Certification of Breadth-First Algorithms by Extraction

Dominique Larchey-Wendling, Ralph Matthes

► To cite this version:

Dominique Larchey-Wendling, Ralph Matthes. Certification of Breadth-First Algorithms by Extraction. 13th International Conference on Mathematics of Program Construction, MPC 2019, Oct 2019, Porto, Portugal. pp.45-75, 10.1007/978-3-030-33636-3_3 . hal-02333423

HAL Id: hal-02333423

<https://hal.science/hal-02333423>

Submitted on 18 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Certification of Breadth-First Algorithms by Extraction

Dominique Larchey-Wendling¹ and Ralph Matthes²

¹ Université de Lorraine, CNRS, LORIA, Vandœuvre-lès-Nancy, France
`dominique.larchey-wendling@loria.fr`

² Institut de Recherche en Informatique de Toulouse (IRIT),
CNRS and University of Toulouse, Toulouse, France
`Ralph.Matthes@irit.fr`

Abstract. By using pointers, breadth-first algorithms are very easy to implement efficiently in imperative languages. Implementing them with the same bounds on execution time in purely functional style can be challenging, as explained in Okasaki’s paper at ICFP 2000 that even restricts the problem to binary trees but considers numbering instead of just traversal. Okasaki’s solution is modular and factors out the problem of implementing queues (FIFOs) with worst-case constant time operations. We certify those FIFO-based breadth-first algorithms on binary trees by extracting them from fully specified Coq terms, given an axiomatic description of FIFOs. In addition, we axiomatically characterize the strict and total order on branches that captures the nature of breadth-first traversal and propose alternative characterizations of breadth-first traversal of forests. We also propose efficient certified implementations of FIFOs by extraction, either with pairs of lists (with amortized constant time operations) or triples of lazy lists (with worst-case constant time operations), thus getting from extraction certified breadth-first algorithms with the optimal bounds on execution time.

Keywords: Breadth-first algorithms · Queues in functional programming · Correctness by extraction · Coq

1 Introduction

Breadth-first algorithms form an important class of algorithms with many applications. The distinguishing feature is that the recursive process tries to be “equitable” in the sense that all nodes in the graph with “distance” k from the starting point are treated before those at distance $k + 1$. In particular, with infinite (but finitely branching) structures, this ensures “fairness” in that all possible branches are eventually pursued to arbitrary depth, in other words, the

The first author got financial support by the TICAMORE joint ANR-FWF project. The second author got financial support by the COST action CA15123 EUTYPES.

recursion does not get “trapped” in an infinite branch.¹ This phenomenon that breadth-first algorithms avoid is different from a computation that gets “stuck”: even when “trapped”, there may be still steady “progress” in the sense of producing more and more units of output in finite time. In this paper, we will not specify or certify algorithms to work on infinite structures although we expect the lazy reading of our extracted programs (e. g., if we choose to extract towards the Haskell language) to work properly for infinite input as well and thus be fair—however without any guarantees from the extraction process. Anyway, as mentioned above, breadth-first algorithms impose a stronger, quantitative notion of “equity” than just abstract fairness to address the order of traversal of the structure.

When looking out for problems to solve with breadth-first algorithms, plain traversal of a given structure is the easiest task; to make this traversal “observable,” one simply prints the visited node labels (following the imperative programming style), or one collects them in a list, in the functional programming paradigm, as we will do in this paper (leaving filtering of search “hits” aside). However, in the interest of efficiency, it is important to use a first-in, first-out queue (FIFO) to organize the waiting sub-problems (see, e. g., Paulson’s book [15] on the ML language). Here, we are also concerned with functional languages, and for them, constant-time (in the worst case) implementations of the FIFO operations were a scientific challenge, solved very elegantly by Okasaki [13].

Breadth-first traversal is commonly [4] used to identify all nodes that are reachable in a graph from a given start node, and this allows creating a tree that captures the subgraph of reachable nodes, the “breadth-first tree” (for a given start node). In Okasaki’s landmark paper at ICFP 2000 [14], the author proposes to revisit the problem of *breadth-first numbering*: the traversal task is further simplified to start at the root of a (binary) tree, but the new challenge is to rebuild the tree in a functional programming language, where the labels of the nodes have been replaced by the value n if the node has been visited as the n -th one in the traversal. The progress achieved by Okasaki consists in separating the concerns of implementing breadth-first numbering from an efficient implementation of FIFOs: his algorithm works for any given FIFO and inherits optimal bounds on execution time from the given FIFO implementation. Thus, breadth-first numbering can be solved as efficiently with FIFOs as traversal,² and Okasaki reports in his paper that quite some of his colleagues did not come up with a FIFO-based solution when asked to find any solution.

In the present paper, using the Coq proof assistant,³ we formalize and solve the breadth-first traversal problem for finite binary trees with a rich specification in the sense of certified programming, i. e., when the output of an algorithm is a

¹ Meaning that recursion would be pursued solely in that branch.

² Jones and Gibbons [7] solved a variant of the problem with a “hard-wired” FIFO implementation—the one we review in Sect. 7.1—and thus do not profit from the theoretically most pleasing FIFO implementation by Okasaki [13].

³ <https://coq.inria.fr>, we have been using the current version 8.9.0, and the authoritative reference on Coq is <https://coq.inria.fr/distrib/current/refman/>.

dependent pair (v, C_v) composed of a value v together with a proof C_v certifying that v satisfies some (partial correctness) property.⁴ The key ingredient for its functional correctness are two equations already studied by Okasaki [14], but there as definitional device. Using the very same equations, we formalize and solve the breadth-first numbering problem, out of which the Coq extraction⁵ mechanism [9] can extract the same algorithm as Okasaki’s (however, we extract code in the OCaml⁶ language), but here with all guarantees concerning the non-functional property of termination and the functional/partial correctness, both aspects together constituting its *total correctness*, i. e., the algorithm indeed provides such a breadth-first numbering of the input tree.

As did Okasaki (and Gibbons and Jones for their solution [7]), we motivate the solution by first considering the natural extension to the problem of traversing or numbering a *forest*, i. e., a finite list of trees. The forests are subsequently replaced by an abstract FIFO structure, which is finally instantiated for several implementations, including the worst-case constant-time implementation following Okasaki’s paper [13] that is based on three lazy lists (to be extracted from coinductive lists in Coq for which we require as invariant that they are finite).

The breadth-first numbering problem can be slightly generalized to *breadth-first reconstruction*: the labels of the output tree are not necessarily natural numbers but come from a list of length equal to the number of labels of the input tree, and the n -th list element replaces the n -th node in the traversal, i. e., a minor variant of “breadth-first labelling” considered by Jones and Gibbons [7]. A slight advantage of this problem formulation is that a FIFO-based solution is possible by structural recursion on that list argument while the other algorithms were obtained by recursion over a measure (this is not too surprising since the length of that list coincides with the previously used measure).

In Sect. 2, we give background type theoretical material, mostly on list notation. Then we review existing tools for reasoning or defining terms by recursion/induction on a decreasing measure which combines more than one argument. We proceed to the short example of a simple interleaving algorithm that swaps its arguments when recurring on itself. In particular, we focus on how existing tools like **Program Fixpoint** or **Equations** behave in the context of extraction. Then we describe an alternative of our own—a tailored **induction-on** tactic—and we argue that it is more transparent to extraction than, e.g., **Equations**. Section 3 concentrates on background material concerning breadth-first traversal (specification, naive algorithm), including original material on a characterization of the breadth-first order, while Sect. 4 revolves around the mathematics of breadth-first traversal of forests that motivates the FIFO-based breadth-first algorithms. In Sect. 5, we use an abstractly specified datatype of FIFOs and deal with the different problems mentioned above: traversal,

⁴ For Coq, this method of rich specifications has been very much advocated in the Coq’Art, the first book on Coq [3].

⁵ The authors of the current version are Filliâtre and Letouzey, see <https://coq.inria.fr/refman/addendum/extraction.html>.

⁶ <http://ocaml.org>.

numbering, reconstruction. In Sect. 6, we elaborate on a level-based approach to numbering that thus is in the spirit of the naive traversal algorithm. Section 7 reports on the instantiation of the FIFO-based algorithms to two particular efficient FIFO implementations. Section 8 concludes.

The full Coq development, that also formalizes the theoretical results (the characterizations) in addition to the certification of the extracted algorithms in the OCaml language, is available on [GitHub](https://github.com/DmxLarchey/BFE):

<https://github.com/DmxLarchey/BFE>

In Appendix A, we give a brief presentation of these Coq vernacular⁷ files.

2 Preliminaries

We introduce some compact notations to represent the language of constructive type theory used in the proof assistant Coq. We describe the method called “certification by extraction” and illustrate how it challenges existing tools like **Program Fixpoint** or even **Equations** on the example of a simple interleaving algorithm. Then we introduce a tailored method to justify termination of fixpoints (or inductive proofs) using measures over, e.g., two arguments. This method is encapsulated into an **induction-on** tactic that is more transparent to extraction than the above-mentioned tools.

The type of propositions is denoted **Prop** while the type (family) of types is denoted **Type**. We use the inductive types of Booleans ($b : \mathbb{B} := 0 \mid 1$), of natural numbers ($n : \mathbb{N} := 0 \mid Sn$), of lists ($l : \mathbb{L} X := [] \mid x :: l$ with $x : X$) over a type parameter X . When $l = x_1 :: \dots :: x_n :: []$, we define $|l| := n$ as the length of l and we may write $l = [x_1; \dots; x_n]$ as well. We use $++$ for the concatenation of two lists (called the “append” function). The function $\text{rev} : \forall X, \mathbb{L} X \rightarrow \mathbb{L} X$ implements list reversal and satisfies $\text{rev} [] = []$ and $\text{rev}(x :: l) = \text{rev } l ++ x :: []$.

For a (heterogeneous) binary relation $R : X \rightarrow Y \rightarrow \mathbf{Prop}$, we define the lifting of R over lists $\forall^2 R : \mathbb{L} X \rightarrow \mathbb{L} Y \rightarrow \mathbf{Prop}$ by the following inductive rules, corresponding to the Coq standard library **Forall2** predicate:

$$\frac{}{\forall^2 R [] []} \qquad \frac{R \ x \ y \qquad \forall^2 R \ l \ m}{\forall^2 R \ (x :: l) \ (y :: m)}$$

Thus $\forall^2 R$ is the smallest predicate closed under the two above rules.⁸ Intuitively, when $l = [x_1; \dots; x_n]$ and $m = [y_1; \dots; y_p]$, the predicate $\forall^2 R \ l \ m$ means $n = p \wedge R \ x_1 \ y_1 \wedge \dots \wedge R \ x_n \ y_n$.

⁷ Vernacular is a Coq idiom for *syntactic sugar*, i.e., a human-friendly syntax for type theoretical notation—the “Gallina” language.

⁸ This is implemented in Coq by two constructors for the two rules together with an induction (or elimination) principle that ensures its smallestness.

2.1 Certification by Extraction

Certification by extraction is a particular way of establishing the correctness of a *given implementation* of an algorithm. In this paper, algorithms are given as programs in the OCaml language.

Hence, let us consider an OCaml program t of type $X \rightarrow Y$. The way to certify such a program by extraction is to implement a Coq term

$$f_t : \forall x : X_t, \text{pre}_t x \rightarrow \{y : Y_t \mid \text{post}_t x y\}$$

where $\text{pre}_t : X_t \rightarrow \mathbf{Prop}$ is the *precondition* (i.e., the domain of use of the function) and $\text{post}_t : X_t \rightarrow Y_t \rightarrow \mathbf{Prop}$ is the (dependent) *postcondition* which (possibly) relates the input with the output. The precondition could be tighter than the actual domain of the program t and the postcondition may characterize some aspects of the functional behavior of t , up to its full correctness, i.e., when $\text{pre}_t/\text{post}_t$ satisfy the following: for any given x such that $\text{pre}_t x$, the value $f_t x$ is the unique y such that $\text{post}_t x y$ holds.

