



Saturating Sorting without Sorts

Pamina Georgiou, Márton Hajdu, and Laura Kovács

TU Wien, Austria

Abstract

We present a first-order theorem proving framework for establishing the correctness of functional programs implementing sorting algorithms with recursive data structures. We formalize the semantics of recursive programs in many-sorted first-order logic and integrate sortedness/permutation properties within our first-order formalization. Rather than focusing on sorting lists of elements of specific first-order theories, such as integer arithmetic, our list formalization relies on a sort parameter abstracting (arithmetic) theories and hence concrete sorts. We formalize the permutation property of lists in first-order logic so that we automatically prove verification conditions of such algorithms purely by superposition-based first-order reasoning. Doing so, we adjust recent efforts for automating induction in saturation. We advocate a compositional approach for automating proofs by induction required to verify functional programs implementing and preserving sorting and permutation properties over parameterized list structures. Our work turns saturation-based first-order theorem proving into an automated verification engine by (i) guiding automated inductive reasoning with manual proof splits and (ii) fully automating inductive reasoning in saturation. We showcase the applicability of our framework over recursive sorting algorithms, including Mergesort and Quicksort.

1 Introduction

Sorting algorithms are integrated parts of any modern programming language, hence ubiquitous in computing, which naturally triggers the demand of validating the functional correctness of sorting routines. Such algorithms typically implement recursive/iterative operations over potentially unbounded data structures, for instance lists or arrays, combined with arithmetic manipulations of numeric data types, such as naturals, integers or reals. Automating the formal verification of sorting routines therefore brings the challenge of automating recursive/inductive reasoning in extensions and combinations of first-order theories, while also addressing the reasoning burden arising from design choices made for the purpose of efficient sorting. Most notably, `Quicksort` [11] is known to be easily implemented when making use of recursive function calls, for example, as given in Figure 1, whereas purely procedural implementations of `Quicksort` require additional recursive data structures such as stacks. While `Quicksort` and other sorting routines have been proven correct by means of manual efforts [7], proof assistants [20, 24, 3], abstract interpreters [9], or model checkers [12], to the best of our knowledge such correctness proofs so far have not been fully automated with saturation-based automated reasoning.

```

1 datatype 'a list = nil | cons('a, ('a list))
2
3 quicksort :: 'a list → 'a list
4 quicksort(nil) = nil
5 quicksort(cons(x, xs)) =
6   append(
7     quicksort(filter<(x, xs)) ,
8     cons(x, quicksort(filter≥(x, xs))))
9
10 append :: 'a list → 'a list → 'a list
11 append(nil, xs) = xs
12 append(cons(x, xs), ys) = cons(x, append(xs, ys))

```

Figure 1: Recursive algorithm of Quicksort, using the recursive function definitions `append`, `filter<` and `filter≥` over lists of sort a .

In this paper we aim to verify the partial correctness of functional programs with recursive data structures, in an automated manner by using saturation-based first-order theorem proving. To achieve this, we turn the automated first-order reasoner into a complementary approach to interactive proof assistance: (i) we rely on manual guidance in splitting inductive proof goals into subgoals (Sections 5 and 6), but (ii) fully automate inductive proofs in saturation-based reasoning (Section 4). The crux of our approach is a compositional reasoning setting based on superposition-based first-order theorem proving [15] with native support for induction [10] and first-order theories of recursively defined data types [14]. We extend this setting to support the first-order theory of lists parameterized by an abstract background theory/sort a and advocate *computation induction* for induction on recursive function calls. As such, our framework allows us to automatically discharge manually split verification conditions that require inductive proofs, without requiring manually proven or a priori given inductive annotations such as loop invariants, nor user input to perform proofs by induction. Doing so, we automatically derive induction axioms during *saturation* to establish the functional correctness of the recursive implementation of Quicksort from Figure 1 by means of automated first-order reasoning. In a nutshell, we proceed as follows.

(i) We formalize the *semantics of functional programs* in extensions of the first-order theory of lists (Section 3), allowing us to quantify over lists. Rather than focusing on lists with a specific background theory, such as integers/naturals, our formalization relies on a parameterized sort/type a abstracting specific (arithmetic) theories. To this end, we impose that the sort a has a linear order \leq . Doing so, we remark that one of the major reasoning burdens towards establishing the correctness of sorting algorithms comes with formalizing permutation properties, for example that two lists are permutations of each other. Universally quantifying over permutations of lists is, however, not a first-order property since it requires quantification over predicates. Hence, reasoning about list permutations demand higher-order logic. Further, while counting and comparing the number of list elements is a viable option to formalize permutation equivalence in first-order logic, the necessary arithmetic reasoning adds an additional burden to the underlying prover. We overcome this challenge by introducing an effective first-order formalization of *permutation equivalence* over parameterized lists. Our permutation equivalence property encodes *multiset* operations over lists, eliminating the need of counting list elements, and therefore arithmetic reasoning, or fully axiomatizing (higher-order) permutations.

(ii) We revise inductive reasoning in first-order theorem proving (Section 4) and introduce

computation induction as a means to tackle *divide-and-conquer* algorithms. We, therefore, extend the first-order reasoner with an inductive inference based on the *computation induction schema* and outline its necessity for recursive sorting routines.

(iii) We leverage *first-order theorem proving for compositional proofs* of recursive parameterized sorting algorithms (Section 5), in particular of **Quicksort** from Figure 1. By embedding the application of induction directly in saturation-based proving, we automatically discharge manually split proof obligations. Each such condition represents a first-order lemma, and hence a proof step. We emphasize that the only manual effort in our framework comes with splitting formulas into multiple lemmas (Section 6.1); each lemma is established automatically by means of automated theorem proving with built-in induction. That is, all our lemmas/verification conditions are automatically proven by means of structural and/or computation induction during the saturation process. We do not rely on user-provided inductive properties, nor on user guidance to perform proofs by induction, but generate inductive hypotheses/invariants via inductive inferences automatically as logical consequences of our program semantics.

(iv) We note that sorting algorithms often follow a divide-and-conquer approach (see Figure 2). We, thus, apply our approach on other sorting routines and investigate a generalized set of manual proof splits/lemmas that is applicable to verify functional sorting algorithms on recursive data structures (Section 6).

(v) We demonstrate our findings (Section 7) by implementing our approach on top of the **VAMPIRE** theorem prover [15], providing thus a fully automated tool support towards validating the functional correctness of sorting algorithms.

2 Preliminaries

We assume familiarity with standard first-order logic (FOL) and briefly introduce saturation-based proof search in first-order theorem proving [15].

