Improvements to `ltlsynt`

Florian Renkin
LRDE/EPITA
Le Kremlin-Bicêtre, France
renkin@lrde.epita.fr

Philipp Schlehuber
LRDE/EPITA
Le Kremlin-Bicêtre, France
philipp@lrde.epita.fr

Alexandre Duret-Lutz
LRDE/EPITA
Le Kremlin-Bicêtre, France
adl@lrde.epita.fr

Adrien Pommellet
LRDE/EPITA
Le Kremlin-Bicêtre, France
adrien@lrde.epita.fr

Abstract

We summarize `ltlsynt`'s evolution since 2018.

1 Introduction and History

The tool `ltlsynt`, distributed in the Spot library [6] since version 2.5, was originally developed by Thibaud Michaud and Maximilien Colange. They submitted it to the 2017 and 2018 [10] editions of SYNTCOMP. This short document summarizes the improvements brought to `ltlsynt` since then.

While both original authors left the project mid-2018 Maximilien had started working on an alternative approach called LAR (described below) that was eventually included in the Spot 2.7 release. Without any submission of `ltlsynt` to SYNTCOMP'19, the organizers installed the latest version distributed with Spot 2.7.4 themselves, and uncovered a bug caused by an incorrect optimization in the LAR approach.

This optimization was reverted in Spot 2.8, and we started working on an optimized reimplementation of LAR for Spot 2.9, adding other improvement to `ltlsynt` along the way.

A quick summary of all versions submitted to SYNTCOMP over the years is given in Table 1. Note that since the release calendar of Spot is not aligned with SYNTCOMP, most submissions are development version containing unreleased patches applied to a previous release.

2 Technical Details

We describe `ltlsynt`'s general approach in Figure 1. Let us ignore the *decompose* box and the bypass above the blue area for now. The main step of the synthesis process is to convert the LTL specification constraining the input and output signals into a deterministic parity automaton (DPA) where transitions labeled by Boolean combinations of input signals are followed by transitions labeled by Boolean combinations of output signals. This step is shown in the blue-colored box and discussed in Section 2.1. This DPA uses a transition-based max-odd parity acceptance. We then interpret this DPA as a game between two players (the *environment* playing the input signals and the *controller* playing the output signals) and search a winning strategy for the *controller* using a transition-based version of Zielonka's algorithm [14], then encode this strategy as an AIGER circuit.

2.1 Determinization pipelines

The algorithm used by `ltlsynt` to convert the LTL input into a DPA suitable for game solving depends on the `--algo` command-line argument. The first two options, `ds` and `sd`, correspond to pipelines that appeared in `ltlsynt`'s very first release. With `--algo=ds`, LTL inputs are first converted to non-deterministic Büchi automata, then determinized to DPA using a variant of Safra. At this point, transitions are labeled by a mix of input and output signals, so transitions of the form $\bigcirc \xrightarrow{i_1\land i_2\land o_1\land o_2} \bigcirc$ are split into $\bigcirc \xrightarrow{i_1\land i_2} \bigcirc \xrightarrow{o_1\land o_2} \bigcirc$. To preserve determinism, we ensure that multiple transitions sharing the same inputs end up sharing the same new intermediate state. In the pipeline `--algo=sd`, we perform this split before determinizing the automata. Intuitively, this choice may be explained by realizing that the determinization function, in order to compute the successors of a given state, has to consider all compatible assignments of the atomic propositions used by transitions leaving said state: in `ds`, there might be up to $2^{|I|+|O|}$ assignments to consider, whereas in `sd` a given state has either $2^{|I|}$ or $2^{|O|}$ successors at most.

