



Proving Termination of C Programs with Lists

Jera Hensel^(✉) and Jürgen Giesl^(✉)

LuFG Informatik 2, RWTH Aachen University, Aachen, Germany
{hensel,giesl}@informatik.rwth-aachen.de

Abstract. There are many techniques and tools to prove termination of C programs, but up to now these tools were not very powerful for fully automated termination proofs of programs whose termination depends on recursive data structures like lists. We present the first approach that extends powerful techniques for termination analysis of C programs (with memory allocation and explicit pointer arithmetic) to lists.

1 Introduction

In [11, 16, 17, 25], we introduced an approach for automatic termination analysis of C that also handles programs whose termination relies on the relation between allocated memory addresses and the data stored at such addresses. This approach is implemented in our tool AProVE [14]. Instead of analyzing C directly, AProVE compiles the program to LLVM code using Clang [9]. Then it constructs a (finite) symbolic execution graph (SEG) such that every program run corresponds to a path through the SEG. AProVE proves memory safety during the construction of the SEG to ensure absence of undefined behavior (which would also allow non-termination). Afterwards, the SEG is transformed into an integer transition system (ITS) such that all paths through the SEG (and hence, the C program) are terminating if the ITS is terminating. To analyze termination of the ITS, AProVE applies standard techniques and calls the tools T2 [7] and LoAT [12, 13] to detect non-termination of ITSs. However, like other termination tools for C, up to now AProVE supported dynamic data structures only in a very restricted way.

In this paper, we introduce a novel technique to analyze C programs on lists. In the program on the right, `nondet_uint` returns a random unsigned integer. The `for` loop creates a list of `n` random numbers if `n > 0`. The `while` loop traverses this list via pointer arithmetic: Starting with `tail`, it computes the address of the `next` field of the

```
struct list {
    unsigned int value;
    struct list* next;
};

int main() {
    // initialize length
    unsigned int n = nondet_uint();
    // initialize list of length n
    struct list* tail = NULL;
    struct list* curr;
    for (unsigned int k = 0; k < n; k++) {
        curr = malloc(sizeof(struct list));
        curr->value = nondet_uint();
        curr->next = tail;
        tail = curr;
    }
    // traverse list
    struct list* ptr = tail;
    while(ptr != NULL) {
        ptr = *((struct list**)((void*)ptr +
            offsetof(struct list, next)));
    }
```

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 235950644 (Project GI 274/6-2).

© The Author(s) 2023

B. Pientka and C. Tinelli (Eds.): CADE 2023, LNAI 14132, pp. 266–285, 2023.

https://doi.org/10.1007/978-3-031-38499-8_16

current element by adding the offset of the `next` field within a `list` to the address of the current `list` and dereferencing the computed address (i.e., the content of the `next` field). This is done by `offsetof`, defined in the C library `stddef.h`.¹ Since the list is acyclic and the `next` pointer of its last element is the null pointer, list traversal always terminates. Of course, the `while` loop could also traverse the list via `ptr = ptr->next`, but in C, memory accesses can be combined with pointer arithmetic. This example contains both the access via `curr->next` (when initializing the list) and pointer arithmetic (when traversing the list).

We present a new general technique to infer *list invariants* via symbolic execution, which express all properties that are crucial for memory safety and termination. In our example, the list invariant contains the information that dereferencing the `next` pointer in the `while` loop is safe and that one finally reaches the null pointer. In general, our novel list invariants allow us to abstract from detailed information about lists (e.g., about their intermediate elements) such that abstract states with “similar” lists can be merged and generalized during the symbolic execution in order to obtain finite SEGs. At the same time, list invariants express enough information about the lists (e.g., their length, their start address, etc.) such that memory safety and termination can still be proved.

We define the abstract states used for symbolic execution in Sect. 2. In Sect. 3, after recapitulating the construction of SEGs, we adapt our techniques for merging and generalizing states from [25] to infer list invariants. Moreover, we adapt those rules for symbolic execution that are affected by introducing list invariants. Section 4 discusses the generation of ITSs and the soundness of our approach. Section 5 gives an overview on related work. Moreover, we evaluate the implementation of our approach in the tool AProVE using benchmark sets from *SV-COMP* [3] and the *Termination Competition* [15]. All proofs can be found in [18].

Limitations. To ease the presentation, in this paper we treat integer types as unbounded. Moreover, we assume that a program consists of a single non-recursive function and that values may be stored at any address. Our approach can also deal with bitvectors, data alignments, and programs with arbitrary many (possibly recursive) functions, see [11, 16, 25] for details. However, so far only lists without sharing can be handled by our new technique. Extending it to more general recursive data structures is one of the main challenges for future work.

2 Abstract States for Symbolic Execution

The LLVM code for the `for` loop is given on the next page. It is equivalent to the code produced by Clang without optimizations on a 64-bit computer. We explain it in detail in Sect. 3. To ease readability, we omitted instructions and keywords that are irrelevant for our presentation, renamed variables, and wrote `list`

¹ Note that `ptr + n` increases `ptr` by `n` times the size of the type `*ptr`. As we want to increase `ptr` by a number of bytes and `ptr` is not an `i8` pointer, we first cast `ptr` to `void*`. Then `((void*)ptr + offsetof(struct list, next))` contains the `next` pointer, so we cast our computed address to `struct list**` before dereferencing it.

instead of `struct.list`. Moreover, we gave the C instructions (in gray) before the corresponding LLVM code. The code consists of several *basic blocks* including `cmpF` and `bodyF` (corresponding to the loop comparison and body).

We now recapitulate the *abstract states* of [25] used for symbolic execution and extend them by a component *LI* for list invariants, i.e., they have the form $((b, i), LV, AL, PT, LI, KB)$. The first component is a *program position* (b, i) , indicating that instruction i of block b is executed next. $Pos \subseteq (Blks \times \mathbb{N})$ is the set of all program positions, and $Blks$ are all basic blocks.

The second component is a partial injective function $LV: \mathcal{V}_{\mathcal{P}} \rightarrow \mathcal{V}_{sym}$, which maps *local program variables* $\mathcal{V}_{\mathcal{P}}$ of the program $\mathcal{P}$ to an infinite set $\mathcal{V}_{sym}$ of symbolic variables with $\mathcal{V}_{sym} \cap \mathcal{V}_{\mathcal{P}} = \emptyset$. We identify LV with the set of equations $\{x = LV(x) \mid x \in domain(LV)\}$ and we often extend LV to a function from $\mathcal{V}_{\mathcal{P}} \uplus \mathbb{N}$ to $\mathcal{V}_{sym} \uplus \mathbb{N}$ by defining $LV(n) = n$ for all $n \in \mathbb{N}$.

The third component of each state is a set AL of (bitwise) allocations $\llbracket v_1, v_2 \rrbracket$ with $v_1, v_2 \in \mathcal{V}_{sym}$, which indicate that $v_1 \leq v_2$ and that all addresses between v_1 and v_2 have been allocated. We require any two entries $\llbracket v_1, v_2 \rrbracket$ and $\llbracket w_1, w_2 \rrbracket$ from AL with $v_1 \neq w_1$ or $v_2 \neq w_2$ to be disjoint.

