



Frame Inference for Inductive Entailment Proofs in Separation Logic

Quang Loc Le¹ , Jun Sun² , and Shengchao Qin¹ 

¹ Teesside University, Middlesbrough, UK
 {Q.Le,S.Qin}@tees.ac.uk

² Singapore University of Technology and Design, Singapore, Singapore
 sunjun@sutd.edu.sg

Abstract. Given separation logic formulae A and C , frame inference is the problem of checking whether A entails C and simultaneously inferring residual heaps. Existing approaches on frame inference do not support inductive proofs with general inductive predicates. In this work, we present an automatic frame inference approach for an expressive fragment of separation logic. We further show how to strengthen the inferred frame through predicate normalization and arithmetic inference. We have integrated our approach into an existing verification system. The experimental results show that our approach helps to establish a number of non-trivial inductive proofs which are beyond the capability of all existing tools.

1 Introduction

Separation logic (SL) [20, 37] has been well established for reasoning about heap-manipulating programs (like linked-lists and trees). Often, SL is used in combination with inductive predicates to precisely specify data structures manipulated by a program. In the last decade, a large number of SL-based verification systems have been developed [1, 3, 6, 8, 13, 18, 19, 24, 29, 33, 36]. In these systems, SL is typically used to express assertions about program states. The problem of validating these assertions can be reduced to the *entailment* problem in SL, i.e., given two SL formulas Δ_a and Δ_c , to check whether $\Delta_a \models \Delta_c$ holds. Moreover, SL provides the frame rule [20], one prominent feature to enable compositional (a.k.a. modular) reasoning in the presence of the heap:

$$\text{FRAME RULE } \frac{\{P\}c\{Q\}}{\{P*F\}c\{Q*F\}}$$

where c is a program, P , Q and F are SL formulas, and $*$ is the separating conjunction in SL. Intuitively, $P*F$ states that P and F hold in disjoint heaps. This conjunction allows the frame rule to guarantee that F is unchanged under the action of c . This feature of SL is essential for scalability [6, 21, 44] as it allows the proof of a program to be decomposed (and reused) into smaller ones, e.g., proofs of procedures. To automate the application of the frame rule, SL-based

proof systems rely on a generalized form of the entailment, which is referred to as frame inference [1, 8, 12, 33, 39]. That is, given Δ_a and Δ_c , to check whether Δ_a entails Δ_c and simultaneously generate the residual heap, which is a satisfiable frame Δ_f capturing properties of the memory in Δ_a that is not covered by Δ_c . This problem, especially if Δ_a and Δ_c are constituted by general inductive predicates, is highly non-trivial as it may require inductive reasoning. Existing approaches [1, 33] are limited to specific predicates e.g., linked lists and trees. The systems reported in [8, 12, 39] do not adequately support the frame inference problem for inductive entailments in separation logic with predicate definitions and arithmetic.

In this work, we propose a sound approach for frame inference which aims to enhance modular verification in an expressive SL fragment with general inductive predicates and Presburger arithmetic. Intuitively, given an entailment $\Delta_a \models \Delta_c$, our goal is to infer a satisfiable frame axiom Δ_f such that $\Delta_a \models \Delta_c * \Delta_f$ holds. Our approach works as follows. We first augment the entailment checking with an *unknown* second-order variable $U_f(\bar{t})$ as a place-holder of the frame, where \bar{t} is a set of pointer-typed variables common in Δ_a and Δ_c . That is, the entailment checking becomes $\Delta_a \models \Delta_c * U_f(\bar{t})$. Afterwards, the following two steps are conducted. Firstly, we invoke a novel proof system to derive a *cyclic* proof for $\Delta_a \models \Delta_c * U_f(\bar{t})$ whilst inferring a predicate which U_f must satisfy so that the entailment is valid. We show that the cyclic proof is valid if this predicate is satisfiable. Secondly, we strengthen the inferred frame with shape normalization and arithmetic inference.

For the first step, we design a new cyclic proof system (e.g., based on [2, 3]) with an automated *cut rule* so as to effectively infer the predicate on U_f . A cyclic proof is a derivation tree whose root is the given entailment checking and whose edges are constructed by applying SL proof rules. A derivation tree of a cyclic proof may contain virtual *back-links*, each of which links a (leaf) node back to an ancestor. Intuitively, a back-link from a node l to an internal node i means that the proof obligation at l is induced by that at i . Furthermore, to avoid potentially unsound cycles (i.e., self-cycles), a global soundness condition must be imposed upon these derivations to qualify them as genuine proofs. In this work, we develop a sequent-based cyclic proof system with a *cyclic cut* rule so as to form back-links effectively and check the soundness condition eagerly. Furthermore, we show how to extract lemmas from the proven cyclic proofs and reuse them through lemma application for an efficient proof system. These synthesized lemmas work as *dynamic cuts* in the proposed proof system.

For the second step, we strengthen the inferred predicate on the frame $U_f(\bar{t})$ so that it becomes more powerful in establishing correctness of certain programs. In particular, the inferred frame is strengthened with predicate normalization and arithmetic inference. The normalization includes predicate split (i.e., to expose the spatial separation of the inferred frame) and predicate equivalence (i.e., to relate the inferred frame with user-supplied predicates). The arithmetic inference discovers predicates on pure properties (size, sum, height, content and bag) to support programs which require induction reasoning on both shape and data properties.

Lastly, we have implemented the proposal and integrated it into a modular verification engine. Our experiments show that our approach infers strong frames which enhances the verification of heap-manipulating programs.

2 Preliminaries

In this section, we present the fragment of SL which is used as the assertion language in this work. This fragment, described in Fig. 1, is expressive enough for specifying and verifying properties of a variety of data structures [24–26, 35, 41]. We use \bar{t} to denote a sequence of terms and occasionally use a sequence (i.e., \bar{t}) to denote a set when there is no ambiguity. A formula Φ in our language is a disjunction of multiple clauses Δ , each of which is a conjunction of a spatial predicate κ and a pure (non-heap) constraint π . The spatial predicate κ captures properties of the heap whereas π captures properties of the data. κ can be an empty heap \mathbf{emp} , or a points-to predicate $r \mapsto c(\bar{v})$ where c is a data structure, or a user-defined predicate $P(\bar{t})$ or a spatial conjunction $\kappa_1 * \kappa_2$. \mathbf{null} is a special heap location. A pure constraint π is in the form of (dis)equality α (on pointers) and Presburger arithmetic ϕ . We write $v_1 \neq v_2$ and $v \neq \mathbf{null}$ for $\neg(v_1 = v_2)$ and $\neg(v = \mathbf{null})$, respectively. We often omit the pure part of a formula Φ when it is \mathbf{true} . For standardizing the notations, we use uppercase letters for unknown (to-be-inferred) predicates, (e.g., $P(\bar{t})$) and lowercase letters (e.g., $p(\bar{t})$) for known predicates.

Formula	$\Phi ::= \Delta \mid \Phi_1 \vee \Phi_2$	$\Delta ::= \exists \bar{v}. (\kappa \wedge \pi)$
Spatial	$\kappa ::= \mathbf{emp} \mid x \mapsto c(\bar{v}) \mid P(\bar{t}) \mid \kappa_1 * \kappa_2$	
Pure	$\pi ::= \mathbf{true} \mid \mathbf{false} \mid \alpha \mid \phi \mid p(\bar{t}) \mid \exists v. \pi \mid \forall v. \pi \mid \neg \pi \mid \pi_1 \wedge \pi_2 \mid \pi_1 \vee \pi_2$	
Equality	$\alpha ::= v_1 = v_2 \mid v = \mathbf{null}$	
Pres. arith.	$\phi ::= a \mid \phi_1 = \phi_2 \mid \phi_1 \leq \phi_2$ $a ::= k^{\mathbf{int}} \mid v \mid k^{\mathbf{int}} \times a \mid a_1 + a_2 \mid -a \mid \max(a_1, a_2) \mid \min(a_1, a_2)$	
Lemma	$\iota ::= P(\bar{v}) \wedge \pi \bowtie \exists \bar{w}. \kappa \wedge \pi$	$\bowtie ::= \rightarrow \mid \leftarrow \mid \leftrightarrow$

Fig. 1. Syntax

A user-defined (inductive) predicate $P(\bar{v})$ with parameters \bar{v} is defined in the form of a disjunction, i.e., $\mathbf{pred} P(\bar{v}) \equiv \Phi$, where each disjunct in Φ is referred to as a branch. In each branch, variables that are not in \bar{v} are implicitly existentially-quantified. We use function $\mathbf{unfold}(P(\bar{t}))$ to replace an occurrence of inductive predicates by the disjuncts in the definition of P with actual/formal parameters renaming. For example, the following predicates \mathbf{lseg} and \mathbf{lsegn} are defined to express list segments where every node contains the same value 1, given data structure $\mathbf{node}\{\mathbf{int} \mathbf{val}; \mathbf{node} \mathbf{next}; \}$.

```

pred lseg(root,l) $\equiv \mathbf{emp} \wedge \mathbf{root} = l \vee \exists q. \mathbf{root} \mapsto \mathbf{node}(1, q) * \mathbf{lseg}(q, l);$ 
pred lsegn(root,l,n) $\equiv \mathbf{emp} \wedge \mathbf{root} = l \wedge n = 0 \vee \exists q. \mathbf{root} \mapsto \mathbf{node}(1, q) * \mathbf{lsegn}(q, l, n - 1);$ 
    
```

where \mathbf{root} is the head, l the end of the segment and n the length of the segment.

