



Cyclic Program Synthesis

Shachar Itzhaky

Technion

Israel

shachari@cs.technion.ac.il

Hila Peleg

University of California, San Diego

USA

hpeleg@eng.ucsd.edu

Nadia Polikarpova

University of California, San Diego

USA

nadia.polikarpova@ucsd.edu

Reuben N. S. Rowe

Royal Holloway, University of London

United Kingdom

reuben.rowe@rhul.ac.uk

Ilya Sergey

Yale-NUS College

National University of Singapore

Singapore

ilya.sergey@yale-nus.edu.sg

Abstract

We describe the first approach to automatically synthesizing heap-manipulating programs with auxiliary recursive procedures. Such procedures occur routinely in data structure transformations (e.g., flattening a tree into a list) or traversals of composite structures (e.g., n -ary trees). Our approach, dubbed *cyclic program synthesis*, enhances deductive program synthesis with a novel application of *cyclic proofs*. Specifically, we observe that the machinery used to form cycles in cyclic proofs can be reused to systematically and efficiently abduce recursive auxiliary procedures.

We develop the theory of cyclic program synthesis by extending Synthetic Separation Logic (SSL), a logical framework for deductive synthesis of heap-manipulating programs from Separation Logic specifications. We implement our approach as a tool called CYPRESS, and showcase it by automatically synthesizing a number of programs manipulating linked data structures using recursive auxiliary procedures and mutual recursion, many of which were beyond the reach of existing program synthesis tools.

CCS Concepts: • Software and its engineering → Automatic programming.

Keywords: Program Synthesis, Separation Logic, Cyclic Proofs

ACM Reference Format:

Shachar Itzhaky, Hila Peleg, Nadia Polikarpova, Reuben N. S. Rowe, and Ilya Sergey. 2021. Cyclic Program Synthesis. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '21)*, June 20–25, 2021, Virtual, Canada. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3453483.3454087>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PLDI '21, June 20–25, 2021, Virtual, Canada

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8391-2/21/06.

<https://doi.org/10.1145/3453483.3454087>

1 Introduction

Consider the task of flattening a binary tree into a linked list, which is typically solved by writing a recursive data traversal program. The promise of *program synthesis* is to automate such tedious data manipulation tasks by generating programs automatically from high-level, declarative specifications. Several recent synthesizers [1, 19, 27–29] are indeed capable of generating *recursive* programs given only the top-level description of their behavior. For example, SuSLIK [29] can generate a provably correct, recursive C-like program that *deallocates* a binary tree given as input the following specification in Separation Logic (SL) [26, 34]:

$$\{\text{tree}(x, s)\} \text{treefree}(x) \{\text{emp}\} \quad (1)$$


This specification says that initially the heap contains a binary tree rooted at address x with set of elements s , and that after executing `treefree` the heap must be empty. Inspired by the success with `treefree`, the programmer might try to synthesize tree flattening from the specification below:

$$\{r \mapsto x * \text{tree}(x, s)\} \text{flatten}(r) \{r \mapsto y * \text{sll}(y, s)\} \quad (2)$$

where `sll` describes a singly-linked list, and the location r initially stores the root of the input tree, and eventually the output list as computed by `flatten`. Much to the programmer's frustration, however, SuSLIK fails to synthesize an implementation: it times out, without producing any useful output. In fact, given only the specification (2) or a similar one, and no additional hints, this example is out of reach for all other state-of-the-art synthesizers for recursive programs.

Challenge: recursive auxiliaries. This failure can be more readily understood when considering the expected solution. One such solution, depicted below, begins with two recursive calls which flatten the immediate subtrees of the input tree, obtaining two linked lists. Now, the synthesizer is at an impasse: how to combine the two lists to create the output list? This step requires an operation that *appends* lists, which is in itself a recursive program (and one which cannot be implemented by another call to `flatten`, for example). This synthesis task illustrates a fundamental limitation of existing approaches to synthesis of recursion: tree flattening

cannot be accomplished without a *recursive auxiliary function*.¹ Existing synthesizers can generate flatten given the specification of append as a *hint* from the user [29]: using this auxiliary specification as a stepping stone, they can then generate code for both flatten and append as two synthesis tasks. It is, however, often non-trivial for the programmer to come up with this kind of hints.



Inferring auxiliary specifications *automatically* has been a long-standing open challenge in program synthesis. Since the hypothesis space is very large, simply allowing the synthesizer to conjecture any arbitrary auxiliary definition would be impractical, as that would make the search space explode. The state of the art is the work by Eguchi *et al.* [13], which attacks this problem in the context of functional programming by assuming a set of *predefined recursion templates* (e.g. a fold-right over lists). While syntactic templates help curb the search space explosion, they also limit the applicability of the technique, and in particular, preclude traversals of arbitrary user-defined data structures.

Cyclic program synthesis. We present a new synthesis technique, dubbed *cyclic program synthesis*, capable of automatically discovering recursive auxiliaries without the need for built-in templates or additional hints from the user. In particular, given the specification (2), our technique synthesizes a provably correct tree-flattening program, automatically discovering a recursive auxiliary that appends two lists.

Our technique draws inspiration from—and owes its name to—*cyclic proofs*, a powerful reasoning mechanism from the line of work on automated theorem proving [3, 9, 11, 25, 36, 37, 39]. In cyclic proofs, derivation “trees” can have *backlinks* from non-axiomatic leaves (called *buds*) to identical internal nodes (called *companions*). These backlinks enable an automation-friendly approach to proofs by induction: instead of conjecturing an induction hypothesis *a-priori*, the method exploits a similarity between a current goal and a companion goal elsewhere in the derivation, essentially turning the latter into an induction hypothesis *on demand*.

The key insight of this work is to apply the cyclic proof approach to automated deductive synthesis of recursive programs: instead of conjecturing auxiliary specifications *a priori*, one can *abduce* them from the repeated goal patterns encountered during the main program derivation. For example, in `flatten` the synthesizer starts generating the code of `append inline`, and later recognizes that the program can be completed by extracting the code into a function and adding a recursive call.

¹The choice of auxiliary is not unique, e.g. one may propose an auxiliary that uses a list accumulator.

Our contributions. Realizing the idea of cyclic synthesis requires a carefully designed deductive system, in which the repeated goals observed in program derivations can be used to form backlinks. Our first contribution is the design of such a deductive system, which we dub Cyclic SSL (SSL_⊆). This system builds on top of Synthetic Separation Logic (SSL), the deductive synthesis framework underlying SuSLiK [29]. SSL_⊆ features a new set of rules for synthesizing procedure calls, incorporating an expressive trace-based termination checking mechanism from cyclic proofs [7, 8]. This mechanism enables SSL_⊆ to derive a wider range of recursive programs, including programs with auxiliaries and *non-structural* recursion.

Our second contribution is an implementation of SSL_⊆ in a new synthesis tool CYPRESS, which extends and subsumes SuSLiK. The addition of cyclic reasoning requires an efficient mechanism for detecting potential companions for the current goal during search. CYPRESS achieves this via a *call abduction oracle*, which matches up goals that are not syntactically equal, adjusting them accordingly. CYPRESS also features more efficient theory reasoning via a mechanism we dub *unification modulo theories*, as well as new *best-first* proof search, guided by the size and shape of the goal.

We evaluated CYPRESS on 46 synthesis benchmarks. Our evaluation shows that CYPRESS is able to solve a number of challenging tasks requiring nested traversals of linked structures (e.g., sorting or de-duplicating a linked list) or traversals of mutually-recursive data structures (e.g., *n*-ary trees). To the best of our knowledge, these programs are beyond reach of any existing approaches to automated hint-free synthesis from declarative specifications.

Paper outline. The following sections provide a brief primer on synthesis via SSL, outlining this work’s innovations via a series of examples (Sec. 2); give a description of SSL_⊆ and its meta-theory (Sec. 3); describe the synthesis algorithm (Sec. 4); and report on our evaluation (Sec. 5).

2 Cyclic Program Synthesis, by Example

This section presents the main ideas of Synthetic Separation Logic (SSL), followed by an overview of the scenarios enabled by our new rules and cyclic proofs principles. All code examples shown in this section are synthesized automatically using our novel synthesis tool CYPRESS.

2.1 Background: SSL and Its Limitations

Specifications. Deductive synthesis based on SSL takes as input a pair of Hoare-style pre- and postconditions. For instance, recall the specification (1) for deallocating a tree:

$$\{\text{tree}(x, s)\} \text{treefree}(x) \{\text{emp}\}$$

Here the precondition `tree(x, s)` states that `treefree` may assume that it starts from a heap containing a binary tree rooted at address `x` with payload set `s`; the postcondition

`emp` states that `treefree` must guarantee that the heap is empty upon its termination.² Note that the tree root `x` also appears as a parameter to `treefree`, and hence is a *program variable*, i.e., can be mentioned in the synthesized program; the payload set `s`, on the other hand, is a *logical variable* and must not appear in the program. In the rest of this section, we distinguish program variables from logical variables by using monotype font for the former.

In general, in a specification $\{\mathcal{P}\} f(\dots) \{Q\}$, assertions \mathcal{P}, Q have the form $\phi; P$, where the *spatial* part P describes the shape of the heap, while the *pure* part ϕ is a plain first-order formula that states the relations between variables (in (1) the trivial pure part `true` is omitted from both pre- and postcondition). For the spatial part, SSL employs the standard *symbolic heap* fragment of Separation Logic [26, 34]. Informally, a symbolic heap is a set of atomic formulas called *heaplets* joined with *separating conjunction* ($*$). The simplest kind of heaplet is a *points-to* assertion $x \mapsto e$, which describes a single memory location with address x and payload e . For example, the formula $x \mapsto 5 * y \mapsto 10$ describes a heap with two memory locations, x and y , which store values 5 and 10, and are *distinct*, as per the semantics of the $*$ connective.

To capture linked data structures, such as lists and trees, SSL specifications make extensive use of *inductive heap predicates*, which are standard in Separation Logic. For instance, the tree predicate from (1) is inductively defined as follows:

$$\begin{aligned} \text{tree}(x, s) \triangleq & x = 0 \Rightarrow \{s = \emptyset; \text{emp}\} \\ & | x \neq 0 \Rightarrow \{s = \{v\} \cup s_l \cup s_r; \\ & [x, 3] * x \mapsto v * \langle x, 1 \rangle \mapsto l * \langle x, 2 \rangle \mapsto r * \\ & \text{tree}(l, s_l) * \text{tree}(r, s_r)\} \end{aligned} \quad (3)$$

This definition consists of two guarded *clauses*: the first one describes the empty tree (and applies when the root pointer x is null), and the second one describes a non-empty tree. In the second clause, a tree node is represented by a three-element record starting at address x . The first field of the record stores the payload v , while the other two store the addresses l and r of the left and right subtrees, correspondingly. Records are represented using a generalized form of the points-to assertion with an *offset*: for example, the heaplet $\langle x, 1 \rangle \mapsto l$ describes a memory location at the address $x + 1$. The *block* assertion $[x, 3]$ is an artifact of C-style memory management: it represents a memory block of size three at address x that has been dynamically allocated by `malloc` (and hence can be de-allocated by `free`). The two disjoint heaps $\text{tree}(l, s_l)$ and $\text{tree}(r, s_r)$ store the two subtrees. Finally, the pure part of the second clause indicates that the payload of the whole tree consist of v and the subtree payloads, s_l and s_r .