The fourth and fifth components PT and LI model the memory contents. PT contains “points-to” entries of the form $v_1 \hookrightarrow_{ty} v_2$ where $v_1, v_2 \in \mathcal{V}_{sym}$ and ty is an LLVM type, meaning that the address v_1 of type ty points to v_2 . In contrast, the set LI of *list invariants* (which is new compared to [25]) does not describe pointwise memory contents but contains invariants $v_{ad} \xrightarrow{v_\ell}_{ty} [(off_i : ty_i : v_i.. \hat{v}_i)]_{i=1}^n$ where $n \in \mathbb{N}_{>0}$, $v_{ad}, v_\ell, v_i, \hat{v}_i \in \mathcal{V}_{sym}$, $off_i \in \mathbb{N}$ for all $1 \leq i \leq n$, ty and ty_i are LLVM types for all $1 \leq i \leq n$, and there is exactly one “recursive field” $1 \leq j \leq n$ such that $ty_j = ty*$.² Such an invariant represents a **struct** ty with n fields that corresponds to a recursively defined list of length v_ℓ . Here, v_{ad} points to the first list element, the i -th field starts at address $v_{ad} + off_i$ (i.e., with offset off_i)³ and has type ty_i , and the values of the i -th fields of the first and last list element are v_i and $\hat{v}_i$, respectively. For example, the following list invariant (1) represents all lists of length x_ℓ and type `list` whose elements store a 32-bit integer in their first field and the pointer to the next element in their second field

```
list = type { i32, list* }

define i32 @main() { ...
cmpF:
  k < n
  0: k = load i32, i32* k_ad
  1: kltn = icmp ult i32 k, n
  2: br i1 kltn, label bodyF, label initPtr
bodyF:
  curr = malloc(sizeof(struct list));
  0: mem = call i8* @malloc(i64 16)
  1: curr = bitcast i8* mem to list*
  curr->value = nondet_uint();
  2: nondet = call i32 @nondet_uint()
  3: curr_val = getelementptr list,
               list* curr, i32 0, i32 0
  4: store i32 nondet, i32* curr_val
  curr->next = tail;
  5: tail = load list*, list** tail_ptr
  6: curr_next = getelementptr list,
               list* curr, i32 0, i32 1
  7: store list* tail, list** curr_next
  tail = curr;
  8: store list* curr, list** tail_ptr
  k++;
  9: kinc = add i32 k, 1
  10: store i32 kinc, i32* k_ad
  11: br label cmpF
  ...
}
```

² Soundness of our approach is not affected if there are other recursive fields, but our symbolic execution technique for list traversal on list invariants in Sect. 3.2.2 can only be applied if the traversal is done along field j .

³ The field offsets can be computed using the data layout string in the LLVM program.

with offset 8. The first list element starts at address x_{mem} , the second starts at address x_{next} , and the last element contains the null pointer. Moreover, the first element stores the integer value x_{nd} and the last list element stores the integer $\hat{x}_{\text{nd}}$.

$$x_{\text{mem}} \xrightarrow{x_\ell}_{\text{list}} [(0 : \text{i32} : x_{\text{nd}}.. \hat{x}_{\text{nd}}), (8 : \text{list*} : x_{\text{next}}..0)] \quad (1)$$

For example, this invariant represents the list with the allocation $\llbracket x_{\text{mem}}, x_{\text{mem}}+15 \rrbracket$, where the first four bytes store the integer 5 and the last eight bytes store the pointer x_{next} , and the allocation $\llbracket x_{\text{next}}, x_{\text{next}}+15 \rrbracket$, where the first four bytes store the integer 2 and the last eight bytes store the null pointer (i.e., the address 0). Here, we have $x_\ell = 2$. Section 3.2.2 will show that the expressiveness of our list invariants is indeed needed to prove termination of programs that traverse a list.

The last component of a state is a *knowledge base* KB of quantifier-free first-order formulas that express integer arithmetic properties of $\mathcal{V}_{\text{sym}}$. We identify *sets* of first-order formulas $\{\varphi_1, \dots, \varphi_m\}$ with their conjunction $\varphi_1 \wedge \dots \wedge \varphi_m$.

A special state *ERR* is reached if we cannot prove absence of undefined behavior (e.g., if memory safety might be violated by dereferencing the null pointer).

As an example, the following abstract state (2) represents concrete states at the beginning of the block `cmpF`, where the program variable `curr` is assigned the symbolic variable x_{mem} , the allocation $\llbracket x_{\text{k.ad}}, x_{\text{k.ad}}^{\text{end}} \rrbracket$ consisting of 4 bytes stores the value x_{kinc} , and x_{mem} points to the first element of a list of length x_ℓ (equal to x_{kinc}) that satisfies the list invariant (1). (This state will later be obtained during the symbolic execution, see State *O* in Fig. 3 in Sect. 3.1.)

$$\boxed{(\text{cmpF}, 0), \{\text{curr} = x_{\text{mem}}, \text{kinc} = x_{\text{kinc}}, \dots\}, \{\llbracket x_{\text{k.ad}}, x_{\text{k.ad}}^{\text{end}} \rrbracket, \dots\}, \{x_{\text{k.ad}} \xrightarrow{\text{i32}} x_{\text{kinc}}, \dots\}, \{x_{\text{mem}} \xrightarrow{x_\ell}_{\text{list}} [(0 : \text{i32} : x_{\text{nd}}.. \hat{x}_{\text{nd}}), (8 : \text{list*} : x_{\text{next}}..0)]\}, \{x_{\text{k.ad}}^{\text{end}} = x_{\text{k.ad}} + 3, x_\ell = x_{\text{kinc}}, \dots\}} \quad (2)$$

A state $s = (p, LV, AL, PT, LI, KB)$ is *represented by a formula* $\langle s \rangle$ which contains KB and encodes AL , PT , and LI in first-order logic. This allows us to use standard SMT solving for all reasoning during the construction of the SEG. Moreover, $\langle s \rangle$ is also used for the generation of the ITS afterwards. The encoding of AL and PT is as in [25], see [18]: $\langle s \rangle$ contains formulas which express that allocated addresses are positive, that allocations represent disjoint memory areas, that equal addresses point to equal values, and that addresses are different if they point to different values. For each element of LI , we add the following new formulas to $\langle s \rangle$ which express that the list length v_ℓ is ≥ 1 and the address v_{ad} of the first element is not null. If $v_\ell = 1$, then the values v_i and $\hat{v}_i$ of the fields in the first and the last element are equal. If $v_\ell \geq 2$, then the `next` pointer v_j in the first element must not be null. Finally, if there is a field whose values v_k and $\hat{v}_k$ differ in the first and the last element, then the length v_ℓ must be ≥ 2 .

$$\begin{aligned} & \{v_\ell \geq 1 \wedge v_{\text{ad}} \geq 1 \mid (v_{\text{ad}} \xrightarrow{v_\ell}_{\text{ty}} [(off_i : \text{ty}_i : v_i.. \hat{v}_i)]_{i=1}^n) \in LI\} \cup \\ & \{\bigwedge_{i=1}^n v_i = \hat{v}_i \mid (v_{\text{ad}} \xrightarrow{v_\ell}_{\text{ty}} [(off_i : \text{ty}_i : v_i.. \hat{v}_i)]_{i=1}^n) \in LI \text{ and } \models \langle s \rangle \Rightarrow v_\ell = 1\} \cup \\ & \{v_j \geq 1 \mid (v_{\text{ad}} \xrightarrow{v_\ell}_{\text{ty}} [(off_i : \text{ty}_i : v_i.. \hat{v}_i)]_{i=1}^n) \in LI \text{ with } \text{ty}_j = \text{ty*} \text{ and } \models \langle s \rangle \Rightarrow v_\ell \geq 2\} \cup \\ & \{v_\ell \geq 2 \mid (v_{\text{ad}} \xrightarrow{v_\ell}_{\text{ty}} [(off_i : \text{ty}_i : v_i.. \hat{v}_i)]_{i=1}^n) \in LI \text{ and } \exists k \in \mathbb{N}_{>0}, k \leq n, \text{ s.t. } \models \langle s \rangle \Rightarrow v_k \neq \hat{v}_k\} \end{aligned}$$

In *concrete* states c , all values of variables and memory contents are determined uniquely. To ease the formalization, we assume that all integers are




Fig. 1. SEG for the First Iteration of the for Loop

unsigned and refer to [16] for the general case. So for all $v \in \mathcal{V}_{sym}(c)$ (i.e., all $v \in \mathcal{V}_{sym}$ occurring in c) we have $\models \langle c \rangle \Rightarrow v = n$ for some $n \in \mathbb{N}$. Moreover, here PT only contains information about allocated addresses and $LI = \emptyset$ since the abstract knowledge in list invariants is unnecessary if all memory contents are known.

For instance, all concrete states $((\text{cmpF}, 0), LV, AL, PT, \emptyset, KB)$ represented by the state (2) contain ℓ allocations of 16 bytes for some $\ell \geq 1$, where in the first four bytes a 32-bit integer is stored and in the last eight bytes the address of the next allocation (or 0, in case of the last allocation) is stored.

