

Spot 2.0 — A Framework for LTL and ω -Automata Manipulation

Alexandre Duret-Lutz^(✉), Alexandre Lewkowicz, Amaury Fauchille,
Thibaud Michaud, Étienne Renault, and Laurent Xu

LRDE, EPITA, Kremlin-Bicêtre, France
`spot@lrde.epita.fr`

Abstract. We present Spot 2.0, a C++ library with Python bindings and an assortment of command-line tools designed to manipulate LTL and ω -automata in batch. New automata-manipulation tools were introduced in Spot 2.0; they support arbitrary acceptance conditions, as expressible in the Hanoi Omega Automaton format. Besides being useful to researchers who have automata to process, its Python bindings can also be used in interactive environments to teach ω -automata and model checking.

1 Introduction

Spot is a C++ library of model-checking algorithms that was first presented in 2004 [15]. It contains algorithms to perform the usual tasks in the automata-theoretic approach to LTL model checking [36]. It was purely a library until Spot 1.0, when we started distributing command-line tools for LTL manipulation [13] and translation of LTL to some generalizations of Büchi Automata.

Spot 2.0 is a very large rewrite of the core of the library, in C++11, with a focus on supporting automata with arbitrary acceptance conditions as described in the Hanoi Omega Automata format (HOA) [6]. Those acceptance conditions are expressed as positive Boolean formulas over terms such as $\text{Inf}(n)$ and $\text{Fin}(n)$, which indicate respectively that some set S_n of states or transitions should be visited infinitely or finitely often. Traditional acceptance conditions look as follows in this formalism:

Büchi:	$\text{Inf}(0)$	generalized-Büchi:	$\text{Inf}(0) \wedge \text{Inf}(1) \wedge \text{Inf}(2) \wedge \dots$
co-Büchi:	$\text{Fin}(0)$	generalized-co-Büchi:	$\text{Fin}(0) \vee \text{Fin}(1) \vee \text{Fin}(2) \vee \dots$
Rabin:	$(\text{Fin}(0) \wedge \text{Inf}(1)) \vee (\text{Fin}(2) \wedge \text{Inf}(3)) \vee \dots$		
Streett:	$(\text{Fin}(0) \vee \text{Inf}(1)) \wedge (\text{Fin}(2) \vee \text{Inf}(3)) \wedge \dots$		

Parity acceptance, generalized-Rabin [5, 24], and any Boolean combination of the above can be expressed as well. The use of HOA as default format makes it easy to chain Spot's command-line tools, and interact with other tools that implement HOA, regardless of the actual acceptance condition used.

Additionally, Spot 2.0 ships with Python bindings usable in interactive environments such as IPython/Jupyter [29], easing development, experimentation, and teaching.

Spot is a free software and can be obtained from <https://spot.lrde.epita.fr/>. The reader who wants to try Spot without installing it is invited to visit <http://spot-sandbox.lrde.epita.fr/> where a live installation of Jupyter and Spot allows all examples (command lines or Python) of this paper to be replayed.

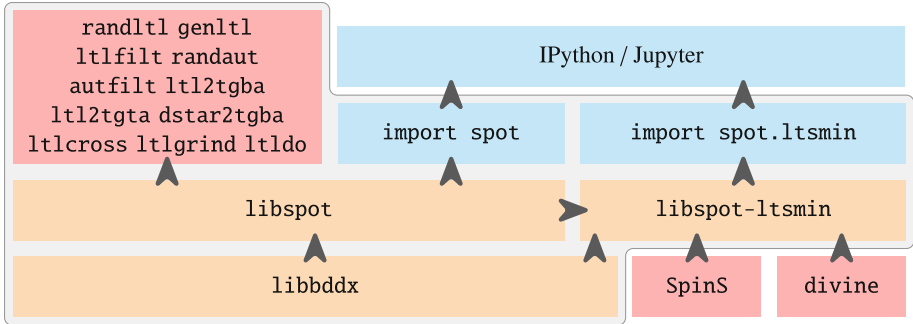


Fig. 1. Architecture of Spot. C++ libraries are in orange boxes, binaries in red, and Python packages in blue. The outlined area is what Spot distributes. (Color figure online)

Figure 1 shows that Spot is actually split in three libraries. `libbddx` is a customized version of BuDDy [26] for representing Binary Decision Diagrams [10] which we use to label transitions in automata, and to implement a few algorithms [4, 14]. `libspot` is the main library containing all data structures and algorithms. `libspot-ltsmin` contains code to interface with state-spaces generated as shared libraries by LTSmin [20].