We consider that t is certified when the postcondition faithfully represents the intended behavior of t and when f_t extracts to t : the extracted term $\text{extract}(f_t)$ but also the extracted types $\text{extract}(X_t)$ and $\text{extract}(Y_t)$ have to match “exactly” their respectively given OCaml definitions, i.e., we want $t \equiv \text{extract}(f_t)$, $X \equiv \text{extract}(X_t)$ and $Y \equiv \text{extract}(Y_t)$. The above identity sign “ $\equiv$ ” should be read as *syntactic identity*. This does not mean character by character equality between source codes but between the abstract syntax representations. Hence some slight differences are allowed, typically the name of bound variables which cannot always be controlled during extraction. Notice that $\text{pre}_t x$ and $\text{post}_t x y$ being of sort $\mathbf{Prop}$, they carry only logical content and no computational content. Thus they are erased by extraction.

As a method towards certified development, extraction can also be used without having a particular output program in mind, in which case it becomes a tool for writing programs that are correct by construction. Of course, getting a clean output might also be a goal and thus, the ability to finely control the computational content is important. But when we proceed from an already written program t , this fine control becomes critical and this has a significant impact on the tools which we can use to implement f_t (see upcoming Sect. 2.3).

2.2 Verification, Certification and the Trusted Computing Base

“Certification” encompasses “verification” in the following way: it aims at producing a certificate that can be checked independently of the software which is used to do the certification—at least in theory. While verification implies trusting the verifying software, certification implies only trusting the software that is used to verify the certificate, hence in the case of Coq, possibly an alternative type-checker. Notice that one of the goals of the MetaCoq⁹ project [1] is precisely to

⁹ <https://github.com/MetaCoq/metacoq>.

produce a type-checker for (a significant fragment of) Coq, a type-checker which will itself be certified by Coq.

Extraction in Coq is very straightforward once the term has been fully specified and type-checked. Calling the command

```
Recursive Extraction some_coq_term
```

outputs the extracted OCaml program and this part is fully automatic, although it is very likely that the output is not of the intended shape on the first attempt. So there may be a back-and-forth process to fine-tune the computational content of Coq terms until their extraction is satisfactory. The method can scale to larger developments because of the principle of *compositionality*. Indeed, provided a development can be divided into manageable pieces, Coq contains the tools that help at composing small certified bricks into bigger ones. Verifying or certifying large monolithic projects is generally hard—whatever tools are involved—because guessing the proper invariants becomes humanly unfeasible.

Considering the Trusted Computing Base (TCB) of certified extraction in Coq, besides trusting a type-checker, it also requires trusting the extraction process. In his thesis [10], P. Letouzey gave a mathematical proof of correctness (w.r.t. syntactic and semantic desiderata) of the extraction principles that guide the currently implemented **Extraction** command. Arguably, there is a difference between principles and an actual implementation. The above-cited MetaCoq project also aims at producing a certified “extraction procedure to untyped lambda-calculus accomplishing the same as the Extraction plugin of Coq.” Hence, for the moment, our work includes **Extraction** in its TCB but so do many other projects such as, e.g., the CompCert compiler.¹⁰ Still concerning verifying Coq extraction in Coq, we also mention the Euf^{11} project [12], but there is work of similar nature also in the Isabelle community [6]. Notice that we expect no or little change in the resulting extracted OCaml programs (once MetaCoq or one of its competitors reaches its goal) since these must respect the computational content of Coq terms. As the principle of “certification by extraction” is to obtain programs which are correct by construction directly from Coq code, we consider certifying extraction itself to be largely orthogonal to the goal pursued here.

2.3 Extraction of Simple Interleaving with Existing Tools

This section explains some shortcomings of the standard tools that can be used to implement recursive schemes in Coq when the scheme is more complicated than just structural recursion. We specifically study the **Function**, **Program Fixpoint**, and **Equations** commands. After discussing their versatility, we focus on how they interact with the **Extraction** mechanism of Coq.

¹⁰ <http://compcert.inria.fr>.

¹¹ <http://oeuf.uwplse.org>.

As a way to get a glimpse of the difficulty of the method, we begin with the case of the following simple interleaving function of type $\mathbb{L}X \rightarrow \mathbb{L}X \rightarrow \mathbb{L}X$ that merges two lists into one

$$[l_1; \dots; l_n], [m_1; \dots; m_n] \mapsto [l_1; m_1; l_2; m_2; \dots; l_n; m_n]$$

by alternating the elements of both input lists. The algorithm we want to certify is the following one:

$$\text{let rec itl } l \ m = \text{match } l \text{ with } [] \rightarrow m \mid x :: l \rightarrow x :: \text{itl } m \ l \quad (1)$$

where l and m switch roles in the second equation/match case for `itl`. Hence neither of these two arguments decreases structurally¹² in the recursive call and so, this definition cannot be used as such in a Coq **Fixpoint** definition. Notice that this algorithm has been identified as a challenge for program extraction in work by McCarthy *et al.* [11, p. 58], where they discuss an OCaml program of a function `cinterleave` that corresponds to our `itl`.

We insist on that specific algorithm—it is however trivial to define a function with the same functional behavior in Coq by the following equations:

$$\text{itl}_{\text{triv}} [] \ m = m \quad \text{itl}_{\text{triv}} l \ [] = l \quad \text{itl}_{\text{triv}} (x :: l) (y :: m) = x :: y :: \text{itl}_{\text{triv}} l \ m$$

Indeed, these equations correspond to a *structurally decreasing Fixpoint* which is accepted nearly as is by Coq.¹³ However, the algorithm we want to certify proceeds through the two following equations `itl [] m = m` and `itl (x :: l) m = x :: itl m l`. While it is not difficult to show that `itltriv` satisfies this specification, in particular by showing

$$\text{Fact itl_triv_fix_1} : \forall x \ l \ m, \text{itl}_{\text{triv}} (x :: l) m = x :: \text{itl}_{\text{triv}} m \ l$$

by nested induction on l and then m , extraction of `itltriv`, however, does not respect the expected code of `itl`, see Eq. (1).

While there is no structural decrease in Eq. (1), there is nevertheless an obvious decreasing measure in the recursive call of `itl`, i.e., $l, m \mapsto |l| + |m|$. We investigate several ways to proceed using that measure and discuss their respective advantages and drawbacks. The comments below are backed up by the file [interleave.v](#) corresponding to the following attempts. The use of Coq 8.9 is required for a recent version of the **Equations** package described below.

¹² Structural recursion is the built-in mechanism for recursion in Coq. It means that **Fixpoints** are type-checked only when Coq can determine that at least one of the arguments of the defined fixpoint (e.g., the first, the second...) decreases structurally on each recursive call, i.e., it must be a strict sub-term in an inductive type; and this must be the same argument that decreases for each recursive sub-call. This is called the *guard condition* for recursion and it ensures termination. On the other hand, it is a very restrictive form a recursion and we study more powerful alternatives here.

¹³ As a general rule in this paper, when equations can be straightforwardly implemented by structural induction on one of the arguments, we expect the reader to be able to implement the corresponding **Fixpoint** and so we do not further comment on this.

- Let us first consider the case of the **Function** command which extends the primitive **Fixpoint** command. It allows the definition of fixpoints on decreasing arguments that may not be structurally decreasing. However the **Function** method fails very quickly because it only allows for the decrease of *one of the arguments*, and with `itl`, only the decrease of a combination of both arguments can be used to show termination. We could of course pack the two arguments in a pair but then, this will modify the code of the given algorithm, a modification which will undoubtedly impact the extracted code;
- Let us now consider the case of the **Program Fixpoint** command [17]. We can define `itlprefix` via a fully specified term:

Program Fixpoint `itlprefixfull l m {measure (|l| + |m|)} : {r | r = itltriv l m} := ...`

basically by giving the right-hand side of Eq. (1) in Coq syntax, and then apply first and second projections to get the result as `itlprefix l m :=  $\pi_1$ (itlprefixfull l m)` and the proof that it meets its specification `itlprefix l m = itltriv l m` as `$\pi_2$ (itlprefixfull l m)`. The **Program Fixpoint** command generates proof obligations that are easy to solve in this case. Notice that defining `itlprefix` without going through the fully specified term is possible but then it is not possible to establish the postcondition `itlprefix l m = itltriv l m`;

- Alternatively, we can use the **Equations** package [18] which could be viewed as an extension/generalization of **Program Fixpoint**. In that case, we can proceed with a weakly specified term:

Equations `itleqs l m :  $\mathbb{L} X$  by wf (|l| + |m|) < := ...`

then derive the two equations that are used to define `itleqs`, i. e., `itleqs [] m = m` and `itleqs (x :: l) m := x :: itleqs m l`. These equations are directly obtained by making use of the handy **simp** tactic provided with **Equations**. With these two equations, it is then easy to show the identity `itleqs l m = itltriv l m`.

To sum up, considering the Coq implementation side only: **Function** fails because no single argument decreases; **Program Fixpoint** succeeds but via a fully specified term to get the postcondition; and **Equations** succeeds directly and generates the defining equations that are accessible through the **simp** tactic.

However, when we consider the extraction side of the problem, both **Program Fixpoint** and **Equations** do not give us the intended OCaml term of Eq. (1). On the left side of Fig. 1, we display the extracted code for `itlprefix` that was developed using **Program Fixpoint**, and on the right side the extracted code for `itleqs` implemented using **Equations**.¹⁴ Both present two drawbacks in our eyes. First, because “full” (resp. “eqs₀”) are not made globally visible, it is impossible to inline their definitions with the **Extraction Inline** directive although this would be an obvious optimization. But then, also more problematic from an algorithmic point of view, there is the packing of the two arguments into a pair

¹⁴ The actual extracted codes for both `itlprefix` and `itleqs` are not that clean but we simplified them a bit to single out two specific problems.

<pre> let itl_{fix} l m = let rec loop p = let l₀ = fst p in let m₀ = snd p in let full = fun l₂ m₂ → loop (l₂, m₂) in match l₀ with [] → m₀ x :: l₁ → x :: full m₀ l₁ in loop (l, m) </pre>	<pre> let itl_{eqs} l m = let rec loop p = let m₀ = snd p in let eqs₀ = fun l₂ m₂ _ → loop (l₂, m₂) in match fst p with [] → m₀ x :: l₁ → x :: eqs₀ m₀ l₁ _ in loop (l, m) </pre>
---	--

Fig. 1. Extraction of itl_{fix} (Program Fixpoint) and itl_{eqs} (Equations).

(l, m) prior to the call to the “loop” function that implements itl_{fix} (resp. itl_{eqs}). As a last remark, the two extracted codes look similar except that for itl_{eqs} , there is the slight complication of an extra *dummy* parameter $_$ added to the above definition of “eqs₀” that is then instantiated with a dummy argument $_$.

To sum up our above attempts, while the **Function** commands fails to handle itl because of the swap between arguments in Eq. (1), both **Program Fixpoint** and **Equations** succeed when considering the specification side, i. e., defining the function (which implicitly ensures termination) and proving its postcondition. However, extraction-wise, both generate artifacts, e. g., the pairing of arguments in this case, and which are difficult or impossible to control making the “certification by extraction” approach defined in Sect. 2.1 fail.