See [18] for a formal definition to determine which concrete states are represented by a state s . To this end, as in [25] we define a *separation logic* formula $\langle s \rangle_{SL}$ which also encodes the knowledge contained in the memory components of states. To extend this formula to list invariants, we use a fragment similar to *quantitative separation logic* [4], extending conventional separation logic by list predicates. For any state s , we have $\models \langle s \rangle_{SL} \Rightarrow \langle s \rangle$, i.e., $\langle s \rangle$ is a weakened version of $\langle s \rangle_{SL}$ that we use for symbolic execution and the termination proof.

3 Symbolic Execution with List Invariants

We first recapitulate the construction of SEGs. Then, Sect. 3.1 extends the technique for *merging* and generalization of states from [25] to infer list invariants. Finally, we adapt the rules for symbolic execution to list invariants in Sect. 3.2.

Our symbolic execution starts with a state A at the first instruction of the first block (called **entry** in our example). Figure 1 shows the first iteration of the

for loop. Dotted arrows indicate that we omitted some symbolic execution steps. For every state, we perform symbolic execution by applying the corresponding inference rule as in [25] to compute its successor state(s) and repeat this until all paths end in return states. We call an SEG with this property *complete*.

As an example, we recapitulate the inference rule for the `load` instruction in the case where a value is loaded from allocated and initialized memory. It loads the value of type `ty` that is stored at the address `ad` to the program variable `x`. Let $\text{size}(\text{ty})$ denote the size of `ty` in bytes for any LLVM type `ty`. If we can prove that there is an allocation $\llbracket v_1, v_2 \rrbracket$ containing all addresses $LV(\text{ad}), \dots, LV(\text{ad}) + \text{size}(\text{ty}) - 1$ and there exists an entry $(w_1 \hookrightarrow_{\text{ty}} w_2) \in PT$ such that w_1 is equal to the address $LV(\text{ad})$ loaded from, then we transform the state s at position $p = (b, i)$ to a state s' at position $p^+ = (b, i + 1)$. In s' , a fresh symbolic variable w is assigned to `x` and $w = w_2$ is added to KB . We write $LV[x := w]$ for the function where $LV[x := w](x) = w$ and $LV[x := w](y) = LV(y)$ for all $y \neq x$.


load from initialized allocated memory ($p: \text{"x = load ty, ty* ad", x, ad} \in \mathcal{V}_P$)	
$s = (p, LV, AL, PT, LI, KB)$	
$s' = (p^+, LV[x := w], AL, PT, LI, KB \cup \{w = w_2\})$	if $w \in \mathcal{V}_{\text{sym}}$ is fresh and
• there is $\llbracket v_1, v_2 \rrbracket \in AL$ with $\models \langle s \rangle \Rightarrow (v_1 \leq LV(\text{ad}) \wedge LV(\text{ad}) + \text{size}(\text{ty}) - 1 \leq v_2)$ • there are $w_1, w_2 \in \mathcal{V}_{\text{sym}}$ with $\models \langle s \rangle \Rightarrow (LV(\text{ad}) = w_1)$ and $(w_1 \hookrightarrow_{\text{ty}} w_2) \in PT$	

In our example, the `entry` block comprises the first three lines of the C program and the initialization of the pointer to the loop variable `k`: First, a non-deterministic unsigned integer is assigned to `n`, i.e., $(n = v_n) \in LV^B$, where v_n is not restricted. Moreover, memory for the pointers `tail_ptr` and `k_ad` is allocated and they point to `tail = NULL` and `k = 0`, respectively ($\text{tail_ptr} = v_{\text{tp}}$ and $\text{k_ad} = v_{\text{k_ad}}$ with $(v_{\text{tp}} \hookrightarrow_{\text{list}*} 0), (v_{\text{k_ad}} \hookrightarrow_{\text{i32}} 0) \in PT^B$). For simplicity, in Fig. 1 we use concrete values directly instead of introducing fresh variables for them. Since we assume a 64-bit architecture, `tail_ptr`'s allocation contains 8 bytes. For the integer value of `k`, only 4 bytes are allocated. Alignments and pointer sizes depend on the memory layout and are given in the LLVM program.

State C results from B by evaluating the `load` instruction at $(\text{cmpF}, 0)$, see the above `load` rule. Thus, there is an *evaluation edge* from B to C .

The next instruction is an `integer comparison` whose Boolean return value depends on whether the `unsigned` value of `k` is `less` than the one of `n`. If we cannot decide the validity of a comparison, we refine the state into two successor states. Thus, the states D and E (with $(v_n > 0) \in KB^D$ and $(v_n \leq 0) \in KB^E$) are reached by *refinement edges* from State C . Evaluating D yields `k1tn = 1` in F . Therefore, the `branch` instruction leads to the block `bodyF` in State G . State E is evaluated to a state with `k1tn = 0`. This path branches to the block `initPtr` and terminates quickly as `tail_ptr` points to an empty list.

The instruction at $(\text{bodyF}, 0)$ allocates 16 bytes of memory starting at v_{mem} in State H . The next instruction casts the pointer to the allocation from `i8*` to `list*` and assigns it to `curr`. Now the allocated area can be treated as a list element. Then `nondet_uint()` is invoked to assign a 32-bit integer to `nondet`.



are allocated in both states (often this is only the first element). Then elements which are the null pointer in one but not in the other state are lost. Hence, proving memory safety (and thus, also termination) fails when the list is traversed afterwards, since now there might be `next` pointers to non-allocated memory.

We solve this problem by introducing *list invariants*. In our example, we will infer an invariant stating that `curr` points to a list of length $x_\ell \geq 1$. This invariant also implies that all struct fields are allocated and that there is no sharing.

To this end, we adapt the merging heuristic from [25]. To merge two states s and s' at the same program position with $\text{domain}(LV^s) = \text{domain}(LV^{s'})$, we introduce a fresh symbolic variable x_{var} for each program variable `var` and use instantiations μ_s and $\mu_{s'}$ which map x_{var} to the corresponding symbolic variables of s and s' . For the merged state $\bar{s}$, we set $LV^{\bar{s}}(\text{var}) = x_{\text{var}}$. Moreover, we identify corresponding variables that only occur in the memory components and extend μ_s and $\mu_{s'}$ accordingly. In a second step, we check which constraints from the memory components and the knowledge base hold in both states in order to find invariants that we can add to the memory components and the knowledge base of $\bar{s}$. For example, if $\llbracket \mu_s(x), \mu_s(x^{\text{end}}) \rrbracket \in AL^s$ and $\llbracket \mu_{s'}(x), \mu_{s'}(x^{\text{end}}) \rrbracket \in AL^{s'}$ for $x, x^{\text{end}} \in \mathcal{V}_{\text{sym}}$, then $\llbracket x, x^{\text{end}} \rrbracket$ is added to $AL^{\bar{s}}$. To extend this heuristic to lists, we have to regard several memory entries together. If there is an `ad` $\in \mathcal{V}_{\mathcal{P}}$ such that $\mu_s(x_{\text{ad}}) = v_1^{\text{start}}$ and $\mu_{s'}(x_{\text{ad}}) = w_1^{\text{start}}$ both point to lists of type `ty` but of different lengths $\ell_s \neq \ell_{s'}$ with $\ell_s, \ell_{s'} \geq 1$, then we create a list invariant.

For a state s we say that v_1^{start} *points to a list of type ty with n fields and length ℓ_s with allocations $\llbracket v_k^{\text{start}}, v_k^{\text{end}} \rrbracket$ and values $v_{k,i}$ (for $1 \leq k \leq \ell_s$ and $1 \leq i \leq n$)* if the following conditions (a)–(d) hold:

- (a) `ty` is an LLVM struct type with subtypes `tyi` and field offsets $\text{off}_i \in \mathbb{N}$ for all $1 \leq i \leq n$ such that there exists exactly one $1 \leq j \leq n$ with `tyj = ty*`.
- (b) There exist pairwise different $\llbracket v_k^{\text{start}}, v_k^{\text{end}} \rrbracket \in AL^s$ for all $1 \leq k \leq \ell_s$ and $\models \langle s \rangle \Rightarrow v_k^{\text{end}} = v_k^{\text{start}} + \text{size}(\text{ty}) - 1$.
- (c) For all $1 \leq k \leq \ell_s$ and $1 \leq i \leq n$ there exist $v_{k,i}^{\text{start}}, v_{k,i} \in \mathcal{V}_{\text{sym}}$ with $\models \langle s \rangle \Rightarrow v_{k,i}^{\text{start}} = v_k^{\text{start}} + \text{off}_i$ and $(v_{k,i}^{\text{start}} \hookrightarrow_{\text{ty}_i} v_{k,i}) \in PT^s$.
- (d) For all $1 \leq k < \ell_s$ we have $\models \langle s \rangle \Rightarrow v_{k,j} = v_{k+1}^{\text{start}}$.

