEnum) =
Map {
  (#changed, #other),
  (#female, #female),
  (#male, #male),
  (#unknown, #unknown),
  (#undefined, #other)
};

```

Listing 6: Mapping ETL enumeration literals

```

var genderMap = Map {
KMEHR!CDSEXvalues#changed
  = FHIR!AdministrativeGenderEnum#other,
KMEHR!CDSEXvalues#female
  = FHIR!AdministrativeGenderEnum#female,
KMEHR!CDSEXvalues#male
  = FHIR!AdministrativeGenderEnum#male,
KMEHR!CDSEXvalues#unknown
  = FHIR!AdministrativeGenderEnum#unknown,
KMEHR!CDSEXvalues#undefined
  = FHIR!AdministrativeGenderEnum#other
};

```

For instance, as shown in Listing 3 the ATL script had a base non-abstract `SUMEHRTRANSACTION` rule, extended by `SUMEHRTRANSACTIONWITHAUTHOR` and `SUMEHRTRANSACTIONWITHCUSTODIAN`. In the ETL script (as shown in Listing 4), these three rules were combined into one, where the body contained additional `if` blocks implementing the behaviour of the subclasses.

2.5. Enumeration literals

ATL has a more convenient syntax for specifying enumeration literals: where the type is known from the context, it is possible to simply use `#changed` to refer to that literal. The Epsilon languages lack this capability, requiring mentioning at least the name of the enumeration that contains the literal (e.g. `CDSEXvalues#changed`). This is awkward when writing complex expressions with one-off use of enumeration literals in ETL, as often the modeller has to stop writing and navigate the metamodels to find out the relevant enumeration to mention. It also introduces unnecessary repetition when specifying mappings between enumeration literals in the source and target metamodels.

The author has filed a pull request adapting this capa-

bility to the Epsilon languages⁵, but due to time limitations it has not yet been reviewed and merged into the main branch. Once it is merged, it should be possible to further simplify the ETL script in this regard.

2.6. Java wrapper and test oracle

In order to integrate with the benchmark framework and simplify its use, the transformation has been wrapped into a Java class called `TRANSFORMATION`, similarly to the reference solution. The class provides a `run()` method for launching the transformation, as well as a method to specify if a trace model is desired (c.f. Section 3). The `DRIVER` class from the reference solution (which handles the environment variables from the Python `run.py` script) was largely adopted as-is, only changing the ATL-specific parts to use the new `TRANSFORMATION` class.

In addition, the correctness of the transformation was checked with JUnit 5 tests that performed line-based differencing between the expected models and the obtained models using the `java-diff-utils` library⁶, while discarding changes due to UUIDs. EMF Compare was initially considered for comparing the expected and obtained models, but was discarded as it produced too many spurious differences in its standard configuration.

3. Generation and visualisation of transformation traces

Having ported the transformation to ETL, the next goal was to help users understand the relationships between the original KMEHR model and the generated FHIR model. It was decided to use Epsilon Picto [2] to develop an interactive visualisation of the traces of the model transformation, which relate source and target objects, and rules.

3.1. Trace metamodel

Whereas ATL can automatically generate a trace model if a “trace” model is added to its launch configurations (based on a pre-existing trace metamodel), ETL only keeps an internal data structure (the `TRANSFORMATIONTRACE`) and leaves its serialisation into a model up to the user (e.g. via a post-execution `post` block)⁷. It was necessary to design and implement a trace metamodel, as well as the code needed to populate trace models from the internal `TRANSFORMATIONTRACES`.

