



# A Verified Implementation of the Bounded List Container

Raphaël Cauderlier and Mihaela Sighireanu<sup>(✉)</sup> 

IRIF, University Paris Diderot and CNRS, Paris, France  
{raphael.cauderlier,mihaela.sighireanu}@irif.fr

**Abstract.** This paper contributes to the trend of providing fully verified container libraries. We consider an implementation of the bounded doubly linked list container which manages the list in a fixed size, heap allocated array. The container provides constant time methods to update the list by adding, deleting, and changing elements, as well as cursors for list traversal and access to elements. The library is implemented in C, but we wrote the code and its specification by imitating the ones provided by GNAT for the standard library of Ada 2012. The proof of functional correctness is done using VeriFast, which provides an auto-active verification environment for Separation Logic extended with algebraic data types. The specifications proved entail the contracts of the Ada library and include new features. The verification method we used employs a precise algebraic model of the data structure and we show that it facilitates the verification and captures entirely the library contracts. This case study may be of interest for other verification platforms, thus we highlight the intricate points of its proof.

## 1 Introduction

Standard libraries of programming languages provide efficient implementations for common data containers. The details of these implementations are abstracted away by generic interfaces which are specified in terms of well understood mathematical structures such as sets, multisets, sequences, and partial functions. The intensive use of container libraries makes important their formal verification.

However, the functional correctness of these libraries is challenging to verify for several reasons. Firstly, their implementation is highly optimized: it employs complex data structures and manages the memory directly through pointers/references or specific memory allocators. Secondly, the specification of containers is rarely formal. Notable exceptions are, e.g., Eiffel [28] and SPARK [11]; recently, [1] provided a specification of the Ada 2012 container library. The formal specifications are very important when the library employs constructs that are out of the scope of the underlying mathematical structure. A typical example of such constructs are iterators. For example, Java iterators are generic and can exist independently of the container; Ada 2012 iterators, called cursors, are part of the container. Thirdly, the specification of the link between the low level implementation and the mathematical specification requires hybrid logics that are

able to capture both low level and high level specifications of the container. For verification purposes, these logics shall be supported by efficient solvers.

This work focuses on the functional verification of the *bounded doubly linked lists* container, which is a GNAT implementation [12] of the doubly linked lists container in the standard library of Ada 2012 [1]. This container is currently used by client programs [2] written in SPARK [22], a subset of Ada targeted at safety- and security-critical applications. The lists have bounded capacity, fixed at the list creation, and thus avoid dynamic memory allocation during the container use. This feature is required in critical code, where it is necessary to supply formal guarantees on the maximal amount of memory used by the running code.

The container implementation is original compared with other implementations of linked lists inside arrays. It employs an array of fixed size in which it manages (i) the occupied array cells inside a doubly linked list representing the content of the container and (ii) a singly linked list of free array cells. The operations provided are classic for lists. The *amortized constant time complexity* is preserved by the implementation of insert and delete operations. The list elements are designated using (bi-directional) cursors, also used to traverse the list. In conclusion, the code of this container was designed to ensure efficiency of operations and not its verification, and therefore it provides a realistic test for the automated verification.

Thanks to the introduction of formal contracts in Ada 2012, the container has been fully specified recently based on a previous specification in Why3 by Dross et al. [11]. The specification given is “meant to facilitate the formal verification of code using this container” [12], and it is presently used to prove the clients written in SPARK. The container is specified in terms of a model representing a functional implementation of bounded vectors, also written in Ada. This kind of specification is a substitute for the algebraic data types, not supported by Ada. It has the advantage of being executable, which enables the run-time verification of the implementation. An important feature of these contracts is their completeness [27] with respect to the models considered for the container and the cursors. This aspect is a challenge for the state of the art verification tools. The formal verification of these contracts can not be done by GNATprove, the deductive verification environment for SPARK, because the code employs language constructs out of its scope.

The goal of our study is to apply on-the-shelf verification tools to prove the full functional correctness and the memory safety for this implementation, without simplifying the code or the specification. To open the case study to more verification platforms, we choose to write this library in C, because C may capture all the features of the container implementation, except the strong typing and the generic types of Ada. The C implementation mimics the Ada code. The functional specification of the C code translates the container contracts from Ada based on (i) *representation predicates* that relate heap regions with algebraic models using inductively defined predicates, (ii) algebraic lists and maps, and (iii) inductively defined predicates and functions on the algebraic models. The logic including these features is undecidable in general. Therefore, we have to help the prover to obtain push button verification. The auto-active verification [20] environments are helpful in such tasks.

The invariant properties of the implementation and the features exhibited by the specification guided us towards deductive verification platforms that support Separation Logic [31] (SL) and algebraic data types. Consequently, we choose the VeriFast [15] auto-active verification tool, which provides means for (a) the specification of representation predicates in the style introduced for SL by O’Hearn et al. [25], (b) the definition of polymorphic algebraic data types, predicates and functions, and (c) the definition of user lemmas to help verification. Using these features, we employ a verification methodology based on the refinement of the original specification. The refined specification not only captures accurately the contracts, but also eases the deductive verification process, i.e., the writing of lemmas. For example, we employ a style of writing representation predicates in SL that leads to simpler lemmas for list segment composition.

To summarize, we verified the C implementation of a bounded doubly linked list container against its functional specification. In addition, we verified the safety of memory accesses using Separation Logic. For this, we annotated the C code and we extended the library for algebraic polymorphic lists of VeriFast with new predicates and lemmas. These logic development may be used in other verification tools based on induction.

The paper begins by presenting the case study in Sect. 2. Then, we highlight in Sect. 3 the main ingredients of the verification approach used and the challenges we faced. Section 4 presents the experimental results. We compare this work with other approaches for verification of containers and complex data structures in Sect. 5.

