

The TTC 2019 TT2BDD Case (v1.1)

Antonio García-Domínguez
Aston University
B4 7ET, Birmingham, United Kingdom
a.garcia-dominguez@aston.ac.uk

May 26, 2019

Model transformation tools have reached a considerable level of maturity in the core features, and are currently developing in many directions. Some tools are focusing on providing higher performance for large models or complex transformations. Others focus on bidirectionality, visualisation, traceability, or verifiability, among other research directions. Whereas past cases in TTC have focused on specific research directions, this case study presents a well-known simple transformation and welcomes researchers to apply their research to it. The aim of this case is to serve as a showcase of the various directions that model transformation research is going towards at the moment.

1 Introduction

Past editions of the Transformation Tool Contest have focused on a variety of topics:

- In 2018, the Quality-based Software Selection and Hardware-Mapping case [4] discussed optimisation-oriented model transformations (with a combination of performance and solution quality). The Social Media Live Case [7] considered performance in updating model views as models changed (with a strong preference for approaches supporting incrementality).
- In 2017, the Smart Grid case focused on incrementality [6], the Families to Persons case discussed bidirectional transformations [1], State Elimination focused on performance and the live case on Transformation Reuse [5] discussed mechanisms to share complex logic across multiple versions of a transformation.
- In 2016, optimisation-oriented model transformations were discussed in considerable breadth through the Class Responsibility Assignment case [2], and an alter-

native dataflow-based notation for model transformation was evaluated in the live case study [3].

While these were notable examples of realistic transformations, they were narrowly focused on a specific topic, and their inherent complexity discouraged some attendees from trying their hand with their own research agenda on the transformation.

This year, we propose a broader contest that welcomes all active lines of work on model transformation. It is based on a simpler, well-known transformation from the ATL Zoo [8]: TT2BDD (Truth Tables to Binary Decision Diagrams). Striving for raw performance is an option, but the case welcomes approaches that focus on other attributes of interest of a high-quality model transformation: for example, verifiability, traceability, bidirectionality, or understandability. Solution providers are welcome to propose their own attributes of interest. In general, this case is proposed as a showcase of the current variety of model transformation tools.

All resources for this case are available on Github¹. Please follow the link in the footnote and create a pull request with your own solution after you have submitted your description to EasyChair.

The rest of the document is structured as follows: Section 2 describes the TT2BDD transformation. Section 3 several tasks of interest that should be tackled in a solution (authors are free to propose their own tasks of interest). Section 4 mentions the benchmark framework for those solutions that focus on raw performance. Finally, Section 5 mentions an outline of the initial audience-based evaluation across all solutions, and the approach that will be followed to derive additional prizes depending on the attributes targeted by the solutions.

2 Transformation Description

This section introduces the Truth Tables to Binary Decision Diagram transformation, with a description of the input and output metamodels, and an outline of an implementation.

2.1 Input Metamodel: Truth Tables

The input metamodel is shown on Figure 1. The TRUTHTABLE class is as the root of the model, and contains a collection of PORTS and ROWS. PORTS come in two types: INPUTPORTS and OUTPUTPORTS. ROWS contain sequences of CELLS, which assign values to the INPUTPORTS and OUTPUTPORTS of the table.

Automated EMF opposite references are used liberally in the metamodel to simplify the specification of the transformation. *owner* is used to access the container from several child objects (e.g. the TRUTHTABLE of a ROW), and it is possible to see which *cells* referenced a specific PORT.