Saturation. Rather than using arbitrary first-order formulae G , most first-order theorem provers rely on a clausal representation of G . The task of first-order theorem proving is to establish that a formula/goal G is a logical consequence of a set \mathcal{A} of clauses, including assumptions. Doing so, first-order provers classify the negation $\neg G$ of G and derive that the set $S = \mathcal{A} \cup \{\neg G\}$ is unsatisfiable¹. To this end, first-order provers *saturate* S by computing all logical consequences of S with respect to some sound inference system \mathcal{I} . A sound inference system \mathcal{I} derives a clause D from clauses C such that $C \rightarrow D$. The saturated set of S w.r.t. \mathcal{I} is called the *closure* of S w.r.t. \mathcal{I} , whereas the process of deriving the closure of S is called *saturation*. By soundness of \mathcal{I} , if the closure of S contains the empty clause \square , the original set S of clauses is unsatisfiable, implying the validity of $\mathcal{A} \rightarrow G$; in this case, we established a *refutation* of $\neg G$ from \mathcal{A} , hence a proof of validity of G .

The *superposition calculus* is a common inference system used by saturation-based provers for FOL with equality [21]. The superposition calculus is sound and *refutationally complete*: for any unsatisfiable formula $\neg G$, superposition-based saturation derives the empty clause \square as a logical consequence of $\neg G$.

Parameterized Lists. We use the first-order theory of recursively defined datatypes [14]. In particular, we consider the list datatype with two constructors `nil` and `cons(x, xs)`, where `nil` is the empty list and x and xs are respectively the head and tail of a list. We introduce a type parameter a that abstracts the sort/background theory of the list elements. Here, we impose the restriction that the sort a has a linear order $<$, that is, a binary relation which

¹for simplicity, we denote by $\neg G$ the classified form of the negation of G

is reflexive, antisymmetric, transitive and total. For simplicity, we also use \geq and \leq as the standard ordering extensions of $<$. We write xs_a, ys_a, zs_a to mean that the lists xs, ys, zs are parameterized by sort a ; that is, their elements are of sort a . Similarly, we use x_a, y_a, z_a to mean that the list elements x, y, z are of sort a . Whenever it is clear from the context, we omit specifying the sort a .

Function definitions. We make the following abuse of notation. For some function \mathbf{f} in some program P , we use the notation $\mathbf{f}(\mathbf{arg}_1, \dots)$ to refer to function definitions/calls appearing in the input algorithm, while the mathematical notation $f(\mathit{arg}_1, \dots)$ refers to its counterpart in our logical representation as per our first-order semantics introduced in Section 3.

3 First-Order Semantics of Functional Sorting Algorithms

We outline our formalization of recursive sorting algorithms in the full first-order theory of parameterized lists.

3.1 Recursive Functions in First-Order Logic

We investigate recursive algorithms written in a functional coding style and defined over lists using list constructors. That is, we consider recursive functions \mathbf{f} that manipulate the empty list nil and/or the list $\mathit{cons}(x, xs)$.

Many recursive sorting algorithms, as well as other recursive operations over lists, implement a *divide-and-conquer* approach: let \mathbf{f} be a function following such a pattern, \mathbf{f} uses (i) a *partition function* to divide $list_a$, that is a *list* of sort a , into two smaller sublists upon which \mathbf{f} is recursively applied to, and (ii) calls a *combination function* that puts together the result of the recursive calls of \mathbf{f} . Figure 2 shows such a divide-and-conquer pattern, where the partition function $\mathit{partition}$ uses an invertible operator \circ , with \circ^{-1} being the complement of \circ ; \mathbf{f} is applied to the results of \circ before these results are merged using the combination function $\mathit{combine}$.

Note that the recursive function \mathbf{f} of Figure 2 is defined via the declaration $f :: 'alist \rightarrow \dots \rightarrow 'alist$, where \dots denotes further input parameters. We formalize the first-order semantics of \mathbf{f} via the function $f : (list_a \times \dots) \mapsto list_a$, by translating the inductive function definitions \mathbf{f} to the following first-order formulas with parameterized lists (in first-order logic, function definitions can be considered as universally quantified equalities):

```

1 f :: 'a list → ... → 'a list
2 f (nil, ...) = nil
3 f (cons(y, ys), ...) = combine (
4   f (partition ◦ (cons(y, ys))),
5   f (partition ◦-1 (cons(y, ys))))
6

```

Figure 2: Recursive divide-and-conquer approach.

$$\begin{aligned}
f(\mathit{nil}) &= \mathit{nil} \\
\forall x_a, xs_a. f(\mathit{cons}(x, xs)) &= \mathit{combine}(f(\mathit{partition}_\circ(\mathit{cons}(x, xs))), \\
&\quad f(\mathit{partition}_{\circ^{-1}}(\mathit{cons}(x, xs)))).
\end{aligned} \tag{1}$$

The recursive divide-and-conquer pattern of Figure 2, together with the first-order semantics (1) of \mathbf{f} , is used in Sections 5-6 for proving correctness of the **Quicksort** algorithm (and other sorting algorithms), as well as for applying lemma generalizations for divide-and-conquer list

operations. We next introduce our first-order formalization for specifying that \mathbf{f} implements a sorting routine.

3.2 First-Order Specification of Sorting Algorithms

We consider a specific function instance of \mathbf{f} implementing a sorting algorithm, expressed through $sort :: 'a\ list \rightarrow 'a\ list$. The functional behavior of $sort$ needs to satisfy two specifications implying the functional correctness of $sort$: (i) sortedness and (ii) permutations equivalence of the list computed by $sort$.

(i) Sortedness: *The list computed by the sort function must be sorted w.r.t. some linear order \leq over the type a of list elements.* We define a parameterized version of this sortedness property using an inductive predicate $sorted$ as follows:

$$\begin{aligned} sorted(\text{nil}) &= \top \\ \forall x_a, xs_a. sorted(\text{cons}(x, xs)) &= (elem_{\leq}list(x, xs) \wedge sorted(xs)), \end{aligned} \tag{2}$$

where $elem_{\leq}list(x, xs)$ specifies that $x \leq y$ for any element y in xs . Proving correctness of a sorting algorithm $sort$ thus reduces to proving the validity of:

$$\forall xs_a. sorted(sort(xs)). \tag{3}$$

(ii) Permutation Equivalence: *The list computed by the sort function is a permutation of the input list to the sort function.* In other words the input and output lists of $sort$ are permutations of each other, in short permutation equivalent.

Axiomatizing permutations requires quantification over relations and is thus not expressible in first-order logic [17]. A common approach to prove permutation equivalence of two lists is to count the occurrences of each element in both lists respectively and compare these numbers per list element. Yet, counting adds a burden of arithmetic reasoning over naturals to the underlying prover, calling for additional applications of mathematical induction. We overcome these challenges of expressing permutation equivalence as follows. We introduce a family of functions $filter_Q$ manipulating lists, with the function $filter_Q$ being parameterized by a predicate Q and as given in Figure 3.

```

1 filter_Q :: 'a → 'a list → 'a list
2 filter_Q(x, nil) = nil
3 filter_Q(x, cons(y, ys)) =
4   if (Q(y, x)) {
5     cons(y, filter_Q(x, ys))
6   } else {
7     filter_Q(x, ys)
8   }
```

Figure 3: Functions $filter_Q$ filtering elements of a list, by using a predicate $Q(y, x)$ over list elements x, y .