## 2 Dynamic Bounded Doubly-Linked Lists

This section presents the container code and its functional specification.

### 2.1 Overview

*Implementation:* The code is written in a very simple fragment of C, which may be easily translated to most imperative programming languages. It uses records and pointers to records, dynamic memory allocation, classic accesses to record fields and array elements, basic integer type and its operations. Like in the original code, the container does not support concurrency and has been written to obtain efficient operations and not to ease the verification. The container elements are designated through *cursors*, which represent valid positions in the list; they may be moved forward and backward in the list. The container interface includes 30 operations including classic operations (creation, copy, size access, clearing and deallocation, equality test, searching) and a rich set of utilities (inserting or deleting bunches of elements at some position, searching from the end, merging lists, swapping elements or links, reversing in place, sorting).

*Specification:* The functional specification is *model based* [28]. Two mathematical models are used: *algebraic lists* (i.e., finite sequences) to represent the content of the list and *finite partial maps* to model the set of valid cursors (see Sect. 2.3 for details). The contracts employ operations on these mathematical models that are beyond their classic usage. For example, the test of inclusion between the set of elements of two sequences, or the test that the domain of a partial mapping has been truncated from a given value. For this reason, we enriched the library of mathematical models provided by our prover with such operations and the corresponding axiomatizations (see Sect. 3.2).

An important feature of our functional specification is the usage of a *refined abstraction* for the list to ease the proof that the operations satisfy their contracts. We introduce a *precise model* for the list, which is an algebraic list of abstract cells, storing container values together with the links between the cells. This precise model is mapped to the abstract model (sequence of values) using a catamorphic mapping [35], called `model`. Moreover, the precise model is used to compute the (abstract) model of cursors, based on a catamorphic mapping, called `positions`. The use of the precise model facilitates the verification effort for proving that implementations of operations satisfy their contracts (see Sect. 3.1).

The functional specification is complete in the sense given by [27]: the post-condition of each operation uniquely defines its result and the side effect on the model of the container and of its cursors. However, it does not check for memory overflow at the container creation.

For the syntax of specifications, we employ in the following the specification language of VeriFast, which extends the normalized specification language for C, ACSL [3], with shorthand notations and operators for Separation Logic. Therefore, we employ ‘?’ to introduce existentially quantified variables, ‘&\*&’ for both classic conjunction and the separating one, ‘|->’ for the points-to operator that defines the content (right operand) of an allocated memory cell (left operand), and `emp` for the empty heap. Algebraic lists of VeriFast have type `list` and are polymorphic; the operations on lists have classic names. The definition of new logic types (and functions) is introduced by the keyword `inductive` (resp. `fixpoint`).

## 2.2 List Container

*List Elements:* The data stored in the list container is typed by an abstract type `Element_Type`, defined as an alias to the integer type in our code. This coding is sound for the proof of the functional correctness of the container implementation because the container assumes only that values of `Element_Type` may be compared for equality.

*List Cell:* Also called *node* in the following, the list cell encapsulates the container element together with links to the next and previous cell in the list. A node is also an element of the array allocated for the container.

```
struct Node_Type { int prev; int next; Element_Type elem; };
```

```

1 inductive pnode = pnode(int, int, Element_Type);
2 fixpoint int pprev(pnode pn) {switch(n) {case pnode(pp, pn, pe): return pp;}}
3 fixpoint int pnext(pnode pn) ...
4 fixpoint int pelem(pnode pn) ...
5
6 predicate node(struct Node_Type* n, int capacity; pnode pn) =
7   malloc_block_Node_Type(n) &&&
8   n->prev|->?iprev &&& n->next|->?inext &&&
9   n->elem|->?pelem &&&
10  inext>=0 &&& inext<=capacity &&&
11  pn==pnode(iprev, inext, pelem);
12
13 fixpoint bool is_free(pnode n) { return pprev(n)==-1; }
14 predicate free_node(struct Node_Type* n, int capacity; int inext) =
15   node(n, capacity, ?pn) &&&
16   is_free(pn)==true &&& inext==pnext(pn);
17
18 fixpoint bool is_occupied(pnode n) { return pprev(n)>=0; }
19 predicate occupied_node(struct Node_Type* n, int capacity; pnode pn) =
20   node(n, capacity, pn) &&& is_occupied(pn)==true;
21
22 predicate bdll(struct Node_Type * tab, int capacity,
23               int iprev, int ifrom, int ilast, int ito; list<pnode> m) =
24   ifrom==ito ?
25     (iprev==ilast &&& values==nil<pnode>)
26   : (occupied_node(tab+ifrom, capacity, ?p) &&&
27     pprev(p)==iprev &&&
28     bdll(tab, capacity, ifrom, pnext(p), ilast, ito, ?mtl) &&&
29     m==cons(p, mtl));
30
31 predicate uninit_free(struct Node_Type* tab, int capacity,
32                       int ifrom, int ito; list<int> model) =
33   ifrom==ito ? model==nil
34   : (ifrom<ito &&& free_node(tab+ifrom, capacity, 0) &&&
35     uninit_free(tab, capacity, ifrom+1, ito, ?mtl) &&&
36     model==cons(ifrom, mtl));
37
38 predicate init_free(struct Node_Type* tab, int capacity,
39                     int ifrom, int ito; list<int> model) =
40   ifrom==ito ? model==nil
41   : (ifrom>0 &&& free_node(tab+ifrom, capacity, ?inext) &&&
42     init_free(tab, capacity, inext, ito, ?mtl) &&&
43     model==cons(ifrom, mtl));
44
45 predicate free_nodes(struct Node_Type* tab, int cap,
46                      int free, int size; list<int> fmodel) =
47   free>=0 ? (init_free(tab, cap, free, 0, ?M) &&&
48             fmodel==M &&& length(M)+size==cap)
49   : (uninit_free(tab, cap, abs(free), cap+1, ?M) &&&
50     fmodel==M &&& length(M)+size==cap);
51
52 predicate list_inv(struct List* L;
53                   struct Node_Type* tab, int cap,
54                   int free, list<int> free_model,
55                   int head, int last, list<pnode> m) =
56   malloc_block_List(L) &&&
57   L->nodes|->tab &&& L->capacity|->cap &&& cap > 0 &&&
58   malloc_block(tab, sizeof(struct Node_Type)*(cap+1)) &&&
59   node(tab, cap, pnode(-1,0,_)) &&&
60   L->first|->head &&& head>=0 &&& head<=cap &&&
61   L->last|->last &&& last>=0 &&&
62   bdll(tab, cap, 0, head, last, 0, m) &&& L->size|->length(vs) &&&
63   L->free|->free &&& free<=cap &&&
64   free_nodes(tab, free, length(m), cap, free_model);

