CEUR-WS.org/Vol-2550/paper4.pdf

Transforming Truth Tables to Binary Decision Diagrams using Relational Reference Attribute Grammars

Johannes Mey, René Schöne, Christopher Werner and Uwe Aßmann
{first.last}@tu-dresden.de

Software Technology Group
Technische Universität Dresden

Abstract

The TTC 2019 case describes the computation of a decision diagram from a truth table. In this paper, we present a solution which uses relational reference attribute grammars to represent both input and output model. To transform the former into the latter, we implemented different strategies and present in detail an algorithm using a reduced order binary decision diagram utilizing higher-order attributes. We further present transformation times and the number of decision nodes in the result model showing the feasibility of our approach.

1 Introduction

In the TTC 2019 case, a given truth table has to be transformed into a logically equivalent binary decision diagram. Both truth table and binary decision diagram (BDD) are two different ways to represent a function over Boolean tuples, where on the one hand truth tables can be seen as a set of assignments, each given as a row of the table. On the other hand, a binary decision diagram represents this function as a directed acyclic graph (DAG) with a single source. The nodes of the DAG are partitioned into decision nodes and leaf nodes. Decision nodes refer to input variables and have edges annotated with either 0 or 1, representing the corresponding decision. Leaf nodes are annotated with an output tuple.[1]

2 Background

Attribute Grammars [Knu68] use a context-free grammar for the declarative definition of an abstract syntax tree (AST) and attributes to define analyses on this tree. In [Hed00], this concept was extended to Reference Attribute Grammars (RAGs), which allow attributes to compute nodes of the AST, effectively enabling the definition of graphs instead of trees. Our solution uses Relational Reference Attribute Grammars [MSH+18], which in turn extend RAGs with first-class relations, i.e, edges between nodes in the AST, to provide an easy definition of those relations and to enable bidirectional relations. Furthermore, we use an advanced feature of RAGs, nonterminal attributes [VSK89]. These are attributes computing a new subtree, which after computation is treated like a nonterminal. Using them, implicit knowledge of the tree can be manifested and later be reused. In our implementation, we use *JastAdd* [EH07], a RAG system to define both grammar and attributes. The grammar is specified using a BNF syntax with inheritance and relations. Every nonterminal defined there is compiled to a Java class with accessors for its children, attributes, and relations. Attribute definitions are specified using a Java-based DSL and are woven into the Java class of the nonterminal they are defined in. As described in [MSH+18], a preprocessor transforms the relations into basic grammar rules and special accessors.

[1] BDDs are usually specified for one output variable. Otherwise, *shared BDDs* with multiple source nodes can be used [MT12].

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

In: A. Garcia-Dominguez, G. Hinkel and F. Krikava (eds.): Proceedings of the 12th Transformation Tool Contest, Eindhoven, The Netherlands, 19-07-2019, published at http://ceur-ws.org.

```
1  LocatedElement ::= <Location:String>;
2  TruthTable : LocatedElement ::= <Name:String> Port:Port* Row:Row* ;  // The start symbol of this grammar
3
4  abstract Port : LocatedElement ::= <Name:String>;                    // Port is an abstract node that cannot be instantiated
5  InputPort : Port;                                                    // InputPort and OutputPort are a special ports, i.e., inheriting
6  OutputPort : Port;                                                   //  name and location from it
7  Row : LocatedElement ::= Cell:Cell*;
8  Cell : LocatedElement ::= <Value:Boolean>;                           // The cell has a intrinsic attribute, in this case
9                                                                       //  the boolean value of the cell
10 rel Port.Cell* <-> Cell.Port;                                        // A bidirectional relation stating, that a cell refers to a port
```

Listing 1: Grammar for a truth table in *JastAdd* syntax

```
1  BDT ::= <Name:String> Port:BDT_Port* Tree:BDT_Tree;
2  abstract BDT_Tree;
3  BDT_Leaf:BDT_Tree ::= Assignment:BDT_Assignment*;
4  BDT_Subtree:BDT_Tree ::= TreeForZero:BDT_Tree TreeForOne:BDT_Tree;
5
6  abstract BDT_Port ::= <Name:String>;
7  BDT_InputPort : BDT_Port;
8  BDT_OutputPort : BDT_Port;
9  BDT_Assignment ::= <Value:boolean>;
10
11 rel BDT_InputPort.Subtree* <-> BDT_Subtree.Port;
12 rel BDT_OutputPort.Assignment* <-> BDT_Assignment.Port;
13
14 // relations to TruthTable model
15 rel BDT.TruthTable -> TruthTable;
16 rel BDT_InputPort.TruthTableInputPort -> InputPort;
17 rel BDT_OutputPort.TruthTableOutputPort -> OutputPort;
18 rel BDT_Leaf.Row* -> Row;
```