Deductive synthesis. Given a pre-/postcondition pair \mathcal{P}, Q , deductive synthesis proceeds by constructing a derivation

$$\begin{array}{c} \text{EMP} \\ \frac{\vdash \phi \Rightarrow \psi}{\{\phi; \text{emp}\} \rightsquigarrow \{\psi; \text{emp}\} | \text{skip}} \quad \text{FRAME} \\ \frac{\{\phi; P\} \rightsquigarrow \{\psi; Q\} | c}{\{\phi; P * R\} \rightsquigarrow \{\psi; Q * R\} | c} \\ \text{FREE} \\ \frac{R = [x, n] * *_{0 \leq i \leq n} (\langle x, i \rangle \mapsto e_i) \quad \{\phi; P\} \rightsquigarrow \{Q\} | c}{\{\phi; P * R\} \rightsquigarrow \{Q\} | \text{free}(x); c} \\ \text{READ} \\ \frac{y \text{ is fresh} \quad [y/a]\{\phi; \langle x, i \rangle \mapsto a * P\} \rightsquigarrow [y/a]\{Q\} | c}{\{\phi; \langle x, i \rangle \mapsto a * P\} \rightsquigarrow \{Q\} | \text{let } y = *(x + i); c} \\ \text{WRITE} \\ \frac{\text{Vars}(e) \subseteq \text{PV} \quad \{\phi; \langle x, i \rangle \mapsto e * P\} \rightsquigarrow \{\psi; \langle x, i \rangle \mapsto e * Q\} | c}{\{\phi; \langle x, i \rangle \mapsto e' * P\} \rightsquigarrow \{\psi; \langle x, i \rangle \mapsto e * Q\} | *(x + i) = e; c} \end{array}$$

Figure 1. Selected SSL rules (simplified).

of the SSL judgment $\{\mathcal{P}\} \rightsquigarrow \{Q\} | c$ for some program c . Intuitively, this judgment has the same meaning as the Hoare triple $\{\mathcal{P}\} c \{Q\}$ (different syntax is used to emphasise that operationally the program c is “the output” rather than “the input”). The derivation is constructed by applying inference rules, a subset of which is presented in Fig. 1.

Inference rules gradually simplify the initial synthesis goal $\{\mathcal{P}\} \rightsquigarrow \{Q\}$, until symbolic heaps in both pre- and postconditions are empty, at which point the *terminal* rule EMP concludes the derivation and “emits” a trivial program `skip`. The FRAME rule reduces the synthesis goal to a smaller one by removing matching symbolic heaps in its pre- and postcondition. The FREE rule eliminates a dynamically-allocated block of memory from the precondition by emitting a `free` statement. The rules READ and WRITE synthesize reads and writes of heap locations, correspondingly. The role of the WRITE rule is to “equalise” points-to heaplets that have the same address but different payloads, so that they can be subsequently “trimmed” by FRAME. The role of the READ rule is to turn a logical variable a into a program variable y , which might enable subsequent application of FREE or WRITE. Note that reading from the heap always creates a fresh program variable (hence the `let` syntax), and variables, unlike heap locations, are never re-assigned.

To deal with inductive predicates, SSL features rules OPEN and CLOSE (elided from this overview), which unfold predicate definitions in the pre- and the postcondition, respectively. Finally, the CALL rule synthesizes a recursive call if some part of the current goal’s precondition matches the precondition \mathcal{P} of the top-level, user-provided specification.

Deriving `treefree`. Let us illustrate how all those rules work in tandem to synthesize the implementation of `treefree` shown on the right from the specification (1). We start by unfolding the definition of `tree` in the top-level goal $\{\text{tree}(x, s)\} \rightsquigarrow \{\text{emp}\}$, which generates two sub-goals (one for each clause of the predicate):

```
1 void treefree(x) {
2   if (x == 0) {
3   } else {
4     let l = *(x + 1);
5     let r = *(x + 2);
6     free(x);
7     treefree(l);
8     treefree(r);
9   }}
```

²This specification also implicitly guarantees that `treefree` always terminates and executes without memory errors (e.g., null-pointer dereferencing).

$$\{x = 0 \wedge s = \emptyset; \text{emp}\} \rightsquigarrow \{\text{emp}\} \mid c_1 \quad (4)$$

$$\{x \neq 0; [x, 3] * x \mapsto v * \langle x, 1 \rangle \mapsto l * \langle x, 2 \rangle \mapsto r * \text{tree}(l, s_l) * \text{tree}(r, s_r)\} \rightsquigarrow \{\text{emp}\} \mid c_2 \quad (5)$$

The programs c_1 and c_2 emitted by these subgoals will be conjoined via the statement **if** $(x = 0) \{c_1\}$ **else** $\{c_2\}$. The first subgoal (4) is trivially solved by the rule EMP, resulting in a program skip. In the second subgoal (5), the two grayed fragments enable two subsequent applications of the rule READ, adding two reads statements, from $*(x + 1)$ and $*(x + 2)$, correspondingly, creating two new program-level bindings, l and r and transforming the current goal into

$$\{x \neq 0; [x, 3] * x \mapsto v * \langle x, 1 \rangle \mapsto l * \langle x, 2 \rangle \mapsto r * \text{tree}(l, s_l) * \text{tree}(r, s_r)\} \rightsquigarrow \{\text{emp}\} \quad (6)$$

The FREE rule now applies to the grayed fragment of the goal's precondition, emitting $\text{free}(x)$, and simplifying the synthesis goal to:

$$\{\dots; \text{tree}(l, s_l) * \text{tree}(r, s_r)\} \rightsquigarrow \{\text{emp}\} \quad (7)$$

To complete the synthesis, we notice that each of the two heaplets in the precondition matches the top-level specification (1), and hence can trigger the rule CALL, synthesizing a procedure call with a suitable argument.

Limitations. The CALL rule of SSL imposes significant limitations on the kinds of recursive functions it can derive. The first limitation is that it only allows using the *top-level* synthesis goal provided by the user as the specification for the callee; this precludes synthesis of recursive auxiliary functions, as required, for example, to flatten a tree into a list, as explained in the introduction.

A second, somewhat subtler limitation arises from the way SSL enforces termination of synthesized programs. To avoid generating trivial non-terminating solutions—such as $\text{treefree}(x)$ immediately calling $\text{treefree}(x)$ again—SSL restricts synthesized programs to be *structurally recursive*.³ More precisely, synthesis starts by picking a *single inductive predicate* in the precondition that the program will “recurse on”, and only allows a recursive call once this predicate has been unfolded at least once. This limitation makes it impossible to synthesize programs with more complex recursion patterns: even something as simple as a (helper-free) function that *deallocates two trees*, as such a function would need to traverse both of those trees recursively in a single run.

In the remainder of this section we will demonstrate how SSL_{\cup} overcomes both of these limitations by harnessing the cyclic proof methodology to enhance the CALL rule.

³Proof assistants like Coq [10] impose similar restrictions on recursion.

$$\begin{array}{c} \text{CALL} \\ \frac{\{\phi_f; P\} \rightsquigarrow \{\psi_f; S\} \mid f(\bar{x}_i) \quad \vdash \phi \Rightarrow [\sigma]\phi_f \quad \frac{\{\phi \wedge [\sigma]\psi_f; [\sigma]S * R\} \rightsquigarrow \{Q\} \mid c}{\{\phi; [\sigma]P * R\} \rightsquigarrow \{Q\} \mid f(\sigma(\bar{x}_i)); c}}{\{\phi; P\} \rightsquigarrow \{\psi; Q\} \mid c} \\ \text{PROC} \end{array}$$

Figure 2. Rules for calls and definitions in SSL_{\cup} (simplified).

2.2 Recursive Programs from Cyclic Proofs

The primary difference between SSL and SSL_{\cup} , is that the latter synthesizes a program by searching for a *cyclic* derivation, with recursive calls arising from *backlinks* in the proof.

New rules for function definitions and calls. To cater for synthesizing recursive calls using cyclic proofs, SSL_{\cup} introduces a new, more general version of the CALL rule, shown in Fig. 2. As before, this rule synthesizes a function call $f(\sigma(\bar{x}_i))$ (where σ substitutes actuals for formals) whenever some part of the current symbolic heap can be unified with f 's spatial precondition P . The main difference is that the new rule no longer requires f 's pre-/postcondition to be the top-level specification provided by the user. Instead—perhaps surprisingly—this rule simply requires that this pre-/post can *synthesize* the call $f(\bar{x}_i)$. At a first glance, this rule is not very useful: in order to synthesize a function call, we (seemingly) need to synthesize a function call! The intention of this rule is to use its first premise as a *bud*, *i.e.*, link it back to an identical *companion* goal earlier in the derivation.

But where do we find a companion goal whose emitted program is $f(\bar{x}_i)$? This is where the second new rule of SSL_{\cup} , PROC, comes into play (Fig. 2). The PROC rule should be understood as a way to “label” a node in a program derivation with a function definition $f(\bar{x}_i)$. According to this rule, one can synthesize a definition for the function $f(\bar{x}_i)$, if one can synthesize the body c for this function out of the goal with same pre- and postconditions. Note that unlike all the rules we've seen so far, PROC does not propagate the code emitted by its premise (the function body) to the conclusion; instead it abstracts this code away into an *identity call* $f(\bar{x}_i)$, *i.e.*, a call whose actual parameters coincide with the formals. The most common occurrence of the rule PROC is at the root of the derivation, with the top-level goal being its conclusion. However, we will see in Sec. 2.3 how this rule also makes it possible to synthesize auxiliary recursive functions.

Cyclic derivation of *treefree*. We now demonstrate the application of these new rules to build a cyclic derivation of *treefree*, shown in Fig. 3. For the sake of brevity, we omit the payload set variables and pure constraints related to them (*e.g.*, $s = \{v\} \cup s_l \cup s_r$) from the goals, as they don't play any role in our example. The instances of the *tree* predicate are annotated with *cardinality* variables α , α_l , and α_r , which play a role in the termination argument, as we will explain shortly. The proof's root goal serves as a conclusion for an application of PROC, with the “conjectured” function f being *treefree*. Most of the subsequent applications follow the sequence of rewritings via standard SSL rules outlined in

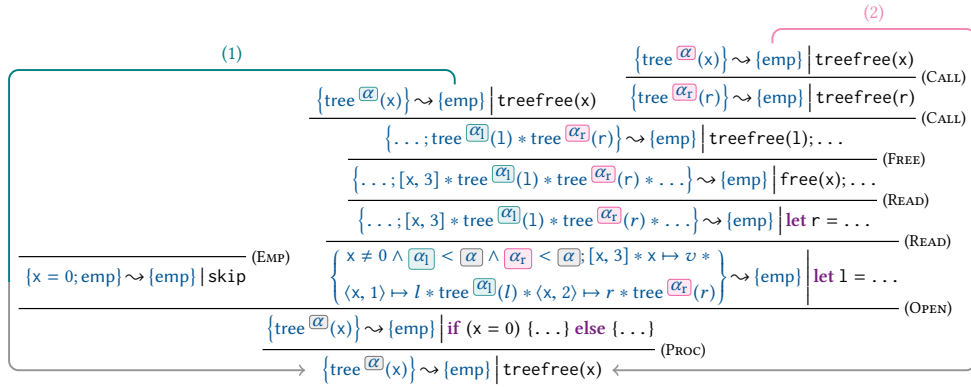


Figure 3. A derivation of $\text{treefree}(x)$ with backlinks. Some pure assertions are omitted.

Sec. 2.1, with some evident parts of the goals elided. An important twist happens at the two inner derivation nodes corresponding to the premise of applying of the rule CALL. The goal in both nodes matches precisely the conclusion of the top-level application of PROC. These pairs of matching nodes create *backlinks* in the program derivation, connecting the premises of two CALL applications to the top-level goal.

Proving termination. Cyclic reasoning is not valid in general: cyclic proofs must satisfy a well-formedness condition in order to preclude infinite derivations. Intuitively, one can view the derivation as a directed graph with edges pointing from conclusions to premises and from buds to companions. The well-formedness condition requires that along every infinite path in this graph, some well-founded measure decreases infinitely often [8]. In SSL_{\cup} , infinite paths in the derivation correspond to *potential* infinite traces of the program’s execution. The well-formedness condition ensures that no such infinite executions exist since this would entail an infinitely decreasing chain in a well-founded set.