Condition (a) states that `ty` is a list type with n fields, where the pointer to the next element is in the j -th field. In (b) we ensure that each list element has a unique allocation of the correct size where v_1^{start} is the start address of the first allocation. Condition (c) requires that for the k -th element, the initial address plus the i -th offset points to a value $v_{k,i}$ of type `tyi`. Finally, (d) states that the recursive field of each element indeed points to the initial address of the next element.

Then, for fresh $x_\ell, x_i, \hat{x}_i \in \mathcal{V}_{\text{sym}}$, we add the following list invariant to $LI^{\bar{s}}$.

$$x_{\text{ad}} \xrightarrow{x_\ell} \text{ty} \llbracket (\text{off}_i : \text{ty}_i : x_i \dots \hat{x}_i) \rrbracket_{i=1}^n \quad (3)$$

To ensure that the allocations expressed by the list invariant are disjoint from all allocations in $AL^{\bar{s}}$, we do not use the list allocations $\llbracket v_k^{\text{start}}, v_k^{\text{end}} \rrbracket$ to infer generalized allocations in $AL^{\bar{s}}$. Similarly, to create $PT^{\bar{s}}$, we only use entries




Fig. 3. Merging of States

$v \hookrightarrow_{\text{ty}} w$ from PT^s and $PT^{s'}$ where v is disjoint from the list addresses, i.e., where $\models \langle s \rangle \Rightarrow v < v_k^{\text{start}} \vee v > v_k^{\text{end}}$ holds for all $1 \leq k \leq \ell_s$ and analogously for s' . Moreover, we add formulas to $KB^{\bar{s}}$ stating that (A) the length x_ℓ of the list is at least the smaller length of the merged lists, (B) x_ℓ is equal to all variables x which result from merging variables v and w that are equal to the lengths ℓ_s and $\ell_{s'}$ in s and s' , and (C) the symbolic variable x_i for the value of the i -th field of the first list element is equal to all variables x with $\mu_s(x) = v_{1,i}$ and $\mu_{s'}(x) = w_{1,i}$ where $v_{1,i}$ and $w_{1,i}$ are the values of the i -th field of the first list element in s and s' (and analogously for the values $\hat{x}_i$ of the last list element):

- (A) $\min(\ell_s, \ell_{s'}) \leq x_\ell$
- (B) $\bigwedge_{x \in \mu_s^{-1}(v) \cap \mu_{s'}^{-1}(w)} x_\ell = x$ for all $v, w \in \mathcal{V}_{\text{sym}}$ with $\models \langle s \rangle \Rightarrow v = \ell_s$ and $\models \langle s' \rangle \Rightarrow w = \ell_{s'}$
- (C) $\bigwedge_{x \in \mu_s^{-1}(v_{1,i}) \cap \mu_{s'}^{-1}(w_{1,i})} x_i = x$ and $\bigwedge_{x \in \mu_s^{-1}(v_{\ell_s,i}) \cap \mu_{s'}^{-1}(w_{\ell_{s'},i})} \hat{x}_i = x$ for all $1 \leq i \leq n$

To identify the variables in the list invariant (3) of $\bar{s}$ with the corresponding values in s and s' , the instantiations μ_s and $\mu_{s'}$ are extended such that $\mu_s(x_\ell) = \ell_s$, $\mu_{s'}(x_\ell) = \ell_{s'}$, $\mu_s(x_i) = v_{1,i}$, $\mu_{s'}(x_i) = w_{1,i}$, $\mu_s(\hat{x}_i) = v_{\ell_s,i}$, and $\mu_{s'}(\hat{x}_i) = w_{\ell_{s'},i}$ for all $1 \leq i \leq n$. Similarly, if there already exist list invariants in s and s' , for each pair of corresponding variables a new variable is introduced and mapped to its origin by μ_s and $\mu_{s'}$. This adaption of the merging heuristic only concerns the result of merging but not the rules *when* to merge two states. Thus, the same reasoning as in [25] can be used to prove soundness and termination of merging.

In our example, L and M contain lists of length $\ell_L = 1$ and $\ell_M = 2$. To ease the presentation, we re-use variables that are known to be equal instead of introducing fresh variables. If x_{mem} is the variable for the program variable **curr**, we have $\mu_L(x_{\text{mem}}) = v_{\text{mem}}$ and $\mu_M(x_{\text{mem}}) = w_{\text{mem}}$. Indeed, v_{mem} resp. w_{mem} points to a list with values $v_{k,i}$ resp. $w_{k,i}$ as defined in (a)–(d): For the type **list** with $n=2$, $\text{ty}_1 = \text{i32}$, $\text{ty}_2 = \text{list*}$, $\text{off}_1 = 0$, $\text{off}_2 = 8$, and $j=2$ (see (a)), we have $\llbracket v_{\text{mem}}, v_{\text{mem}}^{\text{end}} \rrbracket \in AL^L$ and $\llbracket v_{\text{mem}}, v_{\text{mem}}^{\text{end}} \rrbracket, \llbracket w_{\text{mem}}, w_{\text{mem}}^{\text{end}} \rrbracket \in AL^M$, all consisting of $\text{size}(\text{list}) = 16$ bytes, see (b). We have $(v_{\text{mem}} \hookrightarrow_{\text{i32}} v_{\text{nd}}, (v_{\text{cn}} \hookrightarrow_{\text{list*}} 0) \in PT^L$ with $(v_{\text{cn}} = v_{\text{mem}} + 8) \in KB^L$ and $(v_{\text{mem}} \hookrightarrow_{\text{i32}} v_{\text{nd}}, (v_{\text{cn}} \hookrightarrow_{\text{list*}} 0), (w_{\text{mem}} \hookrightarrow_{\text{i32}} w_{\text{nd}}, (w_{\text{cn}} \hookrightarrow_{\text{list*}} v_{\text{mem}}) \in PT^M$ with $(v_{\text{cn}} = v_{\text{mem}} + 8), (w_{\text{cn}} = w_{\text{mem}} + 8) \in KB^M$ (see (c)), so the first list element in M points to the second one (see (d)). Therefore, when merging L and M to a new state O (see Fig. 3), the lists are merged to a list invariant of variable length x_ℓ and we add the formulas (A) $1 \leq x_\ell$ and (B) $x_\ell = x_{\text{kinc}}$ to KB^O . By (C), the **i32** value of the first element is identified with x_{nd} , since $\mu_L(x_{\text{nd}})$ is equal to the first value of the first list element in L and $\mu_M(x_{\text{nd}})$ is equal to the first value of the first list element in M . Similarly, the values of the last list elements are identified with 0, as in L and M .

After merging s and s' to a generalized state $\bar{s}$, we continue symbolic execution from $\bar{s}$. The next time we reach the same program position, we might have to merge the corresponding states again. As described in [25], we use a heuristic for constructing the SEG which ensures that after a finite number of iterations, a state is reached that only represents concrete states that are also represented by an *already existing* (more general) state in the SEG. Then symbolic execution can continue from this more general state instead. So with this heuristic, the construction always ends in a complete SEG or an SEG containing the state *ERR*.

We formalized the concept of “generalization” by a symbolic execution rule in [25]. Here, the state $\bar{s}$ is a generalization of s if the conditions (g1)–(g6) hold.