Listing 2: *JastAdd* grammar for Binary Decision Tree

```
1  BDD ::= <Name:String> Port:BDD_Port* Tree:BDD_Tree*;
2  abstract BDD_Tree;
3  BDD_Leaf:BDD_Tree ::= Assignment:BDD_Assignment*;
4  BDD_Subtree:BDD_Tree;
5  abstract BDD_Port ::= <Name:String>;
6  BDD_InputPort : BDD_Port;
7  BDD_OutputPort : BDD_Port;
8  BDD_Assignment ::= <Value:boolean>;
9  rel BDD.Root -> BDD_Tree;
10 rel BDD_Subtree.TreeForZero <-> BDD_Tree.OwnerSubtreeForZero*;
11 rel BDD_Subtree.TreeForOne <-> BDD_Tree.OwnerSubtreeForOne*;
12 rel BDD_InputPort.Subtree* <-> BDD_Subtree.Port;
13 rel BDD_OutputPort.Assignment* <-> BDD_Assignment.Port;
14 // relations to TruthTable model
15 rel BDD.TruthTable -> TruthTable;
16 rel BDD_InputPort.TruthTableInputPort -> InputPort;
17 rel BDD_OutputPort.TruthTableOutputPort -> OutputPort;
18 rel BDD_Leaf.Row* -> Row;
```

Listing 3: Grammar for a Binary Decision Diagram

```
1  PortOrder;
2  rel TruthTable.PortOrder -> PortOrder ;
3  rel PortOrder.Port* -> InputPort;
```

Listing 4: Grammar for *PortOrder*

```
1  // TruthTable has two NTA defined:
2  syn PortOrder TruthTable.getNaturalPortOrder() = //...
3  syn PortOrder TruthTable.getHeuristicPortOrder() = //...
```

Listing 5: NTA definition for *PortOrder*

3 Computing a Binary Decision Diagrams with Relational RAGs

In this section, we show how truth tables, BDTs, and BDDs can be modelled using relational RAGs and, in particular, how relations support the transformations and traceability between the models. Initially, the case description demanded a binary decision *tree* (BDT) and not a binary decision diagram (BDD). Therefore, we present different configurable algorithms to transform a truth table into either a BDT or a BDD. While the outputs of all provided transformations are semantically equivalent for one given input, they have different characteristics regarding both the runtime of the transformation and the properties of the resulting diagrams.

3.1 Describing Conceptual Models

Many conceptual models can be described by relational RAGs. This is possible, because many model specification languages, e.g., Ecore, require models to have a *containment hierarchy*, and thus a spanning tree. Listing 1 shows the grammar of the truth table that corresponds to the provided model. Since the grammar specification rules of *JastAdd* use concepts like inheritance and abstract types, all given Ecore models can easily be transformed into relational RAGs. The resulting BDT and BDD grammars are shown in Listings 2 and 3. While some concepts like *Port* potentially could be reused between the models, we refrained from doing so, because those concepts do not have exactly the same definition in all models.[2]

Important aspects when dealing with several trees at the same time defined by relational RAGs are the modularity of their specification and the reachability between those trees. RAGs as specified by *JastAdd* are inherently modular: Both grammar and attributes can be split into aspect files. Therefore, both truth table, decision tree and -diagram models are described by one relational grammar, but specified in different modules. Since the truth table grammar shown in Listing 1 has no relations to the other models, it can be used independently. On the other hand, the diagram models have a relation to the truth table, as shown in Listing 2. Note that this is not required, but a design decision, since those relations enable traceability between the models.

[2]For example, a *Port* within a truth table is a *LocatedElement*, while one in a BDD is not. In other cases, the types may be the same, but the relations between them are not.