In our framework, we may have lemmas to assist program verification. A lemma ι of the form $\Delta_l \rightarrow \Delta_r$, which means that the entailment $\Delta_l \models \Delta_r$ holds. We write $A \leftrightarrow B$, a short form of $A \rightarrow B$ and $B \rightarrow A$, to denote a two-way lemma. If $A \leftrightarrow B$, A is semantically equivalent to B . We use E and F to denote an entailment problem.

In the following, we discuss semantics of the SL fragment. Concrete heap models assume a fixed finite collection *Node*, a fixed finite collection *Fields*, a disjoint set *Loc* of locations (i.e., heap addresses), a set of non-address values *Val* such that $\text{null} \in \text{Val}$ and $\text{Val} \cap \text{Loc} = \emptyset$. The semantics is given by a satisfaction relation: $s, h \models \Phi$ that forces the stack s and heap h to satisfy the constraint Φ where $h \in \text{Heaps}$, $s \in \text{Stacks}$, and Φ is a formula. *Heaps* and *Stacks* are defined as follows.

$$\begin{aligned} \text{Heaps} &\stackrel{\text{def}}{=} \text{Loc} \rightarrow_{\text{fin}} (\text{Node} \times \text{Fields} \rightarrow \text{Val} \cup \text{Loc}) \\ \text{Stacks} &\stackrel{\text{def}}{=} \text{Var} \rightarrow \text{Val} \cup \text{Loc} \end{aligned}$$

The details of semantics of this SL fragment follow the one in [25].

```

1 node append(node x, node y)
2   requires ll_n(x,i)*ll_n(y,j)^i>0
3   ensures ll_n(res,i+j);
4   { // ll_n(x,i)*ll_n(y,j)^i>0
5     node t=last(x);
6     // ll_last(x,t,i)*ll_n(y,j)^i>0
7     t->next=y;
8     // alpha:l_segn(x,t,i-1)*t->node(y)*ll_n(y,j)^i>0
9     return x; }
8 node last(node x)
9   requires ll_n(x,n)^n>0
10  ensures ll_last(x,res,n);
11  { if(!x->next) return x;
12    else return last(x->next); }

```

Fig. 2. Code of `append`.

3 Illustrative Example

In the following, we first discuss the limitation of the existing entailment procedures [1, 8] to the frame inference problem. Given an entailment, these procedures deduce it until the following subgoal is obtained: $\Delta_a \vdash \text{emp} \wedge \text{true}$. Then, they conclude that Δ_a is the residual frame. However, these approaches provide limited support for proofs of induction. While [1] provides inference rules as a sequence of inductive reasoning for hardwired lists and trees, our previous work [8] supports inductive proofs via user-supplied lemmas [30]. Hence, it is very hard for these procedures to automatically infer the frame for the entailments which require proofs of induction.

We illustrate our approach via the verification of the `append` method shown in Fig. 2, which appends a singly-linked list referred to by y to the end of the singly-linked list referred to by x . It uses the auxiliary procedure `last` (lines 8–12) to obtain the pointer referring to the last node in the list. Each node object x has a data value $x \rightarrow \text{data}$ and a next pointer $x \rightarrow \text{next}$. For simplicity, we assume

that every node in the x list and the y list has data value 1. The correctness of `append` and `last` is specified using our fragment of SL with a pre-condition (`requires`) and a post-condition (`ensures`). The auxiliary variable `res` denotes the return value of the procedure. Note that these specifications refer to the user-provided predicates `lln` and `ll_last`, which are defined as follows.

$$\begin{aligned} \text{pred } \text{lln}(\text{root}, n) &\equiv \text{emp} \wedge \text{root} = \text{null} \wedge n = 0 \vee \exists q. \text{root} \mapsto \text{node}(1, q) * \text{lln}(q, n-1); \\ \text{pred } \text{ll_last}(\text{root}, l, n) &\equiv 1 \mapsto \text{node}(1, \text{null}) \wedge \text{root} = l \wedge n = 1 \\ &\vee \exists q. \text{root} \mapsto \text{node}(1, q) * \text{ll_last}(q, l, n-1); \end{aligned}$$

Intuitively, the predicate `lln(root, n)` is satisfied if `root` points to a singly-linked list with n nodes. The predicate `ll_last(t, p, n)` is satisfied if t points to a list segment with last element p and length n . In our framework, we provide a library of commonly used inductive predicates (and the corresponding lemmas), including for example the definitions for list segments `lseg` and `lsegn` introduced earlier. Given these specifications, we automatically deduce predicates on the intermediate program states (using existing approaches [8]), shown as comments in Fig. 2, as well as the following three entailment checks that must be established in order to verify the absence of memory errors and the correctness of the method `append`.

$$\begin{aligned} E_1: & \text{lln}(x, i) * \text{lln}(y, j) \wedge i > 0 \vdash \exists n_1. \text{lln}(x, n_1) \wedge n_1 > 0 \\ E_2: & \text{ll_last}(x, t, i) * \text{lln}(y, j) \wedge i > 0 \vdash \exists q, v. t \mapsto \text{node}(v, q) \\ E_3: & \text{lsegn}(\text{res}, t, i-1) * t \mapsto \text{node}(1, y) * \text{lln}(y, j) \wedge i > 0 \vdash \text{lln}(\text{res}, i+j) \end{aligned}$$

E_1 aims to establish a local specification at line 5 which we generate automatically. E_2 must be satisfied so that no null-dereference error would occur for the assignment to `t` \rightarrow `next` at line 6. E_3 aims to establish that the postcondition is met. Frame inference is necessary in order to verify the program. In particular, frame inference for E_2 is crucial to construct a precise heap state after line 6, i.e., the state α in the figure, which is necessary to establish E_3 . Furthermore, the frame of E_3 (which is inferred as `emp`) helps to show that this program does not leak memory. As the entailment checks E_2 and E_3 require both induction reasoning and frame inference, they are challenging for existing SL proof systems [3, 8, 9, 12, 15, 31, 36, 40]. In what follows, we illustrate how our system establishes a cyclic proof with frame inference for E_2 .

Frame Inference. Our frame inference starts with introducing an unknown predicate (a second-order variable) $U_1(x, t, q, v, y)$ ¹ as the initial frame, which is a placeholder for a heap predicate on variables x , t , q and y (i.e., variables referred to in E_2). That is, E_2 is transformed to the following entailment checking problem:

$$F_2: \text{ll_last}(x, t, i) * \text{lln}(y, j) \wedge i > 0 \vdash_{L_0} \exists q, v. t \mapsto \text{node}(v, q) * U_1(x, t, q, v, y)$$

where L_0 is a set of induction hypotheses and sound lemmas. This set is accumulated automatically during the proof search and used for constructing cyclic

¹ In implementation, we add `#` annotation into instantiated variables and non-heap variables to guide proof search which are not shown here.

proofs and lemma application. If a hypothesis is proven, it becomes a lemma and may be applied later during the proof search. In this example, initially $L_0 = \emptyset$. The proposed proof system derives a cyclic proof for the entailment problem and, at the same time, infers a set of constraints \mathcal{R} for $\mathbb{U}_1(x, t, q, v, y)$ such that the proof is valid if the system \mathcal{R} is satisfiable. Each constraint in \mathcal{R} has the form of logical implication i.e., $\Delta_b \Rightarrow \mathbb{U}(\bar{v})$ where Δ_b is the body and $\mathbb{U}(\bar{v})$ is the head (a second-order variable). For \mathbb{F}_2 , the following two constraints are inferred, denoted by σ_1 and σ_2 .

$$\begin{aligned}\sigma_1: \text{lln}(y, j) \wedge t = x \wedge q = \text{null} \wedge v = 1 &\Rightarrow \mathbb{U}_1(x, t, q, v, y) \\ \sigma_2: x_2 \mapsto \text{node}(1, x) * \mathbb{U}_1(x, t, q, v, y) &\Rightarrow \mathbb{U}_1(x_2, t, q, v, y)\end{aligned}$$

We then use a decision procedure (e.g., **S2SAT_{SL}** [25, 26] or [4]) to check the satisfiability of $\sigma_1 \wedge \sigma_2$. Note that we write a satisfiable definition of $(\Delta_1 \Rightarrow \mathbb{U}(\bar{v})) \wedge (\Delta_2 \Rightarrow \mathbb{U}(\bar{v}))$ in the equivalent form of $\mathbb{U}(\bar{v}) \equiv \Delta_1 \vee \Delta_2$. For instance, the above constraints are written as:

$$\begin{aligned}\mathbb{U}_1(\text{root}, t, q, v, y) &\equiv \text{lln}(y, j) \wedge \text{root} = t \wedge q = \text{null} \wedge v = 1 \\ &\vee \exists q_1. \text{root} \mapsto \text{node}(1, q_1) * \mathbb{U}_1(q_1, t, q, v, y);\end{aligned}$$

Note that, in the above definition of \mathbb{U}_1 , the separation of those heap-lets referred to by **root**, y and q is not explicitly captured. Additionally, relations over the sizes are also missing. Such information is necessary in order to establish the left-hand side of \mathbb{E}_3 . The successful verification of \mathbb{E}_3 in turn establishes the post-condition of method **append**. In the following we show how to strengthen the inferred frame.