Condition (g1) prevents cycles consisting only of refinement and generalization edges in the graph. Condition (g2) states that the instantiation $\mu: \mathcal{V}_{sym}(\bar{s}) \rightarrow \mathcal{V}_{sym}(s) \cup \mathbb{Z}$ maps symbolic variables from the more general state $\bar{s}$ to their counterparts from the more specific state s such that they correspond to the same program variable. Conditions (g3)–(g6) ensure that all knowledge present in $\overline{KB}$, $\overline{AL}$, $\overline{PT}$, and $\overline{LI}$ still holds in s with the applied instantiation.

generalization with instantiation μ

$$\frac{s = (p, LV, AL, PT, LI, KB)}{\bar{s} = (p, \overline{LV}, \overline{AL}, \overline{PT}, \overline{LI}, \overline{KB})} \quad \text{if}$$

(g1) s has an incoming evaluation edge

(g2) $\text{domain}(LV) = \text{domain}(\overline{LV})$ and $LV(\text{var}) = \mu(\overline{LV}(\text{var}))$ for all $\text{var} \in \mathcal{V}_{\mathcal{P}}$ where LV and $\overline{LV}$ are defined

(g3) $\models \langle s \rangle \Rightarrow \mu(\overline{KB})$

(g4) if $\llbracket x_1, x_2 \rrbracket \in \overline{AL}$, then $\llbracket v_1, v_2 \rrbracket \in AL$ with $\models \langle s \rangle \Rightarrow v_1 = \mu(x_1) \wedge v_2 = \mu(x_2)$

(g5) if $(x_1 \hookrightarrow_{\text{ty}} x_2) \in \overline{PT}$,

then $(v_1 \hookrightarrow_{\text{ty}} v_2) \in PT$ with $\models \langle s \rangle \Rightarrow v_1 = \mu(x_1) \wedge v_2 = \mu(x_2)$

(g6) if $(x_{\text{ad}} \xrightarrow{x_{\ell}}_{\text{ty}} [(off_i : \text{ty}_i : x_i.. \hat{x}_i)]_{i=1}^n) \in \overline{LI}$,

then either $(v_{\text{ad}} \xrightarrow{v_{\ell}}_{\text{ty}} [(off_i : \text{ty}_i : v_i.. \hat{v}_i)]_{i=1}^n) \in LI$ with

- $\models \langle s \rangle \Rightarrow v_{\text{ad}} = \mu(x_{\text{ad}}) \wedge v_{\ell} = \mu(x_{\ell})$ and
- $\models \langle s \rangle \Rightarrow v_i = \mu(x_i) \wedge \hat{v}_i = \mu(\hat{x}_i)$ for all $1 \leq i \leq n$,

or v_1^{start} points to a list of type ty and length ℓ with allocations $\llbracket v_k^{\text{start}}, v_k^{\text{end}} \rrbracket$ and values $v_{k,i}$ (for $1 \leq k \leq \ell, 1 \leq i \leq n$) such that

- $\models \langle s \rangle \Rightarrow v_1^{\text{start}} = \mu(x_{\text{ad}}) \wedge \ell = \mu(x_{\ell})$,
- $\models \langle s \rangle \Rightarrow v_{1,i} = \mu(x_i) \wedge v_{\ell,i} = \mu(\hat{x}_i)$ for all $1 \leq i \leq n$, and
- if $(z_1 \hookrightarrow_{\text{ty}} z_2) \in \overline{PT}$,
then $\models \langle s \rangle \Rightarrow \mu(z_1) < v_k^{\text{start}} \vee \mu(z_1) > v_k^{\text{end}}$ for all $1 \leq k \leq \ell$.

Condition (g6) is new compared to [25] and takes list invariants into account. So for every list invariant $\bar{l}$ of $\bar{s}$ there is either a corresponding list invariant l in s such that lists represented by l in s are also represented by $\bar{l}$ in $\bar{s}$, or there is a concrete list in s that is represented by $\bar{l}$ in $\bar{s}$. The last condition of the latter case ensures that disjointness between the memory domains of $\overline{PT}$ and $\overline{LI}$ is preserved. See [18] for the soundness proof of the extended generalization rule, i.e., that every concrete state represented by s is also represented by $\bar{s}$.

Our merging technique always yields generalizations according to this rule, i.e., the edges from L and M to O in Fig. 3 are generalization edges. Here, one




Fig. 4. Extending a List Invariant

chooses μ_L and μ_M such that $\mu_L(x_{\text{mem}}) = v_{\text{mem}}$, $\mu_L(x_\ell) = 1$, $\mu_L(x_{\text{nd}}) = v_{\text{nd}}$, $\mu_L(\hat{x}_{\text{nd}}) = v_{\text{nd}}$, $\mu_L(x_{\text{next}}) = 0$, $\mu_M(x_{\text{mem}}) = w_{\text{mem}}$, $\mu_M(x_\ell) = 2$, $\mu_M(x_{\text{nd}}) = w_{\text{nd}}$, $\mu_L(\hat{x}_{\text{nd}}) = v_{\text{nd}}$, and $\mu_M(x_{\text{next}}) = v_{\text{mem}}$. In both cases, all conditions of the second case of (g6) with $\ell_L = 1$ and $\ell_M = 2$ are satisfied. With $\mu_L(x_{\text{kinc}}) = 1$ resp. $\mu_M(x_{\text{kinc}}) = 2$, we also have $\models \langle L \rangle \Rightarrow \mu_L(x_\ell) = \mu_L(x_{\text{kinc}})$ resp. $\models \langle M \rangle \Rightarrow \mu_M(x_\ell) = \mu_M(x_{\text{kinc}})$.

3.2 Adapting List Invariants


To handle and modify list invariants, three of our symbolic execution rules have to be changed. Section 3.2.1 presents a variant of the **store** rule where the list invariant is *extended* by an element. In Sect. 3.2.2, we adapt the **load** rule to load values from the first list element and we present a variant of the **getelementptr** rule for list *traversal*. Soundness of our new rules is proved in [18]. For all other instructions, the symbolic execution rules from [25] remain unchanged.

3.2.1 List Extension

After merging L and M , symbolic execution continues from the more general state O in Fig. 3. Here, the values of $\mathbf{k}$ and $\mathbf{kinc}$ and the length of the list are not concrete but any positive (resp. non-negative) value with $x_\ell = x_{\text{kinc}} = x_{\mathbf{k}} + 1$. The symbolic execution of O is similar to the steps from B to J in Sect. 3 (see Fig. 1). First, the value x_{kinc} stored at $\mathbf{k_ad}$ is loaded to $\mathbf{k}$. To distinguish whether $\mathbf{k} < \mathbf{n}$ still holds, the next state is refined. From the refined state with $\mathbf{k} < \mathbf{n}$, we enter the loop body again. A new block $\llbracket y_{\text{mem}}, y_{\text{mem}}^{\text{end}} \rrbracket$ of 16 bytes is allocated and y_{mem} is assigned to $\mathbf{mem}$ and $\mathbf{curr}$. Then, a new unknown value y_{nd} is assigned to $\mathbf{nondet}$. The address of the $\mathbf{i32}$ value of the current element (equal to y_{mem}) is computed by the first **getelementptr** instruction of the loop and the value y_{nd} of $\mathbf{nondet}$ is stored at it. The second **getelementptr** instruction computes the address y_{cn} of the recursive field and results in State P in Fig. 4, where $y_{\text{cn}} = y_{\text{mem}} + 8$ is added to KB^P . Now, **store** sets the address of the $\mathbf{next}$ field to the head of the list created in the previous iteration. Since this instruction extends the list by an element, instead of adding $y_{\text{cn}} \hookrightarrow_{\text{list}*} x_{\text{mem}}$ to PT^Q , we extend the list invariant: The length is set to y_ℓ and identified with $x_\ell + 1$ in KB^Q . The pointer x_{mem} to the first element is replaced by y_{mem} , while the first recursive field in the list gets the value x_{mem} . Since $(y_{\text{mem}} \hookrightarrow_{\mathbf{i32}} y_{\text{nd}}) \in PT^P$, y_{nd} is the value of the first $\mathbf{i32}$ integer in the list. We remove all entries from PT^Q that are already contained in the new list invariant, e.g., $y_{\text{mem}} \hookrightarrow_{\mathbf{i32}} y_{\text{nd}}$.

To formalize this adaption of list invariants, we introduce a modified rule for **store** in addition to the one in [25]. It handles the case where there is a concrete list at some address v^{start} , **pa** points to the m -th field of this list's first element, one wants to store a value t at the address **pa**, and one already has a list invariant l for the “tail” of the list in the j -th field (if $m \neq j$) resp. for the list at the address t (if $m = j$). In all other cases, the ordinary **store** rule is applied.