To trace well-founded measures, SSL_{\cup} annotates inductive predicates with *cardinality* variables, which can be seen as *sizes of the heap models* of the corresponding predicate. We automatically instrument all predicate definitions with cardinality information; for example, the instrumentation of the tree predicate (3) is highlighted via gray boxes below:

$$\text{tree}^{\alpha}(x, s) \triangleq x = 0 \Rightarrow \{s = \emptyset; \text{emp}\} \\ \left| \begin{array}{l} x \neq 0 \Rightarrow \left\{ \begin{array}{l} \alpha_l < \alpha \wedge \alpha_r < \alpha \wedge s = \{v\} \cup s_l \cup s_r; \\ [x, 3] * x \mapsto v * \langle x, 1 \rangle \mapsto l * \langle x, 2 \rangle \mapsto r \\ * \text{tree}^{\alpha_l}(l, s_l) * \text{tree}^{\alpha_r}(r, s_r) \end{array} \right\} \end{array} \right.$$

Note the highlighted *cardinality constraints* we add to the second (recursive) clause, which state that cardinalities α_l and α_r of the left and right subtrees are strictly smaller than α , that of the enclosing tree. Given this instrumentation, the cardinality constraints end up in the goal’s precondition upon unfolding the predicate (via OPEN), as shown in Fig. 3.

To show that the derivation in Fig. 3 is well-formed, we must pick a sequence of cardinality variables to *trace* along each infinite path and prove that this sequence is strictly decreasing. All infinite paths in this derivation consist of

arbitrarily alternating cycles (1) and (2). In each cycle, we will start by tracing α , and then switch to either α_l or α_r , depending on which cycle is being traversed. In either case, when we make the switch, the traced cardinality strictly decreases (following the cardinality constraints in the predicate definition), while elsewhere along the cycle it stays unchanged. Hence we have shown that along each infinite path the cardinality must strictly decrease infinitely often.

As prior work has shown [35], given cardinality constraints, the appropriate sequence of cardinalities to trace can be inferred *automatically* using automata-theoretic tools—an observation that is crucial for automating synthesis via SSL_{\cup} . This mechanism subsumes termination measures based on maximum and/or lexicographic ordering of multiple arguments, and enable SSL_{\cup} to synthesize non-structurally recursive programs, for example, one that deallocates *two trees* as part of the same traversal (see technical report [16]).

2.3 Synthesizing Auxiliary Recursive Functions

We now move on to the core of our contribution and illustrate how cyclic synthesis handles the motivating example from the introduction: the procedure that flattens a tree into a linked list, whose specification (2) we repeat for convenience:

$$\{r \mapsto x * \text{tree}(x, s)\} \rightsquigarrow \{r \mapsto y * \text{sll}(y, s)\}$$

We omit the definition of the singly-linked list predicate sll , since it is standard for Separation Logic [34] and analogous to the tree predicate (3). The specification enforces that the payload set of the output list be the same as that of the input tree by using the same logical variable s in the pre- and postcondition. For the sake of brevity, we omit the set-related reasoning in the derivations shown below.

Synthesized code. The program synthesized from this specification is shown in Fig. 5. We have given more descriptive names (in lieu of automatically-generated ones) to local variables and the auxiliary function. The main procedure `flatten` calls itself recursively twice to produce the lists corresponding to the left and right subtrees of the node x . The head of the first list, returned from the call in line 9 via r , is then stored in y_1 , while the head of the second list is stored in r (through an extra level of indirection) after the call in line

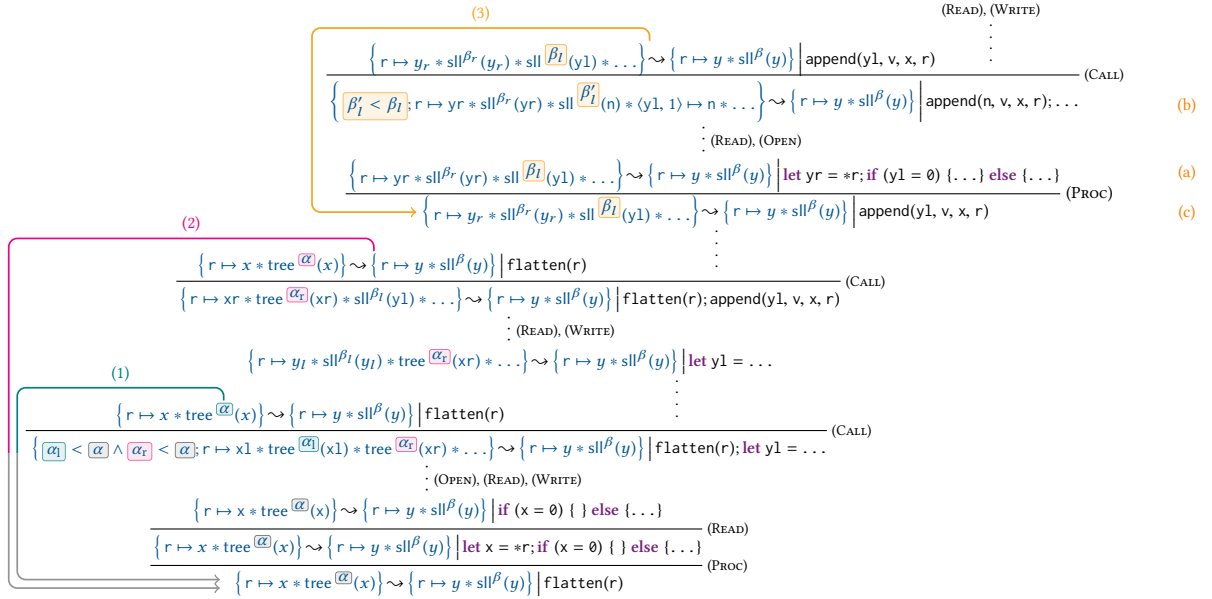


Figure 4. A derivation of `flatten` and its recursive auxiliary `append`.

```

1 void flatten(r) {
2   let x = *r;
3   if (x = 0) {
4   } else {
5     let v = *x;
6     let x1 = *(x + 1);
7     let xr = *(x + 2);
8     *r = x1;
9     flatten(r);
10    let y1 = *r;
11    *r = xr;
12    flatten(r);
13    append(y1, v, x, r);
14  }
15 }
16 void append(y1, v, x, r) {
17   let yr = *r;
18   if (y1 = 0) {
19     let y = malloc(2);
20     free(x);
21     *y = v;
22     *(y + 1) = yr;
23     *r = y;
24   } else {
25     let n = *(y1 + 1);
26     append(n, v, x, r);
27     let y = *r;
28     *(y1 + 1) = y;
29     *r = y1;
30   }

```

Figure 5. Tree flattening program synthesized by CYPRESS.

12. As its last statement, `flatten` calls the recursive auxiliary function `append`, passing it the pointers to the two lists, `y1` and `r`, as well as the parent tree node `x` and its payload `v`.

The auxiliary procedure `append` concatenates the two lists, `y1` and `r`, inserting a new element with payload `v` in the middle (hence, computing an *in-order* unfolding of the tree). To this end, it traverses the first list, `y1`, recursively. Once it reaches the base case where `y1` is empty (lines 19–23), it frees the tree node `x`⁴, allocates a new list node `y` with payload `v`, and prepends it to the second list, storing the result in `r`. Its inductive case (lines 25–29) calls `append` recursively on the tail of `y1`, after which it adjusts the tail pointer (`y1 + 1`) appropriately and stores the result again in `r`.

Cyclic derivation of `flatten`. To understand how the auxiliary `append` has been discovered, let us take a look at the SSL_{\cup} derivation of `flatten`, shown in Fig. 4. The first half

⁴Instead of passing `x` to `append`, it would have been more natural to deallocate it immediately in `flatten`. Although SSL_{\cup} is capable of deriving either program, our implementation makes the less natural choice, which we discuss further in Sec. 5.4.

of this derivation is uneventful and quite similar to what we have already seen for `treefree` in Fig. 3: the two applications of `CALL` induce two backlinks (1) and (2) to the top-level goal and correspond to the recursive calls to `flatten(r)` in lines 9 and 12 of Fig. 5.

After these two calls, however, we find ourselves in the node (a) of the derivation⁵, with the following synthesis goal:

$$\left\{ \begin{array}{l} s = \{v\} \cup s_l \cup s_r; r \mapsto y_r * sll(y1, s_l) * sll(y_r, s_r) * \\ [x, 3] * x \mapsto v * \langle x, 1 \rangle \mapsto _ * \langle x, 2 \rangle \mapsto _ \\ \sim \{r \mapsto y * sll(y, s)\} \end{array} \right\} \quad (8)$$

Here, we are given two lists, `y1` and `yr`, and we need to obtain a single list `y` that contains all of their elements plus `v`. How should the synthesizer proceed to solve this goal? It proceeds just like it would for any goal that contains inductive predicates in the precondition: by unfolding one of their instances—in this case `sll(y1, sl)`—via `OPEN`. The recursive premise of `OPEN` is depicted in the derivation node (b), whose precondition again features two lists: one of them, `yr`, is the same as before, and the other one, `n`, is the tail of `y1`.

Abducing the auxiliary. At this point the synthesizer realizes that the `CALL` rule is applicable again; this time, however, the companion goal is *not* the top-level specification (2), but rather the goal (8) from node (a). Indeed, a sub-heap of the current precondition (the one containing everything but the head of `y1`) can be *unified* with the precondition of (8) by substituting `n` for `y1`. In order to become an eligible companion, however, a node must emit a procedure call, which is clearly not the case for (a). To bridge this mismatch, the synthesizer retroactively inserts an application of `PROC` just below (a) and creates a new node (c) as its conclusion. The

⁵For now, imagine that node (c) and the adjacent application of `PROC` are not there. We will come back to them shortly.

Variable	x, y	Alpha-numeric identifiers \in PV	Logical variable	v, ω	Cardinality variable	α
Size, offset	n, ι	Non-negative integers	\mathcal{T} -term	κ	$::= v \mid e \mid \dots$	
Expression	e	$::= 0 \mid \text{true} \mid x \mid e = e \mid e \wedge e \mid \neg e \mid d$	Pure logic term	ϕ, ψ, χ	$::= \kappa \mid \phi = \phi \mid \phi \wedge \phi \mid \neg \phi$	
\mathcal{T} -expr.	d	$::= n \mid x \mid d + d \mid n \cdot d \mid \{ \} \mid \{d\} \mid \dots$	Symbolic heap	P, Q, R	$::= \text{emp} \mid \langle e, \iota \rangle \mapsto e \mid [e, \iota] \mid p^\alpha(\overline{\phi_i}) \mid P * Q$	
Command	c	$::= \text{let } x = *(x + \iota) \mid *(x + \iota) = e \mid$ $\text{let } x = \text{malloc}(n) \mid \text{free}(x) \mid \text{error}$ $\mid f(\overline{e_i}) \mid c; c \mid \text{if } (e) \{c\} \text{ else } \{c\}$	Heap predicate	\mathcal{D}	$::= p^\alpha(\overline{x_i}) : e_j \Rightarrow \exists \overline{y}. \{ \chi_j; R_j \}$	
Program	Π	$::= \overline{f(\overline{x_i})} \{ c \}; c$	Assertion	\mathcal{P}, \mathcal{Q}	$::= \{ \phi; P \}$	
			Environment	Γ	$::= \forall \overline{x_i}. \exists \overline{y_j}. \quad \text{Context } \Sigma ::= \overline{\mathcal{D}}$	

Figure 6. SSL_{\cup} programming language and assertion syntax.

new node has the same synthesis goal as (a) and an identity call to a fresh procedure—here $\text{append}(y1, v, x, r)$ —as the emitted code.

This *lazy* application of PROC corresponds to abducting the recursive auxiliary on demand. The parameters of the new procedure are all the program variables of (a); its pre- and postcondition are defined by (8), the synthesis goal of (a); finally, its body is the code emitted by (a) (derived by the SSL_{\cup} sub-derivation rooted at (a)).