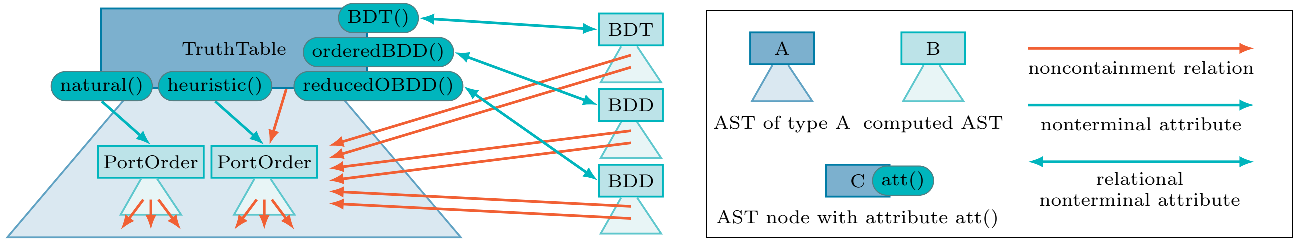

Figure 1: Elements of the transformation and their relations

3.2 Computing Models

A model transformation can be seen as the computation of a target model using information of a source model. RAGs provide a mechanism for that: attributes that compute a tree of the grammar they are defined in, called *higher-order* or *nonterminal attributes (NTAs)*. There are two types of NTAs, those that describe a *subtree* of an AST and those that describe a new tree. In both cases, non-containment relations added by *relational* RAGs provide the means to link these (sub-)trees to the original tree. Figure 1 shows the structure of the trees used for the case. The only given tree is the *TruthTable* tree, containing two subtrees of the type *PortOrder* shown in Listing 4, defined by NTAs from Listing 5. The results of the transformation are stored in separate trees of type *BDT* and *BDD*, computed by the relational NTAs *BDT()*, *orderedBDD()* and *reducedOBDD()*. Note that different NTAs may return different instances of the same *type*. In addition to the relations computed by NTAs, there may be other *intrinsic* relations, such as the relation that selects one of the two provided *PortOrder*s. These relations can also link nodes in computed subtrees to nodes in the originating tree.[3]

In the following, we focus on one transformation into reduced ordered BDDs using the NTA *reducedOBDD()*.

3.3 Computing an Ordered Binary Decision Diagram

Constructing an optimal BDD is a computationally hard, but, since there are many real world applications that use large BDDs, appropriate simplifications and efficient heuristics exist. One commonly used simplification of BDDs are ordered BDDs (OBDD), in which the order of input variables[4] in all paths is identical, allowing a layering of the graph. Given an OBDD, two simple reduction rules are applied to reduce the number of nodes; the result of such a reduction process is a reduced OBDD. In an OBDD, the minimal diagram for a given order can be computed efficiently, however, the computation of the optimal order is still NP-hard [MT12, p. 139ff.].

We follow the definition of an OBDD given in [MT12] extending the problem to several output variables as mentioned in Section 1 as follows:

- An OBDD with m input and n output variables has at most 2^n leaves, each with a different assignment function $f : V_o \to \{0, 1\}$ with V_o as the set of output variables.
- There is an order $<_\pi$ for input variables, such that on an edge from one node referring to input variable x to a node referring to variable y it holds that $x <_\pi y$.

Conceptually, we split the construction of a reduced OBDD in three stages: the computation of a port order, the construction of a perfect tree, and its reduction. Listing 6 shows the synthesized attribute that performs this process. First, a BDD node is constructed and the relations to the truth table and its variables are established (lines 2–5). Then, the leaves are constructed by a helper function and added to the diagram (line 6). Afterwards, a port order is retrieved by accessing the noncontainment relation to a *PortOrder* pointing to one of the NTAs that computes a port order. These are defined in a separate grammar aspect shown in Listing 4. Besides the default order defined in line 2 of Listing 5, we provide a heuristic ordering defined in line 3. Since most heuristics in the literature rely on a logical formula as an input, we have chosen to use a simple metric based on the correlation of an input variable to the output vector in the given truth table. Using this port order, the tree is constructed iteratively by adding the path for each input row (Listing 6, lines 9–13). Finally, the reduction is performed in line 14. The employed algorithm and its properties can be studied in [MT12, p. 96ff.]. For a given truth table, there is exactly one minimal OBDD, computed by the algorithm in $\mathcal{O}(d \log d)$ for d decision nodes.

Besides computing and reducing an OBDD, we have also implemented other approaches as shown in Section 5. How to apply those approaches to the given case will be described in the next section.

[3] These relations must not be bidirectional, since the direction from the source model to the computed NTA model would violate the rule that attribute computations may not alter the tree — other than adding the result of the computation.