More precisely, let the list invariant l describe a list of length v_l at the address v_{ad} . Then l is replaced by a new list invariant l' which describes the list at the address v^{start} after storing t at the address **pa**. Irrespective of whether $m \neq j$ or $m = j$, the resulting list at v^{start} has the list at v_{ad} as its “tail” and thus, its length v'_l is $v_l + 1$. We prevent sharing of different elements by removing the allocation $\llbracket v^{start}, v^{end} \rrbracket$ of the list and all points-to information of pointers in $\llbracket v^{start}, v^{end} \rrbracket$.



list extension (p : “store ty t , ty* **pa**”, $t \in \mathcal{V}_{\mathcal{P}} \cup \mathbb{N}$, $\mathbf{pa} \in \mathcal{V}_{\mathcal{P}}$)

$s = (p, LV, AL, PT, LI, KB)$

$s' = (p^+, LV, AL \setminus \llbracket v^{start}, v^{end} \rrbracket, PT', LI \setminus \{l\} \cup \{l'\}, KB')$ if

- there is $l = (v_{ad} \xrightarrow{v_l} \text{ty} [(off_i : \text{lt}_i : w_i.. \hat{w}_i)]_{i=1}^n) \in LI$ with $\text{lt}_j = \text{lt}_y^*$
- there is $\llbracket v^{start}, v^{end} \rrbracket \in AL$ with $\models \langle s \rangle \Rightarrow v^{end} = v^{start} + \text{size}(\text{lt}_y) - 1$
- there exists $1 \leq m \leq n$ such that $\text{ty} = \text{lt}_m$ and $\models \langle s \rangle \Rightarrow LV(\mathbf{pa}) = v^{start} + \text{off}_m$
- $\models \langle s \rangle \Rightarrow v_{ad} = v_j$ if $m \neq j$ and $\models \langle s \rangle \Rightarrow v_{ad} = LV(t)$ if $m = j$
- for all $1 \leq i \leq n$ with $i \neq m$ there exist $v_i^{start}, v_i \in \mathcal{V}_{sym}$ with $\models \langle s \rangle \Rightarrow v_i^{start} = v^{start} + \text{off}_i$ and $(v_i^{start} \xrightarrow{\text{lt}_i} v_i) \in PT$
- $PT' = \{(x_1 \xrightarrow{\text{sy}} x_2) \in PT \mid \models \langle s \rangle \Rightarrow (v^{end} < x_1) \vee (x_1 + \text{size}(\text{sy}) - 1 < v^{start})\}$
- $l' = (v^{start} \xrightarrow{v'_l} \text{ty} [(off_i : \text{lt}_i : v_i.. \hat{w}_i)]_{i=1}^n)$
- $KB' = KB \cup \{v_m = LV(t), v'_l = v_l + 1\}$, where v_m, v'_l are fresh

3.2.2 List Traversal

After the current element y_{mem} is stored at x_{tp} and the value x_{kinc} of $\mathbf{k}$ is incremented to y_{kinc} and stored at $x_{\text{k.ad}}$, we reach a state R at position $(\text{cmpF}, 0)$ by the branch instruction. However, our already existing state O is more general than R , i.e., we can draw a generalization edge from R to O using the generalization rule with the instantiation μ_R where $\mu_R(x_{\text{mem}}) = y_{\text{mem}}$, $\mu_R(x_{\text{nd}}) = y_{\text{nd}}$, $\mu_R(x_{\text{cn}}) = y_{\text{cn}}$, $\mu_R(x_{\text{k}}) = x_{\text{kinc}}$, $\mu_R(x_{\text{kinc}}) = y_{\text{kinc}}$, $\mu_R(x_{\text{el}}) = y_{\text{el}}$, $\mu_R(\hat{x}_{\text{nd}}) = \hat{x}_{\text{nd}}$, and $\mu_R(x_{\text{next}}) = x_{\text{mem}}$. Thus, the cycle of the first loop closes here.





Fig. 5. Traversing a List Invariant

As mentioned, in the path from O to R there is a state at position $(\text{cmpF}, 1)$ which is refined (similar to State C). If $k < n$ holds, we reach R . The other path with $k \not< n$ leads out of the for loop to the block `initPtr` followed by the while loop (see State S and the corresponding LLVM code on the side). Here, block `cmpW` loads the value x_{mem} at address `tail_ptr` is loaded to `tail'` and stored at a new pointer variable `ptr`. State T is reached after the first iteration of the while loop body. Here, block `cmpW` loads the value x_{mem} stored at `ptr` to `str`. Since it is not the null pointer, we enter `bodyW`, which corresponds to the body of the while loop. First, x_{mem} is cast to an `i8` pointer. Then `getelementptr` computes a pointer x_{np} to the next element by adding 8 bytes to x_{mem} . After another cast back to a `list*` pointer, we load the content of the new pointer to `next`. To this end, we need the following new variant of the load rule to load values that are described by a list invariant.




load from list invariant ($p : \text{"x = load ty, ty* ad_i", } x, \text{ad_i} \in \mathcal{V}_{\mathcal{P}}$)	
$s = (p, LV, AL, PT, LI, KB)$	
$s' = (p^+, LV[x := w], AL, PT, LI, KB \cup \{w = v_i\})$	if $w \in \mathcal{V}_{sym}$ is fresh and
• there is $l = (v_{ad} \xrightarrow{v_\ell} \text{ty} [(off_i : \text{ty}_i : v_i.. \hat{v}_i)]_{i=1}^n) \in LI$ • there exists $1 \leq i \leq n$ such that $\text{ty} = \text{ty}_i$ and $\models \langle s \rangle \Rightarrow LV(\text{ad_i}) = v_{ad} + off_i$	

With this new **load** rule, the content of the new pointer is identified as x_{next} . It is loaded to **next** and stored at x_{ptr} . Then we return to the block **cmpw** (State T). Merging T with its predecessor at the same program position is not possible yet since the domains of the respective LV functions do not coincide. Now, x_{next} is loaded to **str** and compared to the null pointer. Since we do not have information about x_{next} , T 's successor state is refined to a state with $x_{\text{next}} = 0$ (which starts a path out of the loop to a return state), and to a state with $x_{\text{next}} \geq 1$, which reaches U after a few evaluation steps, see Fig. 5. Now, **getelementptr** computes the pointer $x'_{\text{np}} = x_{\text{next}} + 8$ to the third element of the list, which is assigned to **next_ptr**. $\langle U \rangle$ contains $x_\ell \geq 2$ since the first and the last pointer value are known to be different ($x_{\text{next}} \neq 0$). This information is crucial for creating a new list invariant starting at x_{next} , which is used in the next iteration of the loop. Therefore, if our list invariant did not contain variables for the first and the last pointer, we could not prove termination of the program. In such a case where the pointer to the third element of a list invariant is computed and the length of the list is at least two, we *traverse* the list invariant to retain the correspondence between the computed pointer x'_{np} and the new list invariant. In the resulting state V , we represent the first list element by an allocation $\llbracket x_{\text{mem}}, x_{\text{mem}}^{\text{end}} \rrbracket$ and preserve all knowledge about this element that was encoded in the list invariant ($x_{\text{mem}}^{\text{end}} = x_{\text{mem}} + 15$, $x_{\text{mem}} \hookrightarrow_{i32} x_{\text{nd}}$, $x_{\text{np}} \hookrightarrow_{\text{list*}} x_{\text{next}}$). Moreover, we adapt the list invariant such that it now represents the list at x_{next} (i.e., without its first element) starting with the value x'_{nd} . We also relate the length of the new list invariant to the length of the former one ($x'_\ell = x_\ell - 1$).

Thus, in addition to the rule for **getelementptr** in [25], we now introduce rules for list traversal via **getelementptr**. The rule below handles the case where the address calculation is based on the type **i8** and the **getelementptr** instruction adds the number of bytes given by the term t to the address **pa**. Here, the offsets in our list invariants are needed to compute the address of the accessed field. We also have similar rules for list traversal via field access (i.e., where the next element is accessed using **curr'→next** as in the **for** loop) and for the case where we cannot prove that the length v_ℓ of the list is at least 2, see [18].