In the previously cited work by McCarthy *et al.* [11], the authors report that they successfully avoided “verification residues” as far as their generation of obligations for the certification of running time was concerned. We are heading for the absence of such residues from the extracted code. Let us hence now consider a last approach based on a finely tuned tactic which is both user-friendly and inlines inner fixpoints making it transparent to extraction. This is the approach we will be using for the breadth-first algorithms considered in here.

2.4 Recursion on Measures

We describe how to implement definitions of terms by induction on a measure of, e. g., two arguments. Let us consider two types $X, Y : \text{Type}$, a doubly indexed family $P : X \rightarrow Y \rightarrow \text{Type}$ of types and a measure $\|\cdot, \cdot\| : X \rightarrow Y \rightarrow \mathbb{N}$. We explain how to build terms or proofs of type $t : P x y$ by induction on the measure $\|x, y\|$. Hence, to build the term t , we are allowed to use instances of the induction hypothesis:

$$IH : \forall x' y', \|x', y'\| < \|x, y\| \rightarrow P x' y'$$

i. e., the types $P x' y'$ with smaller x'/y' arguments are (recursively) inhabited. The measures we use here are limited to $\mathbb{N}$ viewed as well-founded under the “strictly less” order. Any other type with a well-founded order would work as well.

To go beyond measures and implement a substantial fragment of general recursion, we invite the reader to consult work by the first author and Monin [8].

In the file [wf_utils.v](#), we prove the following theorem while carefully crafting the computational content of the proof term, so that extraction yields an algorithm that is clean of spurious elements:

Theorem `measure_double_rect` ($P : X \rightarrow Y \rightarrow \text{Type}$) :

$$(\forall x y, (\forall x' y', \|x', y'\| < \|x, y\| \rightarrow P x' y') \rightarrow P x y) \rightarrow \forall x y, P x y.$$

It allows building a term of type $P x y$ by simultaneous induction on x and y using the decreasing measure $\|x, y\|$ to ensure termination. To ease the use of theorem `measure_double_rect` we define a *tactic notation* that is deployed as follows:

`induction on x y as IH with measure \|x, y\|`

However, if the `induction-on` tactic was just about applying an instance of the term `measure_double_rect`, it would still leave artifacts in the extracted code much like the `_` (resp. `_`) dummy parameter (resp. argument) in the `itl_eqs` example of Fig. 1. So, to be precise in our description, the `induction-on` tactic actually builds the needed instance of the proof term `measure_double_rect` on a per-use basis, i.e., the proof term is inlined by the tactic itself. This ensures perfect extraction in the sense of the result being free of any artifacts. We refer to the file [wf_utils.v](#) for detailed explanations on how the inlining works.

From the point of view of specification, our `induction-on` tactic gives the same level of flexibility when compared to `Program Fixpoint` or `Equations`, at least when termination is grounded on a decreasing measure. However, when considering extraction, we think that it offers a finer control over the computational content of Coq terms, a feature which can be critical when doing certification by extraction. We now illustrate how to use the `induction-on` tactic from the user's point of view, on the example of the simple interleaving algorithm.

2.5 Back to Simple Interleaving

To define and specify `itl_on` using the `induction-on` tactic, we first implement a fully specified version of `itl_on` as the *functional* equivalent of `itl_triv`.¹⁵ We proceed with the following script, whose first line reads as: we want to define the list r together with a proof that r is equal to `itl_triv l m`—a typical rich specification.

¹⁵ Of course, it is not operationally equivalent.

Definition $\text{itl}_{\text{on}}^{\text{full}} l m : \{r : \mathbb{L} X \mid r = \text{itl}_{\text{triv}} l m\}$.

Proof.

induction on $l m$ as loop with measure $(|l| + |m|)$.

revert loop; refine (match l with

| nil $\mapsto$ fun _ $\mapsto$ exist _ $m \mathbb{O}_1^?$
| $x :: l' \mapsto$ fun loop $\mapsto$ let $(r, H_r) := \text{loop } m l' \mathbb{O}_2^?$
in exist _ $(x :: r) \mathbb{O}_3^?$ end).

★ trivial.

(* proof of $\mathbb{O}_1^?$ *)

★ simpl; omega.

(* proof of $\mathbb{O}_2^?$ *)

★ subst; rewrite itl_triv_fix_1; trivial.

(* proof of $\mathbb{O}_3^?$ *)

Defined.

The code inside the **refine**(...) tactic outputs terms like **exist** _ $s \mathbb{O}_s^?$ where _ is recovered by unification, and $\mathbb{O}_s^?$ is left open to be solved later by the user.¹⁶ The constructor **exist** packs the pair $(s, \mathbb{O}_s^?)$ as a term of dependent type $\{r : \mathbb{L} X \mid r = \text{itl}_{\text{triv}} l m\}$ and thus $\mathbb{O}_s^?$ remains to be realized in the type $s = \text{itl}_{\text{triv}} l m$. This particular use of **refine** generates three *proof obligations* (also denoted PO)

$$\begin{aligned} \mathbb{O}_1^? & // \dots \vdash m = \text{itl}_{\text{triv}} [] m \\ \mathbb{O}_2^? & // \dots \vdash |m| + |l'| < |x :: l'| + |m| \\ \mathbb{O}_3^? & // \dots, H_r : r = \text{itl}_{\text{triv}} m l' \vdash x :: r = \text{itl}_{\text{triv}} (x :: l') m \end{aligned}$$

later proved with their respective short proof scripts. Notice that our newly introduced **induction-on** tactic could alternatively be used to give another proof of **itl_triv_fix_1** by measure induction on $|l| + |m|$. We remark that PO $\mathbb{O}_2^?$ is of different nature than $\mathbb{O}_1^?$ and $\mathbb{O}_3^?$. Indeed, $\mathbb{O}_2^?$ is a precondition, in this case a *termination certificate* ensuring that the recursive call occurs on smaller inputs. On the other hand, $\mathbb{O}_1^?$ and $\mathbb{O}_3^?$ are postconditions ensuring the functional correctness by type-checking (of the logical part of the rich specification). As a general rule in this paper, providing termination certificates will always be relatively easy because they reduce to proofs of strict inequations between arithmetic expressions, usually solved by the **omega** tactic.¹⁷

Considering POs $\mathbb{O}_2^?$ and $\mathbb{O}_3^?$, they contain the following hypothesis (hidden in the dots) which witnesses the induction on the measure $|l| + |m|$. It is called **loop** as indicated in the **induction-on** tactic used at the beginning of the script:

$$\text{loop} : \forall l_0 m_0, |l_0| + |m_0| < |x :: l'| + |m| \rightarrow \{r \mid r = \text{itl}_{\text{triv}} l_0 m_0\}$$

It could also appear in $\mathbb{O}_1^?$ but since we do not need it to prove $\mathbb{O}_1^?$, we intentionally cleared it using **fun** _ $\mapsto$... Actually, **loop** is not used in the proofs

¹⁶ In the Coq code, $\mathbb{O}_s^?$ is simply another _ hole left for the **refine** tactic to either fill it by unification or postpone it. Because $\mathbb{O}_1^?$, $\mathbb{O}_2^?$ and $\mathbb{O}_3^?$ are postponed in this example, we give them names for better explanations.

¹⁷ Because termination certificates have a purely logical content, we do not care whether **omega** produces an “optimal” proof (b.t.w., it never does), but we appreciate its efficiency in solving those kinds of goals which we do not want to spend much time on, thus the gain is in time spent on developing the certified code. Efficiency of code execution is not touched since these proofs leave no traces in the extracted code.

of $\mathbb{O}_2^?$ or $\mathbb{O}_3^?$ either but it is necessary for the recursive call implemented in the `let ... := loop ... in ...` construct.

This peculiar way of writing terms as a combination of *programming style* and *combination of proof tactics* is possible in Coq by the use of the “swiss-army knife” tactic `refine`¹⁸ that, via unification, allows the user to specify only parts of a term leaving holes to be filled later if unification fails to solve them. It is a major tool to finely control computational content while allowing great tactic flexibility on purely logical content.

We continue the definition of itl_{on} as the first projection

Definition $\text{itl}_{\text{on}} l m := \pi_1(\text{itl}_{\text{on}}^{\text{full}} l m)$.

and its specification using the projection π_2 on the dependent pair $\text{itl}_{\text{on}}^{\text{full}} l m$.

Fact $\text{itl_on_spec } l m : \text{itl}_{\text{on}} l m = \text{itl}_{\text{triv}} l m$.

Notice that by asking the extraction mechanism to inline the definition of $\text{itl}_{\text{on}}^{\text{full}}$, we get the following extracted OCaml code for itl_{on} , the one that optimally reflects the original specification of Eq. (1):

```
let rec itl_on l m = match l with [] → m | x :: l → x :: itl_on m l
```

Of course, this outcome of the (automatic) code extraction had to be targeted when doing the proof of $\text{itl}_{\text{on}}^{\text{full}}$: a trivial proof would have been to choose as r just $\text{itl}_{\text{triv}} l m$, and the extracted code would have been just that. Instead, we did pattern-matching on l and chose in the first case m and in the second case $x :: r$, with r obtained as first component of the recursive call `loop m l'`. The outer induction took care that `loop` stood for the function we were defining there.

Henceforth, we will take induction or recursion on measures for granted, assuming they correspond to the use of the tactic `induction-on`.

3 Traversal of Binary Trees

We present mostly standard material on traversal of binary trees that will lay the ground for the original contributions in the later sections of the paper. After defining binary trees and basic notions for them including their branches (Sect. 3.1), we describe mathematically what constitutes breadth-first traversal by considering an order on the branches (Sect. 3.2) that we characterize axiomatically (our Theorem 1). We then look at breadth-first traversal in its most elementary form (Sect. 3.3), by paying attention that it indeed meets its specification in terms of the order of the visited branches.

¹⁸ The `refine` tactic was originally implemented by J.-C. Filliâtre.

3.1 Binary Trees

We use the type of binary trees $(a, b : \mathbb{T} X := \langle x \rangle \mid \langle a, x, b \rangle)$ where x is of the argument type X .¹⁹ We define the $\text{root} : \mathbb{T} X \rightarrow X$ of a tree, the subtrees $\text{subt} : \mathbb{T} X \rightarrow \mathbb{L}(\mathbb{T} X)$ of a tree, and the size $\|\cdot\| : \mathbb{T} X \rightarrow \mathbb{N}$ of a tree by:

$$\begin{aligned} \text{root } \langle x \rangle &:= x & \text{subt } \langle x \rangle &:= [] & \|\langle x \rangle\| &:= 1 \\ \text{root } \langle -, x, - \rangle &:= x & \text{subt } \langle a, -, b \rangle &:= [a; b] & \|\langle a, x, b \rangle\| &:= 1 + \|a\| + \|b\| \end{aligned}$$

and we extend the measure to lists of trees by $\|[t_1; \dots; t_n]\| := \|t_1\| + \dots + \|t_n\|$, hence $\|[]\| = 0$, $\|[t]\| = \|t\|$ and $\|l + m\| = \|l\| + \|m\|$. We will not need more complicated measures than $\|\cdot\|$ to justify the termination of the breadth-first algorithms to come.