Frame Strengthening. We strengthen \mathbb{U}_1 with spatial separation constraints on the pointer variables **root**, y and q . To explicate the *spatial* separation among these pointers, our system generates the following equivalent lemma and splits \mathbb{U}_1 into two disjoint heap regions (with $*$ conjunction):

$$\mathbb{U}_1(\text{root}, t, q, v, y) \equiv \mathbb{U}_2(\text{root}, t) * \text{lln}(y, j) \wedge q = \text{null} \wedge v = 1$$

where \mathbb{U}_2 is a new auxiliary predicate with an inferred definition:

$$\mathbb{U}_2(\text{root}, t) \equiv \text{emp} \wedge \text{root} = t \vee \exists q_1. \text{root} \mapsto \text{node}(1, q_1) * \mathbb{U}_2(q_1, t)$$

Next, our system detects that \mathbb{U}_2 is equivalent to the user-defined predicate **lseg**, and generates the lemma: $\mathbb{U}_2(\text{root}, t) \leftrightarrow \text{lseg}(\text{root}, t)$. Relating \mathbb{U}_2 to **lseg** enhances the understanding of the inferred predicates. Furthermore, as shown in [9], this relation helps to reduce the requirements of induction reasoning among equivalent inductive predicates with different names. Substituting \mathbb{U}_2 with the equivalent **lseg**, \mathbb{U}_1 becomes:

$$\mathbb{U}_1(\text{root}, t, q, v, y) \equiv \text{lseg}(\text{root}, t) * \text{lln}(y, j) \wedge q = \text{null} \wedge v = 1$$

This definition states that frame \mathbb{U}_1 holds in two disjoint heaps: one list segment pointed to by **root** and a list pointed to by y . After substitution the entailment \mathbb{F}_2 becomes

$$\text{ll_last}(x, t, i) * \text{lln}(y, j) \wedge i > 0 \vdash_{L_0} t \mapsto \text{node}(1, \text{null}) * \text{lseg}(x, t) * \text{lln}(y, j)$$

Next, we further strengthen the frame with pure properties, which is necessary to successfully establish the left hand side of E_3 . In particular, we generate constraints to capture that the numbers of allocated heaps in the left hand side and the right hand side of F_2 are identical. Our system obtains these constraints through two phases. First, it automatically augments an argument for each inductive predicate in F_2 to capture its size property. Concretely, it detects that while predicates `ll_last` and `lln` have such size argument already, the shape-based frame `lseg` has not. As so, it extends `lseg(root,t)` to obtain the predicate `lsegn(root,t,m)` where the size property is captured by parameter m . Now, we substitute the `lsegn` into F_2 to obtain:

$$\text{ll_last}(x,t,i)*\text{lln}(y,j)\wedge i>0 \vdash_{L_0} \exists k.t \mapsto \text{node}(1,\text{null})*\text{lsegn}(x,t,k)*\text{lln}(y,j)$$

After that, we apply the same three steps of frame inference to generate the size constraint: constructing unknown predicates, proving entailment and inferring a set of constraints and checking satisfiability. For the first step, the above entailment is enriched with one unknown (pure) predicate: $P_1(i,j,k)$ which is the place-holder for arithmetical constraints among size variables i , j and k . The augmented entailment checking is:

$$\begin{aligned} & \text{ll_last}(x,t,i)*\text{lln}(y,j)\wedge i>0 \\ & \vdash_{L_0} \exists k.\text{lsegn}(x,t,k)*t \mapsto \text{node}(1,\text{null})*\text{lln}(y,j)\wedge P_1(i,j,k) \end{aligned}$$

Secondly, our system successfully derives a proof for the above entailment under condition that the following disjunctive set of two constraints is satisfiable.

$$\begin{aligned} \sigma_3: i=1\wedge k=0 & \Rightarrow P_1(i,j,k) \\ \sigma_4: i_1=i-1\wedge k_1=k-1\wedge i>0\wedge P_1(i_1,j,k_1) & \Rightarrow P_1(i,j,k) \end{aligned}$$

$$\begin{array}{c} \text{EMP} \frac{\pi_a \implies \pi_c}{\text{emp}\wedge\pi_a \vdash_L \text{emp}\wedge\pi_c \rightsquigarrow \text{true}} \quad \text{CCUT} \frac{\begin{array}{c} \Delta_l \rightarrow \Delta_r \in L \quad \Delta_{a_1} \rho \vdash_L \Delta_l \rightsquigarrow \mathcal{R}_1 \\ \Delta_r \rho * \Delta_{a_2} \vdash_L \Delta_c \rightsquigarrow \mathcal{R}_2 \end{array}}{\Delta_{a_1} * \Delta_{a_2} \vdash_L \Delta_c \rightsquigarrow \mathcal{R}_1 \wedge \mathcal{R}_2} \text{gsc} \\ \\ \text{PRED-M} \frac{\Delta_a \vdash_L \Delta_c \wedge \text{freeEQ}([\bar{v}/\bar{w}]) \rightsquigarrow \mathcal{R}}{P(r,\bar{v}) * \Delta_a \vdash_L P(r,\bar{w}) * \Delta_c \rightsquigarrow \mathcal{R}} \quad \text{M} \frac{\Delta_a \vdash_L \Delta_c \wedge \text{freeEQ}([\bar{v}/\bar{w}]) \rightsquigarrow \mathcal{R}}{x \mapsto c(\bar{v}) * \Delta_a \vdash_L x \mapsto c(\bar{w}) * \Delta_c \rightsquigarrow \mathcal{R}} \\ \\ \text{RU} \frac{\begin{array}{c} (x \mapsto c(\bar{w}) \text{ or } x = \text{null} \text{ or } x = q) \in \Delta_a \\ \text{unfold}(P(x,\bar{v})) = \{\Delta_1, \dots, \Delta_n\} \quad \exists i \in \{1 \dots n\} \cdot \Delta_a \vdash_L \Delta_i * \Delta_c \rightsquigarrow \mathcal{R}_i \end{array}}{\Delta_a \vdash_L P(x,\bar{v}) * \Delta_c \rightsquigarrow \mathcal{R}_i} \\ \\ \text{LU} \frac{\begin{array}{c} (x \mapsto c(\bar{w}) \text{ or } x = \text{null} \text{ or } x = q) \in \Delta_c \quad \text{unfold}(P(x,\bar{v})) = \{\Delta_1, \dots, \Delta_n\} \\ L'_i = L \cup \{(P(x,\bar{v}) * \Delta_a \rightarrow \Delta_c)\} \quad \forall i \in \{1 \dots n\} \cdot \Delta_i * \Delta_a \vdash_{L'_i} \Delta_c \rightsquigarrow \mathcal{R}_i \end{array}}{P(x,\bar{v}) * \Delta_a \vdash_L \Delta_c \rightsquigarrow \bigwedge \mathcal{R}_i} \end{array}$$

Fig. 3. Basic inference rules for entailment procedure (where gsc is global soundness condition)

Lastly, to check whether the $\sigma_3 \wedge \sigma_4$ is satisfiable, we automatically compute the closure form for $\sigma_3 \wedge \sigma_4$ as: $P_1(i, j, k) \equiv k = i - 1 \wedge i > 0$. This formula is satisfiable and substituted into the frame as: $\mathbf{1segn}(x, t, k) * \mathbf{1ln}(y, j) \wedge q = \mathbf{null} \wedge v = 1 \wedge k = i - 1 \wedge i > 0$.

4 Frame Inference

In this section, we present our approach for frame inference in detail. Given an entailment $\Delta_a \vdash \Delta_c$, where Δ_a is the antecedent (LHS) and Δ_c is the consequence (RHS), our system attempts to infer a frame Δ_f such that when a frame is successfully inferred, the validity of the entailment $\Delta_a \vdash \Delta_c * \Delta_f$ is established at the same time.

Our approach has three main steps. Firstly, we enrich RHS with an unknown predicate in the form of $\mathbf{U}(\bar{v})$ to form the entailment $\Delta_a \vdash_L \Delta_c * \mathbf{U}(\bar{v})$ where \bar{v} includes all free pointer-typed variables of Δ_a and Δ_c and L is the union of a set of user-supplied lemmas and a set of induction hypotheses (initially \emptyset). Among these, the parameters are annotated with $\#$ following the principle that instantiation (and subtraction) must be done before inference. The detail is as follows: (i) all common variables of Δ_a and Δ_c are $\#$ -annotated; (ii) points-to pointers of Δ_c are $\#$ -annotated; (iii) the remaining pointers are not $\#$ -annotated. In the implementation, inference of frame predicates is performed incrementally such that shape predicates are inferred prior to pure ones. Secondly, we construct a proof of the entailment and infer a set of constraints \mathcal{R} for $\mathbf{U}(\bar{v})$. Thirdly, we check the satisfiability of \mathcal{R} using the decision procedure in [25, 26].