In the rest of this article, we highlight some of the features of Spot by presenting the command-line tools and the Python bindings built on top of these libraries. The reader should keep in mind that everything that we illustrate as shell command or in Python can be performed directly in C++; in fact our web site gives several examples of tasks implemented with each of these three interfaces.

2 Command-Line Tools

Spot 2.0 installs the following eleven command-line tools, that are designed to be combined as traditional Unix tools.

[13]	}	<code>randltl</code>	generates random LTL/PSL formulas
		<code>genltl</code>	generates LTL formulas from scalable patterns
		<code>ltlfilt</code>	filter, converts, and transforms LTL/PSL formulas
		<code>ltl2tgba</code>	translates LTL/PSL formulas into generalized Büchi automata [14], or deterministic parity automata (new in 2.0)
		<code>ltl2tgta</code>	translates LTL/PSL formulas into Testing automata [8]
		<code>ltlcross</code>	cross-compares LTL/PSL-to-automata translators to find bugs (works with arbitrary acceptance conditions since Spot 2.0)
	<code>ltlgrind</code>	mutates LTL/PSL formulas to help reproduce bugs on smaller ones	
	<code>dstar2tgba</code>	converts <code>ltl2dstar</code> automata into Generalized Büchi automata [1]	
	<code>randaut</code>	generates random ω -automata	
	<code>autfilt</code>	filters, converts, and transforms ω -automata	
<code>ltldo</code>	runs LTL/PSL formulas through other translators, providing uniform input and output interfaces		

The first six tools were introduced in Spot 1.0 [13], and have since received several updates. For instance `ltl2tgba` now uses better simulation reductions and degeneralization [4], and it now provides a way to output deterministic automata using transition-based parity acceptance; `ltlfilt` has learned to decide stutter-invariance of any LTL/PSL formula using an automaton-based check that is independent on the actual logic used [27]; and `ltlcross` can now perform precise equivalence checks of automata in addition to supporting arbitrary acceptance conditions—it has been used by the authors of `ltl3dra` [5], `ltl2dstar` [21,22], and Rabinizer 3 [23] to test recent releases of their respective tools.

The `dstar2tgba` tool was introduced in Spot 1.2 while working on the minimization of deterministic generalized Büchi automata using a SAT-solver [1]. It implements algorithms that translate deterministic Rabin automata into Büchi automata, preserving determinism if possible [25], as well as conversion from Streett to generalized Büchi. These two different kinds of input correspond to the possible outputs of `ltl2dstar`. In Spot 2.0, these specialized acceptance conversions have been preserved, but they are supplemented with more general transformations that input automata with arbitrary acceptance conditions, and transform them into automata with “Fin-less” acceptance, or with (Generalized) Büchi acceptance. These acceptance transformations are essential to a few core algorithms that cannot cope with arbitrary acceptance: for instance currently Spot can only check the emptiness of automata with Fin-less acceptance (all SCC-based emptiness-checks [11,12,31] are compatible with that), so more complex acceptances are transformed when needed.

All these acceptance transformations, as well as other automata transformations are available through the `autfilt` tool. This command can input a stream of automata in 4 different formats (HOA [6], LBTT’s format [33], never claims [19], or `ltl2dstar`’s format [22]), and can output automata, maybe after

filtering or transformation, in some other format (including GraphViz’s dot format [17] for display).

As an example of transformation and format conversion, consider:

```
% spin -f' []<a' | autfilt --complement --dot=abr | dot -Tpng >aut.png
```

This command translates the LTL formula GFa into a Büchi automaton using `spin` [19], the resulting never claim is then fed into `autfilt` for complementation, and the complemented automaton is output into GraphViz’s format for graphical rendering with `dot`. The arguments `a`, `b`, and `r` passed to `--dot` cause the acceptance condition to be displayed, and the acceptance marks to be shown as colored bullets.

In the above example the input to `autfilt` happens to be a deterministic Büchi automaton, so the complementation is as simple as changing the acceptance condition into co-Büchi. If a Büchi output is desired instead, the above command should be changed to `autfilt --complement --ba` and will output a non-deterministic Büchi automaton. This of course works with arbitrary acceptance conditions as input.

