

SPEN: A Solver for Separation Logic

Constantin Enea¹, Ondřej Lengál²(✉), Mihaela Sighireanu¹,
and Tomáš Vojnar²

¹ Univ. Paris Diderot, IRIF CNRS UMR 8243, Paris, France

² FIT, Brno University of Technology, IT4Innovations Centre of Excellence,
Brno, Czech Republic
lengal@fit.vutbr.cz

Abstract. SPEN is a solver for a fragment of separation logic (SL) with inductively-defined predicates covering both (nested) list structures as well as various kinds of trees, possibly extended with data. The main functionalities of SPEN are deciding the satisfiability of a formula and the validity of an entailment between two formulas, which are essential for verification of heap manipulating programs. The solver also provides models for satisfiable formulas and diagnosis for invalid entailments. SPEN combines several concepts in a modular way, such as boolean abstractions of SL formulas, SAT and SMT solving, and tree automata membership testing. The solver has been successfully applied to a rather large benchmark of various problems issued from program verification tools.

1 Introduction

For analyzing programs with dynamic memory, *separation logic* (SL) is an established and fairly popular logic introduced by Reynolds et al. [11]. The high expressivity of SL, its ability to generate compact proofs, and its support for local reasoning motivated development of many tools for automatic reasoning about programs with complex dynamic linked data structures. These tools aim at establishing memory safety properties and/or inferring shape properties of the heap. The tools often build on (semi-)decision procedures for checking *satisfiability* and *entailment* problems in SL.

Our tool SPEN¹ provides (semi-)decision procedures for the most commonly considered *symbolic heaps* fragment of SL, extended with user-defined inductive predicates to specify data structures of an unbounded size. Because unrestricted definitions of inductive predicates make the entailment problem for the fragment undecidable [3], only semi-decision procedures have been proposed, e.g., in [2, 4]. Iosif *et al.* [10] identified a rather large class of inductive definitions for which the entailment problem is decidable, although with a high complexity. SPEN focuses on a smaller class of inductive definitions that is, however, expressive enough to specify complex dynamic data structures, such as skip lists, lists of circular lists, AVL trees, or binary search trees.

¹ <https://github.com/mihasighi/spen>.

The chosen class of inductive definitions enables the design of efficient (semi-) decision procedures for satisfiability and entailment [6, 8]. The key idea used for satisfiability checking in SPEN is to exploit the semantics of restricted inductive definitions and of separating conjunction to build an equisatisfiable boolean abstraction of the formula. For entailment checking, the idea is to reduce the problem of checking $\varphi \Rightarrow \psi$ to the problem of checking a set of *simple entailments* where the right-hand side is an inductive predicate atom. The compositionality of this reduction leads to high efficiency (the simple entailments can be checked independently) and to a capability to provide fine diagnosis for invalid entailments.

The current version of SPEN improves on the ones reported in [6, 8] in several directions. First, we introduced caching of constructions and results obtained from checking simple entailments in order to increase its efficiency. Second, the wrappers calling the SAT and SMT solvers have been refined to generate smaller formulas and to exploit the incrementality feature of underlying solvers. Third, we improved the diagnosis produced by SPEN. For satisfiability checking, SPEN now provides either a model of a satisfiable formula or an unsatisfiable core; for entailment checking, SPEN provides a proof witness for valid entailments and a diagnostic information otherwise.

SPEN has been successfully tested on a quite large benchmark. The first version of SPEN participated in the SL-COMP'14 contest [15] where it won one of its divisions and was second in another one. The later extensions now allow SPEN to handle a richer fragment than those considered in the competition. Moreover, the improvements above lead to better execution times (e.g., by 10% within the SL-COMP'14 division won by the first version of SPEN and by 30% on the division where SPEN was the second).

SPEN is not the only solver for SL. The existing solvers differ in the fragment considered (CYCLIST [2], SLIDE [9]) and/or the techniques used (ASTERIX [12], DRYAD [14], GRASSHOPPER [13], SLEEK [4]). A detailed comparison with these solvers is beyond the scope of this paper—we refer the reader to the survey in [6, 8, 15].