In particular, given an element x and a list ys , the functions $filter_{=}$, $filter_{<}$, and $filter_{\geq}$ compute the maximal sublists of ys that contain only equal, resp. smaller and greater-or-equal elements to x . Analogously to counting the multiset multiplicity of x in ys via counting functions, we compare lists given by $filter_{=}$, avoiding the need to count the number of occurrences of x and

hence prolific axiomatizations of arithmetic. Thus, to prove that the input/output lists of $sort$ are permutation equivalent, we show that, for every list element x , the results of applying $filter_{=}$ to the input/output list of $sort$ are the same over all elements. Formally, we have the following first-order property of permutation equivalence:

$$\forall x_a, xs_a. filter_{=}(x, xs) = filter_{=}(x, sort(xs)). \tag{4}$$

4 Computation Induction in Saturation

In this section, we describe our reasoning extension to saturation-based first-order theorem proving, in order to support inductive reasoning for recursive sorting algorithms as introduced in Section 3. Our key reasoning ingredient comes with a structural induction schema of *computation induction*, which we directly integrate in the saturation proving process.

Inductive reasoning has recently been embedded in saturation-based theorem proving [10], by extending the superposition calculus with a new inference rule based on *induction axioms*:

$$(\text{Ind}) \frac{\neg L[t] \vee C}{\text{cnf}(\neg F \vee C)} \quad \text{where} \quad \begin{array}{l} (1) L[t] \text{ is a quantifier-free (ground) literal,} \\ (2) F \rightarrow \forall x.L[x] \text{ is a valid } \textit{induction axiom}, \\ (3) \text{cnf}(\neg F \vee C) \text{ is the clausal form of } \neg F \vee C. \end{array}$$

An *induction axiom* refers to an instance of a valid induction schema. In our work, we use structural and computational induction schemata. In particular, we use the following *structural induction* schema over lists:

$$(L[\text{nil}] \wedge \forall x, ys.(L[ys] \rightarrow L[\text{cons}(x, ys)])) \rightarrow \forall zs.L[zs] \quad (5)$$

Then, considering the induction axiom resulting from applying schema (5) during saturation to L , we obtain the following *Ind* instance for lists:

$$\frac{\neg L[t] \vee C}{\begin{array}{l} \neg L[\text{nil}] \vee L[\sigma_{ys}] \vee C \\ \neg L[\text{nil}] \vee \neg L[\text{cons}(\sigma_x, \sigma_{ys})] \vee C \end{array}}$$

where t is a ground term of sort list, $L[t]$ is ground, and σ_x and σ_{ys} are fresh constant symbols, yielding two clauses as conclusions.

Sorting algorithms, however, often require induction axioms that are more complex than instances of structural induction (5). Such axioms are typically instances of the computation/recursion induction schema, arising from divide-and-conquer strategies as introduced in Section 3.1. Particularly, the complexity arises due to the two recursive calls on different parts of the original input list produced by the *partition* function that have to be taken into account by the induction schema. We therefore use the following *computation induction* schema over lists:

$$\left(L[\text{nil}] \wedge \forall x, ys. \left(\left(\frac{L[\text{partition}_o(x, ys)] \wedge L[\text{partition}_{o-1}(x, ys)]}{L[\text{cons}(x, ys)]} \right) \right) \right) \rightarrow \forall zs.L[zs] \quad (6)$$

yielding the following instance of *Ind* that is applied during saturation:

$$\frac{\neg L[t] \vee C}{\begin{array}{l} \neg L[\text{nil}] \vee L[\text{partition}_o(\sigma_x, \sigma_{ys})] \vee C \\ \neg L[\text{nil}] \vee L[\text{partition}_{o-1}(\sigma_x, \sigma_{ys})] \vee C \\ \neg L[\text{nil}] \vee \neg L[\text{cons}(\sigma_x, \sigma_{ys})] \vee C \end{array}}$$

where t is a ground term of sort list, $L[t]$ is ground, σ_x and σ_{ys} are fresh constant symbols, and partition_o and its complement refer to the functions that partition lists into sublists within the actual sorting algorithms.

5 Proving Recursive Quicksort

We now describe our approach towards proving the correctness of the recursive parameterized version of **Quicksort**, as given in Figure 1. Note that **Quicksort** recursively sorts two sublists

that contain respectively smaller and greater-or-equal elements than the pivot element x of its input list. We reduce the task of proving the functional correctness of `Quicksort` to the task of proving the (i) sortedness property (3) and (ii) the permutation equivalence property (4) of `Quicksort`. As mentioned in Section 3.2, a similar reasoning is needed for most sorting algorithms, which we evidence in Sections 6–7.

5.1 Proving Sortedness for Quicksort

Given an input list xs , we prove that `Quicksort` computes a sorted list by considering the property (3) instantiated for `Quicksort`. That is, we prove:

$$\forall xs_a. \text{sorted}(\text{quicksort}(xs)) \quad (7)$$

The sortedness property (7) of `Quicksort` is proved via *compositional reasoning* over (7). Namely, we enforce the following two properties that together imply (7):

(S1) By using the linear order \leq of the background theory a , for any element y in the sorted list $\text{quicksort}(\text{filter}_{<}(x, xs))$ and any element z in the sorted list $\text{quicksort}(\text{filter}_{\geq}(x, xs))$, we have $y \leq x \leq z$.

(S2) The functions $\text{filter}_{<}$ and filter_{\geq} of Figure 3 are correct. That is, filtering elements from a list that are smaller, respectively greater-or-equal, than an element x results in sublists only containing elements smaller than, respectively greater-or-equal, than x .

Similarly to (2) and to express property (S2), we introduce the inductively defined predicates $\text{elem}_{\leq} \text{list} :: 'a \rightarrow 'a \text{list} \rightarrow \text{bool}$ and $\text{list}_{\leq} \text{list} :: 'a \text{list} \rightarrow 'a \text{list} \rightarrow \text{bool}$:

$$\begin{aligned} \forall x_a. \text{elem}_{\leq} \text{list}(x, \text{nil}) &= \top \\ \forall x_a, y_a, ys_a. \text{elem}_{\leq} \text{list}(x, \text{cons}(y, ys)) &= x \leq y \wedge \text{elem}_{\leq} \text{list}(x, ys), \end{aligned} \quad (8)$$

and

$$\begin{aligned} \forall ys_a. \text{list}_{\leq} \text{list}(\text{nil}, ys) &= \top \\ \forall x_a, xs_a, ys_a. \text{list}_{\leq} \text{list}(\text{cons}(x, xs), ys) &= (\text{elem}_{\leq} \text{list}(x, ys) \wedge \text{list}_{\leq} \text{list}(xs, ys)). \end{aligned} \quad (9)$$

Thus, for some element x and lists xs, ys , we express that x is smaller than or equal to any element of xs by $\text{elem}_{\leq} \text{list}(x, xs)$. Similarly, $\text{list}_{\leq} \text{list}(xs, ys)$ states that every element in list xs is smaller than or equal to any element in ys .