A branch in a binary tree is described as a list of Booleans in $\mathbb{L}\mathbb{B}$ representing a list of left/right choices (0 for left and 1 for right).²⁰ We define the predicate $\text{btb} : \mathbb{T} X \rightarrow \mathbb{L}\mathbb{B} \rightarrow \mathbf{Prop}$ inductively by the rules below where $t \parallel l \downarrow$ denotes $\text{btb } t \ l$ and tells whether l is a branch in t :

$$\frac{}{t \parallel [] \downarrow} \quad \frac{a \parallel l \downarrow}{\langle a, x, b \rangle \parallel 0 :: l \downarrow} \quad \frac{b \parallel l \downarrow}{\langle a, x, b \rangle \parallel 1 :: l \downarrow}$$

We define the predicate $\text{bpn} : \mathbb{T} X \rightarrow \mathbb{L}\mathbb{B} \rightarrow X \rightarrow \mathbf{Prop}$ inductively as well. The term $\text{bpn } t \ l \ x$, also denoted $t \parallel l \rightsquigarrow x$, tells whether *the node identified by l in t is decorated with x* :

$$\frac{}{t \parallel [] \rightsquigarrow \text{root } t} \quad \frac{a \parallel l \rightsquigarrow r}{\langle a, x, b \rangle \parallel 0 :: l \rightsquigarrow r} \quad \frac{b \parallel l \rightsquigarrow r}{\langle a, x, b \rangle \parallel 1 :: l \rightsquigarrow r}$$

We show the result that l is a branch of t if and only if it is decorated in t :

Fact btb_spec $(t : \mathbb{T} X) (l : \mathbb{L}\mathbb{B}) : t \parallel l \downarrow \leftrightarrow \exists x, t \parallel l \rightsquigarrow x$.

By two inductive rules, we define $\sim_{\mathbb{T}} : \mathbb{T} X \rightarrow \mathbb{T} Y \rightarrow \mathbf{Prop}$, *structural equivalence of binary trees*, and we lift it to structural equivalence of lists of binary trees $\sim_{\mathbb{L}\mathbb{T}} : \mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L}(\mathbb{T} Y) \rightarrow \mathbf{Prop}$

$$\frac{}{\langle x \rangle \sim_{\mathbb{T}} \langle y \rangle} \quad \frac{a \sim_{\mathbb{T}} a' \quad b \sim_{\mathbb{T}} b'}{\langle a, x, b \rangle \sim_{\mathbb{T}} \langle a', y, b' \rangle} \quad l \sim_{\mathbb{L}\mathbb{T}} m := \forall^2(\sim_{\mathbb{T}}) \ l \ m$$

i. e., when $l = [a_1; \dots; a_n]$ and $m = [b_1; \dots; b_p]$, the equivalence $l \sim_{\mathbb{L}\mathbb{T}} m$ means $n = p \wedge a_1 \sim_{\mathbb{T}} b_1 \wedge \dots \wedge a_n \sim_{\mathbb{T}} b_n$ (see Sect. 2).

Both $\sim_{\mathbb{T}}$ and $\sim_{\mathbb{L}\mathbb{T}}$ are equivalence relations, and if $a \sim_{\mathbb{T}} b$ then $\|a\| = \|b\|$, i. e., structurally equivalent trees have the same size, and this holds for structurally equivalent lists of trees as well.

¹⁹ In his paper [14], Okasaki considered unlabeled leaves. When we compare with his findings, we always tacitly adapt his definitions to cope with leaf labels, which adds only a small notational overhead.

²⁰ We here intend to model branches from the root to nodes, rather than from nodes to leaves. It might have been better to call the concept paths instead of branches.

3.2 Ordering Branches of Trees

A *strict order* $<_R$ is an irreflexive and transitive (binary) relation. We say it is *total* if the associated partial order $<_R \cup =$ is total in the usual sense of being connex, or equivalently, if $\forall x y, \{x <_R y\} \vee \{x = y\} \vee \{y <_R x\}$. It is *decidable* and *total* if:

$$\forall x y, \{x <_R y\} + \{x = y\} + \{y <_R x\}$$

where, as usual in type theory, a proof of the sum $A + B + C$ requires either a proof of A , of B or of C together with the information whether the first, second or third summand has been proven.

The *dictionary order* of type $\prec_{\text{dic}} : \mathbb{L}\mathbb{B} \rightarrow \mathbb{L}\mathbb{B} \rightarrow \mathbf{Prop}$ is the lexicographic product on lists of 0's or 1's defined by

$$\frac{}{\square \prec_{\text{dic}} b :: l} \quad \frac{}{0 :: l \prec_{\text{dic}} 1 :: m} \quad \frac{l \prec_{\text{dic}} m}{b :: l \prec_{\text{dic}} b :: m}$$

The *breadth-first order* on $\mathbb{L}\mathbb{B}$ of type $\prec_{\text{bf}} : \mathbb{L}\mathbb{B} \rightarrow \mathbb{L}\mathbb{B} \rightarrow \mathbf{Prop}$ is defined by

$$\frac{|l| < |m|}{l \prec_{\text{bf}} m} \quad \frac{|l| = |m| \quad l \prec_{\text{dic}} m}{l \prec_{\text{bf}} m}$$

i. e., the lexicographic product of *shorter* and *dictionary order* (if equal length).

Lemma 1. $\prec_{\text{dic}}$ and $\prec_{\text{bf}}$ are decidable and total strict orders.

We characterize $\prec_{\text{bf}}$ with the following four axioms:

Theorem 1. Let $<_R : \mathbb{L}\mathbb{B} \rightarrow \mathbb{L}\mathbb{B} \rightarrow \mathbf{Prop}$ be a relation s. t.

- (A₁) $<_R$ is a strict order (irreflexive and transitive)
- (A₂) $\forall x l m, l <_R m \leftrightarrow x :: l <_R x :: m$
- (A₃) $\forall l m, |l| < |m| \rightarrow l <_R m$
- (A₄) $\forall l m, |l| = |m| \rightarrow 0 :: l <_R 1 :: m$

Then $<_R$ is equivalent to $\prec_{\text{bf}}$, i. e., $\forall l m, l <_R m \leftrightarrow l \prec_{\text{bf}} m$. Moreover the relation $\prec_{\text{bf}}$ satisfies (A₁)–(A₄).

3.3 Breadth-First Traversal, a Naive Approach

The *zipping* function $\text{zip} : \mathbb{L}(\mathbb{L}X) \rightarrow \mathbb{L}(\mathbb{L}X) \rightarrow \mathbb{L}(\mathbb{L}X)$ is defined by

$$\text{zip } \square :: m \quad \text{zip } l \square := l \quad \text{zip } (x :: l) (y :: m) := (x ++ y) :: \text{zip } l m$$

The level-wise function $\text{niv} : \mathbb{T}X \rightarrow \mathbb{L}(\mathbb{L}X)$ — *niv* refers to the French word “niveaux” — is defined by²¹

$$\text{niv } \langle x \rangle := [x] :: \square \quad \text{niv } \langle a, x, b \rangle := [x] :: \text{zip } (\text{niv } a) (\text{niv } b)$$

²¹ The function *niv* is called “levelorder traversal” by Jones and Gibbons [7].

The $(n + 1)$ -th element of $\text{niv } t$ contains the labels of t in left-to-right order that have distance n to the root. We can then define $\text{bft}_{\text{std}} t := \text{concat}(\text{niv } t)$.²² We lift breadth-first traversal to branches instead of decorations, defining $\text{niv}_{\text{br}} : \mathbb{T} X \rightarrow \mathbb{L}(\mathbb{L}(\mathbb{L} \mathbb{B}))$ by

$$\begin{aligned} \text{niv}_{\text{br}} \langle x \rangle &:= [] :: [] \\ \text{niv}_{\text{br}} \langle a, x, b \rangle &:= [] :: \text{zip}(\text{map}(\text{map}(0 :: \cdot))(\text{niv}_{\text{br}} a))(\text{map}(\text{map}(1 :: \cdot))(\text{niv}_{\text{br}} b)) \end{aligned}$$

and then $\text{bft}_{\text{br}} t := \text{concat}(\text{niv}_{\text{br}} t)$, and we show the two following results:

Theorem `niveaux_br_niveaux` $t : \forall^2 (\forall^2 (\text{bpn } t)) (\text{niv}_{\text{br}} t) (\text{niv } t)$.

Theorem `bft_br_std` $t : \forall^2 (\text{bpn } t) (\text{bft}_{\text{br}} t) (\text{bft}_{\text{std}} t)$.

Hence $\text{bft}_{\text{br}} t$ and $\text{bft}_{\text{std}} t$ traverse the tree t in the same order, except that bft_{std} outputs decorating values and bft_{br} outputs branches.²³ We moreover show that $\text{bft}_{\text{br}} t$ lists the branches of t in $\prec_{\text{bf}}$ -ascending order.

Theorem 2. *The list $\text{bft}_{\text{br}} t$ is strictly sorted w. r. t. $\prec_{\text{bf}}$.*

4 Breadth-First Traversal of a Forest

We lift `root` and `subt` to lists of trees and define $\text{roots} : \mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L} X$ and $\text{subtrees} : \mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L}(\mathbb{T} X)$ by

$$\text{roots} := \text{map root} \qquad \text{subtrees} := \text{flat_map subt}$$

where `flat_map` is the standard list operation given by $\text{flat_map } f [x_1; \dots; x_n] := f x_1 ++ \dots ++ f x_n$. To justify the upcoming fixpoints/inductive proofs where recursive sub-calls occur on `subtrees` l , we show the following

Lemma `subtrees_dec` $l : l = [] \vee \|\text{subtrees } l\| < \|l\|$.

Hence we can justify termination of a recursive algorithm f l with formal argument $l : \mathbb{L}(\mathbb{T} X)$ that is calling itself on $f(\text{subtrees } l)$, as soon as the case $f []$ is computed without using recursive calls (to f). For this, we use, for instance, recursion on the measure $l \mapsto \|l\|$ but we may also use the binary measure $l, m \mapsto \|l ++ m\|$ when f has two arguments instead of just one.

4.1 Equational Characterization of Breadth-First Traversal

We first characterize `bftf`—breadth-first traversal of a forest—with four equivalent equations:

²² Where $\text{concat} := \text{fold } (\cdot ++ \cdot) []$, i. e., $\text{concat } [l_1; \dots; l_n] = l_1 ++ \dots ++ l_n$.

²³ The `bpn` t relation, also denoted $l, x \mapsto t // l \rightsquigarrow x$, relates branches and decorations.

Theorem 3 (Characterization of breadth-first traversal of forests – recursive part). *Let bft_f be any term of type $\mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L} X$ and consider the following equations:*

- (1) $\forall l, \text{bft}_f l = \text{roots } l \mathbin{++} \text{bft}_f (\text{subtrees } l);$
- (2) $\forall l m, \text{bft}_f (l \mathbin{++} m) = \text{roots } l \mathbin{++} \text{bft}_f (m \mathbin{++} \text{subtrees } l);$
- (3) $\forall t l, \text{bft}_f (t :: l) = \text{root } t :: \text{bft}_f (l \mathbin{++} \text{subt } t);$
- (Oka1) $\forall x l, \text{bft}_f (\langle x \rangle :: l) = x :: \text{bft}_f l;$
- (Oka2) $\forall a b x l, \text{bft}_f (\langle a, x, b \rangle :: l) = x :: \text{bft}_f (l \mathbin{++} [a; b]).$

We have the equivalence: $(1) \leftrightarrow (2) \leftrightarrow (3) \leftrightarrow (\text{Oka1} \wedge \text{Oka2}).$

Proof. Equations (1) and (3) are clear instances of Eq. (2). Then $(\text{Oka1} \wedge \text{Oka2})$ is equivalent to (3) because they just represent a case analysis on t . So the only difficulty is to show $(1) \rightarrow (2)$ and $(3) \rightarrow (2)$. Both inductive proofs alternate the roles of l and m . So proving (2) from e.g. (1) by induction on either l or m is not possible. Following the example of the simple interleaving algorithm of Sect. 2.5, we proceed by induction on the measure $\|l \mathbin{++} m\|$. $\square$

Equations (Oka1) and (Oka2) are used by Okasaki [14] as defining equations, while (3) is calculated from the specification by Jones and Gibbons [7]. We single out Eq. (2) above as a smooth gateway between *subtrees*-based breadth-first algorithms and FIFO-based breadth-first algorithms. Unlocking that “latch bolt” enabled us to show correctness properties of refined breadth-first algorithms.