⁵<https://github.com/eclipse/epsilon/pull/42>

⁶<https://github.com/java-diff-utils/java-diff-utils/>

⁷Incidentally, the `ETLTRANSFORMATIONTRACE` was internally revamped while developing this solution, replacing its original data structure (a flat list) with a faster, lookup-based data structure: <https://github.com/eclipse/epsilon/pull/44>

The metamodel for ETL traces is shown in Figure 1, and has been designed specifically for this problem. The root of the model is a `TRACE` from a source model to a target model (both represented by their Unique Resource Identifiers or URIs). These contain a forest of `SOURCEOBJECTS`, a forest of `TARGETOBJECTS`, and a list of `TRANSFORMATIONRULES`. `SOURCEOBJECTS` and `TARGETOBJECTS` are both `MODELOBJECTS`, for which we keep their package URI, their EClass name, and their URI fragment within their model. For `TRANSFORMATIONRULES`, their name, location, and whether they are lazy or not is kept. Finally, we keep track of all the `TRANSFORMATIONS` of every `SOURCEOBJECT`, and the `TRANSFORMATION` that produced a given `TARGETOBJECT`.

The use of forests to reproduce the containment trees of the source and target models is to allow the trace model to be entirely standalone for visualisation. This was mostly due to practical reasons: the KMEHR and FHIR model parsers had been provided as Maven artifacts rather than Eclipse plugins, and so it was not possible within the allotted time to browse KMEHR and FHIR models using standard EMF tree-based model viewers. Given more time, it would have been possible to produce EMF tree-based model viewers for KMEHR and FHIR models, and only keep the relationships between source objects, target objects, and rules (similarly to `TRACELINKS` in ATL’s predefined trace metamodel).

3.2. Trace generation

Rather than using an `post` block in the ETL script to populate a trace model, it was decided to write this code as part of the Java `TRANSFORMATION` class, to make it easier to integrate into future versions of Epsilon (after generalising its trace model to handling multiple source/target models).

The algorithm largely consists of these major steps:

1. Create the forest of `SOURCEOBJECTS`, starting from the direct contents of the source model. If the source object has been involved in a transformation, ensure the appropriate `TRANSFORMATIONRULE` exists and is populated from the ETL rule, create and populate its `TRANSFORMATION` object, and temporarily associate its target objects to the `TRANSFORMATION` object (e.g. via a throwaway `IDENTITYHASHMAP`).
2. Create the forest of `TARGETOBJECTS`, starting from the direct contents of the target model. If the `TARGETOBJECT` has been involved in a transformation (which we should know from the associations in the prior step), have it refer to the transformation.
3. To manage the size of the source and target object forests, these are now pruned: an object is

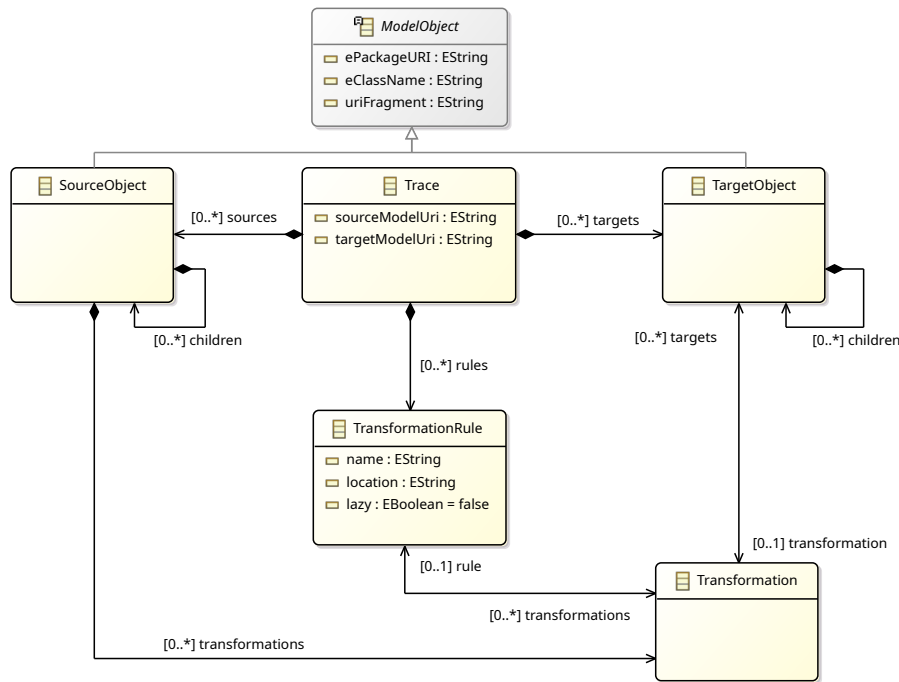


Figure 1: Trace metamodel for ETL

Listing 7: Sample .picto file

```

1 <?nsuri picto?>
2 <picto format="egx" transformation="platform:/resource/etl/src/main/egx/txtrace.egx">
3 </picto>

```

only kept if itself or one of its descendants was involved in a transformation. This removes many of the objects that are unrelated to an ETL rule (e.g. FHIRSTRINGS), and helps reduce the amount of information shown to the user.