In the following, we present our entailment checking procedure with a set of proof rules shown in Figs. 3 and 4. For each rule, the obligation is at the bottom and its reduced form is on the top. In particular, the rules in Fig. 3 are used for entailment proving (i.e., to establish a cyclic proof) and the rules in Fig. 4 are used for predicate inference.

$$\begin{array}{c}
 \sigma \equiv \mathbf{R}(r, \bar{t}) * \mathbf{U}_i(r_i, \bar{t}_i \#) * \mathbf{U}_f(\bar{w}, \bar{t}', \bar{z} \#, r \#) \wedge \nabla(\bar{w} \cup \{r, r_i\}, \pi_1) \Rightarrow \mathbf{U}(r, \bar{w}, \bar{z} \#) \\
 r_i \in \bar{t} \quad \bar{t}' = \bar{t} \setminus (\bar{w} \cup \bar{z} \cup \{r\}) \quad \kappa_1 \wedge \pi_1 \vdash_L \mathbf{U}_f(\bar{w}, \bar{t}', \bar{z} \#, r \#) * \kappa_2 \wedge \pi_2 \rightsquigarrow \mathcal{R} \\
 \text{AF} \frac{}{\mathbf{R}(r, \bar{t}) * \mathbf{U}_i(r_i, \bar{t}_i \#) * \kappa_1 \wedge \pi_1 \vdash_L \mathbf{U}(r, \bar{w}, \bar{z} \#) * \kappa_2 \wedge \pi_2 \rightsquigarrow \sigma \wedge \mathcal{R}} \\
 \\
 \sigma \equiv \mathbf{U}(r, \bar{w}, \bar{z} \#) \Rightarrow \mathbf{R}(r, \bar{t}) * \mathbf{U}_i(r_i, \bar{t}_i \#) * \mathbf{U}_f(\bar{w}, \bar{t}', \bar{z} \#, r \#) \wedge \nabla(\bar{w} \cup \{r, r_i\}, \pi_1) \\
 r_i \in \bar{t} \quad \bar{t}' = \bar{t} \setminus (\bar{w} \cup \bar{z} \cup \{r\}) \quad \mathbf{U}_f(\bar{w}, \bar{t}', \bar{z} \#, r \#) * \kappa_1 \wedge \pi_1 \vdash_L \kappa_2 \wedge \pi_2 \rightsquigarrow \mathcal{R} \\
 \text{AF-F} \frac{}{\mathbf{U}(r, \bar{w}, \bar{z} \#) * \kappa_1 \wedge \pi_1 \vdash_L \mathbf{R}(r, \bar{t}) * \mathbf{U}_i(r_i, \bar{t}_i \#) * \kappa_2 \wedge \pi_2 \rightsquigarrow \sigma \wedge \mathcal{R}}
 \end{array}$$

Fig. 4. Inference rules with predicate synthesis.

Given an entailment check in the form of $\Delta_a \vdash_L \Delta_c$, the rules shown in Fig. 3 are designed to subtract the heap (via the rules $\underline{\mathbf{M}}$ and $\underline{\mathbf{PRED-M}}$) on both sides until their heaps are empty. After that, it checks the validity for the implication of two pure formulas by using an SMT solver, like Z3 [27], as shown in rule $\underline{\mathbf{EMP}}$. Algorithmically, this entailment checking is performed as follows.

- **Matching.** The rules $\boxed{\mathbf{M}}$ and $\boxed{\mathbf{PRED-M}}$ are used to match up identified heap chains. Starting from identified root pointers, the procedure keeps matching all their reachable heaps. It unifies corresponding fields of matched roots by using the following auxiliary function $\text{freeEQ}(\rho)$: $\text{freeEQ}([u_i/v_i]_{i=1}^n) = \bigwedge_{i=1}^n \{u_i = v_i\}$.
- **Unfolding.** The rules $\boxed{\mathbf{LU}}$ and $\boxed{\mathbf{RU}}$ are used to derive alternative heap chains. While rule $\boxed{\mathbf{LU}}$ presents the unfolding in the antecedent, $\boxed{\mathbf{RU}}$ in the consequent.
- **Applying Lemma.** Rule $\boxed{\mathbf{CCUT}}$ derives yet other alternative heap chains. For LHS which has at least one \mathbf{UD} predicate, we attempt to apply a lemma as an alternative search using $\boxed{\mathbf{CCUT}}$ rule. We notice that as we assume that a lemma which is supplied by the user is valid, applying this lemma does not requires the global condition.

Cyclic Proof. The proof rules in Fig. 3 are designed to establish cyclic proofs. In the following, we briefly describe a cyclic proof technique enhancing the proposal in [2].

Definition 1 (Pre-proof). A pre-proof of entailment E is a pair $(\mathcal{T}_i, \mathcal{L})$ where \mathcal{T}_i is a derivation tree and \mathcal{L} is a back-link function such that: the root of \mathcal{T}_i is E ; for every edge from E_i to E_j in \mathcal{T}_i , E_i is a conclusion of an inference rule with a premise E_j . There is a back-link between E_c and E_l if there exists $\mathcal{L}(E_l) = E_c$ (i.e., $E_c = E_l \theta$ with some substitution θ); and for every leaf E_l , E_l is an axiom rule (without conclusion).

If $\mathcal{L}(E_l) = E_c$, E_l (resp. E_c) is referred as a bud (resp. companion).

Definition 2 (Trace). Let $(\mathcal{T}_i, \mathcal{L})$ be a pre-proof of $\Delta_a \vdash_L \Delta_c$; $(\Delta_{a_i} \vdash_{L_i} \Delta_{c_i})_{i \geq 0}$ be a path of \mathcal{T}_i . A trace following $(\Delta_{a_i} \vdash_{L_i} \Delta_{c_i})_{i \geq 0}$ is a sequence $(\alpha_i)_{i \geq 0}$ such that each α_i (for all $i \geq 0$) is an instance of the predicate $\mathbf{P}(\bar{t})$ in the formula Δ_{a_i} , and either:

- α_{i+1} is the subformula containing an instance of $\mathbf{P}(\bar{t})$ in $\Delta_{a_{i+1}}$;
- or $\Delta_{a_i} \vdash_{L_i} \Delta_{c_i}$ is the conclusion of an unfolding rule, α_i is an instance predicate $\mathbf{P}(\bar{t})$ in Δ_{a_i} and α_{i+1} is a subformula $\Delta[\bar{t}/\bar{v}]$ which is a definition rule of the inductive predicate $\mathbf{P}(\bar{v})$. i is a progressing point of the trace.

To ensure that a pre-proof is sound, a global *soundness condition* must be imposed to guarantee well-foundedness.

Definition 3 (Cyclic proof). A pre-proof $(\mathcal{T}_i, \mathcal{L})$ of $\Delta_a \vdash_L \Delta_c$ is a cyclic proof if, for every infinite path $(\Delta_{a_i} \vdash_{L_i} \Delta_{c_i})_{i \geq 0}$ of \mathcal{T}_i , there is a tail of the path $p = (\Delta_{a_i} \vdash_{L_i} \Delta_{c_i})_{i \geq n}$ such that there is a trace following p which has infinitely progressing points.

Brotherston *et al.* proved [2] that $\Delta_a \vdash \Delta_c$ holds if there is a cyclic proof of $\Delta_a \vdash_{\emptyset} \Delta_c$ where Δ_a and Δ_c do not contain any unknown predicate.

In the following, we explain how cyclic proofs are constructed using the proof rules shown in Fig. 3. $\boxed{\mathbf{LU}}$ and $\boxed{\mathbf{CCUT}}$ are the most important rules for forming

back-links and then pre-proof construction. While rule $\llbracket \text{LU} \rrbracket$ accumulates possible companions and stores them in historical sequents L , $\llbracket \text{CCUT} \rrbracket$ links a bud with a companion using some substitutions as well as checks the global soundness condition eagerly. Different to the original cyclic system [3], our linking back function only considers companions selected in the set of historical sequents L . Particularly, $\Delta_l \rightarrow \Delta_r \in L$ is used as an intelligent cut as follows. During proof search, a subgoal (i.e., $\Delta_{a_1} * \Delta_{a_2} \vdash_L \Delta_c$) may be matched with the above historical sequent to form a cycle and close the proof branch using the following principle. First, $\Delta_l \vdash \Delta_r$ is used as an induction hypothesis. As so, we have $\Delta_l \rho * \Delta_{a_2} \models \Delta_r \rho * \Delta_{a_2}$ where ρ are substitutions including those for avoiding clashing of variables between Δ_r and Δ_{a_2} . If both $\Delta_{a_1} * \Delta_{a_2} \vdash_L \Delta_l \rho * \Delta_{a_2}$ and $\Delta_r \rho * \Delta_{a_2} \vdash_L \Delta_c$ are proven, then we have:

$$\Delta_{a_1} * \Delta_{a_2} \Longrightarrow \Delta_l \rho * \Delta_{a_2} \Longrightarrow \Delta_r \rho * \Delta_{a_2} \Longrightarrow \Delta_c.$$

Thus, the subgoal $\Delta_{a_1} * \Delta_{a_2} \vdash_L \Delta_c$ holds. We remark that if a hypothesis is proven, it can be applied as a valid lemma subsequently.