The inductively defined predicates of (8)–(9) allow us to express necessary lemmas over list operations preserving the sortedness property (7), for example, to prove that appending sorted lists yields a sorted list.

Proving properties (S1)–(S2), and hence deriving the sortedness property (7) of `Quicksort`, requires *three first-order lemmas* in addition to the first-order semantics (1) of `Quicksort`. Each of these lemmas is automatically proven by saturation-based theorem proving using the structural and/or computation induction schemata of (5) and (6); hence, by compositionality, we obtain (S1)–(S2) implying (7). We next discuss these three lemmas and outline the complete (compositional) proof of the sortedness property (7) of `Quicksort`.

- In support of (S1), lemma (10) expresses that for two *sorted* lists xs, ys and a list element x , such that $\text{elem}_{\leq} \text{list}(x, xs)$ holds and all elements of the constructed list $\text{cons}(x, xs)$ are greater than or equal to all elements in ys , the result of concatenating ys and $\text{cons}(x, xs)$ yields a sorted

list. Formally, we have

$$\begin{aligned} \forall x_a, xs_a, ys_a. & \quad (sorted(xs) \wedge sorted(ys) \wedge elem_{\leq}list(x, xs) \wedge \\ & \quad list_{\leq}list(ys, cons(x, xs))) \\ & \rightarrow sorted(append(ys, cons(x, xs))) \end{aligned} \quad (10)$$

• In support of **(S2)**, we need to establish that filtering greater-or-equal elements for some list element x results in a list whose elements are greater-or-equal than x . In other words, the inductive predicate of **(8)** is invariant over sorting and filtering operations over lists.

$$\forall x_a, xs_a. elem_{\leq}list(x, quicksort(filter_{\geq}(x, xs))). \quad (11)$$

• Lastly and in further support of **(S1)–(S2)**, we establish that all elements of a list xs are “covered” with the filtering operations $filter_{\geq}$ and $filter_{<}$ w.r.t. a list element x of xs . Intuitively, a call of $filter_{<}(x, xs)$ results in a list containing all elements of xs that are smaller than x , while the remaining elements of xs are those that are greater-or-equal than x and hence are contained in $cons(x, filter_{\geq}(x, xs))$. By applying **Quicksort** over the input list xs , we get:

$$\forall x_a, xs_a. list_{\leq}list(quicksort(filter_{<}(x, xs)), cons(x, quicksort(filter_{\geq}(x, xs)))). \quad (12)$$

The first-order lemmas **(10)–(12)** guide saturation-based proving to instantiate structural/-computation induction schemata and automatically derive the following induction axiom necessary to prove **(S1)–(S2)**, and hence sortedness of **Quicksort**:

$$\begin{aligned} & \left(sorted(quicksort(nil)) \wedge \right. \\ & \quad \left. \forall x_a, xs_a. \left(\begin{array}{l} sorted(quicksort(filter_{\geq}(x, xs))) \wedge \\ sorted(quicksort(filter_{<}(x, xs))) \end{array} \right) \rightarrow sorted(quicksort(cons(x, xs))) \right) \\ & \rightarrow \forall xs_a. sorted(quicksort(xs)), \end{aligned} \quad (13)$$

where axiom **(13)** is automatically obtained during saturation from the computation induction schema **(6)**. Intuitively, the prover replaces F by $sorted(quicksort(t))$ for some term t , and uses $filter_{<}$ and $filter_{\geq}$ as $partition_{\circ}$ and $partition_{\circ^{-1}}$ respectively to find the necessary computation induction schema. We emphasize that this step is fully automated during the saturation run.

The first-order lemmas **(10)–(12)**, together with the induction axiom **(13)** and the first-order semantics **(1)** of **Quicksort**, imply the sortedness property **(4)** of **Quicksort**; this proof can automatically be derived using saturation-based reasoning. Yet, the obtained proof assumes the validity of each of the lemmas **(10)–(12)**. To eliminate this assumption, we propose to also prove lemmas **(10)–(12)** via saturation-based reasoning. Yet, while lemma **(10)** is established by saturation with structural induction **(5)** over lists, proving lemmas **(11)–(12)** requires further first-order formulas. In particular, for proving lemmas **(11)–(12)** via saturation, we use four further lemmas, as follows.

• Lemmas **(14)–(15)** indicate that the order of $elem_{\leq}list$ and $list_{\leq}list$ is preserved under $quicksort$, respectively. That is,

$$\forall x_a, xs_a. elem_{\leq}list(x, xs) \rightarrow elem_{\leq}list(x, quicksort(xs)) \quad (14)$$

and

$$\forall xs_a, ys_a. list_{\leq}list(ys, xs) \rightarrow list_{\leq}list(quicksort(ys), xs). \quad (15)$$

• Proving lemmas (14)–(15), however, requires two further lemmas that follow from saturation with built-in computation and structural induction, respectively. Namely, lemmas (16)–(17) establish that $elem_{\leq}list$ and $list_{\leq}list$ are also invariant over appending lists. That is,

$$\begin{aligned} \forall x_a, y_a, xs_a, ys_a. (y \leq x \wedge elem_{\leq}list(y, xs) \wedge elem_{\leq}list(y, ys)) \\ \rightarrow elem_{\leq}list(y, append(cons(x, ys), xs)) \end{aligned} \quad (16)$$

and

$$\begin{aligned} \forall xs_a, ys_a, zs_a. (list_{\leq}list(ys, xs) \wedge list_{\leq}list(zs, xs)) \\ \rightarrow list_{\leq}list(append(ys, zs), xs) \end{aligned} \quad (17)$$

With lemmas (14)–(17), we automatically prove lemmas (10)–(12) via saturation-based reasoning. The complete automation of proving properties (S1)–(S2), and hence deriving the sortedness property (7) of `Quicksort` in a compositional manner, requires thus *altogether seven lemmas* in addition to the first-order semantics (1) of `Quicksort`. *Each of these lemmas is automatically established via saturation with built-in induction.* Hence, unlike interactive theorem proving, compositional proving with first-order theorem provers can be leveraged to eliminate the need to a priori specifying necessary induction axioms.