Complementation of non-deterministic automata is done via determinization. Our determinization algorithm inputs transition-based Büchi automata (so we may have some preprocessing to do if the input has a different acceptance), and outputs automata with transition-based Parity acceptance. It mixes the construction of Redziejowski [30] with some optimizations of `ltl2dstar` [21,22] and a few of our own.

The `ltldo` command wraps third-party LTL translators and provides them with inputs and outputs that are compatible with the Spot tool-suite. In particular it allows using “single-shot” translators in a pipeline. For instance `spin` can only translate one formula at a time to produce a never claim. The command `ltldo spin` will process multiple formulas (in any syntax supported by Spot [13]), translate them all using `spin`, and output all results in any supported automaton format (HOA by default). For instance the following command uses Spin to translate 10 random LTL formulas into Büchi automata in the HOA format:

```
% randltl -n 10 a b | ltldo spin --name=%f
```

Option `--name=%f` requests input formulas to be used as the “`name:`” field in the HOA format. This field could then be used to retrieve the original formula after further processing: `autfilt --stats=%M` can be used to print the name of each input automaton.

As a more complex example, the following pipeline finds 10 formulas for which `ltl3ba` [3] produces a deterministic Büchi automaton, but `ltl2ba` [18] does not.

```
% randltl -n -1 a b |
  ltldo ltl3ba --name=%f | autfilt --is-deterministic --stats=%M |
  ltldo ltl2ba --name=%f | autfilt -v --is-deterministic --stats=%M -n 10
```

This creates an infinite (`-n -1`) stream of LTL formulas over atomic propositions a and b , translates them using `ltl3ba`, retains those that were translated to deterministic automata, translate them with `ltl2ba` and retains the non-deterministic ones (`-v` inverts matches, as with `grep`). With the final `-n 10`, the pipeline is eventually killed once the last command has found 10 matches.

The `autfilt` tool provides access to other ω -automata algorithms such as product, emptiness checks, language inclusion or equivalence, language-preserving simplifications of automata, refinement of labels [9], strength-based decompositions [32], SAT-based minimization of deterministic automata with arbitrary input and output acceptance [2], or conversion from transition-based acceptance to state-based acceptance. Most algorithms work with arbitrary acceptance conditions, except a few (emptiness checks, determinization) that *currently* have to reduce the acceptance conditions upfront.

3 The Python Interface

Similar tasks can be performed in a more “algorithm-friendly” environment using the Python interface. Combined with the IPython/Jupyter notebook [29] (a web application for interactive programming), this provides a nice environment for experiments, where automata and formulas are automatically displayed. Figure 2 shows two examples that we used in a practical lecture on model checking with students from EPITA.

The first example illustrates how LTL formulas can be parsed (`spot.formula()`), and then translated (using `translate()`) into automata with transition-based generalized Büchi acceptance. Using product, negation, and emptiness check, a student can define a procedure to test the equivalence of two LTL formulas and then use it to explore her understanding of LTL.

The second example illustrates the classical automata-theoretic approach to explicit LTL model checking [36]. Spot can read the shared-libraries used to represent state spaces in the LTSmin project [20]. Those can be compiled from Promela models using SpinS [35], or from DiVinE models using LTSmin’s modified version of DiVinE 2 [7]. In this example the `%dve` keyword is used to specify a short DiVinE model called `adding` (this model comes from the BEEM database [28]) which is immediately compiled and loaded as a shared library. Printing the `adding` Python variable reveals that it is an object using the LTSmin interface, and lists the variables that can be used to build atomic proposition on this model. A Kripke structure can be instantiated from the model by providing a list of atomic propositions that should be valuated on each state. Displaying large Kripke structures is of course not very practical: by default Spot displays only the 50 first states (this can be changed using for instance the `max.states` argument in the first cell). With this interface, we can now easily write a `model_check()` procedure that inputs a model and a formula, instantiates a Kripke structure from the model using all the atomic propositions that appear in the formula, translates the negation of the formula into an automaton, and tests the emptiness of the product between the Kripke structure and this automaton. Note that