2 Logic Fragment

SPEN deals with decision problems in a fragment of SL, denoted as SL^{ID} , that combines the *symbolic heaps* fragment of SL [1] with user-defined *inductive predicates* describing various kinds of *lists* (possibly nested, cyclic, or equipped with skip links) or *trees*, possibly extended with data constraints.

Syntax: We write X, Y, Z to denote *location variables*, d to denote *data variables*, and x, y, z for both kinds of variables. We use the vector notation \vec{x} to abbreviate tuples. We denote by ρ the tuples built from pairs of *field labels* and variables that specify structured values. We assume a finite set $\mathcal{P} = \{P_1, \dots, P_n\}$ of *predicate symbols*, each with an associated arity, and a special location variable nil .

A *symbolic heap formula* ψ is a formula of the form $\exists \vec{x} \cdot \Pi \wedge \Sigma$ where Π is a *pure formula* and Σ is a *spatial formula* with the following syntax:

$$\Pi ::= X = Y \mid X \neq Y \mid \Delta \mid \Pi \wedge \Pi \qquad \Sigma ::= \mathbf{emp} \mid X \mapsto \rho \mid P(X, \vec{x}) \mid \Sigma * \Sigma$$

Here, Δ is a constraint over data variables. We let it unspecified, though SPEN presently supports the first-order theory over multisets of integers with integer linear constraints. The spatial atoms (i.e., the empty heap, the heap cell allocated at X , resp. the heap region shaped by some predicate $P \in \mathcal{P}$) are composed by the separating conjunction “*”. An SL^{ID} formula φ is a set of symbolic heaps interpreted as a disjunction $\vee_i \psi_i$.

Predicates $P \in \mathcal{P}$ are defined by a set of *inductive rules* of the form $\psi \Rightarrow P(X, \vec{x})$ where (X, \vec{x}) is a tuple of distinct variables including all free variables in the symbolic heap ψ (the rule body). X is called the root node of the heap segment defined by P . A rule is called a *base rule* if its spatial part is \mathbf{emp} , i.e., an empty heap; otherwise, it is an *inductive rule*.

Fragments: SPEN considers a restricted class of inductive rules such that the defined predicates specify (possibly empty) heap segments connecting (by location fields) the root location X with all locations in the heap or nil. The restrictions have been defined formally in [6,8]. They mainly require, for each inductive predicate P , the presence of a unique base rule and inductive rules where the root X points to a memory cell that contains at least one field from which another heap specified by P starts. The fragment defined in [6], called SL_L^{ID} , can describe various kinds of lists that can be singly- or doubly-linked, cyclic, nested, and can have skip links. It does not permit data constraints and inductive tree structures. On the other hand, the fragment SL_D^{ID} defined in [8] permits data constraints and can describe tree structures of bounded width, such as sorted list segments, AVL trees, binary search trees, but not nested cyclic lists.

Decision Problems: For both fragments above, SPEN considers the problems of checking satisfiability of a formula, i.e., checking whether $\models \varphi$ holds, and the validity of an entailment $\varphi \Rightarrow \varphi'$ where the symbolic heaps of φ' can be quantified only over data variables. A simple example of an entailment problem in SL_L^{ID} considered by SPEN is:

$$\exists Y, W. \mathbf{X} \neq \mathbf{Z} \wedge \mathbf{X} \mapsto \{(\mathbf{next}, Y)\} * \mathbf{sll}(Y, W) * W \mapsto \{(\mathbf{next}, Z)\} \stackrel{?}{\Rightarrow} \mathbf{sll}(\mathbf{X}, Z),$$

which, intuitively, checks whether a composition of two memory cells specified by the *points-to atoms* $\mathbf{X} \mapsto \{(\mathbf{next}, Y)\}$ and $W \mapsto \{(\mathbf{next}, Z)\}$ and the predicate atom $\mathbf{sll}(Y, W)$ describes a set of heaps that are all also models of the predicate $\mathbf{sll}(\mathbf{X}, Z)$ defining an acyclic singly-linked list segment between \mathbf{X} and \mathbf{Z} .