[4] In literature such as [MT12, RK08], ports are called variables.

```
1  syn BDD TruthTable.reducedOBDD() {
2    BDD bdd = this.asBDD();                                  // Construct an empty BDD with trace link to truth table
3    bdd.setName(getName());                                  //   and copy the name from the truth table
4    for (Port port: this.getPortList())                      // For each port
5      bdd.addPort(port.asBDDPort());                         //   the according BDD variant of the port is created and linked back
6    for (BDD_Leaf leaf: constructLeaves()) bdd.addTree(leaf); // Add leaves
7    PortOrder portOrder = getPortOrder();                    // Obtain the port order by accessing a non-containment relation
8                                                             //   to the selected nonterminal attribute
9    BDD_Subtree root = new BDD_Subtree();                    // Create root node, set its port to the first in the port order,
10   root.setPort(bdd.bddInputPort(portOrder.getPortList().get(0))); //   add it to the BDD, and specify it as the root
11   bdd.addTree(root);
12   bdd.setRoot(root);
13   for (Row row : getRowList()) insertRow(bdd, root, row, 0); // Fill the BDD by adding all defined paths using a helper function
14   bdd.reduce();                                            // Perform the reduction by calling a helper function
15   return bdd;
16 }
```

Listing 6: Relational non-terminal attribute to create an ordered BDD

```
1 syn int BDT_Tree.decisionNodeCount();
2 eq BDT_Subtree.decisionNodeCount() = 1 + getTreeForZero().decisionNodeCount() + getTreeForOne().decisionNodeCount();
3 eq BDT_Leaf.decisionNodeCount() = 0;
```

Listing 7: Computation of the number of decision nodes in the BDT

4 A Dynamic Transformation Toolchain

To perform the transformation, we follow the process outlined in Figure 2. We will show, how our approach supports variability that enables reuse and the combination of different transformation approaches and how relational RAGs support traceability between models and help in analysing these models.

As *JastAdd* is not based on Ecore, the meta model of EMF [SBMP08], we can not directly use the given input models. Instead, we translated the given metamodel into the grammar shown in Listing 1 and built a hand-written XMI parser constructing an AST according to this grammar. However, relational RAGs provide some mechanisms to reduce the implementation overhead and simplify this process.

While parsing an XMI file is rather straightforward, the resolution of the XMI references is not. In attribute grammars, name analysis is well-supported and a frequently used application. Relational RAGs provide an additional method to defer the name resolution after the parsing while still allowing the result of the attribute-computed name resolution to be stored as a non-containment relation. Once the truth table is parsed, the transformation is performed. The first step is to create an additional subtree with a *PortOrder*, which can later be used to create ordered BDTs and BDDs. This is also the first configurable step of the process, since the algorithm for the computation of the *PortOrder* can be switched. Afterwards, BDTs or BDTs are created by different attributes. In the case of the OBDD, the reduction can be seen as an additional step of the computation. Since all variants are independent trees, it is possible to create several variants at the same time. Each created variant contains trace links into the truth table model that are created during construction (cf. Listing 2, lines 15 to 18). Then, the results are validated and different metrics described in Section 5 are computed. One example for a metric is the number of decision nodes in Listing 7. Finally, the resulting BDT or BDD is serialized to XMI. Again, this step requires attributes to compute the references within the file, i.e., XMI path expressions. This whole process is embedded into the provided benchmarking framework[5] and evaluated in the next section.

[5] The implementation can be found at `https://git-st.inf.tu-dresden.de/ttc/bdd`