In our system, often a lemma includes universally quantified variables. We thus show a new mechanism to instantiate those lemmas that include universally quantified variables. We denote constraints with universal variables as universal guards $\forall G$. A universal guard $\forall G$ is equivalent to an *infinite* conjunction $\bigwedge_{\rho} G[\rho]$. Linking a leaf with universal guards is not straightforward. For illustration, let us consider the following bud B_0 and the universally quantified companion/lemma $C_0 \in L$.

$$\begin{aligned} B_0: & \text{lsegn}(\text{root}, \text{null}, n) \wedge n = 10 \vdash_L \exists r. \text{lsegn}(\text{root}, r, 3) * \text{lsegn}(r, \text{null}, 7) \\ C_0: & \forall a, b. \text{lsegn}(\text{root}, \text{null}, n) \wedge n = a + b \wedge a \geq 0 \wedge b \geq 0 \\ & \rightarrow \exists r. \text{lsegn}(\text{root}, r, a) * \text{lsegn}(r, \text{null}, b) \end{aligned}$$

As shown in rule $\llbracket \text{CCUT} \rrbracket$, to link B_0 back to C_0 , the LHS of these two entailments must be implied through some substitution. To obtain that, we propose *lemma instantiation*, a sound solution for universal lemma application. Based on the constraints in the LHS of the bud, our technique instantiates a universally quantified guard (of the selected companion/lemma) before linking it back. Concretely, we replace the universal guard by a *finite* set of its instances; an instantiation of a formula $\forall \bar{v} G(\bar{t})$ is $G(\bar{t})[\bar{w}/\bar{v}]$ for some vector of terms \bar{w} . These instances are introduced based on the instantiations in both LHS and RHS of the corresponding bud e.g., $n = 10 \wedge a = 3 \wedge b = 7$ in B_0 .

Frame Inference. The two inference rules shown in Fig. 4 are designed specifically to infer constraints for frame. In these rules, $\nabla(\bar{w}, \pi)$ is an auxiliary function that existentially quantifies free variables in π that are not in the set \bar{w} . This function extracts relevant arithmetic constraints to define data contents of the unknown predicates. $\mathbf{R}(r, \bar{t})$ is either $r \mapsto c(\bar{t})$ or a known (defined) predicate $\mathbf{P}(r, \bar{t})$, or an unknown predicate $\mathbf{U}'(r, \bar{t}, \bar{w}\#)$. The $\#$ in the unknown predicates is used to guide inference and proof search. We only infer on pointers without $\#$ -annotation. $\mathbf{U}_{\#}(\bar{w}, \bar{t}')$ is another unknown predicate which is generated to infer the shape of

pointers \bar{w} . Inferred pointers are annotated with $\#$ to avoid double inference. A new unknown predicate \mathbf{U}_r is generated only if there exists at least one parameter not to be annotated with $\#$ (i.e., $\bar{w} \cup t' \neq \emptyset$). To avoid conflict between the inference rules and the other rules (e.g., unfolding and matching), root pointers of a heap formula must be annotated with $\#$ in unknown predicates. For example, in our system while $x \mapsto_{\mathbf{c}_1}(y) * \mathbf{U}_1(x\#, y)$ is legal, $x \mapsto_{\mathbf{c}_1}(y) * \mathbf{U}_1(x, y)$ is illegal. Our system applies subtraction on the heap pointed to by x rather than inference for the following check: $x \mapsto_{\mathbf{c}_1}(\text{null}) \vdash_{L} x \mapsto_{\mathbf{c}_1}(y) * \mathbf{U}_1(x\#, y)$.

Soundness. The soundness of the inference rules in Fig. 3 has been shown in unfold-and-match systems for general inductive predicates [3, 8]. In the following, we present the soundness of the inference rules in Fig. 4. We introduce the notation $\mathcal{R}(\Gamma)$ to denote a set of predicate definitions $\Gamma = \{\mathbf{U}_1(\bar{v}_1) \equiv \Phi_1, \dots, \mathbf{U}_n(\bar{v}_n) \equiv \Phi_n\}$ satisfying the set of constraints \mathcal{R} . That is, for all constraints $\Delta_l \Rightarrow \Delta_r \in \mathcal{R}$, (i) Γ contains states (s_i, h_i) , a predicate definition for each unknown predicate appearing in Δ_l and Δ_r ; (ii) by interpreting all unknown predicates according to Γ , then it is provable that Δ_l implies Δ_r (i.e., there exists $s_i \subseteq s, h_i \subseteq h$ for $i \in \{1..n\}$, and $s, h \models \Delta_l$ implies $s, h \models \Delta_r$), written as $\Gamma : \Delta_l \vdash \Delta_r$.

Lemma 1. *Given the entailment judgement $\Delta_a \vdash_{\{\}} \Delta_c \rightsquigarrow \mathcal{R}$, if there is Γ such that $\mathcal{R}(\Gamma)$, the entailment $\Gamma : \Delta_a \vdash \Delta_c$ holds.*

The soundness of the predicate synthesis requires that if definitions generated for unknown predicates are satisfiable, then the entailment is valid.

Theorem 1. *Given the entailment judgement $\Delta_a \vdash_{\emptyset} \Delta_c \rightsquigarrow \mathcal{R}$ $\Delta_a(\Gamma) \vdash \Delta_c(\Gamma)$ holds if there exists a solution Γ of \mathcal{R} .*

Theorem 1 follows from the soundness of the rules in Fig. 3 and Lemma 1.

5 Extensions

In this section, we present two ways to strengthen the inferred frame, by inferring pure properties and by normalizing inductive predicates.

Pure Constraint Inference. The inferred frame is strengthened with pure constraints following two phases. We first enrich the shape-base frame with pure properties such as size, height, sum, set of addresses/values, and their combinations. After that, we apply the same three steps in Sect. 4 to infer relational assumptions on the new pure properties. Lastly, we check satisfiability of these assumptions using FixCalc [34].

In the following, we describe how to infer size properties given a set of dependent predicates. We can similarly infer properties on height, set of addresses and values properties. We first extend an inductive predicate with a size function to capture size properties. That is, given an inductive predicate $\mathbf{P}(\bar{v}) \equiv \bigvee \Delta_i$, we generate a new predicate \mathbf{P}_n with a new size parameter n as: $\mathbf{P}_n(\bar{v}, n) \equiv \bigvee (\Delta_i \wedge n = \text{size}F(\Delta_i))$ where function $\text{size}F$ is inductively defined as follows.

$$\begin{aligned}
sizeF(r \mapsto c(\bar{t})) &= 1 & sizeF(\exists \bar{v}. \kappa \wedge \pi) &= sizeF(\kappa) \\
sizeF(\mathbf{emp}) &= 0 & sizeF(\kappa_1 * \kappa_2) &= sizeF(\kappa_1) + sizeF(\kappa_2) \\
sizeF(\mathbf{P}(\bar{t})) &= t_s & \text{where } t_s \in \bar{t} \text{ and } t_s & \text{ is a size parameter}
\end{aligned}$$

To support pure properties, we extend the proposed cyclic proof system with bi-abduction for pure constraints which was presented in [43]. In particular, we adopt the abduction rules to generate relational assumptions over the pure properties in LHS and RHS. These rules are applied exhaustively until no more unknown predicates occur.

Normalization. We aim to relate the inferred frame to existing user-provided predicates if possible as well as to explicate the heap separation (a.k.a. pointer non-aliasing) which may be implicitly constrained through predicates. Particularly, we present a lemma synthesis mechanism to explore relations between inductive predicates. Our system processes each inductive predicate in four steps. First, it generates *heap-only* conjectures (with quantifiers). Secondly, it enriches these conjectures with unknown predicates. Thirdly, it invokes the proposed entailment procedure to prove these conjectures, infer definitions for the unknown predicates and synthesize the lemmas. Last, it strengthens the inferred lemma with pure inference.