3 Satisfiability Checking

Given a set of inductive definitions \mathcal{P} and a symbolic heap ψ , the procedures for checking satisfiability in SPEN follow the workflow given in Fig. 1. The satisfiability checking of an SL^{ID} formula φ makes a classic use of this basic procedure. The crux of the procedures for both fragments is the definition of a boolean formula $B[\psi]$, called *boolean abstraction*, such that the data-free part of ψ is satisfiable iff $B[\psi]$ is satisfiable [6, 7].

Once the boolean abstraction $B[\psi]$ is computed, SPEN queries a SAT solver (currently, MINISAT²) for the satisfiability of $B[\psi]$. If $B[\psi]$ is unsatisfiable, SPEN can return an unsatisfiable core of ψ , deduced from an unsatisfiable core of $B[\psi]$. If $B[\psi]$ is satisfiable and $\psi \in \text{SL}^{\text{ID}}$, SPEN has the option of returning a model of ψ obtained from a model of $B[\psi]$ by unfolding predicate atoms corresponding to non-empty heap segments. The unfolding of predicate atoms is done twice to emphasize the non-emptiness of the segment. For $\psi \in \text{SL}^{\text{ID}}$, the satisfiability checking continues by constructing a formula Δ_ψ that conjuncts the data part of ψ with the data parts obtained by unfolding the non-empty heap segments given by the model of $B[\psi]$. To check the satisfiability of Δ_ψ , SPEN queries an SMT solver for the theory of multisets with integer data (currently, SPEN implements a wrapper for the UFLIA theory of z3 [5]).

If the boolean abstraction $B[\psi]$ is satisfiable, it is then used to normalize the spatial part of ψ , which is a step used by entailment checking too. This process saturates the pure part of ψ with (dis-)equalities between locations variables and removes predicate atoms that correspond to empty heap segments, producing a normalized formula ψ' .

4 Entailment Checking

To check the validity of an entailment $\varphi_1 \Rightarrow \varphi_2$, SPEN uses a sound procedure to deal with disjunctive formulas: it checks that for every disjunct ψ_1 in φ_1 , there is a disjunct ψ_2 of φ_2 such that $\psi_1 \Rightarrow \psi_2$. The procedure for deciding the validity of entailments between symbolic heaps follows the workflows given in Figs. 2 and 3 (the theoretical foundations were established in [6, 8]). The two formulas are first checked for satisfiability and normalized using the procedures from Sect. 3.

² Available at <http://minisat.se>.

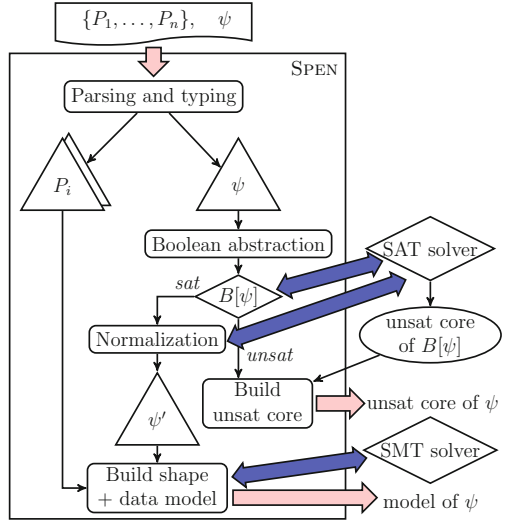


Fig. 1. SPEN workflow for satisfiability checking

If one of the two formulas is unsatisfiable, then the validity of the entailment can be already determined, e.g., if ψ_1 is unsatisfiable then the entailment is valid. When both formulas are satisfiable, SPEN offers two different procedures tuned for each fragment of SL^{ID} .