```

Fig. 1. Logic definitions for the BDLL container

The values of the C type are abstracted by the ghost type `pnode`, defined at line 1 in Fig. 1, which records the values of node fields. The logic functions `pprev`, `pnext`, and `pelem` to access first, second, resp. third component of a `purenode` value.

The predicate `node` (line 6 in Fig. 1) relates a node `n` allocated in the heap with its model `pn`. The allocation property is expressed by the predefined predicate `malloc_block_Node_Type`. The values of the fields are bound to existentially quantified variables and used to build the model of the node. The predicate `node` constrains the fields `prev` and `nxt` to be indexes in an array starting at index 0 and ending at index `capacity`.

There are two kinds of nodes in the array managed by the container: nodes occupied by list elements and nodes not yet used in the list, i.e., free. Free nodes have the `prev` field at `-1` and the `elem` field is irrelevant. They are specified by the predicate `free_node` (line 14 in Fig. 1), which also constrains the parameter `inext` to be equal to the value of the field `next`. Occupied nodes have the `prev` field set to a non-negative integer and the `elem` field is relevant. The predicate `occupied_node` (line 19 in Fig. 1) relates the node with its abstract model.

*Acyclic Doubly Linked List:* The container stores the doubly linked list (BDLL) into an array of fixed capacity, which is given at the container creation. The number of elements stored in the list can not exceed the container capacity. The nodes of the BDLL are stored starting from the index 1; index 0 plays the role of the null reference. The type of the list container is given by the following record:

```
struct List {
  int capacity; struct Node_Type* nodes; int size;
  int free; int first; int last;
};
```

The length of the list is given by the field `size`. The first and the last cells of the lists are stored at indexes `first` resp. `last`. Field `free` denotes the start of the list registering the free nodes. The operation creating the container allocates the array `nodes` and sets at free all nodes in the array. The fields denoting the size and the extreme cells of the doubly linked list are set to 0. The initialization of the `free` field is detailed in the next paragraph.

The representation predicate of the BDLL formed by the occupied nodes, `bdll`, is defined at line 22 in Fig. 1 as a doubly linked list segment starting by the node at index `ifrom`, ending by the node at index `ilast`; the starting node stores as previous node `iprev`, and the ending node stores as next node `ito`. The predicate definition is classic in Separation Logic [24], except the bound constraint on the node indexes (locations). If the source `ifrom` and target `ilast` indexes are equal, the list is empty; otherwise an occupied node is present at index `ifrom` and it is linked to the previous node and the remainder of the list. Notice the use of pointer arithmetics to access the node at index `ifrom`. The predicate `bdll` relates the heap specification with the mathematical model of the list, i.e., the sequence of abstract nodes. We employ the polymorphic algebraic type `list` available in VeriFast mathematical library and we instantiate it with



**Fig. 2.** Two doubly linked lists of capacity 4 and length 2

the logic type `pnode`. This precise model of the list content is mapped by the inductively defined ghost function `model` to the abstract model, sequence of values of `Element_type` stored.

*Acyclic List of Free Nodes:* The free nodes are organized in a singly linked list, called the *free-list*. The start of this list is given by the field `free` of the type `List`. If `free` is negative, the list is built from all nodes stored between `-free` and `capacity` included; this permits a fast initialization of the free-list at the container creation. If `free` is positive, the free-list starts at index `free`, uses as successor relation the `next` field, and ends at index 0. Figure 2 illustrates the two kinds of free list. The representation predicate `uninit_free` (line 31 in Fig. 1) is used when `free` is negative. It collects in the parameter `model` the sequence of the indexes of free nodes. For the second case, we define the predicate `init_free` (line 38 in Fig. 1). The two kinds of free-list are combined in the predicate `free_nodes` (line 45 of Fig. 1) that calls the correct predicate depending on the sign of `free`. Notice that the constraints required by this predicates (the relation between container capacity, BDLL and free-list sizes) are all present in the Ada 2012 specification [12].

*Representation Predicate:* The invariants of the container are collected in the predicate `list_inv` (line 52 in Fig. 1) which mainly specifies that the container is allocated in the heap (predefined `malloc_block_List` predicate), and its field `tab` is also allocated as a block containing `capacity + 1` records of type `Node_Type`. The first node of this array (at address `tab`) has its `prev` and `next` fields set to `-1` resp. `0`. The set of remaining nodes is split between the lists specified by the `bdll` and `free_nodes` predicates due to the separating conjunction. The size of the BDLL is exactly the one of its model and stored in the field `size`.