2. Baarir, S., Duret-Lutz, A.: SAT-based minimization of deterministic ω -automata. In: Davis, M., Fehnker, A., McIver, A., Voronkov, A. (eds.) LPAR 2015. LNCS, vol. 9450, pp. 79–87. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-48899-7_6](https://doi.org/10.1007/978-3-662-48899-7_6)
3. Babiak, T., Křetínský, M., Řehák, V., Strejček, J.: LTL to Büchi automata translation: fast and more deterministic. In: Flanagan, C., König, B. (eds.) TACAS 2012. LNCS, vol. 7214, pp. 95–109. Springer, Heidelberg (2012)
4. Babiak, T., Badie, T., Duret-Lutz, A., Křetínský, M., Strejček, J.: Compositional approach to suspension and other improvements to LTL translation. In: Bartocci, E., Ramakrishnan, C.R. (eds.) SPIN 2013. LNCS, vol. 7976, pp. 81–98. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-39176-7_6](https://doi.org/10.1007/978-3-642-39176-7_6)
5. Babiak, T., Blahoudek, F., Křetínský, M., Strejček, J.: Effective translation of LTL to deterministic Rabin automata: beyond the (F,G)-fragment. In: Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 24–39. Springer, Heidelberg (2013)
6. Babiak, T., Blahoudek, F., Duret-Lutz, A., Klein, J., Křetínský, J., Müller, D., Parker, D., Strejček, J.: The Hanoi Omega-Automata format. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 479–486. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-21690-4_31](https://doi.org/10.1007/978-3-319-21690-4_31). <http://adl.github.io/hoaf/>
7. Barnat, J., Brim, L., Rockai, P.: DiVinE 2.0: high-performance model checking. In: HiBi 2009, pp. 31–32. IEEE Computer Society Press (2009)
8. Ben Salem, A.-E., Duret-Lutz, A., Kordon, F.: Model checking using generalized testing automata. In: Jensen, K., Aalst, W.M., Ajmone Marsan, M., Franceschinis, G., Kleijn, J., Kristensen, L.M. (eds.) Transactions on Petri Nets and Other Models of Concurrency VI. LNCS, vol. 7400, pp. 94–122. Springer, Heidelberg (2012)
9. Blahoudek, F., Duret-Lutz, A., Rujbr, V., Strejček, J.: On refinement of Büchi automata for explicit model checking. In: Fischer, B., Geldenhuys, J. (eds.) SPIN 2015. LNCS, vol. 9232, pp. 66–83. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-23404-5_6](https://doi.org/10.1007/978-3-319-23404-5_6)
10. Bryant, R.E.: Graph-based algorithms for boolean function manipulation. *IEEE Trans. Comput.* **35**(8), 677–691 (1986)
11. Couvreur, J.-M.: On-the-fly verification of linear temporal logic. In: Wing, J.M., Woodcock, J. (eds.) FM 1999. LNCS, vol. 1708, pp. 253–271. Springer, Heidelberg (1999)
12. Couvreur, J.-M., Duret-Lutz, A., Poitrenaud, D.: On-the-fly emptiness checks for generalized Büchi automata. In: Godefroid, P. (ed.) SPIN 2005. LNCS, vol. 3639, pp. 169–184. Springer, Heidelberg (2005). doi:[10.1007/11537328_15](https://doi.org/10.1007/11537328_15)
13. Duret-Lutz, A.: Manipulating LTL formulas using Spot 1.0. In: Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 442–445. Springer, Heidelberg (2013). doi:[10.1007/978-3-319-02444-8_31](https://doi.org/10.1007/978-3-319-02444-8_31)
14. Duret-Lutz, A.: LTL translation improvements in Spot 1.0. *Int. J. Crit. Comput. Based Syst.* **5**(1–2), 31–54 (2014)
15. Duret-Lutz, A., Poitrenaud, D.: SPOT: an extensible model checking library using transition-based generalized Büchi automata. In: MASCOTS 2004, pp. 76–83. IEEE Computer Society Press (2004)
16. Fronc, L., Duret-Lutz, A.: LTL model checking with Neco. In: Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 451–454. Springer, Heidelberg (2013). doi:[10.1007/978-3-319-02444-8_33](https://doi.org/10.1007/978-3-319-02444-8_33). <https://github.com/Lvyn/neco-net-compiler>
17. Gansner, E.R., North, S.C.: An open graph visualization system and its applications to software engineering. *Softw. Pract. Exp.* **30**(11), 1203–1233 (2000)