For the fragment SL_L^{ID} , SPEN reduces the entailment problem $\psi_1 \Rightarrow \psi_2$ to a set of entailment queries of the form $\psi_1[a] \Rightarrow a$, called *simple entailments*, where $\psi_1[a]$ is a sub-formula of ψ_1 and a is a (points-to or inductive) spatial atom of ψ_2 (there will be one such entailment for each spatial atom a in ψ_2). Intuitively, the sub-formula $\psi_1[a]$ describes the region of a heap modelled by ψ_1 that should satisfy a . The procedures for computing $\psi_1[a]$ and testing simple entailments use an intermediary graph representation of symbolic heap formulas, called an *SL-graph* and denoted $G[\psi]$. Basically, nodes of $G[\psi]$ represent sets of aliased variables according to the pure part of ψ , and edges represent dis-equalities and spatial atoms of ψ , e.g., a spatial atom $P(X, Y, \vec{x})$ is represented by a directed edge from X to Y labeled by $P(\vec{x})$. Thus, when a is a predicate atom $P(X, Y, \vec{x})$, $\psi_1[a]$ is obtained from the SL-graph of ψ_1 by selecting the edges reachable from X and co-reachable from Y . The graph selected for $\psi_1[a]$ is transformed into a tree t_1 , which is tested for membership in the language of a tree automaton built from the rules defining P for the atom $a = P(X, Y, \vec{x})$.

For the fragment SL_D^{ID} , SPEN implements a proof search strategy for the entailment problem $\psi_1 \Rightarrow \exists \vec{d}. \psi_2$. The strategy computes a sequence of formulas $\exists \vec{d}^1. \psi_1^1, \dots, \exists \vec{d}^n. \psi_1^n$ such that (1) $\exists \vec{d}^i. \psi_1^i \Rightarrow \exists \vec{d}^{i+1}. \psi_1^{i+1}$ and (2) $\exists \vec{d}^n. \psi_1^n$ is syntactically equivalent to $\exists \vec{d}. \psi_2$. The entailments in point (1) are obtained by applying the inductive rules and lemmas obtained automatically thanks to restriction required on inductive definitions. The procedure requires to check