In the following, we present two types of normalization. This first type is to generate equivalence lemmas. This normalization equivalently matches a new generated predicate to an existing predicate in a given predicate library. Under the assumption that a library of predicates is provided together with advanced knowledge (i.e., lemmas in [1]) to enhance completeness. This normalization helps to reuse this knowledge for the new synthesized predicates, and potentially enhance the completeness of the proof system. Intuitively, given a set \mathbf{S} of inductive predicates and another inductive predicate \mathbf{P} (which is not in \mathbf{S}), we identify all predicates in \mathbf{S} which are equivalent to \mathbf{P} . Heap-only conjecture is generated to explore the equivalent relation between two predicates, e.g., in the case of $\mathbf{P}(x, \bar{v})$ and $\mathbf{Q}(x, \bar{w})$: $\forall \bar{v}. \mathbf{P}(\mathbf{root}, \bar{v}) \rightarrow \exists \bar{w}. \mathbf{Q}(\mathbf{root}, \bar{w})$. The shared root parameter x has been identified by examining all permutations of root parameters of the two predicates. Moreover, our system synthesizes lemmas incrementally for the combined domains of shape and pure properties. For example, with `lln` and `lsegn`, our system generates the following lemma afterwards: `lsegn(root, null, n) ↔ lln(root, n)`.

The other type of normalization is to generate separating lemmas. This normalization aims to expose hidden separation of heaps in inductive definitions. This paragraph explores parallel or consequence separate relations over inductive predicates parameters. Two parameters of a predicate are *parallel* separating if they are both root parameters e.g., r_1 and r_2 of the predicate `zip2` as follows.

$$\begin{aligned}
\mathbf{zip2}(r_1, r_2, n) &\equiv \mathbf{emp} \wedge r_1 = \mathbf{null} \wedge r_2 = \mathbf{null} \wedge n = 0 \\
&\vee r_1 \mapsto c_1(q_1) * r_2 \mapsto c_1(q_2) * \mathbf{zip2}(q_1, q_2, n-2);
\end{aligned}$$

Two arguments of a predicate are *consequence* separating if one is a root parameter and another is reachable from the root in all base formulas derived by

unfolding the predicate (e.g., those of the predicate `ll_last`). We generate these separating lemmas to explicate separation globally. As a result, the separation of actual parameters is externally visible to analyses. This visible separation enables strong updates in a modular heap analysis or frame inference in modular verification. Suppose r_1, r_2 are consequence or parallel parameters in $Q(r_1, r_2, \bar{w})$, heap conjecture is generated as:

$$Q(r_1, r_2, \bar{w}) \rightarrow Q_1(r_1) * Q_2(r_2) * Q_3(\bar{w})$$

This technique could be applied to synthesize split/join lemmas to transform predicates into the fragment of linearly compositional predicates [14, 15]. For example, our system splits the predicate `zip2` into two separating singly-lined lists through the following equivalent lemma: `zip2(root, r2, n) ↔ ll_n(root, n) * ll_n(r2, n)`.

6 Implementation and Experiments

We have implemented the proposed ideas into a procedure called `S2ENT` for entailment checking and frame inference, based on the `SLEEK` [8]. `S2ENT` relies on the SMT solver `Z3` [27] to check satisfiability of arithmetical formulas. We have also integrated `S2ENT` into the verifier `S2` [24]. We have conducted two sets of experiments to evaluate the effectiveness and efficiency of `S2ENT`. The first set of experiments are conducted on a set of inductive entailment checking problems gathered from previous publications [1, 5, 9]. We compare `S2ENT` with the state-of-the-art tools to see how many of these problems can be solved. In the second set of experiments, we apply `S2ENT` to conduct modular verification of a set of non-trivial programs. The experiments are conducted on a machine with the Intel i3-M370 (2.4 GHz) processor and 3 GB of RAM.

Entailment Proving. In Table 1, we evaluate `S2ENT` on a set of 36 *valid* entailment problems that require induction reasoning techniques. In particular, Ent 1–5 were taken from `Smallfoot` [1], Ent 6–19 from `CyclicSL` [3, 5], Ent 20–28 from [9], and Ent 29–36 were generated by us. We evaluate `S2ENT` against the existing proof systems presented for user-defined predicates. While the tools reported in [8, 12, 36] could handle a subset of these benchmarks if users provide auxiliary lemmas/axioms, [15] was designed neither for those inductive predicates in Ent 6–28 nor frame problems in Ent 29–36. The only two tools which we can compare `S2ENT` with are `CyclicSL` [3] and `songbird` [40].

The experimental results are presented in Table 1. The second column shows the entailment problems. Column *bl* captures the number of back-links in cyclic proofs generated by `S2ENT`. We observe that most of problems require only one back-link in the cyclic proofs, except that Ent 4 requires two back-links and Ent 13–15 of mutual inductive odd-even singly linked lists require three back-links. The last three columns show the results of `CyclicSL`, `songbird` and `S2ENT` respectively. Each cell shown in these columns is either CPU times (in seconds) if the tool proves successfully, or TO if the tool runs longer than 30 s, or X if the

Table 1. Inductive entailment checks

Ent	Proven	bl	[3]	[40]	Ours
1	$\text{lseg}(x,t)*\text{lseg}(t,\text{null}) \vdash \text{lseg}(x,\text{null})$	1	0.03	0.04	0.06
2	$\text{lseg}(x,t)*t \rightarrow c_1(y)*\text{lseg}(y,\text{null}) \vdash \text{lseg}(x,\text{null})$	1	0.03	0.04	0.08
3	$\text{lseg}(x,t)*\text{lseg}(t,y)*y \rightarrow c_1(\text{null}) \vdash \text{lseg}(x,\text{null})$	1	0.03	1.36	0.11
4	$\text{lseg}(x,t)*\text{lseg}(t,y)*\text{bt}(y) \wedge y \neq \text{null} \vdash \text{lseg}(x,y)*\text{bt}(y)$	2	TO	0.12	0.21
5	$\text{lseg}(x,t)*\text{lseg}(t,y)*\text{lseg}(y,z) \wedge y \neq z \vdash \text{lseg}(x,y)*\text{lseg}(y,z)$	1	3.00	0.48	0.57
6	$x \rightarrow c_1(y)*\text{rlseg}(y,z) \vdash \text{rlseg}(x,z)$	1	0.02	0.04	0.10
7	$\text{nlseg}(x,z)*z \rightarrow c_1(y) \vdash \text{nlseg}(x,y)$	1	0.02	0.04	0.04
8	$\text{nlseg}(x,z)*\text{nlseg}(z,y) \vdash \text{nlseg}(x,y)$	1	0.03	0.04	0.06
9	$\text{glseg}(x,z)*z \rightarrow c_1(y) \vdash \text{glseg}(x,y)$	1	0.02	0.04	0.04
10	$\text{glseg}(x,z)*\text{glseg}(z,y) \vdash \text{glseg}(x,y)$	1	0.02	0.04	0.04
11	$\text{dlseg}(u,v,x,y) \vdash \text{glseg}_2(u,v)$	1	0.07	0.04	0.04
12	$\text{dlseg}(w,v,x,z)*\text{dlseg}(u,w,z,y) \vdash \text{dlseg}(u,v,x,y)$	1	0.04	0.03	0.11
13	$\text{listo}(x,z)*\text{listo}(z,\text{null}) \vdash \text{liste}(x,\text{null})$	3	0.06	0.04	0.06
14	$\text{liste}(x,z)*\text{liste}(z,\text{null}) \vdash \text{liste}(x,\text{null})$	3	0.18	0.04	0.12
15	$\text{listo}(x,z)*\text{liste}(z,y) \vdash \text{listo}(x,y)$	3	4.33	0.03	0.88
16	$\text{binPath}(x,z)*\text{binPath}(z,y) \vdash \text{binPath}(x,y)$	1	0.03	0.05	0.06
17	$\text{binPath}(x,y) \vdash \text{binTreeSeg}(x,y)$	1	0.12	X	0.08
18	$\text{binTreeSeg}(x,z)*\text{binTreeSeg}(z,y) \vdash \text{binTreeSeg}(x,y)$	1	0.20	0.07	0.66
19	$\text{binTreeSeg}(x,y)*\text{binTree}(y) \vdash \text{binTree}(x)$	1	0.06	0.05	0.03
20	$\text{tmp}(x,\text{size}) \vdash \exists y. \text{ls}(x,y,\text{size})$	1	NA	0.73	0.39
21	$\text{sort11}(x,\text{min}) \vdash \text{ll}(x)$	1	NA	0.1	0.05
22	$\text{sort11n}(x,\text{min},\text{size}) \vdash \text{lln}(x,\text{size})$	1	NA	0.55	0.12
23	$\text{sort11n}(x,\text{min},\text{size}) \vdash \text{sort11}(x,\text{min})$	1	NA	TO	0.08
24	$\text{lsegn}(x,y,sz_1)*\text{lsegn}(y,z,sz_2) \vdash \text{lsegn}(x,z,sz_1+sz_2)$	1	NA	10.49	0.12
25	$\text{lsegn}(x,y,\text{size}_1)*\text{lsn}(y,\text{size}_2) \vdash \text{lsn}(x,\text{size}_1+\text{size}_2)$	1	NA	10.59	0.10
26	$\text{lsegn}(x,\text{ll},n_1)*t_1 \rightarrow c_1(y)*\text{lseg1}(y,\text{ty},n_2) \vdash \text{lseg1}(x,\text{ty},n_1+n_2+1)$	1	NA	16.76	0.10
27	$\text{avl}(x,\text{size},\text{height},\text{bal}) \vdash \text{btn}(x,\text{size})$	1	NA	X	0.08
28	$\text{t11}(x,\text{ll},\text{lr},\text{size}) \vdash \text{btn}(x,\text{size})$	1	NA	0.13	0.07
29	$\text{ll.last}(x,y) \vdash y \rightarrow c_1(\text{null})$	1	NA	NA	0.18
30	$\text{ll.last.size}(x,y,\text{size}) \vdash y \rightarrow c_1(\text{null})$	1	NA	NA	0.23
31	$\text{zip2}(x,y) \vdash \text{ll}(x)$	1	NA	NA	0.23
32	$\text{zip2n}(x,y,\text{size}) \vdash \text{lln}(x,\text{size})$	1	NA	NA	0.46
33	$\text{zip3}(x,y) \vdash \text{ll}(y)$	1	NA	NA	0.39
34	$\text{sort11}(x,\text{min})*\text{sort11}(y,\text{min}_1) \vdash \text{ll}(x)$	1	NA	NA	0.13
35	$\text{sort11n}(x,\text{min1},\text{size1})*\text{slln}(y,\text{min2},\text{size2}) \vdash \text{sll}(x,\text{min1})$	1	NA	NA	0.27
36	$\text{t11}(x,\text{ll},\text{lr},\text{size})*\text{ll}(y) \vdash \text{btn}(x,\text{size})$	1	NA	NA	0.56