Termination. As before, we perform a global well-formedness check on the entire derivation. Note, however, that in this case infinite paths can either follow some combination of backlinks (1) and (2) or always follow the backlink (3); in other words, in the absence of mutual recursion, the termination arguments for the two procedures are entirely disjoint. Termination follows from the highlighted cardinality constraints: $\alpha_l < \alpha$, $\alpha_r < \alpha$, and $\beta'_l < \beta_l$.

2.4 More Examples

We conclude this section by outlining two examples that showcase unique capabilities of SSL_{\cup} and CYPRESS and go beyond the state of the art in synthesis with auxiliaries [13]. More details and synthesized code for these examples can be found in our technical report [16].

Flattening a tree in-place. Our first example leverages the imperative nature of our underlying language, allowing CYPRESS to flatten a binary tree into a *doubly-linked list in-place*: $\{\text{tree}(x, s)\} \text{flatten_to_dll}(x) \{\text{dll}(x, z, s)\}$. Note that here we require that the root of the input tree and the head of the output list be located at the *same* address x . We omit the standard definition of the dll predicate, but note that in-place flattening is possible because both a tree node and a doubly-linked list node are represented as a three-element record, so one can be reinterpreted as the other.

Mutual recursion. Our final example shows the ability of CYPRESS to synthesize manipulations with mutually-recursive data structures. An n -ary tree, *a.k.a. rose tree* [23], can be implemented by storing the children of a tree node in a linked list. Its definition in Separation Logic uses a pair of mutually recursive predicates: rtree , representing a tree, and children , representing a linked list of trees. Given these predicates and a specification $\{\text{rtree}(x, s)\} \text{rtree_free}(x) \{\text{emp}\}$,

CYPRESS is able to generate a pair of *mutually recursive* functions that deallocate a rose tree.

3 Cyclic Program Synthesis, Formally

In this section, we give a formal presentation of declarative rules of SSL_{\cup} and describe the underlying metatheory.

3.1 Programs and Assertions

Programming language. The target language of SSL_{\cup} is an imperative, C-like fragment with dynamic memory allocation, deallocation, store and load (Fig. 6, left). Values include at least booleans and integers, and a special type `loc` designates pointer variables. Pointers are isomorphic to unsigned integers, but there is only a single pointer constant, \emptyset (null). Expressions include at least variables, literal constants, equality check and logical connectives. Additional *theory-specific expressions* are allowed depending on the underlying theory used for checking entailment in derivations; our implementation supports linear integer arithmetic and sets. The language allows pointer arithmetic in the form $x + \iota$, but other arithmetic operations are disabled for pointers. The statement $f(\overline{e_i})$ denotes a procedure call with actual parameters $\overline{e_i}$. Procedures do not have a return value: the behavior of `return` is emulated by passing in an address of the heap location where the result should be stored. There are no variable re-assignments and no `while` loops. A program is a sequence of procedure definitions, followed by a statement.

Assertion language. The assertion language comprises pure assertions in the underlying theory, and SL assertions with inductively defined predicates (Fig. 6, right). Its semantics have been explored throughout Sec. 2. The set of pure logic terms (ϕ, ψ, χ) is a superset of program expressions e . The logic is sorted, and pure parts $(\phi \text{ in } \{ \phi; P \})$ are ensured to be Boolean expressions via simple type checking. Predicate instances are annotated with cardinality variables α and β .

Assertions are interpreted in an environment Γ in which some of the variables are universally quantified and others existentially quantified, with a prefix of the form $\forall \overline{x}. \exists \overline{y}$. Program variables are *always* included in the universal prefix. Logical variables are split between universal (also called *ghost* variables) and existential. We denote $\text{Vars}(\Gamma) = \{ \overline{x}, \overline{y} \}$ for all quantified variables, and $\text{PV}(\Gamma) = \{ \overline{x} \} \cap \text{PV}$, $\text{GV}(\Gamma) = \{ \overline{x} \} \setminus \text{PV}$, $\text{EV}(\Gamma) = \{ \overline{y} \}$ for program variables, ghost variables and existentials of Γ , respectively. We also use $e[\Gamma]$ for the set of all

R0 Terminal rules	
$\frac{\text{EMP} \quad \vdash \phi \Rightarrow \psi}{\Gamma; \{\phi; \text{emp}\} \rightsquigarrow \{\psi; \text{emp}\} \mid \text{skip}}$	$\frac{\text{INCONSISTENCY} \quad \vdash \phi \Rightarrow \perp}{\Gamma; \{\phi; P\} \rightsquigarrow Q \mid \text{error}}$
R1 Pointer operations	
$\frac{\text{READ} \quad \forall y. \Gamma; \{\phi \wedge y = e; \langle x, t \rangle \mapsto e * P\} \rightsquigarrow Q \mid c \quad x \in PV \quad y \in PV \setminus \text{Vars}(\Gamma)}{\Gamma; \{\phi; \langle x, t \rangle \mapsto e * P\} \rightsquigarrow Q \mid \text{let } y = *(x + t); c}$	
$\frac{\text{WRITE} \quad \Gamma; \{\phi; \langle x, t \rangle \mapsto e * P\} \rightsquigarrow \{\psi; \langle x, t \rangle \mapsto e * Q\} \mid c \quad \text{Vars}(e) \subseteq PV}{\Gamma; \{\phi; \langle x, t \rangle \mapsto e' * P\} \rightsquigarrow \{\psi; \langle x, t \rangle \mapsto e * Q\} \mid *(x + t) = e; c}$	
$\frac{\text{ALLOC} \quad \Gamma; \left\{ \begin{array}{l} \phi; [y, n] * \left(\langle y, t \rangle \mapsto t_i \right)_{0 \leq i < n} * P \\ \psi; [x, n] * \left(\langle x, t \rangle \mapsto e_i \right)_{0 \leq i < n} * Q \end{array} \right\} \mid c \quad x \in \text{EV}(\Gamma)}{\Gamma; \{\phi; P\} \rightsquigarrow \left\{ \psi; [x, n] * \left(\langle x, t \rangle \mapsto e_i \right)_{0 \leq i < n} * Q \right\} \mid \text{let } y = \text{malloc}(n); c}$	
$\frac{\text{FREE} \quad \Gamma; \{\phi; P\} \rightsquigarrow Q \mid c \quad x \in PV}{\Gamma; \{\phi; [x, n] * \left(\langle x, t \rangle \mapsto e_i \right)_{0 \leq i < n} * P\} \rightsquigarrow Q \mid \text{free}(x); c}$	
R2 Predicate unfolding	
$\frac{\text{OPEN} \quad \Gamma \cup \forall \overline{\omega}_{jk}; \left[\overline{t}_i / \overline{v}_i \right] \left\{ \phi \wedge e_j \wedge \chi_j; R_j * P \right\} \rightsquigarrow Q \mid c_j \quad \text{for all } j=1..r}{p^\alpha(\overline{v}_i) : e_j \Rightarrow \exists \overline{\omega}_{jk}. \left\{ \chi_j; R_j \right\}_{j=1..r} \in \Sigma \quad \omega_{jk} \notin \text{Vars}(\Gamma), \text{GV}(\overline{t}_i) = \emptyset}}{\Gamma; \left\{ \phi; p^\alpha(\overline{t}_i) * P \right\} \rightsquigarrow Q \mid \begin{array}{l} \text{if } \left(\left[\overline{t}_i / \overline{v}_i \right] e_1 \right) \{c_1\} \\ \text{else if } \left(\left[\overline{t}_i / \overline{v}_i \right] e_2 \right) \{c_2\} \text{ else } \dots \end{array}}$	
$\frac{\text{CLOSE} \quad \Gamma \cup \exists \overline{\omega}_{jk}; P \rightsquigarrow \left[\overline{t}_i / \overline{v}_i \right] \left\{ \phi \wedge e_j \wedge \chi_j; R_j * Q \right\} \mid c_j \quad \text{for some } j=1..r}{p^\alpha(\overline{v}_i) : e_j \Rightarrow \exists \overline{\omega}_{jk}. \left\{ \chi_j; R_j \right\}_{j=1..r} \in \Sigma \quad \omega_{jk} \notin \text{Vars}(\Gamma)}}{\Gamma; P \rightsquigarrow \left\{ \phi; p^\alpha(\overline{t}_i) * Q \right\} \mid c}$	
R3 Procedures	
$\frac{\text{PROC} \quad \Gamma; P \rightsquigarrow Q \mid c \quad \{\overline{x}_i\} = \text{PV}(\Gamma)}{\Gamma; P \rightsquigarrow Q \mid f(\overline{x}_i)}$	
$\frac{\text{CALL} \quad \forall \overline{x}_i, \overline{v}_j, \exists \overline{\omega}_k; \left\{ \phi'; P \right\} \rightsquigarrow \left\{ \psi'; S \right\} \mid f(\overline{x}_i)}{\Gamma \cup \forall \sigma(\overline{\omega}_i); \left\{ \left[\sigma \right] \psi' \wedge \phi; \left[\sigma \right] S * R \right\} \rightsquigarrow Q \mid c \quad \text{dom}(\sigma) = \{\overline{x}_i, \overline{v}_j, \overline{\omega}_k\}}{\vdash \phi \Rightarrow \left[\sigma \right] \phi' \quad \sigma(x_i) \in e[\Gamma]; \sigma(v_j) \in \kappa[\Gamma]}}$	
$\Gamma; \{\phi; \left[\sigma \right] P * R\} \rightsquigarrow Q \mid f(\sigma(\overline{x}_i)); c$	
$\frac{\text{CALLSETUP} \quad \Gamma; \{\phi; P\} \rightsquigarrow \{\psi; S\} \mid c_1 \quad \Gamma \cup \text{BV}(c_1); \{\psi; S * R\} \rightsquigarrow Q \mid f(\overline{e}_i); c_2}{\Gamma; \{\phi; P * R\} \rightsquigarrow Q \mid c_1; f(\overline{e}_i); c_2}$	

Figure 7. Synthesis rules of SSL_{\cup} .

expressions that can be constructed using $\text{PV}(\Gamma)$, and $\kappa[\Gamma]$ —all logical terms that can be constructed with $\text{Vars}(\Gamma)$. When joining environments, the quantifier pattern is preserved, and we use the notation $\forall \overline{x}. \exists \overline{y}. \exists \overline{y}' = \forall \overline{x}, \overline{x}'. \exists \overline{y}, \overline{y}'$. All variable sets $\overline{x}, \overline{x}', \overline{y}, \overline{y}'$ must be disjoint.

A clause $e \Rightarrow \exists \overline{y}. \{\chi; R\}$ in a definition of a heap predicate $p^\alpha(\overline{x}_i)$ consists of a *selector* expression e and an assertion $\{\chi; R\}$. As explained in Sec. 2.2, the pure part χ contains a *cardinality constraint* of the form $\alpha_k < \alpha$ for each predicate instance $q^{\alpha_k}(\dots)$ inside R . These are not part of the surface syntax, but are added automatically during pre-processing.

3.2 Proof Rules

Fig. 7 lists all synthesis-related SSL_{\cup} rules in a declarative manner. The rules operate on *transforming entailment* judgments $\Gamma; P \rightsquigarrow Q \mid c$, which informally means that any state that satisfies the assertion P can be transformed into some state that satisfies Q by a statement c [29]. Within rules, for clarity, we follow the convention of using lower latin letters x, y for program variables, e, t for program-level terms (of the syntactic class e in Fig. 6), greek letters ν, ω for logical variables, and ϕ, ψ, χ for logical formulas. We will now describe the rules and their effects on the synthesis.

R0 Terminal rules. The rules EMP and INCONSISTENCY form the leaves of a derivation by emitting a `skip` for a trivial goal or `error` for a vacuous goal (unsatisfiable precondition). Here $\vdash \phi \Rightarrow \psi$ denotes entailment between pure formulas in the underlying theory; our implementation uses an SMT solver to discharge these premises.