5.2 Proving Permutation Equivalence for Quicksort

In addition to establishing the sortedness property (7) of `Quicksort`, the functional correctness of `Quicksort` also requires proving the permutation equivalence property (4) for `Quicksort`. That is, we prove:

$$\forall x_a, xs_a. filter_{=}(x, xs) = filter_{=}(x, quicksort(xs)). \quad (18)$$

In this respect, we follow the approach introduced in Section 3.2 to enable first-order reasoning over permutation equivalence (18). Namely, we use $filter_{=}$ to filter elements x in the lists xs and $quicksort(xs)$, respectively, and build the corresponding multisets containing as many x as x occurs in xs and $quicksort(xs)$. By comparing the resulting multisets, we implicitly reason about the number of occurrences of x in xs and $quicksort(xs)$, yet, without the need to explicitly count occurrences of x . In summary, we reduce the task of proving (18) to *compositional reasoning* again, namely to proving following *two properties given as first-order lemmas* which, by compositionality, imply (18):

(P1) List concatenation commutes with $filter_{=}$, expressed by the lemma:

$$\forall x_a, xs_a, ys_a. filter_{=}(x, append(xs, ys)) = append(filter_{=}(x, xs), filter_{=}(x, ys)). \quad (19)$$

(P2) Appending the aggregate of both $filter$ -operations results in the same multisets as the unfiltered list, that is, permutation equivalence is invariant over combining complementary reduction operations. This property is expressed via:

$$\begin{aligned} \forall x_a, y_a, xs_a. filter_{=}(x, xs) = append(filter_{=}(x, filter_{<}(y, xs)), \\ filter_{=}(x, filter_{\geq}(y, xs))). \end{aligned} \quad (20)$$

Similarly as in Section 5.1, we prove lemmas (P1)–(P2) by saturation-based reasoning with built-in induction. In particular, investigating the proof output shows that lemma (P1) is established using the structural induction schema (5) in saturation, while the validity of lemma (P2) is obtained by applying the computation induction schema (6) in saturation.

```

1 mergesort :: 'a list → 'a list
2 mergesort(nil) = nil
3 mergesort(xs) = merge(mergesort(take((xslength div 2), xs)),
   mergesort(drop((xslength div 2), xs)))
4
5 merge :: 'a list → 'a list → 'a list
6 merge(nil, ys) = ys
7 merge(xs, nil) = xs
8 merge(cons(x, xs), cons(y, ys)) =
9   if (x ≤ y) {
10    cons(x, merge(xs, cons(y, ys)))
11  } else {
12    cons(y, merge(cons(x, xs), ys))
13  }
14

```

Figure 4: Recursive Mergesort over lists of sort a .

By proving lemmas **(P1)**–**(P2)**, we thus establish validity of permutation equivalence (18) for Quicksort. Together with the sortedness property (7) of Quicksort proven in Section 5.1, we conclude the functional correctness of Quicksort in a fully automated and compositional manner, using saturation-based theorem proving with built-in induction and *altogether nine first-order lemmas* in addition to the first-order semantics (1) of Quicksort.

6 Compositional Reasoning and Lemma Generalizations

Establishing the functional correctness of Quicksort in Section 5 uses nine first-order lemmas that express inductive properties over lists in addition to the first-order semantics (1) of Quicksort. While each of these lemmas is proved by saturation using structural/computation induction schemata, coming up with proper inductive lemmas remains crucial in reasoning about inductive data structures, and, so far, dependent on user guidance. We thus discuss the intuition on manual proof splitting in Section 6.1 and generalize our efforts for sorting algorithms in Section 6.2.

6.1 Guiding Proof Splitting

Contrary to automated approaches that use inductive annotations to alleviate inductive reasoning, our approach synthesizes the correct induction axioms automatically during saturation runs to prove properties and lemmas correct. However, a manual limitation remains, namely deciding when a lemma is necessary or helpful for the automated reasoner.

Splitting the proof into multiple lemmas is necessary to guide the prover to find the right terms to apply the inductive inferences of Section 4. This is particularly the case when input problems, such as sorting algorithms, contain calls to multiple recursive functions – each of which has to be shown to preserve the property that is to be verified.

In case a proof fails, we investigate the synthesized induction axioms, manually strengthen the property and add any additional assumptions as proof obligations whose validity is in turn again verified with the theorem prover and built-in induction. That is, we do not simply assume inductive lemmas but also provide a formal argument of their validity.

We illustrate and examine the need for proof splitting using lemma (10).

Example 1 (Compositional reasoning over sortedness in saturation). Consider the following stronger version of lemma (10) in the proof of `Quicksort`:

$$\forall x_a, xs_a, ys_a. (sorted(xs) \wedge sorted(ys)) \rightarrow sorted(append(ys, cons(x, xs))). \quad (21)$$

This formula was automatically be derived by saturation with computation induction (6) while trying to prove sortedness of the algorithm. However, formula (21) is not helpful for the proof of `Quicksort` since it is not inductive with regards to the specification and thus cannot be resolved and used during proof search. The prover needs additional information to verify sortedness. Therefore, the assumptions $elem_{\leq}list(x, xs)$ and $list_{\leq}list(ys, cons(x, xs))$ are needed in addition to (21), resulting in lemma (10). Yet, lemma (10) from Section 5 can be automatically derived via saturation based on computation induction (6). That is, we manually split proof obligations based on missing information in the saturation runs: we derive (21) from (6) via saturation, strengthen the hypotheses of (21) with missing necessary conditions $elem_{\leq}list(x, xs)$ and $list_{\leq}list(ys, cons(x, xs))$, and prove their validity via saturation, yielding (10).

6.2 Lemma Generalizations for Sorting

The lemmas from Section 5 represent a number of common proof splits that can be applied to various list sorting tasks. In the following we generalize their structure and apply them to two other sorting algorithms, namely `Mergesort` and `Insertionsort`.

Common Patterns of Inductive Lemmas for Sorting Algorithms.

Consider the computation induction schema (6). When using (6) for proving the sortedness (7) and permutation equivalence (18) of `Quicksort`, the inductive formula F of (6) is, respectively, instantiated with the predicates $sorted$ from (7) and $filter_{=}$ from (18). The base case $F[[]]$ of schema (6) is then trivially proved by saturation for both properties (7) and (18) of `Quicksort`.

Proving the induction step case of schema (6) is however challenging as it relies on $partition$ -functions which are further used by $combine$ functions within the divide-and-conquer patterns of Figure 2.

Intuitively this means, that proving the induction step case of schema (6) for the sortedness (7) and permutation equivalence (18) properties requires showing that applying $combine$ functions over $partition$ functions preserve sortedness (7) and permutation equivalence (18), respectively. For divide-and-conquer algorithms of Figure 2, the step case of schema (6) requires thus proving the following lemma:

$$\left(\forall x_a, ys_a. \left(combine \left(\begin{array}{l} L[partition_{\circ}(x, ys)], \\ L[partition_{\circ^{-1}}(x, ys)] \end{array} \right) \rightarrow L[cons(x, ys)] \right) \right). \quad (22)$$

We next describe generic instances of lemmas to be used to prove such step cases and hence functional correctness of sorting algorithms, and exemplify our findings from `Quicksort` by

```

1  insertsort :: 'a list → 'a list
2  insertsort nil = nil
3  insertsort (cons(x, xs)) = insert(x,
      insertsort xs)
4
5  insert :: 'a → 'a list → 'a list
6  insert(x, nil) = cons(x, nil)
7  insert(x, cons(y, ys)) =
8    if (x ≤ y) {
9      cons(x, cons(y, ys))
10   } else {
11     cons(y, insert(x, ys))
12   }
13

```

Figure 5: Recursive algorithm of `Insertionsort`.

application to recursive versions of `Mergesort` and `Insertionsort` given in Figures 4 and 5, respectively.