Theorem 4 (Full characterization of breadth-first traversal of forests). *Adding equation $\text{bft}_f [] = []$ to any one of the equations of Theorem 3 determines the function bft_f uniquely.*

Proof. For any bft_1 and bft_2 satisfying both $\text{bft}_1 [] = []$, $\text{bft}_2 [] = []$ and e.g. $\text{bft}_1 l = \text{roots } l \mathbin{++} \text{bft}_1 (\text{subtrees } l)$ and $\text{bft}_2 l = \text{roots } l \mathbin{++} \text{bft}_2 (\text{subtrees } l)$, we show $\text{bft}_1 l = \text{bft}_2 l$ by induction on $\|l\|$. $\square$

Notice that one should not confuse the uniqueness of the function—which is an extensional notion—with the uniqueness of an algorithm implementing such a function, because there are hopefully numerous possibilities.²⁴

4.2 Direct Implementation of Breadth-First Traversal

We give a definition of $\text{forest}_{\text{dec}} : \mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L} X \times \mathbb{L}(\mathbb{T} X)$ such that, provably, $\text{forest}_{\text{dec}} l = (\text{roots } l, \text{subtrees } l)$, but using a simultaneous computation:

$$\begin{aligned} \text{forest}_{\text{dec}} [] &:= ([], []) & \text{forest}_{\text{dec}} (\langle x \rangle :: l) &:= (x :: \alpha, \beta) \\ \text{forest}_{\text{dec}} (\langle a, x, b \rangle :: l) &:= (x :: \alpha, a :: b :: \beta) & \text{where } (\alpha, \beta) &:= \text{forest}_{\text{dec}} l. \end{aligned}$$

²⁴ Let us stress that extensionality is not very meaningful for algorithms anyway.

Then we show one way to realize the equations of Theorem 3 into a Coq term:

Theorem 5 (Existence of breadth-first traversal of forests). *One can define a Coq term bft_f of type $\mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L} X$ s. t.*

1. $\text{bft}_f [] = []$
2. $\forall l, \text{bft}_f l = \text{roots } l \mathbin{++} \text{bft}_f (\text{subtrees } l)$

and s. t. bft_f extracts to the following OCaml code:²⁵

```
let rec bft_f l = match l with [] -> [] | _ -> let alpha, beta = forest_dec l in alpha @ bft_f beta.
```

Proof. We define the graph $\rightsquigarrow_{\text{bft}}$ of the algorithm bft_f as binary relation of type $\mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L} X \rightarrow \text{Prop}$ with the two following inductive rules:

$$\frac{}{[] \rightsquigarrow_{\text{bft}} []} \quad \frac{l \neq [] \quad \text{subtrees } l \rightsquigarrow_{\text{bft}} r}{l \rightsquigarrow_{\text{bft}} \text{roots } l \mathbin{++} r}$$

These rules follow the intended algorithm. We show that the graph $\rightsquigarrow_{\text{bft}}$ is functional/deterministic, i. e.

Fact $\text{bft}_f.\text{fun}$: $\forall l r_1 r_2, l \rightsquigarrow_{\text{bft}} r_1 \rightarrow l \rightsquigarrow_{\text{bft}} r_2 \rightarrow r_1 = r_2$.

By induction on the measure $\|l\|$ we define a term $\text{bft}_f.\text{full } l : \{r \mid l \rightsquigarrow_{\text{bft}} r\}$ where we proceed as in Sect. 2.5. We get bft_f by the first projection $\text{bft}_f l := \pi_1(\text{bft}_f.\text{full } l)$ and derive the specification

Fact $\text{bft}_f.\text{spec}$: $\forall l, l \rightsquigarrow_{\text{bft}} \text{bft}_f l$.

with the second projection π_2 . Equations 1 and 2 follow straightforwardly from $\text{bft}_f.\text{fun}$ and $\text{bft}_f.\text{spec}$. $\square$

Hence we see that we can use the specifying Eqs. 1 and 2 of Theorem 5 to define the term bft_f . In the case of breadth-first algorithms, termination is not very complicated because one can use induction on a measure to ensure it. In the proof, we just need to check that recursive calls occur on smaller arguments according to the given measure, and this follows from Lemma `subtrees_dec`.

Theorem 6 (Correctness of breadth-first traversal of forests). *For all $t : \mathbb{T} X$, we have $\text{bft}_f[t] = \text{bft}_{\text{std}} t$.*

Proof. We define $\rightsquigarrow_{\text{niv}} : \mathbb{L}(\mathbb{T} X) \rightarrow \mathbb{L}(\mathbb{L} X) \rightarrow \text{Prop}$ by

$$\frac{}{[] \rightsquigarrow_{\text{niv}} []} \quad \frac{l \neq [] \quad \text{subtrees } l \rightsquigarrow_{\text{niv}} ll}{l \rightsquigarrow_{\text{niv}} \text{roots } l :: ll}$$

and we show that $l \rightsquigarrow_{\text{niv}} ll \rightarrow m \rightsquigarrow_{\text{niv}} mm \rightarrow l \mathbin{++} m \rightsquigarrow_{\text{niv}} \text{zip } ll \text{ } mm$ holds, a property from which we deduce $[t] \rightsquigarrow_{\text{niv}} \text{niv } t$. We show $l \rightsquigarrow_{\text{niv}} ll \rightarrow l \rightsquigarrow_{\text{bft}} \text{concat } ll$ and we deduce $[t] \rightsquigarrow_{\text{bft}} \text{concat } (\text{niv } t)$ hence $[t] \rightsquigarrow_{\text{bft}} \text{bft}_{\text{std}} t$. By $\text{bft}_f.\text{spec}$ we have $[t] \rightsquigarrow_{\text{bft}} \text{bft}_f[t]$ and we conclude with $\text{bft}_f.\text{fun}$. $\square$

²⁵ Notice that list append is denoted $@$ in OCaml and $\mathbin{++}$ in Coq.

4.3 Properties of Breadth-First Traversal

The *shape* of a tree (resp. forest) is the structure that remains when removing the values on the nodes and leaves, e.g., by mapping the base type X to the singleton type `unit`. Alternatively, one can use the $\sim_{\mathbb{T}}$ (resp. $\sim_{\mathbb{LT}}$) equivalence relation (introduced in Sect. 3.1) to characterize trees (resp. forests) which have identical shapes. We show that on a given forest shape, breadth-first traversal is an injective map:

Lemma `bft_f_inj` $l \sim_{\mathbb{LT}} m \rightarrow \text{bft}_f l = \text{bft}_f m \rightarrow l = m$.

Proof. By induction on the measure $l, m \mapsto \|l \mathbin{++} m\|$ and then case analysis on l and m . We use (Oka1&Oka2) from Theorem 3 to rewrite `bftf` terms. The shape constraint $l \sim_{\mathbb{LT}} m$ ensures that the same equation is used for l and m . $\square$

Hence, on a given tree shape, `bftstd` is also injective:

Corollary 1. *If $t_1 \sim_{\mathbb{T}} t_2$ are two trees of type $\mathbb{T}X$ (of the same shape) and `bftstd  $t_1$  = bftstd  $t_2$` then $t_1 = t_2$.*

Proof. From Lemma `bft_f_inj` and Theorem 6. $\square$

4.4 Discussion

The algorithm described in Theorem 5 can be used to compute breadth-first traversal as a replacement for the naive `bftstd` algorithm. We could also use other equations of Theorem 3, for instance using `bftf [] = []`, (Oka1) and (Oka2) to define another algorithm. The problem with equation

$$\text{(Oka2)} \quad \text{bft}_f (\langle a, x, b \rangle :: l) = x :: \text{bft}_f (l \mathbin{++} [a; b])$$

is that it implies the use of $\mathbin{++}$ to append two elements at the tail of l , which is a well-known culprit that transforms an otherwise linear-time algorithm into a quadratic one. But Equation (Oka2) hints at replacing the list data-structure with a first-in, first-out queue (FIFO) for the argument of `bftf` which brings us to the central section of this paper.

5 FIFO-Based Breadth-First Algorithms

Here come the certified algorithms in the spirit of Okasaki's paper [14]. They have the potential to be efficient, but this depends on the later implementation of the axiomatic datatype of FIFOs considered here (Sect. 5.1). We deal with traversal, numbering, reconstruction—in breadth-first order.

```

fifo : Type → Type
f2l  : ∀ X, fifo X → ℒ X
emp  : ∀ X, {q | f2l q = []}
enq  : ∀ X q x, {q' | f2l q' = f2l q ++ [x]}
deq  : ∀ X q, f2l q ≠ [] → {(x, q') | f2l q = x :: f2l q'}
void : ∀ X q, {b : ℤ | b = 1 ↔ f2l q = []}

```

Fig. 2. An axiomatization of first-in first-out queues.

5.1 Axiomatization of FIFOs

In Fig. 2, we give an axiomatic description of polymorphic first-in, first-out queues (a.k.a. FIFOs) by projecting them to lists with $\text{f2l } \{X\} : \text{fifo } X \rightarrow \mathbb{L} X$ where the notation $\{X\}$ marks X as an *implicit argument*²⁶ of f2l . Each axiom is fully specified using f2l : **emp** is the empty queue, **enq** the queuing function, **deq** the dequeuing function which assumes a non-empty queue as input, and **void** a Boolean test of emptiness. Notice that when q is non-empty, **deq** q returns a pair (x, q') where x is the dequeued value and q' the remaining queue.

A clean way of introducing such an abstraction in Coq that generates little overhead for program extraction towards OCaml is the use of a *module type* that collects the data of Fig. 2. Coq developments based on a hypothetical implementation of the module type are then organized as *functors* (i.e., modules depending on typed module parameters). Thus, all the Coq developments described in this section are such functors, and the extracted OCaml code is again a functor, now for the module system of OCaml, and with the module parameter that consists of the operations of Fig. 2 after stripping off the logical part. In other words, the parameter is nothing but a hypothetical implementation of a FIFO signature, viewed as a module type of OCaml.

Of course, the FIFO axioms have several realizations (or refinements), including a trivial and inefficient one where f2l is the identity function (and the inefficiency comes from appending the new elements at the end of the list with **enq**). In Sect. 7, we refine these axioms with more efficient implementations following Okasaki’s insights [13] that worst-case $\mathcal{O}(1)$ FIFO operations are even possible in an elegant way with functional programming languages.

5.2 Breadth-First Traversal

We use the equations which come from Theorem 3 and Theorem 5:

$$\text{bft}_f [] = [] \quad \text{bft}_f (\langle x \rangle :: l) = x :: \text{bft}_f l \quad \text{bft}_f (\langle a, x, b \rangle :: l) = x :: \text{bft}_f (l ++ [a; b])$$

²⁶ When a parameter X is marked implicit using the $\{X\}$ notation, it usually means that Coq is going to be able to infer the value of the argument by unification from the constraints in the context. In the case of $\text{f2l } l$, it means X will be deduced from the type of l which should unify with $\text{fifo } X$. While not strictly necessary, the mechanism of implicit arguments greatly simplifies the readability of Coq terms. In particular, it avoids an excessive use of dummy arguments ..