R1 Rules for atomic operations. Each of the pointer operations of the target language has a corresponding rule that describes when to emit it: READ and WRITE for memory access, ALLOC and FREE for dynamic memory management. These rules are more restricted versions of the corresponding symbolic execution rules of standard Separation Logic [5]; the additional restrictions are required to guide rule applications in the context of synthesis, where the program is not available. For example, the WRITE rule uses the heaplet $\langle x, t \rangle \mapsto e$ in the goal's *postcondition* to determine what should be written into the address x .

R2 Rules for inductive predicates. The two rules OPEN and CLOSE unfold definitions of inductive predicates in the pre- and the postcondition, respectively. These rules have dual effects on the program in the sense that OPEN *eliminates* a predicate instance and hence performs a case split on its clauses, while CLOSE *introduces* an instance and hence picks a single clause non-deterministically. As a result, CLOSE does not emit any code, while OPEN emits a conditional statement with one branch per clause, using clause selectors as guards.

R3 Rules for procedures. These rules are the main innovation of SSL_{\cup} as compared with SSL. The rules PROC and CALL were discussed in detail in Sec. 2.2. The rule PROC delineates procedure boundaries by emitting a call to a procedure whose body is emitted by its single premise.

Procedure calls are generated via the rule CALL. It bundles together Separation Logic-style framing, with R as the frame, and substitution of actual into formal parameters via σ , which is also applied to the postcondition $\{\psi'; S\}$. Existential variables in the procedure's environment are remapped to fresh ghost variables in the second premise of the rule. Formal parameters x_i are mapped to program expressions $e[\Gamma]$ using program variables of Γ , and ghosts v_j are mapped to logical terms using any variables of Γ .

CALLSETUP handles the situations when some code has to be emitted before the call can be made. To see why this rule

is required, consider the recursive call `flatten(r)` in line 9 of Fig. 5. To enable this call we first need to write the root of the left subtree `x1` into the return location `r`; but to trigger this write, we need to have `r ↦ x1` in the goal's postcondition! `CALLSETUP` enables that: it performs *sequential decomposition* of the synthesis goal with the call's precondition $\{\psi; S\}$ as the intermediate state. As a result, `r ↦ x1` ends up in the postcondition of its first subgoal, triggering the synthesis of the write in line 8. Although the declarative presentation in Fig. 7 makes it look like the assertion $\{\psi; S\}$ is chosen non-deterministically, in Sec. 4 we explain how `CYPRESS` implements this rule efficiently via a *call abduction oracle*.

Logical rules. The remainder of SSL_{\cup} rules are *logical rules*, which do not emit new code, but instead transform the goal in a way that would ultimately allow the application of other (operational) rules. Our logical rules are mostly standard for SL-based theorem provers [5], so we relegate them to the technical report [16] to save space. They include the `FRAME` rule we have seen in Sec. 2, as well as rules for pure reasoning, e.g. eliminating existentials. In Sec. 4 we discuss how `CYPRESS` implements such pure reasoning efficiently.

3.3 Cyclic Program Derivations

An SSL_{\cup} derivation is a tree consisting of goals $\Gamma; \mathcal{P} \rightsquigarrow \mathcal{Q} \mid c$, and constructed using the inference rules in Fig. 7. In contrast to the standard notion of a proof as a finite derivation tree, SSL_{\cup} derivations are permitted to be regular, non-well-founded (*i.e.*, infinitely tall) trees. Regularity ensures that an SSL_{\cup} derivation tree always has a finite representation as a (possibly) cyclic graph. Concretely, we represent SSL_{\cup} derivations as finite trees, along with a set of *backlinks* connecting each non-terminal leaf node to a syntactically identical ancestor node. In other words, we allow goals from the middle of the proof to be used again as premises higher up the derivation tree. Formally, each backlink denotes the infinite unfolding of the path connecting the leaf—called *bud*—with its associated internal node—called *companion*.

We call SSL_{\cup} derivations *pre-proofs*, since they do not necessarily derive *terminating* programs. To ensure termination, we require that pre-proofs satisfy an additional global property, defined in terms of traces of cardinality variables.

Definition 3.1 (Trace pairs). Let \mathcal{G} and \mathcal{G}' be, respectively, the conclusion and a premise of an inference rule r , and let α and β be cardinality variables occurring universally in \mathcal{G} and \mathcal{G}' , respectively. We say that (α, β) is a *trace pair* for $(\mathcal{G}, \mathcal{G}')$ when: either r is `CALL` with substitution σ , \mathcal{G}' is the left-hand premise, and $\alpha = \sigma(\beta)$; or $\vdash \phi \Rightarrow \beta \leq \alpha$ holds, where ϕ is the pure precondition of \mathcal{G} . If $\vdash \phi \Rightarrow \beta < \alpha$ also holds, we say that the trace pair is *progressing*.

A *path* in a pre-proof P is a sequence \mathcal{G}_i ($i \geq 0$) of nodes in P such that each node \mathcal{G}_i is the parent of \mathcal{G}_{i+1} in the infinite derivation corresponding to P .

Definition 3.2 (Traces). A *trace* is an infinite sequence α_i ($i \geq 0$) of cardinality variables. We say that a trace follows a path \mathcal{G}_i ($i \geq 0$) in a pre-proof P when (α_i, α_{i+1}) is a trace pair for $(\mathcal{G}_i, \mathcal{G}_{i+1})$ for each $i \geq 0$. When (α_i, α_{i+1}) is progressing, we say that the trace *progresses at i* . A trace is called *infinitely progressing* if it progresses at infinitely many points.

Definition 3.3 (Proofs). A pre-proof P is said to satisfy the global trace condition when every infinite path in P is followed by an infinitely progressing trace.

The global trace condition is an ω -regular property, so it is decidable when a pre-proof is a proof, *cf.* [9, Prop. 7.4]. We write $\vdash \mathcal{G}$ when there is a proof deriving \mathcal{G} .

3.4 Soundness

SSL_{\cup} inherits the memory model and operational semantics from traditional SL [5]. In the interest of space we only state the soundness theorem here. Its proof and the rest of the metatheory are relegated to the technical report [16].

Theorem 3.4 (Soundness). *If $\vdash \Gamma; \mathcal{P} \rightsquigarrow \mathcal{Q} \mid c$, then for any heap/stack pair $\langle h, s \rangle$ that satisfies \mathcal{P} , there exists a heap/stack pair $\langle h', s' \rangle$ that satisfies \mathcal{Q} , such that executing c from state $\langle h, s \rangle$ terminates in state $\langle h', s' \rangle$.*

4 Cyclic Program Synthesis, Pragmatically

We implemented cyclic program synthesis in a new synthesizer called `CYPRESS` [17]. `CYPRESS` takes as input a synthesis goal $\{\mathcal{P}\} f(\bar{x}_i) \{\mathcal{Q}\}$ together with the definitions of all inductive predicates it mentions, and performs backtracking *proof search* for an SSL_{\cup} derivation of the judgment $\mathcal{P} \rightsquigarrow \mathcal{Q} \mid f(\bar{x}_i)$. Once a proof has been constructed, *extracting* the synthesized program is straightforward: we simply consider each application of the `PROC` rule in isolation; each such application gives rise to a procedure, whose signature and body are the code emitted by the conclusion and the premise of `PROC`, respectively. For example, the two applications of `PROC` in Fig. 4 give rise to two procedures, `flatten` and `append`.

In the rest of this section we outline the mechanisms that make our proof search tractable. Since `CYPRESS` builds upon `SUSLIK`—the original implementation of SSL —we only focus on the new mechanisms, which make `CYPRESS` more general and/or more efficient than its predecessor. `CYPRESS` also leverages existing proof search features of `SUSLIK`, such as early failure rules, phased proof search, and branch abduction (which enables it to synthesize conditionals beyond those in predicate selectors); we refer the reader to [29] for a detailed account of these features.

Best-first search. One difference in the overall search algorithm is that `CYPRESS` uses memoizing *best-first search* (inspired by [19]), instead of `SUSLIK`'s naïve depth-first search. The search is guided by a cost function that assigns a cost to each heaplet in the goal's pre- and postcondition, with predicate instances growing more expensive as they get unfolded

$$\begin{array}{c}
\text{UNIFY} \\
\frac{\Gamma; \{\phi; P * [\sigma_1]R\} \rightsquigarrow \{\psi \wedge \psi'; Q * [\sigma_1]R\} | c}{\psi' = \bigwedge_{v \in V} (\sigma_1(v) = \sigma_2(v)) \quad \text{dom}(\sigma_1) = \text{dom}(\sigma_2) = V} \\
\frac{\Gamma; \{\phi; P * [\sigma_1]R\} \rightsquigarrow \{\psi; Q * [\sigma_2]R\} | c}{\text{SOLVE-}\exists} \\
\frac{\Gamma; \{\phi; P\} \rightsquigarrow [\sigma]\{\psi; Q\} | c}{\vdash \phi \Rightarrow [\sigma]\psi \quad \text{dom}(\sigma) = \{\bar{w}\}; \text{EV}(\Gamma) = \emptyset} \\
\frac{\Gamma \cup \exists \bar{w}; \{\phi; P\} \rightsquigarrow \{\psi; Q\} | c}{}
\end{array}$$

Figure 8. Two algorithmic rules for pure reasoning.

or go through a call. This cost function prevents the search from getting stuck in a branch that performs infinitely many unfoldings or calls, and encourages it to focus on smaller (and therefore, hopefully, easier to solve) goals.

4.1 Synthesizing Calls

To understand how CYPRESS applies the rules CALL and CALLSETUP in practice, let us revisit the tree flattening program from Fig. 5. After the initial sequence of READ and OPEN applications, which generate lines 1–7, we find ourselves with a pre-heap $r \mapsto x * \text{tree}(x1) * \text{tree}(xr) * \dots$ (where ellipsis stands for the node at x and its payload, and we omit cardinality and set parameters for brevity). To decide whether a call can be synthesized from this goal, CYPRESS considers all *candidate companion* goals, *i.e.* all ancestor goals separated from the current goal by at least one application of OPEN. OPEN is the only rule where cardinality traces can progress, hence any well-formed cycle in the derivation must include an application of OPEN. In our case, the top-level goal with the pre-heap $r \mapsto x * \text{tree}(x)$ is such a candidate.

Call abduction oracle. Given a candidate companion, synthesizing a call involves guessing (1) the *substitution* σ of formals into actuals (*e.g.*, $[x \mapsto x1]$); (2) the *frame* R , *i.e.*, the part of the precondition untouched by the call (*e.g.*, $\text{tree}(xr) * \dots$); and (3) the *setup statement* c required to satisfy the companion’s precondition (here $*r = x1$). CYPRESS finds all these three components at once, using a mechanism we dub *call abduction oracle*. This oracle is a separate synthesis problem that attempts to “bridge the gap” between the current goal’s precondition and that of the companion. In our example, the oracle attempts to derive:

$$\{r \mapsto x * \text{tree}(x1) * \text{tree}(xr) * \dots\} \rightsquigarrow \{r' \mapsto x' * \text{tree}(x')\} | c$$

Note that all variables in the companion goal are replaced with fresh existentials. The call abduction oracle only uses the subset of SSL_{\cup} rules triggered by the post-condition (*e.g.* WRITE and ALLOC), since the rest of rules (*e.g.* READ, OPEN, CALL) could already fire before the oracle was invoked. It also uses a modified version of the EMP rule, which allows the pre-heap to remain non-empty. Upon successful completion, the remaining pre-heap becomes the frame R , the code emitted during this derivation becomes the setup statement c , and the σ comprises all existential substitutions from this derivation.

$$\begin{array}{c}
\text{EMP} \frac{}{\{\text{emp}\} \rightsquigarrow \{s \cup \{a\} = \{a\} \cup s; \text{emp}\}} \\
\text{SOLVE-}\exists \frac{}{\{\text{emp}\} \rightsquigarrow \{s \cup \{a\} = \{a\} \cup w; \text{emp}\}} \\
\text{FRAME} \frac{}{\{\text{sll}(x, s \cup \{a\})\} \rightsquigarrow \{s \cup \{a\} = \{a\} \cup w; \text{sll}(x, s \cup \{a\})\}} \\
\text{UNIFY} \frac{}{\{\text{sll}(x, s \cup \{a\})\} \rightsquigarrow \{\text{sll}(x, \{a\} \cup w)\}}
\end{array}$$

Figure 9. An example derivation with rules from Fig. 8.