(i) **Combining sorted lists preserves sortedness.** For proving the inductive step case (22) of the sortedness property (3) of sorting algorithms, we require the following generic lemma (23):

$$\forall xs_a, ys_a. (sorted(xs) \wedge sorted(ys)) \rightarrow sorted(combine(xs, ys)), \quad (23)$$

ensuring that combining sorted lists results in a sorted list. Lemma (23) is used to establish property (S1) of `Quicksort`, namely used as lemma (10) for proving the preservation of sortedness under the `append` function.

We showcase the generality of lemma (23) with `Mergesort` as given in Figure 4. The sortedness property (3) of `Mergesort` can be proved by using saturation with lemma 23; note that the `merge` function acts as a `combine` function of (23). We thus establish sortedness via the following instance of (23):

$$\forall xs_a, ys_a. sorted(xs) \wedge sorted(ys) \rightarrow sorted(merge(xs, ys))$$

Finally, lemma (23) is not purely restricted to divide-and-conquer routines. When proving the sortedness property (3) of the recursive `Insertionsort` algorithm of Figure 5, we apply lemma (23) on `insert` to establish preservation of sortedness with saturation:

$$\forall x_a, xs_a. sorted(xs) \rightarrow sorted(insert(x, xs))$$

(ii) **Combining partitions preserves permutation equivalence.** Similarly to Section 5.2, proving permutation equivalence (4) over divide-and-conquer sorting algorithms of Figure 2 is established via the following two properties:

- As in (P1) for `Quicksort`, we require that `combine` commutes with `filter=`:

$$\forall x_a, xs_a, ys_a. filter_{=}(x, combine(xs, ys)) = combine(filter_{=}(x, xs), filter_{=}(x, ys)) \quad (24)$$

- Similarly to (P2) for `Quicksort`, we ensure that, by combining (complementary) `partition` functions, we preserve (4). That is,

$$\forall x_a, xs_a. filter_{=}(x, xs) = combine(filter_{=}(x, partition_o(xs)), filter_{=}(x, partition_{o-1}(xs))) \quad (25)$$

Note that lemmas (P1) and (P2) for `Quicksort` are instances of (24) and (25) respectively, as the `append` function of `Quicksort` acts as a `combine` function and the `filter<` and `filter≥` functions are the `partition` functions of Figure 2.

To prove the permutation equivalence (4) property of `Mergesort`, we use the functions `take` and `drop` as the `partition` functions of lemmas (24)–(25). Doing so, we embed a natural number argument n in lemmas (24)–(25), with n controlling how many list elements are *taken* and *dropped*, respectively, in `Mergesort`. As such, the following instances of lemmas (24)–(25) are adjusted to `Mergesort`:

$$\forall x_a, xs_a, ys_a. filter_{=}(x, merge(xs, ys)) = append(filter_{=}(x, xs), filter_{=}(x, ys)),$$

and

$$\forall x_a, n_{\mathbb{N}}, xs_a. filter_{=}(x, xs) = append(filter_{=}(x, take(n, xs)), filter_{=}(x, drop(n, xs))),$$

PermEq				Sortedness			
Benchm.	Pr.	T	Required lemmas	Benchm.	Pr.	T	Required lemmas
IS-PE	✓	0.02	{IS-PE-L1}	IS-S	✓	0.01	{IS-S-L1}
IS-PE-L1	✓	0.13	∅	IS-S-L1	✓	8.28	-
MS-PE	✓	0.06	{MS-PE-L1, MS-PE-L2}	MS-S	✓	0.08	∅
MS-PE-L1	✓*	0	-	MS-S-L1	✓*	0	-
MS-PE-L2	✓	0.03	∅	MS-S-L2	✓	0.02	∅
MS-PE-L3	✓	0.15	∅				
QS-PE	✓	0.5	{QS-PE-L1, QS-PE-L2}	QS-S	✓	0.09	{QS-S-L1, QS-S-L2, QS-S-L3}, {QS-S-L1, QS-S-L3, QS-S-L4}
QS-PE-L1	✓	0.05	∅	QS-S-L1	✓	0.27	∅
QS-PE-L2	✓	0.09	∅	QS-S-L2	✓	0.04	{QS-S-L4}
				QS-S-L3	✓	11.82	{QS-S-L4, QS-S-L5}
				QS-S-L4	✓	8.28	{QS-S-L6}
				QS-S-L5	✓	0	{QS-S-L7}
				QS-S-L6	✓	0.02	∅
				QS-S-L7	✓	0.02	∅

Table 1: Experimental evaluation of proving properties of sorting algorithms, using a time limit of 5 minutes on machine with AMD Epyc 7502, 2.5 GHz CPU with 1 TB RAM, using 1 core and 16 GB RAM per benchmark.

IS, MS and QS correspond to Insertionsort, Mergesort and Quicksort; S and PE respectively denote sortedness (3) and permutation equivalence (4), and Li stands for the i -th lemma of the problem.

with both lemmas being proved via saturation, thus establishing permutation equivalence (4).

Finally, the generality of lemmas (24)–(25) naturally pays off when proving the permutation equivalence property (4) of Insertionsort. Here, we only use a simplified instance of (24) to prove (4) is preserved by the auxiliary function `insert`. That is, we use the following instance of (24):

$$\forall x_a, y_a, ys_a. filter_=(x, cons(y, ys)) = filter_=(x, insert(y, ys)),$$

which is automatically derivable by saturation with computation induction (6).

We conclude by emphasizing that compositional reasoning in saturation with computation induction enables us to prove challenging sorting algorithms in a newly automated manner, replacing the manual effort in carrying out proofs by induction.

7 Implementation and Experiments

Implementation. Our work on saturation with induction in the first-order theory of parameterized lists is implemented in the first-order prover VAMPIRE [15]. In support of parameterization, we extended the SMT-LIB parser of VAMPIRE to support parametric data types from SMT-LIB [2] – version 2.6. In particular, using the `par` keyword, our parser interprets (`par (a1 ... an) ...`) similar to universally quantified blocks where each variable a_i is a type parameter.

Appropriating a generic saturation strategy, we adjust the simplification orderings (LPO) for efficient equality reasoning/rewrites to our setting. For example, the precedence of function `quicksort` is higher than of symbols `nil`, `cons`, `append`, `filter<` and `filter≥`, ensuring that `quicksort` function terms are expanded to their functional definitions. We further apply recent results of encompassment demodulation [5] to improve equality reasoning within saturation (`-drc encompass`). We use induction on data types (`-ind struct`), including complex data type terms (`-indoct on`).