¹<https://github.com/TransformationToolContest/ttc2019-tt2bdd>

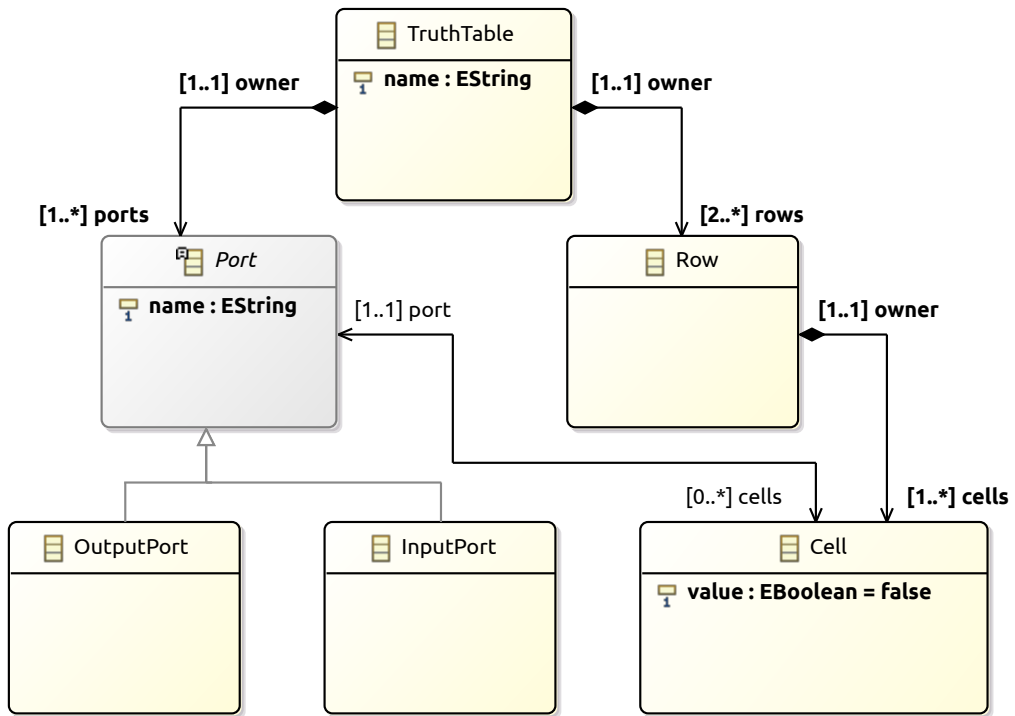


Figure 1: Class diagram for the input Truth Tables metamodel

<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>S</i>
0	0	-	-	0
0	1	0	0	1
0	1	0	0	1
0	1	0	1	0
0	1	1	-	0
1	0	0	0	0
1	0	1	0	1
1	-	-	1	0
1	1	0	0	1
1	1	1	0	0

Table 1: Example truth table: *A* to *D* are input ports and *S* is an output port. “-” means “ignored”.

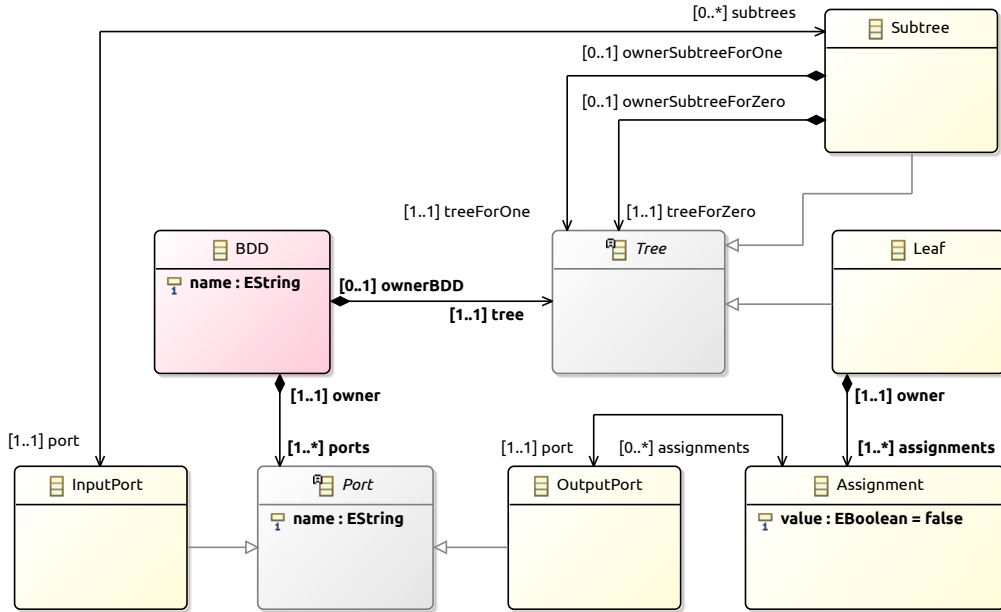


Figure 2: Class diagram for the output Binary Decision Diagram metamodel, in the original tree-based version.

A sample model is shown on Table 1. The truth table has four INPUTPORTS named *A* to *D*, and one OUTPUTPORT named *S*. The first ROW contains only 3 CELLS, specifying that if *A* and *B* are 0, then *S* should be 0.