4. The TRANSFORMATIONRULES in the trace are sorted by name, for easier navigation.

3.3. Trace visualisation

Once the trace models were available, Epsilon Picto was used to visualise them in an interactive manner. Picto is implemented as an Eclipse plugin which provides a user interface divided into two parts: a left part with a tree of selectable views, and a right part with an embedded web browser to show the visualisation generated from the selected view.

The view tree is generated by executing an EGX (EGL Co-Ordination Language) script with a number of rules, specified by the user by adding a file with the same name

as the model plus a .picto suffix. An example for one of the traces in this solution is shown in Listing 7.

An example EGX rule for Picto is shown in Listing 8. It indicates that every SOURCEOBJECT should be listed as an item in the view tree under a certain path, with an icon dependent on whether it participated in a transformation or not. If it were selected, Picto would use the EGL `source2graphviz.egl` to generate a file which would be rendered into a visualisation by using the `circo` algorithm in the Graphviz⁸ tool. It also specifies an (optional) set of layers which can be toggled by the user: in this case, there is a layer for controlling whether lazy rules are shown or not (these layers need to be implemented in the invoked EGL scripts). Finally, there are some additional keys (`focus`, `source`, and `trace`) which are made available as local variables to the EGL script.

Similar rules have been defined for TARGETOBJECT and TRANSFORMATIONRULE. Figure 2 shows what the Picto visualisation for a given rule looks like. The RULE

⁸<https://graphviz.org/>

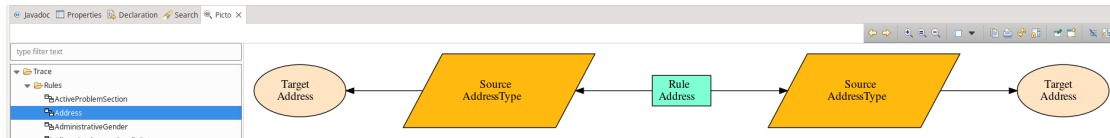


Figure 2: Screenshot of Picto visualising a TRANSFORMATIONRULE

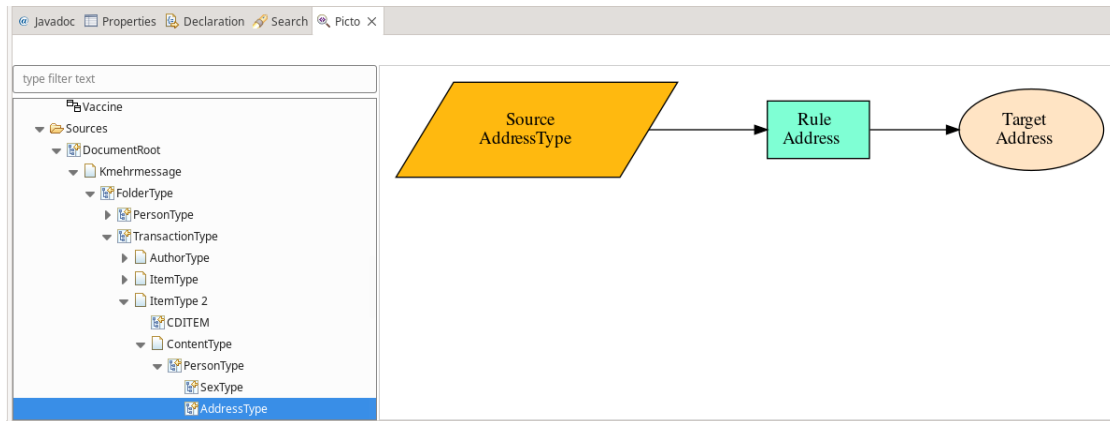


Figure 3: Screenshot of Picto visualising a SOURCEOBJECT

is shown as an aquamarine rectangle, SOURCEOBJECTS are orange parallelograms, and TARGETOBJECTS are light orange (“bisque”) ellipses. Figure 3 and Figure 4 are visualisations of the source and target objects shown in Figure 2. It is important to note that these visualisations are interactive: users can click on a rule, source object or target object in the graphical visualisation, and it will automatically select and visualise the relevant view.