Experimental Evaluation. We evaluated our approach over challenging recursive sorting algorithms taken from [20], namely `Quicksort`, `Mergesort`, and `Insertionsort`. We show that the functional correctness of these sorting routines can be verified automatically by means of saturation-based theorem proving with induction, as summarized in Table 1.

We divide our experiments according to the specification of sorting algorithms: the first column `PermEq` shows the experiments of all sorting routines w.r.t. permutation equivalence (4), while `Sortedness` refers to the sortedness (3) property, together implying the functional correctness of the respective sorting algorithm. Here, the inductive lemmas of Sections 5–6 are proven in separate saturation runs of `VAMPIRE` with structural/computation induction; these lemmas are then used as input assumptions to `VAMPIRE` to prove validity of the respective benchmark.² A benchmark category `SA-PR[-Li]` indicates that it belongs to proving the property `PR` for sorting algorithm `SA`, where `PR` is one of `S` (sortedness (3)) and `PE` (permutation equivalence (4)) and `SA` is one of `IS` (`Insertionsort`), `MS` (`Mergesort`) and `QS` (`Quicksort`). Additionally, an optional `Li` indicates that the benchmark corresponds to the i -th lemma for proving the property of the respective sorting algorithm.

For our experiments, we ran all possible combinations of lemmas to determine the minimal lemma dependency for each benchmark. For example, the sortedness property of `Quicksort` (`QS-S`) depends on seven lemmas (see Section 5.1), while the third lemma for this property (`QS-S-L3`) depends on four lemmas (see Section 5.2). The second column `Pr.` indicates that `VAMPIRE` solved the benchmark by using a minimal subset of needed lemmas given in the fourth column. The third column `T` shows the running time in seconds of the respective saturation run using the first solving strategy identified during portfolio mode.

To identify the successful configuration, we ran `VAMPIRE` in a portfolio setting for 5 minutes on each benchmark, with strategies enumerating all combinations of options that we hypothesized to be relevant for these problems. In accordance with Table 1, `VAMPIRE` compositionally proves permutation equivalence of `Insertionsort` and `Quicksort` and sortedness of `Mergesort` and `Quicksort`. Note that sortedness of `Mergesort` is proven without any lemmas, hence lemma `MS-S-L1` is not needed. The lemmas `MS-PE-L1` for the permutation equivalence of `Mergesort` and `IS-S-L1` for the sortedness of `Insertionsort` could be proven separately by more tailored search heuristics in `VAMPIRE` (hence \checkmark^*), but our cluster setup failed to consistently prove these in the portfolio setting.

Generated Inductive Inference during Proof Search. For all conjectures and lemmas that were proved in portfolio mode, we summarized the applications of inductive inferences with structural and computation induction schemata in Table 2. Specifically, Table 2 compares the number of inductive inferences performed during proof search (column `IndProofSearch`) with the number of used inductive inferences as part of each benchmark’s proof (column `IndProof`). While most safety properties and lemmas required less than 50 inductive inferences, thereby using mostly one or two of them in the proof, some lemma proofs exceeded this by far. Most notably `IS-S-L1` and `QS-S-L1`, `Insertionsort`’s and `Quicksort`’s first lemma respectively, depended on many more inductive inferences until the right axiom was found. Such statistics point to areas where the prover still has room to be finetuned for software verification and quality assurance purposes, here especially towards establishing correctness of functional programs.

²Benchmarks and instructions to run the experiments can be found at https://github.com/mina1604/sorting_wo_sorts.

Benchmark	IndProofSearch	IndProof
IS-S	4	1
IS-S-L1	339	2
IS-PE	5	1
IS-PE-L1	34	1
MS-S	8	1
MS-S-L2	22	1
MS-PE	14	1
MS-PE-L2	16	1
MS-PE-L3	136	3
QS-S	10	2
QS-S-L1	510	2
QS-S-L2	9	1
QS-S-L3	130	2
QS-S-L4	183	3
QS-S-L5	0	0
QS-S-L6	26	1
QS-S-L7	16	2
QS-PE	12	1
QS-PE-L1	10	1
QS-PE-L2	42	4

Table 2: Applications of structural induction in proof search and proofs.

8 Related Work

While `Quicksort` has been proven correct on multiple occasions, first and foremost in the famous 1971 pen-and-paper proof by Foley and Hoare [7], not many have investigated a fully automated proof of the algorithm. A partially automated proof of `Quicksort` relies on `Dafny` [18], where loop invariants are manually provided [4]. While [4] claims to prove some of the lemmas/invariants, not all invariants are proved correct (only assumed to be so). Similarly, the `Why3` framework [6] has been leveraged to prove sortedness and permutation equivalence of `Mergesort` [19] over parameterized lists and arrays. These proofs also rely on manual proof splitting with the additional overhead of choosing the underlying prover for each subgoal as `Why3` is interfaced with both automated and interactive provers.

The work of [24] reports on the verification of functional implementations of multiple sorting algorithms with `VeriFun` [23]. Specifically, the correctness of the sortedness property of `Quicksort` is established with the help of 13 auxiliary lemmas while also establishing the permutation property of `Mergesort` by comparing the number of elements, thus requiring additional arithmetic reasoning. In contrast, our proofs involve less auxiliary lemmas, avoid the overhead of arithmetic theories through our formalization of the permutation property over set equivalence and prove functional implementations with arbitrary sorts permitting a linear order.

The approach of [22] establishes the correctness of permutation equivalence for multiple sorting algorithms based on separation logic through inductive lemmas. However, [22] does not address the correctness proofs of the sortedness property. We instead automate the correctness proofs of sorting algorithms via compositional first-order reasoning in the theory of parameterized lists.

Verifying functional correctness of sorting routines has also been explored in the abstract

interpretation and model-checking communities, by investigating array-manipulating programs [9, 12]. In [9], the authors automatically generate loop invariants for standard sorting algorithms of arrays of fixed length; the framework is, however, restricted solely to inner loops and does not handle recursive functions. Further, in [12] a priori given invariants/interpolants are used in the verification process. Unlike these techniques, we do not rely on a user-provided inductive invariant, nor are we restricted to inner loops.

There are naturally many examples of proofs of sorting algorithms using interactive theorem proving, see e.g. [13, 16]. The work of [13] establishes correctness of insertion sort. Similarly, the setting of [16] proves variations of `Introsort` and `Pdqsort` – both using `Isabelle/HOL` [25]. However, interactive provers rely on user guidance to provide induction schemes, a burden that we eliminate in our approach.