*Examples of Container Contracts:* We illustrate the usage of representation predicates defined above by presenting some contracts specifying container operations. For example, the contract of the constructor is:

```
struct List* List(int capacity);
/*@ requires capacity > 0;
  @ ensures list_inv(result,?t,capacity,-1,_,0,0,?m) &&& length(model(m))=0;
```

It states that the resulting container (denoted by the ghost variable `result`) is a well formed but empty bounded doubly linked list (its abstract model is the empty list) with `capacity` free nodes. As said before, the above contract (like in Ada 2012 specification), does not consider the case of memory shortage.

The contract of `is_equal` illustrates how the catamorphism `model` is used to obtain the abstract contract on the sequence of values from the precise models (given by variables `mL` and `mR` for each list parameter):

```
bool is_equal(struct List* L, struct List* R);
/*@ requires list_inv(L, ?tL, ?cL, ?fL, ?fmL, ?hL, ?lL, ?mL) &&&
        list_inv(R, ?tR, ?cR, ?fR, ?fmR, ?hR, ?lR, ?mR); @*/
/*@ ensures list_inv(L, tL, cL, fL, fmL, hL, lL, mL) &&&
        list_inv(R, tR, cR, fR, fmR, hR, lR, mR) &&&
        result == (model(mL) == model(mR)); @*/
```

The operation `clear` frees all occupied nodes. Its contract only constrains the content of the doubly linked list and leaves unspecified the free list.

```
void clear(struct List* L);
/*@ requires list_inv(L, ?t, ?c, ?f, ?fm, ?h, ?l, ?m);
/*@ ensures list_inv(L, t, c, ?f1, ?fm1, 0, 0, ?m1) &&& length(model(m1))=0;
```

## 2.3 Cursors

Following the Ada 2012 semantics [1], “a cursor designates a particular node within a list (...). A cursor keeps designating the same node (...) as long as the node is part of the container, even if the node is moved in the container. [...] If [a cursor] is not otherwise initialized, it is initialized to [...] `No_Element`.” Therefore, a cursor is a record storing an *array index*. The special cursor `No_Element` is defined as a global constant storing the index 0, indeed invalid for a list node (recall that valid nodes are stored from index 1).

```
struct Cursor { int current; };
const struct Cursor No_Element = { 0 };
```

The logic type `cursor` abstracts the cursor implementation (line 1 in Fig. 3). The representation predicate for cursors, `valid_cursor_or_noelem` (line 3 in Fig. 3), checks that the cursor content, `index`, corresponds to an occupied node in the list using the precise model `m` of the BDLL (see line 13). Moreover, the predicate computes from `m` and `first`, the BDLL starting index, the segments `before` and `after`, into which the cursor `pc` splits `m`.

Given a BDLL container, the model of valid cursors for this container is defined (following Ada 2012 specification) as the finite bijection between the set of abstract cursors and the positions (from 1 to the size) in the list. We encode the mathematical type map by an association list, using the polymorphic type provided in the libraries of `VeriFast`. The cursor model is computed by the logic function `positions` (line 19 in Fig. 3) from the container model, the index of the first node in the BDLL, and the first position in the list.



```

1 inductive cursor = NoElem | Valid(int);
2
3 predicate valid_cursor_or_noelem(struct Cursor* C,
4     int index, int first, list<pnode> m;
5     cursor pc,
6     list<pnode> before, list<pnode> after) =
7     C->current|->index &*&
8     pre_valid_cursor_or_noelem(index, first, vs, pc, before, after);
9
10 predicate pre_valid_cursor_or_noelem(int index, int first, list<pnode> vs;
11     cursor pc, list<pnode> before, list<pnode> after) =
12     pc == (index == 0 ? NoElem : Valid(index)) &*&
13     index == first ?
14     (before == nil &*& after == vs) :
15     (vs != nil &*&
16     pre_valid_cursor_or_noelem(index, pnext(head(vs)), tail(vs), pc, bef, after) &*&
17     before == cons(head(vs), bef));
18
19 fixpoint list<pair<cursor, int>> positions (list<purenode> values,
20     int index, int position) { ... }

```

**Fig. 3.** Logic definitions for cursors

Notice that this manner of specifying cursor model is coherent with the sequence model of the container, because the access to the elements of a sequence is based on positions. However, this specification choice does not combine well with inductive reasoning and induces additional work for the proof (see Sect. 3). We have to enrich the inductive list model with operations using positions. For example, we define the operation `M_Element(m,p)` which returns the  $p$ th element of the list  $m$ . We also defined operations `P_Has_Key` and `P_Get` on association lists to test if an abstract cursor is in the domain of the map resp. to obtain the value to which it is bound.

An example of a contract using cursors is the operation `element`, which returns the value stored at the position in the list given by the cursor `C`:

```

Element_Type element(struct List* L, struct Cursor* C)
/*@ requires list_inv(L,?tab,?capacity,?free,?fm,?head,?last,?m) &*&
    valid_cursor_or_noelem(C,?index,head,m,?pc,?bef,?aft) &*&
    P_Has_Key(positions(m,head,1),pc)==true; @*/
/*@ ensures list_inv(L,tab,capacity,free,fm,head,last,m) &*&
    valid_cursor_or_noelem(C,index,head,m,pc,bef,aft) &*&
    result==M_Element(model(m),P_Get(positions(m,head,1),pc)); @*/