Termination checking. Whenever the call abduction oracle succeeds, CYPRESS adds appropriate applications of CALL and CALLSETUP, inserts an application of PROC below the companion candidate (using all its program variables as formals and a fresh procedure name), and forms a backlink from the first premise of CALL to the conclusion of PROC. Every time a backlink is formed, CYPRESS checks whether the pre-proof constructed so far satisfies the trace condition from Sec. 3.3. To this end, it builds a graph of the current pre-proof with edges labeled with all available trace pairs, according to Def. 3.1. It then invokes CYCLIST [35], an off-the-shelf cyclic theorem prover, which uses an automata-theoretic algorithm to check whether every infinite path in the graph contains an infinitely progressing trace.

4.2 Pure Reasoning

Until now we have focused on synthesis rules, which directly emit code; the success of synthesis, however, also crucially relies on logical rules, which transform the goal into a form where synthesis rules can fire. Logical reasoning in SSL_{\cup} is far from straightforward because it needs to support arbitrary SMT-decidable theories in its pure formulas.

To illustrate the challenges of pure reasoning, consider the following (simplified but representative) synthesis goal:

$$\forall x, s, a. \exists w. \{\text{sll}(x, s \cup \{a\})\} \rightsquigarrow \{\text{sll}(x, \{a\} \cup w)\} \quad (9)$$

The simplest—and hence most desirable—solution to this goal is skip, which informally can be obtained by simply framing away the sll heaplets; but formally, before FRAME can apply, we need to transform the post-heap to be syntactically identical to the pre-heap, which requires: (a) exploiting commutativity of set union, and (b) instantiating the existential w with s . Fig. 8 shows two novel logical rules implemented in CYPRESS, which perform such transformations efficiently; Fig. 9 demonstrates how these rules work together to solve the synthesis goal (9) (we omit the emitted program skip and highlight new parts in each sub-goal for readability).

Unification modulo theories. The traditional approach to exploiting equational theories—such as commutativity of union in our example—is to eagerly normalize the goal after every proof step, by computing all implied equalities between its subterms [4]. For example, if a specification mentions both $s \cup \{a\}$ and $\{a\} \cup s$, normalization would replace one with the other. The normalization approach has two downsides: first, it is rather inefficient if you assume only blackbox access to the SMT solver; second, it does not help with the goal like (9), because here $s \cup \{a\} = \{a\} \cup$

w is not a logical necessity, but rather a possibility, which happens to lead to the shortest solution. To circumvent these limitations, CYPRESS implements a new, lazy approach to equational reasoning via the rule UNIFY in Fig. 8, which we dub *unification modulo theories*. UNIFY looks for a pair of heaplets $[\sigma_1]R$ and $[\sigma_2]R$ in the pre- and postcondition that only differ in pure subterms, and speculatively unifies them, adding equalities between mismatched subterms, $\sigma_1(v) = \sigma_2(v)$, as proofs obligation to the pure postcondition. In Fig. 9, UNIFY is used to unify the pre- and post-heap of the goal (9), producing the proof obligation $s \cup \{a\} = \{a\} \cup w$. Even in the absence of existentials, this approach is more efficient than eager normalization, because the SMT solver only needs to check equalities between terms that appear in matching positions inside a heaplet.

Pure synthesis. The other major challenge is to find appropriate instantiations for existential variables. To this end, CYPRESS includes the rule SOLVE- \exists in Fig. 8, which picks a substitution σ from existentials to universals that validates the pure specification. To find such a substitution, CYPRESS needs to solve the constraint $\exists \sigma. \forall \bar{x}. \phi \Rightarrow [\sigma]\psi$, which is itself a synthesis problem in the pure subset of our logic. Although such pure synthesis is generally a challenging task, it has been the subject of much prior work [2, 21, 32]. CYPRESS out-sources pure synthesis queries to the CVC4 synthesizer [33].

5 Evaluation

We evaluated CYPRESS empirically along three axes: (1) *generality*: its ability to synthesize programs with complex recursion; (2) *efficiency*: the time it takes to synthesize programs; and (3) *utility*: the size of the input specification compared to the size of the generated programs, and the quality of generated programs.

5.1 Benchmarks and Setup

For our empirical evaluation we have assembled a suite of 46 synthesis benchmarks for pointer-manipulating programs. Each benchmark is defined by a top-level synthesis goal expressed as separation-logic specification (and optionally, specifications of library functions the code is allowed to invoke). We collected these benchmarks from three sources:

1. State of the art in synthesis *with recursive auxiliaries* [13]. We include eight out of their nine benchmarks. The remaining one, `merge_sort`, has a specification identical to another one (`sort`), but a different program template, which forces their tool to synthesize merge sort instead of insertion sort; CYPRESS does not use program templates, so the difference between these two benchmarks does not make sense in our setting. The tool [13] synthesizes functional programs from refinement types, so we manually translated these benchmarks into Separation Logic.
2. State of the art in synthesis of *heap-manipulating programs*: SuSLIK [29] and IMPSYNT [31]. We include all 22

benchmarks from SuSLIK, and the 11 recursive benchmarks from IMPSYNT (these two sets overlap); the excluded five IMPSYNT benchmarks are iterative versions of their recursive benchmarks, and similarly require sketches.

3. We supplement this set with 14 *new benchmarks*, half of which, to the best of our knowledge, cannot be solved by any existing synthesis tools. These benchmarks involve the interplay between auxiliary functions and heap manipulation (e.g., an in-place tree flattening) or operate on cyclic or mutually recursive structures (e.g., rose trees).

Out of the 46 benchmarks, 19 exercise *complex recursion*, i.e., they either require a non-trivial termination measure or a recursive auxiliary procedure (not given as a library function). We refer to this benchmark set as `complex` and use it as a primary focus of our empirical evaluation. All complex benchmarks are by construction out of reach for SuSLIK. The remaining 27 benchmarks only exercise *simple (structural) recursion*. We refer to this benchmark set as `simple`; although not the focus of our evaluation, we use these benchmarks to demonstrate the versatility of CYPRESS.

Experiment setup. For our main experiment, we ran CYPRESS on the `complex` benchmark set, and measured synthesis time and size of the generated code. For our second experiment, we ran CYPRESS on the `simple` benchmark suite and compared the synthesis times with SuSLIK. The purpose of this experiment is to confirm that searching a larger program space does not lead to significant degradation in performance. All experiments were conducted on a commodity laptop (2.7 GHz Intel Core i7 Lenovo Thinkpad with 16GB RAM), and CYPRESS was run as a single-threaded process. Timeout for all experiments was set to one hour.

5.2 Results

Experiment results on the `complex` benchmarks are shown in Tab. 1, and on the `simple` benchmarks in Tab. 2. All specifications and generated code can be found in the technical report [16].

5.2.1 Generality. The results in Tab. 1 confirm that CYPRESS is able to synthesize a variety of functions with complex recursion. Benchmarks 10–13 were discussed in Sec. 2. Out of all benchmarks in Tab. 1, only 10 and 17 can be solved without auxiliaries, but they require a complex termination metric; interestingly, for 17—sorted list merge—CYPRESS generates an auxiliary anyway (see discussion in Sec. 5.3) All the other benchmarks require one or even two auxiliaries; for example, benchmark 14 flattens a rose tree into a list, and needs one auxiliary for flattening the list of children and another one for appending two lists (a flattened tree and a flattened list of children).

Finally, four of the synthesized programs feature *mutual recursion*. Two of them operate on rose trees, where mutual recursion is expected, while the other two—flattening a tree

Table 1. Benchmarks with complex recursion; *all of these* are out of reach for SuSLiK. We report the number of *Procedures* generated, total number *Stmt* of statement in those procedures, the ratio *Code/Spec* of code to specification (in AST nodes), and the synthesis *Time* (in seconds).

Group	Id	Description	Proc	Stmt	Code/Spec	Time
Singly Linked List	1	deallocate two	2	9	6.2x	0.3
	2	append three	2	14	2.3x	1.2
	3	non-destructive append	2	21	3.0x	5.2
	4	union	2	24	5.9x	9.6
	5	intersection ¹	3	33	7.3x	95.6
	6	difference ¹	2	22	5.5x	8.1
	7	deduplicate ¹	2	23	7.8x	6.2
List of Lists	8	deallocate	2	11	10.7x	0.3
	9	flatten ¹	2	19	4.8x	0.8
Binary Tree	10	deallocate two	1	16	11.8x	0.3
	11	flatten	2	24	7.4x	1.5
	12	flatten to dll in place	2 [†]	15	9.6x	2.7
Rose Tree	13	deallocate	2 [†]	9	12.0x	0.3
	14	flatten	3 [†]	25	8.0x	12.6
Sorted list	15	reverse ¹	2	11	3.3x	1.1
	16	sort ¹	2	12	3.6x	1.9
	17	merge ²	2	23	2.2x	33.6
BST	18	from list ¹	2	27	5.0x	11.5
	19	to sorted list ¹	2 [†]	35	6.4x	10.2

¹ From [13] ² From [31] [†] Mutually-recursive

into doubly-linked list (12) and flattening a binary search tree (BST) into a sorted list (19)—came as a surprise.

Comparison with other tools. CYPRESS was able to synthesize *all eight* selected benchmarks from [13]. Unlike their tool, CYPRESS does not use function templates as hints, and also targets pointer-manipulating programs instead of functional programs, which is arguably a harder problem. To the best of our knowledge, CYPRESS generates code with similar recursive structure to their tool: for example, list reversal, deduplication, and sorting all have the structure of two nested right-folds (and hence quadratic complexity); one exception is list intersection, discussed below, where CYPRESS generates an overly complex solution.

CYPRESS can also handle *all 22* original SuSLiK benchmarks (which are all part of the *simple* set). Finally, CYPRESS was also able to solve *all 11* recursive benchmarks from IMPSYNT [31] (one of which is in the *complex* set and the rest are in the *simple* set). Unlike IMPSYNT, we do not require the user to provide program sketches, and additionally, our synthesis times on a commodity laptop are at least an order of magnitude faster than the run times reported in [31] on a server with 10 cores and 96GB of RAM.

On the other hand, SuSLiK cannot handle any of the benchmarks in the *complex* set—because of complex recursion—and also fails on five benchmarks in the *simple* set. All five failures are due to restrictions on predicate unfolding, which SuSLiK had to impose due to its ad-hoc termination checking mechanism, which in CYPRESS has been replaced with the more powerful cardinality-based mechanism from Sec. 3.3.

Table 2. Comparison with SuSLiK on benchmarks with simple recursion. We report the number of statements *Stmt* and synthesis *Time* in seconds for both CYPRESS and SuSLiK. Each benchmark generates a single procedure.