They suggest the definition of an algorithm for breadth-first traversal where lists are replaced with queues (FIFOs) so that (linear-time) append at the end $(\dots ++ [a; b])$ is turned into two primitive queue operations $(\text{enq } (\text{enq } \dots a) b)$. Hence, we implement FIFO-based breadth-first traversal.

Theorem 7. *There exists a fully specified Coq term*

$$\text{bft_fifo}_f : \forall q : \text{fifo } (\mathbb{T} X), \{l : \mathbb{L} X \mid l = \text{bft}_f (\text{f2l } q)\}$$

s. t. bft_fifo_f extracts to the following OCaml code:

```
let rec bft_fifo_f q =
  if void q then []
  else let t, q' = deq q
       in match t with
         | <x>      → x :: bft_fifo_f q'
         | <a, x, b> → x :: bft_fifo_f (enq (enq q' a) b).
```

Proof. We proceed by induction on the measure $q \mapsto \|\text{f2l } q\|$ following the method exposed in the interleave example of Sect. 2.5. The proof is structured around the computational content of the above OCaml code. Termination POs are easily solved by ω . Postconditions for correctness are proved using the above equations. $\square$

Corollary 2. *There is a Coq term $\text{bft}_{\text{fffo}} : \mathbb{T} X \rightarrow \mathbb{L} X$ s.t. $\text{bft}_{\text{fffo}} t = \text{bft}_{\text{std}} t$ holds for any $t : \mathbb{T} X$. Moreover, bft_{fffo} extracts to the following OCaml code:*

```
let bft_fffo t = bft_fifo_f (enq emp t).
```

Proof. From a tree t , we instantiate bft_fifo_f on the one-element FIFO $(\text{enq emp } t)$ and thus derive the term $\text{bft_fifo_full } (t : \mathbb{T} X) : \{l : \mathbb{L} X \mid l = \text{bft}_f [t]\}$. The first projection $\text{bft}_{\text{fffo}} t := \pi_1(\text{bft_fifo_full } t)$ gives us bft_{fffo} and we derive $\text{bft}_{\text{fffo}} t = \text{bft}_{\text{std}} t$ from the combination of the second projection $\text{bft}_{\text{fffo}} t = \text{bft}_f [t]$ with Theorem 6. $\square$

5.3 Breadth-First Numbering

Breadth-first numbering was the challenge proposed by Okasaki to the community and which led him to write his paper [14]. It consists in redecorating a tree with numbers in breadth-first order. The difficulty was writing an efficient algorithm in purely functional style. We choose the easy way to specify the result of breadth-first numbering of a tree: the output of the algorithm should be a tree $t : \mathbb{T} \mathbb{N}$ preserving the input shape and of which the breadth-first traversal of $\text{bft}_{\text{std}} t$ is of the form $[1; 2; 3 \dots]$.

As usual with those breadth-first algorithms, we generalize the notions to lists of trees.

Definition $\text{is.bfn } n \ l := \exists k, \text{bft}_f l = [n; \dots; n + k]$.

Lemma 2. *Given a fixed shape, the breadth-first numbering of a forest is unique, i. e., for any $n : \mathbb{N}$ and any $l, m : \mathbb{L}(\mathbb{T}\mathbb{N})$,*

$$l \sim_{\mathbb{L}\mathbb{T}} m \rightarrow \text{is_bfn } n \ l \rightarrow \text{is_bfn } n \ m \rightarrow l = m.$$

Proof. By Lemma `bft_f_inj` in Sect. 4.3. □

We give an equational characterization of breadth-first numbering of forests combined with list reversal. In the equations below, we intentionally consider a `bff` function that outputs the *reverse* of the numbering of the input forest, so that, when viewed as FIFOs of trees (instead of lists of trees), both the input FIFO over $\mathbb{T}X$ and the output FIFO over type $\mathbb{T}\mathbb{N}$ correspond to left dequeuing and right enqueueing.²⁷ That said, Eqs. (E2) and (E3) correspond to (Oka1) and (Oka2) respectively, augmented with an extra argument used for keeping track of the numbering.

Lemma 3. *Let $\text{bf}_f : \mathbb{N} \rightarrow \mathbb{L}(\mathbb{T}X) \rightarrow \mathbb{L}(\mathbb{T}\mathbb{N})$ be a term. Considering the following conditions:*

- (E1) $\forall n, \text{bf}_f \ n \ [] = []$;
- (E2) $\forall n \ x \ l, \text{bf}_f \ n \ (\langle x \rangle :: l) = \text{bf}_f \ (1 + n) \ l ++ [\langle n \rangle]$;
- (E3) $\forall n \ a \ x \ b \ l, \exists a' \ b' \ l',$
 $\text{bf}_f \ (1 + n) \ (l ++ [a; b]) = b' :: a' :: l' \wedge \text{bf}_f \ n \ (\langle a, x, b \rangle :: l) = l' ++ [\langle a', n, b' \rangle]$;
- (Bfn1) $\forall n \ l, l \sim_{\mathbb{L}\mathbb{T}} \text{rev}(\text{bf}_f \ n \ l)$;
- (Bfn2) $\forall n \ l, \text{is_bfn } n \ (\text{rev}(\text{bf}_f \ n \ l)).$

We have the equivalence: $(\text{E1} \wedge \text{E2} \wedge \text{E3}) \leftrightarrow (\text{Bfn1} \wedge \text{Bfn2}).$

Proof. From right to left, we essentially use Lemma 2. For the reverse direction, we proceed by induction on the measure $i, l \mapsto \|l\|$ in combination with Theorem 5 and Theorem 3—Equations (Oka1) and (Oka2). □

Although not explicitly written in Okasaki’s paper [14], these equations hint at the use of FIFOs as a replacement for lists for both the input and output of `bff`. Let’s see this informally for (E3): `bff n` is to be computed on a non-empty list viewed as a FIFO, and left dequeuing gives a composite tree $\langle a, x, b \rangle$ and the remaining list/FIFO l . The subtrees a and b are enqueued to the right of l and `bff (1 + n)` called on the resulting list/FIFO. (E3) guarantees that a' and b' can be dequeued to the left from the output list/FIFO. Finally, $\langle a', n, b' \rangle$ is enqueued to the right to give the correct result, thanks to (E3). This construction will be formalized in Theorem 8.

We define the specification `bff_fifo_f_spec` corresponding to breadth-first numbering of FIFOs of trees

$$\text{Definition } \text{bf}_f\text{-fifo-f_spec } n \ q \ q' := \text{f2l } q \sim_{\mathbb{L}\mathbb{T}} \text{rev}(\text{f2l } q') \wedge \text{is_bfn } n \ (\text{rev}(\text{f2l } q'))$$

and we show the inhabitation of this specification.

²⁷ The fact that input and output FIFOs operate in mirror to each other was already pointed out by Okasaki in [14]. Using reversal avoids defining two types of FIFOs or bi-directional FIFOs to solve the issue.

Theorem 8. *There exists a fully specified Coq term*

$$\text{bfn_fifo}_f : \forall (n : \mathbb{N}) (q : \text{fifo } X), \{q' \mid \text{bfn_fifo_f_spec } n \ q \ q'\}$$

which extracts to the following OCaml code:

```
let rec bfn_fifo_f n q =
  if void q then emp
  else let t, q0 = deq q in match t with
    | ⟨_⟩ → enq (bfn_fifo_f (1 + n) q0) ⟨n⟩
    | ⟨a, -, b⟩ → let b', q1 = deq (bfn_fifo_f (1 + n) (enq (enq q0 a) b)) in
                  let a', q2 = deq q1
                  in enq q2 ⟨a', n, b'⟩.
```

Proof. To define bfn_fifo_f , we proceed by induction on the measure $n, q \mapsto \|\text{f2l } q\|$ where the first parameter does not participate in the measure. As in Sect. 2.5, we implement a proof script which mixes tactics and programming style using the `refine` tactic. We strictly follow the above algorithm to design bfn_fifo_f . Of course, proof obligations like termination certificates or postconditions are generated by Coq and need to be addressed. As usual for these breadth-first algorithms, termination certificates are easy to solve thanks to the `omega` tactic. The only difficulties lie in the postcondition POs but these correspond to the (proofs of the) equations of Lemma 3. $\square$

Corollary 3. *There is a Coq term $\text{bfn_fifo} : \mathbb{T} X \rightarrow \mathbb{T} \mathbb{N}$ s. t. for any tree $t : \mathbb{T} X$:*

$$t \sim_{\mathbb{T}} \text{bfn_fifo } t \quad \text{and} \quad \text{bft}_{\text{std}}(\text{bfn_fifo } t) = [1; \dots; \|t\|]$$

and bfn_fifo extracts to the following OCaml code:

```
let bfn_fifo t = let t', _ = deq (bfn_fifo_f 1 (enq emp t)) in t'.
```

Proof. Obvious consequence of Theorem 8 in conjunction with Theorem 6. $\square$

5.4 Breadth-First Reconstruction

Breadth-first reconstruction is a generalization of breadth-first numbering—see the introduction for its description. For simplicity (since all our structures are finite), we ask that the list of labels that serves as extra argument has to be of the right length, i. e., has as many elements as there are labels in the input data-structure, while the “breadth-first labelling” function considered by Jones and Gibbons [7] just required it to be long enough so that the algorithm does not get stuck.

We define the specification of the breadth-first reconstruction of a FIFO q of trees in $\mathbb{T} X$ using a list $l : \mathbb{L} Y$ of labels

Definition $\text{bfr_fifo_f_spec } q \ l \ q' := \text{f2l } q \sim_{\mathbb{L} \mathbb{T}} \text{rev } (\text{f2l } q') \wedge \text{bft}_f(\text{rev } (\text{f2l } q')) = l.$

We can then define breadth-first reconstruction by structural induction on the list l of labels:

Fixpoint $\text{bfr_fifo}_f q l \{ \text{struct } l \} : \|f2l\ q\| = |l| \rightarrow \{q' \mid \text{bfr_fifo_f_spec } q l q'\}.$

Notice the precondition $\|f2l\ q\| = |l|$ stating that l contains as many labels as the total number of nodes in the FIFO q .²⁸ Since we use structural induction, there are no termination POs. There is however a precondition PO (easily proved) and postcondition POs are similar to those of the proof of Theorem 8. Extraction to OCaml outputs the following:

```
let rec bfr_fifo_f q = function
| []      → emp
| y :: l →
  let t, q0 = deq q in match t with
  | ⟨-⟩      → enq (bfr_fifo_f q0 l) ⟨y⟩
  | ⟨a, -, b⟩ → let b', q1 = deq (bfr_fifo_f (enq (enq q0 a) b) l) in
                 let a', q2 = deq q1
                 in enq q2 ⟨a', y, b'⟩.
```

Notice the similarity with the code of bfn_fifo_f of Theorem 8.