```

Contracts of functions changing the positions in the list (e.g., insert or delete) are complete with respect to the model of cursors. For example, consider the post-condition of operation `delete_first` given below, which deletes first `Count` elements of the list. It uses a conditional expression (syntax like in C) to specify two contract cases. The first case corresponds to an input container with size less than `Count`. In the second case, the container preserved its content after position `Count` (predicate `M_Range_Shifted`) and the positions of valid cursors in the new container (of model `nvs`) are shifted by `Count` (predicate `P_Positions_Shifted`) with respect to the old container.

```

void delete_first(struct List* L, int Count);
/*@ requires list_inv(L,?tab,?cap,?free,?fm,?head,?last,?m) &*&

```

```

    Count >= 0; @*/
/*@ ensures list_inv(L,tab,cap,?nfree,?nfm,?nhead,?nlast,?nm) &&&
    length(m) <= Count ? length(nm) == 0
    : length(nm) == length(m) - Count &&&
    M_Range_Shifted(model(nm),model(m),1,length(nm),Count) &&&
    P_Positions_Shifted(positions(nm,nhead,1),
                        positions(m,head,1),1,Count); @*/

```

### 3 Verification Approach

We employ an auto-active verification approach [20], supported by the tool VeriFast [15]. The auto-active approach provides more automation of the verification process based on the ability given to the user to help the prover by adding annotations and lemmas and the efficient use of back-end solvers. This section highlights the methodology applied to conduct auto-active verification for this case study. This methodology is independent of the specific tool used. We also comment on the advantages and difficulties encountered with the tool used. Notice that we did not have prior experience with VeriFast.

#### 3.1 Model-Based Specification for Verification

The contracts provided for our container are in a first order logic over sequences and maps, which employs recursive logic functions. This theory is undecidable so we have to provide lemma to help the prover to tackle verification conditions.

*Usage of a precise model* is the solution we found to ease the writing of lemmas. It consists in refining the abstract model used for the container specification into a model that captures more details on the container organization. The abstract model is obtained from the refined one using a catamorphic mapping. This methodology is required by the gap between the abstract model and the lower level implementation of the container.

Let us explain why this methodology leads to efficient verification in our case. Consider the specification where (i) the model for the container is the sequence of the values stored and (ii) the model for the cursors is the mapping of occupied nodes to list positions. To capture these models with the representation predicate for the heap, i.e., the predicate `bdll` defined at line 22 in Fig. 1, we have to replace the model `m` by the sequence of values `vs` and the map of cursors `mc`. The verification of iterative operations on the list requires to provide a lemma that allows to compose “well linked” list segments into a new list segment, i.e.,

```

lemma void bdll_concat(struct Node_Type * t,
                      int p, int f, int l1, int n1, int l, int n)
requires bdll(t,?c,p,f,l1,n1,?vs1,?mc1) &&&
         bdll(t,c,l1,n1,l,n,?vs2,?mc2) &&& node(t+n, c, ?pn);
ensures bdll(t,c,p,f,l,n,append(vs1,vs2),
          append_maps(mc1,mc2)) &&& node(t+n, c, pn);
{ ... }

```

This lemma employs an operation `append_maps`, that concatenates two models of counters `mc1` and `mc2` such that the positions associated with counters in `mc1` are shifted by the size of the domain of `mc1`. This operation is more difficult to axiomatize than list concatenation. Moreover, all invariant proofs require to keep together the two loosely related models (sequence and map) which leads to less modular proofs. Our solution to this problem is to employ the precise model of the list segment represented by the `bdll` predicate, as has been presented in Sect. 2.2. The composition lemma for `bdll` predicate is simpler because it avoids the reasoning on the model of cursors.

The catamorphism mappings used to obtain the abstract model of the container and the model of valid cursors have good inductive definitions and enable efficient decision procedures [35]. However, these decision procedures are not available in our verification tool; this work may be a motivation to add them.

*Specification of user types* by representation predicates mapping them to inductive types is classical in Separation Logic. We encode the invariant of the BDLL data structure in the predicate `list_inv`. The adoption of `C` for the implementation keeps us away from the problems of verifying object-related properties, for example. However, this choice leads to an overburden in annotations because we have to specify that parameters of type `'struct List*'` satisfy the invariant.

Additional annotation have been supplied to axiomatize global constants (like the `No_Element` record in Fig. 1) and arrays of user-defined structures (like `nodes` in `List`).

*Contract cases* are intensively used in the considered GNAT library. We got around the absence of contract cases in VeriFast using conditional expressions and logic predicate and functions that relate two models (old and new). We do not observe any expressivity or performance problems with this method of encoding contracts.

### 3.2 Support for Specification Types

Specification of model types is done based on the mathematical models sequence (or inductive polymorphic list) and map (or inductive polymorphic association list). The VeriFast libraries including these models (mainly `list.*`) contain 9 predicates and 20 lemmas, and are not enough for the operations on models required in our specifications. We added tens of lemma and predicates. They are useful not only for the container proof but also for the verification of client programs with inductive back-end solvers. (Nowadays, these proofs are done by GNATprove by calling SMT solvers with quantifiers support.)

More problematic was the lack of support for finite maps and automation of inductive reasoning. VeriFast does not provide sets and finite maps as primitives. The encoding of cursor models by association lists renders more complex the lemmas needed on cursor models. For example, the map inclusion property is defined as follows:

```

1 predicate P_le<t>(list<pair<t,int> > Left, list<pair<t,int> > Right) =
2   switch(Left) {
3     case nil: return true;
4     case cons(p, m):
5       return P_Has_Key(Right, fst(p)) == true &&&
6         snd(p) == P_Get(Right, fst(p)) &&&
7         P_le(m, Right);
8   };

```

This definition is not as easy to reason about as we might expect. In particular, some properties of this definition of inclusion such as reflexivity are only provable under the additional assumption that the keys are distinct.