Group	Id	Description	Stmt	Code/Spec	Time	
					CYPRESS	SuSLiK
Integers	20	swap two ³	4	1.0x	0.2	< 0.1
	21	min of two ^{3,4}	3	1.1x	1.5	0.4
Singly Linked List	22	length ^{2,3}	6	1.2x	1.1	1.1
	23	max ^{2,3}	7	1.9x	0.7	0.7
	24	min ^{2,3}	7	1.9x	0.6	0.7
	25	singleton ^{3,4}	4	0.9x	0.3	< 0.1
	26	dispose ³	4	5.5x	0.2	< 0.1
	27	initialize ³	4	1.6x	0.6	0.1
	28	copy ^{3,5}	11	2.7x	0.8	0.3
	29	append ^{3,5}	6	1.1x	0.5	0.4
	30	delete ^{3,5}	12	2.6x	1.6	0.4
	Sorted list	31	prepend ^{2,3}	4	0.5x	0.3
32		insert ^{2,3}	25	2.6x	4.4	5.2
33		insertion sort ^{2,3}	7	1.0x	1.2	1.4
Tree	34	size ³	9	2.5x	0.7	0.3
	35	dispose ³	6	8.0x	0.2	< 0.1
	36	copy ³	16	3.8x	2.8	0.7
	37	flatten w/append ³	19	5.4x	0.4	0.7
	38	flatten w/acc ³	12	2.1x	0.7	0.7
BST	39	insert ^{2,3}	19	1.9x	9.8	36.9
	40	rotate left ^{2,3}	5	0.2x	6.2	23.9
	41	rotate right ^{2,3}	5	0.2x	4.8	9.1
	42	delete root ²	29	1.7x	1304.3	-
Doubly Linked List	43	copy	22	4.3x	7.3	-
	44	append ⁵	10	1.6x	2.3	-
	45	delete ⁵	19	3.7x	4.7	-
	46	single to double	21	5.5x	1.3	-

²From [31] ³From SuSLiK [29] ⁴From [22] ⁵From [30]

Although in principle IMPSYNT is capable of solving complex benchmarks that require recursive auxiliaries using nested loops instead, none of the results reported in [31] contain nested loops, and in any case, IMPSYNT relies on program sketches. To our knowledge, no existing synthesizer (for either functional or heap-manipulating programs) is able to generate mutually-recursive programs.

5.2.2 Efficiency. Our experiments show that CYPRESS is efficient in synthesizing a variety of programs: all 19 complex benchmarks were synthesized within two minutes, and all but two of them take less than fifteen seconds. Our comparison of CYPRESS with SuSLiK on the simple benchmarks demonstrates that despite searching a larger space of programs, CYPRESS remains efficient. It is slightly slower on easy benchmarks: of the 18 benchmarks that take less than five seconds for SuSLiK, the average time is 0.5 seconds for SuSLiK compared to 0.8 seconds for CYPRESS; the remaining four hard benchmarks is where CYPRESS’s new search strategy pays off: the average time for those four benchmarks is 18.8 seconds for SuSLiK and only 6.3 seconds for CYPRESS.

5.2.3 Utility. To quantify synthesis utility, we measured the ratio between the size of synthesized code and the input specification (taken as AST sizes). For all benchmarks from

the complex set, the generated code is larger than the specification by at least 2.2x (sorted list merge) and at most 12x (deallocate rose tree). This confirms our intuition that for simple, boilerplate data structure manipulations, like deallocation and copying, synthesis offers a good trade-off, since their specifications are very simple, while the code can be quite tricky. We only include pre- and postconditions in the specification size and omit predicate definitions, since those are reused between benchmarks. Anecdotally, the new benchmarks created for this paper were quick and easy to write: most predicate definitions were either reused from `SUSLIK` or were standard SL predicates. In both cases, the pre- and postconditions were very straightforward. Note that the code-to-spec ratios for the simple benchmarks are lower (between 0.2x and 8.0x), which serves as evidence that deductive synthesis for heap-manipulating programs delivers larger pay-off for complex traversals.

5.3 Notable Benchmarks

Merge. One benchmark where `CYPRESS` surprised us is merging two sorted lists (benchmark 17). We initially thought that this problem required an extension to the tool, because the traditional recursive implementation—without auxiliaries—has to unfold both input lists in order to compare their heads, but then fold one of the lists back again to pass it to the recursive call; this “folding back” is something that `CYPRESS` does not explicitly support. To our surprise, `CYPRESS` was instead able to solve this problem by inventing an auxiliary that merges two lists *one of which is non-empty*. This implementation is even slightly more efficient, since it eliminates a redundant emptiness check for one of the lists.

Sort. An important difference between the functional setting of [13] and our setting is that SL specifications give the the user more fine-grained control over the relationship between input and output data structures in memory. For sorting (benchmark 16), we took the liberty of requiring the sort to be *in-place*, by using the following specification:

$$\{\text{sll}(x, n, l, h)\} \rightsquigarrow \{\text{srtl}(x, n, l, h)\}$$

Here `sll` is a list rooted at `x` with length `n` and lower and upper bounds on the elements `l` and `h`, and `srtl` is a sorted list with the same parameters (notably, rooted at the same address `x`). Such in-place sorting is not expressible using refinement types from [13]. Given this specification—which contains no insight as to which algorithm to use to sort the list—`CYPRESS` synthesizes a rather peculiar version of in-place insertion sort: while traditionally insertion sort on linked lists performs insertion by switching `next` pointers, `CYPRESS` chose instead to do it by swapping elements in out-of-order list cells (similarly to a typical insertion sort on an array).

Intersection. Perhaps the most curious case is benchmark 5, which computes the intersection of two sets represented as linked lists with unique elements (denoted by the inductive

predicate `ul`). The simplest specification for intersection is:

$$\{r \mapsto x * \text{ul}(x, s_1) * \text{ul}(y, s_2)\} \rightsquigarrow \{r \mapsto z * \text{ul}(z, s_1 \cap s_2)\}$$

With this specification, however, `CYPRESS` fails to find a solution due to a limitation on the class of auxiliaries it can generate. Unlike the original functional setting, there is no simple, fold-like solution for this specification, because it is *destructive*: once it computes the intersection `z'` of `y` and the tail of `x`, the list `y` is lost, and we cannot decide whether the head of `x` should be added to `z'`. A more sophisticated solution would require an auxiliary that tests membership, which is beyond the capabilities of cyclic synthesis.

To circumvent this issue, we added `*ul(y, s2)` to the postcondition to preserve the input list `y`. Even with this change, however, `CYPRESS` fails to generate the simple solution with a single auxiliary (one that searches for the head of `x` in `y`, and if found, prepends it to `z'`), because doing so requires *weakening* of the auxiliary’s pure precondition, which cyclic synthesis currently does not support. Surprisingly, `CYPRESS` finds another solution, which uses a second auxiliary to *append* the head of `x` to `z'` instead of *prepending* it. This solution is so unusual that we worried we found a soundness issue, until we checked it with an external program verifier: the solution is indeed correct (albeit inefficient), and both of the inferred auxiliary specifications are inductive.

5.4 Limitations

When does `CYPRESS` fail? The main limitation of `SSL∪` is that it can only derive a certain class of auxiliary functions. Intuitively, it can only extract auxiliary specifications from the goals in the main derivation, and cannot invent them “out of thin air”, so, for example, the auxiliary cannot take extra parameters or return a new data structure, which the main specification does not mention. For this reason, `CYPRESS` cannot synthesize linear-time implementations of list reversal or tree flattening, which would require conjecturing an extra accumulator parameter. As explained in Sec. 5.3, `SSL∪` also lacks support for generalizing the pure part of the auxiliary specification, which prevents it from deriving simpler version of set intersection. It is difficult to precisely characterize the class of synthesis problems that cyclic synthesis can and cannot solve, because, as illustrated above, it often finds alternative solutions that circumvent the limitations. Such characterization is an interesting direction for future work; another important and challenging direction is extending `SSL∪` with additional rules to support a wider class of programs, while continuing to strike a careful balance between expressiveness and tractability of proof search.

Quality of solutions. A `CYPRESS` solution might have suboptimal performance *even if* a more efficient program is derivable via `SSL∪`. The core issue is that `SSL∪` has no means of analyzing program efficiency, and hence no reason to prefer a more efficient solution. A promising approach for producing efficient heap-manipulating would be to adopt a

flavor of Separation Logic with *time credits* [15], combining it with techniques for resource-aware synthesis for pure functional programs [20].

We also noticed that the division of code between the procedures in CYPRESS solutions is sometimes unnatural: for example, the tree flattening program in Fig. 5 deallocates tree nodes inside the auxiliary `append` function instead of the main function; in terms of abstraction boundaries, it would be better if `append` were only concerned with appending lists and had no knowledge of tree nodes. The division of code depends on the order in which CYPRESS tries applying synthesis rules (in this case, the relative order of `CALL` and `FREE`); currently the order is optimized for efficiency of proof search rather than for optimal abstraction boundaries. We leave the investigation of improving program quality to future work.

Loop support. SSL_{\cup} currently has no support for loops, which are often a more natural and efficient alternative to recursion in imperative programs. There are several existing techniques for deductive *verification* of loops using cyclic proofs [7, 43]; hence we believe the general cyclic synthesis technique could be extended to also handle them. To keep proof search tractable, however, we might not want to support both recursion and loops in our target language; a better idea might be to synthesize *tail-recursive* programs and then translate them into loops using standard techniques. The challenge, however, is that most tail-recursive programs require invented accumulator parameters (equivalently: most loops require temporary variables). Hence adding support for loops goes hand-in-hand with the extending the logic to support a wider class of auxiliaries.

6 Related Work

Deductive program synthesis. Our work on SSL_{\cup} extends a line of research on synthesizing programs from logical specifications [22, 31, 38, 40, 44], and, in particular, on *deductive synthesis* [12, 19, 24, 28], which implements a search in the space of *proofs* of program correctness (rather than in the space of programs). Our work is the first to combine deductive synthesis with cyclic proof to generate provably correct and terminating programs with complex recursion patterns. Our technical contributions build primarily on the work by Polikarpova and Sergey [29], extending their logic SSL and the tool SuSLIK with ideas from cyclic proofs. The best-first search algorithm of CYPRESS is inspired by the LEON system for deductive program verification, synthesis, and repair [18, 19], but tailored to Separation Logic.

Our work advances the state of the art for deductive synthesis tools with regard to establishing termination of synthesized programs. Specifically, it enables automated derivation of a termination argument along with the program being synthesized without being subject to restrictions of previous approaches: (1) requiring the user to provide an explicit termination measure [28]; (2) restricting the recursion to syntactic structural one [27, 29]; and (3) determining the

termination measure by instantiating one of the pre-defined recursion/looping schemas [19, 31].

Synthesis of auxiliary functions. Eguchi et al. [13] implement a technique for automatically synthesizing implementations of *pure functional programs* with recursive helper functions from refinement types [45]. Their approach infers specifications for recursive helper functions by trying a number of *predefined templates* for the “main” function, which exercise different flavours of recursion (structural folds or divide-and-conquer) on the input data structures. Our work is complementary in that it considers imperative heap-manipulating programs, and also takes a fundamentally different *proof-driven* (rather than *template-driven*) approach to identify recursion patterns. By doing so, our technique removes the need to conjecture recursion principles upfront, yet it is capable of discovering many of those automatically, including non-trivial ones, such as traversals of a rose tree.

Several other techniques cannot synthesize recursive auxiliaries directly, but can synthesize nested folds [14] or nested loops [31, 38, 40]. Out of these techniques, λ^2 [14] is restricted to a set of predefined templates, and also provides no correctness guarantees beyond a finite set of input-output examples. Sketching-based approaches [31, 38, 40] are very flexible, but require the programmer to provide an extensive program sketch in order to make synthesis with nested loops tractable.

Cyclic proofs. The techniques and metatheory of cyclic proofs originates in the logic and proof theory community. Most related to our current work is the application of cyclic proof to reasoning about program correctness [7, 35, 43], as well as to proving pure entailments of Separation Logic with inductively defined predicates [6, 41, 42]. In particular, our use of cardinalities with Separation Logic coincides with the approach used by Rowe and Brotherston [35].

7 Conclusion

We have demonstrated that cyclic proofs, already known as an automation-friendly method for reasoning about program termination, can also be an effective tool for automatically synthesizing programs with recursive auxiliary functions. By investigating this synergy between verification and synthesis, we have once again witnessed that the ideas developed by the logic community for scalable reasoning about heap-manipulating procedures can be instrumental in practical synthesis of correct-by-construction programs.