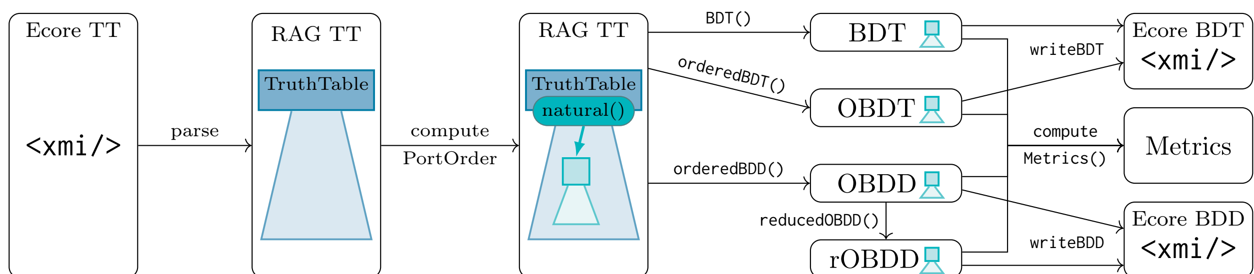

Figure 2: Transformation process and temporary artefacts

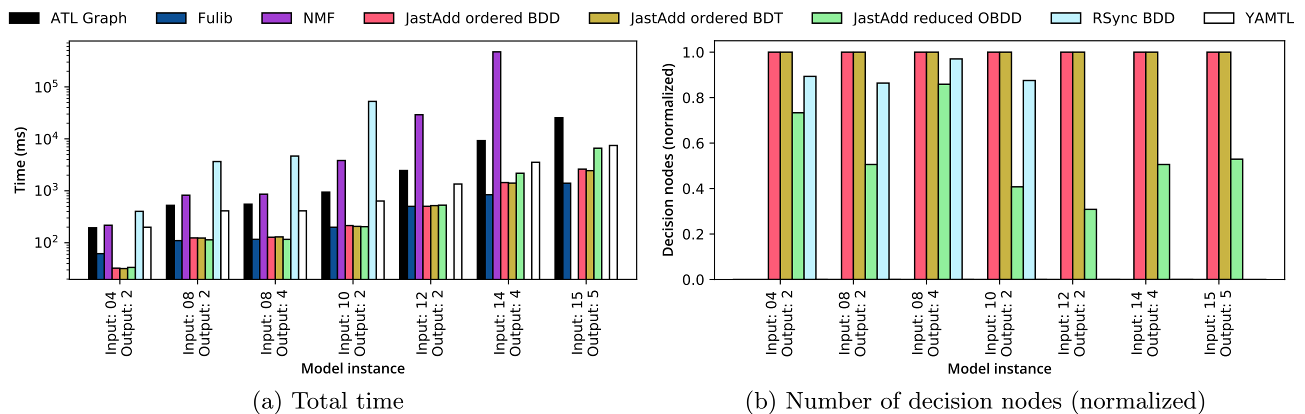

(a) Total time

(b) Number of decision nodes (normalized)

Figure 3: Evaluation results

5 Evaluation

The evaluation consists of two parts. First, the transformation and its properties are discussed with a focus on conciseness, modularity, and reuse, then its performance and the quality of the target models are described.

5.1 Properties of the Transformation

Using relational RAGs, the presented models can be specified concisely. As shown in Listings 1 to 4, the grammars for truth table, BDT and BDD comprise few lines of code. The specification of the transformation using the attributes of *JastAdd* proves to be a good combination of imperative code (which is beneficial for such complex transformation algorithms as, e.g., for the OBDD reduction) and efficient attribute-supported tree navigation. Additionally, attributes can help by checking both the *correctness* and different *metrics* of the obtained model.

An important aspect of the presented solution is modularity and the opportunities for reuse that stem from it. Even though, technically, all models are combined in one large grammar, this grammar can be used in independent modules consisting of grammar fragments and accompanying attributes. The example of the *PortOrder* extension shows how additional analysis modules can simply be added to the grammar. Using relations, entire new models can be integrated and used by existing models. On the attribute level, the definition of attributes in *aspects* also helps to combine and reuse separate parts of the transformation system. Furthermore, relational RAGs allow traceability. Relations between models can be established that, e.g., show which rows have taken part in the creation of a *Leaf* in a BDD.

5.2 Performance and Quality

To evaluate the runtime performance, the measurements of the TTC organizers were considered. All measurements were performed on a Google Cloud Compute Engine `c1-standard-4` machine with 16GiB of RAM and an 30GB SSD, using a Docker image.[6] Each configuration was run ten times. In addition to time measurements, we added a number of result quality metrics computed by attributes, namely the number of nodes of the two different types and the minimum, maximum, and average path length.

Figure 3 shows our evaluation results for the seven provided input models depicting total runtime and the number of decision nodes, i.e., instances of type *Subtree*, normalized to the greatest number among all variants.[7] For the measurement, we included the following variants of our approach: An ordered BDT variant generating the tree iteratively (*ordered BDT*), an ordered BDD generated iteratively (*ordered BDD*) and a reduced variant of the latter, using the algorithm described in Section 3.3 (*reduced OBDD*). For the last variant, we used the order determined by the heuristic. These solutions are compared to the other contributions to the TTC that integrate into the benchmark infrastructure. As shown in Figure 3a, all variants display a good performance compared to the other solutions; the only solution that is considerably faster for larger problems is *Fulib*, which benefits from being a pure Java solution optimized for the input data of the benchmark. However, relational RAGs do not require any initialization other than loading the required Java classes, resulting in only a small overhead compared to *Fulib* that we assume is due to the larger memory footprint of *JastAdd*. Comparing the execution times of our approaches, the non-reduced implementations show a very similar performance, while the