tool returns a false positive, or NA if the entailment is beyond the capability of the tool. *In summary, out of the 36 problems, Cyclic_{SL} solves 18 (with one TO - Ent 4); songbird solves 25 (with two false positive - Ent 17 and 27 and one TO - Ent 23); and S2ENT solves all 36 problems.*

In Table 1, each entailment check in Ent 1–19 has emp as frame axioms (their LHS and RHS have the same heaps). Hence, they may be handled by existing inductive proof systems like [3, 9, 15, 40]. In particular, Ent 1–19 include shape-only predicates. The results show that Cyclic_{SL} and songbird ran a bit faster than S2ENT in most of the their successful cases. It is expected as S2ENT requires additional steps for frame inference. Each entailment check in Ent 20–28 includes inductive predicates with pure properties (e.g., size and sortedness). While Cyclic_{SL} can provide inductive reasoning for arithmetic and heap domains separately [5], there is no system proposed for cyclic proofs in the combined domain. Hence, these problems are beyond the capability of Cyclic_{SL} . Ent 20 which requires mutual induction reasoning is the motivating example of songbird (agumented with size property) [40]. In particular, sort11 represents a sorted list with smallest value min , and t11 is a binary tree whose nodes

point to their parents and leaves are linked by a linked list [19, 24]. **S2ENT** solves each entailment incrementally: shape-based frame and then pure properties. The results show that **S2ENT** was more effective and efficient than **songbird**.

Each entailment check in Ent 29–36 requires both inductive reasoning and frame inference. These checks are beyond the capability of all existing entailment procedures for SL. **S2ENT** generates frame axioms for inductive reasoning. The experiments show that the proposed proof system can support efficient and effective reasoning on both shape and numeric domains as well as inductive proofs and frame inference.

Modular Verification for Memory Safety. We enhance the existing program verifier **S2** [24] with **S2ENT** to automatically verify a range of heap-manipulating programs. We evaluate the enhanced **S2** on the C library Glib open source [16] which includes non-GUI code from the GTK+ toolkit and the GNOME desktop environment. We conduct experiments on heap-manipulating files, i.e., singly-linked lists (`gslis.c`), doubly-linked lists (`glist.c`), balanced binary trees (`gtree.c`) and N-ary trees (`gnode.c`). These files contain fairly complex algorithms (e.g., sortedness) and the data structures used in `gtree.c` and `gnode.c` are very complex.

Table 2. Experiments on Glib library

	LOC	#Pr	wo.		w.		
			#√	Sec.	#syn	#√	Sec.
<code>gslis.c</code>	698	52	41	8.93	126	47	12.47
<code>glist.c</code>	784	51	39	19.41	132	46	30.01
<code>gtree.c</code>	1204	40	36	57.31	96	36	60.88
<code>gnode.c</code>	1128	65	52	37.78	174	53	53.40

Some procedures of `gslis.c` and `glist.c` were evaluated by tools presented in [9, 31, 36] where the user had to manually provide a large number of lemmas to support the tool. Furthermore, the verification in [9] is semi-automatic, i.e., verification conditions were manually generated. Besides the tool in [9], tools in [31, 36] were no longer available for comparison.

In Table 2 we show, for each file the number of lines of code (excluding comments) LOC and the number of procedures #Pr. We remark that these procedures include tail-recursive procedures which are translated from loops. The columns (#√) (and sec.) show the number of procedures (and time in seconds) for which **S2** can verify memory safety without (wo.) and with (w.) **S2ENT**. Column #syn shows the number of synthesized lemmas that used the technique in Sect. 5. With the lemma synthesis, the number of procedures that can be successfully verified increases from 168 (81%) to 182 (88%) with a time overhead of 28% (157s/123s).

A closer look shows that with **S2ENT** we are able to verify a number of challenging methods in `gslis.c` and `glist.c`. By generating separating lemmas, **S2ENT** successfully infers shape specifications of methods manipulating the last element of lists (i.e., `g_slist_concat` in `gslis.c` and `g_list_append` in `glist.c`). By generating equivalence lemmas, matching a newly-inferred inductive predicate with predefined predicates in **S2** is now extended beyond the shape-only domain. Moreover, the experimental results also show that the enhanced **S2** were able to

verify 41/52 procedures in `gslis.c` and 39/51 procedures in `glis.c`. In comparison, while the tool in [9] could semi-automatically verify 11 procedures in `gslis.c` and 6 procedures in `glis.c`, with user-supplied lemmas the tool in [31] could verify 22 procedures in `gslis.c` and 10 procedures in `glis.c`.

7 Related Work and Conclusion

This work is related to three groups of work. The first group are those on entailment procedures in SL. Initial proof systems in SL mainly focus on a decidable fragment combining linked lists (and trees) [1, 7, 11, 13, 14, 17, 22, 29, 32, 33]. Recently, Iosif *et al.* extend the decidable fragment to restricted inductive predicates [19]. Timos *et al.* [42] present a comprehensive summary on computational complexity for entailments in SL with inductive predicates. Smallfoot [1] and GRASShopper [33] provide systematic approaches for frame inference but with limited support for (general) inductive predicates. Extending these approaches to support general inductive predicates is non-trivial. GRASShopper is limited to a GRASS-reducible class of inductive predicates. While Smallfoot system has been designed to allow the use of general inductive predicates, the inference rules in Smallfoot are hardwired for list predicates only and a set of new rules must be developed for a proof system targeting general inductive predicates. SLEEK [8] and jStar [12] support frame inference with a soundness guarantee for general inductive predicates. However, they provide limited support for induction using user-supplied lemmas [12, 30]. Our work, like [8, 36], targets an undecidable SL fragment including (arbitrary) inductive predicates and numerical constraints; we trade completeness for expressiveness. In addition to what are supported in [8, 36], we support frame inference with inductive reasoning in SL by providing a system of cyclic proofs.

The second group is work on inductive reasoning. Lemmas are used to enhance the inductive reasoning of heap-based programs [5, 12, 30]. They are used as alternative unfoldings beyond predicates' definitions [5, 30], external inference rules [12], or intelligent generalization to support inductive reasoning [3]. Unfortunately, the mechanisms in these systems require users to supply those additional lemmas that might be needed during a proof. SPEN [15] synthesizes lemmas to enhance inductive reasoning for some inductive predicates with bags of values. However, it is designed to support some specific classes of inductive predicates and it is difficult to extend it to cater for general inductive predicates. For a solution to inductive reasoning in SL, Smallfoot [1, 3, 5] presents subtraction rules that are consequent from a set of lemmas of lists and trees. Brotherston *et al.* propose cyclic proof system for the entailment problem [2, 3]. Similarly, the circularity rule has been introduced in matching logic [38], Constraint Logic Programming [9] and separation logic combined with predicate definitions and arithmetic [40]. Furthermore, work in [39] supports frame inference based on an ad-hoc mechanism, using a simple unfolding and matching. Like [3, 9, 40], our system also uses historical sequents at case split steps as induction hypotheses. Beyond these systems [3, 9, 15, 40], S2ENT infers frames for inductive

proofs systematically; and thus it gives a better support for modular verification of heap-manipulating programs. Moreover, we show how we can incrementally support inductive reasoning for the combination of heap and pure domains. In contrast, there are no formalized discussions in [5, 9, 40] about inductive reasoning for the combined domains; while [5] supports these domains separately, [9, 40] only demonstrates their support through experimental results.

The third group is on lemma synthesis. In inductive reasoning, auxiliary lemmas are generated to discover theorems (e.g. [10, 23, 28]). The key elements of these techniques are heuristics used to generate equivalent lemmas for sets of given functions, constants and datatypes. In our work, we introduce lemma synthesis to strengthen the inductive constraints. To support theorem discovery, we synthesize equivalent and separating lemmas. This mechanism can be extended to other heuristics to enhance the completeness of modular verification.