18. Gastin, P., Oddoux, D.: Fast LTL to Büchi automata translation. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, pp. 53–65. Springer, Heidelberg (2001). doi:[10.1007/3-540-44585-4_6](https://doi.org/10.1007/3-540-44585-4_6)
19. Holzmann, G.J.: The Spin Model Checker: Primer and Reference Manual. Addison-Wesley, Boston (2003)
20. Kant, G., Laarman, A., Meijer, J., van de Pol, J., Blom, S., van Dijk, T.: LTSmin: high-performance language-independent model checking. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 692–707. Springer, Heidelberg (2015)
21. Klein, J., Baier, C.: Experiments with deterministic ω -automata for formulas of linear temporal logic. Theoret. Comput. Sci. **363**(2), 182–195 (2006)
22. Klein, J., Baier, C.: On-the-fly stuttering in the construction of deterministic ω -automata. In: Holub, J., Ždárek, J. (eds.) CIAA 2007. LNCS, vol. 4783, pp. 51–61. Springer, Heidelberg (2007)
23. Komárková, Z., Křetínský, J.: Rabinizer 3: safralless translation of LTL to small deterministic automata. In: Cassez, F., Raskin, J.-F. (eds.) ATVA 2014. LNCS, vol. 8837, pp. 235–241. Springer, Heidelberg (2014)
24. Křetínský, J., Esparza, J.: Deterministic automata for the (F,G)-fragment of LTL. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 7–22. Springer, Heidelberg (2012)
25. Krishnan, S.C., Puri, A., Brayton, R.K.: Deterministic ω -automata vis-a-vis deterministic Büchi automata. In: Du, D.-Z., Zhang, X.-S. (eds.) ISAAC 1994. LNCS, vol. 834, pp. 378–386. Springer, Heidelberg (1994)
26. Lind-Nielsen, J., Cohen, H.: BuDDy: Binary Decision Diagram Package. <https://sourceforge.net/projects/buddy/>. Accessed 2 Apr 2014
27. Michaud, T., Duret-Lutz, A.: Practical stutter-invariance checks for ω -regular languages. In: Fischer, B., Geldenhuys, J. (eds.) SPIN 2015. LNCS, vol. 9232, pp. 84–101. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-23404-5_7](https://doi.org/10.1007/978-3-319-23404-5_7)
28. Pelánek, R.: BEEM: benchmarks for explicit model checkers. In: Bošnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 263–267. Springer, Heidelberg (2007)
29. Pérez, F., Granger, B.E.: IPython: a system for interactive scientific computing. Comput. Sci. Eng. **9**(3), 21–29 (2007). <http://ipython.org>
30. Redziejewski, R.: An improved construction of deterministic omega-automaton using derivatives. Fundam. Informaticae **119**(3–4), 393–496 (2012)
31. Renault, E., Duret-Lutz, A., Kordon, F., Poitrenaud, D.: Three SCC-based emptiness checks for generalized Büchi automata. In: McMillan, K., Middeldorp, A., Voronkov, A. (eds.) LPAR 2013. LNCS, vol. 8312, pp. 668–682. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-45221-5_44](https://doi.org/10.1007/978-3-642-45221-5_44)
32. Renault, E., Duret-Lutz, A., Kordon, F., Poitrenaud, D.: Strength-based decomposition of the property Büchi automaton for faster model checking. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013. LNCS, vol. 7795, pp. 580–593. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-36742-7_42](https://doi.org/10.1007/978-3-642-36742-7_42)
33. Tauriainen, H., Heljanko, K.: Testing LTL formula translation into Büchi automata. STTT **4**(1), 57–70 (2002)
34. Thierry-Mieg, Y.: Symbolic model-checking using ITS-tools. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 231–237. Springer, Heidelberg (2015)
35. van der Berg, F.I., Laarman, A.W.: SpinS: extending LTSmin with Promela through SpinJa. In: PDMC 2012. ENTCS, vol. 296, pp. 95–105. Elsevier (2012)
36. Vardi, M.Y.: An automata-theoretic approach to linear temporal logic. In: Moller, F., Birtwistle, G. (eds.) LNCS, vol. 1043, pp. 238–266. Springer, Heidelberg (1996). doi:[10.1007/3-540-60915-6_6](https://doi.org/10.1007/3-540-60915-6_6)