2.2 Output Metamodel: Binary Decision Trees

The original output metamodel from the ATL Zoo version is shown on Figure 2. The BDD class is the root of the model, and contains the root of the TREE and a collection of PORTS. Similarly to the Truth Tables metamodel, there are INPUTPORTS and OUTPUTPORTS.

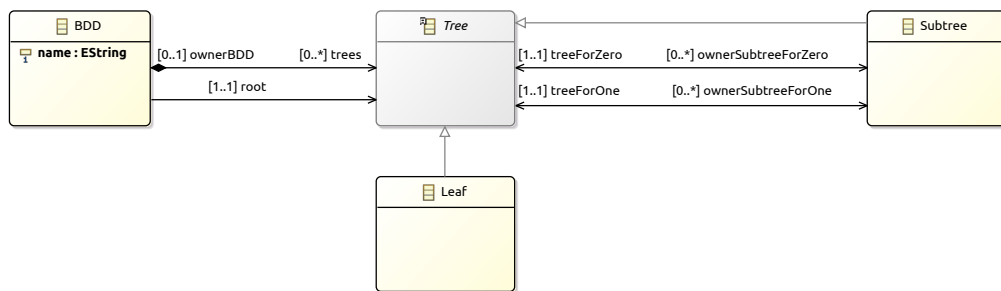


Figure 3: Changed classes in the graph-based variant of the Binary Decision Diagram metamodel (2019-05-26). Omitted classes remain the same.

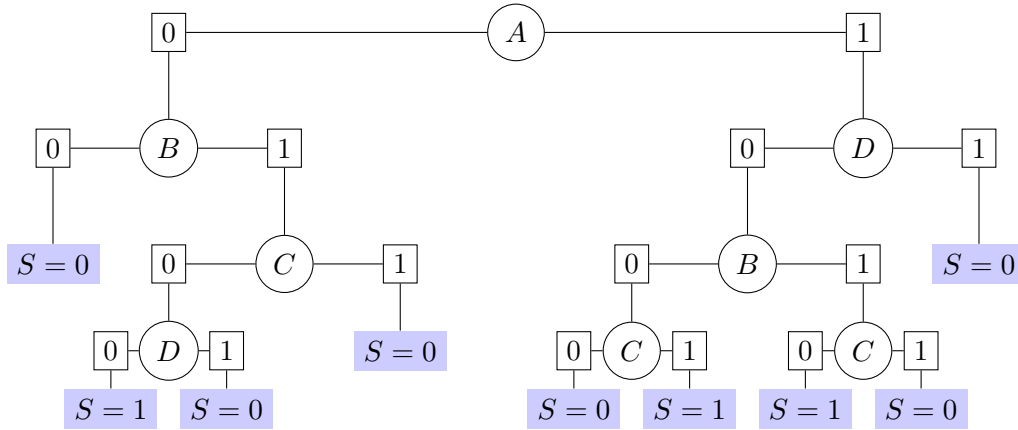


Figure 4: Equivalent BDD for the truth table on Table 1

Figure 3 shows the changes introduced into a revised version of the original metamodel, which allow the BDD to be a rooted acyclic graph and not just a tree. This would allow the BDD to reuse common subtrees, which would produce more compact circuits in some situations. The case provides both versions as separate metamodels: solution implementers are free to choose either version.

TREE is the common superclass for any node in the tree. Inner nodes check the value of an INPUTPORT: if it is a false value, evaluation will proceed through the *treeForZero* TREE; otherwise, evaluation will go through the *treeForOne* TREE. Leaf nodes are LEAF objects, which provide an ASSIGNMENT of a boolean value to each of the available OUTPUTPORTS.

The equivalent BDD for the truth table on Table 1 is shown on Figure 4. SUBTREES are represented by the circle referencing an INPUTPORT and their subtrees for when the port takes a 0 or 1 value. ASSIGNMENTS are represented by the highlighted nodes that provide values to the OUTPUTPORTS.

2.3 Process Outline

The transformation essentially needs to construct a (preferably minimal) binary decision diagram that produces the same values for the output ports, given the same values in the input ports. Given the high interest about this problem in circuit design, many algorithms have been proposed in the literature.

Solution authors are welcome to implement a more optimal algorithm in their transformation tool. This section will only outline a simple approach that can be readily implemented without too much complication.

There are some basic mappings which are immediately obvious:

- Each TRUTHTABLE object should correspond to a BDD object, with the same name and equivalent PORTS.