$\text{list traversal } (p : \text{“pb = getelementptr i8, i8* pa, i m t”}, t \in \mathcal{V}_{\mathcal{P}} \cup \mathbb{N}, \text{pa, pb} \in \mathcal{V}_{\mathcal{P}})$	
$s = (p, LV, AL, PT, LI, KB)$	
$s' = (p^+, LV[\text{pb} := w_j^{start}], AL \cup \llbracket v^{start}, v^{end} \rrbracket, PT', LI \setminus \{l\} \cup \{l'\}, KB')$	if
• there is $l = (v_{ad} \xrightarrow{v_\ell} \text{ty}_j [(off_i : \text{ty}_i : v_i.. \hat{v}_i)]_{i=1}^n) \in LI$ with $\text{ty}_j = \text{ty}^*$, $\models \langle s \rangle \Rightarrow LV(\text{pa}) = v_j$, $\models \langle s \rangle \Rightarrow LV(t) = off_j$, and $\models \langle s \rangle \Rightarrow v_\ell \geq 2$ • $PT' = PT \cup \{(v_i^{start} \hookrightarrow_{\text{ty}_i} v_i) \mid 1 \leq i \leq n\}$ • $l' = (w^{start} \xrightarrow{w_\ell} \text{ty}_j [(off_i : \text{ty}_i : w_i.. \hat{v}_i)]_{i=1}^n)$ • $KB' = KB \cup \{v^{start} = v_{ad}, v^{end} = v^{start} + \text{size}(\text{ty}) - 1, w^{start} = v_j, w_\ell = v_\ell - 1, w_j^{start} = w^{start} + off_j\} \cup \{v_i^{start} = v_{ad} + off_i \mid 1 \leq i \leq n\}$ • $v^{start}, v^{end}, v_1^{start}, \dots, v_n^{start}, w^{start}, w_\ell, w_j^{start}, w_1, \dots, w_n \in \mathcal{V}_{sym}$ are fresh	

We continue the symbolic execution of State V in our example and finally obtain a complete SEG with a path from a state W at the position $(\text{cmpW}, 0)$ to the next state W' at this position, and a generalization edge back from W' to W using an instantiation $\mu_{W'}$. Both W and W' contain a list invariant similar to T where instead of the length x_ℓ in T , we have the symbolic variables z_ℓ and z'_ℓ in W and W' , where $\mu_{W'}(z_\ell) = z'_\ell$ (see [18] for more details).



4 Proving Termination

To prove termination of a program $\mathcal{P}$, as in [25] the cycles of the SEG are translated to an integer transition system whose termination implies termination of $\mathcal{P}$. The edges of the SEG are transformed into ITS transitions whose application conditions consist of the state formulas $\langle s \rangle$ and equations to identify corresponding symbolic variables of the different states. For evaluation and refinement edges, the symbolic variables do not change. For generalization edges, we use the instantiation μ to identify corresponding symbolic variables. In our example, the ITS has cyclic transitions of the following form:

$$\begin{aligned}
 O(x_n, x_k, x_{\text{kinc}}, \dots) &\rightarrow^+ R(x_n, x_k, x_{\text{kinc}}, \dots) & | & \quad x_{\text{kinc}} = x_k + 1 \wedge x_n > x_k \wedge \dots \\
 R(x_n, x_k, x_{\text{kinc}}, \dots) &\rightarrow O(x_n, x_{\text{kinc}}, \dots) \\
 W(z_\ell, z'_\ell, \dots) &\rightarrow^+ W'(z_\ell, z'_\ell, \dots) & | & \quad z_\ell = z'_\ell - 1 \wedge z_\ell \geq 1 \wedge \dots \\
 W'(z_\ell, z'_\ell, \dots) &\rightarrow W(z'_\ell, \dots)
 \end{aligned}$$

The first cycle resulting from the generalization edge from R to O terminates since k is increased until it reaches n . The generalization edge yields a condition identifying x_{kinc} in R with x_k in O , since $\mu_R(x_k) = x_{\text{kinc}}$. With the conditions $x_{\text{kinc}} = x_k + 1$ and $x_n > x_k$ (from KB^O), the resulting transitions of the ITS are terminating. The second cycle from the generalization edge from W' to W

terminates since the length of the list starting with `curr'` decreases. Although there is no program variable for the length, due to our list invariants the states contain variables for this length, which are also passed to the ITS. Thus, the ITS contains the variable z_ℓ (where z_ℓ in W is identified with z'_ℓ in W' due to $\mu_{W'}(z_\ell) = z'_\ell$). Since the condition $z'_\ell = z_\ell - 1$ is obtained on the path from W to W' and $z_\ell \geq 1$ is part of $\langle W \rangle$ due to the list invariant with length z_ℓ in LI^W , the resulting transitions of the ITS clearly terminate. Analogous to [25, Cor. 11 and Thm. 13], we obtain the following theorem. To prove that a complete SEG represents all program paths, in [25] we used the LLVM semantics defined by the Vellvm project [26]. One now also has to prove soundness of those symbolic execution rules which were modified due to the new concept of list invariants (i.e., generalization, list extension, and list traversal), see [18].

Theorem 1 (Memory Safety and Termination). *Let $\mathcal{P}$ be a program with a complete SEG $\mathcal{G}$. Since a complete SEG does not contain *ERR*, $\mathcal{P}$ is memory safe for all concrete states represented by the states in $\mathcal{G}$.⁴ If the ITS corresponding to $\mathcal{G}$ is terminating, then $\mathcal{P}$ is also terminating for all states represented by $\mathcal{G}$.*

5 Conclusion, Related Work, and Evaluation

We presented a new approach for automated proofs of memory safety and termination of C/LLVM-programs on lists. It first constructs a symbolic execution graph (SEG) which overapproximates all program runs. Afterwards, an integer transition system (ITS) is generated from this graph whose termination is proved using standard techniques. The main idea of our new approach is the extension of the states in the SEG by suitable *list invariants*. We developed techniques to infer and modify list invariants automatically during the symbolic execution.

During the construction of the SEG, the list invariants abstract from a concrete number of memory allocations to a list of allocations of variable length while preserving knowledge about some of the contents (the values of the fields of the first and the last element) and the list shape (the start address of the first element, the list length, and the content of the last recursive pointer which allows us to distinguish between cyclic and acyclic lists). They also contain information on the memory arrangement of the list fields which is needed for programs that access fields via pointer arithmetic. The symbolic variables for the list length and the first and last values of list elements are preserved when generating an ITS from the SEG. Thus, they can be used in the termination proof of the ITS (e.g., the variables for list length can occur in ranking functions).

In [5, 6, 22] we developed a technique for termination analysis of Java, based on a program transformation to *integer term rewrite systems* instead of ITSs. This approach does not require specific list invariants as recursive data structures on the heap are abstracted to terms. However, these terms are unsuitable for

⁴ Our approach can only *prove* but not *disprove* memory safety, i.e., a SEG with the state *ERR* just means that we failed in showing memory safety.

C, since they cannot express memory allocations and the connection to their contents.

Separation logic predicates for termination of list programs were also used in [1], but their list predicates only consider the list length and the recursive field, but no other fields or offsets. The tools *Cyclist* [24] and *HipTNT+* [19] are integrated in separation logic systems which also allow to define heap predicates. However, they require annotations and hints which parameters of the list predicates are needed as a termination measure. The tool *2LS* [20] also provides basic support for dynamic data structures. But all these approaches are not suitable if termination depends on the contents or the shape of data structures combined with pointer arithmetic. In [10], programs can be annotated with arithmetic and structural properties to reason about termination. In contrast, our approach does not need hints or annotations, but finds termination arguments fully automatically.

We implemented our approach in *AProVE* [25]. While C programs with lists are very common, existing tools can hardly prove their termination. Therefore, the current benchmark collections for termination analysis contain almost no list programs. In 2017, a benchmark set⁵ of 18 typical C-programs on lists was added to the *Termination* category of the *Competition on Software Verification (SV-COMP)* [3], where 9 of them are terminating. Two of these 9 programs do not need list invariants, because they just create a list without operating on it afterwards. The remaining seven terminating programs create a list and then traverse it, search for a value, or append lists and compute the length afterwards. Only few tools in *SV-COMP* produced correct termination proofs for programs from this set: *HipTNT+* and *2LS* failed for all of them. *CPAchecker* [2] and *PeSCo* [23] proved termination and non-termination for one of these programs in 2020. *UAutomizer* [8] proved termination for two and non-termination for seven programs. The termination proofs of *CPAchecker*, *PeSCo*, and *UAutomizer* only concern the programs that just create a list. Our new version of *AProVE* is the only termination prover⁶ that succeeds if termination depends on the shape or contents of a list after its creation. Note that for non-termination, a proof is a single non-terminating program path, so here list invariants are less helpful.