Option `--algo=lar.old` in Spot 2.9 used to be called `--algo=lar` in versions 2.7 and 2.8, and relies on Spot's ability to translate LTL formulas into automata with Emerson-Lei acceptance condition (i.e., any acceptance condition). To do so, this algorithm decomposes the input LTL formula on Boolean operators, translates sub-formulas into deterministic automata (possibly using algorithms specialized for a particular class of formulas), recombines the resulting automata using synchronous products, then applies the relevant Boolean operations on the acceptance conditions. If we are lucky enough, we may avoid Safra-based determinization entirely. However, the resulting deterministic automaton may feature some arbitrary conditions that have yet to be paritized. Therefore, we use a transition-based adaptation of the *state appearance record* algorithm, typically used to convert state-based Muller acceptance to state-based parity. This option was named LAR as a reference to the *latest appearance record* family of algorithms to which SAR belongs (some variants of SAR are often called LAR). We called our variant CAR, for *color appearance record* as it tracks only the colors but neither the states nor the transitions.

Year	Spot version	Main news in `ltlsynt`
2017	pre-2.4? + patches	first implementation
2018	2.5.3 + patches	optimizations to determ., and game solving; incr. determ. approach
2019	2.7.4	(bogus) LAR; improved LTL translation; incr. determ. removed
2020	2.9 + patches	reimplemented LAR, split, and game solving; parity minimization
2021	2.9.7 + patches	input decomposition; strategy simplification; specialized strategy construction

Table 1. Versions of Spot on which `ltlsynt` submissions to SYNTCOMP were based.

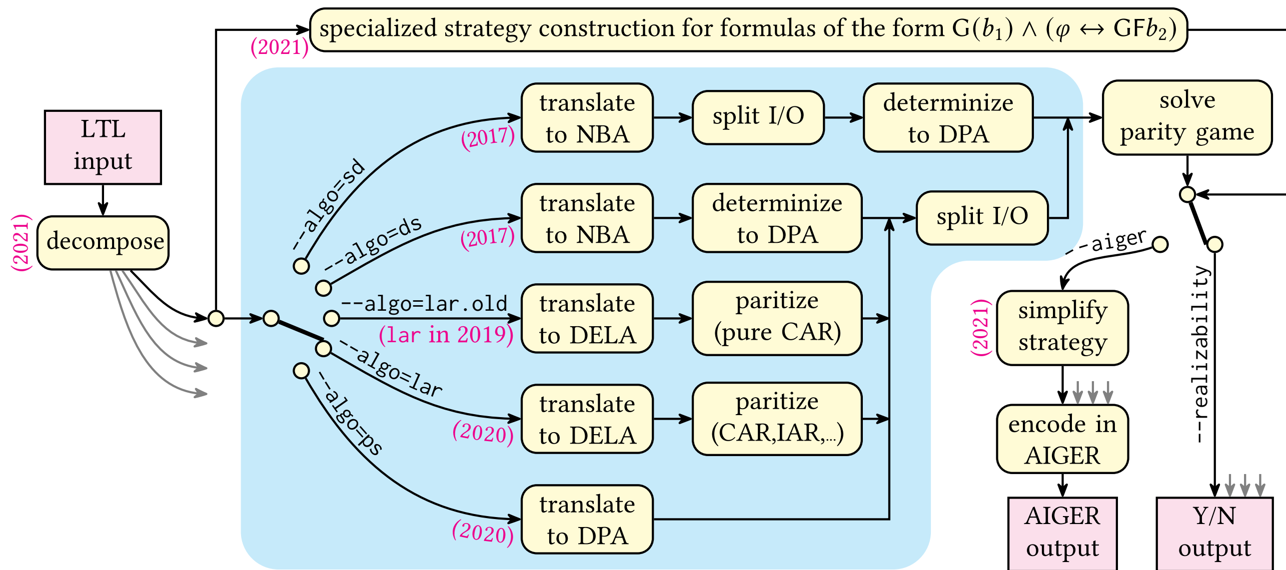

Figure 1. Architecture of `ltlsynt`. The blue zone shows different pipelines for building a parity game, selected by option `--algo`. For some types a formulas, a strategy can be constructed directly from a DBA, bypassing the game construction. If the input is decomposed in multiple conjuncts, recomposition occurs during AIGER encoding.