Conclusion. We have presented a novel approach to frame inference for inductive entailments in SL with inductive predicates and arithmetic. The core of our proposal is the system of lemma synthesis through cyclic proofs in which backlinks are formed using the cut rule. Moreover, we have presented two extensions to strengthen the inferred frames. Our evaluation indicates that our system is able to infer frame axioms for inductive entailment checking that are beyond the capability of the existing systems.

Acknowledgements. This research is partially supported by project T2MOE1704 from Ministry of Education, Singapore.

References

1. Berdine, J., Calcagno, C., O’Hearn, P.W.: Symbolic execution with separation logic. In: Yi, K. (ed.) APLAS 2005. LNCS, vol. 3780, pp. 52–68. Springer, Heidelberg (2005). https://doi.org/10.1007/11575467_5
2. Brotherston, J.: Cyclic proofs for first-order logic with inductive definitions. In: Beckert, B. (ed.) TABLEAUX 2005. LNCS (LNAI), vol. 3702, pp. 78–92. Springer, Heidelberg (2005). https://doi.org/10.1007/11554554_8
3. Brotherston, J., Distefano, D., Petersen, R.L.: Automated cyclic entailment proofs in separation logic. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE 2011. LNCS (LNAI), vol. 6803, pp. 131–146. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22438-6_12
4. Brotherston, J., Fuhs, C., Gorogiannis, N., Pérez, J.N.: A decision procedure for satisfiability in separation logic with inductive predicates. In: Proceedings of CSL-LICS (2014)
5. 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). https://doi.org/10.1007/978-3-642-35182-2_25
6. Calcagno, C., Distefano, D., O’Hearn, P.W., Yang, H.: Compositional shape analysis by means of bi-abduction. In: Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 289–300. ACM (2009)

7. Chen, T., Song, F., Wu, Z.: Tractability of separation logic with inductive definitions: beyond lists. In: Meyer, R., Nestmann, U. (eds.) 28th International Conference on Concurrency Theory (CONCUR 2017), vol. 85, pp. 37:1–37:17 (2017)
8. 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. *SCP* **77**(9), 1006–1036 (2012)
9. Chu, D.-H., Jaffar, J., Trinh, M.-T.: Automatic induction proofs of data-structures in imperative programs. In: *PLDI 2015* (2015)
10. Claessen, K., Johansson, M., Rosén, D., Smallbone, N.: Automating inductive proofs using theory exploration. In: Bonacina, M.P. (ed.) *CADE 2013*. LNCS (LNAI), vol. 7898, pp. 392–406. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38574-2_27
11. Cook, B., Haase, C., Ouaknine, J., Parkinson, M., Worrell, J.: Tractable reasoning in a fragment of separation logic. In: Katoen, J.-P., König, B. (eds.) *CONCUR 2011*. LNCS, vol. 6901, pp. 235–249. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23217-6_16
12. Distefano, D., Parkinson, M.: jStar: towards practical verification for java. In: *OOPSLA 2008*, pp. 213–226. ACM, New York (2008)
13. 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). https://doi.org/10.1007/978-3-319-12736-1_17
14. Enea, C., Lengál, O., Sighireanu, M., Vojnar, T.: Compositional entailment checking for a fragment of separation logic. *Formal Meth. Syst. Des.* **51**(3), 575–607 (2017)
15. 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). https://doi.org/10.1007/978-3-319-24953-7_7
16. Glib. version 2.38.2 (2013). <https://developer.gnome.org/glib/>. Accessed 13 Oct 2017
17. Gu, X., Chen, T., Wu, Z.: A complete decision procedure for linearly compositional separation logic with data constraints. In: Olivetti, N., Tiwari, A. (eds.) *IJCAR 2016*. LNCS (LNAI), vol. 9706, pp. 532–549. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40229-1_36
18. Holik, L., Lengál, O., Rogalewicz, A., Šimáček, J., Vojnar, T.: Fully automated shape analysis based on forest automata. In: Sharygina, N., Veith, H. (eds.) *CAV 2013*. LNCS, vol. 8044, pp. 740–755. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_52
19. 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). https://doi.org/10.1007/978-3-642-38574-2_2
20. Ishtiaq, S., O’Hearn, P.W.: BI as an assertion language for mutable data structures. In: *ACM POPL*, London, pp. 14–26, January 2001
21. Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F.: VeriFast: a powerful, sound, predictable, fast verifier for C and Java. In: Bobaru, M., Havelund, K., Holzmann, G.J., Joshi, R. (eds.) *NFM 2011*. LNCS, vol. 6617, pp. 41–55. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20398-5_4

22. Jansen, C., Katelaan, J., Matheja, C., Noll, T., Zuleger, F.: Unified reasoning about robustness properties of symbolic-heap separation logic. In: Yang, H. (ed.) ESOP 2017. LNCS, vol. 10201, pp. 611–638. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54434-1_23
23. Johansson, M., Dixon, L., Bundy, A.: Conjecture synthesis for inductive theories. *J. Autom. Reason.* **47**(3), 251–289 (2011)
24. Le, Q.L., Gherghina, C., Qin, S., Chin, W.-N.: Shape analysis via second-order bi-abduction. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 52–68. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08867-9_4
25. Le, Q.L., Sun, J., Chin, W.-N.: Satisfiability modulo heap-based programs. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9779, pp. 382–404. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41528-4_21
26. Le, Q.L., Tatsuta, M., Sun, J., Chin, W.-N.: A decidable fragment in separation logic with inductive predicates and arithmetic. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10427, pp. 495–517. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63390-9_26
27. de Moura, L., Björner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_24
28. McCasland, R., Bundy, A., Serge, A.: Automated discovery of inductive theorems. *Stud. Logic Grammar Rhetor.* **10**(23), 135–149 (2007)
29. Navarro Pérez, J.A., Rybalchenko, A.: Separation logic modulo theories. In: Shan, C. (ed.) APLAS 2013. LNCS, vol. 8301, pp. 90–106. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-03542-0_7
30. Nguyen, H.H., Chin, W.-N.: Enhancing program verification with lemmas. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 355–369. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-70545-1_34
31. Pek, E., Qiu, X., Madhusudan, P.: Natural proofs for data structure manipulation in C using separation logic. In: Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation, pp. 440–451. ACM, New York (2014)
32. Pérez, J.A.N., Rybalchenko, A.: Separation logic + superposition calculus = heap theorem prover. In: PLDI, pp. 556–566 (2011)
33. 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). https://doi.org/10.1007/978-3-642-39799-8_54
34. Popeea, C., Chin, W.-N.: Inferring disjunctive postconditions. In: Okada, M., Satoh, I. (eds.) ASIAN 2006. LNCS, vol. 4435, pp. 331–345. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-77505-8_26
35. Qin, S., He, G., Chin, W.-N., Craciun, F., He, M., Ming, Z.: Automated specification inference in a combined domain via user-defined predicates. *Sci. Comput. Program.* **148**(C), 189–212 (2017)
36. Qiu, X., Garg, P., Ștefănescu, A., Madhusudan, P.: Natural proofs for structure, data, and separation. In: PLDI, pp. 231–242. ACM, New York (2013)
37. Reynolds, J.C.: Separation logic: a logic for shared mutable data structures. In: Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science, LICS 2002, pp. 55–74. IEEE Computer Society, Washington, DC (2002)
38. Rosu, G., Stefanescu, A.: Checking reachability using matching logic. In: OOPSLA 2012, pp. 555–574. ACM, New York (2012)
39. Rowe, R.N.S., Brotherston, J.: Automatic cyclic termination proofs for recursive procedures in separation logic. In: CPP (2017)

40. Ta, Q.-T., Le, T.C., Khoo, S.-C., Chin, W.-N.: Automated mutual explicit induction proof in separation logic. In: Fitzgerald, J., Heitmeyer, C., Gnesi, S., Philippou, A. (eds.) FM 2016. LNCS, vol. 9995, pp. 659–676. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48989-6_40
41. Tatsuta, M., Le, Q.L., Chin, W.-N.: Decision procedure for separation logic with inductive definitions and Presburger arithmetic. In: Igarashi, A. (ed.) APLAS 2016. LNCS, vol. 10017, pp. 423–443. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47958-3_22
42. Antonopoulos, T., Gorogiannis, N., Haase, C., Kanovich, M., Ouaknine, J.: Foundations for decision problems in separation logic with general inductive predicates. In: Muscholl, A. (ed.) FoSSaCS 2014. LNCS, vol. 8412, pp. 411–425. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54830-7_27
43. Trinh, M.-T., Le, Q.L., David, C., Chin, W.-N.: Bi-abduction with pure properties for specification inference. In: Shan, C. (ed.) APLAS 2013. LNCS, vol. 8301, pp. 107–123. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-03542-0_8
44. Yang, H., Lee, O., Berdine, J., Calcagno, C., Cook, B., Distefano, D., O’Hearn, P.: Scalable shape analysis for systems code. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 385–398. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-70545-1_36

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