For the *Termination Competition* [15] 2022, we submitted 18 terminating C programs on lists⁷ (different from the ones at *SV-COMP*), where two of them just create a list. Three traverse it afterwards (by a loop or recursion), and ten search for a value, where for nine, also the list contents are relevant for termination. Three programs perform common operations like inserting or deleting an element. *UAutomizer* proves termination for a program that just creates a list but not for programs operating on the list afterwards. With our approach, *AProVE* succeeds on 17 of the 18 programs. Overall, *AProVE* and *UAutomizer* were the two

⁵ <https://github.com/sosy-lab/sv-benchmarks/tree/master/c/termination-memory-linkedlists>.

⁶ We did not compare with the tool *VeriFuzz* [21], since it does not prove termination but only tests for non-termination and thus, it is unsound for inferring termination.

⁷ <https://github.com/TermCOMP/TPDB/tree/master/C/Hensel.22>.

most powerful tools for termination of C in *SV-COMP* 2022 and the *Termination Competition* 2022, with UAutomizer winning the former and AProVE winning the latter competition. To download AProVE, run it via its web interface, and for details on our experiments, see https://aprove-developers.github.io/recursive_structs.

	SV-C T.	SV-C Non-T.	TermCmp T.
AProVE	7 (of 9)	5 (of 9)	17 (of 18)
UAutomizer	2 (of 9)	7 (of 9)	1 (of 18)

References

1. Berdine, J., Cook, B., Distefano, D., O’Hearn, P.W.: Automatic termination proofs for programs with shape-shifting heaps. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 386–400. Springer, Heidelberg (2006). https://doi.org/10.1007/11817963_35
2. Beyer, D., Keremoglu, M.E.: CPAchecker: a tool for configurable software verification. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 184–190. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_16
3. Beyer, D.: Progress on software verification: SV-COMP 2022. In: Fisman, D., Rosu, G. (eds.) TACAS 2022. LNCS, vol. 13244, pp. 375–402. Springer, Cham (2022). <http://sv-comp.sosy-lab.org/2022/>. https://doi.org/10.1007/978-3-030-99527-0_20
4. Bozga, M., Iosif, R., Perarnau, S.: Quantitative separation logic and programs with lists. J. Aut. Reason. **45**(2), 131–156 (2010). <http://dx.doi.org/10.1007/s10817-010-9179-9>
5. Brockschmidt, M., Otto, C., Giesl, J.: Modular termination proofs of recursive Java Bytecode programs by term rewriting. In: Schmidt-Schauss, S. (ed.) RTA 2011, LIPIcs 10, pp. 155–170 (2011). <http://dx.doi.org/10.4230/LIPIcs.RTA.2011.155>
6. Brockschmidt, M., Musiol, R., Otto, C., Giesl, J.: Automated termination proofs for Java programs with cyclic data. In: Madhusudan, P., Seshia, Sanjit A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 105–122. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31424-7_13
7. Brockschmidt, M., Cook, B., Ishtiaq, S., Khlaaf, H., Piterman, N.: T2: temporal property verification. In: Chechik, M., Raskin, J.-F. (eds.) TACAS 2016. LNCS, vol. 9636, pp. 387–393. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49674-9_22
8. Chen, Y.-F., et al.: Advanced automata-based algorithms for program termination checking. In: Foster, J.S., Grossman, D. (eds.) PLDI 2018, pp. 135–150 (2018). <http://dx.doi.org/10.1145/3192366.3192405>
9. Clang: <http://clang.llvm.org>
10. David, C., Kroening, D., Lewis, M.: Propositional reasoning about safety and termination of heap-manipulating programs. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 661–684. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46669-8_27
11. Emrich, F., Hensel, J., Giesl, J.: AProVE: Modular termination analysis of memory-manipulating C programs. CoRR, abs/2302.02382 (2023). <http://dx.doi.org/10.48550/arXiv.2302.02382>
12. Frohn, F., Giesl, J.: Proving non-termination via loop acceleration. In: Barrett, C.W., Yang, J. (eds.) FMCAD 2019, pp. 221–230 (2019). <http://dx.doi.org/10.23919/FMCAD.2019.8894271>

13. Frohn, F., Giesl, J.: Proving non-termination and lower runtime bounds with LoAT (system description). In: Blanchette, J., Kovács, L., Pattinson, D. (eds.) *Automated Reasoning. IJCAR 2022*. LNCS, vol. 13385, pp. 712–722. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-10769-6_41
14. Giesl, J., et al.: Analyzing program termination and complexity automatically with AProVE. *J. Autom. Reason.* **58**(1), 3–31 (2017). <http://dx.doi.org/10.1007/s10817-016-9389-x>
15. Giesl, J., Rubio, A., Sternagel, C., Waldmann, J., Yamada, A.: The termination and complexity competition. In: Beyer, D., Huisman, M., Kordon, F., Steffen, B. (eds.) *TACAS 2019*. LNCS, vol. 11429, pp. 156–166. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17502-3_10. http://termination-portal.org/wiki/Termination_Competition_2022
16. Hensel, J., Giesl, J., Frohn, F., Ströder, T.: Termination and complexity analysis for programs with bitvector arithmetic by symbolic execution. *J. Log. Algebr. Methods Program.* **97**, 105–130, (2018). <http://dx.doi.org/10.1016/j.jlamp.2018.02.004>
17. Hensel, J., Mensendiek, C., Giesl, J.: AProVE: non-termination witnesses for C programs (competition contribution). In: Fisman, D., Rosu, G. (eds.) *TACAS 2022*. LNCS, vol. 13244, pp. 403–407. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-99527-0_21
18. Hensel, J., Giesl, J.: Proving termination of C programs with lists. *CoRR*, abs/2305.12159 (2023). <http://dx.doi.org/10.48550/arXiv.2305.12159>
19. Le, T.C., Qin, S., Chin, W.N.: Termination and non-termination specification inference. In: Grove, D., Blackburn, S.M. (eds.) *PLDI 2015*, pp. 489–498 (2015). <http://dx.doi.org/10.1145/2737924.2737993>
20. Malík, V., Martiček, Š., Schrammel, P., Srivas, M., Vojnar, T., Wahlang, J.: 2LS: memory safety and non-termination. In: Beyer, D., Huisman, M. (eds.) *TACAS 2018*. LNCS, vol. 10806, pp. 417–421. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89963-3_24
21. Metta, R., Yeduru, P., Karmarkar, H., Medicherla, R.K.: VeriFuzz 1.4: checking for (non-)termination (competition contribution). In: Sankaranarayanan, S., Sharygina, N. (eds.) *TACAS 2023*. LNCS, vol. 13994, pp. 594–599. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30820-8_42
22. Otto, C., Brockschmidt, M., von Essen, C., Giesl, J.: Automated termination analysis of Java Bytecode by term rewriting. In: Lynch, C. (ed.) *RTA 2011, LIPIcs*, vol. 6, pp. 259–276 (2010). <http://dx.doi.org/10.4230/LIPIcs.RTA.2010.259>
23. Richter, C., Wehrheim, H.: PeSCo: predicting sequential combinations of verifiers. In: Beyer, D., Huisman, M., Kordon, F., Steffen, B. (eds.) *TACAS 2019*. LNCS, vol. 11429, pp. 229–233. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17502-3_19
24. Rowe, R.N.S., Brotherston, J.: Automatic cyclic termination proofs for recursive procedures in separation logic. In: Bertot, Y., Vafeiadis, V. (eds.) *CPP 2017*, pp. 53–65 (2017). <http://dx.doi.org/10.1145/3018610.3018623>
25. Ströder, T., et al.: Automatically proving termination and memory safety for programs with pointer arithmetic. *J. Autom. Reason.* **58**(1), 33–65 (2017). <http://dx.doi.org/10.1007/s10817-016-9389-x>
26. Zhao, J., Nagarakatte, S., Martin, M.M.K., Zdancewic, S.: Formalizing the LLVM intermediate representation for verified program transformations. In: Field, J., Hicks, M. (eds.) *POPL 2012*, pp. 427–440 (2012). <http://dx.doi.org/10.1145/2103656.2103709>

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

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