The CAR implementation in Spot 2.7 featured an optimization that reduced the number of colors tracked by computing the classes of symmetric colors in the acceptance condition (two colors are symmetric if swapping them in the acceptance formula results in an equivalent formula). The intent was to keep track of a smaller number of acceptance classes instead of colors, but this optimization was found to be incorrect during SYNTCOMP'19. This optimization was removed from Spot 2.8 for correctness sake, then replaced by many new optimizations in Spot 2.9 [12].

Option `--algo=lar` now triggers the new implementation of the paritization procedure [12]. It combines CAR (a generic transformation to parity) with IAR (a conversion of Rabin-like or Streett-like acceptance conditions to parity) as well as a partial-degeneralization (in order to reduce conjunctions of `Inf` or disjunctions of `Fin` that occur in the acceptance condition to a single term, as intended by the original symmetry-based optimization) and multiple simplifications of the acceptance conditions. All these transformations are performed on each SCC separately, and it may for instance happen that one SCC is paritized using CAR while another SCC is partially degeneralized to produce an acceptance condition that can be paritized with IAR. Our benchmarks performed on data from SYNTCOMP'17 suggest that the option `--algo=lar` often produces significantly smaller DPAs than `--algo=lar.old` [12].

A new option available since Spot 2.9.1 is `--algo=ps`. This is a close variant of `--algo=ds`, that relies on the translation code that powers `ltl2tgba -P -D` to obtain a DPA. This procedure splits the top-level LTL formulas on Boolean operators in order to translate subformulas corresponding to obligations formulas separately. The remaining subformulas are separately translated to NBA, determinized using Safra if needed, then combined back with the obligation part. Our preliminary experiments showed this option to be inferior to the other methods, and since we had to pick three configurations for this year's competition, we excluded this procedure. In the future it could be improved by tagging the subformulas based on their corresponding acceptance conditions, as performed by Strix [9].

2.2 Various optimizations

We now discuss other optimizations introduced since 2018.

LTL decomposition If the input specification can be seen as a conjunction $\psi_1 \wedge \psi_2 \wedge \ldots \wedge \psi_n$ of subformulas with disjoint output variables, then a strategy for each ψ_i can be computed separately, as suggested by Finkbeiner et al. [8]. Unlike in their experiments, we recompose the different strategies during the AIGER encoding, in case they may share gates.

Translation Since Spot 2.7, the LTL translation engine (that stands behind the "translate to xxA" boxes in Figure 1) learned to split the input formula on Boolean operators in order to separately translate parts to automata then combine these to produce the desired result. This is similar to the process used by the `delag` tool [11], but we use slightly improved algorithms. Extracting *obligations* subformulas is beneficial because those can be converted to minimal weak deterministic automata [5]. Subformulas of the form GF(*guarantee*) or FG(*safety*) can be converted to DBA or DCA using dedicated algorithms (our implementation is a crossover between two different works [7, 11]). Finally, the products combining the resulting automata handle weak-automata and suspendable properties [2] specifically. The heuristics used depend on the type of automata to produce. For instance, in order to generate NBA or DBA, we only split the LTL formula on conjunctions. The post-Spot-2.9 version submitted to SYNTCOMP also deals with xor and equivalence operators while converting to DELA (following Strix's footsteps [9]).

Parity minimization Spot 2.8 features a function that minimizes the number of colors in a DPA [4], now called in `ltlsynt` once a DPA is produced, before merging states with identical successors.

Split — from automata to arenas The split operation described above transforms an automaton into a two-player arena. Even though this step is merely a technicality, benchmarks on Spot 2.9 have shown that it can consume up to 20% of the total run time. In the submitted version, this process has been optimized thanks to caching operations, as labels are often shared among multiple transitions. Moreover, the number of edges and states has been reduced by sharing the introduced intermediate states.

Solving the game Since 2020, the parity game solver of `ltlsynt` is a transition-based adaptation of the one from van Dijk [13]. It supports (non-recursive) SCC decomposition, parity compression (a.k.a. priority compression) and detection of sub-arenas having a single parity. In the majority of the SYNTCOMP benchmarks, solving the parity game is not the bottleneck of `ltlsynt`, but nonetheless remains a crucial step as it also determines the strategy which directly influences the size of the resulting AIGER circuit.