[6] `https://github.com/TransformationToolContest/ttc2019-tt2bdd/blob/master/Dockerfile`

[7] Since a variant is included that creates a perfect tree, this number is $2^m - 1$ where m is the number of input variables.

reduced OBDD takes up more time for the largest model, due to the final reduction step.While the benchmark only measured resident memory before and after the transformation, the figures for this stayed below 1GB for all variants and scenarios and are thus lower than for all other solutions except *Fulib*. Looking at the number of decisions in the final result, there are three classes of approaches. All non-reduced OBDD variants generate the largest number of decision nodes, as they always produce the full tree. The algorithm used in the case description reduces the nodes in the tree by 2% to 14% compared to the full tree. When applying the reduction algorithm to an ordered BDD, 14% to 69% of the decision nodes in the full tree are removed.

6 Conclusion and Future Work

We have shown how to apply Relational Reference Attribute Grammars to the problem of transforming truth tables into (ordered) binary decision diagrams. This included defining a suitable grammar for both truth tables and BDDs, parsing the given input models, computing a BDD in different ways while reusing intermediate results, and finally printing the result to the required XMI format. We used the *JastAdd* system to implement our solution and were able to create several configurable transformations with good performance compared to most other solutions. As the focus of this contribution was not to create new transformation algorithms, we show how the implementation of those algorithms, metrics, and tracing relations were improved by relational RAGs.

The presented approach demonstrates a manual bidirectional transformation from Ecore-based models to *JastAdd* ASTs and back. While the grammar as well as the parser and printer were hand-written, there is no obvious reason why this could not be automated. Having an automated transformation from Ecore meta-models to grammars and their instances to ASTs would instantly make a huge set of existing models available for efficient RAG-based analysis and is therefore an important direction of research.

Acknowledgements
This work has been funded by the German Research Foundation within the Research Training Group "Role-based Software Infrastructures for continuous-context-sensitive Systems" (GRK 1907), the research project “Rule-Based Invasive Software Composition with Strategic Port-Graph Rewriting” (RISCOS) and by the German Federal Ministry of Education and Research within the project “OpenLicht’.

References

[EH07] Torbjörn Ekman and Görel Hedin. The JastAdd system—modular extensible compiler construction. *Science of Computer Programming*, 69(1):14–26, 2007.

[GDH19] Antonio Garcia-Dominguez and Georg Hinkel. Truth Tables to Binary Decision Diagrams. In Antonio Garcia-Dominguez, Georg Hinkel, and Filip Krikava, editors, *Proceedings of the 12th Transformation Tool Contest, a part of the Software Technologies: Applications and Foundations (STAF 2019) federation of conferences*, CEUR Workshop Proceedings. CEUR-WS.org, July 2019.

[Hed00] Görel Hedin. Reference attributed grammars. *Informatica (Slovenia)*, 24(3), 2000.

[Knu68] Donald E. Knuth. Semantics of context-free languages. *Mathematical systems theory*, 2(2), 1968.

[MSH+18] Johannes Mey, René Schöne, Görel Hedin, Emma Söderberg, Thomas Kühn, Niklas Fors, Jesper Öqvist, and Uwe Aßmann. Continuous Model Validation Using Reference Attribute Grammars. In *Proceedings of the 11th ACM SIGPLAN International Conference on Software Language Engineering*, SLE 2018, pages 70–82. ACM, 2018.

[MT12] Christoph Meinel and Thorsten Theobald. *Algorithms and Data Structures in VLSI Design: OBDD-foundations and applications.* Springer Science & Business Media, 2012.

[RK08] Michael Rice and Sanjay Kulhari. A survey of static variable ordering heuristics for efficient bdd/mdd construction. Technical report, 2008.

[SBMP08] Dave Steinberg, Frank Budinsky, Ed Merks, and Marcelo Paternostro. *EMF: Eclipse Modeling Framework.* Pearson Education, 2008.

[VSK89] H. H. Vogt, S. D. Swierstra, and M. F. Kuiper. Higher order attribute grammars. In *Proceedings of the ACM SIGPLAN 1989 Conference on Programming Language Design and Implementation*, PLDI ’89, pages 131–145, New York, NY, USA, 1989. ACM.