We proved that the models of cursors fulfill the constraint `distinct_keys` (defined also in VeriFast libraries) because keys are index positions in the array used to denote separated cells.

```

1 lemma void positions_distinct_keys(int index, list<purenode> m, int pos)
2 requires bdl1(?tab, ?cap, ?fst, index, ?last, ?z, m) &&& node(tab+z, cap, ?tab0);
3 ensures bdl1(tab, cap, fst, index, last, z, m) &&& node(tab+z, cap, tab0) &&&
4   distinct_keys(positions(m, index, pos)) == true;

```

Notice that these proofs are not necessary for provers with support for finite maps and sets. Although VeriFast supports as back-end solver Z3 [9], it does not use it for such theories. The inductive theories are supported by other back-end solvers, e.g., CVC4 [30] that are not connected to VeriFast.

### 3.3 Annotations Load

The annotations required by the proof of our library belong to two categories: (1) mandatory annotations including function contracts and predicates employed by these contracts and (2) auxiliary annotations including loop invariants, open/-close of predicates, definitions and calls of lemmas.

The prover VeriFast includes all this annotation burden, since we can not direct the prover in the usage of these annotations. VeriFast provides two mechanisms to limit the burden of the auxiliary annotations: (i) lemmas can be marked as automated which means they will be given to the backend solver on all problems, (ii) inductive predicate definitions can be automatically folded and unfolded when used with computed parameters.

We introduce few automated lemmas and call them in order to lighten the prover load. We don't observe performance problems by including all these annotations and despite the absence of modular proofs. The frame reasoning rule of Separation Logic seems to play an important role in this good behavior.

We found useful the two ways of specifying inductive predicates in VeriFast: by case on the model or by case over the aliasing of heap locations. We started with the first style, but finally chose the second to bring advantages of automatic folding and unfolding of computed predicates.

### 3.4 Challenges Dealt

To resume, we faced the following challenges during the verification process:

- We considered a functional specification which is already in use in client code. Therefore, we can not adapt this specification to ease the verification. Instead, we propose a method based on a refined specification based on a precise model of the container that eases the verification and allows to obtain the initial specification with minimal cost.
- The specification we received is complete with respect to the model of containers and cursors. This requires to specify logic functions and predicates that are more complex than the usual ones.
- The code has been designed to obtain efficient container implementation and does not focus on verification. Therefore, the verification task has been more difficult compared with previous work verifying functional specification of container libraries [28,39] designed with verification in mind.
- Only specifications of contracts for public operations on the container were provided. We had to annotate the code and the internal operations. This implied an additional cost in annotations because some internal operations break the data structure invariants.
- Having in mind the extension of this verification effort to other bounded container libraries (for sets or maps), we propose reusable logic libraries and suggest some improvements for the auto-active verification tool in use.

## 4 Verification Results

*Bugs Found:* We did not find spectacular bugs in the code, which is normal for a library that has been used for years. We only detected a potential arithmetic overflow in the computation of the memory to be allocated and a potential memory shortage. The last problem is in fact dealt for the SPARK clients using tools that measure the memory allocated by the program.

*Complete Specifications:* We also fix some minor completeness problems with the original specifications. Our verification effort leads to a complete functional specifications for all operations, including non public operations.

*Specification Load:* We have coded, specified, and verified 27 functions out of the 39 provided by the container library including equality and emptiness tests, clear, assign and copy, getting and setting one element, manipulating the cursors, inserting and deleting at some cursor, finding an element before and after a cursor. Most of the 12

**Table 1.** Statistics on the proof

File	#pred	#fix-points	#lemma	lines	
				annot	code
<code>vflist.gh</code>	2	8	9	234	–
<code>vfseq.gh</code>	14	10	34	482	–
<code>vfmap.gh</code>	13	9	64	1251	–
<code>cfdlll.h</code>	4	10	0	407	36
<code>cfdlll.c</code>	14	5	64	2684	506
Total	47	42	171	5058	542

remaining functions deal with sorting. The size of our development is given in Table 1. To obtain a specification close to the Ada 2012 one, we wrote two files of logic definitions for models (`vfseq.gh` and `vfmap.gh`) extending the VeriFast libraries. Additional fixpoint functions and lemmas required on VeriFast lists are written in file `vflist.gh`. The rate between source code and annotations is about 1 to 9. The required annotations (i.e., data structure invariant, pre/post conditions, and logic predicate and function used directly in them) represent a quarter of all annotations (including also loop invariants and lemmas). In Ada 2012 container, the rate between source code and contracts is already of about 1 to 3.

*Verification Performance:* We run VeriFast on a machine with 16 GB RAM, Intel core i5, and 2.70 GHz, installed with Linux. The back-end solver of VeriFast was `redux`. The verification takes 1.3s for the full container.

## 5 Related Work

The verification of individual data structures has received special attention. General safety properties (i.e., absence of out of array bounds accesses, null dereferences, division by zero, arithmetic overflow) may be verified automatically with low load of annotations using *static analysis methods*, e.g. [13, 17, 19, 21]. More complex properties like reachability of locations in the heap and shape of the data structures could also be proved with static analysis methods based on shape analysis, e.g., [4, 6, 10, 32]. These automatic techniques have been applied to linked lists coded in arrays [34]. These methods concern limited properties and may be used in the early stages of the library development to infer internal invariant properties. Extension of fully automatic techniques to cover functional specification abstractions like sets or bags are based either on *shape analysis*, e.g., [7, 14] or on logic fragments supported by SMT decision procedures [16, 18, 37, 38]. These functional specifications capture essential mathematical properties of the data structure but do not deal with properties of iterators over them.

At the opposite end of the spectrum of verification techniques, *interactive provers* have been used to obtain detailed specifications about data structures based on powerful theories, e.g., [8, 23, 29], but they require expertise and great amount of proof scripting.