Theorem 9. *There is a Coq term $\text{bfr_fifo} : \forall (t : \mathbb{T} X) (l : \mathbb{L} Y), \|t\| = |l| \rightarrow \mathbb{T} Y$ such that for any tree $t : \mathbb{T} X$, $l : \mathbb{L} Y$ and $H : \|t\| = |l|$ we have:*

$$t \sim_{\mathbb{T}} \text{bfr_fifo } t l H \quad \text{and} \quad \text{bft}_{\text{std}} (\text{bfr_fifo } t l H) = l.$$

Moreover, bfr_fifo extracts to the following OCaml code:

```
let bfr_fifo t l = let t', _ = deq (bfr_fifo_f (enq emp t) l) in t'.
```

Proof. Direct application of $\text{bfr_fifo}_f (\text{enq emp } t) l$. □

6 Numbering by Levels

Okasaki reports in his paper [14] on his colleagues' attempts to solve the breadth-first numbering problem and mentions that most of them were level-oriented, as is the original traversal function bft_{std} . In his Sect. 4, he describes the “cleanest” of all those solutions, and this small section is devoted to get it by extraction (and thus with certification through the method followed in this paper).

We define $\text{children}_f : \forall \{K\}, \mathbb{L}(\mathbb{T} K) \rightarrow \mathbb{N} \times \mathbb{L}(\mathbb{T} K)$ such that, provably, $\text{children}_f l = (|l|, \text{subtrees } l)$ but using a more efficient simultaneous computation:

```
children_f [] := (0, [])
children_f (⟨-⟩ :: l) := (1 + n, m)
children_f (⟨a, -, b⟩ :: l) := (1 + n, a :: b :: m)
where (n, m) := children_f l
```

²⁸ This condition could easily be weakened to $\|f2l\ q\| \leq |l|$ but in that case, the specification bfr_fifo_f_spec should be changed as well.

and $\text{rebuild}_f : \forall \{K\}, \mathbb{N} \rightarrow \mathbb{L}(\mathbb{T} K) \rightarrow \mathbb{L}(\mathbb{T} \mathbb{N}) \rightarrow \mathbb{L}(\mathbb{T} \mathbb{N})$

$\text{rebuild}_f n [] := []$
 $\text{rebuild}_f n (\langle _ \rangle :: t_s) c_s := \langle n \rangle :: \text{rebuild}_f (1 + n) t_s c_s$
 $\text{rebuild}_f n (\langle _, _, _ \rangle :: t_s) (a :: b :: c_s) := \langle a, n, b \rangle :: \text{rebuild}_f (1 + n) t_s c_s$
 and otherwise $\text{rebuild}_f _ _ _ := []$.

Since we will need to use both children_f and rebuild_f for varying values of the type parameter K , we define them as fully polymorphic here.²⁹ We then fix $X : \text{Type}$ for the remainder of this section.

The algorithms children_f and rebuild_f are (nearly) those defined in [14, Figure 5] but that paper does not provide a specification for rebuild_f and thus cannot show the correctness result which follows. Instead of a specification, Okasaki offers an intuitive explanation of rebuild_f [14, p. 134]. Here, we will first rephrase the following lemma in natural language: the task is to obtain the breadth-first numbering of list t_s , starting with index n . We consider the list $\text{subtrees } t_s$ of all immediate subtrees, hence of all that is “at the next level” (speaking in Okasaki’s terms), and assume that c_s is the result of breadth-first numbering of those, but starting with index $|t_s| + n$, so as to skip all the roots in t_s whose number is $|t_s|$. Then, $\text{rebuild}_f n t_s c_s$ is the breadth-first numbering of t_s . In view of this description, the first three definition clauses of rebuild_f are unavoidable, and the last one gives a dummy result for a case that never occurs when running algorithm bfn_level_f in Theorem 10 below.

Lemma 4. *The function rebuild_f satisfies the following specification: for any $n, t_s : \mathbb{L}(\mathbb{T} X)$ and $c_s : \mathbb{L}(\mathbb{T} \mathbb{N})$, if both $\text{subtrees } t_s \sim_{\mathbb{L} \mathbb{T}} c_s$ and $\text{is_bfn } (|t_s| + n) c_s$ hold then $t_s \sim_{\mathbb{L} \mathbb{T}} \text{rebuild}_f n t_s c_s$ and $\text{is_bfn } n (\text{rebuild}_f n t_s c_s)$.*

Proof. First we show by structural induction on t_s that for any $n : \mathbb{N}$ and $t_s : \mathbb{L}(\mathbb{T} \mathbb{N})$, if $\text{roots } t_s = [n; n + 1; \dots]$ then $\text{rebuild}_f n t_s (\text{subtrees } t_s) = t_s$. Then we show that for any $Y, Z : \text{Type}$, $n : \mathbb{N}$, $t_s : \mathbb{L}(\mathbb{T} Y)$, $t'_s : \mathbb{L}(\mathbb{T} Z)$ and any $c_s : \mathbb{L}(\mathbb{T} \mathbb{N})$, if $t_s \sim_{\mathbb{L} \mathbb{T}} t'_s$ then $\text{rebuild}_f n t_s c_s = \text{rebuild}_f n t'_s c_s$. This second proof is by structural induction on proofs of the $t_s \sim_{\mathbb{L} \mathbb{T}} t'_s$ predicate. The result follows using Lemma 2. $\square$

The lemma suggests the recursive algorithm contained in the following theorem: in place of c_s as argument to rebuild_f , it uses the result of the recursive call on bfn_level_f with the index shifted by the number of terms in t_s (hence the number of roots in t_s —which is different from Okasaki’s setting with labels only at inner nodes) and the second component of $\text{children}_f t_s$.

Theorem 10. *There is a fully specified Coq term*

$$\text{bfn_level}_f : \forall (i : \mathbb{N}) (l : \mathbb{L}(\mathbb{T} X)), \{m \mid l \sim_{\mathbb{L} \mathbb{T}} m \wedge \text{is_bfn } i m\}$$

²⁹ Hence the $\forall \{K\}$ where K is declared as implicit.

which extracts to the following OCaml code:

```
let rec bfn_levelf i ts = match ts with
| [] → []
| _ → let n, ss = childrenf ts in rebuildf i ts (bfn_levelf (n + i) ss).
```

Proof. By induction on the measure $i, l \mapsto \|l\|$. The non-trivial correctness PO is a consequence of Lemma 4. $\square$

Corollary 4. *There is a Coq term $\text{bfn}_{\text{level}} : \mathbb{T} X \rightarrow \mathbb{T} \mathbb{N}$ s. t. for any tree $t : \mathbb{T} X$:*

$$t \sim_{\mathbb{T}} \text{bfn}_{\text{level}} t \quad \text{and} \quad \text{bft}_{\text{std}} (\text{bfn}_{\text{level}} t) = [1; \dots; \|t\|]$$

and $\text{bfn}_{\text{level}}$ extracts to the following OCaml code:

```
let bfn_level t = match bfn_levelf 1 [t] with t' :: _ → t'.
```

Proof. Direct application of Theorems 6 and 10. $\square$

7 Efficient Functional FIFOs

We discuss the use of our breadth-first algorithms that are parameterized over the abstract FIFO datatype with the implementations of efficient and purely functional FIFOs. As described in Sect. 5.1, the Coq developments of Sects. 5.2, 5.3 and 5.4 take the form of (Coq module) functors and their extracted code is structured as (OCaml module) functors. Using these functors means instantiating them with an implementation of the parameter, i. e., a module of the given module type. Formally, this is just application of the functor to the argument module. In our case, we implement Coq modules of the module type corresponding to Fig. 2 and then apply our (Coq module) functors to those modules. The extraction process yields the application of the extracted (OCaml module) functors to the extracted FIFO implementations. This implies a certification that the application is justified logically, i. e., that the FIFO implementation indeed satisfies the axioms of Fig. 2.

7.1 FIFOs Based on Two Lists

It is an easy exercise to implement our abstract interface for FIFOs based on pairs (l, r) of lists, with list representation $\text{f2l} (l, r) := l \mathbin{++} \text{rev } r$. The **enq** operation adds the new element to the *front* of r (seen as the tail of the second part), while the **deq** operation “prefers” to take the elements from the *front* of l , but if l is empty, then r has to be carried over to l , which requires reversal. It is well-known that this implementation still guarantees amortized constant-time operations if list reversal is done efficiently (in linear time). As before, we obtain the implementation by automatic code extraction from constructions with the rich specifications that use measure induction for **deq**.

We can then instantiate our algorithms to this specific implementation, while Jones and Gibbons [7] calculated a dedicated breadth-first traversal algorithm for this implementation from the specification.

We have the advantage of a more modular approach and tool support for the code generation (once the mathematical argument in form of the rich specification is formalized in Coq). Moreover, we can benefit from a theoretically yet more efficient and still elegant implementation of FIFOs, the one devised by Okasaki [13], to be discussed next.

7.2 FIFOs Based on Three Lazy Lists

While amortized constant-time operations for FIFOs seem acceptable—although imperative programming languages can do better—Okasaki showed that also functional programming languages allow an elegant implementation of worst-case constant time FIFO operations [13].

The technique he describes relies on lazy evaluation. To access those data structures in terms extracted from Coq code, we use coinductive types, in particular finite or infinite streams (also called “colists”):

CoInductive $\mathbb{S}X := \langle \rangle : \mathbb{S}X \mid _ \# _ : X \rightarrow \mathbb{S}X \rightarrow \mathbb{S}X$.

However, this type is problematic just because of the infinite streams it contains: since our inductive arguments are based on measures, we cannot afford that such infinite streams occur in the course of execution of our algorithms. Hence we need to guard our lazy lists with a purely logical *finiteness predicate* which is erased by the extraction process.

Inductive $\text{lfin} : \mathbb{S}X \rightarrow \text{Prop} :=$
 $\mid \text{lfin_nil} : \text{lfin } \langle \rangle$
 $\mid \text{lfin_cons} : \forall x s, \text{lfin } s \rightarrow \text{lfin } (x \# s)$.

We can then define the type of (finite) lazy lists as:

Definition $\mathbb{L}_l X := \{s : \mathbb{S}X \mid \text{lfin } s\}$.

Compared to regular lists $\mathbb{L}X$, for a lazy list $(s, Hs) : \mathbb{L}_l X$, on the one hand, we can also do pattern-matching on s but on the other hand, we cannot define **Fixpoints** by structural induction on s . We replace it with structural induction on the proof of the predicate $H_s : \text{lfin } s$. Although a bit cumbersome, this allows working with such lazy lists as if they were regular lists, and this practice is fully compatible with extraction because the guards of type $\text{lfin } s : \text{Prop}$, being purely logical, are erased at extraction time. Here is the extraction of the type $\mathbb{L}_l$ in OCaml:

```
type  $\alpha$  llist =  $\alpha$  _llist Lazy.t
and  $\alpha$  _llist = Lnil | Lcons of  $\alpha$  *  $\alpha$  llist
```

Here, we see the outcome of the effort: the (automatic) extraction process instructs OCaml to use lazy lists instead of standard lists (a distinction that does not even exist in the lazy language Haskell).

Okasaki [13] found a simple way to implement simple FIFOs³⁰ efficiently using triples of lazy lists. By efficiently, he means where enqueue and dequeue operations take *constant time* (in the worst case). We follow the proposed implementation using our own lazy lists $\mathbb{L}_l X$. Of course, his proposed code, being of beautiful simplicity, does not provide the full specifications for a correctness proof. Some important invariants are present, though. The main difficulty we faced was to give a correct specification for his intermediate `rotate` and `make` functions.

We do not enter more into the details of this implementation which is completely orthogonal to breadth-first algorithms. We invite the reader to check that we do get the exact intended extraction of FIFOs as triples of lazy lists; see the files `llists.v`, `fifo_3llists.v` and `extraction.v` described in Appendix A.

7.3 Some Remarks About Practical Complexity

However, our experience with the extracted algorithms for breadth-first numbering in OCaml indicate for smaller (with size below 2k nodes) randomly generated input trees that the FIFOs based on three lazy lists are responsible for a factor of approximately 5 in execution time in comparison with the 2-list-based implementation. For large trees (with size over 64k nodes), garbage collection severely hampers the theoretic linear-time behaviour. A precise analysis is out of scope for this paper.