- Each INPUTPORT and OUTPUTPORT should be mapped to an object of the BDD type with the same name, and should be assigned the same *name*.
- Each ROW should become a LEAF node: the CELLS for the OUTPUTPORTs will become ASSIGNMENTS.

The complexity is in deriving the inner nodes: the SUBTREE objects. One simple approach is to find a TT INPUTPORT which is (ideally) defined in all the ROWs, and turn it into an inner node (a SUBTREE) which points to the equivalent BDD INPUTPORT and has two TREES:

- The zero subtree, produced from the ROW(s) where the port was 0. This will be a SUBTREE if there are at least two rows in that situation: the transformation should proceed recursively in this case with those rows, excluding the input ports that have already been considered. If there is only one such row, this would simply point to the LEAF created above.
- The one subtree, produced in a parallel way to the zero subtree.

This simple approach does not necessarily ensure a minimal subtree, as in some points there may be multiple ports to choose from. It may require improvements for cases where there are no input ports which are defined in all available rows. Authors are welcome to implement a more efficient or optimal approach in their solutions.

3 Main Task

The obvious main task in this case is to implement the TT2BDD transformation in your favourite model transformation language, ensuring that the resulting BDDs are equivalent to the original truth table. To simplify the work involved, the case includes an Eclipse Modelling Framework implementation of the metamodels in the `ttc2019.metamodels` project within the `metamodels` folder, and a set of sample XMI input models conforming to the metamodels in the `models` folder.

The `models` folder includes two additional tools, which can be run with `java -jar tool.jar`:

- `generator.jar` can produce truth tables of an arbitrary number of input and output ports, given seed. Solution implementers may want to be careful with the values given, as truth tables grow exponentially as more input ports are added.
- `validator.jar` can validate if a BDD model is equivalent to a source truth table model, by evaluating all the rows in the TT model through the BDD and comparing the values of the OUTPUTPORTs. This tool is required to check that your transformation is correct. The tool accepts both tree-based and graph-based BDD models.

Solutions are free to focus on any specific quality attribute of the transformation beyond optimality and performance. For instance, it may be useful to be able to prove that the transformation does indeed produce a BDD which is equivalent to the TT (even if suboptimal). It may also be interesting to be able to visualize the BDD being built from the TT, or to be able to trace which subtree came from which collection of rows.

The conciseness of the transformation may also be of interest, given that its recursive nature may be difficult to map to some existing transformation languages.

4 Benchmark Framework

If focusing on performance, the solution authors should integrate their solution with the provided benchmark framework. It is based on that of the TTC 2017 Smart Grid case [6], and supports the automated build and execution of solutions. For this specific case study, the visualisation of the results is currently disabled.

The benchmark consists of three phases:

1. **Initialization**, which involves setting up the basic infrastructure (e.g. loading metamodels). These measurements are optional.
2. **Load**, which loads the input models.
3. **Run**, which runs the transformation, produces the BDD and saves the file.

4.1 Solution requirements

Solutions should be forks of the main Github project², and should be submitted as pull requests after the descriptions have been submitted to EasyChair. For details about the submission format, please consult the call for solutions³.

Each solution wishing to use the benchmarking framework should print to the standard output a line with the following fields, separated by semicolons (“;”):

- **Tool**: name of the tool.
- **Model**: name of the model being transformed.
- **RunIndex**: the index of the repetition of the transformation.
- **Phase**: the name of the phase.
- **MetricName**: the name of the metric. It may be the used **Memory** in bytes, or the wall clock **Time** spent in integer nanoseconds.

To enable automatic execution by the benchmark framework, solutions should add a subdirectory to the `solutions` folder of the benchmark with a `solution.ini` file stating how the solution should be built and how it should be run. As an example, the

²<https://github.com/TransformationToolContest/ttc2019-tt2bdd>

³<https://www.transformation-tool-contest.eu/cfs.html>

Listing 1: `solution.ini` file for the reference ATL solution

```
1 [build]
2 default=./gradlew installDist
3 skipTests=./gradlew installDist
4
5 [run]
6 cmd=JAVA_OPTS="-Xms4g" build/install/EMFSolutionATL/bin/EMFSolutionATL
```

`solution.ini` file for the reference ATL solution is shown on Listing 1. In the `build` section, the `default` option specifies the command to build and test the solution, and the `skipTests` option specifies the command to build the solution while skipping unit tests. In the `run` section, the `cmd` option specifies the command to run the solution.