Bypassing the game For inputs of the form $\mathsf{G}(b_1) \wedge (\varphi \leftrightarrow \mathsf{GF}b_2)$, where b_1 is a synthetizable Boolean formula, φ is a DBA-realizable property (a.k.a. recurrence) using only input variables, and b_2 is a Boolean formula using only output variables, a strategy can be constructed directly from the DBA by "anding" each transition with: $b_1 \wedge \neg b_2$ if the transition *can* belong to a rejecting cycle, $b_1 \wedge b_2$ if it *always* belong to an accepting cycle, or b_1 if it cannot belong to any cycle. (If a "false" transition is created, the formula is unrealizable.)

Strategy simplification The winning strategy of a game can be seen an incompletely specified Mealy machine (ISMM): the value of output variables might be unspecified when it does not matter. We implement two algorithms for the simplification of such ISMM. One is a variant of Spot's simulation-based reduction based on BDD signatures [2], where to states whose signature are equivalent up-to-unspecified outputs, can be merged. A second is our own reimplementation of MeMin's SAT-based minimization algorithm for ISMM [1].

Optimizing the output circuit For the synthesis track submission, the AIGER output of `ltlsynt` is run through `abc` for simplification [3]. This is done in the driver script for `starexec`, not by `ltlsynt` itself.

References

[1] A. Abel and J. Reineke. MeMin: SAT-based exact minimization of incompletely specified Mealy machines. In *ICCAD'15*, pp. 94–101. IEEE Press, 2015.

[2] T. Babiak, T. Badie, A. Duret-Lutz, M. Křetínský, and J. Strejček. Compositional approach to suspension and other improvements to LTL translation. In *SPIN'13*, *LNCS* 7976, pp. 81–98. Springer, 2013.

[3] R. Brayton and A. Mishchenko. Abc: An academic industrial-strength verification tool. In *CAV'10*, pp. 24–40. Springer, 2010.

[4] O. Carton and R. Maceiras. Computing the Rabin index of a parity automaton. *Informatique théorique et applications*, 33(6):495–505, 1999.

[5] C. Dax, J. Eisinger, and F. Klaedtke. Mechanizing the powerset construction for restricted classes of ω-automata. In *ATVA'07*, *LNCS* 4762. Springer, 2007.

[6] A. Duret-Lutz, A. Lewkowicz, A. Fauchille, T. Michaud, E. Renault, and L. Xu. Spot 2.0 — a framework for LTL and ω-automata manipulation. In *ATVA'16*, *LNCS* 9938, pp. 122–129. Springer, 2016.

[7] J. Esparza, J. Křetínský, and S. Sickert. One theorem to rule them all: A unified translation of LTL into ω-automata. In *LICS'18*, pp. 384–393. ACM, 2018.

[8] B. Finkbeiner, G. Geier, and N. Passing. Specification decomposition for reactive synthesis. In *NFM'21*, 2021. To appear. https://arxiv.org/abs/2103.08459.

[9] M. Luttenberger, P. J. Meyer, and S. Sickert. Practical synthesis of reactive systems from LTL specifications via parity games. *Acta Informatica*, 57:3—-36, 2020.

[10] T. Michaud and M. Colange. Reactive synthesis from LTL specification with Spot. In *SYNT'18*, 2018. URL http://www.lrde.epita.fr/dload/papers/michaud.18.synt.pdf.

[11] D. Müller and S. Sickert. LTL to deterministic Emerson-Lei automata. In *GandALF'17*, vol. 256 of *EPTCS*, pp. 180–194, 2017.

[12] F. Renkin, A. Duret-Lutz, and A. Pommellet. Practical "paritizing" of Emerson-Lei automata. In *ATVA'20*, *LNCS* 12302, pp. 127–143. Springer, 2020.

[13] T. van Dijk. Oink: An implementation and evaluation of modern parity game solvers. In *TACAS'18*, pp. 291–308. Springer, 2018.

[14] W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200 (1):135–183, 1998.