With this configuration of Picto, it is possible for users to see at a glance which objects participated in a transfor-

mation (from the icons), and then drill down into which objects they produced / were produced from, and which rules were involved. In fact, it was shortly after using this visualisation that the author noticed some “orphan” objects were being left out of the DOCUMENTROOT of the target model, found the rules involved in their creation, and fixed the issues producing these orphan objects. Picto is designed with the idea that a visualisation quickly becomes overwhelming when the entirety of the model is displayed at once, and instead focuses on the immediate

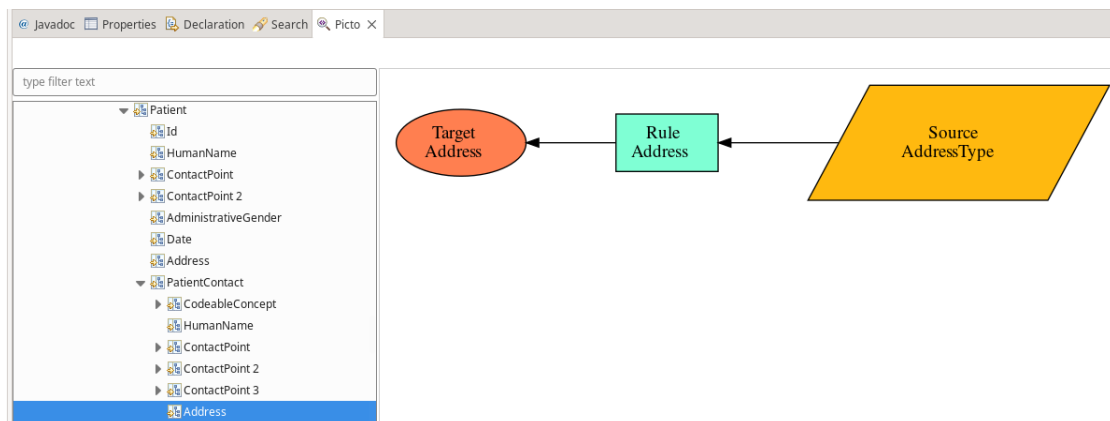


Figure 4: Screenshot of Picto visualising a TARGETOBJECT

Listing 8: Sample EGX rule for visualising TARGET-OBJECTS

```

1 rule SourceObject2Graphviz
2   transform sob: SourceObject
3   {
4     template : "source2graphviz.egl"
5     parameters : Map {
6       "path" = sob.path(),
7       "icon" = sob.transformations.isEmpty()
8         ? "document" : "tree_left",
9       "format" = "graphviz-circo",
10      "layers" = Sequence {
11        Map { "id"="lazy", "title"="Show lazy rules",
12          "active"=true }
13      },
14      "focus" = sob,
15      "source" = sob,
16      "trace" = trace::Trace.all.first
17    }
18  }

```

neighbours of the selected element (as in Figures 2 to 4).

This interactivity is achieved by having the EGL scripts produce HTML links that trigger the `showView` JavaScript function contributed by Picto into the embedded web browser, which takes the tree path of the view to be selected. There was some brief experimentation with the `showElement` JavaScript function which allows for selecting a model element in an appropriate editor (e.g. to directly jump to the KMEHR/FHIR model element in question), but this did not work as expected due to the fact that the KMEHR/FHIR Maven artifacts were not contributing any tree-based model editors to Eclipse (as mentioned at the end of Section 3.1).

