

Demiurge 1.2.0: A SAT-Based Synthesis Tool

Robert Könighofer¹ and Martina Seidl²

¹Institute for Applied Information Processing and Communications, Graz University of Technology, Austria

²Institute for Formal Models and Verification, Johannes Kepler University, Linz, Austria.

Abstract—This document describes the synthesis tool **Demiurge**, version 1.2.0, as submitted to **SyntComp 2015**. **Demiurge** is an open-source tool for synthesizing reactive systems from safety specifications using decision procedures for the satisfiability of quantified and unquantified Boolean formulas. **Demiurge** can also be seen as an extendable framework: New synthesis algorithms and optimizations can easily be integrated in new back-ends, reusing existing infrastructure like the parser, interfaces to SAT- and QBF solvers, existing procedures to extract circuits from strategies, etc. We describe the basic architecture and sketch the working principle of the existing back-ends. Details to the algorithms, optimizations, and first experiments can be found in [2], [1]. More extensive experiments on the **SyntComp 2014** benchmarks can be found in the downloadable archive.

I. INTRODUCTION

Demiurge follows the traditional game-based approach to the synthesis of reactive systems from safety specifications. The specification is seen as a game between two players: the environment player controls the inputs, and the system player controls the outputs of the system. The goal of the system player is to satisfy the specification, i.e., visit only safe states, independent of the environment behavior. In a first step, a so-called *winning region* is computed. The winning region is the set of all states from which the system player can enforce to satisfy the specification. In a second step, a *winning strategy* is derived from the winning region. For every (current) state and input, the winning strategy defines a set of outputs that are okay for satisfying the specification. The last step is to implement this strategy in a circuit, where a concrete choice for the outputs has to be made for every state and input.

In order to achieve acceptable scalability, it is important to implement synthesis algorithms symbolically, i.e., by manipulating formulas instead of enumerating states. In synthesis, these symbolic algorithms are, in turn, usually implemented with Binary Decision Diagrams (BDDs). One reason is that solving games inherently involves dealing with quantifier alternations, and BDDs offer both kinds of quantification. However, BDDs also have their scalability issues. On the other hand, there have been enormous performance improvements in decision procedures for the satisfiability of formulas over the last years and decades. This has led to efficient tools like SAT- and QBF (Quantified Boolean Formulas) solvers. Demiurge leverages this development by implementing symbolic synthesis algorithms using such SAT- and QBF solvers.

This work was supported in part by the Austrian Science Fund (FWF) through projects RiSE (S11406-N23 and S11408-N23) and QUAIN (I774-N23), and by the European Commission through project STANCE (317753) and IMMORTAL (644905).

Demiurge implements various synthesis algorithms in different back-ends. Back-ends can be run in different modes (optimizations and heuristics enabled or disabled) and with various solvers. The *learning-based back-end* computes the winning region with computational learning using either a QBF- or a SAT solver. It also implements heuristics to exploit reachability information and to expand quantifiers partially. The *template-based back-end* uses a QBF- or a SAT solver to compute the winning region as instantiation of a template for a CNF formula over the state variables. Further back-ends include a re-implementation of [5] and an approach based on a reduction to Effectively Propositional Logic (EPR) [2]. A *parallel back-end* runs different methods with different solvers and optimizations in different threads. Details to the back-ends can be found in [2]. All back-ends can be used with several methods to extract circuits from the computed winning region. This includes methods based on QBF certification as well as computational learning using SAT- or QBF solvers. ABC [3] is used in a post-processing step to reduce the circuit size. Details to our circuit synthesis methods can be found in [1].

The modular architecture makes Demiurge easily extendable with new algorithms and optimizations. A lot of infrastructure like interfaces to solvers and entire steps of the synthesis procedure (like extracting a circuit from a winning region) can be reused. At the moment, Demiurge contains uniform interfaces to the APIs of Minisat, Lingeling, PicoSat, and DepQBF (with and without the QBF preprocessor Bloqqer [6]). The interface to DepQBF also supports incremental QBF solving [4]. Interfaces to SAT- and QBF solvers supporting the (Q)DIMACS format are available as well. Furthermore, Demiurge interfaces ABC [3] for circuit minimization. Demiurge is written in C++. The source code is available¹ under the GNU Lesser General Public License version 3. The downloadable archive also contains extensive experimental results on the SyntComp 2014 benchmarks and scripts to reproduce them.

In the following, we outline the architecture of Demiurge and then briefly explain the different back-ends.

II. ARCHITECTURE

The architecture of Demiurge is outlined in Fig. 1. The input is a safety specification in AIGER format. The AIG2CNF module parses it into CNF formulas representing the transition relation and the set of safe states. Next, the back-end selected by the user is executed. The back-ends mostly differ in

¹http://www.iaik.tugraz.at/content/research/design_verification/demiurge/

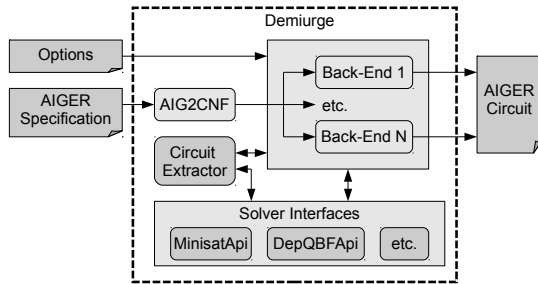


Fig. 1. The architecture of Demiurge.

their method for computing the winning region, and can be parameterized with a method for computing a circuit from the winning region. Both the computation of the winning region and the extraction of circuits rely on external solvers like SAT- and QBF solvers. The resulting circuits are optimized with ABC [3] and dumped in AIGER format again.