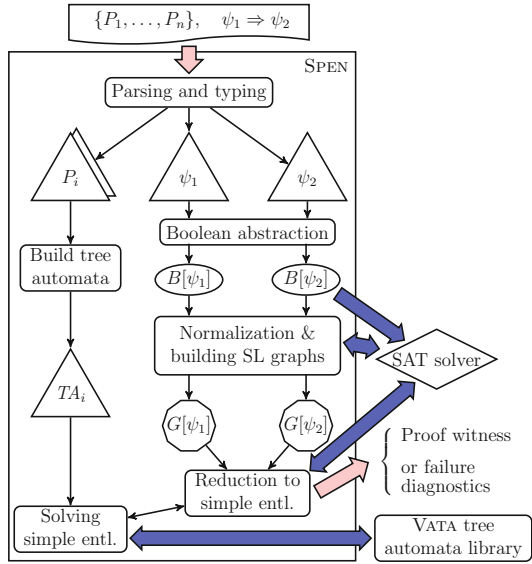


Fig. 2. SPEN workflow for entailment in SL_L^{ID}

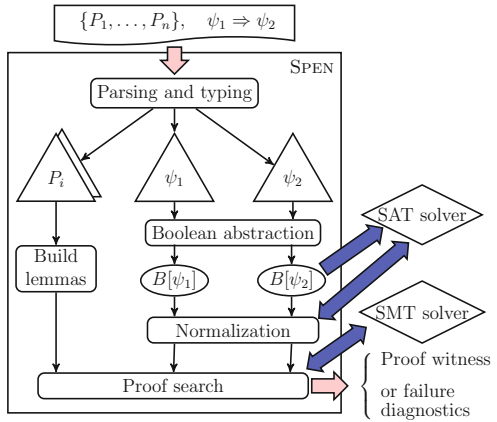


Fig. 3. SPEN workflow for entailment in SL_D^{ID}

Table 1. Experimental results on an Intel(R) Core(TM) i7-2600 CPU at 1.60 GHz

Fragments		Benchmark	Size	Time [s]		SL-COMP'14 results
SL _L ^{ID}	SL _D ^{ID}			SPEN _L	SPEN _D	Time [s] StarExec/solver
✓	✓	<code>s110_sat</code>	110	11.20	11.28	(I) 1.06/Asterix, (II) 3.27/SPEN
✓	✓	<code>s110_ent1</code>	292	34.45	34.94	(I) 2.98/Asterix, (II) 7.58/SPEN
✓	✓	<code>FDB_ent1</code>	43	1.08	1.00	(I) 0.61/SPEN, (II) 43.65/SLEEK
✓		<code>FDB_ent1</code> ⁺	55	0.65	—	—

entailments between data constraints, which is done using the previously mentioned wrapper to the SMT solver.

For both procedures, when the input entailment $\psi_1 \Rightarrow \exists \vec{d}. \psi_2$ holds, SPEN has the option of providing a proof witness that either indicates the fact that ψ_1 is unsatisfiable or it consists of the normalized forms of ψ_1 and ψ_2 and the mapping of sub-formulas in ψ_1 to atoms of ψ_2 . When the input entailment is not valid and the procedure terminates, SPEN provides a diagnosis that explains why the entailment fails.

5 Experimental Results

SPEN has been applied to a benchmark of 578 problems (available in the repository), 90% obtained from verification conditions of iterative programs on complex dynamic data structures. The remaining problems are crafted to test the capabilities of the solver. Tables 1 and 2 provide an overview of results obtained by SPEN on this benchmark.

The benchmark of SL_L^{ID} problems includes three divisions of SL-COMP'14: satisfiability and entailment problems for acyclic singly linked lists (`s110_sat` resp. `s110_ent1`), and entailment checking for formulas describing more complicated types of linked lists, e.g., doubly-linked lists, skip lists, and nested lists (`FDB_ent1`). SPEN spends less than 0.05s on 90% of the problems with the maximum time of 0.5s; these times include calls to a SAT solver. The benchmark `FDB_ent1`⁺ includes the problems not in the SL-COMP'14 benchmark (e.g., formulas describing lists of cyclic lists). The reported times in the last column have been obtained in 2014 on the StarExec³ platform.

The benchmark of SL_D^{ID} problems (see Table 2) includes verification conditions for proving the correctness of iterative procedures (delete, insert, search) over recursive data structures storing integer data: sorted lists, binary search trees, AVL trees, and red-black trees. SPEN spends less than 0.4s on each problem,

Table 2. Results for SL_D^{ID}

Benchmark	Size	Time [s]
<code>s110_sorted</code>	16	0.45
BST	45	1.67
AVL	22	1.21
RBT	21	3.61

³ www.starexec.org, an Intel(R) Xeon(R) CPU E5-2609 at 2.40 GHz of and 10 MB cache.

including calls to SAT and SMT solvers. The first three lines of Table 1 demonstrate that the two approaches implemented in SPEN (based on tree automata—column “SPEN_L”—and on proof search—column “SPEN_D”) are not only complementary but also comparable on the common fragment. The improvements discussed in this paper reduce the execution times by 10% within the division `s110_ent1` and by 30% within `FDB_ent1` w.r.t. the old version [6].

Acknowledgement. This work was supported by the French ANR project Vecolib, the Czech Science Foundation (project 17-12465S), the BUT FIT project FIT-S-17-4014, the IT4IXS: IT4Innovations Excellence in Science project (LQ1602), and by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 678177).