8 Final Remarks

This paper shows that, despite their simplicity, breadth-first algorithms for finite binary trees present an interesting case for algorithm certification, in particular when it comes to obtain certified versions of efficient implementations in functional programming languages, as those considered by Okasaki [14].

Contributions. For this, we used the automatic extraction mechanism of the Coq system, whence we call this “breadth-first extraction.” Fine-tuning the proof constructions so that the extraction process could generate the desired code—without any need for subsequent code polishing—was an engineering challenge, and the format of our proof scripts based on a hand-crafted measure induction tactic (expressed in the Ltac language for Coq [5] that is fully usable as user of Coq—as opposed to Coq developers only) should be reusable for a wide range of algorithmic problems and thus allow their solution with formal certification by program extraction.

³⁰ He also found a simple solution for double-ended FIFOs.

We also considered variations on the algorithms that are not optimized for efficiency but illustrate the design space and also motivate the FIFO-based solutions. And we used the Coq system as theorem prover to formally verify some more abstract properties in relation with our breadth-first algorithms. This comprises as original contributions an axiomatic characterization of relations on paths in binary trees to be the breadth-first order (Theorem 1) in which the paths are visited by the breadth-first traversal algorithm. (Sect. 3.3). This also includes the identification of four different but logically equivalent ways to express the recursive behaviour of breadth-first traversal on forests (Theorem 3) and an equational characterization of breadth-first numbering of forests (Lemma 3). Among the variations, we mention that breadth-first reconstruction (Sect. 5.4) is amenable to a proof by structural recursion on the list of labels that is used for the relabeling while all the other proofs needed induction w. r. t. measures.

Perspectives. As mentioned in the introduction, code extraction of our constructions towards lazy languages such as Haskell would yield algorithms that we expect to work properly on infinite binary trees (the forests and FIFOs would still contain only finitely many elements, but those could then be infinite). The breadth-first nature of the algorithms would ensure fairness (hinted at also in the introduction). However, our present method does not certify in any way that use outside the specified domain of application (in particular, the non-functional correctness criterion of productivity is not guaranteed). We would have to give coinductive specifications and corecursively create their proofs, which would be a major challenge in Coq (cf. the experience of the second author with coinductive rose trees in Coq [16] where the restrictive guardedness criterion of Coq had to be circumvented in particular for corecursive constructions).

As further related work, we mention a yet different linear-time breadth-first traversal algorithm by Jones and Gibbons [7] that, as the other algorithms of their paper, is calculated from the specification, hence falls under the “algebra of programming” paradigm. Our methods should apply for that algorithm, too. And there is also their breadth-first reconstruction algorithm that relies on lazy evaluation of streams—a version of it which is reduced to breadth-first numbering has been discussed by Okasaki [14] to relate his work to theirs. To obtain such kind of algorithms would be a major challenge for the research in certified program extraction.

Another Related Research Direction. We mentioned directly related work throughout the paper, and we discussed certification of the program extraction procedure in Sect. 2.2. Let’s briefly indicate a complementary approach. We state in our Theorems 5, 7, 8, 9 and 10 the OCaml code we wanted to obtain by extraction (and that we then got), but there is no tool support to start with that code and to work towards the (fully specified) Coq terms. The `hs-to-coq`³¹ tool [19] transforms Haskell code (in place of OCaml that we chose to use) into Coq code and provides means to subsequent verification provided the Haskell code does not exploit non-termination.

³¹ <https://github.com/antalsz/hs-to-coq>.

Acknowledgments. We are most grateful to the anonymous reviewers for their thoughtful feedback that included numerous detailed suggestions for improvement of the presentation.

A Code Correspondence

Here, we briefly describe the Coq vernacular files behind our paper that is hosted at <https://github.com/DmxLarchey/BFE>. Besides giving formal evidence for the more theoretical characterizations, it directly allows doing program extraction, see the README section on the given web page.

We are here presenting 24 Coq vernacular files in useful order:

- **list_utils.v**: One of the biggest files, all concerning list operations, list permutations, the lifting of relations to lists (Sect. 2) and segments of the natural numbers – auxiliary material with use at many places.
- **wf_utils.v**: The subtle tactics for measure recursion in one or two arguments with a $\mathbb{N}$ -valued measure function (Sect. 2.4) – this is crucial for smooth extraction throughout the paper.
- **l1list.v**: Some general material on coinductive lists, in particular proven finite ones (including append for those), but also the rotate operation of Okasaki [13], relevant in Sect. 7.2.
- **interleave.v**: The example of interleaving with three different methods in Sects. 2.3 (with existing tools—needs Coq v8.9 with package **Equations**) and Sect. 2.5 (with our method).
- **zip.v**: Zipping with a rich specification and relations with concatenation – just auxiliary material.
- **sorted.v**: Consequences of a list being sorted, in particular absence of duplicates in case of strict orders – auxiliary material for Sect. 3.2.
- **increase.v**: Small auxiliary file for full specification of breadth-first traversal (Sect. 3.3).
- **bt.v**: The largest file in this library, describing binary trees (Sect. 3.1), their branches and orders on those (Sect. 3.2) in relation with breadth-first traversal and structural relations on trees and forests (again Sect. 3.1).
- **fifo.v**: the module type for abstract FIFOs (Sect. 5.1).
- **fifo_triv.v**: The trivial implementation of FIFOs through lists, mentioned in Sect. 5.1.
- **fifo_2lists.v**: An efficient implementation that has amortized $\mathcal{O}(1)$ operations (see, e.g., the paper by Okasaki [13]), described in Sect. 7.1.
- **fifo_3llists.v**: The much more complicated FIFO implementation that is slower but has worst-case $\mathcal{O}(1)$ operations, invented by Okasaki [13]; see Sect. 7.2.
- **bft_std.v**: Breadth-first traversal naively with levels (specified with the traversal of branches in suitable order), presented in Sect. 3.3.
- **bft_forest.v**: Breadth-first traversal for forests of trees, paying much attention to the recursive equations that can guide the definition and/or verification (Sect. 4.1).

- **bft_inj.v**: Structurally equal forests with the same outcome of breadth-first traversal are equal, shown in Sect. 4.3.
- **bft_fifo.v**: Breadth-first traversal given an abstract FIFO, described in Sect. 5.2.
- **bfn_spec_rev.v**: Characterization of breadth-first numbering, see Lemma 3.
- **bfn_fifo.v**: The certified analogue of Okasaki’s algorithm for breadth-first numbering [14], in Sect. 5.3.
- **bfn_trivial.v**: Just the instance of the previous with the trivial implementation of FIFOs.
- **bfn_level.v**: A certified reconstruction of **bfn** on page 134 (Sect. 4 and Fig. 5) of Okasaki’s article [14]. For its full specification, we allow ourselves to use breadth-first numbering obtained in **bfn_trivial.v**.
- **bfr_fifo.v**: Breadth-first reconstruction, a slightly more general task (see next file) than breadth-first numbering, presented in Sect. 5.4.
- **bfr_bfn_fifo.v**: Shows the claim that breadth-first numbering is an instance of breadth-first reconstruction (although they have been obtained with different induction principles).
- **extraction.v**: This operates extraction on-the-fly.
- **benchmarks.v**: Extraction towards .ml files.

References

1. Anand, A., Boulrier, S., Cohen, C., Sozeau, M., Tabareau, N.: Towards certified meta-programming with typed TEMPLATE-COQ. In: Avigad, J., Mahboubi, A. (eds.) ITP 2018. LNCS, vol. 10895, pp. 20–39. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-94821-8_2
2. Andronick, J., Felty, A.P. (eds.): Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2018, Los Angeles, CA, USA, 8–9 January 2018. ACM (2018). <http://dl.acm.org/citation.cfm?id=3176245>
3. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development. Coq’Art: The Calculus of Inductive Constructions. Texts in Theoretical Computer Science. Springer, Heidelberg (2004). <https://doi.org/10.1007/978-3-662-07964-5>
4. Cormen, T.H., Leiserson, C.E., Rivest, R.L.: Introduction to Algorithms. The MIT Press and McGraw-Hill Book Company (1989)
5. Delahaye, D.: A proof dedicated meta-language. Electr. Notes Theor. Comput. Sci. **70**(2), 96–109 (2002). [https://doi.org/10.1016/S1571-0661\(04\)80508-5](https://doi.org/10.1016/S1571-0661(04)80508-5)
6. Hupel, L., Nipkow, T.: A verified compiler from Isabelle/HOL to CakeML. In: Ahmed, A. (ed.) ESOP 2018. LNCS, vol. 10801, pp. 999–1026. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89884-1_35
7. Jones, G., Gibbons, J.: Linear-time breadth-first tree algorithms: an exercise in the arithmetic of folds and zips. Technical report, No. 71, Department of Computer Science, University of Auckland, May 1993
8. Larchey-Wendling, D., Monin, J.F.: Simulating induction-recursion for partial algorithms. In: Espírito Santo, J., Pinto, L. (eds.) 24th International Conference on Types for Proofs and Programs, TYPES 2018, Abstracts. University of Minho, Braga (2018). http://www.loria.fr/~larchey/papers/TYPES_2018_paper_19.pdf

9. Letouzey, P.: A new extraction for Coq. In: Geuvers, H., Wiedijk, F. (eds.) TYPES 2002. LNCS, vol. 2646, pp. 200–219. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-39185-1_12
10. Letouzey, P.: Programmation fonctionnelle certifiée - L'extraction de programmes dans l'assistant Coq. Ph.D. thesis, Université Paris-Sud, July 2004. https://www.irif.fr/~letouzey/download/these_letouzey_English.pdf
11. McCarthy, J.A., Fetscher, B., New, M.S., Feltey, D., Findler, R.B.: A Coq library for internal verification of running-times. *Sci. Comput. Program.* **164**, 49–65 (2018). <https://doi.org/10.1016/j.scico.2017.05.001>
12. Mullen, E., Pernsteiner, S., Wilcox, J.R., Tatlock, Z., Grossman, D.: (Euf: minimizing the Coq extraction TCB. In: Andronick and Felty [2], pp. 172–185. <https://doi.org/10.1145/3167089>
13. Okasaki, C.: Simple and efficient purely functional queues and dequeues. *J. Funct. Program.* **5**(4), 583–592 (1995)
14. Okasaki, C.: Breadth-first numbering: lessons from a small exercise in algorithm design. In: Odersky, M., Wadler, P. (eds.) Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP 2000), pp. 131–136. ACM (2000)
15. Paulson, L.C.: ML for the Working Programmer. Cambridge University Press, Cambridge (1991)
16. Picard, C., Matthes, R.: Permutations in coinductive graph representation. In: Pattinson, D., Schröder, L. (eds.) CMCS 2012. LNCS, vol. 7399, pp. 218–237. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32784-1_12
17. Sozeau, M.: Subset coercions in COQ. In: Altenkirch, T., McBride, C. (eds.) TYPES 2006. LNCS, vol. 4502, pp. 237–252. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74464-1_16
18. Sozeau, M.: Equations: a dependent pattern-matching compiler. In: Kaufmann, M., Paulson, L.C. (eds.) ITP 2010. LNCS, vol. 6172, pp. 419–434. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14052-5_29
19. Spector-Zabusky, A., Breitner, J., Rizkallah, C., Weirich, S.: Total Haskell is reasonable Coq. In: Andronick and Felty [2], pp. 14–27. <https://doi.org/10.1145/3167092>