The repetition of executions as defined in the benchmark configuration is done by the benchmark. For 5 runs, the specified command will be called 5 times, passing any required information (e.g. run index, or input model name) through environment variables. Solutions must not save intermediate data between different runs: each run should be entirely independent.

The name and absolute path of the input model, the run index and the name of the tool are passed using environment variables `Model`, `ModelPath`, `RunIndex` and `Tool`. Solution authors are suggested to study the example ATL solution on how to use these values to run their transformation.

4.2 Running the benchmark

The benchmark framework only requires Python 3.3 to be installed. Furthermore, the solutions may imply additional frameworks. We would ask solution authors to explicitly note dependencies to additional frameworks necessary to run their solutions.

If all prerequisites are fulfilled, the benchmark can be run using Python with the command `python scripts/run.py`. Additional options can be queried using the option `--help`. The benchmark framework can be configured through the `config/config.json` file: this includes the input models to be evaluated (some of which have been excluded by default due to their high cost with the sample solution), the names of the tools to be run, the number of runs per tool+model, and the timeout for each command in milliseconds.

5 Evaluation

Given the call for a broader set of research interests in this transformation, the evaluation will operate on two dimensions:

- Is it a high-quality model transformation? The transformation should be complete, correct, easy to understand, efficient, and produce optimal results.

The `validator.jar` will be used to check the produced BDDs against the source truth tables, and the authors will need to provide a convincing argument about the correctness of the solution.

Attendees to the contest will evaluate the understandability of the solution, and the benchmark framework will provide independent measurements of memory and time usage.

Tree sizes for the BDDs will be considered for the optimality of the transformations: the smaller, the better.

- Does it highlight a promising research direction? Although the transformation may not be entirely complete or may be harder to understand, it may serve as an example of an active research area within model transformations that the community may wish to showcase.

This may include aspects such as incrementality, bidirectionality, traceability, verifiability, or the ability to visualize interactively the transformation, among other areas of interest in the field.

Depending on the submitted solutions, the awards for this dimension may change. If there are no common areas of interest in the solutions, an audience-driven “Most Promising” award will be awarded. If there are common themes between solutions, dedicated 1st/2nd/3rd place awards for those research areas may be awarded.

References

- [1] Anthony Anjorin, Thomas Buchmann, and Bernhard Westfechtel. The Families to Persons Case. In *Proceedings of the 10th Transformation Tool Contest*, volume 2026, pages 27–34, Marburg, Germany, July 2017. CEUR-WS.org.
- [2] Martin Fleck and Javier Troya. The Class Responsibility Assignment Case. In *Proceedings of the 9th Transformation Tool Contest*, volume 1758, pages 1–8, Vienna, Austria, July 2016. CEUR-WS.org.
- [3] Antonio García-Domínguez. TTC’16 live contest case study: execution of dataflow-based model transformations. <https://www.transformation-tool-contest.eu/2016/livecontest.html>, July 2016. Last accessed on 2019-05-10. Archived on <http://archive.is/gHEys>.
- [4] Sebastian Götz, Johannes Mey, René Schöne, and Uwe Aßmann. Quality-based Software-Selection and Hardware-Mapping as Model Transformation Problem. In *Proceedings of the 11th Transformation Tool Contest*, volume 2310, pages 3–11, Toulouse, France, June 2018. CEUR-WS.org.
- [5] Georg Hinkel. The TTC 2017 live contest on transformation reuse in the presence of multiple inheritance and redefinitions. <https://www.>

transformation-tool-contest.eu/2017/solutions_livecontest.html, July 2017. Last accessed on 2019-05-10. Archived on <http://archive.is/gHEys>.

- [6] Georg Hinkel. The TTC 2017 Outage System Case for Incremental Model Views. In *Proceedings of the 10th Transformation Tool Contest*, volume 2026, pages 3–12, Marburg, Germany, July 2017. CEUR-WS.org.
- [7] Georg Hinkel. The TTC 2018 Social Media Case. In *Proceedings of the 11th Transformation Tool Contest*, volume 2310, pages 39–43, Toulouse, France, June 2018. CEUR-WS.org.
- [8] Guillaume Savaton. Truth Tables to Binary Decision Diagrams, ATL Transformations. <https://www.eclipse.org/atl/atlTransformations/#TT2BDD>, February 2006. Last accessed on 2019-05-14. Archived on <http://archive.is/HdoHM>.