At the intermediate level of automation, functional verification tools have been used to tackle the verification of specific data structures (e.g., Dafny [33], GRASShoper [26], VeriFast [15], or Why3 [36]) but we are not aware of any experiment on bounded lists.

The full functional correctness of container libraries has been considered in [28, 39]. They consider complex data structures in imperative and object oriented languages that require to verify special properties and may benefit from modular verification thanks to inheritance. In both cited works, a special effort has been deployed to improve the prover to call solvers for different theories or to generate verification conditions that may be dealt with efficiently. This efforts

lead to a low annotation overhead, especially in [28]. We use an on-the-shelf auto-active verification tool but improve its performances by employing a refinement method which leads to more automation but a more important annotation overhead. None of these works consider the container of bounded list.

## 6 Conclusion

We apply auto-active verification provided by the `VeriFast` tool to prove the functional specifications of the bounded doubly linked list container. The implementation we consider is in C, but it mimics the GNAT library [12], which is used in SPARK client programs. The functional specification is model-based and uses sequence and map mathematical models in a specific way to model the content of the list and its valid cursors. Our main contributions are (i) the improvement of the logic libraries of `VeriFast` to deal with such specific models and (ii) the use of a refinement based methodology to ease the proof automation.

This case study provides a motivation for the development of inductive solvers and their connection with auto-active provers like `VeriFast`. This experiment is another demonstration of the known fact (see [28]) that proving functional specifications of real world containers is more difficult than proving functional specification of data structures. The support for automation of these proofs is of an utmost importance to scale the verification to a full library of containers.

**Data Availability Statement and Acknowledgements.** The annotated code of the library and the sources used by its proof are available in the FIGSHARE repository [5]. The file-set also includes the distribution of the `VeriFast` tool running in the virtual machine provided at <https://doi.org/10.6084/m9.figshare.5896615>.

This work has been supported by the French national research organization ANR (grant ANR-14-CE28-0018). We thank Claire Dross and Yannick Moy from AdaCore for guiding us through the Ada standard library and for supplying the latest version of its specification. We thank Samantha Dihm for the first C version of the Ada containers.

## References

1. Ada Europe: Ada Reference Manual - Language and Standard Libraries, Chapter A.18.3 The Generic Package Containers. `Doubly_Linked_Lists` Norm ISO/IEC 8652:2012(E) (2012). [http://www.adaic.org/resources/add\\_content/standards/12rm/html/RM-TTL.html](http://www.adaic.org/resources/add_content/standards/12rm/html/RM-TTL.html)
2. AdaCore: SPARK verification gallery. <https://github.com/AdaCore/spark2014/tree/master/testsuite/gnatprove/tests>
3. Baudin, P., Filliâtre, J.C., Hubert, T., Marché, C., Monate, B., Moy, Y., Prevosto, V.: ACSL: ANSI C Specification Language (preliminary design V1.2), preliminary edition, May 2008
4. Calcagno, C., Distefano, D., O’Hearn, P.W., Yang, H.: Compositional shape analysis by means of bi-abduction. *J. ACM* **58**(6), 26:1–26:66 (2011)
5. Cauderlier, R., Sighireanu, M.: A verified implementation of the bounded list container. In: TACAS 2018 (2018). Figshare. <https://doi.org/10.6084/m9.figshare.5919145.v1>

6. Chang, B.E., Rival, X.: Relational inductive shape analysis. In: Proceedings of POPL, pp. 247–260. ACM (2008)
7. Chin, W.-N., David, C., Nguyen, H.H., Qin, S.: Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Sci. Comput. Program.* **77**(9), 1006–1036 (2012). The Programming Languages track at the 24th ACM Symposium on Applied Computing (SAC 2009)
8. Chlipala, A., Malecha, J.G., Morrisett, G., Shinnar, A., Wisnesky, R.: Effective interactive proofs for higher-order imperative programs. In: Proceedings of ICFP, pp. 79–90. ACM (2009)
9. de Moura, L.M., Bjørner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-78800-3\\_24](https://doi.org/10.1007/978-3-540-78800-3_24)
10. Drăgoi, C., Enea, C., Sighireanu, M.: Local shape analysis for overlaid data structures. In: Logozzo, F., Fähndrich, M. (eds.) SAS 2013. LNCS, vol. 7935, pp. 150–171. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-38856-9\\_10](https://doi.org/10.1007/978-3-642-38856-9_10)
11. Dross, C., Filiâtre, J.-C., Moy, Y.: Correct code containing containers. In: Gogolla, M., Wolff, B. (eds.) TAP 2011. LNCS, vol. 6706, pp. 102–118. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-21768-5\\_9](https://doi.org/10.1007/978-3-642-21768-5_9)
12. GNU Foundation: GNAT library components in gcc 7.1. [https://sourceware.org/svn/gcc/tags/gcc-7.1.0.release/gcc/ada/files/a-cfd11i.ad\\*](https://sourceware.org/svn/gcc/tags/gcc-7.1.0.release/gcc/ada/files/a-cfd11i.ad*)
13. Halbwachs, N., Péron, M.: Discovering properties about arrays in simple programs. In: Proceedings of PLDI, pp. 339–348. ACM (2008)
14. Itzhaky, S., Bjørner, N., Reps, T., Sagiv, M., Thakur, A.: Property-directed shape analysis. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 35–51. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08867-9\\_3](https://doi.org/10.1007/978-3-319-08867-9_3)
15. Jacobs, B., Smans, J., Piessens, F.: A quick tour of the VeriFast program verifier. In: Ueda, K. (ed.) APLAS 2010. LNCS, vol. 6461, pp. 304–311. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-17164-2\\_21](https://doi.org/10.1007/978-3-642-17164-2_21)
16. Jacobs, S., Kuncak, V.: Towards complete reasoning about axiomatic specifications. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 278–293. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-18275-4\\_20](https://doi.org/10.1007/978-3-642-18275-4_20)
17. Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., Yakobowski, B.: Frama-C: a software analysis perspective. *FAC* **27**(3), 573–609 (2015)
18. Kuncak, V., Piskac, R., Suter, P., Wies, T.: Building a calculus of data structures. In: Barthe, G., Hermenegildo, M. (eds.) VMCAI 2010. LNCS, vol. 5944, pp. 26–44. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-11319-2\\_6](https://doi.org/10.1007/978-3-642-11319-2_6)
19. Laviro, V., Logozzo, F.: SubPolyhedra: a family of numerical abstract domains for the (more) scalable inference of linear inequalities. *STTT* **13**(6), 585–601 (2011)
20. Leino, K.R.M., Moskal, M.: Usable auto-active verification (2010). <http://www.fm.csl.sri.com>
21. Liu, J., Rival, X.: Abstraction of arrays based on non contiguous partitions. In: D’Souza, D., Lal, A., Larsen, K.G. (eds.) VMCAI 2015. LNCS, vol. 8931, pp. 282–299. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46081-8\\_16](https://doi.org/10.1007/978-3-662-46081-8_16)
22. McCormick, J.W., Chapin, P.C.: Building High Integrity Applications with SPARK. Cambridge University Press, Cambridge (2015)
23. Nanevski, A., Morrisett, G., Shinnar, A., Govereau, P., Birkedal, L.: Ynot: dependent types for imperative programs. In: Proceeding of ICFP, pp. 229–240. ACM (2008)
24. O’Hearn, P.W., Reynolds, J.C., Yang, H.: Local reasoning about programs that alter data structures. In: Fribourg, L. (ed.) CSL 2001. LNCS, vol. 2142, pp. 1–19. Springer, Heidelberg (2001). [https://doi.org/10.1007/3-540-44802-0\\_1](https://doi.org/10.1007/3-540-44802-0_1)