References

1. Berdine, J., Calcagno, C., O’Hearn, P.W.: A decidable fragment of separation logic. In: Lodaya, K., Mahajan, M. (eds.) FSTTCS 2004. LNCS, vol. 3328, pp. 97–109. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-30538-5_9](https://doi.org/10.1007/978-3-540-30538-5_9)
2. Brotherston, J., Gorogiannis, N., Petersen, R.L.: A generic cyclic theorem prover. In: Jhala, R., Igarashi, A. (eds.) APLAS 2012. LNCS, vol. 7705, pp. 350–367. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35182-2_25](https://doi.org/10.1007/978-3-642-35182-2_25)
3. Calcagno, C., Yang, H., O’Hearn, P.W.: Computability and complexity results for a spatial assertion language for data structures. In: Hariharan, R., Vinay, V., Mukund, M. (eds.) FSTTCS 2001. LNCS, vol. 2245, pp. 108–119. Springer, Heidelberg (2001). doi:[10.1007/3-540-45294-X_10](https://doi.org/10.1007/3-540-45294-X_10)
4. Chin, W.-N., David, C., Nguyen, H.H., Qin, S.: Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Sci. Comput. Program.* **77**(9), 1006–1036 (2012). Elsevier
5. De Moura, L., Bjørner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-78800-3_24](https://doi.org/10.1007/978-3-540-78800-3_24)
6. Enea, C., Lengál, O., Sighireanu, M., Vojnar, T.: Compositional entailment checking for a fragment of separation logic. In: Garrigue, J. (ed.) APLAS 2014. LNCS, vol. 8858, pp. 314–333. Springer, Cham (2014). doi:[10.1007/978-3-319-12736-1_17](https://doi.org/10.1007/978-3-319-12736-1_17)
7. Enea, C., Saveluc, V., Sighireanu, M.: Compositional invariant checking for overlaid and nested linked lists. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 129–148. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-37036-6_9](https://doi.org/10.1007/978-3-642-37036-6_9)
8. Enea, C., Sighireanu, M., Wu, Z.: On automated lemma generation for separation logic with inductive definitions. In: Finkbeiner, B., Pu, G., Zhang, L. (eds.) ATVA 2015. LNCS, vol. 9364, pp. 80–96. Springer, Cham (2015). doi:[10.1007/978-3-319-24953-7_7](https://doi.org/10.1007/978-3-319-24953-7_7)
9. Iosif, R., Rogalewicz, A., Vojnar, T.: Deciding entailments in inductive separation logic with tree automata. In: Cassez, F., Raskin, J.-F. (eds.) ATVA 2014. LNCS, vol. 8837, pp. 201–218. Springer, Cham (2014). doi:[10.1007/978-3-319-11936-6_15](https://doi.org/10.1007/978-3-319-11936-6_15)
10. Iosif, R., Rogalewicz, A., Simacek, J.: The tree width of separation logic with recursive definitions. In: Bonacina, M.P. (ed.) CADE 2013. LNCS (LNAI), vol. 7898, pp. 21–38. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-38574-2_2](https://doi.org/10.1007/978-3-642-38574-2_2)

11. O'Hearn, P., Reynolds, J., Yang, H.: Local reasoning about programs that alter data structures. In: Fribourg, L. (ed.) CSL 2001. LNCS, vol. 2142, pp. 1–19. Springer, Heidelberg (2001). doi:[10.1007/3-540-44802-0_1](https://doi.org/10.1007/3-540-44802-0_1)
12. Pérez, J.A.N., Rybalchenko, A.: Separation logic modulo theories. In: Shan, C. (ed.) APLAS 2013. LNCS, vol. 8301, pp. 90–106. Springer, Cham (2013). doi:[10.1007/978-3-319-03542-0_7](https://doi.org/10.1007/978-3-319-03542-0_7)
13. Piskac, R., Wies, T., Zufferey, D.: Automating separation logic using SMT. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 773–789. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-39799-8_54](https://doi.org/10.1007/978-3-642-39799-8_54)
14. Qiu, X., Garg, P., Stefanescu, A., Madhusudan, P.: Natural proofs for structure, data, and separation. In: Proceedings of PLDI 2013. ACM Press (2013)
15. Sighireanu, M., Cok, D.: Report on SL-COMP'14. JSAT **9**, 173–186 (2014)