A verified a real-world implementation of `Quicksort` is given in [3]. Here, Java’s inbuilt dual pivot `Quicksort` class is verified with the semi-automatic `KeY` prover [1]. Additionally, the `KeY` prover has also been leveraged to analyze industrial implementations of `Radixsort` and `Countingsort` [8]. By relying on inductive method annotations, such as loop invariants or method contracts, and asking the user to guide the proof rule application during the verification process, the work faces similar limitations as the ones using `Dafny`, `Why3` or interactive proving. While we manually split our proofs into multiple steps, our lemmas are proved automatically thanks to saturation-based theorem proving with structural/computation induction. As such, we do not require guidance on rule application or inductive annotations.

When it comes to the landscape of automated saturation-based reasoning, we are not aware of other techniques enabling the fully automated verification of such sorting routines, with or without compositional reasoning.

9 Conclusion and Future Work

We present an integrated formal approach to establish program correctness over recursive programs based on saturation-based theorem proving. We automatically prove recursive sorting algorithms, particularly the `Quicksort` algorithm, by formalizing program semantics in the first-order theory of parameterized lists. Doing so, we expressed the common properties of sortedness and permutation equivalence in an efficient way for first-order theorem proving. By leveraging common structures of divide-and-conquer sorting algorithms, we advocate compositional first-order reasoning with built-in structural/computation induction.

We believe the implications of our work are twofold. First, integrating inductive reasoning in automated theorem proving to prove (sub)goals during interactive theorem proving can significantly alleviate the use of proof obligations to be shown manually, since automated theorem proving from our work can synthesize induction hypotheses to verify these conditions. Second, finding reasonable strategies to automatically split proof obligations on input problems can tremendously enhance the degree of automation in proofs that require heavy inductive reasoning. We hope that our work opens up future directions in combining interactive and automated reasoning by further decreasing the amount of manual work in proof splitting, allowing superposition frameworks to be better applicable to a wider range of recursive algorithms. Proving further recursive sorting/search algorithms in future work, with improved compositionality, is therefore an interesting challenge to investigate.

Acknowledgements. This research was partially supported by the ERC Consolidator Grant ARTIST 101002685, the TU Wien SecInt Doctoral College, the FWF SFB project SpyCoDe F8504, the WWTF ICT22-007 grant ForSmart, and the Amazon Research Award 2023 QuAT.

References

- [1] Ahrendt, W., Baar, T., Beckert, B., Bubel, R., Giese, M., Hähnle, R., Menzel, W., Mostowski, W., Roth, A., Schlager, S., et al.: The KeY tool: integrating object oriented design and formal verification. *Software & Systems Modeling* **4**, 32–54 (2005)
- [2] Barrett, C., Fontaine, P., Tinelli, C.: The Satisfiability Modulo Theories Library (SMT-LIB). www.SMT-LIB.org (2016)
- [3] Beckert, B., Schiffel, J., Schmitt, P.H., Ulbrich, M.: Proving jdk’s dual pivot quicksort correct. In: *International Conference on Verified Software. Theories, Tools, and Experiments (VSTTE)*. pp. 35–48. Springer (2017)
- [4] Certezeanu, R., Drossopoulou, S., Egelund-Muller, B., Leino, K.R.M., Sivaraman, S., Wheelhouse, M.: Quicksort revisited: Verifying alternative versions of quicksort. *Theory and Practice of Formal Methods: Essays Dedicated to Frank de Boer on the Occasion of His 60th Birthday* pp. 407–426 (2016)
- [5] Duarte, A., Korovin, K.: Ground joinability and connectedness in the superposition calculus. In: *International Joint Conference on Automated Reasoning (IJCAR)*. pp. 169–187. Springer (2022)
- [6] Filliâtre, J.C., Paskevich, A.: Why3 — where programs meet provers. In: *European Symposium on Programming (ESOP)*. pp. 125–128. Springer (2013)
- [7] Foley, M., Hoare, C.A.R.: Proof of a recursive program: Quicksort. *The Computer Journal* **14**(4), 391–395 (1971)
- [8] de Gouw, S., de Boer, F.S., Rot, J.: Verification of counting sort and radix sort. *Deductive Software Verification—The KeY Book: From Theory to Practice* pp. 609–618 (2016)
- [9] Gulwani, S., McCloskey, B., Tiwari, A.: Lifting Abstract Interpreters to Quantified Logical Domains. In: *Symposium on Principles of Programming Languages (POPL)*. pp. 235–246. ACM (2008)
- [10] Hajdu, M., Hozzová, P., Kovács, L., Reger, G., Voronkov, A.: Getting saturated with induction. In: *Principles of Systems Design: Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday*, pp. 306–322. Springer (2022)
- [11] Hoare, C.A.: Quicksort. *The Computer Journal* **5**(1), 10–16 (1962)
- [12] Jhala, R., McMillan, K.L.: Array abstractions from proofs. In: *International Conference on Computer-Aided Verification (CAV)*. pp. 193–206. Springer (2007)
- [13] Jiang, D., Zhou, M.: A comparative study of insertion sorting algorithm verification. In: *IEEE Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*. pp. 321–325. IEEE (2017). <https://doi.org/10.1109/ITNEC.2017.8284998>
- [14] Kovács, L., Robillard, S., Voronkov, A.: Coming to Terms with Quantified Reasoning. In: *Symposium on Principles of Programming Languages (POPL)*. pp. 260–270. ACM (2017)
- [15] Kovács, L., Voronkov, A.: First-Order Theorem Proving and Vampire. In: *International Conference on Computer-Aided Verification (CAV)*. pp. 1–35. Springer (2013)
- [16] Lammich, P.: Efficient verified implementation of introsort and pdqsort. In: *International Joint Conference on Automated Reasoning (IJCAR)*. pp. 307–323. Springer (2020)
- [17] Laneve, C., Montanari, U.: Axiomatizing permutation equivalence. *Mathematical Structures in Computer Science* **6**(3), 219–249 (1996)
- [18] Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: *International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR)*. pp. 348–370. Springer (2010)
- [19] Lévy, J.J.: Simple proofs of simple programs in Why3. *Essays for the Luca Cardelli Fest* p. 177 (2014)
- [20] Nipkow, T., Blanchette, J., Eberl, M., Gómez-Londoño, A., Lammich, P., Sternagel, C., Wimmer, S., Zhan, B.: *Functional Algorithms, Verified* (2021)

- [21] Robinson, A.J., Voronkov, A.: Handbook of automated reasoning, vol. 1. Elsevier (2001)
- [22] Safari, M., Huisman, M.: A generic approach to the verification of the permutation property of sequential and parallel swap-based sorting algorithms. In: International Conference on Integrated Formal Methods (IFM). pp. 257–275. Springer (2020)
- [23] Walther, C., Schweitzer, S.: About VeriFun. In: International Conference on Automated Deduction (CADE). pp. 322–327. Springer (2003)
- [24] Walther, C., Schweitzer, S.: Verification in the classroom. *Journal of Automated Reasoning* **32**, 35–73 (2004)
- [25] Wenzel, M., Paulson, L.C., Nipkow, T.: The Isabelle Framework. In: Theorem Proving in Higher Order Logics (TPHOLs). pp. 33–38. Springer (2008)