4. Benchmark results

Having developed the transformations and visualised the results, the last part of the solution was to compare its time and space (memory) usage and compare it against the ATL reference solution. For performance comparisons, the Python script included with the reference solution was used, except for producing plots comparing the results. Since the R-based reporting scripts did not seem to have been updated for the KMEHR to FHIR case study, a new set of R scripts (using Packrat for reproducible R package management) were created in the `src/main/viz` directory of the ETL solution. Each combination of tool and model size (1, 10, 100 and 1000) was run 3 times, and execution times and memory usages were averaged.

The results are shown in Figures 5 and 6. ETL did not outperform ATL in any scenario: while it took less time and memory in its Initialization and Load phases, it was approximately 10 times slower in the two smallest model sizes, 6 times slower with model size 100, and over 2 times slower with the largest model size. Looking at the graph, there may be some indications that the performance gap may be covered as models become larger, but this would require additional experimentation. As mentioned above, ETL had to switch to a slower execution strategy due to the use of lazy rules, as ported over from the ATL reference solution.

On the other hand, both tools used similar amounts of memory, with ETL using slightly more (around 20% more for the largest model size) but still only reaching about 142.50MiB. Memory does not seem to be a bottleneck in this scenario.

5. Conclusions and future work

In this paper, a port of the reference implementation to the Epsilon Transformation Language was presented and evaluated. ETL was extended with code that produced trace model similar in spirit to those produced by ATL, with the difference that the source and target model containment forests are represented in the trace model, after some pruning to remove subtrees that do not involve any transformations. The trace model was used with Picto to provide interactive visualisations of the transformation trace, which was helpful in detecting and correcting some of the last remaining defects in this transformation.

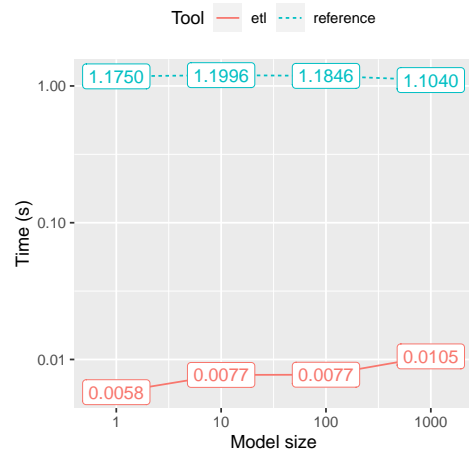
ETL did not outperform ATL in this scenario: it is believed that a major contributing factor may be the use of lazy rules, as ported from the original ATL script. It would be interesting to re-engineer the transformation to avoid lazy rules, and compare again the performance of ETL against ATL.

One line of work would be to take the generation of trace models in this solution and merge it into the Epsilon core, after some generalisation of the trace metamodel. An appropriate Picto visualisation could be bundled together with this feature, inspired by this solution.

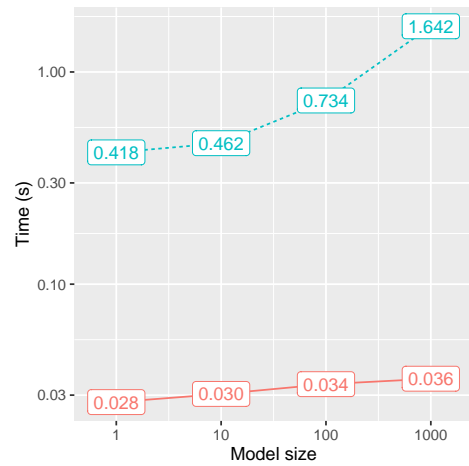
Finally, there were some improvements done to Epsilon during the development of this solution: the internal ETL `TRANSFORMATIONTRACE` data structure was optimised, and a pull request for adding ATL-style partial enumerations is currently under review in the Eclipse Epsilon project. There may be other opportunities to take advantage of more advanced data structures within Epsilon and improve performance (e.g. Guava caches, multimaps, and tables).