III. BACK-ENDS

A. Learning-Based Back-End

The learning-based back-end computes a CNF representation of the winning region W in an iterative manner. It starts with the set of all safe states. In each iteration, it computes a state within the current version F of the winning region from which the environment can enforce to leave F . Obviously, such a state cannot be part of the final winning region W . Hence, the algorithm refines F by removing this state. The state is represented as a cube over the state variables, so removing it from F amounts to adding a clause. By dropping literals from the cube as long as it only contains states that must be excluded from the winning region, the algorithm generalizes the state into a larger region before removing it from the winning region. The detailed algorithm can be found in [2].

For SyntComp 2015, we use the following configuration. Instead of a QBF solver, we use two competing SAT solvers to compute and generalize states to be removed from the winning region (algorithm LEARN SAT from [2] with optimization RG enabled, but optimization RC disabled). As a difference to the SyntComp 2014 submission (version 1.1.0), we also apply partial universal quantifier expansion to reduce the number of iterations. Minisat version 2.2.0 is the underlying SAT solver.

B. Template-Based Back-End

In order to obtain a winning region, this back-end constructs a parameterized CNF formula over the state variables: different concrete values for the (Boolean) parameters induce a different concrete CNF formula over the state variables. This way, the search for a formula over the state variables (the winning region) is reduced to a search for Boolean constants (the template parameter values) [2]. While Demiurge 1.1.0 could only use a QBF solver for finding a template instantiation, version 1.2.0 can also use a SAT solver in a Counterexample-Guided Inductive Synthesis (CEGIS) approach.

For SyntComp 2015, this back-end is not run separately, but only as one thread in our parallelization.

C. Parallel Back-End

The parallel back-end is a playground for combining different methods that refine a CNF representation of the winning region iteratively with additional clauses. Several threads compute and add additional clauses in parallel.

For SyntComp 2015, we use 3 threads: One thread executes the learning-based back-end, one the template-based back-end (alternating between QBF- and SAT solving in 20 second turns), and one our re-implementation of [5]. Minisat 2.2.0 is used as SAT solver. DepQBF 3.04 with our extension of the QBF preprocessor Bloqqer [6] is used for QBF solving in the template-based thread. Note that our parallelization is not just a portfolio approach. The different threads share clauses of the winning region as soon as they are discovered such that other threads can immediately benefit from this information.

D. Circuit Extraction

Demiurge provides several methods for computing circuits from the winning region [1]. One uses QBF Cert to compute Skolem functions for the output signals in a QBF that asserts completeness of the strategy relation. The second one uses computational learning to compute circuits for one output after the other. A third method is based on interpolation.

For SyntComp 2015, we use the learning approach (method SL from [1]) with Lingeling ayv as SAT solver. In our parallelization, we use 3 threads. The first two execute the learning approach in two variants (SL and SLN from [1]). The third thread executes the learning approach using incremental QBF solving (method QL from [1]) with DepQBF 3.04.

E. Other Back-Ends

Demiurge contains more back-ends that are either experimental or did not turn out to be particularly competitive. The *EPR back-end* [2] reduces the synthesis problem to Effectively Propositional Logic (EPR) and uses iProver to solve the formulas. It suffers from high memory consumption. The *IFM back-end* is a re-implementation of [5]. It performs well on certain benchmarks, but is outperformed on many others [2].

IV. CONCLUSION

Demiurge is an open-source synthesis tool for safety specifications. It implements several synthesis algorithms based on SAT- and QBF solving [2], [1]. Demiurge is also an extendable framework for implementing new synthesis algorithms, thereby reducing the entry barrier for new research on SAT- and QBF-based synthesis algorithms and optimizations.

REFERENCES

- [1] R. Bloem, U. Egly, P. Klampf, R. Könighofer, and F. Lonsing. SAT-based methods for circuit synthesis. In *FMCAD'14*. IEEE, 2014.
- [2] R. Bloem, R. Könighofer, and M. Seidl. SAT-based synthesis methods for safety specs. In *VMCAI'14*. Springer, 2014.
- [3] R. K. Brayton and A. Mishchenko. ABC: An academic industrial-strength verification tool. In *CAV'10*. Springer, 2010.
- [4] F. Lonsing and U. Egly. Incremental QBF solving. In *CP'14*, 2014.
- [5] A. Morgenstern, M. Gesell, and K. Schneider. Solving games using incremental induction. In *IFM'13*. Springer, 2013.
- [6] M. Seidl and R. Könighofer. Partial witnesses from preprocessed quantified boolean formulas. In *DATE'14*, pages 1–6. IEEE, 2014.