Acknowledgments

The authors would like to thank the anonymous reviewers and our shepherd, Kwangkeun Yi, for their valuable feedback on earlier drafts of this paper. This work was supported by the National Science Foundation under Grant No. 1911149, by Singapore MoE Tier 1 Grant No. IG18-SG102, by the Israeli Science Foundation (ISF) Grants No. 243/19 and 2740/19, and by the United States-Israel Binational Science Foundation (BSF) Grant No. 2018675.

References

- [1] Aws Albarghouthi, Sumit Gulwani, and Zachary Kincaid. 2013. Recursive Program Synthesis. In *CAV (Part II) (LNCS, Vol. 9780)*. Springer, 934–950. https://doi.org/10.1007/978-3-642-39799-8_67
- [2] Rajeev Alur, Rastislav Bodík, Garvit Juniwal, Milo M. K. Martin, Mukund Raghothaman, Sanjit A. Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and Abhishek Udupa. 2013. Syntax-guided synthesis. In *FMCAD*. IEEE, 1–8. <https://doi.org/10.1145/3434311>
- [3] David Baelde, Amina Doumane, and Alexis Saurin. 2016. Infinitary Proof Theory: the Multiplicative Additive Case. In *CSL (LIPIcs, Vol. 62)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 42:1–42:17. <https://doi.org/10.4230/LIPIcs.CSL.2016.42>
- [4] Josh Berdine and Nikolaj Bjørner. 2014. Computing All Implied Equalities via SMT-Based Partition Refinement. In *IJCAR*. 168–183. https://doi.org/10.1007/978-3-319-08587-6_12
- [5] Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. 2005. Symbolic Execution with Separation Logic. In *APLAS (LNCS, Vol. 3780)*. Springer, 52–68. https://doi.org/10.1007/11575467_5
- [6] James Brotherston. 2007. Formalised Inductive Reasoning in the Logic of Bunched Implications. In *SAS (LNCS, Vol. 4634)*. Springer, 87–103. https://doi.org/10.1007/978-3-540-74061-2_6
- [7] James Brotherston, Richard Bornat, and Cristiano Calcagno. 2008. Cyclic proofs of program termination in separation logic. In *POPL*. ACM, 101–112. <https://doi.org/10.1145/1328897.1328453>
- [8] James Brotherston, Nikos Gorogiannis, and Rasmus Lerchedahl Petersen. 2012. A Generic Cyclic Theorem Prover. In *APLAS (LNCS, Vol. 7705)*. Springer, 350–367. https://doi.org/10.1007/978-3-642-35182-2_25
- [9] James Brotherston and Alex Simpson. 2011. Sequent Calculi for Induction and Infinite Descent. *J. Log. Comput.* 21, 6 (2011), 1177–1216. <https://doi.org/10.1093/logcom/exq052>
- [10] Coq Development Team. 2020. *The Coq Proof Assistant Reference Manual - Version 8.11*. Available at <http://coq.inria.fr/>.
- [11] Anupam Das and Damien Pous. 2018. Non-Wellfounded Proof Theory For (Kleene+Action)(Algebras+Lattices). In *CSL (LIPIcs, Vol. 119)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 19:1–19:18. <https://doi.org/10.4230/LIPIcs.CSL.2018.19>
- [12] Benjamin Delaware, Clément Pit-Claudel, Jason Gross, and Adam Chlipala. 2015. Fiat: Deductive Synthesis of Abstract Data Types in a Proof Assistant. In *POPL*. ACM, 689–700. <https://doi.org/10.1145/2775051.2677006>
- [13] Shingo Eguchi, Naoki Kobayashi, and Takeshi Tsukada. 2018. Automated Synthesis of Functional Programs with Auxiliary Functions. In *APLAS (LNCS, Vol. 11275)*. Springer, 223–241. https://doi.org/10.1007/978-3-030-02768-1_13
- [14] John K. Feser, Swarat Chaudhuri, and Isil Dillig. 2015. Synthesizing data structure transformations from input-output examples. In *PLDI*. ACM, 229–239. <https://doi.org/10.1145/2737924.2737977>
- [15] Armaël Guéneau, Arthur Charguéraud, and François Pottier. 2018. A Fistful of Dollars: Formalizing Asymptotic Complexity Claims via Deductive Program Verification. In *ESOP (LNCS, Vol. 10801)*. Springer, 533–560. https://doi.org/10.1007/978-3-319-89884-1_19
- [16] Shachar Itzhaky, Hila Peleg, Nadia Polikarpova, Reuben N. S. Rowe, and Ilya Sergey. 2021. *Cyclic Program Synthesis*. Technical Report. University of California, San Diego. <https://cseweb.ucsd.edu/~npolikarpova/publications/cypress-extended.pdf>
- [17] Shachar Itzhaky, Hila Peleg, Nadia Polikarpova, Reuben N. S. Rowe, and Ilya Sergey. 2021. *Cypress (PLDI 2021 Artifact): Code and Benchmarks*. <https://doi.org/10.5281/zenodo.4679743>
- [18] Etienne Kneuss, Manos Koukoutos, and Viktor Kuncak. 2015. Deductive Program Repair. In *CAV (LNCS, Vol. 9207)*. Springer, 217–233. https://doi.org/10.1007/978-3-319-21668-3_13
- [19] Etienne Kneuss, Ivan Kuraj, Viktor Kuncak, and Philippe Suter. 2013. Synthesis modulo recursive functions. In *OOPSLA*. ACM, 407–426. <https://doi.org/10.1145/2544173.2509555>
- [20] Tristan Knoth, Di Wang, Nadia Polikarpova, and Jan Hoffmann. 2019. Resource-guided program synthesis. In *PLDI*. ACM, 253–268. <https://doi.org/10.1145/3314221.3314602>
- [21] Viktor Kuncak, Mikael Mayer, Ruzica Piskac, and Philippe Suter. 2010. Complete Functional Synthesis. In *PLDI*. 316–329. <https://doi.org/10.1145/1809028.1806632>
- [22] K. Rustan M. Leino and Aleksandar Milicevic. 2012. Program Extrapolation with Jennisys. In *OOPSLA*. ACM, 411–430. <https://doi.org/10.1145/2398857.2384646>
- [23] Grant Malcolm. 1990. Data Structures and Program Transformation. *Sci. Comput. Program.* 14, 2-3 (1990), 255–279. [https://doi.org/10.1016/0167-6423\(90\)90023-7](https://doi.org/10.1016/0167-6423(90)90023-7)
- [24] Zohar Manna and Richard J. Waldinger. 1980. A Deductive Approach to Program Synthesis. *ACM Trans. Program. Lang. Syst.* 2, 1 (1980), 90–121. <https://doi.org/10.1145/357084.357090>
- [25] Damian Niwiński and Igor Walukiewicz. 1996. Games for the μ -Calculus. *Theor. Comput. Sci.* 163, 1&2 (1996), 99–116. [https://doi.org/10.1016/0304-3975\(95\)00136-0](https://doi.org/10.1016/0304-3975(95)00136-0)
- [26] Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. 2001. Local Reasoning about Programs that Alter Data Structures. In *CSL (LNCS, Vol. 2142)*. Springer, 1–19. https://doi.org/10.1007/3-540-44802-0_1
- [27] Peter-Michael Osera and Steve Zdancewic. 2015. Type-and-example-directed program synthesis. In *PLDI*. ACM, 619–630. <https://doi.org/10.1145/2737924.2738007>
- [28] Nadia Polikarpova, Ivan Kuraj, and Armando Solar-Lezama. 2016. Program synthesis from polymorphic refinement types. In *PLDI*. ACM, 522–538. <https://doi.org/10.1145/2980983.2908093>
- [29] Nadia Polikarpova and Ilya Sergey. 2019. Structuring the Synthesis of Heap-Manipulating Programs. *Proc. ACM Program. Lang.* 3, POPL (2019), 72:1–72:30. <https://doi.org/10.1145/3290385>
- [30] Xiaokang Qiu, Pranav Garg, Andrei Stefanescu, and Parthasarathy Madhusudan. 2013. Natural proofs for structure, data, and separation. In *PLDI*. ACM, 231–242. <https://doi.org/10.1145/2499370.2462169>
- [31] Xiaokang Qiu and Armando Solar-Lezama. 2017. Natural synthesis of provably-correct data-structure manipulations. *PACMPL* 1, OOPSLA (2017), 65:1–65:28. <https://doi.org/10.1145/3133889>
- [32] Andrew Reynolds, Morgan Deters, Viktor Kuncak, Cesare Tinelli, and Clark W. Barrett. 2015. Counterexample-Guided Quantifier Instantiation for Synthesis in SMT. In *CAV (Part II) (LNCS, Vol. 9207)*. Springer, 198–216. https://doi.org/10.1007/978-3-319-21668-3_12
- [33] Andrew Reynolds, Viktor Kuncak, Cesare Tinelli, Clark W. Barrett, and Morgan Deters. 2019. Refutation-based synthesis in SMT. *Formal Methods Syst. Des.* 55, 2 (2019), 73–102. <https://doi.org/10.1007/s10703-017-0270-2>
- [34] John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *LICS*. IEEE Computer Society, 55–74. <https://doi.org/10.1109/LICS.2002.1029817>
- [35] Reuben N. S. Rowe and James Brotherston. 2017. Automatic cyclic termination proofs for recursive procedures in separation logic. In *CPP*. ACM, 53–65. <https://doi.org/10.1145/3018610.3018623>
- [36] Luigi Santocanale. 2002. A Calculus of Circular Proofs and Its Categorical Semantics. In *FoSSaCS (LNCS, Vol. 2303)*. Springer, 357–371. https://doi.org/10.1007/3-540-45931-6_25
- [37] Alex Simpson. 2017. Cyclic Arithmetic Is Equivalent to Peano Arithmetic. In *FoSSaCS (LNCS, Vol. 10203)*. Springer, 283–300. https://doi.org/10.1007/978-3-662-54458-7_17
- [38] Armando Solar-Lezama. 2013. Program sketching. *STTT* 15, 5-6 (2013), 475–495. <https://doi.org/10.1007/s10009-012-0249-7>
- [39] Christoph Sprenger and Mads Dam. 2003. On the Structure of Inductive Reasoning: Circular and Tree-Shaped Proofs in the μ -Calculus. In *FoSSaCS (LNCS, Vol. 2620)*. Springer, 425–440. https://doi.org/10.1007/3-540-36576-1_27

- [40] Saurabh Srivastava, Sumit Gulwani, and Jeffrey S. Foster. 2010. From program verification to program synthesis. In *POPL*. ACM, 313–326. <https://doi.org/10.1145/1706299.1706337>
- [41] Quang-Trung Ta, Ton Chanh Le, Siau-Cheng Khoo, and Wei-Ngan Chin. 2016. Automated Mutual Explicit Induction Proof in Separation Logic. In *FM (LNCS, Vol. 9995)*, John S. Fitzgerald, Constance L. Heitmeyer, Stefania Gnesi, and Anna Philippou (Eds.). Springer, 659–676.
- [42] Quang-Trung Ta, Ton Chanh Le, Siau-Cheng Khoo, and Wei-Ngan Chin. 2018. Automated lemma synthesis in symbolic-heap separation logic. *Proc. ACM Program. Lang.* 2, POPL (2018), 9:1–9:29. <https://doi.org/10.1145/3158097>
- [43] Gadi Tellez and James Brotherston. 2020. Automatically Verifying Temporal Properties of Pointer Programs with Cyclic Proof. *J. Autom. Reasoning* 64, 3 (2020), 555–578. https://doi.org/10.1007/978-3-319-63046-5_30
- [44] Emina Torlak and Rastislav Bodík. 2014. A lightweight symbolic virtual machine for solver-aided host languages. In *PLDI*. ACM, 530–541. <https://doi.org/10.1145/2666356.2594340>
- [45] Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon L. Peyton Jones. 2014. Refinement types for Haskell. In *ICFP*. ACM, 269–282. <https://doi.org/10.1145/2628136.2628161>