References

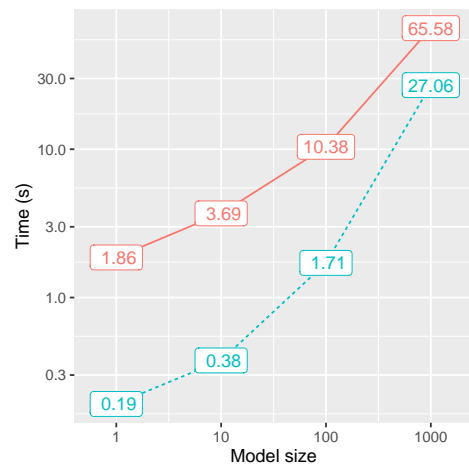
- [1] D. S. Kolovos, R. Paige, F. Polack, The Epsilon Transformation Language, in: Theory and Practice of Model Transformations, First International Conference, ICMT 2008, Zürich, Switzerland, July 1-2, 2008, Proceedings, 2008, pp. 46–60. URL: http://dx.doi.org/10.1007/978-3-540-69927-9_4. doi:10.1007/978-3-540-69927-9_4.
- [2] D. Kolovos, A. de la Vega, J. Cooper, Efficient generation of graphical model views via lazy model-to-text transformation, in: Proceedings of the 23rd ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, MODELS '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 12–23. URL: <https://doi.org/10.1145/3365438.3410943>. doi:10.1145/3365438.3410943.



(a) Initialization phase

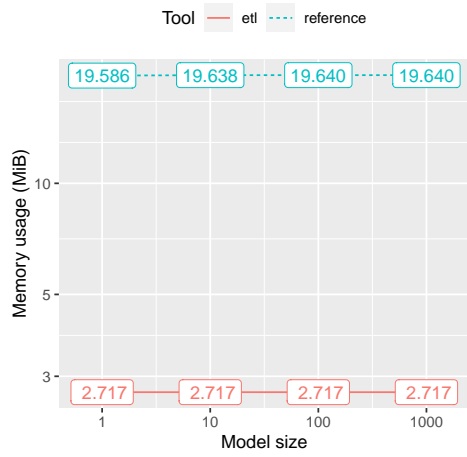


(b) Load phase

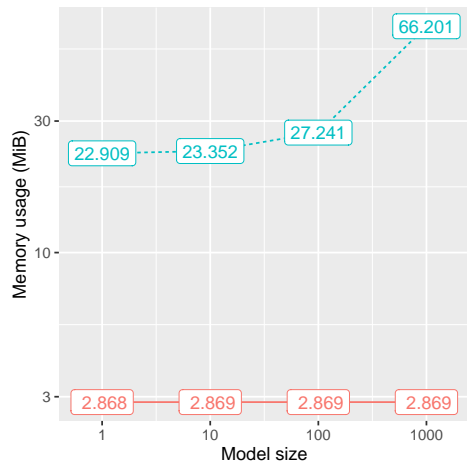


(c) Run phase

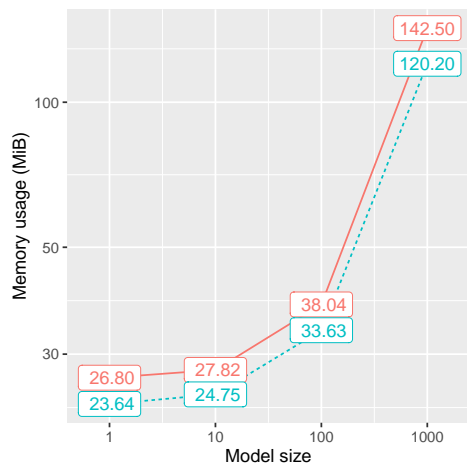
Figure 5: Average execution times in seconds, by phase



(a) Initialization phase



(b) Load phase



(c) Run phase

Figure 6: Average memory usage in MiB, by phase