25. O’Hearn, P.W., Yang, H., Reynolds, J.C.: Separation and information hiding. In: Proceedings of POPL, pp. 268–280. ACM (2004)
26. Piskac, R., Wies, T., Zufferey, D.: Automating Separation logic with trees and data. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 711–728. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08867-9\\_47](https://doi.org/10.1007/978-3-319-08867-9_47)
27. Polikarpova, N., Furia, C.A., Meyer, B.: Specifying reusable components. In: Leavens, G.T., O’Hearn, P.W., Rajamani, S.K. (eds.) VSTTE 2010. LNCS, vol. 6217, pp. 127–141. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-15057-9\\_9](https://doi.org/10.1007/978-3-642-15057-9_9)
28. Polikarpova, N., Tschannen, J., Furia, C.A.: A fully verified container library. In: Bjørner, N., de Boer, F. (eds.) FM 2015. LNCS, vol. 9109, pp. 414–434. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-19249-9\\_26](https://doi.org/10.1007/978-3-319-19249-9_26)
29. Pottier, F.: Verifying a hash table and its iterators in higher-order separation logic. In: Proceedings of CPP, pp. 3–16. ACM (2017)
30. Reynolds, A., Kuncak, V.: Induction for SMT solvers. In: D’Souza, D., Lal, A., Larsen, K.G. (eds.) VMCAI 2015. LNCS, vol. 8931, pp. 80–98. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46081-8\\_5](https://doi.org/10.1007/978-3-662-46081-8_5)
31. Reynolds, J.C.: Separation logic: a logic for shared mutable data structures. In: Proceedings of LICS, pp. 55–74. IEEE (2002)
32. Rinetzky, N., Sagiv, M., Yahav, E.: Interprocedural shape analysis for cutpoint-free programs. In: Hankin, C., Siveroni, I. (eds.) SAS 2005. LNCS, vol. 3672, pp. 284–302. Springer, Heidelberg (2005). [https://doi.org/10.1007/11547662\\_20](https://doi.org/10.1007/11547662_20)
33. Rustan, K., Leino, M.: Main Microsoft research Dafny web page. <http://research.microsoft.com/en-us/projects/dafny>
34. Sotin, P., Rival, X.: Hierarchical shape abstraction of dynamic structures in static blocks. In: Jhala, R., Igarashi, A. (eds.) APLAS 2012. LNCS, vol. 7705, pp. 131–147. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-35182-2\\_10](https://doi.org/10.1007/978-3-642-35182-2_10)
35. Suter, P., Dotta, M., Kuncak, V.: Decision procedures for algebraic data types with abstractions. In: Proceedings of POPL, pp. 199–210. ACM (2010)
36. Why3-Team: Why3 verification gallery. <http://toccata.lri.fr/gallery/>
37. Wies, T., Muñoz, M., Kuncak, V.: An efficient decision procedure for imperative tree data structures. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE 2011. LNCS (LNAI), vol. 6803, pp. 476–491. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-22438-6\\_36](https://doi.org/10.1007/978-3-642-22438-6_36)
38. Wies, T., Muñoz, M., Kuncak, V.: Deciding functional lists with sublist sets. In: Joshi, R., Müller, P., Podelski, A. (eds.) VSTTE 2012. LNCS, vol. 7152, pp. 66–81. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-27705-4\\_6](https://doi.org/10.1007/978-3-642-27705-4_6)
39. Zee, K., Kuncak, V., Rinard, M.C.: Full functional verification of linked data structures. In: Proceedings of PLDI, pp. 349–361. ACM (2008)

**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.

